



HAL
open science

Numerical methods for Gaussian discretizations in electronic structure theory problems

Ioanna-Maria Lygatsika

► **To cite this version:**

Ioanna-Maria Lygatsika. Numerical methods for Gaussian discretizations in electronic structure theory problems. Theoretical and/or physical chemistry. Sorbonne Université, 2024. English. NNT : 2024SORUS149 . tel-04693444

HAL Id: tel-04693444

<https://theses.hal.science/tel-04693444v1>

Submitted on 10 Sep 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Sorbonne Université

Méthodes numériques pour les discrétisations Gaussiennes des problèmes en structure électronique

par **Ioanna-Maria Lygatsika**

Thèse de doctorat de Mathématiques appliquées
dirigée par Yvon Maday et Jean-Philip Piquemal

École Doctorale de Sciences Mathématiques de Paris Centre (ED386)
Laboratoire Jacques-Louis Lions

Présentée et soutenue publiquement le 20 mars 2024

devant le jury composé de

<i>Rapporteurs</i>	Virginie Ehrlacher Filippo Lipparini	Professeure, École des Ponts & Inria Professeur adjoint, University of Pisa
<i>Présidente du jury</i>	Julia Contreras-Garcia	Directrice de recherche, CNRS & Sorbonne Université
<i>Examineurs</i>	Geneviève Dusson Antoine Levitt Marc Torrent	Chargée de recherche, CNRS & Université Bourgogne Franche-Comté Professeur junior, Université Paris-Saclay Ingénieur-chercheur, CEA Bruyères-le-Châtel
<i>Directeurs de thèse</i>	Yvon Maday Jean-Philip Piquemal	Professeur, Sorbonne Université Professeur, Sorbonne Université



Sorbonne Université

Numerical methods for Gaussian discretizations in electronic structure theory problems

by **Ioanna-Maria Lygatsika**

Ph.D thesis in Applied Mathematics
supervised by Yvon Maday and Jean-Philip Piquemal

Doctoral School of Mathematical Sciences of Central Paris (ED386)
Laboratoire Jacques-Louis Lions

Publicly defended on March 20th 2024

with a jury composed of

<i>Reviewers</i>	Virginie Ehrbacher Filippo Lipparini	Professor, École des Ponts & Inria Associate Professor, University of Pisa
<i>President of jury</i>	Julia Contreras-Garcia	Director of Research, CNRS & Sorbonne Université
<i>Examiners</i>	Geneviève Dusson Antoine Levitt Marc Torrent	Research Fellow, CNRS & Université Bourgogne Franche-Comté Junior Professor, Université Paris-Saclay Research Engineer, CEA Bruyères-le-Châtel
<i>Ph.D Advisors</i>	Yvon Maday Jean-Philip Piquemal	Professor, Sorbonne Université Professor, Sorbonne Université

Résumé

La simulation moléculaire est l'un des outils les plus courants de la chimie moderne. Les calculs réalisés au cours de ces simulations présentent souvent des difficultés, qui conduisent à une réduction de leurs performances lorsque les systèmes simulés sont des larges molécules composées de plusieurs atomes. Cette thèse se focalise sur les limitations liées à l'utilisation de fonctions de base centrées sur les atomes des molécules pour la discrétisation des équations de type Schrödinger, qui est un type de discrétisation très populaire en chimie quantique. Nous adoptons une approche d'analyse numérique pour formuler et traiter ces limitations. Le présent travail aborde deux questions fondamentales liées aux éléments de base de type gaussien centrés sur les atomes, à savoir l'évaluation des intégrales moléculaires sur les fonctions de base et la génération des éléments de cette base. Ces deux points ont un impact sur le coût de calcul et les exigences en mémoire des simulations moléculaires. Notre objectif principal est de concevoir des nouvelles méthodes mathématiques ainsi que des nouveaux algorithmes efficaces qui améliorent les simulations moléculaires modernes. Les principales contributions de cette thèse sont les deux suivantes : premièrement, l'accélération de l'évaluation des intégrales moléculaires sur les fonctions de base centrées sur les atomes et, deuxièmement, la proposition d'estimateurs d'erreur a posteriori pour les discrétisations centrées sur les atomes des problèmes linéaires à valeurs propres. Pour le premier objectif, nous avons développé une nouvelle méthode d'ajustement de densité ("Density Fitting" en anglais) pour l'approximation de la densité électronique, complétant les méthodes existantes dans la littérature, qui vise à réduire le coût de calcul en utilisant des approximations de rang faible et creuses, basées sur l'élimination des dépendances linéaires et la décomposition de Cholesky avec pivot. Notre schéma est présenté en utilisant un nouveau formalisme d'optimisation discrète et de recherche du plus court chemin sur les graphes. En outre, nous avons généralisé nos techniques en développant un nouveau schéma d'ajustement de densité indépendant des positions atomiques, en utilisant la méthode des bases réduites. La performance numérique de nos méthodes est démontrée par les résultats numériques des calculs d'énergie d'interaction intermoléculaire basée sur les densités en chimie. Pour le second objectif, notre travail constitue une extension de la théorie de l'estimation de l'erreur a posteriori basée sur les résidus des discrétisations gaussiennes sur des domaines non bornés. Un tel cadre, qui est couramment utilisé en chimie, n'a pas fait l'objet d'études théoriques dans la littérature mathématique jusqu'à présent. Notre contribution dans ce domaine permet la génération adaptative et automatique des bases centrées sur les atomes. Nous présentons des résultats numériques préliminaires des calculs en structure électronique afin d'illustrer un exemple d'application de nos estimateurs d'erreur a posteriori. En résumé, les bases de discrétisation centrées sur les atomes sont largement utilisées dans les simulations moléculaires. Les conclusions de cette thèse contribuent à la compréhension de telles bases du point de vue numérique, tout en proposant des solutions qui permettent l'amélioration des simulations moléculaires en chimie.

Mots clés : analyse numérique, structure électronique, orbitales gaussiennes atomiques, intégrales moléculaires, ajustement de densité, estimateurs d'erreur a posteriori

Abstract

Molecular simulation is among the most common tools in modern chemistry. Such simulations often suffer from several computational bottlenecks that reduce their performance when applied to large systems of molecules or atoms. This thesis primarily focuses on the limitations arising from the use of atom-centered basis functions for the discretization of Schrödinger-type equations for molecules, which is a popular type of discretization in quantum chemistry applications. We adopt a numerical analysis approach to formulate and tackle such limitations. The present work addresses two of the most impactful issues related to Gaussian-type atom-centered basis sets, namely, the evaluation of integrals on the basis functions and the generation of such basis sets. Both issues significantly affect the computational cost and memory requirements of molecular simulations. Our main goal is to design novel mathematical methods as well as new efficient low-complexity algorithms improving modern molecular simulations. The main contributions of this thesis are twofold: first, accelerating the evaluation of high-dimensional integrals on atom-centered basis functions, and, second, establishing a posteriori error estimators for atom-centered discretizations of linear eigenvalue problems. For the first purpose, we developed a new density fitting method for approximating the one-electron density, beyond the existing classical and robust density fitting methods of the literature, achieving tunable cost reduction via sparse low-rank approximation based on linear dependency elimination and the pivoted Cholesky decomposition. Our scheme is presented using a novel formalism of discrete optimization and shortest path search on graphs. In addition, we generalized our main techniques by developing a new atomic-position-independent density fitting scheme using the reduced basis method. The numerical performance of our methods is demonstrated by numerical results of an application to density-based intermolecular electrostatic interaction energy calculations in chemistry. For the second purpose, our work constitutes an extension of residual-based a posteriori error estimation theory to Gaussian discretizations over unbounded domains. Such a setting, which is routinely used in chemistry, was lacking theoretical investigation in the mathematical literature up to now. Our contribution on this domain paves the way towards adaptive and automatic generation of atom-centered basis sets. As numerical evidence, we present preliminary numerical results of an application to electronic structure theory calculations. To sum up, atom-centered Gaussian basis sets are widely used in molecular simulations. The conclusions of this thesis provide insights to the numerical analysis as well as to the computational aspects of the use of such basis sets in practice, while numerically demonstrating the ability of our methodologies to improve realistic simulations in chemistry.

Keywords: numerical analysis, electronic structure theory, gaussian-type atomic orbitals, molecular integrals, density fitting, a posteriori error estimators

Remerciements

Mes premiers remerciements sont adressés à mes directeurs de thèse, Yvon et Jean-Philip. Merci de m'avoir fait confiance et m'avoir donné l'opportunité de travailler au sein de votre équipe interdisciplinaire dans un milieu unique, celui du projet EMC2. C'était une chance pour moi d'avoir appris sous votre encadrement l'importance de repérer et de comprendre les liens entre mathématiques et calculs en chimie dans mon travail. La liberté que vous m'avez aussi accordée m'a permis d'avancer à mon rythme et de développer mes idées dans les meilleures conditions. Je n'aurais pas pu progresser et avoir accès à autant de sources de connaissance, ni m'exposer à des sujets si variés et stimulants, sans votre encouragement, votre soutien et nos discussions. Je souhaite également remercier Mi-Song Dupuy, tant pour ses conseils que pour sa disponibilité. Enfin, mon orientation vers ces thématiques mathématiques n'aurait pas pu s'effectuer sans notre collaboration avec Julia Contreras García et David Gontier : merci de m'avoir introduite au monde de la chimie quantique.

Merci à Virginie Ehrlacher et Filippo Lipparini d'avoir accepté de rapporter ma thèse, ainsi qu'à Julia Contreras García, Geneviève Dusson, Antoine Levitt et Marc Torrent de faire partie du jury. J'ai eu le privilège de me pencher sur vos recherches, ou même de collaborer avec une partie d'entre vous, et je suis honorée de vous savoir évaluateurs de cette thèse.

Je tiens à exprimer ma reconnaissance à Alston Misquitta de sa disponibilité, d'avoir éclairci ma compréhension des interactions intermoléculaires à travers ses explications pédagogiques et m'avoir aidée à rendre mon code *dfwpy* plus efficace grâce à ses propositions. Une autre pensée pour Geneviève Dusson et Benjamin Stamm qui ont eu l'ouverture et la patience de m'écouter exposer mes problèmes de recherche, la gentillesse de partager leur point de vue, inspirant ou aboutissant à des parties du présent manuscrit. Merci également à Susi Lehtola de nos discussions instructives au sujet de density fitting. Il me faut également remercier Louis Lagardère pour son accompagnement bienveillant et encourageant pendant ma découverte de logiciels en chimie. Pour avoir assuré mon comité de suivi, merci à Gabriel Stoltz et Mathieu Lewin.

Je tiens aussi à exprimer ma gratitude à toutes les personnes qui auront apporté leur aide à mes démarches administratives et logistiques : Malika Larcher, Erika Loyson, Salima Lounici, Corentin Maday mais aussi les membres de l'école doctorale, Valeria Banica, Corentin Lacombe et Jean-François Venuti. Merci également à Antoine Le Hyaric et Pierre-Henri Tournier pour leur support sur les machines de calcul *pascal* et *leto*, sans lesquels les résultats numériques de cette thèse n'auraient pas pu voir le jour.

Pour les nombreuses discussions intéressantes et tous ces mémorables moments, j'aimerais remercier les membres du Laboratoire de Chimie Théorique, en particulier ceux du couloir 33-34 : Diata (team muffin), Thomas, Théo, Pier, Nastasia, César, Margaret, Cassandra, Zoe, Bruno, ainsi qu'Olivier, notamment pour sa disponibilité et les discussions instructives et fort pédagogiques sur le calcul haute performance (et ses gâteaux); mais aussi ceux du couloir 12-13 : Trinidad, Andrea, Jessica, Matías, Timothée. Pour nos moments au-delà des couloirs de Jussieu et ma découverte de la Bretagne, j'adresse un grand merci aux autres membres de la team Roscoff : Siwar (team crevettes), Laurent, Gaspard, Etienne, Hassan, Edouard, Eloïse, Alfred, Solal, Baptiste, Alicia, Julien, Yanfei, Mathieu, Mohammad, Jean-Guillaume. Mes pensées vont également aux membres du Laboratoire Jacques-Louis Lions, en particulier ceux du bureau 15-16-221 mais aussi Yipeng et Gong pour leurs conseils, Rui, Chourouk, Roxane, Matthieu,

Agustin et Lise pour les soirées cinéma, ainsi que toutes les personnes que j'ai eu l'occasion de croiser au labo ou aux pauses café : Nicolai, Pauline, Fabrice, Thomas, Rémi, Jules, María, Pierre, Lucas, Elena, Emma, Jesus, Ramon, Fatima.

Rien de tout cela n'aurait été possible sans l'aide de mes parents, Hélène et Zénon. Un grand merci à ma famille, Ioanna, Christina, Vassilis, Kostas, Simos et ma grand-mère Ioanna pour leur soutien, ainsi qu'à mes amis de longue date pour leur présence quand il le fallait : Alexandra, Eirini, Manos, Constantinos et Tonia. Je tiens aussi à exprimer ma gratitude à la communauté des soeurs pour leur soutien inconditionnel, ainsi qu'aux membres du foyer, en particulier à Elena pour ses nombreux conseils.

Pour finir, je remercie Igor de ses conseils sur la rédaction du présent manuscrit ainsi que de son accompagnement pendant ces années de thèse.

Introduction

Context

The present thesis focuses on numerical problems arising in the context of molecular systems modeling. Molecular systems consist of atoms held together by chemical bonds. Since Dalton's discovery of the atom in 1808 [92], followed by Thomson's discovery of the electron in 1897 [318], these systems are the object of extensive study towards the understanding of matter.

During the early 20th century, the theoretical study of molecular systems, primarily the contributions of Bohr [46, 48, 47] and Rutherford's discovery of the atomic nucleus in 1911 [286], paved the way for understanding atoms and chemical bonds. The early models were based on concepts from atomic and nuclear physics as well as classical mechanics, and were only attempting to study small atoms, for example Hydrogen. The understanding of spectroscopy only became possible with the advent of quantum mechanics. Following the postulates proposed by Planck, Pauli and de Broglie on the quantum state of matter, Schrödinger proposed the equation that carries his name in 1925. This equation describes the physical properties of matter at the atomic and subatomic scale, and is the basis of modern quantum mechanics. The mathematical formulation of quantum mechanics was developed about one century ago, by Dirac, von Neumann [329], Hilbert and Weyl. Since then, new problems and concepts that lie at the intersection of mathematics and physics have been revealed on the topic of quantum mechanics, attracting great interest in the mathematical community to this day.

Computer simulations applied to condensed matter systems (typically, systems in solid or liquid phase) began in the 1950's [6]. Their aim is to help computing system properties that are predicted by numerical models, validating or sometimes replacing time-consuming and expensive experimental methods altogether. This active field of research is at the intersection of various disciplines, including chemistry, physics, biology, mathematics and computer science. The study of molecular systems contributes to the understanding of properties of materials, climate systems, planet formation, biochemical processes such as photosynthesis [216], biotechnological processes such as bioenergy and biofuels [254], and biophysical processes such as protein folding and RNA structure [98]. It lies at the core of knowledge about living organisms and biophysical systems, discovery of new materials and drugs.

The 1950's development of numerical simulation tools inspired new research in applied mathematics, including the study of numerical models for quantum mechanical systems. The mathematical tools that were developed for the analysis of such models, specifically the Hartree-Fock model [147, 116], led to remarkable mathematical concepts that find applications in other areas of mathematics as well, notably the concentration-compactness theory of P.L. Lions [205, 204], who received the Fields Medal for his work in 1994.

The evolution of molecular simulations is marked by numerical methods that allowed to extend the size of the molecular system under study, starting from the first molecular dynamics simulation on an idealized atomic system of 32 hard spheres in 1957 [5]. It was not until 1977 when the first molecular dynamics simulation of a small protein (882 atoms [224]) appeared [227]. Since the 1980's, the rapid advance of computer technology has gradually lead to a shift from conventional research to computed assisted research, namely computer aided drug design for drug discovery [15, 111, 209]. The modern advances on hardware architecture (parallel

computing, graphics processing units) is fully exploited by molecular simulation software packages [4, 140, 291], in order to perform high-performance scalable simulations on supercomputers and carry out breakthrough discoveries, as the multiscale simulation of SARS-CoV-2 virus' main protease (10^5 atoms including solvation) in 2021 [167]. Extreme-scale molecular dynamics is able to model systems of millions and even billions of atoms [176, 244, 254].

For practical application purposes, new kinds of models have been developed for different modeling scales, striking a balance between accuracy and computational efficiency. At the level of electrons, during the mid-50's has emerged a family of methods for numerically approximating solutions to electronic structure theory problems, namely the Hartree-Fock problem, based on Gaussian functions, suggested by Boys [54]. This greatly improved quantum mechanics simulations applied to molecular systems. However, biomolecular systems, made up of millions of atoms, still remain too large to be efficiently simulated at the quantum mechanics level using today's available computational resources. Besides, processes of interest in biomolecular systems often take place in multiple timescales [98, 75]. In this respect, force field models [80] are designed to break computational barriers and manage to compute the forces between atoms in biomolecular systems, leading to the computation of positions of interacting atoms with respect to time using molecular dynamics [6].

Present work

One of the main challenges in force fields is the computational workload required for evaluating quantities of interest, at any given timestep, which scale non-linearly with respect to system size [121]. This complexity can be a bottleneck for large systems or for processes requiring multiple simulation timesteps. An example is the design of computationally efficient algorithms for forces computation, that is still nowadays a hard task. Indeed, it has been reported using the high-performance Tinker-HP package [4] that it takes more than 95% of a timestep to evaluate forces when running a molecular dynamics simulation sequentially. In particular, 90% of this portion is the intermolecular forces between molecules.

In the present work, we focus on one of the existing methods to compute intermolecular interactions. During the mid-2000's has emerged a method known as the Gaussian Electrostatic Model (GEM) [257], aiming at computationally efficient intermolecular interaction evaluation without sacrificing the accuracy of electronic structure methods. The use of GEM presents various advantages in force field applications [320], which is the reason why it has gathered increasing interest during recent years. Among other features, GEM proposes to reduce the complexity of intermolecular forces, in particular their so-called *frozen core* energy part, by employing the Density Fitting method [340], designed to provide fast evaluation of pairwise interactions. In this respect, GEM relies upon Gaussian or Hermite-Gaussian discretizations for approximating atomic orbitals of isolated fragments at their ground state. However, the accuracy of density fitting using such discretizations for calculating frozen core energies in GEM is neither yet fully controlled nor previously studied in the mathematical literature. Systematic methods that allow to tune density fitting in order to reach a target level of accuracy will extend the applicability of GEM to larger systems. Figure 1 shows an example of a target system of application of GEM, that is a water solvent. Liquid water was first simulated using molecular dynamics in 1971 [275], but remains a challenge until today.

Contributions

To address such problems, the present thesis introduces various contributions.

Pairwise interaction problem formulation. We introduce a new general framework for treating frozen core energy computation problems in the context of GEM using finite discretization basis sets. This framework allows us to carry out a detailed comparison and numerical analysis of existing and new methods.




Figure 1: Visualization of interacting water fragments (H_2O) of an explicit water solvent system in three-dimensional space. Any given fragment is interacting pairwise with all other fragments. All pairwise interactions are summed at every timestep $t > 0$ of a molecular dynamics simulation, to obtain the intermolecular forces of the system. Such energy computation has quadratic complexity in the total number of fragments, creating a bottleneck.

Fast frozen core energy for rigid fragments. We reduce the complexity of pairwise interaction calculations using density fitting methods and extensions, suitable for rigid fragments (i.e. subunits of a given molecular system with fixed bond lengths and angles during the course of a molecular dynamics simulation). In practice, overlap matrices over Gaussian discretization bases underlying interactions of interest have certain properties, such as low-rank or block structure. We exploit these properties as well as basis-specific aspects. Our main contributions include, first, a new Auxiliary Basis Set (ABS) generation method for use in density fitting, and, second, a new class of sparse summation methods based on density fitting, named SDF, for Sparse Density Fitting, designed to precalculate the sparse structure of interactions, extending the capabilities of standard density fitting. SDF introduces several novel optimizations to further reduce the number of pairwise interactions.

Fast frozen core energy for flexible fragments. We reduce the cost of frozen core energy calculations on flexible fragments (intramolecular bond lengths and angles change during molecular dynamics) by introducing novel precomputation steps that can be used in simulations. Our main contribution is a new Density Fitting scheme coupled with Reduced Basis methods [270], named DF-RB. Our method employs a greedy auxiliary basis construction for use in density fitting, as well as multivariate interpolation for auxiliary coefficient computation.

A new library for computing frozen core energies. We identified the need to develop an easy-to-use library for computing frozen core energies between molecules for prototyping purposes. The aim of this library is validating a set of density fit algorithms on water clusters in particular, while exploring the full capabilities of our methods. Existing codes, namely `GEM_fit` [138], may lack this feature for two reasons. First, the Gaussian bases implemented in `GEM_fit` are limited to orbital *spd*-type for efficiency in the molecular integral evaluation [84]. Second, only two-center molecular integrals were implemented, not allowing to have access to the four-center reference ones for true error assesement. To deal with these issues, we implemented a new library for numerical testing purposes, named `dfwpy` for Density Fitting methods for Water molecules in Python. Our library interfaces with `PySCF` [312] to access all basis types. We also reimplemented

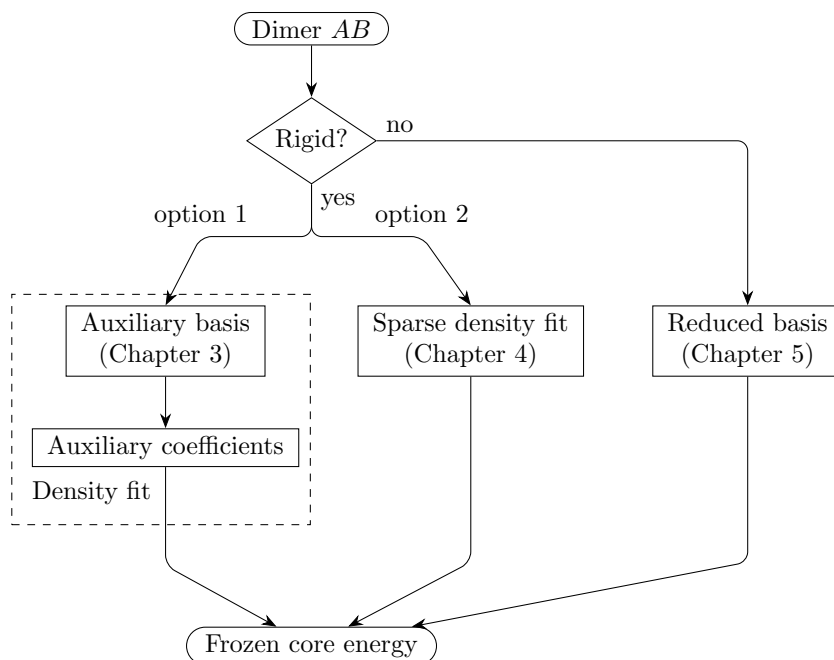


Figure 2: Flowchart describing the main framework for computing a single frozen core interaction energy using contributions of this thesis.

the intermolecular interaction calculation based on integration routines of `Libcint` [310] library, developed and optimized for Gaussian integrals. Our library thus provides numerical results on frozen core energies, both reference and density-fitted ones. A careful numerical comparison between our density fit methods and the state-of-the-art ones, using `dfwpy`, allows to discuss the different methodological choices.

***A posteriori* error estimators for atom-centered discretizations.** During this thesis, we also became interested in atomic orbital bases as an object of study on their own. This led us to the following work, that is not related to the frozen core energies of the remaining thesis, but rather focuses on discretizations from a mathematical perspective. One of the main practical results of this work consists of adaptive strategies, that help to choose on which atomic center one should add more basis functions, for a given molecule, in order to reduce the discretization error during an electronic structure calculation. To this end, we study linear Schrödinger-type eigenvalue problems in electronic structure theory. Our main contribution is a new residual-based practically computable *a posteriori* estimator for atom-centered discretization errors.

Overview

The present manuscript is organised as follows. Chapter 1 is dedicated to state-of-the-art numerical methods of the thesis, issued from electronic structure theory. Chapter 2 focuses on the formulation of the pairwise interaction problem, arising in the GEM frozen core energy computation. Figure 2 schematically presents the link between Chapters 3 to 5 in the goal of solving this problem. The remaining Chapter 6 deals with *a posteriori* error estimation. The conclusion and perspectives for future work are presented in Chapter 7. Implementation details on our `dfwpy` library are gathered in Appendix A. As a supplementing material, Appendix B contains our proof on Cartesian multipole expansions, used in applications of GEM to force fields, Appendix C summarizes practical details on Gaussian molecular integral evaluation and Appendix D introduces preliminary notions on model error estimation for equilibrium geometries.

Contents

Introduction	iii
Glossary of notations	xi
1 State of the art	1
1.1 Numerical methods in computational chemistry	2
1.1.1 Molecular dynamics	2
1.1.2 Electronic structure theory	3
1.1.3 The Hartree-Fock model	5
1.1.4 Linear combination of atomic orbitals	7
1.1.5 The electronic density function	12
1.1.6 The theory of intermolecular interactions	13
1.2 Large-scale problems for molecular integrals	19
1.2.1 The Coulomb matrix assembly	20
1.2.2 The frozen core energy	20
1.2.3 Acceleration methods	22
1.3 Molecular systems	26
1.3.1 Explicit water models	26
1.3.2 Case study: water clusters	27
1.4 Computational chemistry software	29
1.5 Summary	29
2 The pairwise interaction problem	31
2.1 Problem formulation	32
2.1.1 The pairwise interaction problem	32
2.1.2 From the Gaussian electrostatic model to interaction problems	32
2.2 Method overview	39
2.2.1 Existing work	39
2.2.2 Present work	40
2.2.3 Contributions	42
2.3 The density fitting method	44
2.3.1 Preliminaries	44
2.3.2 Definition of fitted density	46
2.3.3 Application to the pairwise interaction problem	47
2.4 Proofs	51
2.5 Conclusion	54
3 Auxiliary basis generation	55
3.1 Introduction	56
3.2 Theory	57
3.2.1 Problem formulation	57
3.2.2 From basis generation to submatrix selection	59

3.2.3	The case of electronic density functions	61
3.2.4	Auxiliary basis generation schemes	65
3.2.5	Summary of Auxiliary Basis Set (ABS) generation scheme	69
3.3	Numerical results	70
3.3.1	Frozen core energy of rigid fragments	71
3.3.2	Low-rank structure of Gram matrices	74
3.3.3	Auxiliary basis set detailed comparison	78
3.4	Proofs	81
3.5	Conclusion and perspectives	82
4	The sparse density fitting method	83
4.1	Introduction	84
4.2	Theory	84
4.2.1	Problem formulation	85
4.2.2	Sparse density fitting schemes	91
4.2.3	Summary of sparse density fitting schemes	96
4.3	Numerical results	101
4.3.1	Comparison of sparsification methods	102
4.3.2	Comparison of density fitting methods	105
4.4	Proofs	107
4.5	Conclusion and perspectives	110
5	Reduced basis methods for density fitting	111
5.1	Introduction	112
5.2	Theory	113
5.2.1	Setting	113
5.2.2	Reduced-auxiliary basis generation	114
5.2.3	Auxiliary coefficient interpolation	115
5.2.4	Summary of Reduced Basis Density Fitting (RB-DF) scheme	117
5.2.5	Computational complexity	117
5.3	Numerical results	119
5.3.1	Flexible water clusters	120
5.3.2	Reduced-auxiliary basis accuracy	120
5.3.3	Interpolation error	122
5.3.4	Pairwise interaction errors	124
5.4	Conclusion and perspectives	126
6	Adaptive methods for atomic basis sets	129
6.1	Introduction	131
6.2	Theory	132
6.2.1	Setting	132
6.2.2	Variational approximation	132
6.2.3	Preliminaries	133
6.3	Results	135
6.3.1	Guaranteed <i>a posteriori</i> error estimates	135
6.3.2	Practical estimates	135
6.3.3	Adaptive refinement strategy	137
6.4	Numerical results	138
6.4.1	1D diatomic molecules	138
6.4.2	Quality of the error estimator	139
6.4.3	Influence of numerical parameters	140
6.4.4	Adaptive basis sets	142
6.5	Proofs	142
6.5.1	Preliminaries	142

6.5.2	Source problem	144
6.5.3	Eigenvalue problem - any single eigenvalue	145
6.6	Conclusion	148
7	Conclusion and perspectives	151
	Appendices	155
A	dfwpy: Density Fitting for Water molecules in Python	157
A.1	Algorithmic design	157
A.2	Linear algebra	157
A.3	Water fragments	157
A.3.1	Local to global transformation	158
A.4	Molecular integrals	158
A.5	Sparsification	159
A.6	Rigid fragments	159
A.6.1	Operators	159
A.6.2	Optimizations	159
A.7	Flexible fragments	162
A.7.1	Transformation between frames	162
A.8	Multipole moments	162
B	Force fields	163
B.1	The Cartesian multipole expansion	163
B.1.1	Related numerical aspects	165
B.2	Application to chemistry	166
B.2.1	Modeling electrostatics of charge distributions	166
B.2.2	Force fields	166
B.2.3	The AMOEBA polarizable force field	169
B.2.4	Generation of multipoles from GEM	171
B.3	Summary	173
C	Molecular integral evaluation	175
C.1	Preliminaries	175
C.1.1	McMurchie-Davidson scheme	176
C.2	Gaussian multipole expansion	177
D	Error bars for equilibrium geometries	181
	Bibliography	182

Glossary of notations

N_e electron number

N_a atom number

$\mathbf{r} = (x, y, z)$ electron coordinates

$\mathbf{R} = (R_x, R_y, R_z)$ nuclear coordinates

α Gaussian orbital exponent

$\ell = l + m + n$ angular momentum quantum number

N_b atomic orbital basis set size

$\mu, \nu, \kappa, \lambda$ indices of atomic orbital basis functions

$\chi_\mu(\mathbf{r})$ atomic Gaussian-type orbital basis function

$(\mu\nu|\kappa\lambda)$ bielectronic integral on atomic orbitals

$\mathbf{D} = (D_{\mu\nu})_{1 \leq \mu, \nu \leq N_b}$ density matrix

$\rho(\mathbf{r})$ electronic density function

M auxiliary basis set size

i, j indices of auxiliary basis functions

$\xi_i(\mathbf{r})$ atomic Gaussian-type auxiliary basis function

$\tilde{\rho}(\mathbf{r})$ auxiliary density function

F, A, B molecular fragments

$K(\mathbf{r}, \mathbf{r}')$ interaction kernel function

\mathbb{H} real Hilbert space

$\mathbb{H}^1(\mathbb{R}^d)$ first order Sobolev space over \mathbb{R}^d

Chapter 1

State of the art

Contents

1.1 Numerical methods in computational chemistry	2
1.1.1 Molecular dynamics	2
1.1.2 Electronic structure theory	3
1.1.3 The Hartree-Fock model	5
1.1.4 Linear combination of atomic orbitals	7
1.1.5 The electronic density function	12
1.1.6 The theory of intermolecular interactions	13
1.2 Large-scale problems for molecular integrals	19
1.2.1 The Coulomb matrix assembly	20
1.2.2 The frozen core energy	20
1.2.3 Acceleration methods	22
1.3 Molecular systems	26
1.3.1 Explicit water models	26
1.3.2 Case study: water clusters	27
1.4 Computational chemistry software	29
1.5 Summary	29

The present chapter introduces several methods to compute interaction energies between atoms or molecules of large systems in chemistry. Among the various existing methods for evaluating interaction energies, we will focus on the Gaussian Electrostatic Model (GEM) in particular, designed to use Gaussian-type orbital basis sets. Unfortunately, the use of such basis sets can cause a computational bottleneck when it comes to interaction computations. In order to motivate this problem, we review basic notions and concepts related to it. In Section 1.1, we first present the chemistry context and our understanding in this domain, as an attempt to illustrate the range of applications of the problems studied in this thesis. In Section 1.2, we describe how the problem of calculating intermolecular interactions using Gaussian-type orbitals leads to a computationally intensive summation and discuss existing solutions, as well as related well-studied problems, namely the Coulomb matrix assembly. Finally, in Section 1.3, we provide useful background on test case data and computer software used throughout the thesis.

1.1 Numerical methods in computational chemistry

The present section introduces numerical methods to solve problems in chemistry. The concept of molecular dynamics is presented (see Section 1.1.1), that involves methods to compute forces between atoms, followed by a formulation of electronic structure theory for describing the behavior of electrons (see Section 1.1.2). We address a popular discretization family used for electronic problems, which is the linear combination of atomic orbitals (see Section 1.1.4). The section ends with the definition of the frozen density approximation for calculating interaction energies between molecules (see Section 1.1.6), motivating one of the main approaches used in this thesis for improving molecular simulations.

1.1.1 Molecular dynamics

We focus on one of the two cornerstones of molecular simulation methods, that is Molecular Dynamics (MD) [80], allowing to simulate the movement of molecular systems evolving in time, the other one, not considered in this thesis, being Monte Carlo sampling [231].

Suppose that we have a molecular system, consisting of nuclei and electrons. Fundamental to our study is the separation of electronic and nuclear motion provided by the Born-Oppenheimer (BO) approximation [49]. It allows to assume that nuclei are moving within a potential created by electrons confined to a specific electronic state. This assumption is supported by the physical observation that electrons ($m_e = 5 \times 10^{-31}$ kg) are much lighter than nuclei, made up of protons and neutrons ($m_p = m_n = 10^{-27}$ kg). As a result, electrons react instantaneously to the motion of nuclei. A second level of approximation consists of treating nuclei as classical point particles that follow classical dynamics. This leads to the following equations of motion.

Consider a molecular system of N_a atoms. For any atom indexed by i , for all $1 \leq i \leq N_a$, we denote by $\mathbf{R}_i \in \mathbb{R}^3$ its position in Cartesian coordinates, $\mathbf{v}_i \in \mathbb{R}^3$ its velocity and m_i its mass. In MD, the positions and velocities are evolved according to the Newton's equation of motion, that writes as a system of first order differential equations (ODEs): for all $1 \leq i \leq N_a$,

$$\mathbf{F}_i = m_i \frac{d\mathbf{v}_i}{dt}, \quad (1.1)$$

$$\mathbf{v}_i = \frac{d\mathbf{R}_i}{dt}, \quad (1.2)$$

where $\mathbf{F}_i \in \mathbb{R}^3$ is the force acting on atom i , that is a vector given by

$$\mathbf{F}_i = -\nabla_i U(\mathbf{R}_1, \dots, \mathbf{R}_{N_a}). \quad (1.3)$$

Eq. (1.3) derives from classical mechanics, according to which the conservative force on atom i can be written as the negative gradient in the i -th coordinate of a function of all atomic positions, called *potential energy*, denoted by U . It is assumed that a potential energy model provides this quantity in the form of a well-defined function

$$U : \mathbb{R}^{3N_a} \rightarrow \mathbb{R}. \quad (1.4)$$

The system of Eq. (1.1) and Eq. (1.2) is solved using time discretization schemes. The computation of atomic positions and velocities can be decomposed into two main stages. One is the update of atomic positions and velocities at each time step, using Verlet or multiple time-scale symplectic integrators [26] or their extensions [192]. The other stage is the calculation of forces on each atom at each timestep. For this, energy gradients must be evaluated. Both computational steps rely on the potential energy model that assigns an expression to U and captures the effect of electrons. Yet, the construction of energy models is a nontrivial task. The model influences directly the accuracy of the dynamics, as it gives the interaction between all atoms of the system. This work focuses on potential energy models designed for MD simulations of large systems, consisting of thousands of atoms, for timescales of the order of nanoseconds

(10^{-9} s) [75]. Note that this timescale is still too short to capture several processes, including folding of proteins or RNAs [75], taking place on the longer scale of microseconds (10^{-6} s) [98]. Lastly, a typical MD timestep is of the order of femtoseconds (10^{-15} s). Since many steps are required to cover the simulation time, the cost of evaluating the potential energy at each timestep highly affects the total computational cost of MD. The computationally low-cost potential energy models is the topic this work.

Note that, the assumption that nuclei can be treated as classical particles has certain limitations [67, 304], which may have an impact on biomolecular simulations [252]. Alternative approaches that account for quantum mechanical description of nuclear coordinates include path integral molecular dynamics (PIMD) [222] under the BO approximation. Such approaches can be computationally much more expensive than classical MD.

Potential energy surface

The potential energy has an interest within MD, as well as an object of study on its own.

In general, useful insight on properties of a system can be obtained by studying the graph of its potential energy function in Eq. (1.4), known as Potential Energy Surface (PES). PES is a $3N_a$ -dimensional object. It provides information on dynamics, as well as structure and thermodynamics. In particular, the local minima of U correspond to mechanically stable configurations, where the forces of the system vanish. Elsewhere, the non-vanishing forces determine all the classical dynamical properties. The thermodynamical properties depend upon U via ensemble averages [330]. Lastly, PES can be used for describing how the chemical reactivity of a system varies with structural changes. Recall that PES is defined within BO approximation. Cases where the BO approximation breaks down can be studied using a variation known as exact time-dependent PES [2, 1, 191].

PES is a high-dimensional object whose number of stationary points varies with system size and also depends on the system, for instance on its underlying interactions that may be weak or strong. As an example, ten stationary points have been identified in the PES of the water dimer (6 atoms) [300, 163]. For proteins ($> 10^3$ atoms), the number of stationary points may increase exponentially with system size [306, 330]. Among them, several local minima provide useful information as they are associated with intermediate conformations of protein folding [287].

1.1.2 Electronic structure theory

Electronic structure models the behavior of electrons in atoms and molecules. After defining the main objects of this theory, we discuss its application to computing potential energies.

Consider a molecular system of N_e electrons and N_a nuclei. Atomic units are assumed, in which the electronic charge and mass are equal to one. Within the previously introduced notion of BO approximation, nuclei are treated as classical particles, with positions $(\mathbf{R}_i)_{1 \leq i \leq N_a} \in \mathbb{R}^{3N_a}$ and nuclear charges $(z_i)_{1 \leq i \leq N_a} \in \mathbb{R}^{N_a}$. The electrons are treated as quantum particles. The electronic state is described by a function of the electronic coordinates $(\mathbf{r}_i)_{1 \leq i \leq N_e} \in \mathbb{R}^{3N_e}$, denoted by $\Psi_e : \mathbb{R}^{3N_e} \rightarrow \mathbb{R}$, called *electronic wave function*. Such level of description of the system derives from first principles (*ab initio*). The N_e -electron wave function Ψ_e belongs to the so-called *Fock space*, denoted by \mathcal{F}_{N_e} , which is defined as the subspace of the tensor product $\otimes_{i=1}^{N_e} \mathbb{H}^1(\mathbb{R}^3)$ of N_e single-particle Hilbert spaces, whose elements satisfy the antisymmetry property with respect to the exchange of electronic coordinates. This antisymmetry property for electrons, known as Pauli exclusion principle, reads, for any two distinct indices $1 \leq i \neq j \leq N_e$, as

$$\Psi_e(\mathbf{r}_1, \dots, \mathbf{r}_i, \dots, \mathbf{r}_j, \dots, \mathbf{r}_{N_e}) = -\Psi_e(\mathbf{r}_1, \dots, \mathbf{r}_j, \dots, \mathbf{r}_i, \dots, \mathbf{r}_{N_e}). \quad (1.5)$$

We write

$$\Psi_e \in \mathcal{F}_{N_e} := \bigwedge_{i=1}^{N_e} \mathbb{H}^1(\mathbb{R}^3),$$

where \wedge denotes the antisymmetric tensor product of vector spaces. The electronic Hamiltonian specifying the electronic energy of the system is an operator acting on the Fock space, defined as

$$\mathcal{H}_e^{(\mathbf{R}_i)} = \sum_{i=1}^{N_e} -\frac{1}{2}\Delta_{\mathbf{r}_i} - \sum_{i=1}^{N_e} \sum_{j=1}^{N_a} \frac{z_j}{|\mathbf{r}_i - \mathbf{R}_j|} + \sum_{1 \leq i < j \leq N_e} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|}, \quad (1.6)$$

where the first term is the electronic kinetic energy, the second term is the Coulomb attraction between electron and nuclei, and the third term is the Coulomb repulsion between electrons.

1.1.2.1 Minimization problem

For fixed nuclear configuration, the electronic structure of the system is described by the stationary Schrödinger equation

$$\mathcal{H}_e^{(\mathbf{R}_i)} \Psi_e = E \Psi_e,$$

that is an eigenvalue problem. The eigenvalue E is the electronic potential energy of the system and the corresponding eigenfunction Ψ_e is the electronic state. Note that, typically, the wave function is normalized, because the probability of finding an electron in the entire space is one. From a mathematical point of view, the eigenvalue $E \in \mathbb{R}$ is the Lagrange multiplier associated to the normalization constraint in the minimization of the energy, that we recall: the energy is a measurable quantity expressed as the expectation value of the Hamiltonian operator, that is equal to the Rayleigh quotient

$$\begin{aligned} E_0 = E(\Psi_0; \mathbf{R}_1, \dots, \mathbf{R}_{N_a}) &:= \frac{\langle \Psi_0, \mathcal{H}_e^{(\mathbf{R}_i)} \Psi_0 \rangle_{L^2(\mathbb{R}^{3N_e})}}{\langle \Psi_0, \Psi_0 \rangle_{L^2(\mathbb{R}^{3N_e})}} \\ &= \min_{\substack{\Psi_e \in \mathcal{F}_{N_e} \\ \Psi_e \neq 0}} \frac{\langle \Psi_e, \mathcal{H}_e^{(\mathbf{R}_i)} \Psi_e \rangle_{L^2(\mathbb{R}^{3N_e})}}{\langle \Psi_e, \Psi_e \rangle_{L^2(\mathbb{R}^{3N_e})}}, \end{aligned} \quad (1.7)$$

where $\langle \cdot, \cdot \rangle_{L^2(\mathbb{R}^{3N_e})}$ denotes the standard L^2 -inner product for real-valued square-integrable functions defined over \mathbb{R}^{3N_e} . The smallest signed eigenvalue of the electronic Hamiltonian corresponds to the electronic *ground state energy* of the system (more details can be found in [60, §1.2]).

This is a constraint optimization problem on the infinite-dimensional Fock space. The quantity E_0 as a function of the nuclear configuration corresponds to the nuclear potential energy contribution due to the presence of the negatively charged electron cloud, assuming that electrons are confined to their ground state. One can then define the total potential energy of the system as

$$U(\mathbf{R}_1, \dots, \mathbf{R}_{N_a}) := E_0(\Psi_0; \mathbf{R}_1, \dots, \mathbf{R}_{N_a}) + \sum_{1 \leq i < j \leq N_a} \frac{z_i z_j}{|\mathbf{R}_i - \mathbf{R}_j|}, \quad (1.8)$$

where the last term is the classical Coulomb repulsion between nuclei contributing to the nuclear potential energy.

1.1.2.2 Potential energy model from first principles

The evaluation of potential energies from true first principles is still computationally unfeasible using modern computers, even for small systems. Such limitation is related to dimensionality, discussed here.

The potential energy model of Eq. (1.8), obtained from first principles, indeed depends explicitly on the result of the minimization problem of Eq. (1.7). It is known that the latter suffers from the curse of dimensionality, which is described as follows. Suppose we apply a spatial discretization for the space of electronic coordinates using a grid of d points per Cartesian direction. Since we have three directions and N_e electronic coordinates, the total number of

degrees of freedom, i.e. the number of possible values that Ψ_e can take on the grid, is d^{3N_e} . The exponential scaling with respect to the electron number N_e is computationally prohibitive even for small-sized systems with more than one electron [14]. An example of small system, water H_2O (3 atoms) has $N_e = 10$ electrons. One of the smallest proteins, Deca-alanine $\text{C}_{30}\text{H}_{52}\text{N}_{10}\text{O}$ (93 atoms) has $N_e = 310$ electrons. Lastly, one of the realistic biomolecular systems, the CD2-CD58 protein (97.594 atoms), has a number of electrons of the order of 10^4 . For such reasons, in this work we work with simplifications of the ground state problem of Eq. (1.7) that allow to compute approximate potential energies from first principles for large systems.

1.1.3 The Hartree-Fock model

After introducing one of the most popular approximations for breaking the curse of dimensionality of the electronic ground state problem, namely the Hartree-Fock approximation, we describe how to compute ground state energies numerically. Note that the Hartree-Fock model also serves as a toy problem for introducing general electronic structure theory notions, useful for the rest of the thesis, two of the most important ones being *molecular orbitals* and the *Coulomb matrix*.

In order to treat the curse of dimensionality, Hartree-Fock theory [147, 116] replaces the Fock space with a subspace achieving separation in N_e variables, N_e being the number of electrons. This idea is based on introducing a special type of wave function called *Slater determinant*, defined as

$$\Psi_e(\mathbf{r}_1, \dots, \mathbf{r}_{N_e}) = \frac{1}{\sqrt{N_e!}} \det(\phi_i(\mathbf{r}_j))_{1 \leq i, j \leq N_e},$$

where each ϕ_i is a normalized one-electron function of $\text{H}^1(\mathbb{R}^3)$ referred to as *Molecular Orbital* (MO). A Slater determinant is an antisymmetric function that belongs to the Fock space. The space of Slater determinants is denoted by

$$\mathcal{S}_{N_e} = \left\{ \Psi_e(\mathbf{r}_1, \dots, \mathbf{r}_{N_e}) = (N_e!)^{-1/2} \det(\phi_i(\mathbf{r}_j))_{1 \leq i, j \leq N_e} : (\phi_i)_{1 \leq i \leq N_e} \in \mathcal{W}_{N_e} \right\}, \quad (1.9)$$

and the space of configurations of N_e molecular orbitals by

$$\mathcal{W}_{N_e} = \left\{ \Phi = (\phi_i)_{1 \leq i \leq N_e} : \phi_i \in \text{H}^1(\mathbb{R}^3), \int_{\mathbb{R}^3} \phi_i \phi_j = \delta_{ij}, \quad 1 \leq i, j \leq N_e \right\}.$$

Evaluation of the electronic potential energy for a Slater determinant $\Psi_e \in \mathcal{S}_{N_e}$ associated to the molecular orbital configuration $\Phi \in \mathcal{W}_{N_e}$ yields that the Hartree-Fock electronic energy is given by

$$\begin{aligned} E^{\text{HF}} &:= E^{\text{HF}}(\Phi; \mathbf{R}_1, \dots, \mathbf{R}_{N_a}) \\ &= \langle \Psi_e, \mathcal{H}_e^{(\mathbf{R}_i)} \Psi_e \rangle_{L^2(\mathbb{R}^{3N_e})} \\ &= \frac{1}{2} \sum_{i=1}^{N_e} \int_{\mathbb{R}^3} |\nabla \phi_i|^2 - \sum_{i=1}^{N_e} \sum_{j=1}^{N_a} z_j \int_{\mathbb{R}^3} \frac{|\phi_i(\mathbf{r})|^2}{|\mathbf{r} - \mathbf{R}_j|} \, \text{d}\mathbf{r} \\ &\quad + \frac{1}{2} \sum_{i=1}^{N_e} \sum_{j=1}^{N_e} \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{|\phi_i(\mathbf{r}) \phi_j(\mathbf{r}')|^2}{|\mathbf{r} - \mathbf{r}'|} \, \text{d}\mathbf{r} \, \text{d}\mathbf{r}' - \frac{1}{2} \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{\left| \sum_{i=1}^{N_e} \phi_i(\mathbf{r}) \phi_i(\mathbf{r}') \right|^2}{|\mathbf{r} - \mathbf{r}'|} \, \text{d}\mathbf{r}' \, \text{d}\mathbf{r}, \end{aligned}$$

where the first term of the sum is the kinetic energy of molecular orbitals, the second term is the electrostatic interaction between nuclei and electrons, the third term is the Coulomb repulsion between electrons, that can be interpreted as the classical Coulomb energy of electrons, and the last term is the exchange term coming from the antisymmetry property of the wave function. Finding the electronic ground state energy then comes down to solving the electronic Hartree-Fock ground state problem

$$E_0^{\text{HF}}(\mathbf{R}_1, \dots, \mathbf{R}_{N_a}) := \min_{\Phi \in \mathcal{W}_{N_e}} E^{\text{HF}}(\Phi; \mathbf{R}_1, \dots, \mathbf{R}_{N_a}). \quad (1.10)$$

The Hartree-Fock energy is a non-linear and non-convex functional with respect to Φ . The existence of a minimiser on \mathcal{W}_{N_e} was demonstrated by P.L. Lions [206], and Lieb and Simon [202], to which the following result is attributed (see [60, Proposition 5.1]), leading to the mathematical foundations of Hartree-Fock theory.

Theorem 1.1.1 (Compactness of the Hartree-Fock model). *Let us suppose that the total nuclear charge $Z = \sum_{i=1}^{N_a} z_i$ satisfies $Z > N_e - 1$. All minimising sequences of the Hartree-Fock problem of Eq. (1.10) are relatively compact in $(H^1(\mathbb{R}^3))^{N_e}$ and converge up to the extraction of a subsequence, in particular there exists a minimum.*

To summarize, the Hartree-Fock model breaks the curse of dimensionality. In particular, the number of degrees of freedom of Slater determinants scales linearly with the number of electrons. For example, using d number of points in each Cartesian direction of electronic coordinates, the total number of degrees of freedom is $3dN_e$. Using the Hartree-Fock approximation, we are able to compute ground state energies of systems containing hundreds and even thousands of atoms [346]. Note that, still, such computations remain too expensive in terms of computing time for many purposes. In this thesis, we will use the Hartree-Fock model for small systems with few atoms, such as water, for which it provides a computationally accessible ground state approximation. A weakness of the Hartree-Fock model, for systems of any size, is the approximation accuracy. Fortunately, this can be improved using post-Hartree-Fock techniques, such as Coupled Cluster [87] and Configuration Interaction [193] methods. The idea of these techniques is to write the wave function as a sum of Slater determinants.

The mathematical foundations of Hartree-Fock theory have led to the development of several numerical approximation schemes, discussed in the following section.

1.1.3.1 Discretization of the Hartree-Fock problem

The basic step for solving the Hartree-Fock problem using a computer is to approximate molecular orbitals in a finite basis set, presented as follows.

The Hartree-Fock problem of Eq. (1.10) is solved numerically using a Galerkin approximation. The main idea is to constrain the one-electron molecular orbitals

$$\forall 1 \leq i \leq N_e, \quad \phi_i \in H^1(\mathbb{R}^3)$$

to a finite-dimensional subspace of $H^1(\mathbb{R}^3)$, denoted by \mathcal{V} . Then, we introduce the discretized problem

$$\min_{\Phi \in \mathcal{W}_{N_e}(\mathcal{V})} E^{\text{HF}}(\Phi; \mathbf{R}_1, \dots, \mathbf{R}_{N_a}), \quad (1.11)$$

where $\mathcal{W}_{N_e}(\mathcal{V})$ is the subspace of \mathcal{W}_{N_e} defined as

$$\mathcal{W}_{N_e}(\mathcal{V}) = \left\{ \Phi = (\phi_i)_{1 \leq i \leq N_e} : \phi_i \in \mathcal{V}, \quad \int_{\mathbb{R}^3} \phi_i \phi_j = \delta_{ij}, \quad 1 \leq i, j \leq N_e \right\}.$$

Problem of Eq. (1.11) is solved using a variational method. Let $\{\chi_\mu\}_{1 \leq \mu \leq N_b}$ be a finite basis set of \mathcal{V} of size N_b . To every MO configuration $\Phi = (\phi_i) \in \mathcal{W}_{N_e}(\mathcal{V})$ we associate the matrix $\mathbf{C} \in \mathbb{R}^{N_b \times N_e}$ gathering the coefficients of ϕ_i in the basis $\{\chi_\mu\}$:

$$\forall 1 \leq i \leq N_e, \quad \phi_i = \sum_{\mu=1}^{N_b} C_{\mu i} \chi_\mu \in \mathcal{V}. \quad (1.12)$$

Let us write down the Hartree-Fock problem of Eq. (1.11) in matrix form. First, the *density matrix*, denoted by \mathbf{D} , is defined as

$$\mathbf{D} := \mathbf{C} \mathbf{C}^\top. \quad (1.13)$$

The *Fock matrix*, denoted by \mathbf{F} , is defined as

$$\mathbf{F}(\mathbf{D}) := \mathbf{T} + \mathbf{V} + \frac{1}{2}\mathbf{J}(\mathbf{D}) - \frac{1}{2}\mathbf{K}(\mathbf{D}), \quad (1.14)$$

where \mathbf{T} is the kinetic energy matrix, \mathbf{V} is the external potential, $\mathbf{J}(\mathbf{D})$ is the Coulomb matrix and $\mathbf{K}(\mathbf{D})$ is the exchange matrix. The explicit definition of such matrices can be obtained by writing the Hartree-Fock energy in the finite basis of molecular orbitals (more details can be found in [60, §6.2.3]). In Section 1.1.4.4 we provide the explicit definition of the Coulomb matrix in particular. The total Hartree-Fock energy then writes in matrix form as

$$\begin{aligned} E^{\text{HF}}(\mathbf{D}) &= \text{Tr}(\mathbf{F}(\mathbf{D})\mathbf{D}) \\ &= \text{Tr}(\mathbf{T}\mathbf{D}) + \text{Tr}(\mathbf{V}\mathbf{D}) + \frac{1}{2}\text{Tr}(\mathbf{J}(\mathbf{D})\mathbf{D}) - \frac{1}{2}\text{Tr}(\mathbf{K}(\mathbf{D})\mathbf{D}). \end{aligned} \quad (1.15)$$

The computation cost required to evaluate this energy expression is discussed in later sections.

1.1.3.2 Roothaan equations

The coefficients \mathbf{C} in the approximation of Eq. (1.12) can be computed as follows. The Euler-Lagrange method for the minimization problem of Eq. (1.11), for the energy expressed in matrix form as Eq. (1.15), lead to a generalized eigenvalue problem, reading

$$\mathbf{F}(\mathbf{D})\mathbf{C} = \mathbf{S}\mathbf{C}\mathbf{E},$$

where \mathbf{S} is the overlap matrix of the basis set, with entries

$$\forall 1 \leq \mu, \nu \leq N_b, \quad S_{\mu\nu} = \int_{\mathbb{R}^3} \chi_\mu(\mathbf{r})\chi_\nu(\mathbf{r}) \, d\mathbf{r}, \quad (1.16)$$

and \mathbf{E} is a diagonal matrix of eigenenergies. Here, \mathbf{E} plays the role of the Lagrange multiplier associated to the orthonormality condition of the molecular orbitals, writing as the constraint

$$\mathbf{C}^\top \mathbf{S} \mathbf{C} = \mathbf{I},$$

with \mathbf{I} the N_e -by- N_e identity matrix. Notice that the basis functions χ_μ do not have to be orthonormal. One method for finding the coefficients \mathbf{C} comes down to solving the so-called *Roothaan equations*, reading

$$\begin{cases} \mathbf{F}(\mathbf{D})\mathbf{C} = \mathbf{S}\mathbf{C}\mathbf{E} \\ \mathbf{C}^\top \mathbf{S} \mathbf{C} = \mathbf{I} \\ \mathbf{D} = \mathbf{C}\mathbf{C}^\top. \end{cases} \quad (1.17)$$

An overview of existing methods for finding the coefficients \mathbf{C} can be found in [61].

1.1.4 Linear combination of atomic orbitals

In the previous section, we used the Hartree-Fock problem as a toy example for introducing a discretization of MOs of the form of Eq. (1.12). This type of discretization appears in more accurate models as well, including the Kohn-Sham model [21, 200]. In all such models, the most commonly used method for approximating MOs is the Linear Combination of Atomic Orbitals (LCAO), according to which MOs are expanded on basis sets of Atomic Orbitals (AOs). Even though AO bases have various advantages, such as that their form is motivated by chemistry, alternatives include fully numerical approaches, e.g. finite element shape functions [248]. We refer to [197] for a useful review on various choices for the basis sets, as well as their advantages and disadvantages.

Here, we define atomic orbital basis sets, focusing on Gaussian-type ones, and discuss their properties as well as their advantages concerning molecular integration.

1.1.4.1 Atomic orbital basis sets

The AO basis sets are defined as follows [62]. To every chemical element of the periodic table, we associate a finite collection of functions of $H^1(\mathbb{R}^3)$, that define the atomic orbitals associated with the particular chemical type. To construct the AO basis for a given molecule, we consider the union of all AOs associated to the chemical types of the atoms of the molecule. AOs are centered on atomic positions, yielding the general basis form:

$$\{\chi_\mu\} = \left\{ \xi_{1,1}(\mathbf{r} - \mathbf{R}_1), \dots, \xi_{1,n_1}(\mathbf{r} - \mathbf{R}_1); \dots; \xi_{N_a,1}(\mathbf{r} - \mathbf{R}_{N_a}), \dots, \xi_{N_a,n_{N_a}}(\mathbf{r} - \mathbf{R}_{N_a}) \right\},$$

where each $\xi_{i,j} \in H^1(\mathbb{R}^3)$ is centered at zero. As an example, the AO basis associated to the anion OH^- , which is used as a test case in Figure 1.2, is obtained as the union of the basis sets that are associated with the element O and the element H.

Due to physical prerequisites, AOs are exponentially decaying functions [185]. Existing explicit forms include Gaussian-type orbitals [53], Slater-type orbitals [347, 299] and numerical atomic orbitals [16]. In the present thesis, we use Gaussian-type orbitals due to their advantages in molecular integration, presented in Section 1.1.4.4.

1.1.4.2 Gaussian-type orbitals

We define Gaussian-type orbitals in two different forms, namely Cartesian and spherical ones, and then discuss the classification of orbitals regarding rotational symmetries.

Cartesian form. A Cartesian Gaussian-Type Orbital (GTO) [150] centered on the origin is defined as the function

$$\forall \mathbf{r} \in \mathbb{R}^3, \quad G_{lmn}(\mathbf{r}, \alpha) := x^l y^m z^n \exp(-\alpha |\mathbf{r}|^2), \quad (1.18)$$

where l, m, n are positive integers and α is a positive orbital exponent empirically defined. The value $\ell := l + m + n$ is the angular momentum quantum number (or *angular degree*, in short) of the orbital. Note that a Cartesian GTO essentially consists of a monomial times a radial part. It is often useful to choose each AO to be a *contracted* GTO, defined as a finite linear combination of GTOs called *primitives*:

$$\chi_\mu(\mathbf{r}) = \sum_{k=1}^{d_\mu} c_k G_{l_k m_k n_k}(\mathbf{r}, \alpha_k), \quad (1.19)$$

where d_μ is the contraction degree equal to the number of primitives, α_k is a fixed empirical parameter, $l_k + m_k + n_k$ is a fixed angular degree and c_k are the contraction coefficients. Contracted GTOs generally allow to reduce the number of degrees of freedom in the AO basis. Existing methods for defining GTO bases include empirical methods, collected in the Basis Set Exchange database [267], and the even-tempered ansatz [18].

Further notice that Cartesian GTOs are separable in the space variable $\mathbf{r} = (x, y, z) \in \mathbb{R}^3$, yielding the convenient tensorized expression

$$G_{lmn}(\mathbf{r}, \alpha) = G_l(x, \alpha) G_m(y, \alpha) G_n(z, \alpha), \quad (1.20)$$

where for example $G_l(x, \alpha) = x^l \exp(-\alpha x^2)$. This property is extensively used for computing integrals of GTOs (we refer to Appendix C for more details).

Spherical form. The angular-radial decomposition [113] of a Gaussian-type orbital in spherical form, for a given angular degree ℓ and orbital exponent α , reads as

$$\forall \mathbf{r} \in \mathbb{R}^3, \quad G(\mathbf{r}; \alpha) = Y_\ell^m(\theta, \phi) R_\ell(r), \quad (1.21)$$

where $Y_\ell^m : S^2 \rightarrow \mathbb{R}$ is a real spherical harmonic [122] of fixed degree ℓ and order m , with $-\ell \leq m \leq \ell$ and $S^2 := \{\mathbf{r} \in \mathbb{R}^3 : |\mathbf{r}| = 1\}$ being the unit 2-sphere, the two symbols θ and ϕ are




Figure 1.1: Radial part of Gaussian-type orbitals, normalized to one, for fixed orbital exponent $\alpha = 1$ and varying angular degrees. Normalization factor closed-form formulas can be found in [293]. Notice that higher angular degrees lead to more spatially diffuse functions.

used to denote angular spherical coordinates, $r := |\mathbf{r}|$ and R_ℓ is the radial part, given for any $r \in \mathbb{R}^+$ as $R_\ell(r) := B(\ell, \alpha)r^\ell \exp(-\alpha r^2)$ where $B(\ell, \alpha) \in \mathbb{R}$ is a scaling factor. Representations of the radial part are given in Figure 1.1.

Orbital types. The angular degree determines the orbital type. In particular, an orbital of angular degree $\ell = 0$ is called *s*-type for *sharp*, $\ell = 1$ is *p*-type for *principal*, $\ell = 2$ is *d*-type for *diffuse* and $\ell = 3$ is *f*-type for *fundamental*. Each class has a different type of symmetry: *s*-type has spherical symmetry, *p*-type has axial symmetry along an axis, etc. The number of Cartesian GTO components of angular degree ℓ is equal to $(\ell + 1)(\ell + 2)/2$. It is useful to mention that, for spherical GTOs, however, the number of components is smaller and equal to $2\ell + 1$. Increasing the size of the basis typically implies including GTOs of higher orbital type. Generally, increasing the angular degree ℓ also increases symmetries and captures angular information more accurately, as the set of real spherical harmonics is complete. In particular, it is well-known [122, §7.3] that for every function f belonging to the space $L^2(S^2)$ of square-integrable functions on the 2-sphere, there is a unique sequence of spherical harmonics so that

$$f = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} a_{m\ell} Y_\ell^m.$$

where $a_{m\ell} := \langle f, Y_\ell^m \rangle_{L^2(S^2)}$ and $\langle \cdot, \cdot \rangle_{L^2(S^2)}$ refers to the natural inner product in $L^2(S^2)$. Moreover, $\langle Y_\ell^m, Y_{\ell'}^{m'} \rangle_{L^2(S^2)} = \delta_{\ell\ell'} \delta_{mm'}$, for $\ell \geq 0, -\ell \leq m \leq \ell$ and $\ell' \geq 0, -\ell' \leq m' \leq \ell'$ [60, §6.1.2].

This is why GTOs come in inseparable blocks regrouping orbital components belonging to the same type. For example, all *d*-type components of a given radial part will be included in the basis as individual AOs. In most empirical basis sets, ℓ varies between zero and three, while in some cases it can reach $\ell = 4$ (*g*-type). Note that individual angular components do not need to be stored explicitly. They are typically recovered on-the-fly, for example during molecular integration, which is the topic of the upcoming paragraph.

1.1.4.3 Properties of atomic GTO basis sets

Polynomial times Gaussian basis functions generally offer better accuracy for fewer degrees of freedom compared to polynomial spaces, e.g. finite elements [55], making them a suitable




Figure 1.2: Condition number of the GTO basis overlap matrix for the anion OH^- , at varying atomic separations and contracted Gaussian-type empirical bases, in spherical (**sph**) and Cartesian (**cart**) form. Note that, in the water molecule, the bond O–H has length equal to 0.9584\AA . The overlap matrix has been calculated using PySCF [312].

discretization for larger molecules. Unfortunately, atomic orbital bases also come with properties that may cause numerical issues, as we describe in the present section.

First of all, atomic orbital functions are typically normalized, i.e. the overlap matrix \mathbf{S} defined in Eq. (1.16) is filled with ones on the diagonal. However, GTO basis sets over unbounded domains, used in quantum chemistry, are generally not orthogonal [200]:

$$\forall 1 \leq \mu \neq \nu \leq N_b, \quad S_{\mu\nu} \neq 0.$$

Another characteristic of GTO basis sets is redundancy. GTOs, defined in Eq. (1.18) are expressed in terms of a Gaussian $e^{-\alpha r^2}$, where α is a fixed empirical parameter that generally differs from one AO to another. Linear dependencies typically come from this Gaussian radial part of the basis function. The numerical linear dependence may be stronger or weaker depending on the coordinate system on which the basis set is expressed, namely the Cartesian or the spherical one. In the latter, AOs, defined in Eq. (1.21), are given in terms spherical harmonics, which have convenient orthogonality properties. As a result, the spherical form has less linear dependencies than the Cartesian one. Numerically, linear dependencies are manifested through overlap matrices that have poor conditioning, as presented in Figure 1.2 for the case of contracted bases. As expected, for uncontracted bases, the condition is even worse. Note that, even in the case of spherical coordinates, the radial part of functions centered on different atoms may generate linear dependencies in the basis. Results show that these dependencies become stronger when the overlap of the radial parts increases, which is when nuclei are close to each other.

The source of ill-conditioning of GTO overlap matrices is the so-called *overcompleteness* property of GTO discretizations. This property occurs when adding new basis functions that are numerically linearly dependent to the ones already belonging to the basis, i.e. GTOs centered on a nearby atom. Poorly conditioned basis sets may influence [312] the convergence of the generalized eigenvalue solver used to solve the Roothaan equations of Eq. (1.17). This is why, in practice, it is preferable to remove linear dependencies and treat overcompleteness by applying either an orthonormalization procedure designed for ill-conditioned matrices [200], or a pivoted Cholesky decomposition to the overlap matrix [196], before solving the Roothaan equations. At this point, it may be useful to illustrate that, generally adding more GTOs into the basis tends to better reproduce the nuclear cusp, the inner nuclear region, as well as the decay behaviour far from the nuclei, compared to the exact solution. Doing so, one can achieve improvements to the approximate solution defined over unbounded domains, as shown in Figure 1.3, using only few basis functions. Concurrently, however, the GTO basis acquires redundancies, which inevitably




Figure 1.3: Wave functions reproducing the nuclear cusp for a one-electron diatomic molecule with Dirac nuclear potential in 1D. Approximate wave functions are discretized on s -type GTO nested sets of varying size N . Nuclear positions are $x = \pm 1.1$. The orbital exponent takes values (in this order) $(0.5, 0.88, 3, 10, 20, 10^2, 10^3, 0.02, 0.6)$. The overlap matrix condition number explodes from 10^2 ($N = 1$) to 10^6 ($N = 9$).

arise from the linear dependence of Gaussians with similar orbital exponents, i.e. spatial widths. To summarize, fast GTO discretization error convergence comes at the price of overcompleteness.

1.1.4.4 Molecular integrals

The main advantage of GTOs concerning this thesis is integration. A necessary step in the resolution of the Roothaan equations of Eq. (1.17) is the Fock matrix assembly, that requires the evaluation of molecular integrals, i.e. high-dimensional integrals evaluated on atomic orbital basis functions. An example is the overlap integral introduced in Eq. (1.16). Let us introduce some notation for a specific integral of increased interest, that will be further discussed in the following sections. Let $\mu, \nu, \kappa, \lambda$ denote indices of elements of the AO basis χ , referred to as *centers* of the integral. The Electron Repulsion Integral (ERI), also known as *bielectronic* integral, is defined as the four-center two-electron integral, given by the formula

$$(\mu\nu|\kappa\lambda) := \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{\chi_\mu(\mathbf{r})\chi_\nu(\mathbf{r})\chi_\kappa(\mathbf{r}')\chi_\lambda(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r} d\mathbf{r}'. \quad (1.22)$$

The Coulomb-type singularity of the above integral can make conventional integration schemes especially difficult to apply. The main reason that Cartesian GTOs are easy to handle is the tensorization property of Eq. (1.20) and the Gaussian product rule (see Theorem C.1.1), that states that the product of two Cartesian GTOs can be expressed in terms of a single Gaussian function. Deriving integration schemes for this type of basis becomes possible using these two properties. Note that spherical GTOs can be converted to Cartesian ones using available transformations [293]. For this reason, GTO basis sets have become very popular in electronic structure calculations [119].

Molecular integrals appear in the Fock matrix assembly as follows. The entries of the Fock matrix of Eq. (1.14), as well as of the overlap matrix of Eq. (1.16), are given as molecular integrals. In particular, the Coulomb matrix, that will be an object of study in the following section due to its computationally expensive evaluation, is defined for any matrix $\mathbf{D} \in \mathbb{R}^{N_b \times N_b}$ as

$$\forall 1 \leq \mu, \nu \leq N_b, \quad [\mathbf{J}(\mathbf{D})]_{\mu\nu} := \sum_{\kappa, \lambda=1}^{N_b} (\mu\nu|\kappa\lambda) D_{\kappa\lambda}. \quad (1.23)$$

In the present thesis, we focus on the evaluation of an energy of interest related to the bielectronic integral, involving density functions discretized on atomic orbital bases. In the next section we introduce the notion of an electronic density.

1.1.5 The electronic density function

In quantum mechanics, the charge distribution of electrons in space is described by the electronic density function. We define this object, that will be abundantly used throughout the document.

1.1.5.1 Electronic density from first principles

Within the Born-Oppenheimer approximation, consider an electronic wave function denoted by $\Psi_e \in \mathcal{F}_{N_e}$, as introduced in Section 1.1.2. The *electronic density function* associated to Ψ_e refers to the real-valued function defined as

$$\forall \mathbf{r} \in \mathbb{R}^3, \quad \rho_{\Psi_e}(\mathbf{r}) := N_e \int_{\mathbb{R}^{3(N_e-1)}} |\Psi_e(\mathbf{r}, \mathbf{r}_2, \dots, \mathbf{r}_{N_e})|^2 d\mathbf{r}_2 \cdots d\mathbf{r}_{N_e}. \quad (1.24)$$

Contrary to other objects defined up to now in electronic structure (wave function, Hamiltonian operators), the electronic density ρ_{Ψ_e} has the property of being measurable experimentally, for example using X-ray diffraction techniques [60]. Notice that ρ_{Ψ_e} is non-negative and if $\|\Psi_e\|_{L^2(\mathbb{R}^{3N_e})} = 1$, then $\|\rho_{\Psi_e}\|_{L^1(\mathbb{R}^3)} = N_e$.

1.1.5.2 Electronic density from molecular orbitals

Given a Slater determinant $\Psi_e \in \mathcal{S}_{N_e}$ associated to a molecular orbital configuration $\Phi \in \mathcal{W}_{N_e}$, the electronic density function introduced in Eq. (1.24) admits the simple expression in terms of molecular orbitals

$$\forall \mathbf{r} \in \mathbb{R}^3, \quad \rho_{\Psi_e}(\mathbf{r}) = \rho_{\Phi}(\mathbf{r}) = \sum_{i=1}^{N_e} |\phi_i(\mathbf{r})|^2. \quad (1.25)$$

LCAO approximation. In LCAO discretization methods for solving the electronic structure problem of Eq. (1.11), MOs are expanded in a basis set $\{\chi_{\mu}\}_{1 \leq \mu \leq N_b}$ of AOs, yielding the following expression of the electronic density in terms of AOs:

$$\begin{aligned} \rho_{\Phi} &= \sum_{i=1}^{N_e} \left(\sum_{\mu=1}^{N_b} C_{\mu i} \chi_{\mu} \right) \left(\sum_{\nu=1}^{N_b} C_{\nu i} \chi_{\nu} \right) \\ &= \sum_{\mu, \nu=1}^{N_b} \left(\sum_{i=1}^{N_e} C_{\mu i} C_{\nu i} \right) \chi_{\mu} \chi_{\nu} \\ &= \sum_{\mu, \nu=1}^{N_b} D_{\mu\nu} \chi_{\mu} \chi_{\nu}, \end{aligned} \quad (1.26)$$

where \mathbf{D} is the density matrix defined previously in Eq. (1.13). In the following, we will use the notation $\rho \equiv \rho_{\Phi}$ when this does not lead to a confusion.

MO to AO transformation. The last important notion of this section is the so-called molecular orbital (MO) to atomic orbital (AO) transformation, which is a relation that expresses the bielectronic integral on MOs in terms of ERIs on AOs, introduced previously in Eq. (1.22), reading

$$\sum_{i=1}^{N_e} \sum_{j=1}^{N_e} \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{|\phi_i(\mathbf{r}) \phi_j(\mathbf{r}')|^2}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r} d\mathbf{r}' = \sum_{\mu, \nu, \kappa, \lambda=1}^{N_b} (\mu\nu | \kappa\lambda) D_{\mu\nu} D_{\kappa\lambda}. \quad (1.27)$$

The right-hand side is equal to the scalar quantity $\text{Tr}(\mathbf{J}(\mathbf{D})\mathbf{D})$ which is one of the energy terms in the Hartree-Fock energy in matrix form of Eq. (1.15). Note that we may additionally use the expansion of the electronic density in terms of MOs, given by Eq. (1.25), to obtain the bielectronic integral of a given electronic density in terms of ERI tensor entries, reading

$$\iint_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{\rho(\mathbf{r})\rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r} d\mathbf{r}' = \sum_{\mu,\nu,\kappa,\lambda=1}^{N_b} (\mu\nu|\kappa\lambda) D_{\mu\nu} D_{\kappa\lambda}. \quad (1.28)$$

A formula of similar type, involving densities of different molecules, will be the object of study in the following section.

1.1.6 The theory of intermolecular interactions

The intermolecular interaction energy is a quantity of interest in chemistry. The aim of the present section is to provide a presentation of the Gaussian Electrostatic Model (GEM), which is a method for computing intermolecular interactions with priority to computational efficiency. Our interest is to address the *frozen core* intermolecular interaction in particular, which is one of the interaction components. For the sake of generality, we first introduce the theoretical context for computing the frozen core energy, namely perturbation theory, in Section 1.1.6.2, before explicitly defining this quantity in Section 1.1.6.3 and, finally, addressing the techniques employed by GEM for accelerating its computation in Section 1.1.6.4. A very useful introduction to intermolecular forces can be found in [306].

1.1.6.1 The supermolecule approach

Let us begin by defining the intermolecular interaction. Suppose that we are given a molecular system, referred to as *supermolecule*, made up of two subsystems, referred to as *fragments*, denoted by A and B . It is further assumed that fragments are *closed-shell* systems, meaning that they have an even number of electrons. According to the supermolecule method [306, §5.6], the *intermolecular interaction* of a pair of molecules AB can be defined as the difference between the ground state energy of the supermolecule and that of the isolated molecules:

$$E_{\text{inter}} := E_{AB} - (E_A + E_B).$$

Direct evaluation of this scalar quantity can lead to numerical difficulties (we refer to [306, §5.6.2] for more details). The topic that concerns this work is alternative models for approximating this quantity. We will explain how electronic density functions are computed in these models. Note that, for obtaining the total energy, we still have to add the intramolecular energy part in E_A and E_B , obtained for example using the force field method (see Appendix B).

1.1.6.2 Perturbation theory

The present section contains preliminary theory that enables to define tools for computing intermolecular interactions.

A traditional method for computing intermolecular interactions is perturbation theory [277, 11]. The idea is to obtain parts of the intermolecular interaction by gradually perturbing the system of isolated fragments. These perturbations can be either explicitly defined in terms of perturbed operators or can be induced by relaxing the interacting molecular orbital groups, depending on the level of theory in question. To name a few, the Energy Decomposition Analysis (EDA) [237, 90] family of methods, including Kitaura-Morokuma (KM) EDA [187] and Constrained Space Orbital Variation (CSOV) [258, 22], is based on gradually mixing different groups of molecular orbitals, associated to core or valence electrons, during a self-consistent field procedure for minimizing the total energy. On the other hand, in Symmetry-Adapted Perturbation Theory (SAPT) [171], corrections are explicitly defined in terms of interaction operators and antisymmetrization operators. In the following, we focus on defining the operators that enable us to obtain a perturbation expansion.

Unperturbed operator. The first step in discussing the interaction between two systems is to set up wave functions in the case where there is no interaction [211]. If the molecules are far enough apart that the overlap between their wavefunctions can be ignored, the theory becomes relatively simple and dates back to the 1930's [210]. Let us identify a set of N_e^A electrons and N_a^A nuclei, as belonging to molecule A , and define an electronic Hamiltonian \mathcal{H}_A for molecule A in terms of these electrons and nuclei, of the form of Eq. (1.6). Similarly, the Hamiltonian \mathcal{H}_B for molecule B is defined in terms of the set of N_e^B electrons and N_a^B nuclei of B . The number of fragment electrons and nuclei sum to those of the supermolecule, that is $N_e = N_e^A + N_e^B$ and $N_a = N_a^A + N_a^B$. The unperturbed Hamiltonian for the combined system is then

$$\mathcal{H}_0 := \mathcal{H}_A + \mathcal{H}_B. \quad (1.29)$$

Due to the fact that electrons and nuclei of A do not see those of B in the definition of the unperturbed Hamiltonian, ordinary Rayleigh-Schrödinger perturbation theory [179] gives the ground state energy $E_0 \in \mathbb{R}$ of the total unperturbed system as the sum of the ground state energies of fragments A and B :

$$E_0 = E_A + E_B, \quad E_A := \langle \Psi_A, \mathcal{H}_A \Psi_A \rangle_{L^2(\mathbb{R}^{3N_e^A})}, \quad E_B := \langle \Psi_B, \mathcal{H}_B \Psi_B \rangle_{L^2(\mathbb{R}^{3N_e^B})}, \quad (1.30)$$

where $\Psi_A \in \mathcal{F}_{N_e^A}$ and $\Psi_B \in \mathcal{F}_{N_e^B}$ are the electronic wave functions of fragments, defined as solutions to two separate electronic structure problems, reading

$$\mathcal{H}_A \Psi_A = E_A \Psi_A, \quad \mathcal{H}_B \Psi_B = E_B \Psi_B. \quad (1.31)$$

As a consequence, one may obtain the electronic density function of the supermolecule, denoted by ρ , as the sum of the two fragment densities, using a so-called affine decomposition:

$$\rho \approx \rho_A + \rho_B, \quad (1.32)$$

where $\rho_A := \rho_{\Psi_A}$ and $\rho_B := \rho_{\Psi_B}$ are the electronic densities of the two isolated fragments, referred to as *frozen densities*. The equation Eq. (1.32) is referred to as *frozen density approximation* [23].

Note that the unperturbed operator \mathcal{H}_0 of Eq. (1.29) cannot describe strong chemical bond forms between open-shell atoms or covalent chemical bonds [133]. However, between closed-shell atoms or ions when strong covalent bonds do not happen, it provides a good enough approximation. Otherwise, we need to include corrective terms obtained by a perturbation expansion. The operator that then describes the interaction between system entities is described below.

Interaction operator. Interaction energy terms are obtained as higher-order corrections to the ground state given by Eq. (1.30) of the unperturbed system of Hamiltonian \mathcal{H}_0 . Recall that both A and B are closed-shell molecules in their ground state. The total Hamiltonian of the supermolecule AB is given as

$$\mathcal{H} = \mathcal{H}_A + \mathcal{H}_B + \mathcal{W},$$

where \mathcal{H}_A and \mathcal{H}_B are the Hamiltonians for the isolated molecules and \mathcal{W} is the interaction term. The interaction operator is explicitly defined as follows. The Rayleigh-Schrödinger perturbation applied to \mathcal{H}_0 is the electrostatic interaction between electrons and nuclei of fragment A and those of B :

$$\begin{aligned} \mathcal{W} := & \sum_{i=1}^{N_e^A} \sum_{j=1}^{N_e^B} \frac{1}{|\mathbf{r}_i^A - \mathbf{r}_j^B|} \\ & - \sum_{i=1}^{N_e^A} \sum_{j=1}^{N_a^B} \frac{z_i^B}{|\mathbf{r}_i^A - \mathbf{R}_j^B|} - \sum_{i=1}^{N_e^B} \sum_{j=1}^{N_a^A} \frac{z_i^A}{|\mathbf{r}_i^B - \mathbf{R}_j^A|} + \sum_{i=1}^{N_a^A} \sum_{j=1}^{N_a^B} \frac{z_i^A z_j^B}{|\mathbf{R}_i^A - \mathbf{R}_j^B|}, \end{aligned} \quad (1.33)$$

where \mathbf{r}_i^A, z_i^A and \mathbf{R}_i^A denote the electronic coordinates, the nuclear charges and the nuclear coordinates of A , respectively, similarly for B . Eq. (1.33) is made up of four terms in total. The first term is the Coulomb repulsion between electrons of A and those of B , the second term is the Coulomb attraction between electrons of A and nuclei of B , similarly for the third term. Lastly, the fourth term is a constant representing the nuclear repulsion between A and B .

Antisymmetrization operator. Notice that the unperturbed operator \mathcal{H}_0 of Eq. (1.29) does not include all electrons in a fully symmetric way, thus does not preserve the antisymmetry property of wave functions. The role of the antisymmetrization operator is to properly impose the correct permutational symmetry, namely the Pauli exclusion principle of Eq. (1.5), to the wave function of the supermolecule. The explicit definition of this operator is specific to the theory, e.g. we refer to [172, 171, 247] for SAPT.

1.1.6.3 The frozen core energy

The present paragraph focuses on the definition of the intermolecular interaction up to first order, known as *frozen core* energy, using the previously introduced operators of Section 1.1.6.2. The idea is to derive this energy by *freezing* the wave functions of the isolated interacting entities.

The concept of intermolecular interaction up to first order is described as follows. At the unperturbed stage, fragments are supposed to be far apart, no interaction takes place. The first-order correction is then obtained by asking what the energy could be if two infinitely separated molecules are brought together without relaxing their electronic wave functions, i.e. without mixing their electrons. This correction [159] corresponds to translating the electron densities of the subsystems along with their respective nuclei from infinitely far away to the supersystem geometry, resulting in a supersystem electron density, that is the sum of subsystem densities confined to the ground state of their fragments, as in Eq. (1.32). According to [343], the Frozen Core Energy (FCE) refers to the energy change associated with bringing the fragments from infinitely apart to their positions in the supermolecule, without relaxing the electron densities of fragments. At this point it is useful to mention that the computation of higher-order terms is obtained by separating the energy change from the frozen density to the fully relaxed density into *polarization*, *charge transfer* and *induction*, which will not further concern this thesis (we refer to [306, 171] for definition of these terms).

The frozen core energy is made up of the following two components. Due to the antisymmetry property of the supermolecule wave function, the frozen core energy is decomposed into a long-range contribution, where electrons are far apart, and a short-range one, where electronic clouds overlap. Each separation is treated using a different perturbation operator, namely the interaction operator of Eq. (1.33) treats the long-range, while the antisymmetrization operator treats the short-range. This leads to the definition of the frozen core energy (E_{FCE}) as the sum of the *electrostatic* (V_{AB}) and the *exchange-repulsion* (X_{AB}), also known as Pauli repulsion, terms:

$$E_{\text{inter}} = E_{\text{FCE}} + \dots, \quad E_{\text{FCE}} := V_{AB} + X_{AB}.$$

Before providing the explicit definition of the frozen core energy, we quickly summarize the physical description of each of the terms:

- **Coulomb electrostatic energy:** it accounts for Coulomb forces at long-range separations, where penetration between electronic charge distributions is negligible.
- **Exchange-repulsion energy:** it accounts for quantum mechanical corrections of the error of classical electrostatics theory in short-range separations, where the electron exchange between molecules becomes possible [306, 290].

Note that we restrict our presentation to two-fragment systems for simplicity. For an arbitrary number of fragments, the computation is performed on pairs, by summing pairwise fragment

interactions

$$E_{\text{FCE}} = \frac{1}{2} \sum_{A \neq B} V_{AB} + \frac{1}{2} \sum_{A \neq B} X_{AB}, \quad (1.34)$$

where we used the fact that the electrostatics and exchange-repulsion terms are symmetric up to permutation of A and B .

In the following, we explicitly define the frozen core energy accounting for the interaction \mathcal{W} of Eq. (1.33) up to first order, while respecting the antisymmetry property of wave functions.

Long-range separation: Electrostatics. A two-fold assumption is made for this first-order term, namely that the interaction is purely electrostatic, meaning that it results from the attraction between positive or negative charges as given by the operator \mathcal{W} of Eq. (1.33), and that the two subsystems are sufficiently widely separated for electron exchange to be negligible [211]. This implies that each molecule's wave function distorts under the influence of the other's electrostatic field, but with no mixing of the orbitals of one molecule into the orbitals of the other. For this reason, the supermolecular electronic wave function is a simple Hartree product defined as

$$\Psi^{(1)} := \Psi_A \Psi_B \in \mathcal{F}_{N_e}, \quad (1.35)$$

where Ψ_A is the unperturbed electronic wavefunction of electrons and nuclei of A , similarly for Ψ_B of B , given as solutions to the system of electronic structure problems of Eq. (1.31). Note that $\Psi^{(1)}$ does not satisfy the Pauli exclusion principle of Eq. (1.5). As a result, this correction appears classical because it discounts electronic wave function antisymmetry. We obtain the electrostatic interaction energy as a first-order long-range approximation to the interaction energy, defined as

$$V_{AB} := \langle \Psi^{(1)}, \mathcal{W} \Psi^{(1)} \rangle_{L^2(\mathbb{R}^{3N_e})}. \quad (1.36)$$

Let us rigorously derive a closed-form expression of this term in terms of frozen densities. We focus on the electron-electron term. Then the electron-nuclei terms are treated similarly. For any $1 \leq i \leq N_e^A$ and $1 \leq j \leq N_e^B$, the interaction between the i th and j th electrons is defined as

$$v_{ij} := \langle \Psi^{(1)}, \frac{1}{|\mathbf{r}_i^A - \mathbf{r}_j^B|} \Psi^{(1)} \rangle_{L^2(\mathbb{R}^{3N_e})}$$

For convenience, let us change the indices and set $j' := j + N_e^A$ the index of the electron of B in the electrons of the complex AB , with $N_e^A < j' \leq N_e$. Recall that $N_e = N_e^A + N_e^B$. For any given $1 \leq i \leq j' \leq N_e$ there holds

$$\begin{aligned} v_{ij'} &= \int_{\mathbb{R}^{3N_e}} |\Psi_A(\mathbf{r}_1, \dots, \mathbf{r}_i, \dots, \mathbf{r}_{N_e^A})|^2 \frac{1}{|\mathbf{r}_i - \mathbf{r}_{j'}|} |\Psi_B(\mathbf{r}_{N_e^A+1}, \dots, \mathbf{r}_{j'}, \dots, \mathbf{r}_{N_e})|^2 \mathbf{d}\mathbf{r}_1 \cdots \mathbf{d}\mathbf{r}_{N_e} \\ &= \int_{\mathbb{R}^{3N_e}} |\Psi_A(\mathbf{r}_i, \dots, \mathbf{r}_1, \dots, \mathbf{r}_{N_e^A})|^2 \frac{1}{|\mathbf{r}_i - \mathbf{r}_{j'}|} |\Psi_B(\mathbf{r}_{j'}, \dots, \mathbf{r}_{N_e^A+1}, \dots, \mathbf{r}_{N_e})|^2 \mathbf{d}\mathbf{r}_1 \cdots \mathbf{d}\mathbf{r}_{N_e} \\ &= \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \left[\int_{\mathbb{R}^{3(N_e^A-1)}} |\Psi_A(\mathbf{r}_i, \mathbf{r}_2, \dots, \mathbf{r}_1, \dots, \mathbf{r}_{N_e^A})|^2 \mathbf{d}\mathbf{r}_1 \cdots \mathbf{d}\mathbf{r}_{i-1} \mathbf{d}\mathbf{r}_{i+1} \cdots \mathbf{d}\mathbf{r}_{N_e^A} \right] \frac{1}{|\mathbf{r}_i - \mathbf{r}_{j'}|} \\ &\quad \times \left[\int_{\mathbb{R}^{3(N_e^B-1)}} |\Psi_B(\mathbf{r}_{j'}, \mathbf{r}_{N_e^A+2}, \dots, \mathbf{r}_{N_e^A+1}, \dots, \mathbf{r}_{N_e})|^2 \mathbf{d}\mathbf{r}_{N_e^A+1} \cdots \mathbf{d}\mathbf{r}_{j'-1} \mathbf{d}\mathbf{r}_{j'+1} \cdots \mathbf{d}\mathbf{r}_{N_e} \right] \mathbf{d}\mathbf{r}_i \mathbf{d}\mathbf{r}_{j'} \\ &= \frac{1}{N_e^A N_e^B} \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{\rho_A(\mathbf{r}_i) \rho_B(\mathbf{r}_{j'})}{|\mathbf{r}_i - \mathbf{r}_{j'}|} \mathbf{d}\mathbf{r}_i \mathbf{d}\mathbf{r}_{j'}. \end{aligned}$$

Note that the Pauli exclusion principle of Eq. (1.5) for the fragment wave functions has been evoked in the above calculations. Finally, the total electron-electron interaction between electrons

of A and those of B admits the expression

$$\begin{aligned} \sum_{i=1}^{N_e^A} \sum_{j=1}^{N_e^B} v_{ij} &= \frac{1}{N_e^A N_e^B} \sum_{i=1}^{N_e^A} \sum_{j=1}^{N_e^B} \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{\rho_A(\mathbf{r}_i) \rho_B(\mathbf{r}_j)}{|\mathbf{r}_i - \mathbf{r}_j|} \mathrm{d}\mathbf{r}_i \mathrm{d}\mathbf{r}_j \\ &= \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{\rho_A(\mathbf{r}) \rho_B(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} \mathrm{d}\mathbf{r} \mathrm{d}\mathbf{r}'. \end{aligned}$$

To summarize, the electronic electrostatic interaction is represented as a Coulomb interaction between two unperturbed monomer charge distributions [83]. The classical Coulomb electrostatic interaction energy V_{AB} between molecules A and B may be expressed as

$$\begin{aligned} V_{AB} &= \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{\rho_A(\mathbf{r}) \rho_B(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} \mathrm{d}\mathbf{r} \mathrm{d}\mathbf{r}' \\ &\quad - \sum_{i=1}^{N_a^A} z_i^A \int_{\mathbb{R}^3} \frac{\rho_B(\mathbf{r})}{|\mathbf{r} - \mathbf{R}_i^A|} \mathrm{d}\mathbf{r} - \sum_{i=1}^{N_a^B} z_i^B \int_{\mathbb{R}^3} \frac{\rho_A(\mathbf{r})}{|\mathbf{r} - \mathbf{R}_i^B|} \mathrm{d}\mathbf{r} + \sum_{i=1}^{N_a^A} \sum_{j=1}^{N_a^B} \frac{z_i^A z_j^B}{|\mathbf{R}_i^A - \mathbf{R}_j^B|}, \end{aligned} \quad (1.37)$$

Note that this is the exact classical interaction energy of the two molecular charge distributions in a form that does not depend on the multipole expansion, typically employed in applications, see for example [84]. For this reason, the exact Coulomb energy is well-defined for all intermolecular distances.

