



HAL
open science

A Coordinate System associated to a Point Cloud issued from a Manifold: Definition, Properties and Applications

Julia Flötotto

► **To cite this version:**

Julia Flötotto. A Coordinate System associated to a Point Cloud issued from a Manifold: Definition, Properties and Applications. Computational Geometry [cs.CG]. Université Nice Sophia Antipolis, 2003. English. NNT: . tel-00832487

HAL Id: tel-00832487

<https://theses.hal.science/tel-00832487>

Submitted on 10 Jun 2013

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.

THÈSE

Présentée pour obtenir le titre de

Docteur en SCIENCES
de l'Université de Nice-Sophia Antipolis

Spécialité : **INFORMATIQUE**

par

Julia FLÖTOTTO

A coordinate system associated to a point cloud issued from a manifold: definition, properties and applications

Un système de coordonnées associé à un échantillon de points d'une variété:
définition, propriétés et applications

Directeur de thèse M. **Jean-Daniel BOISSONNAT**

Soutenue publiquement le 22. Septembre 2003 devant le jury composé de :

Président	M. Pierre	BERNHARD
Rapporteurs	M. Bernard	LACOLLE
	M. Jean-Marie	MORVAN
	M. Günter	ROTE
Examineurs	M. Jean-Daniel	BOISSONNAT
	M. André	LIEUTIER

Résumé

Dans de nombreux domaines d'applications, une variété plongée dans l'espace euclidien est souvent représentée par un échantillon de points. Nous définissons dans cette thèse un système de coordonnées associé à un tel échantillon sur la variété qui généralise les coordonnées naturelles définies par Sibson. Nous exhibons ses propriétés mathématiques fondamentales ainsi que son application à l'interpolation d'une fonction définie sur la variété. Nous introduisons la notion d'atlas de Voronoï, défini comme un ensemble de cellules approximant le diagramme de Voronoï restreint à la variété et montrons son application à la reconstruction de surface et au remaillage. Enfin, nous étendons les propriétés des coordonnées naturelles aux diagrammes de puissance et proposons une synthèse des méthodes d'interpolation par coordonnées naturelles. Cette dernière détaille des preuves omises dans les articles originaux.

Mots clés: nuage de points, surface échantillonnée, diagramme de Voronoï, coordonnées naturelles, interpolation sur une surface, reconstruction 3D, remaillage, modélisation géométrique, CAO.

Abstract

Surfaces and manifolds represented by a set of discrete points are encountered in various application areas. In this thesis, we define a coordinate system on the manifold associated to such a point set which is a generalization of Sibson's natural neighbor coordinates. We show its fundamental mathematical properties as well as its application to scattered data interpolation on manifolds. Furthermore, we introduce the notion of Voronoi atlas defined as a collection of Voronoi cells that approximate the Voronoi diagram restricted to the manifold. We describe its application in surface reconstruction and re-meshing. In addition, we show the basic properties of natural neighbor coordinates in power diagrams and we survey the interpolation methods based on natural neighbor coordinates. This survey details some proofs that are omitted in the original papers.

Keywords: point set surface, Voronoi diagram, natural neighbor coordinates, scattered data interpolation, surface on surface problem, 3D reconstruction, re-meshing, geometric modeling, CAD

Thèse préparée à l'INRIA dans le projet PRISME.

Remerciements

Je tiens à remercier Jean-Daniel Boissonnat de m'avoir accueillie au sein de son équipe et d'avoir dirigé ma thèse, ainsi que pour sa disponibilité et sa patience pendant ce travail, Frédéric Cazals pour toutes les discussions animées que nous avons eues, et enfin tous les membres de mon jury qui m'ont permise par leurs remarques d'améliorer ce manuscrit de thèse : Pierre Bernhard, Bernard Lacolle, Jean-Marie Morvan, Günter Rote et André Lieutier.

Je remercie également les membres et les anciens membres du projet Prisme pour les rapports professionnels et amicaux, Frédéric Chazal pour son aide précieuse concernant les questions de topologie, les collègues de la Faculté des Sciences, ainsi que des autres projets de l'INRIA Sophia-Antipolis et du centre de documentation, mes amis d'Antibes, de Nice, du DEA ARAVIS, de Molveno 1997 et de Hambourg, toute ma famille pour leur soutien et last but not least Philippe.

*Für meine Eltern und meine Schwestern
À Philippe*

Contents

1	Introduction	1
2	Voronoi and power diagrams	13
2.1	Voronoi diagrams	13
2.2	Power diagrams	15
2.3	Space of spheres	17
2.4	Power diagrams and sections of Voronoi diagrams	21
3	Natural neighbor coordinates in power diagrams	23
3.1	Introduction	23
3.2	Definition of the coordinate function and basic properties	25
3.3	Further continuity properties of the coordinate function	34
3.4	Gradient of the coordinate function	41
3.5	Non-convex natural neighbor coordinates	52
4	Scattered data interpolation with natural neighbor coordinates	57
4.1	Sibson's natural neighbor interpolants	58
4.2	Farin's C^1 interpolant	64
4.3	Clarkson's K^0 interpolant	68
4.4	Quadratic precision interpolant	68
5	Sampled manifolds: definitions and results	71
5.1	Topological notions, medial axis and local feature size	71
5.2	Voronoi diagram restricted to a k -manifold	73
5.3	Properties of well sampled manifolds	75
6	A coordinate system on a k-manifold of \mathbb{R}^d	83
6.1	Introduction	83
6.2	Manifold neighbors	85


6.2.1	Definition of \mathcal{T}_x -neighbors	85
6.2.2	Locality of the \mathcal{T}_x -neighbors	89
6.3	Manifold coordinate system	91
6.3.1	Definition of the \mathcal{T}_x -coordinate function	92
6.3.2	Continuity properties of the coordinate function	92
6.3.3	Bounding the distance to the weighted barycenter	93
6.3.4	Visualization of the coordinate function	95
6.4	Gradient of the \mathcal{T} -coordinate function defined on a surface in \mathbb{R}^3	95
6.5	Non-convex coordinate system for convex surfaces with barycentric property	99
7	The Voronoi atlas of a point cloud on a k-manifold of \mathbb{R}^d	107
7.1	Definition and basic properties	108
7.1.1	The Voronoi atlas	108
7.1.2	The dual of the Voronoi atlas	110
7.2	Enlargement for hypersurfaces	112
7.2.1	The \mathcal{T} -restricted Voronoi diagram	113
7.2.2	The \mathcal{T} -restricted Delaunay triangulation	117
7.3	Natural-neighbor coordinates defined in the Voronoi atlas	118
7.3.1	Definition of the coordinate function μ_i	118
7.3.2	Continuity properties and gradient of the coordinate function μ_i	121
8	Applications	123
8.1	Interpolating a function on a surface S	123
8.1.1	Introduction and related work	123
8.1.2	Definition and properties of interpolants	125
8.1.3	Experimental results	126
8.2	Delaunay-based surface and curve reconstruction	135
8.2.1	Previous work	135
8.2.2	Reconstruction with \mathcal{T} -neighbors: the basic idea	143
8.2.3	\mathcal{T} -neighbor based curve reconstruction	144
8.3	Centroidal Voronoi diagrams on a surface	147
8.3.1	Introduction and previous work	147
8.3.2	\mathcal{T} -restricted centroidal Voronoi diagrams	149
8.3.3	Experimental results of re-meshing with Centroidal Voronoi diagrams	149
9	Implementation of \mathcal{T}-coordinates in CGAL	153
9.1	Introduction	153

9.2	Main classes and their functionalities	154
9.2.1	Voronoi sections	154
9.2.2	Surface and curve neighbors	158
9.2.3	Function interpolation on a surface	161
	Conclusion	163



Chapter 1

Introduction



Dans ce chapitre, nous décrivons le problème abordé dans cette thèse: la définition d'un système de coordonnées et d'un voisinage local associé à un nuage de points répartis sur une variété lisse. Nous présentons ensuite quelques applications de ce nouveau concept. L'ensemble des chapitres ainsi que les contributions nouvelles de cette thèse sont résumés dans la suite.

Définition du problème

Le diagramme de Voronoï qui constitue l'objet d'étude principal de cette thèse est une structure de données fondamentale de la géométrie algorithmique. Étant donné un ensemble de points discrets dans l'espace Euclidien, le *diagramme de Voronoï* est la partition de l'espace en cellules qui contiennent tous les points ayant le même plus proche voisin parmi les points de données.¹ Cette structure particulière que Franz Aurenhammer appelle *The Universal Spatial Data Structure*² a attiré beaucoup d'attention dans le dernier siècle. Nous renvoyons le lecteur à l'ouvrage [14] du même auteur pour une mise en perspective historique des divers aspects des diagrammes de Voronoï. Pour celui-ci, l'importance du diagramme de Voronoï est liée à trois facteurs principaux: son occurrence fréquente dans les processus naturels comme par exemple en biologie ou en cristallographie, ses propriétés mathématiques riches et intéressantes et enfin, ses applications informatiques nombreuses.

La *triangulation de Delaunay* est une structure de données toute aussi fameuse qui est très étroitement liée au diagramme de Voronoï. Plus particulièrement, la triangulation de Delaunay d'un ensemble de points donnés se définit comme le dual du diagramme de Voronoï dans le sens suivant: deux points de données sont liés par une arête dans la triangulation de Delaunay si et seulement si leur cellule de Voronoï partagent une frontière commune. Un triangle existe dans la triangulation si trois cellules de Voronoï partagent une

¹Les points de données sont parfois appelés les *sommets*.

²Une citation après la page internet de C. Gold sur les diagrammes de Voronoï <http://www.voronoi.com/>.

frontière commune etc. Les propriétés géométriques et combinatoires de ces structures, diagramme de Voronoï et triangulation de Delaunay, sont complémentaires et le calcul pratique de chacune de celles-ci se déduit facilement l'un de l'autre.

L'intérêt porté par la géométrie algorithmique pour le diagramme de Voronoï depuis l'émergence de la discipline dans les années soixante-dix se traduit aujourd'hui par de nombreux résultats pratiques et théoriques. Parmi les questions théoriques examinées on peut citer, par exemple, l'étude de la complexité combinatoire des diagrammes de Voronoï, la découverte des relations entre diagrammes de Voronoï et enveloppes convexes, ou encore la généralisation des diagrammes de Voronoï à d'autres métriques que la métrique Euclidienne. D'un point de vue beaucoup plus pratique, ces travaux ont permis la conception d'algorithmes très efficaces pour le calcul de diagrammes de Voronoï et de triangulations de Delaunay et l'on dispose aujourd'hui d'implantations robustes trouvant de nombreux champs d'applications dans des domaines divers.

Diagrammes de Voronoï sur une variété

Dans ce manuscrit de thèse, nous considérons le cas particulier où l'ensemble des points de données est issu d'une variété de dimension k plongée dans l'espace Euclidien. Ce cas est extrêmement intéressant en pratique où une variété est souvent représentée par un ensemble de points discrets, comme c'est le cas, par exemple, pour de nombreuses applications en graphisme, en conception assisté par ordinateur (CAO) et ingénierie inverse, en traitement d'image ou calcul scientifique.

Pour ces applications, la variété sous-jacente n'est pas forcément connue. Lorsque celle-ci est connue, le diagramme de Voronoï peut être défini par rapport à la distance géodésique sur la variété, on parle alors de *diagramme de Voronoï géodésique*, ou par rapport à la distance Euclidienne dans l'espace ambiant. Les diagrammes de Voronoï géodésiques sur des surfaces paramétrées sont spécifiés dans [73] par Kunze, Wolter et Rausch. Leibon et Letscher [74] formulent des conditions d'échantillonnage pour les variétés Riemanniennes garantissant l'existence d'une unique triangulation de Delaunay et d'un diagramme de Voronoï dual possédant les mêmes propriétés que celles d'un diagramme de Voronoï dans l'espace Euclidien. Cependant, bien qu'apparaissant comme plus adéquate que la distance euclidienne de par son caractère intrinsèque à la variété, l'utilisation de la distance géodésique rend le calcul pratique du diagramme de Voronoï très difficile. Pour illustrer ces difficultés, prenons pour exemple un des éléments de base du diagramme de Voronoï, le *bissecteur* de deux points de données, c. à. d. l'ensemble des points sur la variété qui sont équidistants aux deux points de données. Alors que celui-ci est défini par un hyperplan en métrique Euclidienne, il peut être courbé et déconnecté pour la métrique géodésique.

Le *diagramme de Voronoï restreint* spécifié par Edelsbrunner et Shah [54] est défini quant à lui comme la restriction du diagramme de Voronoï Euclidien à la variété, une cellule du diagramme de Voronoï restreint correspondant à l'intersection de la cellule de Voronoï avec la variété. Bien que dépendant du plongement de la variété dans l'espace \mathbb{R}^d , il est possible de démontrer que la triangulation de Delaunay

duale au diagramme ainsi calculé est homéomorphe à la variété si certaines conditions d'échantillonnage (qui peuvent être vérifiées) sont satisfaites. Comme pour le cas du diagramme de Voronoï géodésique, le principal inconvénient du diagramme de Voronoï restreint est que celui-ci est impossible à calculer lorsque la variété sous-jacente est connue uniquement sur les points de données et que son calcul reste excessivement coûteux lorsque la variété sous-jacente est connue.

Nous proposons dans cette thèse une approche très pratique pour approximer une cellule du diagramme de Voronoï dans le cas où seulement un nuage de points répartis sur la variété est donné. L'idée de base consiste à approximer localement la variété autour d'un point de données par son espace tangent en ce point et à calculer le diagramme restreint à cette approximation locale. Il existe dans la littérature plusieurs méthodes pour estimer, à partir d'un nuage de points issu d'une variété, l'espace tangent en un point donné. Le chapitre 6 montre que l'espace tangent et la variété sont très proches à l'intérieur d'une cellule du diagramme Euclidien si la variété est bien échantillonnée relativement à la courbure et à l'épaisseur locale de la variété.

Le diagramme ainsi calculé correspond à la restriction du diagramme de Voronoï défini dans \mathbb{R}^d à un sous-espace linéaire de dimension k et est en fait un diagramme de puissance k -dimensionnel. Ses propriétés sont par conséquent bien connues et le diagramme lui-même est facilement calculable avec les méthodes existantes.

On appelle *atlas de Voronoï* la collection des cellules du diagramme de Voronoï restreintes à l'espace tangent de chaque sommet (Chapitre 7). Les principales propriétés d'un tel atlas résident dans le fait que chacune de ses cellules est homéomorphe à la cellule correspondante sur la variété et que celui-ci approxime linéairement et de façon optimale la variété sous-jacente (chacun des points de la variété étant approximé par l'espace tangent de son plus proche sommet). Les bords des cellules de l'atlas de Voronoï ne sont toutefois pas en général communs avec d'autres cellules.

La triangulation duale à l'atlas de Voronoï est un graphe dans lequel deux sommets sont reliés si l'espace tangent à un des deux points intersecte la facette de Voronoï commune aux deux points. Bien que la triangulation duale à l'atlas de Voronoï n'est pas une triangulation globale des points de données (la relation de voisinage n'est pas nécessairement symétrique), celle-ci permet, néanmoins, d'obtenir autour de chaque sommet une petite triangulation qui est homéomorphe à un voisinage du sommet sur la variété (pourvu qu'une certaine condition d'échantillonnage soit respectée). Le principal avantage d'une telle triangulation est que, pour un sommet donné, l'ensemble des sommets adjacents dans la triangulation définit un voisinage local autour du point.

Interpolation naturelle

L'interpolation par coordonnées naturelles, introduite en 1980 par Sibson [90][91], est une méthode d'interpolation de données discrètes multivariées qui constitue l'une des applications pratiques les plus importantes des diagrammes de Voronoï. Les coordonnées naturelles constituent un système de coordonnées

barycentriques défini à partir du diagramme de Voronoï. Étant donné un point de requête, on associe un poids à chacun de ses sommets voisins de manière à ce que le point de requête soit leur barycentre pondéré.

Les méthodes d'interpolation basées sur les coordonnées naturelles sont particulièrement intéressantes du fait qu'elles s'adaptent facilement à des données non uniformes et anisotropes, ce qui les rend particulièrement adaptées aux problèmes issus de la modélisation géophysique et des systèmes d'information géographique (SIG). De plus, ces méthodes d'interpolation ne dépendent pas du maillage utilisé et sont par conséquent indépendantes d'une partition artificielle et a priori de l'espace. Une synthèse des différentes méthodes d'interpolation basées sur les coordonnées naturelles est donnée dans le Chapitre 4.

Il est possible de généraliser la définition des coordonnées naturelles aux diagrammes de puissance et, par conséquent, aux diagrammes de Voronoï restreints à un espace tangent. Défini sur une variété, un tel système de coordonnées est très utile pour l'interpolation de données multivariées sur une surface. On note cependant que cette solution est seulement approximative, le système de coordonnées n'étant barycentrique que par rapport aux points de données projetés sur l'espace tangent. Quelques propriétés importantes de ce système de coordonnées sont développées dans cette thèse. On montre notamment que ce système de coordonnées est local sur la variété, qu'il est C^1 continu presque partout sur celle-ci et on détermine son gradient.

Dans le cas où la variété n'est pas connue, nous définissons un système de coordonnées basé sur l'atlas de Voronoï. Dans ce nouveau système, les coordonnées ne sont plus définies dans l'espace tangent du point requête mais par rapport aux espaces tangents des sommets voisins au point requête. Bien que ne disposant pas d'autant de bonnes propriétés que le système de coordonnées précédent, celui-ci a l'avantage d'être applicable directement à partir d'un échantillon de points extrait de la variété.

Applications

Dans ce contexte, nous étudions trois applications pratiques de l'atlas de Voronoï et du diagramme de Voronoï restreint à un espace tangent. La première est l'interpolation de données éparpillées sur une surface: étant donné un ensemble de valeurs discrètes d'une fonction définie sur une surface, on souhaite interpoler cette fonction sur la surface. Nous montrons dans cette thèse que les méthodes d'interpolation par coordonnées naturelles se généralisent directement aux diagrammes de Voronoï restreints aux espaces tangents et constituent une solution adéquate pour l'interpolation sur des surfaces.

La seconde application aborde un problème ayant attiré beaucoup d'attention ces vingt dernières années à savoir la reconstruction de surface ou de courbe à partir d'un nuage de points. Nous proposons ici une solution qui est très semblable aux méthodes de filtrages par Voronoï où la variété est approximée par un sous-ensemble des faces de la triangulation de Delaunay.

Finalement, j'ai développé un prototype pour le calcul d'une approximation des diagrammes de Voronoï centroïdal sur des surfaces. Un diagramme de Voronoï centroïdal est un diagramme de Voronoï dans lequel chaque sommet coïncide avec le centre de masse de la cellule qu'il définit. Pour le cas d'une

surface, des *diagrammes de Voronoï centroïdales contraintes* ont été définis par Du, Gunzburger et Ju [50]. La méthode la plus efficace pour calculer ce type de diagramme est basée sur le calcul du diagramme de Voronoï restreint à une surface. Comme nous l'avons vu précédemment, le calcul d'un diagramme de Voronoï restreint à une surface est cependant très coûteux et seul le calcul de tel diagramme sur des cas particuliers, comme par exemple la sphère, est envisageable. La nouvelle idée présentée dans la Section 8.3 consiste à calculer le diagramme de Voronoï centroïdal à partir de l'atlas de Voronoï, ce qui revient à approximer la surface à l'intérieur de chaque cellule par le plan tangent à la surface au sommet qui génère la cellule. Le diagramme de Voronoï est alors calculé à partir du diagramme approximé d'une manière itérative. Nous présentons des résultats expérimentaux très convaincants, notamment en ce qui concerne la régularité de la triangulation duale par exemple.

Résumé des chapitres

Ce manuscrit est organisé de la manière suivante.

1. Après l'introduction, le deuxième chapitre présente les concepts de diagrammes de Voronoï et de diagrammes de puissance. Leur interprétation dans l'espace des sphères est également rappelée. Enfin, il est montré que l'intersection d'un diagramme de puissance avec un sous-espace affine est elle-même un diagramme de puissance. Cette observation est le cœur de notre travail sur les systèmes de coordonnées sur des variétés.
2. Le troisième chapitre généralise le concept de coordonnées naturelles aux diagrammes de puissances. Après un état de l'art sur les coordonnées naturelles dans des diagrammes de Voronoï et la définition des fonctions de coordonnées généralisées, les propriétés de continuité et le gradient de la fonction de coordonnée sont montrés. De plus, un nouveau système de coordonnées proposé par Clarkson dans [35] est présenté. Ce dernier possède quelques propriétés supplémentaires par rapport à celles des coordonnées naturelles sur lequel il est basé.
3. Le quatrième chapitre contient une synthèse des différentes méthodes d'interpolation basées sur les coordonnées naturelles. Cet état de l'art contient, en plus, quelques améliorations de méthodes existantes et il complète certaines preuves qui n'ont pas été fournies dans les articles originaux.
4. Le cinquième chapitre introduit des notions topologiques et géométriques concernant les variétés échantillonnées. Nous définissons ensuite le diagramme de Voronoï restreint à une variété et nous spécifions une condition d'échantillonnage supposée être vérifiée dans la suite. La dernière partie rappelle plusieurs résultats relatifs aux propriétés locales d'une variété échantillonnée et à l'estimation de certaines quantités de la variété, comme par exemple les normales, à partir des points d'échantillonnage et en supposant la condition d'échantillonnage vérifiée.

5. Le coeur de ce travail est présenté dans le sixième et le septième chapitre. Étant donné un nuage de points issu d'une variété, le sixième chapitre contient la définition d'un voisinage et d'un système de coordonnées pour chaque point de la variété. Il est montré sous quelles conditions ces notions sont locales sur la variété. En outre, les propriétés de continuité des fonctions de coordonnées et leurs gradients (pour des surface uniquement) sont examinés.
6. Le septième chapitre définit l'atlas de Voronoï d'un nuage de points sur une variété. Quelques propriétés important de l'atlas de Voronoï et de son dual sont prouvées. Pour des variétés de dimension $d - 1$, nous proposons une définition alternative pouvant être généralisée afin de connecter les différentes cellules de Voronoï. Enfin, un système de coordonnées est défini par rapport à l'atlas de Voronoï et quelques unes de ses propriétés sont démontrées.
7. Quelques applications de ces concepts sont présentés dans le Chapitre 8. Les trois applications traitées en détails sont l'interpolation de données répartie sur une variété, la reconstruction de surface et de courbes ainsi qu'une technique de remaillage par diagramme de Voronoï centroïdal. Ce chapitre donne enfin quelques résultats expérimentaux.
8. Finalement, le dernier chapitre abordent les questions relatives à l'implantation pratique de ces travaux.

Contributions

À partir d'un nuage de points répartis sur une variété, la nouvelle approche proposée dans cette thèse consiste à approximer linéairement une cellule du diagramme de Voronoï restreint à la variété afin d'en déduire un voisinage et un système de coordonnées sur la variété associé au nuage de points (Chapitre 6). Enfin, nous introduisons la notion d' *atlas de Voronoï* (Chapitre 7), et détaillons trois applications potentielles.

D'autres contributions se trouvent dans la preuve des propriétés des coordonnées naturelles généralisées au cas des diagrammes de puissance et dans l'amélioration des méthodes d'interpolation par coordonnées naturelles.

Introduction



This chapter motivates the problems addressed in this thesis: the definition of a coordinate system and a local neighborhood associated with a set of points scattered on a smooth manifold. We present some applications of this new concept and discuss its further impact. Then, an overview of the organization of this thesis is given and the contributions of this thesis are summarized.

Problem definition

The Voronoi diagram is one of the most fundamental data structures in computational geometry. It is also the main object of attention of this thesis. Given a set of discrete data points in Euclidean space, the *Voronoi diagram* is a partition of the space in cells that contain all points having the same nearest neighbor among the data points. This concept has attracted a lot of attention in the last century. To underlay the generality of Voronoi diagrams, we cite Franz Aurenhammer who called it *The Universal Spatial Data Structure*.³ In [14], Aurenhammer overviews different aspects of Voronoi diagrams including the historical perspective. He gives three main reasons for the importance of Voronoi diagrams, namely, its frequent occurrence in natural processes for example in crystallography, its interesting mathematical properties and, at last, the numerous applications that are based on Voronoi diagrams.

Another prominent data-structure that is closely related to Voronoi diagrams is the *Delaunay triangulation*. The Delaunay triangulation is a graph on the data points that is dual to the Voronoi diagram in the following sense: Two data points are connected by an edge in the Delaunay triangulation if and only if their Voronoi cells are adjacent, i.e. they share a common boundary. A triangle exists in the Delaunay triangulation if three Voronoi cells are adjacent one to another and so on. The geometric and combinatorial properties of both structures, Voronoi diagrams and Delaunay triangulations, are complementary and most properties translate easily from one structure to the other. By this, the study of both concepts is related in theoretical aspects as well as for practical issues concerning, for example, the computation of Voronoi diagrams and Delaunay triangulations.

Detached from a particular application, the computational geometry community started to investigate theoretical and practical viewpoints of Voronoi diagrams since its emergence in the seventies. In the

³A citation after C. Gold's Voronoi website <http://www.voronoi.com/>.

eighties, it has become a major interest of this field. The theoretical questions include the generalization of Voronoi diagrams to different kinds of metrics, the study of the combinatorial complexity of Voronoi diagrams, the discovery of a relation between Voronoi diagrams and convex hulls, etc. On the other hand, practical algorithms have been developed and implemented to compute Voronoi diagrams and Delaunay triangulations. This permitted the development of various applications in all kinds of areas.

Voronoi diagrams on manifolds

In this thesis, we are interested in the particular case where the data points are issued from a k -dimensional manifold embedded in Euclidean space. This case is particularly important in practice where surfaces and general k -manifolds represented by a set of unordered sample points are encountered in many application areas such as computer graphics, computer aided design (CAD) and reverse engineering, image processing, and scientific computation. In some cases, the underlying manifold might be known, in others not.

If it is known, one can either resort to the *geodesic Voronoi diagram* that is defined by the metric on the manifold or to the Voronoi diagram that is determined by the Euclidean metric in the embedding space. Geodesic Voronoi diagrams on parametric surfaces are specified by Kunze, Wolter and Rausch in [73]. In [74], Leibon and Letscher formulate a sampling condition for Riemann manifolds that guarantees the existence of a unique Delaunay triangulation and a dual Voronoi diagram with the same properties as in Euclidean space. While the geodesic distance might be more adequate as it is intrinsic to the manifold, it has the drawback that it is a lot more difficult to compute. As an example consider the *bisector* of two points, i.e. the set of points that are at equal distance to the points. It is a basic element in the computation of Voronoi diagrams. While in Euclidean metric, the bisector is a hyperplane, using geodesic distance, it is curved and might not even be connected.

Using Euclidean distance, Edelsbrunner and Shah [54] define a *restricted Voronoi diagram* as the restriction of the Euclidean Voronoi diagram to the manifold. A cell of the restricted Voronoi diagram is then the intersection of the cell with the manifold. Even though this Voronoi diagram depends on the particular embedding of the manifold in \mathbb{R}^d , one can, for example, show that the dual Delaunay triangulation is homeomorphic to the surface if some (checkable) conditions are satisfied. However, even this diagram might be too costly to compute.

If the underlying manifold is known only by the data points, it is, actually, impossible to compute the restricted Voronoi diagram. For this reason, this thesis proposes a very practical approach to approximate a cell of the restricted Voronoi diagram – even if only a point cloud issued from the manifold is known. The main idea consists of approximating the manifold locally, inside a Voronoi cell, by the tangent space to the manifold at the data point that generates the cell. If the manifold is well sampled with respect to the curvature and to the local thickness of the manifold, the tangent space and the manifold are close inside the cell. This is explained in detail in Chapter 6. We need to mention that there exist several methods in the literature to estimate the tangent spaces from a point cloud.

Furthermore, the restriction of a Voronoi diagram defined in \mathbb{R}^d to a k -dimensional linear subspace is, in fact, a k -dimensional power diagram, i.e. a Voronoi diagram for weighted points. Consequently, it satisfies all properties known for power diagrams, and existing algorithms can be easily adopted. The complexity of a cell of the diagram (concerning the number of faces) depends on the dimension k of the manifold and not on the dimension of the embedding space. The Voronoi diagram restricted to the tangent space of the manifold at a point x is called the \mathcal{T}_x -restricted Voronoi diagram.

The collection of Voronoi cells restricted to the tangent space of each sample point is the so-called \mathcal{T} -restricted Voronoi atlas defined in Chapter 7. By the Voronoi atlas, the manifold is approximated linearly in a best possible way with respect to the given data: a point on the manifold is approximated by the tangent space of the closest data point. If the sampling is sufficiently dense, each cell of the Voronoi atlas is homeomorphic to the corresponding cell on the manifold. However, the cells of the Voronoi atlas have, in general, no common boundary.

The dual triangulation of the Voronoi atlas is the graph in which two data points are connected if one of their tangent spaces intersects the dual Voronoi face. This does not yield a global triangulation of the data points because the neighborhood relation is not necessarily symmetric. Nevertheless, it allows to obtain a small triangulation around each data point, which is homeomorphic to a small neighborhood of that point on the manifold – provided that some sampling condition is satisfied. For each sample point, the set of adjacent sample points in the triangulation defines a local neighborhood around the point. This is interesting in a variety of applications.

Natural neighbor interpolation

One important application of Voronoi diagrams is *natural neighbor interpolation*. This is a multi-variate scattered data interpolation method that has been introduced by Sibson in 1980 [90][91]. From the Voronoi diagram, Sibson defines a barycentric coordinate system – the so-called natural neighbor coordinates. Given a query point, barycentric coordinates associate a weight to each of the neighboring data points such that the query point is their weighted barycenter. Sibson shows that the natural neighbor coordinates fulfill this definition.

Interpolation methods based on natural neighbor coordinates are particularly interesting because they adapt easily to non-uniform and highly anisotropic data. For this reason, they have been widely used in geophysical modeling and GIS applications. Other important properties are the locally bounded support of the coordinate functions and the meshless character of the methods. No artifacts are introduced due to a priori and often arbitrary subdivision of the space. Apart from Sibson's methods, other interpolation methods exist that are based on natural neighbor coordinates. An overview is given in this thesis.

The definition of natural neighbor coordinates can be generalized to power diagrams, and, consequently, to \mathcal{T}_x -restricted Voronoi diagrams. The resulting coordinate system – defined on the manifold – responds to the demand of a definition of barycentric coordinates on general manifolds, notably, for scattered data in-

terpolation on surfaces. Note that the solution is only approximate, so the coordinate system is barycentric only with respect to each tangent plane. Some important properties of the resulting \mathcal{T}_x -coordinate system are developed in this thesis. Notably, it is shown that it is local on the manifold, that it is C^1 continuous almost everywhere on the manifold, and its gradient is determined.

If the manifold is not known, another natural neighbor coordinate system is defined with respect to the Voronoi atlas. While the \mathcal{T}_x -coordinate system is defined in the tangent space of a point on the manifold, the natural neighbor coordinates in the Voronoi atlas are defined in the tangent spaces of the sample points. Its definition domain is a small neighborhood of the manifold. Thus, even if its properties are weaker than those of the \mathcal{T}_x -coordinate system, it has the advantage that it can be applied without knowing the manifold.

Applications

In this context, I consider three applications of the \mathcal{T}_x -restricted Voronoi diagram and the Voronoi atlas in more detail. The first is scattered data interpolation on surfaces. Given discrete values of a function that is defined on a manifold, the problem is to interpolate this function elsewhere on the manifold. Natural neighbor interpolation can be generalized in a straightforward manner to the \mathcal{T}_x -restricted Voronoi diagram. In this thesis, we show that this yields an adequate solution to interpolation on surfaces.

The second application is concerned with reconstructing a surface or a curve from a set of scattered points. This problem has received a lot of attention in the last twenty years. Our solution is closely related to existing so-called Voronoi-filtering methods where the manifold is approximated by a subset of faces of the Delaunay triangulation.

At last, I developed a prototype for approximating centroidal Voronoi diagrams on manifolds. A centroidal Voronoi diagram is a Voronoi diagram such that each point defining a Voronoi cell coincides with the mass centroid of that cell. For the surface case, so-called *Constrained Centroidal Voronoi diagrams* have been specified by Du, Gunzburger and Ju [50]. Their method is not practical for general surfaces because, so far, there exist no software for computing the restricted Voronoi diagram (even with Euclidean metric). The new idea presented in Section 8.3 is compute the centroidal Voronoi diagram with respect to the Voronoi atlas. This means to linearly approximate the surface inside each Voronoi cell by the tangent plane at the point that generates the cell. The centroidal Voronoi diagram is computed with respect to the approximated diagram by an iterative method. Some experimental results are presented.

Outline of this thesis

This thesis is organized as follows.

1. After this introduction, the second chapter introduces the concepts of Voronoi and power diagrams which are the basic underlying concepts to this work. It recalls also their description in the space of

spheres. Finally, it shows that the intersection of a power diagram with an affine subspace is itself a power diagram which is, in some sense, the key to our work about a coordinate system on surfaces.

2. The third chapter generalizes the concept of natural neighbor coordinates from Voronoi diagrams to power diagrams. After a state of art about natural neighbor coordinates in Voronoi diagrams and the definition of the generalized coordinate system, the continuity properties and the gradient of the coordinate function are proven. Furthermore, a new coordinate system is presented that is based on natural neighbor coordinates but which has additional properties. It has been proposed by Clarkson in [35].
3. A survey about scattered data interpolation methods that are based on natural neighbor coordinate systems is presented in the fourth chapter.
4. The fifth chapter introduces the necessary topological and geometric notations concerning the sampled manifold. Then, it defines how a Voronoi diagram can be restricted to a manifold, and it specifies a sampling condition that we suppose to be satisfied in the sequel. At last, it re-calls several results about local properties of a sampled manifold and about the estimation of certain quantities from the sample points supposing the sampling condition is met.
5. The heart of this work is presented in the sixth and in the seventh chapter. Given a set of points issued from a manifold, the sixth chapter contains the definition of a neighborhood as well as a coordinate system for each point on the manifold. It is shown under which conditions both notions are local on the manifold. Furthermore, the continuity properties of the coordinate functions as well as their gradients are examined (the latter for surfaces only).
6. The seventh chapter defines the \mathcal{T} -restricted Voronoi atlas of a point cloud on a manifold. Some properties of the Voronoi atlas and its dual are proven. For $(d - 1)$ -manifolds, we present a different definition which allows to connect the Voronoi cells between the tangent spaces. A coordinate system with respect to the Voronoi atlas is specified and some of its properties are shown.
7. The applications of these concepts are presented in Chapter 8. The three applications that are investigated in detail are scattered data interpolation of functions on general surfaces, surface and curve reconstruction and re-meshing with centroidal Voronoi diagrams. This chapter presents also experimental results.
8. Finally, the last chapter treats some implementation issues to demonstrate the practical feasibility of this work.

Contributions

The main contribution of this thesis is the definition of a neighborhood as well as a coordinate system on a manifold associated with a point cloud scattered over the manifold in Chapter 6 and the assembly of the local neighborhoods to the Voronoi atlas described in Chapter 7. These chapters contain exclusively original work. The proof of Lemma 6.3 is inspired from [9]. The content of Chapter 6 is the subject of two publications [24] [25].

In Chapter 6, it is shown under which conditions the neighborhood and the coordinate system are local on the manifold. Furthermore, the continuity properties of the coordinate functions as well as their gradients are examined (the latter for surfaces only). The approximation of a Voronoi diagram restricted to a manifold by a \mathcal{T} -restricted Voronoi atlas is defined as well as its dual collection of Delaunay simplices. The alternative definition of the Voronoi atlas in Section 7.2 is inspired from the work of Nullans [83]. A coordinate system is determined with respect to the Voronoi atlas.

Three applications of the coordinate system and the atlas are presented. For scattered data interpolation of functions defined on surfaces, I adopted the existing methods of natural neighbor interpolation. The methods are implemented and some experimental results are presented.

The second application is surface and curve reconstruction. It appears that this work is very close to the work of Dey et al. on the *COCONE*-algorithm (e.g. [7]). The reconstruction algorithm for surfaces is only partially implemented.

The third application is remeshing with centroidal Voronoi diagrams. I realized a prototype that produces convincing results. Implementation issues are also discussed in the practical part of the thesis, namely in Chapter 9.

Another important contribution of this thesis consists of the generalization of the concept of natural neighbor coordinates from Voronoi diagrams to power diagrams. The definition of this generalized coordinate system is known from [12], however, to my knowledge, there exists no further analysis of its properties. Chapter 3 contains, notably, a proof of the continuity properties (Lemma 3.6 and Lemma 3.7) and a proof of the gradient of the coordinate function (Lemma 3.8 and Lemma 3.11). Also, the behavior of the coordinate function close to a data point is examined in detail in Lemma 3.14 and Lemma 3.15.

At last, this thesis provides a synthesis about scattered data interpolation methods that are based on natural neighbor coordinate systems. It completes omissions of the original papers, notably, the proof of the C^1 continuity of Sibson's interpolant was not contained in the original paper. A new interpolant is defined that is designed specifically for interpolation of quadratic functions. Clarkson's interpolation method that is described in section 4.3 is so far unpublished. Clarkson mentioned it in a talk [35].

Chapter 2

Voronoi and power diagrams

One of the main topics of this thesis are Voronoi and power diagrams. In this section, we give a short introduction and the necessary definitions concerning both concepts.

2.1 Voronoi diagrams

We define a k -polytope to be the intersection of a finite number of closed halfspaces in \mathbb{R}^d of dimension k . It can be bounded or unbounded. A bounded k -polytope can equivalently be defined as the convex hull of a finite set of points in \mathbb{R}^d . A *cell complex* C is a set of polytopes such that any face of a polytope is itself a polytope in C and such that the intersection of two polytopes of C is either empty or it is a polytope of smaller dimension which is their common face of maximal dimension. If a bounded k -polytope is the convex hull of $k + 1$ (affinely independent) points it is called a k -simplex. A cell complex containing only simplices is called a *simplicial complex*.

Let $\mathcal{P} = \{\mathbf{p}_1, \dots, \mathbf{p}_n\}$ be a set of points in \mathbb{R}^d . Without real loss of generality, we can assume that no $d + 2$ points lie on the same sphere. The Voronoi cell of \mathbf{p}_i is the locus of points that are closer to \mathbf{p}_i than to any other point of \mathcal{P} :

$$V(\mathbf{p}_i) = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x} - \mathbf{p}_i\| \leq \|\mathbf{x} - \mathbf{p}_j\| \quad \forall j = 1, \dots, n\}$$

where $\|\mathbf{x} - \mathbf{y}\|$ denotes the Euclidean distance between points $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$. Let the bisector hyperplane of two points \mathbf{p}_i and \mathbf{p}_j be called \mathcal{H}_{ij} , and \mathcal{H}_{ij}^i be the halfspace limited by \mathcal{H}_{ij} that contains \mathbf{p}_i . Then, $V(\mathbf{p}_i)$ is the intersection of the halfspaces $\mathcal{H}_{ij}^i, j \neq i$. It is a non-empty convex polytope that contains \mathbf{p}_i and that is unbounded if \mathbf{p}_i is a vertex of the convex hull $\text{conv}(\mathcal{P})$ of \mathcal{P} .

One can show that the Voronoi cells and their faces form a cell complex, see e.g. [26, chapter 17]. This leads to the following definition:

Definition 2.1 *The cell complex of Voronoi cells $V(\mathbf{p}_i), i = 1, \dots, n$, is called the Voronoi diagram of \mathcal{P} or $\text{Vor}(\mathcal{P})$.*

Let \mathcal{P}' be a subset of points of \mathcal{P} whose Voronoi cells have a non-empty intersection. The convex hull $\text{conv}(\mathcal{P}')$ is called a Delaunay face. The Delaunay faces form a cell complex called the Delaunay triangulation of \mathcal{P} , denoted $\text{Del}(\mathcal{P})$. The Delaunay triangulation is *dual* to the Voronoi diagram which means that for any k -face of $\text{Del}(\mathcal{P})$, $k \leq d$, there exists a dual $(d - k)$ -face in $\text{Vor}(\mathcal{P})$ that is exactly the intersection of the Voronoi cells of the $k + 1$ generators of the Delaunay face. In other words, if two Voronoi cells share a common $(d - 1)$ -face then the points generating the Voronoi cells are connected by an edge in the Delaunay triangulation, and vice versa. If three cells share a common $(d - 2)$ -face then their generators are connected by a triangle, etc. See also Figure 2.1.




Figure 2.1: (left) a Voronoi diagram, (right) the dual Delaunay triangulation.

Each point on a $(d - k)$ -dimensional Voronoi face is equidistant from the $k + 1$ vertices of the dual Delaunay face and is further away from all other vertices (since it is on the boundary of each of their Voronoi cells). Consequently, there exist an open ball centered at the point that has the $k + 1$ generators on its boundary and that is empty of other points from \mathcal{P} . We call such a ball an *empty sphere*. In the dual, this means that for each Delaunay face there is an empty sphere passing through its generators that is centered on the dual Voronoi face. This property is called the *empty-sphere* property of the Delaunay triangulation. Under the general position assumption, the d -dimensional faces of $\text{Del}(\mathcal{P})$ are simplices and $\text{Del}(\mathcal{P})$ is a simplicial complex. Hence the name triangulation.

Each vertex of the Voronoi diagram is generated by the intersection of $(d + 1)$ -Voronoi cells. Its dual is a d -dimensional face of $\text{Del}(\mathcal{P})$. The Voronoi vertex is therefore the center of the circumsphere of the Delaunay simplex. These empty spheres, that are centered on Voronoi vertices and circumscribe a Delaunay face, are called *Delaunay spheres*.

The *arrangement of Delaunay spheres* is the partition of \mathbb{R}^d into cells such that all points of one cell lie inside the same set of Delaunay spheres. The points of one cell have the following in common: if they were added to \mathcal{P} , all Voronoi vertices of $\text{Vor}(\mathcal{P})$ whose Delaunay spheres cover the cell would be

destroyed and so would their dual Delaunay faces. We say that a point \mathbf{x} is inserted into $\text{Vor}(\mathcal{P})$ if $\text{Vor}(\mathcal{P})$ is replaced by $\text{Vor}(\mathcal{P} \cup \{\mathbf{x}\})$ (resp. for $\text{Del}(\mathcal{P})$). A point is *in conflict* with a Delaunay sphere if it lies inside the sphere. It is also *in conflict* with the Voronoi vertex that is the center of the Delaunay sphere and that disappears at its insertion into $\text{Vor}(\mathcal{P})$. After its insertion into $\text{Del}(\mathcal{P})$, the new point is incident to the vertices belonging to the destroyed Delaunay faces.

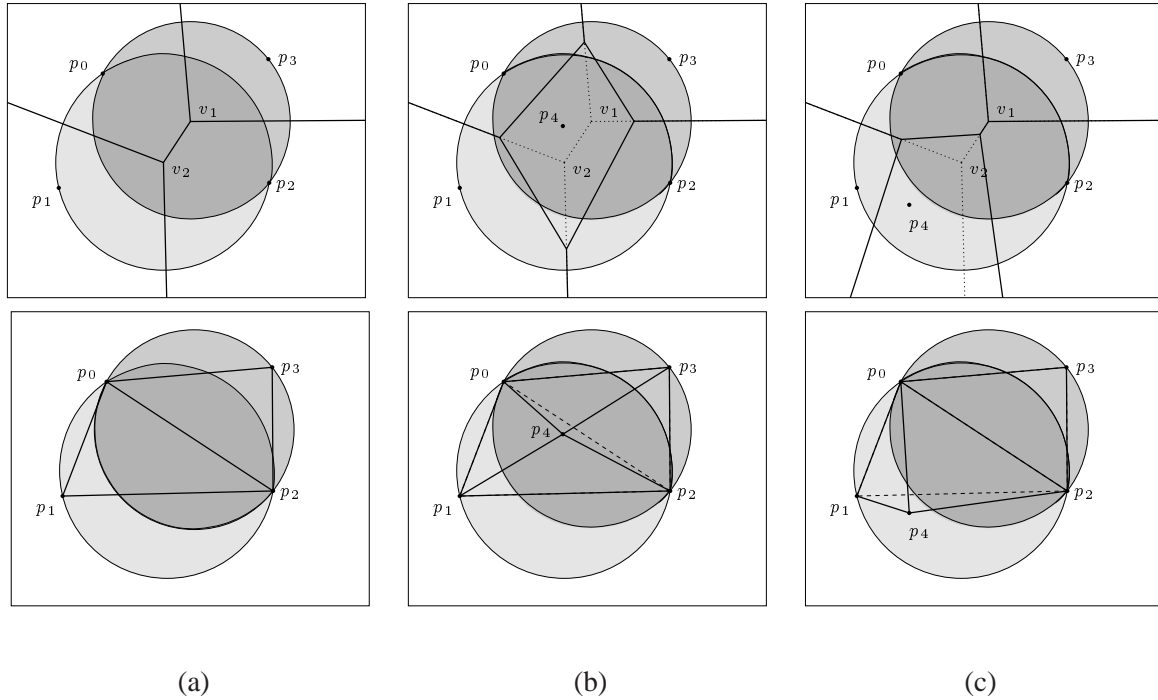


Figure 2.2: The arrangement of Delaunay spheres (top (a)) with the Voronoi diagram, (bottom (a)) with the Delaunay triangulation. Insertion of \mathbf{p}_4 (b) two Delaunay spheres are in conflict with \mathbf{p}_4 , (c) one Delaunay sphere is in conflict with \mathbf{p}_4 .

Figure 2.2 shows an arrangement of two Delaunay spheres. The figures in the middle show that two Voronoi vertices (top figure) and resp. two Delaunay triangles (bottom figure) are destroyed after the insertion of \mathbf{p}_4 into the dark grey cell of the arrangement. At its insertion onto the light grey cell, only one vertex and resp. one triangle is destroyed (right figures).

2.2 Power diagrams

Next, we generalize the definitions of the previous section to power diagrams, i.e. weighted Voronoi diagrams. These diagrams are also called *Laguerre diagrams*. Let $\mathcal{B} = \{B_1, B_2, \dots, B_n\}$ be a set of weighted points (or sites) in \mathbb{R}^d , $B_i = (\mathbf{p}_i, w_i) \in \mathbb{R}^d \times \mathbb{R}$. A weighted point B_i can also be considered as a sphere with center \mathbf{p}_i and radius $\sqrt{w_i}$. Notice, that w_i might be negative and the associated sphere

imaginary. The *power distance* or *Laguerre distance* $\Gamma(B_i, B_j)$ between two weighted points is defined as

$$\Gamma(B_i, B_j) := \|\mathbf{p}_i - \mathbf{p}_j\|^2 - w_i - w_j.$$

A point $\mathbf{x} \in \mathbb{R}^d$ has power distance $\Gamma_i(\mathbf{x})$ with respect to B_i , where

$$\Gamma_i(\mathbf{x}) := \|\mathbf{p}_i - \mathbf{x}\|^2 - w_i.$$

A weighted point B_i (or the corresponding sphere) is *orthogonal* to another weighted point B_j (or sphere) if their power distance is zero. Without real loss of generality, we can assume that the weighted points are in general position. This means that no $d + 2$ weighted points are orthogonal to the same sphere.

The power cell of B_i is the locus of points whose power distance wrt. B_i is less than its power distance wrt. any other site $B_j \in \mathcal{B}$:

$$P(B_i) = \{\mathbf{p} \in \mathbb{R}^d : \Gamma_i(\mathbf{p}) \leq \Gamma_j(\mathbf{p}) \quad \forall j = 1, \dots, n\}.$$

Let the bisector hyperplane of two weighted points B_i and B_j be the hyperplane $\mathcal{H}_{ij} = \{\mathbf{p} \in \mathbb{R}^d : \Gamma_i(\mathbf{p}) = \Gamma_j(\mathbf{p})\}$, and \mathcal{H}_{ij}^i be the halfspace limited by \mathcal{H}_{ij} such that for all $\mathbf{p} \in \mathcal{H}_{ij}^i : \Gamma_i(\mathbf{p}) \leq \Gamma_j(\mathbf{p})$. Again, $P(B_i)$ is the intersection of the halfspaces $\mathcal{H}_{ij}^i, j \neq i$, but, in the weighted case, $P(B_i)$ might be empty and must not necessarily contain \mathbf{p}_i . However, if it is not empty, it is a convex polytope and one can show that the power cells and their faces form a cell complex.

Definition 2.2 *The cell complex of power cells $P(B_i), i = 1, \dots, n$, is called the power diagram of \mathcal{B} or $\text{Pow}(\mathcal{B})$.*

The dual of the power diagram is called the *regular triangulation* or $\text{Reg}(\mathcal{B})$. The vertices of $\text{Reg}(\mathcal{B})$ belong to $\{\mathbf{p}_1, \dots, \mathbf{p}_n\}$ and, under the general position assumption, the d -dimensional faces of $\text{Reg}(\mathcal{B})$ are simplices and $\text{Reg}(\mathcal{B})$ is a simplicial complex. For $k \leq d$, each k -face $F(B_{i_0}, \dots, B_{i_k})$ of $\text{Reg}(\mathcal{B})$ (with vertices $\mathbf{p}_{i_0}, \dots, \mathbf{p}_{i_k}$) is dual to a $(d - k)$ -face of $\text{Pow}(\mathcal{B})$ that is exactly the intersection of the power cells of the $k + 1$ corresponding sites, and, consequently, the intersection of the bisector hyperplanes $\mathcal{H}_{ij}, i, j \in \{i_0, \dots, i_k\}$. We define

$$P(B_{i_0}, \dots, B_{i_k}) = \bigcap_{j=0 \dots k} P(B_{i_j}), 0 \leq k \leq d,$$

and we say that B_{i_0}, \dots, B_{i_k} are the *generators* of $P(B_{i_0}, \dots, B_{i_k})$. In particular, the $(d - 1)$ -dimensional power face dual to an edge of $\text{Reg}(\mathcal{B})$ is contained in the bisector hyperplane of its generators.

A sphere that is orthogonal to $k + 1$ weighted points is called the *orthosphere* of the $k + 1$ sites. It is *empty* if it has positive or zero power distance to all sites of \mathcal{B} . For any k -face of $\text{Reg}(\mathcal{B})$ there exist an empty orthosphere of its $k + 1$ generators. Its center lies on the dual power face. This property is equivalent to the empty-sphere property of the Delaunay triangulation. We say that a weighted point is *in conflict with an orthosphere* if its power distance to the orthosphere is negative.

The *regular orthosphere* of a d -face of $\text{Reg}(\mathcal{B})$ is the empty orthosphere of its $d + 1$ generators. Its center is a vertex of $\text{Pow}(\mathcal{B})$. If a weighted point X is in conflict with a regular orthosphere, we may also say that it is *in conflict with the vertex* of $\text{Pow}(\mathcal{B})$ that is its center.

The set of regular orthospheres is called \mathcal{O} . The *arrangement of regular orthospheres* is the generalized concept of the arrangement of Delaunay spheres. We say that a weighted point X lies in a cell of the arrangement if it is in conflict with the regular orthospheres that cover the cell. Thus, at the insertion of X , i.e. when X is added to \mathcal{B} , all faces of $\text{Reg}(\mathcal{B})$ whose regular orthospheres are in conflict with X are destroyed and also all vertices of $\text{Pow}(\mathcal{B})$ that are the centers of these regular orthospheres. However, the arrangement of regular orthospheres is better defined by its equivalent in a dual transform called the space of spheres which is introduced in the next section. See e.g. [26] for further properties of power diagrams and regular triangulations.

2.3 Space of spheres

It is well known that a power diagram can be obtained as a projection of the facets of a $(d + 1)$ -dimensional polytope onto \mathbb{R}^d . See e.g. [26, chapter 17] and [39]. This polytope is defined by the well known transform from spheres of \mathbb{R}^d to hyperplanes in \mathbb{R}^{d+1} . In the sequel, the $(d + 1)$ th coordinate axis is called *vertical axis*, and a point is above (or below) another point if its $(d + 1)$ th coordinate is bigger (or smaller) than the other's. The *vertical projection* is the projection onto the hyperplane $x_{d+1} = 0$.

Let $B_i = (\mathbf{p}_i, w_i)$ be the sphere (or the weighted point as described in the preceding section) of equation

$$(\mathbf{p}_i - \mathbf{x})^t(\mathbf{p}_i - \mathbf{x}) - w_i = 0.$$

Its power distance to the origin $\mathbf{0}$ is $\Gamma_i(\mathbf{0}) = \mathbf{p}_i^t \mathbf{p}_i - w_i$. The lifting of B_i to the space of spheres is defined by the bijection

$$\phi : B_i = (\mathbf{p}_i, w_i) \in \mathbb{R}^d \times \mathbb{R} \longrightarrow \phi(B_i) = (\mathbf{p}_i, \Gamma_i(\mathbf{0})) \in \mathbb{R}^{d+1}.$$

We consider the image by ϕ of all spheres orthogonal to B_i . Let $X = (\mathbf{x}, w_x)$ be such a sphere and Γ_x the power distance wrt. X . The center \mathbf{x} and the squared radius w_x must satisfy the following equation:

$$(\mathbf{x} - \mathbf{p}_i)^t(\mathbf{x} - \mathbf{p}_i) - w_x - w_i = 0 \iff \mathbf{x}^t \mathbf{x} - w_x = 2 \mathbf{x}^t \mathbf{p}_i - (\mathbf{p}_i^t \mathbf{p}_i - w_i) \iff \Gamma_x(\mathbf{0}) = 2 \mathbf{x}^t \mathbf{p}_i - \Gamma_i(\mathbf{0}).$$

Since the $(d + 1)$ th coordinate of $\phi(X)$ is $\Gamma_x(\mathbf{0})$, we deduce that the spheres orthogonal to B_i are mapped by ϕ to the $(d + 1)$ -dimensional hyperplane H_i of equation

$$x_{d+1} = 2 \mathbf{x}^t \mathbf{p}_i - \Gamma_i(\mathbf{0}).$$

We argue in the sequel that H_i is the hyperplane *dual* to $\phi(B_i)$ in the space of spheres. Thus, it is also called $\phi(B_i)^*$. By H_i^+ (equiv. $\phi(B_i)^{*+}$) we refer to the positive halfspace delimited by H_i , $x_{d+1} \geq 2 \mathbf{p}_i^t \mathbf{x} - \Gamma_i(\mathbf{0})$. Respectively, by H_i^- (equiv. $\phi(B_i)^{* -}$), we refer to the negative halfspace.

A *duality* is any bijection that reverses inclusion relationships. The bijection from $\phi(B_i)$ to H_i is, indeed, a duality because

$$\begin{aligned}\phi(B_i) \in H_j^+ &\iff \Gamma_i(\mathbf{0}) \geq 2\mathbf{p}_i^t \mathbf{p}_j - \Gamma_j(\mathbf{0}) \iff \Gamma_j(\mathbf{0}) \geq 2\mathbf{p}_j^t \mathbf{p}_i - \Gamma_i(\mathbf{0}) \iff \phi(B_j) \in H_i^+ \\ \phi(B_i) \in H_j^- &\iff \Gamma_i(\mathbf{0}) \leq 2\mathbf{p}_i^t \mathbf{p}_j - \Gamma_j(\mathbf{0}) \iff \Gamma_j(\mathbf{0}) \leq 2\mathbf{p}_j^t \mathbf{p}_i - \Gamma_i(\mathbf{0}) \iff \phi(B_j) \in H_i^-\end{aligned}$$

Notice that, unless $w_i = 0$, B_i is not orthogonal to itself, hence, $\phi(B_i) \notin \phi(B_i)^*$. Instead, the conjugate weighted point (or sphere) $\overline{B}_i := (\mathbf{p}_i, -w_i)$ is orthogonal to B_i and $\phi(\overline{B}_i)$ lies in $\phi(B_i)^*$.

An unweighted point or sphere with zero radius is mapped by ϕ to the unit d -paraboloid Ψ with axis x_{d+1} : $\phi((\mathbf{x}, 0)) = (\mathbf{x}, \mathbf{x}^t \mathbf{x})$. Its dual is the tangent plane to Ψ at $(\mathbf{x}, \mathbf{x}^t \mathbf{x})$. In general, the image by Φ of weighted points with the same weight w is the paraboloid $\Psi(w)$. $\Psi(w)$ is the vertical translation of Ψ such that $\Psi(w)$ has the same axis than Ψ but is centered on a point $(\mathbf{0}, -w)$.¹ Indeed, $\phi((\mathbf{x}, w)) = (\mathbf{x}, \mathbf{x}^t \mathbf{x} - w)$. The dual hyperplane $\phi^*(X)$ is tangent to $\Psi(-w)$ at $\phi((\mathbf{x}, -w))$. The normal of $\phi^*(X)$ is independent of the weight: it is the normal of Ψ at $(\mathbf{x}, \mathbf{x}^t \mathbf{x})$.




Figure 2.3: (a) $\phi(B_i)^* \cap \Psi$ projects onto the sphere B_i (b) $\phi(B_i)^* \cap \phi(B_j)^*$ projects onto the bisector hyperplane \mathcal{H}_{ij} .

The intersection of $\phi(B_i)^*$ with the paraboloid Ψ vertically projects onto the sphere B_i :

$$(\mathbf{p}_i - \mathbf{x})^t (\mathbf{p}_i - \mathbf{x}) - w_i = 0 \iff \mathbf{x}^t \mathbf{x} = 2\mathbf{x}^t \mathbf{p}_i - \Gamma_i(\mathbf{0}) \iff \phi((\mathbf{x}, 0)) \in \phi(B_i)^*.$$

See Figure 2.3 (a). A point \mathbf{x} is inside the sphere B_i if $\phi((\mathbf{x}, 0))$ is below $\phi(B_i)^*$ and outside B_i if $\phi((\mathbf{x}, 0))$ lies above $\phi(B_i)^*$. The analogue works for the power distance: a weighted point X has negative power distance to B_i if $\phi(X)$ lies below $\phi(B_i)^*$ and it has positive power distance to B_i if $\phi(X)$ is above

¹ $\mathbf{0}$ denotes the origin in \mathbb{R}^d

$\phi(B_i)^*$. More precisely, the signed vertical distance of $\phi(X)$ to $\phi(B_i)^*$ is, in fact, the power distance $\Gamma(B_i, X)$: since the $(d + 1)$ th coordinate of $\phi(X)$ is $\Gamma_x(\mathbf{0})$ and the vertical line through \mathbf{x} intersects $\phi(B_i)^*$ at $x_{d+1} = 2 \mathbf{x}^t \mathbf{p}_i - \Gamma_i(\mathbf{0})$, we get

$$\begin{aligned} \Gamma_x(\mathbf{0}) - (2 \mathbf{x}^t \mathbf{p}_i - \Gamma_i(\mathbf{0})) &= \mathbf{x}^t \mathbf{x} - w_x - 2 \mathbf{x}^t \mathbf{p}_i + \mathbf{p}_i^t \mathbf{p}_i - w_i \\ &= (\mathbf{p}_i - \mathbf{x})^t (\mathbf{p}_i - \mathbf{x}) - w_i - w_x = \Gamma(B_i, X). \end{aligned}$$

Consequently, if a point \mathbf{x} has less power wrt. B_i than wrt. B_j then the vertical line through \mathbf{x} intersects $\phi(B_i)^*$ above $\phi(B_j)^*$. The $(d - 1)$ -dimensional bisector hyperplane \mathcal{H}_{ij} is, therefore, the vertical projection of the intersection of $\phi(B_i)^*$ and $\phi(B_j)^*$. See Figure 2.3(b).

We define the following polytope:

$$\mathcal{P}(\mathcal{B}) := H_1^+ \cap \dots \cap H_n^+ \quad (2.1)$$

where H_i^+ is the positive halfspace delimited by $\phi(B_i)^*$. One can see that the faces of $\mathcal{P}(\mathcal{B})$ vertically project onto $\text{Pow}(\mathcal{B})$: a point of the facet F_i of $\mathcal{P}(\mathcal{B})$ carried by the hyperplane $\phi(B_i)^*$ lies above all other hyperplanes $\phi(B_j)^*$, so, its vertical projection has less power wrt. B_i than wrt. any other site B_j . See Figure 2.4.



Figure 2.4: $\text{Pow}(\mathcal{B})$ is the vertical projection of the faces of $\mathcal{P}(\mathcal{B})$.

The empty-sphere property of the regular triangulation translates as follows in the space of spheres: the regular orthosphere of a simplex of the triangulation is orthogonal to the sites of the simplex and has positive power distance to all other sites of \mathcal{B} . This means that its dual hyperplane in the space of spheres

contains the image by ϕ of its generators and is below the image by ϕ of all other sites. More precisely, let $O = (\mathbf{o}, w_o)$ be orthogonal to B_1, \dots, B_{d+1} , H_o be the dual of $\phi(O)$, and, H_o^+ the positive halfspace defined by H_o . $F(B_1, \dots, B_{d+1})$ is a face of $\text{Reg}(\mathcal{B})$ if and only if

$$\phi(B_i) \in H_o, \quad i = 1, \dots, d+1, \quad \text{and} \quad \phi(B_j) \in H_o^+, \quad j = d+2, \dots, n.$$

We deduce that the lower convex hull of $\{\phi(B_i), i = 1, \dots, n\}$ projects vertically onto $\text{Reg}(\mathcal{B})$. Refer to Figure 2.5. The general position assumption is satisfied if no $d+2$ points of $\phi(\mathcal{B})$ lie in the same d -dimensional hyperplane.




Figure 2.5: The lower convex hull of $\phi(\mathcal{B})$ projects onto $\text{Reg}(\mathcal{B})$.

The arrangement of regular orthospheres corresponds, in the space of spheres, to the arrangement of hyperplanes that are dual to the image of the regular orthospheres $O \in \mathcal{O}$. We call it $\text{arr}(\phi(\mathcal{O})^*)$. Of course, the hyperplanes of $\text{arr}(\phi(\mathcal{O})^*)$ are equivalently defined by the simplices of $\text{Reg}(\mathcal{B})$: for each simplex there exist a hyperplane in the arrangement that contains the image of the simplex's generators. If O is the simplex's regular orthosphere, the hyperplane is $\phi^*(O)$. All points within a cell of this arrangement of hyperplanes have in common that their preimage by ϕ is in conflict with the same regular orthospheres, and the same simplices of $\text{Reg}(\mathcal{B})$ are destroyed at their insertion into $\text{Reg}(\mathcal{B})$.

Expression for a vertex of the power diagram

A vertex of the power diagram is a rational functions of its generators. Re-call that $P(B_{i_1}, \dots, B_{i_{d+1}})$ is the vertex of the power diagram dual to the face $F(B_{i_1}, \dots, B_{i_{d+1}})$ of $\text{Reg}(\mathcal{B})$. It is the center of the regular orthosphere O of this face. If $\phi^*(O)$ is the dual hyperplane of $\phi(O)$, it must contain $\phi(B_{i_1}), \phi(B_{i_2}), \dots, \phi(B_{i_{d+1}})$.

Thus, we know the equation of $\phi^*(O)$ and we can determine the coordinates of $\phi(O)$ from it. The vertex of the power diagram is the projection of $\phi(O)$ on $x_{d+1} = 0$. We develop the expression only for $d = 2$ because this facilitates the exposition. The same approach applies in any dimension.

In the two-dimensional case with $B_1 = ((x_1, y_1), w_1)$, $B_2 = ((x_2, y_2), w_2)$, $B_3 = ((x_3, y_3), w_3)$, and $(B_i)^2 := x_i^2 + y_i^2 - w_i$, the equation of $\phi^*(O)$ is

$$\begin{vmatrix} 1 & 1 & 1 & 1 \\ x_1 & x_2 & x_3 & x \\ y_1 & y_2 & y_3 & y \\ (B_1)^2 & (B_2)^2 & (B_3)^2 & z \end{vmatrix} = 0 \quad (2.2)$$

On the other hand, taking $\phi(O) = (x_o, y_o, z_o)$, we know that $\phi^*(O)$ has the equation

$$z = 2(xx_o + yy_o) - z_o \iff xx_o + yy_o - \frac{1}{2}z - \frac{1}{2}z_o = 0.$$

We develop Equation (2.2) and identify the terms. In particular, we need to obtain the factor $\frac{1}{2}$ for the z -coordinate in Equation (2.3), thus, we divide Equation (2.2) by $2D$ with

$$D := \begin{vmatrix} 1 & 1 & 1 \\ x_1 & x_2 & x_3 \\ y_1 & y_2 & y_3 \end{vmatrix}.$$

Identifying the factors of the x - and the y - coordinate, we find

$$P(B_1, B_2, B_3) = (x_o, y_o) = \frac{1}{2D} \left(- \begin{vmatrix} 1 & 1 & 1 \\ y_1 & y_2 & y_3 \\ (B_1)^2 & (B_2)^2 & (B_3)^2 \end{vmatrix}, \begin{vmatrix} 1 & 1 & 1 \\ x_1 & x_2 & x_3 \\ (B_1)^2 & (B_2)^2 & (B_3)^2 \end{vmatrix} \right). \quad (2.3)$$

2.4 Power diagrams and sections of Voronoi diagrams

The following lemma proves that the intersection of a d -dimensional power diagram with a k -flat \mathcal{H} of \mathbb{R}^d with $k < d$ is a power diagram of dimension k . A k -flat of \mathbb{R}^d is a k -dimensional affine subspace of \mathbb{R}^d .

We say that \mathcal{H} is in *general position* if the intersection of an $(d - l)$ -dimensional face of the power diagram, $l \leq k$, with \mathcal{H} is either empty or has dimension $k - l$, and if the intersection of \mathcal{H} with a $(d - l)$ -dimensional face of the power diagram with $l > k$ is always empty. Notice, that a small perturbation of \mathcal{H} always removes a degenerate position. Without loss of generality, we suppose the general position of \mathcal{H} in order to ensure that the general position assumption for the resulting k -dimensional power diagram is satisfied.

Lemma 2.3 *The intersection of the power diagram $\text{Pow}(\mathcal{B})$ with a k -flat \mathcal{H} of \mathbb{R}^d in general position is a k -dimensional power diagram $\overline{\text{Pow}}(\overline{\mathcal{B}})$ defined in \mathcal{H} where $\overline{\mathcal{B}} = \{\overline{B}_i = (\overline{\mathbf{p}}_i, \overline{w}_i), i = 1, \dots, n\}$ such that $\overline{\mathbf{p}}_i$ is the orthogonal projection of \mathbf{p}_i onto \mathcal{H} and $\overline{w}_i = w_i - \|\mathbf{p}_i - \overline{\mathbf{p}}_i\|^2$ for each $B_i \in \mathcal{B}$.*

Proof: Because $\overline{\mathbf{p}}_i$ is the orthogonal projection of \mathbf{p}_i onto \mathcal{H} , we know that

$$\forall \mathbf{x} \in \mathcal{H} : \|\mathbf{x} - \mathbf{p}_i\|^2 = \|\mathbf{x} - \overline{\mathbf{p}}_i\|^2 + \|\mathbf{p}_i - \overline{\mathbf{p}}_i\|^2.$$

Consequently, if $\overline{\Gamma}_i(\mathbf{x})$ is the power of \mathbf{x} wrt. the site $\overline{B}_i = (\overline{\mathbf{p}}_i, w_i - \|\mathbf{p}_i - \overline{\mathbf{p}}_i\|^2)^2$,

$$\begin{aligned} \forall \mathbf{x} \in \mathcal{H} : \Gamma_i(\mathbf{x}) \leq \Gamma_j(\mathbf{x}) &\iff \|\mathbf{x} - \overline{\mathbf{p}}_i\|^2 + \|\mathbf{p}_i - \overline{\mathbf{p}}_i\|^2 - w_i \leq \|\mathbf{x} - \overline{\mathbf{p}}_j\|^2 + \|\mathbf{p}_j - \overline{\mathbf{p}}_j\|^2 - w_j \\ &\iff \overline{\Gamma}_i(\mathbf{x}) \leq \overline{\Gamma}_j(\mathbf{x}). \quad \square \end{aligned}$$

The preceding lemma implies in particular that the intersection of a Voronoi diagram with \mathcal{H} is a power diagram. Let the Voronoi diagram $\text{Vor}(\mathcal{P})$ be defined as before and $\overline{\text{Vor}}(\mathcal{P})$ be the intersection of the Voronoi diagram with \mathcal{H} , $\overline{\text{Vor}}(\mathcal{P}) := \text{Vor}(\mathcal{P}) \cap \mathcal{H}$. For all $\mathbf{p}_i \in \mathcal{P}$, $\overline{V}(\mathbf{p}_i) := V(\mathbf{p}_i) \cap \mathcal{H}$.

Observation 2.4 *$\overline{\text{Vor}}(\mathcal{P})$ is the power diagram of the points $\overline{\mathbf{p}}_i$ that are the projection of the sample points $\mathbf{p}_i \in \mathcal{P}$ onto \mathcal{H} weighted with $w_i = -\|\mathbf{p}_i - \overline{\mathbf{p}}_i\|^2$.*

²The notation is the same as for conjugate spheres. However, we always mention specifically when the conjugate sphere is referred to.

Chapter 3

Natural neighbor coordinates in power diagrams

3.1 Introduction

Natural neighbor interpolation has been introduced by Sibson [91] to interpolate multivariate scattered data. Given a set of points $\mathcal{P} = \{\mathbf{p}_1, \dots, \mathbf{p}_n\}$, the natural neighbor coordinate system associated to \mathcal{P} is defined from the Voronoi diagram of \mathcal{P} . Various papers ([90], [56], [84], [29],[71]) show that it satisfies the following definition by Brown [29].

Definition 3.1 ([29]) A system of coordinates over $\mathcal{U} \subseteq \mathbb{R}^d$ associated with \mathcal{P} is a set of continuous functions $s_i : \mathcal{U} \rightarrow \mathbb{R}, i = 1, \dots, n$, such that for all $\mathbf{x} \in \mathcal{U}$,

(i) $\mathbf{x} = \sum_{i=1}^n s_i(\mathbf{x})\mathbf{p}_i$ (barycentric coordinate property or BCP).

(ii) For any $i \leq n, s_i(\mathbf{p}_j) = \delta_{ij}$, where δ_{ij} is the Kronecker symbol.

(iii) $\sum_{i=1}^n s_i(\mathbf{x}) = 1$ (partition of unity property).

In our context, it is useful to enlarge this definition. We say that the coordinate system is *convex* if $s_i(\mathbf{x}) \geq 0$ for all $\mathbf{x} \in \mathcal{U}$, otherwise it is not convex. Let the *weighted barycenter* of a point $\mathbf{x} \in \mathcal{U}$ be

$$\mathbf{b}(\mathbf{x}) = \sum_{i=1}^n s_i(\mathbf{x})\mathbf{p}_i.$$

With the barycentric coordinate property, $\mathbf{x} = \mathbf{b}(\mathbf{x})$. If the barycentric coordinate property is not satisfied but $e(\mathbf{x}) = \|\mathbf{x} - \mathbf{b}(\mathbf{x})\|$ is small, we say that we are concerned with a coordinate system that is *almost barycentric*. $e(\mathbf{x})$ is the error concerning the barycentric coordinate property.

The coordinate system is *local* if the support of each coordinate function s_i is contained in a small topological ball centered on \mathbf{p}_i .¹ We will see in this chapter that the natural neighbor coordinate system

¹See Section 5.1 for the definition of a topological ball.

defined in a power diagram is a convex barycentric coordinate system. It is local whenever the weights of the points are bounded.

Related Work

Sibson introduces the natural neighbor coordinate system in [90] and proves the barycentric coordinate property in a generalized manner. (He calls this property the *local coordinate property*.) In a second paper, he describes its application to scattered data interpolation over a two-dimensional domain [91]. He claims that the coordinate function is continuously differentiable except on the data sites but the proof is omitted. He gives an algorithm to compute the natural neighbor coordinates in two dimensions. In 1988, Aurenhammer enlarges the definition of natural neighbor coordinates to power diagrams and subsets of several points in order to define various linear combinations among points of the defining set \mathcal{P} [12]. Later, he shows the relationship between natural neighbor coordinates and the Gale transform [13]. Notably, he proves that the so-called *power matrix* represents a Gale transform of \mathcal{P} . This matrix contains the coordinate $s_i(\mathbf{p}_j)$ of the diagram for $\mathcal{P} \setminus \{\mathbf{p}_j\}$ for each point \mathbf{p}_j in the interior of the convex hull of \mathcal{P} .

Farin [56] discusses the question of support, differentiability and other properties of the natural neighbor coordinate function in more detail than Sibson. Farin's paper refers also only to the two-dimensional case.