Short-range separation: Exchange-repulsion. This second term corrects the short-range, where electrons of the same spin cannot be at the same place. This is achieved by forming a properly antisymmetrized supermolecule wave function from the unrelaxed wave functions of each fragment. To get the exchange-repulsion, one writes the wave function as an antisymmetrized Hartree product $\mathcal{A}\Psi_A\Psi_B$ of the isolated monomer wave functions, where \mathcal{A} is the antisymmetrization operator. Applying this operator to the dimer wave function $\Psi^{(1)}$ given by Eq. (1.35) will force the correct permutational symmetry to $\Psi^{(1)}$. The exchange-repulsion energy is obtained as a first-order short-range approximation to the interaction energy when perturbing \mathcal{H}_0 of Eq. (1.33) by the antisymmetrization operator [145]. A closed-form expression for the exchange-repulsion can then be derived, that we do not explicitly provide here but we refer to [171, Eq. 112].

We are concerned by exchange-repulsion approximations. An empirical model for computing exchange-repulsion is the density overlap model of Wheatly and Price [338, 272]. According to this empirical model, the exchange-repulsion interaction energy between two molecules A and B , with frozen electronic density functions ρ_A and ρ_B , is defined as

$$X_{AB}^{\mathrm{WP}} := K_{AB} \Omega_{AB}, \quad \Omega_{AB} = \int_{\mathbb{R}^3} \rho_A(\mathbf{r}) \rho_B(\mathbf{r}) \mathrm{d}\mathbf{r}. \quad (1.38)$$

The constant K_{AB} is an empirical parameter and Ω_{AB} is the density overlap. This model is supported with experimental evidence [186, 165] and has been parametrized for various applications [103, 94, 246].

Frozen core energy from LCAO approximations. We use a numerical method for computing frozen densities, namely the LCAO approximation previously presented in Section 1.1.4. Thanks to the affine decomposition of Eq. (1.32), electronic structure calculations can now be performed separately on each isolated fragment. Let us denote by $\{\chi_\mu^A\}_\mu$ and $\{\chi_\mu^B\}_\mu$ the basis of Gaussian orbitals centered on atoms of A and atoms of B , respectively, and $\mathbf{D}_A, \mathbf{D}_B$ the corresponding density matrices. The density matrix \mathbf{D}_A is obtained by numerically solving the electronic ground state problem of Eq. (1.11) for the operator \mathcal{H}_A , discretized on the atomic Gaussian orbitals associated to A . Similarly for B . Note that we can use the Hartree-Fock model

or any other model, such as the Kohn-Sham model [21], to obtain the density matrices. Using the discretized frozen densities over atomic orbitals, of the form of Eq. (1.26), the computation of electron–electron term in the electrostatics model of Eq. (1.37) and the exchange–repulsion term in the density overlap model of Eq. (1.38) leads to the following integrals over Gaussian functions:

$$\iint_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{\rho_A(\mathbf{r})\rho_B(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} \mathrm{d}\mathbf{r} \mathrm{d}\mathbf{r}' = \sum_{\mu,\nu,\kappa,\lambda} D_{\mu\nu}^A D_{\kappa\lambda}^B \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{\chi_\mu^A(\mathbf{r})\chi_\nu^A(\mathbf{r})\chi_\kappa^B(\mathbf{r}')\chi_\lambda^B(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} \mathrm{d}\mathbf{r} \mathrm{d}\mathbf{r}' \quad (1.39)$$

$$\Omega_{AB} = \sum_{\mu,\nu,\kappa,\lambda} D_{\mu\nu}^A D_{\kappa\lambda}^B \int_{\mathbb{R}^3} \chi_\mu^A(\mathbf{r})\chi_\nu^A(\mathbf{r})\chi_\kappa^B(\mathbf{r})\chi_\lambda^B(\mathbf{r}) \mathrm{d}\mathbf{r}. \quad (1.40)$$

These two quantities yield the most computationally expensive integrals one needs to evaluate during a frozen core energy computation. The first quantity is a sum of two-electron four-center integrals, with the first electron belonging to A and the second to B . The second quantity is a sum of one-electron four-center integrals, again expressed on crossed orbitals centered at A and B . These four-center integrals lead to a quartic complexity with regard to the number of atomic orbital basis functions and will be one of the main topics of the present thesis. Note that the electron–nuclei term of the electrostatics admits an expansion similar to the one of Eq. (1.39), but involving one-electron two-center integrals, it is therefore less expensive. Lastly, we precise that all such quantities admit integral evaluation schemes over Gaussian orbitals, similar to the ones of the molecular integrals (we refer to Appendix C for more details).

1.1.6.4 The Gaussian electrostatic model

The Gaussian Electrostatic Model (GEM) [257, 83, 84, 256] belongs to the energy decomposition analysis family of methods for computing intermolecular interactions between molecules. Due to its computational efficiency, it finds applications to molecular dynamics [105]. One of the main ideas in GEM is to accelerate the computation of frozen core energy integrals, namely those of the form of Eq. (1.39) and Eq. (1.40), using a method known as Density Fitting (DF) [108]. Here we quickly mention how DF is used within GEM, postponing the description of the general DF method to Section 1.2.3 for the sake of clarity. We will provide a definition of the approximated frozen core energy part of GEM. The definition of remaining components of GEM can be found in [256] and will not be further discussed in this thesis.

GEM uses the reference analytical expression of the interaction operator \mathcal{W} of Eq. (1.33) for computing electrostatics. It employs the density overlap model of Eq. (1.38) for computing the exchange–repulsion, in favour of computational cost [302]. A differentiating feature of GEM, with respect to other methods, is that it employs a simplified representation of the frozen densities. In particular, the reference frozen densities ρ_A and ρ_B of Eq. (1.32) are obtained by quantum mechanics (QM) calculations and are referred to as *QM densities*. GEM introduces approximations $\tilde{\rho}_A$ and $\tilde{\rho}_B$ to fragment QM densities, referred to as *fitted densities*, obtained by DF, namely the variational Coulomb fitting method [106]. The general form of fitted densities can be considered as

$$\tilde{\rho}_A := \sum_{i=1}^{M_A} c_i^A \xi_i^A, \quad \tilde{\rho}_B := \sum_{i=1}^{M_B} c_i^B \xi_i^B,$$

where c_i^A and c_i^B are real expansion coefficients. The functions ξ_i^A and ξ_i^B are Hermite–Gaussian orbitals (defined in Appendix C.1.0.2). Note that GEM does not restrict to positive expansion coefficients [257, §A.1]. An additional constant charge constraint is enforced, according to which $\int_{\mathbb{R}^3} \tilde{\rho}_A = N_e^A$ and $\int_{\mathbb{R}^3} \tilde{\rho}_B = N_e^B$. The exact electrostatics V_{AB} of Eq. (1.37) is then approximated

as

$$V_{AB}^{\text{GEM}} := \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{\tilde{\rho}_A(\mathbf{r})\tilde{\rho}_B(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} \quad (1.41)$$

$$- \sum_{i=1}^{N_a^A} z_i^A \int_{\mathbb{R}^3} \frac{\tilde{\rho}_B(\mathbf{r})}{|\mathbf{r} - \mathbf{R}_i^A|} - \sum_{i=1}^{N_a^B} z_i^B \int_{\mathbb{R}^3} \frac{\tilde{\rho}_A(\mathbf{r})}{|\mathbf{r} - \mathbf{R}_i^B|} + \sum_{i=1}^{N_a^A} \sum_{j=1}^{N_a^B} \frac{z_i^A z_j^B}{|\mathbf{R}_i^A - \mathbf{R}_j^B|}.$$

The exchange-repulsion X_{AB}^{WP} of Eq. (1.38) is approximated as

$$X_{AB}^{\text{GEM}} := K_{AB} \tilde{\Omega}_{AB}, \quad \tilde{\Omega}_{AB} = \int_{\mathbb{R}^3} \tilde{\rho}_A(\mathbf{r})\tilde{\rho}_B(\mathbf{r}) \, \text{d}\mathbf{r}.$$

The parameter K_{AB} is obtained by linear regression using as reference the exchange-repulsion value obtained by the CSOV method [86]. Each GEM term tries to reproduce its SAPT (see [242]) or CSOV (see [84]) counterpart.

Using DF, the evaluation of four-center integrals over Gaussians is reduced to a two-center integral evaluation over Hermite-Gaussian orbitals:

$$\iint_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{\tilde{\rho}_A(\mathbf{r})\tilde{\rho}_B(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} \, \text{d}\mathbf{r} \, \text{d}\mathbf{r}' = \sum_{i,j} c_i^A c_j^B \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{\xi_i^A(\mathbf{r})\xi_j^B(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} \, \text{d}\mathbf{r} \, \text{d}\mathbf{r}' \quad (1.42)$$

$$\tilde{\Omega}_{AB} = \sum_{i,j} c_i^A c_j^B \int_{\mathbb{R}^3} \xi_i^A(\mathbf{r})\xi_j^B(\mathbf{r}) \, \text{d}\mathbf{r}. \quad (1.43)$$

The two-center expansions in the evaluation of the reduced integrals break the computational complexity of the initial expansions of Eq. (1.39) and of Eq. (1.40), namely from quartic to quadratic with regard to the number of atomic orbital basis functions. Let us further comment that Gaussian-type orbitals can be transformed to Hermite-Gaussians, using a mapping defined in Appendix C.1.0.3. Molecular integrals over orbitals of these two types admit the same integration schemes up to this transformation. GEM employs reciprocal space methods [84] for approximating the above Coulomb-type integral, as an additional approximation method on top of DF, to gain further computational efficiency. As a result, there is an accumulation of DF errors and integral evaluation errors in GEM.

1.2 Large-scale problems for molecular integrals

In the present section, we review algorithmic aspects that commonly create a bottleneck in molecular simulations, from a computational point of view, and come down to the problem of storing a large matrix entry-wise. This problem occurs in the molecular integration over atom-centered Gaussian basis with various applications, including electronic structure theory and intermolecular interaction theory (the frozen core energy computation in particular). Different areas of application often employ different tools to solve this problem. In this respect, we adopt a unified approach and first present the mathematically better-studied problem of the Coulomb matrix assembly in conventional electronic structure problems in Section 1.2.1. We then introduce our newly studied application context, namely the density-based frozen core energy computation, that also leads to this kind of problems, as presented in Section 1.2.2. We then highlight general acceleration methods for integrals over atom-centered Gaussian basis sets in Section 1.2.3, independently of the application context. The aim of our high-level approach is to enable the frozen core energy methods to take advantage of existing electronic structure methods.

1.2.1 The Coulomb matrix assembly

Ground state calculations in electronic structure theory (e.g. using Hartree-Fock or Kohn-Sham density functional theory) generally come down to the resolution of a problem whose matrix form is not very far from the one of Roothaan equations in Eq. (1.17). In practice the Roothaan equations are solved using a method known as Self-Consistent Field (SCF) [200]. This is an iterative procedure for approximating the density matrix \mathbf{D} by a sequence of approximations $(\mathbf{D}_n)_n$ until a self-consistent convergence criterion is reached. At each iteration n of SCF, the matrix terms $\mathbf{J}(\mathbf{D}_n)$ and $\mathbf{K}(\mathbf{D}_n)$ of the Fock matrix in Eq. (1.14) must be assembled. Unfortunately, this assembly is the bottleneck of the entire energy minimization process, explained as follows.

The computation of the Fock matrix in Eq. (1.14) is dominated by the non-linear terms with respect to the density matrix, namely the Coulomb and exchange matrices. The entries of the linear terms can be precomputed. On the contrary, at any given SCF iteration n , one needs to reassemble the Coulomb and exchange matrices, i.e. $\mathbf{J}(\mathbf{D}_{n-1})$ and $\mathbf{K}(\mathbf{D}_{n-1})$, given the already computed density matrix \mathbf{D}_{n-1} . Let us focus on the first term involving the Coulomb matrix and illustrate the main difficulty of assembling it, by counting the computational complexity cost. Each entry of the Coulomb matrix, defined in Eq. (1.23), is expressed in terms of a sum of ERIs on Gaussian-type orbitals. ERIs can be precomputed and stored entry-wise in a fourth-order tensor of dimension $N_b \times N_b \times N_b \times N_b$, with entries given as bielectronic integrals $(\mu\nu|\kappa\lambda)$, forming the ERI tensor. Then, every entry of the Coulomb matrix is evaluated by accessing the ERI tensor entry-wise. Due to this inevitable entry-wise data access, the assembly of the Coulomb matrix $\mathbf{J}(\mathbf{D})$ scales as

$$\mathcal{O}(N_b^4) \text{ flops.} \quad (1.44)$$

This is the bottleneck of the entire Hartree-Fock computation [51] and is one of the major bottlenecks of computational quantum chemistry [31]. Note that the prefactor of the scaling depends among other things on the implementation of the tensor entry-wise assembly. A naive four-order loop of full size N_b assembles the ERI tensor with a worst case loop count, analogous to the following nested sum:

$$\sum_{\mu=1}^{N_b} \sum_{\nu=1}^{N_b} \sum_{\kappa=1}^{N_b} \sum_{\lambda=1}^{N_b} (\mu\nu|\kappa\lambda).$$

A smarter assembly implementation utilizes the 8-fold permutational symmetry of the bielectronic integral with respect to the indices of basis functions, that is the property

$$(\mu\nu|\kappa\lambda) = (\mu\nu|\lambda\kappa) = (\nu\mu|\kappa\lambda) = (\nu\mu|\lambda\kappa) = (\kappa\lambda|\mu\nu) = (\kappa\lambda|\nu\mu) = (\lambda\kappa|\mu\nu) = (\lambda\kappa|\nu\mu).$$

Unfortunately, this optimization only decreases the loop count by a small factor, without drastically reducing the complexity of the problem. As a result, accelerating the ERI assembly is a non-trivial task that requires more sophisticated techniques, exploiting the ERI tensor structure (see Section 1.2.3 for discussion). Note that storing ERIs can require a significant amount of memory, for instance of the order of a gigabyte for a system of 100 electrons [348]. It is therefore important to avoid explicitly constructing and storing ERIs.

1.2.2 The frozen core energy

This relevant topic concerns the repeated entry-wise assembly of Coulomb-like matrices, arising in the context of intermolecular interaction theory. What follows is a high-level description of a density-based family of methods for computing the frozen core energy (FCE) part of intermolecular interactions, which include the Gaussian electrostatic model (see Section 1.1.6.4), the effective fragment potential [39], energy decomposition analysis [76] and symmetry-adapted perturbation theory [172, 232, 171].

Suppose that we are given a molecular system decomposed into subsystems, referred to as *fragments*, denoted by $\{A, B, C, \dots\}$. Our goal is to evaluate the total frozen core energy of interacting fragments, defined by Eq. (1.34). In this respect, we first solve the discrete ground

state problem of Eq. (1.11) separately for any given fragment A , to obtain its frozen electronic density, denoted by ρ_A . From the frozen core approximation of Eq. (1.32) follows that the total density ρ of the system can be approximated by an affine decomposition, i.e.

$$\rho \approx \sum_A \rho_A, \quad (1.45)$$

where the sum runs over the set of all fragments, of total number N_f . The computation of the total frozen core intermolecular interaction energy consists in summing all pairwise frozen density interactions. In this respect, a single pairwise interaction between any two given densities ρ_A and ρ_B , with regard to a given interaction *kernel* function $K : \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{R}^3$, describing the interaction between two points in space, is given by the *intermolecular integral*, defined as

$$E_{AB}^K := \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \rho_A(\mathbf{r}) K(\mathbf{r}, \mathbf{r}') \rho_B(\mathbf{r}') \, d\mathbf{r} \, d\mathbf{r}'. \quad (1.46)$$

We recover the electron–electron term of Eq. (1.39) for the Coulomb kernel $K(\mathbf{r}, \mathbf{r}') = |\mathbf{r} - \mathbf{r}'|^{-1}$ and the exchange–repulsion term of Eq. (1.40) for the Dirac kernel $K(\mathbf{r}, \mathbf{r}') = \delta(\mathbf{r} - \mathbf{r}')$, where δ is the Dirac delta function. A kernel explicit frozen core energy component is defined as

$$E_{\text{FCE}}^K := \frac{1}{2} \sum_{A \neq B} E_{AB}^K, \quad (1.47)$$

where the double sum runs over fragment pairs. There exist two computational difficulties related to this quantity. The first is the quadratic scaling with respect to fragment number. Large-scale fast summation methods (we refer to the following Section 1.2.3) can improve this scaling to linear or sublinear, by exploiting range-separation of interactions and accelerating the long-range ones. The second issue, primarily concerning this thesis, is the cost of computing a single pairwise interaction E_{AB}^K . Recalling Eq. (1.46), this term is given as a high-dimensional kernel explicit integral, that can be computed by following the principle of the transformation of Eq. (1.28) under the LCAO approximation, i.e. by expanding the molecular density into a given atomic orbital basis. In particular, let us denote by $\{\chi_\mu^A\}_{1 \leq \mu \leq N_b}$ the atomic orbital basis set centered on atoms of the fragment A and $\mathbf{D}_A = (D_{\mu\nu}^A)_{1 \leq \mu, \nu \leq N_b}$ the density matrix obtained by an *ab initio* calculation on the isolated fragment A . The electronic density being expanded in terms of the AO basis, as in Eq. (1.26), finally yields an orbital-wise evaluation of the interaction energy, reading

$$E_{AB}^K = \sum_{\kappa, \lambda, \mu, \nu} \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \chi_\mu^A \chi_\nu^A(\mathbf{r}) K(\mathbf{r}, \mathbf{r}') \chi_\kappa^B \chi_\lambda^B(\mathbf{r}') \, d\mathbf{r} \, d\mathbf{r}', \quad (1.48)$$

where the fourth-order sum runs over indices of atomic orbital basis functions, of total number N_b . Now, the summation of Eq. (1.48) needs to be evaluated, stored and accessed entry-wise. Unfortunately, the fourth-order sum on the approximation basis size N_b leads to a scaling $\mathcal{O}(N_b^4)$ for a single pairwise interaction, similar to the one reported in the Coulomb matrix assembly earlier in Section 1.2.1. Summing up, the total operation count for computing the total frozen core energy of Eq. (1.47) is

$$\mathcal{O}(N_b^4 N_f^2) \text{ flops.}$$

Typically in biomolecular systems, the fragment number varies between 10^2 and 10^5 , while the number of basis functions per fragment is about 10^2 . This leads to a computational effort that can be a limitation in various applications, including molecular dynamics simulations, where the quantity E_{FCE}^K needs to be evaluated at every timestep. The calculation of E_{FCE}^K is the bottleneck that concerns the present thesis. Efficient methods should focus on improving both the quartic and the quadratic scaling with respect to problem parameters, as discussed further in Section 1.2.3.

1.2.2.1 Comparison with the Coulomb matrix assembly

It is useful to identify two differences between the calculation of Eq. (1.47) and the one of the Coulomb matrix from Section 1.2.1. In order to establish a link between the two, let us denote the underlying interaction matrix, in the evaluation of a single pairwise frozen core intermolecular interaction between two given fragments A and B , by \mathbf{G}_{AB} . This interaction matrix is defined as the Gram matrix under the prescribed kernel K , with entries given as

$$\forall 1 \leq i, j \leq N_b^2, \quad G_{ij}^{AB} := \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \chi_{\mu_i}^A \chi_{\nu_i}^A(\mathbf{r}) K(\mathbf{r}, \mathbf{r}') \chi_{\kappa_j}^B \chi_{\lambda_j}^B(\mathbf{r}') \, \mathrm{d}\mathbf{r} \, \mathrm{d}\mathbf{r}', \quad (1.49)$$

essentially obtained by folding the first two and the last two dimensions of a four-index tensor between different fragments. The four-center integral of the right-hand side of Eq. (1.49) is the so-called intermolecular integral, mentioned earlier in Eq. (1.46). In matrix form, the main calculation of Eq. (1.48) writes as a bilinear form, i.e.

$$E_{AB}^K = \mathbf{d}_A^\top \mathbf{G}_{AB} \mathbf{d}_B.$$

Here \mathbf{d}_A is obtained by reshaping the density matrix \mathbf{D}_A of fragment A to a vector of size N_b^2 , similarly for B .

A first difference to notice is that, in the present setting, the interaction matrix \mathbf{G}_{AB} is used only for energy calculation, where the density matrix, consequently \mathbf{d}_A , is *a priori* available. On the contrary, in the electronic structure theory setting, the Coulomb interaction matrix is used both in the ground state energy calculation of Eq. (1.15), as well as in matrix-matrix products during SCF iterations, where essentially the density matrix is updated on-the-fly. Let us mention a second difference between the quantities of interest in the aforementioned settings, i.e. ground state energy minimization and frozen core energy computations. The first context targets the approximation of a matrix, while the second targets the approximation of a scalar quantity (the energy). In the first case, the matrix approximation accuracy is assessed using a matrix norm, so that the matrix product between $\mathbf{J}(\mathbf{D})$ and \mathbf{D} is well-approximated. On the contrary, the scalar value evaluation can be subject to error cancellation. In particular, this effect occurs when numerical errors of opposite signs cancel each other, during summation of different energy components, eventually lowering the approximation error on the total energy. This can be further exploited by approximation methods. For instance, the density fitting method (see Section 1.2.3.1) is known to produce significant cancellation errors in the electronic structure setting, reported to cause around two digits of accuracy improvements on ground state energy approximations [345, Table III]. The error cancellation could be expected to appear when summing different components to obtain the approximated Coulomb electrostatics of Eq. (1.41).

In the following section, we present a short overview on molecular integral evaluation algorithms applied to ERI matrices. The techniques proposed in this thesis, targeting arbitrary kernel-explicit intermolecular interaction matrices, make heavy use of ideas proposed in existing acceleration techniques for ERI matrix assembly.

1.2.3 Acceleration methods

Related acceleration techniques for solving large-scale molecular integral computation problems may be regrouped and summarized as follows. A helpful review of existing methods may also be found in [278, 71]. Without being exhaustive, we may list two classes of acceleration methods, in order of execution during the course of a simulation:

- **Stage I - Simplifying the electronic density function:** density fitting [106, 335, 340], real-space partitioning [203, 34, 19], integral screening [148, 264, 316].
- **Stage II - Evaluating the interaction:** direct computation [228, 110], cut-off method [316, 317], reciprocal space method [84, 221, 96], continuous fast multipole method [339, 70, 69, 253, 309], hierarchical matrix representation [344, 345], pivoted Cholesky decomposition

[31, 8, 146], grid-based compression [183, 185], Chebyshev interpolation [20], particle mesh Ewald [93, 288] and variations [79].

The main difference between the two stages is the acquired input data. In particular, the first one requires knowledge on quantities associated to isolated parts of the system, such as fragment or atom densities, i.e. *a priori* known quantities before the start of an electronic structure or molecular dynamics simulation. Note that no knowledge on molecular pairs is required at this level. Quantities approximated at this stage can be, for example, the isolated fragment electronic density function or fragment geometry parameters. On the contrary, the second class of methods requires access to pairwise fragment interactions, available on-the-fly, and includes evaluations that must be launched at each given timestep. The main purpose then is to compress the interaction data at a given MD timestep or at a given SCF iteration. The output of this process is an approximated interaction matrix (or energy).

1.2.3.1 Existing work

The following is a succinct presentation of some of the existing methods belonging to the two aforementioned groups.

Direct computation. The most naive way for evaluating molecular integrals is direct computation. Note that the bielectronic integral cannot be evaluated analytically. Nevertheless, it admits closed-form expressions that reduce to one-dimensional integrals, whose numerical evaluation is straightforward [149]. One of the most common methods over Cartesian Gaussian functions is the McMurchie-Davidson recursions [150, 228]. Recurrence relations can be accelerated with either the use of *L-tree codes*, as it has been done for example in [173, 70], or high-performance implementations on Graphical Processing Units (GPUs) [13]. Alternative Rys quadrature algorithms are available [110] and implemented on GPUs [12]. When no other method is available, the direct computation still remains an option for numerically evaluating molecular integrals. Note that, the theoretical complexity of a direct computation remains $\mathcal{O}(N_b^4)$, where N_b is the size of the Gaussian basis.

Density fitting. Given an electronic density ρ expanded on an atomic orbital basis of size N_b , the Density Fitting (DF) method refers to the approximation of ρ by the function

$$\tilde{\rho} = \sum_{i=1}^M c_i \xi_i,$$

where $\mathbf{c} = (c_i)_{1 \leq i \leq M} \in \mathbb{R}^M$ is a vector of *auxiliary coefficients* and $\boldsymbol{\xi} = \{\xi_i\}_{1 \leq i \leq M}$ is a given *auxiliary basis* of prescribed size M . Typically, the target size verifies $M \ll N_b^2$. The auxiliary basis type is restricted most often to Gaussian orbitals and sometimes to Hermite-Gaussian ones [84]. For fixed $\boldsymbol{\xi}$, the vector \mathbf{c} is defined as solution to the least-squares problem

$$\mathbf{c} := \arg \min_{\mathbf{x} \in \mathbb{R}^M} \left\| \rho - \sum_{i=1}^M x_i \xi_i \right\|^2,$$

for a prescribed error norm $\|\cdot\|$ on a function space. The DF method, originally introduced for the Coulomb integral [340] and the exchange integral [21], can be used to approximate various molecular integrals. A DF method is built upon a target integral in mind. Specific molecular integrals lead to particular auxiliary basis choices and error norms. This makes the DF method a good candidate for arbitrary kernel-based intermolecular integrals of the form of Eq. (1.46). A commonly treated case is the Coulomb integral in Hartree-Fock calculations [335]. Then DF schemes lead to a fast approximated ERI matrix assembly that scales as $\mathcal{O}(N_b^2)$ or $\mathcal{O}(N_b^3)$ [108]. It also exists a family of linear scaling density fitting methods [289, 51], including the following

two strategies. One of them is to fit the density using an appropriate metric that achieves sparsity in the interactions [106, 177]. Alternatively, the density is partitioned into spatially localized parts, that are fitted separately [120, 117, 301], introducing the concept of local fitting [262], with applications to post-Hartree-Fock methods [295].

Real-space partitioning. The goal of real-space partition methods is to spatially redistribute a given electronic density function into a small number of *a priori* fixed sites, which may be atomic or non-atomic. This is achieved by introducing a spatial weight. In particular, given a molecular fragment of M fixed sites and electronic density ρ , the real-space family of methods constructs a *distributed density* $\tilde{\rho}$ from ρ , defined in terms of a spatial partition $\{w_i\}_{1 \leq i \leq M}$, with $w_i : \mathbb{R}^3 \rightarrow \mathbb{R}$ and $0 \leq w_i \leq 1$, for all $1 \leq i \leq M$, and $\sum_{i=1}^M w_i = 1$, as

$$\forall \mathbf{r} \in \mathbb{R}^3, \quad \tilde{\rho}(\mathbf{r}) = \sum_{i=1}^M w_i(\mathbf{r})\rho(\mathbf{r}).$$

Each weight function is associated to a site and its role is to spatially localize the density on a significant region of that site [19]. In practice, the weight function approximately sums to the identity, hence $\tilde{\rho}$ approximates ρ . There is a variety of methods for constructing weight functions, namely fuzzy functions proposed by Becke [30] using Voronoi cells centered on sites, or Hirshfeld-like partitioning schemes [157], including Iterative Stockholder Atoms (ISA) [203], which employ an iterative procedure for constructing site-densities by minimizing an information-theoretic entropy error functional [34]. An advantage of this method is that it provides physically meaningful charge densities, distributed and spatially localized on sites [235, 234]. By doing so, the expansion of the density scales as $\mathcal{O}(M)$, instead of the *ab initio* scaling $\mathcal{O}(N_b^2)$. However, in the present thesis, we may not use real-space partition schemes due to the following reasons. First, the error functional used in the minimization procedure is generally less flexible than the one used in DF methods, hence making the procedure less suitable for kernel-driven optimization. Moreover, the optimization procedure for allocating contributions to non-atomic sites often admits multiple local minima or gives rise to weight functions for which there is an increased difficulty in evaluating the molecular integrals [314]. Besides, there is no restriction on the type of basis functions because quadrature rules are typically used in such cases. Note that the atomic decomposition has applications to the Coulomb electrostatics evaluation using the multipole expansion [234].

Integral screening. The idea is to use easy-to-check numerical thresholding criteria for discarding insignificant molecular integral contributions. The screening may be based on *a priori* estimators for pairwise primitive interactions given by Cauchy-Schwartz bounds on the bielectronic integral [148, 264, 340, 344]. In its more simple form, this bound reads:

$$|(\mu\nu|\kappa\lambda)| \leq (\mu\nu|\mu\nu)^{1/2}(\kappa\lambda|\kappa\lambda)^{1/2}. \quad (1.50)$$

Orbital tuples are then skipped if the bound calculated from orbital pairs is below a prescribed threshold value. Note that the bound does not depend on the distance between molecules, but simply on the magnitude of integrals centered on a single molecule. Other criteria use simplifications that arise from Hermite-Gaussian orbitals [70, 84, 69], in order to prescreen small-valued interactions of primitive orbitals. Screening methods can also be utilized in high-performance optimizations of molecular integral evaluations, such as vectorization [161].

Cut-off methods. Another type of thresholding uses locality of Gaussians and spatial cut-offs, for exploiting separation and small overlap of distributions based on the distance between atomic centers [316, 317]. The idea is to deduce a decay behaviour of integral values, with respect to the distance between centers, from the multipole expansion, due to the asymptotic decay of the Coulomb kernel (i.e. $1/r$). This family of methods then simply consists of neglecting contributions using distance criteria.

Reciprocal space methods. The reciprocal space approach performs numerical integration in the Fourier space using the Fast Fourier Transform [221, 84], for treating long-range interactions. This method can be incorporated into Particle Mesh Ewald methods for periodic systems. Such combined techniques are used in GEM [84, 257].

Continuous fast multipole multipole method. The Continuous Fast Multipole Method (CFMM) separates the interaction contributions in the ERI matrix into direct and well-separated ones, using a locality criterion. Well-separated ones are computed relying on multipole expansions, which is an analytic tool for approximations of interaction kernels (see Appendix B for more details). The remaining interactions are evaluated directly through explicit integration of orbitals. Such separation techniques exploit the fast decay of the Coulomb potential. Range-separation approaches may be combined with other analytic methods, such as Particle Mesh Ewald (PME) for periodic systems [104, 20, 84, 4, 79]. Overall, such methods can achieve linear scaling [79, 236, 51, 239, 176] with respect to the number of atoms. This is due to the use of multipole expansions to represent interactions between well-separated distributions into compressed form. However, notice that the prefactor of this method will be analogous to the number of distributions that are not well-separated, i.e. direct interactions. It has been reported [344] that CFMM is not as efficient because in practice a large number of distributions overlap, thus the number of interactions that must be evaluated directly is large [309]. Apparently, a weakness is that, direct computations, that scale in a quartic manner, dominate the computational cost of CFMM.

Pivoted Cholesky decomposition. Pivoted Cholesky Decomposition (PCD) is a matrix factorization technique [130] useful in solving rank deficient systems of equations. A related fact is that the ERI matrix has numerical linear dependence among its columns [31]. PCD can be used to construct low-rank representations of the ERI matrix [146]. The Cholesky decomposition without pivoting of the ERI matrix was first studied in 1977 [31]. The pivoted version has become quite a standard tool since [8, 43]. An advantage of PCD is that it naturally relates to density fitting as well as the auxiliary basis generation procedure, as we discuss in Chapter 3.

Tensor methods. It can be numerically justified that the ERI matrix has numerical rank $\mathcal{O}(N_b)$ [345]. The main goal of algebraic tensor methods is to exploit the low-rank structure of the ERI tensor for calculating low-rank approximations. On a purely algebraic level, one may use tensor format compression techniques [183, 185, 184], or hierarchical low-rank representations [345]. An advantage of the tensor-based approach is that it offers systematic on-the-fly methods. In tensor-based schemes, the analytical integration of densities may be replaced by a tensor-structure numerical quadrature [182], which makes such methods mostly suitable for numerical densities evaluated on tensor-grids. It has been reported [135] that such algebraic techniques for compressing interactions between point charges can be more efficient (in terms of the rank of the approximation and the range of applicability) than analytic techniques like multipole expansions.

Hybrid techniques. A direct compression technique has been reported in [344], based on low-rank approximations in the form of a block density fitting, that allows to compress far more interactions than could be compressed using multipole expansions, resulting in far fewer interactions that must be computed directly. The interesting point of this method, concerning the present thesis, is that it successfully combines a form of density fitting and hierarchical matrix compression techniques. It does so by combining a range-separation criterion and a matrix compression, that takes the algebraic form of a density fitting technique. The overall procedure can be interpreted as selectively applying density fitting to interactions using the range-separation criterion. According to this approach, density fitting is applied to small blocks of the interaction matrix. As it has been pointed out by the authors [344], the forementioned point is the main difference compared with classical DF, which applies to the full ERI matrix. This block-DF is computed using rank-revealing QR. Such hierarchical representation of the ERI

matrix has been interpreted as a generalization of DF, that locally and hierarchically applies DF to certain pairs of subsets of basis function products.

1.2.3.2 Motivations of present work

The work presented in this manuscript focuses on density fitting and integral screening methods and the reason can be motivated as follows. Recall that existing techniques fall into two categories. On the one category, interaction evaluation acceleration methods have already been widely studied in the mathematical literature, as they are related to well-known problems solved by a variety of approaches. These are traditionally based on ideas of the fast multipole method for breaking the quadratic scaling with respect to the interacting particle number. On the other category, density function simplification methods target scaling improvements on the explicit formula of the electronic density with respect to the basis function size. To our knowledge, these methods are significantly less documented in the mathematical literature, with the exception of the recent mathematical analysis of the ISA method [34]. Nevertheless, the major advantage of density-based methods is that they can precede any existing interaction evaluation, substantially improving its prefactor.

Numerical results are presumably in favour of this direction. Focusing on the density fitting (DF) method, it has been numerically reported that large-scale molecular simulations can greatly benefit, in terms of computational cost as well as storage requirements, by combining DF with fast summation techniques. In particular, the Gaussian electrostatic model couples DF with PME fast summation [83]. Using DF+PME, a Coulomb energy calculation on a cluster of 10^3 water molecules requires a CPU time of 40 seconds [84, 105]. However, based on available results [344], we may estimate that CFMM could require more than 10^3 seconds in an analogous energy calculation, even for systems smaller than the previous one, i.e. with fewer atoms by a factor of 1.5. Based on these encouraging numerical results, we believe that density-based methods merit further investigation. In this respect, the present thesis aims at filling this gap in the mathematical literature, by conducting a numerical study of density fitting and integral screening techniques in particular and trying to improve or couple them. The main objective is to design algorithms for solving the kernel-explicit frozen core energy computation problem of Section 1.2.2.

1.3 Molecular systems

In the present section, we provide a presentation of molecular systems under study, generally divided into two families according to their conformation geometry features. Although we use water models as concrete examples for introducing notions, note that our framework is not limited to water molecules and generally holds for any molecule. Motivation and background from chemistry is presented with a focus on polar molecules and their role in biochemistry.

1.3.1 Explicit water models

A polar molecule is a chemical species in which the spatial distribution of electrons between the covalently bonded atoms is not uniform. Many biomolecular processes and reactions take place in solution. Solvated biosystems remain a challenge due to computational efforts required to model multi-million atom systems. The solvent typically consists of polar molecules, that significantly contribute to polarization of biomolecules. Polarization, induced either by water solution, or by ions, such as magnesium or potassium ions, has been found to be critical to model various systems in the condensed phase [36], in folded RNA molecular structure [269], as well as ligand binding to proteins for drug design [276, 342, 95]. Since 1933, an extensively studied solvent is liquid water [37], the solution in which all biomolecular processes take place. Besides, H_2O is polar, i.e. has non-uniform electron distribution, because oxygen has eight electrons, while each of the two hydrogens has one. Due to its importance in biomolecular simulations, water will be used as a test case for numerical simulations in this thesis.

We focus on explicit water solvation models, based on individual solvent molecule description. Explicit water models are essentially force fields and can be regrouped into two categories, based on the way they assign atomic charges:

- *Rigid* water models assign fixed point charges at atomic or non-atomic sites. The geometric configuration of each water molecule is kept rigid during simulation, equal to the experimental equilibrium geometry of H₂O at gas phase. Due to its rigidity, the fixed point charge model fails to accurately capture some part of polarization [153]. An example is the rigid water model TIP3P (transferable intermolecular potential with three points) [174].
- *Flexible* water models assign flexible charges, moving along with atoms. An example is the AMOEBA polarizable force field for water [281, 283]. The AMOEBA water model assigns polarizable higher order atomic charges of multiple moment type (see Appendix B.2.3) to each one of the three atoms of water. Most importantly, charge centers are flexible, i.e. they move along the current atomic configuration, according to local coordinate frames assigned to atoms [263]. The AMOEBA water has been successfully used to study solvation dynamics in proteins and other complex systems [298, 137, 208]. The AMOEBA parameters for water are fitted to reproduce liquid properties, most often at room temperature. Different parametrizations of water are being developed over time, starting with AMOEBA03 [281], its improvement AMOEBA14 [195] and the latest Q-AMOEBA [225], that reproduces nuclear quantum effects.

A challenge for explicit water models is that the particle number can increase fast, as water makes up a big part of the system. As a typical example, SARS-COV-2 virus has 10⁴ atoms alone, while, after inclusion of water as well as ions, which act as neutralizers, the total number of atoms grows to 10⁵ [167]. Note that alternatives, generally yielding a less accurate solvent description in favour of computational efficiency, include continuous models, based on treating the solvent as a continuous medium [188, 230].

1.3.2 Case study: water clusters

In the present thesis, we assess the numerical performance of our methods on water clusters, i.e. explicit water solvents. We consider clusters made up of rigid or flexible water fragments. We are particularly concerned by studying how the intramolecular geometry class, i.e. rigid or flexible, affects the performance of numerical methods. Our goal is to design methods that efficiently handle both geometry classes. In this respect, we first describe the two classes here.

Consider a cluster consisting of water fragments and suppose we perform a molecular dynamics (MD) simulation. The following two fragment conformation geometry classes are treated separately in our study:

- Clusters of *rigid* fragments: Intramolecular bond angles and bond lengths of each fragment stay *fixed* during all timesteps of MD simulation. This type of cluster is studied in Chapters 3 and 4. TIP3P water intramolecular geometry has been utilized in numerical experiments.
- Clusters of *flexible* fragments: Intramolecular bond angles and bond lengths of each fragment are *a priori* different at any timestep. This type of geometry is treated using local atomic frames throughout Chapter 5. AMOEBA03 parametrization has been used in numerical experiments.

Figure 1.4 illustrates the two geometry types, where the rigid cluster also consists of identical fragments. Note that we study clusters comprising of water molecules for simplicity, but this framework could be generalized to clusters comprised of many different molecules.

1.3.2.1 Acquired data

Throughout the present thesis, we used or generated the following data on water molecules. All water test cases fall into one of the two categories of Figure 1.4.

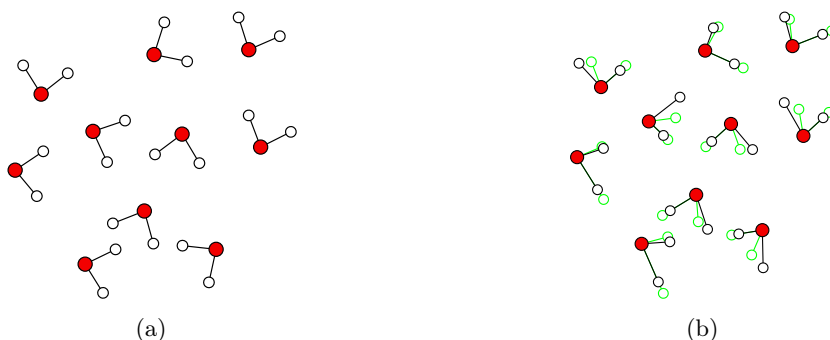


Figure 1.4: Simplified representation of water clusters: non-polarizable/rigid (left panel) and polarizable/flexible (right panel). The superposition of polarizable fragments at the current timestep and the previous one (in green) shows that bond lengths and angles change.

name	size	geometry data
dimerx	2	[273]
dimerx_rigid	2	[325]
watertiny	27	MD [274]
waterdot	365	
waterbox	500	
waterbig	1600	
waterglobe	6400	

Table 1.1: Summary of water clusters used as test cases. The size is equal to the total number of water fragments. For dimers, the symbol $x=01,02,\dots,10$ corresponds to the ten Smith dimers.

We use empirical atomic Gaussian-type basis parameter data optimized for water, with polarization functions on heavy atoms, such as 6-31gs [100] and Dunning-type basis sets [109], i.e. aug-cc-pvdz and aug-cc-pvtz, as the basis types used in the Gaussian electrostatic model [84].

In chemistry literature [279], it is common to use as test case the water dimers, for evaluating the accuracy of existing force fields and their ability to accurately reproduce the different components of the interaction energy. Available data has been used when possible. Concerning the rigid geometry configuration data, we use the ten standardized Smith water dimers, taken by the supplementary information of [325]. Data is given as a list of Cartesian coordinates in Angstrom, for the ten stationary point geometries of the water dimer as determined by the PES of Smith et al. [300]. In particular, the monomer geometries have been standardized at equilibrium geometries, namely HOH bond angle equal to 104.52 degrees and OH bond length equal to 0.957 Angstrom. TIP3P rigid water model uses exactly the same geometry parameters. For polarizable dimers, the ten Smith water dimers data have been used, found in the supporting information in [273]. For water clusters of flexible fragments, we use AMOEBA03 water model parametrization with bond length 0.9572 Angstrom and bond angle 108.50 degrees. Notice that the angle is slightly larger than the equilibrium one. All water clusters (rigid or flexible) with more than two molecules have been equilibrated using molecular dynamics, ran in Tinker [274]. The simulation setting used is 300 K temperature, 1 femtoseconds, sampling from the NVT (constant number of particles, velocity and temperature) thermodynamic ensemble [126]. We performed 100 timesteps for AMOEBA03 and 1000 timesteps for TIP3P. Table 1.1 summarizes the data used as molecular geometries, some of them taken from the Tinker test cases.

1.4 Computational chemistry software

The numerical results presented in Chapters 1, 3, 4 and 5 are obtained with Python programming language (version 3.10.12) [326]. Results of Chapter 6 are obtained with Julia programming language (version 1.8.3) [41]. The chemistry software used to perform numerical experiments throughout the present thesis consists of four open-source codes, for the following purposes:

- PySCF, Python-based Simulations of Chemistry Framework [312]: *ab initio* calculations, treatment of atomic Gaussian-type orbital basis sets, Wigner D-matrices [292].
- Libcint, general Gaussian-type orbital integrals for quantum chemistry [310]: molecular integral evaluation.
- Tinker molecular modeling package [274]: molecular dynamics simulations.
- vmd, visual molecular dynamics [164]: molecular visualization.

1.5 Summary

Among the families of methods for computing intermolecular interactions, energy decomposition analysis methods aim at describing the frozen core energy part of interactions using the frozen density approximation. The Gaussian electrostatic model is an energy decomposition method designed to use atom-centered Gaussian orbital basis sets for discretizing frozen electronic density functions of fragments. The computation of a single pairwise interaction among electronic densities discretized using N_b orbitals leads to a $\mathcal{O}(N_b^4)$ operation count for a naive implementation. This prohibitive numerical cost quickly becomes a bottleneck in many applications, including calculations of the frozen core energy part of the Gaussian electrostatic model on large biomolecular systems or water solvents. The aim of interaction acceleration methods is to break this complexity. There exist a lot of methods to reduce the number of interactions in large systems and we are particularly concerned by the ones applying to frozen fragment electronic densities. As opposed to the pairwise interaction compression methods, the studied ones are based on function approximations, namely density fitting. Systems under study are molecular clusters whose fragments fall into two categories, treated separately: rigid and flexible ones.

The next chapter aims at presenting the density fitting method, through the prism of solving fast summation problems for interacting densities.

Chapter 2

The pairwise interaction problem

Contents

2.1 Problem formulation	32
2.1.1 The pairwise interaction problem	32
2.1.2 From the Gaussian electrostatic model to interaction problems	32
2.2 Method overview	39
2.2.1 Existing work	39
2.2.2 Present work	40
2.2.3 Contributions	42
2.3 The density fitting method	44
2.3.1 Preliminaries	44
2.3.2 Definition of fitted density	46
2.3.3 Application to the pairwise interaction problem	47
2.4 Proofs	51
2.5 Conclusion	54

The computation of all pairwise interactions among N functions with regard to a specific inner product is a computationally intensive task leading to a $\mathcal{O}(N^2)$ operation count for a naive implementation. This prohibitive numerical cost quickly becomes a limitation in many applications such as in frozen core energy calculations. We are particularly concerned by breaking this complexity using a family of projection methods known as density fitting.

The present chapter is dedicated to the presentation of the density fitting method arising in the context of interacting densities. We first define in Section 2.1 the general summation problem under study and describe how the frozen core part of the Gaussian electrostatic model leads to this kind of problems. The Gaussian electrostatic model can benefit from improvements that may be generalized to other density-based problems (e.g. calculation of multipole moments [82], real-space partitioning methods [34], continuous fast summation [339]). However, the lack of a rigorous framework allowing to describe densities of rigid molecules makes these generalizations sometimes non-trivial. We want to handle all common aspects of density functions using a general theoretical framework, independently of the explicit function approximation method or the inner product used for interactions. The theoretical framework we describe uses a high-level approach and applies to various density-based methods.

A careful comparison between families of density fitting methods is presented in Section 2.2, that enables us to discuss the behavior of these methods according to the fragment molecular conformation. Then, in Section 2.3 we define the density fitting method with several advantages in our application context. We also present a set of general *a priori* error estimation results on existing density fit methods.

2.1 Problem formulation

Among the families of numerical methods for calculating frozen core energies in molecular dynamics applications, the Gaussian electrostatic model (see Section 1.1.6.4), designed to systematically approximate *ab initio* electronic density functions, has various applications to biomolecular simulations. Recall that this family of methods uses frozen fragment densities to represent the electronic charge distribution of large clusters. Hence, the cluster being fragmented into frozen fragments, its numerical treatment requires performing electronic structure calculations to smaller systems separately. However, as opposed to conventional electrostatics that employ point charges or multipoles, in density-based electrostatics a single pairwise intermolecular interaction is expensive to compute, as it is given as a high-dimensional molecular integral on an atomic orbital basis set (see Section 1.2.2). Fortunately, there exists a relation between interaction energy errors and best approximation errors on densities, as explained in Section 2.3. Before describing this link, we formulate in Section 2.1.1 the definition of a general pairwise molecular interaction summation problem and a set of kernels involved in common interactions. For the sake of clarity, we first announce the general summation problem and we specify details later, when treating aspects regarding electronic density functions and atomic orbitals in Section 2.1.2.2.

2.1.1 The pairwise interaction problem

Let \mathbb{H} be a real Hilbert space, which is by definition a complete normed vector space, defined over the real numbers, whose norm is induced by an inner product, denoted by $\langle \cdot, \cdot \rangle$. Let us fix $N \in \mathbb{N}^*$ the size of the problem we are looking at. Given two finite sets of N functions, denoted by $\mathbf{f} = \{f_i\}_{1 \leq i \leq N} \subseteq \mathbb{H}$ and $\mathbf{g} = \{g_i\}_{1 \leq i \leq N} \subseteq \mathbb{H}$, the *pairwise interaction problem* between \mathbf{f} and \mathbf{g} (with the inner product $\langle \cdot, \cdot \rangle$ describing the interaction between any two functions) refers to the problem of computing the scalar quantity $E \in \mathbb{R}$ defined by:

$$E := \sum_{i=1}^N \sum_{j=1}^N \langle f_i, g_j \rangle. \quad (2.1)$$

It is assumed that evaluating a single pairwise interaction between any two functions f_i and g_i has an operation count equal to a constant number $\kappa \in \mathbb{N}^*$, expressed in flops and defined as

$$\forall 1 \leq i, j \leq N, \quad \kappa := \text{cost}(\langle f_i, g_j \rangle). \quad (2.2)$$

Note that κ is not necessarily $\mathcal{O}(1)$, as it depends both on the expression of the inner product and the analytical form of functions in the sets \mathbf{f} and \mathbf{g} . For example, the inner product might be defined in terms of a high-dimensional integral and functions should have an analytic form that admits integration schemes. The cost κ then actually counts the cost of applying the integration scheme. Finally, the computation of E has a total operation count given by

$$\mathcal{O}(\kappa N^2) \text{ flops,}$$

for a naive implementation.

Particular inner products and analytical forms of families \mathbf{f} and \mathbf{g} , have features that allow a fast approximated evaluation of Eq. (2.1). We mainly restrict ourselves to the context of interest, which is the Gaussian electrostatic model.

2.1.2 From the Gaussian electrostatic model to interaction problems

The pairwise interaction problem arises during the calculation of quantities of interest in the Gaussian electrostatic model (GEM). In this model, one of the important quantities of interest is the frozen core energy part of the intermolecular interaction of a given set of fragments defined in Eq. (1.47), obtained by summing pairwise interacting frozen electronic density functions. The aim is to explicitly define the summation problem of Eq. (2.2) in the context of this model, starting from a given molecular system.

2.1.2.1 From molecules to density functions

Suppose that we want to perform a frozen core energy calculation on a given molecular system. First, GEM uses a specific type of function to describe the electronic structure of individual fragments and a set of operators to process them. We introduce these tools here.

Rigid and flexible fragments. The starting point are two important notions, that we will define following [91]: *isometry* and *congruence*. We focus on real coordinate spaces of dimension three, where atomic positions live. Let us consider the three-dimensional Euclidean metric space, defined as the vector space \mathbb{R}^3 equipped with the Euclidean distance between any two points, defining the metric

$$\forall(\mathbf{r}, \mathbf{r}') \in \mathbb{R}^3 \times \mathbb{R}^3, \quad d(\mathbf{r}, \mathbf{r}') := |\mathbf{r} - \mathbf{r}'|.$$

Definition 2.1.1 (Isometry in \mathbb{R}^3). A map $f : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ is called an *isometry* if for any two points $\mathbf{r}, \mathbf{r}' \in \mathbb{R}^3$, there holds

$$d(\mathbf{r}, \mathbf{r}') = d(f(\mathbf{r}), f(\mathbf{r}')).$$

Definition 2.1.2 (Congruence). Any two sets of points $X, Y \subseteq \mathbb{R}^3$ are called *congruent* if there exists an isometry $f : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ with $f(X) = Y$.

As an example, two triangles are congruent if one triangle can be repositioned by rotation, translation and reflection (but not resized) so as to coincide with the other triangle. According to the classification of isometries over \mathbb{R}^3 [66, 219], any isometry is either a rigid motion (translation or rotation), or a composition of a rigid motion and a reflection. Note that rotations are orientation-preserving transformations, while reflections are orientation-reversing ones, i.e. they can mirror an object.

Since the conformation of any molecular fragment is determined by the set of points corresponding to its atomic positions, these notions can be used for fragments. Now, in a molecular dynamics simulation, the atomic positions of any fragment are constantly updated. The set of atomic positions of a fragment at given time $t > 0$ is denoted by

$$X_a(t) := \{\mathbf{R}_1(t), \dots, \mathbf{R}_{N_a}(t)\} \subseteq \mathbb{R}^3,$$

where N_a is the atom number. A *rigid* fragment refers to a fragment for which, at any two timesteps $t_1, t_2 > 0$ of a molecular dynamics simulation, $X_a(t_1)$ and $X_a(t_2)$ are congruent point sets through translation or rotation. Notice that reflecting a fragment is not physically possible during the evolution of molecular dynamics, for this reason we restrict to orientation-preserving isometries. An example of a rigid fragment is provided in Figure 2.1. In this setting, a *flexible* fragment refers to a fragment that is not rigid.




Figure 2.1: Snapshots of a rigid H_2O fragment at three timesteps t_1, t_2, t_3 of a molecular dynamics simulation, expressed in a global reference frame. Identical markers indicate equality of segments or angles. Note that all three conformations are congruent through orientation-preserving isometries.

Reference frames. For any given fragment, rigid or flexible, we can construct a (non-unique) reference coordinate system, referred to as *local reference frame*, that is typically centered on the fragment. To distinguish between frames, the common coordinate system used to calculate forces in Eq. (1.2) is referred to as *global reference frame*. From now on, the notation A refers to a fragment expressed in the global frame and A' refers to the same fragment, expressed in the local one. Given a fragment, rigid or flexible, its local and global frame expressions are assumed to be congruent through either a translation and rotation, or a translation and symmetry. A feature of any given rigid fragment is that at any timestep, we can find a local frame in which the fragment coordinates, given by its atomic positions, stay fixed during the time of simulation.

Frozen fragment library. The electronic density function of a molecular system is calculated using the method of frozen densities defined in Eq. (1.45). According to this method, given any rigid fragment A , we first define an atomic Gaussian-type orbital basis set $\{\chi_\mu^{A'}\}_{1 \leq \mu \leq N_b}$, centered on atomic coordinates of the fragment A' expressed in its local reference frame. Then the discretized electronic ground state problem of Eq. (1.11) is solved on this basis set, for the system of interacting electrons and nuclei belonging to the fragment A only. The solution of this problem yields the density matrix $\mathbf{D}_{A'} = (D_{\mu\nu}^{A'})_{\mu,\nu} \in \mathbb{R}^{N_b \times N_b}$ defined in Eq. (1.13). This leads to a discretized electronic density function admitting the expansion

$$\forall \mathbf{r} \in \mathbb{R}^3, \quad \rho_{A'}(\mathbf{r}) = \sum_{\mu=1}^{N_b} \sum_{\nu=1}^{N_b} D_{\mu\nu}^{A'} \chi_\mu^{A'}(\mathbf{r}) \chi_\nu^{A'}(\mathbf{r}). \quad (2.3)$$

Notice that, for rigid fragments, this function is independent of the simulation time, as it is expressed on the local reference frame, where atomic positions of the rigid fragment stay fixed. The function $\tilde{\rho}_{A'}$ refers to an arbitrary approximation of $\rho_{A'}$, also expressed in the local reference frame of A . The *frozen fragment library* refers to a set of molecules expressed in their local reference frames, along with their electronic densities and their approximations.

Operators. Before describing the frozen core energy calculation in the Gaussian electrostatic model, we quickly provide the general terminology for the molecular coordinate transformations. The explicit definitions of these operators are given in Section 2.1.2.3.

1. **F2L (Fragment-to-Local)**: takes a molecular fragment expressed in a global reference frame and expresses it in its local reference frame.
2. **L2G (Local-to-Global)**: converts the frozen density expression of Eq. (2.3) from the local reference frame to the global reference frame.

Note that these two operators are of different nature. The first one is a simple change of coordinate frames, whereas the second one acts on functions, namely the densities.

During a molecular dynamics simulation, the **F2L** operator allows to visit the frozen fragment library where frozen densities are stored. Then approximated frozen densities are mapped back to the global reference frame, using the **L2G** operator, where they eventually interact. The operators **F2L/L2G** essentially allow to circulate information to/from the frozen fragment library and between reference frames. A schematic representation of the action of these operators is provided in Figure 2.2. Concerning computational complexity, the cost of computing the operator **F2L** scales as the number of atoms in the fragment. For the operator **L2G**, the cost depends on the maximal angular degree of the atomic orbital basis used in Eq. (2.3) [84].

2.1.2.2 The interaction problem in the Gaussian electrostatic model

We show how the computation of the frozen core intermolecular interaction between any two given rigid fragments, using the Gaussian electrostatic model, leads to a summation problem, belonging to the general family of summation problems of Eq. (2.1).

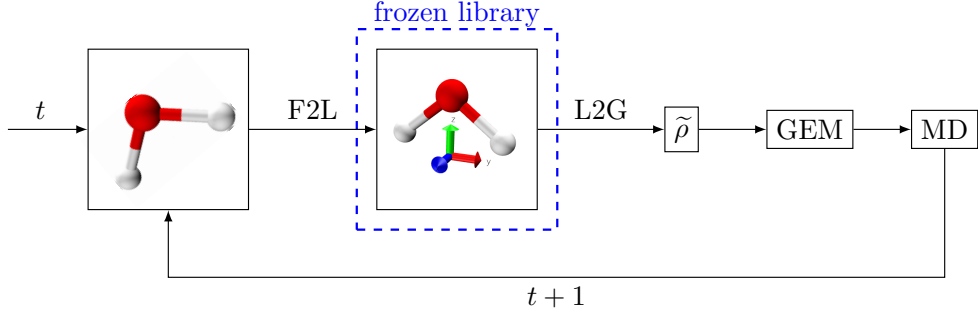


Figure 2.2: Execution of a molecular dynamics (MD) simulation using GEM. In order of application: for a given time t , recovery of fragment coordinates in global reference frame, application of **F2L** operator, application of **L2G** operator to obtain approximated density $\tilde{\rho}$, GEM interaction computation in global reference frame, launch MD, update timestep.

In the Gaussian electrostatic model, interacting function sets are defined as follows. For any given rigid fragment A expressed in a global reference frame, we can find a congruent fragment A' (i.e. the set of atomic positions of A is congruent to that of A') belonging to the frozen fragment library. It is additionally assumed that A and A' are linked through either a translation and rotation, or a translation and symmetry. We will denote by T_A the **L2G** operator converting A' to A and by $\rho_{A'}$ the frozen density of A' expressed in its local reference frame. Then the frozen density of A expressed in the global frame is defined as

$$\rho_A := T_A \rho_{A'}. \quad (2.4)$$

Let us define a finite set of functions, denoted by $\rho_A := \{\rho_i^A\}_{1 \leq i \leq N}$, with elements

$$\forall 1 \leq i \leq N, \quad \rho_i^A := T_A D_{\mu\nu}^{A'} \chi_\mu^{A'} \chi_\nu^{A'},$$

with $(\mu, \nu) := \pi(i)$ for π an enumeration of pairs $1 \leq \mu, \nu \leq N_b$ and $N := N_b^2$. Assuming for the moment that T_A is an affine mapping (this will be proved for specific cases in Section 2.1.2.3), there holds

$$\rho_A = \sum_{i=1}^N \rho_i^A. \quad (2.5)$$

Let us consider a cluster of rigid fragments expressed in a fixed global reference frame, denoted by $\{A, B, C, \dots\}$. The computation of the frozen core intermolecular interaction between any two fragments A and B , under an arbitrary kernel function K , as defined by Eq. (1.46), leads to the summation problem of pairwise interaction between sets ρ_A and ρ_B :

$$\begin{aligned} E_{AB}^K &= \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \rho_A(\mathbf{r}) K(\mathbf{r}, \mathbf{r}') \rho_B(\mathbf{r}') \, \mathrm{d}\mathbf{r} \, \mathrm{d}\mathbf{r}' \\ &= \sum_{i=1}^N \sum_{j=1}^N \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \rho_i^A(\mathbf{r}) K(\mathbf{r}, \mathbf{r}') \rho_j^B(\mathbf{r}') \, \mathrm{d}\mathbf{r} \, \mathrm{d}\mathbf{r}'. \end{aligned} \quad (2.6)$$

The integral takes two densities and computes the interaction between them. This computation is done directly (i.e. using direct kernel evaluations on elements of the sets ρ_A and ρ_B). Note that this step is an intermediary calculation, since the main calculation is the total frozen core energy, obtained by summing pairwise fragments as defined by Eq. (1.48). Also note that the cost of a single interaction defined in Eq. (2.2) is equal to the cost of computing a molecular integral over atomic Gaussian-type basis functions. It should be noted that this problem does not assume rigid fragments. It holds for any type of fragments.

It remains to introduce Hilbert spaces and their corresponding inner products. The definitions of these Hilbert spaces are specific to each interaction kernel. Explicit definitions for the Dirac and Coulomb kernels are given in Section 2.1.2.3.

2.1.2.3 Mathematical foundations of kernel explicit Hilbert spaces

An explicit form of inner products is introduced. Then we look for a characterization of the Hilbert spaces equipped with given kernel explicit inner products of interest, being the Dirac and Coulomb kernels in particular. We refer to [24, 57, 243] for more details on functional analysis of Hilbert spaces.

Let us begin by defining the following main notion.

Definition 2.1.3 (Kernel explicit Hilbert space). To any given *kernel* function $K : \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{R}$, we associate the so-called *kernel explicit Hilbert space*, denoted by \mathbb{H}_K , defined as the Hilbert space of functions over \mathbb{R}^3 equipped with the inner product associated to the bilinear form $\langle \cdot, \cdot \rangle_K : \mathbb{H}_K \times \mathbb{H}_K \rightarrow \mathbb{R}$ defined as

$$\forall (f, g) \in \mathbb{H}_K \times \mathbb{H}_K, \quad \langle f, g \rangle_K := \iint_{\mathbb{R}^3 \times \mathbb{R}^3} f(\mathbf{r}) K(\mathbf{r}, \mathbf{r}') g(\mathbf{r}') \, d\mathbf{r} \, d\mathbf{r}'. \quad (2.7)$$

It is thus assumed that the kernel K is such that the bilinear form is positive definite: $\langle f, f \rangle_K > 0$ for every $f \in \mathbb{H}_K, f \neq 0$.

The form $\langle \cdot, \cdot \rangle_K$ is referred to as *K-explicit* inner product over \mathbb{H}_K and the induced norm is denoted by $\| \cdot \|_K := \langle \cdot, \cdot \rangle_K^{1/2}$. Notice that the operator induced by the bilinear form of Eq. (2.7) is a generalization to Hilbert-Schmidt operators (see [60, Eq. B.8]) for which the kernel must belong to $L^2(\mathbb{R}^3 \times \mathbb{R}^3)$. Here we do not impose any regularity constraint on the kernel. This allows for instance to consider $K(\mathbf{r}, \mathbf{r}') = \delta(\mathbf{r} - \mathbf{r}')$, where δ denotes the Dirac delta distribution. This is a valid choice as it is easy to see that the kernel explicit Hilbert space associated to the Dirac kernel is $\mathbb{H}_K = L^2(\mathbb{R}^3)$.

Kernel examples. Throughout this work, we will focus on concrete kernel examples, that we define here explicitly. Typically in chemistry, one considers an interaction kernel induced by

$$K(\mathbf{r}, \mathbf{r}') = W(\mathbf{r} - \mathbf{r}'),$$

with W appropriate weight function. We focus on two common kernels:

- Coulomb kernel: induced by the Coulomb operator $W(\mathbf{r}) = |\mathbf{r}|^{-1}$. It gives the Coulomb energy [340] or generally the electron-electron interaction of Eq. (1.39).
- Dirac kernel: induced by the Dirac delta function $W(\mathbf{r}) = \delta(\mathbf{r})$. It gives the overlap energy [21], or generally the exchange-repulsion interaction of Eq. (1.38).

As opposed to smooth kernels, such as the Gaussian one induced, for fixed parameter σ , by $W(\mathbf{r}) = \exp(-\frac{1}{2}|\mathbf{r}|^2/\sigma^2)$, the Coulomb and Dirac kernels are singular around points satisfying $\mathbf{r} = \mathbf{r}'$, meaning that their expression tends to $+\infty$ around this singularity. Moreover, note that the Coulomb kernel is radial, as the Dirac one is not. Other choices of weights, not studied in the present work, are the attenuated Coulomb operator $W(\mathbf{r}) = \text{erfc}(\omega|\mathbf{r}|)/|\mathbf{r}|$ [177] and the Yukawa fitting $W(\mathbf{r}) = \exp(-\gamma|\mathbf{r}|)/|\mathbf{r}|$, for which by varying the shielding parameter γ , one can move smoothly from the Coulomb potential ($\gamma = 0$) to a tall spike potential ($\gamma \gg 1$) [127].

Hilbert space for the Coulomb kernel. Identifying the explicit Hilbert space for the Coulomb kernel, i.e. for $K(\mathbf{r}, \mathbf{r}') = |\mathbf{r} - \mathbf{r}'|^{-1}$, requires a bit more work. As it have been pointed out by previous works, the Coulomb operator defines an inner product [128]. Here we introduce the homogeneous Sobolev space of order one and show that the Coulomb inner product coincides with the norm of the dual of this space.

Let $d \in \mathbb{N}^*$ be the dimension of the problem we are looking at.

Definition 2.1.4. The Schwartz space $\mathcal{S}(\mathbb{R}^d)$ is the set of smooth functions u on \mathbb{R}^d such that for any $k \in \mathbb{N}$ we have

$$\|u\|_{k,\mathcal{S}} := \sup_{|\alpha| \leq k, x \in \mathbb{R}^d} (1 + |x|)^k |\partial^\alpha u(x)| < \infty.$$

We remind that the space of smooth compactly supported functions on \mathbb{R}^d is dense in $\mathcal{S}(\mathbb{R}^d)$. Note that the Schwartz space of rapidly decaying functions is a natural space for Gaussian-type orbitals, that decay exponentially at infinity.

Definition 2.1.5. A tempered distribution u on \mathbb{R}^d is any continuous linear functional on $\mathcal{S}(\mathbb{R}^d)$, such that there exist a constant C and an integer k for which there holds

$$|\langle u, \phi \rangle| \leq C \|\phi\|_{k,\mathcal{S}} \quad \forall \phi \in \mathcal{S}(\mathbb{R}^d).$$

Definition 2.1.6. Let s be in \mathbb{R} . The homogeneous Sobolev space $\dot{H}^s(\mathbb{R}^d)$ is the space of tempered distributions u over \mathbb{R}^d , the Fourier transform of which belongs to $L^1_{loc}(\mathbb{R}^d)$ and satisfies

$$\|u\|_{\dot{H}^s(\mathbb{R}^d)}^2 := \int_{\mathbb{R}^d} |\xi|^{2s} |\hat{u}(\xi)|^2 d\xi < \infty.$$

We may now prove the following known identity [32], relating the bielectronic integral and the norm of the space $\dot{H}^{-1}(\mathbb{R}^d)$ for $d = 3$.

Proposition 2.1.1.

$$\|u\|_{\dot{H}^{-1}(\mathbb{R}^3)}^2 = \pi \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{|u(\mathbf{r})u(\mathbf{r}')|}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r} d\mathbf{r}'.$$

Proof. See [proof](#) in Section 2.4. □

The following result shows that the function space equipped with the Coulomb kernel explicit inner product is a Hilbert space.

Proposition 2.1.2. $\dot{H}^s(\mathbb{R}^3)$ is a Hilbert space for $s = -1$.

Proof. See [proof](#) Section 2.4. □

Summary of kernel explicit Hilbert spaces. Recalling the Gaussian electrostatic energy terms from Section 1.1.6.4, we have established the following correspondance between intermolecular interaction energies of interest and kernel explicit Hilbert spaces.

GEM term	Kernel	Hilbert space
electron-electron	$K(\mathbf{r}, \mathbf{r}') = \frac{1}{ \mathbf{r} - \mathbf{r}' }$	$\dot{H}^{-1}(\mathbb{R}^3)$
exchange-repulsion	$K(\mathbf{r}, \mathbf{r}') = \delta(\mathbf{r} - \mathbf{r}')$	$L^2(\mathbb{R}^3)$

Operators. The explicit definition of operators used in the frozen core part of the Gaussian electrostatic model for rigid fragments is provided. First, the **F2L** operator is an isometry, acting as a change of basis between coordinate frames defined over \mathbb{R}^3 . Matrices of rotations or reflections are elements of the orthogonal group, defined for any dimension $n \in \mathbb{N}^*$ as the group of n -by- n orthogonal matrices:

$$O(n) := \{\mathbf{Q} \in \mathbb{R}^{n \times n} : \mathbf{Q}\mathbf{Q}^\top = \mathbf{Q}^\top\mathbf{Q} = \mathbf{I}\}.$$

This group, along with translations that are arbitrary vectors, can be used to characterize any isometry over \mathbb{R}^n . We focus here on the case $n = 3$.




Figure 2.3: Snapshots of a rigid fragment in a global reference frame at timesteps t_1, t_2 and t_3 . Given the set of points $X_i := X_a(t_i)$ for $i = 1, 2, 3$, the global frame is mapped to the local one using the **F2L** operator (an isometry), inducing the **L2G** operator T_i for $i = 1, 2, 3$ acting on Hilbert spaces, that is undoing the change of basis and mapping back to the global frame.

Definition 2.1.7. (3D **F2L**) Given any matrix $\mathbf{Q} \in O(3)$ and vector $\mathbf{t} \in \mathbb{R}^3$, the **F2L** operator is the mapping $\mathbf{F2L}[\mathbf{Q}, \mathbf{t}] : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ such that

$$\forall \mathbf{r} \in \mathbb{R}^3, \quad (\mathbf{F2L}[\mathbf{Q}, \mathbf{t}])(\mathbf{r}) := \mathbf{Q}(\mathbf{r} - \mathbf{t}).$$

Definition 2.1.8. (3D **L2G**) Let H be a real Hilbert space of functions defined over \mathbb{R}^3 . Given any matrix $\mathbf{Q} \in O(3)$ and vector $\mathbf{t} \in \mathbb{R}^3$, the **L2G** operator is the mapping $\mathbf{L2G}[\mathbf{Q}, \mathbf{t}] : H \rightarrow H$ such that

$$\forall f \in H, \quad (\mathbf{L2G}[\mathbf{Q}, \mathbf{t}] f)(\mathbf{r}) := f(\mathbf{Q}(\mathbf{r} - \mathbf{t})). \quad (2.8)$$

Note that a **F2L** operator maps vectors to vectors, while a **L2G** operator maps functions to functions. The operator **F2L**, acting as a change of basis between coordinate systems, and the operator **L2G**, undoing the action of **F2L** and mapping back to the global frame, are represented in Figure 2.3. We notice that a **L2G** operator is linear.

Isometries over Hilbert spaces. Definition 2.1.1 concerns isometries on metric spaces. Now we define isometries on Hilbert spaces.

Definition 2.1.9 (Isometry). Given a real Hilbert space H equipped with the inner product $\langle \cdot, \cdot \rangle$, an endomorphism $T : H \rightarrow H$ is said to be an *isometry* if it satisfies the property

$$\forall (f, g) \in H \times H, \quad \langle Tf, Tg \rangle = \langle f, g \rangle. \quad (2.9)$$

An isometry is a distance-preserving transformation, since a direct consequence of its definition is that there holds

$$\forall f \in H, \quad \|Tf\| = \|f\|, \quad (2.10)$$

where $\|\cdot\| := \langle \cdot, \cdot \rangle^{1/2}$. We now connect kernel explicit Hilbert spaces and isometries. To begin with, a kernel is said to be *isometry invariant*, i.e. the following two properties hold:

- $K(\mathbf{r} - \mathbf{t}, \mathbf{r}' - \mathbf{t}) = K(\mathbf{r}, \mathbf{r}')$ for every $\mathbf{t} \in \mathbb{R}^3$,
- $K(\mathbf{Q}\mathbf{r}, \mathbf{Q}\mathbf{r}') = K(\mathbf{r}, \mathbf{r}')$ for every $\mathbf{Q} \in O(3)$.

Notice that the first property is satisfied by Coulomb and Dirac kernels, and more generally by all kernels of the form $K(\mathbf{r}, \mathbf{r}') = W(\mathbf{r} - \mathbf{r}')$, for W a weight function. The second property is also satisfied by the Coulomb kernel, due to the invariance of the Euclidean distance under orthogonal matrices, and by the Dirac kernel, due to the homogeneity of the Dirac delta function.

The following result shows that the **L2G** operator is an isometry in kernel explicit Hilbert spaces for isometry invariant kernels.

Lemma 2.1.1. *Let K be an isometry invariant kernel. Consider the kernel explicit Hilbert space \mathbb{H}_K equipped with the K -explicit inner product $\langle \cdot, \cdot \rangle_K$. For any **L2G** operator over \mathbb{H}_K , denoted by T , there holds*

$$\forall (f, g) \in \mathbb{H}_K \times \mathbb{H}_K, \quad \langle Tf, Tg \rangle_K = \langle f, g \rangle_K.$$

Proof. See **proof** in Section 2.4. □

This simple but convenient property of the **L2G** operator is widely used in approximation methods for rigid fragments, as further detailed in Section 2.3.3.2.

2.2 Method overview

We are concerned by the density fitting (DF) [257] method for solving the summation problem of the Gaussian electrostatic model, defined in Eq. (2.6). The idea of the DF method is to approximate each of the interacting densities using a projection technique. There are various strategies for applying density fitting in the Gaussian electrostatic model, that we review here. Their main difference is the way they handle precomputation.

2.2.1 Existing work

We present here a set of conventional methods based on density fitting. One of their differences lies in their online interaction evaluation. According to the type of density used, i.e. the exact one or its DF approximation, we distinguish three types: *QM-QM* for interacting QM densities (exact interaction); *DF-DF* for interacting DF densities; *QM-DF* for hybrid interacting densities. Existing strategies for performing the intermolecular interaction approximation based on offline density fitting can be summarized as follows.

2.2.1.1 Direct application of density fitting

Rigid fragments. Density fitting with given empirical basis sets. Fit is performed by solving the normal equations of a least-squares problem, that is generally well-conditioned in practice. Robust DF [108] is an extension of density fitting including QM-DF corrective terms. This method will be explicitly defined later in Eq. (2.30). Notice that the form of QM-DF terms is more expensive to compute than conventional DF-DF. Moreover, the main limitation of RDF is that it requires to have access to QM densities during the online phase, which is prohibitive in terms of storage cost in the context of GEM.

Flexible fragments. The most naive way is to directly apply density fitting. Generally, the most naive and most accurate way for performing density fitting consists in computing the QM and DF densities online. This results in *ab initio* complexity at the level of fragment during the molecular dynamics simulation, which quickly becomes prohibitive when trying to deal with large systems. When no other method is available, the online QM and DF density computation remains an option that we consider in this work for the purpose of numerical validation only.

2.2.1.2 Gaussian electrostatic model

GEM divides the simulation time into two phases and employs density fitting to one of the two.

Splitting simulation time. Suppose the following simulation time division, employed by molecular dynamics applications, illustrated in Figure 2.4. Two distinct phases make up the simulation time, in order of execution:

1. *Offline* phase: electronic densities are evaluated on the local reference frame of each fragment belonging to the frozen fragment library. An unlimited number of computational resources is available.

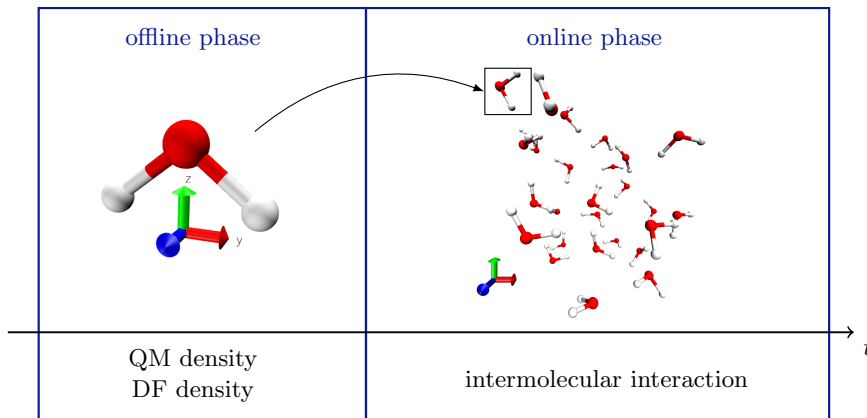


Figure 2.4: Simulation time of a frozen core GEM calculation on a water cluster. The 3D coordinate systems are the local reference frame (offline phase) and the global reference frame (online phase). The mapping corresponds to the **L2G** operator.

2. *Online* phase: Fragment atomic positions are updated at every timestep of molecular dynamics, in a global reference frame. The computation of the intermolecular interaction given by Eq. (1.46) is performed at each timestep, using electronic densities expressed on the global frame.

As an example, the frozen fragment library of a rigid water cluster, comprised of identical and congruent molecules through translation and rotation or symmetry, as the one of Figure 1.4a, contains a single element.

Rigid fragments. The Gaussian electrostatic model is essentially coupled to density fitting, in a method that we refer to as GEM-DF. GEM applies density fitting using a post-processed given empirical auxiliary basis set [84]. First, it discards orbitals of angular degree higher than two from a given empirical auxiliary basis set. Then, the basis is converted to *spd*-type, by adding orbitals, if needed, to form orbitals of one *s*, one *p*, and one *d*-type Gaussian with the same radial parts, centered on atoms. The main purpose of *spd* auxiliary basis is to reproduce the multipoles of the QM density (see Appendix B.2.4). Due to this post-processing, the resulting least-squares fit is often ill-conditioned. For this reason, normal equations are performed using truncated SVD, also using the additional electron charge constraint (see Section 2.3.1.3). Different choices of error metrics can be used to fit separate energy terms [242]. Interaction is evaluated in DF-DF form using DF scheme (2.29).

Flexible fragments. They are treated by averaging densities. GEM-DF extends the notion of local reference frames to flexible fragments and defines an approximate **L2G** operator. As an example, GEM-DF uses a frozen fragment library containing a single molecule to describe flexible water clusters, such as the one of Figure 1.4b. Then the operator **F2L** is obviously not an isometry.

2.2.2 Present work

The common framework of all DF-based methods applied to the Gaussian electrostatic model includes, first, the DF method itself and then the access to the frozen fragment library for precalculating a large part of the expensive operations.

The idea is to solve the summation problem of computing the quantity of interest E , defined in Eq. (2.1), by constructing two finite sets of functions, denoted by $\tilde{\mathbf{f}} = \{\tilde{f}_i\}_{1 \leq i \leq M} \subseteq \mathbf{H}$ and

$\tilde{\mathbf{g}} = \{\tilde{g}_i\}_{1 \leq i \leq M} \subseteq \mathbf{H}$. Then the quantity E can be approximately computed as

$$\tilde{E} := \sum_{i=1}^M \sum_{j=1}^M \langle \tilde{f}_i, \tilde{g}_j \rangle. \quad (2.11)$$

Assuming that the pairwise interaction between any two functions \tilde{f}_i and \tilde{g}_j has an operation count equal to a constant number $\tilde{\kappa} \in \mathbb{N}^*$, defined as

$$\forall 1 \leq i, j \leq M, \quad \tilde{\kappa} := \text{cost}(\langle \tilde{f}_i, \tilde{g}_j \rangle), \quad (2.12)$$

the computation of \tilde{E} has a total operation count given by

$$\mathcal{O}(\tilde{\kappa}M^2) \text{ flops.}$$

This complexity can be used to improved the one for computing the reference quantity E .

2.2.2.1 Density fitting

Recalling Section 1.2.3.1, given a density ρ , the density fitting method constructs the approximation

$$\rho \approx \tilde{\rho} = \sum_{i=1}^M c_i \xi_i. \quad (2.13)$$

Using notation from (2.11), density fitting allows to construct the family of functions $\tilde{\boldsymbol{\rho}} := \{\tilde{\rho}_i\}_{1 \leq i \leq M}$, with $\tilde{\rho}_i := c_i \xi_i$ for any $1 \leq i \leq M$. There are two main subproblems related to this approximation, handled in this order:

P.1 *Auxiliary basis generation*: given ρ , construct the set of functions $\boldsymbol{\xi} := \{\xi_i\}_{1 \leq i \leq M}$.

P.2 *Auxiliary coefficient computation*: given ρ and $\boldsymbol{\xi}$, compute coefficients $\mathbf{c} = (c_i)_{1 \leq i \leq M}$.

The criteria that $\boldsymbol{\xi}$ and \mathbf{c} should satisfy in order to provide a good approximation will be further studied in Section 2.3. Existing methods differ as to their actual criteria. Recalling the Gaussian electrostatics model terminology of Section 1.1.6.4, electronic density functions, obtained by quantum mechanics (QM), are referred to as *QM densities*, and fitted densities, obtained by density fitting (DF), are referred to as *DF densities*.

Focusing on the specific form of the summation problem in the context of the frozen core energy part of the Gaussian electrostatic model (GEM) (see Section 2.1.2.2), one of the simplest ways one can benefit from precomputation is by solving Problems **P.1** or **P.2** (or both) at the level of the frozen fragment library.

2.2.2.2 Integral screening

We add an additional improvement by solving the following problem.

P.3 *Integral screening*: for any two given functions f and g , find an explicit criterion for discarding the interaction $\langle f, g \rangle$.

2.2.2.3 Precomputation

We are concerned by methods that efficiently exploit the offline-online simulation framework. In this respect, in the context of the frozen core part of GEM, we introduce precomputation steps suitable to the fragment type.

Rigid fragments. The framework of this thesis for applying density fitting to rigid fragments can be described as follows. Intermolecular interactions between any two given rigid fragments can be approximated, for any two given isometries $T, T' : \mathbb{H} \rightarrow \mathbb{H}$, as

$$E(Tf, T'g) \approx E(T\tilde{f}, T'\tilde{g}), \quad (2.14)$$

with $E(f, g) := \sum_{i=1}^N \sum_{j=1}^N \langle f_i, g_j \rangle$. This will allow us to compute E for rigid fragments using offline schemes for precomputing \tilde{f} and \tilde{g} using density fitting.

Flexible fragments. The precomputation is much more complicated for flexible fragments. The issue is that for flexible fragments, however, the library depends on the definition of **F2L/L2G** operators. The definition of these operators is not straightforward and is the topic of Chapter 5. The approximation Eq. (2.14) is not well-defined in this context. We also considered the same problem, under flexible fragments. We distinguish two difficulties when treating flexible fragments. A first issue is that the explicit definition of operators **F2L/L2G** cannot be based on the notion of isometry, since bond angles and bond lengths are free to change. Second, in the rigid fragment case, the size of the frozen fragment library is finite, since for any given fragment at all timesteps, there exists exactly one congruent representative in the library. In the flexible case, however, if we add in the library all online fragments expressed in their local reference frames, we end up with an infinite library size. Overall, flexible fragments can be treated within the offline-online framework, either by relaxing the notion of **F2L/L2G** beyond isometries, or by allowing an infinite size of frozen fragment library. The first strategy captures the flexibility in bond angles and bond lengths, by allowing scaling of distance by a factor. However, there is no clear strategy that allows to define a computable closed-form expression of the operator **L2G** in that case. While the second naive strategy attempts to precompute all possible flexible fragment conformations, it quickly becomes prohibitive in practice.

2.2.3 Contributions

The work presented in this manuscript focuses on acceleration methods for pairwise intermolecular interaction computations based on electronic density functions discretized over Gaussian bases. The aim is to bridge the gap between density fitting methods in the frozen core part of GEM and numerical methods in mathematics. Hence, we explored the state-of-the-art methods and tried to improve them. There are various features to take into account in order to efficiently compute interactions: integral evaluation schemes, orbital symmetries, molecular conformation types, ability of precomputation etc. One of the main objectives of this thesis is to finely tune all aspects and end up with methods able to systematically handle any type of rigid or flexible molecular fragment geometry. To address such problems, the present thesis introduces various contributions. We give here a quick review of these, as well as of the challenges that have motivated our methods. Difficulties are specific to molecular conformation classes, treated separately.

Concerning rigid fragments, the goal of this work is to improve the direct application of density fitting, by constructing new auxiliary basis sets. The limitation of existing empirical auxiliary basis sets is that they are already reduced and do not allow further accuracy improvements. Our main contribution for rigid fragments consists in improving existing methods, by proposing, first, an auxiliary basis construction method with tunable accuracy and, second, an interaction sparsification method based on densities, where a large part of the interaction computation is omitted. Figure 2.5 summarizes the main differences regarding existing and new methods. Our methods for rigid fragments can be summarized as follows.

Auxiliary basis generation. We tackle the Problem P.1 for rigid fragments. Our method is pivoted Cholesky decomposition coupled to DF (PCD-DF). The idea is to apply classical density fitting with new auxiliary basis sets generated with our method, presented in Chapter 3.

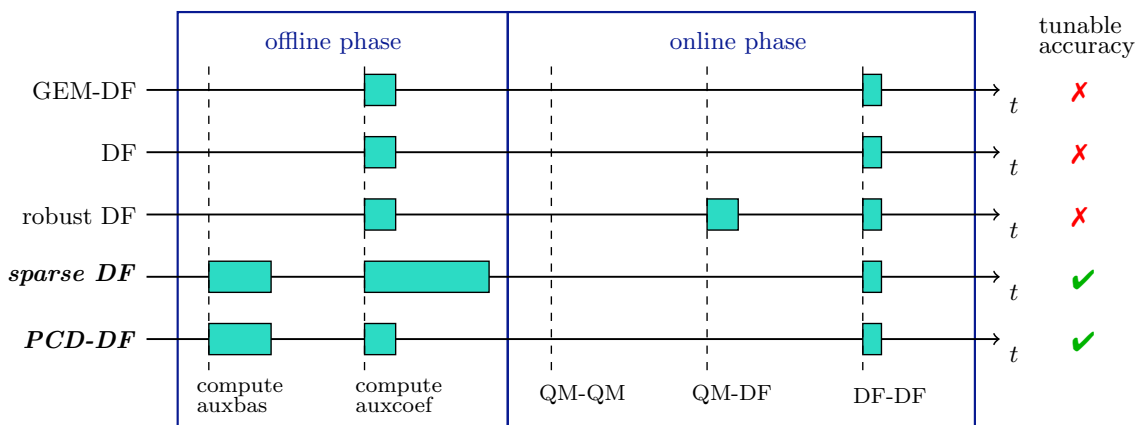


Figure 2.5: Simulation time and accuracy for rigid fragments (diagram inspired by [141]).

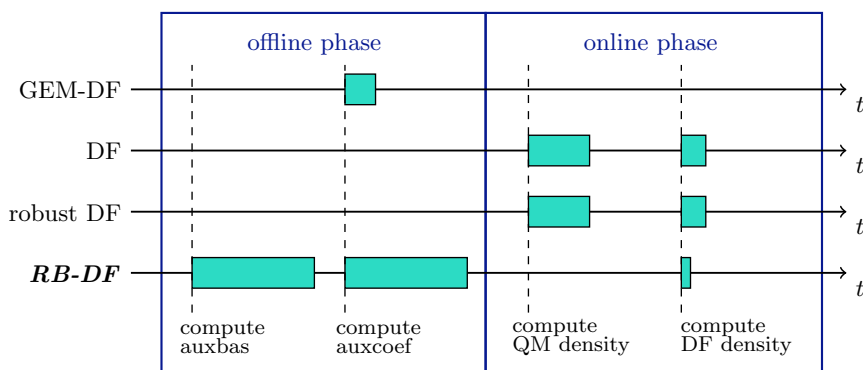


Figure 2.6: Simulation time and accuracy for flexible fragments (diagram inspired by [141]).

Reducing the number of pairwise interactions. We tackle the Problem **P.3** for rigid fragments. Our novel method for computing interaction energies exploiting sparsity of interactions, presented in Chapter 4, is referred to as Sparse Density Fitting (SDF).