The smoothness properties of the coordinate function in arbitrary dimension have been properly proven by Piper [84]. It is shown that the coordinate function is continuous on the convex hull $\text{conv}(\mathcal{P})$ of \mathcal{P} and continuously differentiable on $\text{conv}(\mathcal{P})$ except at the data points \mathcal{P} themselves. Fixing the point \mathbf{x} , the coordinate function $s_i(\mathbf{x})$ is also continuously differentiable with respect to a data point $\mathbf{p}_j \in \mathcal{P}$. Furthermore, Piper shows a formula for the gradient of the coordinate function on $\text{conv}(\mathcal{P}) \setminus \mathcal{P}$. The gradient formula gives rise to another proof of the barycentric coordinate property.

Brown [29] pursues a more algebraic than geometric approach: a general method to produce barycentric coordinate systems is defined which is based on the weighted combination of barycentric coordinates with respect to a triangle. The proofs of the major properties of Definition 3.1 become immediate. As an example, he gives the weight function leading to Sibson's coordinates. The work is written for the two-dimensional case but generalizable to higher dimensions.

The most recent work on natural neighbor coordinates by Hiyoshi and Sugihara [71] spent some effort to smoothen the coordinate function not on the data points but on the Delaunay spheres. In fact, as observed by Farin, the coordinate function is C^∞ in $\text{conv}(\mathcal{P})$ except on the Delaunay spheres where it is only C^1 . In [71], a generalized coordinate systems is defined with arbitrary smoothness on the Delaunay spheres except on the data points. This work leads to yet another proof of the barycentric coordinate property. The same idea would be generalizable to natural neighbor coordinates in power diagrams, however, we do not further pursue this direction.

Last but not least, Clarkson proposes the construction of a new system of coordinates that is based on

natural neighbor coordinates. It is an unpublished idea that he describes in a talk [35]. The new coordinate system fulfills the barycentric coordinate property, i.e. a point is the linear combination of its neighbors, but in addition, the squares of the point is the linear combination of the square of its neighbors. This property is interesting in scattered data interpolation. However, Clarkson's coordinate system is not convex which means that the coordinate functions may have negative values.

Outline of the chapter

In the sequel, we define a natural neighbor coordinate system from power diagrams. Of course, the definition also applies to Voronoi diagrams which are special power diagrams – namely, when all sites have equal weight. After the definition of the natural neighbor coordinates, we prove some of their basic properties, notably, the properties of Definition 3.1, the continuity properties and the gradient of the coordinate function.

Aurenhammer has defined the generalized natural neighbor coordinates in [12] and he gives a proof for the barycentric coordinate property. The contribution of this thesis is to develop further properties of the coordinate functions by generalizing notably the work of Piper [84]. We also present a new proof of the barycentric coordinate property.

3.2 Definition of the coordinate function and basic properties

Let \mathcal{B} be a set of n weighted points. The natural neighbor coordinate function of a weighted point X with respect to a site $B_i \in \mathcal{B}$ is defined as the proportion of the power cell of X that has been 'stolen' from the site B_i at the insertion of X into the power diagram $\text{Pow}(\mathcal{B})$. The natural neighbors are the sites that have passed a part of their cells to the cell of X . For this to be well defined, the cell of X has to be neither empty nor infinite. Figure 3.1 shows an example of natural neighbors in a Voronoi diagram: the cell of X is highlighted, and its partition with respect to the Voronoi diagram without X is indicated by dotted lines.

The major difference between the coordinates in Voronoi diagrams and in power diagrams is due to the fact that a site of a power diagram might have an empty cell. This cannot happen in Voronoi diagrams. It implies that the cell of some sites may disappear when inserting a new site X . However, this does not change the definition of the coordinate function. Simply, X 'stole' the entire cell of the site B_i at its insertion. In the sequel, we give a more formal definition. Let $X = (\mathbf{x}, w_x) \in \mathbb{R}^d \times \mathbb{R}$. We define $\text{Pow}^+(\mathcal{B}) := \text{Pow}(\mathcal{B} \cup \{X\})$ and $\text{Reg}^+(\mathcal{B}) := \text{Reg}(\mathcal{B} \cup \{X\})$. ($\text{Pow}(\mathcal{B})$ and $\text{Reg}(\mathcal{B})$ are defined in Section 2.2.) $P^+(X)$ is the power cell of X in $\text{Pow}^+(\mathcal{B})$ and the subcell $P_i(X) := P^+(X) \cap P(B_i)$ is the part of the power cell $P^+(X)$ that has been 'stolen' from $P(B_i)$ at the insertion of X . Remark that the cell of B_i can be entirely contained in the cell of X , so that $P_i(X) = P(B_i)$. The volume of $P_i(X)$ is denoted by $\pi_i(X)$ and the volume of $P(X)$ is denoted by $\pi(X)$. The $(d - k)$ -faces of $P^+(X)$ are called $P^+(X, B_{i_1}, \dots, B_{i_k})$. The definition domain of the natural coordinate functions is the subset




Figure 3.1: x has five natural neighbors p_1, \dots, p_5 .

$\mathcal{U} = \{X \in \mathbb{R}^d \times \mathbb{R} \mid \pi(X) \neq \infty \wedge \pi(X) \neq 0\}$ of $\mathbb{R}^d \times \mathbb{R}$. Further in this section, we explain its geometric interpretation in the space of spheres.

We define the natural neighbor coordinate functions formally and prove later in this section that they fulfill Definition 3.1.

Definition 3.2 *In the natural neighbor coordinate system associated to a set of weighed points \mathcal{B} , the (natural neighbor) coordinate function $\lambda_i(X), i = 1, \dots, n$, of a point $X \in \mathcal{U} \setminus \mathcal{B}$ is $\lambda_i(X) = \frac{\pi_i(X)}{\pi(X)}$ and $\lambda_i(B_j) = \delta_{ij}$ for all $B_j \in \mathcal{B}$, where δ_{ij} is the Kronecker symbol.*

Definition 3.3 *Given a set of weighted points \mathcal{B} , the natural neighbors of a point $X \in \mathcal{U}$ are the sites $B_i \in \mathcal{B}$ with $\lambda_i(X) \neq 0, i = 1, \dots, n$.*

In Voronoi diagrams, the natural neighbors of x are exactly the neighbors of x in $\text{Del}^+(\mathcal{P})$, i.e. vertices connected to x by an edge of $\text{Del}^+(\mathcal{P})$. In power diagrams, the natural neighbors of X include not only the neighbors of X in $\text{Reg}^+(\mathcal{B})$ but also the vertices that disappear from $\text{Reg}(\mathcal{B})$ when inserting X .

The support of the natural neighbor coordinate λ_i is $\Delta_i = \{X \in \mathcal{U} \mid \lambda_i(X) \neq 0\}$. It is the set of points that are in conflict with at least one vertex of the cell $P(B_i)$. Remember that this means that the power distance between X and the regular orthosphere centered on the vertex is negative. In this case, $\lambda_i(X) \neq 0$ because the vertex is 'stolen' from $P(B_i)$ at the insertion of X .

With respect to the space of spheres, the insertion of the point $X \in \mathcal{U}$ into $\text{Pow}(\mathcal{B})$ corresponds to inserting the hyperplane $H_x = \phi(X)^*$ into the space of spheres. The corresponding halfspace H_x^+ intersects the polytope $\mathcal{P}(\mathcal{B})$ "cutting off" the polytope \mathcal{P}_x which is the intersection of the halfspaces H_i^+ with the halfspace H_x^- described by $x_{d+1} \leq 2 \cdot \mathbf{x}^t \cdot \mathbf{x} - \Sigma_x(\mathbf{0})$. See Figure 3.2. Precisely,

$$\mathcal{P}_x := \bigcap_{i=1, \dots, n} H_i^+ \cap H_x^- \quad (3.1)$$

The polytope \mathcal{P}_x projects vertically onto $P^+(X)$ and its lower facets $F_i(X)$ project onto $P_i(X)$. The volume of $F_i(X)$ in the projection is $\pi_i(X)$ and the volume of the projection of the upper facet is $\pi(X)$. Hence, $\lambda_i(X)$ can alternatively be defined in the space of spheres. The cell of X is empty if and only if




Figure 3.2: (a vertical cut) (a) The polytope $\mathcal{P}_x \neq \emptyset$. \iff (b) X is in conflict with a cell of $\text{arr}(\phi(\mathcal{O})^*)$.

$\phi(X)$ is above the lower convex hull of $\phi(\mathcal{B})$. In this case, X is in conflict with no regular orthosphere. It is infinite if and only if the center x lies outside the convex hull of the centers p_i of the sites $B_i \in \mathcal{B}$. This is independent of the weight of X . Thus, X has a finite and non-empty cell iff $\phi(X)$ lies in the (unbounded) polytope whose upper facets are the lower convex hull of $\phi(B_i)$ and whose remaining facets are delimited by vertical rays from the lower convex hull to $x_{d+1} = -\infty$. This region is called $\phi(\mathcal{U})$ since \mathcal{U} is its preimage by ϕ . Refer to Figure 3.3 (a). The support Δ_i corresponds to the set of points $\phi(\Delta_i)$ that




Figure 3.3: (a vertical cut) (a) $\phi(\mathcal{U})$ is unbounded in direction $x_{d+1} = -\infty$. (b) The support $\phi(\Delta_i)$ is highlighted.

lie outside $\text{conv}(\phi(\mathcal{B}))$, the convex hull of the points $\phi(\mathcal{B})$, and that “see” $\phi(B_i)$ (in the sense that the line from the point to $\phi(B_i)$ does not intersect $\text{conv}(\phi(\mathcal{B}))$). In other words, the points of $\phi(\Delta_i)$ lie below at least one hyperplane that supports a facet of the lower convex hull of $\phi(\mathcal{B})$ incident to $\phi(B_i)$. The dual of such a hyperplane H_o is the image by ϕ of a regular orthosphere whose center is a vertex of $P(B_i)$. If $\phi(X)$ lies below H_o , X is in conflict with this vertex. See Figure 3.3(b).

To be more precise, we define the cone at $\phi(B_i)$ to be the intersection of the positive halfspaces defined by the hyperplanes that support the facets of the lower convex hull of $\phi(\mathcal{B})$ incident to $\phi(B_i)$,

$$\text{cone}(\phi(B_i)) := \cap \{H_o^+ \mid H_o \in \phi(\mathcal{O})^* \text{ and } \phi(B_i) \in H_o\}. \quad (3.2)$$

With this definition,

$$\phi(\Delta_i) = \mathcal{U} \setminus \text{cone}(\phi(B_i)).$$

An edge of $\text{cone}(\phi(B_i))$ is the line through $\phi(B_i)$ and $\phi(B_j)$ where B_j is a neighbor of B_i in $\text{Reg}(\mathcal{B})$.

The cell of B_i disappears at the insertion of X if $\phi(B_i)$ is not a vertex of the convex hull $\text{conv}(\phi(\mathcal{B}) \cup \phi(X))$. In this case, $\phi(X)$ lies below all hyperplanes of $\phi(\mathcal{O})^*$ containing $\phi(B_i)$. This part of $\phi(\Delta_i)$ is indicated by the darkest region in Figure 3.3(b).

Notice what happens as X changes: Changing the weight of X induces a vertical movement of $\phi(X)$ and of the hyperplane H_x : if the weight grows, $\phi(X)$ moves downwards while H_x moves upwards (and the opposite if the weight diminishes). The normal of H_x is independent of the weight. Translating the center x by some horizontal vector \vec{v} , $\phi(X)$ stays on the paraboloid $\Psi(w_x)$ while the horizontal component of the movement is \vec{v} . See Figure 3.4.




Figure 3.4: Changing (a) the weight and (b) the center of X .

A finite cell $P(X)$ can never become infinite only by raising the weight w_x because the center x stays in the interior of the convex hull of the centers of the other sites. Consequently, the cell is finite.

The interpretation in the space of spheres makes it easy to see that the volume function π_i is continuous as long as $X \neq B_i$. Indeed, the facet $F_i(X)$ is contained in H_i and delimited by H_x and some hyperplanes $H_j, j = 1, \dots, n, j \neq i$,

$$F_i(X) = H_i \cap_{j \neq i} H_j^+ \cap H_x^-.$$

As H_x changes continuously with X , the volume of $F_i(X)$ (and, thus, $\pi_i(X)$) changes continuously as long as $H_x \neq H_i$, i.e. as long as $X \neq B_i$. With $\lambda_i(X) = \frac{\pi_i(X)}{\sum_{i=1}^n \pi_i(X)}$, we see that λ_i is continuous in \mathcal{U} as long as $X \neq B_i \in \mathcal{B}$. Before we show that λ_i is continuous even at $B_i \in \mathcal{B}$, we consider what happens as X approaches the boundary of \mathcal{U} and as $x \rightarrow \mathbf{p}_i$ with $w_x > w_i$.

Behavior of π_i at the border of the definition domain

As X approaches the boundary of \mathcal{U} , two situations might occur: either $\pi(X) \rightarrow 0$ or $\pi(X) \rightarrow \infty$.




Figure 3.5: (a vertical cut) $\pi(X) = 0$: (a) H_x does not intersect $\mathcal{P}(\mathcal{B})$, (b) $\phi(X)$ is on $\text{conv}(\phi(\mathcal{B}))$.




Figure 3.6: X approaches B_i while $w_x < w_i$. (a) $X \cap B_i \neq \emptyset$, (b) X is contained in B_i , (c) $\pi(X) = 0$.

In the first case, $\pi(X) \rightarrow 0$, this means in the space of spheres that $\phi(X)$ approaches the lower convex hull of $\phi(\mathcal{B})$ (Figure 3.5). In $\mathbb{R}^d \times \mathbb{R}$, this means that w_x diminishes and/or that the center of X moves towards the center of a site B_i that has bigger weight than X .

Figures 3.6, 3.7 and 3.8 depict the different power diagrams as X changes. X corresponds to the black circle and the cell of X is indicated by the dark (blue) lines.

In Figure 3.6, we show what happens in the power diagram as the center x approaches the center of B_i with bigger weight than X : the bisector hyperplane of B_i and X approaches the center p_i as long as B_i does not completely contain X (Figure 3.6(a)). When X is contained in B_i , the bisector hyperplane moves away from p_i (while $x \rightarrow p_i$) until the cell of X disappears, $\pi(X) = 0$. At $x = p_i$, the bisector hyperplane of X and B_i is at infinity. However, the volume functions $\pi(X)$ as well as $\pi_i(X)$ change continuously with X as long as $\pi(X) \neq 0$.

They are also continuous as $\pi(X) \rightarrow \infty$. Figure 3.7 shows how the cell of X becomes unbounded as x approaches the convex hull of $\{p_i | B_i \in \mathcal{B}\}$.




Figure 3.7: The cell of X becomes unbounded.

Behavior of λ_i and π_i as $x \rightarrow p_i$ with $w_x > w_i$

In a similar way $\lambda_i(X)$ is well defined and continuous when the center of X approaches the center of B_i while the weight of X is bigger than the weight of B_i . This situation is depicted in Figure 3.8. At some point, the cell of B_i is entirely contained in the cell of X , $P_i(X) = P(B_i)$. In this case, $\pi_i(X)$ is constant with respect to X and the coordinate function $\lambda_i(X) = \frac{\pi_i(X)}{\pi(X)}$ varies with $\pi(X)$. Although $\pi_i(X)$ is constant, $\lambda_i(X)$ is continuous with $\pi(X) = \sum_j \pi_j(X)$ as long as $X \neq B_i$ and $\pi(X) > 0$. Remember that, if $P_i(X) = P(B_i)$, in the space of spheres, $\phi(X)$ lies below all hyperplanes of $\phi(\mathcal{O})^*$ containing $\phi(B_i)$ as indicated by the darkest part in Figure 3.3(b).

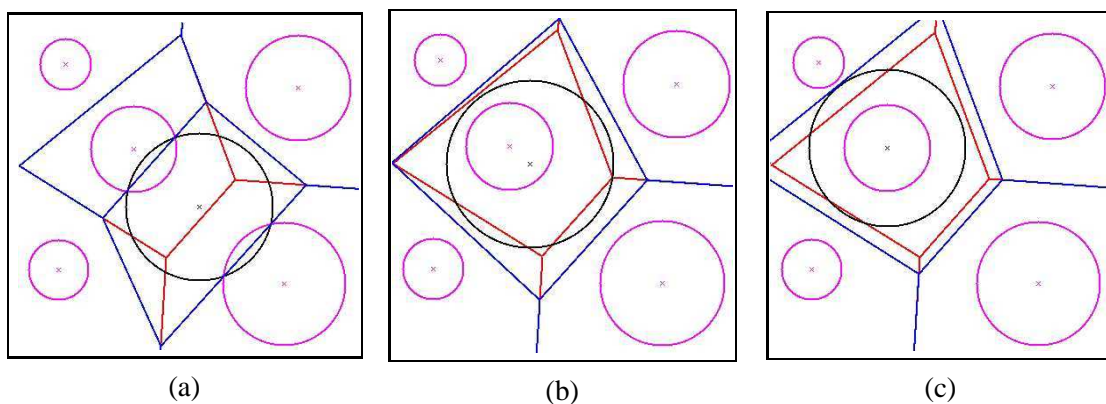


Figure 3.8: X approaches B_i while $w_x > w_i$. (a) $\pi_i(X) < \pi(B_i)$, (b) $\pi_i(X) = \pi(B_i)$, (c) $x = p_i$.

Verifying the conditions of Definition 3.1

As already mentioned, the natural neighbor coordinates satisfy the conditions of Definition 3.1.

Property 3.1(ii), namely $\lambda_i(B_j) = \delta_{ij}$ is satisfied by definition. Property 3.1(iii), namely $\sum_i \lambda_i(X) = 1$, follows immediately from the definition of the coordinate functions. Indeed, since the coordinate function is normalized by the total volume of the power cell, it naturally sums to one. Before we overview the different proofs of the barycentric coordinate property 3.1(i), we show that the coordinate functions are continuous. This is also a condition of Definition 3.1.

We argued before that the natural neighbor coordinate function is continuous in $\mathcal{U} \setminus \mathcal{B}$ since, elsewhere than on \mathcal{B} , the volume functions $\pi_i(X)$ are continuous. It remains to show the continuity of λ_i at $X = B_j$, i.e. $\lim_{X \rightarrow B_j} \lambda_i(X) = \delta_{ij}$.

As shown in Figure 3.9 and Figure 3.10, the volume $\pi(X)$ is not continuous at B_i unless $\mathbf{x} = \mathbf{p}_i$ as $X \rightarrow B_i$. Nevertheless, we can use an argument of extension by continuity: $\lim_{X \rightarrow B_i} \lambda_j(X) = \delta_{ij}$ because $\lim_{X \rightarrow B_i} \pi_j(X) = 0$ for all $j \neq i$ and $\sum_i \lambda_i(X) = \sum_i \frac{\pi_i(X)}{\pi(X)} = 1$ while $\pi(X)$ is finite. It follows that $\lim_{X \rightarrow B_i} \lambda_i(X) = 1$ and $\lim_{X \rightarrow B_i} \lambda_j(X) = 0$ for all $j \neq i$.




Figure 3.9: The limit as X approaches B_5 (in the power diagram), (a) the general case, (b) $\mathbf{x} = \mathbf{p}_5$.

Furthermore, the coordinate system is convex, i.e. $\forall X \in \mathcal{U}, \lambda_i(X) \geq 0, i = 1, \dots, n$, because the coordinate functions are defined by the ratio of two volumes. It is local whenever there is an upper bound for the weight. Otherwise, $\phi(\Delta_i)$ is unbounded (recall Figure 3.3).

Proofs for the barycentric coordinate property

The proof of Property 3.1(i), also called the barycentric coordinate property or BCP, requires more insight, and there exist several different proofs in the literature –most of them written for the unweighted case. We sketch the different approaches in this section –for the unweighted case. Then, we give, in detail, a proof of the BCP for the natural neighbor coordinate system defined in power diagrams. This proof is related to the space of spheres. (For Voronoi diagrams, we adopt the same notations as for power diagrams. The superscript $+$ indicates that we refer to the Voronoi diagram (resp. the Delaunay triangulation) of $\mathcal{P} \cup \{\mathbf{x}\}$.)




Figure 3.10: (vertical cut) The limit as X approaches B_5 (in the space of spheres), (a) the general case, (b) $\mathbf{x} = \mathbf{p}_5$.

The original proof by Sibson [90] is valid for a generalization of Voronoi diagrams, namely, a higher dimensional space is considered that embeds the d -dimensional space. The generalized d -dimensional Voronoi diagram is the intersection of an ordinary Voronoi diagram in the higher-dimensional space by the d -dimensional space. From Observation 2.4, we know that this generalized diagram is a power diagram. Sibson proves the case $d = 1$, and by an integration argument, the barycentric coordinate property for natural neighbor coordinates in this generalized diagram is shown. A very similar proof is given by Aurenhammer [12].

Piper proves the BCP by showing that the function $\mathcal{H}(X) = \sum_i \lambda_i(X)(\mathbf{p}_i - \mathbf{x})$ is constantly zero. He deduces from the expression for the gradient that the derivative of \mathcal{H} is zero. Then, he shows that the limit $\lim_{\mathbf{x} \rightarrow \mathbf{p}_i} \mathcal{H}(X) = 0$ and concludes the BCP.

Brown shows how to define Sibson's coordinate as weighted average of barycentric coordinates over some Delaunay triangles. The proof of the BCP becomes trivial.




Figure 3.11: The Laplace coordinate function $\lambda_i^0(\mathbf{x})$.

In [71], the proof of the BCP by Hiyoshi and Sugihara is based on a proof of the BCP for the Laplace

coordinate function defined as follows:

Definition 3.4 *In the Laplace coordinate system associated to \mathcal{P} , the (Laplace) coordinate function $\lambda_i^0(\mathbf{x})$ of a point $\mathbf{x} \in \text{conv}(\mathcal{P})$ is defined as*

$$\lambda_i^0(\mathbf{x}) = \frac{\frac{v_{ix}}{\|\mathbf{p}_i - \mathbf{x}\|}}{\sum_i \frac{v_{ix}}{\|\mathbf{p}_i - \mathbf{x}\|}}, i = 1, \dots, n.$$

where v_{ix} denotes the volume of the $(d - 1)$ -dimensional Voronoi face dual to the edge $F^+(\mathbf{p}_i, \mathbf{x})$.

The BCP for Laplace coordinates is easily shown with the help of a theorem by Minkowski stated as follows (a proof can be found in [66, chapter 15.3]):

Theorem 3.5 *The sum of the normal vectors of the facets of a convex polytope, where the length of the vector is the area of the corresponding facet, is zero.*

In our case, consider the Voronoi cell $V^+(\mathbf{x})$ which is always convex. The facet of $V^+(\mathbf{x})$ that is dual to $F^+(\mathbf{p}_i, \mathbf{x})$ is part of the bisector \mathcal{H}_{ix} between \mathbf{x} and \mathbf{p}_i . Of course, the unit normal of the facet is the vector $\frac{1}{\|\mathbf{p}_i - \mathbf{x}\|} \overrightarrow{\mathbf{x}\mathbf{p}_i}$, and the volume of the facet is v_{ix} . See also Figure 3.11. We get

$$\sum_i \frac{v_{ix}}{\|\mathbf{p}_i - \mathbf{x}\|} \overrightarrow{\mathbf{x}\mathbf{p}_i} = 0 = \sum_i \lambda_i^0(\mathbf{x}) \overrightarrow{\mathbf{x}\mathbf{p}_i}.$$

The authors observe that Sibson's coordinate function can be obtained as some integration of the Laplace coordinate. Since the BCP is valid during the process of integration, it is also valid for the result.

A proof of the BCP using Minkowski's theorem in the lifting map

There exists a direct proof of the BCP for power diagrams based on Minkowski's theorem applied to the polytope \mathcal{P}_x that is defined in the space of spheres. See also the beginning of the chapter and, in particular Equation (3.1).

Let $f_i(X)$ denote the volume of the face $F_i(X)$ of \mathcal{P}_x that is contained in $H_i = \phi(B_i)^*$ and f_x denotes the volume of the upper facet of \mathcal{P}_x (that is contained in H_x). $\vec{n}_i = (-2\mathbf{p}_i, 1)$ is the normal vector of H_i directed to the negative halfspace H_i^- and $\vec{n}_x = (-2\mathbf{x}, 1)$ is defined in the same manner for H_x . Minkowski's theorem applied to \mathcal{P}_x writes

$$\sum_i f_i(X) \frac{\vec{n}_i}{\|\vec{n}_i\|} - f_x \frac{\vec{n}_x}{\|\vec{n}_x\|} = 0. \quad (3.3)$$

Since the polytope \mathcal{P}_x projects vertically onto $P^+(X)$ and its lower facets $F_i(X)$ project onto $P_i(X)$, we deduce that

$$\pi_i(X) = f_i(X) \frac{\vec{n}_i}{\|\vec{n}_i\|} \cdot \vec{\mathbf{i}}_{d+1} = \frac{f_i(X)}{\|\vec{n}_i\|},$$

where $\vec{\mathbf{i}}_{d+1}$ denotes the unit vector $(\mathbf{0}, 1)$ along the vertical axis.² Similarly, $\pi(X) = f_x \frac{\vec{n}_x}{\|\vec{n}_x\|} \cdot \vec{\mathbf{i}}_{d+1} = f_x \frac{1}{\|\vec{n}_x\|}$. Equation 3.3 becomes

$$\sum_i \pi_i(X) \vec{n}_i - \pi(X) \vec{n}_x = 0.$$

Finally, by using the expression for the normals and by projecting onto $x_{d+1} = 0$, this yields the BCP:

$$\sum_i -\pi_i(X) \mathbf{p}_i + \pi(X) \mathbf{x} = 0.$$

3.3 Further continuity properties of the coordinate function

The following two sections treat essentially the generalization of Piper's work [84] about natural neighbor coordinates in Voronoi diagrams to natural neighbor coordinates in power diagrams. We assume \mathcal{B} to be given. We recall that $\mathcal{U} = \{X \in \mathbb{R}^d \times \mathbb{R} \mid \pi(X) \neq \infty \wedge \pi(X) \neq 0\}$.

Continuity properties with respect to X

Lemma 3.6 *The natural neighbor coordinate $\lambda_i(X)$ of X is C^0 continuous over \mathcal{U} and C^1 continuous over \mathcal{U} except at a finite set of lines in \mathcal{U} .*

Proof: As Farin noticed in [56], the natural neighbor coordinate in a Voronoi diagram is a piecewise rational function of \mathbf{x} as long as the set of natural neighbors of \mathbf{x} does not change. This is also the case in power diagrams. The coordinates of the vertices of the sub-cell $P_i(X)$ are rational functions of X . In the two-dimensional case, the coordinates of a vertex are given by Equation (2.3) with $B_1 = X$ and B_2 and B_3 are natural neighbors of X . Therefore, $\pi_i(X)$, which is the volume of $P_i(X)$, is a piecewise rational function of X . Since by assumption, the volume function $\pi(X) = \sum \pi_i(X)$ does not vanish, the differentiability of the normalized coordinate λ_i follows from the differentiability of the π_i .

In the space of spheres, the natural neighbors of X are the same as long as $\phi(X)$ remains in one cell of $\text{arr}(\phi(\mathcal{O})^*)$, the arrangement of hyperplanes which are dual to the image by ϕ of the regular orthospheres \mathcal{O} . This is explained in Section 2.3. The rational function $\pi_i(X)$ changes when $\phi(X)$ traverses the boundary of a cell of $\text{arr}(\phi(\mathcal{O})^*)$ – because either X has a new natural neighbor or it loses a neighbor. We consider the function $\pi_i \circ \phi^{-1}$ defined in the space of spheres. Because $\phi(X)$ is differentiable and bijective, π_i is differentiable if $\pi_i \circ \phi^{-1}$ is differentiable.

In order to show the differentiability of $\pi_i \circ \phi^{-1}$ on the boundary of the cells, it is sufficient to show, for each point on a hyperplane $H_o \in \phi(\mathcal{O})^*$, the differentiability of $\pi_i \circ \phi^{-1}$ restricted to a (one-dimensional) curve that crosses H_o transversally. Inside each cell of the arrangement, $\pi_i \circ \phi^{-1}$ is C^∞ continuous.

²We denote the scalar product of two vectors by $\vec{p} \cdot \vec{q}$.

Hence, by showing the existence of some cross boundary derivative, we establish the differentiability on the boundary.

In the unweighted case, Piper restricted π_i to rays emanating from \mathbf{p}_i because, on such a ray, π_i has a particularly simple form. The same approach works in power diagrams if the center of X is constrained to such a ray. With respect to the space of spheres, we may restrict $\pi_i \circ \phi^{-1}$ to an ascending ray whose source is on the vertical line through $\phi(B_i)$ below all hyperplanes of $\phi(\mathcal{C})^*$. Hence, the ray intersects any hyperplane transversally, and the derivative of the restriction of $\pi_i \circ \phi^{-1}$ to the ray is a valid cross boundary derivative. Refer to Figure 3.12. Observe that the source of the ray lies in the support $\phi(\Delta_i)$ of $\lambda_i \circ \phi^{-1}$




Figure 3.12: A ray with source $L(0)$ intersects any hyperplane H_o transversally. It is parameterized by distance to \mathbf{p}_i in the vertical projection.

because it lies below $\phi(B_i)$. (Re-call from Figure 3.3(b) that the support $\phi(\Delta_i) = \mathcal{U} \setminus \text{cone}(\phi(B_i))$.) For the sake of the proof, the ray should leave the support $\phi(\Delta_i)$ before it leaves the support $\phi(\mathcal{U})$. Notice that the source can be chosen as low as necessary for the rays to still intersect any hyperplane of $\phi(\mathcal{C})^*$.

Consider such a ray L in \mathbb{R}^{d+1} and suppose, at first, that L is not vertical. (This case is treated later.) We parameterize L by the distance to \mathbf{p}_i in the vertical projection, so that $L(t)$ projects vertically onto some point $l(t)$ with $\|\mathbf{p}_i - l(t)\| = t$ and $L(0)$ projects vertically onto \mathbf{p}_i . Let $X(t) = \phi^{-1}(L(t))$. In the sequel, we consider the point $X = X(t)$ to be fixed.

Let L^* be the set of hyperplanes dual to the points of L . By duality, they intersect in a $(d - 2)$ -flat A . The hyperplane $H_i = \phi(B_i)^*$ is a vertical translation of $L(0)^*$ because the dual points $\phi(B_i)$ and $L(0)$ differ only by their height. (Their centers project vertically onto \mathbf{p}_i (Fig. 3.12); refer also to Section 2.3.)

The intersection of H_i with any hyperplane of $L^*(u)$, $u > 0$, is, therefore, parallel to the $(d - 2)$ flat A . As u varies from 0 to ∞ , a part of the hyperplane H_i is swept by $L(u)^* \cap H_i$. Refer to Figure 3.13.




Figure 3.13: Let $A(u) = L(u)^* \cap H_i$. As $u \rightarrow \infty$, $L(u)^*$ approaches the vertical hyperplane containing A normal to the vertical projection of L (cf. Section 2.3 and Section 3.2).

Recall that F_i is the facet of the polytope $\mathcal{P}(\mathcal{B})$ (defined in Equation (2.1)) carried by the hyperplane H_i . $F_i(X)$ is the part of F_i that is “cut off” by the negative halfspace H_x^- delimited by $H_x = \phi(X)^*$. Thus, $F_i(X)$ is the facet of the polytope \mathcal{P}_x carried by H_i . (\mathcal{P}_x is defined in Equation (3.1). It is highlighted in Figure 3.14.) $F_i(X)$ projects vertically onto $P_i(X)$, thus, the volume $\pi_i(X)$ is alternatively defined by the volume of the projection of $F_i(X)$. Let $f_i(X)$ be the $(d$ -dimensional) volume of $F_i(X)$. It is clear that the differentiability of f_i implies the differentiability of π_i .

Consider Figure 3.14. By definition, with $X = X(t)$, $F_i(X) = F_i \cap L(t)^{* -}$. With the hyperplanes L^* , we can sweep the facet $F_i(X)$. Consider the $(d - 1)$ -dimensional volume $v_i(u)$ of $F_i \cap L(u)^*$. We obtain the volume of $F_i(X)$ by integrating $v_i(u)$ from $u = t$ to ∞ . There exists an interval (c_0, c_1) on L such that $v_i(u) \neq 0$ for all $u \in (c_0, c_1)$ and $v_i(u) = 0$ outside this interval, i.e. for all $u \in]c_0, c_1[$. Of course, if $t \leq c_0$, the volume is zero at the beginning of the sweep until $u = c_0$. (The bounds c_0 and c_1 are characterized in more detail later.)

The volume function f_i at $X(t)$ with $X(t) \in \mathcal{U} \setminus \mathcal{B}$ is

$$f_i(X(t)) = \begin{cases} \int_{\max(t, c_0)}^{c_1} v_i(u) du & \text{if } t \in (0, c_1), \\ 0 & \text{otherwise.} \end{cases} \quad (3.4)$$

By the fundamental theorem of calculus, the differentiability of $f_i \circ \phi^{-1}$ restricted to L follows from the continuity of v_i in (c_0, c_1) . If v_i was continuous in (c_0, c_1) but discontinuous at a bound c of the interval ($c = c_0$ or $c = c_1$), $f_i \circ \phi^{-1}$ would be differentiable on L except at $t = c$.




Figure 3.14: With $\phi(X) = L(t)$: sweeping $F_i(X)$ by the hyperplanes dual to $L(u)$ for $u \in (t, c_1)$.

If L contains no point of discontinuity of v_i , the derivative is

$$\frac{df_i(X(t))}{dt} = \begin{cases} -v_i(t) & \text{if } t \in (c_0, c_1), \\ 0 & \text{if } 0 < t \leq c_0 \text{ or } t \geq c_1. \end{cases} \quad (3.5)$$

We will now determine the interval bounds c_0 and c_1 and we show that v_i is continuous on L except at some isolated points that correspond to a bound c of the interval (c_0, c_1) .

The two bounds c_0 and c_1 correspond to the first and the last intersection of L with the hyperplanes of $\phi(\mathcal{O})^*$ incident to $\phi(B_i)$. The intersections are ordered by increasing value of u . The upper bound c_1 corresponds to the parameter at which L leaves the support $\phi(\Delta_i)$ of $\lambda_i \circ \phi^{-1}$. For all $u > c_1$, $L(u)$ lies in $\text{cone}(\phi(B_i))$. In Section 3.2, we explained that, in this case, $\phi^{-1}(L(u)) \notin \Delta_i$. For all u smaller than the lower bound c_0 , $L(u)$ is sufficiently low so that $\phi(B_i)$ is no longer on the lower convex hull of $\phi(\mathcal{B}) \cup L(u)$. This means, with respect to the dual hyperplanes, that the halfspace $L(u)^{-}$ contains the entire facet F_i . Consequently, $L(u)^*$ does not intersect F_i . (Refer again to Figure 3.14.)

v_i is, indeed, a continuous function in (c_0, c_1) because the boundary of F_i are portions of hyperplanes and $v_i(u)$ is the volume of $L(u)^* \cap F_i$. It is discontinuous at c ($c = c_0$ or $c = c_1$) only if $L(c)^*$ contains a $(d - 1)$ -dimensional facet of F_i . This $(d - 1)$ -dimensional facet is the intersection of H_i with some hyperplane $H_j = \phi(B_j)^*$ where B_j is a neighbor of B_i in $\text{Reg}(\mathcal{B})$.

It is easier to explain this situation in the power diagram. The intersection of H_i and $L(u)^*$ projects onto the $((d - 1)$ dimensional) bisector hyperplane of the two weighted points B_i and $X(u)$, called $\mathcal{H}(u)$. As u varies, $\mathcal{H}(u)$ sweeps the sub-cell $P_i(X)$. $v_i(u)$ is discontinuous at c only if $\mathcal{H}(c)$ coincides with the bisector hyperplane \mathcal{H}_{ij} as depicted in Figure 3.15.³ If the vertical projection of L contains the center of some site B_j that is a neighbor of B_i in $\text{Reg}(\mathcal{B})$, there must be a point $L(u)$ such that $\mathcal{H}(u) = \mathcal{H}_{ij}$ (because the hyperplanes $\mathcal{H}(u)$ are parallel to \mathcal{H}_{ij} and they sweep the entire facet). This must be either $L(c_0)$ or $L(c_1)$ because \mathcal{H}_{ij} contains the boundary of the cell F_i .




Figure 3.15: $P_i(X(t))$ is swept by parallel hyperplanes as u varies from t to c_1 . (a) The volume function $v(u)$ is continuous over $[t, c_1]$. (b) $v(u)$ is discontinuous at c_1 because $\mathcal{H}(c_1) = \mathcal{H}_{ij}$.

From duality, we can derive that if v_i is discontinuous at c , then $L(c)$ lies on the line through $\phi(B_i)$ and $\phi(B_j)$: the hyperplanes $\phi(B_i)^*$, $\phi(B_j)^*$ and $L(c)^*$ have a common $(d - 1)$ -dimensional intersection, thus, by duality, $\phi(B_i)$, $\phi(B_j)$ and $L(c)$ must lie on one line. The lines through $\phi(B_i)$ and $\phi(B_j)$ where B_j is a neighbor of B_i in $\text{Reg}(\mathcal{B})$ are exactly the edges of $\text{cone}(\phi(B_i))$.

The differentiability of f_i implies the differentiability of π_i . We conclude that π_i restricted to $\phi^{-1}(L)$ is differentiable as long as the point $L(t)$ does not lie on an edge of $\text{cone}(\phi(B_i))$. So far, we excluded the case that L is a vertical ray. However, the proof works in the same manner since the facet $F_i(X)$ can also be swept by the hyperplanes L^* which are, this time, parallel. In Lemma 3.10, we determine the gradient of π_i in this direction. This will complete the proof.

We can now conclude. Let \mathcal{E}_i be the set $X \in \Delta_i$ such that $\phi(X)$ lies on an edge of $\text{cone}(\phi(B_i))$. Obviously, \mathcal{E}_i is a subset of $\mathbb{R}^d \times \mathbb{R}$ with measure 0. The C^1 continuity of the restriction of π_i to $\phi^{-1}(L)$ implies the differentiability of π_i in all $\mathcal{U} \setminus \mathcal{E}_i$ as we argued in the beginning. The differentiability of λ_i

³Both drawings are schematic so there no coherence between the distances.

follows from the differentiability of π_i . The C^0 continuity of λ_i has been shown in the preceding section. At the points $\mathcal{E}_i \setminus \mathcal{B}$ it is obvious because λ_i is defined by volume functions which are C^0 continuous. At $B_j \in \mathcal{B}$, the C^0 continuity of $\lambda_i(X)$ is extended by continuity. \square

Continuity properties with respect to a data site B_j

In Chapter 6, we make use of the partial derivative of $\pi_i(X)$ with respect to a sample point B_j assuming X to be fixed. The proof of the following lemma about the continuity properties of $\pi_i(X)$ with respect to a sample point B_j is useful to understand the formula for the gradient. The result is the same as in unweighted Voronoi diagrams.

Lemma 3.7 *The natural neighbor coordinate $\lambda_i(X)$ of $X \in \mathcal{U}$ is a C^0 continuous function of $B_j \in \mathcal{B}$ and a C^1 continuous function of $B_j \in \mathcal{B}$ except at a finite set of lines in \mathcal{U} .*




Figure 3.16: (a) $P_i^{-j}(X)$ with $\text{Pow}(\mathcal{B}^{-j})$ (dotted lines) and $\text{Pow}(\mathcal{B}^{-j+})$ (slashed lines) (b) $P_i^{-j}(B_j)$ with $\text{Pow}(\mathcal{B}^{-j})$ (dotted) and $\text{Pow}(\mathcal{B})$ (solid) (c) $P_i^{-j+}(B_j)$ with $\text{Pow}(\mathcal{B}^{-j+})$ (slashed) and $\text{Pow}^+(\mathcal{B})$ (small slashed lines).




Figure 3.17: (a) $P_i(X)$ (b) dark grey: $P_i^{-j+}(B_j)$, middle grey: $P_i^{-j}(B_j) \setminus P_i^{-j+}(B_j)$, light grey: $P_i(X)$.

Proof: Piper's proof [84, section 3] generalizes directly to power diagrams: Let $\mathcal{B}^{-j} = \mathcal{B} \setminus \{B_j\}$ and $\mathcal{B}^{-j+} = \mathcal{B} \setminus \{B_j\} \cup \{X\}$. In the sequel, the superscript $-j$ (resp. $-j+$) indicates that $\text{Pow}(\mathcal{B}^{-j})$ (resp.

$\text{Pow}(\mathcal{B}^{-j+})$) is referred to. Similar to Piper [84], the sub-cell $P_i(X)$ is obtained as a composition of several power cells from $\text{Pow}(\mathcal{B}^{-j})$ and $\text{Pow}(\mathcal{B}^{-j+})$:

$$P_i(X) = P_i^{-j}(X) \setminus [P_i^{-j}(B_j) \setminus P_i^{-j+}(B_j)], i \neq j \quad (3.6)$$

See Figures 3.16 and 3.17 for a geometric proof of the formula in two dimensions. In the sequel we give a formal proof that is valid in any dimension. It consists of writing the equations of each power cell as done in [84].

The subcells $P_i^{-j}(B_j)$ and $P_i^{-j+}(B_j)$ consists of the points with less power distance to B_i than to any site except B_j . In $P_i^{-j+}(B_j)$, the points have also less power to X than to B_i . The power distance of a point \mathbf{p} with respect to X is denoted $\Gamma_x(\mathbf{p})$. Its power distance with respect to B_j is $\Gamma_j(\mathbf{p})$.

$$P_i^{-j}(B_j) = \{\mathbf{p} \in \mathbb{R}^d : \Gamma_i(\mathbf{p}) \leq \Gamma_k(\mathbf{p}), i \neq k, j, \Gamma_j(\mathbf{p}) \leq \Gamma_i(\mathbf{p})\}.$$

$$P_i^{-j+}(B_j) = \{\mathbf{p} \in \mathbb{R}^d : \Gamma_i(\mathbf{p}) \leq \Gamma_k(\mathbf{p}), i \neq k, j, \Gamma_i(\mathbf{p}) \leq \Gamma_x(\mathbf{p}), \Gamma_j(\mathbf{p}) \leq \Gamma_i(\mathbf{p})\}.$$

$P_i^{-j}(B_j)$ without its subcell $P_i^{-j+}(B_j)$ is

$$P_i^{-j}(B_j) \setminus P_i^{-j+}(B_j) = \{\mathbf{p} \in \mathbb{R}^d : \Gamma_i(\mathbf{p}) \leq \Gamma_k(\mathbf{p}), i \neq k, j, \Gamma_i(\mathbf{p}) > \Gamma_x(\mathbf{p}), \Gamma_j(\mathbf{p}) \leq \Gamma_i(\mathbf{p})\}.$$

In the power diagram without B_j , the subcell $P_i^{-j}(X)$ has the equation

$$P_i^{-j}(X) = \{\mathbf{p} \in \mathbb{R}^d : \Gamma_i(\mathbf{p}) \leq \Gamma_k(\mathbf{p}), i \neq k, j, \Gamma_x(\mathbf{p}) \leq \Gamma_i(\mathbf{p})\}.$$

We gather the equations to obtain

$$\begin{aligned} P_i^{-j}(X) \setminus [P_i^{-j}(B_j) \setminus P_i^{-j+}(B_j)] &= \{\mathbf{p} \in \mathbb{R}^d : \Gamma_i(\mathbf{p}) \leq \Gamma_k(\mathbf{p}), i \neq k, j, \Gamma_x(\mathbf{p}) \leq \Gamma_i(\mathbf{p}), \Gamma_j(\mathbf{p}) > \Gamma_i(\mathbf{p})\} \\ &= \{\mathbf{p} \in \mathbb{R}^d : \Gamma_i(\mathbf{p}) \leq \Gamma_k(\mathbf{p}), i \neq k, \Gamma_x(\mathbf{p}) \leq \Gamma_i(\mathbf{p})\} \\ &= P_i(X). \end{aligned}$$

This is exactly the subcell $P_i(X)$ which proves Equation 3.6.

With Equation 3.6 and since $P_i^{-j+}(B_j)$ is included in $P_i^{-j}(B_j)$ while $P_i^{-j}(B_j) \setminus P_i^{-j+}(B_j)$ is included in $P_i^{-j}(X)$, the volume $\pi_i(X)$ can be re-written

$$\pi_i(X) = \pi_i^{-j}(X) - [\pi_i^{-j}(B_j) - \pi_i^{-j+}(B_j)], i \neq j. \quad (3.7)$$

A similar re-writing is possible for the case $i = j$. See also [84].

$$i = j : \quad \pi_i(X) = \sum_{k \neq i} \pi_k^{-i}(B_i) - [\sum_{k \neq i} \pi_k^{-i+}(B_i) + \pi_x^{-i+}(B_i)]. \quad (3.8)$$

It follows that the continuity and differentiability of $\lambda_i(X) = \frac{\pi_i(X)}{\sum_i \pi_i(X)}$ wrt. B_j depends on the continuity properties of $\pi_i^{-j}(B_j)$ and $\pi_i^{-j+}(B_j)$. For both functions, Lemma 3.6 applies as long as the cell of B_j is not empty in $\text{Pow}(\mathcal{B})^+$, i.e. $\pi^{-j+}(B_j) = 0$. However, as the cell of B_j disappears, i.e. $\pi^{-j+}(B_j) \rightarrow 0$, also its subcells disappear, so that $\pi_i(X) \rightarrow \pi_i^{-j}(X)$. The latter function is constant wrt. B_j . \square

3.4 Gradient of the coordinate function

In this section, we derive the formula for the gradient of the volume function π_i . With $\lambda_i(X) = \frac{\pi_i(X)}{\sum_{i=1}^n \pi_i(X)}$, the gradient of the coordinate function can easily be determined from $\nabla \pi_i(X)$. We first state the lemma. The proof is given in two parts because we compute the partial derivatives with respect to \mathbf{x} and w_x separately. The one-dimensional case is proven apart. Notice that $\nabla \pi_i(X) = \mathbf{0}$ if $P_i(X) = \emptyset$ or $P_i(X) = P(B_i)$ because, in both cases, $\pi_i(X)$ is constant with respect to X . The gradient of π_i is well defined and different from zero only inside the set \mathcal{U}_i with

$$\mathcal{U}_i := \{X \in \mathcal{U} \setminus \mathcal{E}_i \mid P_i(X) \neq \emptyset \text{ and } P_i(X) \neq P(B_i)\}.$$

Lemma 3.8 *Let $X \in \mathcal{U} \setminus \mathcal{E}_i$. The gradient of the volume function $\pi_i(X)$ which is the volume of $P_i(X)$ in $\text{Pow}^+(\mathcal{B})$ is*

$$\nabla \pi_i(X) = \begin{cases} \frac{v_{ix}}{\|P_i - \mathbf{x}\|} \begin{pmatrix} \mathbf{c}_{ix} - \mathbf{x} \\ \frac{1}{2} \end{pmatrix} & \text{if } X \in \mathcal{U}_i, \\ \mathbf{0} & \text{otherwise,} \end{cases}$$

where v_{ix} is the $(d-1)$ -dimensional volume of the face $P_i(X)$ of $\text{Pow}^+(\mathcal{B})$, for $d > 1$, and \mathbf{c}_{ix} is the centroid of $P_i(X)$,

$$\mathbf{c}_{ix} = \frac{1}{v_{ix}} \int_{\mathbf{q} \in P_i(X)} \mathbf{q} d\mathbf{q}.$$

For $d = 1$, $v_{ix} = 1$.

The gradient in the one-dimensional case

We prove at first the one-dimensional case. Assume that $X \in \mathcal{U}_i$. Otherwise $\nabla \pi_i(X) = \mathbf{0}$ because $\pi_i(X)$ is constant. For $d = 1$, c_{ix} is the vertex $P^+(X, B_i)$ of $\text{Pow}^+(\mathcal{B})$ and

$$\pi_i(X) = \|P^+(X, B_i) - P(B_j, B_i)\|.$$

See Figure 3.18. Solving $\Gamma_i(P^+(X, B_i)) = \Gamma_x(P^+(X, B_i))$, we obtain the equation of the vertex $P^+(X, B_i)$:

$$P^+(X, B_i) = \frac{x + p_i}{2} + \frac{w_x - w_i}{2(p_i - x)}.$$

Let $Y = (y, w_y)$. The volume $\pi_i(Y)$ and the vertex $P^+(Y, B_i)$ are defined in the same way as for X . If Y is close enough to X (so that $P^+(X, B_i)$ lies on the same side of $P(B_j, B_i)$ as $P^+(Y, B_i)$), the sign of $P^+(X, B_i) - P(B_j, B_i)$ and $P^+(Y, B_i) - P(B_j, B_i)$ is the same. It is determined by the sign of $p_i - x$ (cf. Figure 3.18). We obtain

$$\|\pi_i(X) - \pi_i(Y)\| = \|P^+(X, B_i) - P^+(Y, B_i)\| = \left\| \frac{x - y}{2} + \frac{w_x - w_i}{2(p_i - x)} - \frac{w_y - w_i}{2(p_i - y)} \right\|.$$

We compute the two directional derivatives with respect to x and to w_x . In order to obtain the derivative




Figure 3.18: $P_i(X)$ is delimited by $P(B_j, B_i)$ and $P(X, B_i)$.

wrt. x , we take $Y = (x + h, w_x)$ and get

$$\frac{\partial \pi(X)}{\partial x} = \lim_{h \rightarrow 0} \frac{\|P^+(X, B_i) - P^+(Y, B_i)\|}{h} \frac{(p_i - x)}{\|p_i - x\|} = \left\| \frac{1}{2} + \frac{w_x - w_i}{2(p_i - x)^2} \right\| \frac{(p_i - x)}{\|p_i - x\|} = \frac{1}{\|p_i - x\|} (c_{ix} - x),$$

because $c_{ix} = P^+(X, B_i)$ and $P^+(X, B_i) - x = \frac{p_i - x}{2} + \frac{w_x - w_i}{2(p_i - x)}$.

To compute the derivative with respect to the weight w_x , we take $Y = (x, w_x + h)$:

$$\frac{\partial \pi(X)}{\partial w_x} = \lim_{h \rightarrow 0} \frac{\pi_i(X) - \pi_i(Y)}{h} = \frac{1}{2\|p_i - x\|}.$$

This proves Lemma 3.8 for $d = 1$ and $X \in \mathcal{U}_i$:

$$d = 1 : \quad \nabla \pi_i(X) = \frac{1}{\|p_i - x\|} \begin{pmatrix} c_{ix} - x \\ \frac{1}{2} \end{pmatrix}.$$

Partial derivative with respect to \mathbf{x}

Piper derived a formula for the gradient of the volume function $\pi_i(X)$ in the unweighted Voronoi diagram. His proof can be adapted to the weighted case. For dimensions $d > 2$, however, we show a simpler proof by induction over the dimension. We consider the d -dimensional derivative of the volume function $\pi_i(X)$ as a function of the point location \mathbf{x} while the weight w_x is constant,

$$\frac{\partial}{\partial \mathbf{x}} \pi_i(X) := \left(\frac{\partial \pi_i(X)}{\partial x_1}, \frac{\partial \pi_i(X)}{\partial x_2}, \dots, \frac{\partial \pi_i(X)}{\partial x_d} \right).$$

Lemma 3.9 *Let $X \in \mathcal{U} \setminus \mathcal{E}_i$. The gradient of the volume function $\pi_i(X)$ as a function of \mathbf{x} with fixed weight w_x is*

$$\frac{\partial}{\partial \mathbf{x}} \pi_i(X) = \begin{cases} \frac{v_{ix}}{\|p_i - \mathbf{x}\|} \overrightarrow{\mathbf{x}c_{ix}} & \text{if } X \in \mathcal{U}_i, \\ \mathbf{0} & \text{otherwise,} \end{cases}$$

where v_{ix} and c_{ix} are defined as in Lemma 3.8.

Proof: Suppose that $X \in \mathcal{U}_i$. Otherwise $\nabla \pi_i(X) = \mathbf{0}$ because $\pi_i(X)$ is constant. First, we consider the two-dimensional case, $d = 2$. The proof follows all steps of Piper's proof with some small differences. Let the vertices of the edge $P^+(X, B_i)$ be called \mathbf{a}_1 and \mathbf{a}_2 . We suppose that the vertex \mathbf{a}_1 is dual to the triangle $F^+(B_i, X, B_j) \in \text{Reg}^+(\mathcal{B})$. The gradient is determined as \mathbf{x} moves on a circle centered on the




Figure 3.19: Replacing x by $r(t)$ changes a_{i_2} but not a_{i_1} .

vertex a_{i_1} (thus, we compute the scalar product of the gradient with the vector $\overrightarrow{xa_{i_1}}^\perp$ that is orthogonal to $\overrightarrow{xa_{i_1}}$). Contrary to the unweighted case, this circle does not necessarily pass through the centers of the two other generators of a_{i_1} . However, as x moves on this circle, a_{i_1} remains the vertex dual to the triangle $F^+(B_i, X, B_j)$. Only the second vertex a_{i_2} changes as shown in Figure 3.19. We proceed as Piper and




Figure 3.20: The area of the highlighted triangle is $g(t) - g(0)$.

assume the circle to be parameterized counterclockwise by $r(t)$ such that $r(0) = x$ and t is the angle $\angle xa_{i_1}r(t)$. Define $X(t) = (r(t), w_x)$ and $g(t) = \pi_i(X(t))$. We want to compute the derivative $g'(0)$ given by $g'(0) = \lim_{t \rightarrow 0} \frac{g(t) - g(0)}{t}$. Let $a_{i_2}(t)$ be the vertex of the edge $P^+(X(t), B_i)$ other than a_{i_1} – see Figure 3.20. For small enough t , the difference between the power cells at $t_0 = 0$ and t is a triangle with area $\frac{1}{2} \sin(\theta(t))v_{ix}v_{ix}(t)$ where $\theta(t)$ is the angle $\angle xp_i r(t)$ and, $v_{ix}(t)$ is the volume of the face between

$X(t)$ and B_i . (The triangle is highlighted in Figure 3.20.) We determine $\lim_{t \rightarrow 0} \frac{\sin(\theta(t))}{t}$: Let α be the




Figure 3.21: As $t \rightarrow 0$, $r(t)$ lies on the tangent line.

angle at x in the triangle $(xr(t)p_i)$. With law of sines, we can rewrite $\sin(\theta(t))$:

$$\frac{\sin(\theta(t))}{\|x - r(t)\|} = \frac{\sin \alpha}{\|p_i - r(t)\|} \Rightarrow \sin(\theta(t)) = \frac{\sin \alpha \|x - r(t)\|}{\|p_i - r(t)\|}. \quad (3.9)$$

For $\lim t \rightarrow 0$, $r(t)$ lies on the line perpendicular to the line $(a_{i_1}x)$ containing x . Therefore, we approximate

$$\lim_{t \rightarrow 0} \|x - r(t)\| = \lim_{t \rightarrow 0} \tan(t) \|x - a_{i_1}\| = \lim_{t \rightarrow 0} t \|x - a_{i_1}\|. \quad (3.10)$$

Furthermore,

$$\lim_{t \rightarrow 0} \|p_i - r(t)\| = \|p_i - x\|. \quad (3.11)$$

With Equations (3.9), (3.10), and (3.11), we obtain:

$$\lim_{t \rightarrow 0} \frac{\sin(\theta(t))}{t} = \frac{\sin(\alpha) \|x - r(t)\|}{t \|p_i - r(t)\|} = \frac{\sin(\alpha) \|x - a_{i_1}\|}{\|x - p_i\|} = \frac{\|x - m\|}{\|x - p_i\|},$$

where m is the orthogonal projection of a_{i_1} onto the line (xp_i) . The last equality is due to the fact that \overrightarrow{xm} is the orthogonal projection of $\overrightarrow{xa_{i_1}}$ onto $\overrightarrow{xp_i}$ (thus, $\|x - m\| = \cos(\frac{\Pi}{2} - \alpha) \|x - a_{i_1}\|$).

We have shown that

$$\frac{\partial}{\partial x} \pi_i(X) \overrightarrow{xa_{i_1}}^\perp = g'(0) = \frac{v_{ix}^2 \|x - m\|}{2 \|x - p_i\|} \quad (3.12)$$

By a similar argument, the gradient can be determined in a second direction $\overrightarrow{xa_{i_2}}^\perp$ where a_{i_2} denotes the second vertex of $P^+(X, B_i)$. We get $\frac{\partial}{\partial x} \pi_i(X) \overrightarrow{xa_{i_2}}^\perp = -\frac{v_{ix}^2 \|x - m\|}{2 \|x - p_i\|}$.

Hence, we only need to verify that the two directional derivatives correspond to the formula in the lemma.

Notice that $\overrightarrow{xc_{ix}} = \frac{1}{2}(\overrightarrow{xa_{i_1}} + \overrightarrow{xa_{i_2}})$.

$$\frac{\partial}{\partial x} \pi_i(X) \overrightarrow{xa_{i_1}}^\perp = \frac{v_{ix}}{2 \|p_i - x\|} (\overrightarrow{xa_{i_1}} + \overrightarrow{xa_{i_2}}) \overrightarrow{xa_{i_1}}^\perp = \frac{v_{ix}}{2 \|p_i - x\|} \overrightarrow{xa_{i_2}} \overrightarrow{xa_{i_1}}^\perp$$

Piper remarks that $\overrightarrow{\mathbf{x}\mathbf{a}_{i_2}}\overrightarrow{\mathbf{x}\mathbf{a}_{i_1}}^\perp$ is twice the area of the triangle $(\mathbf{x}\mathbf{a}_{i_1}\mathbf{a}_{i_2})$. The same quantity can be written as $v_{ix}\|\mathbf{x} - \mathbf{m}\|$ since $\|\mathbf{x} - \mathbf{m}\|$ is the triangle's height and $v_{ix} = \|\mathbf{a}_{i_1} - \mathbf{a}_{i_2}\|$. Hence, replacing $\overrightarrow{\mathbf{x}\mathbf{a}_{i_2}}\overrightarrow{\mathbf{x}\mathbf{a}_{i_1}}^\perp = v_{ix}\|\mathbf{x} - \mathbf{m}\|$, we obtain the directional derivative of Equation (3.12). Similar arguments hold for the second direction.

The generalization to higher dimensions is obtained by an inductive argument over the dimension. The volume $\pi_i(X)$ in a d -dimensional power diagram can be obtained by sweeping the cell $P_i(X)$ with parallel hyperplanes and by integrating the volume of the intersection of $P_i(X)$ with the sweeping plane. We denote the sweeping hyperplane by $\mathcal{H}(t)$ where $t \in [t_0, t_1]$ such that

$$\mathcal{H}(t) \cap P(B_i, X) \neq \emptyset \quad \forall t \in [t_0, t_1] \quad \text{and} \quad \mathcal{H}(t) \cap P(B_i, X) = \emptyset, \text{ otherwise.}$$

From Lemma 2.3, we know that the intersection $\mathcal{H}(t) \cap \text{Pow}(\mathcal{B})$ is a $(d-1)$ -dimensional power diagram – assuming that $\mathcal{H}(t)$ is in general position. However, it is sufficient to prove the gradient in d linearly independent directions so that this is no constraint. Although the sweeping plane depends on the direction for which we compute the gradient, degenerate positions can be avoided.

Let $\overline{\mathbf{p}}_i(t)$ be the projection of \mathbf{p}_i onto $\mathcal{H}(t)$ and $\overline{w}_i(t) = w_i - \|\mathbf{p}_i - \overline{\mathbf{p}}_i(t)\|^2$. We use the same notation as in Lemma 2.3 and, for each $B_i \in \mathcal{B}$, we denote the site corresponding to B_i in $\mathcal{H}(t) \cap \text{Pow}(\mathcal{B})$ by $\overline{B}_i(t) = (\overline{\mathbf{p}}_i(t), \overline{w}_i(t))$ (the same for $\overline{X}(t)$). Of course, $\pi_i(\overline{X}(t))$ is the volume of $P_i(\overline{X}(t)) = \mathcal{H}(t) \cap P_i(X)$.

The gradient of $\pi_i(X)$ is the integral of the gradients $\frac{\partial}{\partial \overline{\mathbf{x}}(t)}\pi_i(\overline{X}(t))$ over the interval $[t_0, t_1]$:

$$\pi_i(X) = \int_{t_0}^{t_1} \pi_i(\overline{X}(t)) dt \quad \implies \quad \frac{\partial}{\partial \mathbf{x}}\pi_i(X) = \int_{t_0}^{t_1} \frac{\partial}{\partial \mathbf{x}}\pi_i(\overline{X}(t)) dt.$$

It is computed in a fixed, arbitrary direction $\overrightarrow{\mathbf{v}}$ other than $\overrightarrow{\mathbf{x}\mathbf{p}}_i$. We choose the sweeping plane to be parallel to $\overrightarrow{\mathbf{v}}$ and to $\overrightarrow{\mathbf{x}\mathbf{p}}_i$. Thus, during the sweep, we maintain $\|\mathbf{p}_i - \mathbf{x}\| = \|\overline{\mathbf{p}}_i(t) - \overline{\mathbf{x}}(t)\|$. Observe that for any vector $\overrightarrow{\mathbf{x}\mathbf{y}}$:

$$\overrightarrow{\mathbf{x}(t)\mathbf{y}(t)} \cdot \overrightarrow{\mathbf{v}} = \overrightarrow{\mathbf{x}\mathbf{y}} \cdot \overrightarrow{\mathbf{v}}$$

because $\overrightarrow{\mathbf{v}}$ is in the sweeping plane $\mathcal{H}(t)$ onto which we project.

For $d = 2$, we know the gradient from the first part of the proof. In direction $\overrightarrow{\mathbf{v}}$, by replacing the expression for the centroid, it writes as follows:

$$d = 2: \quad \frac{\partial}{\partial \mathbf{x}}\pi_i(X) \cdot \overrightarrow{\mathbf{v}} = \frac{1}{\|\mathbf{p}_i - \mathbf{x}\|} \int_{q \in P_i(X)} \overrightarrow{\mathbf{x}\mathbf{q}} \cdot \overrightarrow{\mathbf{v}} dq.$$

By inductive hypothesis, we suppose that we know the gradient for dimension $d-1$, $d > 2$, in the diagram $\mathcal{H}(t) \cap \text{Pow}(\mathcal{B})$:

$$\frac{\partial}{\partial \mathbf{x}}\pi_i(\overline{X}(t)) \cdot \overrightarrow{\mathbf{v}} = \frac{1}{\|\overline{\mathbf{p}}_i(t) - \overline{\mathbf{x}}(t)\|} \int_{q \in P_i(\overline{X}(t))} \overrightarrow{\mathbf{x}(t)\mathbf{q}} \cdot \overrightarrow{\mathbf{v}} dq. \quad (3.13)$$

We show that

$$\begin{aligned}
\frac{\partial}{\partial \mathbf{x}} \pi_i(X) \cdot \vec{\mathbf{v}} &= \left(\int_{t_0}^{t_1} \frac{\partial}{\partial \mathbf{x}} \pi_i(\overline{X}(t)) dt \right) \cdot \vec{\mathbf{v}} = \int_{t_0}^{t_1} \left(\frac{1}{\|\overline{\mathbf{p}_i}(t) - \overline{\mathbf{x}}(t)\|} \int_{q \in P_i(\overline{X}(t))} \overrightarrow{\overline{\mathbf{x}}(t)\mathbf{q}} \cdot \vec{\mathbf{v}} d\mathbf{q} \right) dt \\
&= \frac{1}{\|\mathbf{p}_i - \mathbf{x}\|} \int_{t_0}^{t_1} \int_{q \in P_i(\overline{X}(t))} \overrightarrow{\overline{\mathbf{x}}(t)\mathbf{q}} \cdot \vec{\mathbf{v}} d\mathbf{q} dt \\
&= \left(\frac{1}{\|\mathbf{p}_i - \mathbf{x}\|} \int_{q \in P_i(X)} \overrightarrow{\mathbf{x}\mathbf{q}} d\mathbf{q} \right) \cdot \vec{\mathbf{v}}.
\end{aligned}$$

□

Derivative as function of w_x

Lemma 3.10 *The partial derivative of the volume function $\pi_i(X)$, $X \in \mathcal{U} \setminus \mathcal{E}_i$, with respect to w_x is*

$$\frac{\partial}{\partial w_x} \pi_i(X) = \begin{cases} \frac{v_{ix}}{2\|\mathbf{p}_i - \mathbf{x}\|} & \text{if } X \in \mathcal{U}_i, \\ 0 & \text{otherwise,} \end{cases}$$

where v_{ix} is defined as in Lemma 3.8.

Proof: Suppose that $X \in \mathcal{U}_i$. Otherwise $\nabla \pi_i(X) = 0$ because $\pi_i(X)$ is constant. The volume of $P_i(X)$ can be obtained by sweeping the cell with a hyperplane $\mathcal{H}(t)$ parallel to the bisector hyperplane $\mathcal{H}_{ix} = \mathcal{H}(0)$ of B_i and X . Let h be the maximum distance of a vertex of $P_i(X)$ to $\mathcal{H}(0)$, and $\mathcal{H}(t)$ be at distance t from $\mathcal{H}(0)$, $t \in [0 \dots h]$. Furthermore, $l_i(t)$ is the volume $\mathcal{H}(t) \cap P_i(X)$. We write

$$\pi_i(X) = \int_0^h l_i(t) dt. \tag{3.14}$$

Another way to obtain this sweep is to decrease the weight w_x until the cell $P_i(X)$ disappears and to consider the volume $v_i((\mathbf{x}, w))$ of the intersection of the bisector hyperplane of (\mathbf{x}, w) and B_i with $P(B_i)$. Refer to Figure 3.22(a). We can write t as a function of w with $l_i(t) = v_i((\mathbf{x}, w))$ for $t = f(w)$. Changing the variable in Equation 3.14, we obtain

$$\pi_i(X) = \int_{f^{-1}(0)}^{f^{-1}(h)} -l_i(f(w)) f'(w) dw.$$

We need to determine the function $f(w)$, its derivate $f'(w)$ and the new boundaries $f^{-1}(0)$ and $f^{-1}(h)$. From the formula for the power distance from Chapter 2, we deduce the distance from \mathbf{x} to the closest point $\mathbf{y} \in \mathcal{H}_{ix}$ by solving $\Gamma_i(\mathbf{y}) = \Gamma_x(\mathbf{y})$:

$$\|\mathbf{x} - \mathbf{y}\| = \frac{1}{2} \|\mathbf{p}_i - \mathbf{x}\| + \frac{w_x - w_i}{2\|\mathbf{p}_i - \mathbf{x}\|}.$$

The distance from \mathbf{x} to the closest point $\mathbf{y}(t) \in \mathcal{H}(t)$ (see Figure 3.22) is

$$\|\mathbf{x} - \mathbf{y}(t)\| = \frac{1}{2} \|\mathbf{p}_i - \mathbf{x}\| + \frac{w - w_i}{2\|\mathbf{p}_i - \mathbf{x}\|}.$$




Figure 3.22: (a) Sweeping $P_i(\mathbf{x})$ by changing w_x . (b) The distance between y and $y(t)$ is $t = \frac{w_x - w}{2\|\mathbf{p}_i - \mathbf{x}\|}$.

Hence, t , the distance between \mathcal{H}_{ix} and $\mathcal{H}(t)$, is

$$t = f(w) = \|\mathbf{x} - \mathbf{y}\| - \|\mathbf{x} - \mathbf{y}(t)\| = \frac{w_x - w}{2\|\mathbf{p}_i - \mathbf{x}\|} \quad \text{with derivative } f'(w) = -\frac{1}{2\|\mathbf{p}_i - \mathbf{x}\|}.$$

Let $w_\emptyset = f^{-1}(h)$ so that $h = \frac{w_x - w_\emptyset}{2\|\mathbf{p}_i - \mathbf{x}\|}$. w_\emptyset is the largest weight with $v_i((\mathbf{x}, w_\emptyset)) = 0$. With these results (and replacing $l_i(f(w)) = v_i((\mathbf{x}, w))$), Equation 3.4 writes

$$\pi_i(X) = \int_{w_\emptyset}^{w_x} \frac{v_i((\mathbf{x}, w))}{2\|\mathbf{p}_i - \mathbf{x}\|} dw.$$

Notice that f is continuous over the interval w_\emptyset to w_x and $l_i(t)$ is continuous over $f(w_\emptyset)$ to $f(w_x)$ —except when \mathcal{H}_{ix} contains a bisector hyperplane of \mathbf{p}_i and one of its neighbors. This is excluded on $\mathcal{U} \setminus \mathcal{E}_i$ as explained in the proof of Lemma 3.6. The result follows:

$$\frac{\partial}{\partial w_x} \pi_i(X) = \frac{\partial}{\partial w_x} \left(\int_{w_\emptyset}^{w_x} \frac{v_i((\mathbf{x}, w))}{2\|\mathbf{p}_i - \mathbf{x}\|} dw \right) = \frac{v_i((\mathbf{x}, w_x))}{2\|\mathbf{p}_i - \mathbf{x}\|} = \frac{v_{ix}}{2\|\mathbf{p}_i - \mathbf{x}\|}.$$

□

With Lemma 3.9 and the preceding lemma, we complete the proof of Lemma 3.8.

Remark: It is important to notice that, except if $X \in \mathcal{E}_i$, the volume v_{ix} of the face $P^+(X, B_i)$ approaches zero as $P_i(X) \rightarrow \emptyset$ or $P_i(X) \rightarrow P(B_i)$. Consequently, the gradient is continuous in $\mathcal{U} \setminus \mathcal{E}_i$. If $X \rightarrow Y \in \mathcal{E}_i$, the bisector of X and B_i contains a $((d-1)$ -dimensional) face of $P(B_i)$, so $\lim_{X \rightarrow Y} v_{ix} > 0$. This is explained in the proof of Lemma 3.6. Refer to Figure 3.15.

Gradient as function of a data site B_j

Let $P_{ij}(X)$ be the part of the face $P(B_i, B_j)$ that is “stolen” from it at the insertion of X into $\text{Pow}(B)$,

$$P_{ij}(X) := P(B_i, B_j) \cap P^+(X).$$

We are interested in the gradient of the volume $\pi_i(X)$ of the subcell $P_i(X)$ as a function of a data site B_j . $P_i(X)$ depends on B_j whenever $P_{ij}(X) \neq \emptyset$. We define

$$\mathcal{U}_{ij} := \{X \in \mathcal{U} \setminus \mathcal{E}_i \mid P_{ij}(X) \neq \emptyset\}.$$

Only if $X \in \mathcal{U}_{ij}$, the subcell $P_i(X)$ exists and is influenced by B_j . Outside \mathcal{U}_{ij} , $P_i(X)$ is constant as B_j changes. The gradient of the volume function $\pi_i(X)$ is therefore zero.

Lemma 3.11 *The gradient of the volume function $\pi_i(X)$, $X \in \mathcal{U} \setminus \mathcal{E}_i$, with respect to $B_j \in \mathcal{B}$ is*

(i) for $j \neq i$:

$$\frac{\partial}{\partial B_j} \pi_i(X) = \begin{cases} \frac{-v_{ij}(X)}{\|\mathbf{p}_i - \mathbf{p}_j\|} \begin{pmatrix} \mathbf{c}_{ij}(X) - \mathbf{p}_j \\ \frac{1}{2} \end{pmatrix} & \text{if } X \in \mathcal{U}_{ij}, \\ \mathbf{0} & \text{otherwise,} \end{cases}$$

(ii) else,

$$\frac{\partial}{\partial B_i} \pi_i(X) = \begin{cases} \sum_{k \neq i} \frac{v_{ik}(X)}{\|\mathbf{p}_i - \mathbf{p}_k\|} \begin{pmatrix} \mathbf{c}_{ik}(X) - \mathbf{p}_i \\ \frac{1}{2} \end{pmatrix} - \frac{v_{ix}}{\|\mathbf{p}_i - \mathbf{x}\|} \begin{pmatrix} \mathbf{c}_{ix} - \mathbf{p}_i \\ \frac{1}{2} \end{pmatrix} & \text{if } X \in \mathcal{U}_i, \\ \mathbf{0} & \text{otherwise,} \end{cases}$$

where $v_{ij}(X)$ is the volume of the face $P_{ij}(X) = P_i(X) \cap P_j(X)$ and $\mathbf{c}_{ij}(X)$ is its centroid. v_{ix} and \mathbf{c}_{ix} are defined as in Lemma 3.9.

Proof: For $X \notin \mathcal{U}_{ij}$, $\pi_i(X)$ is constant with respect to B_j , so the gradient is $\mathbf{0}$.