Building frozen fragment libraries for flexible fragments. We tackle Problems **P.1** and **P.2** for flexible fragments. Our method performs a novel coupling between the Reduced Basis (RB) method [270] and the density fitting method, to define a new method known as RB-DF. In Chapter 5, we develop a strategy that allows to select a *finite* number of fragments for composing the frozen fragment library, and define local reference frames, for flexible fragments.

Figure 2.6 schematically summarizes the comparison between new and existing methods. Note that the main difference between existing methods is the phase chosen for performing DF. In practice, GEM beats the computational cost issue by introducing an error and placing the entire DF and QM computations in the offline phase. The advantage is the speed up, however the limitation is the systematically uncontrollable error due to flexible atomic positions. We consider that DF and RDF do not suffer from this error and compute DF and QM online at every timestep instead, that of course comes at a prohibitive cost price. To tackle this issue, we propose a new strategy based on reduced basis methods.

Our contribution is based on computing auxiliary coefficients online and precomputing auxiliary functions offline, as presented in Chapter 5. Our method performs offline precomputations that allow to construct suitable auxiliary basis sets offline and auxiliary coefficients at small cost online. Note that our method has a relatively expensive precomputation step with respect to other methods.

2.3 The density fitting method

In the present section, we first introduce the density fitting method, which is a projection technique. Finally, we focus on applications of density fitting for evaluating interaction energies and solving the general summation problem of Eq. (2.1).

2.3.1 Preliminaries

Before defining density fitting, we first quickly introduce the mathematical background and basic concepts on orthogonal projections in finite-dimensional vector spaces, best approximations and underlying least-squares problems. We adopt the bra-ket notation (see Section 2.3.1.1), offering a concise way of expressing orthogonal projections over finite basis sets in an abstract operator theory setting. The use of bra-kets for the purpose of density fitting in chemistry has been initiated in 1970 by Löwdin [212], see also [215, 213].

2.3.1.1 Bra-ket notation

The bra-ket notation consists of defining for any element f of H , the *ket*, denoted by $|f\rangle$, as being the element f itself,

$$|f\rangle := f \in H,$$

and the *bra*, denoted by $\langle f|$, as being the element of the dual space H^* , i.e. the linear functional which maps H to the real space of scalars, acting as

$$\langle f| : H \rightarrow \mathbb{R}, \quad \langle f| := \langle f, \cdot \rangle.$$

The existence and uniqueness of such linear functional $\langle f|$ is a consequence of the Riesz representation theorem [57]. As an example for vectors, if $H = \mathbb{R}^3$ then $|u\rangle$ is a column vector and $\langle u|$ is a row vector, while there holds $|u\rangle = \langle u|^\top$. To summarize, kets are elements and bras are linear forms.

For all $(f, g) \in H \times H$, the product of a bra and a ket is a real scalar, written as $\langle f|g\rangle$ and equal to the inner product $\langle f, g\rangle$, and the product of a ket and a bra is an operator mapping elements of H to H , written as $|f\rangle\langle g|$. Notice that $\langle f|f\rangle = \|f\|^2$. For any operator $T : H \rightarrow H$, the quantity

$$\langle f|T|f\rangle := \langle f|T(f)\rangle$$

essentially acts as a quadratic form.

Let us introduce the bra-ket notation for direct products of (possibly infinite-dimensional) vector spaces. Let $n \in \mathbb{N}$ and $H^n := \prod_{i=1}^n H$ denote the direct product of n copies of H . An element $\mathbf{f} \in H^n$ can be interpreted as a finite collection of n elements of H , i.e. $\mathbf{f} = (f_1, f_2, \dots, f_n)$ with $f_i \in H$ for all $1 \leq i \leq n$. Adopting quasi-matrix notation [154], the bra of \mathbf{f} is a functional acting on H^n , i.e. a column quasi-matrix, while the ket of \mathbf{f} is an element of the same space, i.e. a row quasi-matrix:

$$\langle \mathbf{f}| := [\langle f_1| \quad \langle f_2| \quad \dots \quad \langle f_n|]^\top, \quad |\mathbf{f}\rangle := [|f_1\rangle \quad |f_2\rangle \quad \dots \quad |f_n\rangle].$$

Note that each row of $\langle \mathbf{f}|$ is a function in H , and thus, whereas $\langle \mathbf{f}|$ is discrete as usual in the vertical direction, it is continuous in the horizontal direction [322]. Now, applying $|\mathbf{f}\rangle$ to $\langle \mathbf{f}|$ yields a real matrix. This allows to express Gram matrices in a compact way. Indeed, for any pair $(\mathbf{f}, \mathbf{g}) \in H^n \times H^n$, the *cross-Gram* matrix, with (i, j) -th entry equal to $\langle f_i, g_j\rangle$, is defined as

$$\langle \mathbf{f}|\mathbf{g}\rangle := \begin{bmatrix} \langle f_1| \\ \langle f_2| \\ \vdots \\ \langle f_n| \end{bmatrix} \begin{bmatrix} |g_1\rangle & |g_2\rangle & \dots & |g_n\rangle \end{bmatrix} = \begin{pmatrix} \langle f_1|g_1\rangle & \dots & \langle f_1|g_n\rangle \\ \langle f_2|g_1\rangle & \dots & \langle f_2|g_n\rangle \\ \vdots & \ddots & \vdots \\ \langle f_n|g_1\rangle & \dots & \langle f_n|g_n\rangle \end{pmatrix} \in \mathbb{R}^{n \times n}.$$

It is easy to see that the Gram matrix $\langle \mathbf{f}|\mathbf{f}\rangle$ is symmetric and positive definite, due to the positive definiteness of the inner product.

2.3.1.2 Best approximation

We refer to [97] for more details on best approximations. Let us introduce the orthogonal projection operator. Let X be a closed subspace of a real Hilbert space H . The orthogonal projection operator on X , denoted by P_X , is defined as the operator mapping H to X , satisfying the following property, for all $\rho \in H$,

$$\forall \sigma \in X, \quad \langle \sigma, \rho - P_X \rho \rangle = 0. \quad (2.15)$$

In particular, P_X maps any $\rho \in H$ to an element of X called the *best approximation* of ρ on X , defined as the solution to the least-squares problem

$$\tilde{\rho} := \arg \min_{\sigma \in X} \|\rho - \sigma\|^2. \quad (2.16)$$

Notice that the minimum is indeed attained due to X being closed. In other words, $\tilde{\rho} = P_X \rho$.

2.3.1.3 Resolution of least-squares problem

Let H be a real Hilbert space and $X \subseteq H$ a finite-dimensional subspace, of dimension M . Consider a basis $\boldsymbol{\xi} = \{\xi_i\}_{1 \leq i \leq M}$ of X , i.e. a linearly independent set of functions spanning X , not necessarily orthonormal. Finding the best approximation of any element of the Hilbert space on $\boldsymbol{\xi}$ breaks down to solving the least-squares problem of Eq. (2.16), as we describe in the following. For any $\rho \in H$, there exists a vector of coefficients $\mathbf{c} = (c_i)_{1 \leq i \leq M} \in \mathbb{R}^M$, satisfying

$$P_X \rho = \sum_{i=1}^M c_i \xi_i, \quad (2.17)$$

where P_X denotes the orthogonal projection operator on the subspace X . The vector \mathbf{c} is obtained by minimizing a least-squares error functional, defined as

$$\forall \mathbf{x} \in \mathbb{R}^M, \quad \mathcal{E}(\mathbf{x}) := \left\| \rho - \sum_{i=1}^M x_i \xi_i \right\|^2,$$

i.e. by solving the problem

$$\mathbf{c} := \arg \min_{\mathbf{x} \in \mathbb{R}^M} \mathcal{E}(\mathbf{x}). \quad (2.18)$$

Using the Euler method, the minimizer \mathbf{c} is obtained as solution to the equations

$$\forall 1 \leq i \leq M, \quad \frac{\partial \mathcal{E}}{\partial x_i}(\mathbf{c}) = 0. \quad (2.19)$$

The error functional admits the closed-form expression

$$\mathcal{E}(\mathbf{x}) = \langle \rho | \rho \rangle - 2 \sum_{i=1}^M x_i \langle \rho | \xi_i \rangle + \sum_{i=1}^M \sum_{j=1}^M x_i x_j \langle \xi_i | \xi_j \rangle.$$

Evaluating partial derivatives of this expression, one finds that the zeros \mathbf{c} of the gradient in Eq. (2.19) satisfy the linear system of equations

$$\forall 1 \leq i \leq M, \quad \sum_{j=1}^M \langle \xi_i | \xi_j \rangle c_j = \langle \xi_i | \rho \rangle,$$

or, equivalently, in matrix form, the *normal equations*, i.e.

$$\langle \boldsymbol{\xi} | \boldsymbol{\xi} \rangle \mathbf{c} = \langle \boldsymbol{\xi} | \rho \rangle, \quad (2.20)$$

where $\langle \boldsymbol{\xi} | \boldsymbol{\xi} \rangle$ is the Gram matrix of the family $\boldsymbol{\xi}$ and the right-hand side is given by the vector

$$\langle \boldsymbol{\xi} | \rho \rangle := (\langle \xi_1 | \rho \rangle \quad \cdots \quad \langle \xi_M | \rho \rangle)^T \in \mathbb{R}^M.$$

Note that the normal equations is indeed a well-posed system of equations, because, by the linear independence of $\boldsymbol{\xi}$, the Gram matrix is invertible. It may be relevant to mention at this point that, in density fitting methods where ρ is an electronic density function, additional constraints may be imposed using the Euler-Lagrange multipliers, such as the constant electron charge constraint, reading $\int_{\mathbb{R}^3} P_X \rho = N_e$. This constraint may increase the least-squares error but it will make the fit variational, since any electronic density function ρ satisfies this constraint (see Section 1.1.5.1).

Summing up, by solving the normal system of equations (2.20) to find \mathbf{c} by inverting the Gram matrix, then substituting to (2.17) and using that kets commute with scalars, i.e. writing

$$\forall \rho \in \mathbb{H}, \quad P_X |\rho\rangle = \mathbf{c} | \boldsymbol{\xi} \rangle = | \boldsymbol{\xi} \rangle \langle \boldsymbol{\xi} | \boldsymbol{\xi} \rangle^{-1} \langle \boldsymbol{\xi} | \rho \rangle,$$

one finds that the orthogonal projection operator on the space spanned by $\boldsymbol{\xi}$ admits the decomposition

$$P_X = | \boldsymbol{\xi} \rangle \langle \boldsymbol{\xi} | \boldsymbol{\xi} \rangle^{-1} \langle \boldsymbol{\xi} |. \quad (2.21)$$

Notice that P_X is orthogonal, i.e. $P_X^2 = P_X$.

2.3.2 Definition of fitted density

Essentially, density fitting is an orthogonal projection method. The projection is obtained by solving a least-squares fitting problem (see Section 2.3.1.2). This fitting finds numerous applications to electronic density functions in chemistry, hence the term *density fitting* for this kind of application. Before proceeding to the use of density fitting in chemistry, let us announce the general definition and terminology of this method.

Given a function $\rho \in \mathbb{H}$, consider a finite set of M linearly independent elements of \mathbb{H} , denoted by $\boldsymbol{\xi} = \{\xi_i\}_{1 \leq i \leq M} \subseteq \mathbb{H}$, that is referred to as *auxiliary basis set*. Let us denote by $X := \text{Span}(\boldsymbol{\xi})$ the vector space spanned by the auxiliary basis. Density Fitting (DF) proposes to approximate the function ρ by the element of X given by orthogonally projecting ρ on the auxiliary basis, i.e.

$$\rho \approx \tilde{\rho} := P_X \rho = \sum_{i=1}^M c_i \xi_i, \quad (2.22)$$

where $\mathbf{c} = (c_i)_{1 \leq i \leq M} \in \mathbb{R}^M$ and the orthogonal projection operator P_X has been previously defined in Section 2.3.1.2. Finding the coefficients of the vector \mathbf{c} , known as *auxiliary coefficients*, breaks down to solving a least-squares problem (see Section 2.3.1.3). The function $\tilde{\rho}$, called *fitted density*, is essentially the best approximation of ρ on the space spanned by the auxiliary basis set.

Remark. *In chemistry literature, density fitting is also known as the resolution of identity approximation. In operator theory, this method consists of decomposing the identity operator on a finite or infinite sequence of operators. We refer to [243, Definition 6.8.1] for mathematical background. If \mathbb{H} is a separable Hilbert space admitting a countable orthonormal basis, i.e. there exists a complete orthonormal system of elements of \mathbb{H} , denoted by $\mathbf{e} = \{e_i\}_{i \in \mathbb{N}} \subseteq \mathbb{H}$, then the Resolution of Identity (RI) reads*

$$\text{Id} = \sum_{i \in \mathbb{N}} |e_i\rangle \langle e_i|,$$

where Id denotes the identity operator over \mathbb{H} . This means that every element of \mathbb{H} admits an exact representation on the basis \mathbf{e} . Note that, if the set \mathbf{e} is not complete in \mathbb{H} , then we obtain an RI approximation, equivalent to DF. Thus, in operator formulation, we may explain that DF proposes to approximately factorize the identity operator using an auxiliary set $\boldsymbol{\xi}$, typically incomplete, by the decomposition of Eq. (2.21), i.e. $\text{Id} \approx | \boldsymbol{\xi} \rangle \langle \boldsymbol{\xi} | \boldsymbol{\xi} \rangle^{-1} \langle \boldsymbol{\xi} |$.

2.3.3 Application to the pairwise interaction problem

One of the main chemistry applications of density fitting concerns energy calculations. Here we present approximate solutions to the pairwise interaction summation problem defined in Eq. (2.1) using density fitting. The results are organized per interacting densities type. We first present the special case of self-interacting density, before proceeding to the general case of distinct interacting densities and concluding with the case of isometries. Each case will be studied in terms of *a priori* error bounds on interactions. The main idea of this kind of application is to use fitted densities, instead of exact ones, to approximate density-based interactions.

In a general form, we study density-based interactions, for molecular systems whose charge distributions are described by functions. An intermolecular interaction is defined as the scalar value equal to the inner product between any two functions, denoted by

$$\forall (f, g) \in \mathbb{H} \times \mathbb{H}, \quad E(f, g) := \langle f, g \rangle.$$

2.3.3.1 Self-interaction approximation

The *self-interaction* energy of a molecular system described by a density $\rho \in \mathbb{H}$ refers to the scalar quantity

$$E(\rho, \rho) = \langle \rho | \rho \rangle = \|\rho\|^2. \quad (2.23)$$

Assume that ρ is given as a finite expansion over a set of N basis functions, denoted by $\boldsymbol{\rho} = \{\rho_i\}_{1 \leq i \leq N} \subseteq \mathbb{H}$, reading

$$\rho = \sum_{i=1}^N \rho_i. \quad (2.24)$$

Such general form includes as a special case the frozen electronic density functions in the Gaussian electrostatic model, defined in Eq. (2.5). Recall that, in that case, the basis functions are products of atomic orbital products weighted by the density matrix. In a general setting, the self-interaction computation based on the following summation

$$E(\rho, \rho) = \sum_{i=1}^N \sum_{j=1}^N \langle \rho_i | \rho_j \rangle \quad (2.25)$$

scales as $\mathcal{O}(N^2)$, assuming that the sum is evaluated entry-wise and that a single pairwise interaction $\langle \rho_i | \rho_j \rangle$ can be evaluated in $\mathcal{O}(1)$.

Note that the problem of computing the self-interaction using Eq. (2.25) falls into the general family of summation problems of Eq. (2.1). We propose to use density fitting to solve the self-interaction computation problem. In this respect, we expand the density ρ in terms of a smaller auxiliary basis $\boldsymbol{\xi}$ using the density fitting approximation of Eq. (2.22), defined as $\tilde{\rho} := P_X \rho$ with $X = \text{Span}(\boldsymbol{\xi})$, to obtain the self-interaction approximation reading

$$\begin{aligned} E(\rho, \rho) &\approx E(\tilde{\rho}, \tilde{\rho}) \\ &= \langle P_X \rho | P_X \rho \rangle \\ &= \langle \rho | P_X | \rho \rangle \\ &= \langle \rho | \boldsymbol{\xi} \rangle \langle \boldsymbol{\xi} | \boldsymbol{\xi} \rangle^{-1} \langle \boldsymbol{\xi} | \rho \rangle, \end{aligned}$$

where the orthogonal property of the projection operator P_X , i.e. $P_X^2 = P_X$, and the factorization of P_X given by Eq. (2.21) have been evoked. Further using the expansion of the fitted density $\tilde{\rho}$ on the auxiliary basis, given by Eq. (2.22), the approximate self-energy can be evaluated using the reduced sum

$$E(\tilde{\rho}, \tilde{\rho}) = \sum_{i=1}^M \sum_{j=1}^M c_i \langle \xi_i | \xi_j \rangle c_j, \quad (2.26)$$

which scales as $\mathcal{O}(M^2)$. By choosing $M = \mathcal{O}(\sqrt{N})$, one can calculate a self-interaction approximation in linear cost, i.e. $\mathcal{O}(N)$, hence improving the scaling of self-interaction by one order of magnitude.

The error on self-interaction approximation due to density fitting can be quantified as follows, using properties of the orthogonal projection operator. To this respect, the following statement links the self-interaction energy error to the best approximation error due to density fitting.

Theorem 2.3.1 (Density fit error – self-interaction). *Consider a real Hilbert space $(\mathbb{H}, \langle \cdot, \cdot \rangle)$ with induced norm $\| \cdot \| := \langle \cdot, \cdot \rangle^{1/2}$. Let $\boldsymbol{\xi}$ be a finite family of functions in \mathbb{H} and $\mathbb{X} = \text{Span}(\boldsymbol{\xi})$. For any $\rho \in \mathbb{H}$, let $\tilde{\rho} = P_{\mathbb{X}}\rho \in \mathbb{X}$ be the density fitting approximation of ρ . The self-interaction approximation based on density fitting satisfies*

$$E(\rho, \rho) - E(\tilde{\rho}, \tilde{\rho}) = \min_{\sigma \in \mathbb{X}} \|\rho - \sigma\|^2.$$

Proof. See [proof](#) in Section 2.4. □

According to this result, the convergence of the self-interaction error is quadratic in terms of the best approximation error on densities. Moreover, the sign of the self-interaction error is positive. To summarize, such statement shows that the self-interaction energy error caused by density fitting is entirely controlled by the best approximation accuracy of the density on the given subspace \mathbb{X} .

Remark. *Computation time can be gained by performing the fit once, with a given error metric, and reusing the obtained fitted density for calculating interactions given by other inner products. This reduces the storage requirements, as a single fitted density is stored per molecule, instead of storing as many fitted densities as the number of target interaction kernels. This is possible thanks to norm equivalence. The key is that a density ρ given by Eq. (2.24) belongs to the finite-dimensional vector space $V := \text{Span}(\rho_i)_{1 \leq i \leq N}$. It is known that all norms on finite-dimensional vector spaces are equivalent [57]. It is hence straightforward to obtain the desired density fit transferability from the norm equivalence between the Dirac and Coulomb metrics:*

$$\forall \rho \in V, \quad C_1 \|\rho\|_{L^2(\mathbb{R}^3)} \leq \|\rho\|_{\dot{H}^{-1}(\mathbb{R}^3)} \leq C_2 \|\rho\|_{L^2(\mathbb{R}^3)},$$

where the positive constants C_1 and C_2 depend on the dimension of V .

2.3.3.2 Intermolecular interaction approximation

The general family of summation problems of Eq. (2.1) occurs when computing intermolecular pairwise interactions, between any two distinct densities. In the present section, we apply density fitting separately to each of the two densities, to solve this problem.

The *intermolecular interaction* energy between any two given functions $\rho_A, \rho_B \in \mathbb{H}$ is given as the scalar quantity

$$E(\rho_A, \rho_B) = \langle \rho_A | \rho_B \rangle. \quad (2.27)$$

We further assume that densities ρ_A and ρ_B are given as finite expansions over sets of N_A and N_B basis functions, respectively, reading:

$$\rho_A = \sum_{i=1}^{N_A} \rho_i^A, \quad \rho_B = \sum_{i=1}^{N_B} \rho_i^B,$$

where $\boldsymbol{\rho}_A := \{\rho_i^A\}_{1 \leq i \leq N_A} \subseteq \mathbb{H}$ and $\boldsymbol{\rho}_B := \{\rho_i^B\}_{1 \leq i \leq N_B} \subseteq \mathbb{H}$ are two given finite sets of functions. For simplicity in the notation, we consider that N_A and N_B have the same order of magnitude, i.e. there exist N such that $\mathcal{O}(N) = \mathcal{O}(N_A) = \mathcal{O}(N_B)$. The computation of intermolecular interaction of Eq. (2.27) by direct summation of pairwise interactions, i.e.

$$E(\rho_A, \rho_B) = \sum_{i=1}^{N_A} \sum_{j=1}^{N_B} \langle \rho_i^A | \rho_j^B \rangle, \quad (2.28)$$

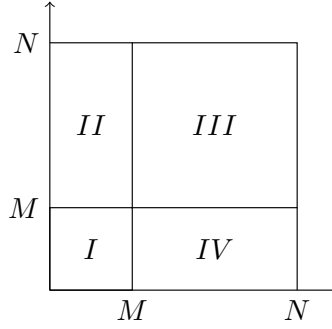


Figure 2.7: Basis set size diagram. Partition *I* corresponds to classical DF summation domain. The union of partitions *I*, *II* and *IV* corresponds to robust DF method. The union of all four partitions corresponds to the exact evaluation of $\langle \rho | \rho \rangle$.

scales as $\mathcal{O}(N^2)$.

We plan to reduce the quadratic cost using density fitting. In this respect, we introduce two auxiliary basis sets $\xi_A, \xi_B \subseteq \mathbb{H}$, whose sizes are equal to fixed positive integers M_A and M_B , respectively, assuming there exist M such that $\mathcal{O}(M) = \mathcal{O}(M_A) = \mathcal{O}(M_B)$ for simplicity. Now, applying the density fitting approximation of Eq. (2.22) separately on the density ρ_F , for any $F \in \{A, B\}$, yields the fitted density $\tilde{\rho}_F$, defined as the best approximation of ρ_F on the space $X_F := \text{Span}(\xi_F)$. There exist two available schemes for defining intermolecular interaction approximations based on fitted densities:

- **Density Fitting.** Directly generalizing the self-interaction case, one approximates the reference quantity of Eq. (2.27) by

$$E^{\text{DF}}(\tilde{\rho}_A, \tilde{\rho}_B) := E(\tilde{\rho}_A, \tilde{\rho}_B). \quad (2.29)$$

The computational complexity of this evaluation is $\mathcal{O}(M^2)$. By building auxiliary bases for small targeted sizes M , namely $M = \mathcal{O}(\sqrt{N})$, one can achieve linear cost $\mathcal{O}(N)$ in practice.

- **Robust density fitting.** The Robust Density Fitting (RDF) [108, 107, 106, 229, 115] method is designed to include corrections to DF, by proposing to evaluate the intermolecular interaction as

$$E^{\text{RDF}}(\rho_A, \rho_B, \tilde{\rho}_A, \tilde{\rho}_B) := \langle \rho_A | \tilde{\rho}_B \rangle + \langle \tilde{\rho}_A | \rho_B \rangle - \langle \tilde{\rho}_A | \tilde{\rho}_B \rangle. \quad (2.30)$$

The computational complexity of this quantity is $\mathcal{O}(MN)$ or $\mathcal{O}(N^{3/2})$, by appropriately choosing $M = \mathcal{O}(\sqrt{N})$.

In the above energy definitions, arguments will be omitted when this does not lead to a confusion.

Note that the complexity of RDF is worse than that of DF, due to cross-terms between reference densities and fitted ones, that are unfortunately dominating the computational and storage cost. Figure 2.7 schematically represents the basis set interaction sizes on the same diagram for comparison. Nevertheless, the gain of RDF lies in its accuracy. To see this, first notice that RDF is obtained by adding corrective terms to DF:

$$E^{\text{RDF}} = E^{\text{DF}} + \langle \rho_A - \tilde{\rho}_A | \tilde{\rho}_B \rangle + \langle \tilde{\rho}_A | \rho_B - \tilde{\rho}_B \rangle. \quad (2.31)$$

The purpose of corrective terms can be explained as follows. To begin with, it is easy to observe that the following identity holds:

$$\begin{aligned} \langle \rho_A | \rho_B \rangle - \langle \tilde{\rho}_A | \tilde{\rho}_B \rangle &= \langle \rho_A - \tilde{\rho}_A | \rho_B - \tilde{\rho}_B \rangle + \langle \rho_A - \tilde{\rho}_A | \tilde{\rho}_B \rangle + \langle \tilde{\rho}_A | \rho_B - \tilde{\rho}_B \rangle \\ &= \langle \rho_A - \tilde{\rho}_A | \rho_B - \tilde{\rho}_B \rangle + E^{\text{RDF}} - E^{\text{DF}}. \end{aligned} \quad (2.32)$$

In the case where $\xi_A = \xi_B$, then corrective terms in (2.31) vanish, since fitted densities belong to the same subspace, i.e. $\tilde{\rho}_A \in X_A$ and $\tilde{\rho}_B \in X_B$ with $X_A = X_B$, and orthogonal projection errors vanish on subspaces, as given by property (2.15). Therefore, in this case, RDF and DF coincide. By using this observation within identity (2.32) and applying Cauchy-Schwartz inequality to remaining terms, one obtains that the intermolecular interaction error scales quadratically in terms of best approximation errors. However, in the general case where $\xi_A \neq \xi_B$, RDF and DF do not coincide. In particular, non-vanishing terms in equation (2.32) result in linear scaling of the intermolecular interaction error due to DF with respect to best approximation errors, as seen again by application of the Cauchy-Schwartz inequality. The role of RDF is then to recover quadratic error convergence, which is achieved by adding appropriate corrective terms to DF. For this reason, RDF can be used in applications where accuracy is a priority, despite of computational cost.

The following two results are generalizations to Whitten's *a priori* error estimators, established for the special case of electron repulsion interactions in [168, 169], here extended to interactions induced by any given inner product.

Theorem 2.3.2 (Density fit error). *Consider a real Hilbert space $(H, \langle \cdot, \cdot \rangle)$ with induced norm $\|\cdot\| := \langle \cdot, \cdot \rangle^{1/2}$. For $F \in \{A, B\}$, let ξ_F be a finite family of functions in H and $X_F = \text{Span}(\xi_F)$. For any $\rho_F \in H$, let $\tilde{\rho}_F := P_{X_F} \rho_F \in X_F$ be the density fitting approximation of ρ_F . The intermolecular interaction approximation of Eq. (2.29) based on density fitting satisfies*

$$|E(\rho_A, \rho_B) - E(\tilde{\rho}_A, \tilde{\rho}_B)| \leq \|\rho_B\| \min_{\sigma_A \in X_A} \|\rho_A - \sigma_A\| + \|\rho_A\| \min_{\sigma_B \in X_B} \|\rho_B - \sigma_B\|.$$

Proof. See proof in Section 2.4. □

Theorem 2.3.3 (Robust density fit error). *Under the same assumptions as in Theorem 2.3.2, the intermolecular interaction approximation of Eq. (2.30) based on robust density fitting satisfies*

$$|E(\rho_A, \rho_B) - E^{\text{RDF}}(\rho_A, \rho_B, \tilde{\rho}_A, \tilde{\rho}_B)| \leq \min_{\sigma_A \in X_A} \|\rho_A - \sigma_A\| \min_{\sigma_B \in X_B} \|\rho_B - \sigma_B\|.$$

Proof. See proof in Section 2.4. □

Note that DF error converges linearly, while RDF recovers quadratic convergence thanks to corrective terms. Another remark is that, contrary to the positive self-interaction error of Theorem 2.3.1, now the error on energies can have any sign. This may lead to error cancellation when summing different energy terms (we refer to discussion of Section 1.2.2.1). To summarize, estimators show that the intermolecular interaction error is entirely controlled by the best approximation error on densities, due to the use of auxiliary basis sets.

Densities under the action of isometries. In the present section, we study density fitting under the action of isometries. This context has applications to the Gaussian electrostatic model for rigid fragments (see Section 2.1.2.2). Our results show that error estimators on energies, which are calculated from interacting densities transformed by isometries, remain unchanged by the action of isometries.

As a follow-up to Section 2.3.3.2, the problem of calculating Eq. (2.28) is now studied in the special case where densities undergo the action of isometries. The following two statements show that intermolecular interaction error estimators of Theorems 2.3.2 and 2.3.3 are preserved, under the action of isometries on both reference and fitted densities. The main reason behind is that the best approximation error is isometry-independent. Before stating the results, it is important to clarify that they can be easily adjusted to the frozen densities of Eq. (2.4) by applying them using local frames for functions, namely $\rho_{F'}$, $\tilde{\rho}_{F'}$ and $\xi_{F'}$ for any $F \in \{A, B\}$. Since the local frame is the global frame up to isometry for rigid fragments, the heavy notation, i.e. primes for local frames, will be omitted without this affecting the generality of the results.

Corollary 2.3.3.1 (Density fit error – rigid fragments). *Under the same assumptions as in Theorem 2.3.2, let $T_F : \mathbb{H} \rightarrow \mathbb{H}$ be an isometry for any given $F \in \{A, B\}$. The intermolecular interaction approximation based on density fitting, introduced in Eq. (2.29), satisfies*

$$|E(T_A \rho_A, T_B \rho_B) - E(T_A \tilde{\rho}_A, T_B \tilde{\rho}_B)| \leq \|\rho_B\| \min_{\sigma_A \in X_A} \|\rho_A - \sigma_A\| + \|\rho_A\| \min_{\sigma_B \in X_B} \|\rho_B - \sigma_B\|.$$

Proof. See [proof](#) in Section 2.4. □

Corollary 2.3.3.2 (Robust density fit error – rigid fragments). *Under the same assumptions as in Corollary 2.3.3.1, the intermolecular interaction approximation based on robust density fitting, introduced in Eq. (2.30), can be decomposed as*

$$|E(T_A \rho_A, T_B \rho_B) - E^{\text{RDF}}(T_A \rho_A, T_B \rho_B, T_A \tilde{\rho}_A, T_B \tilde{\rho}_B)| \leq \min_{\sigma_A \in X_A} \|\rho_A - \sigma_A\| \min_{\sigma_B \in X_B} \|\rho_B - \sigma_B\|.$$

Proof. See [proof](#) in Section 2.4. □

Note that the above estimators show that the intermolecular interaction error may be minimized by optimizing isometry-independent best approximation errors. It thus suffices to calculate fitted densities once for a single density pair and then simply transform them by an isometry, in order to generate fitted densities for density pairs at different orientations. Additionally, if each fragment density needs to be fitted separately from others, it is suitable to consider X_A and X_B independent of each other. Then best approximation errors can also be optimized separately on different isolated fragments.

The following statement shows that the intermolecular interaction based on density fitting can be decomposed into isometry-independent and -dependent components. Such decomposition is useful since it allows to detect quantities that can be precomputed in advance for all isometries. Moreover, the result provides an approximated solution to the summation problem of Section 2.1.2.2 and justifies the offline precomputation (see Section 2.2.1.2) of fitted densities in the Gaussian electrostatic model for rigid fragments.

Theorem 2.3.4. *Under the same assumptions as in Corollary 2.3.3.1, let us further consider, for any given $F \in \{A, B\}$, the family of functions $\{T_F \xi_i^F\}_{1 \leq i \leq M_F} \subseteq \mathbb{H}$, denoted by an abuse of notation by $T_F \xi_F$. Then the intermolecular energy of Eq. (2.29) based on density fitting satisfies the formula*

$$E(T_A \tilde{\rho}_A, T_B \tilde{\rho}_B) = \langle \rho_A | \xi_A \rangle \langle \xi_A | \xi_A \rangle^{-1} \langle T_A \xi_A | T_B \xi_B \rangle \langle \xi_B | \xi_B \rangle^{-1} \langle \xi_B | \rho_B \rangle. \quad (2.33)$$

Proof. See [proof](#) in Section 2.4. □

Figure 2.8 represents a dimension diagram for the matrix decomposition of Eq. (2.33). Note that the only isometry-dependent quantity is the cross-Gram matrix $\langle T_A \xi_A | T_B \xi_B \rangle$. All remaining quantities are isometry-independent, they can therefore be precalculated. Note that some of these quantities have in fact a much larger dimension than the small red matrix, since they depend on reference densities. Others involve a matrix inversion that requires computational effort. Hence the precomputation gain is significant.

2.4 Proofs

Proof of Proposition 2.1.1. The following calculation on the inverse Laplacian operator Δ^{-1} for the case $d = 3$ is originally due to [313]. When f is a plane wave such as $f(\mathbf{r}) = e^{2\pi i \mathbf{r} \cdot \boldsymbol{\xi}}$, then there is a relationship

$$\Delta e^{2\pi i \mathbf{r} \cdot \boldsymbol{\xi}} = -4\pi^2 |\boldsymbol{\xi}|^2 e^{2\pi i \mathbf{r} \cdot \boldsymbol{\xi}}$$

that we can multiply by Δ^{-1} to obtain

$$\Delta^{-1}(e^{2\pi i \mathbf{r} \cdot \boldsymbol{\xi}}) = -\frac{|\boldsymbol{\xi}|^{-2}}{4\pi^2} e^{2\pi i \mathbf{r} \cdot \boldsymbol{\xi}}.$$

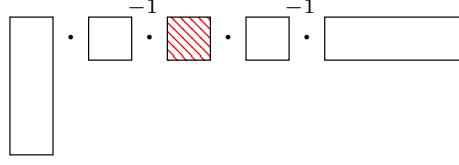


Figure 2.8: Matrix decomposition of intermolecular interaction. The isometry-dependent component (in red) is a small M_A -by- M_B matrix, evaluated entry-wise, whose dimension depends on the auxiliary basis sizes M_A and M_B .

We thus have a formula involving the Fourier transform for the inverse Laplacian of a general function:

$$\begin{aligned}\Delta^{-1}(f)(\mathbf{r}) &= \Delta^{-1} \int_{\mathbb{R}^3} \hat{f}(\boldsymbol{\xi}) e^{2\pi i \mathbf{r} \cdot \boldsymbol{\xi}} \, d\boldsymbol{\xi} \\ &= \int_{\mathbb{R}^3} \hat{f}(\boldsymbol{\xi}) \Delta_x^{-1}(e^{2\pi i \mathbf{r} \cdot \boldsymbol{\xi}}) \, d\boldsymbol{\xi} \\ &= -\frac{1}{4\pi^2} \int_{\mathbb{R}^3} |\boldsymbol{\xi}|^{-2} \hat{f}(\boldsymbol{\xi}) e^{2\pi i \mathbf{r} \cdot \boldsymbol{\xi}} \, d\boldsymbol{\xi}.\end{aligned}$$

Now, we can write the \dot{H}^{-1} norm of u as

$$\begin{aligned}\|u\|_{\dot{H}^{-1}(\mathbb{R}^3)}^2 &= \int_{\mathbb{R}^3} |\boldsymbol{\xi}|^{-2} |\hat{u}(\boldsymbol{\xi})|^2 \, d\boldsymbol{\xi} \\ &= \int_{\mathbb{R}^3} \int_{\mathbb{R}^3} |\boldsymbol{\xi}|^{-2} \hat{u}(\boldsymbol{\xi}) e^{-2\pi i \mathbf{r} \cdot \boldsymbol{\xi}} u(\mathbf{r}) \, d\mathbf{r} \, d\boldsymbol{\xi} \\ &= -4\pi^2 \int_{\mathbb{R}^3} \Delta^{-1}(u)(\mathbf{r}) u(\mathbf{r}) \, d\mathbf{r} \\ &= \pi \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{|u(\mathbf{r}) u(\mathbf{r}')|}{|\mathbf{r} - \mathbf{r}'|} \, d\mathbf{r} \, d\mathbf{r}'\end{aligned}$$

and this last quantity is called Coulomb metric of u (up to a constant π), induced by the Coulomb kernel. We deduced the last equality from the fact that the Coulomb potential satisfies the Poisson equation

$$-(4\pi)^{-1} \Delta \left(\int_{\mathbb{R}^3} \frac{|u(\mathbf{r}')|}{|\mathbf{r} - \mathbf{r}'|} \, d\mathbf{r}' \right) (\mathbf{r}) = u(\mathbf{r})$$

or by inverting,

$$\Delta^{-1}(u)(\mathbf{r}) = -(4\pi)^{-1} \int_{\mathbb{R}^3} \frac{|u(\mathbf{r}')|}{|\mathbf{r} - \mathbf{r}'|} \, d\mathbf{r}'.$$

This concludes the proof of Proposition 2.1.1. □

Proof of Proposition 2.1.2. See [24, Proposition 1.34]. □

Proof of Lemma 2.1.1. By definition of the given **L2G** operator T , one may write

$$\begin{aligned}\langle Tf, Tg \rangle_K &= \langle f(\mathbf{Q}(\cdot - \mathbf{t})), g(\mathbf{Q}(\cdot - \mathbf{t})) \rangle_K \\ &= \iint f(\mathbf{Q}(\mathbf{r} - \mathbf{t})) K(\mathbf{r}, \mathbf{r}') g(\mathbf{Q}(\mathbf{r}' - \mathbf{t})) \, d\mathbf{r} \, d\mathbf{r}'.\end{aligned}$$

By employing the change of variables $\mathbf{r} \mapsto \mathbf{Q}^{-1}\mathbf{r}$ and $\mathbf{r}' \mapsto \mathbf{Q}^{-1}\mathbf{r}'$ the above integral becomes

$$\iint f(\mathbf{r} - \mathbf{t}) K(\mathbf{Q}^{-1}\mathbf{r}, \mathbf{Q}^{-1}\mathbf{r}') g(\mathbf{r}' - \mathbf{t}) |\det(\mathbf{Q}^{-1})|^2 \, d\mathbf{r} \, d\mathbf{r}'. \quad (2.34)$$

The matrix \mathbf{Q}^{-1} belongs to $O(3)$, thus its determinant in absolute value is equal to one. The isometry invariance property of kernel K yields $K(\mathbf{Q}^{-1}\mathbf{r}, \mathbf{Q}^{-1}\mathbf{r}') = K(\mathbf{r} - \mathbf{t}, \mathbf{r}' - \mathbf{t})$. Finally a change of variable in Eq. (2.34) by the affine transformations $\mathbf{r} \mapsto \mathbf{r} - \mathbf{t}$ and $\mathbf{r}' \mapsto \mathbf{r}' - \mathbf{t}$ yields the result. This completes the proof of Lemma 2.1.1. \square

Proof of Theorem 2.3.1. One evaluates the orthogonal projection error for $\tilde{\rho} = P_X\rho$ as

$$\begin{aligned} \|\rho - P_X\rho\|^2 &= \langle \rho - P_X\rho, \rho - P_X\rho \rangle \\ &= \langle \rho, \rho - P_X\rho \rangle - \underbrace{\langle P_X\rho, \rho - P_X\rho \rangle}_{=0} \\ &= \langle \rho, \rho - P_X\rho \rangle + \langle P_X\rho, \rho - P_X\rho \rangle \\ &= \langle \rho, \rho \rangle - \langle \rho, P_X\rho \rangle + \langle P_X\rho, \rho \rangle - \langle P_X\rho, P_X\rho \rangle, \end{aligned}$$

where the term vanished due to the orthogonal projection definition of Eq. (2.15). Lastly, evoking the symmetry of the inner product, symmetric terms vanish and one obtains the desired expression of Theorem 2.3.1. \square

Proof of Theorem 2.3.2. By adding and subtracting terms, one evaluates the intermolecular interaction error as

$$\begin{aligned} \langle \rho_A, \rho_B \rangle - \langle \tilde{\rho}_A, \tilde{\rho}_B \rangle &= \langle \rho_A, \rho_B \rangle - \langle \rho_A, \tilde{\rho}_B \rangle + \langle \rho_A, \tilde{\rho}_B \rangle - \langle \tilde{\rho}_A, \tilde{\rho}_B \rangle \\ &= \langle \rho_A, \rho_B - \tilde{\rho}_B \rangle + \langle \rho_A - \tilde{\rho}_A, \tilde{\rho}_B \rangle. \end{aligned}$$

First employing Cauchy-Schwartz inequality, then using that $\tilde{\rho}_B$ is equal to the orthogonal projection, hence $\|\tilde{\rho}_B\| \leq \|\rho_B\|$, lastly majorizing with the best approximation error, we obtain the bound

$$\begin{aligned} |\langle \rho_A, \rho_B \rangle - \langle \tilde{\rho}_A, \tilde{\rho}_B \rangle| &\leq \|\rho_A\| \|\rho_B - \tilde{\rho}_B\| + \|\tilde{\rho}_B\| \|\rho_A - \tilde{\rho}_A\| \\ &\leq \|\rho_A\| \min_{\sigma_B \in X_B} \|\rho_B - \sigma_B\| + \|\rho_B\| \min_{\sigma_A \in X_A} \|\rho_A - \sigma_A\|. \end{aligned}$$

This concludes the proof of Theorem 2.3.2. \square

Proof of Theorem 2.3.3. The interaction error expression for RDF approximation is equal to

$$\langle \rho_A, \rho_B \rangle - E^{\text{RDF}} = \langle \rho_A - \tilde{\rho}_A, \rho_B - \tilde{\rho}_B \rangle.$$

Employing Cauchy-Schwartz inequality, we obtain the bound

$$|\langle \rho_A, \rho_B \rangle - E^{\text{RDF}}| \leq \|\rho_A - \tilde{\rho}_A\| \|\rho_B - \tilde{\rho}_B\|.$$

Lastly, by majorizing errors by best approximation errors, similarly to the proof of Theorem 2.3.2, one obtains the desired result of Theorem 2.3.3. \square

Proof of Corollary 2.3.3.1. We employ Theorem 2.3.2 and the distance-preserving property of Eq. (2.10). \square

Proof of Corollary 2.3.3.2. We employ Theorem 2.3.3 and the distance-preserving property of Eq. (2.10). \square

Proof of Theorem 2.3.4. We first perform single-fragment density fitting, as given by Eq. (2.22), to obtain a fitted density for any given $F \in \{A, B\}$, defined as

$$\tilde{\rho}_F = \sum_{i=1}^{M_F} c_i^F \xi_i^F,$$

where $\mathbf{c}_F = (c_i^F)_{1 \leq i \leq M_F} \in \mathbb{R}^{M_F}$ is the vector of orthogonal projection coefficients, obtained by solving the normal equations, as explained in Section 2.3.1.3, reading

$$\mathbf{c}_F = \langle \boldsymbol{\xi}_F | \boldsymbol{\xi}_F \rangle^{-1} \langle \boldsymbol{\xi}_F | \rho_F \rangle. \quad (2.35)$$

Using that any isometry that is an endomorphism over a real Hilbert space is affine, that is a consequence of the Mazur-Ulam theorem [226], we have that for any given $F \in \{A, B\}$ and the isometry T_F in particular there holds

$$T_F \tilde{\rho}_F = \sum_{i=1}^{M_F} c_i^F T_F \xi_i^F.$$

The intermolecular interaction based on density fitting then reads, in matrix form, as

$$\begin{aligned} \langle \tilde{\rho}_A | \tilde{\rho}_B \rangle &= \sum_{i=1}^{M_A} \sum_{j=1}^{M_B} c_i^A \langle T_A \xi_i^A | T_B \xi_j^B \rangle c_j^B \\ &= \mathbf{c}_A^\top \langle T_A \boldsymbol{\xi}_A | T_B \boldsymbol{\xi}_B \rangle \mathbf{c}_B. \end{aligned}$$

Substituting the vector \mathbf{c}_F using Eq. (2.35), one obtains the desired result of Theorem 2.3.4. \square

2.5 Conclusion

In the present document, we revisit the density fitting method, also known as resolution of identity. Compared with conventional applications of density fitting in electronic structure theory, the major difficulty in carrying out the formulation of density fitting in molecular dynamics context is inscribing the constantly changing molecular fragment orientation. In order to proceed, we introduce a formulation that allows to track coordinate systems and operations between them, in the case of rigid fragments. The result is a complete mathematical framework for describing the Gaussian electrostatic model frozen core intermolecular interaction computation for rigid fragments, under the Dirac and Coulomb kernels.

Our analysis focuses on the frozen core of the Gaussian Electrostatic Model coupled with density fitting. In order to describe this coupling in more detail, we split the computation time into offline-online phases. From a theoretical point of view, our analysis contributes in mathematically validating the foundations of this density fitting application for calculating arbitrary kernel explicit frozen core intermolecular interactions between rigid fragments pairs. By demonstrating *a priori* convergence properties, we justify the use of density fitting on each fragment separately, during the offline frozen fragment phase, and its online reutilization for rigid fragments. In general, based on our results, we can conclude that the rigid fragment case generally mimics the standard single-fragment one.

We plan to carry out our analysis on flexible fragments. One difficulty to consider while testing fragments of flexible intramolecular geometry is that the offline density fit should somehow take into account electronic densities at all possible molecular geometries. The hope is that significant information on electronic densities at arbitrary geometry can be detected and precomputed offline. Another difficulty lies in the fact that the coordinate transformation mapping offline fragments to flexible ones is no longer an isometry, and in general little is known about its closed-form expression. This case should then be handled carefully, since transforming density approximations using shape-non-preserving mappings has the risk of introducing untractable errors. In order to tackle such problems and extract as much information as possible during the offline phase, Chapter 5 employs a reduced basis approach.

Chapter 3

Auxiliary basis generation

Contents

3.1 Introduction	56
3.2 Theory	57
3.2.1 Problem formulation	57
3.2.2 From basis generation to submatrix selection	59
3.2.3 The case of electronic density functions	61
3.2.4 Auxiliary basis generation schemes	65
3.2.5 Summary of Auxiliary Basis Set (ABS) generation scheme	69
3.3 Numerical results	70
3.3.1 Frozen core energy of rigid fragments	71
3.3.2 Low-rank structure of Gram matrices	74
3.3.3 Auxiliary basis set detailed comparison	78
3.4 Proofs	81
3.5 Conclusion and perspectives	82

In the present chapter, we formulate the problem of generating optimal auxiliary basis sets for density fitting. We plan to design a novel auxiliary basis generation algorithm that will automatically provide arbitrary precision basis sets, made up of atomic Gaussian-type orbitals, from a required accuracy on the density fit quality. We define a new auxiliary basis from any given electronic density function discretized over atomic orbital basis sets. Our tools are based on Gram matrix row and column selection, using a greedy selection procedure starting from the full Gram matrix of orbital products. Particular attention is paid to meeting the requirements of atomic Gaussian-type orbital basis sets. In this respect, we introduce a block treatment of the Gram matrix. Generated basis sets are numerically tested in self-interaction and frozen core intermolecular interaction computations. Comparisons with other state-of-the-art empirical auxiliary basis sets are provided.

3.1 Introduction

The problem of auxiliary basis construction exists since the beginnings of density fitting. Existing auxiliary basis sets include most often Gaussian-type orbitals and in some cases (in the Gaussian electrostatic model in particular, defined in Section 1.1.6.4) Hermite-Gaussian type orbitals [84]. Orbital centers are typically restricted to atomic positions [158]. Additional sites are sometimes placed on midbond centers for more accurate representation of the density [257, 233].

There exist two principal strategies for generating auxiliary basis parameters in the literature. One is empirically precalculating auxiliary basis sets for every atom chemical type, creating a library of atomic basis sets. This parameter tuning is usually application-specific and is carried out for approximating various molecular integrals. Commonly used families of this kind are the JKFIT family [335], optimized for Coulomb and Hartree-Fock exchange integrals, and the RI family [337], with applications, among others, to second-order Møller-Plesset perturbation theory integrals [38, 151]. The difficulties of using empirical rules for auxiliary basis generation have been discussed in chemistry literature [324, 336]. The main drawback of empirical basis sets is that they often provide functions that do not accurately span the desired space. Another limitation is that they can neither be tuned to achieve lower or higher accuracy, nor be optimized for new energy terms.

Alternatives include systematic strategies based on matrix approximations. Much attention has been paid to auxiliary basis sets systematically resulting from atomic orbital product basis sets. A relevant class of methods exploits the linear dependencies of rows (and columns) of Gram matrices on Gaussian orbitals. The goal is to remove linear dependencies of the atomic orbital product space [8], by creating redundancy-free product-basis sets. This is achieved by re-expanding atomic orbital products in an auxiliary basis set [324]. It has been numerically observed that the full space may be accurately approximated using only a linearly scaling number of auxiliary functions with respect to the number of atomic orbital basis functions [198, 31]. Several methods have been employed for selecting such auxiliary functions. A popular family is based on a pivoted Cholesky-selected auxiliary functions [10, 9, 198], known as *unbiased* auxiliary sets, generated on-the-fly for a given molecular system. Such methods use the pivots of the Cholesky decomposition to discard rows and columns under a prescribed tolerance criterion. This method can be combined with other low-rank procedures, for instance in [178] the authors perform a Löwdin approximation (truncated eigendecomposition) on the Cholesky factors and select auxiliary basis functions using a singular value criterion. This post-processing allows to systematically further reduce the size of a given auxiliary basis. One of the main drawbacks of current methods is that the product selection procedure is performed starting from the orbital basis specified for an atom. In this work we present a method to generate more accurate auxiliary basis sets using fewer auxiliary functions, starting from the orbital basis specified for an entire molecule and exploiting additional knowledge on the available density matrix. This method finds application to pairwise intermolecular interactions for frozen densities in particular.

In the present work, we focus on building auxiliary basis sets that achieve low computational complexity, in a constraint optimization point of view. To achieve this, we employ a matrix row and column selection method using greedy criteria. Moreover, we analyze the density fit errors induced by our generated auxiliary basis sets and discuss the computational cost of an auxiliary basis in computing molecular integrals. The remaining sections are organized as follows. The auxiliary basis construction problem is introduced in Section 3.2.1, followed by its formulation in the case of density functions. Density fit error analysis caused by our generated auxiliary basis sets is presented in Section 3.2.2. New approximation schemes are first presented in Section 3.2.4 and then numerically tested in Section 3.3. All proofs of mathematical results are postponed to Section 3.4.

3.2 Theory

In the present chapter, we focus on methods for generating auxiliary basis sets of target size respecting a prescribed level of accuracy. An accuracy criterion is the error of projecting a given density on a given auxiliary basis set, which governs the interaction energy error caused by density fitting, as shown in Section 2.3.3. As a consequence of the underlying least-squares problem of density fitting, the auxiliary basis accuracy can be assessed in matrix form. For the sake of clarity, we first announce the overall targeted problem in Section 3.2.1 and we specify details later, when treating specific aspects regarding electronic density functions and atomic orbitals in Section 3.2.3.

3.2.1 Problem formulation

We tackle Problem P.1, which is a preliminary step of the density fitting method and is recalled here: given a density ρ , construct the Auxiliary Basis Set (ABS) $\xi = \{\xi_i\}_{1 \leq i \leq M}$ used in the density fitting approximation of Eq. (2.13). To this end, we first introduce our notion of optimal ABS and then formulate the problem in question using a novel constraint optimization framework. The idea is to minimize the accuracy under a complexity constraint. Before giving the definition of this minimization problem, let us quickly mention the underlying notions. First, the accuracy of a given auxiliary basis set can be assessed in terms of the best approximation error on the density, defined by Eq. (2.16). Second, the complexity of an auxiliary basis set is measured in terms of the total number of pairwise interactions in the summation of Eq. (2.26), that scales as $\mathcal{O}(M^2)$, where M is the auxiliary basis set size. We also take into account the operation count for computing a single pairwise interaction, that is related to the analytic form of admissible auxiliary functions.

The problem formulation starts from the most general case to end up with a restricted one that is easier to solve in practice. In particular, Section 3.2.1.1 describes the optimization problem in its more general form. In Section 3.2.1.2, we then focus on a heuristic optimization suitable for the special case of electronic density functions, that leads to the problem that we solve in our work.

3.2.1.1 The full ABS generation problem

Let \mathbb{H} be a Hilbert space endowed with the inner product $\langle \cdot, \cdot \rangle$ and the induced norm $\|\cdot\| := \langle \cdot, \cdot \rangle^{1/2}$. Let \mathcal{G} denote a given admissible set of auxiliary basis functions that are elements of \mathbb{H} . The operation count for computing a single interaction between any two elements in \mathcal{G} is assumed to be a fixed number, denoted by $\kappa_{\mathcal{G}} \in \mathbb{N}^*$ and satisfying

$$\forall(\xi, \xi') \in \mathcal{G} \times \mathcal{G}, \quad \kappa_{\mathcal{G}} = \text{cost}(\langle \xi, \xi' \rangle). \quad (3.1)$$

Assuming this fixed cost, our strategy consists of minimizing the size of the ABS, which is made up of elements in \mathcal{G} . The problem of achieving accurate interaction approximation with few auxiliary functions belonging to \mathcal{G} may be formulated as an optimization problem. Given any $\rho \in \mathbb{H}$ and a target size M , the *ABS generation problem* refers to the problem of finding a finite family $\xi \subseteq \mathcal{G}$ of size at most M , such that the best approximation error of ρ on the space spanned by ξ is minimized, i.e.

$$\min_{\xi \subseteq \mathcal{G}} \min_{\sigma \in \text{Span}(\xi)} \|\rho - \sigma\|^2 \quad (3.2)$$

$$\text{subject to} \quad \text{card}(\xi) \leq M, \quad (3.3)$$

where $\text{card}(\cdot)$ denotes the size of a given set. A solution to ABS generation then attains the smallest best approximation error on ρ over approximation spaces spanned by at most M admissible functions.

Concerning the nature of our proposed problem, let us emphasize that our target size M concerns the number of admissible elements. In other words, our problem is picking individual functions in \mathcal{G} whose total number is equal to M , rather than finding a basis of M elements in $\text{Span}(\mathcal{G})$. The main reason is the complexity of computing a single interaction. The point to be handled carefully can be highlighted as follows. Let $\{\xi_i\}_{1 \leq i \leq M}$ be a given basis of a subspace of $\text{Span}(\mathcal{G})$. Then performing density fitting using this basis yields the approximation $\rho \approx \sum_{i=1}^M c_i \xi_i$. For two given fragments A and B and corresponding isometries T_A (resp. T_B) sending A' to A (resp. B' to B), the pairwise interaction is given by $\langle \rho_A, \rho_B \rangle \approx \sum_{i,j=1}^M c_i^{A'} c_j^{B'} \langle T_A \xi_i^{A'}, T_B \xi_j^{B'} \rangle$. Each one of the interactions $\langle T_A \xi_i^{A'}, T_B \xi_j^{B'} \rangle$ cannot be precomputed and costs $n^2 \kappa_{\mathcal{G}}$, if $\xi_i^{A'}$ and $\xi_j^{B'}$ are assumed to be expressed as linear combinations of $n \in \mathbb{N}^*$ elements in \mathcal{G} . However, this online cost can become untractable if n is arbitrarily large. For this reason, we impose $n = 1$. Hence, by picking individual functions in \mathcal{G} instead of basis elements, we assure that the complexity of a single interaction computed online is equal to $\kappa_{\mathcal{G}}$ and remains controlled.

A difficulty in solving the minimization problem of Eq. (3.2) under the constraint of Eq. (3.3) may be identified. The simple example where the set \mathcal{G} is equal to the set of primitive Cartesian Gaussian-type orbitals (GTOs), of the form of Eq. (1.18), allows to illustrate this. One may think of GTO parameters as continuous variables to be optimized. Unfortunately, optimizing GTOs with respect to their four parameters, i.e. orbital exponent and monomial degrees, can be a hard combinatorial problem, connected to NP-hard selection problems in optimization [52]. Second, the best approximation problem over arbitrary GTO sets can be highly ill-conditioned (see discussion in Section 1.1.4.3). Best approximation methods [154] on spaces spanned by ill-conditioned basis sets might become necessary.

3.2.1.2 The restricted problem

Our approach is to avoid difficulties caused by the search over the space \mathcal{G} in Eq. (3.2) and propose a restricted search, based on heuristics defined for a specific form of ρ , allowing to significantly narrow down the search space. In the present paragraph, our special ABS choice is proposed for simplifying the problem.

Consider a given function $\rho \in \mathbb{H}$ defined as

$$\rho := \sum_{i=1}^N \rho_i, \quad (3.4)$$

where $\boldsymbol{\rho} := \{\rho_i\}_{1 \leq i \leq N} \subseteq \mathbb{H}$ is a given finite set of functions. Recall that this is the function form of interest that arises in our main problem previously introduced in Section 2.1.1. We propose to carry out the ABS construction by looking at subsets of $\boldsymbol{\rho}$ of the form

$$\forall S \subseteq \{1, \dots, N\}, \quad \boldsymbol{\xi}_S := \{\rho_i\}_{i \in S}. \quad (3.5)$$

This consideration leads to the following naive but convenient technique. Assuming that $\boldsymbol{\rho}$ contains admissible ABS elements (e.g. $\boldsymbol{\rho} \subset \mathcal{G}$), our goal is to find an optimal selection of indices $S \subseteq \{1, \dots, N\}$, such that the following density fitting approximation

$$\rho \approx \rho_S := \sum_{i \in S} c_i \rho_i, \quad (3.6)$$

has an accuracy that is assessed as $\|\rho - \rho_S\|$ and is smaller than a prescribed tolerance value. Here, the vector $\mathbf{c} := (c_i)_{i \in S}$ is defined as the solution to the least-squares problem of Eq. (2.18), for orthogonally projecting ρ on the space $\text{Span}(\rho_i)_{i \in S}$. This allows to restrict the full space search from the admissible set \mathcal{G} to discrete subsets of indices.

For the special case of a function $\rho \in \mathbb{H}$ defined as Eq. (3.4), the ABS generation problem for bases of the form Eq. (3.5) can be formulated as follows. We can obtain an auxiliary basis set

by solving the following discrete optimization problem, of finding subsets $S \subseteq \{1, \dots, N\}$ that minimize the associated best approximation error, subject to a user-defined size constraint, i.e.

$$\min_{S \subseteq \{1, \dots, N\}} \min_{\sigma \in \text{Span}(\boldsymbol{\xi}_S)} \|\rho - \sigma\|^2. \quad (3.7)$$

$$\text{subject to } \text{card}(S) \leq M. \quad (3.8)$$

Note that the first minimization in problem (3.7) is performed over discrete sets, contrary to the continuous one in the full problem (3.2). Hence, we now look for subsets of a finite set of size N , instead of a possibly uncountable set \mathcal{G} . For these reasons, the restricted ABS problem is expected to be algorithmically easier to solve than the full one.

3.2.2 From basis generation to submatrix selection

In the present paragraph, we focus on deriving a matrix expression of the best approximation error for a given index selection. The main idea is to exploit the underlying least-squares problem of density fitting and use the normal equations of Eq. (2.20). Let \mathbf{G} denote the Gram matrix of the finite set of functions $\boldsymbol{\rho}$, defined as

$$\mathbf{G} := \langle \boldsymbol{\rho} | \boldsymbol{\rho} \rangle. \quad (3.9)$$

Note that there holds

$$\langle \rho, \rho \rangle = \sum_{i=1}^N \sum_{j=1}^N G_{ij}.$$

For any subset $S \subseteq \{1, \dots, N\}$ of M indices, we denote by $\mathbf{G}(S, S)$ the associated $M \times M$ submatrix of \mathbf{G} with indices in S . Similarly, $\mathbf{G}(S, :)$ and $\mathbf{G}(:, S)$ are submatrices of dimension $M \times N$ and $N \times M$, respectively. The best approximation error on the density ρ admits a matrix expression, in terms of submatrices of the Gram matrix, given by the following statement.

Corollary 3.2.0.1 (Error in matrix form). *Let $S \subseteq \{1, \dots, N\}$ and assume that the submatrix $\mathbf{G}(S, S)$ is invertible. The following expression holds*

$$\min_{\sigma \in \text{Span}(\boldsymbol{\xi}_S)} \|\rho - \sigma\|^2 = \mathbf{1}^\top [\mathbf{G} - \mathbf{G}(:, S)\mathbf{G}^{-1}(S, S)\mathbf{G}(S, :)] \mathbf{1}, \quad (3.10)$$

where $\mathbf{1} \in \mathbb{R}^N$ is the vector of entries $\mathbf{1}_i = 1$ for all $1 \leq i \leq N$.

Proof. See [proof](#) in Section 3.4. □

This statement connects the best approximation error to a Gram submatrix expression. Note that if S is the full set of indices, then the matrix difference on the right-hand side of Eq. (3.10) is equal to zero.

3.2.2.1 The Gram matrix row and column selection problem

We deduce from Corollary 3.2.0.1 that the restricted problem of Eq. (3.7) under the constraint of Eq. (3.8) can be equivalently written as follows:

$$\min_{S \subseteq \{1, \dots, N\}} \mathbf{1}^\top [\mathbf{G} - \mathbf{G}(:, S)\mathbf{G}^{-1}(S, S)\mathbf{G}(S, :)] \mathbf{1}, \quad (3.11)$$

$$\text{subject to } \text{card}(S) \leq M. \quad (3.12)$$

Note that this is not a matrix approximation problem. It is simply a row and column selection problem under an error functional. The matrix error functional gives the quadratic form realization for the vector consisting of ones, i.e. it is defined as the mapping

$$\forall \mathbf{A} = (A_{ij})_{1 \leq i, j \leq N} \in \mathbb{R}^{N \times N}, \quad \mathbf{A} \mapsto \mathbf{1}^\top \mathbf{A} \mathbf{1} = \sum_{i=1}^N \sum_{j=1}^N A_{ij}. \quad (3.13)$$

Let us emphasize that this mapping does not define a matrix norm. The objective error functional is less general than a matrix norm, which aims at optimizing the algebraic error of the matrix when applied to an arbitrary vector. In our case, however, we only look at the output of the matrix realization measured for a single vector, i.e. $\mathbf{1}$. As a result, matrix approximations being optimal for matrix norms may not necessarily be suitable for the error functional of Eq. (3.13). Inversely, when selecting the subset $S \subseteq \{1, \dots, N\}$, we have the freedom to commit algebraic errors as long as the error realization on the vector $\mathbf{1}$ remains small.

3.2.2.2 Related methods

We review related methods for selecting row and column indices in Gram matrices. Let $n \in \mathbb{N}^*$.

Adaptive cross approximation. Given a square matrix \mathbf{A} of dimension n -by- n , existing numerical algebra methods construct cross approximations [27, 28] of the form

$$\mathbf{A} \approx \mathbf{A}(:, C)\mathbf{A}^{-1}(R, C)\mathbf{A}(R, :), \quad (3.14)$$

where C is a subset of column indices and R is a subset of row indices, with $\text{card}(C) = \text{card}(R)$. The Adaptive Cross Approximation (ACA) [35] proposes to construct sets C and R in (3.14), yielding submatrices of approximately maximal volume, i.e. maximal determinant in absolute value. ACA serves as an alternative to truncated singular value decomposition for obtaining low-rank matrix approximations, in favour of computational complexity cost. Standard ACA algorithms [349] follow a process analogous to Gaussian elimination with complete pivoting, for greedy selection of row and column indices [223]. Note that the selection of rows and columns is not necessarily symmetric. For symmetric positive semidefinite matrices, however, it is known that the submatrix of maximum volume is attained for a principal submatrix [223, 89], therefore the constraint $R = C$ can be added.

Nyström approximation. The Nyström approximation [220] is based on performing a submatrix eigendecomposition for approximating a given matrix \mathbf{A} . It finds applications to discrete data kernel matrices, i.e. Gram matrices with (i, j) -entry equal to $K(\mathbf{x}_i, \mathbf{x}_j)$, for $\{\mathbf{x}_i\}_{1 \leq i \leq n} \subseteq \mathbb{R}^d$ data points living in dimension $d \in \mathbb{N}^*$ and $K : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ a kernel function. For this type of Gram matrix, the Nyström method yields a cross approximation of Eq. (3.14), for $S := R = C$. Existing methods for selecting the index set S are based on the so-called *landmark points*, i.e. subsets of the discrete dataset, that employ randomized or geometric criteria for point sampling [341, 350].

Pivoted Cholesky decomposition. The Pivoted Cholesky Decomposition (PCD) is a method that allows to construct low-rank approximations based on greedy row and column selection. We refer to [156, 214, 143] for more information on this method. If $\mathbf{A} \in \mathbb{R}^{n \times n}$ is a symmetric positive semi-definite matrix, Cholesky factorization with complete pivoting consists of finding a permutation matrix $\mathbf{P} \in \mathbb{R}^{n \times n}$ such that

$$\mathbf{P}^\top \mathbf{A} \mathbf{P} \approx \mathbf{L} \mathbf{L}^\top, \quad (3.15)$$

where $\mathbf{L} \in \mathbb{R}^{n \times r}$ is lower triangular with positive diagonal elements and $r \leq \text{rank}(\mathbf{A})$ a rank estimation under a given tolerance value. Note that [207] PCD yields the exact decomposition

$$\mathbf{P}_r^\top \mathbf{A} \mathbf{P}_r = \mathbf{L}_r \mathbf{L}_r^\top, \quad (3.16)$$

where \mathbf{P}_r (resp. \mathbf{L}_r) is a matrix of dimension n -by- r (resp. r -by- r) resulting from discarding the columns (resp. rows) of \mathbf{P} (resp. \mathbf{L}) indexed by $r + 1, \dots, n$. PCD determines numerical estimates r for the rank of \mathbf{A} using a stopping criterion. At the j th iteration of PCD, let $\mathbf{A}^{(j)}$ denote the computed matrix, with initialization $\mathbf{A}^{(0)} = \mathbf{A}$. Pivoting is introduced by finding the

largest possible entry from the remaining $n - j + 1$ diagonal elements of $\mathbf{A}^{(j-1)}$ and using it as the pivot in $\mathbf{A}^{(j)}$. At j th iteration, if the pivot satisfies an appropriate condition, then we set the trailing matrix to zero and the computed rank is $r = j - 1$. In particular, the algorithm we consider (see [214, Algorithm 5.1]) is stopped on the j th step if the following stopping criterion on the pivot is satisfied

$$\max_{j \leq i \leq n} A_{ii}^{(j-1)} \leq \varepsilon, \quad (3.17)$$

where ε is a prescribed zero-threshold. The standard form of pivoting is defined as the largest diagonal entry of the current trailing matrix:

$$\forall 1 \leq j \leq r, \quad s_j := \arg \max \{A_{ii}^{(j-1)} : j \leq i \leq n\}. \quad (3.18)$$

Note that this is a greedy pivot selection. In particular, the diagonal elements of the matrix \mathbf{L} are the pivot values. Due to the stopping criterion (3.17), all pivot values are greater than ε . The non-zero elements of the permutation matrix corresponding to pivot indices are such that

$$\forall 1 \leq j \leq r, \quad P_{s_j, j} = 1.$$

Truncated singular value decomposition. The Truncated Singular Value Decomposition (TSVD) [144] is a low-rank approximation method. It does not allow to select row and column indices, it can however construct *optimal* subspaces of given dimension approximating a given space. We mention TSVD as it can be used for comparison to our methods, regarding the projection accuracy using a given number of functions. TSVD consists of first calculating a full Singular Value Decomposition (SVD) for a given matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ as

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top,$$

with $\mathbf{\Sigma} = \text{diag}(\sigma_i)_{1 \leq i \leq n}$ diagonal matrix, $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n$ singular values of \mathbf{A} and \mathbf{U}, \mathbf{V} orthogonal matrices. Then, truncating $\mathbf{\Sigma}$ yields low-rank approximations of \mathbf{A} that are known to be optimal in Frobenius and spectral norms, as stated by the Eckart-Young theorem [156, 130]. TSVD can also be used to approximately solve a possibly ill-posed given linear system $\mathbf{A}\mathbf{x} = \mathbf{b}$ (e.g. ill-posed normal equations of a least-square problem). The idea is to discard singular values that are numerically zero and define the pseudoinverse of the matrix $\mathbf{\Sigma}$, for a fixed rank $r \leq n$, denoted by $\mathbf{\Sigma}^\dagger$, as

$$\forall 1 \leq i \leq n, \quad \Sigma_{ii}^\dagger := \begin{cases} \sigma_i^{-1}, & i \leq r, \\ 0, & \text{otherwise.} \end{cases} \quad (3.19)$$

An approximate solution $\mathbf{x} \approx \mathbf{x}_r$ using TSVD can then be obtained by letting

$$\mathbf{x}_r := \mathbf{V}\mathbf{\Sigma}^\dagger\mathbf{U}^\top\mathbf{b}. \quad (3.20)$$

This method can be interpreted as a regularization, where singular values of the Moore-Penrose pseudoinverse \mathbf{A}^\dagger of \mathbf{A} are truncated if they are below a threshold [154, 88]. The low-rank solution \mathbf{x}_r to $\mathbf{A}\mathbf{x} = \mathbf{b} \Leftrightarrow \mathbf{x} = \mathbf{A}^\dagger\mathbf{b}$ obtained by TSVD is optimal [50], since $\mathbf{V}\mathbf{\Sigma}^\dagger\mathbf{U}^\top$ is the best rank- r approximation of \mathbf{A}^\dagger .

3.2.3 The case of electronic density functions

We are particularly concerned by the application context of the Gaussian electrostatic model presented in Section 2.1.2.2. In this context, reference functions are electronic densities, interactions are given by intermolecular integrals and the convention is to use Gaussian-type orbital auxiliary basis sets, due to molecular integration schemes (see Section 1.1.4.4). For simplicity, we plan to construct ABSs made of orbitals centered on atoms. Our methodology then directly generalizes to more general center choices, including midbonds. The context of interest can thus be summarized as follows:

- ρ is an electronic density function discretized using a linear combination of atomic Gaussian-type orbitals, in the form of Eq. (1.26).
- ξ is an atomic Gaussian-type orbital (GTO) basis set, which is a family of discretizations defined in Section 1.1.4.

Our schemes will be proposed in this specific setting. First, we explain in more detail the important points to take into account in this application context. We begin by discussing in Section 3.2.3.1 how to define the admissible set of ABS elements, such that the cost of evaluating a single interaction, defined by Eq. (3.1), is fixed. We then explain in Section 3.2.3.2 how to derive an initial auxiliary basis set of desired form from a given electronic density function. This is a starting point for obtaining ABSs of desired form when solving the restricted ABS generation problem of Eq. (3.7) under the constraint of Eq. (3.8).

3.2.3.1 Basis set requirements

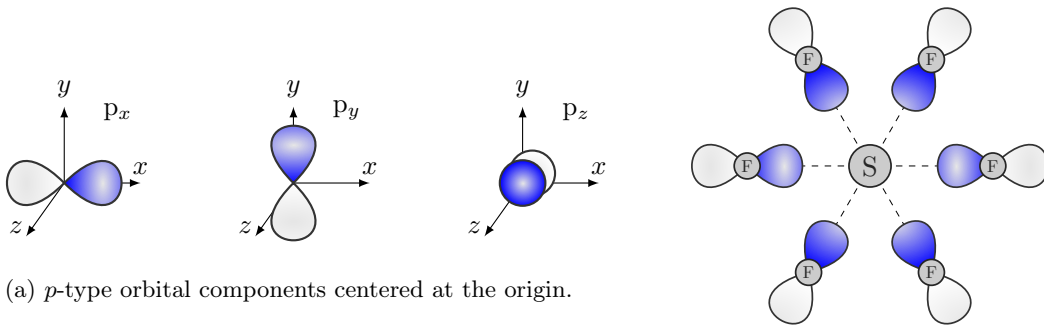
We are particularly concerned by minimizing the cost of a single pairwise interaction defined in Eq. (2.2). In this respect, we focus on interactions given as intermolecular integrals. We mention aspects relative to Gaussian-type orbitals and molecular integration. Our constructed auxiliary basis sets will respect the following properties, regarding their analytic form.

Atom-wise consistency. The entire concept of atomic Gaussian-type orbital basis sets is built upon assigning basis sets per atom type. In particular, recalling how atomic basis sets have been defined in Section 1.1.4, to every chemical type we associate a set of orbitals centered at zero. Orbital centers are assigned at a second level, given the nuclear coordinates of an atom in a molecule. Assigning an atomic center to an orbital χ_μ , centered at the origin, is performed by translation in the Euclidean space, i.e. by translating the origin to a given coordinate $\mathbf{R} \in \mathbb{R}^3$ as $\chi_\mu(\cdot - \mathbf{R})$. All atoms of the same chemical type, independently of their position, will have the same orbital types and exponents, which is a property that we refer to as *atom-wise consistency*. This property assures compatibility with available formats in data bases, as well as efficient storage and processing of orbital bases in quantum chemistry codes.

Symmetry. A property closely related to atom-wise consistency is orbital symmetry with respect to the atomic conformation. The starting point is the angular-radial decomposition of a Gaussian-type orbital, given by Eq. (1.21). Essentially the angular part, given by the spherical harmonic, is independent of the atomic position and acquires spherical symmetries. Figure 3.1a shows the angular part of p -type orbitals ($\ell = 1$). Notice the rotational symmetry of orbital components with respect to the real axes. The angular part is expected to point to the appropriate direction, according to neighbour atoms and bond directions. This can be achieved by rotating or reflecting orbitals. Figure 3.1b shows an example of six fluorine atoms in the SF_6 molecule, for which atomic orbitals of p -type have six different axes of symmetry, each one obtained by rotation around the sulfur atom. Yet, the reason why a p -type orbital family is actually rotated "naturally", i.e. by taking linear combinations of atomic orbitals, is the following. Consider a rotation $\mathbf{R} \in \mathbb{R}^{2 \times 2}$ around the origin in the Euclidean space, sending the unit vector $\mathbf{r} \in S^2$ to $\mathbf{r}' \in S^2$. The following decomposition holds [113, Equation 4.1.4]:

$$Y_\ell^m(\mathbf{r}') = \sum_{m'=-\ell}^{\ell} b_{mm'} Y_\ell^{m'}(\mathbf{r}), \quad (3.21)$$

where $b_{mm'}$ is a coefficient that depends on the rotation \mathbf{R} , admitting an expression in terms of Wigner D-matrices [292]. Moreover, for a given degree ℓ , all $m = -\ell, \dots, \ell$ components are linearly independent. In other words, rotated spherical harmonics write as an expansion of $2\ell + 1$ spherical harmonics of the same degree. Lastly, spherical harmonics remain the same up to a sign under the action of reflections. For these reasons, the set of p -type orbitals in Figure 3.1b,



(a) p -type orbital components centered at the origin.

(b) Translated and rotated p -type orbitals.

Figure 3.1: Sulfur hexafluoride, with formula SF_6 , is a compound with six fluorine atoms symmetrically arranged around a central sulfur atom. Sulfur shares two electrons with each fluorine. Orbitals of sulfur have been omitted.

translated to coordinates of fluorine atoms, can capture all bond directions when used as a discretization basis for calculating the electronic ground state. To summarize, an orbital basis can be rotated in any direction or reflected only if all components of a given angular degree ℓ are included.

Orbital type. The orbital type affects the computational cost of direct molecular integral computations. In Cartesian coordinates, for instance, the higher the orbital type, the deeper the recurrence relations used in the McMurchie-Davidson scheme for Gaussian-type orbitals [228]. These methods are based on Hermite-Gaussian intermediary conversion of GTOs. In this case, we should also take into account the cost of converting an auxiliary basis function from Gaussian to Hermite-Gaussian form. This cost depends on the angular momentum of the Gaussian. From the McMurchie-Davidson recursion, the number of Hermite coefficients, denoted by $H_{tuv}^n \in \mathbb{R}$, we need to generate given a GTO of angular momentum ℓ in spherical form, must be such that $0 \leq n + t + u + v \leq 2\ell + 1$. This number can be explicitly computed as [288]

$$\sum_{j=0}^{2\ell+1} \binom{j+3}{3} = \binom{(2\ell+1)+4}{4}.$$

From this point of view, spd orbitals, used in the Gaussian electrostatic model [84] are optimal since thanks to the inclusion of all lower types of the same orbital exponent, the McMurchie-Davidson recursion is greatly simplified, helping to reduce the time spent on direct integration (see Appendix C.2.0.3 for more details). For these reasons, a reasonable requirements is to impose a given maximal orbital degree on the ABS, so as to narrow down these recursions.

3.2.3.2 Decomposing electronic density functions

We propose a method to define an initial set of admissible auxiliary basis functions, given an electronic density function ρ , discretized over a set of atomic Gaussian-type orbitals $\{\chi_\mu\}_{1 \leq \mu \leq N_b}$, as in Eq. (1.26), that we recall here:

$$\forall \mathbf{r} \in \mathbb{R}^3, \quad \rho(\mathbf{r}) := \sum_{\mu=1}^{N_b} \sum_{\nu=1}^{N_b} D_{\mu\nu} \chi_\mu(\mathbf{r}) \chi_\nu(\mathbf{r}), \quad (3.22)$$

with $\mathbf{D} = \mathbf{C}\mathbf{C}^\top$ and $\mathbf{C}^\top\mathbf{S}\mathbf{C} = \mathbf{I}$. This electronic density is a special case of the general finite sum decomposition of Eq. (3.4), reading

$$\rho = \sum_{i=1}^N \rho_i, \quad (3.23)$$

for $N := \binom{N_b}{2} + N_b$ the number of orbital pairs up to permutational symmetry, i.e. using $\chi_\mu\chi_\nu = \chi_\nu\chi_\mu$, and $\boldsymbol{\rho} := \{\rho_i\}_{1 \leq i \leq N}$ the family of functions defined as

$$\forall 1 \leq i \leq N, \quad \rho_i = 2^{1-\delta_{\mu\nu}} D_{\mu\nu} \chi_\mu \chi_\nu, \quad (3.24)$$

with $(\mu, \nu) := \pi(i)$ for π an enumeration of indices, i.e. a bijective map between index sets

$$\pi : \{1, \dots, N\} \rightarrow \{(\mu, \nu) : 1 \leq \mu \leq \nu \leq N_b\}. \quad (3.25)$$

Note that any element of the family $\boldsymbol{\rho}$ is a product of atomic orbitals, multiplied by some scaling coefficient. The scaling depends on the density matrix, it also treats the permutational symmetry of pairwise products. The finite sum decomposition of the electronic density given by Eq. (3.23) allows to apply the restricted search strategy and look for auxiliary basis sets as defined in Eq. (3.5).

Atomic and non-atomic products. Orbital products can be used as auxiliary basis [324]. Indeed, among orbital products, atomic orbitals can be extracted directly from the definition of Eq. (3.22), as follows. The electronic density admits a decomposition consisting of two parts:

$$\rho = \left(\sum_{\substack{\mu, \nu=1 \\ \mathbf{r}_\mu = \mathbf{r}_\nu}}^{N_b} D_{\mu\nu} \chi_\mu \chi_\nu \right) + \left(\sum_{\substack{\mu, \nu=1 \\ \mathbf{r}_\mu \neq \mathbf{r}_\nu}}^{N_b} D_{\mu\nu} \chi_\mu \chi_\nu \right), \quad (3.26)$$

where $\mathbf{r}_\mu \in \mathbb{R}^3$ denotes the atomic center of the orbital χ_μ . The first part contains orbital products centered on the same atom, while in the second one, product components are centered on different atoms. The Gaussian product theorem (see Theorem C.1.1) states that a GTO product can be written as a Gaussian overlap distribution with a composite center. If $\mathbf{r}_\mu = \mathbf{r}_\nu$, we find that the product $\chi_\mu\chi_\nu$ is an atomic GTO, centered on the common center. On the other hand, if $\mathbf{r}_\mu \neq \mathbf{r}_\nu$, then the product $\chi_\mu\chi_\nu$ can be written in terms of a Gaussian, whose center lies on the segment joining the atomic centers \mathbf{r}_μ and \mathbf{r}_ν , times a polynomial that is not a monomial, however. Such products, not respecting the GTO definition, are thus non-admissible ABS elements. This simple observation shows that admissible ABS elements can be found among orbital products, precisely in the first part of the decomposition of Eq. (3.26).

Picking atomic orbital products in order to reach a target ABS size lies at the core of our methodology. Atom-centered GTOs in the finite set $\boldsymbol{\rho}$, with elements defined as in Eq. (3.24), are precisely indexed by the set

$$I_{\text{atm}} := \{\pi^{-1}(\mu, \nu) : \mathbf{r}_\mu = \mathbf{r}_\nu, \quad 1 \leq \mu \leq \nu \leq N_b\}, \quad (3.27)$$

with π an enumeration defined as in Eq. (3.25). Note that elements in I_{atm} are not necessarily consecutive indices, as we made an index selection that essentially discards off-atom products. We propose to discard the second part of the decomposition of Eq. (3.26) altogether and narrow down the ABS search by looking for subsets $S \subseteq I_{\text{atm}}$, that solve problem of Eq. (3.7) under the constraint of Eq. (3.8). We use this idea as a starting point for building ABS generation schemes, presented in the following.

3.2.4 Auxiliary basis generation schemes

The selection of the index set $S \subseteq I_{\text{atm}}$, yielding the solution to the restricted ABS generation problem of Eq. (3.7) under the constraint of Eq. (3.8), respecting the additional basis set requirements of Section 3.2.3.1, can be carried out using the scheme proposed in the present section.

Our method consists of a series of matrix operations that allows to obtain auxiliary basis sets systematically. The calculation of the auxiliary functions hinges upon the selection of GTO products in the decomposition of Eq. (3.26), for a prescribed level of accuracy on the error of Eq. (3.10). This can be translated into a selection procedure in matrix form thanks to our method. The requirements on the auxiliary basis form, previously presented in Section 3.2.3.1, first lead us to a new method for defining a suitable type of Gram matrix. Then, we apply a greedy procedure for selecting Gram matrix rows and columns based on the pivoted Cholesky decomposition. Before describing our algorithm, we first provide a quick overview of its main steps:

1. Initialization of a matrix, denoted by \mathbf{A} , representing the overlap of admissible orbitals under a given inner product. It is obtained by block treatment of an appropriate matrix built from a given electronic density function.
2. Row (and column) selection for the matrix \mathbf{A} .
3. Post-treatment of selected orbitals.

The series of preparatory operations of Step 1, allowing to assemble the matrix on which we perform row and column selection at the later stage of Section 3.2.4.2, is first described in Section 3.2.4.1. The goal of this preparation, as well as of the post-processing in Step 3, is to assure that selected functions will respect a desired analytical form. At the end of Section 3.2.4.3, the result of our algorithm is an auxiliary basis made up of primitive atomic Gaussian-type orbitals, respecting the atom-wise consistency and the orbital type symmetry requirements.

Note that the use of pivoted Cholesky decomposition for constructing auxiliary basis sets is not a new idea [251]. However, compared to existing methods, the originality of our method is that, first, the Gram matrix on which we apply this selection is systematically defined from a given electronic density and second, we use the density matrix that is available due to the specific context of the Gaussian electrostatic model (see Section 2.1.2.2). Overall, we avoid empirical rules in the creation of admissible auxiliary basis sets, such as the high angular momentum pruning [199] and the effective Gaussian exponent [196].

3.2.4.1 Block treatment of Gram matrix

Let us consider as input an AO basis set $\{\chi_\mu\}_{1 \leq \mu \leq N_b}$ and a density matrix $\mathbf{D} \in \mathbb{R}^{N_b \times N_b}$. The following treatment allows to initialize a Gram matrix containing orbital product data in admissible form. Before detailing the steps, let us briefly state that the first step weights the AO Gram matrix by the density matrix and expands the possibly contracted AO basis to primitives. The second and third steps perform post-processing of matrix blocks based on angular degrees and atomic centers. The goal is to end up with a matrix for which any selection of rows and columns yields admissible auxiliary basis functions of desired form.

Step I - Basis decontraction and precondition. If the given AO basis consists of contracted GTOs, as it is commonly the case, then we propose to additionally expand the contractions before proceeding. This allows to gain systematic control over primitives. Note that most conventional auxiliary basis sets existing in the literature are made up of primitive orbitals. In the present step, we illustrate how to decontract the product basis set in matrix form. Simultaneously, we weight primitives using the density matrix, initially given on the contracted basis. Note that such step assumes that the density matrix has been evaluated by solving the ground state problem of

Eq. (1.17). This is the case for example in the Gaussian electrostatic model (see Section 2.1.2.2). Otherwise, if this is not the case and the density matrix is not available, weighting by the density matrix should be omitted altogether.

Let N_b denote the contracted AO basis size and N_p the primitive (uncontracted) one. There holds $N_b \leq N_p$, since the primitive basis consists of individual components in AO contractions, obtained by decontracting summations of the form of Eq. (1.19). For example, the contracted basis set $\{2s, p\}$ uncontracts to $\{s, s, p\}$. We can find a change of basis matrix, denoted by $\mathbf{B} \in \mathbb{R}^{N_p \times N_b}$, that allows to pass from the contracted basis, made up of functions χ_μ 's, to the primitive basis. The Gram tensor of weighted primitive AO pairwise products, denoted by $\widehat{\mathbf{G}} \in \mathbb{R}^{N_p \times N_p \times N_p \times N_p}$, has entries defined as

$$\widehat{G}_{pqrs} := \sum_{\mu, \nu, \kappa, \lambda=1}^{N_b} B_{p\mu} B_{q\nu} D_{\mu\nu} \langle \chi_\mu \chi_\nu, \chi_\kappa \chi_\lambda \rangle D_{\kappa\lambda} B_{r\kappa} B_{s\lambda}. \quad (3.28)$$

Every dimension of $\widehat{\mathbf{G}}$ corresponds to a set of primitive GTOs.

Step II - Orbital type block contraction. This step concerns the block contraction of the primitive Gram tensor defined in Eq. (3.28), for blocks defined by orbital types.

To begin with, we consider that any two given AOs of the same atomic center and orbital exponent belong to the same orbital type if they have the same angular degree. For Cartesian GTOs of Eq. (1.18), components of an angular degree ℓ are obtained by varying $(l, m, n) \in \mathbb{N}^*$ such that $l + m + n = \ell$. For spherical GTOs of Eq. (1.21), this is achieved by varying the angular order $-\ell \leq m \leq \ell$. Recalling Section 1.1.4.2, the component number for a given angular number ℓ is $(\ell + 1)(\ell + 2)/2$ for the Cartesian form and $2\ell + 1$ for the spherical one. By convention, any given atomic orbital basis contains all components of a given orbital type. We want to ensure that the same property holds for auxiliary basis sets. This allows to achieve the desired rotational symmetry, as given by Eq. (3.21).

The idea is to contract atomic orbitals belonging to the same orbital type at this stage. This ensures that all components of an orbital type are treated as inseparable blocks during the selection procedure that comes later on. A block contraction, applied to the tensor $\widehat{\mathbf{G}}$, that allows to achieve this goal can be introduced as follows. Let $\mathcal{B} = \{B_1, \dots, B_{N_p^*}\}$ be a partition of the set of indices $\{1, \dots, N_p\}$, representing a grouping of the family of N_p primitive AOs into blocks of fixed total number $N_p^* \in \mathbb{N}^*$ with $N_p^* \leq N_p$. In our case, each block is defined as a set of primitive AOs belonging to the same orbital type. For example, if the given orbital basis in primitive form is $\{s, p_x, p_y, p_z\}$ centered on an atom, then the same basis regrouped by orbital type is $\{s, p\}$. The set of primitive indices is $\{1, 2, 3, 4\}$ and the orbital group partition is $\mathcal{B} = \{B_1, B_2\}$ with $B_1 = \{1\}$ and $B_2 = \{2, 3, 4\}$.

Define the contraction matrix, denoted by \mathbf{M} , that contracts blocks of primitive orbitals belonging to the same partition, as the N_p^* -by- N_p matrix with entries

$$\forall 1 \leq i \leq N_p^*, \quad M_{ij} = 1, \quad \forall j \in B_i,$$

and zero otherwise. In practice, this matrix \mathbf{M} is a block diagonal matrix, with each i -block being a row matrix of dimension 1-by- $\text{card}(B_i)$. Applying this block contraction to the tensor $\widehat{\mathbf{G}}$ yields a new tensor, denoted by $\mathbf{T} \in \mathbb{R}^{N_p^* \times N_p^* \times N_p^* \times N_p^*}$, with entries defined as

$$T_{ijkl} := \sum_{p, q, r, s=1}^{N_b} M_{ip} M_{jq} \widehat{G}_{pqrs} M_{kr} M_{ls}. \quad (3.29)$$

Indeed, by considering each partition in \mathcal{B} to be an orbital type, we manage to consider the corresponding primitive atomic orbitals as inseparable blocks, as shown in the example of Figure 3.2a. This example concerns the simpler but similar case of Gram matrix of a given atomic orbital basis, with two-center integrals for simplicity in the representation. For the case of four-center integrals, which is the one of interest appearing in the tensor \mathbf{T} , one needs to add two extra dimensions to this example.

	atom 1			atom 2		
atom 1	ss	ss	sp	ss	ss	sp
	ss	ss	sp	ss	ss	sp
atom 2	ps	ps	pp	ps	ps	pp
	ps	ps	pp	ps	ps	pp

(a)

	atom 1	off-atom	atom 2
atom 1	11, 11	11, 12	11, 21
	11, 22	11, 12	11, 21
off-atom	12, 11	12, 12	12, 21
	21, 11	21, 12	21, 21
atom 2	22, 11	22, 12	22, 21
	22, 22	22, 12	22, 21

(b)

Figure 3.2: Left: Block structure of Gram matrix of atomic orbital basis for a diatomic system with identical orbitals {s, s, p} centered on each atom. Each block within panel (a) contains orbitals of the same angular type. Right: Block structure of Gram matrix of atomic orbital products, with entries of the form $\langle \chi_\mu \chi_\nu, \chi_\kappa \chi_\lambda \rangle$ (with pairwise folded dimensions) for a given atomic orbital basis $\{\chi_\mu\}_\mu$. The block label corresponds to the four atoms on which each of the four orbitals is centered. Blocks coloured green are atomic blocks, the remaining ones are discarded.

Step III - Atomic block partition. The operation described in the present paragraph concerns discarding non-atomic orbital products. To achieve this, we first fold pairwise the dimensions of the tensor \mathbf{T} defined in Eq. (3.29), in order to recover orbital products. Let $\pi : a \mapsto (i, j)$ be an enumeration, i.e. a bijective function mapping indices $1 \leq a \leq N_p^{*2}$ to pairs of indices $1 \leq i, j \leq N_p^*$. Then we introduce the N_p^{*2} -by- N_p^{*2} matrix \mathbf{G} with entries defined as

$$G_{ab} = T_{ijkl}, \quad (i, j) := \pi(a), \quad (k, l) := \pi(b). \quad (3.30)$$

Every row (and column) of \mathbf{G} corresponds to a product between two orbital types. The full set of product indices $\{1, \dots, N_p^{*2}\}$ can be partitioned into blocks according to the orbital center, that is atomic or off-atomic, of the associated product. The atomic center associated to a primitive belonging to a block B_i in the partition \mathcal{B} should be denoted by \mathbf{r}_i . Given the set of atomic product indices, belonging to

$$I_{\text{atm}}^* := \{a_1, \dots, a_{N_c}\} = \{\pi^{-1}(i, j) : \mathbf{r}_i = \mathbf{r}_j, \quad 1 \leq i \leq j \leq N_p^*\} \quad (3.31)$$

of total number $N_c \leq N_p^{*2}$, we may associate a restriction matrix, denoted by \mathbf{R} , defined as the matrix of dimension N_c -by- N_p^{*2} with non-zero entries

$$\forall 1 \leq i \leq N_c, \quad R_{i, a_i} = 1. \quad (3.32)$$

Note that the above definition of I_{atm}^* in Eq. (3.31) is an alternative to Eq. (3.27), that consists of further taking into account orbital type symmetry. The aim of the restriction matrix \mathbf{R} is to discard blocks centered on midbonds. An example of such restriction is shown in Figure 3.2b. The Gram matrix of N_c atom-centered products can then be recovered as

$$\mathbf{A} := \mathbf{RGR}^\top. \quad (3.33)$$

Every row (or column) of \mathbf{A} corresponds to a group of primitives belonging to a given orbital type, centered on an atom.

3.2.4.2 A new selection procedure

After the initialization and preliminary block treatment of the Gram matrix of orbital products, resulting in the final matrix \mathbf{A} , defined in Eq. (3.33) the last step is a symmetric row and column selection that allows to give rise to auxiliary basis sets in a systematic way. Here, we present our method to construct the index set $S \subseteq \{1, \dots, N_c\}$ for solving the minimization problem of Eq. (3.11) under the constraint of Eq. (3.12).

It is crucial to select products in a way such that the spanned space accurately reproduces the full space of products, under the error functional of Eq. (3.10). To this end, our main idea lies in exploiting that the full set of GTO products suffers from strong linear dependencies, as numerically observed in the literature [62, 196, 198]. Due to this property, the underlying Gram matrix can have a low-rank structure, that we exploit to extract auxiliary basis subsets. Note that such methods also apply to any kernel matrix, measuring the similarity between pairs of discrete or continuous data. Just as in the discrete data case [154, 77], in the continuous case we have that if a given basis set ξ contains redundancy in the form of linear dependence, then the Gram matrix $\langle \xi | \xi \rangle$ is ill-conditioned.

Remark. *Let us make a point on auxiliary basis sizes for the sake of clarity. The number N_c of rows and columns of \mathbf{A} counts the number of admissible auxiliary functions regrouped by orbital type. Given a target $k \in \mathbb{N}^*$, our algorithm provides a selection of k rows (and columns). The point is that k is smaller or equal to the total number of individual components of selected orbitals, up to now denoted M , obtained by expanding orbital types. For example, a selected p-type orbital needs to be expanded in individual components p_x, p_y, p_z to obtain the auxiliary basis of size M . Hence, we shall keep in mind that our proposed algorithm directly controls only the number k .*

Step IV - Greedy row and column selection. Our auxiliary basis generation is based on a greedy procedure, for symmetrically selecting row and column indices, i.e. submatrices, from the matrix \mathbf{A} defined in Eq. (3.33). This will allow us to construct subsets of the index set $\{1, \dots, N_c\}$ of target size. Our idea is to use low-rank approximations of \mathbf{A} . Algebraic motivation lies in numerical evidence demonstrating the low-rank structure of the electron repulsion integral matrix [31], that is the Gram matrix of orbital products for the Coulomb kernel, as well as of the overlap matrix, see Section 1.1.4.3, related to the Dirac interaction kernel.

Our method to construct index sets S is based on the pivoted Cholesky decomposition. Employing notation from Section 3.2.4, let us fix a tolerance value $\varepsilon > 0$ and compute PCD for the matrix \mathbf{A} , obtaining the decomposition of Eq. (3.15), with estimated rank denoted by $r \in \mathbb{N}^*$. Given a target number $M_* \in \mathbb{N}^*$ of auxiliary functions grouped by orbital type and assuming that $M_* \leq r$, our selection procedure consists of defining the index set S , solution to the restricted ABS generation problem, as the set consisting of the M_* -first pivots obtained by PCD, denoted by s_i and defined in Eq. (3.18), i.e.

$$S := S_{M_*} = \{s_1, \dots, s_{M_*}\} \subseteq \{1, \dots, N_c\}.$$

Recall that indices s_i 's are obtained by looking at M_* -first rows and columns of the PCD matrix \mathbf{L} while inverting the index permutation \mathbf{P} to recover initial rows and columns of \mathbf{A} . This way, the index set S_{M_*} is a collection of row indices of \mathbf{A} sorted in pivot decreasing order, according to the stopping criterion in Eq. (3.17). Hence, the set S_{M_*} can be interpreted as the set containing the most linearly independent set of M_* rows and columns of \mathbf{A} . Such selection then allows to recover auxiliary functions by looking at partitions \mathcal{B} of orbital types, defined during the Step II. In particular, the retained auxiliary function set, solution to the restricted ABS problem, is the set of primitive atomic orbital products, defined, for any fixed $1 \leq M_* \leq r$, as

$$\xi_{M_*} := \{\widehat{\chi}_p \widehat{\chi}_q : (p, q) \in B_i \times B_j, \quad (i, j) = \pi(a_s), \quad s \in S_{M_*}\}, \quad (3.34)$$

where $\{\widehat{\chi}_p\}_{1 \leq p \leq N_p}$ denotes the primitive AO basis, obtained from the given (possibly contracted) AO basis $\{\chi_\mu\}_{1 \leq \mu \leq N_b}$. Note that, at this stage, the recovered basis contains primitives that respect the orbital type symmetry. Finally, $M_0 := \text{card}(\xi_{M_*})$. Note that $M_0 \geq M_*$.

3.2.4.3 Post-treatment

A post-processing step must be taken into account, described as follows. For physical reasons we should consider the same auxiliary basis for all atoms of the same chemical type, a property described as atom-wise consistency in previous Section 3.2.3.1. A result of this requirement is the following. If during the selection procedure of Step IV, a function is discarded on one atom but included on another atom with the same chemical type, then this function will be included back again in the final auxiliary basis to make the basis set atom-wise consistent.

This procedure can be formally described as follows. Let us consider a molecule of N_a atoms, with atomic positions $\{\mathbf{R}_i\}_{1 \leq i \leq N_a}$ and atomic charges $\{z_i\}_{1 \leq i \leq N_a}$. Let $\boldsymbol{\xi}_{M_*}$ denote the ABS obtained by the Step IV of our algorithm applied to the given molecule, of the form of Eq. (3.34). It is assumed that the ABS is of Cartesian GTO form and contains all components of a given orbital type defined as follows. The set of Cartesian GTOs of the same orbital type is defined, given an orbital exponent α , center $\mathbf{R} \in \mathbb{R}^3$ and angular degree ℓ , as

$$\mathcal{G}[\alpha, \mathbf{R}, \ell] := \{G_{lmn}(\mathbf{r} - \mathbf{R}, \alpha) : (l, m, n) \in \mathbb{N}^3, \quad l + m + n = \ell\}, \quad (3.35)$$

where G_{lmn} is a GTO in Cartesian form, centered at the origin, defined by Eq. (1.18), of angular degree $l + m + n$. Now, we may regroup orbital types belonging to $\boldsymbol{\xi}_{M_*}$ by atomic centers. For every nucleus, uniquely defined by its position \mathbf{R}_i and atomic type z_i , we define the ABS associated to this nucleus, but centered at an arbitrary point, as the set of size N_i , with elements the orbital types

$$\forall \mathbf{R} \in \mathbb{R}^3, \quad \boldsymbol{\xi}_i[\mathbf{R}] := \{\mathcal{G}[\alpha_j^i, \mathbf{R}, \ell_j^i] : 1 \leq j \leq N_i\},$$

where α_j^i denotes the orbital exponent of the j -orbital type centered on the i -atom obtained from the given ABS $\boldsymbol{\xi}_{M_*}$, similarly for ℓ_j^i . The above definition allows to center a GTO basis on arbitrary point \mathbf{R} .

The final ABSs will be assigned by chemical type. In this respect, atomic charges can be used to identify elements in the periodic tables, e.g. $z = 1$ corresponds to H, $z = 2$ to He, etc. The final ABS can be obtained by taking the union of for all chemical types.

$$\boldsymbol{\xi} := \bigcup_{i=1}^{N_a} \bigcup_{\substack{j=1 \\ z_j=z_i}}^{N_a} \boldsymbol{\xi}_j[\mathbf{R}_i].$$

A small example of dihydrogen molecule. Suppose that the ABS for the first atom is $\mathcal{G}(\alpha_1, \mathbf{R}_1, \ell_1)$ and for the second atom is $\mathcal{G}(\alpha_2, \mathbf{R}_2, \ell_2)$, with $\alpha_1 \neq \alpha_2$ and $\ell_1 \neq \ell_2$. Since $z_1 = z_2$, the final ABS of the molecule will be $\{\mathcal{G}(\alpha_1, \mathbf{R}_1, \ell_1), \mathcal{G}(\alpha_2, \mathbf{R}_1, \ell_2), \mathcal{G}(\alpha_1, \mathbf{R}_2, \ell_1), \mathcal{G}(\alpha_2, \mathbf{R}_2, \ell_2)\}$. Notice that in the final ABS, both nuclei are assigned to the same basis function parameters, i.e. $\{(\alpha_1, \ell_1), (\alpha_2, \ell_2)\}$. Lastly, note that the final number of auxiliary basis functions, denoted by $M := \text{card}(\boldsymbol{\xi})$, can be (slightly) larger than $M_0 = \text{card}(\boldsymbol{\xi}_{M_*})$ constructed during the previous Step IV. Recall that our method has direct control only on M_* .

3.2.5 Summary of Auxiliary Basis Set (ABS) generation scheme

The present section summarizes the overall procedure for generating auxiliary basis sets for density fitting using our method.

Algorithm 1 summarizes the main result of the present work, which as a new ABS generation scheme. Recall that, our algorithm essentially controls the number of auxiliary functions regrouped by orbital type, denoted by M_* . The auxiliary basis set generated with our method will be denoted by PCD- M_* and corresponds to an ABS of rank M_* , containing the best M_* atomic orbital products from a given primitive atomic orbital basis set, grouped by orbital type. From M_* given functions grouped by orbital type, one then obtains the M auxiliary functions by counting individual components included in each orbital type. The M selected functions respect orbital symmetry and atom-wise consistency.

Lastly, let us make a comment on angular degrees and introduce a related modification to our method. In previous Section 3.2.3.1, we saw that high orbital types are related to less cost-efficient molecular integral computations. At the same time, the angular degree of product bases is expected to be higher than the one of standard orbital bases. In order to lower the angular degree within our method, we may include an additional step during the block-treatment of the Gram matrix, consisting of discarding rows and columns corresponding to orbitals with higher type than the desired one. Alternatively, one could manually modify the orbital type of products, to lower it if needed. This idea could allow to create for example bases respecting a given maximal angular degree.

Algorithm 1: Auxiliary Basis Set (ABS) generation yielding the ABS PCD- M_* .

- Input: AO basis set $\{\chi_\mu\}_{1 \leq \mu \leq N_b}$, $\mathbf{D} \in \mathbb{R}^{N_b \times N_b}$, $M_* \in \mathbb{N}^*$, inner product $\langle \cdot, \cdot \rangle$
Output: AO ABS of target size M_* (by orbital type) optimized for best approximations of the density $\rho = \sum_{\mu, \nu=1}^{N_b} D_{\mu\nu} \chi_\mu \chi_\nu$ under the norm $\|\cdot\| := \langle \cdot, \cdot \rangle^{1/2}$
- 1: Compute uncontraction matrix $\mathbf{B} \in \mathbb{R}^{N_p \times N_b}$, with N_p the number of AO primitives.
 - 2: Weight by \mathbf{D} and uncontract by \mathbf{B} the AO product four-index tensor with entries $\langle \chi_\mu \chi_\nu, \chi_\kappa \chi_\lambda \rangle$, to obtain the weighted primitive four-index tensor $\widehat{\mathbf{G}}$.
 - 3: Assembly the orbital type block contraction matrix \mathbf{M} and contract orbital types in each dimension of $\widehat{\mathbf{G}}$ to form the tensor \mathbf{T} .
 - 4: Fold pairwise dimensions of \mathbf{T} to form the overlap matrix \mathbf{G} of products.
 - 5: Assembly the atomic restriction matrix \mathbf{R} and discard off-atom orbital products to form the matrix $\mathbf{A} = \mathbf{RGR}^\top$.
 - 6: Perform PCD to \mathbf{A} with tolerance value $\varepsilon = 10^{-14}$ to obtain r -rank estimation and pivot indices sorted in pivot decreasing order, denoted by $\{s_i: 1 \leq i \leq r\}$.
 - 7: If $M_* > r$ then **return** error.
 - 8: Recover AOs associated to rows indexed by $\{s_i: 1 \leq i \leq M_*\}$ in \mathbf{A} .
 - 9: Post-process recovered AOs to force atom-wise consistency.
 - 10: **return** ABS.

3.3 Numerical results

In the present section, we show how our auxiliary basis generation method numerically compares to conventional ones. All results presented in this section are generated using our `dfwpy` code in Python (see Appendix A for implementation details). Target molecular systems used for testing are rigid water fragments.

Before discussing results, let us quickly mention the methods under comparison. In particular, our generated ABSs obtained by Algorithm 1 are compared to existing empirical ones, found in the literature, namely the RI [337] and JKFIT [335] auxiliary basis sets. Table 3.1 shows the sizes of state-of-the-art auxiliary basis sets, counted by orbital type and by individual components. Recall that the number of individual components in spherical form is obtained by counting all orders $-\ell \leq m \leq \ell$ for a given ℓ . Results show that the auxiliary size M is roughly between $2N_b$ and $5N_b$, where N_b is the AO basis size. Note that all compared auxiliary bases are made up of primitive orbitals. The aim of our numerical study, presented in this section, is to see if our method can achieve fewer auxiliary basis functions than state-of-the-art methods with low accuracy loss on frozen core energies.

Lastly, let us mention that all *ab initio* density matrices of water fragments used in experiments are obtained by solving a Restricted Kohn-Sham problem using the B3LYP hybrid exchange-correlation functional [303], that breaks down to the use of a self-consistent field algorithm for finding the ground state solution [200].

basis	AO		RI		JKFIT	
	N_b^*	N_b	M_*	M	M_*	M
aug-cc-pvdz	18	41	40	118	52	150
aug-cc-pvtz	31	92	56	198	58	196

Table 3.1: H₂O molecule. Spherical basis size given by orbital type (starred) and by individual components (N_b and M), for AO bases and state-of-the-art auxiliary bases.

3.3.1 Frozen core energy of rigid fragments

The performance of methods is assessed in terms of frozen core energy errors as well as of timings, throughout this section. We use three test systems, for evaluating pairwise rigid fragment interactions: (i) ten standardized Smith dimers, (ii) a small water cluster of 27 molecules, (iii) large-scale water clusters up to 10^4 molecules. For a summary of test cases, we refer to Table 1.1. Let us recall that in all such systems, all fragments are identical. This means that the frozen library contains a single water fragment, whose *ab initio* density is computed offline.

Frozen core energies are electron-electron (denoted by EE), whose exact formula is defined in Eq. (1.39) and its approximation in Eq. (1.42), and exchange-repulsion (denoted by XR), whose exact formula is defined in Eq. (1.40) and its approximation in Eq. (1.43). For each energy term, we used a kernel explicit metric to perform the density fit on a single water fragment belonging to the frozen library. Namely, the approximate density for the EE energy has been fitted using the Coulomb metric while for the XR energy using the Dirac one. Recall from Section 2.1.2.3 that these metrics are optimal for their respective choices of energy terms. The frozen core energies computed from auxiliary bases using density fitting are given by the formula of Theorem 2.3.4.

Frozen core energy errors are assessed using the following metric. Consider a cluster of $N_f \in \mathbb{N}^*$ fragments, denoted by A_1, A_2, \dots, A_{N_f} . Let j be an index used to enumerate a given fragment pair (A_{i_1}, A_{i_2}) with $1 \leq i_1 \neq i_2 \leq N_f$, $1 \leq j \leq n$ and $n := \binom{N_f}{2}$ the number of pairs of distinct fragments. If $E_j \in \mathbb{R}$ is the intermolecular interaction energy of the j -th fragment pair computed using reference *ab initio* fragment densities and $\tilde{E}_j \in \mathbb{R}$ the same energy term computed using approximated densities, then the total error of the cluster is measured as the root mean square error (RMSE) of all errors of its fragment pairs, defined as

$$\text{RMSE} = \sqrt{\sum_{j=1}^n \frac{(E_j - \tilde{E}_j)^2}{n}}. \quad (3.36)$$

Note that this is an absolute error measure, used throughout the present work.

3.3.1.1 Ten Smith dimers

We first assess our auxiliary basis set performance for the ten standardized Smith dimers. More information on the exact energies used as reference measured on this cluster can be found in Table 3.2. Note that, results show that the order of magnitude of the exchange-repulsion is very small and is generally comparable to the target chemical accuracy, that is equal to 10^{-3} Hartree. For this reason, relative errors would not be suitable for assessing the density fit accuracy on energies. Taking this consideration into account, we prefer the use of the absolute root mean square error of Eq. (3.36), that does not have this issue.

Figure 3.3 shows the absolute root mean square errors on frozen core energies for ten Smith dimers, using optimal choices of kernel explicit metrics to perform the density fit. Obtained results allow to compare our auxiliary basis generation method (PCD- M_*) with the state-of-the-art ones (JKFIT and RI). Results on XR energy demonstrate that our method can perform roughly the same or better than conventional auxiliary basis sets in terms of accuracy, while using less auxiliary functions. In particular, the maximal size gain due to our method is the reduction of

energy	min.	max.	mean	std
EE	16.002	19.111	18.133	0.926
XR	23E-4	16E-3	83E-4	47E-4

Table 3.2: Statistics of reference frozen core interaction energies (in Hartree) of ten standardized Smith dimers, using the AO basis set aug-cc-pvdz of H₂O.

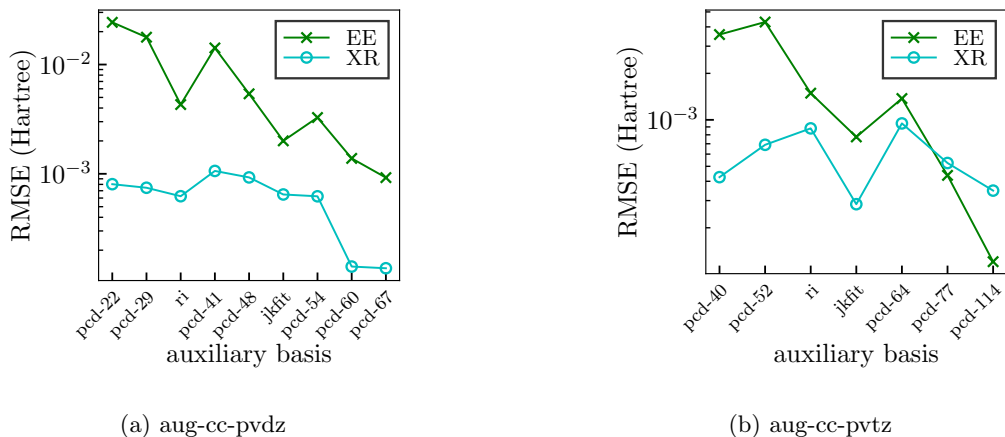


Figure 3.3: Ten standardized Smith dimers. Errors on exchange-repulsion (XR) and electron-electron (EE) energies using DF densities on various auxiliary basis sets (x -axis) of H₂O, which appear in order of increasing basis size and are generated from two different AO basis sets, shown in panel (a) and (b). On the x -axis, the auxiliary basis size is M_* in the label PCD- M_* , counted by orbital type in spherical form. The root mean square error has been obtained from Eq. (3.36) with $n = 10$ for each one of the frozen core energy terms.

52 (JKFIT, see Table 3.1) auxiliary functions to 22 (PCD-22), for the aug-cc-pvdz AO basis. For this case, our auxiliary basis is slightly less precise than JKFIT, with a loss of accuracy less than a factor of 2 measured on the XR RMSE, which can be interpreted as negligible loss. On the other hand, results on EE energy do not show an advantage of our method for reaching a target accuracy. They show, however, that our method allows to systematically increase or decrease the accuracy by tuning the number of auxiliary basis functions. This can be interesting for applications where speed is a priority over accuracy. The most accurate auxiliary bases obtained with our method can achieve up to one order of magnitude smaller errors than existing methods. Lastly, note that a given ABS PCD- M_* is theoretically included in the PCD- $(M_* + 1)$ one. However, we observe a non-monotonous decay behaviour of errors. This may be attributed to error cancellation, appearing when summing errors of opposite sign associated to a given fragment pair. Recall that the measured error is on energies, which are scalar quantities given as a sum of terms, possibly subject to this kind of numerical phenomenon.

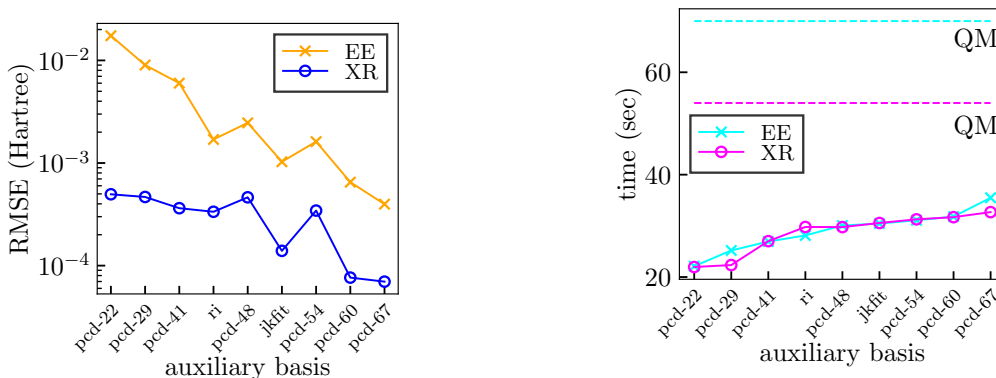
At this point we should recall that the requested accuracy in molecular simulation applications is 1 kcal/mol, equal to 10^{-3} Hartree. We may note that our method allows to generate auxiliary basis sets with errors that are below this threshold.

3.3.1.2 A small water cluster

Table 3.3 recapitulates the exact energies used as reference for our second test case, which is a cluster of 27 rigid water molecules. Results on errors due to density fitting are summarized in Figure 3.4a. Results show that the EE error of our auxiliary basis PCD- M_* decreases when adding more basis functions. The same tendency is observed for the XR errors, with two exceptions at

energy	min.	max.	mean	std	time
EE	4.147	19.955	9.561	3.873	69.192
XR	3E-18	4E-3	1E-4	6E-4	54.464

Table 3.3: Statistics of reference frozen core interaction energies (in Hartree) of a cluster of 27 water fragments, using the aug-cc-pvdz AO basis of H₂O. Timings (in seconds) are averaged over two executions for calculating the total energy.



(a) Root mean square error of DF with respect to the energies obtained from *ab initio* densities.

(b) Elapsed time in seconds (not averaged). The quantum mechanics (QM) time is obtained from Table 3.3.

Figure 3.4: Cluster of 27 water fragments. Evaluation of total exchange-repulsion (XR) and electron-electron (EE) frozen core energy parts using the aug-cc-pvdz AO basis set of H₂O and auxiliary bases. Same *x*-axis as in Figure 3.3.

PCD-48 and PCD-54. Similarly to the ten Smith dimers, we may attribute the non-monotonous error decay behaviour to error cancellation. Concerning the XR error, we can achieve roughly the same accuracy with the state-of-the-art RI method (40 basis functions by orbital type, see Table 3.1), using almost half the number of auxiliary functions with our method PCD-22.

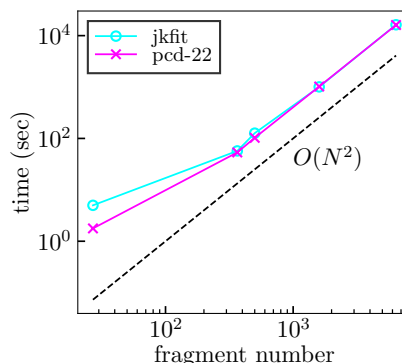
Next, results on the execution time spent on the total frozen core interaction of the 27 fragments are shown in Figure 3.4b. Note that we used the default OpenMP [44] parallelism implemented in PySCF for assembling molecular integrals when obtaining these timings, while the loop on fragment pairs is a sequential naive loop. These algorithms were launched on a home machine of 4 cores with 2 threads each. DF timings are compared to the exact reference computation obtained using electronic densities from *ab initio* methods. Results show that our schemes enable to reduce the execution time in comparison with state-of-the-art auxiliary basis sets. With roughly no accuracy loss, our PCD-22 method has a speed gain over the existing RI method, reducing the DF execution time approximately by a factor 1/3. Lastly, the speed of our larger generated auxiliary basis is comparable to that of JKFIT method. Overall, we reported that the use of DF method achieves a gain of at least 4/7 for EE and 2/5 for XR energies, with respect to the use of *ab initio* methods. This shows the important impact of DF schemes in the efficient evaluation of frozen core energies.

3.3.1.3 Large water clusters

The test cases that lastly concern us is a range of rigid clusters of varying size with minimum 27 water molecules, namely the water complexes given in Table 1.1. These are larger systems that are closer to real-world cases appearing in chemistry simulations. The aim of our present

nb.	jkfit	pcd-22	speedup
27	4.996	1.775	2.814
365	57.169	53.420	1.070
500	126.662	102.707	1.233
1600	1,005.946	1,023.038	0.983
6400	16,266.771	16,102.587	1.010

(a) The speedup factor is obtained as the ratio of timing of `jkfit` divided by timing of `pcd-22`.



(b) Time with respect to fragment number for each of the five test cases.

Figure 3.5: Large water clusters. Timings (in seconds) of EE frozen core energy total computation, using DF densities on two different auxiliary bases generated from the aug-cc-pvdz AO basis for H₂O, namely `jkfit` (52 basis functions by orbital type, see Table 3.1) and `pcd-22` (22 basis functions by orbital type).

study is to put our method to a test with respect to timings when it comes to large systems. Note that we did not attempt the calculation of the exact energy for assessing DF errors. The exact energy calculation based on *ab initio* densities is expected to be too expensive to compute. Besides, conclusions on the frozen core energy errors can be drawn by studying smaller systems.

Concerning simulation details, the fit has been performed using the Coulomb kernel explicit inner product. Timings were obtained by averaging two executions for each cluster, on a 256 CPU machine (2 sockets, 64 cores each, 2 threads per core), with 3.30 Ghz CPU frequency. The loop of fragment pairs has been parallelized as a naive attempt to gain efficiency. We refer to Appendix A.6.2.2 for more details on our employed parallelization strategy. In Figure 3.5, we compare the time performance of our smaller basis set obtained from aug-cc-pvdz, namely PCD-22, and the JKFIT state-of-the-art auxiliary basis set. Results show that our ABS has a better time performance with respect to the existing ABS. Obtained results of panel (a) show that our ABS enables computational gain, as the speedup factor is overall larger than one. The speedup is reported to be larger for the smallest cluster. In particular, for 27 fragments, we can gain roughly 30% of execution time by using our smaller auxiliary basis, compared to the state-of-the-art one. However, for larger clusters, the gain is smaller. Our interpretation is that this may be caused by the loop on fragment pairs, that eventually dominates the entire computation for larger numbers of fragments. Due to this loop, the gain obtained by DF is less noticeable. Recall that the speedup obtained by DF is a prefactor to the entire calculation, as discussed in previous Section 1.2.2. Indeed, our naive loop parallelization is not enough to break the quadratic complexity in terms of fragment number. Note that our naive parallel loop achieves a speedup by a numerical factor of 26 in this machine, equal to the ratio of timing without parallelization divided by timing with parallelization, measured for the cluster of 27 fragments. However, this might not be enough speedup in practice. This demonstrates the need to employ more sophisticated methods to accelerate the frozen core computation on large water clusters, which is typically the case for practical codes dealing with such systems [84].

3.3.2 Low-rank structure of Gram matrices

In this section, we present a study that enables to numerically describe the low-rank structure of the Gram matrix obtained by pairwise products of a given set of atomic orbitals. We first study the exact Gram matrix in Section 3.3.2.1 and then its restriction to atomic products in

AO basis	N_b^*	L	N_p^{*2}	$\text{cond}(\mathbf{G}) = \mathcal{O}(10^x)$				$\text{rank}(\mathbf{G})$				$\ \rho\ _K$	
				unweight.		weight.		unweight.		weight.		J	S
				J	S	J	S	J	S	J	S		
6-31g**	12	2	625	36	39	40	40	389	379	224	243	93.850	81.131
aug-cc-pvdz	18	2	961	37	39	39	41	595	548	337	360	93.279	81.261
aug-cc-pvtz	31	3	1,936	37	39	40	41	1,276	1,189	551	592	93.416	81.143

Table 3.4: Full Gram matrix \mathbf{G} of Eq. (3.30) from GTO basis sets of H₂O, for the K -explicit inner product. $K = J$ is the Coulomb kernel and $K = S$ is the Dirac one. The basis size N_b^* is given by spherical orbital type. The maximal angular degree of the basis is denoted by L . Scaling by the density matrix in Eq. (3.28) has been either used (weight.) or omitted (unweight.). The rank of \mathbf{G} is approximated as the number of singular values greater than 10^{-16} .

Section 3.3.2.2. The latter matrix is the starting point for constructing auxiliary basis functions using our method.

3.3.2.1 Full Gram matrix

Given an interaction kernel K and an electronic density ρ , the full Gram matrix is the N_p^{*2} -by- N_p^{*2} matrix defined in Eq. (3.30) and denoted by \mathbf{G} , built from ρ and the K -explicit inner product. The K -explicit norm of ρ is then equal to

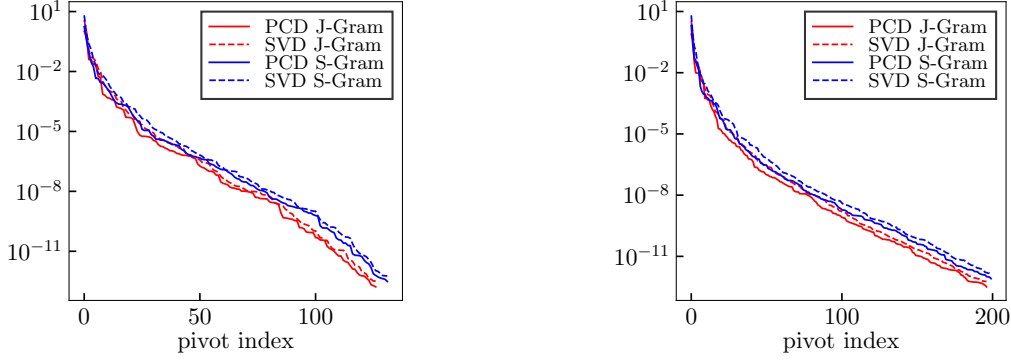
$$\|\rho\|_K = \sum_{i=1}^{N_p^{*2}} \sum_{j=1}^{N_p^{*2}} G_{ij}.$$

Analyzing the rank of the matrix \mathbf{G} can give insights on the structure we approximate. First, we demonstrate the numerical advantage of preconditioning with the density matrix from Eq. (3.28). Table 3.4 shows information about the full Gram matrices. Two types of Gram matrices have been considered, the first being weighted by the density matrix and the second is not weighted. Results show that it is preferable to weight the Gram matrix by the density matrix, in order to reduce its rank. Lower rank means that a smaller number of rows spans the full row space, therefore smaller auxiliary bases can be created. In the same table, the condition number of the weighted Gram matrix is also higher. This means that rows of weighted Gram have stronger linear dependencies, that we shall take them into account.

3.3.2.2 Atomic Gram matrix

The atomic Gram matrix of dimension N_c -by- N_c is defined in Eq. (3.33) and is denoted by \mathbf{A} . We recall that it was obtained by restricting the full Gram matrix to atomic products only. Note that our ABS method, described in Algorithm 1, selects rows and columns of \mathbf{A} .

A method to assess the low-rank structure of a matrix is to study the decay behaviour of its singular values. Figure 3.6 shows the singular values obtained by singular value decomposition (SVD) and the pivot values obtained by pivoted Cholesky decomposition (PCD) for the matrix \mathbf{A} . PCD has been launched with a tolerance equal to 10^{-15} . Results show a fast singular value decay around chemical accuracy, equal to 10^{-3} , demonstrating the low-rank structure of the matrix that we wish to approximate. In particular, for an accuracy equal to 10^{-3} , the approximate rank of the matrix in question is at least 10 times smaller than its maximal approximate rank, obtained by counting the number of singular values greater than 10^{-12} . The obtained results further show that the decay of Cholesky pivots follows the decay behaviour of SVD, that is the optimal low-rank approximation method. This means that Cholesky pivots provide a good indicator of the approximate rank for the cases of the tested matrices.



(a) aug-cc-pvdz. $N_c = 387$. Maximal numerical rank is 127 for J -Gram and 132 for S -Gram.

(b) aug-cc-pvtz. $N_c = 726$. Maximal numerical rank is 197 for J -Gram and 200 for S -Gram.

Figure 3.6: Singular values of SVD and pivot values of PCD for atomic Gram matrix \mathbf{A} of Eq. (3.33), for the K -explicit inner product. $K = J$ is the Coulomb kernel and $K = S$ is the Dirac one. PCD pivots are equal to $\text{diag}(\mathbf{L})^2$, where \mathbf{L} is the matrix in Eq. (3.15). All values are truncated to the maximal numerical rank of the matrix, equal to the number of singular values greater than 10^{-12} .

3.3.2.3 Quality of the approximation space

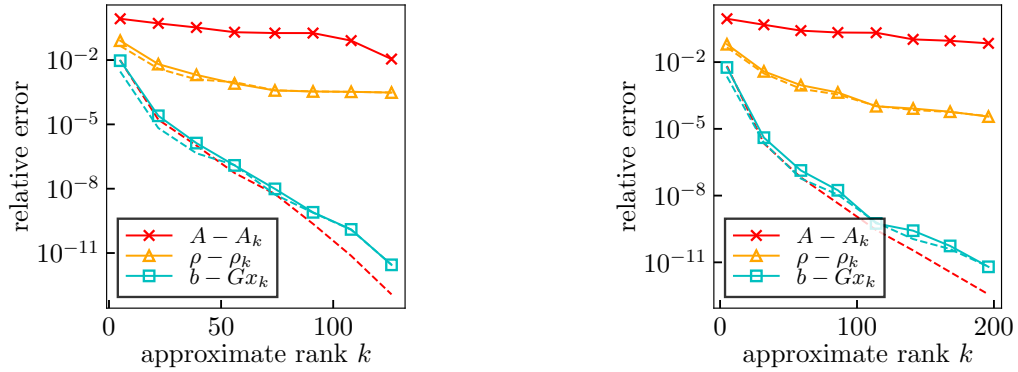
In the present section, we use various approximation methods to construct bases of the space spanned by the rows (and columns) of the atomic Gram matrix \mathbf{A} , previously defined in Eq. (3.33). Recall that constructing a basis of target size is not a goal in our ABS method, for the reason discussed previously in Section 3.2.1.1. However, the present study enables to compare our ABS with optimal bases of the same size, in terms of approximation accuracy in spanning the full space. The idea is to use our study to numerically examine if our ABS method leads to optimal best approximation errors for a given rank.

In particular, we numerically compare our method with the optimal method to generate basis sets, which is the Truncated Singular Value Decomposition (TSVD) (see Section 3.2.2.2). Our interest is to assess the accuracy of the space spanned by the PCD truncated at rank k and that of TSVD, truncated at the same rank. Our comparison is organized as follows. Three different methods are employed to measure the accuracy of the spanned space. First, we assess the ability of a low-rank approximation scheme to approximate the atomic restricted Gram matrix \mathbf{A} from Eq. (3.33). To this end, we use the spectral matrix norm, denoted by $\|\cdot\|_2$. Second, we examine the orthogonal projection error caused by the use of the basis of rank k in the density fit. Finally, we consider the system underlying density fit, that reads $\mathbf{G}\mathbf{x} = \mathbf{b}$, where $\mathbf{b} = \mathbf{G}\mathbf{1}$, $\mathbf{1}$ is the vector filled with ones of dimension N_p^{*2} and \mathbf{G} is the full Gram matrix from Eq. (3.30). Notice that this linear system is generally ill-posed, due to the ill-conditioning of the matrix \mathbf{G} , observed previously in Table 3.4. An approximate solution $\mathbf{x} \approx \mathbf{x}_k$ for a given rank k will be obtained using one of the two low-rank methods under comparison applied to \mathbf{A} , summarized as follows. Based on the TSVD scheme of Eq. (3.20), we derive, for any $k := M_* \leq N_c$

$$\mathbf{A} \approx \mathbf{U}\mathbf{\Sigma}_k\mathbf{V}^\top, \quad \mathbf{x} \approx \mathbf{x}_k := \mathbf{V}\mathbf{\Sigma}_k^\dagger\mathbf{U}^\top\mathbf{b}, \quad (3.37)$$

where $\mathbf{\Sigma}_k^\dagger$ is the pseudoinverse of the $\mathbf{\Sigma}$ matrix for a rank k , previously defined in Eq. (3.19). This scheme corresponds to approximating the full space of orbital products by a space generated by linear combinations of rows and columns generated by SVD. Next, based on PCD scheme of Eq. (3.16) of rank r (obtained for a tolerance value equal to 10^{-15} in experiments) we study the following approximations, for any $k := M_* \leq r$

$$\mathbf{P}_r^\top\mathbf{A}\mathbf{P}_r \approx \widehat{\mathbf{L}}_k\widehat{\mathbf{L}}_k, \quad \mathbf{P}_k^\top\mathbf{A}\mathbf{P}_k\widehat{\mathbf{x}}_k = \mathbf{P}_k^\top\mathbf{b}, \quad \mathbf{x} \approx \mathbf{x}_k := \mathbf{P}_k\widehat{\mathbf{x}}_k, \quad (3.38)$$



(a) aug-cc-pvdz: $N_c = 387$, $r = 139$ (PCD)

(b) aug-cc-pvtz: $N_c = 726$, $r = 241$ (PCD)

Figure 3.7: H_2O molecule. Comparison of TSVD (dashed lines) and PCD (solid lines) in constructing approximation spaces. The value $k := M_* \leq r$ approximates the rank of the atomic Gram matrix \mathbf{A} of Eq. (3.33). All Gram matrices are built for the Coulomb metric. Relative errors are $\|\mathbf{b} - \mathbf{G}\mathbf{x}_k\|_2 \|\mathbf{b}\|_2^{-1}$ (in blue) and $\|\rho - \rho_k\|_J \|\rho\|_J^{-1}$ (in orange) for ρ_k obtained by density fitting on the subspace of rank k , $\|\mathbf{A} - \mathbf{A}_k\|_2 \|\mathbf{A}\|_2^{-1}$ (in red). Explicit definitions are given for TSVD in Eq. (3.37) and for PCD in Eq. (3.38). Clarification: the reference matrix of PCD is $\mathbf{P}_r^\top \mathbf{A} \mathbf{P}_r$ but we write \mathbf{A} in plots for convenience.

where $\widehat{\mathbf{L}}_k$ is the r -by- r matrix obtained from \mathbf{L}_r by filling the rows and columns indexed by $k + 1, \dots, r$ with zeros. Notice that, we don't approximate the matrix \mathbf{A} but rather its symmetric permutation associated to \mathbf{P}_r . This scheme corresponds to approximating the full space of orbital products by a subspace generated by a selection of rows and columns indexed by PCD pivots. Lastly, we are concerned by the evaluation of the orthogonal projection error. This is a quantity of interest appearing in the minimization of the ABS problem in Eq. (3.7). To evaluate this quantity, we employ the formula of Theorem 2.3.1, according to which the orthogonal projection error is equal to the energy difference. The exact energy is equal to $\|\rho\|_K^2$ while the approximate energy is equal to $\mathbf{x}_k^\top \mathbf{A} \mathbf{x}_k$.

Figure 3.7 summarizes the results of our comparison between PCD and TSVD on the approximation space accuracy. TSVD is expected to provide the best low-rank matrix approximation, due to the Eckart-Young theorem. Results verify this fact, as the dashed red line of TSVD is far below the solid one of PCD. The interesting point to mention is that the lack of optimality of PCD in terms of low-rank errors does not appear to affect the numerical precision in other types of errors. Indeed, concerning the approximation of the solution to the ill-posed problem $\mathbf{G}\mathbf{x} = \mathbf{b}$, we see that the two methods perform equally well in terms of numerical precision. This is also true for the orthogonal projection error on the electronic densities. We interpret this as an encouraging result that demonstrates the capabilities of the PCD method in providing quasi-optimal approximation spaces for the specific setting we considered (i.e. specific choice of \mathbf{b} and \mathbf{G}). For instance, we may attribute this quasi-optimality to an artefact caused by the right-hand side \mathbf{b} of the problem. Note that this does not mean that our method is better than TSVD, but only that its numerical precision is comparable to that of TSVD. To summarize, the presented results tend to validate the theoretical error estimates and explicit comparison with TSVD leads to measured practical quasi-optimality on the target quantities for our problem setting.

Finally, we tested the stability of our least-squares fit by assessing the Euclidean norm of the approximate solution \mathbf{x}_k to normal equations $\mathbf{G}\mathbf{x} = \mathbf{b}$. We do not report complete results here, but we briefly outline that for both methods, i.e. PCD and TSVD, obtained results vary between 10^1 and 10^3 for different values of k . Additionally, the Euclidean norm of \mathbf{x}_k increases with the

rank. This means that adding more basis functions makes the linear system more unstable.

3.3.3 Auxiliary basis set detailed comparison

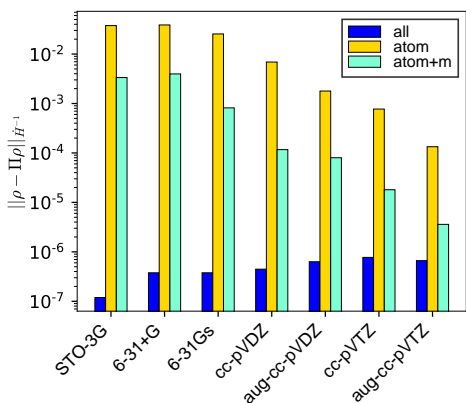
The present section contains a detailed side-by-side comparison of auxiliary basis sets generated by Algorithm 1 and conventional ones, namely RI and JKFIT. We focus on the methods’ performance on a single molecule of fixed geometry. Various aspects are discussed, such as projection errors, ABS sizes, maximal angular degrees and indicative generation timings. We are also interested in testing the ability of our ABS, generated for a given kernel metric, in approximating densities under other kernel metrics, which is a property referred to as *transferability* of ABS.

In this numerical experiment, the error is measured in terms of the orthogonal projection error under a given kernel-explicit norm. Table 3.5 presents the results on ABSs generated from AO bases for the H₂O molecule at equilibrium geometry. Results of our method are obtained for varying the input parameter, which is the target size of the ABS counted by orbital type. A first observation is that, for the AO basis aug-cc-pvdz, we attain the projection error values of state-of-the-art methods while using about half their number of basis functions. For the AO basis aug-cc-pvtz, this gap between the two compared methods is smaller. Additionally, for a target fixed size equal to 150, we achieve orthogonal projection errors that are 17.5% smaller than those of state-of-the-art methods. Overall, we believe that the main advantage of our method lies in the efficient reduction of ABS sizes, to obtain smaller or equal ones to those of state-of-the-art methods. Indeed, results on projection errors, as well as on frozen core errors in Section 3.3.1, demonstrate that one can achieve a balance between accuracy and complexity by using our ABSs belonging to such size range. However, our method generally lowers the projection error by adding orbitals of higher angular momentum to the ABS. This may have computational limitations due to the fact that the number of primitives increases and molecular integration of higher orbital types takes more time. Further reducing our ABS orbital types at the lowest accuracy level on projection errors (e.g. 10^{-4} and below) could improve the balance between accuracy and complexity.

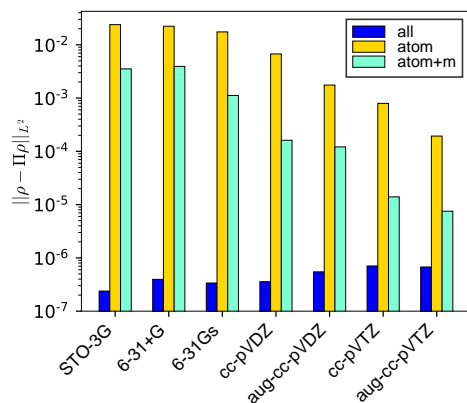
Further discussing Table 3.5, the next observation concerns the control of ABS sizes. Note that, the number of ABS elements is counted either by orbital type or by primitives. The fact that the ABS size by orbital type is almost identical, up to a difference smaller than 3 elements, to the input parameter used as a target, demonstrates that the post-processing step of Section 3.2.4.3, responsible for forcing atom-wise consistency, does not significantly add more functions to the ABS. We conclude that our method attains efficient control over the number of ABS elements counted by orbital type.

Concerning the transferability of our ABS, results show that changing the kernel metric, used withing the ABS generation, causes a minor loss on errors, that is less than 10% error increase. This means that we could interchange fit metrics and metrics used in the generation of ABSs when computing EE or XR energies. This has the advantage of requiring to only store ABS data and fit densities once for one metric, to obtain an approximate density that can be reused for all energies. We may notice in our results, however, that ABSs optimized for the Coulomb metric overall induce a larger loss on orthogonal projection errors when used for the XR energy, than the ones induced by ABSs optimized for Dirac when used for the EE energy. Hence, it might be preferable, in terms of orthogonal projection errors, to use ABSs generated from the Dirac metric to compute EE and XR energies.

Up next, let us discuss execution time and memory requirements. Timings are obtained only for the selection procedure of our method that depends on the target input parameter, that is Step IV of Section 3.2.4.2. Results show that the ABS selection time remains small, slightly increasing as the ABS size increases. As a supplementary result, we report that the total memory allocated during the generation of our ABSs from aug-cc-pvdz is 317 MB and from aug-cc-pvtz is 4685 MB. These values are dominated by the storage of the full Gram matrix \mathbf{G} of Eq. (3.30) and are thus constant for all ABS input parameters. Moreover, the assembly of the atomic Gram matrix \mathbf{A} of Eq. (3.33) and the *ab initio* computation are estimated to have a total execution



(a) Coulomb kernel norm.



(b) Dirac kernel norm.

Figure 3.8: H_2O molecule. Orthogonal projection errors on electronic density function for different choice of centers of the approximation basis. **atom**: atomic centers, **atom+m**: atomic centers and midbonds, **all**: unrestricted, all centers of the original orbital products are retained. Normal equations of the least-squares problem underlying the orthogonal projection have been solved using PCD with a tolerance equal to 10^{-14} .

time of about 3 seconds for aug-cc-pvdz and 10 seconds for aug-cc-pvtz. Note that this time should be added to the individual time of ABS generation, to obtain the total execution time of our method. Recall that, in frozen core energy computations, our method is executed offline.

Lastly, a numerical remark is that the projection error is generally decreasing, however is not strictly monotonous, even though the basis of size M_* is theoretically included to the one of size $M_* + 1$. This is caused by numerical instabilities of the PCD selection procedure, that make the pivot order slightly different along different executions (up to 10 entries difference). As a result, a given basis is not always included in bigger ones in practice. The impact on the projection errors is manifested with a small error difference, that is less than 10, hence we consider this effect as a minor numerical instability.

In the general case, the main reasons of our gains on the accuracy are both the weight by the density matrix and the ABS optimization based on molecules, i.e. from a given molecular geometry, instead of atoms as it is the case for state-of-the-art methods.

3.3.3.1 Selecting orbital centers

The present section is concerned by methods to attain projection error values beyond the ones reported in Table 3.5. The maximal accuracy of our ABS method (i.e. 10^{-5}) is attributed to discarding off-atom GTO products during Step III of our algorithm. Relaxing the restriction on orbital centers leads to higher precision.

Indeed, Figure 3.8 shows the density fit accuracy for different choices of centers in the approximation basis, obtained from the full Gram matrix \mathbf{G} of orbital products defined in Eq. (3.30). Three different methods are compared. Note that, we forced the midbond center to products localized in-between atoms. To achieve this, off-atom-centered products are replaced with a midbond-centered product of the same orbital parameters. Results show that adding midbond products can improve the accuracy of density fitting. Overall, results show that atomic-based density fitting is limited due to discarding non-atomic contributions.

Figure 3.9 shows the absolute value of pointwise errors due to density fitting, for the electronic density projected in the two-dimensional plane defined by three atoms of water molecule. Results show that the pointwise errors are larger on the heaviest atom, which is water. It is interesting to see larger errors on the bond segments OH as well, which may come from the fact that we used

Auxiliary basis	Input	M_*	M	L	J-fit error	S-fit error	Time (sec)
AO basis set: aug-cc-pvdz, $N_b^* = 18$, $N_b = 41$							
RI	–	40	118	3	9E-2	7E-3	–
JKFIT	–	52	150	3	4E-3	7E-4	–
PCD(J)	16	16	38	2	1E-2	7E-2	6E-3
	22	22	52	2	8E-3	6E-2	3E-3
	29	29	75	3	5E-3	1E-2	5E-3
	41	41	109	3	2E-3	1E-2	7E-3
	48	49	135	3	9E-4	3E-3	7E-3
	54	54	150	3	7E-4	3E-3	8E-3
	60	60	172	3	5E-4	3E-3	6E-3
	67	68	214	4	3E-4	3E-3	1E-2
118	118	372	4	2E-4	4E-4	1E-2	
PCD(S)	16	17	37	2	6E-2	4E-2	6E-3
	29	29	65	2	7E-3	8E-3	8E-3
	35	35	77	3	4E-3	5E-3	8E-3
	41	42	96	3	3E-3	2E-3	1E-2
	48	49	133	3	1E-3	1E-3	7E-3
	67	67	185	3	6E-4	9E-4	6E-3
	73	74	218	4	4E-4	6E-4	1E-2
	118	119	373	4	2E-4	4E-4	1E-2
AO basis set: aug-cc-pvtz, $N_b^* = 31$, $N_b = 92$							
RI	–	56	198	4	1E-3	2E-2	–
JKFIT	–	58	196	4	4E-4	4E-3	–
PCD(J)	40	41	105	2	2E-3	1E-2	6E-3
	52	52	142	3	1E-3	3E-3	1E-2
	64	65	189	4	3E-4	2E-3	9E-3
	77	77	231	4	2E-4	2E-3	1E-2
	126	129	509	4	8E-5	2E-4	1E-2
PCD(S)	40	40	96	3	3E-3	3E-3	7E-3
	52	52	128	3	1E-3	2E-3	1E-2
	64	64	172	4	7E-4	7E-4	9E-3
	77	77	227	4	4E-4	4E-4	8E-3
	114	116	396	4	9E-5	2E-4	2E-2
	126	129	465	4	8E-5	2E-4	1E-2

Table 3.5: Comparison of auxiliary basis sets generated from given AO bases in spherical form for H₂O molecule. Our method PCD(K) is obtained using the kernel K during the generation procedure, for $K \in \{J, S\}$. The K -fit error is equal to the orthogonal projection relative error $\|\rho - \tilde{\rho}\|_K \|\rho\|_K^{-1}$, with $\tilde{\rho}$ fitted to ρ using the K -norm. The ABS size is M_* by orbital type and M by primitives. M is lower (in green) or higher (in orange) than state-of-the-art values. L denotes the maximal angular degree in the ABS. Timings are averaged over 30 executions.

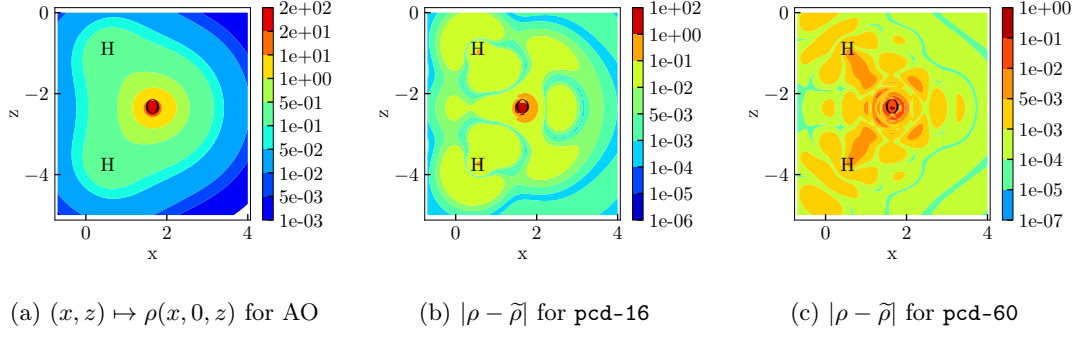


Figure 3.9: Electronic density function pointwise errors for water molecule. AO basis is aug-cc-pvdz. Isosurfaces in two-dimensional plane, coplanar with three atoms. Auxiliary basis sets are obtained with Algorithm 1 and are optimized for the Coulomb metric.

auxiliary GTOs centered on atomic positions only, neglecting off-atom centered orbital products. Moreover, we note that, increasing the auxiliary basis angular degree, i.e. from 2 to 3, improves the midbond errors by one order. This may be relevant to the numerical observation [198] that diffuse higher angular momentum orbitals improve the accuracy by better representing midbond orbital products.

3.4 Proofs

Proof of Corollary 3.2.0.1. We will use Theorem 2.3.1 and then solve the normal equations underlying density fitting on the auxiliary basis ξ_S to derive the desired error expression in terms of Gram matrices. In particular, in Theorem 2.3.1, we have derived an expression of the best approximation error on the density ρ in terms of self-interaction energies, according to which there holds

$$\min_{\sigma \in \text{Span}(\xi_S)} \|\rho - \sigma\|^2 = \langle \rho | \rho \rangle - \langle \tilde{\rho}_S | \tilde{\rho}_S \rangle, \quad (3.39)$$

with $\tilde{\rho}_S$ defined as the best approximation of ρ on $\text{Span}(\xi_S)$. Energies on the right-hand side of Eq. (3.39) can be expressed in terms of Gram matrices, as follows. First, the coefficients \mathbf{c} of $\tilde{\rho}_S$ on ξ_S can be found by solving a least-squares problem defined in Eq. (2.18). The normal equations for finding the coefficients \mathbf{c} read

$$\begin{aligned} \langle \xi_S | \xi_S \rangle \mathbf{c} &= \langle \xi_S | \rho \rangle \\ \Leftrightarrow \mathbf{G}(S, S) \mathbf{c} &= \mathbf{G}(S, :)\mathbf{1}. \end{aligned} \quad (3.40)$$

Problem (3.40) is well-posed due to the assumption of $\mathbf{G}(S, S)$ being non-singular. Its solution can be obtained by inverting the matrix, as

$$\mathbf{c} = \mathbf{G}^{-1}(S, S) \mathbf{G}(S, :)\mathbf{1}. \quad (3.41)$$

Now, using the expression of Eq. (3.41) for \mathbf{c} , we find

$$\begin{aligned} \langle \tilde{\rho}_S | \tilde{\rho}_S \rangle &= \mathbf{c}^\top \langle \xi_S | \xi_S \rangle \mathbf{c} \\ &= \mathbf{c}^\top \mathbf{G}(S, S) \mathbf{c} \\ &= \mathbf{1}^\top \mathbf{G}(:, S) \mathbf{G}^{-1}(S, S) \mathbf{G}(S, S) \mathbf{G}^{-1}(S, S) \mathbf{G}(S, :)\mathbf{1} \\ &= \mathbf{1}^\top \mathbf{G}(:, S) \mathbf{G}^{-1}(S, S) \mathbf{G}(S, :)\mathbf{1}. \end{aligned}$$

Finally, we employ Eq. (3.39) and the simple relation

$$\langle \rho | \rho \rangle = \mathbf{1}^\top \mathbf{G} \mathbf{1}$$

to conclude the proof of Corollary 3.2.0.1. \square

3.5 Conclusion and perspectives

In the present work, we construct auxiliary basis sets for use in density fitting problems in chemistry. Compared with empirical precomputed auxiliary basis sets of fixed size, the major difficulty in carrying out the auxiliary basis construction is guaranteeing that a prescribed basis size is attained. In order to proceed, we develop systematic tools for regrouping atomic orbitals and exploiting the low-rank structure of the orbital product Gram matrix. The result is a greedy algorithm for generating atomic orbital auxiliary basis sets of prescribed size. As a perspective, we plan to carry out numerical evaluation of the electron-nuclei interaction term of the frozen core electrostatics of Eq. (1.37) using our auxiliary basis set. Moreover, we will set up efficient parameter tuning strategies for using a single fitted density, optimized for one energy term, to different energy terms.

Besides solving the restricted ABS problem, it could be interesting to explore solutions to full auxiliary basis generation problems obtained by optimizing the spatial width of Gaussians, i.e. the orbital exponent. Such optimization could be based on sampling exponents over a bounded continuous interval, using random pointwise sampling techniques [102]. This direction could lead to generating optimal auxiliary basis sets. Relevant work on constructive schemes for selecting the propagating direction of plane wave basis can be found in [125]. One could also investigate the inclusion of additional non-atomic sites as Gaussian centers, using a relevant approach on enriched approximation spaces [154].

The numerical application of interest is the frozen core energy computation based on density fitting using our basis sets. From a practical point of view, results on density fitting errors using our auxiliary bases offer a greater range of accuracy than the ones obtained with fixed size state-of-the-art auxiliary bases. We demonstrate the practical use of our auxiliary basis set for water dimers and clusters of rigid fragments, in calculating frozen core intermolecular electron-electron and exchange-repulsion interaction energies. We also plan to apply our auxiliary basis generation strategy to other molecules. In addition to standard density fitting, we plan to investigate the performance of our auxiliary basis sets in the use of density fitting with additional constant electron charge constraints. The main difficulty is that we expect the least-squares error to increase when taking into account this constraint. A strategy that balances accuracy and complexity may become necessary in this case. Another possible direction to investigate is the link between orbital type, projection errors and size of the ABS. This could lead to new techniques designed to generate ABSs of low orbital type without sacrificing accuracy, by balancing all parameters.

Besides frozen core energies, another possible application of our auxiliary basis sets is the multipole moment generation. Such moments can then be used to calculate multipole expansions for electrostatics calculations in polarizable force fields (we refer to Appendix B for more details and discussion on the possible issues). The difficulty in such case is that often fitted densities are very diffuse, i.e. contain orbitals of high angular momentum, posing issues in the multipole expansion convergence. As a preliminary remark towards this direction, our generated bases contain only s and p orbitals at low prescribed accuracy, making them suitable for multipole moment calculations. For instance, the AMOEBA polarizable force field [263] uses up to dipole moments, that are generated from p -type orbitals. A second attempt, that we may interpret as promising, is that we may easily impose an additional constraint on maximum angular momentum within our auxiliary basis construction, in order to discard orbitals of high angular momentum. The effect of maximal orbital type constraints to the ABS size and accuracy is a topic left for future investigation. Note that the available `GEM_fit` implementation [138] allows the calculation of multipoles from our auxiliary basis input and could be used to conduct numerical testing. However, this objective requires performing numerical tests on geometry-specific water cases and carefully studying the effect of the atomic configuration to the result of the ABS generation, that could be studied in the future.

Chapter 4

The sparse density fitting method

Contents

4.1	Introduction	84
4.2	Theory	84
4.2.1	Problem formulation	85
4.2.2	Sparse density fitting schemes	91
4.2.3	Summary of sparse density fitting schemes	96
4.3	Numerical results	101
4.3.1	Comparison of sparsification methods	102
4.3.2	Comparison of density fitting methods	105
4.4	Proofs	107
4.5	Conclusion and perspectives	110

In the present chapter, we formulate our main contributions to density fitting methods for approximating frozen core intermolecular energies of large systems of rigid monomers. Our approach combines density fitting and integral screening techniques. First, we set up a Cholesky-based auxiliary basis set generation strategy performed entirely at the monomer level. Second, we derive a new density fitting scheme that decreases the number of effective interactions by exploiting the low-rank structure of interactions as well as their numerical sparsity. We present a discrete optimization perspective to the integral prescreening problem and then propose new tools to solve this problem. To achieve this, we employ a novel heuristic for threshold-based sparsification of the interaction matrix of the monomer under an arbitrary metric, that is numerically observed to achieve an optimal sparsity pattern in this setting. Our perspective specifically targets applications in the context of pairwise frozen core interaction energies of a large number of identical molecular fragments in arbitrary position in space, where precomputation of quantities presents computational advantages. The numerical performance of our schemes is compared to conventional density fitting schemes. Results show that our strategy is able to accelerate molecular simulations up to two orders of magnitude, with respect to system size, while respecting chemical accuracy.

4.1 Introduction

As we saw in the previous chapter, there is a need for accelerating the intermolecular interaction energy computation for large systems. The systems of rigid monomers, whose interatomic angles and distances remain unchanged between monomers, is the subject of the present chapter. Existing methods in the literature directly apply state-of-the-art density fitting schemes to the context of rigid monomers. Such state-of-the-art schemes use empirically optimized auxiliary basis sets associated to atoms. An advantage of such methods is that the auxiliary basis set is precomputed entirely for every atom. On the other hand, a limitation is that the accuracy of the auxiliary basis set is not tunable neither to a target fit metric nor to a target auxiliary basis size.

Basic integral screening methods rely on the Cauchy-Schwartz inequality (see Corollary 4.2.0.1). Sparsification methods for ERI matrices, based on entry-wise truncation, have been explored in [334]. This work used algebraic criteria for truncating matrices. In the present work, we focus on energy criteria, which are related to truncation of sums of entries instead of individual entries. This approach is targeting intermolecular interaction energy calculation applications. Sorting methods have been useful before applying integral screening [161], based on the order of magnitude. Alternatives to element dropping include the fast multiplication technique of matrices with decay [68], for the treatment of matrices whose elements decay exponentially.

In the present manuscript, we propose to revisit the existing methods and further adapt them to our context of application, under a greedy optimization point of view. We have devised algorithms that perform the density fitting computation offline at the monomer level, without any *a priori* knowledge of the positions of monomers inside the cluster. We have reduced the computational cost of existing density fitting schemes in two main ways. The first improvement allows the user to tune the size of the auxiliary basis during the auxiliary basis generation step, while the second one uses the sparsity of the interactions to further reduce the time spent on the interaction calculation. We also note that our algorithm allows for the treatment of all forms of pairwise interactions, including electron-electron, electron-nuclei and exchange-repulsion terms. Our first result is a new application of the Cholesky-based auxiliary basis generation method particularly adapted to the creation of orthonormal basis sets. The second result is a new *sparse* density fitting scheme based on the sparsification of the interaction matrix using a prescribed tolerance value. For the first method, we employ ideas based on low-rank approximation of the space spanned by Gaussian-type orbital products introduced in [198, 196]. For the second method, we use ideas of standard integral screening based on Cauchy-Schwartz bounds [317, 316, 264, 148] and combine them with a novel selection heuristic for maximal sparsification. Relevant topics on the sparsity of the interaction matrix have been previously discussed in [334].

The remaining sections are organized as follows. The theory of the main methods is formulated in an abstract setting in Section 4.2.1. Section 4.2.1.3 focuses on the new auxiliary basis construction as well as the sparsification of the Gram matrix, based on coupling density fitting and integral screening. Error bounds for density fitting schemes are formulated in this section. Section 4.2.1.4 introduces a computational complexity model to estimate the theoretical cost in terms of number of operations of a density fitting scheme. Our main contributions are presented in algorithms of Section 4.2.3. Lastly, Section 4.3 presents a numerical comparison of our methods with alternative state-of-the-art methods, in terms of accuracy, theoretical complexity as well as execution time.

4.2 Theory

We first announce the general problem of accelerating the evaluation of intermolecular interaction energies. Our approximation scheme is presented and analysed as for *a priori* errors, before focusing on ABS construction. In Section 4.2.1, we first formulate the summation problem under study and introduce sparse summations. Then in Section 4.2.1.3, we analyze *a priori* errors on interactions, that leads us to a thresholding strategy. In this respect, in Section 4.2.1.4 we present a constraint optimization problem that allows us to obtain solutions to the initial sparse

summation problem in a heuristic way. This problem is then solved using novel schemes in Section 4.2.2.

4.2.1 Problem formulation

The computation of intermolecular interactions in the Gaussian electrostatic model breaks down to solving a new summation problem defined in Eq. (2.1). Among the families of numerical methods solving this problem, density fitting allows to reduce the number of interactions by introducing a double summation over a approximation basis (see Section 2.3.3.2). However, among all interactions, not all of them contribute equally to the energy. Skipping sum terms can have great benefit in order to spend less computation time on entry-wise evaluation of molecular integrals. While density fitting construct fast summations by reducing the approximation basis size, we now propose a new family of methods, referred to as *sparse density fitting*, that achieve fast summation by reducing the actual summation domain geometrically. Using this strategy, we hope to beat the quadratic and cubic scaling on the number of interactions, obtained by state-of-the-art density fitting methods. In this respect, we define sparse density fitting schemes as solutions to the problem of selecting a minimal number of pairwise interactions given a summation and a prescribed level of accuracy.

4.2.1.1 The sparse summation problem

Let $(\mathbb{H}, \langle \cdot, \cdot \rangle)$ denote a real Hilbert space. Suppose that we want to solve the summation problem of Eq. (2.1), that consisted of computing the scalar quantity

$$E(f, g) = \sum_{i=1}^N \sum_{j=1}^N \langle f_i, g_j \rangle,$$

for any two given functions $f := \sum_{i=1}^N f_i$ and $g := \sum_{i=1}^N g_i$ expanded on finite basis sets of functions $\mathbf{f} = \{f_i\}_{1 \leq i \leq N} \subseteq \mathbb{H}$ and $\mathbf{g} = \{g_i\}_{1 \leq i \leq N} \subseteq \mathbb{H}$, respectively. In this respect, we first propose to employ the summation scheme of Eq. (2.11) to obtain two finite families of functions, denoted by $\tilde{\mathbf{f}} = \{\tilde{f}_i\}_{1 \leq i \leq M} \subseteq \mathbb{H}$ and $\tilde{\mathbf{g}} = \{\tilde{g}_i\}_{1 \leq i \leq M} \subseteq \mathbb{H}$, of fixed size $M \in \mathbb{N}^*$. Now, given a tolerance value $\delta > 0$, the *sparse summation problem* refers to the problem of constructing the set of indices $I_\delta \subseteq \{1, \dots, M\} \times \{1, \dots, M\}$ such that the scalar quantity $E_\delta \in \mathbb{R}$, defined as

$$E_\delta(\tilde{\mathbf{f}}, \tilde{\mathbf{g}}) := \sum_{(i,j) \in I_\delta} \langle \tilde{f}_i, \tilde{g}_j \rangle, \quad (4.1)$$

with $\tilde{f} := \sum_{i=1}^M \tilde{f}_i$ and $\tilde{g} := \sum_{i=1}^M \tilde{g}_i$, satisfies

$$|E - E_\delta| < \delta.$$

Recalling that the cost of computing a single pairwise interaction between any two functions \tilde{f}_i and \tilde{g}_j is given by $\tilde{\kappa}$, defined in Eq. (2.12), the computation of E_δ has a total operation count

$$\mathcal{O}(\tilde{\kappa} \text{card}(I_\delta)) \text{ flops.}$$

This is the expected theoretical complexity for evaluating a sparse summation, but we will show that it can be reduced in practice. The idea is to use E_δ to obtain fast summations of E .

From the Gaussian electrostatic model to sparse summations. The particular context of rigid fragments (recall Section 2.2.2.3) suggests to introduce the general problem of constructing I_δ such that for any two given isometries $T, T' : \mathbb{H} \rightarrow \mathbb{H}$ there holds

$$|E(Tf, T'g) - E_\delta(T\tilde{f}, T'\tilde{g})| < \delta.$$

This is the problem of finding a single pair of (\tilde{f}, \tilde{g}) and a single I_δ that works for all isometries.

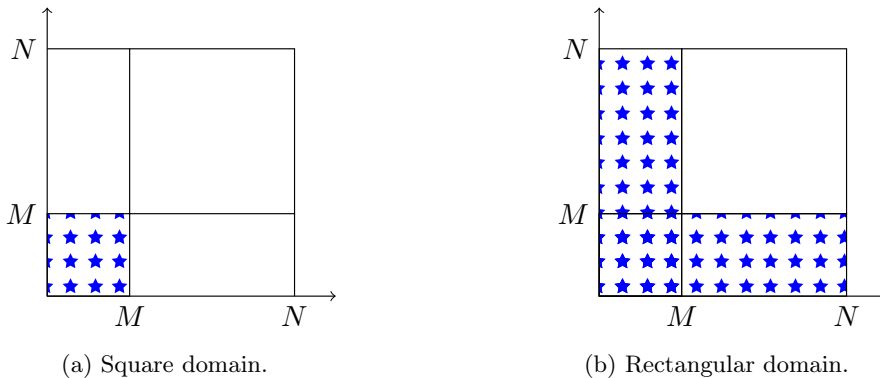


Figure 4.1: Comparison of different fast summation sizes, for the double summation of Eq. (2.1). The energy approximation is obtained by summing over retained indices marked in blue. Density fitting constructs square or rectangular domains given a size M .

4.2.1.2 Existing methods

Density fitting. This constructs squared domains. Conventional density fitting scales as $\mathcal{O}(M^2)$ while robust density fitting scales as $\mathcal{O}(MN)$. The limitation is that these methods impose a restriction on the domain pattern, i.e. either squared or rectangular, as seen in Figure 4.1.

Entry-wise screening. The most naive method to solve this problem is to directly use the Cauchy-Schwartz bound. This method, known as entry-wise screening [148], discards the sum entry indexed by (i, j) is discarded if $\|f_i\| < \delta'$ and $\|g_j\| < \delta'$, for some empirical value $\delta' > 0$ typically smaller than the target $\delta > 0$. The limitation is that the final error is not rigorously bounded by the target value δ .

Maximum norm screening. This consists of constructing the entire matrix with entries $A_{ij} = \langle f_i, g_j \rangle$ and then sparsifying it. Matrix norms may be used for truncating sum entries, such as the maximum norm truncation [334], which consist of skipping entries indexed by J if $\max_{(i,j) \in J} |A_{ij}| < \varepsilon$. Although easy to apply in practice, since it suggests discarding all entries smaller than a threshold, this criterion overestimates the error of discarding a single entry.

4.2.1.3 The sparse density fitting method

We propose to solve the sparse summation problem using density fitting. In this respect, we will construct an auxiliary basis set and projection coefficients on this set.

A priori error analysis. Our scheme is obtained by coupling density fitting and sum truncation. The coupling is explained as follows. First, we prove an *a priori* error bound that theoretically shows the error relation between two methods. We show that the intermolecular interaction approximation error can be decomposed into two parts, that are not independent. Then, we focus on each of the two parts and propose a method to balance errors.

Another interaction acceleration technique that can be studied under the prism of isometries is integral screening (see Section 1.2.3.1). We discuss screening techniques based on discarding isometry-independent quantities. First, notice that the isometry definition of Eq. (2.9) relates pairwise density interactions at different orientations. This has an application to screening based on spatial overlap, according to which one discards spatial distributions whose support overlap only on a small region. Applying an isometry will preserve the spatial overlap. The following

statement concerns another type of screening, based on discarding small contributions. It shows that small contributions can be discarded independently of density pairs and density orientations.

Corollary 4.2.0.1 (Integral screening). *Consider a real Hilbert space $(\mathbb{H}, \langle \cdot, \cdot \rangle)$ with induced norm $\|\cdot\| := \langle \cdot, \cdot \rangle^{1/2}$ and isometries $T_i : \mathbb{H} \rightarrow \mathbb{H}$ for $i = 1, 2$. There holds*

$$\forall (f_1, f_2) \in \mathbb{H} \times \mathbb{H}, \quad |\langle T_1 f_1, T_2 f_2 \rangle| \leq \|f_1\| \|f_2\|.$$

Proof. See [proof](#) in Section 4.4. □

Note that the upper bound is isometry-independent. Moreover, the bound can be computed using isolated densities, independently of their relative orientation.

The following statement yields an *a priori* error bound for sparse density fit methods, obtained by coupling a density fitting method with a sparse summation, based on discarding summation terms using integral screening.

Theorem 4.2.1 (Sparse density fit error). *Consider a real Hilbert space $(\mathbb{H}, \langle \cdot, \cdot \rangle)$ with induced norm $\|\cdot\| := \langle \cdot, \cdot \rangle^{1/2}$. For any fragment $F \in \{A, B\}$ of a given dimer AB , let $T_F : \mathbb{H} \rightarrow \mathbb{H}$ be an isometry. Given two electronic densities $\rho_A, \rho_B \in \mathbb{H}$ assume we perform the following density fitting scheme for each one of them (omitting the index):*

$$\rho \approx \tilde{\rho} = \sum_{i=1}^M \tilde{\rho}_i, \quad \tilde{\rho}_i := c_i \xi_i. \quad (4.2)$$

Let us fix $\delta > 0$. Given an index set $I_\delta \subseteq \{1, \dots, M_A\} \times \{1, \dots, M_B\}$, the intermolecular interaction approximation defined in equation (4.1) satisfies

$$|E(T_A \rho_A, T_B \rho_B) - E_\delta(T_A \tilde{\rho}_A, T_B \tilde{\rho}_B)| \leq \|\rho_B\| \|\rho_A - \tilde{\rho}_A\| + \|\rho_A\| \|\rho_B - \tilde{\rho}_B\| + \sum_{(i,j) \in I_\delta^c} \|\tilde{\rho}_i^A\| \|\tilde{\rho}_j^B\|,$$

with $I_\delta^c := (\{1, \dots, M_A\} \times \{1, \dots, M_B\}) \setminus I_\delta$.

Proof. See [proof](#) in Section 4.4. □

Remark. The error bound decomposes into two terms, reading

$$|E - E_\delta| \leq |E - E^{\text{DF}}| + |E^{\text{DF}} - E_\delta| \quad (4.3)$$

with E^{DF} denoting an intermediary Density Fitting (DF) approximation, defined as

$$E^{\text{DF}} := \langle T_A \tilde{\rho}_A, T_B \tilde{\rho}_B \rangle. \quad (4.4)$$

This bound is made up of two contributions. The first is the best approximation error due to density fitting and the second one is the sparse summation error, related to discarding interactions. The two errors are not independent of each other. Assuming that the best approximation errors are small, approximating E by E^{DF} induces an error smaller than the truncation threshold δ . The errors should be balanced. To this end, M_A and M_B should be small otherwise the retained summation terms have a larger number than the ones obtained for smaller M_A and M_B . A last remark is that the error bound is independent of the local to global transformations T_A and T_B . This means that we can construct a scheme completely offline, that guarantees convergence in intermolecular interactions for any orientation.

Given a tolerance $\delta > 0$, we define this index set, denoted by I_δ , as follows. The summation uses a truncation criterion of the following form. The set of retained indices I_δ is characterized by the set satisfying

$$\text{Find } I_\delta \text{ such that } \sum_{(i,j) \in I_\delta^c} \|\tilde{\rho}_i^A\| \|\tilde{\rho}_j^B\| < \delta. \quad (4.5)$$

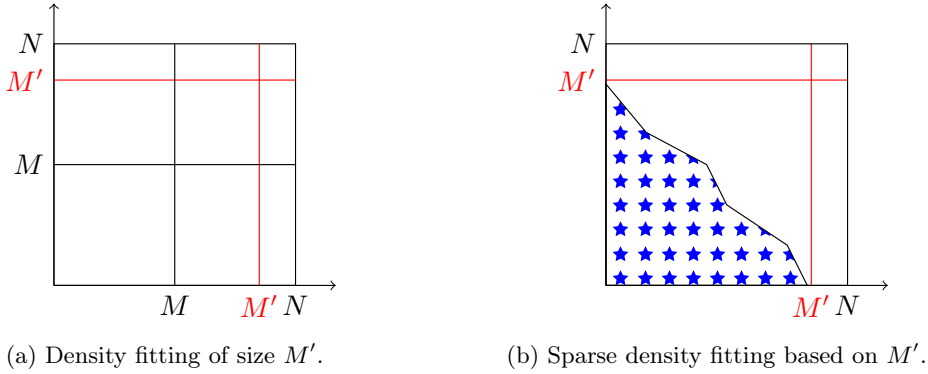


Figure 4.2: Sparse density fitting is obtained as a two-step procedure. First, we perform conventional density fitting with larger auxiliary bases than usual, i.e. M' with $M \leq M'$. Second, we truncate the square domain to achieve the desired complexity reduction. Extracted indices are marked in blue.

This truncation criterion is an alternative to the conventional truncation method, according to which a single function $\tilde{\rho}_i^A$ is skipped if $\|\tilde{\rho}_i^A\|$ is smaller than a prescribed threshold.

Our error balancing technique, between the two error components of Eq. (4.3), can be described as follows. Compared to conventional density fitting, our method consists of first fixing the auxiliary basis sizes M_A and M_B and then truncating the M_A -by- M_B square. Note that our strategy first fixes the density fit error, that shall be small enough so that one can afford a truncation error on top of it, without exceeding the target tolerance value overall. To this end, our auxiliary basis sizes are typically larger than conventional ones, namely the ones of Chapter 3. However, the complexity of the summation is then reduced by truncating the square, namely by solving the problem of Eq. (4.5). Figure 4.2 schematically illustrates the main idea of our method.

4.2.1.4 From density fitting to sparse density fitting

Let us begin by density fitting. If we further assume that the auxiliary functions, ξ_i 's, used in density fitting of Eq. (4.2) are normalized, then the truncation criterion of Eq. (4.5) reads as

$$\sum_{(i,j) \in I_\delta^c} |c_i^A c_j^B| < \delta. \quad (4.6)$$

Our goal is to construct an auxiliary basis that minimizes the cardinal of the retained index set I_δ . First, there are various ways to obtain the retained indices.

Truncation methods. If the auxiliary coefficients c_i^A and c_j^B are arbitrary, then a naive method to truncate the full sum

$$\sum_{i=1}^{M_A} \sum_{j=1}^{M_B} |c_i^A c_j^B|, \quad (4.7)$$

under the given tolerance value equal to δ , is to form all possible partial sums and find the one with the largest number of terms whose sum is smaller than δ . A smarter way to accelerate the search is to sort coefficients $|c_i^A c_j^B|$ in decreasing order and create partial sums incrementally. Hence, an alternative method is to set up an auxiliary basis such that the coefficients c_i^A and c_j^B are bounded in absolute value by a decreasing sequence that tends to zero. This enables to deduce that if a given $|c_i^A c_j^B|$ is small, then all couples greater than (i, j) , i.e. couples (k, l) with $k \geq i$ and $l \geq j$, will contribute to the sum even less than the couple (i, j) , they therefore have

more chances to be truncated. This observation is helpful for reducing the number of retained entries. However, in order to improve the decay behaviour of auxiliary coefficients, one should allow mixing of orbitals, contrary to the spirit of Chapter 3, increasing the complexity of the computation. A naive mixing makes the complexity of a single interaction equal to

$$\tilde{\kappa} = N_A N_B \kappa, \quad (4.8)$$

where N_A and N_B are the finite basis sizes of the reference densities ρ_A and ρ_B , used for mixing. Our mixing technique will propose alternative complexity costs based on mixing fewer functions.

Density fitting with mixed orbitals. We first describe the general framework that allows to construct auxiliary basis sets for which the auxiliary coefficients have a decay behaviour. Our motivation lies in the known fact that orthonormalizing before least-square fitting can improve the approximation error decay [3]. To achieve this, we assume that a reference density is given in a finite approximation basis of N functions, denoted by (the fragment index is omitted for simplicity)

$$\rho = \sum_{i=1}^N \rho_i. \quad (4.9)$$

The functions ρ_i 's will be referred to as *primitive orbitals* (or pure ones), because they determine the complexity cost κ in Eq. (4.8). The mixing that precisely allows us to reduce the cost of Eq. (4.8) is the one obtained by the following Gram-Schmidt procedure applied to functions $\{\rho_{s_i}\}_{1 \leq i \leq M}$, with $M \leq N$ and $\{s_1, \dots, s_M\}$ an appropriate permutation of a selection of M indices in $\{1, \dots, N\}$, whose construction is discussed later, as it affects the decay behaviour of the auxiliary coefficients. The entire process is referred to as *modified* Gram-Schmidt (MGS), due to the preliminary preordering step for functions to be orthonormalized. In this framework, the aim is to construct a permutation of primitive orbital indices that improves the decay behaviour of auxiliary coefficients.

The resulting mixed orbitals from MGS are the auxiliary basis functions, defined as follows. Let us fix $M \in \mathbb{N}^*$ with $M \leq N$. Given indices $\{s_1, \dots, s_M\} \subseteq \{1, \dots, N\}$ and a lower triangular matrix of coefficients $\mathbf{Y} = (Y_{ki})_{1 \leq k, i \leq M} \in \mathbb{R}^{M \times M}$, referred to as *mixing* matrix, we define the auxiliary basis set $\boldsymbol{\xi} = \{\xi_i\}_{1 \leq i \leq M}$ consisting of functions

$$\forall 1 \leq i \leq M, \quad \xi_i := \sum_{k=1}^i Y_{ki} \rho_{s_k}. \quad (4.10)$$

We then deduce that $\text{Span}(\xi_i)_{1 \leq i \leq M} = \text{Span}(\rho_{s_k})_{1 \leq k \leq M}$. The coefficients \mathbf{Y} are precisely orthonormalization coefficients, i.e. there holds

$$\forall 1 \leq i, j \leq M, \quad \langle \xi_i, \xi_j \rangle = \delta_{ij}.$$

Applying density fitting to ρ using $\boldsymbol{\xi}$ we obtain the approximation $\tilde{\rho}$, defined as $\tilde{\rho} = \sum_{i=1}^M \tilde{\rho}_i$ with $\tilde{\rho}_i := c_i \xi_i$ and \mathbf{c} being the auxiliary coefficients found by solving a set of normal equations (see Section 2.3.1.3). Note that if \mathbf{Y} is diagonal, then we end up with conventional truncation schemes without mixing orbitals, since auxiliary functions are primitive orbitals.

Cost of mixed interactions. The mixing matrix \mathbf{Y} of Eq. (4.10) being lower triangular in particular, conveniently enables precomputation in the case of auxiliary coefficients with decay behaviour. The main observation is that if a single mixed interaction, denoted by $\langle \xi_i^A, \xi_j^B \rangle$, is precomputed, then the computation of mixed interactions for couples (k, l) with $k \leq i$ and $l \leq j$ has zero additional cost. Moreover, the computation of the couple (k', l') , with $k' \geq i$ and $l' \geq j$, has an additional cost of computing $\langle \rho_{s_a}^A, \rho_{s_b}^B \rangle$ with $i+1 \leq a \leq k'$ and $j+1 \leq b \leq l'$. At the same time, as coefficients in the summation of Eq. (4.7) have decay behaviour, if (i, j) is retained, then

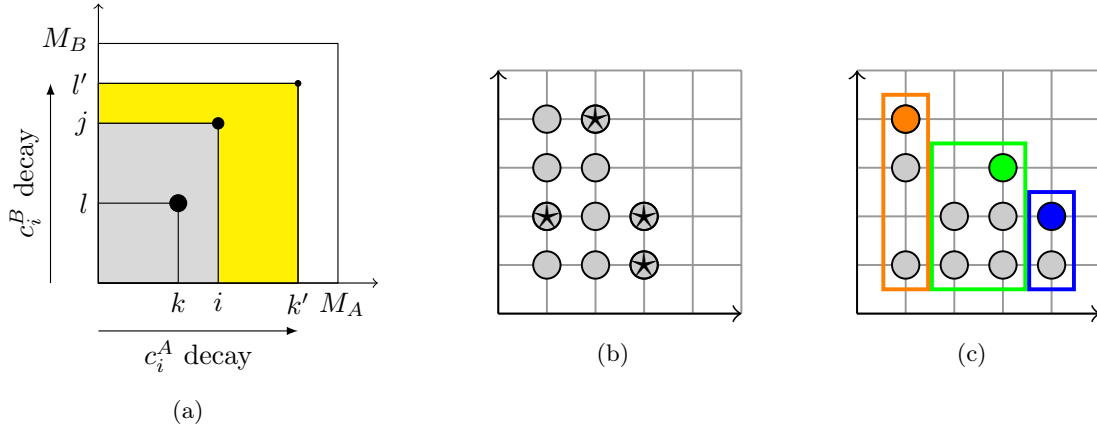


Figure 4.3: Primitive interactions involved in the computation of single mixed interactions in 4.3a, with precomputed couples (in gray) and additional couples (in yellow). In panel 4.3b, we shade couples (in gray) whose complexity is zero under the assumption that the starred couples have been precomputed. In panel 4.3c, boundary (three colored nodes) of a summation domain. The horizontal boundary points are obtained as the upper right corners of rectangles covering the domain.

it makes sense to retain all (k, l) with $k \leq i$ and $l \leq j$ as well, since their coefficients are likely larger than the one of the couple (i, j) . Indeed, retaining them will improve accuracy without cost, since all scalar products of associated primitive orbitals have been precomputed already. The other couple (k', l') , requiring additional computation, has less chance of being retained, as it corresponds to smaller coefficients. Such sets can be interpreted as two-dimensional patterns. Examples are provided in Figures 4.3a and 4.3b. This justifies how the lower triangular structure of the mixing matrix naturally relates to decaying coefficients.

With the forementioned consideration in mind, we can explicitly compute the cost of summation over mixed interactions. Given a summation denoted by (isometries are omitted for simplicity)

$$s := \sum_{i=1}^{M_A} \sum_{j=1}^{M_B} c_i^A c_j^B \langle \xi_i^A, \xi_j^B \rangle, \quad (4.11)$$

let us denote by $\kappa_s(I)$ the number of primitive interactions, i.e. $\langle \rho_{s_i^A}^A, \rho_{s_j^B}^B \rangle$'s, required to obtain the value s when skipping interactions indexed by I^c , for any $I \subseteq \{1, \dots, M_A\} \times \{1, \dots, M_B\}$. This cost allows to count the complexity in terms of primitive interactions using an alternative formula to that of Eq. (4.8). The starting point is the following related notion.