Suppose $X \in \mathcal{U}_{ij}$. The proof follows from Lemma 3.7 and Lemma 3.8. For the case $i \neq j$:

$$\frac{\partial}{\partial B_j} \pi_i(X) = \frac{\partial}{\partial B_j} \left(\pi_i^{-j}(X) - [\pi_i^{-j}(B_j) - \pi_i^{-j+}(B_j)] \right) = \frac{\partial}{\partial B_j} \left(-\pi_i^{-j}(B_j) + \pi_i^{-j+}(B_j) \right),$$

because $\pi_i^{-j}(X)$ is constant wrt. B_j .

Lemma 3.8 applied to B_j in the power diagram $\text{Pow}^{-j}(\mathcal{B})$ yields the gradient $\frac{\partial}{\partial B_j} \pi_i^{-j}(B_j)$. We write the partial derivatives with respect to \mathbf{p}_j and w_j separately and replace the expression for the centroid. v_{ij} denotes the volume of the face $P(B_i, B_j)$.

$$\frac{\partial}{\partial \mathbf{p}_j} \pi_i^{-j}(B_j) = \frac{1}{\|\mathbf{p}_i - \mathbf{p}_j\|} \int_{q \in P(B_i, B_j)} \overrightarrow{\mathbf{p}_j \mathbf{q}} \, d\mathbf{q} \quad (3.15)$$

$$\frac{\partial}{\partial w_j} \pi_i^{-j}(B_j) = \frac{v_{ij}}{2\|\mathbf{p}_i - \mathbf{p}_j\|}. \quad (3.16)$$

The gradient $\frac{\partial}{\partial B_j} \pi_i^{-j+}(B_j)$ is very similar. In $\text{Pow}^{-j+}(\mathcal{B})$, the face between B_i and B_j is $P^+(B_i, B_j)$ and not $P(B_i, B_j)$. To obtain $\frac{\partial}{\partial B_j} \pi_i^{-j+}(B_j)$, we replace, in the preceding equations, $P(B_i, B_j)$ by $P^+(B_i, B_j)$ and v_{ij} by the volume v_{ij}^+ of $P^+(B_i, B_j)$.




Figure 3.23: The simplification of the gradient formula.

From Figure 3.23, it is easy to tell that $P(B_i, B_j) \setminus P^+(B_i, B_j) = P_{ij}(X)$. This simplifies the expression of the gradient. The partial derivatives with respect to \mathbf{p}_j and w_j reduce to

$$\begin{aligned}
\frac{\partial}{\partial \mathbf{p}_j} \pi_i(X) &= -\frac{1}{\|\mathbf{p}_i - \mathbf{p}_j\|} \int_{q \in P(B_i, B_j)} \overrightarrow{\mathbf{p}_j \mathbf{q}} \, dq + \frac{1}{\|\mathbf{p}_i - \mathbf{p}_j\|} \int_{q \in P^+(B_i, B_j)} \overrightarrow{\mathbf{p}_j \mathbf{q}} \, dq \\
&= -\frac{1}{\|\mathbf{p}_i - \mathbf{p}_j\|} \int_{q \in P(B_i, B_j) \setminus P^+(B_i, B_j)} \overrightarrow{\mathbf{p}_j \mathbf{q}} \, dq \\
&= -\frac{1}{\|\mathbf{p}_i - \mathbf{p}_j\|} \int_{q \in P_{ij}(X)} \overrightarrow{\mathbf{p}_j \mathbf{q}} \, dq = \frac{-v_{ij}(X)}{\|\mathbf{p}_i - \mathbf{p}_j\|} \overrightarrow{\mathbf{p}_j \mathbf{c}_{ij}(X)}, \text{ and} \\
\frac{\partial}{\partial w_j} \pi_i(X) &= \frac{-v_{ij} + v_{ij}^+}{2\|\mathbf{p}_i - \mathbf{p}_j\|} = \frac{-v_{ij}(X)}{2\|\mathbf{p}_i - \mathbf{p}_j\|}.
\end{aligned}$$

For the case $j = i$, we simplify Equation 3.8 in the same way.

$$\begin{aligned}
\frac{\partial}{\partial B_i} \pi_i(X) &= \frac{\partial}{\partial B_i} \left(\sum_{k \neq i} \pi_k^{-i}(B_i) - \sum_{k \neq i} \pi_k^{-i+}(B_i) - \pi_X^{-i+}(B_i) \right) \\
&= \sum_{k \neq i} \frac{\partial}{\partial B_i} (\pi_k^{-i}(B_i) - \pi_k^{-i+}(B_i)) - \frac{\partial}{\partial B_i} \pi_X^{-i+}(B_i) \\
&= \sum_{k \neq i} \frac{v_{ik}(X)}{\|\mathbf{p}_i - \mathbf{p}_k\|} \begin{pmatrix} \mathbf{c}_{ik}(X) - \mathbf{p}_i \\ \frac{1}{2} \end{pmatrix} - \frac{v_{ix}}{\|\mathbf{p}_i - \mathbf{x}\|} \begin{pmatrix} \mathbf{c}_{ix} - \mathbf{p}_i \\ \frac{1}{2} \end{pmatrix} \square
\end{aligned}$$

Limit behavior of $\lambda_i(X)$ and $\nabla \lambda_i(X)$ at B_j

Since the coordinate function is not differentiable on the sites of \mathcal{B} , it is particularly interesting to look at the behavior of λ_i as X approaches a site B_j . Remember that only if $w_x \geq w_j$, X is in the definition domain \mathcal{U} as $\mathbf{x} \rightarrow \mathbf{p}_j$. The critical case occurs if the weights approach faster than the point positions, i.e. $\frac{w_x - w_j}{\|\mathbf{x} - \mathbf{p}_j\|} \rightarrow 0$ as $X \rightarrow B_j$. Otherwise, $\nabla \pi_i(X) = \mathbf{0}$ (as $X \rightarrow B_j$). We show that the coordinate function is Lipschitz at B_j and that the norm of the gradient is bounded. By $\|B_j - X\|$, we denote the Euclidean distance between X and B_j in $\mathbb{R}^d \times \mathbb{R}$.

At first, we show two corollaries concerning the volume function π_i .

Corollary 3.12 *There exist some real constants c_1 and c_2 such that for any unit vector $\vec{\mathbf{u}} \in \mathbb{R}^d \times \mathbb{R}$ and $X \in \mathcal{U}$,*

1. *if $\|X - B_j\|$ sufficiently small and with $i \neq j$, $|\nabla \pi_i(X) \cdot \vec{\mathbf{u}}| \leq c_1$,*
2. *$|\nabla \pi_j(X) \cdot \vec{\mathbf{u}}| \leq \frac{c_2}{\|\mathbf{x} - \mathbf{p}_j\|}$.*

Proof: From Lemma 3.9, we know the gradient of $\pi_i(X)$.

Case $i \neq j$: if $\|X - B_j\|$ is sufficiently small, $\nabla \pi_i(X)$ has a bounded norm for all $i \neq j$ because $\|\mathbf{p}_i - \mathbf{x}\|$ and $\|\mathbf{c}_{ix} - \mathbf{x}\|$ are bounded and v_{ix} is bounded or approaches zero. Thus,

$$|\nabla \pi_i(X) \cdot \vec{\mathbf{u}}| = \left| \frac{v_{ix}}{\|\mathbf{p}_i - \mathbf{x}\|} \begin{pmatrix} \mathbf{c}_{ix} - \mathbf{x} \\ \frac{1}{2} \end{pmatrix} \cdot \vec{\mathbf{u}} \right| \leq c_1, \text{ if } i \neq j.$$

In fact, v_{ix} becomes 0 as the cell $P_i(X)$ disappears and $\pi_i(X) \rightarrow 0$ or as $P_i(X)$ is completely contained in $P^+(X)$ and $\pi_i(X) \rightarrow \pi(B_i)$. The exception to this behavior occurs if $X \rightarrow Y \in \mathcal{E}_i$.

Case $i = j$: Since v_{jx} and $\|\mathbf{c}_{jx} - \mathbf{x}\|$ are bounded,

$$|\nabla \pi_j(X) \cdot \vec{\mathbf{u}}| = \left| \frac{v_{jx}}{\|\mathbf{p}_j - \mathbf{x}\|} \begin{pmatrix} \mathbf{c}_{jx} - \mathbf{x} \\ \frac{1}{2} \end{pmatrix} \cdot \vec{\mathbf{u}} \right| \leq \frac{c_2}{\|\mathbf{p}_j - \mathbf{x}\|}. \square$$

Corollary 3.13 *There exists some real constant c_3 , such that for $i \neq j$,*

$$\pi_i(X) \leq c_3 \|X - B_j\| \quad \text{with } \|X - B_j\| \text{ small enough.}$$

Proof: We write the Taylor expansion of $\pi_i(B_j)$ at X :

$$\pi_i(B_j) = \pi_i(X) + \nabla \pi_i(X) (B_j - X) + O(\|X - B_j\|^2).$$

Since $\pi_i(B_j) = 0$, we can bound $\pi(X)$ by

$$\pi_i(X) = \nabla \pi_i(X) (X - B_j) + O(\|X - B_j\|^2) \leq c_3 \|X - B_j\|,$$

with Corollary 3.12 (1) since $i \neq j$ and $\|X - B_j\|$ is small. □

Lemma 3.14 *The coordinate function $\lambda_i(X)$ is Lipschitz at $X = B_j$, i.e. there exists a constant $c_4 \in \mathbb{R}$ such that*

$$|\lambda_i(X) - \lambda_i(B_j)| \leq c_4 \|X - B_j\| \quad \text{with } \|X - B_j\| \text{ small enough..}$$

Proof: Let c_6 and c_7 denote real constants.

Case $i \neq j$: Because $\lambda_i(X) = \frac{\pi_i(X)}{\pi(X)}$ and $\lim_{X \rightarrow B_j} \pi(X) = c_6$, the lemma is proven with Corollary 3.13. Indeed,

$$\lim_{X \rightarrow B_j} |\lambda_i(X) - \lambda_i(B_j)| = \lim_{X \rightarrow B_j} |\lambda_i(X)| = \lim_{X \rightarrow B_j} \frac{1}{c_6} \pi_i(X) \leq c_7 \|X - B_j\|, \quad (3.17)$$

Case $i = j$: With $\lambda_j(X) = 1 - \sum_{k \neq j} \lambda_k(X)$, we have shown the lemma because

$$\begin{aligned} \lim_{X \rightarrow B_j} |\lambda_j(X) - \lambda_j(B_j)| &= \lim_{X \rightarrow B_j} |1 - \sum_{k \neq j} \lambda_k(X) - 1| = \lim_{X \rightarrow B_j} | - \sum_{k \neq j} \lambda_k(X) | \\ &\leq \sum_{k \neq j} c_7 \|X - B_j\| \leq c_4 \|X - B_j\|. \square \end{aligned}$$

Lemma 3.15 *The norm of the gradient $\nabla \lambda_i(X)$ is bounded, i.e. there exists a constant real value c_5 such that for any unit vector $\vec{\mathbf{u}}$ in $\mathbb{R}^d \times \mathbb{R}$*

$$|\nabla \lambda_i(X) \cdot \vec{\mathbf{u}}| \leq c_5.$$

Proof: Again, let $c_k, k = 1, \dots$ denote real constants.

With $\pi(X) = \sum_k \pi_k(X)$, the gradient of the coordinate function $\lambda_i(X) = \frac{\pi_i(X)}{\pi(X)}$ can be written as

$$\nabla \lambda_i(X) = \frac{\nabla \pi_i(X) \pi(X) - \pi_i(X) \nabla \pi(X)}{\pi(X)^2} = \frac{1}{\pi(X)} \left(\nabla \pi_i(X) - \lambda_i(X) \sum_k \nabla \pi_k(X) \right)$$

At first, we suppose that $\|X - B_j\|$ is sufficiently small and consider the case $i \neq j$. We make three observations:

1. For $i \neq j$, we know from Lemma 3.14 that $\lambda_i(X) \leq c_4 \|X - B_j\|$.
2. For $k \neq j$, we know from Corollary 3.12 that $|\nabla \pi_k(X) \cdot \vec{\mathbf{u}}| \leq c_1$. Therefore,

$$\lambda_i(X) |\nabla \pi_k(X) \cdot \vec{\mathbf{u}}| \leq c_8 \|X - B_j\|.$$

3. For $k = j$, Corollary 3.12 yields $|\nabla \pi_j(X) \cdot \vec{\mathbf{u}}| \leq \frac{c_2}{\|\mathbf{x} - \mathbf{p}_j\|}$. Consequently, with Lemma 3.14,

$$\lambda_i(X) |\nabla \pi_j(X) \cdot \vec{\mathbf{u}}| \leq \frac{c_2 c_4 \|X - B_j\|}{\|\mathbf{x} - \mathbf{p}_j\|} \leq \frac{c_2 c_4 (\|\mathbf{x} - \mathbf{p}_j\| + |w_x - w_j|)}{\|\mathbf{x} - \mathbf{p}_j\|} \leq c_9,$$

because either $|w_x - w_j| \leq \|\mathbf{x} - \mathbf{p}_j\|$, or $\lambda_i(X)|\nabla \pi_j(X) \cdot \vec{\mathbf{u}}| = 0$ as $\mathbf{x} \rightarrow \mathbf{p}_j$.

Using again Corollary 3.12 (to confirm that $|\nabla \pi_i(X) \cdot \vec{\mathbf{u}}| \leq c_1$), we obtain

$$\begin{aligned} |\nabla \lambda_i(X) \cdot \vec{\mathbf{u}}| &\leq \frac{1}{\pi(X)} \left(|\nabla \pi_i(X) \cdot \vec{\mathbf{u}}| + \lambda_i(X) |\nabla \pi_j(X) \cdot \vec{\mathbf{u}}| + \sum_{k \neq j} \lambda_i(X) |\nabla \pi_k(X) \cdot \vec{\mathbf{u}}| \right) \\ &\leq c_{10}. \end{aligned} \tag{3.18}$$

For the case $i = j$, we exploit the following relation which is true for every normalized coordinate system:

$$\sum_{i=1}^n \lambda_i(X) = 1 \quad \implies \quad \sum_{i=1}^n \nabla \lambda_i(X) = \mathbf{0}.$$

With the right equation, we can express $\nabla \lambda_j(X) = -\sum_{i=1}^n \nabla \lambda_i(X)$ and the result follows immediately from the bound of Equation (3.18):

$$|\nabla \lambda_j(X) \cdot \vec{\mathbf{u}}| = \left| -\sum_{i=1}^n \nabla \lambda_i(X) \cdot \vec{\mathbf{u}} \right| \leq n c_{10} \leq c_5.$$

□

3.5 Non-convex natural neighbor coordinates

In this section, we define a system of coordinates that is not convex, i.e. the coordinate functions may be negative, but it has an interesting additional property. It is an unpublished idea of Clarkson that is described in [35]. By combining natural neighbor coordinates, a coordinate function is constructed such that \mathbf{x} is a linear combination of its natural neighbors, and, at the same time, $\mathbf{x}^2 = \mathbf{x}^t \mathbf{x}$, the square of the norm of \mathbf{x} , is a linear combination of the square norms of the neighbor points. However, to obtain this goal, we give up the convexity condition and allow negative coordinates.

As in the original work, our exposition uses natural neighbor coordinates for Voronoi diagrams. Re-call that $\mathcal{P} = \{\mathbf{p}_1, \dots, \mathbf{p}_n\}$ is a set of n points in \mathbb{R}^d and $\text{Vor}(\mathcal{P})$ its Voronoi diagram. Since the weight is constantly zero, we denote the coordinate function by $\lambda_i(\mathbf{x})$ and the volume function by $\pi_i(\mathbf{x})$.

Definition of the coordinate function γ_j^i

We now describe how to define the coordinate function γ_j^i that has the above properties. Let $\lambda_x^+(\mathbf{p}_i)$ be the natural coordinate of \mathbf{p}_i with respect to \mathbf{x} in the power diagram $\text{Vor}(\mathcal{P} \cup \{\mathbf{x}\} \setminus \{\mathbf{p}_i\})$, and $\lambda_j^+(\mathbf{p}_i)$ be the natural coordinate of \mathbf{p}_i with respect to \mathbf{p}_j in the same diagram. They are defined by the corresponding volume functions $\pi_x^+(\mathbf{p}_i)$, $\pi_j^+(\mathbf{p}_i)$, and $\pi^+(\mathbf{p}_i)$, the volume of the cell of \mathbf{p}_i in $\text{Vor}^+(\mathcal{P})$. See also Figure 3.24(a). We state precisely that

$$\lambda_x^+(\mathbf{p}_i) = \frac{\pi_x^+(\mathbf{p}_i)}{\pi^+(\mathbf{p}_i)} \quad \text{and} \quad \lambda_j^+(\mathbf{p}_i) = \frac{\pi_j^+(\mathbf{p}_i)}{\pi^+(\mathbf{p}_i)}.$$

We denote by $I_1(\mathbf{x})$ the set of indices of the natural neighbors of \mathbf{x} and by $J_i(\mathbf{x})$ the set of indices of the Delaunay neighbors of \mathbf{p}_i in $\text{Del}^+(\mathcal{P})$ other than \mathbf{x} . Let $I_2(\mathbf{x}) = I_1(\mathbf{x}) \cup_{i \in I_1(\mathbf{x})} J_i(\mathbf{x})$ be the natural neighbors of \mathbf{x} as well as the neighbors in $\text{Del}^+(\mathcal{P})$ of the natural neighbors of \mathbf{x} (other than \mathbf{x}). Notice, $I_1(\mathbf{x}) \subseteq I_2(\mathbf{x})$. See Figure 3.24(b). The first observation is that the barycentric coordinate property is also




Figure 3.24: (a) The volume functions $\pi_j^+(\mathbf{p}_i)$ and $\pi_x^+(\mathbf{p}_i)$ defined from $\text{Vor}(\mathcal{P} \cup \{\mathbf{x}\} \setminus \{\mathbf{p}_i\})$. $\pi^+(\mathbf{p}_i)$ is the total volume of the highlighted zones. (b) The first and second ring neighbors of \mathbf{x} .

true for any natural neighbor \mathbf{p}_i of \mathbf{x} wrt. λ_j^+ , i.e. $\mathbf{p}_i = \lambda_x^+(\mathbf{p}_i) \mathbf{x} + \sum_j \lambda_j^+(\mathbf{p}_i) \mathbf{p}_j$. The equation for the barycentric coordinate property of \mathbf{p}_i can be re-arranged in order to express \mathbf{x} as a linear combination of \mathbf{p}_i and its neighbors:

$$\mathbf{p}_i = \lambda_x^+(\mathbf{p}_i) \mathbf{x} + \sum_{j \in J_i(\mathbf{x})} \lambda_j^+(\mathbf{p}_i) \mathbf{p}_j \iff \mathbf{x} = \frac{1}{\lambda_x^+(\mathbf{p}_i)} (\mathbf{p}_i - \sum_{j \in J_i(\mathbf{x})} \lambda_j^+(\mathbf{p}_i) \mathbf{p}_j). \quad (3.19)$$

We also know that

$$\mathbf{x} = \sum_{i \in I_1(\mathbf{x})} \lambda_i(\mathbf{x}) \mathbf{p}_i. \quad (3.20)$$

Since $f(\mathbf{x}) = \mathbf{x}^2$ is a convex function, we obtain from Equations 3.19 and 3.20 the following two inequalities when taking the square of the points:⁴

$$\begin{aligned} \mathbf{p}_i^2 \leq \lambda_x^+(\mathbf{p}_i) \mathbf{x}^2 + \sum_{j \in J_i(\mathbf{x})} \lambda_j^+(\mathbf{p}_i) \mathbf{p}_j^2 &\iff \mathbf{x}^2 \geq \frac{1}{\lambda_x^+(\mathbf{p}_i)} (\mathbf{p}_i^2 - \sum_{j \in J_i(\mathbf{x})} \lambda_j^+(\mathbf{p}_i) \mathbf{p}_j^2) \\ \mathbf{x}^2 = \left(\sum_{i \in I_1(\mathbf{x})} \lambda_i(\mathbf{x}) \mathbf{p}_i \right)^2 &\leq \sum_{i \in I_1(\mathbf{x})} \lambda_i(\mathbf{x}) \mathbf{p}_i^2. \end{aligned} \quad (3.21)$$

⁴Jensen's inequality about the convex combination of a convex function f says that $f(\sum_i s_i x_i) \leq \sum_i s_i f(x_i)$ where s_i are positive real numbers with $\sum_i s_i = 1$.

The errors of both inequalities with respect to \mathbf{x}^2 are

$$e_x = \sum_{i \in I_1(\mathbf{x})} \lambda_i(\mathbf{x}) \mathbf{p}_i^2 - \mathbf{x}^2, \quad \iff \quad \mathbf{x}^2 = \sum_{i \in I_1(\mathbf{x})} \lambda_i(\mathbf{x}) \mathbf{p}_i^2 - e_x$$

$$e_i = \mathbf{x}^2 - \frac{1}{\lambda_x^+(\mathbf{p}_i)} (\mathbf{p}_i^2 - \sum_{j \in J_i(\mathbf{x})} \lambda_j^+(\mathbf{p}_i) \mathbf{p}_j^2) \quad \iff \quad \mathbf{x}^2 = \frac{1}{\lambda_x^+(\mathbf{p}_i)} (\mathbf{p}_i^2 - \sum_{j \in J_i(\mathbf{x})} \lambda_j^+(\mathbf{p}_i) \mathbf{p}_j^2) + e_i.$$

From Inequalities 3.21, we know that e_x and e_i are positive. We combine the equalities for \mathbf{x}^2 to obtain

$$\mathbf{x}^2 = \frac{e_i}{e_i + e_x} \left(\sum_{i \in I_1(\mathbf{x})} \lambda_i(\mathbf{x}) \mathbf{p}_i^2 \right) + \frac{e_x}{e_i + e_x} \left(\frac{1}{\lambda_x^+(\mathbf{p}_i)} (\mathbf{p}_i^2 - \sum_{j \in J_i(\mathbf{x})} \lambda_j^+(\mathbf{p}_i) \mathbf{p}_j^2) \right) =: \sum_{j \in I_1(\mathbf{x}) \cup J_i(\mathbf{x})} \gamma_j^i(\mathbf{x}) \mathbf{p}_j^2. \quad (3.22)$$

The coordinate function γ_j^i , $i \in I_1(\mathbf{x}), j \in J_i(\mathbf{x}) \cup I_1(\mathbf{x})$, is obtained by re-arranging all terms concerning the same sample point \mathbf{p}_j :

$$\gamma_i^i(\mathbf{x}) = \frac{e_i}{e_i + e_x} \lambda_i(\mathbf{x}) + \frac{e_x}{e_i + e_x} \frac{1}{\lambda_x^+(\mathbf{p}_i)} \quad (3.23)$$

$$\gamma_j^i(\mathbf{x}) = \frac{e_i}{e_i + e_x} \lambda_j(\mathbf{x}) - \frac{e_x}{e_i + e_x} \frac{\lambda_j^+(\mathbf{p}_i)}{\lambda_x^+(\mathbf{p}_i)}, \quad i \neq j. \quad (3.24)$$

Notice that $\lambda_j(\mathbf{x}) = 0$ if \mathbf{p}_j is not a natural neighbor of \mathbf{x} , in particular if $j \in J_i(\mathbf{x}) \setminus I_1(\mathbf{x})$, and $\lambda_j^+(\mathbf{p}_i) = 0$ if \mathbf{p}_j is not a Delaunay neighbor of \mathbf{p}_i in $\text{Del}^+(\mathcal{P})$, hence for $j \in I_1(\mathbf{x}) \setminus J_i(\mathbf{x})$.

The sum of the coordinate functions is one because, it is a convex combination of two systems of coordinates (which satisfy Definition 3.1):

$$\begin{aligned} \sum_{j \in J_i(\mathbf{x}) \cup I_1(\mathbf{x})} \gamma_j^i(\mathbf{x}) &= \frac{e_i}{e_i + e_x} \left(\sum_{j \in I_1(\mathbf{x})} \lambda_j(\mathbf{x}) \right) + \frac{e_x}{e_i + e_x} \frac{1}{\lambda_x^+(\mathbf{p}_i)} (1 - \sum_{j \in J_i(\mathbf{x})} \lambda_j^+(\mathbf{p}_i)) \\ &= \frac{e_i}{e_i + e_x} + \frac{e_x}{e_i + e_x} \frac{1}{\lambda_x^+(\mathbf{p}_i)} (1 - (1 - \lambda_x^+(\mathbf{p}_i))) = 1. \end{aligned}$$

For the same reason, the barycentric coordinate property is respected. It is a convex combination of Equations 3.19 and 3.20:

$$\sum_{j \in I_1(\mathbf{x}) \cup J_i(\mathbf{x})} \gamma_j^i(\mathbf{x}) \mathbf{p}_j = \frac{e_i}{e_i + e_x} \left(\sum_{j \in I_1(\mathbf{x})} \lambda_j(\mathbf{x}) \mathbf{p}_j \right) + \frac{e_x}{e_i + e_x} \left(\frac{1}{\lambda_x^+(\mathbf{p}_i)} (\mathbf{p}_i - \sum_{j \in J_i(\mathbf{x})} \lambda_j^+(\mathbf{p}_i) \mathbf{p}_j) \right) = \mathbf{x}.$$

We show that, at $\mathbf{x} = \mathbf{p}_k \in \mathcal{P}$, the coordinate function evaluates to

$$\gamma_j^i(\mathbf{p}_k) = \delta_{jk} \quad \forall i \in I_1(\mathbf{x}), j \in J_i(\mathbf{x}) \cup I_1(\mathbf{x}) :$$

Since $\lim_{\mathbf{x} \rightarrow \mathbf{p}_k} e_x = 0$ while $\lim_{\mathbf{x} \rightarrow \mathbf{p}_k} e_i > 0$ for $i \neq k$ and with $\lambda_j(\mathbf{p}_k) = 0$ and $\lambda_k(\mathbf{p}_k) = 1$, we confirm the result for $i \neq k$.

For $i = k = j$, $\lambda_k(\mathbf{p}_k) = 1$ and $\lim_{\mathbf{x} \rightarrow \mathbf{p}_k} \frac{1}{\lambda_x^+(\mathbf{p}_k)} = 1$, so $\gamma_k^k(\mathbf{p}_k) = 1$. For $i = k$ and $j \neq k$, $\lambda_j(\mathbf{p}_k) = 0$ and $\lim_{\mathbf{x} \rightarrow \mathbf{p}_k} \frac{\lambda_j^+(\mathbf{p}_k)}{\lambda_x^+(\mathbf{p}_k)} = 0$, so $\gamma_j^k(\mathbf{p}_k) = 0$.

For the final result, we combine the coordinate functions γ_j^i for every natural neighbor \mathbf{p}_i of \mathbf{x} , $i \in I_1(\mathbf{x})$, weighted with $\lambda_i(\mathbf{x})$. The following equalities are satisfied:

$$\begin{aligned}
\sum_{i \in I_1(\mathbf{x}), j \in I_2(\mathbf{x})} \lambda_i(\mathbf{x}) \gamma_j^i(\mathbf{x}) &= \sum_{i \in I_1(\mathbf{x})} \lambda_i(\mathbf{x}) \left(\sum_{j \in I_2(\mathbf{x})} \gamma_j^i(\mathbf{x}) \right) = \sum_{i \in I_1(\mathbf{x})} \lambda_i(\mathbf{x}) = 1 \\
\sum_{i \in I_1(\mathbf{x}), j \in I_2(\mathbf{x})} \lambda_i(\mathbf{x}) \gamma_j^i(\mathbf{x}) \mathbf{p}_j &= \mathbf{x} \\
\sum_{i \in I_1(\mathbf{x}), j \in I_2(\mathbf{x})} \lambda_i(\mathbf{x}) \gamma_j^i(\mathbf{x}) \mathbf{p}_j \mathbf{p}_j^t &= \mathbf{x}^t \mathbf{x}. \quad (3.25)
\end{aligned}$$

The definition of the non-convex natural neighbor coordinates depends only on the local coordinate property of the natural neighbor coordinates. The same techniques can be applied to natural neighbor coordinates in power diagrams because they also satisfy this property as shown in Section 3.2.

Chapter 4

Scattered data interpolation with natural neighbor coordinates

In this chapter, we describe several methods for scattered data interpolation that are based on natural neighbor coordinates. The methods are originally defined using natural neighbor coordinates from Voronoi diagrams. They are based on the fact that natural neighbor coordinates are local and constitute a coordinate system as described in Definition 3.1. Again, our exposition uses natural neighbor coordinates for Voronoi diagrams as in the original works. The same techniques can be applied to natural neighbor coordinates in power diagrams since we will see in the sequel that the methods are based on the barycentric coordinate property (Definition 3.1 (i)) and on the continuity and differential properties of the coordinate function. In Section 3.2 and in Lemma 3.6, we have shown that these properties are equally satisfied by the natural neighbor coordinate system defined in power diagrams.

We re-call that $\mathcal{P} = \{\mathbf{p}_1, \dots, \mathbf{p}_n\}$ is a set of n point in \mathbb{R}^d and $\text{Vor}(\mathcal{P})$ its Voronoi diagram. Since the weight is constantly zero, we denote the coordinate function by $\lambda_i(\mathbf{x})$, the volume function by $\pi_i(\mathbf{x})$, and the gradient from Lemma 3.9 by $\nabla \pi(\mathbf{x}) = \frac{\partial}{\partial \mathbf{x}} \pi((\mathbf{x}, 0))$. Let Φ be a scalar function defined on the convex hull of \mathcal{P} . We assume that the function values are known at the points of \mathcal{P} , i.e. to each $\mathbf{p}_i \in \mathcal{P}$, we associate $z_i = \Phi(\mathbf{p}_i)$. Sometimes, the gradient of Φ is also known at \mathbf{p}_i . We denote it $\mathbf{g}_i = \nabla \Phi(\mathbf{p}_i)$. The interpolation is carried out for a point \mathbf{x} in the convex hull of \mathcal{P} .

We overview the results presented in this chapter. We define two linear interpolants, Sibson's Z^0 interpolant (in Section 4.1) and Clarkson's K^0 interpolant (in Section 4.3). For all other interpolants, the gradients \mathbf{g}_i are assumed to be part of the input. Two interpolants, notably Sibson's Z^1 interpolant and Farin's interpolant, are C^1 continuous. Sibson's Z^1 interpolant is described in Section 4.1. It re-produces exactly spherical quadrics.¹ The interpolant is C^1 continuous by definition everywhere except at the data points \mathcal{P} (because all of its components including the natural coordinate functions are C^1 continuous). The first terms of the Taylor expansion of Z^1 at a data point \mathbf{p}_j yield the proof of the C^1 continuity at $\mathbf{p}_j \in \mathcal{P}$.

¹A spherical quadric has the form $\Phi(\mathbf{x}) = a + \mathbf{b}^t \mathbf{x} + \gamma \mathbf{x}^t \mathbf{x}$.

Farin's interpolant is summarized in Section 4.2. It is based on Bernstein-Bézier techniques. In opposition to Sibson's interpolant, it re-produces general quadratic functions. Again, the C^1 continuity is obvious everywhere apart from the data points. There, it is obtained by a careful choice of the control net of the Bézier simplex. At last, we present an interpolant called I^1 that is designed to re-produce quadratic functions. We bound its error with respect to the interpolation of general functions by looking at the first terms of the Taylor expansion of I^1 around the point \mathbf{x} .




Figure 4.1: The example data set used in this chapter.

Throughout this chapter, we use the same data set to illustrate the results of the different interpolation methods. It contains four non-zero function values and nine zero function values. The data set is depicted in Figure 4.1. All methods are applied with gradient $\mathbf{0}$ at the data points. In the function graph, the data points correspond to the highest and the lowest points. We do not indicate them further because we want to show the C^1 continuity or discontinuity of the interpolant at these points. A second example is depicted if an interpolation method relies on the gradient. We apply the same data set given the gradients that are indicated by the horizontal lines in Figure 4.1. In this case, the data points are indicated in the graph.

4.1 Sibson's natural neighbor interpolants

Sibson's natural neighbor interpolants have found applications in a variety of areas such as geology, geophysics [89], and partial differential equations [27]. Dave Watson implemented Sibson's interpolants [99], [100] and his code has, for example, been integrated into the NCAR Graphics Software package for scientific visualization and contouring.²

Linear precision interpolation

Sibson [91] defines a very simple interpolant that re-produces linear functions exactly. The interpolation of $\Phi(\mathbf{x})$ is given as the linear combination of the function values at the natural neighbors weighted by the

²<http://ngwww.ucar.edu/ng4.3/index.html>

coordinates:

$$Z^0(\mathbf{x}) = \sum_i \lambda_i(\mathbf{x}) z_i.$$

Indeed, if $z_i = a + \mathbf{b}^t \mathbf{p}_i$ for all natural neighbors of \mathbf{x} , we have

$$Z^0(\mathbf{x}) = \sum_i \lambda_i(\mathbf{x}) (a + \mathbf{b}^t \mathbf{p}_i) = a + \mathbf{b}^t \mathbf{x}$$

by the barycentric coordinate property.

Sibson's Z^1 interpolant

The so-called Z^1 interpolant proposed by Sibson [91] is C^1 continuous with gradient \mathbf{g}_i at \mathbf{p}_i . It reproduces spherical quadrics of the form $\Phi(\mathbf{x}) = a + \mathbf{b}^t \mathbf{x} + \gamma \mathbf{x}^t \mathbf{x}$ exactly. The proof relies on the barycentric coordinate property of the natural neighbor coordinates and assumes that the gradient of Φ at the data points is known or approximated from the function values as described in [91]. However, in this context, we suppose that the function value z_i as well as the gradient \mathbf{g}_i of Φ is known at $\mathbf{p}_i \in \mathcal{P}$.

Sibson's Z^1 interpolant is a combination of the linear interpolant Z^0 and an interpolant ξ . The interpolant ξ is the weighted sum of the first degree functions

$$\xi_i(\mathbf{x}) = z_i + \mathbf{g}_i^t (\mathbf{x} - \mathbf{p}_i), \quad \xi(\mathbf{x}) = \frac{\sum_i \frac{\lambda_i(\mathbf{x})}{\|\mathbf{x} - \mathbf{p}_i\|} \xi_i(\mathbf{x})}{\sum_i \frac{\lambda_i(\mathbf{x})}{\|\mathbf{x} - \mathbf{p}_i\|}}.$$

Sibson observed that the combination of Z^0 and ξ reconstructs exactly a spherical quadric if they are mixed as follows:

$$Z^1(\mathbf{x}) = \frac{\alpha(\mathbf{x}) Z^0(\mathbf{x}) + \beta(\mathbf{x}) \xi(\mathbf{x})}{\alpha(\mathbf{x}) + \beta(\mathbf{x})} \text{ where } \alpha(\mathbf{x}) = \frac{\sum_i \lambda_i(\mathbf{x}) \|\mathbf{x} - \mathbf{p}_i\|}{\sum_i \frac{\lambda_i(\mathbf{x})}{\|\mathbf{x} - \mathbf{p}_i\|}} \text{ and } \beta(\mathbf{x}) = \sum_i \lambda_i(\mathbf{x}) \|\mathbf{x} - \mathbf{p}_i\|^2.$$

Indeed, suppose $\Phi(\mathbf{x}) = a + \mathbf{b}^t \mathbf{x} + \gamma \mathbf{x}^t \mathbf{x}$ and $z_i = a + \mathbf{b}^t \mathbf{p}_i + \gamma \mathbf{p}_i^t \mathbf{p}_i$. Each of Z^0 and ξ yields an error with respect to Φ . We will show that

$$Z^0(\mathbf{x}) - \Phi(\mathbf{x}) = \gamma \beta(\mathbf{x}) \quad \text{and} \quad \Phi(\mathbf{x}) - \xi(\mathbf{x}) = \gamma \alpha(\mathbf{x}) \quad (4.1)$$

In Z^1 , by the combination of Z^0 and ξ , the error terms cancel out and $Z^1(X) = \Phi(X)$.