Definition 4.2.1. The *horizontal boundary* of any set $I \subseteq \{1, \dots, M_A\} \times \{1, \dots, M_B\}$ is the index set consisting of maximal row and column entries, i.e.

$$\partial^H I := \{(i, j) \in I : (k, j) \notin I, \quad \forall k > i, \quad (i, l) \notin I, \quad \forall l > j\}.$$

An example is provided in Figure 4.3c. Using this horizontal boundary we can count the number of entries. For the following, we assume for the sake of simplicity that fragments A and B are identical up to an isometry. This enables the use of the same auxiliary coefficients for both A and B , i.e. $\mathbf{c} := \mathbf{c}_A = \mathbf{c}_B$, since the auxiliary basis of A is simply the image of B under an isometry. We then end up with symmetric summation domains. Otherwise, for the non-symmetric case, the corresponding formula can be derived in an analogous manner.

Theorem 4.2.2 (*A priori complexity bound*). Consider the summation of Eq. (4.11), for ξ_i^A 's and ξ_j^B 's defined by the MGS of Eq. (4.10) on two fragments A and B , respectively. We further

assume that A and B are congruent, $M := M_A = M_B$ and $\mathbf{c} := \mathbf{c}_A = \mathbf{c}_B$ in. For any symmetric index set $I \subseteq \{1, \dots, M\} \times \{1, \dots, M\}$, let us denote by $\partial^H I := \{(i_k, j_k)\}_{1 \leq k \leq K}$ the set of horizontal boundary points, sorted in increasing order with respect to the first coordinate, i.e. $i_k < i_{k+1}$ for all $1 \leq k \leq K - 1$. There holds

$$\kappa_s(I) = \sum_{\substack{k=1 \\ i_k \leq j_k}}^K (i_k - i_{k-1})(j_k - i_k + 1),$$

with $i_0 := 0$.

Proof. See [proof](#) in Section 4.4. □

The sparse summation problem as an optimization problem. For a given tolerance value $\delta > 0$, let us fix $M = M(\delta) \in \mathbb{N}^*$. We assume that one has constructed a desired permutation of indices $\{s_1, \dots, s_M\}$ that allows to perform density fitting with auxiliary basis set defined as Eq. (4.10) and auxiliary coefficients denoted by $c_1, \dots, c_M \in \mathbb{R}$. Up next, we propose to solve the following problem:

$$\min_{I \subseteq \{1, \dots, M\} \times \{1, \dots, M\}} \kappa_s(I) \tag{4.12}$$

$$\text{subject to } \sum_{(i,j) \in I^c} |c_i c_j| < \delta. \tag{4.13}$$

The solution to this problem, denoted by $I_\delta = I_\delta(M)$, will provide solutions to Eq. (4.5). Note that M depends on δ and I_δ on both M and δ . The geometrical pattern of I_δ that is beneficial to respect when solving such optimization problem will be discussed in the following section. Optimal patterns are considered in the sense that adding more couples does not increase the complexity but improves the accuracy.

Equivalent matrix sparsification problem. In matrix form, our problem is also equivalent to a sparse approximation problem. Let us define the matrix $\mathbf{S} \in \mathbb{R}^{M \times M}$ with entries

$$S_{ij} := |c_i c_j|, \tag{4.14}$$

referred to as *screening matrix*. We sparsify this matrix using the 1, 1-norm, defined as

$$\|\mathbf{S}\|_{1,1} := \sum_{i=1}^M \sum_{j=1}^M |S_{ij}|.$$

Since \mathbf{S} has positive entries, this matrix norm is related to sum truncation as follows. To any summation index set $I \subseteq \{1, \dots, M\} \times \{1, \dots, M\}$ we may associate a sparse matrix $\widehat{\mathbf{S}} \in \mathbb{R}^{M \times M}$, with non-zero entries indexed by I , obtained by skipping entries of the screening matrix \mathbf{S} . The truncation of Eq. (4.5) can then be reformulated as

$$\|\mathbf{S} - \widehat{\mathbf{S}}\|_{1,1} < \delta.$$

4.2.2 Sparse density fitting schemes

In the present section, we present a two-step construction of sparse density fitting scheme. First, we build an auxiliary basis set based on mixed orbitals, then we specify how to select auxiliary function indices by solving a constraint optimisation problem. In the previous section we saw how to couple density fitting and sum truncation schemes. Here we propose methods to perform density fitting and sum truncation in practice. Before describing the algorithm, let us quickly introduce the main steps. For a given tolerance value $\delta > 0$, let us fix $M = M(\delta) \in \mathbb{N}^*$. Then, in order of execution:

1. Construct index permutation $\{s_1, \dots, s_M\}$.
2. Construct mixing matrix $\mathbf{Y} \in \mathbb{R}^{M \times M}$.
3. Construct screening matrix $\mathbf{S} \in \mathbb{R}^{M \times M}$.
4. Truncate screening matrix following appropriate geometric patterns.

4.2.2.1 Step I - Density fitting scheme

Modified Gram-Schmidt orthonormalization. This paragraph treats the first error component of Theorem 4.2.1 concerning density fitting. The goal is to end up performing density fitting using an orthonormal auxiliary basis set, for every fragment, with improved decay behaviour on auxiliary coefficients.

For simplicity we omit the fragment index to add it later. Consider a given finite family of functions $\boldsymbol{\rho} = \{\rho_i\}_{1 \leq i \leq N}$ such that $\rho := \rho_1 + \rho_2 + \dots + \rho_N$ defines a reference function to approximate, as in Eq. (4.9). For the sake of generality, we assume that the set $\boldsymbol{\rho}$ yields an admissible auxiliary basis set. Further basis set requirements, such as atomic center restrictions, orbital symmetries, etc., will be addressed later in Section 4.2.3, as they only affect the initialization of the admissible auxiliary basis set.

We compute the number $M \in \mathbb{N}^*$, the set of indices $\{s_1, \dots, s_M\}$ and the matrix \mathbf{Y} , for use in Eq. (4.10), as follows. A Gram-Schmidt process [322] is a systematic method for orthonormalizing a set of vectors or functions in a space equipped with an inner product. Our proposed Modified Gram-Schmidt (MGS) orthonormalization procedure is based on the Pivoted Cholesky Decomposition (PCD), previously introduced in Section 3.2.4.2. Let us define the Gram matrix $\mathbf{G} := \langle \boldsymbol{\rho} | \boldsymbol{\rho} \rangle$ of the family $\boldsymbol{\rho}$. PCD yields a decomposition of a permuted submatrix of \mathbf{G} , that reads

$$\mathbf{P}_M \mathbf{G} \mathbf{P}_M^\top = \mathbf{L}_M \mathbf{L}_M^\top, \quad (4.15)$$

with $\mathbf{P}_M \in \mathbb{R}^{M \times n}$ a row permutation matrix, $\mathbf{L}_M \in \mathbb{R}^{M \times M}$ lower triangular and $M \leq N$ the estimated rank of \mathbf{G} , for a given prescribed tolerance threshold. The resulting selection and permutation of indices is obtained by the first M pivots, defined in Eq. (3.18) and denoted by $\{s_1, \dots, s_M\}$. Let us assume that \mathbf{L}_M is invertible. The first M PCD-selected rows and columns of \mathbf{G} can be orthonormalized using the following matrix, denoted by \mathbf{Y} and referred to as *mixing matrix*, defined as

$$\mathbf{Y} := \mathbf{L}_M^{-1}. \quad (4.16)$$

Notice that this matrix verifies the property

$$\mathbf{Y} (\mathbf{P}_M \mathbf{G} \mathbf{P}_M^\top) \mathbf{Y}^\top = \mathbf{I},$$

with \mathbf{I} the M -by- M identity matrix, meaning that it orthonormalizes the set $\{\rho_{s_i}\}_{1 \leq i \leq M}$ where notice that functions are sorted in pivot order. Now, the coefficient matrix \mathbf{Y} being lower triangular, this orthonormalization defines a modified Gram-Schmidt scheme for the function set $\{\rho_{s_i}\}_{1 \leq i \leq M}$.

In analytic form, our MGS scheme corresponds to constructing an orthonormal basis of the space $\text{Span}(\rho_{s_i})_{1 \leq i \leq M}$, by defining the set of functions $\boldsymbol{\xi}$ of the form Eq. (4.10) for the choice of \mathbf{Y} given by Eq. (4.16). Recall that the orthonormalization order is entirely determined by PCD, which is greedy in terms of linear dependence. To summarize, our MGS scheme can be interpreted as a Gram-Schmidt procedure applied to a Cholesky pivot preordered function set, truncated to rank M for avoiding common numerical instabilities caused by the standard Gram-Schmidt procedure. Let us recall that PCD yields the M most linearly independent functions that span the full space. Intuitively, this is expected (we do not provide theoretical evidence) to enable a decay behaviour on the absolute value of the auxiliary coefficients.

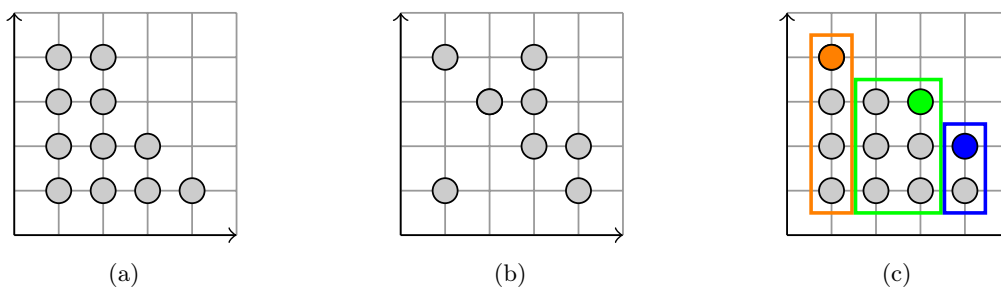


Figure 4.4: An example of L-shaped pattern 4.4a and a counter-example 4.4b. Schematically, a domain should fill lower left blocks of the first quadrant of the lattice incrementally without gaps. In 4.4c, a given horizontal boundary (three colored points) uniquely defines an L-pattern.

Lastly, the auxiliary coefficients on the orthonormalized basis can be found as follows. According to Eq. (2.18), least-squares normal equations simplify to

$$\mathbf{c} = \mathbf{Y}\mathbf{P}_M\mathbf{G}\mathbf{1}, \quad \mathbf{1} := (1, 1, \dots, 1) \in \mathbb{R}^M, \quad (4.17)$$

since the Gram matrix of orthonormal basis is the identity. Note this vector can be obtained by backward substitution, i.e.

$$\mathbf{L}_M\mathbf{c} = \mathbf{P}_M\mathbf{G}\mathbf{1},$$

for avoiding the explicit matrix inversion in favour of numerical stability. Let us emphasize that the vector \mathbf{c} may be evaluated entirely offline.

4.2.2.2 Step II - Sum truncation scheme

Once the auxiliary basis set has been constructed, we can compute the screening matrix \mathbf{S} defined in Eq. (4.14). Then, we may solve the optimization problem of Eq. (4.12) under the constraint Eq. (4.13). Here, we will propose an equivalent formulation based on graph theory notions. First, we introduce the target pattern that we choose to construct.

Target pattern. A pattern adapted to our problem should include as many terms as possible, in order to improve accuracy, while exploiting precomputations due to the lower triangular structure of our mixing matrix. This pattern may be formally defined as follows.

Definition 4.2.2 (L-pattern). Subsets of the integer lattice are said to have the *L-shaped* pattern (or *L-pattern*) if they belong to the following family

$$\mathcal{L} := \{I \subseteq \mathbb{Z}^2: \{\mathbf{j}: 1 \leq \mathbf{j} \leq \mathbf{i}\} \subseteq I, \quad \forall \mathbf{i} \in I\},$$

where \mathbf{i}, \mathbf{j} are multi-indices and the order relation \leq is defined component-wise, namely $(i, j) \leq (k, l)$ if and only if $i \leq k$ and $j \leq l$.

Sets satisfying this property are connected in the topological sense, as the examples illustrated in Figure 4.4. Note that a horizontal boundary uniquely defines an L-pattern, obtained by taking the union of points \mathbf{i} such that $\mathbf{i} \leq \mathbf{j}$ for \mathbf{j} boundary point. An illustration can be found in Figure 4.4c. Moreover, for any index set I , we can find an L-shaped domain $L \in \mathcal{L}$ such that $\kappa_s(I) = \kappa_s(L)$ and $I \subseteq L$. An example is provided in Figure 4.3c. In that sense, L-shaped domains achieve better accuracy than their subsets, since they contain more elements, while achieving the same complexity. For this reason, we seek to construct index sets that belong to the family \mathcal{L} , for maximizing accuracy for fixed operation count. This is actually an optimal sparsity pattern for our problem and we derived it *a priori*, based on the known structure of \mathbf{Y} .

Graph formulation. In the present paragraph, we construct sparse density fitting schemes from L-shaped patterns. Our main construction of L-patterns is based on graphs associated to summation domains. Using graphs, we introduce a novel discrete optimization strategy for sparse pattern construction. We define the main objects in the following. For the remaining chapter, fragments A and B are congruent, thus $\mathbf{c} := \mathbf{c}_A = \mathbf{c}_B$ are the auxiliary coefficients defined in Eq. (4.17).

The case of distinct fragments can be deduced in a similar manner. The graph building procedure is based on summation graphs. Consider an arbitrary number $n \in \mathbb{N}^*$. Let us motivate the summation graph construction. Before defining the directed graph, let us quickly summarize our terminology.

- *Vertices* are elements of the integer lattice $\Lambda(n, n) := \{1, \dots, n\} \times \{1, \dots, n\}$, corresponding to double summation index sets, or, in other words, to orbital pair indices. Note that, by the screening matrix \mathbf{S} symmetry, it suffices to restrict vertices to the integer lattice upper left corner.
- *Arcs* (south or east) are directed edges of unit length. Traversing a south arc means by convention skipping its starting node (i, j) from the summation. On the other hand, traversing an east arc means including its starting node in the summation. A *cost* has to be paid for traversing an arc and a *resource* has to be used along the arc. The cost corresponds to the computational complexity contribution to the summation, while the resource corresponds to the truncation entry error.
- *Paths* are ordered sequences of south or east arcs, with a source and a destination node.

The question one may ask is how to choose a minimal number of indices of the matrix \mathbf{S} under the constraint of Eq. (4.13). Such a choice of set of indices over the two-dimensional integer lattice can quickly become a tricky decision problem, due to the large number of combinations of indices one is allowed to consider. In order to deal with this issue, our method proposes to first construct horizontal boundaries and then recover L-shaped patterns defined from these boundaries. An elementary but useful remark is that, by Definition 4.2.1, any two given horizontal boundary points (i_1, j_1) and (i_2, j_2) with $i_1 \leq i_2$ necessarily satisfy $j_1 \geq j_2$. This allows to recover horizontal boundaries from consecutive south or east steps. Here, we define the notion of appropriate paths consisting of this kind of steps, referred to as *truncation path*. Then the main idea of our construction scheme is that any given truncation path on a summation graph defines a horizontal boundary of some domain. Since a given horizontal boundary uniquely defines an L-shaped domain, constructing truncation paths could lead to a solution for problem of Eq. (4.12) under the constraint Eq. (4.13). Before presenting the decision procedure for constructing such paths, let us first introduce graph notions that allow us to formulate the problem in this context. Let us quickly emphasize that our goal is to end up with a graph on which we may search for shortest paths under resource constraints.


Let us first define the summation graph on which we will solve an optimization problem, defined later on.

Definition 4.2.3 (Summation graph). Let $n \in \mathbb{N}^*$ and $\mathbf{S} = (S_{ij}) \in \mathbb{R}^{n \times n}$ a symmetric matrix. The *summation graph*, denoted by $G(\mathbf{S})$, is defined as the directed graph (V, A) with set of vertices

$$V = \{(i, j) : 0 \leq j \leq n, \quad 1 \leq i \leq j + 1\}$$

and set of arcs (directed edges) $A = A_{\text{east}} \cup A_{\text{south}} \cup A_{\text{diag}}$, where each set contains the following ordered pairs of vertices:

$$\begin{aligned} A_{\text{east}} &= \{((i, j), (i + 1, j)) : \quad 1 \leq i \leq j, \quad 1 \leq j \leq n\} \\ A_{\text{south}} &= \{((i, j), (i, j - 1)) : \quad 1 \leq i \leq j, \quad 1 \leq j \leq n\} \\ A_{\text{diag}} &= \{((i + 1, i), (i + 2, i + 1)) : 0 \leq i \leq n - 1\}. \end{aligned}$$



(a) Cost assignment for steps of the summation graph.

(b) Shortest path (in blue) and index set \bar{I}_* (in red). Horizontal boundary points have stars.

Figure 4.5: Steps of Algorithm 3 applied to an index set (in gray) of size 4-by-4. Note that vertices of the summation graph include the first sub-diagonal.

A cost κ_a is associated with each arc $a = ((i_1, j_1), (i_2, j_2)) \in A$, defined as

$$\kappa_a := \begin{cases} j_1 - i_1 + 1 & a \text{ east,} \\ 0 & a \text{ south or diagonal,} \end{cases}$$

together with a quantity of resource r_a associated to each arc a ,

$$r_a := \begin{cases} 2 \sum_{k=i_1}^{j_1} S_{k,j_1} - S_{j_1,j_1} & a \text{ south,} \\ 0 & a \text{ east or diagonal.} \end{cases}$$

Note that the factor 2 in the resource for south steps comes from the symmetry of \mathbf{S} . It is not necessary to include this factor 2 in the cost, however, evoking symmetry of \mathbf{S} that allows reusing computations.

Definition 4.2.4 (Truncation path). Let $G(\mathbf{S}) = (V, A)$ be a summation graph. A *truncation path* on $G(\mathbf{S})$ is defined as an ordered sequence of l arcs, $p = (a_1, \dots, a_l)$ such that $a_q \in A$ for all $q = 1, \dots, l$. Arcs are further assumed to be consecutive, meaning that the end of a_q is the start of a_{q+1} for all $1 \leq q \leq l-1$. The family of truncation paths on the graph $G(\mathbf{S})$ with source node \mathbf{o} and destination node \mathbf{d} is denoted by $\mathcal{P}_{\mathbf{S}}^{\mathbf{o} \rightarrow \mathbf{d}}$.

The resource constrained shortest path problem. The following provides a presentation of our strategy for constructing truncation paths. We propose a sparse pattern construction based on finding shortest paths under resource constraints. To quickly highlight the link between paths on graphs and retained summation indices, adding arcs to a path corresponds to deciding whether the arc starting node is skipped from the summation or not.

The underlying optimization problem is introduced as follows. The Resource Constrained Shortest Path Problem (RCSP) [268, 166] in operational research and in discrete optimization, is widely studied in a variety of settings, such as the truck driver scheduling problem, the vehicle routing problem, etc, in transportation systems. It is defined over a directed graph $G(V, E)$, where V is the set of n vertices and E is the set of m directed edges. A path with source vertex O and destination D is defined as an ordered sequence of l edges, i.e. $p = \langle (A_0, A_1), \dots, (A_{l-1}, A_l) \rangle$ with $A_0 = O$, $A_l = D$ and $(A_i, A_j) \in E$. The aim of RCSP is to find a path p with the smallest cost, such that $w(p) \leq W$, where $w(p)$ is the total resource of p obtained by summing resources associated to its arcs and W is a prescribed maximum amount of available resource.

We propose to apply this technique for solving problem of Eq. (4.12) under the constraint Eq. (4.13). The decision problem of constructing L-shaped patterns can be mapped to a well-defined discrete optimization problem, as described in the following. In this respect, optimized

summation domains can be obtained by solving the following problem. Given a maximum amount of available resource W (with $W = \delta$ in our application), our heuristic L-shaped domain construction is based on solving the following RCSSP

$$\min_{p \in \mathcal{P}_{\mathbf{S}}^{\mathbf{g} \rightarrow \mathbf{d}}} \sum_{a \in p} \kappa_a \quad (4.18)$$

$$\text{subject to } \sum_{a \in p} r_a \leq W. \quad (4.19)$$

We want to define a symmetric summation domain with the L-shaped sparsity pattern, such that the condition Eq. (4.13) is satisfied. This may be obtained by using our following heuristic. Let $\mathbf{S} \in \mathbb{R}^{n \times n}$ symmetric and $\delta > 0$. Let p_* be the solution to RCSSP (4.18) – (4.19) for resource capacity $W = \delta$, source node $\mathbf{o} = (1, n)$ and destination node $\mathbf{d} = (n, n + 1)$. Define the index set

$$I_* := \{(i, j) : (i, j) \leq v_1, (v_1, v_2) \in p_* \cap A_{\text{east}}\}. \quad (4.20)$$

Notice that the horizontal boundary of I_* is obtained from the starting points of east steps in the path p_* . Let us denote by \bar{I}_* the set obtained by symmetrizing I_* , i.e. by adding the corresponding symmetric indices such that $(i, j) \in \bar{I}_*$ if and only if $(j, i) \in \bar{I}_*$. Then \bar{I}_* has the following properties:

- i. $\bar{I}_* \in \mathcal{L}$,
- ii. $\sum_{(i,j) \notin \bar{I}_*} \mathbf{S}_{ij} < \delta$.

Notice that since \bar{I}_* is L-shaped, it is optimal for this δ -sparsity. Figure 4.5 illustrates the main steps of our L-pattern construction. Another remark is that, since the cost and resource of diagonal arcs, i.e. arcs in A_{diag} , is equal to zero, by Definition 4.2.3, we know that diagonal arcs are always chosen when possible. This is made possible precisely when the previous arc is either a diagonal one or it starts from a point belonging to the diagonal. Then the truncation path will always connect this point and the destination node \mathbf{d} through diagonal arcs. At the same time, diagonal arcs do not add points to the index set I_* . Diagonal arcs can hence be interpreted as fictitious arcs that simply allow to terminate the path once the diagonal has been reached.

4.2.3 Summary of sparse density fitting schemes

In the present section, we summarize sparse approximation methods for frozen core intermolecular interaction energy calculations based on density fitting. The entire process is presented within the offline-online simulation framework of Section 2.2.1.2. The core of our method is the sparsification strategy. Prior to that, we discuss various choices of Gram matrices one may use to apply such sparsification schemes.

4.2.3.1 Choice of admissible primitive ABS (offline)

During the Step I of density fitting, presented in Section 4.2.2.1, one has the option of initializing an admissible primitive auxiliary basis set (ABS) satisfying desired assumptions. The explicit definition of the matrix \mathbf{G} , used throughout Step I, starting from Eq. (4.15), is obtained from this primitive ABS. Here, we present and motivate possible choices of \mathbf{G} .

First, an intuitive choice of admissible auxiliary basis set is obtained by restricting the full set of orbital products, appearing in the density decomposition of Eq. (3.26), to orbital products centered on atoms, using methods proposed in Chapter 3. We recall from Section 3.2.3.1 that an important auxiliary basis set requirement is orbital symmetry with respect to angular degree. According to this property, all primitives associated to a given angular degree must be included in an auxiliary basis set. In order to impose this requirement during our sparse density fitting, we choose to contract primitives belonging to the same angular group, defined in Eq. (3.35),

before truncating orbital pairs. This allows to assure that all primitives of a given angular degree will be retained after the truncation step of sparse density fitting. Let us emphasize that a consequence of this contraction is that all primitives of a given angular degree will then have the same auxiliary coefficient. This is a difference that should be kept in mind, when comparing our sparse method with conventional density fitting, where all primitives had different auxiliary coefficients. The requirement is imposed by using the full matrix \mathbf{G} defined in Eq. (3.30), instead of the matrix \mathbf{A} defined in Eq. (3.33), for performing the density fitting in Section 4.2.2.1.

An alternative choice of admissible ABS can be obtained as follows. Recalling our MGS procedure of Section 4.2.2.1, the mixing matrix is computed from the pivoted Cholesky decomposition of the matrix \mathbf{G} , as in Eq. (4.16). A feature that plays an important role in the present study is the preconditioner applied to \mathbf{G} before performing the pivoted Cholesky decomposition. The preconditioner affects the permutation of indices $\{s_1, \dots, s_M\}$, used in Eq. (4.10), as well as the magnitude of auxiliary coefficients of mixed orbitals. In particular, the decay of the absolute value of the auxiliary coefficients is expected to be affected. We then recall that the screening matrix under sparsification is defined in terms of the auxiliary coefficients, as in Eq. (4.14). Our interest is thus to numerically assess the influence of the preconditioner of \mathbf{G} on the number of retained entries, obtained from sparsifying the screening matrix. In order to analyze the impact of the weights during sparsification, we also consider the case where the admissible primitive orbitals are normalized. The motivation behind studying this case lies in the fact that, our MGS procedure may be in practice applied either for a given reference density ρ , or for an approximate density $\tilde{\rho}$ obtained by density fitting. Let us emphasize that these two distinct cases lead to different scaling coefficients on the admissible auxiliary basis set. In particular, in the first case, the auxiliary basis functions used to build the Gram matrix, on which we apply pivoted Cholesky decomposition, will be weighted by the density matrix associated to the given ρ . In the second case, however, the Gram matrix is built from auxiliary functions associated to $\tilde{\rho}$, that are normalized by convention. Note that the density matrix is not used in this case. We expect that normalizing the functions used for building the Gram matrix will affect the decay behaviour of the auxiliary coefficients. Since, in quantum chemistry codes, the auxiliary basis is commonly normalized, assessing the sensitivity of our MGS procedure to this normalization is one of the purposes of the present study. This will enable us to investigate the applicability of our method to a fitted density $\tilde{\rho}$, for sparsifying auxiliary interactions between fitted densities.

A last choice of auxiliary basis set, considered in the present work, is the unrestricted one, obtained by the full set of atomic orbital products (i.e. including those centered on atoms, off-atoms or midbonds). Note that this last choice is considered in this work only to illustrate the influence of adding more centers to the auxiliary basis set. Although accuracy can be improved by adding more centers, recall that adding all centers without any restrictions is prohibitive for physical reasons in chemistry applications.

Let us briefly summarize the explicit definition of various Gram matrices for use in Eq. (4.15), employed in this work throughout our MGS procedure:

- A. The atomic Gram matrix of orbital products, scaled by the density matrix. The explicit definition of this matrix is given by Eq. (3.33). Recall that admissible auxiliary functions are formed by orbital type and are restricted to atomic centers. However, the resulting ABS does not respect atom-wise consistency.
- B. The *preconditioned* atomic Gram matrix of orbital products. Given the matrix \mathbf{A} defined in Eq. (3.33), the idea is to write $\mathbf{A} = \mathbf{D}\hat{\mathbf{A}}\mathbf{D}$ where $\mathbf{D} := \text{diag}(\mathbf{A})^{1/2}$. The result is that $\hat{\mathbf{A}}$ has unit diagonal, i.e. orbitals associated to rows and columns of $\hat{\mathbf{A}}$ are normalized. This matrix $\hat{\mathbf{A}}$ is known to have nearly optimal minimal condition number [156, Corollary 7.6].
- C. The full Gram matrix based on primitives. This can be obtained by folding dimensions of the matrix defined by Eq. (3.28) to obtain a N_p^2 -by- N_p^2 matrix. Note that this admissible ABS does not respect neither the orbital symmetry of Eq. (3.21) nor the atomic center restriction, as all centers are allowed.

4.2.3.2 Sparsification (offline)

Before presenting the explicit description of algorithms for sparsifying frozen core interaction matrices on finite basis sets, we quickly highlight the main features of each algorithm. Let us fix a target tolerance value $\delta > 0$ and assume we are given a finite basis of M primitive orbitals, selected among the admissible orbitals. Our goal is then to solve the problem of Eq. (4.1). We consider two different methods to achieve this:

- **Sparsification of Primitive Orbitals Pairs (SPOP)**. This algorithm decides if individual primitive orbital pairs must be retained or discarded. A direct sparsification technique is employed for solving problem of Eq. (4.5) for generating summation domains of arbitrary shape. Note that the MGS procedure is omitted altogether.
- **Sparsification of Mixed Orbital Pairs from Modified Gram-Schmidt (SMOP-MGS)**. This algorithm decides if individual mixed orbital pairs must be retained or discarded. Primitive orbitals are orthonormalized using our modified Gram-Schmidt procedure to obtain mixed orbitals (see Section 4.2.2.1). Our novel graph-based sparsification technique (see Section 4.2.2.2) is employed for solving problem of Eq. (4.12) under the constraint of Eq. (4.13) for generating L-shaped summation domains.

Let us clarify that a density fitting intermediate step is incorporated on each one of the two algorithms. The finite basis of M primitive or mixed orbitals is used as an auxiliary basis. The M -by- M matrix to be sparsified is obtained from the fitted coefficients. Before describing the explicit forms of the sparsification, we first introduce the algorithms for any given auxiliary basis set, for the sake of clarity. Then, we discuss sparsifications obtained from specific auxiliary basis set generation methods of interest.

Any given pure auxiliary basis set. For a given tolerance value $\delta > 0$, auxiliary basis set $\{\xi_i\}_{1 \leq i \leq M}$ (consisting of pure orbitals) and density ρ , we construct schemes such that the following self-interaction energy approximation

$$\langle \tilde{\rho}, \tilde{\rho} \rangle \approx \sum_{(i,j) \in I_\delta} x_i \langle \xi_i, \xi_j \rangle x_j, \quad (4.21)$$

has an absolute error smaller than δ , where $\tilde{\rho}$ is obtained by density fitting ρ on an appropriate basis, made up of (mixed or pure) orbitals obtained from the auxiliary basis, with \mathbf{x} the projection coefficients. Our schemes then provide I_δ and \mathbf{x} . Recall that the error δ comes on top of density fitting, as discussed in Remark 4.2.1.3. First, Algorithm 2 summarizes our scheme SPOP. The vector \mathbf{x} contains the auxiliary coefficients, since we consider pure orbitals in this scheme. Note that the sparsification, employed in Steps 3-7 of this algorithm, has an operation count given by $\mathcal{O}(M^2)$. Next, Algorithm 3 summarizes our scheme SMOP-MGS. The vector \mathbf{x} now contains the projection coefficients on the orthonormalized MGS basis of mixed orbitals. The sparsification is performed in Steps 5-7 of this algorithm. Notice that the graph assembly, as well as the graph problem of Eq. (4.18) and Eq. (4.19), underlying this sparsification, is much more complex than the one employed in SPOP and the complexity of its solution depends on the explicit implementation. Fortunately, this cost does not affect the online operation count, since we recall that all sparsifications are computed offline.

Specific choices of auxiliary basis sets. Explicit formulas of the sparsification algorithms can be provided for the ABSs of interest, presented in Section 4.2.3.1. In the present section, we describe how to generate auxiliary basis sets and then apply sparse density fitting.

Let us focus on Options A. and B.. Table 4.1 recapitulates formulas used for each method under study. Note that all quantities listed in the table are evaluated offline. The starting point is the matrix \mathbf{G} of dimension $N_p^{*2} \times N_p^{*2}$, obtained from a given AO basis set, as defined in Eq. (3.30). From the set of orbitals corresponding to rows and columns of \mathbf{G} , we obtain

Algorithm 2: Sparsify Primitive Orbital Pairs (SPOP).

Input: reference density ρ , auxiliary basis set $\boldsymbol{\xi} = \{\xi_i\}_{1 \leq i \leq M}$, kernel K , $\delta > 0$.

Output: summation domain I_δ and coefficients \mathbf{x} in Eq. (4.21).

- 1: Assembly auxiliary Gram matrix $\mathbf{G}_M = \langle \boldsymbol{\xi} | \boldsymbol{\xi} \rangle_K$.
- 2: Perform density fitting by solving $\mathbf{G}_M \mathbf{c} = \mathbf{b}$, with $\mathbf{b} := \langle \rho | \boldsymbol{\xi} \rangle_K$.
- 3: Build screening matrix $\mathbf{S} \in \mathbb{R}^{M \times M}$ with entries $S_{ij} := |c_i c_j| \|\xi_i\|_K \|\xi_j\|_K$.
- 4: Change the dimension of the screening matrix by flattening to a vector $\mathbf{v} \in \mathbb{R}^{M^2}$, with entries defined as

$$\forall 1 \leq k \leq M^2, \quad v_k := S_{i_k, j_k},$$

for $k \mapsto (i_k, j_k)$ a given enumeration of index pairs.

- 5: Sort entries of \mathbf{v} in increasing order, by constructing a permutation π of indices in $\{1, \dots, M^2\}$ such that

$$\forall 1 \leq k \leq M^2, \quad v_{\pi(k)} \leq v_{\pi(k+1)}.$$

- 6: Compute the cumulative sum of \mathbf{v} , i.e. the vector $\mathbf{c} \in \mathbb{R}^{M^2}$ with entries defined as

$$\forall 1 \leq k \leq M^2, \quad c_k := \sum_{l=1}^k v_{\pi(l)}.$$

- 7: Find the maximal index $1 \leq k_* \leq M^2$ such that $c_{k_*} < \delta$.
- 8: **return** $I_\delta := \{(i_k, j_k) : k := \pi^{-1}(l), \quad \forall 1 \leq l \leq k_*\}$ and $\mathbf{x} := \mathbf{c}$.

Algorithm 3: Sparsify Mixed Orbital Pairs from Modified Gram-Schmidt (SMOP-MGS).

Input: reference density ρ , auxiliary basis set $\boldsymbol{\xi} = \{\xi_i\}_{1 \leq i \leq M}$, kernel K , $\delta > 0$.

Output: summation domain I_δ and coefficients \mathbf{x} in Eq. (4.21).

- 1: Assembly auxiliary Gram matrix $\mathbf{G}_M = \langle \boldsymbol{\xi} | \boldsymbol{\xi} \rangle_K$.
- 2: Perform Cholesky decomposition $\mathbf{G}_M = \mathbf{L}\mathbf{L}^\top$.
- 3: Density fit on basis orthonormalized by MGS, by solving $\mathbf{L}\mathbf{c} = \mathbf{b}$, with $\mathbf{b} := \langle \rho | \boldsymbol{\xi} \rangle_K$.
- 4: Build screening matrix $\mathbf{S} \in \mathbb{R}^{M \times M}$ with entries $S_{ij} := |c_i c_j|$.
- 5: Assembly the summation graph associated to \mathbf{S} .
- 6: Solve RCSSP of Eq. (4.18) under the constraint of Eq. (4.19) for given δ to obtain the solution path p_δ .
- 7: Define I_δ from p_δ as in Eq. (4.20).
- 8: **return** symmetrized I_δ and $\mathbf{x} := \mathbf{L}^{-1}\mathbf{c}$.

orbital step	pure	mixed
full Gram	$\mathbf{A} = \mathbf{D}\mathbf{R}^\top \mathbf{G}\mathbf{R}\mathbf{D}$	
PCD	$\mathbf{P}_M^\top \mathbf{A} \mathbf{P}_M = \mathbf{L}_M \mathbf{L}_M^\top$	
aux. Gram	$\mathbf{G}_M = \mathbf{P}_M^\top \mathbf{A} \mathbf{P}_M$	
RHS	$\mathbf{b} = \mathbf{D}\mathbf{R}^\top \mathbf{G}\mathbf{1}$	
DF	$\mathbf{G}_M \mathbf{c}_M = \mathbf{P}_M^\top \mathbf{b}$	$\mathbf{L}_M \mathbf{c}_M = \mathbf{P}_M^\top \mathbf{b}$
screen	$\mathbf{S} = \mathbf{c}_M \mathbf{c}_M^\top \odot (\mathbf{w}\mathbf{w}^\top)$	$\mathbf{S} = \mathbf{c}_M \mathbf{c}_M^\top $
coef.	$\mathbf{x} = \mathbf{D}\mathbf{P}_M \mathbf{c}_M$	$\mathbf{x} = \mathbf{D}\mathbf{P}_M \mathbf{L}_M^{-1} \mathbf{c}_M$
energy	$\langle \tilde{\rho}, \tilde{\rho} \rangle = \mathbf{x}^\top \mathbf{R}^\top \mathbf{G}\mathbf{R}\mathbf{x}$	

Table 4.1: Explicit matrix formulas of various sparse DF methods, from a given matrix \mathbf{G} defined in Eq. (3.30). PCD of fixed rank M is given by Eq. (3.16). \mathbf{R} is the atomic restriction matrix defined in Eq. (3.32), \mathbf{w} is the vector of size M with entries $w_i := [\mathbf{P}_M^\top \mathbf{A} \mathbf{P}_M]_{ii}^{1/2}$ (note that $w_i = 1$ for Option B.). \mathbf{D} is a N_c -by- N_c diagonal matrix, defined as the identity for Option A. and $\mathbf{D} = \text{diag}(\mathbf{A})^{-1/2}$ for Option B.. The operation \odot denotes the entry-wise Hadamard product between same-sized matrices.

the set of admissible primitive orbitals by restricting to atomic centers. The selected primitive auxiliary basis set has size $M \leq N_c$, for N_c the number of atomic products for contracted AOs by orbital type and M a fixed rank used in PCD. The auxiliary Gram matrix, of dimension M -by- M , corresponds either to $\mathbf{P}_M^\top \mathbf{A} \mathbf{P}_M$ for pure orbitals, or to the identity for mixed orbitals. Notice that, in the case of pure orbitals, we obtain M auxiliary coefficients, which have the following feature. The Gram matrix used in density fitting is a submatrix of \mathbf{G} . Hence each row (and column) corresponds to a pairwise product of contracted AOs. Each contraction contains primitives belonging to an orbital type and having identical center and exponent. Hence, auxiliary functions are essentially contracted orbitals. This point affects the auxiliary coefficient computation during density fitting. This feature, also mentioned in previous Section 4.2.3.1, distinguishes the present methods, when compared with conventional density fitting schemes of Chapter 2, where auxiliary functions are primitive orbitals.

Lastly, one applies sparsification to the screening matrix \mathbf{S} , using Algorithms 2 and 3, namely SPOP for pure orbitals and SMOP-MGS for mixed orbitals. To this end, the auxiliary basis set of primitives, obtained by PCD of rank M given by Table 4.1, is the input to such algorithms. Note that auxiliary functions must be ordered according to PCD pivots.

4.2.3.3 Pairwise intermolecular interaction (online)

Our sparse density fitting scheme allows to construct solutions to the pairwise interaction problem for computing frozen core energies, of the form of Eq. (2.1). In the present section, we formulate the final scheme in matrix form that allows us to obtain fast summations as solutions to this problem, computed online. Throughout this section, we are placed in the Gaussian electrostatic model framework for computing frozen core energies, previously presented in Section 2.1.2.2.

Let us consider a dimer made up of two rigid identical congruent molecular fragments, denoted by A and B . Recalling notions from Section 2.1.2.1, the frozen fragment library then consists of a single fragment, congruent to A and B , whose frozen electronic density, denoted ρ , is given in

the form of

$$\rho = \sum_{i=1}^{N_b} \sum_{j=1}^{N_b} D_{\mu\nu} \chi_\mu \chi_\nu,$$

where \mathbf{D} is a N_b -by- N_b density matrix obtained by solving an *ab initio* ground state problem, of the form of Eq. (1.11), discretized over the atomic orbital basis $\{\chi_\mu\}_{1 \leq \mu \leq N_b}$, centered on atomic positions of the frozen fragment expressed in the local frame. Note that the decomposition of Eq. (4.9), reading $\rho = \rho_1 + \rho_2 + \dots + \rho_N$ then corresponds to taking $N = N_b^2$ and each ρ_i is a function obtained as the pairwise product of two AOs, weighted by the density matrix.

For any two given positions of fragments A and B in the global fragment frame, we may find two isometries, denoted by T_A and T_B , which map the frozen density to its expression in the global frame. The exact cross-Gram matrix underlying the interaction between *ab initio* frozen densities of fragments A and B can be defined as the four-index tensor \mathbf{G}_{AB} , with entries

$$\forall 1 \leq \mu, \nu, \kappa, \lambda \leq N_b, \quad [\mathbf{G}_{AB}]_{\mu\nu\kappa\lambda} := \langle T_A(\chi_\mu \chi_\nu), T_B(\chi_\kappa \chi_\lambda) \rangle_K. \quad (4.22)$$

Notice that each entry corresponds to a (possibly) contracted atomic orbital.

Let us begin by Options **A.** and **B.** The vector \mathbf{x} , obtained by Table 4.1, is then left and right multiplied with the matrix \mathbf{A}_{AB} , of dimension N_c -by- N_c , that is obtained from the tensor \mathbf{G}_{AB} defined in Eq. (4.22), by applying the block treatment given by Steps I - III of Section 3.2.4.1. We recall that such steps perform primitive uncontraction, weight by the density matrix, contract by orbital type and restrict products to atomic centers. The frozen core energy approximation, given by the intermediary DF approximation of Eq. (4.4), is given by

$$\langle T_A \rho, T_B \rho \rangle_K \approx E^{\text{DF}} = \mathbf{x}^\top \mathbf{A}_{AB} \mathbf{x}.$$

For fixed sparsity tolerance value $\delta > 0$, the final sparse approximation is obtained as

$$E^{\text{DF}} \approx E_\delta = \sum_{(i,j) \in I_\delta} x_i x_j [\mathbf{A}_{AB}]_{ij},$$

where I_δ is obtained by sparsifying the screening matrix \mathbf{S} under the tolerance value δ , using either Algorithm 2 for primitive orbitals or Algorithm 3 for mixed orbitals.

Lastly, for Option **C.**, our density fitting scheme has the following matrix form, that we explicitly write down here since it presents a particular interest. Let us use the hat notation over matrices to denote that they are obtained from the given AO basis, by uncontracting primitives, weighting by the density matrix and then folding dimensions of products, i.e. by turning the tensor of Eq. (3.28) into a square matrix. Let $\hat{\mathbf{G}}$ be the Gram matrix on the local frame and $\hat{\mathbf{G}}_{AB}$ the cross-Gram matrix on the global frame in which A and B live. Assume that S denotes the index set selected by PCD under a given tolerance value. In the SPOP method, the cross-Gram matrix between A and B is approximated, due to conventional density fitting, as

$$\hat{\mathbf{G}}_{AB} \approx \hat{\mathbf{G}}(:, S) \hat{\mathbf{G}}(S, S)^{-1} \hat{\mathbf{G}}_{AB}(S, S) \hat{\mathbf{G}}(S, S)^{-1} \hat{\mathbf{G}}(S, :), \quad (4.23)$$

It may be interesting to comment that our matrix decomposition of Eq. (4.23) is known as a nested adaptive cross-approximation [29, 139]. The SMOP-MGS method, applied to this kind of Gram matrix, uses the screening matrix

$$S_{ij} = |x_i x_j|, \quad \mathbf{Lx} = \hat{\mathbf{G}}(S, :)\mathbf{1}, \quad \hat{\mathbf{G}}(S, S) = \mathbf{L}\mathbf{L}^\top.$$

4.3 Numerical results

The present section concerns numerical results and is organized as follows. We first compare two methods of Gram matrix scaling and two methods of sparsification, for each scaling. The total of four methods is summarized in Table 4.1. The results of this comparison enable us to

deduce the best strategy of sparse density fitting, among the compared ones. For this chosen strategy, we then compare it with various methods for computing frozen core energies existing in the literature, namely conventional density fitting and robust density fitting.

Concerning simulation details, all results presented in this section are generated using our `dfwpy` code in Python (see Appendix A for implementation details). The discrete optimization problem in SMOP-MGS scheme of Algorithm 3 is solved using an exact method, presented in Appendix A.5. Target molecular systems used for testing are rigid water fragments.

4.3.1 Comparison of sparsification methods

The present section concerns quantifying the impact of the sparsification on the number of orbital pairs, using different scaling methods. Before presenting the comparison, we first introduce a notion that enables us to assess the efficiency of each algorithm, measured in terms of discarded orbitals pairs. Then, the comparison is conducted based on various methods, including the sparsification effectiveness, frozen core energy errors and timings.

4.3.1.1 Sparsity index

The comparison of sparsification schemes is carried out based on a quantitative indicator of the sparsification effectiveness, inspired by [334]. Given a threshold value $\delta > 0$ and a matrix \mathbf{M} of dimension n -by- n , the *sparsity index* of \mathbf{M} is defined as the fraction of non-zero elements:

$$\frac{|\{(i, j): |M_{ij}| < \delta\}|}{n^2}.$$

Zero sparsity index means that the matrix is numerically dense, while index equal to one means that the matrix is numerically zero. The sparsity index basically indicates how fast we can evaluate the matrix \mathbf{M} entry-wise, under the assumption that entries below δ are numerically zero. Larger sparsity index means that the underlying sparsification scheme is more effective.

In order to compare sparsification methods, the idea is to first fix an admissible Gram matrix \mathbf{A} of dimension N_c -by- N_c , defined in Table 4.1, and a fixed rank $M \leq N_c$. Recall that \mathbf{A} is scaled accordingly, depending on either Option A. or Option B.. For the auxiliary Gram matrix of size M -by- M obtained by PCD for M , denoted by \mathbf{G}_M , one then applies the scheme SPOP, given by Algorithm 2, and the scheme SMOP-MGS, given by Algorithm 3, to obtain the retained entries of \mathbf{G}_M under a tolerance δ . Notice that, by fixing the auxiliary Gram matrix \mathbf{G}_M , we assure that all algorithms are applied to the same auxiliary basis, up to the scaling coefficients. In this numerical experiment, we compute the sparsity index for the sparse approximation of \mathbf{G}_M obtained for various tolerance values δ given as input in either SPOP or SMOP-MGS. We thus obtain four methods under comparison.

Figure 4.6 presents the results of this comparison. Results show that for small target threshold values, all sparsification methods perform equally. However, for larger threshold values, including the target chemical accuracy at 10^{-3} , it is observed that SMOP-MGS is more effective, especially when applied to the scaling of Option A., corresponding to density matrix coefficients. Another observation is that the sparsity index of a given scheme seems to decrease, as we increase the AO basis size. This can be seen by comparing any given curve point in panels (a) and (b).

4.3.1.2 Frozen core energy

Up next, we assess the effective sparsification error on frozen core energies. For each sparsification method under comparison, we vary the rank of the PCD used to obtain the auxiliary basis of density fitting. The goal is to find the sparsification method that achieves greater accuracy on frozen core using the smallest number of orbital pairs, for fixed rank.

Table 4.2 illustrates the results obtained for ten standardized Smith dimers, using the root mean square error (RMSE) metric of Eq. (3.36). This table shows the approximation errors and the number of entries that each sparsification scheme considers as nonzero, for varying Gram

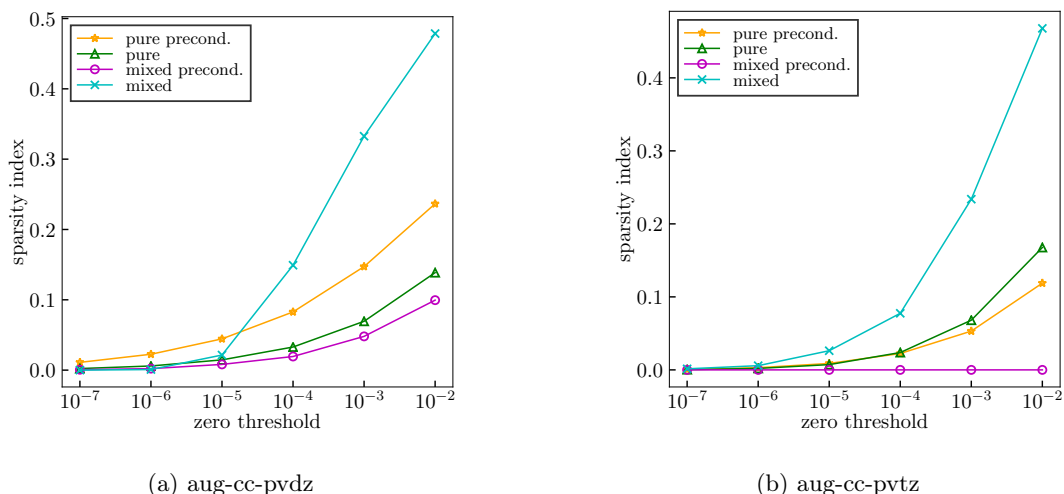


Figure 4.6: Fixed $M = 120$. Sparsity index of sparsification methods applied to auxiliary Gram matrix \mathbf{G}_M of dimension M -by- M , given by Table 4.1, for H_2O and the Coulomb kernel. SPOP (pure) and SMOP-MGS (mixed) sparsifications, for varying zero-threshold δ , are applied to \mathbf{G}_M scaled either as Option A. or as Option B. (precond.).

matrices on different AO bases. The auxiliary bases used in density fit are truncated to several lower ranks using PCD. We observe that each sparsification method yields a different number of retained entries. In particular, results show that, by using SMOP-MGS and the density matrix scaling on the auxiliary Gram matrix, i.e. Option A., we achieve better overall accuracy on frozen core energies using fewer orbital pairs than other schemes. This holds true for various ranks M used in PCD. From the obtained results on nonzero entries, we may also conclude that the minimal condition number of the auxiliary Gram matrix \mathbf{G}_M , achieved for the Option B. as expected from theory [156], does not seem to improve the sparsification efficiency. Regarding RMSE errors, it is often observed that the effective error on frozen core energies is lower than the sparsification tolerance. This may indicate that error cancellation effects might be taking place when summing orbital pairs, but these effects are much harder to master.

On the same experiment, the varying rank M used in PCD allows to study the balance between density fitting and sparsification, as addressed in Remark 4.2.1.3 following our main error estimator of Theorem 4.2.1. Recall that, for every M , we sparsify the screening matrix of size M -by- M , with an increasing number of entries as M increases. Hence, the number of retained entries is expected to increase with M , for fixed tolerance value used in the sparsification, as results show. A parameter tuning strategy may become necessary in order to determine the best size M one should use in order to achieve a target RMSE value on energies.

Lastly, concerning computational complexity, the exact frozen core energy computation used as a reference involves the summation of $\mathcal{O}(N_b^4)$ orbital pairs, while sparsification methods only involve the number of nonzero entries. Further commenting on results of Table 4.2, we can see that the number of nonzero entries of the Gram matrix, precisely after applying the best overall sparsification scheme SMOP-MGS, is linear in the total number of rows, which is equal to N_b^2 . The prefactor of the linear scaling is estimated to be roughly equal to 10. Such observation holds true for various rank values and AO basis sets. Hence, overall, it may be interpreted as an indicator showing that linear-time (i.e. quadratic in N_b) \mathbf{G}_M assembly schemes are possible for the specific matrices of our problem setting.

AO basis	N_b^2	M	\mathbf{G}_M nonzero entries ($\delta = 10^{-3}$)				RMSE on EE (Hartree)			
			Option B.		Option A.		Option B.		Option A.	
			prim.	mix.	prim.	mix.	prim.	mix.	prim.	mix.
6-31g**	625	25	608	625	580	624	8E-1	8E-1	1E-2	1E-2
		40	1,550	1,600	1,580	1,521	3E-1	3E-1	3E-3	3E-3
		60	3,211	3,560	3,400	3,198	4E-3	4E-3	4E-3	4E-3
aug-cc-pvdz	961	50	2,143	2,500	2,435	2,415	5E-1	5E-1	4E-3	4E-3
		80	5,971	6,396	6,216	5,524	2E-2	2E-2	1E-3	1E-3
		100	8,925	9,850	9,359	7,324	3E-3	3E-3	1E-3	1E-3
		120	12,184	13,950	13,400	9,610	2E-4	2E-4	3E-4	3E-4
aug-cc-pvtz	1,936	90	7,611	8,096	8,100	6,929	2E-1	2E-1	4E-4	4E-4
		120	13,755	14,395	14,400	11,033	8E-2	8E-2	3E-4	3E-4
		150	21,199	22,486	20,526	16,250	1E-2	1E-2	1E-4	1E-4

Table 4.2: Ten standardized Smith dimers. Sparsification methods SPOP (prim.) and SMOP-MGS (mix.) for fixed $\delta = 10^{-3}$, applied to M -by- M Gram matrices defined in Table 4.1, using different scaling options. The Gram matrix is obtained for H_2O under the Coulomb kernel. RMSE is given for the electron-electron (EE) frozen core energy.

$M \backslash \delta$	10^{-2}	10^{-3}	10^{-4}
80	0.0632	0.0624	0.0630
100	0.0676	0.0662	0.0662
120	0.0682	0.0653	0.0658

(a) SPOP (primitive) of Algorithm 2.

$M \backslash \delta$	10^{-2}	10^{-3}	10^{-4}
80	20.30	1.85	0.31
100	164.00	30.51	2.05
120	502.07	132.38	12.29

(b) SMOP-MGS (mixed) of Algorithm 3.

Figure 4.7: Timings (in seconds) for sparsifying M -by- M Gram matrix \mathbf{G}_M (Option A.) obtained from aug-cc-pvdz for H_2O , using a tolerance value δ . We averaged execution time of 3 runs.

4.3.1.3 Timings

We compare timings of different sparsification methods. Table 4.7 presents execution times for the two sparsification schemes under study, applied to matrices of several rank values. Results clearly show that the sparsification of SMOP-MGS is much more expensive than SPOP, by a factor of 1000 on average. This is expected, because SMOP-MGS solves a discrete optimization problem. Moreover, the larger the target tolerance, the longer it takes to run, since the search depends on the path size. Note that the high time cost is not prohibitive in applications where the sparsification can be performed entirely offline, which is the case for rigid fragments.

4.3.1.4 Auxiliary coefficient decay

In the present section, we propose a sparsity analysis attempting to explain the reason behind the better sparsification performance of the SMOP-MGS scheme. We focus on the auxiliary coefficient decay, which we recall, from Section 4.2.3.1, that affects the magnitude of entries in the screening matrix.

To this end, we plot the vector of coefficients \mathbf{v} for which there holds $\mathbf{S} = \mathbf{v}\mathbf{v}^\top$, for \mathbf{S} the screening matrix defined in four methods under comparison, we refer to Table 4.1 for definitions. Figure 4.8 presents the results on the magnitude of entries in \mathbf{v} . Results show that, the SMOP-MGS algorithm based on orthonormalization, applied to both types of scalings, generally creates entries that are more structured than SPOP obtained ones. In particular, the points marked with blue crosses and the ones marked with purple circles, corresponding to mixed orbitals, present a characteristic decay behaviour. This behaviour may be interpreted as encouraging, as it provides numerical evidence that our MGS scheme, which orthonormalizes orbitals preordered

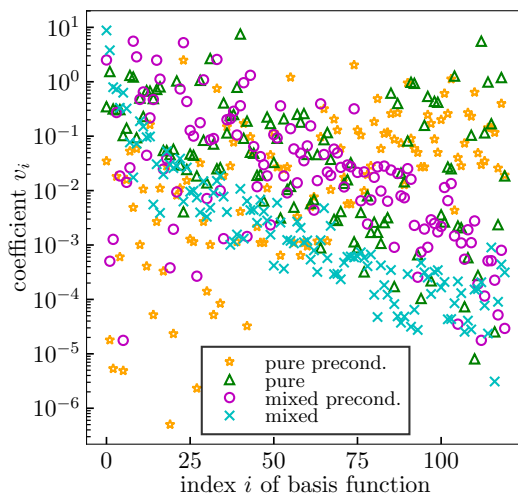


Figure 4.8: aug-cc-pvdz AO basis of H_2O , Gram matrix under the Coulomb metric. Fixed $M = 120$. Entry-wise plot of the vector \mathbf{v} , for the four methods (same labels as in Figure 4.6). The screening matrix \mathbf{S} admits the decomposition $\mathbf{S} = \mathbf{v}\mathbf{v}^\top$, with $v_i := |[c_M]_i w_i|$ (see Table 4.1 for definitions), where any $1 \leq i \leq M$ corresponds to a (mixed or pure) auxiliary basis function.

by PCD pivots, ensures the desired decay on coefficients. This demonstrates the effectiveness of our greedy orthonormalization procedure based on PCD.

4.3.2 Comparison of density fitting methods

The present section is dedicated to the comparison of existing and new methods for the computation of frozen core energies. Namely, our sparse density fitting (SDF) method of Algorithm 3 is compared with conventional (DF) as well as robust (RDF) density fitting. We refer to Section 2.3.3.2 for a definition of state-of-the-art methods. The goal is to present how our sparse density fitting method allows to improve existing methods. Two different admissible auxiliary basis sets have been studied, with different requirements, notably concerning the orbital center restrictions.

4.3.2.1 Atomic auxiliary basis set

In the present experiment, we consider the sparse density fitting scheme SMOP-MGS obtained for the scaling of Option A., which concerned scaling with the density matrix. Recall from Section 4.3.1 that this scheme has provided the best results overall, among the schemes under comparison, in terms of sparsification effectiveness. Let us note that the resulting auxiliary basis set exclusively uses atomic positions as centers, thus this method has physical applications.

Figure 4.9 presents a plot of the RMSE (see Eq. (3.36)) on electron-electron frozen core energy as a function of the number of orbital pairs used in the summation that allows to obtain the energy. Each compared method leads to a different number of retained orbital pairs, obtained by varying input parameters of the given method. For example, for DF and RDF, we vary the target auxiliary basis size M , while for SDF, we vary both M and the sparsity tolerance value δ . Recall that the exact energy used as a reference uses $\mathcal{O}(N_b^4)$ orbital pairs, with $N_b = 18$ (counting by orbital type, see Table 3.1) using the AO basis aug-cc-pvdz for H_2O in the present numerical experiment. Overall, density fitting allows to reduce the number of orbital pairs, as shown by results. In particular, results verify that robust density fitting achieves the lowest errors, that is theoretically expected due to the corrective terms, previously introduced in Eq. (2.31). The cost

however increases significantly, as shown by the high number of orbital pairs. As a result, RDF improves the initial number of orbital pairs only by a factor varying from 5 to 10.

On the other hand, conventional and sparse density fitting methods allow to reach chemical accuracy, i.e. 10^{-3} Hartree, with a low number of orbital pairs than that of robust density fitting. In particular, the use of sparse DF allows to reduce the number of orbital pairs of DF by a numerical factor roughly equal to 10, while reaching the same accuracy as DF, and the total number of orbital pairs used in the exact calculation by a factor varying from 100 to 1000. We may interpret this feature as an encouraging result in favour of sparse density fitting, demonstrating the capabilities of our new method. Comparing ranks used in SDF, the choices $M = 40$ and $M = 80$ present a monotonous decay on orbital pairs with decreasing δ . The choice $M = 120$, however, is shown to give results without significant improvements for lower values of δ . This may mean that we need to further lower values of δ below 10^{-3} to see an improvement.

Up next, more details are provided on the sparsification of a given Gram matrix of dimension M -by- M , by explicitly illustrating its retained and discarded entries. We focus again on SMOP-MGS for scaling Option A.. Figure 4.10 shows entry-wise sparsification results, for varying sizes M and sparsification tolerance values δ . Moreover, it displays the magnitude of the screening matrix entries used in the sparsification. One may first observe an accumulation of large entries on the lower left block of the screening matrix. Such structure is geometrically compatible with the L-shaped sparse pattern, given by Definition 4.2.2. Indeed, the decrease of the sparsification tolerance value causes to expand the boundary of the L-shaped pattern, in a gradual and incremental fashion. We may interpret this geometrical observation as a numerical demonstration of the convergence of our sparsification scheme, introduced in Section 4.2.2.2, with respect to δ , for constructing L-shaped patterns for a fixed M . As a last remark, note that effective range of δ depends on the size M of the screening matrix. For instance, results show that for $M = 40$ the sparsification parameter cannot be set below $\delta = 10^{-3}$, without causing to retain all entries. However, for $M = 120$ one notices that lower values of δ can be possible. This may be a consideration one needs to take into account when performing parameter tuning, for achieving a desired balance between values of M and δ .

4.3.2.2 Auxiliary basis including off-atom centers

The admissible auxiliary basis sets, studied in the present section, concern an unrestricted choice of orbital centers. Note that such consideration may have computational limitations due to the large number of centers, however we consider it here only for testing purposes.

Before discussing results, let us provide the simulation details. In this numerical experiment, we used the SMOP-MGS method with Option C. to perform the sparse density fitting. Recall that, in this setting of the Gram matrix, auxiliary functions are primitives obtained by orbital products, centered on atom or off-atom centers. Moreover, sparsification is applied without any pre-processing of orbital products, thus orbital types may have missing components, resulting in a violation of the orbital symmetry of Eq. (3.21). Note that this setting may have computational limitations when used along with molecular integration schemes applying operations on orbital types. At the same time, due to the fact that orbitals are primitives, the size M of the screening matrix can quickly increase. Hence, we are limited to smaller AO basis sets, than the ones used in Section 4.3.2.1, due to the time consuming execution of SMOP-MGS algorithm for high values of M . We used the 6-31Gs AO basis set for water, with $N_b = 36$. According to Option C., M is the approximate rank obtained by PCD to a N_p^2 -by- N_p^2 Gram matrix (here with tolerance value fixed to 10^{-14}), where $N_p \geq N_b$ is the number of primitives in the AO basis. Our test system is a cluster of 27 rigid water molecules (`watertiny` test in Table 1.1).

Figure 4.11 shows density fitting results, using different methods, on the frozen core RMSE errors (see Eq. (3.36)), for varying number of summed entries used to compute the energy. The sparse summation domain of SDF is obtained for varying tolerance values. Results show that the sparsification allows to reach the accuracy levels of robust DF, while using only the number of orbital pairs of DF. In particular, the binary plots of panel (a) show that the sparsification




Figure 4.9: RMSE on electron-electron (EE) interaction due to density fitting schemes, for ten standardized Smith dimers, using the aug-cc-pvdz AO basis set and the Coulomb kernel metric. The x -axis gives the number of retained orbital pairs used in the EE summation, for a single fragment pair, that is N_b^4 for the exact method, M^2 for DF and $N_b^2 M$ for RDF. Each point on DF and RDF curves corresponds to an auxiliary basis set, in this order: PCD-22, PCD-29, PCD-41, RI, PCD-48, JKFIT, PCD-54, PCD-60, PCD-67. SDF corresponds to SMOP-MGS with scaling Option A., for varying M ranks. Each point on the SDF curves is obtained by varying sparsification tolerance values δ , in this order: 10^1 , 1, 10^{-1} , 10^{-2} , 10^{-3} .

may be interpreted as a reduction of corrective terms of RDF that do not contribute to the total energy. Results of panel (b) then verify that the sparsification does not significantly degrade the accuracy of RDF. This shows the effectiveness of the sparsification. We may conclude that SDF significantly improves state-of-the-art methods in this setting. Lastly, contrary to the setting of Section 4.3.2.1, where we recall that SDF mainly improved DF without reaching the accuracy levels of RDF, now in the present setting, SDF actually improves RDF, by achieving a number of retained orbital pairs roughly 10 times smaller than that of RDF, without accuracy loss.

4.4 Proofs

Proof of Corollary 4.2.0.1. A direct consequence of Cauchy-Schwartz inequality and the distance-preserving property of Eq. (2.10). \square

Proof of Theorem 4.2.1. One regroups terms, then applies the triangle inequality, to write

$$\begin{aligned}
 |E(T_A \rho_A, T_B \rho_B) - E_\delta(T_A \tilde{\rho}_A, T_B \tilde{\rho}_B)| &= \left| \langle T_A \rho_A, T_B \rho_B \rangle - \sum_{(i,j) \in I_\delta} \langle T_A \tilde{\rho}_i^A, T_B \tilde{\rho}_j^B \rangle \right| \\
 &\leq |\langle T_A \rho_A, T_B \rho_B \rangle - \langle T_A \tilde{\rho}_A, T_B \tilde{\rho}_B \rangle| + \\
 &\quad \left| \langle T_A \tilde{\rho}_A, T_B \tilde{\rho}_B \rangle - \sum_{(i,j) \in I_\delta} \langle T_A \tilde{\rho}_i^A, T_B \tilde{\rho}_j^B \rangle \right|.
 \end{aligned}$$

This bound decomposes into two terms. The first term can be estimated using Corollary 2.3.3.1 as

$$|\langle T_A \rho_A, T_B \rho_B \rangle - \langle T_A \tilde{\rho}_A, T_B \tilde{\rho}_B \rangle| \leq \|\rho_B\| \min_{\sigma_A \in X_A} \|\rho_A - \sigma_A\| + \|\rho_A\| \min_{\sigma_B \in X_B} \|\rho_B - \sigma_B\|.$$

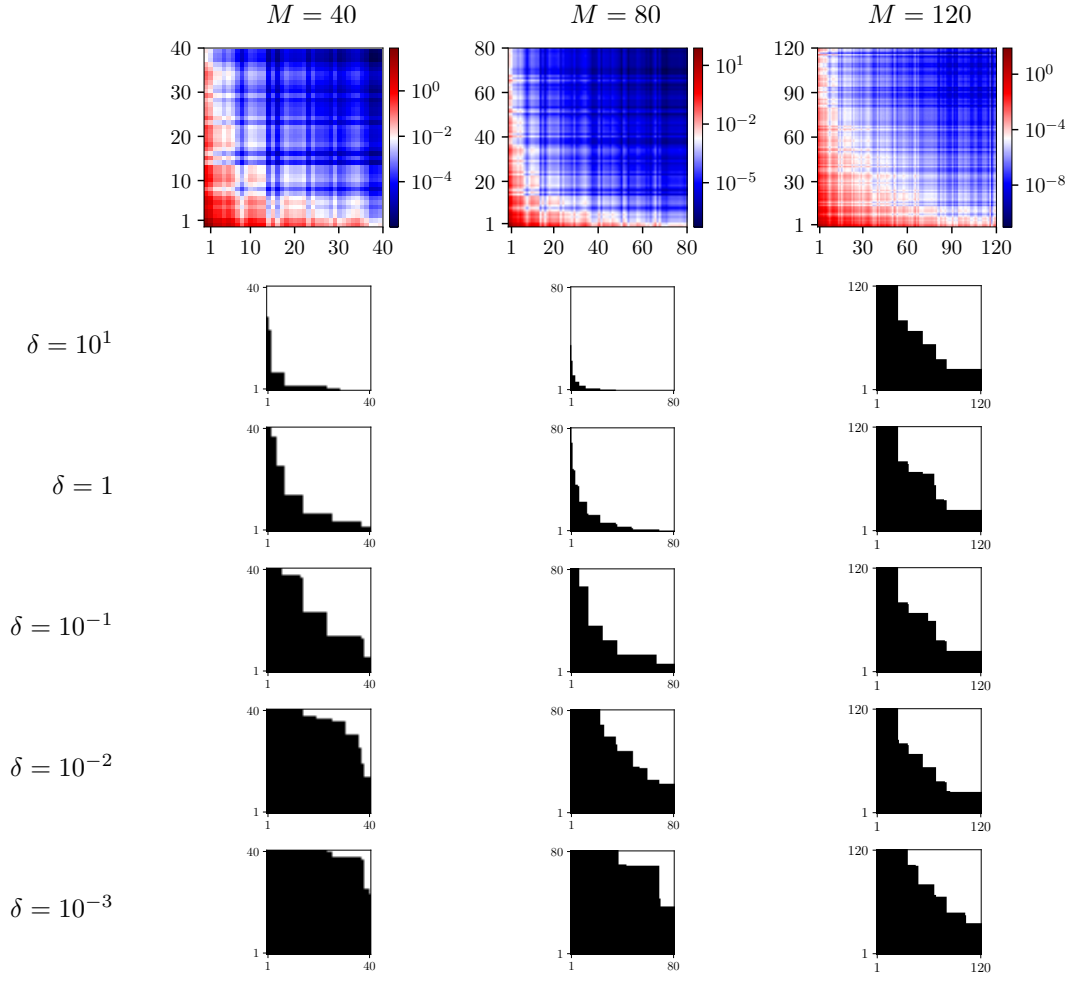


Figure 4.10: aug-cc-pvdz AO basis set for water molecule. M -by- M screening matrix sparsification using SMOP-MGS algorithm with Option **A.**, for varying tolerance threshold δ values. Each binary plot corresponds to a point in the SDF curve of Figure 4.9.

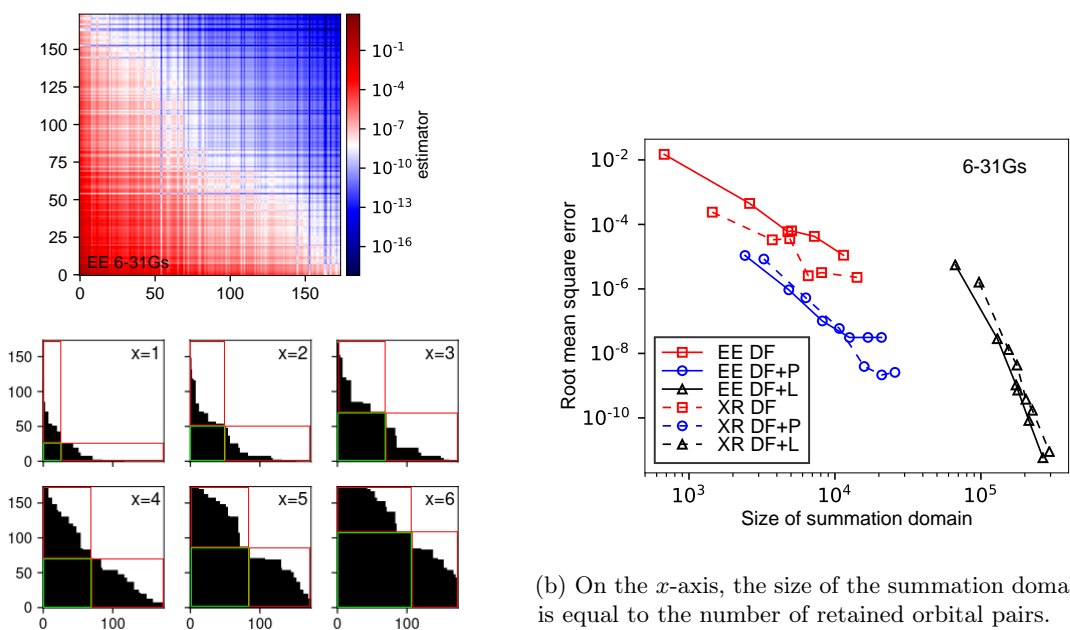
For the second term, we let

$$C := \left| \langle T_A \tilde{\rho}_A, T_B \tilde{\rho}_B \rangle - \sum_{(i,j) \in I_\delta} \langle T_A \tilde{\rho}_i^A, T_B \tilde{\rho}_j^B \rangle \right|.$$

By introducing summations, we may write

$$\begin{aligned} C &= \left| \sum_{i=1}^{M_A} \sum_{j=1}^{M_B} \langle T_A \tilde{\rho}_i^A, T_B \tilde{\rho}_j^B \rangle - \sum_{(i,j) \in I_\delta} \langle T_A \tilde{\rho}_i^A, T_B \tilde{\rho}_j^B \rangle \right| \\ &= \left| \sum_{(i,j) \in I_\delta^c} \langle T_A \tilde{\rho}_i^A, T_B \tilde{\rho}_j^B \rangle \right| \\ &\leq \sum_{(i,j) \in I_\delta^c} \|\tilde{\rho}_i^A\| \|\tilde{\rho}_j^B\| \end{aligned}$$

where we applied Cauchy-Schwartz inequality and employed the integral screening result of Lemma 2.1.1 based on isometries. This concludes the proof of Theorem 4.2.1. \square



(a) Bottom: binary plots for M -by- M matrix. Maximal M' -by- M' square is shown in green and corrective rectangles of RDF in red.

Figure 4.11: RMSE on electron-electron (EE) and exchange-repulsion (XR) frozen core interactions due to density fitting schemes, for 27 water cluster. Methods are tuned in this order: SDF (DF+P) first generates an L-pattern from a M -by- M screening matrix, using the SMOP-MGS method with Option C. ($M = 175$ for EE, $M = 192$ for XR); DF (DF) then uses the maximal M' -by- M' square included in the L-pattern; RDF (DF+L) uses same as DF plus additional corrective terms. Each point of the SDF curve corresponds to a binary plot with tolerance $\delta = 10^{-x}$. The density fit for each energy is obtained using the optimal kernel (Coulomb for EE and Dirac for XR).

Proof of Theorem 4.2.2. We propose a proof by induction. The main enumeration argument of the proof is illustrated in Figure 4.12. If I is a given summation domain, then we start counting from the first column indexed by i_1 . The number of elements in this column is exactly j_1 . By symmetry, we do not need to count elements of the row i_1 . At the next iteration, we move to column i_2 . There are j_2 row elements in this column. By symmetry, we have already computed the first row and column indexed by i_1 and j_1 respectively. In order to not double count elements, all elements of the i_2 column below the diagonal will not be counted. The number of entries of the i_2 column that are above the diagonal, including the diagonal entry, is equal to $j_2 - i_2 + 1$. We stop counting when we hit the diagonal since all terms will have been enumerated by then. Note that the $j_2 - i_2 + 1$ entries must be computed $i_2 - i_1$ times, one time for each vertical column indexed by the integer k for all $i_1 < k \leq i_2$. Such an example is provided in Figure 4.3c. Focusing on the green rectangle associated to the second boundary point in this example, one has $i_1 = 1$ and $i_2 = 3$, hence we must take into account the two columns indexed by $k = 2, 3$ when counting terms associated to i_2 . This argument yields the desired closed-form expression of Theorem 4.2.2. \square




Figure 4.12: Cost enumeration pattern at iteration k , for dotted symmetric domain. Nodes are marked in blue: leading node of the current iteration belonging to the horizontal boundary; red: nodes enumerated in current iteration; gray: already counted nodes. For $k = 4$ the diagonal leading node has zero contribution.

4.5 Conclusion and perspectives

In this work, we develop sparse factorizations of atomic orbital Gram matrices, for use in frozen core interaction energy calculations. Compared with standard entry-wise truncation methods, the major characteristic of this specific use is allowing to sparsify based on sum of entries truncation. In order to proceed, we develop heuristic techniques for increasing sparsity of matrices based on a modified Gram-Schmidt orthonormalization procedure and a Cauchy-Schwartz bound for prescreening sums of entries. We also construct sparsity patterns that minimize the number of retained entries. The result of this sparsification procedure is a sparse intermolecular interaction matrix. We plan to carry out a numerical verification of the complexity gain of this sparse interaction assembly in practice. One difficulty to consider while implementing sparse interaction matrix assembly is that the sparsification might not be compatible with caching, i.e. reusing computed quantities, during integration. Lastly, notice that we didn't respect all atomic orbital requirements in this section for convenience, namely our auxiliary basis stored does not respect atom-wise consistency. This point is left to be handled for future practical implementations.

The entry-wise assembly of the Gram matrix based on orbital products is then studied. From a practical point of view, the results on sparsification of interaction matrices show that linear-time assembly schemes are possible, in terms of the total number of orbital products, using orbital bases obtained by our modified Gram-Schmidt technique. We interpret this as a promising result for constructing orbital bases of maximal sparsity. We demonstrate the practical use of such interaction sparsification for the ten standardized Smith dimers. Besides the standard interaction assembly, we plan to explore sparse assembly techniques for future implementations. Careful implementations may lead to further computational efficiency. Lastly, it could be interesting to investigate if our proposed framework enables to choose admissible auxiliary basis functions that lead to a sparse structure of the Gram matrix underlying the interaction.

Chapter 5

Reduced basis methods for density fitting

Contents

5.1	Introduction	112
5.2	Theory	113
5.2.1	Setting	113
5.2.2	Reduced-auxiliary basis generation	114
5.2.3	Auxiliary coefficient interpolation	115
5.2.4	Summary of Reduced Basis Density Fitting (RB-DF) scheme	117
5.2.5	Computational complexity	117
5.3	Numerical results	119
5.3.1	Flexible water clusters	120
5.3.2	Reduced-auxiliary basis accuracy	120
5.3.3	Interpolation error	122
5.3.4	Pairwise interaction errors	124
5.4	Conclusion and perspectives	126

The present chapter deals with the topic of applying density fitting methods to a given *ab initio* electronic density, expanded on an atomic orbital basis of *a priori* unknown centers. Such centers are assumed to take arbitrary values during the course of a molecular dynamics simulation. The main difficulty in this case is that interatomic distances and angles within a molecular fragment take *a priori* unknown values. We are interested in methods that allow to reduce this new complexity, which is the intramolecular geometric conformation space, and still provide fitted densities at a low computational cost. We explore reduced basis approaches for constructing auxiliary basis sets, as well as associated coefficients. Within the offline-online simulation framework, our auxiliary basis is precomputed entirely offline, while auxiliary coefficients are recovered online by performing Lagrange interpolation on a small number of snapshots. The performance of our method is illustrated by running tests on ten Smith dimers, with novel promising results, showing that, thanks to the reduced basis approach, one can successfully extend density fitting methods to non-rigid fragments.

5.1 Introduction

Molecular simulations using polarizable force fields update geometrical conformation of fragments at every timestep, according to the evolution of dynamics (see Appendix B.2). This creates a difficulty in density fitting, since the fragment *ab initio* density is expressed on basis sets centered on atomic positions. By updating conformations, interatomic distances become *a priori* unknown during the offline stage, where density fitting is typically computed (we refer to Section 2 for the standard density fitting procedure). In other words, the local to global transformation, introduced in Section 2.1.2.3, is not an isometry, as it was the case in the previous Chapters 3 and 4.

Thus density fitting needs to be recalculated online for every updated fragment geometry. This is computationally prohibitive since it would require knowledge of the *ab initio* density function. In the present chapter, our goal is to extend density fitting methods to treat non-rigid fragments of flexible geometry, such as AMOEBA water clusters (see Section 1.3.1).

Existing methods are based on empirical considerations. In practical molecular simulations within the Gaussian electrostatic model framework (we refer to Section 1.1.6.4 for more details), auxiliary densities are *averaged* over a large number of fragment geometrical conformations, before being used on AMOEBA fragments of different interatomic geometries at every MD timestep. In particular, in [241, Section II], the fitted densities are obtained by averaging molecular densities of 500 water monomers extracted by dimers PES. A similar strategy has been employed in [84, Section G], where coefficients calculated from *ab initio* density matrices of the same fragments in 99 random conformations were averaged, before performing density fitting. Note that such approaches attempt to achieve *transferability* of density fitting for force field applications (see Section B.2.2). The limitation of this approach is that in reality not all conformations have equal significance as representative samples. A more suitable technique would be to account for a weighted average, where the weights depend on the importance of the configuration. A systematic method to achieve such technique is the reduced basis method.

The Reduced Basis (RB) method [270, 155] has been introduced in the context of partial differential equations (PDE). It applies to linear or non-linear PDEs that depend on several parameters, such as physical constants. RB accounts as a model order reduction technique for lowering the computational cost of repeated resolution of PDEs. In particular, let \mathcal{S} be the solution manifold consisting of all PDE solutions associated to all parameter values. The idea is to select few parameters μ_1, \dots, μ_n and approximate any other solution $u(\mu)$ in \mathcal{S} by a linear combination of functions in the set $\{u(\mu_i)\}_{1 \leq i \leq n}$, which forms a so-called *reduced basis* of the solution space. This approach greatly reduces the computational cost of repeated PDE resolutions, as the calculation of the solution for arbitrary parameter is obtained at very low computational cost from the reduced basis, avoiding the expensive full resolution of the PDE for the given parameter.

In the present work, the reduced basis method is applied to the repeated calculations of density fitting. A relative work on the reduced basis approach applied to chemistry is [260, 261], that uses empirical interpolation for improving the density matrix initial guess in self-consistent field calculations to accelerate *ab initio* molecular dynamics simulations, based on reusing converged density matrices at different geometries of several PES points.

The remaining content of the present chapter is organized as follows. After introducing the problem setting in Section 5.2.1, the framework of reduced-auxiliary basis set is presented in Section 5.2.2. Our main contribution is an extension of our standard auxiliary basis generation method, presented in Chapter 3, to the context of flexible fragments. This contribution is completed with the Lagrange interpolation strategy for reusing auxiliary coefficients at low computational cost, introduced in Section 5.2.3. The summary of our main result, which is a new Reduced Basis Density Fitting (RB-DF) scheme, is given in Section 5.2.4. The computational complexity model of our scheme is deduced in Section 5.2.5. Section 5.3 is dedicated to numerical experiments, showing the capabilities of the RB-DF method applied to the ten Smith dimers. The last part of this chapter contains conclusions and perspectives for future use of our methods,

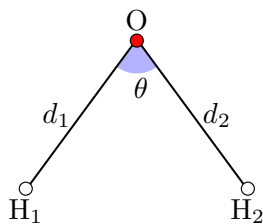


Figure 5.1: Intramolecular coordinates $\boldsymbol{\mu} = (d_1, d_2, \theta)$ of water molecule.

included in Section 5.4.

5.2 Theory

5.2.1 Setting

Consider a cluster consisting of molecular fragments. It is assumed for simplicity that the atom number in fragments is constant. At every timestep of a molecular dynamics (MD) simulation, a fragment's geometrical configuration can be assessed using an inner coordinate parametrization. Let us denote an arbitrary intramolecular geometry parametrization by a real vector of size d , for some $d \in \mathbb{N}^*$, i.e.

$$\boldsymbol{\mu} = (\mu_1, \mu_2, \dots, \mu_d) \in \mathbb{R}^d =: \mathcal{M}.$$

The set \mathcal{M} is called *parameter space* and contains all admissible intramolecular atomic conformations, that, for this work, there is no need to further restrict with additional structural assumptions. Figure 5.1 shows an inner coordinate example for the case of water molecule.

In the context of Gaussian electrostatic model for arbitrary fragments (see Section 2.1.2.1 for more details), the *ab initio* electronic density of a fragment at conformation $\boldsymbol{\mu}$ is computed by solving the discretized electronic ground state problem for the Hamiltonian $\mathcal{H}(\boldsymbol{\mu})$, depending only on the atomic positions of the fragment configuration $\boldsymbol{\mu}$. The electronic density function is written in the atomic orbital discretization basis as

$$\rho(\boldsymbol{\mu}) = \sum_{\mu, \nu=1}^{N_b} \mathbf{D}_{\mu\nu}(\boldsymbol{\mu}) \chi_\mu(\boldsymbol{\mu}) \chi_\nu(\boldsymbol{\mu}),$$

where $\{\chi_\mu(\boldsymbol{\mu})\}_{1 \leq \mu \leq N_b}$ are Gaussian atomic orbitals centered on atomic configuration $\boldsymbol{\mu}$ (represented in the global frame) and $\mathbf{D}(\boldsymbol{\mu}) \in \mathbb{R}^{N_b \times N_b}$ is the density matrix obtained by an *ab initio* calculation on the AO basis associated to $\boldsymbol{\mu}$. For simplicity, it is assumed that all fragments have the same type of atomic orbital basis sets of constant size N_b .

We are interested in the problem of calculating fitted densities approximating the *ab initio* ones, i.e. finding auxiliary basis functions $\xi_1(\boldsymbol{\mu}), \dots, \xi_M(\boldsymbol{\mu})$ and auxiliary coefficients $c_1(\boldsymbol{\mu}), \dots, c_M(\boldsymbol{\mu})$ for a prescribed $M \ll N_b^2$, such that the following approximation holds:

$$\forall \boldsymbol{\mu} \in \mathcal{M}, \quad \rho(\boldsymbol{\mu}) \approx \tilde{\rho}(\boldsymbol{\mu}) = \sum_{i=1}^M c_i(\boldsymbol{\mu}) \xi_i(\boldsymbol{\mu}).$$

The question that concerns us is, how can we calculate as much information on the auxiliary functions and coefficients as possible, without using *a priori* knowledge of the parameter value $\boldsymbol{\mu}$. Such calculation can thus be performed during an offline simulation phase. The goal is then to use this precalculation in order to perform only few additional computationally non-expensive steps to recover the fitted density $\tilde{\rho}(\boldsymbol{\mu})$ online. Note that our auxiliary basis set is assumed to

have a fixed size M for all geometries $\boldsymbol{\mu}$. To summarize, our main interest is the dimensionality reduction of the set

$$\mathcal{S} := \{\tilde{\rho}(\boldsymbol{\mu}) : \boldsymbol{\mu} \in \mathcal{M}\},$$

called the *solution space*.

5.2.2 Reduced-auxiliary basis generation

Let us focus on the dependence of the atomic orbital basis with respect to the parameter $\boldsymbol{\mu} \in \mathcal{M}$. A Gaussian AO depends explicitly on $\boldsymbol{\mu}$, in terms of the Gaussian center. This dependence is highly non-linear and known *a priori*, i.e.

$$\chi_{\mu}(\mathbf{r}; \boldsymbol{\mu}) = \chi_{\mu}(\mathbf{r} - \mathbf{R}(\boldsymbol{\mu})), \quad (5.1)$$

for some atomic center $\mathbf{R}(\boldsymbol{\mu}) \in \mathbb{R}^3$. The same holds for the atomic Gaussian auxiliary basis. An affine decomposition of the fitted density is introduced on such a basis, as follows. The auxiliary basis generation problem in the context of Section 5.2.1 consists of defining, given a target M , a finite parameter set $\mathcal{M}_r = \{\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_r\} \subset \mathcal{M}$ such that the following approximation is well-defined:

$$\forall \boldsymbol{\mu} \in \mathcal{M}, \quad \rho(\boldsymbol{\mu}) \approx \tilde{\rho}(\boldsymbol{\mu}) = \sum_{i=1}^M c_i(\boldsymbol{\mu}; \boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_r) \xi_i(\boldsymbol{\mu}; \boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_r), \quad (5.2)$$

with $\mathbf{c}(\boldsymbol{\mu}; \boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_r) \in \mathbb{R}^M$ a coefficient vector and $\{\xi_i(\boldsymbol{\mu}; \boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_r)\}_{1 \leq i \leq M}$ the reduced-auxiliary basis, which is the *reduced basis* of the solution space. Note that the reduced basis depends on the arbitrary parameter $\boldsymbol{\mu}$ through the *a priori* determined explicit formula of Eq. (5.1) and on the finite parameter set \mathcal{M}_r in a manner left to be determined.

A heuristic method for constructing such approximation of Eq. (5.2) consists of the following main stages. First, we *a priori* fix a finite number of points in the parameter space \mathcal{M} , denoted by $\mathcal{M}_r = \{\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_r\}$. Note that, in the present work, the construction of the finite parameter space \mathcal{M}_r is based on naive preliminary methods, detailed in the following Section 5.3. Then, we construct the reduced basis using the finite subset \mathcal{M}_r , based on the methods detailed in the present section. At a later stage, auxiliary coefficients on this reduced basis are computed by employing the methodologies given in Section 5.2.3.

Let us focus on the reduced-auxiliary basis construction. Before describing our algorithm in detail, let us quickly outline its motivation. The main idea of our method is to use a greedy technique to construct an auxiliary basis set that reproduces the span of the set of *ab initio* densities $\rho(\boldsymbol{\mu})$ associated to all parameter values $\boldsymbol{\mu}$ under some given tolerance value. The main tool to achieve this is the incomplete pivoted Cholesky decomposition (PCD), as formulated in [146]. PCD can be interpreted as a modified Gram-Schmidt orthonormalization that first treats functions that are more linearly dependent. In that sense, it is a greedy method. Our idea is to apply PCD to all Gram matrices associated to parameters belonging to \mathcal{M}_r . Let us emphasize that PCD is hence repeated multiple times, in this setting. The Gram matrices evaluated for all points of the finite parameter set \mathcal{M}_r are called *snapshots*.

In the present work, the reduced basis construction given by Algorithm 4 is essentially obtained as a variation to [146, Algorithm 1], by including an additional minimization step (see step 4). This step performs the greedy selection, which is repeated for each reduced basis function, while taking into account all snapshots of the finite family \mathcal{M}_r . The greedy selection is essentially the repeated Cholesky pivot selection, based on the maximal diagonal term of the residual matrix. However, instead of having a single residual matrix to manipulate, now we have r matrices, one for every point of the finite parameter subset \mathcal{M}_r . At the pivot selection step, the pivot is made significant for *all* r matrices. This is achieved by the minimization step of step 4. The minimization assures that there is no point $\boldsymbol{\mu}_r$ for which the selected pivot is insignificant. It also prevents adding pivots that correspond to zero diagonal entries, if another choice is

possible. This way, at every iteration, we assure that the function added into the reduced basis is the most significant possible for all r . Moreover, we imposed a target size criterion instead of a tolerance threshold, because the tolerance is sensitive to the Gram matrix coefficient scale, which may slightly vary along snapshots. An advantage of the minimization of step 4 is that it is independent of the Gram matrix scale. Note that, in the input of Algorithm 4, it is assumed that GTO-products have the same row (and column) index in all Gram snapshots, so that the operation of step 6 is well-defined.

Algorithm 4: Reduced-auxiliary basis greedy construction (RB-AUX).

Input: Finite parameter set $\mathcal{M}_r = \{\boldsymbol{\mu}_u\}_{1 \leq u \leq r}$, GTO-product Gram matrix snapshots $\mathbf{G}(\boldsymbol{\mu}_u) \in \mathbb{R}^{n \times n}$, for all $1 \leq u \leq r$, and target reduced basis size $M \ll n$.
Output: Reduced-auxiliary basis $\{\xi_1, \dots, \xi_M\}$.