Using the barycentric coordinate property, the error of Z^0 with respect to Φ evaluates to

$$\begin{aligned} Z^0(\mathbf{x}) - \Phi(\mathbf{x}) &= \sum_i \lambda_i(\mathbf{x}) z_i - \Phi(\mathbf{x}) = \sum_i \lambda_i(\mathbf{x}) (a + \mathbf{b}^t \mathbf{p}_i + \gamma \mathbf{p}_i^t \mathbf{p}_i) - (a + \mathbf{b}^t \mathbf{x} + \gamma \mathbf{x}^t \mathbf{x}) \\ &= \sum_i \lambda_i(\mathbf{x}) \gamma \mathbf{p}_i^t \mathbf{p}_i - \gamma \mathbf{x}^t \mathbf{x} = \gamma \left(\sum_i \lambda_i(\mathbf{x}) \mathbf{p}_i^t \mathbf{p}_i \right) + \gamma (\mathbf{x}^t \mathbf{x} - 2 \mathbf{x}^t \mathbf{x}) \\ &= \gamma \sum_i \lambda_i(\mathbf{x}) (\mathbf{p}_i^t \mathbf{p}_i + \mathbf{x}^t \mathbf{x} - 2 \mathbf{x}^t \mathbf{p}_i) = \gamma \sum_i \lambda_i(\mathbf{x}) (\mathbf{x} - \mathbf{p}_i)^t (\mathbf{x} - \mathbf{p}_i) = \gamma \beta(\mathbf{x}). \end{aligned}$$

The error $\xi(\mathbf{x})$ with respect to $\Phi(\mathbf{x})$ depends on the error of $\xi_i(\mathbf{x})$ with respect to $\Phi(\mathbf{x})$,

$$\begin{aligned}\xi_i(\mathbf{x}) &= a + \mathbf{b}^t \mathbf{p}_i + \gamma \mathbf{p}_i^t \mathbf{p}_i + (\mathbf{b} + 2 \gamma \mathbf{p}_i)^t (\mathbf{x} - \mathbf{p}_i) = a + \mathbf{b}^t \mathbf{x} - \gamma \mathbf{p}_i^t \mathbf{p}_i + 2 \gamma \mathbf{p}_i^t \mathbf{x} \\ &= \Phi(\mathbf{x}) - \gamma \mathbf{x}^t \mathbf{x} - \gamma \mathbf{p}_i^t \mathbf{p}_i + 2 \gamma \mathbf{p}_i^t \mathbf{x} = \Phi(\mathbf{x}) - \gamma (\mathbf{x} - \mathbf{p}_i)^t (\mathbf{x} - \mathbf{p}_i).\end{aligned}$$

For each $\xi_i(\mathbf{x})$, we obtain $\Phi(\mathbf{x}) - \xi_i(\mathbf{x}) = \gamma \|\mathbf{p}_i - \mathbf{x}\|^2$, thus the error of $\xi(\mathbf{x})$ with respect to $\Phi(\mathbf{x})$ is $\Phi(\mathbf{x}) - \xi(\mathbf{x}) = \gamma \alpha(\mathbf{x})$. Equations (4.1) are henceforth proven.

Rote [87] pointed out that other weighting factors are possible in order to mix the first degree functions ξ_i . Instead of $\frac{\lambda_i(\mathbf{x})}{\|\mathbf{p}_i - \mathbf{x}\|}$, one might choose any smooth positive function f with $f(\mathbf{0}) = 0$ to weight the functions ξ_i with $\frac{\lambda_i(\mathbf{x})}{f(\|\mathbf{p}_i - \mathbf{x}\|)}$. The slope of f around the origin determines how fast the interpolant approaches ξ_i as $\mathbf{x} \rightarrow \mathbf{p}_i$.

In the general case, $\alpha(\mathbf{x})$ becomes

$$\alpha(\mathbf{x}) = \frac{\sum_i \lambda_i(\mathbf{x}) \frac{\|\mathbf{x} - \mathbf{p}_i\|^2}{f(\|\mathbf{x} - \mathbf{p}_i\|)}}{\sum_i \frac{\lambda_i(\mathbf{x})}{f(\|\mathbf{x} - \mathbf{p}_i\|)}}.$$

Figure 4.3 depicts the result of Sibson's interpolants on a data-set with only four non-zero values and gradient $\mathbf{0}$. These data points correspond to the highest points of each graph. In Figure 4.2, the gradient is directed from the data-point to the origin (the center of the interpolated zone) as shown in Figure 4.1. The data points correspond to the dark points indicated in the graph.

The gradient of Sibson's interpolant

In the sequel, we show that the Z^1 interpolant is indeed C^1 continuous. This has been stated in Sibson's original paper but the proof was omitted. We develop the proof following a suggestion of Rote [87]. Except on the data points themselves, the interpolant is differentiable since all ingredients, namely the coordinate function, the distance function, and ξ_i are differentiable. Therefore, we need to consider the gradient of Z^1 at a data point \mathbf{p}_j .

We determine the first two terms of the Taylor expansion of Z^1 at \mathbf{p}_j . The first order term yields the gradient of Z^1 at \mathbf{p}_j . In fact, we compute the Taylor expansions of Z^0 and ξ at \mathbf{p}_j and we bound $\alpha(\mathbf{x})$ and $\beta(\mathbf{x})$. Putting the approximations that we obtain in Equations (4.2), (4.3) and (4.4) together will yield the result.

In the Taylor expansion of ϕ at \mathbf{p}_j , we write $z_i = z_j + \mathbf{g}_j^t (\mathbf{p}_i - \mathbf{p}_j) + O(\|\mathbf{p}_i - \mathbf{p}_j\|^2)$ and

$$\begin{aligned}Z^0(\mathbf{x}) &= \sum_i \lambda_i(\mathbf{x}) z_i = \sum_i \lambda_i(\mathbf{x}) (z_j + \mathbf{g}_j^t (\mathbf{p}_i - \mathbf{p}_j) + O(\|\mathbf{p}_i - \mathbf{p}_j\|^2)) \\ &= z_j + \mathbf{g}_j^t (\mathbf{x} - \mathbf{p}_j) + \sum_i \lambda_i(\mathbf{x}) O(\|\mathbf{p}_i - \mathbf{p}_j\|^2).\end{aligned}$$

As \mathbf{x} approaches \mathbf{p}_j , the last term behaves as follows: from Lemma 3.14, we know that $\lambda_i(\mathbf{x}) = O(\|\mathbf{x} - \mathbf{p}_j\|)$ while the factor $O(\|\mathbf{p}_i - \mathbf{p}_j\|^2)$ is constant. We deduce that

$$Z^0(\mathbf{x}) = z_j + O(\|\mathbf{x} - \mathbf{p}_j\|), \text{ as } \mathbf{x} \rightarrow \mathbf{p}_j. \quad (4.2)$$

Concerning the interpolant ξ , we observe that it is continuous and differentiable at \mathbf{p}_j because of following two statements.

1. $\xi(\mathbf{p}_j) = \xi_j(\mathbf{p}_j) = z_j$,

because $\forall i \neq j : \lim_{\mathbf{x} \rightarrow \mathbf{p}_j} \frac{\lambda_i(\mathbf{x})}{\|\mathbf{x} - \mathbf{p}_i\|} = 0$, therefore, $\lim_{\mathbf{x} \rightarrow \mathbf{p}_j} \frac{\lambda_j(\mathbf{x})}{\|\mathbf{x} - \mathbf{p}_j\|} / \sum_i \frac{\lambda_i(\mathbf{x})}{\|\mathbf{x} - \mathbf{p}_i\|} = 1$.

2. $\nabla \xi(\mathbf{p}_j) = \mathbf{g}_j$:

we use Lemma 3.14 to show that the partial derivative of ξ at \mathbf{p}_j with respect to the k th coordinate corresponds to the k th coordinate of \mathbf{g}_j . This proves the claim.

Let \mathbf{i}_k , $k = 1, \dots, d$, be the unit vector along the k th coordinate axis, and let $\mathbf{x}(h) = \mathbf{p}_j + h \mathbf{i}_k$. By definition, $h = \|\mathbf{x}(h) - \mathbf{p}_j\|$.

$$\begin{aligned} \frac{\partial \xi}{\partial x_k}(\mathbf{p}_j) &= \lim_{h \rightarrow 0} \frac{\xi(\mathbf{x}(h)) - \xi(\mathbf{p}_j)}{h} = \lim_{h \rightarrow 0} \frac{\xi(\mathbf{x}(h)) - z_j}{h} \\ &= \lim_{h \rightarrow 0} \frac{\sum_i \frac{\lambda_i(\mathbf{x}(h))}{f(h)} \xi_i(\mathbf{x}(h)) - \left(\sum_i \frac{\lambda_i(\mathbf{x}(h))}{f(h)} \right) z_j}{h \sum_i \frac{\lambda_i(\mathbf{x}(h))}{f(h)}} \\ &= \lim_{h \rightarrow 0} \frac{\sum_{i \neq j} \frac{\lambda_i(\mathbf{x}(h))}{f(h)} \xi_i(\mathbf{x}(h)) + \frac{\lambda_j(\mathbf{x}(h))}{f(h)} \xi_j(\mathbf{x}(h)) - \frac{\lambda_j(\mathbf{x}(h))}{f(h)} \left(1 + \frac{f(h)}{\lambda_j(\mathbf{x}(h))} \sum_{i \neq j} \frac{\lambda_i(\mathbf{x}(h))}{f(h)} \right) z_j}{h \frac{\lambda_j(\mathbf{x}(h))}{f(h)} \left(1 + \frac{f(h)}{\lambda_j(\mathbf{x}(h))} \sum_{i \neq j} \frac{\lambda_i(\mathbf{x}(h))}{f(h)} \right)} \\ &= \lim_{h \rightarrow 0} \frac{\frac{f(h)}{\lambda_j(\mathbf{x}(h))} \sum_{i \neq j} \frac{\lambda_i(\mathbf{x}(h))}{f(h)} \xi_i(\mathbf{x}(h)) + z_j + \mathbf{g}_j^t (\mathbf{x}(h) - \mathbf{p}_j) - z_j \left(1 + \frac{f(h)}{\lambda_j(\mathbf{x}(h))} \sum_{i \neq j} \frac{\lambda_i(\mathbf{x}(h))}{f(h)} \right)}{h \left(1 + \frac{f(h)}{\lambda_j(\mathbf{x}(h))} \sum_{i \neq j} \frac{\lambda_i(\mathbf{x}(h))}{f(h)} \right)} \\ &= \lim_{h \rightarrow 0} \frac{\mathbf{g}_j^t (\mathbf{x}(h) - \mathbf{p}_j) + f(h) O(h)}{h (1 + O(h))} = \lim_{h \rightarrow 0} \mathbf{g}_j^t \mathbf{i}_k + f(h) O(1) = \mathbf{g}_j^t \mathbf{i}_k. \end{aligned}$$

We write the Taylor expansion of ξ at \mathbf{p}_j :

$$\xi(\mathbf{x}) = z_j + \mathbf{g}_j^t (\mathbf{x} - \mathbf{p}_j) + O(\|\mathbf{x} - \mathbf{p}_j\|^2). \quad (4.3)$$

For $\alpha(\mathbf{x})$, we obtain the following result by factorising the denominator and applying Lemma 3.14:

$$\begin{aligned} \alpha(\mathbf{x}) &= \frac{\frac{\lambda_j(\mathbf{x})}{f(\|\mathbf{x} - \mathbf{p}_j\|)} \|\mathbf{x} - \mathbf{p}_j\|^2 + \sum_{i \neq j} \frac{\lambda_i(\mathbf{x})}{f(\|\mathbf{x} - \mathbf{p}_i\|)} \|\mathbf{x} - \mathbf{p}_i\|^2}{\frac{\lambda_j(\mathbf{x})}{f(\|\mathbf{x} - \mathbf{p}_j\|)} \left(1 + \frac{f(\|\mathbf{x} - \mathbf{p}_j\|)}{\lambda_j(\mathbf{x})} \sum_{i \neq j} \frac{\lambda_i(\mathbf{x})}{f(\|\mathbf{x} - \mathbf{p}_i\|)} \right)} \\ &= \frac{\|\mathbf{x} - \mathbf{p}_j\|^2}{1 + f(\|\mathbf{x} - \mathbf{p}_j\|) O(\|\mathbf{x} - \mathbf{p}_j\|)} + \frac{f(\|\mathbf{x} - \mathbf{p}_j\|) O(\|\mathbf{x} - \mathbf{p}_j\|)}{1 + O(\|\mathbf{x} - \mathbf{p}_j\|)} \\ &= O(\|\mathbf{x} - \mathbf{p}_j\|^2) + f(\|\mathbf{x} - \mathbf{p}_j\|) O(\|\mathbf{x} - \mathbf{p}_j\|). \end{aligned}$$

If we assume that f is at least linear, then $\alpha(\mathbf{x}) = O(\|\mathbf{x} - \mathbf{p}_j\|^2)$. With Lemma 3.14, $\beta(\mathbf{x})$ is $O(\|\mathbf{x} - \mathbf{p}_j\|)$. We get

$$\frac{\alpha(\mathbf{x})}{\alpha(\mathbf{x}) + \beta(\mathbf{x})} = O(\|\mathbf{x} - \mathbf{p}_j\|) \text{ and } \frac{\beta(\mathbf{x})}{\alpha(\mathbf{x}) + \beta(\mathbf{x})} = \frac{1}{1 + O(\|\mathbf{x} - \mathbf{p}_j\|)} = 1 + O(\|\mathbf{x} - \mathbf{p}_j\|). \quad (4.4)$$

Putting the approximations of Equations (4.2), (4.3) and (4.4) together yields

$$\begin{aligned} Z^1(\mathbf{x}) &= \frac{\alpha(\mathbf{x})}{\alpha(\mathbf{x}) + \beta(\mathbf{x})} Z^0(\mathbf{x}) + \frac{\beta(\mathbf{x})}{\alpha(\mathbf{x}) + \beta(\mathbf{x})} \xi(\mathbf{x}) \\ &= \frac{\alpha(\mathbf{x})}{\alpha(\mathbf{x}) + \beta(\mathbf{x})} (z_j + O(\|\mathbf{x} - \mathbf{p}_j\|)) + \frac{\beta(\mathbf{x})}{\alpha(\mathbf{x}) + \beta(\mathbf{x})} (z_j + \mathbf{g}_j^t (\mathbf{x} - \mathbf{p}_j) + O(\|\mathbf{x} - \mathbf{p}_j\|^2)) \\ &= z_j + \mathbf{g}_j^t (\mathbf{x} - \mathbf{p}_j) + O(\|\mathbf{x} - \mathbf{p}_j\|^2), \quad \text{as } \mathbf{x} \rightarrow \mathbf{p}_j. \end{aligned}$$

We deduce the gradient at \mathbf{p}_j ,

$$\nabla Z^1(\mathbf{p}_j) = \mathbf{g}_j.$$




Figure 4.2: The result of Sibson's Z^1 interpolants on 4 non-zero values with gradient directed to the center of the figure. From top left to bottom right: Z^1 interpolants with (a) $f(\|\mathbf{x} - \mathbf{p}_i\|) = \sqrt{\|\mathbf{p}_i - \mathbf{x}\|}$, (b) $f(\|\mathbf{x} - \mathbf{p}_i\|) = \|\mathbf{x} - \mathbf{p}_i\|$, the "classical" Z^1 interpolant, (c) $f(\|\mathbf{x} - \mathbf{p}_i\|) = \|\mathbf{x} - \mathbf{p}_i\|^2$, (d) $f(\|\mathbf{x} - \mathbf{p}_i\|) = \|\mathbf{x} - \mathbf{p}_i\|^4$.




Figure 4.3: The result of Sibson's interpolants on 4 non-zero values with gradient $\mathbf{0}$. From top left to bottom right: (a) Z^0 interpolant, Z^1 interpolants with (b) $f(\|\mathbf{x} - \mathbf{p}_i\|) = \sqrt{\|\mathbf{p}_i - \mathbf{x}\|}$, (c) $f(\|\mathbf{x} - \mathbf{p}_i\|) = \|\mathbf{x} - \mathbf{p}_i\|$, the "classical" Z^1 interpolant, (d) $f(\|\mathbf{x} - \mathbf{p}_i\|) = \|\mathbf{p}_i - \mathbf{x}\|^{1.5}$, (e) $f(\|\mathbf{x} - \mathbf{p}_i\|) = \|\mathbf{p}_i - \mathbf{x}\|^2$, (f) $f(\|\mathbf{x} - \mathbf{p}_i\|) = \|\mathbf{p}_i - \mathbf{x}\|^4$.

4.2 Farin's C^1 interpolant

Farin [56] extended Sibson's work and realizes a C^1 continuous interpolant by embedding natural neighbor coordinates in the Bernstein-Bézier representation of a cubic simplex. Bernstein-Bézier patches and related concepts are widely used in the area of computer-aided geometric design and surface approximation. An introduction to this theory can, for instance, be found in [57] or [38]. Farin's interpolant has found applications for example in partial differential equations [92] and in free form deformations [81] used for hand simulation. Farin defines it only for the two-dimensional case.

In Section 3, we noticed that all points within a cell of the arrangement of Delaunay spheres have the same natural neighbors. Farin shows that such a set of natural neighbors can be used as a support for a multivariate Bézier simplex of arbitrary degree. Multivariate Bézier polynomials over higher-dimensional simplices were formally defined by de Boor [37]. According to the degree of the polynomial, a Bézier control net must be built from the vertices of the simplex, i.e. from the natural neighbors in our case. Any location inside the convex hull of the natural neighbors can be expressed in terms of multivariate Bézier polynomials, and the interpolant over a cell of the arrangement of Delaunay spheres is constructed based on these polynomials.

Assume that $\mathbf{x} \in \text{conv}(\mathcal{P})$ has $l + 1$ natural neighbors, $\mathcal{N}_x = \{\mathbf{p}_0, \dots, \mathbf{p}_l\}$, and the associated natural coordinates $\mathbf{u} = (\lambda_0(\mathbf{x}), \lambda_1(\mathbf{x}), \dots, \lambda_l(\mathbf{x}))$. We consider the points of \mathcal{N}_x as the projection of the vertices of an l -simplex σ onto the plane. Each point \mathbf{x} in the plane which has \mathcal{N}_x as natural neighbors corresponds to a unique point \mathbf{u} in \mathbb{R}^l because it can be written as a unique linear combination of the $l + 1$ natural neighbors \mathcal{N}_x . Therefore, we can define a Bernstein-Bézier polynomial over this simplex of \mathbb{R}^l .

We first define the multivariate Bernstein polynomials (after [37]) which are the basic elements of this interpolation method. Let $I = (i_0, i_1, \dots, i_l)$ be a *multi-index*, i.e. an $(l + 1)$ -tuple of indices. $|I| = i_0 + i_1 + \dots + i_l$ denotes its norm, and $I/m = (\frac{i_0}{m}, \dots, \frac{i_l}{m})$. The multi-index with $i_j = m$ and all other entries $i_i = 0, i \neq j$ is denoted I_j . Given the barycentric coordinates \mathbf{u} over an l -dimensional simplex σ (notably, defined by \mathcal{N}_x), the *Bernstein polynomials* of degree m are defined as:

$$B_I^m(\mathbf{u}) = \binom{m}{I} \lambda_0^{i_0} \lambda_1^{i_1} \dots \lambda_l^{i_l}, \text{ where } \binom{m}{I} = \frac{m!}{i_0! i_1! \dots i_l!}.$$

They are monomials of degree m . They have the property $\sum_{|I|=m} B_I^m(\mathbf{u}) = (\lambda_0 + \lambda_1 + \dots + \lambda_l)^m = 1$. Moreover, $B_{I_j}^m(\mathbf{u}) = \lambda_j^m$.

An l -variate *Bézier simplex* (also called a *Bernstein-Bézier polynomial*) interpolates a scalar function (defined in \mathbb{R}^l) that is known at the vertices of the l -dimensional simplex σ . It is determined by the Bernstein polynomials and by its *Bézier control net*. This is a network of $((l + 1)$ -dimensional) *control points* each composed of a (*Bézier*) *base point* (also called *abscissa* in the relevant literature) – that is its projection in \mathbb{R}^l – as well as of a (*Bézier*) *height* (also called *ordinate*) – which is a real value determining

its height over \mathbb{R}^l . The *corners* of the control net are the vertices of the simplex σ lifted by their function value to \mathbb{R}^{l+1} .

A Bézier simplex of degree one is the linear interpolation of the corners of the control net. For a simplex of higher degree, there are additional control points whose base points are linear combinations of the vertices of the simplex.

For each multi-index I with norm $|I| = m$, there exist exactly one control point $I/m = (\mathbf{q}_I, b_I)$ in the control net. If the vertex set of σ is $\{\mathbf{p}_0, \dots, \mathbf{p}_l\}$, the Bézier base point \mathbf{q}_I is the weighted barycenter

$$\mathbf{q}_I = \sum_{j=1}^l \frac{i_j}{m} \mathbf{p}_j. \quad (4.5)$$

For a quadratic polynomial ($m = 2$), the Bézier base points are the vertices plus the midpoints of the edges formed by all pairs of vertices. A cubic polynomial ($m = 3$) has three type of base points which are the weighted barycenter of one, two or three vertices. The base points of a Bézier tetrahedron include $\mathbf{q}_{(3000)} = \mathbf{p}_1$, $\mathbf{q}_{(1200)} = \frac{1}{3}\mathbf{p}_1 + \frac{2}{3}\mathbf{p}_2$ and $\mathbf{q}_{(1110)} = \frac{1}{3}\mathbf{p}_1 + \frac{1}{3}\mathbf{p}_2 + \frac{1}{3}\mathbf{p}_3$. The other base points are obtained by permuting the indices (and replacing the corresponding data points). Figure 4.4 depicts the projection of the control net of a quadratic ($m = 2$) and a cubic ($m = 3$) Bézier tetrahedron onto the plane.




Figure 4.4: $l = 4$ and (a) $m = 2$: projection of the control net of a quadratic Bézier tetrahedron, (b) $m = 3$: projection of the control net of a cubic Bézier tetrahedron (Farin [56]).

The corner that corresponds to the data point \mathbf{p}_j is $I_j/m = (\mathbf{p}_j, z_j)$, thus, its height is $b_{I_j} = z_j$. The way that the heights of the remaining control points are chosen is not fixed a priori. It depends on the properties that are expected from the interpolant, e.g. its continuity properties.

The interpolant of the function Φ for which we know the function values $z_i = \Phi(\mathbf{p}_i)$ for all $\mathbf{p}_i \in \mathcal{P}$ is

the l -variate Bézier simplex f^m constructed over \mathcal{N}_x of degree m . It is defined by

$$f^m(\mathbf{u}) = \sum_{|I|=m} b_I B_I^m(\mathbf{u}),$$

where B_I^m are the Bernstein polynomials of degree m and b_I is the Bézier height of the control point I/m . The one-variate Bézier simplex f^1 corresponds to Sibson's Z^0 interpolant,

$$f^1(\mathbf{u}) = \sum_{i=1}^n \lambda_i(\mathbf{x}) z_i,$$

because $B_{I_j}^1(\mathbf{u}) = \lambda_j(\mathbf{x})$ and $b_{I_j} = z_j$ for $j = 1, \dots, n$.

The interpolant f^m is continuously differentiable except on the data points because the natural neighbor coordinates are not differentiable there. However, Farin remarks in [56] that the directional derivatives of f^m at (\mathbf{p}_i, z_i) are linear combinations of the control points that are directly connected to (\mathbf{p}_i, z_i) in the control net. Thus, f^m is differentiable at (\mathbf{p}_i, z_i) iff all these control points are coplanar.

The m control points connected to (\mathbf{p}_i, z_i) have the i th entry $j_i = m - 1$ and one other entry 1. They form the l -ring neighbors of (\mathbf{p}_i, z_i) . See also Figure 4.4.

Assume we know the gradient \mathbf{g}_i of Φ at \mathbf{p}_i . We might choose the Bézier height of the 1-ring neighbors of (\mathbf{p}_i, z_i) such that they lie in the tangent plane of Φ at \mathbf{p}_i (which is determined by \mathbf{g}_i). The derivative of the Bézier simplex then yields that the gradient of f^m at \mathbf{p}_i is \mathbf{g}_i . The interpolant is overall C^1 continuous, whatever the Bézier heights of the remaining control points are.

For a cubic polynomial ($m = 3$), Farin [56] determines the heights of the remaining control points in order to obtain an interpolant that re-produces quadratic functions. The same idea has been used before by the same author [55] for interpolation over triangular Bézier patches. It is based on so-called *degree elevation* (see also [67, chapter 4.4.6]) which means that a polynomial of degree $m - 1$ is re-written as a polynomial of degree m such that

$$\sum_{|I|=m} c_I B_I^m(\mathbf{u}) = \sum_{|J|=m-1} b_J B_J^{m-1}(\mathbf{u}). \quad (4.6)$$

By multiplying the right hand side of Equation 4.6 with $(\lambda_0 + \lambda_1 + \dots + \lambda_l) = 1$ and by gathering the appropriate terms, we find the new Bézier heights

$$c_I = \frac{1}{m} (i_1 b_{I-ei_1} + i_2 b_{I-ei_2} + \dots + i_n b_{I-ei_n}), \quad (4.7)$$

where $I - ei_j = (i_0, \dots, i_{j-1}, i_j - 1, i_{j+1}, \dots, i_l)$ is the n -tuple of indices I except that i_j is replaced by $i_j - 1$.

For each triplet of sites in \mathcal{N}_x , nine control points of the cubic polynomial are fixed by the data points and their 1-ring neighbors. The height of one control point (whose base point is the barycenter of the

three sites) is not determined. It will be fixed such that f^3 is the re-writing of a quadratic polynomial as in equation 4.6. To do so, Equations (4.7) can be transformed to express the coefficients b_J , $|J| = 2$, of a quadratic polynomial in dependence of the coefficients c_I , $|I| = 3$, of the cubic. The undetermined Bézier height is then, in turn, given by Equation (4.7).

Precisely, the solution is as follows. Let I_{ijk} be the subscript of a control point with undefined Bézier height. Its base point is the barycenter of \mathbf{p}_i , \mathbf{p}_j and \mathbf{p}_k . I_{ijk} has the three 1-entries $i_i = i_j = i_k = 1$ and all other entries are 0. There are nine subscripts with the same 0 entries as I_{ijk} : I_i , I_j and I_k which correspond to data points (with one 3 entry and all others 0) and six other indices which correspond to the 1-ring neighbors of \mathbf{p}_i , \mathbf{p}_j and \mathbf{p}_k (with one 2 and one 1 entry and all other entries 0). Let a be the average of the heights corresponding to the data points, $a = \frac{1}{3}(b_{I_i} + b_{I_j} + b_{I_k})$, and b be the average of the heights corresponding to their 1-ring neighbors. If we choose the Bézier height

$$b_{I_{ijk}} = \frac{3}{2}b - \frac{1}{2}a,$$

f^3 reproduces quadratic functions of the form $\Phi(\mathbf{x}) = a + b^t \mathbf{x} + \mathbf{x}^t \mathcal{Q} \mathbf{x}$ with $\mathcal{Q} = \begin{pmatrix} c & d \\ 0 & f \end{pmatrix}$.

The cubic polynomial f^3 is actually quadratic, it interpolates the function values and the gradients on the data points, and is globally C^1 continuous. Since the six parameters of Φ are determined by three function values and three gradients, the quadratic function f^3 is unique.

Two examples are depicted in Figure 4.5.




Figure 4.5: The result of Farin's interpolants on 4 non-zero values with (a) gradient 0, (b) gradient directed to the center of the image. In the left picture, the non-zero data points correspond to the four peaks.

4.3 Clarkson's K^0 interpolant

Another way to construct a C^1 smooth interpolant using natural neighbor coordinates is based on an unpublished idea of Clarkson that is described in [35]. In this thesis, they are introduced in Section 3.5.

The interpolant defined by this new coordinate system is

$$K^0(\mathbf{x}) = \sum_{i \in I_1(\mathbf{x}), j \in I_2(\mathbf{x})} \lambda_i(\mathbf{x}) \gamma_j^i(\mathbf{x}) z_j.$$

It reproduces spherical quadrics of the form $\Phi(\mathbf{x}) = a + \mathbf{b}^t \mathbf{x} + \gamma \mathbf{x}^t \mathbf{x}$ – without knowing the gradient of Φ at the data points as it was the case in the preceding sections. Indeed, if Φ is a spherical quadric, we write $z_j = a + \mathbf{b}^t \mathbf{p}_j + \gamma \mathbf{p}_j^t \mathbf{p}_j$. With equations 3.25, we obtain

$$K^0(\mathbf{x}) = \sum_{i \in I_1(\mathbf{x}), j \in I_2(\mathbf{x})} \lambda_i(\mathbf{x}) \gamma_j^i(\mathbf{x}) (a + \mathbf{b}^t \mathbf{p}_j + \gamma \mathbf{p}_j^t \mathbf{p}_j) = a + \mathbf{b}^t \mathbf{x} + \gamma \mathbf{x}^t \mathbf{x} = \Phi(\mathbf{x}).$$

Figure 4.6 shows the result of the K^0 interpolant on the example data set (see Figure 4.1). With this interpolant, the gradient is not known.




Figure 4.6: The result of Clarkson's interpolants on 4 non-zero values.

Even though the graph appears smooth at the data points, we were not able to show the global C^1 continuity of this interpolant.

4.4 Quadratic precision interpolant

In this section, we define an interpolant that re-produces quadratic functions, i.e. functions of the form $\Phi(\mathbf{x}) = a + \mathbf{b}^t \mathbf{x} + \mathbf{x}^t \mathbf{Q} \mathbf{x}$. If the interpolated function has higher order terms, say Φ is a C^k continuous function, $k \geq 2$, it is interpolated up to its second order terms when the distance between the sample point tends to zero, i.e. the density of the sampling tends to infinity. As additional input, we assume the gradients

$\mathbf{g}_i = \nabla(\Phi(\mathbf{p}_i))$ of Φ at all sample points $\mathbf{p}_i \in \mathcal{P}$ to be known. The interpolant is defined as follows:

$$I^1(\mathbf{x}) = \sum_i \lambda_i(\mathbf{x}) \left(z_i + \frac{1}{2} \mathbf{g}_i^t (\mathbf{x} - \mathbf{p}_i) \right)$$

We estimate the error of the interpolation by looking at the Taylor expansion of Φ at \mathbf{x} up to the second order terms. Let $\mathcal{H}(\Phi(x))$ denote the Hessian of Φ at x . We write

$$z_i = \Phi(\mathbf{x}) + \nabla(\Phi(\mathbf{x}))^t (\mathbf{p}_i - \mathbf{x}) + \frac{1}{2} (\mathbf{p}_i - \mathbf{x})^t \mathcal{H}(\Phi(\mathbf{x})) (\mathbf{p}_i - \mathbf{x}) + O(\|\mathbf{x} - \mathbf{p}_i\|^3)$$

$$\text{and } \mathbf{g}_i = \nabla(\Phi(\mathbf{x})) + \mathcal{H}(\Phi(\mathbf{x})) (\mathbf{p}_i - \mathbf{x}) + O(\|\mathbf{x} - \mathbf{p}_i\|^2) \mathbf{1}.$$

Applying the barycentric property to I^1 , we obtain

$$\begin{aligned} I^1(\mathbf{x}) &= \sum_i \lambda_i(\mathbf{x}) \left(\Phi(\mathbf{x}) + \frac{1}{2} \nabla(\Phi(\mathbf{x}))^t (\mathbf{p}_i - \mathbf{x}) + \frac{1}{2} (\mathbf{p}_i - \mathbf{x})^t \mathcal{H}(\Phi(\mathbf{x})) (\mathbf{p}_i - \mathbf{x}) \right. \\ &\quad \left. + \frac{1}{2} (\mathbf{p}_i - \mathbf{x})^t \mathcal{H}(\Phi(\mathbf{x})) (\mathbf{x} - \mathbf{p}_i) + O(\|\mathbf{x} - \mathbf{p}_i\|^3) + O(\|\mathbf{x} - \mathbf{p}_i\|^2) \mathbf{1}^t (\mathbf{x} - \mathbf{p}_i) \right) \\ &= \Phi(\mathbf{x}) + \frac{1}{2} \nabla(\Phi(\mathbf{x}))^t (\mathbf{x} - \mathbf{x}) + \sum_i \lambda_i(\mathbf{x}) O(\|\mathbf{x}_0 - \mathbf{p}_i\|^3) \\ &= \Phi(\mathbf{x}) + \sum_i \lambda_i(\mathbf{x}) O(\|\mathbf{p}_i - \mathbf{x}\|^3). \end{aligned}$$

From the above calculations, we deduce that, firstly, a quadratic function is exactly interpolated since the remainders of the Taylor expansion are zero. Secondly, as the sampling density tends to infinity, the second order terms of a general function Φ are exactly interpolated because the distance from \mathbf{x} to its neighbors \mathbf{p}_i vanishes while $\lambda_i(\mathbf{x})$ is finite. The interpolant is not C^1 continuous on \mathcal{P} because of the C^1 discontinuity of the coordinate functions λ_i on \mathcal{P} .

Figure 4.7 depicts the result of the I^1 interpolant on a data-set with only 4 non-zero values. The gradient of Φ at these points is directed from the data-point to the origin (the center of the interpolated zone). Notice, that we obtain respectively Sibson's Z^0 interpolant if the gradient is $\mathbf{0}$.




Figure 4.7: The result of the quadratic precision interpolant I^1 on four non-zero values with gradient directed to the center.

The gradient of the I^1 interpolant

In this section, we compute the gradient of the I^1 interpolant, in order to examine its value at \mathbf{p}_j . Let $h_i(\mathbf{x}) = z_i + \frac{1}{2}\mathbf{g}_i^t(\mathbf{x} - \mathbf{p}_i)$. The gradient of I^1 is as follows:

$$\nabla I^1(x) = \sum_{i=1}^n (\nabla \lambda_i(x) h_i(x) + \lambda_i(x) \nabla h_i(x)) = T_1 + T_2,$$

$$\text{where } T_1 = \sum_{i=1}^n \nabla \lambda_i(\mathbf{x}) h_i(x), \text{ and } T_2 = \sum_{i=1}^n \lambda_i(x) \frac{1}{2} \mathbf{g}_i, \text{ since } \nabla h_i(x) = \frac{1}{2} \mathbf{g}_i.$$

From $\sum_{i=1}^n \lambda_i(x) = 1$ and from $\sum_{i=0}^n \lambda_i(x) \mathbf{p}_i = \mathbf{x}$, we derive the two following equations concerning the gradients $\nabla \lambda_j(x)$ ($\mathbf{0}$ denotes the origin of \mathbb{R}^d):

$$\sum_{i=0}^n \nabla \lambda_i(x) = \mathbf{0}, \quad \sum_{i=0}^n \nabla \lambda_i(\mathbf{x})^t \mathbf{p}_i = 1. \quad (4.8)$$

We use the equations to re-arrange the writing of T_1 :

$$T_1 = \sum_{i=1, i \neq j}^n \nabla \lambda_i(\mathbf{x}) (h_i(\mathbf{x}) - h_j(\mathbf{x})),$$

and to notice that

$$\sum_{i \neq j} \nabla \lambda_i(x)^t (\mathbf{p}_i - \mathbf{p}_j) = 1 - \nabla \lambda_j(\mathbf{x})^t \mathbf{p}_j + \nabla \lambda_j(\mathbf{x})^t \mathbf{p}_j = 1 \quad \left(\text{with } \sum_{i \neq j} \nabla \lambda_i(x) = -\nabla \lambda_j(x) \right).$$

We consider now the case $\mathbf{x} \rightarrow \mathbf{p}_j$: from $\lim_{\mathbf{x} \rightarrow \mathbf{p}_j} \lambda_j(\mathbf{x}) = 1$ and $\lim_{\mathbf{x} \rightarrow \mathbf{p}_j} \lambda_i(\mathbf{x}) = 0, i \neq j$, it follows that $\lim_{\mathbf{x} \rightarrow \mathbf{p}_j} T_2 = \frac{1}{2} \mathbf{g}_j$. The analysis of the T_1 term is a bit more involved. We inject the Taylor expansion of z_i, z_j , and \mathbf{g}_i about \mathbf{x} shown in the preceding section:

$$\begin{aligned} \lim_{\mathbf{x} \rightarrow \mathbf{p}_j} T_1 &= \lim_{\mathbf{x} \rightarrow \mathbf{p}_j} \sum_{i=1, i \neq j}^n \nabla \lambda_i(\mathbf{x}) \left(z_i - z_j + \frac{1}{2} \mathbf{g}_i^t (\mathbf{x} - \mathbf{p}_i) \right) \quad \left(\text{with } \lim_{\mathbf{x} \rightarrow \mathbf{p}_j} (\mathbf{x} - \mathbf{p}_j) = \mathbf{0} \right) \\ &= \lim_{\mathbf{x} \rightarrow \mathbf{p}_j} \sum_{i=1, i \neq j}^n \nabla \lambda_i(\mathbf{x}) \left(\nabla \Phi(\mathbf{x})^t \left(\frac{1}{2} \mathbf{p}_i + \frac{1}{2} \mathbf{x} - \mathbf{p}_j \right) + (\mathbf{p}_i - \mathbf{x})^t \mathcal{H}(\Phi(\mathbf{x})) (\mathbf{p}_i - \mathbf{x}) \right. \\ &\quad \left. + O(\|\mathbf{p}_i - \mathbf{x}\|^3) \right). \end{aligned}$$

With Equations 4.8, we simplify $\lim_{\mathbf{x} \rightarrow \mathbf{p}_j} \sum_{i=1, i \neq j}^n \nabla \lambda_i(\mathbf{x}) \left(\frac{1}{2} \mathbf{p}_i + \frac{1}{2} \mathbf{x} - \mathbf{p}_j \right) = \frac{1}{2} (1 - \nabla \lambda_j(\mathbf{x})) - \frac{1}{2} \nabla \lambda_j(\mathbf{x}) + \nabla \lambda_j(\mathbf{x}) = \frac{1}{2}$. Thus,

$$\lim_{\mathbf{x} \rightarrow \mathbf{p}_j} T_1 = \frac{1}{2} \mathbf{g}_j + \lim_{\mathbf{x} \rightarrow \mathbf{p}_j} \sum_{i=1, i \neq j}^n \nabla \lambda_i(\mathbf{x}) \left((\mathbf{p}_i - \mathbf{x})^t \mathcal{H}(\Phi(\mathbf{x})) (\mathbf{p}_i - \mathbf{x}) + O(\|\mathbf{p}_i - \mathbf{x}\|^3) \right).$$

Because we know from Lemma 3.14 that $\lim_{\mathbf{x} \rightarrow \mathbf{p}_j} \|\nabla \lambda_i(\mathbf{x})\|$ is of order $O(1)$, we deduce that

$$\lim_{\mathbf{x} \rightarrow \mathbf{p}_j} \nabla I^1(\mathbf{x}) = \mathbf{g}_j + O(\|\mathbf{p}_i - \mathbf{x}\|^2) \vec{\mathbf{e}}$$

where \mathbf{p}_i is the furthest natural neighbor of \mathbf{x} and $\vec{\mathbf{e}}$ is a unit vector of \mathbb{R}^d .