- 1: Set $d(\boldsymbol{\mu}_u) := \text{diag}(\mathbf{G}(\boldsymbol{\mu}_u))$, $1 \leq u \leq r$.
- 2: Initialize $\pi := (1, 2, \dots, n)$.
- 3: **for** $1 \leq m \leq M$ **do**
- 4: Set
$$i_0 := \arg \max_{m \leq i \leq n} \min_{1 \leq u \leq r} d_{\pi_i}(\boldsymbol{\mu}_u).$$
- 5: Swap π_m and π_{i_0} .
- 6: Set ξ_m equal to the GTO-orbital function corresponding to the π_m -row index.
- 7: Set $L_{m, \pi_m}(\boldsymbol{\mu}_u) := \sqrt{d_{\pi_m}(\boldsymbol{\mu}_u)}$, $1 \leq u \leq r$.
- 8: **for** $m+1 \leq i \leq n$ **do**
- 9: Compute
$$L_{m, \pi_i}(\boldsymbol{\mu}_u) = \left(G_{\pi_m, \pi_i}(\boldsymbol{\mu}_u) - \sum_{j=1}^{m-1} L_{j, \pi_m}(\boldsymbol{\mu}_u) L_{j, \pi_i}(\boldsymbol{\mu}_u) \right) / L_{m, \pi_m}(\boldsymbol{\mu}_u), \quad 1 \leq u \leq r.$$
- 10: Update $d_{\pi_i}(\boldsymbol{\mu}_u) := d_{\pi_i}(\boldsymbol{\mu}_u) - L_{m, \pi_i}^2(\boldsymbol{\mu}_u)$, $1 \leq u \leq r$.
- 11: **end for**
- 12: **end for**
- 13: **return** functions $\{\xi_m : 1 \leq m \leq M\}$.

5.2.3 Auxiliary coefficient interpolation

In the present section, we focus on the dependence of the auxiliary coefficients with respect to the parameter $\boldsymbol{\mu} \in \mathcal{M}$. Once the reduced-auxiliary basis is available, by applying methods described in Section 5.2.2, one can solve a least-squares problem for every parameter $\boldsymbol{\mu}$ for finding the auxiliary coefficients, i.e.

$$\forall \boldsymbol{\mu} \in \mathcal{M}, \quad \mathbf{c}(\boldsymbol{\mu}) := \arg \min_{\mathbf{c} \in \mathbb{R}^M} \left\| \rho(\boldsymbol{\mu}) - \sum_{i=1}^M c_i \xi_i(\boldsymbol{\mu}) \right\|^2 \quad (5.3)$$

with $\|\cdot\|$ being a user-defined error norm. Then the fitted density is recovered using the expansion of Eq. (5.2). The resolution of this minimization problem is prohibitive during the online stage, because knowledge of the *ab initio* electronic density at conformation $\boldsymbol{\mu}$ is not available without paying the cost of solving a full electronic structure problem, at $\mathcal{O}(N_b^4)$. For this reason, alternative methods are considered. The main approach of the present work is to introduce an intermediary decomposition of auxiliary coefficients, reading

$$\forall \boldsymbol{\mu} \in \mathcal{M}, \quad \forall 1 \leq i \leq M, \quad c_i(\boldsymbol{\mu}) \approx \sum_{u=1}^r a_u(\boldsymbol{\mu}) c_i(\boldsymbol{\mu}_u), \quad (5.4)$$

with $\mathbf{a}(\boldsymbol{\mu}) = (a_1(\boldsymbol{\mu}), \dots, a_r(\boldsymbol{\mu})) \in \mathbb{R}^r$ depending on arbitrary $\boldsymbol{\mu}$ and $c_1(\boldsymbol{\mu}_1), \dots, c_M(\boldsymbol{\mu}_r)$, depending only on the finite parameter set \mathcal{M}_r , under the assumption that $\mathbf{a}(\boldsymbol{\mu})$ is computationally much

cheaper to compute than the true auxiliary coefficients $\mathbf{c}(\boldsymbol{\mu})$. Approximation of Eq. (5.4) can be interpreted as an interpolation on auxiliary coefficient snapshots. Its main purpose is to allow precomputing the finite number of snapshots entirely offline.

Let us introduce a strategy for computing this approximation. We propose to represent each coefficient c_i of the auxiliary vector $\mathbf{c}(\boldsymbol{\mu})$ as a real multivariate continuous function in d variables, i.e.

$$\forall \boldsymbol{\mu} \in \mathcal{M}, \quad \forall 1 \leq i \leq M, \quad c_i(\boldsymbol{\mu}) = c_i(\mu_1, \mu_2, \dots, \mu_d) : \mathbb{R}^d \rightarrow \mathbb{R}. \quad (5.5)$$

A convenient approach for solving the interpolation problem of Eq. (5.4), provided that the interpolation domain is the tensorization of intervals (that is a cuboid), consists in using a tensorized grid formed by 1D interpolation grids along each interval (see [123]). Let $C \subset \mathbb{R}^d$ be a cuboid such that

$$C = [a_1, b_1] \times \dots \times [a_d, b_d],$$

with $a_k \leq b_k$, for all $1 \leq k \leq d$. The tensorized rule is composed of nodes for any multi-index $\mathbf{l} = (l_1, \dots, l_d)$, defined by

$$\boldsymbol{\mu}_{\mathbf{l}} := (\mu_{l_1}^{(1)}, \dots, \mu_{l_d}^{(d)}).$$

Assuming that the admissible geometry parameters are cuboids, these tensorized interpolation rules can be used for each parameter of the chemical space.

Now, for the sake of simplicity, we may use multivariate Lagrange interpolation to approximate the number of M functions in Eq. (5.5). For fixed $d \in \mathbb{N}^*$, let us first consider a given grid of the d -dimensional cuboid C , of size $n_1 \times n_2 \times \dots \times n_d$, with $n_k \in \mathbb{N}^*$ a given number of nodes in the interval $[a_k, b_k]$ for all $1 \leq k \leq d$. The grid is defined as the tensorized set of points

$$\left\{ (\mu_{l_1}^{(1)}, \mu_{l_2}^{(2)}, \dots, \mu_{l_d}^{(d)}) \in \mathbb{R}^d : 1 \leq l_k \leq n_k, \quad 1 \leq k \leq d \right\} \subset C.$$

The multivariate Lagrange interpolation of any given function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ on the grid points is defined as

$$\forall (\mu_1, \dots, \mu_d) \in \mathbb{R}^d, \quad \mathcal{I}[f(\mu_1, \dots, \mu_d)] := \sum_{l_1=1}^{n_1} \sum_{l_2=1}^{n_2} \dots \sum_{l_d=1}^{n_d} \prod_{k=1}^d L_{l_k}(\mu_k) f(\mu_{l_1}^{(1)}, \mu_{l_2}^{(2)}, \dots, \mu_{l_d}^{(d)}), \quad (5.6)$$

where $L_{l_k} : \mathbb{R} \rightarrow \mathbb{R}$ is the Lagrange polynomial of degree l_k in one dimension, given by

$$\forall \mu \in \mathbb{R}, \quad L_{l_k}(\mu) = \prod_{\substack{l=0 \\ l \neq l_k}}^{n_k} \frac{\mu - \mu_l^{(k)}}{\mu_{l_k}^{(k)} - \mu_l^{(k)}}, \quad (5.7)$$

satisfying the property $L_{l_k}(\mu_l^{(k)}) = \delta_{l_k l}$, with $\delta_{l_k l}$ the Kronecker delta. Due to this last property of Lagrange polynomials, there holds

$$\mathcal{I}[f(\boldsymbol{\mu}_{\mathbf{l}})] = f(\boldsymbol{\mu}_{\mathbf{l}}),$$

in other words, the approximation of f by its interpolant $\mathcal{I}[f]$ is exact on the grid points $\boldsymbol{\mu}_{\mathbf{l}}$. Note that more efficient approaches are given by the magic points of the Empirical Interpolation Method (EIM) [217]. The latter are more natural in the context of reduced basis methods but require more investigation.

Applying the above interpolation to every single c_i coefficient of the auxiliary vector, one obtains the following approximation. We propose to use the same finite parameter space for interpolating and for sampling snapshots. Given a finite family \mathcal{M}_r of size $r \in \mathbb{N}^*$, the tensorized interpolation grid, of given degree $n_k \in \mathbb{N}^*$ at the k -th component, can be defined by the full finite parameter space of dimension $n_1 \times n_2 \times \dots \times n_d$, given by

$$\left\{ \boldsymbol{\mu}_u : \boldsymbol{\mu}_u = (\mu_{u_1}^{(1)}, \mu_{u_2}^{(2)}, \dots, \mu_{u_d}^{(d)}) \in \mathcal{M}_r, \quad 1 \leq u \leq r, \quad 1 \leq u_k \leq n_k, \quad 1 \leq k \leq d \right\}. \quad (5.8)$$

The approximation of Eq. (5.5) of the auxiliary coefficient vector then takes the form of an interpolation on snapshots, where the cheaper-to-compute coefficients $a_u(\boldsymbol{\mu})$ are given as a tensorized Lagrange polynomial, i.e.

$$\forall \boldsymbol{\mu} \in \mathcal{M}, \quad \forall 1 \leq i \leq M, \quad c_i(\boldsymbol{\mu}) \approx \mathcal{I}[c_i(\boldsymbol{\mu})] = \sum_{u=1}^r a_u(\boldsymbol{\mu}) c_i(\boldsymbol{\mu}_u), \quad (5.9)$$

with

$$\begin{aligned} \mathcal{I}[c_i(\boldsymbol{\mu})] &= \sum_{u_1=1}^{n_1} \sum_{u_2=1}^{n_2} \cdots \sum_{u_d=1}^{n_d} \prod_{k=1}^d L_{u_k}(\mu_k) c_i(\mu_{u_1}^{(1)}, \mu_{u_2}^{(2)}, \dots, \mu_{u_d}^{(d)}) \\ &= \sum_{u=1}^r a_u(\boldsymbol{\mu}) c_i(\boldsymbol{\mu}_u), \quad a_u(\boldsymbol{\mu}) := \prod_{k=1}^d L_{u_k}(\mu_k), \end{aligned}$$

by flattening the multiple summation into a single sum of degree $r = n_1 n_2 \cdots n_d$. Such an approximation suggests that it suffices to precompute auxiliary coefficient snapshots $c_i(\boldsymbol{\mu}_u)$, for all $1 \leq i \leq M$ and $1 \leq u \leq r$, by solving the least-squares problem of Eq. (5.10) only on the small *a priori* known parameter space \mathcal{M}_r , i.e.

$$\forall 1 \leq u \leq r, \quad \mathbf{c}(\boldsymbol{\mu}_u) := \arg \min_{\mathbf{c} \in \mathbb{R}^M} \left\| \rho(\boldsymbol{\mu}_u) - \sum_{i=1}^M c_i \xi_i(\boldsymbol{\mu}_u) \right\|^2, \quad (5.10)$$

during the offline stage. Then, for arbitrary parameter $\boldsymbol{\mu}$, the approximate auxiliary coefficient can be recovered using Eq. (5.9). Moreover, note that the intermediary coefficient $a_u(\boldsymbol{\mu})$ in this approximation only depends on the grid and on the parameter $\boldsymbol{\mu}$, being independent of the auxiliary basis size M . It therefore needs to be calculated once and then reused for all $1 \leq i \leq M$. The computational cost of the offline precomputation is further discussed in the following section.

5.2.4 Summary of Reduced Basis Density Fitting (RB-DF) scheme

In the present section, we summarize our main contribution, which is a new Reduced Basis Density fitting scheme, with applications to polarizable molecular fragments.

Given an *ab initio* electronic density function ρ associated to a molecular fragment chemical type, our scheme constructs the approximation

$$\rho(\boldsymbol{\mu}) \approx \tilde{\rho}(\boldsymbol{\mu}) = \sum_{i=1}^M \mathcal{I}[c_i(\boldsymbol{\mu})] \xi_i(\boldsymbol{\mu}). \quad (5.11)$$

Let us assume that the target interpolation degree in the auxiliary coefficient generation is n for each parameter dimension. The interpolation grid, given as Eq. (5.8), then has n grid points in each dimension, with a total number of $r = n^d$ grid points, the same as the finite parameter samples in the sample space \mathcal{M}_r . The following Algorithm 5 provides offline precomputations, while Algorithm 6 provides the online approximate density computation for arbitrary parameter value.

5.2.5 Computational complexity

The computational complexity cost of the density fitting scheme, based on the reduced-auxiliary basis generation of Section 5.2.2 and the auxiliary coefficient generation of Section 5.2.3, may be theoretically estimated as follows. For each one of the point samples in \mathcal{M}_r constructed in step 2 of Algorithm 5, a Gram matrix snapshot is required as input in Algorithm 4 and an auxiliary

Algorithm 5: Reduced basis density fitting scheme (RB-DF – OFFLINE PHASE).

Input: *Ab initio* electronic density ρ of fragment chemical type, n target sample size, d intramolecular parametrization size, $\{[a_k, b_k]\}_{1 \leq k \leq d}$ admissible geometry intervals, target auxiliary basis size M .

Output: Finite sample space, auxiliary basis set, auxiliary coefficient snapshots.

- 1: Construct uniform grid $(\mu_{l_k}^{(k)})_{1 \leq l_k \leq n}$ of size n on the interval $[a_k, b_k]$, for all $1 \leq k \leq d$.
- 2: Define finite parameter set $\mathcal{M}_r = \{\boldsymbol{\mu}_u : 1 \leq u \leq r\}$, with $\boldsymbol{\mu}_u = (\mu_{l_1}^{(1)}, \mu_{l_2}^{(2)}, \dots, \mu_{l_d}^{(d)})$ and $1 \leq l_k \leq n$ for all $1 \leq k \leq d$, of total size $r := n^d$.
- 3: Compute Gram matrix snapshots $\mathcal{G}_r = \{\mathbf{G}(\boldsymbol{\mu}_u) : 1 \leq u \leq r\}$ given by Eq. (3.33).
- 4: Call RB-AUX($\mathcal{M}_r, \mathcal{G}_r, M$) of Algorithm 4, to obtain the reduced-auxiliary basis set $\boldsymbol{\xi} := \{\xi_1, \dots, \xi_M\}$.
- 5: Compute *ab initio* electronic density snapshots $\{\rho(\boldsymbol{\mu}_u) : 1 \leq u \leq r\}$.
- 6: Compute auxiliary coefficient snapshots $\mathcal{C}_r = \{\mathbf{c}(\boldsymbol{\mu}_u) : 1 \leq u \leq r\}$ by solving the least-squares problem

$$\forall 1 \leq u \leq r \in \mathcal{M}, \quad \mathbf{c}(\boldsymbol{\mu}_u) := \arg \min_{\mathbf{c} \in \mathbb{R}^M} \left\| \rho(\boldsymbol{\mu}_u) - \sum_{i=1}^M c_i \xi_i(\boldsymbol{\mu}_u) \right\|^2.$$

- 7: **return** parameter samples \mathcal{M}_r , $\boldsymbol{\xi}$ and auxiliary coefficient snapshots \mathcal{C}_r .

Algorithm 6: Reduced basis density fitting scheme (RB-DF – ONLINE PHASE).

Input: Arbitrary parameter $\boldsymbol{\mu} = (\mu_1, \dots, \mu_d) \in \mathcal{M}$, precomputed auxiliary basis $\boldsymbol{\xi}$ and coefficients $\mathcal{C}_r = \{(c_1(\boldsymbol{\mu}_u), \dots, c_r(\boldsymbol{\mu}_u)) : 1 \leq u \leq r\}$.

Output: Approximate density $\tilde{\rho}(\boldsymbol{\mu})$ on auxiliary basis.

- 1: Compute $a_u(\boldsymbol{\mu}) = \prod_{k=1}^d L_{u_k}(\mu_k)$, $1 \leq u \leq r$.
- 2: Interpolate

$$\forall 1 \leq i \leq M, \quad \tilde{c}_i(\boldsymbol{\mu}) := \sum_{u=1}^r a_u(\boldsymbol{\mu}) c_i(\boldsymbol{\mu}_u).$$

- 3: **return** $\tilde{\rho}(\boldsymbol{\mu}) = \sum_{i=1}^M \tilde{c}_i(\boldsymbol{\mu}) \xi_i(\boldsymbol{\mu})$.

coefficient snapshot for the interpolation 5.9. Such a computation cost requirement may be summarized as follows:

$$\begin{aligned}
\text{AO Gram matrix snapshots (offline):} & \quad \mathcal{O}(n^d N_b^4) \\
\text{RB construction (offline):} & \quad \mathcal{O}(n^d M N_b^2) \\
\text{auxiliary coefficient snapshots (offline):} & \quad \mathcal{O}(n^d N_b^4 + n^d M N_b^2 + n^d M^2) \\
\text{auxiliary coefficients (online):} & \quad \mathcal{O}(n^d M + n + d).
\end{aligned}$$

Note that, as it is expected in reduced basis approaches [141], the offline cost is a lot more expensive than several full simulations.

Let us further comment on the source of the estimated cost. First, the AO Gram matrix snapshot cost is dominated by N_b^4 , the number of all four-center interactions. Next, the reduced-auxiliary basis is generated using Algorithm 4. This algorithm cost is dominated by a nested loop on $1 \leq m \leq M$ and $m + 1 \leq i \leq N_b^2$, inside which we perform operations that scale as n^d , yielding the total cost. As for the auxiliary coefficient snapshots, the cost is dominated by the resolution of a least-squares problem. In particular, for each point among the n^d samples, we must solve the least-squares problem of Eq. (5.10). For this, we need to compute the *ab initio* electronic density function, again at a cost of N_b^4 (see Section 1.2.1). Regarding the assembly of the least-squares normal equations, the auxiliary basis Gram matrix of dimension M by M needs to be evaluated entry-wise, as well as the right-hand side term, at a cost of $M N_b^2$. To summarize, the total estimated cost is computed as $n^d(N_b^4 + M N_b^2 + M^2)$.

Concerning the online computation, it is reduced to the interpolation formula evaluation of Eq. (5.9). This formula requires n^d additions for all M . Summation weighting coefficients are obtained by multiplying d values. Such d values are given as Lagrange polynomial evaluations. Due to the factorized form of Lagrange polynomials of Eq. (5.7), a single Lagrange polynomial of degree n can be evaluated at the cost of performing $2n$ additions, n divisions and n multiplications. This is linear cost in the polynomial degree. Moreover, Lagrange polynomial factors may be stored in memory and used recursively for evaluating Lagrange polynomials of any degree $1 \leq m \leq n$. Summing up, the total online cost may be modelled as $n^d M + n + d$, using the fact that the summation weighting coefficients are independent of M , they can thus be computed once and for all.

To conclude, the main advantage of the reduced basis approach is that, the online cost has been reduced from the initial $\mathcal{O}(N_b^4)$ to $\mathcal{O}(n^d M)$. For efficient system fragmentation, the number of intermolecular parameters d stays small. For example, for the water molecule in Figure 5.1, it is $d = 3$. Now, the online cost is linear on the auxiliary basis set size, with a prefactor depending on the parameter space samples. Note also that this may be further reduced by the use of EIM that could lead to a $\mathcal{O}(\tilde{n}M)$ where \tilde{n} is the number of local EIM points.

5.3 Numerical results

The present section concerns numerical results. We focus on results obtained using our method RB-DF, defined in Algorithm 6. Errors due to the use of our auxiliary-reduced basis generation of Algorithm 4 and of interpolation of Eq. (5.9) are studied separately. The goal is to present how our method enables to improve existing density fitting schemes for computing frozen core energies between flexible fragments, at low computational cost. To this end, various metrics are assessed, including orthogonal projection errors on densities, electron-electron frozen core energy errors, interpolation errors on coefficients, online timings.

Our test case is the ten true Smith water dimers (see Table 1.1). We refer to Section 1.3.1 for more details on the flexible clusters used as test cases. Results presented in this section are obtained using our `dfwpy` code for flexible fragments, we refer to Appendix A.

coord.	min.	max.	mean	std
d_1	0.932	0.998	0.961	0.019
d_2	0.928	1.008	0.967	0.023
θ	1.697	2.012	1.875	0.08

Table 5.1: Intramolecular water fragment geometry statistical analysis obtained for the equilibrated cluster of 27 AMOEBA polarizable water fragments (`watertiny` test case). Lengths are in Angstrom and angles in radians.

5.3.1 Flexible water clusters

The present section presents simulation details, concerning the set up of a water parametrization strategy that allows to apply reduced basis methods. We focus on defining the internal parameters used to sample water conformations. Overall, efficient sampling is crucial since it allows us to create snapshots used in the learning part of reduced basis methods. Note that the conformation sampling is relatively simple for the water molecule, since it is a planar object. More elaborate methods may become necessary when studying more complex molecules.

Sampling conformation space. We employ a simple sampling strategy for the conformation space of water molecules in three-dimensional space, using local water frames. The used frame is shown in Figure 5.1. The motivation for this choice of intramolecular parameter space of the water molecule H_2O is based on the three normal modes, which consist of symmetric stretching, asymmetric stretching and bending. These modes can be described by inner parameters, denoted by d_1 for the OH_1 bond length, d_2 for OH_2 bond length and θ for the angle H_1OH_2 . Hence, three degrees of freedom are used to generate configuration space samples of H_2O .

Up next, we propose a method to estimate parameter intervals, based on the statistical study of available clusters. Table 5.1 presents a statistical analysis of the parameter space of a small polarizable AMOEBA water cluster consisting of 27 fragments, after geometry equilibration using molecular dynamics. Recall that each water fragment in this cluster has a different conformation at the end of the geometry equilibration. We refer to Section 1.3.2.1 for practical details concerning such cluster. Obtained results show that the parameters values do not vary a lot with respect to the equilibrium water geometry, which is $d_1 = d_2 = 0.957$ Angstrom and $\theta = 1.82$ rad. We consider that such parameter values are representative of the general image of intramolecular water conformations, attained during the course of a typical molecular dynamics simulation with flexible fragments. Based on this sample, we chose the intervals of the finite parameter space as follows:

$$[d_{\min}, d_{\max}] = [0.9, 1.0], \quad [\theta_{\min}, \theta_{\max}] = [1.5, 2.2].$$

The finite set of sample points $\mathcal{M}_r = \{(d_1^{(l_1)}, d_2^{(l_2)}, \theta^{(l_3)})\}_{l_1, l_2, l_3}$ of size $r \in \mathbb{N}^*$ is then a subset of

$$\mathcal{M}_r \subset [d_{\min}, d_{\max}] \times [d_{\min}, d_{\max}] \times [\theta_{\min}, \theta_{\max}],$$

obtained by discretizing the parameter space using a multivariate uniform grid of constant size $n \in \mathbb{N}^*$ on each dimension, with a total size of $r = n^3$. Note that our tensorized strategy, concerning the grid nodes, induces a cubic scaling of the number of snapshots with respect to the number of nodes used in 1D. This scaling may become prohibitive for large values of n .

5.3.2 Reduced-auxiliary basis accuracy

The present section contains results on the orthogonal projection errors due to the use of our reduced-auxiliary basis, generated by our method RB-AUX given by Algorithm 4), within density fitting. The *ab initio* electronic densities are used as a reference. Additionally, we measure the

errors on self-interaction electron-electron energies. Let us emphasize that all experiments of the present section are performed offline in the frozen fragment library, i.e. on the frozen fragment level using a local coordinate system. We refer to Section A.7.1 for implementation details.

Let us begin by presenting simulation details. Given a cluster of $N_f \in \mathbb{N}^*$ fragments, the metric used to assess various errors, computed at the level of each fragment of the cluster, is defined as follows. Let us assume that a given error is obtained, in the form of a vector $(x_1, \dots, x_{N_f}) \in \mathbb{R}^{N_f}$, where each x_i is the error under study, i.e. $x_i = \|\rho(\mathbf{\mu}_i) - \Pi\rho(\mathbf{\mu}_i)\|/\|\rho(\mathbf{\mu}_i)\|$ for the case of orthogonal projection and $x_i = |E_{\text{self}}(\mathbf{\mu}_i) - \tilde{E}_{\text{self}}(\mathbf{\mu}_i)|/|E_{\text{self}}(\mathbf{\mu}_i)|$ for the case of self-interaction $E_{\text{self}} = \langle \rho, \rho \rangle$, with $\mathbf{\mu}_i$ being the vector of intramolecular parameters of the i -th AMOEBA fragment. All errors presented in this section are equal to the root mean square errors of the set of N_f fragments, defined as

$$\text{RMS} := \sqrt{\frac{1}{N_f} \sum_{i=1}^{N_f} x_i^2},$$

Note that this is a relative error measure. The Coulomb norm ($\dot{H}^{-1}(\mathbb{R}^3)$ inner product) is used for calculating orthogonal projections. Our reduced basis method uses the finite chemical space of the water fragment, defined in Section 5.3.1. The finite parameter space sample, on which snapshots are generated, has a fixed size $r = 3^3$ for $n = 3$.

Concerning admissible auxiliary basis sets, the snapshot Gram matrices are defined using Eq. (3.33), for contracted orbitals by orbital type and using the Coulomb metric. As a result, we directly control the size M_* of auxiliary basis functions by orbital type, measured for spherical GTOs, using our reduced-auxiliary basis generation procedure, just as in our standard ABS generation scheme (see Algorithm 1). Let us recall that the number M_* is smaller or equal to the total number of primitives in the basis, denoted by M . Lastly, atom-wise consistency is forced to all of our auxiliary basis sets, using the post-processing step previously described in Section 3.2.4.3.

5.3.2.1 Offline error assessment

In the present section, we assess conventional density fit errors due to the use of our reduced-auxiliary basis sets. Let us emphasize that the auxiliary basis coefficients are computed by orthogonally projecting a given density on an auxiliary basis set by solving the problem of Eq. (5.10). Hence, no interpolation strategy is employed in this numerical experiment.

Figure 5.2 presents reduced-auxiliary basis errors, for two different AO basis sets given as input to our generation algorithm RB-AUX of Algorithm 4. The cluster used for testing is the AMOEBA flexible water cluster of $N_f = 27$ fragments equilibrated using molecular dynamics (parameter information can be found in Table 5.1). The target reduced-auxiliary basis sizes have been selected as follows. Standard auxiliary basis sets used in chemistry typically contain about $3N_b^*$ to $5N_b^*$ auxiliary functions, where N_b^* is the number of AO basis functions by orbital type. For cc-pvdz, $N_b^* = 11$, therefore suitable auxiliary basis sizes M_* vary in the range of 30 to 60. For aug-cc-pvdz, one has $N_b^* = 18$, therefore M_* takes values in the range of 50 to 90. For each value of M_* , we plot the projection errors on the density and the energy errors. Results show that the error due to our reduced basis, both on the self-energy and on the *ab initio* density function, converges with respect to the reduced basis size, to a value that is lower than the chemical accuracy. Let us emphasize that we used only a few sample points, three per dimension. The fact that few sample points reproduce the *ab initio* electronic density function with such accuracy means that the solution space of auxiliary basis sets has a small Kolmogorov n -width [255]. This is a major argument in favor of the use of reduced basis methods.

Further commenting on results of Figure 5.2, for the moment, it is not clear to us how to further improve the reduced basis accuracy below the current errors. We tested a bigger sample size $n = 5$ for cc-pvdz AO basis, obtaining identical errors as the $n = 3$ case. Therefore increasing the sample size may not always help. Furthermore, we cannot choose, for example for the case

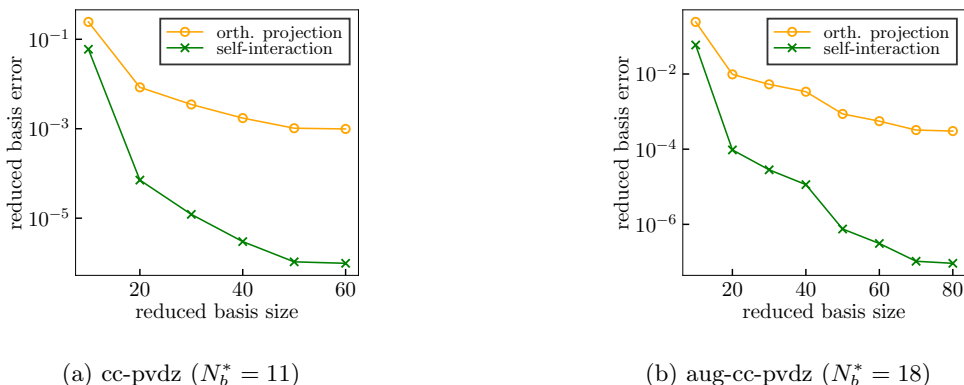


Figure 5.2: Fixed $n = 3$. RMS errors for 27 AMOEBA water fragments, for varying auxiliary basis sizes M_* . Measured errors are relative errors obtained for: the orthogonal projection $\|\rho - \tilde{\rho}\|_K$, with K the Coulomb kernel, and the electron-electron self-interaction $|\langle \rho, \rho \rangle - \langle \tilde{\rho}, \tilde{\rho} \rangle|$, where $\tilde{\rho}$ is obtained by density fitting on our basis of size M_* obtained with Algorithm 4. All basis sizes are measured by orbital type in spherical form.

of cc-pvdz basis sets, basis sizes larger than 110, because then our RB-AUX algorithm fails to choose the pivot. In particular, the issue is that remainder matrix diagonal elements in the PCD procedure become numerically zero, so the pivot selection cannot converge to add new pivots. Precisely, the maximization of step 5 in Algorithm 4, responsible for finding the m -pivot, yields inadmissible pivots with $i < m$. However, this is an issue concerning incomplete PCD algorithms in general. It may occur for any choice of finite sample parameter space snapshots.

5.3.2.2 Offline timings

Up next, we assessed computation timings for the offline reduced-auxiliary basis generation of Algorithm 4. Table 5.2 presents execution times for the RB-AUX generation, as well as for the precalculation of Gram matrix snapshots of the form of Eq. (3.33), for every point of the finite parameter space fixed to a total of $r = 3^3$ points. Results show that the offline calculation time is dominated by the Gram matrix snapshot calculation. In this case of $n = 3$, $r = 27$ Gram matrices must be calculated and stored. For cc-pvdz basis, this precomputation takes about 62 seconds. The reduced basis generation using Algorithm 4 takes significantly less than the snapshot computation, about a factor of 200 to 3, depending on the target reduced-auxiliary basis size. It is expected that the Gram matrix snapshot time calculation will grow with the number of snapshots. For example, for $n = 5$ yielding $r = 3125$ sample points, the Gram snapshot takes 290 seconds. The Gram snapshot calculation is by far the most expensive step of the entire reduced-auxiliary basis generation procedure. Recall that this computation is performed entirely offline.

Additionally, we assessed condition numbers of obtained auxiliary Gram matrices. Results give that condition number that is of the order of $\mathcal{O}(10^x)$. We generally observed that a limitation of the RB-based auxiliary basis is that the condition number of its Gram matrix explodes faster than classical auxiliary basis sets. This is expected because there is no explicit criterion for testing linear dependence of the selected reduced basis set in the Algorithm 4, as it was the case in Algorithm 1.

5.3.3 Interpolation error

The present section concerns the assessment of offline errors, measured on the level of individual fragments, due to the use of our interpolation strategy on coefficients, presented in Section 5.2.3.


RB size	Time	Cond. 10^x
10	0.29	3
30	2.26	7
50	7.45	10
72	11.66	12
97	16.86	16
110	20.13	21

(a) cc-pvdz. Gram precomputation: 62 seconds. AO orbital number 11, spherical component number 24.

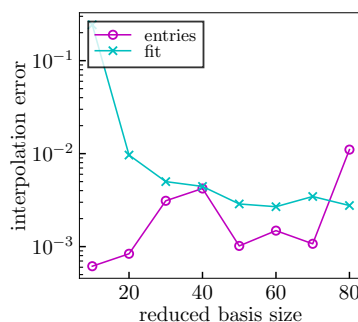
RB size	Time	Cond. 10^x
10	0.44	3
30	6.22	7
50	9.47	9
60	12.10	9
70	16.32	10
80	20.50	11

(b) aug-cc-pvdz. Gram precomputation: 85 seconds. AO orbital number 18, spherical component number 41.

Table 5.2: Fixed $n = 3$. Reduced basis generation timing (in seconds) for various sizes M_* and condition number of M_* -by- M_* auxiliary Gram matrices, with Coulomb metric.



(a) cc-pvdz



(b) aug-cc-pvdz

Figure 5.4: Fixed $n = 3$. One curve (**entries**) is the relative error in L_2 vector norm on the approximation of $\mathbf{c}(\boldsymbol{\mu})$ by $(\mathcal{I}[c_1(\boldsymbol{\mu})], \dots, \mathcal{I}[c_M(\boldsymbol{\mu})])$, for $\boldsymbol{\mu}$ of each fragment. The second curve (**fit**) is the relative error in Coulomb norm due to the density approximation of Eq. (5.11).

We used the same simulation setting as the one of Section 5.3.2.

Before discussing the results, let us quickly recall the explicit definition of our interpolation, for the sake of clarity. We consider a uniform 3D grid of $n \in \mathbb{N}^*$ nodes in each dimension. The interpolation scheme reads, for every $\boldsymbol{\mu} = (d_1, d_2, \theta)$, as

$$\forall 1 \leq i \leq M, \quad \mathcal{I}[c_i(\boldsymbol{\mu})] = \sum_{l_1, l_2, l_3=1}^n L_{l_1}(d_1)L_{l_2}(d_2)L_{l_3}(\theta)c_i(d_1^{(l_1)}, d_2^{(l_2)}, \theta^{(l_3)}). \quad (5.12)$$

Figure 5.4 presents the interpolation error, measured in two different ways, namely, first, the L_2 vector norm on interpolation coefficients written as a vector of size M and, second, the density error in the Coulomb metric. Notice that as more basis functions are added in the basis, the error in the entries in L_2 norm tends to increase. This may be due to the fact that added elements in the reduced basis are more and more linearly independent, causing more numerical instabilities. Generally, the interpolation error stays small.

Number of interpolation nodes. Lastly, we assessed the impact of adding more points to the sample. Results show that the accuracy of the reduced basis is independent of the number of sample points. This may indicate that the choice of the sample points is not suitable. However, the interpolation accuracy can be significantly improved by adding more points, as seen in Figure 5.5, where we increased the sample size to $r = 5^3$. Note that the interpolation degree

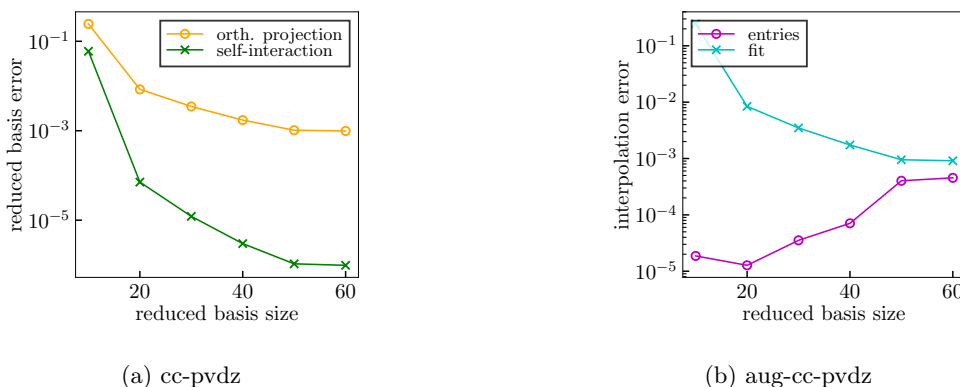


Figure 5.5: Fixed $n = 5$. We refer to Figure 5.2 and Figure 5.4a for the label description.

being equal to the sample size in each dimension, increasing the sample size also increased the interpolation degree.

5.3.4 Pairwise interaction errors

The present section provides results on the online frozen core energy computation, using our offline-online reduced basis scheme of Algorithm 5 and Algorithm 6, abbreviated here RB-INT, for Reduced Basis with Interpolation, in order to distinguish from RB, which consists of Algorithm 4 used in conventional density fitting. A detailed comparison between our method and state-of-the-art methods for flexible fragments is presented. Lastly, timings demonstrate the low computational cost of our method.

Before providing the results, let us quickly present simulation details and summarize the methods under study. We have evaluated the pairwise intermolecular electron-electron interaction for the ten Smith dimers obtained by [300], using four different computation methods for comparison. First, the exact method, used as a reference, is obtained by calculating the *ab initio* electronic density function online and then using such density to compute intermolecular integrals, namely Eq. (1.39) for the electron-electron term under study. Second, the RB approximation method is obtained by online fitting of the *ab initio* electronic density, computed online, using the reduced-auxiliary basis set, computed offline using Algorithm 4. Third, the RB-INT method replaces both the *ab initio* density computation and the online fit by our interpolation strategy of Section 5.2.3, with explicit definition given by Eq. (5.12). Lastly, the state-of-the-art GEM-DF method (recall Section 2.2.1.2) uses a given empirical auxiliary basis, namely RI or JKFIT, to fit once the *ab initio* electronic density offline for the water conformation at equilibrium, with inner parameters $d_1 = d_2 = 0.9572$ Å and $\theta = 1.8242$ rad. The recovered auxiliary coefficients are then reused online for other conformations, using the same empirical basis, repositioned at different atomic centers. Concerning our test case, the ten dimers' geometry is given by the inner parameters of Table 5.3. When comparing to the results of Table 5.1 used to define the parameter intervals, we see that the bond lengths of ten Smith dimers are within the limits. However, the bond angle takes values that are below the minimum considered in the sample size. As we will soon discuss, this doesn't seem to be a major problem.

5.3.4.1 Online frozen core energies

Table 5.4 provides a detailed comparison of the performance of RB-INT (our reduced basis including interpolation) and GEM-DF methods, with respect to the exact one used as a reference. Recall that the GEM-DF method has zero elapsed time for online density fitting, since density fitting is performed entirely offline, i.e. auxiliary coefficients are precomputed. Table columns

quantity	min.	max.	mean	std
d_1	0.9554	0.9576	0.9566	0.0007
d_2	0.9564	0.9644	0.9589	0.0023
θ	1.7775	1.8403	1.8228	0.0201
EE	16.1823	19.0915	18.0111	0.9151

Table 5.3: Inner coordinates of the ten Smith dimers. Exact EE (in Hartree) obtained for AO basis set cc-pvdz on H₂O.

dimer method	01	02	03	04	05	06	07	08	09	10
RI	18.3674	18.2125	18.1186	18.9327	19.0178	19.1060	17.9922	16.1945	17.5684	16.7308
JKFIT	18.3594	18.2034	18.1080	18.9223	19.0054	19.0921	17.9812	16.1838	17.5602	16.7224
RBI-10	21.0102	20.6693	20.4546	20.9492	20.7950	20.6868	19.8904	18.2046	19.5819	18.6179
RBI-28	18.3175	18.1678	18.0768	18.8950	18.9859	19.0778	17.9558	16.1499	17.5332	16.6962
RBI-46	18.3629	18.3629	18.1090	18.9178	18.9972	19.0813	17.9753	16.1858	17.5539	16.7166
RBI-64	18.3544	18.1972	18.1016	18.9191	18.9990	19.0832	17.9747	16.1844	17.5540	16.7163
RBI-82	18.3570	18.2002	18.1048	18.9201	19.0008	19.0855	17.9762	16.1829	17.5559	16.7180
RBI-100	18.3578	18.2011	18.1057	18.9204	19.0013	19.0862	17.9765	16.1825	17.5562	16.7184
exact	18.3558	18.2005	18.1062	18.9213	19.0047	19.0915	17.9771	16.1823	17.5546	16.7172

Table 5.4: Ten Smith dimers using AO basis cc-pvdz ($N_b^* = 11$) on H₂O. Electron-electron frozen core energy (in Hartree) computed using RB-INT- M_* , for M_* target size, and GEM-DF- x where $x \in \{\text{RI}, \text{JKFIT}\}$ is the given empirical auxiliary basis. The auxiliary basis size is $M_* = 30$ for RI and $M_* = 42$ for JKFIT. We mark the most precise energy approximation (in blue) and our RBI-46 energies that are more precise than JKFIT ones (in green).

correspond to dimers and rows to different methods. For each method and dimer, we provide the electron-electron interaction value computed online. Results show the lower errors on electron-electron interactions with our method, marked in blue. The lowest achieved absolute error on electron-electron energy is of the order of 10^{-3} . One may further notice that, for target size $M_* = 46$, our method RBI-46 reaches lower frozen core accuracies than JKFIT ($M_* = 42$), in 30% of the dimer cases, marked in green. Therefore, our RB-INT method improves the accuracy of state-of-the-art methods by adding more functions, which can be seen by the overall higher precision (in 80% of dimers) of our method with M_* greater than state-of-the-art ones. Nevertheless, results show that state-of-the-art methods perform not so bad after all. This may be interpreted as an indicator verifying the chemistry intuition, according to which changes on water conformation geometry do not induce significant changes on the densities. As a result, we may conclude that conformational changes of water do not make significant difference to density fit. Our method captures those changes efficiently, since results show that increasing the reduced-auxiliary basis size overall lowers the error on electron-electron energy, for all dimers.

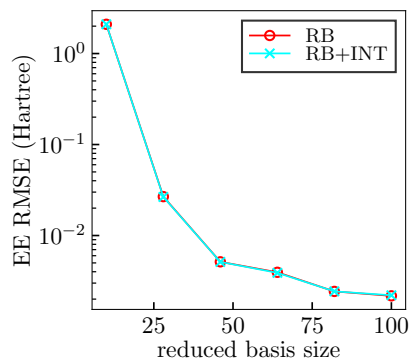
Up next, the following results show that our reduced basis method efficiently learns conformation changes when approximating a given density. Figure 5.6b shows the absolute root mean square error of Eq. (3.36) on electron-electron intermolecular interaction. Results show that interpolating does not introduce an error, as the two error curves coincide. We may interpret this as an encouraging result in favour of the interpolation. Moreover, results numerically demonstrate systematic error convergence on frozen core energies with respect to the reduced-auxiliary basis size.

5.3.4.2 Timings

In the present section we numerically demonstrate the advantage of using interpolation for reducing the computational cost of least-squares fitting, by employing precomputation steps.

RB size	QM	RB	RB-INT
10		4.216	0.016
28		4.250	0.022
46	4.19	4.539	0.060
64		4.624	0.075
82		4.679	0.078
100		4.687	0.100

(a) Online time (seconds).



(b) cc-pvdz

Figure 5.6: Fixed $n = 3$. For the timings (in seconds) we averaged a single run of all ten dimers. AO cc-pvdz has 11 orbitals with 24 spherical components. For comparison with state-of-the-art methods, namely GEM-DF, with RI basis we had an error $1E-2$ and for JKFIT basis $3E-3$.

Timings of our method RB-INT are given in Figure 5.6a. Results contain the time elapsed for the true density calculation (reference QM), the true coefficients on the reduced basis (RB) as well as the interpolated coefficients (RB-INT). Results show that the time spent on online density fitting is significantly reduced thanks to the interpolation. Notably, the interpolation speeds up the time of RB and QM calculations by a factor varying from 40 to 250, depending on the reduced-auxiliary basis size. This is an encouraging result in favour of RB-INT, that is expected since evaluating the auxiliary coefficients only requires a multivariate Lagrange polynomial evaluation in 3D. The 1D Lagrange polynomial has only degree $n = 3$ three, that is computationally not expensive. Note that the RB and QM methods have comparable timings. In particular, the RB online phase is more expensive than the reference QM one because it requires to calculate the density matrix, just as QM, and perform an additional least-squares fit. The density matrix computation dominates and there is zero effective acceleration gain by the use of auxiliary integrals. For this reason, the RB method does not present practical interest.

5.4 Conclusion and perspectives

In this work, we develop density fitting schemes for flexible fragments with applications to polarizable force field simulations. Compared with rigid fragments, the major difficulty in carrying out the fitted density approximation is that the fit cannot be precomputed offline, due to the reference fragment density and intramolecular geometry being updated at every timestep. In order to proceed, we set up tools for exploring the dimension of the solution space and constructing a reduced basis of the problem, able to reproduce all admissible densities. The result of this reduced basis approach is an offline-online greedy procedure for cost-efficient recovery of fitted densities, at linear time online. We plan to investigate and propose methods to decrease the intramolecular geometry sample number, based on greedy techniques such as the Empirical Interpolation Method (EIM) [25, 217], that allows to define proper interpolation points better suited than tensor ones. One difficulty to consider in sampling intramolecular geometries, during the reduced basis construction, is taking into account symmetries of molecules and assuring that the snapshots produce a physically meaningful continuous potential energy surface, i.e. without singularities.

From a numerical point of view, the results on the reduced basis size and accuracy on frozen core energies confirm the chemistry intuition for averaging fitted densities over a large sample of intramolecular geometries. We interpret this as a promising result in generating *a priori*

geometry-flexible auxiliary basis sets, with tunable accuracy. We demonstrate the practical use of the reduced basis on ten Smith dimers. Besides the diatomic water molecules, we also plan to use our reduced basis on larger molecules with richer symmetries. One difficulty in such case is that the intramolecular geometry space requires a large number of parameters to carry out parametrization. The hope is that by appropriate encoding of symmetries, the parameter space size can be decreased.

Besides intramolecular geometries, another possible application of reduced basis is the generation of multipole moments for polarizable force fields. We plan to numerically investigate the correlation of atomic moments for different intramolecular geometries. The hope is that the solution space of atomic moments has small n -width, admitting a reduced basis. Such investigation could lead to the mathematical foundation of permanent electrostatics parametrization for polarizable force fields (we refer to Appendix B for more details).

Chapter 6

Adaptive methods for atomic basis sets

This chapter includes joint work with Mi-Song Dupuy (Laboratoire Jacques-Louis Lions, Sorbonne Université) and Geneviève Dusson (Laboratoire de Mathématiques de Besançon, Université Bourgogne Franche-Comté).

Contents

6.1	Introduction	131
6.2	Theory	132
6.2.1	Setting	132
6.2.2	Variational approximation	132
6.2.3	Preliminaries	133
6.3	Results	135
6.3.1	Guaranteed <i>a posteriori</i> error estimates	135
6.3.2	Practical estimates	135
6.3.3	Adaptive refinement strategy	137
6.4	Numerical results	138
6.4.1	1D diatomic molecules	138
6.4.2	Quality of the error estimator	139
6.4.3	Influence of numerical parameters	140
6.4.4	Adaptive basis sets	142
6.5	Proofs	142
6.5.1	Preliminaries	142
6.5.2	Source problem	144
6.5.3	Eigenvalue problem - any single eigenvalue	145
6.6	Conclusion	148

In the present chapter, we focus on atomic orbital basis sets from a mathematical perspective. Let us recall that this work is not related to the frozen core energy computation treated in the remaining parts of the thesis. We establish guaranteed and practically computable *a posteriori* error estimators for source problems and linear eigenvalue problems, involving linear Schrödinger operators discretized with linear combinations of atomic orbitals. Such problems, arising in electronic structure theory of molecular systems, are indeed commonly discretized using linear combinations of atomic Gaussian-type orbitals. Our main contribution shows that the energy norm of the discretization error can be estimated by the dual energy norm of the residual, that further decomposes into atomic contributions, characterizing the error localized on atoms.

Moreover, we establish a method for practically computing the dual norms of atomic residuals by inverting radial Schrödinger operators associated to atoms. As an application of our result, we formulate a new adaptive refinement strategy for atomic orbital basis sets. Our strategy is adaptive in the sense that it indicates the atom on which one should add more atomic basis functions. Let us clarify that this does not provide the explicit basis functions one has to add. The practical use of such strategy is demonstrated on one-dimensional toy system models and numerical results are further discussed.

6.1 Introduction

We address the problem of bounding the errors in the approximated solution of linear Schrödinger-type problems, using atom-centered discretization basis sets.

Two discretization methods have been mainly studied in the literature. First, problems with periodic boundary conditions are discretized in the periodic space using an orthonormal plane-wave basis. Relevant estimators have been proposed, for the linear Kohn-Sham equations [152] and the nonlinear elliptic problem of Gross-Pitaevskii eigenvalue problem [112]. Existing work on the augmented plane wave method has been reported as well, for linear and nonlinear Kohn-Sham equations [73]. Second, problems posed over bounded domains have been discretized using a discontinuous Galerkin elements, studied for the linear Kohn-Sham equation [63, 180], or using conforming finite elements [64]. In this work, we focus on a different type of basis that is popular in molecular simulations, namely atom-centered discretizations over unbounded domains. Relevant *a priori* estimators on Gaussian atomic orbital discretizations have been investigated in [18]. To our knowledge, one of the few existing *a posteriori* analysis of the Hartree-Fock problem over unbounded domains for generic discretizations is proposed in [218], using invariance properties of the Hartree-Fock energy.

In the present work, we derive error bounds for atom-centered basis sets. From a theoretical point of view, we use the same strategy as that of [74], for constructing a partition of unity that spatially localizes contributions on atomic regions, and that of [64], for bounding the error using the dual norm of the residual. From a numerical point of view, the atomic basis sets we tested are Hermite basis sets, in the one-dimensional setting previously introduced in [62]. The main limitation of our approach is that the quality of the estimator depends on the values of the atomic potential on the region in between atoms. An advantage of our method is that the estimator is localized on atoms and the evaluation of the *a posteriori* bound can be performed in linear time with respect to the number of atoms.

The contribution of this work is two-fold: first, we illustrate the theory of residual-based estimators over unbounded domains and second, we present numerical results as a first attempt to generate adaptive atomic orbital basis sets. As an application of our estimator, we introduce novel atomic error indicators and propose an adaptive refinement strategy for increasing the basis size of each atom iteratively. From a numerical point of view, the results are encouraging. Preliminary results in 1D indicate that the *a posteriori* error estimator captures the inhomogeneity with respect to atoms of the system and therefore gives a promising approach to improving the accuracy of solutions to linear eigenvalue problems using atom-centered basis sets.

We distinguish the following practical limitation, concerning the evaluation of dual norms of residual terms over unbounded domains. Note that, in the case of bounded domains, such evaluation is straightforward using a finite spectral approximation in practice. However, operators defined over unbounded domains possibly have a non-empty essential spectrum. This makes the computation of norms of residual terms less straightforward and in the following paragraphs we address the issue of practically computing the dual norms in this case.

The rest of the chapter is organized as follows. In Section 6.2, we introduce the variational framework for solving linear problems and the construction of the atom-centered basis sets. Section 6.3 is devoted to the derivation of practical residual-based *a posteriori* estimators for atom-centered bases, as well an explanation of the refinement strategy. The effectiveness of the refinement strategy is verified in Section 6.4 by applying the refinement strategy to the solution of the 1D Hartree-Fock problem for the diatomic one-electron LiH^{3+} molecule. The conclusion and discussion of future work on applications of the *a posteriori* error estimator are given in Section 6.6. The details of the proofs used in Section 6.3 are provided in Section 6.5.

6.2 Theory

6.2.1 Setting

For $d \in \{1, 2, 3\}$, let

$$\mathbf{H}^1(\mathbb{R}^d) = \{v \in L^2(\mathbb{R}^d) : \nabla v \in L^2(\mathbb{R}^d)\}$$

be the Sobolev space endowed with the usual inner product on $L^2(\mathbb{R}^d)$, denoted by $\langle \cdot, \cdot \rangle_{\mathbb{R}^d}$ and defined as

$$\langle v, w \rangle_{\mathbb{R}^d} = \int_{\mathbb{R}^d} v(\mathbf{r})w(\mathbf{r}) \, d\mathbf{r}, \quad (6.1)$$

with the induced norm, denoted by $\|\cdot\|_{\mathbb{R}^d}$ and defined as

$$\|v\|_{\mathbb{R}^d} = \langle v, v \rangle_{\mathbb{R}^d}^{1/2}.$$

Given atomic positions $\mathbf{R}_1, \dots, \mathbf{R}_M \in \mathbb{R}^d$, consider the linear Hamiltonian operator of Schrödinger type defined by

$$H = -\Delta + \sum_{k=1}^M V_k + \sigma, \quad (6.2)$$

where $\sigma \in \mathbb{R}$ is a shift factor and $V_k(\mathbf{r}) = V(|\mathbf{r} - \mathbf{R}_k|)$ is a radial atomic potential with $V : \mathbb{R}^+ \rightarrow \mathbb{R}$ continuous in $(0, \infty)$. No additional regularity is imposed on V , allowing for common singular potentials such as the Coulomb potential $V(\mathbf{r}) = -|\mathbf{r}|^{-1}$.

Two problems on unbounded domains will be of interest in this work: first, the source problem

$$\text{Find } u \in \mathbf{H}^1(\mathbb{R}^d) \text{ solution to } Hu = f \text{ in } \mathbb{R}^d, \quad (6.3)$$

and second, the eigenvalue problem

$$\text{Find } (\lambda, u) \in \mathbb{R} \times \mathbf{H}^1(\mathbb{R}^d) \text{ solution to } Hu = \lambda u, \quad \|u\|_{\mathbb{R}^d} = 1. \quad (6.4)$$

6.2.2 Variational approximation

Problems (6.3) and (6.6) are solved using a Galerkin method with the atomic orbital (AO) basis set $\{\chi_\mu\}_{1 \leq \mu \leq N}$ of size N , composed of functions centered on atomic positions:

$$\{\chi_\mu\} = \{\xi_{1,1}(\mathbf{r} - \mathbf{R}_1), \dots, \xi_{1,n_1}(\mathbf{r} - \mathbf{R}_1); \dots; \xi_{M,1}(\mathbf{r} - \mathbf{R}_M), \dots, \xi_{M,n_M}(\mathbf{r} - \mathbf{R}_M)\},$$

with $\xi_{i,j} \in \mathbf{H}^1(\mathbb{R}^d)$ a fast decaying function centered at the origin, called atomic orbital, previously introduced in Chapter 1.1.4. Defining the Hamiltonian matrix

$$\mathbf{H} = (\langle \chi_\mu, H\chi_\nu \rangle_{\mathbb{R}^d})_{1 \leq \mu, \nu \leq N}$$

and the overlap matrix

$$\mathbf{S} = (\langle \chi_\mu, \chi_\nu \rangle_{\mathbb{R}^d})_{1 \leq \mu, \nu \leq N},$$

the discretisation of the source problem (6.3) in the AO basis set χ writes as the linear system

$$\mathbf{H}\mathbf{a} = \mathbf{S}\mathbf{f}, \quad (6.5)$$

with solution the vector of coefficients $\mathbf{a} \in \mathbb{R}^N$, where $\mathbf{f} = (f_\mu^X)_{1 \leq \mu \leq N}$ denotes the vector of orthogonal projection coefficients of the source term f on the AO basis set χ for the natural L^2 inner product. The approximation u_N of u in the AO basis set χ can then be recovered as a Linear Combination of Atomic Orbitals (LCAO), i.e.

$$\forall \mathbf{r} \in \mathbb{R}^d, \quad u_N(\mathbf{r}) = \sum_{\mu=1}^N \mathbf{a}_\mu \chi_\mu(\mathbf{r}).$$

The solution of the eigenvalue problem (6.4) is approximated using the Galerkin variational method as

$$\text{Find } (\lambda, u) \in \mathbb{R} \times \mathbf{H}^1(\mathbb{R}^d) \text{ solution to } \begin{cases} \langle v, Hu \rangle_{\mathbb{R}^d} = \lambda \langle v, u \rangle_{\mathbb{R}^3} & \forall v \in \mathbf{H}^1(\mathbb{R}^d), \\ \|u\|_{\mathbb{R}^d} = 1. \end{cases} \quad (6.6)$$

Using the AO basis set χ as a discretization basis, the variational problem (6.6) writes as the following generalized eigenvalue problem: find $(\mu_i, \mathbf{a}_i) \in \mathbb{R} \times \mathbb{R}^N$ such that

$$\forall 1 \leq i \leq N, \quad \mathbf{H}\mathbf{a}_i = \mu_i \mathbf{S}\mathbf{a}_i, \quad \mathbf{a}_i^\top \mathbf{S}\mathbf{a}_i = \mathbf{I}, \quad (6.7)$$

where \mathbf{I} denotes the N -by- N identity matrix. Similarly to the source problem, eigenfunctions can be recovered as

$$\forall 1 \leq i \leq N, \quad \forall \mathbf{r} \in \mathbb{R}^d, \quad u_{iN}(\mathbf{r}) = \sum_{\mu=1}^N \mathbf{a}_{i\mu} \chi_\mu(\mathbf{r}), \quad \|u_{iN}\|_{\mathbb{R}^d} = 1.$$

The associated eigenvalues are computed using the Rayleigh quotient

$$\forall 1 \leq i \leq N, \quad \lambda_{iN} = \langle u_{iN}, Hu_{iN} \rangle_{\mathbb{R}^d}.$$

6.2.3 Preliminaries

The Hamiltonian operator H acts on $L^2(\mathbb{R}^d)$ with domain $\mathbf{H}^2(\mathbb{R}^d)$. We define the residuals for the linear and eigenvalue problems as follows:

- for the linear problem (6.3)

$$\text{Res}(u_N) = f - Hu_N \quad (6.8)$$

- for the eigenvalue problem (6.6)

$$\text{Res}(\lambda_{iN}, u_{iN}) = \lambda_{iN} u_{iN} - Hu_{iN}. \quad (6.9)$$

Usual a posteriori estimates rely on the computation of dual norms. To reduce the computational cost of these quantities, the main idea is to take advantage of the radial symmetry of the potentials V_k . For $k = 1, \dots, M$, let H_k be defined by

$$H_k = -\Delta + V_k + \sigma_k,$$

for $\sigma_k \in \mathbb{R}$ shift factors. Notice that each operator H_k is radially symmetric with respect to \mathbf{R}_k . Associated to these operators, we introduce a cover $(\Omega_k)_{1 \leq k \leq M+1}$ of \mathbb{R}^d

$$\{\Omega_k\}_{1 \leq k \leq M+1}, \quad \Omega_k \subseteq \mathbb{R}^d, \quad k = 1, \dots, M+1,$$

with the property

$$\bigcup_{k=1}^{M+1} \Omega_k = \mathbb{R}^d, \quad \Omega_k = \bar{B}(\mathbf{R}_k, r_k), \quad k = 1, \dots, M,$$

where $r_k > 0$ and $\bar{B}(\mathbf{R}_k, r_k)$ is the closed ball centered at \mathbf{R}_k of radius r_k .

We consider a nonnegative partition of unity subordinate to the finite cover, denoted by

$$\{p_k\}_{1 \leq k \leq M+1}, \quad p_k \in C^2(\mathbb{R}^d), \quad \text{supp}(p_k) = \Omega_k, \quad p_k \geq 0, \quad k = 1, \dots, M+1,$$

with the property

$$\forall \mathbf{r} \in \mathbb{R}^d, \quad \sum_{k=1}^{M+1} p_k(\mathbf{r}) = 1. \quad (6.10)$$

Note that neighboring subdomains may intersect, i.e. $\Omega_k \cap \Omega_{k'} \neq \emptyset$ when $k \neq k'$, as illustrated in the example of Figure 6.1.

We will make the following assumption on the operators H and $(H_k)_{1 \leq k \leq M+1}$, such that they define natural energy norms.

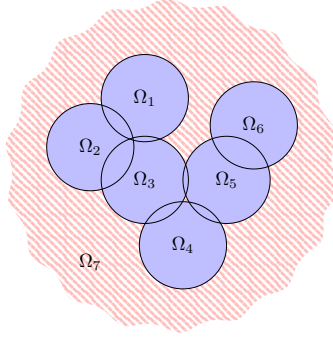


Figure 6.1: Finite covering of \mathbb{R}^d for $d = 2$ using the closed balls $\Omega_1, \dots, \Omega_M$, for $M = 6$ with fixed radius $r \in \mathbb{R}$ (in blue). Their unbounded complement is defined by $\Omega_{M+1} = \mathbb{R}^2 \setminus \cup_{k=1}^M \Omega_k$ (in red).

Assumption 1. The shifts σ and $\sigma_1, \dots, \sigma_{M+1}$ are such that the operators H and H_1, \dots, H_{M+1} are positive definite.

Assumption 2. The shift σ is such that

$$\forall v \in \mathbf{H}^1(\mathbb{R}^d), \quad c_H \|v\|_{H(\mathbb{R}^d)} \geq \|v\|_{\mathbb{R}^d}.$$

The corresponding energy norm for H is defined by

$$\forall v \in \mathbf{H}^1(\mathbb{R}^d), \quad \|v\|_{H(\mathbb{R}^d)} = \langle v, Hv \rangle_{\mathbb{R}^d}^{1/2}. \quad (6.11)$$

The associated dual norm is given by

$$\forall v \in L^2(\mathbb{R}^d), \quad \|v\|_{H^{-1}(\mathbb{R}^d)} = \sup_{w \in \mathbf{H}^1(\mathbb{R}^d)} \frac{\langle v, w \rangle_{\mathbb{R}^d}}{\|w\|_{H(\mathbb{R}^d)}},$$

Likewise, for H_{M+1} under Assumption 1, we can define the energy and the dual norms associated to H_{M+1}

$$\begin{aligned} \forall v \in \mathbf{H}^1(\mathbb{R}^d), \quad \|v\|_{H_{M+1}(\mathbb{R}^d)} &= \langle v, H_{M+1}v \rangle_{\mathbb{R}^d}^{1/2}, \\ \forall v \in L^2(\mathbb{R}^d), \quad \|v\|_{H_{M+1}^{-1}(\mathbb{R}^d)} &= \sup_{w \in \mathbf{H}^1(\mathbb{R}^d)} \frac{\langle v, w \rangle_{\mathbb{R}^d}}{\|w\|_{H_{M+1}(\mathbb{R}^d)}}, \end{aligned} \quad (6.12)$$

For $1 \leq k \leq M$, the operators H_k are acting on $L^2(\Omega_k)$ with domain

$$\mathbf{H}_0^1(\Omega_k) := \{v \in \mathbf{H}^1(\Omega_k) : u = 0 \text{ on } \partial\Omega_k\}.$$

The corresponding energy norm is given by

$$\forall v \in \mathbf{H}_0^1(\Omega_k), \quad \|v\|_{H_k(\Omega_k)} = \langle v, H_k v \rangle_{\Omega_k}^{1/2}, \quad (6.13)$$

where $\langle \cdot, \cdot \rangle_{\Omega_k}$ is the usual L^2 inner product of $L^2(\Omega_k)$. The associated dual norms are

$$\forall v \in L^2(\Omega_k), \quad \|v\|_{H_k^{-1}(\Omega_k)} = \sup_{w \in \mathbf{H}_0^1(\Omega_k)} \frac{\langle v, w \rangle_{\Omega_k}}{\|w\|_{H_k(\Omega_k)}}, \quad k = 1, \dots, M. \quad (6.14)$$

6.3 Results

6.3.1 Guaranteed *a posteriori* error estimates

Let us define the positive constant C by

$$C := 1 + c_H^2 \sup_{\mathbb{R}^d} \left(\sum_{k=1}^{M+1} -\frac{1}{2} \Delta p_k + \frac{(\nabla p_k)^2}{4p_k} + V_k(p_k - 1) + (\sigma_k - \sigma) p_k \right)^+. \quad (6.15)$$

Theorem 6.3.1 (*H*-error estimation for source problem). *Let u be a solution to the source problem (6.3) and u_N be an approximate solution. There holds*

$$\|u - u_N\|_{H(\mathbb{R}^d)} \leq C^{1/2} \left[\sum_{k=1}^{M+1} \|\sqrt{p_k} \text{Res}(u_N)\|_{H_k^{-1}(\Omega_k)}^2 \right]^{1/2}. \quad (6.16)$$

Let us define the gap constants associated to the i -th eigenvalue, for $1 \leq i \leq N$ by

$$\tilde{C}_i := \min_{\substack{j \geq 1 \\ j \neq i}} \left(1 - \frac{\lambda_{iN}}{\lambda_j} \right)^2, \quad (6.17)$$

$$\hat{C}_i := \min_{\substack{j \geq 1 \\ j \neq i}} \left(1 - \frac{\lambda_{iN}}{\lambda_j} \right)^2 \lambda_j. \quad (6.18)$$

Theorem 6.3.2 (*H*-error estimation for eigenvalue problem - single eigenvalue). *Let $1 \leq i \leq N$, (λ_i, u_i) be a solution to (6.6) and (λ_{iN}, u_{iN}) be an approximate solution. Under the assumption $\langle u_i, u_{iN} \rangle_{\mathbb{R}^d} \geq 0$, there holds*

$$\|u_i - u_{iN}\|_{H(\mathbb{R}^d)} \leq \left(C \tilde{C}_i^{-1} r_i + \lambda_i C^2 \hat{C}_i^{-2} r_i^2 \right)^{1/2}, \quad (6.19)$$

where

$$r_i := \sum_{k=1}^{M+1} \|\sqrt{p_k} \text{Res}(\lambda_{iN}, u_{iN})\|_{H_k^{-1}(\Omega_k)}^2.$$

Moreover,

$$0 \leq \lambda_{iN} - \lambda_i \leq C \tilde{C}_i^{-1} r_i. \quad (6.20)$$

6.3.2 Practical estimates

The main interest of Theorem 6.3.1 and Theorem 6.3.2 is that the dual norms appearing in the upper bound are easily computable, either by taking advantage of the radial symmetry or the explicit formula of the Green's function of the free Laplacian on the whole space.

For $1 \leq k \leq M$, since Ω_k is compact, the spectrum of H_k is purely discrete. We denote by $(\varepsilon_j^{(k)}, \psi_j^{(k)}) \in \mathbb{R} \times H_0^1(\Omega_k)$, $j \in \mathbb{N}$, the eigenpairs of H_k defined by

$$\begin{aligned} H_k \psi_j^{(k)} &= \varepsilon_j^{(k)} \psi_j^{(k)}, \\ \int_{\Omega_k} \psi_j^{(k)}(x) \psi_{j'}^{(k)}(x) dx &= \delta_{jj'}, \quad j, j' \geq 1. \end{aligned} \quad (6.21)$$

For any $v \in H_0^1(\Omega_k)$, we have

$$H_k v = \sum_{j=1}^{\infty} \varepsilon_j^{(k)} \psi_j^{(k)} \langle v, \psi_j^{(k)} \rangle_{\Omega_k},$$

hence a direct consequence is the following proposition.

Proposition 6.3.1 (Dual norm decomposition for compact operators). *Let $1 \leq k \leq M$ and $(\varepsilon_j^{(k)}, \psi_j^{(k)}) \in \mathbb{R} \times H_0^1(\Omega_k)$ be solutions to (6.21) for $j \geq 1$. There holds*

$$\forall v \in L^2(\Omega_k), \quad \|v\|_{H_k^{-1}(\Omega_k)}^2 = \sum_{j=1}^{\infty} \frac{1}{\varepsilon_j^{(k)}} \left| \langle v, \psi_j^{(k)} \rangle_{\Omega_k} \right|^2.$$

For all $1 \leq k \leq M$, let us fix a natural number $n = n(k) \in \mathbb{N}$ and introduce the partial expansion of the dual norm on the first n eigenpairs, denoted by $\mathcal{I}_n(\cdot)$, obtained by truncating the infinite expansion of Proposition 6.3.1 at order n , i.e.

$$\forall v \in L^2(\Omega_k), \quad \mathcal{I}_n(v) := \sum_{j=1}^n \frac{1}{\varepsilon_j^{(k)}} \left| \langle v, \psi_j^{(k)} \rangle_{\Omega_k} \right|^2.$$

This quantity appears in the bound of the dual norm as given in the following lemma and is fully computable.

Lemma 6.3.3 (Practical estimate). *Under the same assumptions as Proposition 6.3.1, fix $n = n(k) \in \mathbb{N}$. There holds*

$$\forall v \in L^2(\Omega_k), \quad \mathcal{I}_n(v) \leq \|v\|_{H_k^{-1}(\Omega_k)}^2 \leq \mathcal{I}_n(v) + \frac{1}{\varepsilon_{n+1}^{(k)}} \left(\|v\|_{\Omega_k}^2 - \sum_{j=1}^n \left| \langle v, \psi_j^{(k)} \rangle_{\Omega_k} \right|^2 \right).$$

Lastly, for the dual norm associated to the operator H_k for $k = M + 1$, one may use that the operator $H_{M+1} = -\Delta + \sigma_{M+1}$ is invertible using the explicit Green's function stated in the following lemma.

Lemma 6.3.4. *The Green's function of the linear differential operator $-\Delta + a^2$ is*

$$G(\mathbf{r}, \mathbf{r}') = \begin{cases} \frac{\exp(-a|\mathbf{r} - \mathbf{r}'|)}{2a}, & d = 1, \\ \frac{1}{2\pi} K_0(a|\mathbf{r} - \mathbf{r}'|), & d = 2, \\ \frac{\exp(-a|\mathbf{r} - \mathbf{r}'|)}{4\pi|\mathbf{r} - \mathbf{r}'|}, & d = 3, \end{cases}$$

where K_0 is a modified Bessel function of the second kind.

Our final estimates are guaranteed and practically computable.

Theorem 6.3.5 (Practical estimation for source problem). *Under the same assumptions as Theorem 6.3.1, let us fix a sequence of natural numbers $\mathbf{n} = (n_k)_{1 \leq k \leq M}$. There holds*

$$\|u - u_N\|_{H(\mathbb{R}^d)} \leq C^{1/2} \tilde{r}(\mathbf{n})^{1/2}, \quad (6.22)$$

where

$$\tilde{r}(\mathbf{n}) := \sum_{k=1}^M \left[\mathcal{I}_{n_k}(\sqrt{p_k} \text{Res}(u_N)) + \frac{1}{\varepsilon_{n_k+1}^{(k)}} \left(\|\sqrt{p_k} \text{Res}(u_N)\|_{\Omega_k}^2 - \sum_{j=1}^{n_k} \left| \langle \sqrt{p_k} \text{Res}(u_N), \psi_j^{(k)} \rangle_{\Omega_k} \right|^2 \right) \right] + \langle \sqrt{p_{M+1}} \text{Res}(u_N), g_{M+1} \rangle_{\Omega_{M+1}},$$

with

$$g_{M+1}(\mathbf{r}) := \int_{\mathbb{R}^d} G_{M+1}(\mathbf{r}, \mathbf{r}') \sqrt{p_{M+1}(\mathbf{r}')} \text{Res}(u_N)(\mathbf{r}') \, d\mathbf{r}',$$

G_{M+1} being the Green's function of operator $-\Delta + \sigma_{M+1}$.

Theorem 6.3.6 (Practical estimation for eigenvalue problem - single eigenvalue). *Under the same assumptions as Theorem 6.3.2, let us fix a sequence of natural numbers $\mathbf{n} = (n_k)_{1 \leq k \leq M}$. There holds*

$$\|u_i - u_{iN}\|_{H(\mathbb{R}^d)} \leq \left(C\tilde{C}_i^{-1}\tilde{r}_i(\mathbf{n}) + \lambda_i C^2 \hat{C}_i^{-2} \tilde{r}_i(\mathbf{n})^2 \right)^{1/2}, \quad (6.23)$$

where

$$\begin{aligned} \tilde{r}_i(\mathbf{n}) := & \sum_{k=1}^M \left[\mathcal{I}_{n_k}(\sqrt{p_k} \text{Res}(\lambda_{iN}, u_{iN})) \right. \\ & + \frac{1}{\varepsilon_{n_k+1}^{(k)}} \left(\|\sqrt{p_k} \text{Res}(\lambda_{iN}, u_{iN})\|_{\Omega_k}^2 - \sum_{j=1}^{n_k} \left| \langle \sqrt{p_k} \text{Res}(\lambda_{iN}, u_{iN}), \psi_j^{(k)} \rangle_{\Omega_k} \right|^2 \right) \\ & \left. + \langle \sqrt{p_{M+1}} \text{Res}(\lambda_{iN}, u_{iN}), g_{M+1} \rangle_{\Omega_{M+1}}, \right] \end{aligned}$$

with

$$g_{M+1}(\mathbf{r}) := \int_{\mathbb{R}^d} G_{M+1}(\mathbf{r}, \mathbf{r}') \sqrt{p_{M+1}(\mathbf{r}')} \text{Res}(\lambda_{iN}, u_{iN})(\mathbf{r}') \, d\mathbf{r}',$$

G_{M+1} being the Green's function of operator $-\Delta + \sigma_{M+1}$. Moreover,

$$0 \leq \lambda_{iN} - \lambda_i \leq C\tilde{C}_i^{-1}\tilde{r}_i(\mathbf{n}). \quad (6.24)$$

Note that, the computational cost of evaluating the bound of Theorem 6.3.6 scales linearly in the number of atoms M , as any two residual terms associated to different atoms are independent.

6.3.3 Adaptive refinement strategy

For any atom indexed by $1 \leq k \leq M$, let us define the error indicator

$$\begin{aligned} \eta_k := & \mathcal{I}_{n_k}(\sqrt{p_k} \text{Res}(\lambda_{iN}, u_{iN})) \\ & + \frac{1}{\varepsilon_{n_k+1}^{(k)}} \left(\|\sqrt{p_k} \text{Res}(\lambda_{iN}, u_{iN})\|_{\Omega_k}^2 - \sum_{j=1}^{n_k} \left| \langle \sqrt{p_k} \text{Res}(\lambda_{iN}, u_{iN}), \psi_j^{(k)} \rangle_{\Omega_k} \right|^2 \right), \end{aligned}$$

using the same notation as Theorem 6.3.6. Finally, for any sequence of natural numbers $\mathbf{n} = (n_k)_{1 \leq k \leq M}$, we obtain the total atomic discretization error indicator decomposed on atoms, defined as

$$\eta_i(\mathbf{n}) = \sum_{k=1}^M \eta_k.$$

The refinement is based on atoms, using the k -atom indicator $\eta_k \geq 0$. Let us emphasize that, the eigenpairs of the restricted operators H_k on subdomains for $1 \leq k \leq M$, given as solutions to (6.21), may be precomputed before the start of simulation. Lastly, note that obtaining the value of the error indicator η_k essentially requires evaluating several L^2 inner products. The total number of inner products for each $1 \leq k \leq M$ could be estimated as

$$n_k + 2.$$

To summarize, algorithm 7 presents a method to generate adaptive atom-centered basis sets. Theorem 6.3.6 guarantees that the sequence of approximate solutions $(u_{iN})_{N \in \mathbb{N}}$ obtained with such refinement strategy will converge to the exact u_i at the limit $N \rightarrow \infty$. The evaluation of error indicators in steps 1-3 of Algorithm 7 may be performed in parallel.


Algorithm 7: Adaptive refinement strategy for atom-centered basis sets.Input: system of M atoms, atomic basis of size N_k centered on atom k , for $1 \leq k \leq M$ Output: refined atomic basis of size N'_k , with $N'_k \geq N_k$ for all $1 \leq k \leq M$ 1: **for** $k = 1, \dots, M$ **do**2: Evaluate error indicator η_k .3: **end for**4: Set $k_0 = \arg \max_{1 \leq k \leq M} \eta_k$.5: **return** $N'_{k_0} = N_{k_0} + 1$ and $N'_k = N_k$ for all $1 \leq k \leq M, k \neq k_0$.

Figure 6.2: Partition of unity for two atoms at $\pm R$ in one dimension. Partition function associated to atom $(z_1, -R)$ (in blue), partition function associated to atom (z_2, R) (in red), remaining partition function associated to the complementary domain Ω_3 (in green).

6.4 Numerical results

6.4.1 1D diatomic molecules

Preliminary results are obtained for diatomic molecules with one electron in one dimension. Let us denote by z_1 and z_2 the nuclear charges of two atoms positioned at $-R$ and R in the real line, respectively, for $R > 0$. The computational domain \mathbb{R} is decomposed into partitions $\Omega_1, \Omega_2, \Omega_3$, defined symmetrically for the two atoms as

$$\begin{aligned}\Omega_1 &= [-\mu, \ell] \\ \Omega_2 &= [-\ell, \mu] \\ \Omega_3 &= (-\infty, -\bar{\mu}] \cup [\bar{\mu}, \infty), \quad \bar{\mu} := 2R - \ell.\end{aligned}$$

Note that the overlap of any two partition domains has length 2ℓ . Let us consider partition of unity functions that decay or grow on the overlap of domains based on the function

$$\forall x \in \mathbb{R}, \quad p(x) = \frac{h(x-a)}{h(x-a) + h(b-x)}, \quad h(x) = \exp(-1/x).$$

This function is increasing on $[a, b]$, equal to zero at $x = a$ and equal to one at $x = b$. Moreover, it is smooth on the closed interval $[a, b]$. It can therefore be extended smoothly to a constant function outside $[a, b]$. Such partition of unity based on p is illustrated in Figure 6.2.

The atomic orbital basis set studied is the Hermite basis set (HBS) introduced in [62], defined for any $n \in \mathbb{N}$ as

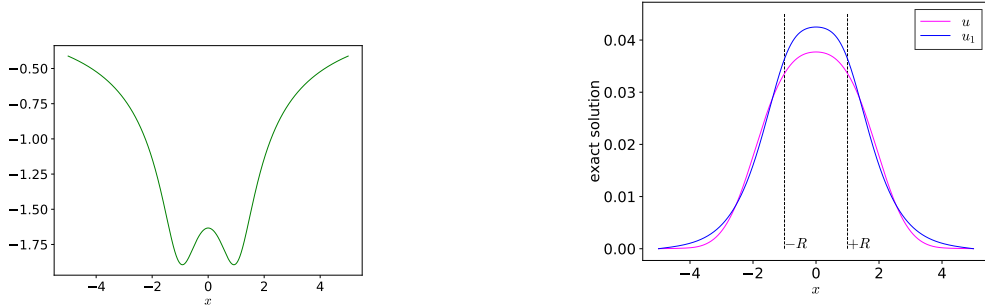
$$\forall x \in \mathbb{R}, \quad h_n(x) = c_n p_n(x) \exp(-x^2/2),$$

where p_n is the Hermite polynomial of degree n and c_n a normalization constant such that $\int_{\mathbb{R}} h_n^2 = 1$. Note that $\{h_n\}_{n \in \mathbb{N}}$ is an orthonormal basis of $L^2(\mathbb{R})$.

We use the soft-core Coulomb potential in 1D [323], obtained by avoiding the singularity of the Coulomb potential by adding a soft parameter α so that the potential becomes

$$V_\alpha(x) = \frac{1}{\sqrt{\alpha^2 + x^2}}.$$

Throughout our simulations, we fixed $\alpha = 0.5$ and set $V = V_\alpha$. The potential for this choice of soft parameter is illustrated in Figure 6.3a. A practical limitation is the lack of analytical



(a) Smeared Coulomb $x \mapsto V_\alpha$ for $\alpha = 0.5$.

(b) Reference solution of source problem (in pink) and first eigenfunction of the eigenvalue problem (in blue).

Figure 6.3: Fixed $z_1 = z_2 = 1$, $R = 1$. Atoms are fixed at positions $\pm R$.

expressions for the exact solution. To overcome this issue, we use a finite difference scheme over a bounded computational domain in 1D as a reference [62]. The code that allowed to calculate such finite difference solution on Hermite basis sets can be found in [181]. We consider a large enough computational domain, i.e. $\Omega = [-5R, 5R]$. If the Laplacian is discretized using a first-order finite difference scheme on a uniform grid $(x_i)_{1 \leq i \leq N_g}$ of step δ and N_g points on Ω , the discrete Hamiltonian matrix is written as

$$\mathbf{H} = -\frac{1}{2\delta^2} \begin{pmatrix} -2 & 1 & & & \\ 1 & -2 & 1 & & \\ & \ddots & \ddots & \ddots & \\ & & & 1 & -2 \end{pmatrix} + \text{diag}(V(x_i) + \sigma)_{1 \leq i \leq N_g}, \quad (6.25)$$

where the missing matrix entries are equal to zero. All integrals are computed exactly on the finite difference grid using the quadrature rule

$$\int_{\mathbb{R}} \chi_\mu \chi_\nu = \int_{\Omega} \chi_\mu \chi_\nu = \frac{1}{\delta^2} \sum_{i=1}^{N_g} \chi_\mu(x_i) \chi_\nu(x_i).$$

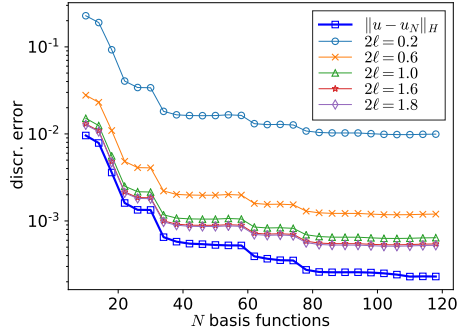
In all simulations we fixed a finite difference grid of $N_g = 2001$ points. Lastly, the right-hand side of the source problem we considered is

$$f(x) = H_1 h_1(x + R) + H_2 h_1(x - R),$$

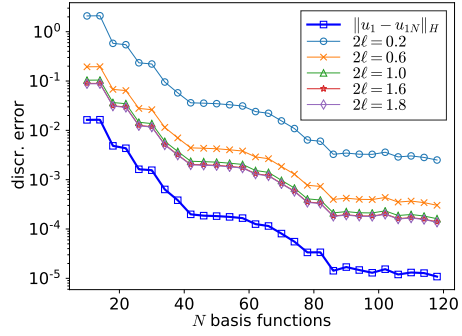
where H_k denotes the atomic Hamiltonian on atom k . For this choice of source we obtain the following reference finite difference solution illustrated in Figure 6.3b. Note that the exact inverse norms of atomic residuals in Theorem 6.3.2 may be computed in this setting by solving problems $H_k w = \text{Res}(\lambda_{iN}, u_{iN})$ on the finite difference grid, then evaluating the quantity $\langle \text{Res}(\lambda_{iN}, u_{iN}), w \rangle_{\Omega_k}$. Instead of performing the exact computation, we use the practical error bounds of Theorem 6.3.6 in the following numerical simulations.

6.4.2 Quality of the error estimator


We fixed a sequence of natural numbers $\mathbf{n} = (K, K)$, $K \in \mathbb{N}$, for evaluating the practical bound of Lemma 6.3.3. The eigenpairs (6.21) of the restricted operators are computed using the Lanczos eigensolver from KrylovKit module in Julia language, using a tolerance threshold equal to 10^{-12} . Note that we set $K = 17$ equal to the maximal number of converged eigenpairs. As it has been observed in [62, Fig. 2], the condition number of HBS blows up for large basis



(a) Estimation of $\|u - u_N\|_H$ for u solution of $Hu = f$.



(b) Estimation of $\|u_1 - u_{1N}\|_H$ for first eigenfunction u_1 of $Hu = \lambda u$.



(c) Practical estimator (see Theorem 6.3.6) of the exact discretization error on first eigenvalue.

Figure 6.4: Practical estimators (see Theorems 6.3.5 and 6.3.6) of the exact discretization error and influence of the overlap parameter ℓ on the estimator. Fixed $z_1 = z_2 = 1$, $R = 1$, $\sigma = 4$, $\sigma_1 = \sigma_2 = 3$, $\sigma_3 = 1$, $K = 17$. On the x -axis is the number of AO basis functions equal to $N = N_1 + N_2$ with $N_1 = N_2 = N/2$ per atom.

sets. This creates a numerical stability issue, that we partially reduced by applying a low-rank approximation to the overlap matrix of the basis using pivoted Cholesky decomposition with a tolerance threshold equal to 10^{-8} .

Figure 6.4 presents the numerical performance of the practical error estimators given by Theorems 6.3.5 and 6.3.6. The estimator is evaluated for several sizes of the partition overlap between $0.2 \leq 2\ell \leq 1.8 \leq 2R$, for fixed R . Results show that estimators become tighter for larger overlap sizes. Figure 6.4c presents the estimator on a single eigenvalue, given by Theorem 6.3.6. Results show that the quality of the eigenvalue error estimator is consistent with the one of the eigenvector error.

6.4.3 Influence of numerical parameters

In Table 6.1 we summarize the parameters of the model.

Spectral basis size. Figure 6.5 shows the influence of the spectral basis size $n_k = K$ in Lemma 6.3.3 on the estimation of the dual norm. Results show that the upper and lower bounds become tighter for increasing truncation degree, which is expected as the spectral approximation improves.

σ	shift factor for Hamiltonian H
σ_1	shift factor for Hamiltonian H_1
σ_2	shift factor for Hamiltonian H_2
σ_3	shift factor for Hamiltonian H_3
z_1	atomic charge of atom 1
z_2	atomic charge of atom 2
$2R$	interatomic distance
N_1	AO basis set size on atom $(z_1, -R)$
N_2	AO basis set size on atom (z_2, R)
K	size of spectral basis in (6.3.3) for all atoms
2ℓ	size of partition overlap

Table 6.1: Summary of parameters for diatomic molecule with atoms at $-R$ and R .




Figure 6.5: Fixed $R = 1$, $z_1 = z_2 = 1$, $\sigma = 4$, $\sigma_1 = \sigma_2 = 3$, $\sigma_3 = 1$, $\ell = 0.8$. Upper and lower bound of Lemma 6.3.3 on the dual norm of the residual error on the first eigenpair of the eigenvalue problem.

Overlap size. The influence of the size of the overlap given by 2ℓ is shown in Figure 6.4. The reason due to which the estimator is tighter for larger overlaps can be explained as follows. The constant C (6.15) of the practical estimator depends on ℓ and on the shift factors. For fixed shifted factors, the constant C tends to explode as the overlap size decreases, as seen in Table 6.2. Notice that the decay of C with respect to ℓ is exponential. Large C values result in less tight estimators.

Choice of shifts. In Table 6.2 we may also study how to optimize the constant C by appropriate choice of parameters. Notice that the atomic shift value $\sigma_1 = \sigma_2$ may not be decreased further without breaking the Assumption 1 on positive-definiteness of atomic operators. The last entry on this matrix could not be obtained for this reason. It is observed that small values of atomic shift factors yield optimal constants $C \geq 1$.

Gap constants. Lastly, we study the gap constants defined in equations (6.17) and (6.18). In practice the magnitude of the two constants for the case of the first eigenvalue depends on the ratio

$$\frac{\lambda_1}{\lambda_2} \leq \frac{\lambda_{1N}}{\lambda_2} \leq 1,$$

where we used the variational property $\lambda_{iN} \geq \lambda_i$. If the ratio is close to one, then the constants will be small. Ratios close to one can be obtained by selecting large values of the shift factor, as

$\ell \backslash \sigma_a$	5.0	3.0	2.0	1.0
0.1	421.38	421.052	420.884	420.716
0.3	6.49	6.155	5.987	5.819
0.5	2.11	1.780	1.612	1.478
0.8	1.68	1.344	1.176	1.000
0.9	1.63	1.298	1.130	–

Table 6.2: Values (up to three decimal digits) of the constant C defined in (6.15) for varying overlap size 2ℓ and atomic shift $\sigma_a = \sigma_1 = \sigma_2$. Fixed $z_1 = z_2 = 1$, $R = 1$, $\sigma = 4$, $\sigma_3 = 1$.

σ	λ_1/λ_2	c_H
3.0	0.746	0.694
4.0	0.833	0.409
5.0	0.875	0.290
6.0	0.900	0.225
7.0	0.917	0.183
8.0	0.929	0.155

Table 6.3: Values (up to three decimal digits) of the gap ratio λ_1/λ_2 as well as the positive-definiteness constant c_H of Assumption 2, for varying shift factor σ . Fixed $z_1 = z_2 = 1$, $R = 1$, $\sigma_1 = \sigma_2 = 1$, $\sigma_3 = 1$.

seen in Table 6.3. At the same time, however, large values of σ lead to sub-optimal values of the constant c_H in Assumption 2 therefore in larger values of the constant C of (6.15). Notice that we cannot go lower on values of σ otherwise the positive-definiteness Assumption 1 on H is violated.

6.4.4 Adaptive basis sets

In Figure 6.6 we tested the performance of the adaptive basis generation of Algorithm 7 for two particles of different nuclear charge. The tested diatomic molecule is LiH^{3+} with one-electron. The adaptive refinement fixes $N_1 = 5$ and varies N_2 only, from 5 to 53. On the contrary, the non-adaptive case varies $N_1 = N_2$ uniformly, from 5 to 29. Figure 6.6. The slopes of the two lines are very close, namely -95×10^{-4} (in green) -97×10^{-4} (in orange). For small basis sizes below $N = 20$, the adaptive basis strategy has no effect. However, it is observed that for a large basis size as $N = 60$, we gain one order of magnitude by using our refinement strategy.