Chapter 5

Sampled manifolds: definitions and results

In this chapter, we consider the case where the sample points \mathcal{P} are taken from a smooth, closed k -manifold \mathcal{M} . In the first section, we introduce some notations concerning the definition of the manifold and the assumptions we make. Then, we define the restriction of a Voronoi diagram to a manifold, and finally we are concerned with the sampling density of \mathcal{M} , and recall several results about the local behavior of the manifold samples. Apart from very few lemmas that we have adapted or derived, this chapter contains no proofs as it recalls known concepts and results. The reader may refer to the original literature for further details.

5.1 Topological notions, medial axis and local feature size

Two topological spaces are called *homeomorphic* if there exists a bijective map from one space to the other that is continuous in both directions. A *neighborhood* of \mathbf{x} in a set \mathcal{M} is an open sub-set of \mathcal{M} that contains \mathbf{x} . A *topological k -ball* is a subset of \mathbb{R}^d that is homeomorphic to \mathbb{R}^k .

A *homotopy* between two functions f and g from a space \mathcal{U} to a space \mathcal{V} is a continuous map $h : \mathcal{U} \times [0, 1] \rightarrow \mathcal{V}$ such that $h(x, 0) = f(x)$ and $h(x, 1) = g(x)$. Two functions f and g are *homotopic*, if there exists a homotopy h between them, i.e. if they can be continuously transformed from one to the other. Two topological spaces \mathcal{U} and \mathcal{V} are *homotopy equivalent* (or of the same *homotopy type*) if one can find two continuous maps $f : \mathcal{U} \rightarrow \mathcal{V}$ and $g : \mathcal{V} \rightarrow \mathcal{U}$ such that $g \circ f : \mathcal{U} \rightarrow \mathcal{U}$ is homotopic to the identity map in \mathcal{U} and $f \circ g : \mathcal{V} \rightarrow \mathcal{V}$ is homotopic to the identity map in \mathcal{V} . \mathcal{U} is *contractible* if it is homotopy equivalent to a point.

$\mathcal{M} \subseteq \mathbb{R}^d$ is a *closed k -manifold* of \mathbb{R}^d if the neighborhood of each point $\mathbf{x} \in \mathcal{M}$ is homeomorphic to \mathbb{R}^k . It is a *k -manifold of \mathbb{R}^d with boundary* if the neighborhood of each point $\mathbf{x} \in \mathcal{M}$ is either homeomorphic to \mathbb{R}^k (an *interior* point) or homeomorphic to a closed half-space of \mathbb{R}^k (a *boundary* point). The *boundary* of a k -manifold \mathcal{F} , i.e. the set of boundary points of \mathcal{F} , is either empty or it is a $(k-1)$ -manifold. Let a $(d-1)$ -manifold be called a *hypersurface*.

The *tangent space* of \mathcal{M} at \mathbf{x} is called \mathcal{T}_x . We define it via a parameterization of \mathcal{M} : Suppose that $f_x = (f_{x_1}, \dots, f_{x_k}) : \mathcal{U} \subset \mathbb{R}^k \rightarrow \mathbb{R}^d$ is a local parameterization of a \mathcal{M} . At the point $\mathbf{x} = f_x(\mathbf{0})$, the tangent space \mathcal{T}_x is the image of the Jacobian of f_x as a linear transformation from \mathbb{R}^k to \mathbb{R}^d .

We call \mathcal{N}_x the *normal space* of \mathcal{M} at \mathbf{x} , i.e. the affine sub-space of \mathbb{R}^d that is spanned by the unit normals $\{\vec{\mathbf{n}}_x^i\}_{i=1, \dots, d-k}$, at \mathbf{x} and that is orthogonal to \mathcal{T}_x . Since \mathcal{M} is a k -manifold of \mathbb{R}^d , \mathcal{T}_x is of dimension k and \mathcal{N}_x of dimension $d - k$.

At last, we need to introduce the notion of *subanalytic* sets. This theory allows for example to define precisely what is usually referred to as piecewise analytic sets. Some important properties of subanalytic sets will be used in Chapter 6 (and notably in Lemma 6.2). Chazal and Soufflet [32] claim that most general manifolds belong to this setting and consider it as a generalization of real semialgebraic geometry.

The definition is as follows (omitting some technical details). If a function is *analytic* in a region \mathcal{R} , it is infinitely differentiable in \mathcal{R} . A set $\mathcal{X} \subset \mathbb{R}^d$ is *semianalytic* if, for all $\mathbf{x} \in \mathcal{X}$, there is an open neighborhood \mathcal{U} of \mathbf{x} such that $\mathcal{X} \cap \mathcal{U}$ is a finite Boolean combination of sets $\{\bar{\mathbf{x}} \in \mathcal{U} | f(\bar{\mathbf{x}}) = 0\}$ and $\{\bar{\mathbf{x}} \in \mathcal{U} | g(\bar{\mathbf{x}}) > 0\}$, where $f, g : \mathcal{U} \rightarrow \mathbb{R}$ are analytic. Thus, such a set is locally defined by analytic equalities and inequalities. $\mathcal{X} \subset \mathbb{R}^d$ is *subanalytic* if, for all $\mathbf{x} \in \mathcal{X}$, there is an open set \mathcal{U} and a bounded semianalytic set $\mathcal{Y} \subset \mathbb{R}^{n+m}$ such that $\mathcal{X} \cap \mathcal{U}$ is the projection of \mathcal{Y} into \mathcal{U} . (We omit the exact definition of the properties of this projection.) This is a generalization of semianalytic sets (and it is, unlike the further, stable under linear projections). See [32] for a short description of the properties of subanalytic geometry.

We suppose in the sequel that \mathcal{M} is closed, compact and sufficiently smooth, i.e. three times differentiable. We suppose further that \mathcal{M} is subanalytic. This property is used in the proof of Lemma 6.2. Let $\mathcal{P} \subset \mathcal{M}$ be a set of sample points from \mathcal{M} . We denote $B(\mathbf{x}, r)$ an open ball centered at \mathbf{x} with radius r .

The following definitions allow to characterize some features of a manifold. See [101] and [9] for basic results on properties of the medial axis.

- Definition 5.1**
1. **(Medial axis)** The medial axis of a manifold \mathcal{M} in \mathbb{R}^d is the closure of the set of points with more than one closest point on \mathcal{M} .
 2. **(Medial ball)** A ball that is centered on the medial axis, tangent to the manifold, and whose interior does not intersect the manifold is called a medial ball.
 3. **(Local feature size)** The local feature size $\text{lfs}(\mathbf{x})$ of a point $\mathbf{x} \in \mathcal{M}$ is the distance from \mathbf{x} to the medial axis of \mathcal{M} .

We show that the local feature size is Lipschitz. The lemma has been first stated in [9, lemma 1].

Lemma 5.2 For any two points $\mathbf{x}, \mathbf{y} \in \mathcal{M}$, $\text{lfs}(\mathbf{x}) \leq \text{lfs}(\mathbf{y}) + \|\mathbf{x} - \mathbf{y}\|$.

Proof: Since the ball $B(\mathbf{y}, \text{lfs}(\mathbf{x}) - \|\mathbf{x} - \mathbf{y}\|)$ is contained in the ball $B(\mathbf{x}, \text{lfs}(\mathbf{x}))$, it does not contain any point of the medial axis. Thus, its radius must be smaller than $\text{lfs}(\mathbf{y})$, i.e. $\text{lfs}(\mathbf{x}) - \|\mathbf{x} - \mathbf{y}\| \leq \text{lfs}(\mathbf{y})$. \square

The following two propositions will be useful in Section 6.2.2 and Section 8.2. They have been stated by Boissonnat and Cazals [19] for the case $k = d - 1$.

Lemma 5.3 [19, Proposition 13]

Let $B = B(\mathbf{x}, r)$ be a ball that intersects \mathcal{M} . If $B \cap \mathcal{M}$ is not a topological ball, then B contains a point of the medial axis of \mathcal{M} .

Of course, the dimension of the topological ball is the dimension of the k -manifold \mathcal{M} . This observation generalizes [19, Proposition 14].

Lemma 5.4 For any $\mathbf{x} \in \mathcal{M}$ and any $r < \text{lfs}(\mathbf{x})$, $\mathcal{M} \cap B(\mathbf{x}, r)$ is a topological k -ball.

Proof: By definition of $\text{lfs}(\mathbf{x})$, $B(\mathbf{x}, r)$ cannot contain a point of the medial axis of \mathcal{M} . Lemma 5.3 implies that $\mathcal{M} \cap B(\mathbf{x}, r)$ is a topological k -ball. \square

At last in this section, we show a lemma that is needed for the proof of Lemma 6.2. It is topological result based on the work of Milnor [79]. The proof follows a suggestion of Chazal [33].

Lemma 5.5 Let \mathcal{M} be a compact k -manifold of \mathbb{R}^d and \mathcal{N}_x the normal space to \mathcal{M} at a point $\mathbf{x} \in \mathcal{M}$. There always exists a second point $\mathbf{x}' \in \mathcal{M} \cap \mathbf{N}_x$ other than \mathbf{x} .

Proof: Consider the orthogonal projection $f : \mathcal{M} \rightarrow \mathcal{T}_x$ from \mathcal{M} onto \mathcal{T}_x . Because \mathcal{M} is compact, there must exist an open subset of \mathcal{T}_x that is not in the image of f . Consequently, there exists a regular value \mathbf{y} of f without pre-image in \mathcal{M} , thus, $\#f^{-1}(\mathbf{y}) = 0$ where $\#f^{-1}(\mathbf{y})$ denotes the number of (distinct) \mathbf{p} such that $f(\mathbf{p}) = \mathbf{y}$. Milnor [79, §4] shows that the residue class modulo 2 of $\#f^{-1}(\mathbf{y})$ is independent of the regular value \mathbf{y} . It is called the *degree mod 2* of f . We deduce that f has degree 0 mod 2.

Since \mathcal{T}_x and \mathcal{N}_x are orthogonal and $\mathbf{x} = \mathcal{T}_x \cap \mathcal{N}_x$, the pre-image $f^{-1}(\mathbf{x})$ of \mathbf{x} are the intersection points of \mathcal{M} with \mathcal{N}_x . We prove by contradiction: Suppose \mathbf{x} were the only point in $\mathbf{N}_x \cap \mathcal{M}$. In this case, it is a regular value of f since \mathcal{M} intersects \mathcal{N}_x transversally in \mathbf{x} . However, by [79, §4], $\#f^{-1}(\mathbf{x}) \bmod 2 = 0$, thus, there must be an even number of points in $\mathbf{N}_x \cap \mathcal{M}$. This contradicts the assumption. \square

5.2 Voronoi diagram restricted to a k -manifold

We first define the Voronoi diagram of a set of points restricted to a k -manifold of \mathbb{R}^d , following previous work by Chew [34] (for the case $d = 3, k = 1, 2$) and Edelsbrunner and Shah [54].

Definition 5.6 Let \mathcal{P} be a set of points. The Voronoi diagram of \mathcal{P} restricted to \mathcal{M} is the (curved) cell complex obtained by intersecting each face of $\text{Vor}(\mathcal{P})$ with \mathcal{M} . We denote it by $\text{Vor}_{\mathcal{M}}(\mathcal{P})$.

We denote by $V_{\mathcal{M}}(\mathbf{p}_i)$ the Voronoi cell of $\text{Vor}_{\mathcal{M}}(\mathcal{P})$ consisting of the points of \mathcal{M} that are closer to $\mathbf{p}_i \in \mathcal{P}$ (for the Euclidean distance) than to any $\mathbf{p}_j \in \mathcal{P}, j \neq i$. A vertex of $V_{\mathcal{M}}(\mathbf{p}_i)$ is the intersection of a $(d - k)$ dimensional face of $V(\mathbf{p}_i)$ with \mathcal{M} . Hence, it is the center of a ball passing through $k + 1$ points of \mathcal{P} and not enclosing other points of \mathcal{P} . The dual of the restricted Voronoi diagram is defined as follows:

Definition 5.7 Let \mathcal{P} be a set of points. The Delaunay triangulation of \mathcal{P} restricted to \mathcal{M} is the subcomplex of $\text{Del}(\mathcal{P})$ consisting of the faces of $\text{Del}(\mathcal{P})$ whose dual Voronoi faces intersect \mathcal{M} . We denote it by $\text{Del}_{\mathcal{M}}(\mathcal{P})$.

Boissonnat and Cazals [19] define the natural neighbors of a point $\mathbf{x} \in \mathcal{M}$ with respect to the Voronoi diagram of \mathcal{P} restricted to \mathcal{M} .

Definition 5.8 Given a set of points $\mathcal{P} \subset \mathcal{M}$,

- (a) the \mathcal{M} -(natural) neighbors of a point \mathbf{x} of \mathcal{M} are the vertices of the facets of $\text{Del}_{\mathcal{M}}(\mathcal{P} \cup \{\mathbf{x}\})$ that are incident to \mathbf{x} .
- (b) Two sample points are \mathcal{M} -neighbors if they are adjacent in $\text{Del}_{\mathcal{M}}(\mathcal{P})$.

The topological concepts and the lemma that are introduced next are employed by Edelsbrunner and Shah [54] to show a sufficient condition under which $\text{Del}_{\mathcal{M}}(\mathcal{P})$ is homeomorphic to \mathcal{M} . If \mathcal{E} is a collection of sets, we denote the union of all sets in \mathcal{E} by $\cup \mathcal{E}$ and their intersection by $\cap \mathcal{E}$.

A subspace $\mathcal{F} \subseteq \mathbb{R}^d$ is *triangulable* if there exist a simplicial complex \mathcal{K} such that $\cup \mathcal{K}$ is homeomorphic to \mathcal{F} . A finite *covering* of \mathcal{F} is a finite collection \mathcal{E} of subsets of \mathcal{F} such that $\mathcal{F} = \cup \mathcal{E}$. The *nerve* of a finite covering \mathcal{E} of \mathcal{F} is a set of collections of sets that consists of all subcollections of \mathcal{E} with non-empty intersection.

$$\text{Nrv}(\mathcal{E}) := \{E \subseteq \mathcal{E} \mid \cap E \neq \emptyset\}.$$

We observe that $E_2 \subset E_1$ and $E_1 \in \mathcal{E}$ implies that $E_2 \in \mathcal{E}$ which is the property of a cell complex as it was introduced in Chapter 2. A *geometric realization* of $\text{Nrv}(\mathcal{E})$ is a simplicial complex \mathcal{K} together with a bijection β from \mathcal{E} to the vertices of \mathcal{K} so that $E \in \text{Nrv}(\mathcal{E})$ iff $\beta(E)$ spans a simplex in \mathcal{K} .

The following theorem by Leray is known under the name *nerve theorem*.

Theorem 5.9 [75] Let \mathcal{E} be a finite closed covering of a triangulable space $\mathcal{F} \subseteq \mathbb{R}^d$ so that for every $E \subseteq \mathcal{E}$, $\cap E$ is either empty or contractible. Let \mathcal{K} be a geometric realization of $\text{Nrv}(\mathcal{E})$, then \mathcal{F} and $\cup \mathcal{K}$ are homotopy equivalent.

In particular, the nerve of a Voronoi diagram of a set of points in general position is its dual Delaunay triangulation, and the nerve of the restricted Voronoi diagram $\text{Vor}_{\mathcal{M}}(\mathcal{P})$ is the restricted Delaunay triangulation $\text{Del}_{\mathcal{M}}(\mathcal{P})$. From Theorem 5.9, we can deduce sufficient conditions under which $\text{Del}_{\mathcal{M}}(\mathcal{P})$ is homeomorphic to the closed k -manifold \mathcal{M} . We deduce directly that they are of the same homotopy type if the intersection of a set of cells of $\text{Vor}_{\mathcal{M}}(\mathcal{P})$ is either empty or contractible. If, in addition, any non-empty intersection of a face of $\text{Vor}(\mathcal{P})$ with \mathcal{M} has the “right” dimension, i.e. the intersection of a $(d - l)$ -dimensional face with \mathcal{M} has dimension $k - l$, then, $\text{Del}_{\mathcal{M}}(\mathcal{P})$ and \mathcal{M} are homeomorphic. For details, see [54].

5.3 Properties of well sampled manifolds

In this section, we are concerned with sample point sets \mathcal{P} that fulfill some sampling condition with respect to the underlying manifold. The sampling condition allows to deduce some properties of the sample points, and, notably, to estimate the normal space. Most results are inspired from Amenta and Bern [9] where they were originally written for 2-manifolds in \mathbb{R}^3 . The work of Dey et al. [46] performs the generalization to k -manifolds in order to define an algorithm that detects the dimension of a manifold from a sample set by looking at the shape of the Voronoi cells. Notably, the concept of a *pole vector* which has been introduced in [9] is enlarged to the general case.

Definition 5.10 *Let \mathcal{P} be a set of sample points of a smooth k -manifold \mathcal{M} .*

1. (**ϵ -sample**) *For $\epsilon < 1$, \mathcal{P} is an ϵ -sample of \mathcal{M} if every point $\mathbf{x} \in \mathcal{M}$ has a sample point at distance at most $\epsilon \text{ lfs}(\mathbf{x})$.*
2. (**(ϵ, δ) -sample**) *Let ϵ and δ such that $\frac{\epsilon}{2} \leq \delta < \epsilon < 1$. \mathcal{P} is an (ϵ, δ) -sample of \mathcal{M} if it is an ϵ -sample of \mathcal{M} and if each sample point $\mathbf{p}_i \in \mathcal{P}$ is at least at distance $\delta \text{ lfs}(\mathbf{x})$ from any other sample point.*

The definition of (ϵ, δ) -sample is necessary if the dimension of the manifold \mathcal{M} is not known in advance [46]: without the requirement of a minimum distance between sample points, the sample points of a surface could, for example, trace a curve on the surface and the two cases, $k = 1$ or $k = 2$, are not distinguishable even if ϵ is small.

Because the local feature size $\text{lfs}(\cdot)$ is Lipschitz, given an ϵ -sample \mathcal{P} , it follows that the maximum distance between $\mathbf{x} \in \mathcal{M}$ and the closest sample point $\mathbf{p}_i \in \mathcal{P}$ is $\frac{\epsilon}{1-\epsilon} \text{lfs}(\mathbf{p}_i)$:

$$\|\mathbf{x} - \mathbf{p}_i\| \leq \epsilon \text{ lfs}(\mathbf{x}) \leq \epsilon (\text{lfs}(\mathbf{p}_i) + \|\mathbf{x} - \mathbf{p}_i\|) \iff \|\mathbf{x} - \mathbf{p}_i\| \leq \frac{\epsilon}{1-\epsilon} \text{lfs}(\mathbf{p}_i). \quad (5.1)$$

The next lemma by Boissonnat and Cazals bounds the maximum distance between \mathbf{x} and its \mathcal{M} -neighbors. It is the same as the maximum distance between two adjacent vertices of $\text{Del}_{\mathcal{M}}(\mathcal{P})$.

Lemma 5.11 [19, Lemma 9 (2)]

Let \mathcal{P} be an ϵ -sample of \mathcal{M} with $\epsilon < \frac{1}{2}$. If \mathbf{p}_i and \mathbf{p}_j are \mathcal{M} -neighbors, i.e. they are adjacent in $\text{Del}_{\mathcal{M}}(\mathcal{P})$, then $\|\mathbf{p}_i - \mathbf{p}_j\| \leq \frac{2\epsilon}{1-\epsilon} \text{lfs}(\mathbf{p}_i)$ and $\|\mathbf{p}_i - \mathbf{p}_j\| \leq \frac{2\epsilon}{1-\epsilon} \text{lfs}(\mathbf{p}_j)$

Proof: If \mathbf{p}_i and \mathbf{p}_j are connected by an edge in $\text{Del}_{\mathcal{M}}(\mathcal{P})$, there exist a point $\mathbf{v} \in \mathcal{M}$ that is part of the Voronoi face $V(\mathbf{p}_i, \mathbf{p}_j)$. The ball centered on \mathbf{v} having \mathbf{p}_i and \mathbf{p}_j on its boundary is empty of other sample points. Consequently, \mathbf{p}_i and \mathbf{p}_j are the closest sample points from \mathbf{v} . With $r = \|\mathbf{v} - \mathbf{p}_i\| = \|\mathbf{v} - \mathbf{p}_j\|$, we deduce from Equation (5.1) that

$$r \leq \frac{\epsilon}{1-\epsilon} \text{lfs}(\mathbf{p}_i) \quad \text{and} \quad r \leq \frac{\epsilon}{1-\epsilon} \text{lfs}(\mathbf{p}_j).$$

The distance $\|\mathbf{p}_i - \mathbf{p}_j\|$ at most twice the distance r between \mathbf{v} and \mathbf{p}_i or \mathbf{p}_j . □

Amenta and Bern were the first to apply Theorem 5.9 to show that, if the sampling is sufficiently dense, the three-dimensional Delaunay triangulation of the sample points contains a subset of triangles that is homeomorphic to the sampled surface \mathcal{M} .

Theorem 5.12 [9, Theorem 2] *Let \mathcal{M} be a 2-manifold in \mathbb{R}^3 . For $\epsilon < 0.1$, the restricted Delaunay triangulation of an ϵ -sample \mathcal{P} of \mathcal{M} is homeomorphic to \mathcal{M} .*

Location of Voronoi vertices

If the underlying manifold \mathcal{M} is a curve ($k = 1$), we can show with Lemma 5.3 that no Voronoi vertex of $\text{Vor}(\mathcal{P})$ is close to \mathcal{M} with respect to the local feature size.

Lemma 5.13 *The minimum distance from $\mathbf{x} \in \mathcal{M}$ to a Voronoi vertex \mathbf{v} of $V(\mathbf{x})$ is an upper bound for $\frac{1}{2}\text{lfs}(\mathbf{x})$, $\frac{1}{2}\text{lfs}(\mathbf{x}) \leq \min_{\mathbf{v} \in V(\mathbf{x})} \|\mathbf{x} - \mathbf{v}\|$.*

Proof: A Voronoi vertex \mathbf{v} is the center of an empty Delaunay ball and it intersects the curve at its $d + 1$ generators. Either its intersection with the curve is a 1-ball which means that it is or it contains a medial ball. In this case, its radius is greater or equal than the local feature size of each of the generators. Or it is not a 1-ball and Lemma 5.3 implies that it must contain a point of the curve's medial axis, thus, the radius of the Delaunay ball must be greater or equal than half the local feature size of each of the generators. □




Figure 5.1: Which side of the sliver tetrahedron is closer to \mathcal{M} : (a) the outer side with T_1 and T_2 or (b) the inner side with T_3 and T_4 ?

The same does not work for higher dimensional manifolds due to so-called *sliver* tetrahedra. A sliver is a very flat simplex of the Delaunay triangulation whose vertices lie almost on a $(d - 1)$ -ball. The center of its Delaunay sphere (of dimension d) might be arbitrary close to the manifold, and, consequently, far from the medial axis of \mathcal{M} . Figure 5.1 shows a sliver in 3D seen from the exterior of the surface \mathcal{M} . Without further input, it is not possible to determine $\text{Del}_{\mathcal{M}}(\mathcal{P})$ even if the sampling is very dense: \mathcal{M}

might intersect the Voronoi edges dual to the “outer” triangles T_1 and T_2 , or those dual to the “inner” triangles, T_3 and T_4 . Knowing only the sample points, the information is not present.

Even though not all Voronoi vertices of an ϵ -sample are close to the medial axis of the manifold, at least some of them are. The definitions and results that are described below are based on this observation.

Normal estimation via poles

The Voronoi diagram of an ϵ -sample consists of cells that are long and skinny in the normal directions at a point $\mathbf{x} \in \mathcal{M}$. Therefore, the normal directions can be estimated from the Voronoi cell $V(\mathbf{x})$ of $\text{Vor}(\mathcal{P} \cup \{\mathbf{x}\})$.¹ For hypersurfaces, i.e. $(d - 1)$ manifolds, this observation led to the definition of a pole by Amenta et al. [9]. The authors show that the vector from the surface point to the furthest point in its Voronoi cell is a good estimate for the normal to the hypersurface at this point. In the general case, the concept of poles has been defined in [46]. To be more specific, we recall the exact definition.

Definition 5.14 ([46]) *The Voronoi subpolytopes for a point $\mathbf{x} \in \mathcal{M}$ are subsets $V^i(\mathbf{x}) \subseteq V(\mathbf{x}), i = 1, \dots, d - k$, (of dimension $d - i + 1$) defined recursively as follows: let $V^1(\mathbf{x}) = V(\mathbf{x})$ and assume $V^i(\mathbf{x})$ to be defined. The pole \mathbf{v}_x^i is the Voronoi vertex of $V^i(\mathbf{x})$ which is furthest from \mathbf{x} , and the i th pole vector $\vec{\mathbf{v}}_x^i$ of \mathbf{x} is the unit vector with direction $\overrightarrow{\mathbf{x}\mathbf{v}_x^i}$. If $V^i(\mathbf{x})$ is unbounded, \mathbf{v}_x^i is taken at infinity and the direction of $\vec{\mathbf{v}}_x^i$ is the average of all directions of unbounded edges of $V^i(\mathbf{x})$. The Voronoi subpolytope $V^{i+1}(\mathbf{x})$ is the minimal polytope containing all points $\{\mathbf{y} \in V^i(\mathbf{x}) \mid \angle \overrightarrow{\mathbf{x}\mathbf{y}} \overrightarrow{\mathbf{v}}_x^i = \frac{\pi}{2}\}$.*

As stated in the next lemma, the pole vectors provide a good approximation of the (non-oriented) direction of the normals to \mathcal{M} at \mathbf{x} . Therefore, we refer to $\tilde{\mathcal{N}}_x$ as the affine space with origin \mathbf{x} spanned by the pole vectors $\vec{\mathbf{v}}_x^i, i = 1, \dots, d - k$. The affine space \tilde{T}_x that contains \mathbf{x} and is orthogonal to $\tilde{\mathcal{N}}_x$ approximates well the tangent space to \mathcal{M} at \mathbf{x} .

Lemma 5.15 [9, lemma 5], [46, lemma 3.1] *There exist a normal vector $\vec{\mathbf{n}}_x^i \in \mathcal{N}_x$ for each pole vector $\vec{\mathbf{v}}_x^i, 1 \leq i \leq d - k$, such that the smaller angle between $\vec{\mathbf{n}}_x^i$ and $\vec{\mathbf{v}}_x^i$ is at most $2 \arcsin(\frac{\epsilon}{1-\epsilon})$.*

The next lemma from [46] shows that for any unit vector $\vec{\mathbf{n}}_x \in \mathcal{N}_x$, there exists a corresponding vector in the approximated normal space $\tilde{\mathcal{N}}_x$ which makes a small angle with $\vec{\mathbf{n}}_x$.

Lemma 5.16 [46, lemma 3.3] *For any unit vector $\vec{\mathbf{n}}_x \in \mathcal{N}_x$, there exists a corresponding vector $\vec{\mathbf{v}}_x \in \tilde{\mathcal{N}}_x$ such that the angle β between both vectors is bounded by $\beta \leq 4\sqrt{d - k} \arcsin(\frac{\epsilon}{1-\epsilon})$.*

Normal estimation via closest neighbors

The following lemma from [9] justifies another method to estimate the normal. It has the advantage of being independent from the Voronoi diagram $\text{Vor}(\mathcal{P})$. This is important for very large sets of sample

¹All results are written with respect to an arbitrary point $\mathbf{x} \in \mathcal{M}$. Notice, that they hold equally for a sample point $\mathbf{p}_i \in \mathcal{P}$.

points and, even more, for manifolds embedded in higher dimensional spaces ($d > 3$). In both cases, it may be too costly to compute $\text{Vor}(\mathcal{P})$. The lemma generalizes directly to k -manifolds in \mathbb{R}^d .

Lemma 5.17 Generalization of [9, lemma 2]

For any two points \mathbf{x} and \mathbf{y} on \mathcal{M} with $\|\mathbf{x} - \mathbf{y}\| \leq \rho \text{lfs}(\mathbf{x})$, the smaller angle between the line segment $[\mathbf{x}\mathbf{y}]$ and any vector $\vec{\mathbf{n}}_x$ from the normal space \mathcal{N}_x to \mathcal{M} at \mathbf{x} is at least $\frac{\pi}{2} - \arcsin(\frac{\rho}{2})$.

Proof: The proof of [9] is independent of the dimensions d and k : The manifold \mathcal{M} is exterior to the tangent balls to \mathcal{M} at \mathbf{x} , so $\mathbf{y} \in \mathcal{M}$ cannot lie inside any tangent ball. The angle θ between a vector in \mathcal{N}_x and the segment $[\mathbf{x}\mathbf{y}]$ is minimized when \mathbf{y} lies on one of the tangent balls. Since the radius of the tangent balls is at least $\text{lfs}(\mathbf{x})$ and $\|\mathbf{x} - \mathbf{y}\| \leq \rho \text{lfs}(\mathbf{x})$, we obtain $\theta \geq \arcsin(\frac{\rho}{2})$. \square

For curves, $k = 1$, we apply this lemma in Section 8.2 in order to estimate the tangent line \mathcal{T}_x at \mathbf{x} by the line through \mathbf{x} and its closest neighbor in \mathcal{P} .

For surfaces in 3D, Dey, Funke and Ramos proposed a normal estimation method in [44] and in [60] based on this lemma: Choose \mathbf{p} to be the nearest neighbor of $\mathbf{x} \in \mathcal{M}$ in \mathcal{P} and \mathbf{q} to be the nearest neighbor of \mathbf{x} among all points of \mathcal{P} such that the angle $\angle \mathbf{p}\mathbf{x}\mathbf{q}$ is between 45° and 135° . The normal to the triangle $(\mathbf{p}\mathbf{x}\mathbf{q})$ approximates the normal $\vec{\mathbf{n}}_x$. Since the distances from \mathbf{x} to \mathbf{p} and to \mathbf{q} are small, Lemma 5.17 applies. They claim

Lemma 5.18 [60, lemma 8.2] [44, lemma 2]

For $\mathbf{x} \in \mathcal{M}$, let \mathbf{p} and \mathbf{q} be determined as above. The angle between $\vec{\mathbf{n}}_x$ and the normal to the triangle $(\mathbf{p}\mathbf{x}\mathbf{q})$ is $O(\epsilon)$.

However, this method is very sensitive to the sampling density and does not work in undersampled zones. Other methods are more stable since they estimate the normal from the k -nearest neighbors of a point for some greater k . For example, Hoppe et al. [72] approximate $\vec{\mathbf{n}}_x$ by the normal to the least-square fitting plane of the k -nearest neighbors of \mathbf{x} . This method is also justified with 5.17 as long as the k -nearest neighbors are sufficiently close to \mathbf{x} . Gopi et al. [65] propose to estimate $\vec{\mathbf{n}}_x$ by the vector that minimizes the variance of the dot product between itself and the vector from \mathbf{x} to its k -nearest neighbors. The authors of [65] affirm that both methods are equivalent.

Normal variation

In the remainder of this section, we restrict ourselves to hypersurfaces, i.e. $k = d - 1$. Considering two points on the hypersurface that are close, Amenta and Bern show that the angle between the normals at the sample points is small. The proof is again based on the argument that the medial balls tangent to \mathcal{M} at a point \mathbf{x} are not intersected by \mathcal{M} . From this lemma, we derive other properties that are useful for some issues in Chapter 6. The generalization to k -manifolds seems possible, however, it would require to introduce additional notations and concepts.

Lemma 5.19 [9, lemma 3]

For any two points \mathbf{x} and \mathbf{y} on \mathcal{M} with $\|\mathbf{x} - \mathbf{y}\| \leq \rho \text{ lfs}(\mathbf{x})$, $\rho < \frac{1}{3}$, the angle between the normals to \mathcal{M} at \mathbf{x} and at \mathbf{y} is at most $\frac{\rho}{1-3\rho}$.

From the preceding lemmas, we can deduce that the angle between the tangent plane to \mathcal{M} at a point \mathbf{x} and the bisector hyperplane of any two sample points that are close to \mathbf{x} cannot be too small. This is shown in the next lemma.

Lemma 5.20 *The angle between the tangent plane \mathcal{T}_x to \mathcal{M} at $\mathbf{x} \in \mathcal{M}$ and the bisector of two sample points \mathbf{p}_i and \mathbf{p}_j that are at distance at most $\rho \text{ lfs}(\mathbf{x})$ from \mathbf{x} , $\rho \leq \frac{1}{3}$, is at least $\pi/2 - \arcsin(\frac{\rho}{1-\rho}) - \frac{\rho}{1-3\rho}$.*




Figure 5.2: Bounding the angle between the normal \mathbf{n}_x and the segment $[\mathbf{p}_i \mathbf{p}_j]$.

Proof: See Figure 5.2 for notations. We derive a lower bound for the angle α between the vector $\overrightarrow{\mathbf{p}_i \mathbf{p}_j}$ and the surface normal $\vec{\mathbf{n}}_i$ at \mathbf{p}_i and an upper bound for the angle β between the normals at \mathbf{p}_i and \mathbf{x} . The angle between $\vec{\mathbf{n}}_x$ and $\overrightarrow{\mathbf{p}_i \mathbf{p}_j}$ is at least $\alpha - \beta$.

In order to bound α , we apply Lemma 5.17. The distance between the sample points is bounded by $\|\mathbf{p}_i - \mathbf{p}_j\| \leq \|\mathbf{p}_i - \mathbf{x}\| + \|\mathbf{p}_j - \mathbf{x}\| \leq 2\rho \text{ lfs}(\mathbf{x})$. Using Lemma 5.2, we get $\text{lfs}(\mathbf{x}) \leq \text{lfs}(\mathbf{p}_i) + \rho \text{ lfs}(\mathbf{x})$ and $\text{lfs}(\mathbf{x}) \leq \frac{1}{1-\rho} \text{lfs}(\mathbf{p}_i)$. Hence, $\|\mathbf{p}_i - \mathbf{p}_j\| \leq \frac{2\rho}{1-\rho} \text{lfs}(\mathbf{p}_i)$. It follows from Lemma 5.17 that α is at least $\pi/2 - \arcsin(\frac{\rho}{1-\rho})$. The angle β between the normals at \mathbf{p}_i and \mathbf{x} is at most $\frac{\rho}{1-3\rho}$ if $\rho < \frac{1}{3}$ (Lemma 5.19). Concluding, we get $\alpha - \beta \geq \pi/2 - \arcsin(\frac{\rho}{1-\rho}) - \frac{\rho}{1-3\rho}$, provided that $\rho < \frac{1}{3}$. \square

From this lemma, we deduce the corollary.