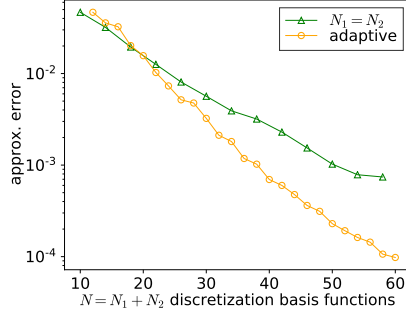
Finally, in Figure 6.7 we test the performance of the refinement strategy for particles of identical charges, for the molecule H_2^+ . The refinement strategy globally reproduces the behaviour of the optimal symmetric case $N_1 = N_2$ for identical charges. Note that the green line in the same figure is obtained by increasing by two both N_1 and N_2 at each refinement iteration.

6.5 Proofs


6.5.1 Preliminaries

Lemma 6.5.1. *There holds*

$$\forall v \in \mathbf{H}^1(\mathbb{R}^d), \quad \left(\sum_{k=1}^{M+1} \|\sqrt{p_k} v\|_{H_k(\Omega_k)}^2 \right)^{1/2} \leq C^{1/2} \|v\|_{H(\mathbb{R}^d)}, \quad (6.26)$$

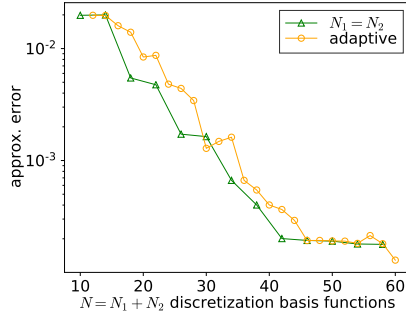


(a) Exact error $\|u_1 - u_{1N}\|_H$ for adaptive basis set versus uniformly refined one with $N_1 = N_2$.




(b) First exact eigenfunction of eigenvalue problem.

Figure 6.6: Adaptive basis sets for particles of different charges. Fixed $R = 1$, $z_1 = 1$, $z_2 = 3$, $\sigma = 7$, $\sigma_1 = 4$, $\sigma_2 = 4$, $\sigma_3 = 1$, $K = 17$, $\ell = 0.3$.



(a) Exact error $\|u_1 - u_{1N}\|_H$ for adaptive basis set (in orange) versus uniformly refined one with $N_1 = N_2$ (in green).



(b) First exact eigenfunction of eigenvalue problem.

Figure 6.7: Adaptive basis sets for particles of identical charges. Fixed $R = 1$, $z_1 = 1$, $z_2 = 1$, $\sigma = 7$, $\sigma_1 = 4$, $\sigma_2 = 4$, $\sigma_3 = 1$, $K = 17$, $\ell = 0.3$.

with C being the positive constant defined in (6.15).

Proof. By definition of the energy norm restricted on the subdomain Ω_k , for $k = 1, \dots, M + 1$, one has the equality

$$\|\sqrt{p_k}v\|_{H_k(\Omega_k)}^2 = \langle \sqrt{p_k}v, H_k(\sqrt{p_k}v) \rangle_{\Omega_k} = \langle \sqrt{p_k}v, H_k(\sqrt{p_k}v) \rangle_{\mathbb{R}^d}, \quad (6.27)$$

recalling that the support of p_k is Ω_k . Evaluating this inner product for $H_k = -\Delta + V_k + \sigma_k$ and integrating by parts, yields

$$\langle \sqrt{p_k}v, H_k(\sqrt{p_k}v) \rangle_{\mathbb{R}^d} = \langle \nabla(\sqrt{p_k}v), \nabla(\sqrt{p_k}v) \rangle_{\mathbb{R}^d} + \langle \sqrt{p_k}v, (V_k + \sigma_k)(\sqrt{p_k}v) \rangle_{\mathbb{R}^d}. \quad (6.28)$$

Expanding the differential operator ∇ gives

$$\begin{aligned} \langle \nabla(\sqrt{p_k}v), \nabla(\sqrt{p_k}v) \rangle_{\mathbb{R}^d} &= \int_{\mathbb{R}^d} \left(\sqrt{p_k} \nabla v + \frac{\nabla p_k}{2\sqrt{p_k}} v \right)^2 \\ &= \int_{\mathbb{R}^d} p_k (\nabla v)^2 + \int_{\mathbb{R}^d} (\nabla p_k) v \nabla v + \int_{\mathbb{R}^d} \frac{(\nabla p_k)^2}{4p_k} v^2 \\ &= \int_{\mathbb{R}^d} p_k (\nabla v)^2 - \frac{1}{2} \int_{\mathbb{R}^d} (\Delta p_k) v^2 + \int_{\mathbb{R}^d} \frac{(\nabla p_k)^2}{4p_k} v^2, \end{aligned}$$

where in the last line we integrated by parts. Using this expression within (6.28) and (6.27), then summing on k for (6.27) and evoking the property (6.10) of the partition of unity, we obtain

$$\sum_{k=1}^{M+1} \|\sqrt{p_k}v\|_{H_k(\Omega_k)}^2 = \|\nabla v\|_{\mathbb{R}^d}^2 + \sum_{k=1}^{M+1} \int_{\mathbb{R}^d} \left(-\frac{1}{2}\Delta p_k + \frac{(\nabla p_k)^2}{4p_k} + (V_k + \sigma_k)p_k \right) v^2. \quad (6.29)$$

Integrating by parts while using the regularity assumption on v , the energy norm of v squared admits the expression

$$\|v\|_{H(\mathbb{R}^d)}^2 = \langle v, (-\Delta + V + \sigma)v \rangle_{\mathbb{R}^d} = \|\nabla v\|_{\mathbb{R}^d}^2 + \sum_{k=1}^{M+1} \langle v, V_k v \rangle_{\mathbb{R}^d} + \sigma \|v\|_{\mathbb{R}^d}^2.$$

Adding and subtracting this quantity to (6.29), then majorizing the integral over the whole space, one obtains the estimate

$$\begin{aligned} \sum_{k=1}^{M+1} \|\sqrt{p_k}v\|_{H_k(\Omega_k)}^2 &= \|v\|_{H(\mathbb{R}^d)}^2 + \sum_{k=1}^{M+1} \int_{\mathbb{R}^d} \left(-\frac{1}{2}\Delta p_k + \frac{(\nabla p_k)^2}{4p_k} + V_k(p_k - 1) + (\sigma_k - \sigma)p_k \right) v^2 \\ &\leq \|v\|_{H(\mathbb{R}^d)}^2 + \|v\|_{\mathbb{R}^d}^2 \sup_{\mathbb{R}^d} \left(\sum_{k=1}^{M+1} -\frac{1}{2}\Delta p_k + \frac{(\nabla p_k)^2}{4p_k} + V_k(p_k - 1) + (\sigma_k - \sigma)p_k \right)^+, \end{aligned}$$

where $f(x)^+ := \max_{x \in \mathbb{R}^d} \{f(x), 0\}$. Lastly evoking $\|v\|_{\mathbb{R}^d} \leq c_H \|v\|_{H(\mathbb{R}^d)}$ of Assumption 2, we prove the desired result. \square

Lemma 6.5.2. *There holds*

$$\forall v \in L^2(\mathbb{R}^d), \quad \|v\|_{H^{-1}(\mathbb{R}^d)} \leq C^{1/2} \left(\sum_{k=1}^{M+1} \|\sqrt{p_k}v\|_{H_k^{-1}(\Omega_k)}^2 \right)^{1/2}. \quad (6.30)$$

Proof. By definition of the dual norm and Lemma 6.5.1,

$$\begin{aligned} \|v\|_{H^{-1}(\mathbb{R}^d)} &= \sup_{w \in H^1(\mathbb{R}^d)} \frac{\langle v, w \rangle_{\mathbb{R}^d}}{\|w\|_{H(\mathbb{R}^d)}} \\ &\leq C^{1/2} \sup_{w \in H^1(\mathbb{R}^d)} \frac{\langle v, w \rangle_{\mathbb{R}^d}}{\left(\sum_{k=1}^{M+1} \|\sqrt{p_k}w\|_{H_k(\Omega_k)}^2 \right)^{1/2}}. \end{aligned}$$

Moreover,

$$\begin{aligned} \langle v, w \rangle_{\mathbb{R}^d} &= \sum_{k=1}^{M+1} \langle H_k(\Omega_k)^{-1/2} \sqrt{p_k}v, H_k(\Omega_k)^{1/2} \sqrt{p_k}w \rangle_{\mathbb{R}^d} \\ &\leq \left(\sum_{k=1}^{M+1} \|\sqrt{p_k}v\|_{H_k^{-1}(\Omega_k)}^2 \right)^{1/2} \left(\sum_{k=1}^{M+1} \|\sqrt{p_k}w\|_{H_k(\Omega_k)}^2 \right)^{1/2}, \end{aligned}$$

hence the result. \square

6.5.2 Source problem

Proof of Theorem 6.3.1. Since u satisfies $Hu = f$, there holds $\text{Res}(u_N) = H(u - u_N)$. Decomposing $u - u_N$ using the partition of unity, one obtains the expression

$$\begin{aligned} \|u - u_N\|_{H(\mathbb{R}^d)}^2 &= \langle u - u_N, H(u - u_N) \rangle_{\mathbb{R}^d} \\ &= \sum_{k=1}^{M+1} \langle \sqrt{p_k}(u - u_N), \sqrt{p_k} \text{Res}(u_N) \rangle_{\mathbb{R}^d}. \end{aligned}$$

Since the support of p_k is Ω_k for every k , we may restrict the inner products on the subdomains and multiply by the identity operator $H_k^{1/2}H_k^{-1/2}$, hence the equality

$$\|u - u_N\|_{H(\mathbb{R}^d)}^2 = \sum_{k=1}^{M+1} \langle H_k^{1/2} \sqrt{p_k}(u - u_N), H_k^{-1/2} \sqrt{p_k} \text{Res}(u_N) \rangle_{\Omega_k}.$$

By the Cauchy–Schwarz inequality on Ω_k applied to the above right-hand side, one obtains the bound

$$\|u - u_N\|_{H(\mathbb{R}^d)}^2 \leq \sum_{k=1}^{M+1} \|\sqrt{p_k}(u - u_N)\|_{H_k(\Omega_k)} \|\sqrt{p_k} \text{Res}(u_N)\|_{H_k^{-1}(\Omega_k)},$$

that can be further estimated, using the discrete Cauchy–Schwarz inequality, as

$$\|u - u_N\|_{H(\mathbb{R}^d)}^2 \leq \left[\sum_{k=1}^{M+1} \|\sqrt{p_k}(u - u_N)\|_{H_k(\Omega_k)}^2 \right]^{1/2} \left[\sum_{k=1}^{M+1} \|\sqrt{p_k} \text{Res}(u_N)\|_{H_k^{-1}(\Omega_k)}^2 \right]^{1/2}.$$

Evoking Lemma 6.5.1 yields the result. \square

6.5.3 Eigenvalue problem - any single eigenvalue

Proposition 6.5.1 (L^2 -error estimation). *Let (λ_i, u_i) be solution to (6.6) associated to the i -th eigenvalue and (λ_{iN}, u_{iN}) be the approximate solution. There holds*

$$\|u_i - u_{iN}\|_{\mathbb{R}^d}^2 - \frac{1}{4} \|u_i - u_{iN}\|_{\mathbb{R}^d}^4 \leq \widehat{C}_i^{-1} \|\text{Res}(\lambda_{iN}, u_{iN})\|_{H^{-1}(\mathbb{R}^d)}^2,$$

where \widehat{C}_i is the positive constant defined in (6.18).

Proof. By definition of the residual, one has the equality

$$\|\text{Res}(\lambda_{iN}, u_{iN})\|_{H^{-1}(\mathbb{R}^d)}^2 = \langle H^{-1/2}(H - \lambda_{iN})u_{iN}, H^{-1/2}(H - \lambda_{iN})u_{iN} \rangle_{\mathbb{R}^d}. \quad (6.31)$$

By the spectral decomposition using the orthonormal eigenvectors $\{u_j\}_{j \geq 1}$ associated to the j -th eigenvalue $\lambda_j - \lambda_{iN}$ of the shifted operator $H - \lambda_{iN}$, where u_j is the eigenvector associated to the j -th eigenvalue λ_j of H , one has

$$(H - \lambda_{iN})u_{iN} = \sum_{j \geq 1} (\lambda_j - \lambda_{iN}) u_j \langle u_{iN}, u_j \rangle_{\mathbb{R}^d}, \quad (6.32)$$

while for the inverse operator $H^{-1/2}$, the spectral decomposition yields

$$\forall w \in L^2(\mathbb{R}^d), \quad H^{-1/2}w = \sum_{j \geq 1} \frac{1}{\lambda_j^{1/2}} u_j \langle w, u_j \rangle_{\mathbb{R}^d}. \quad (6.33)$$

Combining (6.32) and (6.33) in (6.31), while using the orthonormality of the eigenvectors u_j , one obtains the bound

$$\begin{aligned} \|\text{Res}(\lambda_{iN}, u_{iN})\|_{H^{-1}(\mathbb{R}^d)}^2 &= \sum_{j \geq 1} (\lambda_j - \lambda_{iN})^2 \left| \langle H^{-1/2}u_{iN}, H^{-1/2}u_j \rangle_{\mathbb{R}^d} \right|^2 \\ &= \sum_{j \geq 1} \left(1 - \frac{\lambda_{iN}}{\lambda_j}\right)^2 \lambda_j |\langle u_{iN}, u_j \rangle_{\mathbb{R}^d}|^2 \\ &\geq \sum_{\substack{j \geq 1 \\ j \neq i}} \left(1 - \frac{\lambda_{iN}}{\lambda_j}\right)^2 \lambda_j |\langle u_{iN} - u_i, u_j \rangle_{\mathbb{R}^d}|^2. \end{aligned} \quad (6.34)$$

We show that the following inequality is satisfied with the constant \widehat{C}_i defined in (6.18):

$$\|\text{Res}(\lambda_{iN}, u_{iN})\|_{H^{-1}(\mathbb{R}^d)}^2 \geq \widehat{C}_i \sum_{\substack{j \geq 1 \\ j \neq i}} |\langle u_{iN} - u_i, u_j \rangle_{\mathbb{R}^d}|^2. \quad (6.35)$$

Indeed, observe that the function $\mathbb{R}^+ \ni x \mapsto \left(1 - \frac{\lambda_{iN}}{x}\right)^2 x$ reaches its minimum at $x = \lambda_{iN}$ and is decreasing on $(0, \lambda_{iN}]$ and increasing on $[\lambda_{iN}, \infty)$. Thus the constant must be equal to $\min_{j \geq 1, j \neq i} \left(1 - \frac{\lambda_{iN}}{\lambda_j}\right)^2 \lambda_j$ or \widehat{C}_i .

Moreover, employing the scalings $\|u_i\|_{\mathbb{R}^d} = \|u_{iN}\|_{\mathbb{R}^d} = 1$, there holds

$$\begin{aligned} \langle u_i - u_{iN}, u_i \rangle_{\mathbb{R}^d} &= 1 - \langle u_{iN}, u_i \rangle_{\mathbb{R}^d} \\ &= \frac{\|u_i\|_{\mathbb{R}^d}}{2} + \frac{\|u_{iN}\|_{\mathbb{R}^d}}{2} - \langle u_{iN}, u_i \rangle_{\mathbb{R}^d} \\ &= \frac{1}{2} \|u_i - u_{iN}\|_{\mathbb{R}^d}^2. \end{aligned} \quad (6.36)$$

Hence

$$\begin{aligned} \sum_{\substack{j \geq 1 \\ j \neq i}} |\langle u_{iN} - u_i, u_j \rangle_{\mathbb{R}^d}|^2 &= \|u_{iN} - u_i\|_{\mathbb{R}^d}^2 - |\langle u_{iN} - u_i, u_i \rangle_{\mathbb{R}^d}|^2 \\ &= \|u_{iN} - u_i\|_{\mathbb{R}^d}^2 - \frac{1}{4} \|u_{iN} - u_i\|_{\mathbb{R}^d}^4. \end{aligned}$$

Plugging this expression within (6.35) yields the result. \square

Lemma 6.5.3. *Under the assumption $\langle u_i, u_{iN} \rangle_{\mathbb{R}^d} \geq 0$, there holds*

$$\|u_i - u_{iN}\|_{\mathbb{R}^d} \leq (2\widehat{C}_i^{-1})^{1/2} \|\text{Res}(\lambda_{iN}, u_{iN})\|_{H^{-1}(\mathbb{R}^d)}.$$

Proof. Using assumption $\langle u_i, u_{iN} \rangle_{\mathbb{R}^d} \geq 0$ and employing the scaling $\|u_i\|_{\mathbb{R}^d} = \|u_{iN}\|_{\mathbb{R}^d} = 1$, one obtains the inequality

$$\|u_i - u_{iN}\|_{\mathbb{R}^d}^2 = 2 - 2\langle u_i, u_{iN} \rangle_{\mathbb{R}^d} \leq 2. \quad (6.37)$$

Writing

$$\begin{aligned} \|u_i - u_{iN}\|_{\mathbb{R}^d}^2 - \frac{1}{4} \|u_i - u_{iN}\|_{\mathbb{R}^d}^4 &= \|u_i - u_{iN}\|_{\mathbb{R}^d}^2 - \frac{1}{4} \|u_i - u_{iN}\|_{\mathbb{R}^d}^2 \|u_i - u_{iN}\|_{\mathbb{R}^d}^2 \\ &\geq \|u_i - u_{iN}\|_{\mathbb{R}^d}^2 - \frac{1}{2} \|u_i - u_{iN}\|_{\mathbb{R}^d}^2 \\ &= \frac{1}{2} \|u_i - u_{iN}\|_{\mathbb{R}^d}^2 \end{aligned}$$

and evoking Proposition 6.5.1 yields the result. \square

Proposition 6.5.2 (*H-error estimation*). *Under the same assumptions as Proposition 6.5.1, there holds*

$$\|u_i - u_{iN}\|_{H(\mathbb{R}^d)}^2 - \frac{\lambda_i}{4} \|u_i - u_{iN}\|_{\mathbb{R}^d}^4 \leq \widetilde{C}_i^{-1} \|\text{Res}(\lambda_{iN}, u_{iN})\|_{H^{-1}(\mathbb{R}^d)}^2,$$

where \widetilde{C}_i is the constant defined in (6.18).

Proof. We show that the right-hand side of the inequality (6.34) may be bounded as follows, where the bound is satisfied for \widetilde{C}_i :

$$\|\text{Res}(\lambda_{iN}, u_{iN})\|_{H^{-1}(\mathbb{R}^d)}^2 \geq \widetilde{C}_i \sum_{\substack{j \geq 1 \\ j \neq i}} \lambda_j |\langle u_{iN} - u_i, u_j \rangle_{\mathbb{R}^d}|^2. \quad (6.38)$$

Indeed, observe that the function $\mathbb{R}^+ \ni x \mapsto \left(1 - \frac{\lambda_{iN}}{x}\right)^2$ reaches its minimum at $x = \lambda_{iN}$ and is decreasing on $(0, \lambda_{iN}]$ and increasing on $[\lambda_{iN}, \infty)$. Thus the constant must be equal to $\min_{j \geq 1, j \neq i} \left(1 - \frac{\lambda_{iN}}{\lambda_j}\right)^2$ or \tilde{C}_i .

Moreover, employing (6.36), one obtains the expression

$$\begin{aligned} \sum_{\substack{j \geq 1 \\ j \neq i}} \lambda_i |\langle u_{iN} - u_i, u_j \rangle_{\mathbb{R}^d}|^2 &= \|u_{iN} - u_i\|_{H(\mathbb{R}^d)}^2 - \lambda_i |\langle u_{iN} - u_i, u_i \rangle_{\mathbb{R}^d}|^2 \\ &= \|u_{iN} - u_i\|_{H(\mathbb{R}^d)}^2 - \frac{\lambda_i}{4} \|u_{iN} - u_i\|_{\mathbb{R}^d}^4. \end{aligned}$$

Using this equality within (6.38) yields the result. \square

Proposition 6.5.3 (Eigenvalue estimation).

$$\lambda_{iN} - \lambda_i \leq C_i^{-1} \|\text{Res}(\lambda_{iN}, u_{iN})\|_{H^{-1}(\mathbb{R}^d)}^2.$$

Proof. Evoking (6.36), there holds

$$\begin{aligned} \|u_i - u_{iN}\|_{H(\mathbb{R}^d)}^2 &= \lambda_i + \lambda_{iN} - 2\lambda_i \langle u_i, u_{iN} \rangle_{\mathbb{R}^d} \\ &= \lambda_{iN} - \lambda_i + 2\lambda_i (1 - \langle u_i, u_{iN} \rangle_{\mathbb{R}^d}) \\ &= \lambda_{iN} - \lambda_i + \lambda_i \|u_i - u_{iN}\|_{\mathbb{R}^d}^2. \end{aligned} \tag{6.39}$$

Hence

$$\lambda_{iN} - \lambda_i = \|u_i - u_{iN}\|_H^2 - \lambda_i \|u_i - u_{iN}\|^2$$

Since

$$\frac{\|u_i - u_{iN}\|^2}{4} \leq 1,$$

there holds

$$\lambda_{iN} - \lambda_i \leq \|u_i - u_{iN}\|_H^2 - \frac{\lambda_i}{4} \|u_i - u_{iN}\|^4.$$

Using Proposition 6.5.2 ends the proof. \square

Proof of Theorem 6.3.2. Employing Lemma 6.5.3 within Proposition 6.5.2 yields

$$\begin{aligned} \|u_i - u_{iN}\|_{H(\mathbb{R}^d)}^2 &\leq \tilde{C}_i^{-1} \|\text{Res}(\lambda_{iN}, u_{iN})\|_{H^{-1}(\mathbb{R}^d)}^2 + \frac{\lambda_i}{4} \|u_i - u_{iN}\|_{\mathbb{R}^d}^4 \\ &\leq \tilde{C}_i^{-1} \|\text{Res}(\lambda_{iN}, u_{iN})\|_{H^{-1}(\mathbb{R}^d)}^2 + \lambda_i \hat{C}_i^{-2} \|\text{Res}(\lambda_{iN}, u_{iN})\|_{H^{-1}(\mathbb{R}^d)}^4. \end{aligned}$$

Evoking Lemma 6.5.2 yields the result. \square

Proof of Proposition 6.3.1. We apply the spectral theorem to the operator H^{-1} , whose eigenvalues are $1/\varepsilon_j^{(k)}$. Then we use the orthonormality of the eigenvectors to compute the norm. \square

Proof of Lemma 6.3.3. We decompose the dual H_k -norm of an arbitrary v as

$$\|v\|_{H_k^{-1}(\Omega_k)}^2 = \mathcal{I}_n(v) + \sum_{j=n+1}^{\infty} \frac{1}{\varepsilon_j^{(k)}} \left| \langle v, \psi_j^{(k)} \rangle_{\Omega_k} \right|^2. \tag{6.40}$$

Assuming without loss of generality that the eigenvalues are sorted in increasing order, i.e. $\varepsilon_j^{(k)} \leq \varepsilon_{j+1}^{(k)}$ for all $j \geq 1$, there holds

$$\sum_{j=n+1}^{\infty} \frac{1}{\varepsilon_j^{(k)}} \left| \langle v, \psi_j^{(k)} \rangle_{\Omega_k} \right|^2 \leq \frac{1}{\varepsilon_{n+1}^{(k)}} \sum_{j=n+1}^{\infty} \left| \langle v, \psi_j^{(k)} \rangle_{\Omega_k} \right|^2. \tag{6.41}$$

By the spectral theorem, the infinite family $\{\psi_j^{(k)}\}_{j \geq 1}$ of eigenvectors is a Schauder basis of $H^1(\Omega_k)$, therefore

$$v = \sum_{j=1}^{\infty} \langle v, \psi_j^{(k)} \rangle_{\Omega_k} \psi_j^{(k)} \implies \|v\|_{\Omega_k}^2 = \sum_{j=1}^{\infty} \left| \langle v, \psi_j^{(k)} \rangle_{\Omega_k} \right|^2.$$

Decomposing the sum over j into two parts, we obtain the expression

$$\sum_{j=n+1}^{\infty} \left| \langle v, \psi_j^{(k)} \rangle_{\Omega_k} \right|^2 = \|v\|_{\Omega_k}^2 - \sum_{j=1}^n \left| \langle v, \psi_j^{(k)} \rangle_{\Omega_k} \right|^2.$$

Substituting this formula into the bound (6.41), we obtain the final upper estimator for the dual H_k -norm of v ,

$$\|v\|_{H_k^{-1}(\Omega_k)}^2 \leq \mathcal{I}_n(v) + \frac{1}{\varepsilon_{n+1}^{(k)}} \left(\|v\|_{\Omega_k}^2 - \sum_{j=1}^n \left| \langle v, \psi_j^{(k)} \rangle_{\Omega_k} \right|^2 \right).$$

The lower estimator is obtain by noticing that in expression (6.40), all eigenvalues $\varepsilon_j^{(k)}$ are positive due to Assumption 1 on positive-definiteness of H_k . \square

Proof of Theorem 6.3.5. Let $w_k := \sqrt{p_k} \text{Res}(u_N)$ for $k = 1, \dots, M+1$. Applying Proposition 6.3.3 to w_k for $1 \leq k \leq M$ yields the upper bound

$$\|w_k\|_{H_k^{-1}(\Omega_k)}^2 \leq \mathcal{I}_n(w_k) + \frac{1}{\varepsilon_{n+1}^{(k)}} \left(\|w_k\|_{\Omega_k}^2 - \sum_{j=1}^n \left| \langle w_k, \psi_j^{(k)} \rangle_{\Omega_k} \right|^2 \right).$$

Lastly, for $k = M+1$, we use the Green's function of the operator $-\Delta + \sigma_{M+1}$ in order to invert the operator H_{M+1} directly on the whole space \mathbb{R}^d . Notice that since the support of p_k is Ω_k ,

$$\|w_k\|_{H_k^{-1}(\Omega_k)} = \|w_k\|_{H_k^{-1}(\mathbb{R}^d)}$$

for any k and in particular for $k = M+1$. As a consequence, we compute for $k = M+1$,

$$\|w_k\|_{H_k^{-1}(\Omega_k)}^2 = \langle w_k, H_k^{-1} w_k \rangle_{\Omega_k} = \langle w_k, v_k \rangle_{\Omega_k} = \langle w_k, v_k \rangle_{\mathbb{R}^d},$$

with $v_k := H_k^{-1} w_k$, equal to

$$\forall \mathbf{r} \in \mathbb{R}^d, \quad v_k(\mathbf{r}) = \int_{\mathbb{R}^d} G(\mathbf{r}, \mathbf{r}') w_k(\mathbf{r}') \, d\mathbf{r}',$$

obtained using the Green function of the operator H_{M+1} . \square

Proof of Theorem 6.3.6. Similarly as proof of Theorem 6.3.5. \square

6.6 Conclusion

In this work, we develop residual-based *a posteriori* estimators for solving linear Schrödinger-type problems in an atom-centered orbital basis framework. Compared with standard estimators over bounded domains, the major difficulty in carrying out the analysis is that the dual norm of the residual term cannot be directly calculated due to the lack of compactness of the unbounded operator. In order to proceed, we develop tools for practically computing this dual norm based on restricting the unbounded Hamiltonian using a partition of unity defined over compact subdomains. The result of this *a posteriori* error analysis is a practical adaptive refinement

strategy for atom-centered basis sets. We plan to carry out a numerical verification of the validity of the adaptive refinement strategy for realistic three-dimensional systems, Coulomb potentials and Gaussian-type orbital bases commonly used in chemistry. Moreover, we plan to investigate methods of decomposing the physical space into atomic regions and define partitions of unity satisfying smoothness assumptions in 3D. One difficulty to consider while testing realistic models is the Coulomb potential singularity. The hope is that such numerical experiments show that, the singularity of the Coulomb potential can be controlled by appropriate partitions of unity so that the prefactors of the estimators remain small.

From a practical point of view, the results of the adaptive refinement strategy for atom-centered basis sets mimic the chemistry intuition for placing more basis function on the heavier atom. As a first attempt, we interpret this as a promising result in GTO basis optimization studies. We demonstrate the practical use of the estimator for one-dimensional diatomic molecules. Besides the diatomic molecules, we also plan to use the non-uniform refinement to study many-atom systems. One difficulty in that case is to define the overlap regions in 3D and to calculate the dual norms over such regions. In these cases, several atomic domains may overlap at the same time, yielding highly non-uniform overlap regions that we hope can still capture the inhomogeneity of atomic residuals.

Besides the adaptive atom-centered basis sets, another possible application of *a posteriori* error estimation is the construction of error bars for equilibrium geometries given by the geometry optimization problem for molecular geometries (see [60, Equation 1.1]). Such error bars could allow to predict confidence intervals for equilibrium geometries of a given energy model, such as Hartree-Fock (HF), or Local Density Approximation (LDA). The hope is to construct a proof entirely based on *a posteriori* error estimators in order to demonstrate theoretically the experimental observation according to which the LDA underestimates equilibrium geometries while HF overestimates them [250], as discussed in Appendix D. However, this objective requires, first, a much more precise error estimator than the one used in the current setting, and, second, combining the self-consistent field errors and discretization errors, and this could be studied in the future.

Chapter 7

Conclusion and perspectives

The contributions of this thesis are divided into two parts, regarding numerical methods for atomic Gaussian-type orbital basis sets. The first part was dedicated to a high-level approach for the density fit method, with applications to frozen core intermolecular interaction computations for rigid and flexible fragments, and extensions to sparse approximations. In a second part, we described an *a posteriori* analysis to estimate atom-centered Gaussian discretization errors on linear eigenvalue problems of Schrödinger type, with applications to adaptive Gaussian basis set refinement.

High-level approach for frozen core interactions. We developed a general framework, based on the concept of pairwise interaction summations, allowing to represent the frozen core intermolecular interaction computation problem in an abstract way. We came up with novel fast summation solutions to this problem. Because we are mainly interested in density fitting methods, we proposed a high-level presentation of solutions based on projection methods. Our work also introduces the splitting between offline and online simulation time, applied to our problem setting, and motivates the interest in performing offline precalculations. In this respect, we formulate a novel analysis of the density fitting method within the offline-online framework. The main application of this presentation is the general treatment of molecular conformation geometries appearing in molecular clusters, which is widely applied to other parts of this thesis. While the density fitting method is not a new concept, we believe that our contribution made it not kernel specific anymore, allowing to rely directly on assumptions on the Hilbert space rather than on the explicit form of the interaction kernel. It has also made it possible to provide rigorous justification for the possible offline precalculation of density fitting.

Among the perspectives associated to this work, the implementation of a general unified framework for the design of density approximation methods, with extensions to multipole moment generation, based on the tools we presented would be a challenging task.

Rigid fragments in frozen core interactions. To minimize the number of pairwise orbital interactions of a single frozen core intermolecular interaction computation and in order to cover the case of fragments of rigid intramolecular geometry, we proposed to use a new systematically generated auxiliary basis set computed once on the local frame, instead of the state-of-the-art empirical bases, for density fitting. Our method reduces the auxiliary basis size to reach a target size but requires additional treatment to respect the admissible format of orbitals. To tackle this issue, we developed a purely algebraic approach for the block treatment of the matrix form of the set of pairwise products of atomic orbitals. We presented the applications of our method to water clusters. Using an efficient strategy for the approximation of Gram matrices, we significantly reduced the complexity of frozen core interaction summations. We provided several ideas to further improve our method, such as including midbond centers and imposing maximal orbital type constraints.

The basis generation and size reduction approaches are based on quite general ideas and can be applied to other problems. However, the theoretical scaling of our method is the same as conventional density fitting. In order to deal with this issue, we proposed a sparsification technique that is able to further reduce the number of retained orbital pairs within a given pairwise interaction summation. Our method generalizes the ideas of integral screening for discarding small integral contributions. Although selecting significant sum contributions under a prescribed tolerance is not a new idea, our method introduces a novel discrete optimization framework to construct sparse summations. Concerning the applications of our method, it could be interesting to investigate efficient sparse integral assembly implementations, in order to benefit from the gain our method in practice.

The coupling of density fitting with numerical integration methods can become critical to apply the Gaussian electrostatic model to larger systems. With our results we showed that our existing implementation, based on direct integral evaluation, is not enough to reduce the computation time for reasonably large systems of interest. Indeed, the computational and storage of integrals and intermediary integral computation in recursion formulas become prohibitive. The use of numerical integral methods may at least limit the impact of intermediary operations on these two points. Moreover, the loop on fragment pairs plays an important role and density fitting should be coupled to fast summation methods for practical uses. A future plan is to incorporate our methods into the `GEM_fit` code [138], that employs efficient numerical integration methods enabling the frozen core calculation for large-scale molecular clusters.

Flexible fragments in frozen core interactions. To face up the problem of evaluating intermolecular interactions in the flexible intramolecular geometry case, we proposed our own flexible fragment strategy, based on the reduced basis method. We finally proposed a numerical comparison between our strategy and the two most well-known variants of empirical basis sets and density fitting methods, exhibiting important performance gains, especially using precalculations. Moreover, perspectives concerning our work include the further improvement of the reduced basis method, in order to identify more quantities that admit precalculation or attempt more efficient snapshot sampling methods, such as the empirical interpolation method [217].

Application of density fitting in force fields. It is interesting to see our frozen core interaction calculation incorporated into force fields. See Appendix B for more details. To achieve this, we need to compute electrostatics using multipole expansions or other acceleration methods based on range separation [84]. We identify several steps that must be studied to this goal. First, the comparison of existing multipole moment generation methods with the ones based on density fitting. Then, the accuracy of multipoles needs to be assessed up to quadrupoles. To this end, one can first obtain auxiliary basis sets using our Algorithm 1 imposing an additional maximal orbital type restriction. Aspects to be investigated include the convergence of the multipole expansion, the radius of convergence and the impact of the choice of centers.

In order to realize this perspective, one needs to implement our density fitting scheme, or couple our existing `dfwpy` code, within the `GEM_fit` code [138] for generating multipole moments and interfacing with polarizable force fields, namely AMOEBA [320]. This will allow to compute frozen core interactions based on multipoles, as an alternative to the exact evaluation employed in the present thesis.

The contributions of the present thesis can be inscribed into a more general framework. In particular, concerning molecular dynamics simulations, if we wish to use our work to develop new methods within `Tinker-HP` [4], one should implement a code for multipole moments (or use the existing implementation of `GEM_fit`) and test the applications of our methods for large-scale molecules.

***A posteriori* error estimation of atom-centered discretizations.** The problem of formulating an *a posteriori* discretization error analysis for atomic orbital basis sets has been tackled. This led us to new adaptive refinement strategies that allow to systematically choose

the atom on which one should add more orbitals, in order to improve the discretization error of a given atomic orbital basis set. We obtained early promising results in the one-dimensional case using small systems. These results only concern Hermite basis sets. We thus plan to investigate the Gaussian-type orbital basis sets, that represent the chemistry setting in real-world simulations. The difficulty when doing so concerns the recovery of the exact solution to validate our experiments, taking into consideration that a finite element library for three-dimensional molecules with more than three atoms is to available, to our knowledge.

Appendices

Appendix A

dfwpy: Density Fitting for Water molecules in Python

The present chapter contains implementation details on our code, named `dfwpy`, realizing the density fitting methods proposed in this thesis.

A.1 Algorithmic design

Our code is based on a minimal implementation calling `PySCF` [312] routines for water dimers when needed. Although density fitting implementations is not a new thing [138], the novelty of our code consists in providing routines for the computation of exact frozen core energies obtained by *ab initio* densities, for the purpose of error validation of density fitting.

A precomputation strategy is used in our code, in order to reduce simulation time, following Section 2.2.1.2. In this respect, the computation is divided into two stages: the offline and the online one. The offline one consists in defining the molecules on which one can obtain the density matrix (using DFT/B3LYP) and the auxiliary coefficients. The online stage includes frozen core energy computations.

Frozen core energy computation optimization was not a priority at this stage, as most of our test cases are small- to medium-size molecules (less than 100 atoms) used for testing purposes. Note that further optimizations need to be taken into account, in order to use our code for larger systems in practical applications. Overall, it is important to note that no optimization of the code was performed at this time and that the density fitting core program was initially designed for energy decomposition and not for force field calculations [83]. An alternative is `GEM_fit` which targets larger systems, by improving the complexity based on multipole expansions for the Coulomb term, in addition to density fitting.

A.2 Linear algebra

The reference implementation is the `dpstrf` pivoted Cholesky factorization routine [214] implemented in `LAPACK` [7] using level 3 BLAS. The default threshold used in pivoted Cholesky decomposition is 10^{-14} . The singular value decomposition (SVD) of matrices is taken from `numpy` package.

A.3 Water fragments

For any fragment, we read the coordinates, then we define the local frame. By convention, we chose the frame presented in Figure A.1. Then between two local frames we can find




Figure A.1: Local 3D coordinate system defined on water. The x -axis passes from the center of the bond H_1H_2 . The y -axis is co-planar to the molecule and perpendicular to the first axis. The z -axis is perpendicular to the other two.

the rotation matrix around the origin that aligns the three vectors defining the axis, using `scipy.spatial.transform.Rotation.align_vectors`. This gives us a rotation matrix between the local frame and the position of the fragment in the global frame.

A.3.1 Local to global transformation

We implemented the **L2G** transformation using the Wigner matrices [292] obtained from PySCF. In practice, any coefficient computed in the frozen library should be expressed in the global frame using Wigner matrices. PySCF allows to generate the Wigner matrix associated to a given AO basis from a rotation matrix applied to the atomic positions of its centers. This rotation matrix in particular maps the local frame to the global one. Let us denote the Wigner transformation by U . Then the global density matrix can be written as $\hat{\mathbf{D}} = \mathbf{U}^T \mathbf{D} \mathbf{U}$, where \mathbf{D} is the density matrix belonging to the frozen library expressed in the local frame. The auxiliary coefficients on the global frame are similarly expressed as $\hat{\mathbf{c}} = \mathbf{c} \mathbf{U}$, where \mathbf{c} are the auxiliary coefficients on the local frame. Note that the Euclidean rotation, mapping the local to the global frame and denoted by \mathbf{Q} , is theoretically an orthogonal matrix. In practice, however, round-off errors are observed, of the order of 10^{-6} , on individual entries of $\mathbf{Q}^T \mathbf{Q} - \mathbf{I}$. This numerical error comes from computing the transformations that align vectors in \mathbb{R}^3 .

A.4 Molecular integrals

The results are obtained with our implementation in Python. Our code interfaces with PySCF for QM calculations and molecular integrals. PySCF integrals and Wigner are on AO basis sets in spherical form. All AO basis sets are contracted when computing the density matrix. All auxiliary basis sets contain primitive functions only.

All molecular integrals are evaluated using Libcint [310]. We used the latest version in which the bug on four-center one-electron integral has been fixed. The loop for assembling cross-Gram matrix is implemented in the library `gto` of PySCF. This library uses OpenMP [44] to accelerate the assembly. The implementation of this library is in C and the entries of the matrix is performed in parallel using OpenMP and dynamic scheduling. The dynamic scheduling allows to deal with different angular momenta, because every group has different execution time and treating them statically would create a load-balancing issue since threads would wait others to finish. The parallelism paradigm is data-parallelism, that is the same operation is executed on different chunks of data, i.e. on different entries of the Gram matrix.

Remark. *To perform computation required for this thesis, we needed to fix a bug occurring in Libcint version 5.5.0. This bug concerned a segmentation fault error in the array used for cache during the assembly of the four-center one-electron integral matrix, in particular in the `int4c1e` routine. The reason of the memory leak issue was that the allocated space was*

interaction	two-center	four-center
EE	<code>aux.int2e</code>	<code>mol.int2e</code>
XR	<code>mol.int1e_ovlp</code>	<code>mol.int4c1e</code>

Table A.1: Molecular integral routines used for different energy terms: EE (electron-electron) and XR (exchange-repulsion). Note that XR four-center integral requires `Libcint` version at least 5.5.0, where the memory leak issue of the `int4c1e` has been fixed.

too small. The bug was that the arrays were actually stored as complex ones, because the assembly function is also used for spinors, which include an additional variable which is the spin, stored in complex arrays. However the cache size was assuming that the arrays are of type double. We thus adjusted the size and obtained a fix. Details of the correction can be found at <https://github.com/sunqm/libcint/pull/100>. Note that we also deactivated the optimizer in `moleintor` for four-center one-electron integrals in `PySCF` (this is expected to affect for timings).

In practice, we computed intermolecular integrals between two different molecular fragments using the `shell` argument of integrator engines of `PySCF`, interfacing with `Libcint`. The dimensions correspond to different molecules.

A.5 Sparsification

The solution to RCSPP (see Eq. (4.18) and Eq. (4.19)) has been calculated using `cspy` module [321], which is a collection of programs for solving constraint shortest path problems in operational research. There exist several resolutions of this problem, namely exact and metaheuristic algorithms. We opted for the only exact algorithm in the library, which is the bidirectional labeling algorithm with dynamic half-way point [319]. This means that it provides an exact (optimal) solution to RCSPP. For this reason, it takes a long time to run. Note that the available implementation is sequential. Figure A.2 shows the output of our main graph algorithm for a small example summation of random numbers.

A.6 Rigid fragments

A.6.1 Operators

Note that the local to global transformation has a numerical error of the order 10^{-7} , calculated with `spatial.transform.Rotation` module of `scipy` in Python. In practice, the numerical error comes from aligning vectors in 3D. Before calculating the rotation, we first translated the global fragment so that the two Oxygen atoms of water fragments coincide. We used Wigner matrices of `PySCF` to transform the density matrix to the current basis.

A.6.2 Optimizations

In the present section, we provide details on preliminary code optimizations.

A.6.2.1 Sparse summation

We implemented our mask in the C interface of `PySCF`. We could have stayed in the Python interface but the loop for assembling the matrix would not be optimized nor parallel. We used a simple implementation for discarding atomic products. The summation over sparse domains is implemented within this library. We used a COO format for sparse matrices.

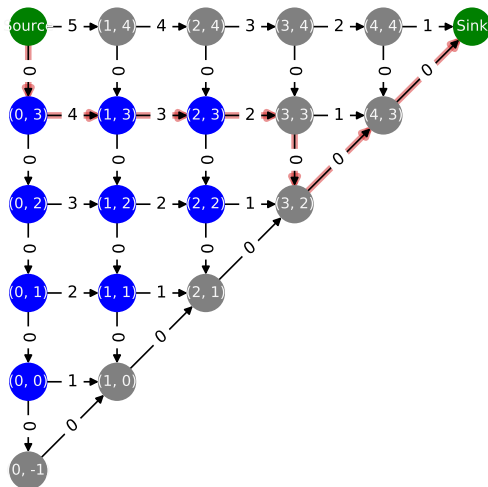


Figure A.2: Truncation of a double sum of degree $n = 5$ with random entries in $[0, 1)$. Paths have a given origin (**Source**) and destination (**Sink**). Resource tolerance is equal to 4, shortest path (in red) has length equal to 9 ($= 4 + 3 + 2$) and resource consumption equal to 3.87 (< 4). Summation domain (in blue) is defined by shortest path.

Mask implementation. We implemented a sparse two-index Gram matrix assembly in `C`. Our implementation is naive but manages to capture execution time improvements. The matrix assembly library used by PySCF is `gto`. Within this library, the file `fill_int2c` contains routines for assembling two-center matrices entry-wise. We propose to use a boolean array, i.e. a mask, to store the indices of the discarded orbital pairs. The mask is stored as a binary matrix. This strategy requires storing the full matrix entry wise, even entries that are zero, and then using the mask to find the discarded entries. We implemented the binary matrix, with entries zero and ones, and added a condition to check if the entry is zero. If it is zero, then we continue the iteration to next entry. Otherwise, we evaluate the intermolecular integral. A naive implementation consists of performing the full loop and inserting at each iteration a condition that skips the iteration if the entry of the binary matrix is equal to zero. Note that the existing matrix assembly loops of the `gto` library are highly optimized. The n -by- n matrix assembly is performed using a single n^2 loop, that is parallelized using OpenMP multithreading [44]. The blocks are scheduled dynamically, meaning that threads do not receive exactly the same iteration range. The scheduling is based on distributing work to different threads without any specific order. This is suitable when iterations require different computational costs and are not balanced. This is the case of molecular integral calculations, where the cost of every integral formula depends on the angular momentum of the orbitals.

In order to connect this implementation with PySCF we added an additional mask argument in the `moleintor` function of PySCF. We deactivated the optimizer when assembling matrices by setting it to null. Our sparse implementation can be activated by giving a value to the optional argument `mask` in `moleintor`. Note that the sparse assembly can be combined with a sparse dot product multiplication in other formats, in order to further accelerate the auxiliary coefficient multiplication. We do not address this acceleration here.

Timings. We discuss optimizations on sparse assembly and provide preliminary timings, using naive methods. We focus on implementations of the interaction sparse assembly. In order to obtain timings, we set the mask to the identity for DF and run all molecular integral calculations with the same algorithm for sparse summation. This ensures that the same assembly implementation

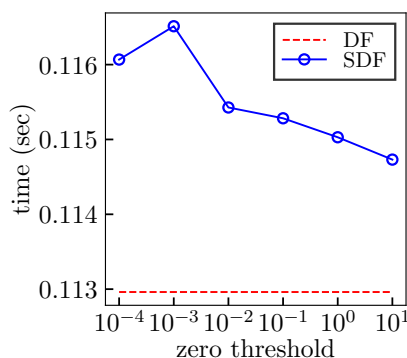


Figure A.3: Execution time (in seconds) of sparse assembly for different sparsity tolerance values. Interaction is electron-electron. Auxiliary basis is aug-cc-pvdz-pcd-120 generated using Algorithm 1. Sparsification has been performed using SPOP given by Algorithm 2, taking as input the Gram matrix of the normalized auxiliary basis.

is used for both compared methods. Figure A.3 presents the execution time averaged over 20 runs, for our naive sparse summation implementation. Results show that the implementation is far from efficient, since it is more expensive to compute the sparse summation than the full one. In particular, the sparse summation timings are worst than the full one, by approximately a factor of 10. This may be due to reasons related to cache or to the sparse format we used. To conclude, efficient implementations of the sparse summation are left for future investigation.

Lastly, note that in this numerical experiment, we compute the sparsification from the fitted density, instead of the *ab initio* one. This has to do with storage convenience, since for the case of the *ab initio* density, the scaling of orbitals becomes more complicated to implement. In this case, the basis of products must be stored as a new basis in PySCF. Recall that the new basis is normalized by default. This should be fixed to match the weights in the *ab initio* density decomposition. In particular, one should first store the weight coefficients of pairwise products of orbitals, along with the density matrix weights, and then impose this coefficients to the stored product basis in PySCF, after undoing its default normalization. The part that may become complicated is storing the correspondance between the coefficients and orbital parameters. We refer to Section 4.2.3.1 for a discussion on the influence of the scaling coefficients to the sparsification.

A.6.2.2 Loop on fragments

The loop on fragment pairs was implemented in Python and parallelized using `pypm` [194], an Open-MP-like functionality for Python. To parallelize the double loop, the strictly upper triangular indices (i, j) with $i < j$ up to n were stored in a one-dimensional array of size $n(n+1)/2$. The loop on elements of this array is run in parallel. Note that our implementation is far from optimal, but gives an idea of the orders of magnitude, as well as the limitations, of energy computations using a naive algorithm. Note that we forced PySCF to use a single thread and deactivated true Open-MP multithreading for molecular integral evaluation on a single fragment pair with `Libcint`, in order to avoid conflicts. The Python module `timeit` was used to measure elapsed time.

A.7 Flexible fragments

A.7.1 Transformation between frames

The **L2G** transformation. Intermediary transformation at any step is that for any given molecular coordinates, we must first recover the local coordinate system. This system is used to calculate L2G.

Local coordinate system. The water local coordinate system, on which density fitting and orthogonal projections are first calculated and then mapped to the global frame via the local to global transformation (see Section 2.1.2.3), is defined in Figure A.1. Note that for small three-atom molecules such as water, it is straightforward to define a local frame and one suffices for the entire molecule. However, for many-atom systems defining atomic local frames can be more complicated. There are multiple methods to define a local reference frame for each atom, notably the ones used by the machine learning community for atomic environment representations and descriptors [245, 266], exploiting rotational symmetries [265]. For simplicity, we focus on water. Wigner D-matrices are used to numerically calculate the local to global transformation.

A.8 Multipole moments

The multipole moments are a quantity of interest in polarizable force fields, used in the computation of the electrostatic intermolecular forces. The role of multipole moments is to describe the charges of interacting charge distributions using multipole expansions. We refer to Appendix B for mathematical details on multipole moments for the Coulomb kernel.

A given auxiliary basis set (ABS), generated with our method described in Chapter 3, has applications to multipole moment generation. Indeed, multipoles generated from the Gaussian electrostatic model ABSs have been previously [320] included to the AMOEBA polarizable force field for the water molecule, with promising results. Before presenting the theory on multipole moments and force fields, we quickly provide practical details on how to interface our code `dfwpy` with the existing `GEM_fit` implementation [138] of multipole moment generation.

Our code can output ABS parameters per atom type in `GEM_fit` internal formal. The main feature of the current `GEM_fit` implementation to take into account is that basis sets are exclusively of *spd*-type, in favour of computational efficiency [84]. We can impose *spd* orbitals in our ABS generation method in Algorithm 1, as follows. We propose to add a new step, right after Step 3, with two goals: discard orbitals of type higher than *d* and extend each *s,p* and *d*-type orbital in order to form inseparable groups of contracted *spd* orbitals. This ensures that the selection procedure is performed for selecting *spd* orbitals. Note that the implementation of such modified Algorithm 1 within `dfwpy` is an ongoing work. First, discarding high orbital types has already been implemented and we are currently implementing the *spd*-type contractions. An important point one should be careful is the use of the same density matrix for ABS generation as the one used in the *ab initio* electronic density of `GEM_fit`. The latter is reading the density matrix from a Gaussian formatted checkpoint file. The same file should be given as an input in our `dfwpy` code, in particular it should replace the current PySCF interface used to obtain the density matrix altogether. Lastly, our code contains an implementation of a dedicated parser for formatting ABS parameters in `GEM_fit` readable format. This setting enables the direct application of our ABS to multipole moment generation based on `GEM_fit`. Using such an implementation, our future work is to compare results on our ABS multipoles to the state-of-the-art ones, found in [82].

Appendix B

Force fields

In the present section, we formulate the theory of multipole expansions for calculating Coulomb electrostatics of charge distributions. We provide our proofs for deriving known formulas in Cartesian form used in AMOEBA (see Section B.2.3), following original theory from [71]. The present chapter is organised into two parts. The first one is theoretical and provides a presentation of our proofs on the derivation of multipole expansions. The second part concerns applications of the Gaussian Electrostatic Model (GEM) (see definition in Chapter 1.1.6.4) to force fields, based on multipole expansions.

Recalling Eq. (1.36), when penetration between two electron densities ρ_A and ρ_B is negligible, the electrostatic energy of their interaction is

$$V_{AB}^{\text{elec}} = \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{\rho_A(\mathbf{r})\rho_B(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r} d\mathbf{r}'. \quad (\text{B.1})$$

The electrostatic potential energy due to the charge distribution ρ_A reads

$$\forall \mathbf{r} \in \mathbb{R}^3, \quad \Phi_A(\mathbf{r}) = \int_{\mathbb{R}^3} \frac{\rho_A(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r}'. \quad (\text{B.2})$$

B.1 The Cartesian multipole expansion

In practice, integrals (B.1) and (B.2) can be evaluated using a multipole expansion, presented here. Consider two expansion centers, not necessarily atomic ones, denoted by $\mathbf{A} = (A_x, A_y, A_z)$ and $\mathbf{B} = (B_x, B_y, B_z)$, associated to molecules A and B , respectively. The basic mathematical tool for analytically evaluating integral formulas for electrostatics is to develop the Coulomb potential, defined as the function

$$\forall \mathbf{r} \in \mathbb{R}^3, \quad v(\mathbf{r}) := |\mathbf{r}|^{-1},$$

in a multipole expansion. Such expansions have been well established for some time, involving the development of $v(\mathbf{r} - \mathbf{r}')$ in either spherical harmonics using spherical polar coordinates [114, Appendix V], or a Taylor series in Cartesian coordinates [58]. We will focus on the Cartesian form.

The multipole expansion mainly relies on the Taylor expansion, that develops a multivariate function $F : \mathbb{R}^3 \rightarrow \mathbb{R}$ by a series with arbitrary expansion center $\mathbf{R} = (R_x, R_y, R_z)$ as

$$\forall \mathbf{r} \in \mathbb{R}^3, \quad F(\mathbf{r}) = \sum_{n_x, n_y, n_z=0}^{\infty} \frac{(x - R_x)^{n_x} (y - R_y)^{n_y} (z - R_z)^{n_z}}{n_x! n_y! n_z!} \left(\frac{\partial^{n_x + n_y + n_z} F}{\partial x^{n_x} \partial y^{n_y} \partial z^{n_z}} \right) (\mathbf{R}).$$

The main ingredients of the multipole expansion are the *multipole moments* of rank ℓ centered on the expansion center \mathbf{A} of the charge distribution ρ_A , defined for every $(l, m, n) \in \mathbb{N}^3$ with

$\ell = l + m + n$ as the scalar quantity

$$Q_{lmn}^A(\mathbf{A}) = \frac{1}{l!m!n!} \int_{\mathbb{R}^3} (x - A_x)^l (y - A_y)^m (z - A_z)^n \rho_A(\mathbf{r}) \, d\mathbf{r}, \quad (\text{B.3})$$

similarly for ρ_B and its expansion center \mathbf{B} , and the *interaction moment tensor* centered at arbitrary point $\mathbf{r} \in \mathbb{R}^3$, defined for every $(l, m, n) \in \mathbb{N}^3$ as the scalar

$$T_{lmn}(\mathbf{r}) = \left(\frac{\partial^\ell v}{\partial x^l \partial y^m \partial z^n} \right) (\mathbf{r}).$$

The following terminology is used to refer to multipole moments of different orders: the multipole moment of order one is called *monopole* (or charge), of order two *dipole*, of order three *quadrupole*, of order four *octupole*, and so on.

We formulate two main results on multipole expansions.

Theorem B.1.1 (Electric scalar potential). *The electric potential Φ_A produced by the charge density ρ_A admits the following multipole expansion with expansion center \mathbf{A} :*

$$\forall \mathbf{r} \in \mathbb{R}^3, \quad \Phi_A(\mathbf{r}) = \sum_{L,M,N=0}^{\infty} Q_{LMN}^A(\mathbf{A}) T_{LMN}(\mathbf{A} - \mathbf{r}). \quad (\text{B.4})$$

Proof. Let us use the notation $\mathbf{r}' = (x', y', z')$. Taylor expanding the Coulomb potential v with expansion center $\mathbf{A} - \mathbf{R} - \mathbf{r}'$, where $\mathbf{R} = (R_x, R_y, R_z)$ arbitrary, and then letting $\mathbf{R} = -\mathbf{r}$, yields

$$v(\mathbf{r}) = \sum_{L,M,N=0}^{\infty} \frac{(x' - A_x)^L (y' - A_y)^M (z - A_z)^N}{L!M!N!} T_{LMN}(\mathbf{A} + \mathbf{r} - \mathbf{r}'). \quad (\text{B.5})$$

Using that $v(-\mathbf{r}) = v(\mathbf{r})$, Φ_A can be expressed in terms of v as

$$\Phi_A(\mathbf{r}) = \int_{\mathbb{R}^3} v(\mathbf{r}' - \mathbf{r}) \rho_A(\mathbf{r}') \, d\mathbf{r}'.$$

Substituting v by the expansion (B.5) evaluated at point $\mathbf{r}' - \mathbf{r}$, yields the result. \square

Theorem B.1.2 (Electrostatic interaction energy). *The electrostatic interaction energy V_{AB}^{elec} between charge densities ρ_A and ρ_B admits the following multipole expansion with expansion centers \mathbf{A} and \mathbf{B} :*

$$V_{AB}^{\text{elec}} = \sum_{LMN}^{\infty} \sum_{L'M'N'}^{\infty} Q_{LMN}^A(\mathbf{A}) T_{L+L',M+M',N+N'}(\mathbf{A} - \mathbf{B}) Q_{L'M'N'}^B(\mathbf{B}).$$

Proof. Taylor expanding Φ_A with expansion center \mathbf{B} yields

$$\Phi_A(\mathbf{r}) = \sum_{L',M',N'=0}^{\infty} \frac{(x - B_x)^{L'} (y - B_y)^{M'} (z - B_z)^{N'}}{L'!M'!N'!} \left(\frac{\partial^{\ell'} \Phi_A}{\partial x^{L'} \partial y^{M'} \partial z^{N'}} \right) (\mathbf{B}). \quad (\text{B.6})$$

We calculate partial derivatives of Φ_A by differentiating its expansion (B.4) evaluated at $\mathbf{r} = \mathbf{B}$, to obtain

$$\left(\frac{\partial^{\ell'} \Phi_A}{\partial x^{L'} \partial y^{M'} \partial z^{N'}} \right) (\mathbf{B}) = \sum_{L,M,N=0}^{\infty} Q_{LMN}^A(\mathbf{A}) T_{L+L',M+M',N+N'}(\mathbf{A} - \mathbf{B}).$$

Substituting this expression into (B.6), and relating Φ_A with V_{AB}^{elec} , using

$$V_{AB}^{\text{elec}} = \int_{\mathbb{R}^3} \Phi_A(\mathbf{r}) \rho_B(\mathbf{r}) \, d\mathbf{r},$$

yields the result. \square

B.1.1 Related numerical aspects

Multipole expansions are widely used techniques for improving the evaluation of long-range separation Coulomb interactions in terms of computational cost [134, 294, 45, 99]. They can be also incorporated to other methods for treating short- and long-range, such as CFMM (see Section 1.2.3). The purpose of multipole expansions is to reduce the computational cost of a direct evaluation of bielectronic integrals. We may distinguish the following difficulties in numerically computing multipole expansions.

Expansion sites. In practice, infinite multipole expansions are truncated up to a finite degree. Thus, in total, the cost of computing a multipole expansion is dominated by the degree of truncation and the number of sites used to expand multipoles. However, the total number of centers of an electronic charge distribution obtained by *ab initio* methods on a Gaussian basis set is at most $N_b(N_b + 1)/2$, due to the Gaussian product theorem (see Theorem C.1.1). For such full set of expansion sites, the infinite series truncation error is zero, since it can be shown that the expansion can be truncated to the maximal angular moment of the charge distribution without loss of accuracy (see Appendix C.2.0.1). The best choice of fewer expansion sites typically coincides with points of large concentrations of Gaussian product centers, e.g. atoms and bond mid-points. Other works consider the full set of expansion centers [271]. While this yields finite expansions at each point, it results in a heavy number of points that are basis set dependent. A related work on the exact multicenter expansion can be found in [328].

Numerical methods are used in practice for reducing the number of centers. Ideally, the number of expansion sites should be linear in the number of atoms of a molecule. Note that the multipole expansion is exact if its order is infinite, independently of the choice of expansion sites. While truncating the expansion, however, the choice of expansion sites becomes important. Generally, the choice of expansion sites greatly affects the accuracy of the truncated multipole expansion. This is due to the fact that infinite series can no longer be truncated without loss of accuracy, as soon as a multipole center is not the true one, i.e. if the density is not localized on the multipole center in the analytical sense. For this reason, there is a need for methods that redistribute the density localization to a small number of sites.

Well-separated distributions. Another aspect to be taken into account is that the choice of expansion sites greatly depends on the separation distance between interacting molecules. The closer the molecules are, the more expansion sites we need to converge. In the case of self-interaction, the sites would be too many, since we cannot spatially separate atomic distributions of the same molecule. In general, the multipole expansion cannot be used when electron densities overlap numerically [345], due to lack of convergence. Hence, for not well-separated interactions, the calculation is not accelerated since direct integral evaluation is necessary. On the other hand, the multipole expansion can be safely employed for well-separated distributions. Due to these reasons, a feature that slows down multipole expansion calculations is the need to identify sets of distributions whose numerical supports do not overlap.

Sources of approximation error. To summarize, we may identify two sources of errors in the approximation electrostatic interaction energy between molecular charge distributions using truncated multipole expansions: (a) the choice of expansion centers, and (b) the truncation order. Both factors affect the accuracy and the convergence of the multipole expansion. Generally, the radius of convergence of the Taylor expansion of the potential Φ_A of molecule A determines the converge of the multipole expansion. Ideally, the molecule B should lie within the convergence radius. The convergence radius depends on the magnitude of the multipole moments, which may be influenced by the choice of expansion centers, as well as by the spatial exponent of the charge density.

Numerical schemes achieve a trade-off between a low number expansion centers and convergence at low order, within a tolerance value [305, 271, 101]. Using such schemes, the truncated

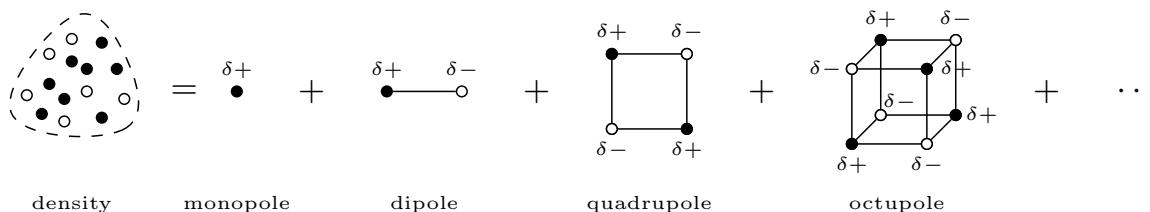


Figure B.1: Approximating a charge distribution using the multipole expansion.

multipole expansion is typically an accurate approximation of electrostatics at long-range separations, but may fail to converge at short-range [306, Section 8.1]. Improved choices of expansion centers based on Distributed Multipole Analysis (DMA) [305] manage to reach convergence at all separations. However, at short-range separation when the molecular charge distributions overlap, the converged value has an error by a constant factor called the *penetration energy*. The physical meaning attributed to this error is that, at short-range, the charge density functions overlap, causing the nuclei of one molecule to be attracted to the electrons of the other molecule [119].

B.2 Application to chemistry

In the present section, we summarize relevant aspects of calculating multipole expansions in force field applications for molecular simulations, as well as practical considerations of the AMOEBA polarizable force field parametrization in Section B.2.3. The following content aims at gathering information about these topics, as well as pointing out useful elements and providing an overview that needs to be completed in future studies.

B.2.1 Modeling electrostatics of charge distributions

The chemistry intuition behind multipole moments can be described as follows. The electrostatics of charge distributions comprising of electrons is given by the formula of Eq. (B.1). However, the exact evaluation of this quantity can be computationally expensive in practice (see Section 1.2). In order to reduce the complexity with respect to the number of charges, an approximation based on multipole moments is used, belonging to a class of analytical techniques. At first order, charge distributions are approximated by fixed point charges, describing electronegativity effects in a chemically intuitive way. At second order, dipole moments arise from overlap of *s* and *p* orbitals and describe lone pairs and other atomic distortions. Higher order terms are quadrupoles, octupoles, etc. Figure B.1 schematically represents the main idea. For example, moments up to octupoles are used in SAPT [247], up quadrupoles in AMOEBA and GEM. Multipole moments and expansion sites can be generated by empirical rules, numerical fit or *ab initio* density-based methods. Note that, in principle, several charges are necessary to capture the inhomogeneities of the charge distribution for polar molecules. The multipole expansion is highly related to the concept of polarization. Summing up, the truncation order of multipole expansion gives rise to different types of force field methods, described in the following.

B.2.2 Force fields

An alternative method for computing potential energy surfaces for MD, other than the previously discussed electronic structure theory of Section 1.1.2, is the method of force fields. In the present section, we present force fields, before focusing on numerical models used in force fields based on Gaussian discretizations, namely the Gaussian electrostatic model [256]. For more details on force fields, we refer to chemistry literature [346, 33, 131, 17, 240].

Force Fields (FF) are potential energy models based on the following functional form, which is expected to be computationally more efficient than any *ab initio* model, due to the absence of electronic coordinates:

$$U \approx U_{\text{FF}} = U_{\text{intra}}(\mathbf{R}_1, \dots, \mathbf{R}_{N_a}; \alpha_1, \dots, \alpha_t) + U_{\text{inter}}(\mathbf{R}_1, \dots, \mathbf{R}_{N_a}; \alpha'_1, \dots, \alpha'_s).$$

Here, N_a is the number of atoms at fixed positions $\mathbf{R}_1, \dots, \mathbf{R}_{N_a}$, $(\alpha_i)_{1 \leq i \leq t}$ is a set of intramolecular parameters and $(\alpha'_i)_{1 \leq i \leq s}$ a set of intermolecular ones. This distinction between two parameter sets is motivated by the standard theory of intermolecular forces. In particular, one group (*intra*) is dedicated to covalent bonds due to sharing of electrons between atoms, holding together atoms to form molecules. A second group (*inter*) concerns weaker forces, holding fragments together to form liquids or solids. Note that the choice of FF parameter sets is not unique, but empirical and heavily system and context specific. In a sense, FF methods can be interpreted as a reparametrization of the *ab initio* energy functional, with the goal of getting rid of the explicit presence of electronic coordinates, responsible for bad scaling. Electronic degrees of freedom are only taken into account implicitly, in favour of computational speedup.

Computational efficiency of force fields mainly lies on notions of system fragmentation and parameter precomputation. Namely, the starting point is to presuppose a decomposition of the total system into smaller subunits of appropriate chemical type, referred to as *fragments* or *monomers*, for which precomputations can be carried out separately [132]. Fragmentation is hence intrinsic to FF parametrization. Before the start of simulation, parameter values are precomputed and assigned by fitting to reference data for isolated individual components of a chemical type, building a frozen fragment library. Reference data is either obtained experimentally, e.g. by infrared spectroscopy or vibrational frequencies, or by high-accuracy *ab initio* data on fragments at equilibrium geometry. Launching the molecular dynamics simulation will evolve the interaction of fragments and update their geometric conformation, based on their parameters. Force fields strike a balance between accuracy and speed. Most importantly, the parametrization needs to be flexible enough to accurately capture changes in molecular conformations. This is a challenge for modern force fields, known as *transferability* [170]. According to this property, frozen parameters on individual fragments shall be capable of predicting conformations for larger systems.

Let us focus on intermolecular parametrization models for force fields. Our study is motivated by the fact that intermolecular forces can be a bottleneck for MD simulations on multi-atom systems. The electrostatic interaction energy is one of the components of the intermolecular energy, along with van der Waals interaction energy. This term is unimportant for non-polar systems such as hydrocarbons and is in these cases often completely neglected. For polar systems, such as water, however, the long-range electrostatic interaction is often the leading contribution. The electrostatic intermolecular energy is defined by the classical Coulomb electrostatic model, presented in the following paragraph. For more details on electrostatics and intermolecular forces, we refer to [306, 58].

B.2.2.1 Classical models

At the lowest approximation, the classical force fields compute electrostatics from point charges assigned to atoms. Such charges are calculated during the parametrization and remain fixed. Note that this parametrization doesn't capture polarizability. The classical electrostatic energy model between two point charges q_i and q_j , separated by a distance r_{ij} , is given by the Coulomb law as

$$(U_{\text{Coul}})_{ij} = \frac{q_i q_j}{r_{ij}}.$$

The total electrostatic energy of the system of N_a atoms is obtained by summing all pairwise interactions,

$$U_{\text{Coul}} = \sum_{1 \leq i < j \leq N_a} (U_{\text{Coul}})_{ij}.$$

Such summation can be accelerated computed using the fast multipole method (FMM) [134] or PME [93] for periodic systems. Popular classical force fields include CHARMM [162], AMBER [332], OPLS [175], GROMOS [296].

The limitations of fixed point charges regarding electrostatics are several. The atomic point charges can be assigned by empirical rules [124], by fitting to *ab initio* electrostatic potentials, or by using LCAO approximations, such as the Mulliken population analysis [238]. However, the fitting is often poorly conditioned [118]. It is well known that the partial atomic charges depend on the molecular geometry [190, 189, 142], but this is ignored in fixed-charge force fields. Moreover, the partial charge model gives a rather crude representation of the electrostatic potential surrounding a molecule, with errors often being in the 10 – 20 kJ/mol range [170]. It also fails to include charge penetration effects [333], that consist of an attractive force between two atoms arising at van der Waals distance due to overlap of their electron clouds. The electrostatic interaction model can be systematically improved by including non-atom-centered charges [101] or higher-order multipole moments [65].

We briefly mention one of the available methods for deriving point charges. The ElectroStatic Potential (ESP) [170], at a point \mathbf{r} , is given by the nuclei $(z_i, \mathbf{R}_i)_{1 \leq i \leq N_a}$ and the electron density ρ as

$$\forall \mathbf{r} \in \mathbb{R}^3, \quad U_{\text{ESP}}(\mathbf{r}) = \sum_{i=1}^{N_a} \frac{z_i}{|\mathbf{r} - \mathbf{R}_i|} - \int \frac{\rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r}'.$$

A method for calculating partial charges consists of fitting charges to the ESP function of periodic and nonperiodic systems [72, 160, 59]. Such method is known as ESP fitting. In particular, the fitting is done by minimizing a least-squares error function of the form

$$\forall \mathbf{q} = (q_i)_{1 \leq i \leq N}, \quad \mathcal{E}(\mathbf{q}) = \frac{1}{m} \sum_{j=1}^m \left(U_{\text{ESP}}(\mathbf{r}_j) - \sum_{i=1}^N \frac{q_i}{|\mathbf{r}_j - \mathbf{p}_i|} \right)^2,$$

for given partial charge positions $(\mathbf{p}_i)_{1 \leq i \leq N}$ and sampling points $(\mathbf{r}_i)_{1 \leq i \leq m}$, under the constraint that the sum of the partial charges is equal to the total molecular charge. The electrostatic potential is typically sampled at a few thousand points in the near vicinity of the molecule [56, 78]. This suffers from highly non-transferable charges [118], i.e. two conformations of the same molecule may give significantly different sets of fitted parameters. This is a main limitation for force fields that rely on the transferability assumption, according to which similar fragments and atoms that are easily interchanged (e.g. by bond rotation) should have identical parameters.

B.2.2.2 The need for polarizability

In practice, not all charge distributions can be considered as point charges. This is the case of polar molecules, i.e. molecules whose electronic charge distribution is uneven. Such concept is described by the notion of polarizability in chemistry, that we introduce with the following example.

The example of HCl. Electronic polarization describes the redistribution of electron density due to an external field and is defined as the scalar quantity $\alpha = |\boldsymbol{\mu}_{\text{ind}}|/|\mathbf{E}|$, where $\boldsymbol{\mu}_{\text{ind}}$ is the induced dipole moment of the molecule to the electric field \mathbf{E} that produced this dipole moment. Consider the diatomic molecule HCl. The seventeen positively charged protons of the chlorine atom attract the single electron strongly bonded to the hydrogen atom. This creates an uneven electron distribution and makes HCl a polar molecule. As a result, a partial negative point charge δ^- is created on Cl and a partial positive point charge δ^+ is created on H. The two point charges define the *permanent dipole* of HCl, that is a vector pointing from the negative to the positive charge. Now suppose that HCl interacts with the non-polar molecule Cl_2 , whose electron cloud is evenly distributed. The negative partial point charge of HCl will repel the electron cloud of Cl_2 , acting as an external electric field. This field will redistribute the electrons




Figure B.2: Creation of an induced dipole.

of Cl_2 unevenly in space, creating an *induced dipole*, as shown in Figure B.2. The induced dipole can then interact with all dipoles in the system via intermolecular forces, or even contribute to the induced dipoles of neighbouring molecules, so it is important to take it into account when modeling electrostatics. The difficulty is that induced dipoles are flexible and highly dependent on molecular conformations.

B.2.2.3 Permanent and induced multipoles

Polarization can be taken into account in the force field by including permanent multipoles of the form (B.3) into the parametrization. The resulting FF is known as a *polarizable* force field. The permanent multipoles are fed into the parametrization method in order to tune the remaining parameters. In that case, all parameters of the force field are build upon the base of the permanent dipoles. Inclusion of higher-order electric moments will improve the representation of the electrostatic potential of a given molecule, but it does not take the geometry dependence into account. The variation of the electrostatic energy with the geometry must be accurately modeled in order for the remaining frozen intramolecular parameters to be sufficiently transferable, and this is difficult to achieve without including polarization [170]. To do so, induced dipoles must be taken into account, calculated from the permanents ones.

B.2.3 The AMOEBA polarizable force field

AMOEBA (atomic multipole optimized energetics for biomolecular applications) is a polarizable force field [280, 282, 263] proposed in 2002 by Ponder and co-workers. Focusing on the the intermolecular energy contribution, it is given as the sum of three terms, namely the non-bonded van der Waals and the Coulomb electrostatic contributions from permanent and induced dipoles, reading

$$U_{\text{inter}}^{\text{AMOEBA}} = U_{\text{vdW}} + U_{\text{Coul}}^{\text{perm}} + U_{\text{Coul}}^{\text{ind}}. \quad (\text{B.7})$$

The AMOEBA permanent electrostatics model is based on point multipole moments (charge, dipole and quadrupole) centered on atoms. Atomic multipoles are derived from *ab initio* calculations on small molecules using Stone’s distributed multipole analysis (DMA) [305], described in more detail in Appendix B.2.3.2. Note that molecules are flexible, in the sense that the multipole expansion sites are allowed to move, and the multipoles also move along the sites’ orientation. The orientation is calculated from local coordinate frames associated to sites [263]. For energy evaluations, the locally defined multipole moments are rotated into the global frame.

The permanent electrostatics term is given by a truncated multipole expansion in matrix form, as explained in the following. Truncating the multipole expansion in low orders is preferred due to the high computational cost of evaluating each multipole moment and interaction tensor entry using recursion relations (see Appendix C.2). Assuming that the expansion centers are the points \mathbf{A} and \mathbf{B} for molecules A and B , respectively, the truncated multipole expansion of order n can be written in matrix form as follows. We only use moments with L, M, N such that

$L + M + N \leq n$, ordered as

$$Q_{000}^A, Q_{100}^A, Q_{010}^A, Q_{001}^A, Q_{200}^A, Q_{110}^A, Q_{101}^A, Q_{110}^A, Q_{020}^A, \dots$$

where the center \mathbf{A} is omitted for brevity, similarly for ρ_B . These moments are gathered into the multipole vector denoted by $\mathbf{M}_A^{(n)}$ for molecule A , similarly for B . Finally, $\mathbf{T}_{AB}^{(n)}$ is the square matrix containing the interaction moments for expansion sites \mathbf{A} and \mathbf{B} , its columns and rows are ordered in the same manner as the multipole vectors. For example, the multipole expansion in AMOEBA is truncated at order $n = 2$. The Coulomb electrostatics between molecules i and j is approximated by a finite expansion, given as

$$(U_{\text{Coul}}^{\text{perm}})_{ij} = \mathbf{M}_i^\top \mathbf{T}_{ij} \mathbf{M}_j, \quad (\text{B.8})$$

The total electrostatic energy of N atoms is obtained by summing up all pairwise interactions

$$U_{\text{Coul}}^{\text{perm}} = \sum_{1 \leq i < j \leq N} (U_{\text{Coul}}^{\text{perm}})_{ij}.$$

To sum up, given N atoms, the evaluation of permanent electrostatics using AMOEBA can be divided into two computational stages. First, permanent electrostatic interaction energy $U_{\text{Coul}}^{\text{perm}}$ is precalculated for all atoms. This allows to recover a permanent multipole \mathbf{M}_i . Then, during MD simulation on pairs of fragments, the interaction matrix \mathbf{T}_{ij} is computed and the matrix-vector product $\mathbf{M}_i^\top \mathbf{T}_{ij} \mathbf{M}_j$ is calculated, to obtain the pairwise energy. The sum of pairwise energies yields the total permanent interaction of the system of N atoms.

The main acceleration consists of precomputing the multipoles and then performing simply an FMM online, as it is done in practical codes [4]. The main idea is to fix the sites. Once we have that, the problem writes as an FMM problem. The main problem is the site expansion. We discuss state of the art methods to expand a molecular density to sites.

AMOEBA takes into account polarization using the induced dipole model [285, 284]. Updating dipole moments at every geometry using an iterative approach [315], that allows to make the charges fully flexible and account for geometry dependence.

B.2.3.1 Distributed multipole analysis

The Distributed Multipole Analysis (DMA) [305] is a method for describing the spatial distribution of electric charge within a molecule, based on continuous densities. Multipoles of charge distributions obtained by DMA are especially adapted to the Coulomb electrostatics evaluation by the multipole expansion. In particular, the radius of convergence of this multipole series is sufficiently small that the relevant series will be convergent when describing two molecules in van der Waals contact.

AMOEBA permanent electrostatics is primarily based on DMA methods, namely the Gaussian version. The Distributed Multipole Analysis for Gaussian wavefunctions (GDMA) method [308, 307] has been proposed by Stone. The main idea of GDMA can be roughly described as follows. First, the electronic density function is derived from *ab initio* or density functional theory calculations using Gaussian basis sets. Then, GTO products are expressed as Gaussian overlap distributions using the Gaussian product rule (see Theorem C.1.1 in Appendix C). The center of the GTO overlap will generally be some site. If it is an atomic site, then the product along with its center is included in the distributed multipole. Otherwise, the site is redistributed to force the product to be atom-centered or mid-bond-centered. The redistributed density can be obtained based on real-space methods (see Section 1.2.3.1), Mulliken charges [238], or empirical methods for assigning spatial weights. As an example of empirical treatment, the origins of multipole moment of each piece of charge density are simply moved to the nearest expansion site. In particular, if the overlap center is not at an atom, one can move the origin of the multipole expansion to the nearest distributed multipole site, re-expressing the series to account for the change of origin.

By redistributing sites along the fragment, DMA is an improvement to the point multipole method that considers only atomic sites. The reason for not distributing sites to atoms only is the following. Assigning a single moment to every fragment has poor convergence properties [40]. For example, if the fragments are too close together, the expansion no longer converges. To solve the convergence issue, a distributed multipole moment is used. This consists of placing a multipole moment for a chosen number of sites distributed throughout each fragment.

Summing up, DMA constructs localized charge distributions along with their centers, that are usually atoms and bond midpoints. Then, multipole moments are calculated using such distributed densities, using DMA centers as expansion sites [40]. AMOEBA obtains the permanent multipole moments on each atom based on DMA, at the MP2/aug-cc-pVTZ level for the experimental equilibrium geometry of the gas-phase monomer. A tolerance equal to 0.5 kcal/mol has been used in the approximation of *ab initio* electron density multipoles by DMA.

B.2.3.2 Parametrization in AMOEBA

The entire AMOEBA parametrization, i.e. intramolecular and multipole moments combined, is obtained using the Poltype 2 method [331], which is a step-by-step automated parametrization procedure based on error correction by error compensation. Roughly speaking, it fits the parameters of the $n + 1$ term to the fit error of the n term. At initialization $n = 1$ we find the distributed multipoles from QM densities of isolated monomers computed using GDMA (see Section B.2.3.1). Next, we move on to intermolecular van der Waals parameters, fitted to the error of the AMOEBA intermolecular energy of the dimer AB with respect to the QM intermolecular energy $E_{AB} - E_A - E_B$. The last treated is intramolecular parameters, torsion energy being the final one, fitted to QM data of the isolated monomer.

B.2.4 Generation of multipoles from GEM

The present section discusses some of the existing methods for generating multipole moments, with force field applications, and briefly mentions how density fitting, which is one of the main methods developed in the present thesis (see Chapter 2), could enable to improve them.

B.2.4.1 Electrostatics in GEM

The present section provides more insight into the background of the Gaussian Electrostatic Model (GEM), introduced in Section 1.1.6.4. GEM also has applications to force fields, yielding another class of force fields is based on densities. The underlying theoretical background of such methods is the Energy Decomposition Analysis (EDA) [187, 237, 90]. This method allows to take into account effects related to the interaction of electrons of different fragments, such as the charge penetration of electron clouds. To do so, it includes physically meaningful energy terms. The term that allows to include corrections to the error due to the fragmentation of the density (1.45) is the *exchange-repulsion* term. To summarize, the two contributions we focus on are the following [42, 171, 306]:

- Electrostatic: electrostatic energy of interaction, calculated classically from the unperturbed charge distributions. It dominates at long-range separations. It is calculated from classical Coulomb electrostatics.
- Exchange-repulsion: the repulsive force resulting from exchange of electrons and the Pauli exclusion principle [249, 272], according to which electrons of the same spin may not be at the same place. It dominates at short-range, where the electron exchange between molecules becomes possible [306, 290]. It is calculated from the density overlap model of Wheatly and Price [338].

At this point, it may be useful to mention that such energy terms can be rigorously defined within perturbation theory [277, 11], in which interaction energy terms are obtained as higher-

order corrections to the ground state of the unperturbed system, when applying as Rayleigh-Schrödinger perturbation [179] the Coulomb interaction operator for electrons and nuclei between monomers (more details can be found in [306, Section 4.1]). In such context, the electrostatic interaction energy is obtained as a first order long-range approximation to the interaction energy. Additionally, within symmetry-adapted perturbation theory (SAPT) [172, 171, 247], one can apply a second perturbation to the unperturbed system, in order to account for short-range interaction, by forcing the correct permutational symmetry (1.5) to the dimer wave function. This operator is known as antisymmetrization operator and the exchange-repulsion energy [145] is obtained as a first order short-range approximation to the interaction energy when applying this antisymmetrization perturbation [171, Equation 112].

Short-range separation: electron exchange correction. The classical interaction V_{AB} (1.37) fails to take into account all possible intermolecular interactions between molecules A and B . In particular, electron clouds of A and B overlap, making quantum effects possible, including electron exchange [306]. Moreover, nuclei of one molecule can be attracted to the electrons of the other molecule, consuming energy which cannot be computed using the classical Coulomb interaction. The energy difference resulting from this increased attraction is known as *charge penetration* [119], also mentioned earlier in Section B.1.1. As a result, the electrostatic energy V_{AB} yields an accurate approximation of the intermolecular interaction energy at long-range, but fails at short-range separation due to charge penetration effects.

Continuous descriptions of molecular charge distributions are known to reduce the charge penetration error [105]. The energy term that allows us to reduce this error is the exchange-repulsion defined in terms of continuous densities. Note that GEM reduces penetration errors.

B.2.4.2 AMOEBA multipoles from GEM

The fitted density generated by GEM, using density fitting, can be used to compute distributed multipole moments, in a scheme known as GEM-DM [82]. Multipole moments generated by GEM-DM have been applied to the AMOEBA polarizable force field. Results show that GEM-DM provides accurate multipoles for several systems, such as ionic liquids [327, 320, 81, 105, 259]. The advantage of GEM-DM is that it attributes polarizable multipole moments centered on atoms using a density-based systematic approximation method. This paves the way towards parametrization of force fields from first principles [257, 82, 129]. In [136, Figure 9], electrostatics from GEM have been compared to partial charges obtained by ESP fitting, namely the Merz-Singh-Kollman (MK) scheme [297], in which atomic charges are fitted to reproduce the ESP at a number of points around the molecule. Results show that for the water molecule, electrostatic potential maps calculated from GEM are better than the ones of MK charges, with respect to reference *ab initio* calculations.

GEM can generally be coupled to QM/MM multi-scale methods [85] and to AMOEBA force field, in order to improve the treatment of electrostatics, in a hybrid density-based force field known as GEM* [256, 105]. In particular, in such force field, the intramolecular and polarization terms are computed with AMOEBA and the remaining intermolecular terms with GEM. Such treatment can reduce charge penetration energy errors, present in DMA-based electrostatics, due to the exchange-repulsion term of GEM.

B.2.4.3 Brief comparison of DMA- and DF-based multipoles

Here, we briefly illustrate one of the possible applications of the Density Fitting (DF) method (see Chapter 2), that we wish to explore further in the future, following the present thesis. In particular, DF can be seen as an alternative to DMA, as if instead of only choosing sites we also update the coefficients by fitting. The existent density fitting methods are atom-centered methods [158], that is, auxiliary functions are centered on the same atom as the composite functions. In particular, the physical constraint imposed on the space of auxiliary functions is that they must be centered uniquely on atomic sites of the system. The reasoning behind such

constraint is related to the number of centers of the multipole expansion, that must remain small for computational complexity reasons.

DMA has found extensive use in crystal structure prediction for small organic molecules, where significant progress can often be made while using rigid molecular structures. It has also been applied to the non-rigid case in the AMOEBA force field. A limitation of DMA are charge penetration energy errors [105]. Let us emphasize the conformation-dependence of the DMA. As a consequence of its high level of accuracy, DMA captures features of the molecular charge distribution that depend strongly on molecular conformation. Thus, DMA-based multipoles would have to be recalculated whenever a molecule underwent a conformational change. The important point is that AMOEBA employs DMA-based multipoles without recalculating multipoles at every timestep. However computationally cheap, this introduces an approximation error. Moreover, DMA generally introduces a multipole expansion error related to the selection of expansion sites [40]. This creates an error because if the expansion site is not the same as the center of the Gaussian, then the truncation accuracy depends on the basis function type, if it is diffuse, if we evaluate the potential at a point sufficiently far from the expansion site, see [40].

Density fitting offers a systematic alternative to this, as the coefficients of GTO products will be modified as to better reproduce the initial density after modifying the sites. A perspective of density fitting is localizing contributions to atomic sites. The mathematical framework that allows to formulate this notion of localized contributions is left for future investigation.

B.3 Summary

In the present section, we saw the multipole expansion, a widely used method for accelerating interaction evaluation. The calculation of multipole moments requires constructing multi-center decompositions of the molecular density. An efficient and accurate density decomposition should satisfy two criteria:

- have a small number of centers,
- minimize an appropriate error functional with respect to the *ab initio* density.

DMA is a standard method for constructing such decompositions, currently used in AMOEBA polarizable force field. Density fitting can be an alternative method for constructing such decompositions, based on systematic density approximations under tunable accuracy.

Appendix C

Molecular integral evaluation

For the sake of completeness, we summarize useful elements of molecular integral evaluation for Gaussian-type orbitals. We gather few known and basic results from chemistry textbooks [149, 150].

C.1 Preliminaries

C.1.0.1 Gaussian overlap distribution

A product of two Gaussians, generally of different centers, is referred to as an *overlap distribution* [150]:

$$\Omega_{ij}(x, a, b, A_x, B_x) = G_i(x - A_x, a)G_j(x - B_x, b).$$

According to the following Gaussian product rule of Theorem C.1.1, a Cartesian overlap distribution may be written in terms of a single Gaussian centered on \mathbf{P} ,

$$\Omega_{ij} = K_{AB}x_A^i x_B^j \exp(-px_P^2).$$

Theorem C.1.1. [Gaussian product rule [150, Section 6.6.7]]

$$\exp(-ax_A^2) \exp(-bx_B^2) = K_{AB} \exp(-px_P^2)$$

where

$$\begin{aligned} p &= a + b, \\ P_x &= \frac{aA_x + bB_x}{p}, \\ q &= \frac{ab}{p}, \\ Q_x &= A_x - B_x, \\ K_{AB} &= \exp(-qQ_x^2). \end{aligned}$$

C.1.0.2 Hermite-Gaussian orbitals

These are intermediate orbitals commonly used for direct evaluation of molecular integrals, that we define here.

A Cartesian Hermite-Gaussian orbital (HGO) of exponent α centered on zero is defined as the function

$$\Lambda_{tuv}(\mathbf{r}, \alpha) = (-1)^{t+u+v} \frac{\partial^{t+u+v}}{\partial x^t \partial y^u \partial z^v} \exp(-\alpha|\mathbf{r}|^2).$$

Like GTOs, these functions are separable in factors such as

$$\Lambda_t(x, \alpha) = (-1)^t \frac{\partial^t}{\partial x^t} \exp(-\alpha x^2).$$

The Hermite polynomials H_t , defined by

$$H_n(x) = (-1)^n \exp(x^2) \frac{d^n}{dx^n} \exp(-x^2),$$

are related to Hermite-Gaussians through the equation

$$\Lambda_t(x, \alpha) = \alpha^t H_t(\sqrt{\alpha}x) \exp(-\alpha x^2). \quad (\text{C.1})$$

This type of basis function is mainly used as intermediate expression for evaluating integrals on GTOs, as illustrated in following chapters. The important point is that Hermite-Gaussians span the space of Gaussian-type orbitals of the same orbital exponent.

C.1.0.3 GTO to HG transformation

The Hermite polynomials satisfy the relation

$$H_L(x) = 2xH_{L-1}(x) - 2(L-1)H_{L-2}(x),$$

which forms the basis of the McMurchie-Davidson two-term recurrence relations [228] for computation of the expansion coefficients.

There is a transformation matrix that allows to pass from Gaussian-type orbitals to Hermite-Gaussian orbitals, thus the two representations span the same space.

Here we explain McMurchie-Davidson recursion and we give the Gaussian product theorem. We use notation from [150, Section 10.3]. We can then transform using the GTO to HG transformation

$$\Omega_{ij} = \sum_{t=0}^{i+j} E_t^{ij} \Lambda_t.$$

Starting coefficient

$$E_0^{00} = K_{AB}$$

and relations

$$\begin{aligned} E_t^{i+1,j} &= \frac{1}{2p} E_{t-1}^{ij} - \frac{qQ_x}{a} E_t^{ij} + (t+1) E_{t+1}^{ij}, \\ E_t^{i,j+1} &= \frac{1}{2p} E_{t-1}^{ij} + \frac{qQ_x}{b} E_t^{ij} + (t+1) E_{t+1}^{ij}. \end{aligned}$$

It is assumed that the coefficient is zero if $t > i$ or $t > j$.

C.1.1 McMurchie-Davidson scheme

Moreover, we use the intermediate integral that admits the following closed-form expression

$$\int_{\mathbb{R}^3} \frac{\Lambda_{tuv}(\mathbf{r})}{|\mathbf{r} - \mathbf{C}|} d\mathbf{r} = \frac{2\pi}{p} R_{tuv}(p\mathbf{R}_{PC}^2),$$

where p and \mathbf{P} are the exponent and the center of the Hermite-Gaussian respectively. The auxiliary function R_{tuv} is defined as

$$R_{tuv} = \left(\frac{\partial}{\partial P_x} \right)^t \left(\frac{\partial}{\partial P_y} \right)^u \left(\frac{\partial}{\partial P_z} \right)^v F_0,$$

where F_0 is the incomplete Gamma function defined by

$$\forall x \geq 0, \quad F_n(x) = \int_0^1 \exp(-xt^2) t^{2n} dt.$$

This integral does not admit an analytic expression but it may be efficiently evaluated using approximation methods. The auxiliary function is computed recursively.

C.1.1.1 Auxiliary integrals

This is the exact method, however it is very slow. It is described in [142] and it is based on the auxiliary function R , introduced in the following, and the McMurchie-Davidson recursions.

Auxiliary integrals combine properties of Hermite-Gaussians and the Boys functions. This is only used for the direct interaction in GEM. However, this computation is not used in quantum chemistry codes due to its high cost. The first step is to transform GTO overlaps to HG functions.

In order to evaluate Coulomb integrals, we will need to set up an auxiliary Hermite Coulomb integral

$$R_{tuv}^n(s, \mathbf{S}, \mathbf{D}) = \left(\frac{\partial}{\partial S_x} \right)^t \left(\frac{\partial}{\partial S_y} \right)^u \left(\frac{\partial}{\partial S_z} \right)^v F_n(s \mathbf{R}_{SD}^2)$$

that handles the Coulomb interaction between a Gaussian charge distribution centered at \mathbf{S} and a nuclei centered at \mathbf{D} . We have set $\mathbf{R}_{SD} = |\mathbf{S} - \mathbf{D}|$ and $F_n(T)$ is the Boys function

$$F_n(T) = \int_0^1 x^{2n} \exp(-Tx^2) dx$$

which is a special case of the Kummer confluent hypergeometric function ${}_1F_1(a, b, x)$

$$F_n(T) = \frac{{}_1F_1(n + \frac{1}{2}, n + \frac{3}{2}, -T)}{2n + 1}$$

implemented in `scipy.special`. The Boys function for $n = 0$ appears due to the fact that

$$\int_{\mathbb{R}^3} \frac{\exp(-sr_S^2)}{r_D} d\mathbf{r} = \frac{2\pi}{s} F_0(s \mathbf{R}_{SD}^2).$$

The Hermite Coulomb integral, like its counterpart E_t^{ij} , is defined recursively:

$$\begin{aligned} R_{t+1,u,v}^n &= t R_{t-1,u,v}^{n+1} + X_{SD} R_{t,u,v}^{n+1} \\ R_{t,u+1,v}^n &= u R_{t,u-1,v}^{n+1} + Y_{SD} R_{t,u,v}^{n+1} \\ R_{t,u,v+1}^n &= v R_{t,u,v-1}^{n+1} + Z_{SD} R_{t,u,v}^{n+1} \\ R_{000}^n &= (-2s)^n F_n(s \mathbf{R}_{SD}^2). \end{aligned}$$

Lastly, according to [288], the number of R_{tuv}^n functions is

$$\sum_{j=0}^{2l+1} \binom{j+3}{3} = \binom{(2l+1)+4}{4}$$

C.2 Gaussian multipole expansion

The fast scheme uses multipole moments of Hermite-Gaussian functions. This is an alternative to exact recursive schemes.

C.2.0.1 Hermite-Gaussian multipole moments

This is related to truncations of the multipolar expansion. In the above equation, we shall assume that the multipole expansion is truncated at order \mathcal{L}_a , the order of the series expansion about \mathbf{A} .

We can generate higher-order moments from those of lower order using the recursion [150, Equation 97]

$$M_t^{e+1} = tM_{t-1}^e + X_{PC}M_t^e + \frac{1}{2p}M_{t+1}^e.$$

In the special case where ρ_A is a Hermite-Gaussian, the multipole expansion is finite, in the following sense. Let us introduce the Hermite moment

$$M_t^e = \int_{-\infty}^{\infty} x_C^e \Lambda_t dx. \quad (\text{C.2})$$

e is the order of the moment and t is the order of the Hermite-Gaussian.

Proposition C.2.1. [see [150, Section 11.1]] M_t^e is zero for $t > e$.

Proof. This can be seen by partial integration of (C.2)

$$\begin{aligned} M_t^e &= \int_{-\infty}^{\infty} x_C^e \left[\frac{\partial^t \exp(-px_P^2)}{\partial P_x^t} \right] dx \\ &= (-1)^t \int_{-\infty}^{\infty} x_C^e \left[\frac{\partial^t \exp(-px_P^2)}{\partial x^t} \right] dx \\ &= \int_{-\infty}^{\infty} \left(\frac{\partial^t x_C^e}{\partial x^t} \right) \exp(-px_P^2) dx = 0, \quad (t > e), \end{aligned}$$

since the derivative inside the last integral vanishes for $t > e$. □

Therefore, we can assume, without introducing an error, that the multipole expansion is truncated at order $\mathcal{L} = \mathcal{L}_a + \mathcal{L}_b$, where \mathcal{L}_a is the order of the series expansion about \mathbf{A} , and \mathcal{L}_b is the order of the series expansion about \mathbf{B} .

C.2.0.2 Computation of the Cartesian multipole tensor

Elements of the Cartesian multipole tensor are

$$T_{lmn}^{AR} = \frac{\partial^\ell}{\partial^\ell |\mathbf{A} - \mathbf{B}|}, \quad \ell = l + m + n,$$

where AB_x, AB_y, AB_z are the Cartesian components of $\mathbf{A} - \mathbf{B}$. This quantity is computed online. The multipole moments are computed offline. The moment is expensive and should not depend on the orientation of A , otherwise it cannot be computed offline.

$$T_{000j}^{PQ} = (-1)^j (2j - 1)!! |\mathbf{P} - \mathbf{Q}|^{-(2j+1)}, \quad j = 0, \dots, \mathcal{L}.$$

3-term recursion relations

$$\begin{aligned} T_{l,0,0,j} &= PQ_x T_{l-1,0,0,j+1} + (l-1) T_{l-1,0,0,j+1}, \\ T_{l,m,0,j} &= PQ_y T_{l,m-1,0,j+1} + (m-1) T_{l,m-2,0,j+1}, \\ T_{l,m,n,j} &= PQ_z T_{l,m,n-1,j+1} + (n-1) T_{l,m,n-1,j+1}, \end{aligned}$$

Notice that in the case of $P = Q$ this simplifies and can be computed completely offline.

C.2.0.3 Acceleration methods for Hermite-Gaussians

In practical implementations, GEM [84] employs constraints to further reduce the computational workload, namely the use of *spd*-basis up to quadrupoles. We give more details and relevant work that motivate this choice in the following.

A possible choice of auxiliary functions is *spd* sets, which consists of one s, one p, and one d-type Gaussian on each atom, with applications to multipole expansion. Relevant work has recently been reported on this topic [311]. This type of basis is also used in GEM [84]. In this application setting, the primary purpose of the auxiliary basis is to reproduce the multipoles of the orbital products. The charge (monopole) effect on each atom is captured by at least one s-type Gaussian function, which is the smallest auxiliary basis set. In order to increase accuracy and capture high-order multipoles, the author in [311] adds polarized functions, namely p-type and d-type Gaussian functions. The same tendency is also followed in GEM.

Another practical acceleration used in GEM is described as follows. Instead of the usual approach where four-center ERIs are generated and then contracted with the density matrix, efficiency can be obtained by first transforming the density to an HGTF basis using expansion coefficients obtained by McMurchie-Davidson scheme. Asymptotically, this step requires effort and memory proportional to N and avoids the costly transformation from HGTF to CGTF inside the four-order loop over atom centers [71]. Representation in an HGTF basis also allows dramatic simplifications when the LCAO basis set shares exponents between CGTFs of different angular symmetry. Examples of such basis sets include the Pople *sp*-type functions, or the even-tempered functions. The expansion of CGTF distributions leads to HGTF families that are all related to a "mother" (s-type) Gaussian by differentiation. This fact has also been exploited in order to accelerate the multipoles in the code of GEM [84], where *sp*- and *spd*-type functions have been used.

Appendix D

Error bars for equilibrium geometries

This proceeding contains preliminary joint work with Julia Contreras-Garcia (Laboratoire de Chimie Théorique, Sorbonne Université), David Gontier (CEREMADE, Université Paris-Dauphine) and Trinidad Novoa (Laboratoire de Chimie Théorique, Sorbonne Université).

Let us consider a one-electron diatomic molecule. It is assumed that both nuclear charges are equal to $z = 1$ for simplicity. The two nuclei are placed on the x -axis at points $-\frac{1}{2}R\mathbf{e}_1$ and $\frac{1}{2}R\mathbf{e}_1$, where \mathbf{e}_1 denotes the unit vector $(1, 0, 0)$. The Coulomb potential generated by the nuclei is

$$V(\mathbf{r}) := V_R(\mathbf{r}) = \frac{1}{|\mathbf{x} - \frac{R}{2}\mathbf{e}_1|} + \frac{1}{|\mathbf{x} + \frac{R}{2}\mathbf{e}_1|}.$$

Let us introduce the one-electron energy functional without exchange

$$E_0[\rho; R] = \frac{1}{8} \int_{\mathbb{R}^3} \frac{|\nabla \rho|^8}{|\rho|} - \int_{\mathbb{R}^3} V_R \rho + \frac{1}{2} \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{\rho(\mathbf{r})\rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|},$$

where ρ is the electronic density of the one-electron system, equal to $\rho = \phi^2$ for $\phi \in L^2(\mathbb{R}^3)$ a normalized one-electron wave function. The Fock exchange is defined in the one-electron case as

$$E_x^{\text{HF}}[\rho] = -\frac{1}{2} \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{\rho(\mathbf{r})\rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|}.$$

Notice that this term cancels out with the electron-electron repulsion (or Hartree term) in E_0 . The Kohn-Sham Local Density Approximation (LDA) [60] exchange is defined as

$$E_x^{\text{LDA}}[\rho] = -C_{\text{LDA}} \int_{\mathbb{R}^3} \rho^{4/3}.$$

We may ask the question of comparing the Hartree-Fock (HF) and the Density Functional Theory (DFT) equilibrium molecular geometries. A hybrid approach consists of including fractions of the exact exchange, that is the exchange of the HF model, in the Kohn-Sham model. Let us introduce the hybrid energy functional

$$E[\rho; t, R] = E_0[\rho; R] + tE_x^{\text{LDA}}[\rho] + (1-t)E_x^{\text{HF}}[\rho], \quad (\text{D.1})$$

where t is a real parameter varying from 0 to 1, so that for $t = 0$ we recover the HF energy and for $t = 1$ the LDA energy. The problem of optimization of geometry reads as the two-step optimization

$$\forall t, \quad R(t) := \arg \min_R E[\rho_{t,R}; t, R] + \frac{1}{R} \quad (\text{D.2})$$

$$\forall (t, R), \quad \rho_{t,R} := \arg \min_{\rho} E[\rho; t, R]. \quad (\text{D.3})$$

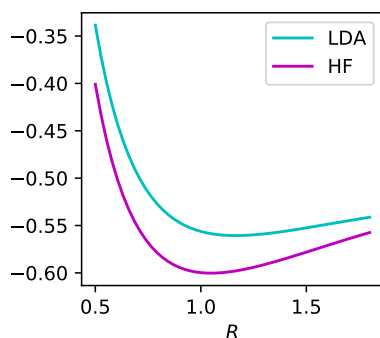


Figure D.1: Unrestricted Kohn-Sham potential energy curve for H_2^+ . The unit of the interatomic distance R is Angstrom.

The quantity $\hat{E}[\rho; t, R] := E[\rho; t, R] + \frac{1}{R}$ gives the clamped nucleus energy. We recover $R_{\text{HF}} := R(0)$ and $R_{\text{LDA}} := R(1)$. Numerical results obtained using PySCF for the unrestricted Kohn-Sham energy functional, show that R_{HF} is about 1.0571 Angstrom and R_{LDA} about 1.1631 Angstrom. This small difference can be seen in Figure D.1.

The problem of theoretically verifying the numerical observation of Figure D.1 can be interesting from a mathematical point of view. Besides, it can help practitioners since it is related to attaching error bars on equilibrium geometries to hybrid energy functionals, for rigorous numerical experiments [250].

Unfortunately, a preliminary study showed that finding the sign of the first derivative $R'(t)$ with respect to t , by directly differentiating the analytic functional formulas, is a non-trivial task. To see this, one may use the Hellmann-Feynman theorem as a starting point, that states that the molecular forces of the system vanish at the equilibrium geometry, i.e.

$$\forall t, \quad \frac{\partial \hat{E}}{\partial R}[\rho; t, R(t)] = 0.$$

To our understanding, this formula cannot be easily solved explicitly with respect to $R(t)$, in order to derive closed-form expressions. In general, it was not possible to express derivatives with respect to t in terms of quantities whose sign is known or can be obtained from existing results. For example, one may start by investigating the variation of the energy (D.1) with respect to t . This functional is equal to the difference in exchange energies. A relative result is the Lieb-Oxford (LO) bound [201] that gives the inequality

$$0 \leq \frac{1}{2} \iint_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{\rho(\mathbf{r})\rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} \leq C_{\text{LO}} \int_{\mathbb{R}^3} \rho^{4/3},$$

with $C_{\text{LO}} := 1.58$. Unfortunately the right-hand side of this estimator is greater than the quality of interest, since

$$C_{\text{LDA}} = \frac{3}{4} \left(\frac{3}{\pi} \right)^{1/3} = 0.738558 \dots$$

To summarize, the problem of mathematically studying the variations of the equilibrium geometry with respect to fractions of the exchange remains open.

Bibliography

- [1] Ali Abedi, Federica Agostini, Yasumitsu Suzuki, and E. K. U. Gross. Dynamical steps that bridge piecewise adiabatic shapes in the exact time-dependent potential energy surface. *Phys. Rev. Lett.*, 110:263001, Jun 2013.
- [2] Ali Abedi, Neepa T. Maitra, and E. K. U. Gross. Exact factorization of the time-dependent electron-nuclear wave function. *Phys. Rev. Lett.*, 105:123002, Sep 2010.
- [3] Ben Adcock and Daan Huybrechs. Frames and numerical approximation. *Siam Review*, 61(3):443–473, 2019.
- [4] Olivier Adjoua, Louis Lagardère, Luc-Henri Jolly, Arnaud Durocher, Thibaut Very, Isabelle Dupays, Zhi Wang, Théo Jaffrelot Inizan, Frédéric Célerse, Pengyu Ren, Jay W. Ponder, and Jean-Philip Piquemal. Tinker-HP: Accelerating Molecular Dynamics Simulations of Large Complex Systems with Advanced Point Dipole Polarizable Force Fields Using GPUs and Multi-GPU Systems. *Journal of Chemical Theory and Computation*, 17(4):2034–2053, 2021. PMID: 33755446.
- [5] B. J. Alder and T. E. Wainwright. Phase Transition for a Hard Sphere System. *The Journal of Chemical Physics*, 27(5):1208–1209, 08 1957.
- [6] Michael P. Allen and Dominic J. Tildesley. *Computer simulation of liquids*. Oxford University Press, 2nd edition, 2017.
- [7] Edward Anderson, Zhaojun Bai, Christian Bischof, L Susan Blackford, James Demmel, Jack Dongarra, Jeremy Du Croz, Anne Greenbaum, Sven Hammarling, Alan McKenney, et al. *LAPACK users' guide*. SIAM, 1999.
- [8] Francesco Aquilante, Linus Boman, Jonas Boström, Henrik Koch, Roland Lindh, Alfredo Sánchez de Merás, and Thomas Bondo Pedersen. Cholesky decomposition techniques in electronic structure theory. In Robert Zalesny, Manthos G. Papadopoulos, Paul G. Mezey, and Jerzy Leszczynski, editors, *Linear-Scaling Techniques in Computational Chemistry and Physics: Methods and Applications*, pages 301–343. Springer Netherlands, Dordrecht, 2011.
- [9] Francesco Aquilante, Laura Gagliardi, Thomas Bondo Pedersen, and Roland Lindh. Atomic cholesky decompositions: A route to unbiased auxiliary basis sets for density fitting approximation with tunable accuracy and efficiency. *The Journal of chemical physics*, 130(15), 2009.
- [10] Francesco Aquilante, Roland Lindh, and Thomas Bondo Pedersen. Unbiased auxiliary basis sets for accurate two-electron integral approximations. *The Journal of chemical physics*, 127(11):114107, 2007.
- [11] P. Arrighini. *Intermolecular Forces and Their Evaluation by Perturbation Theory*, volume 25 of *Lecture Notes in Chemistry*. Springer Berlin, Heidelberg, 1981.

- [12] Andrey Asadchev, Veerendra Allada, Jacob Felder, Brett M Bode, Mark S Gordon, and Theresa L Windus. Uncontracted rys quadrature implementation of up to g functions on graphical processing units. *Journal of Chemical Theory and Computation*, 6(3):696–704, 2010.
- [13] Andrey Asadchev and Edward F Valeev. High-performance evaluation of high angular momentum 4-center gaussian integrals on modern accelerated processors. *arXiv preprint arXiv:2307.03452*, 2023.
- [14] Stephen H Ashworth. *Molecular Quantum Mechanics*, 5th edn, by Peter Atkins and Ronald Friedman: Scope: textbook. Level: undergraduate, 2012.
- [15] Sadybekov A.V. and Katritch V. Computational approaches streamlining drug discovery. *Nature*, 616:673–685, 2023.
- [16] FW Averill and Donald E Ellis. An efficient numerical multicenter basis set for molecular orbital calculations: application to fecl4. *The journal of chemical physics*, 59(12):6412–6418, 1973.
- [17] Félix Aviat. *Development of High Performance Molecular Dynamics : Fast Evaluation of Polarization Forces*. Theses, Sorbonne Université, September 2019.
- [18] Markus Bachmayr, Huajie Chen, and Reinhold Schneider. Error estimates for Hermite and even-tempered Gaussian approximations in quantum chemistry. *Numerische Mathematik*, 128(1):137–165, 2014.
- [19] R. F. W. Bader and P. M. Beddall. Virial Field Relationship for Molecular Charge Distributions and the Spatial Partitioning of Molecular Properties. *The Journal of Chemical Physics*, 56(7):3320–3329, 09 2003.
- [20] Siwar Badreddine, Igor Chollet, and Laura Grigori. Factorized structure of the long-range two-electron integrals tensor and its application in quantum chemistry. working paper or preprint, October 2022.
- [21] E.J. Baerends, D.E. Ellis, and P. Ros. Self-consistent molecular hartree-fock-slater calculations i. the computational procedure. *Chemical Physics*, 2(1):41–51, 1973.
- [22] Paul S Bagus and Francesc Illas. Decomposition of the chemisorption bond by constrained variations: Order of the variations and construction of the variational spaces. *The Journal of chemical physics*, 96(12):8962–8970, 1992.
- [23] PS Bagus and K Hermann. Origin of lone pair binding energy shifts in photoemission from adsorbed molecules: Csov analysis for the cu5co cluster. *Applications of Surface Science*, 22:444–451, 1985.
- [24] H Bahouri, R Danchin, and JY Chemin. Fourier analysis and nonlinear partial differential equations, a series of comprehensive studies in mathematics. *Grundlehren der Mathematischen Wissenschaften*, 343, 2011.
- [25] Maxime Barrault, Yvon Maday, Ngoc Cuong Nguyen, and Anthony T Patera. An ‘empirical interpolation’ method: application to efficient reduced-basis discretization of partial differential equations. *Comptes Rendus Mathématique*, 339(9):667–672, 2004.
- [26] G. Battimelli and G. Ciccotti. Berni Alder and the pioneering times of molecular simulation. *The European Physical Journal H*, 43:303–335, 2018.
- [27] Mario Bebendorf. Approximation of boundary element matrices. *Numerische Mathematik*, 86:565–589, 2000.

- [28] Mario Bebendorf. *Hierarchical matrices*. Springer, 2008.
- [29] Mario Bebendorf, Christian Kuske, and Raoul Venn. Wideband nested cross approximation for Helmholtz problems. *Numerische Mathematik*, 130:1–34, 2015.
- [30] A. D. Becke. A multicenter numerical integration scheme for polyatomic molecules. *The Journal of chemical physics*, 88(4):2547–2553, 1988.
- [31] Nelson HF Beebe and Jan Linderberg. Simplifications in the generation and transformation of two-electron integrals in molecular calculations. *International Journal of Quantum Chemistry*, 12(4):683–705, 1977.
- [32] Jacopo Bellazzini, Marco Ghimenti, and Tohru Ozawa. Sharp lower bounds for coulomb energy. *Mathematical Research Letters*, 23(3):621–632, 2016.
- [33] Robert Benda. *Multi-scale modeling of water quality nanosensors based on carbon nanotubes and conjugated polymers*. Theses, Institut Polytechnique de Paris, December 2021.
- [34] Robert Benda, Eric Cancès, Virginie Ehrlacher, and Benjamin Stamm. Multi-center decomposition of molecular densities: A mathematical perspective. *The Journal of Chemical Physics*, 156(16):164107, 04 2022.
- [35] Michele Benzi, Dario Bini, Daniel Kressner, Hans Munthe-Kaas, Charles Van Loan, Jonas Ballani, and Daniel Kressner. Matrices with hierarchical low-rank structures. *Exploiting Hidden Structure in Matrix Computations: Algorithms and Applications: Cetraro, Italy 2015*, pages 161–209, 2016.
- [36] H. J. C. Berendsen, J. R. Grigera, and T. P. Straatsma. The missing term in effective pair potentials. *The Journal of Physical Chemistry*, 91(24):6269–6271, 1987.
- [37] J. D. Bernal and R. H. Fowler. A Theory of Water and Ionic Solution, with Particular Reference to Hydrogen and Hydroxyl Ions. *The Journal of Chemical Physics*, 1(8):515–548, 11 1933.
- [38] David E Bernholdt and Robert J Harrison. Fitting basis sets for the ri-mp2 approximate second-order many-body perturbation theory method. *The Journal of chemical physics*, 109(5):1593–1600, 1998.
- [39] Colleen Bertoni, Lyudmila V. Slipchenko, Alston J. Misquitta, and Mark S. Gordon. Multipole moments in the effective fragment potential method. *The Journal of Physical Chemistry A*, 121(9):2056–2067, 2017. PMID: 28211686.
- [40] Colleen Bertoni, Lyudmila V. Slipchenko, Alston J. Misquitta, and Mark S. Gordon. Multipole moments in the effective fragment potential method. *The Journal of Physical Chemistry A*, 121(9):2056–2067, 2017. PMID: 28211686.
- [41] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. Julia: A fresh approach to numerical computing. *SIAM review*, 59(1):65–98, 2017.
- [42] Anand Bharti and Tamal Banerjee. Solubility prediction of bio-oil derived chemicals in aqueous media by Localized Molecular Orbital-Energy Decomposition Analysis (LMO-EDA) and COSMO-RS predictions. *Computational and Theoretical Chemistry*, 1067:48–59, 2015.
- [43] Simon Blaschke and Stella Stopkowicz. Cholesky decomposition of complex two-electron integrals over GIAOs: Efficient MP2 computations for large molecules in strong magnetic fields. *The Journal of Chemical Physics*, 156(4), 2022.

- [44] OpenMP Architecture Review Board. OpenMP Application Programming Interface. Available at <https://www.openmp.org/>.
- [45] John A Board Jr, Jeffrey W Causey, James F Leathrum Jr, Andreas Windemuth, and Klaus Schulten. Accelerated molecular dynamics simulation with the parallel fast multipole algorithm. *Chemical Physics Letters*, 198(1-2):89–94, 1992.
- [46] N. Bohr. I. On the constitution of atoms and molecules. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 26(151):1–25, 1913.
- [47] N. Bohr. LXXIII. On the constitution of atoms and molecules. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 26(155):857–875, 1913.
- [48] N. Bohr. XXXVII. On the constitution of atoms and molecules. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 26(153):476–502, 1913.
- [49] Max Born and W Heisenberg. Zur quantentheorie der molekeln. *Original Scientific Papers Wissenschaftliche Originalarbeiten*, pages 216–246, 1985.
- [50] Christos Boutsidis and Malik Magdon-Ismail. Faster SVD-truncated regularized least-squares. In *2014 IEEE International Symposium on Information Theory*, pages 1321–1325. IEEE, 2014.
- [51] David R Bowler and Tsuyoshi Miyazaki. $O(N)$ methods in electronic structure calculations. *Reports on Progress in Physics*, 75(3):036503, 2012.
- [52] Stephen P Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [53] S. F. Boys and Alfred Charles Egerton. Electronic wave functions - i. a general method of calculation for the stationary states of any molecular system. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 200(1063):542–554, 1950.
- [54] S Francis Boys. Electronic wave functions-i. a general method of calculation for the stationary states of any molecular system. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 200(1063):542–554, 1950.
- [55] Moritz Braun. Finite element Hartree–Fock calculations in three dimensions for atoms and small molecules. *Journal of Computational and Applied Mathematics*, 270:100–108, 2014.
- [56] Curt M Breneman and Kenneth B Wiberg. Determining atom-centered monopoles from molecular electrostatic potentials. The need for high sampling density in formamide conformational analysis. *Journal of Computational Chemistry*, 11(3):361–373, 1990.
- [57] Haim Brezis and Haim Brézis. *Functional analysis, Sobolev spaces and partial differential equations*, volume 2. Springer, 2011.
- [58] A. D. Buckingham. Basic theory of intermolecular forces: Applications to small molecules. In B. Pullman, editor, *Intermolecular Interactions: From Diatomics to Biopolymers*, chapter 1. J. Wiley and Sons, 1978.
- [59] Carlos Campaña, Bastien Mussard, and Tom K Woo. Electrostatic potential derived atomic charges for periodic systems using a modified error functional. *Journal of Chemical Theory and Computation*, 5(10):2866–2878, 2009.
- [60] Eric Cancès, Claude Le Bris, and Yvon Maday. *Méthodes mathématiques en chimie quantique: une introduction*, volume 53 of *Mathématiques et Applications*. Springer, 2006.

- [61] Eric Cancès, Mireille Defranceschi, Werner Kutzelnigg, Claude Le Bris, and Yvon Maday. Computational quantum chemistry: a primer. *Handbook of numerical analysis*, 10:3–270, 2003.
- [62] Eric Cancès, Geneviève Dusson, Gaspard Kemlin, and Laurent Vidal. On basis set optimisation in quantum chemistry. *ESAIM: Proceedings and Surveys*, 73:107–129, 2023.
- [63] Eric Cancès, Geneviève Dusson, Yvon Maday, Benjamin Stamm, and Martin Vohralík. Guaranteed and robust a posteriori bounds for Laplace eigenvalues and eigenvectors: a unified framework. *Numerische Mathematik*, 140(4):1033–1079, 2018.
- [64] Eric Cancès, Geneviève Dusson, Yvon Maday, Benjamin Stamm, and Martin Vohralík. Guaranteed and robust a posteriori bounds for Laplace eigenvalues and eigenvectors: conforming approximations. *SIAM Journal on Numerical Analysis*, 55(5):2228–2254, 2017.
- [65] Salvatore Cardamone, Timothy J. Hughes, and Paul L. A. Popelier. Multipolar electrostatics. *Phys. Chem. Chem. Phys.*, 16:10367–10387, 2014.
- [66] Judith Cederberg. *A course in modern geometries*. Springer Science & Business Media, 2004.
- [67] Michele Ceriotti, Wei Fang, Peter G. Kusalik, Ross H. McKenzie, Angelos Michaelides, Miguel A. Morales, and Thomas E. Markland. Nuclear quantum effects in water and aqueous systems: Experiment, theory, and current challenges. *Chemical Reviews*, 116(13):7529–7550, 2016. PMID: 27049513.
- [68] Matt Challacombe and Nicolas Bock. Fast multiplication of matrices with decay. *arXiv preprint arXiv:1011.3534*, 2010.
- [69] Matt Challacombe and Eric Schwegler. Linear scaling computation of the Fock matrix. *The Journal of chemical physics*, 106(13):5526–5536, 1997.
- [70] Matt Challacombe, Eric Schwegler, and Jan Almlöf. Fast assembly of the Coulomb matrix: A quantum chemical tree code. *The Journal of chemical physics*, 104(12):4685–4698, 1996.
- [71] Matt Challacombe, Eric Schwegler, and Jan Almlöf. Modern developments in hartree-fock theory: Fast methods for computing the coulomb matrix. In *Computational Chemistry: Reviews of Current Trends*, pages 53–107. World Scientific Publishing Company, 1996.
- [72] De-Li Chen, Abraham C Stern, Brian Space, and J Karl Johnson. Atomic charges derived from electrostatic potentials for molecular and periodic systems. *The Journal of Physical Chemistry A*, 114(37):10225–10233, 2010.
- [73] Huajie Chen and Reinhold Schneider. Numerical analysis of augmented plane wave methods for full-potential electronic structure calculations. *ESAIM: Mathematical Modelling and Numerical Analysis*, 49(3):755–785, 2015.
- [74] Jingrun Chen and Jianfeng Lu. Analysis of the divide-and-conquer method for electronic structure calculations. *Mathematics of Computation*, 85(302):2919–2938, 2016.
- [75] Shi-Jie Chen and Donald H Burke-Aguero. *Computational methods for understanding riboswitches*. Academic Press, 2015.
- [76] Wei Chen and Mark S Gordon. Energy decomposition analyses for many-body interaction and applications to water complexes. *The Journal of Physical Chemistry*, 100(34):14316–14328, 1996.
- [77] Yifan Chen, Ethan N Epperly, Joel A Tropp, and Robert J Webber. Randomly pivoted Cholesky: Practical approximation of a kernel matrix with few entry evaluations, 2022.

- [78] Lisa Emily Chirlian and Michelle Miller Franci. Atomic charges derived from electrostatic potentials: A detailed study. *Journal of Computational Chemistry*, 8(6):894–905, 1987.
- [79] Igor Chollet, Louis Lagardère, and Jean-Philip Piquemal. ANKH: A Generalized $O(N)$ Interpolated Ewald Strategy for Molecular Dynamics Simulations. *Journal of Chemical Theory and Computation*, 19(10):2887–2905, 2023. PMID: 37134146.
- [80] Giovanni Ciccotti, Christoph Dellago, Mauro Ferrario, Eduardo Hernández, and Mark E. Tuckerman. Molecular simulations: past, present, and future (a topical issue in epjb). *European Physical Journal B*, 95(1), Jan 2022.
- [81] G Andrés Cisneros, Jean-Philip Piquemal, and Thomas A Darden. Quantum mechanics/molecular mechanics electrostatic embedding with continuous and discrete functions. *The journal of physical chemistry B*, 110(28):13682–13684, 2006.
- [82] G. Andrés Cisneros. Application of gaussian electrostatic model (gem) distributed multipoles in the amoeba force field. *Journal of Chemical Theory and Computation*, 8(12):5072–5080, 2012. PMID: 26593198.
- [83] G. Andrés Cisneros, Jean-Philip Piquemal, and Thomas A. Darden. Intermolecular electrostatic energies using density fitting. *The Journal of Chemical Physics*, 123(4):044109, 08 2005.
- [84] G. Andrés Cisneros, Jean-Philip Piquemal, and Thomas A. Darden. Generalization of the Gaussian electrostatic model: Extension to arbitrary angular momentum, distributed multipoles, and speedup with reciprocal space methods. *The Journal of Chemical Physics*, 125(18):184101, 11 2006.
- [85] G. Andrés Cisneros, Jean-Philip Piquemal, and Thomas A. Darden. Quantum mechanics/molecular mechanics electrostatic embedding with continuous and discrete functions. *The Journal of Physical Chemistry B*, 110(28):13682–13684, 2006. PMID: 16836309.
- [86] G.A. Cisneros, D. Elking, J.P. Piquemal, and T.A. Darden. Numerical fitting of molecular properties to Hermite Gaussians. *The Journal of Physical Chemistry A*, 111(47):12049–12056, 2007.
- [87] Fritz Coester and Hermann Kümmel. Short-range correlations in nuclear wave functions. *Nuclear Physics*, 17:477–485, 1960.
- [88] Coppé, Vincent and Huybrechs, Daan and Matthysen, Roel and Webb, Marcus. The AZ algorithm for least squares systems with a known incomplete generalized inverse. *SIAM Journal on Matrix Analysis and Applications*, 41(3):1237–1259, 2020.
- [89] Alice Cortinovis, Daniel Kressner, and Stefano Massei. On maximum volume submatrices and cross approximation for symmetric semidefinite and diagonally dominant matrices. *Linear Algebra and its Applications*, 593:251–268, 2020.
- [90] C.A. Coulson. The hydrogen bond - a review of the present position. *Research*, 10(4):149, 1957.
- [91] Harold Scott Macdonald Coxeter. *Introduction to geometry*. John Wiley & Sons, Inc., 1969.
- [92] John Dalton. *A new system of chemical philosophy*, volume 1. Cambridge University Press, 2010.
- [93] Tom Darden, Darrin York, and Lee Pedersen. Particle mesh Ewald: An $N \cdot \log(N)$ method for Ewald sums in large systems. *The Journal of Chemical Physics*, 98(12):10089–10092, 06 1993.

- [94] Graeme M Day and Sarah L Price. A nonempirical anisotropic atom-atom model potential for chlorobenzene crystals. *Journal of the American Chemical Society*, 125(52):16434–16443, 2003.
- [95] S.B. de Beer, N.P. Vermeulen, and C. Oostenbrink. The role of water molecules in computational drug design. *Current Topics in Medicinal Chemistry*, 10(1):55–66, 2010.
- [96] Ondřej Demel, Michael J Lecours, Richard Habrovský, and Marcel Nooijen. Toward Laplace MP2 method using range separated Coulomb potential and orbital selective virtuals. *The Journal of chemical physics*, 155(15), 2021.
- [97] Frank Deutsch and F Deutsch. *Best approximation in inner product spaces*. Springer, 2001.
- [98] Ken A. Dill and Justin L. MacCallum. The protein-folding problem, 50 years on. *Science*, 338(6110):1042–1046, 2012.
- [99] Hong-Qiang Ding, Naoki Karasawa, and William A Goddard III. The reduced cell multipole method for Coulomb interactions in periodic systems with million-atom unit cells. *Chemical Physics Letters*, 196(1-2):6–10, 1992.
- [100] RHWJ Ditchfield, Warren J Hehre, and John A Pople. Self-consistent molecular-orbital methods. IX. An extended Gaussian-type basis for molecular-orbital studies of organic molecules. *The Journal of Chemical Physics*, 54(2):724–728, 1971.
- [101] Richard W. Dixon and Peter A. Kollman. Advancing beyond the atom-centered model in additive and nonadditive molecular mechanics. *Journal of Computational Chemistry*, 18(13):1632–1646, 1997.
- [102] Matthieu Dolbeault and Albert Cohen. Optimal pointwise sampling for l2 approximation. *Journal of Complexity*, 68:101602, 2022.
- [103] Carmen Domene Nunez, Patrick W. Fowler, M. Wilson, Paul A. Madden, and Richard J. Wheatley. Overlap-model and ab initio cluster calculations of ion properties in distorted environments. *Chemical Physics Letters*, 333(5):403–412, January 2001.
- [104] Robert E Duke and G Andres Cisneros. Ewald-based methods for Gaussian integral evaluation: Application to a new parameterization of GEM. *Journal of molecular modeling*, 25(10):307, 2019.
- [105] Robert E. Duke, Oleg N. Starovoytov, Jean-Philip Piquemal, and G. Andrés Cisneros. GEM*: A Molecular Electronic Density-Based Force Field for Molecular Dynamics Simulations. *Journal of Chemical Theory and Computation*, 10(4):1361–1365, 2014. PMID: 26580355.
- [106] B. I. Dunlap, J. W. D. Connolly, and J. R. Sabin. On some approximations in applications of α theory. *The Journal of Chemical Physics*, 71(8):3396–3402, 1979.
- [107] BI Dunlap. Robust and variational fitting: Removing the four-center integrals from center stage in quantum chemistry. *Journal of Molecular Structure: THEOCHEM*, 529(1-3):37–40, 2000.
- [108] Brett I. Dunlap. Robust and variational fitting. *Phys. Chem. Chem. Phys.*, 2:2113–2116, 2000.
- [109] Thom H Dunning Jr. Gaussian basis sets for use in correlated molecular calculations. I. The atoms boron through neon and hydrogen. *The Journal of chemical physics*, 90(2):1007–1023, 1989.

- [110] Michel Dupuis and Antonio Marquez. The Rys quadrature revisited: A novel formulation for the efficient computation of electron repulsion integrals over Gaussian functions. *The Journal of Chemical Physics*, 114(5):2067–2078, 2001.
- [111] Jacob D. Durrant and J. Andrew McCammon. Molecular dynamics simulations and drug discovery. *BMC Biology*, 9(1), Oct 2011.
- [112] Geneviève Dusson and Yvon Maday. A posteriori analysis of a nonlinear Gross–Pitaevskii-type eigenvalue problem. *IMA Journal of Numerical Analysis*, 37(1):94–137, 2017.
- [113] Alan Robert Edmonds. *Angular momentum in quantum mechanics*. Princeton university press, 1996.
- [114] Henry Eyring, John Walter, and Kimball George E. *Quantum Chemistry*. J. Wiley and Sons, 1st edition, 1944.
- [115] Martin Feyereisen, George Fitzgerald, and Andrew Komornicki. Use of approximate integrals in ab initio theory. An application in MP2 energy calculations. *Chemical physics letters*, 208(5-6):359–363, 1993.
- [116] Vladimir Fock. Näherungsmethode zur Lösung des quantenmechanischen Mehrkörperproblems. *Zeitschrift für Physik*, 61:126–148, 1930.
- [117] C Fonseca Guerra, JG Snijders, G t Te Velde, and E Jan Baerends. Towards an order-N DFT method. *Theoretical Chemistry Accounts*, 99:391–403, 1998.
- [118] Michelle Francl and Lisa Chirlian. The pluses and minuses of mapping atomic charges to electrostatic potentials. *Reviews in Computational Chemistry*, 14:1 – 31, 01 2007.
- [119] Mark A Freitag, Mark S Gordon, Jan H Jensen, and Walter J Stevens. Evaluation of charge penetration between distributed multipolar expansions. *The Journal of Chemical Physics*, 112(17):7300–7306, 2000.
- [120] Roger T Gallant and Alain St-Amant. Linear scaling for the charge density fitting procedure of the linear combination of gaussian-type orbitals density functional method. *Chemical physics letters*, 256(6):569–574, 1996.
- [121] Giulia Galli. Linear scaling methods for electronic structure calculations and quantum molecular dynamics simulations. *Current Opinion in Solid State and Materials Science*, 1(6):864–874, 1996.
- [122] Jean Gallier and Jocelyn Quaintance. Spherical harmonics and linear representations of lie groups. In *Differential Geometry and Lie Groups: A Second Course*, pages 265–360. Springer, 2020.
- [123] Mariano Gasca and Thomas Sauer. *Polynomial interpolation in several variables*, volume 12 of *Advances in Computational Mathematics*. Springer, 2000.
- [124] Johann Gasteiger and Heinz Saller. Calculation of the charge distribution in conjugated systems by a quantification of the resonance concept. *Angewandte Chemie International Edition in English*, 24(8):687–689, 1985.
- [125] Andrew Gibbs, David P Hewett, Daan Huybrechs, and Emile Parolin. Fast hybrid numerical-asymptotic boundary element methods for high frequency screen and aperture problems based on least-squares collocation. *SN Partial Differential Equations and Applications*, 1:1–26, 2020.
- [126] Josiah Willard Gibbs. *Elementary principles in statistical mechanics: developed with especial reference to the rational foundations of thermodynamics*. C. Scribner’s sons, 1902.

- [127] Peter MW Gill, Andrew TB Gilbert, SW Taylor, G Friesecke, and Martin Head-Gordon. Decay behavior of least-squares coefficients in auxiliary basis expansions. *The Journal of chemical physics*, 123(6), 2005.
- [128] Peter MW Gill, Benny G Johnson, and John A Pople. A simple yet powerful upper bound for Coulomb integrals. *Chemical physics letters*, 217(1-2):65–68, 1994.
- [129] Hatice Gokcan, Eric Kratz, Thomas A Darden, Jean-Philip Piquemal, and G Andrés Cisneros. Qm/mm simulations with the gaussian electrostatic model: A density-based polarizable potential. *The journal of physical chemistry letters*, 9(11):3062–3067, 2018.
- [130] Gene H Golub and Charles F Van Loan. *Matrix computations*. The Johns Hopkins University Press, 3rd edition, 2013.
- [131] González, M.A. Force fields and molecular dynamics simulations. *JDN*, 12:169–200, 2011.
- [132] Mark S Gordon and Theresa L Windus. Modern architectures and their impact on electronic structure theory. *Chemical Reviews*, 120(17):9015–9020, 2020.
- [133] Roy G Gordon and Yung Sik Kim. Theory for the forces between closed-shell atoms and molecules. *The Journal of Chemical Physics*, 56(6):3122–3133, 1972.
- [134] Leslie Greengard and Vladimir Rokhlin. A fast algorithm for particle simulations. *Journal of computational physics*, 73(2):325–348, 1987.
- [135] Leslie Greengard and Vladimir Rokhlin. A new version of the fast multipole method for the laplace equation in three dimensions. *Acta numerica*, 6:229–269, 1997.
- [136] Nohad Gresh, G Andrés Cisneros, Thomas A Darden, and Jean-Philip Piquemal. Anisotropic, polarizable molecular mechanics studies of inter- and intramolecular interactions and ligand- macromolecule complexes. A bottom-up strategy. *Journal of chemical theory and computation*, 3(6):1960–1986, 2007.
- [137] Alan Grossfield, Pengyu Ren, and Jay W. Ponder. Ion solvation thermodynamics from simulation with a polarizable force field. *Journal of the American Chemical Society*, 125(50):15671–15682, 2003. PMID: 14664617.
- [138] Cisneros Research Group. `GEM_fit`. Available at https://github.com/CisnerosResearch/GEM_fit.
- [139] Vaishnavi Gujjula and Sivaram Ambikasaran. A new nested cross approximation. *arXiv preprint arXiv:2203.14832*, 2022.
- [140] Andreas W. Götz, Mark J. Williamson, Dong Xu, Duncan Poole, Scott Le Grand, and Ross C. Walker. Routine microsecond molecular dynamics simulations with amber on gpus. 1. generalized born. *Journal of Chemical Theory and Computation*, 8(5):1542–1555, 2012. PMID: 22582031.
- [141] Bernard Haasdonk. Reduced Basis Methods for Parametrized PDEs-A Tutorial Introduction for Stationary and Instationary Problems. In *Model Reduction and Approximation*, Computational Science & Engineering, chapter 2, pages 65–136. Society for Industrial and Applied Mathematics, 2017.
- [142] Thomas A Halgren and Wolfgang Damm. Polarizable force fields. *Current Opinion in Structural Biology*, 11(2):236–242, 2001.
- [143] Sven Hammarling, Nicholas J Higham, and Craig Lucas. LAPACK-style codes for pivoted Cholesky and QR updating. In *International Workshop on Applied Parallel Computing*, pages 137–146. Springer, 2006.

- [144] Per Christian Hansen. Truncated singular value decomposition solutions to discrete ill-posed problems with ill-determined numerical rank. *SIAM Journal on Scientific and Statistical Computing*, 11(3):503–518, 1990.
- [145] Michał Hapka, Michał Przybytek, and Katarzyna Pernal. Symmetry-adapted perturbation theory based on multiconfigurational wave function description of monomers. *Journal of Chemical Theory and Computation*, 17(9):5538–5555, 2021.
- [146] Helmut Harbrecht, Michael Peters, and Reinhold Schneider. On the low-rank approximation by the pivoted cholesky decomposition. *Applied Numerical Mathematics*, 62(4):428–440, 2012. Third Chilean Workshop on Numerical Analysis of Partial Differential Equations (WONAPDE 2010).
- [147] Douglas Rayne Hartree. The wave mechanics of an atom with a non-coulomb central field. Part II. Some results and discussion. *Mathematical Proceedings of the Cambridge Philosophical Society*, 24(1):111–132, 1928.
- [148] Marco Häser and Reinhart Ahlrichs. Improvements on the direct SCF method. *Journal of Computational Chemistry*, 10(1):104–111, 1989.
- [149] Trygve Helgaker, Poul Jørgensen, and Jeppe Olsen. *Molecular Electronic-Structure Theory*. J. Wiley and Sons, 8 2000.
- [150] Trygve Helgaker and Peter R. Taylor. Gaussian basis sets and molecular integrals. In D R Yarkony, editor, *Modern Electronic Structure Theory*, pages 725–856. J. Wiley and Sons, 1995.
- [151] Arnim Hellweg, Christof Hättig, Sebastian Höfener, and Wim Klopper. Optimized accurate auxiliary basis sets for ri-mp2 and ri-cc2 calculations for the atoms rb to rn. *Theoretical Chemistry Accounts*, 117(4):587–597, 2007.
- [152] Michael F Herbst, Antoine Levitt, and Eric Cancès. A posteriori error estimation for the non-self-consistent Kohn–Sham equations. *Faraday Discussions*, 224:227–246, 2020.
- [153] Kersti Hermansson. Many-body effects in tetrahedral water clusters. *The Journal of Chemical Physics*, 89(4):2149–2159, 08 1988.
- [154] Astrid Herremans and Daan Huybrechs. Efficient Function Approximation in Enriched Approximation Spaces, 2023.
- [155] Jan S Hesthaven, Gianluigi Rozza, Benjamin Stamm, et al. *Certified reduced basis methods for parametrized partial differential equations*, volume 590. Springer, 2016.
- [156] Nicholas J Higham. *Accuracy and stability of numerical algorithms*. SIAM, 2002.
- [157] F. L. Hirshfeld. Bonded-atom fragments for describing molecular charge densities. *Theoretica chimica acta*, 44:129–138, 1977.
- [158] David S Hollman, Henry F Schaefer, and Edward F Valeev. Fast construction of the exchange operator in an atom-centred basis with concentric atomic density fitting. *Molecular physics*, 115(17-18):2065–2076, 2017.
- [159] Paul R Horn and Martin Head-Gordon. Alternative definitions of the frozen energy in energy decomposition analysis of density functional theory calculations. *The Journal of chemical physics*, 144(8), 2016.
- [160] Hao Hu, Zhenyu Lu, and Weitao Yang. Fitting molecular electrostatic potentials from quantum mechanical calculations. *Journal of chemical theory and computation*, 3(3):1004–1013, 2007.

- [161] Hua Huang and Edmond Chow. Accelerating quantum chemistry with vectorized and batched integrals. In *SC18: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 529–542. IEEE, 2018.
- [162] Jing Huang and Alexander D. MacKerell. CHARMM36 all-atom additive protein force field: Validation based on comparison to NMR data. *Journal of Computational Chemistry*, 34:2135 – 2145, 2013.
- [163] Xinchuan Huang, Bastiaan J. Braams, and Joel M. Bowman. Ab Initio Potential Energy and Dipole Moment Surfaces of (H₂O)₂. *The Journal of Physical Chemistry A*, 110(2):445–451, 2006. PMID: 16405316.
- [164] William Humphrey, Andrew Dalke, and Klaus Schulten. VMD: visual molecular dynamics. *Journal of molecular graphics*, 14(1):33–38, 1996.
- [165] H Inouye and S Kita. Experimental determination of the repulsive potentials between K⁺ ions and rare-gas atoms. *The Journal of Chemical Physics*, 56(10):4877–4882, 1972.
- [166] Stefan Irnich and Guy Desaulniers. Shortest path problems with resource constraints. In *Column generation*, pages 33–65. Springer, 2005.
- [167] Théo Jaffrelot Inizan, Frédéric Célerse, Olivier Adjoua, Dina El Ahdab, Luc-Henri Jolly, Chengwen Liu, Pengyu Ren, Matthieu Montes, Nathalie Lagarde, Louis Lagardère, Pierre Monmarché, and Jean-Philip Piquemal. High-resolution mining of the sars-cov-2 main protease conformational space: supercomputer-driven unsupervised adaptive sampling. *Chem. Sci.*, 12:4889–4907, 2021.
- [168] JA Jafri and JL Whitten. Electron repulsion integral approximations and error bounds: Molecular applications. *The Journal of Chemical Physics*, 61(5):2116–2121, 1974.
- [169] JA Jafri and JL Whitten. Iterative introduction of integral approximations with error bounds in SCF calculations. *International Journal of Quantum Chemistry*, 10(3):383–391, 1976.
- [170] Frank Jensen. *Introduction to Computational Chemistry*. J. Wiley and Sons, 3rd edition, 2017.
- [171] B. Jeziorski, R. Moszynski, A. Ratkiewicz, S. Rybak, K. Szalewicz, and H.L. Williams. SAPT: A Program for Many-Body Symmetry-Adapted Perturbation Theory Calculations of Intermolecular Interaction Energies. *Methods and Techniques in Computational Chemistry: METECC-94*, B:79, 1993.
- [172] Bogumil Jeziorski, Robert Moszynski, and Krzysztof Szalewicz. Perturbation theory approach to intermolecular potential energy surfaces of van der waals complexes. *Chemical Reviews*, 94(7):1887–1930, 1994.
- [173] Benny G Johnson, Peter MW Gill, and John A Pople. Exact and approximate solutions to the one-center McMurchie–Davidson tree-search problem. *International journal of quantum chemistry*, 40(6):809–827, 1991.
- [174] William L. Jorgensen, Jayaraman Chandrasekhar, Jeffrey D. Madura, Roger Impey, and Michael L. Klein. Comparison of simple potential functions for simulating liquid water. *Journal of Chemical Physics*, 79:926–935, 1983.
- [175] William L. Jorgensen, David S. Maxwell, and Julian Tirado-Rives. Development and Testing of the OPLS All-Atom Force Field on Conformational Energetics and Properties of Organic Liquids. *Journal of the American Chemical Society*, 118(45):11225–11236, 1996.

- [176] Jaewoon Jung, Chigusa Kobayashi, Kento Kasahara, Cheng Tan, Akiyoshi Kuroda, Kazuo Minami, Shigeru Ishiduki, Tatsuo Nishiki, Hikaru Inoue, Yutaka Ishikawa, Michael Feig, and Yuji Sugita. New parallel computing algorithm of molecular dynamics for extremely huge scale biological systems. *Journal of Computational Chemistry*, 42(4):231–241, 2021.
- [177] Yousung Jung, Alex Sodt, Peter MW Gill, and Martin Head-Gordon. Auxiliary basis expansions for large-scale electronic structure calculations. *Proceedings of the National Academy of Sciences*, 102(19):6692–6697, 2005.
- [178] Mihály Kállay. A systematic way for the cost reduction of density fitting methods. *The Journal of chemical physics*, 141(24), 2014.
- [179] Tosio Kato. *Perturbation theory for linear operators*, volume 132. Springer Science & Business Media, 2013.
- [180] Jason Kaye, Lin Lin, and Chao Yang. A posteriori error estimator for adaptive local basis functions to solve Kohn-Sham density functional theory. *arXiv preprint arXiv:1401.0920*, 2014.
- [181] Gaspard Kemlin. 1d basis optimization. Available at https://github.com/gkemlin/1D_basis_optimization.
- [182] Venera Khoromskaia and Boris N Khoromskij. Tensor numerical methods in quantum chemistry: from hartree-fock to excitation energies. *Physical Chemistry Chemical Physics*, 17(47):31491–31509, 2015.
- [183] Venera Khoromskaia, Boris N Khoromskij, and Reinhold Schneider. Tensor-structured factorized calculation of two-electron integrals in a general basis. *SIAM Journal on Scientific Computing*, 35(2):A987–A1010, 2013.
- [184] Boris N Khoromskij and Venera Khoromskaia. Multigrid accelerated tensor approximation of function related multidimensional arrays. *SIAM Journal on Scientific Computing*, 31(4):3002–3026, 2009.
- [185] Boris N Khoromskij, Venera Khoromskaia, Sambasiva Rao Chinnamsetty, and H-J Flad. Tensor decomposition in electronic structure calculations on 3D Cartesian grids. *Journal of computational physics*, 228(16):5749–5762, 2009.
- [186] Yung Sik Kim, Seong Keun Kim, and Won Don Lee. Dependence of the closed-shell repulsive interaction on the overlap of the electron densities. *Chemical Physics Letters*, 80(3):574–575, 1981.
- [187] Kazuo Kitaura and Keiji Morokuma. A new energy decomposition scheme for molecular interactions within the Hartree-Fock approximation. *International Journal of Quantum Chemistry*, 10(2):325–340, 1976.
- [188] A. Klamt and G. Schüürmann. Cosmo: a new approach to dielectric screening in solvents with explicit expressions for the screening energy and its gradient. *J. Chem. Soc., Perkin Trans. 2*, pages 799–805, 1993.
- [189] U. Koch, P.L.A. Popelier, and A.J. Stone. Conformational dependence of atomic multipole moments. *Chemical Physics Letters*, 238(4):253–260, 1995.
- [190] Uwe Koch and Anthony J. Stone. Conformational dependence of the molecular charge distribution and its influence on intermolecular interactions. *J. Chem. Soc., Faraday Trans.*, 92:1701–1708, 1996.
- [191] Lionel Lacombe, Norah M. Hoffmann, and Neepa T. Maitra. Exact potential energy surface for molecules in cavities. *Phys. Rev. Lett.*, 123:083201, Aug 2019.

- [192] Louis Lagardère, Félix Aviat, and Jean-Philip Piquemal. Pushing the limits of multiple-time-step strategies for polarizable point dipole molecular dynamics. *The Journal of Physical Chemistry Letters*, 10(10):2593–2599, 2019.
- [193] Stephen R Langhoff and Ernest R Davidson. Configuration interaction calculations on the nitrogen molecule. *International Journal of Quantum Chemistry*, 8(1):61–72, 1974.
- [194] Christoph Lassner. `pymp`. Available at <https://github.com/classner/pymp>.
- [195] Marie L. Laury, Lee-Ping Wang, Vijay S. Pande, Teresa Head-Gordon, and Jay W. Ponder. Revised parameters for the amoeba polarizable atomic multipole water model. *The journal of physical chemistry. B*, 119 29:9423–9437, 2015.
- [196] Susi Lehtola. Curing basis set overcompleteness with pivoted Cholesky decompositions. *The Journal of Chemical Physics*, 151(24), 12 2019. 241102.
- [197] Susi Lehtola. A review on non-relativistic, fully numerical electronic structure calculations on atoms and diatomic molecules. *International Journal of Quantum Chemistry*, 119(19):e25968, 2019.
- [198] Susi Lehtola. Straightforward and accurate automatic auxiliary basis set generation for molecular calculations with atomic orbital basis sets. *Journal of Chemical Theory and Computation*, 17(11):6886–6900, 2021. PMID: 34614349.
- [199] Susi Lehtola. Automatic generation of accurate and cost-efficient auxiliary basis sets. *arXiv preprint arXiv:2306.11039*, 2023.
- [200] Susi Lehtola, Frank Blockhuys, and Christian Van Alsenoy. An overview of self-consistent field calculations within finite basis sets. *Molecules*, 25(5):1218, 2020.
- [201] Mathieu Lewin, Elliott H Lieb, and Robert Seiringer. Improved Lieb–Oxford bound on the indirect and exchange energies. *Letters in Mathematical Physics*, 112(5):92, 2022.
- [202] E.H. Lieb and B Simon. The Hartree-Fock theory for Coulomb systems. *Communications in Mathematical Physics*, 53:185–194, 1977.
- [203] Timothy C Lillestolen and Richard J Wheatley. Redefining the atom: atomic charge densities produced by an iterative stockholder approach. *Chemical communications*, 45:5909–5911, 2008.
- [204] P.L. Lions. The concentration-compactness principle in the calculus of variations. The locally compact case, part 2. *Annales de l’I.H.P. Analyse non linéaire*, 1(4):223–283, 1984.
- [205] P.L. Lions. The concentration-compactness principle in the calculus of variations. The locally compact case, part 1. *Annales de l’I.H.P. Analyse non linéaire*, 1(2):109–145, 1984.
- [206] P.L. Lions. Solutions of Hartree-Fock equations for Coulomb systems. *Communications in Mathematical Physics*, 109:33–97, 1987.
- [207] Dishu Liu and Hermann G Matthies. Pivoted cholesky decomposition by cross approximation for efficient solution of kernel systems. *arXiv preprint arXiv:1505.06195*, 2015.
- [208] Daniele Loco, Étienne Polack, Stefano Caprasecca, Louis Lagardère, Filippo Lipparini, Jean-Philip Piquemal, and Benedetta Menucci. A QM/MM Approach Using the AMOEBA Polarizable Embedding: From Ground State Energies to Electronic Excitations. *Journal of chemical theory and computation*, 12 8:3654–61, 2016.
- [209] Philip Loftus, Marvin Waldman, and Robert F. Hout. Computer-based approaches to drug design. In Michael Williams and Jeffrey B. Malick, editors, *Drug Discovery and Development*, pages 73–96. Humana Press, Totowa, NJ, 1987.

- [210] Fritz London. Zur theorie und systematik der molekularkräfte. *Zeitschrift für Physik*, 63(3-4):245–279, 1930.
- [211] Hugh Christopher Longuet-Higgins. The electronic states of composite systems. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 235(1203):537–543, 1956.
- [212] Per-Olov Löwdin. Some properties of inner projections. *International Journal of Quantum Chemistry*, 5(S4):231–237, 1970.
- [213] Per-Olov Löwdin. On linear algebra, the least square method, and the search for linear relations by regression analysis in quantum chemistry and other sciences. *Advances in Quantum Chemistry*, 23:83–126, 1992.
- [214] Craig Lucas. LAPACK-style codes for level 2 and 3 pivoted Cholesky factorizations. *LAPACK Working*, 4(5), 2004.
- [215] Per-Olov Löwdin. On the use of resolvents and inner projections in the development of computational quantum chemistry in the future. In *Computational Chemistry And Chemical Engineering-Proceedings Of The Third Unam-cray Supercomputing Confrence*. World Scientific, 1997.
- [216] Karplus M. and McCammon J. Molecular dynamics simulations of biomolecules. *Nature Structural Biology*, 9:646–652, 2002.
- [217] Yvon Maday and Olga Mula. A generalized empirical interpolation method: Application of reduced basis techniques to data assimilation. In Franco Brezzi, Piero Colli Franzone, Ugo Gianazza, and Gianni Gilardi, editors, *Analysis and Numerics of Partial Differential Equations*, pages 221–235. Springer Milan, Milano, 2013.
- [218] Yvon Maday and Gabriel Turinici. Error bars and quadratically convergent methods for the numerical simulation of the Hartree-Fock equations. *Numerische Mathematik*, 94:739–770, 2003.
- [219] George E Martin. *Transformation geometry: An introduction to symmetry*. Springer Science & Business Media, 2012.
- [220] Per-Gunnar Martinsson and Joel A. Tropp. Randomized numerical linear algebra: Foundations and algorithms. *Acta Numerica*, 29:403–572, 2020.
- [221] Glenn J Martyna and Mark E Tuckerman. A reciprocal space based method for treating long range interactions in ab initio and force-field-based calculations in clusters. *The Journal of chemical physics*, 110(6):2810–2821, 1999.
- [222] Dominik Marx and Michele Parrinello. Ab initio path integral molecular dynamics: Basic ideas. *The Journal of chemical physics*, 104(11):4077–4082, 1996.
- [223] Stefano Massei. Some algorithms for maximum volume and cross approximation of symmetric semidefinite matrices. *BIT Numerical Mathematics*, 62(1):195–220, 2022.
- [224] Thierry Matthey. *Framework design, parallelization and force computation in molecular dynamics*. Citeseer, 2002.
- [225] Nastasia Mauger, Thomas Plé, Louis Lagardère, Simon Huppert, and Jean-Philip Piquemal. Improving condensed-phase water dynamics with explicit nuclear quantum effects: The polarizable q-amoeba force field. *The Journal of Physical Chemistry B*, 126(43):8813–8826, 2022. PMID: 36270033.

- [226] Stanisław Mazur and Stanisław Ulam. Sur les transformations isométriques d'espaces vectoriels normés. *CR Acad. Sci. Paris*, 194(946-948):116, 1932.
- [227] J. McCammon, B. Gelin, and M Karplus. Dynamics of folded proteins. *Nature*, 267:585–590, 1977.
- [228] Larry E McMurchie and Ernest R Davidson. One- and two-electron integrals over cartesian gaussian functions. *Journal of Computational Physics*, 26(2):218–231, 1978.
- [229] Daniel Mejía-Rodríguez and Andreas M Köster. Robust and efficient variational fitting of fock exchange. *The Journal of Chemical Physics*, 141(12), 2014.
- [230] B. Mennucci and R. Cammi. *Continuum Solvation Models in Chemical Physics: From Theory to Applications*. J. Wiley and Sons, 2007.
- [231] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of State Calculations by Fast Computing Machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 12 2004.
- [232] A. J. Misquitta, R. Podeszwa, B. Jeziorski, and K. Szalewicz. Intermolecular potentials based on symmetry-adapted perturbation theory with dispersion energies from time-dependent density-functional calculations. *The Journal of chemical physics*, 123(21), 2005.
- [233] Alston J Misquitta and Anthony J Stone. Camcasp 6.0.
- [234] Alston J Misquitta and Anthony J Stone. Isa-pol: Distributed polarizabilities and dispersion models from a basis-space implementation of the iterated stockholder atoms procedure. *Theoretical Chemistry Accounts*, 137:1–20, 2018.
- [235] Alston J Misquitta, Anthony J Stone, and Farhang Fazeli. Distributed multipoles from a robust basis-space implementation of the iterated stockholder atoms procedure. *Journal of Chemical Theory and Computation*, 10(12):5405–5418, 2014.
- [236] Stephan Mohr, Laura E Ratcliff, Luigi Genovese, Damien Caliste, Paul Boulanger, Stefan Goedecker, and Thierry Deutsch. Accurate and efficient linear scaling DFT calculations with universal applicability. *Physical Chemistry Chemical Physics*, 17(47):31360–31370, 2015.
- [237] Keiji Morokuma. Molecular Orbital Studies of Hydrogen Bonds. III. C=O...H–O Hydrogen Bond in H₂CO...H₂O and H₂CO...2H₂O. *The Journal of Chemical Physics*, 55(3):1236–1244, 12 2003.
- [238] Robert S Mulliken. Electronic population analysis on LCAO–MO molecular wave functions. I. *The Journal of chemical physics*, 23(10):1833–1840, 1955.
- [239] Ayako Nakata, Jack S Baker, Shereif Y Mujahed, Jack TL Poulton, Sergiu Arapan, Jianbo Lin, Zamaan Raza, Sushma Yadav, Lionel Truffandier, Tsuyoshi Miyazaki, et al. Large scale and linear scaling DFT with the CONQUEST code. *The Journal of chemical physics*, 152(16), 2020.
- [240] Christophe Narth. *Développement de champs de forces polarisables : vers la dynamique moléculaire SIBFA*. Theses, Université Pierre et Marie Curie - Paris VI, September 2015.
- [241] Sehr Naseem-Khan, Jean-Philip Piquemal, and G Andrés Cisneros. Improvement of the gaussian electrostatic model by separate fitting of coulomb and exchange-repulsion densities and implementation of a new dispersion term. *The Journal of Chemical Physics*, 155(19), 2021.

- [242] Sehr Naseem-Khan, Jean-Philip Piquemal, and G. Andrés Cisneros. Improvement of the Gaussian Electrostatic Model by separate fitting of Coulomb and exchange-repulsion densities and implementation of a new dispersion term. *The Journal of Chemical Physics*, 155(19):194103, 11 2021.
- [243] Arch W Naylor and George R Sell. *Linear operator theory in engineering and science*. Springer Science & Business Media, 1982.
- [244] Van A. Ngo, Rajiv K. Kalia, Aiichiro Nakano, and Priya Vashishta. Supercrystals of dna-functionalized gold nanoparticles: A million-atom molecular dynamics simulation study. *The Journal of Physical Chemistry C*, 116(36):19579–19585, 2012.
- [245] Jigyasa Nigam, Sergey N Pozdnyakov, Kevin K Huguenin-Dumittan, and Michele Ceriotti. Completeness of atomic structure representations. *arXiv preprint arXiv:2302.14770*, 2023.
- [246] I. Nobeli, S.L. Price, and R.J. Wheatley. Use of molecular overlap to predict intermolecular repulsion in N···H–O hydrogen bonds. *Molecular Physics*, 95(3):525–537, 1998.
- [247] Trent M. Parker, Lori A. Burns, Robert M. Parrish, Alden G. Ryno, and C. David Sherrill. Levels of symmetry adapted perturbation theory (SAPT). I. Efficiency and performance for interaction energies. *The Journal of Chemical Physics*, 140(9):094106, 03 2014.
- [248] John E Pask, Barry M Klein, Philip A Sterne, and Chingyao Y Fong. Finite-element methods in electronic-structure theory. *Computer Physics Communications*, 135(1):1–34, 2001.
- [249] Wolfgang Pauli. Über den Zusammenhang des Abschlusses der Elektronengruppen im Atom mit der Komplexstruktur der Spektren. *Zeitschrift für Physik*, 31:765–783, 1925.
- [250] Francesca Peccati, Rubén Laplaza, and Julia Contreras-García. Overcoming distrust in solid state simulations: Adding error bars to computational data. *The Journal of Physical Chemistry C*, 123(8):4767–4772, 2019.
- [251] Thomas Bondo Pedersen, Susi Lehtola, Ignacio Fdez. Galván, and Roland Lindh. The versatility of the cholesky decomposition in electronic structure theory. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, page e1692, 2023.
- [252] Leonid B. Pereyaslavets, Igor V. Kurnikov, Ganesh Kamath, Oleg Butin, Alexey Illarionov, Igor Leontyev, Michael Olevanov, Michael Levitt, Roger D. Kornberg, and Boris Fain. On the importance of accounting for nuclear quantum effects in ab initio calibrated force fields in biological simulations. *Proceedings of the National Academy of Sciences of the United States of America*, 115(36):8878–8882, Aug 2018.
- [253] José M Pérez-Jordá and Weitao Yang. Fast evaluation of the Coulomb energy for electron densities. *The Journal of chemical physics*, 107(4):1218–1226, 1997.
- [254] Juan R. Perilla, Boon Chong Goh, C. Keith Cassidy, Bo Liu, Rafael C. Bernardi, Till Rudack, Hang Yu, Zhe Wu, and Klaus Schulten. Molecular dynamics simulations of large macromolecular complexes. *Current Opinion in Structural Biology*, 31:64–74, Apr 2015.
- [255] Allan Pinkus. *N-widths in Approximation Theory*, volume 7. Springer Science & Business Media, 2012.
- [256] Jean-Philip Piquemal and G. Andrés Cisneros. Status of the Gaussian Electrostatic Model, a Density-Based Polarizable Force Field. In Pengyu Ren Qiang Cui and Markus Meuwly, editors, *Many-body effects and electrostatics*, number 28 in multi-scale computations of Biomolecules, pages 269–299. Pan Stanford Publishing, March 2016.

- [257] Jean-Philip Piquemal, G. Andrés Cisneros, Peter Reinhardt, Nohad Gresh, and Thomas A. Darden. Towards a force field based on density fitting. *The Journal of Chemical Physics*, 124(10):104101, 03 2006.
- [258] Jean-Philip Piquemal, Antonio Marquez, Olivier Parisel, and Claude Giessner-Prettre. A CSOV study of the difference between HF and DFT intermolecular interaction energy values: The importance of the charge transfer contribution. *Journal of computational chemistry*, 26(10):1052–1062, 2005.
- [259] Jean-Philip Piquemal, Lalith Perera, G Andrés Cisneros, Pengyu Ren, Lee G Pedersen, and Thomas A Darden. Towards accurate solvation dynamics of divalent cations in water using the polarizable AMOEBA force field: From energetics to structure. *The Journal of chemical physics*, 125(5), 2006.
- [260] É Polack, A Mikhalev, Geneviève Dusson, B Stamm, and F Lipparini. An approximation strategy to compute accurate initial density matrices for repeated self-consistent field calculations at different geometries. *Molecular Physics*, 118(19-20):e1779834, 2020.
- [261] Étienne Polack, Geneviève Dusson, Benjamin Stamm, and Filippo Lipparini. Grassmann extrapolation of density matrices for Born–Oppenheimer molecular dynamics. *Journal of Chemical Theory and Computation*, 17(11):6965–6973, 2021.
- [262] Robert Polly, Hans-Joachim Werner*, Frederick R Manby, and Peter J Knowles. Fast hartree–fock theory using local density fitting approximations. *Molecular Physics*, 102(21-22):2311–2321, 2004.
- [263] Jay W. Ponder, Chuanjie Wu, Pengyu Y. Ren, Vijay S. Pande, John D. Chodera, Michael J. Schnieders, Imran S. Haque, David L. Mobley, Daniel S. Lambrecht, Robert A. Distasio, Martin Head-Gordon, Gary N. I. Clark, Margaret E. Johnson, and Teresa Head-Gordon. Current status of the AMOEBA polarizable force field. *The journal of physical chemistry. B*, 114 8:2549–64, 2010.
- [264] James D Power and Russell M Pitzer. Inequalities for electron repulsion integrals. *Chemical Physics Letters*, 24(4):478–483, 1974.
- [265] Sergey N Pozdnyakov and Michele Ceriotti. Smooth, exact rotational symmetrization for deep learning on point clouds. *arXiv preprint arXiv:2305.19302*, 2023.
- [266] Sergey N Pozdnyakov, Michael J Willatt, Albert P Bartók, Christoph Ortner, Gábor Csányi, and Michele Ceriotti. Incompleteness of atomic structure representations. *Physical Review Letters*, 125(16):166001, 2020.
- [267] Benjamin P. Pritchard, Doaa Altarawy, Brett Didier, Tara D. Gibson, and Theresa L. Windus. New basis set exchange: An open, up-to-date resource for the molecular sciences community. *Journal of Chemical Information and Modeling*, 59(11):4814–4820, 2019. PMID: 31600445.
- [268] Luigi Di Puglia Pugliese, Francesca Guerriero, and Michael Poss. The resource constrained shortest path problem with uncertain data: a robust formulation and optimal solution approach. *Computers & Operations Research*, 107:140–155, 2019.
- [269] Anna Pyle. Metal ions in the structure and function of rna. *JBIC Journal of Biological Inorganic Chemistry*, 7:679–690, 2002.
- [270] Alfio Quarteroni, Andrea Manzoni, and Federico Negri. *Reduced basis methods for partial differential equations: an introduction*, volume 92. Springer, 2015.

- [271] James R. Rabinowitz, Krishnan Namboodiri, and Harel Weinstein. A finite expansion method for the calculation and interpretation of molecular electrostatic potentials. *International Journal of Quantum Chemistry*, 29(6):1697–1704, 1986.
- [272] Joshua A Rackers and Jay W Ponder. Classical Pauli repulsion: An anisotropic, atomic multipole model. *The Journal of chemical physics*, 150(8), 2019.
- [273] Joshua A Rackers, Roseane R Silva, Zhi Wang, and Jay W Ponder. Polarizable water potential derived from a model electron density. *Journal of chemical theory and computation*, 17(11):7056–7084, 2021.
- [274] Joshua A Rackers, Zhi Wang, Chao Lu, Marie L Laury, Louis Lagardère, Michael J Schnieders, Jean-Philip Piquemal, Pengyu Ren, and Jay W Ponder. Tinker 8: software tools for molecular design. *Journal of chemical theory and computation*, 14(10):5273–5289, 2018.
- [275] Aneesur Rahman and Frank H. Stillinger. Molecular Dynamics Study of Liquid Water. *The Journal of Chemical Physics*, 55(7):3336–3359, 12 2003.
- [276] Thomas D. Rasmussen, Pengyu Ren, Jay W. Ponder, and Frank Jensen. Force field modeling of conformational energies: Importance of multipole moments and intramolecular polarization. *International Journal of Quantum Chemistry*, 107(6):1390–1395, 2007.
- [277] H. Ratajczak and W.J. Orville-Thomas. *Molecular Interactions, Volume 2*. J. Wiley and Sons, 1981.
- [278] Simen Reine, Trygve Helgaker, and Roland Lindh. Multi-electron integrals. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 2(2):290–303, 2012.
- [279] P Reinhardt and J-P Piquemal. New intermolecular benchmark calculations on the water dimer: SAPT and supermolecular post-Hartree–Fock approaches. *International Journal of Quantum Chemistry*, 109(14):3259–3267, 2009.
- [280] Pengyu Ren and Jay W. Ponder. Consistent treatment of inter- and intramolecular polarization in molecular mechanics calculations. *Journal of Computational Chemistry*, 23(16):1497–1506, 2002.
- [281] Pengyu Ren and Jay W. Ponder. Polarizable atomic multipole water model for molecular mechanics simulation. *The Journal of Physical Chemistry B*, 107(24):5933–5947, 2003.
- [282] Pengyu Ren and Jay W Ponder. Polarizable atomic multipole water model for molecular mechanics simulation. *The Journal of Physical Chemistry B*, 107(24):5933–5947, 2003.
- [283] Pengyu Ren and Jay W Ponder. Temperature and pressure dependence of the AMOEBA water model. *The Journal of Physical Chemistry B*, 108(35):13427–13437, 2004.
- [284] J.A.C. Rullmann and P.Th. van Duijnen. A polarizable water model for calculation of hydration energies. *Molecular Physics*, 63(3):451–475, 1988.
- [285] G. Ruocco and M. Sampoli. Molecular dynamics on a water model with polarizability and hyperpolarizability. In Marie-Claire Bellissent-Funel and John C. Dore, editors, *Hydrogen Bond Networks*, pages 73–76. Springer Netherlands, Dordrecht, 1994.
- [286] E. Rutherford. LXXIX. The scattering of α and β particles by matter and the structure of the atom. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 21(125):669–688, 1911.
- [287] Konstantin Röder, Jerelle A. Joseph, Brooke E. Husic, and David J. Wales. Energy landscapes for proteins: From single funnels to multifunctional systems. *Advanced Theory and Simulations*, 2(4):1800175, 2019.

- [288] Celeste Sagui, Lee G Pedersen, and Thomas A Darden. Towards an accurate representation of electrostatics in classical force fields: Efficient implementation of multipolar interactions in biomolecular simulations. *The Journal of chemical physics*, 120(1):73–87, 2004.
- [289] Paweł Sałek, Stinne Høst, Lea Thøgersen, Poul Jørgensen, Pekka Manninen, Jeppe Olsen, Branislav Jansík, Simen Reine, Filip Pawłowski, Erik Tellgren, et al. Linear-scaling implementation of molecular electronic self-consistent field theory. *The Journal of chemical physics*, 126(11), 2007.
- [290] L. Salem. The forces between polyatomic molecules. ii. short-range repulsive forces. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 264(1318):379–391, 1961.
- [291] Romelia Salomon-Ferrer, Andreas W. Götz, Duncan Poole, Scott Le Grand, and Ross C. Walker. Routine microsecond molecular dynamics simulations with amber on gpus. 2. explicit solvent particle mesh ewald. *Journal of Chemical Theory and Computation*, 9(9):3878–3888, 2013. PMID: 26592383.
- [292] Jaime Axel Rosal Sandberg. *New efficient integral algorithms for quantum chemistry*. PhD thesis, Biotechnology, KTH Royal Institute of Technology, 2014.
- [293] H Bernhard Schlegel and Michael J Frisch. Transformation between cartesian and pure spherical harmonic gaussians. *International Journal of Quantum Chemistry*, 54(2):83–87, 1995.
- [294] KE Schmidt and Michael A Lee. Multipole Ewald sums for the fast multipole method. *Journal of Statistical Physics*, 89:411–424, 1997.
- [295] Martin Schütz and Frederick R Manby. Linear scaling local coupled cluster theory with density fitting. Part I: 4-external integrals. *Physical Chemistry Chemical Physics*, 5(16):3349–3358, 2003.
- [296] Walter R. P. Scott, Philippe H. Hünenberger, Ilario G. Tironi, Alan E. Mark, Salomon R. Billeter, Jens Fennen, Andrew E. Torda, Thomas Huber, Peter Krüger, and Wilfred F. van Gunsteren. The gromos biomolecular simulation program package. *The Journal of Physical Chemistry A*, 103(19):3596–3607, 1999.
- [297] U Chandra Singh and Peter A Kollman. An approach to computing electrostatic charges for molecules. *Journal of computational chemistry*, 5(2):129–145, 1984.
- [298] R. E. Skyner, J. L. McDonagh, C. R. Groom, T. van Mourik, and J. B. O. Mitchell. A review of methods for the calculation of solution free energies and the modelling of systems in solution. *Phys. Chem. Chem. Phys.*, 17:6174–6191, 2015.
- [299] John C Slater. Atomic shielding constants. *Physical review*, 36(1):57, 1930.
- [300] Brian J. Smith, David J. Swanton, John A. Pople, III Schaefer, Henry F., and Leo Radom. Transition structures for the interchange of hydrogen atoms within the water dimer. *The Journal of Chemical Physics*, 92(2):1240–1247, 01 1990.
- [301] Alex Sodt, Joseph E Subotnik, and Martin Head-Gordon. Linear scaling density fitting. *The Journal of chemical physics*, 125(19), 2006.
- [302] Olga A Stasyuk, Robert Sedlak, Célia Fonseca Guerra, and Pavel Hobza. Comparison of the DFT-SAPT and canonical EDA schemes for the energy decomposition of various types of noncovalent interactions. *Journal of Chemical Theory and Computation*, 14(7):3440–3450, 2018.

- [303] Philip J Stephens, Frank J Devlin, Cary F Chabalowski, and Michael J Frisch. Ab initio calculation of vibrational absorption and circular dichroism spectra using density functional force fields. *The Journal of physical chemistry*, 98(45):11623–11627, 1994.
- [304] Harry A. Stern and B. J. Berne. Quantum effects in liquid water: Path-integral simulations of a flexible and polarizable ab initio model. *The Journal of Chemical Physics*, 115(16):7622–7628, 10 2001.
- [305] A.J. Stone. Distributed multipole analysis, or how to describe a molecular charge distribution. *Chemical Physics Letters*, 83(2):233–239, 1981.
- [306] Anthony Stone. *The Theory of Intermolecular Forces*. Oxford University Press, 2nd edition, 01 2013.
- [307] Anthony J Stone. Distributed Multipole Analysis of Gaussian wavefunctions GDMA version 2.2. 02. *GDMA University of Cambridge, UK*, 1999.
- [308] Anthony J Stone. Distributed multipole analysis: Stability for large basis sets. *Journal of Chemical Theory and Computation*, 1(6):1128–1132, 2005.
- [309] Matthew C Strain, Gustavo E Scuseria, and Michael J Frisch. Achieving linear scaling for the electronic quantum Coulomb problem. *Science*, 271(5245):51–53, 1996.
- [310] Qiming Sun. Libcint: An efficient general integral library for Gaussian basis functions. *Journal of Computational Chemistry*, 36:1664–1671, 2015.
- [311] Qiming Sun. Long-range density fitting. *arXiv preprint arXiv:2306.12764*, 2023.
- [312] Qiming Sun, Timothy C Berkelbach, Nick S Blunt, George H Booth, Sheng Guo, Zhendong Li, Junzi Liu, James D McClain, Elvira R Sayfutyarova, Sandeep Sharma, et al. Pyscf: the python-based simulations of chemistry framework. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 8(1):e1340, 2018.
- [313] Terence Tao. Fourier transform. Available at <https://www.math.ucla.edu/~tao/preprints/fourier.pdf>.
- [314] Alireza Tehrani, James SM Anderson, Debajit Chakraborty, Juan I Rodriguez-Hernandez, David C Thompson, Toon Verstraelen, Paul W Ayers, and Farnaz Heidar-Zadeh. An information-theoretic approach to basis-set fitting of electron densities and other non-negative functions. *Journal of Computational Chemistry*, 2023.
- [315] B Th Thole. Molecular polarizabilities calculated with a modified dipole interaction. *Chemical Physics*, 59(3):341–350, 1981.
- [316] Travis H Thompson and Christian Ochsenfeld. Distance-including rigorous upper bounds and tight estimates for two-electron integrals over long-and short-range operators. *The Journal of Chemical Physics*, 147(14), 2017.
- [317] Travis H Thompson and Christian Ochsenfeld. Integral partition bounds for fast and effective screening of general one-, two-, and many-electron integrals. *The Journal of Chemical Physics*, 150(4), 2019.
- [318] Joseph John Thomson. Xl. cathode rays. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 44(269):293–316, 1897.
- [319] Christian Tilk, Ann-Kathrin Rothenbächer, Timo Gschwind, and Stefan Irnich. Asymmetry matters: Dynamic half-way points in bidirectional labeling for solving shortest path problems with resource constraints faster. *European Journal of Operational Research*, 261(2):530–539, 2017.

- [320] Hedieh Torabifard, Oleg N Starovoytov, Pengyu Ren, and G Andrés Cisneros. Development of an AMOEBA water model using GEM distributed multipoles. *Theoretical Chemistry Accounts*, 134(8):101, 2015.
- [321] David Torres Sanchez. cspy: A python package with a collection of algorithms for the (resource) constrained shortest path problem. *Journal of Open Source Software*, 5(49):1655, 2020.
- [322] Lloyd N Trefethen and David Bau. *Numerical linear algebra*, volume 181. Siam, 2022.
- [323] Thu DH Truong, Hanh H Nguyen, Hieu B Le, H-M Tran, Nguyen Duy Vy, Tran Duong Anh-Tai, Vinh NT Pham, et al. Soft parameters in Coulomb potential of noble atoms for nonsequential double ionization: Classical ensemble model and simulations. *Computer Physics Communications*, 276:108372, 2022.
- [324] O Vahtras, J Almlöf, and MW Feyereisen. Integral approximations for lcao-scf calculations. *Chemical Physics Letters*, 213(5-6):514–518, 1993.
- [325] JGCM van Duijneveldt-van de Rijdt, WTM Mooij, and FB van Duijneveldt. Testing the quality of some recent water–water potentials. *Physical Chemistry Chemical Physics*, 5(6):1169–1180, 2003.
- [326] Guido Van Rossum, Fred L Drake, et al. *Python reference manual*, volume 111. Centrum voor Wiskunde en Informatica Amsterdam, 1995.
- [327] Erik A Vazquez-Montelongo, G Andres Cisneros, and Hugo M Flores-Ruiz. Multipolar/polarizable molecular dynamics simulations of Liquid–Liquid extraction of benzene from hydrocarbons using ionic liquids. *Journal of Molecular Liquids*, 296:111846, 2019.
- [328] Fabienne Vigné-Maeder and Pierre Claverie. The exact multicenter multipolar part of a molecular charge distribution and its simplified representations. *The Journal of chemical physics*, 88(8):4934–4948, 1988.
- [329] J. von Neumann. *Mathematical Foundations of Quantum Mechanics: New edition*, volume 53. Princeton University Press, 2018.
- [330] David Wales. *Energy Landscapes: Applications to Clusters, Biomolecules and Glasses*. Cambridge Molecular Science. Cambridge University Press, 2004.
- [331] Brandon Walker, Chengwen Liu, Elizabeth Wait, and Pengyu Ren. Automation of AMOEBA polarizable force field for small molecules: Poltype 2. *Journal of Computational Chemistry*, 43(23):1530–1542, 2022.
- [332] Junmei Wang, Romain M. Wolf, James W. Caldwell, Peter A. Kollman, and David A. Case. Development and testing of a general amber force field. *Journal of Computational Chemistry*, 25(9):1157–1174, 2004.
- [333] Qiantao Wang, Joshua A. Rackers, Chenfeng He, Rui Qi, Christophe Narth, Louis Lagardère, Nohad Gresh, Jay W. Ponder, Jean-Philip Piquemal, and Pengyu Y. Ren. General model for treating short-range electrostatic penetration in a molecular mechanics force field. *Journal of Chemical Theory and Computation*, 11:2609 – 2618, 2015.
- [334] Zhenling Wang, Abdulrahman Aldossary, and Martin Head-Gordon. Sparsity of the electron repulsion integral tensor using different localized virtual orbital representations in local second-order Møller–Plesset theory. *The Journal of chemical physics*, 158(6), 2023.
- [335] Florian Weigend. A fully direct RI-HF algorithm: Implementation, optimised auxiliary basis sets, demonstration of accuracy and efficiency. *Physical Chemistry Chemical Physics*, 4(18):4285–4291, 2002.

- [336] Florian Weigend, Marco Häser, Holger Patzelt, and Reinhart Ahlrichs. Ri-mp2: optimized auxiliary basis sets and demonstration of efficiency. *Chemical Physics Letters*, 294(1):143–152, 1998.
- [337] Florian Weigend, Andreas Köhn, and Christof Hättig. Efficient use of the correlation consistent basis sets in resolution of the identity MP2 calculations. *The Journal of chemical physics*, 116(8):3175–3183, 2002.
- [338] Richard J. Wheatley and Sarah L. Price. An overlap model for estimating the anisotropy of repulsion. *Molecular Physics*, 69(3):507–533, 1990.
- [339] Christopher A White, Benny G Johnson, Peter MW Gill, and Martin Head-Gordon. The continuous fast multipole method. *Chemical physics letters*, 230(1-2):8–16, 1994.
- [340] J. L. Whitten. Coulombic potential energy integrals and approximations. *The Journal of Chemical Physics*, 58(10):4496–4501, 1973.
- [341] Christopher Williams and Matthias Seeger. Using the Nyström method to speed up kernel machines. *Advances in neural information processing systems*, 13, 2000.
- [342] Soohaeng Yoo Willow, Bing Xie, Jason Lawrence, Robert S. Eisenberg, and David D. L. Minh. On the polarization of ligands by proteins. *Phys. Chem. Chem. Phys.*, 22:12044–12057, 2020.
- [343] Qin Wu, Paul W Ayers, and Yingkai Zhang. Density-based energy decomposition analysis for intermolecular interactions with variationally determined intermediate state energies. *The Journal of chemical physics*, 131(16), 2009.
- [344] Xin Xing and Edmond Chow. Fast coulomb matrix construction via compressing the interactions between continuous charge distributions. *SIAM Journal on Scientific Computing*, 42(1):A162–A186, 2020.
- [345] Xin Xing, Hua Huang, and Edmond Chow. A linear scaling hierarchical block low-rank representation of the electron repulsion integral tensor. *The Journal of chemical physics*, 153(8), 2020.
- [346] Peng Xu, Emilie B. Guidez, Colleen Bertoni, and Mark S. Gordon. Perspective: Ab initio force field methods derived from quantum mechanics. *The Journal of Chemical Physics*, 148(9):090901, 03 2018.
- [347] Clarence Zener. Analytic atomic wave functions. *Physical Review*, 36(1):51, 1930.
- [348] Chaoqun Zhang, Filippo Lipparini, Stella Stopkowicz, Jürgen Gauss, and Lan Cheng. A cholesky decomposition-based implementation of relativistic two-component coupled-cluster methods for medium-sized molecules. *arXiv preprint arXiv:2309.08934*, 2023.
- [349] Kezhong Zhao, Marinos N Vouvakis, and Jin-Fa Lee. The adaptive cross approximation algorithm for accelerated method of moments computations of emc problems. *IEEE transactions on electromagnetic compatibility*, 47(4):763–773, 2005.
- [350] Shifan Zhao, Tianshi Xu, Edmond Chow, and Yuanzhe Xi. An Adaptive Factorized Nyström Preconditioner for Regularized Kernel Matrices. *arXiv preprint arXiv:2304.05460*, 2023.