Corollary 5.21 *The tangent plane \mathcal{T}_x to \mathcal{M} at a point $\mathbf{x} \in \mathcal{M}$ cannot be parallel to the bisector of two sample points that are at distance at most $\rho \text{ lfs}(\mathbf{x})$ from \mathbf{x} , for $\rho < \frac{1}{3}$.*

Lemma 5.22 *For any $\rho < \frac{\pi}{2+3\pi}$, the parallel projection onto \mathcal{T}_x of the intersection of the ball $B(\mathbf{x}, \rho \text{ lfs}(\mathbf{x}))$ with M is 1-1.*

Proof: By Lemma 5.19, for any point \mathbf{y} at distance at most $\rho \text{ lfs}(\mathbf{x})$ from \mathbf{x} , the angle between the normals to \mathcal{M} at \mathbf{x} and at \mathbf{y} is at most $\frac{\rho}{1-3\rho}$. The projection is 1-1 if $\frac{\rho}{1-3\rho} < \frac{\pi}{2} \iff \rho < \frac{\pi}{2+3\pi} \approx 0.27$. \square

Bounding the local feature size

If the underlying manifold \mathcal{M} is a curve ($k = 1$), we saw in Lemma 5.13 that the minimum distance from $\mathbf{x} \in \mathcal{M}$ to a Voronoi vertex \mathbf{v} of $V(\mathbf{x})$ is an upper bound for $\frac{1}{2}\text{lfs}(\mathbf{x})$. For hypersurface, $k = d - 1$, the local feature size $\text{lfs}(\mathbf{x})$ can be upper bounded by the distance to the furthest Voronoi vertices of $V(\mathbf{x})$.

There is only one pole vector, so we omit the superscript and denote it $\vec{\mathbf{v}}_x^+$. It is the *positive* pole vector in opposition to the *negative* pole vector $\vec{\mathbf{v}}_x^-$: the *negative pole* \mathbf{v}_x^- be the furthest Voronoi vertex of $V(\mathbf{x})$ such that the angle between the positive and the negative pole vector, $\vec{\mathbf{v}}_x^- = \overline{\mathbf{x}\mathbf{v}_x^-}$, is bigger than $\frac{\pi}{2}$.

We show in the sequel that the distance from the poles to \mathbf{x} must be greater than $\text{lfs}(\mathbf{x})$. Intuitively, one can say that the Delaunay spheres centered at the positive and the negative pole approximate the medial balls at \mathbf{x} . Since $\text{lfs}(\mathbf{x})$ is smaller than the smallest radius of the medial balls tangent to \mathcal{M} at \mathbf{x} , the smaller distance from \mathbf{x} to its poles is an estimate for the local feature size. For the proof, we need the following lemma from Amenta and Bern [9].

Lemma 5.23 [9, lemma 5] *Let $\mathbf{x} \in \mathcal{M}$ and \mathbf{v} be any point in $V(\mathbf{x})$ with $\|\mathbf{v} - \mathbf{x}\| \geq \nu \text{lfs}(\mathbf{x})$ for $\nu > 0$. The angle at \mathbf{x} between the vector to \mathbf{v} and the normal $\vec{\mathbf{n}}_x$ (oriented in the same direction) is at most $\arcsin \frac{\epsilon}{\nu(1-\epsilon)} + \arcsin \frac{\epsilon}{(1-\epsilon)}$.*

We can state the lemma. The proof is part of the proof of Theorem 1 from [41].

Lemma 5.24 (after [41]) *Let $k = d - 1$ and $\epsilon < 0.27$. The poles \mathbf{v}_x^+ and \mathbf{v}_x^- of \mathbf{x} are at distance at least $\text{lfs}(\mathbf{x})$ from \mathbf{x} .*

Proof: It is sufficient to prove the lemma for the negative pole because the positive pole is, by definition, further from \mathbf{x} than the negative pole. We choose the medial ball at \mathbf{x} with center \mathbf{m} such that \mathbf{v}_x^+ and \mathbf{m} lie




Figure 5.3: The angle $\beta = \pi - \alpha^+ - \alpha^-$ must be bigger than $\frac{\pi}{2}$.

on different sides of \mathcal{M} , i.e. $\overrightarrow{\mathbf{x}\mathbf{m}}$ makes a large angle greater than $\frac{\pi}{2}$ with the positive pole vector $\overrightarrow{\mathbf{v}_x^+}$. See Figure 5.3. We know that $\|\mathbf{x} - \mathbf{m}\| \geq \text{lfs}(\mathbf{x})$ and $\mathbf{m} \in V(\mathbf{x})$. By convexity of $V(\mathbf{x})$, there exists a Voronoi vertex \mathbf{v} that is the furthest Voronoi vertex in direction $\overrightarrow{\mathbf{x}\mathbf{m}}$ and $\|\mathbf{x} - \mathbf{v}\| \geq \text{lfs}(\mathbf{x})$. Thus, with Lemma 5.23, the angle α^- between $\overrightarrow{\mathbf{x}\mathbf{m}}$ and $\overrightarrow{\mathbf{x}\mathbf{v}}$ is smaller than $2 \arcsin \frac{\epsilon}{(1-\epsilon)}$. This is also the case for the angle α^+ between $\overrightarrow{\mathbf{v}_x^+}$ and $\overrightarrow{\mathbf{n}_x}$ but in the opposite direction (Lemma 5.15). For $\epsilon < 0.27$, the angle between $\overrightarrow{\mathbf{v}_x^+}$ and $\overrightarrow{\mathbf{x}\mathbf{v}}$ is greater than $\frac{\pi}{2}$: $\pi - (\alpha^+ + \alpha^-) = \pi - 4 \arcsin \frac{\epsilon}{(1-\epsilon)} > \frac{\pi}{2}$ for $\epsilon < \frac{\sqrt{2-\sqrt{2}}}{2+\sqrt{2-\sqrt{2}}} < 0.27$. Thus, either $\mathbf{v} = \mathbf{v}_x^-$, or \mathbf{v}_x^- is further away from \mathbf{x} than \mathbf{v} . \square



Chapter 6

A coordinate system on a k -manifold of \mathbb{R}^d

6.1 Introduction

Surfaces and general k -manifolds represented by a set of unordered sample points are encountered in many application areas such as computer graphics, computer aided design (CAD) and reverse engineering, image processing, and scientific computation. Many algorithms that are applied to sampled manifolds rely on the definition of a local neighborhood on the manifold. Reconstructing the manifold from the sample points is one way to respond to this demand. However, it might be unnecessary and also too costly to establish a global approximation of the manifold. Differently, by adapting the definition of natural neighbor coordinates to the manifold case, our method defines, for any point on a sampled manifold, a set of neighboring sample points as well as a set of coordinates associated to them. If the sampling is sufficiently dense, this coordinate system is provably local on the manifold and has good continuity properties. Unless the manifold is locally flat, it is not barycentric but the distance between a point of the manifold and its weighted barycenter is bounded with respect to the sampling density. Moreover, the set of neighbors and the coordinate functions can be computed efficiently because locality allows efficient filtering methods. We do not impose any restriction on the genus of the manifold, the number of connected components, nor any other global features of the manifold. Uniform sampling is neither required, and we allow the sampling density to be related to the local curvature of the manifold.¹

Related Work

The interest in the definition of an easy-to-compute coordinate system on general manifolds is for example witnessed by [82]. The major drawback to applying the natural neighbor coordinate system as they are defined in the d -dimensional Voronoi diagram to points issued from a manifold comes from the fact that its definition is limited to the convex hull of the sample points. To avoid this problem, a common solution consists of adding a box enclosing the object. Obviously, this solution causes problems, e.g. the choice

¹This chapter is the subject of two publications [24] [25].

of the size of the bounding box, the number of sample points taken from it, artifacts arising from the bounding box points, and the augmented computation cost. In [29], Brown has enlarged the coordinate definition outside the convex hull to cover the union of the Delaunay balls, which is still too restrictive in many applications.

A second drawback is that a point is likely to have neighbors that are far away from the point. Boissonnat and Cazals show that the sum of the coordinates associated to those neighbors that are far away tends to zero when the sampling density increases [19]. However, even though the influence of the far neighbors is small, the fact that Sibson's coordinates are not local affects not only the beauty and rigor of the result, but the time required to compute the coordinates and the exactness of an interpolation scheme.

For points issued from a sphere, Brown proposes a solution in [28]. Instead of referring to the d -dimensional Voronoi diagram, Brown defines natural neighbor coordinates with respect to the geodesic Voronoi diagram on the sphere. This definition generalizes Sibson's coordinates in a straightforward manner. Therefore, the basic properties of definition 3.1 are fulfilled, except the barycentric coordinate property (i) which cannot be fulfilled since points on the sphere do not belong to the convex hull of their neighbors. The obvious difficulty in enlarging Brown's approach to general manifolds is that geodesic Voronoi diagrams are much more complicated than Euclidean diagrams and difficult to compute [74]. Moreover, in many applications the manifold is not known and neither is the geodesic Voronoi diagram.

In this thesis, we suggest another system of coordinates for points on a manifold. It is closely related to natural neighbor coordinates, yet instead of considering the geodesic Voronoi diagram on the manifold, as Brown, or the Euclidean d -dimensional Voronoi diagram of the sample points, as Sibson, it is defined in the intersection of the tangent space of each manifold point with the Euclidean Voronoi diagram of the sample set. If the tangent spaces are not given as part of the input, they can be easily estimated from the sample points. The resulting coordinate system is local and it is inherently k -dimensional.

Outline of the chapter

After this introduction, we define in section 6.2 the \mathcal{T}_x -neighbors of a manifold point \mathbf{x} with respect to a sample of the manifold. We show that all \mathcal{T}_x -neighbors of \mathbf{x} lie in a small neighborhood around \mathbf{x} if the manifold is well sampled. In section 6.3, we define the \mathcal{T}_x -coordinate system on the manifold associated to the sample. We show the main properties of the \mathcal{T}_x -coordinates; in particular, we show that the coordinate functions have compact supports and are continuously differentiable almost everywhere on the manifold. Although a point on a manifold cannot, in general, be expressed as a convex combination of other points on the manifold, we show in subsection 6.3.3, that the barycentric coordinate property is approximately satisfied, with an error that depends on the local curvature of the manifold and on the sampling density. For surfaces in $3D$, we determine the gradient of the coordinate function. At last, we define a non-convex coordinate function for surfaces that fulfills the barycentric coordinate property.

6.2 Manifold neighbors

With this section begins the core part of the thesis in which we define a local neighborhood of a point \mathbf{x} of a smooth, compact k -manifold \mathcal{M} without boundary with respect to an ϵ -sample \mathcal{P} of \mathcal{M} . For the proof of Lemma 6.2, we suppose further that \mathcal{M} is subanalytic, see also Section 5.1. In order not to compute the geodesic Voronoi diagram on \mathcal{M} , we approximate \mathcal{M} locally by the tangent space \mathcal{T}_x of \mathbf{x} to \mathcal{M} . We determine the natural neighbors of \mathbf{x} in the Voronoi diagram $\text{Vor}(\mathcal{P})$ restricted to \mathcal{T}_x , and we call them \mathcal{T}_x -neighbors of \mathbf{x} .




Figure 6.1: Natural neighbors \mathcal{N} and \mathcal{T}_x -neighbors \mathcal{T}_x of \mathbf{x}

In Figure 6.1, a schematic drawing illustrates that some of the natural neighbors of \mathbf{x} which are highlighted in the picture can be far away from \mathbf{x} with respect to the geodesic distance on \mathcal{M} . The \mathcal{T}_x -neighbors are the subset of the natural neighbors that are close to \mathbf{x} on \mathcal{M} . In the remainder of this section, we formally define the \mathcal{T}_x -neighbors of a point $\mathbf{x} \in \mathcal{M}$ and prove that they are close to \mathbf{x} .

6.2.1 Definition of \mathcal{T}_x -neighbors

At first, we assume that for each point $\mathbf{x} \in \mathcal{M}$, the set of normals $\{\vec{\mathbf{n}}_x^i\}_{i=1, \dots, d-k}$ to \mathcal{M} at \mathbf{x} is known, and therefore, the tangent space \mathcal{T}_x at \mathbf{x} . Let $\overline{\text{Vor}}(\mathcal{P})$ be the intersection of $\text{Vor}(\mathcal{P})$ with the tangent space \mathcal{T}_x . The Delaunay triangulation restricted to \mathcal{T}_x that consists of the faces of $\text{Del}(\mathcal{P})$ whose dual Voronoi faces intersect \mathcal{T}_x is called $\overline{\text{Del}}(\mathcal{P})$. Alternatively, by Observation 2.4, $\overline{\text{Vor}}(\mathcal{P})$ is the k -dimensional power diagram of the weighted points $B_i = (\overline{\mathbf{p}}_i, w_i)$ where $\overline{\mathbf{p}}_i$ is the projection of the sample point $\mathbf{p}_i \in \mathcal{P}$ onto \mathcal{T}_x and $w_i = -\|\mathbf{p}_i - \overline{\mathbf{p}}_i\|^2$. The $(k-l)$ -faces of $\overline{\text{Vor}}(\mathcal{P})$, $l \leq k$, are $\overline{V}(\mathbf{p}_{i_0}, \dots, \mathbf{p}_{i_l}) = \mathcal{T}_x \cap \left(\bigcap_{j=0, \dots, l} V(\mathbf{p}_{i_j}), \mathbf{p}_{i_j} \in \mathcal{P} \right)$. The superscript $+$ indicates that we refer to $\overline{\text{Vor}}^+(\mathcal{P}) = \overline{\text{Vor}}^+(\mathcal{P} \cup \mathbf{x})$. We define $\overline{V}(\mathbf{x}) = V^+(\mathbf{x}) \cap \mathcal{T}_x$.

Let $\overline{\text{Reg}}(\mathcal{P})$ be the regular triangulation dual to $\overline{\text{Vor}}(\mathcal{P})$. Since two cells of $\overline{\text{Vor}}(\mathcal{P})$ are adjacent iff their corresponding cells in $\text{Vor}(\mathcal{P})$ are adjacent and intersect \mathcal{T}_x , $\overline{\text{Reg}}(\mathcal{P})$ is the projection of $\overline{\text{Del}}(\mathcal{P})$ onto \mathcal{T}_x .

Definition 6.1 (\mathcal{T}_x -neighbor of \mathbf{x}) Given a set of sample points $\mathcal{P} \subset \mathcal{M}$ The \mathcal{T}_x -neighbors of a point $\mathbf{x} \in \mathcal{M}$ are the sample points $\mathbf{p}_i \in \mathcal{P}$ such that their projection $\overline{\mathbf{p}_i}$ is a natural neighbor of \mathbf{x} in $\overline{\text{Vor}}(\mathcal{P})$.

To see that the concept of \mathcal{T}_x -neighbors is well-defined, we make two observations: First, the definitions of Section 2.2 assume general position of the point sites. The case that two sites have the same position and the same weight must be excluded. In our context, this occurs if the bisector of two \mathcal{T}_x -neighbors \mathbf{p}_k and \mathbf{p}_j of \mathbf{x} coincides with \mathcal{T}_x : \mathbf{p}_k and \mathbf{p}_j are projected at the same position, and they have the same weight because they are at the same distance to the tangent space but on opposite sides.

For $k = d - 1$, we can easily show that the angle between \mathcal{T}_x and the bisector hyperplane of \mathbf{p}_k and \mathbf{p}_j is strictly positive, if $\epsilon \leq \frac{1}{9}$: In the remainder of this chapter, we show that \mathbf{p}_k and \mathbf{p}_j are at distance at most $\frac{2\epsilon}{\sqrt{1-2\epsilon}} \text{lfs}(\mathbf{x})$ to \mathbf{x} (Lemma 6.3), thus, Corollary 5.21 applies with $\rho = \frac{2\epsilon}{\sqrt{1-2\epsilon}}$.

In any case, we can assume without real loss of generality that no bisector hyperplane of two sample points is tangent to \mathcal{M} because any small perturbation of \mathcal{P} would remove this situation.

Second, we need the following lemma to show that \mathbf{x} lies in the convex hull of the projection of its \mathcal{T}_x -neighbors. It is equivalent to show that \mathbf{x} lies in the convex hull of its natural neighbors in $\overline{\text{Vor}}(\mathcal{P})$ or to show that $\overline{\text{Vor}}(\mathbf{x})$ is bounded.

Lemma 6.2 $\mathbf{x} \in \mathcal{M}$ belongs to the convex hull of the projection of its \mathcal{T}_x -neighbors on \mathcal{T}_x .

Proof: For a contradiction, assume that $\overline{\text{Vor}}(\mathbf{x})$ is unbounded. Then the interior of $\overline{\text{Vor}}(\mathbf{x})$ contains a point \mathbf{p}_∞ at infinity. Among the points of $\mathcal{P} \cup \{\mathbf{x}\}$, \mathbf{x} is the closest to \mathbf{p}_∞ since \mathbf{p}_∞ lies in the Voronoi cell of \mathbf{x} . Therefore, there is no point of \mathcal{P} in the halfspace \mathcal{H}^+ that contains \mathbf{p}_∞ and that is limited by the hyperplane \mathcal{H} passing through \mathbf{x} and normal to the line $(\mathbf{x}\mathbf{p}_\infty)$. See Figure 6.2.




Figure 6.2: \mathcal{H}^+ must contain sample points.

The proof is easier for $k = d - 1$. In this case, since the hypersurface \mathcal{M} is closed, it delimits a bounded region in \mathbb{R}^d . There is exactly one medial ball tangent to \mathcal{M} at \mathbf{x} that lies inside this region. Suppose its center is \mathbf{c} . Let \mathbf{y} be the intersection point of \mathcal{M} with the ray issued from \mathbf{c} that is orthogonal to \mathcal{H} and

contained in \mathcal{H}^+ . Such a point must always exist since the region delimited by \mathcal{M} is closed. Figure 6.2 shows a view onto the plane containing \mathbf{x} , \mathbf{p}_∞ , \mathbf{y} , and \mathbf{c} .

By definition, \mathbf{c} is a point of the medial axis of \mathcal{M} . We have $\|\mathbf{y} - \mathbf{c}\| \geq \text{lfs}(\mathbf{y})$ and for any $\mathbf{p}_i \in \mathcal{P}$, $\|\mathbf{y} - \mathbf{p}_i\| > \|\mathbf{y} - \mathbf{c}\|$ because \mathcal{H}^+ is empty of sample points. This contradicts the fact that \mathcal{P} is an ϵ -sample with $\epsilon < 1$.

For $k < d - 1$, the proof is slightly more complicated because it is more difficult to construct the point \mathbf{y} that lies in \mathcal{H}^+ and whose orthogonal projection onto \mathcal{H} is the center \mathbf{c} of a medial ball passing through \mathbf{x} . If such a point \mathbf{y} is found, the proof is completed because the arguments from above apply.

By construction, the vector $\overrightarrow{\mathbf{x}\mathbf{p}_\infty}$ is orthogonal to \mathcal{H} and contained in \mathcal{H}^+ . Also, it lies in the tangent space \mathcal{T}_x because $\mathbf{p}_\infty \in \overline{\mathcal{V}}(\mathbf{x})$. Consequently, \mathcal{H} must contain the normal space \mathcal{N}_x . Consider the set of medial balls \mathcal{B} tangent to \mathcal{M} at \mathbf{x} . Their centers \mathcal{C} lie in \mathcal{N}_x . More precisely, they lie in the intersection of \mathcal{N}_x with the medial axis of \mathcal{M} (see Def. 5.1).

We distinguish two cases: If \mathbf{x} lies in the interior of the convex hull of \mathcal{M} , any ray r in \mathcal{N}_x emanating from \mathbf{x} contains exactly one center $\mathbf{c}_r \in \mathcal{C}$ at finite distance from \mathbf{x} : A growing sphere tangent to \mathcal{M} at \mathbf{x} with center on r must intersect \mathcal{M} before degenerating to a halfspace because any hyperplane containing \mathbf{x} must intersect \mathcal{M} . The medial ball center \mathbf{c}_r is the first intersection point of r with the medial axis of \mathcal{M} . Notice that this intersection is in isolated points because a contact point \mathbf{y} of the medial ball B_r other than \mathbf{x} is contained in any ball centered on r tangent to \mathcal{M} at \mathbf{x} with larger radius, thus, these balls cannot be empty. See figure 6.3.




Figure 6.3: The contact point \mathbf{y} of the medial ball B_r (with center \mathbf{c}_r) is necessarily contained in the ball B'_r (with center \mathbf{c}'_r).

Let \mathcal{A} be the union of segments $[\mathbf{x}\mathbf{c}_r]$ for all r , $\mathcal{A} := \bigcup_r [\mathbf{x}\mathbf{c}_r]$. Assuming some generic condition, the boundary $\partial\mathcal{A}$ has dimension $(d - k - 1)$ and it separates \mathcal{N}_x into two disconnected components. In particular, this is true if the medial axis of \mathcal{M} is subanalytic. (In this case, $\partial\mathcal{A}$ is a $(d - k - 1)$ -dimensional

subanalytic set.) Following Chazal and Soufflet [32], this is ensured if \mathcal{M} is subanalytic.²

In the second case, \mathbf{x} lies on the convex hull of \mathcal{M} . Any ray in the *normal cone* of the convex hull of \mathcal{M} , i.e. the cone spanned by the normals to the convex hull of \mathcal{M} , defines a halfspace tangent to \mathcal{M} at \mathbf{x} that does not intersect \mathcal{M} and that is normal to the ray. Such halfspaces correspond to degenerate maximal balls with center at infinity. We consider them to be part of the set of medial balls \mathcal{B} . Any ray r emanating from \mathbf{x} in \mathcal{N}_x that lies outside the normal cone intersects the medial axis of \mathcal{M} at finite distance. Each such ray r contains, therefore, exactly one center $\mathbf{c}_r \in \mathcal{C}$ corresponding to the first intersection of r with the medial axis of \mathcal{M} . Since \mathcal{M} is compact, there must exist some $\mathbf{c} \in \mathcal{C}$ at finite distance and, with the same arguments as above, $\mathcal{N}_x \setminus \mathcal{C}$ has two disconnected components.




Figure 6.4: The construction of \mathbf{y} for a curve in \mathbb{R}^3 .

We determine a set $\partial\mathcal{J}$ by adding an affine translation of the tangent space \mathcal{T}_x to each $\mathbf{c} \in \mathcal{C}$, $\partial\mathcal{J} = \mathcal{C} \times \mathcal{T}_x$. By construction, $\partial\mathcal{J}$ separates \mathbb{R}^d into two components. Let \mathcal{J} be the region determined by $\partial\mathcal{J}$ that contains \mathbf{x} . In $\mathcal{J}' = \mathcal{J} \setminus \mathcal{B}$, the interior of the medial balls \mathcal{B} is excluded from \mathcal{J} . At \mathbf{x} , \mathcal{J}' has no thickness because the medial balls \mathcal{B} are tangent to each other (and to \mathcal{M}). Moreover, $\mathcal{J}' \cap \mathcal{H} = \mathbf{x}$ because the normal space \mathcal{N}_x is contained in \mathcal{H} , the tangent space \mathcal{T}_x is contained in \mathcal{J}' and $\mathcal{N}_x \cap \mathcal{T}_x = \mathbf{x}$. Figure 6.4 indicates the intersection of \mathcal{H} with the medial balls \mathcal{B} as well as the region \mathcal{J} (enclosed by the horizontal rays) in the case of a curve in \mathbb{R}^3 .

In the neighborhood of \mathbf{x} , \mathcal{M} must lie in the interior of \mathcal{J}' because $\mathbf{x} \in \mathcal{J}'$ and \mathcal{M} is exterior to all medial balls \mathcal{B} . By Lemma 5.5, there exists a second point \mathbf{x}' in the intersection of \mathcal{M} and $\mathcal{N}_x \subseteq \mathcal{H}$. However, \mathbf{x}' cannot lie inside \mathcal{J}' because $\mathcal{J}' \cap \mathcal{H} = \mathbf{x}$ and $\mathbf{x}' \neq \mathbf{x}$. Consequently, \mathcal{M} must leave \mathcal{J}' in \mathcal{H}^+ , thus, it must intersect $\partial\mathcal{J}$.

Let \mathbf{y} be any intersection point of \mathcal{M} and $\partial\mathcal{J}$. Its orthogonal projection onto \mathcal{H} is the center of a medial ball $\mathbf{c} \in \mathcal{C}$ because \mathbf{y} lies in $\partial\mathcal{J} = \mathcal{C} \times \mathcal{T}_x$ and \mathcal{H} is orthogonal to \mathcal{T}_x . With the arguments from above, the

² [32, Theorem 2.1] shows that the medial axis of a subanalytic hypersurface is subanalytic. This work can be extended to general subanalytic k -manifolds [33].

lemma follows. \square

If the normals at \mathbf{x} are not known (this may be the case when the manifold is only known at a finite set of points), we can approximate the tangent space \mathcal{T}_x by the space $\tilde{\mathcal{T}}_x$ that passes through \mathbf{x} and is orthogonal to the k pole vectors of \mathbf{x} (see Definition 5.14). We can define $\tilde{\mathcal{T}}_x$ -neighbors in very much the same way as \mathcal{T}_x -neighbors: The only difference is that \mathcal{T}_x is replaced by $\tilde{\mathcal{T}}_x$.

6.2.2 Locality of the \mathcal{T}_x -neighbors

We now derive a bound on the distance between a point $\mathbf{x} \in \mathcal{M}$ and its \mathcal{T}_x -neighbors with respect to the local feature size $\text{lfs}(\mathbf{x})$. The proof of the lemma is closely inspired from [9, lemma 5].

Lemma 6.3 *Let \mathcal{P} be an ϵ -sample of \mathcal{M} and $\mathbf{x} \in \mathcal{M}$.*

(a) *The \mathcal{T}_x -neighbors of \mathbf{x} are all contained in a ball of radius $\frac{2\epsilon}{\sqrt{1-2\epsilon}}\text{lfs}(\mathbf{x})$ centered at \mathbf{x} .*

(b) *The $\tilde{\mathcal{T}}_x$ -neighbors of \mathbf{x} are contained in a ball of radius*

$$\frac{2\epsilon}{1-\epsilon} \sec(\arcsin(\frac{\epsilon}{1-\epsilon}))(1 + 4\sqrt{d-k}) \text{lfs}(\mathbf{x})$$

(c) *In both cases, the \mathcal{T}_x -neighbors and $\tilde{\mathcal{T}}_x$ -neighbors of \mathbf{x} are contained in a ball of radius $2\epsilon(1 + O(\epsilon))\text{lfs}(\mathbf{x})$.*

Proof: (a) Let \mathbf{v} be a vertex of $\bar{V}(\mathbf{x})$. We derive an upper bound on the distance between \mathbf{x} and \mathbf{v} . Because $\bar{V}(\mathbf{x})$ is bounded by the (intersection with \mathcal{T}_x of the) bisector hyperplanes of \mathbf{x} and its \mathcal{T}_x -neighbors, the distance between \mathbf{x} and its \mathcal{T}_x -neighbors is at most twice the distance between \mathbf{x} and \mathbf{v} .

Let B_1 be a ball with radius $\text{lfs}(\mathbf{x})$ that is tangent to \mathcal{M} at \mathbf{x} with center \mathbf{m}_1 such that the line segment $[\mathbf{v}\mathbf{m}_1]$ intersects \mathcal{M} at the point \mathbf{q} . (B_1 is uniquely defined only for $k = 1$ or $k = d - 1$.) Let α be the angle $\angle \mathbf{v}\mathbf{m}_1\mathbf{x}$. We find the same angle $\alpha = \angle \mathbf{v}\mathbf{x}\mathbf{q}'$ where \mathbf{q}' is the orthogonal projection of \mathbf{x} onto $[\mathbf{v}\mathbf{m}_1]$. See Figure 6.5. The ball $B(\mathbf{v}, \|\mathbf{x} - \mathbf{v}\|)$ is empty of sample points because $\mathbf{v} \in V^+(\mathbf{x})$. Because B_1 is also empty of sample points, \mathbf{x} is the point of $\mathcal{P} \cup \{\mathbf{x}\}$ that is closest to \mathbf{q} . It follows that $\|\mathbf{x} - \mathbf{q}\| \leq \frac{\epsilon}{1-\epsilon}\text{lfs}(\mathbf{x})$. On the other hand, $\|\mathbf{x} - \mathbf{q}\| \geq \|\mathbf{x} - \mathbf{q}'\| = \sin \alpha \text{lfs}(\mathbf{x})$. Hence, $\alpha \leq \arcsin(\frac{\epsilon}{1-\epsilon})$.

Since the triangle $(\mathbf{v} \ \mathbf{x} \ \mathbf{m}_1)$ has a right angle at \mathbf{x} and since the radius of B_1 is $\text{lfs}(\mathbf{x})$, we obtain

$$\|\mathbf{x} - \mathbf{v}\| = \tan(\alpha) \text{lfs}(\mathbf{x}) \leq \tan(\arcsin(\frac{\epsilon}{1-\epsilon}))\text{lfs}(\mathbf{x}) = \frac{\epsilon}{\sqrt{1-2\epsilon}}\text{lfs}(\mathbf{x})$$

with $\tan(\arcsin(\frac{\epsilon}{1-\epsilon})) = \frac{\epsilon}{\sqrt{1-2\epsilon}}$.

(b) We consider a vertex $\tilde{\mathbf{v}}$ of $\tilde{V}(\mathbf{x}) = V^+(\mathbf{x}) \cap \tilde{\mathcal{T}}_x(\mathbf{x})$. This time, we define B_1 to be a ball tangent to \mathcal{M} at \mathbf{x} with center \mathbf{m}_1 such that $[\tilde{\mathbf{v}}, \mathbf{m}_1]$ intersects \mathcal{M} in \mathbf{q} , and the angle α is defined by $\alpha = \angle \tilde{\mathbf{v}}\mathbf{m}_1\mathbf{x}$. As in part (a), $B(\tilde{\mathbf{v}}, \|\mathbf{x} - \tilde{\mathbf{v}}\|)$ and B_1 are empty of sample points, so we can apply the same arguments to obtain the bounds $\|\mathbf{x} - \mathbf{q}\| \leq \frac{\epsilon}{1-\epsilon}\text{lfs}(\mathbf{x})$ and $\alpha \leq \arcsin(\frac{\epsilon}{1-\epsilon})$. See Figure 6.6 for notations.




Figure 6.5: For the proof of Lemma 6.3(a)




Figure 6.6: For the proof of Lemma 6.3(b)

Because B_1 is tangent to \mathcal{M} at x , the unit vector \vec{n}_x with direction $\overline{m_1 x}$ lies in the normal space \mathcal{N}_x . From Lemma 5.16, we know that there exists a corresponding vector $\vec{v}_x \in \tilde{N}_x$ such that the angle β between \vec{n}_x and \vec{v}_x is bounded. Either, \vec{v}_x lies in the same plane with \tilde{v} , x , and m_1 (and, consequently, with \vec{n}_x), or we project it orthogonally onto this plane. This does not change the upper bound on the angle β ($\beta = \angle \vec{v}_x \vec{n}_x$).

Let y be the intersection point of the line (\tilde{v}, m_1) with the hyperplane perpendicular to \vec{n}_x containing x . (For $k = d - 1$, this hyperplane is \mathcal{T}_x , so $y \in \overline{V(x)}$. But this is not true in general.) The angle α is also the angle $\angle yxq'$ and β is the angle $\angle \tilde{v}xy$. Refer to Figure 6.6.

With the upper bound for α and β , we can conclude: Let $\gamma = \angle \mathbf{m}_1 \tilde{\mathbf{v}} \mathbf{x} = \frac{\pi}{2} - \alpha - \beta$. From triangle $\tilde{\mathbf{v}} \mathbf{x} \mathbf{q}'$, we know that

$$\|\mathbf{x} - \tilde{\mathbf{v}}\| = \frac{\|\mathbf{q}' - \mathbf{x}\|}{\sin \gamma} \leq \frac{\|\mathbf{x} - \mathbf{q}\|}{\sin(\frac{\pi}{2} - \alpha - \beta)} \leq \frac{\epsilon}{(1 - \epsilon) \cos(\alpha + \beta)}.$$

In the case $k = d - 1$, we know from Lemma 5.15 that $\beta \leq 2 \arcsin(\frac{\epsilon}{1-\epsilon})$, thus,

$$\|\mathbf{x} - \tilde{\mathbf{v}}\| \leq \frac{\epsilon}{1 - \epsilon} \sec(3 \arcsin(\frac{\epsilon}{1 - \epsilon})) \text{ lfs}(\mathbf{x}) = \epsilon(1 + O(\epsilon)) \text{ lfs}(\mathbf{x}).$$

In the general case, Corollary 5.16 shows that $\beta \leq 4\sqrt{d-k} \arcsin(\frac{\epsilon}{1-\epsilon})$ which yields

$$\|\mathbf{x} - \tilde{\mathbf{v}}\| \leq \frac{\epsilon}{1 - \epsilon} \sec(\arcsin(\frac{\epsilon}{1 - \epsilon}))(1 + 4\sqrt{d-k}) \text{ lfs}(\mathbf{x}) = \epsilon(1 + O(\epsilon)) \text{ lfs}(\mathbf{x}).$$

□

The following lemma states that \mathbf{x} and all its \mathcal{T}_x -neighbors are contained in a small topological ball, if $\epsilon < 0.3$. It is a direct consequence of Lemma 6.3.

Lemma 6.4 *Let $B(\mathbf{x}, r)$ be a ball centered at \mathbf{x} with radius $r = \frac{2\epsilon}{\sqrt{1-2\epsilon}} \text{ lfs}(\mathbf{x})$. If $\epsilon < \frac{\sqrt{5}-1}{4}$, $B(\mathbf{x}, r) \cap \mathcal{M}$ is a topological k -ball that contains all \mathcal{T}_x -neighbors of \mathbf{x} .*

Proof: By Lemma 6.3, all \mathcal{T}_x -neighbors of \mathbf{x} are contained in a ball $B(\mathbf{x}, r)$ centered at \mathbf{x} of radius $r = \frac{2\epsilon}{\sqrt{1-2\epsilon}} \text{ lfs}(\mathbf{x})$. With Lemma 5.4, $\mathcal{M} \cap B(\mathbf{x}, r)$ is a topological k -ball if $r < \text{lfs}(\mathbf{x})$, which is true for $\epsilon < \frac{\sqrt{5}-1}{4} \approx 0.3$. □

6.3 Manifold coordinate system

We can now define a coordinate system on \mathcal{M} associated with a set of sample points \mathcal{P} . Similarly to the \mathcal{T}_x -neighbors, it is defined in the tangent space of each manifold point and consequently called \mathcal{T}_x -coordinate system. In the rest of the section, we show some properties of the \mathcal{T}_x -coordinate system: we show that the support of the \mathcal{T}_x -coordinates is local in Subsection 6.3.1, we study their continuity properties in Subsection 6.3.2, and we prove in Subsection 6.3.3 that the barycentric coordinate property is approximately satisfied, with an error that depends on the local curvature of the surface and on the sampling density.

In this section, we suppose given a query point $\mathbf{x} \in \mathcal{M}$ together with the tangent space \mathcal{T}_x to \mathcal{M} at \mathbf{x} . Thus, the manifold including the first derivatives are known or at least query points with this information are provided. This corresponds to the setting in the application of the coordinate system to interpolation on surfaces that is described in Section 8.1.

6.3.1 Definition of the \mathcal{T}_x -coordinate function

The same definitions as in Section 6.2 apply.

Definition 6.5 (*\mathcal{T}_x -coordinate system associated with \mathcal{P}*) *In the \mathcal{T}_x -coordinate system associated with a set of sample points $\mathcal{P} \subset \mathcal{M}$, the \mathcal{T}_x -coordinate $\tau_i(\mathbf{x})$ of a point \mathbf{x} of \mathcal{M} is the natural neighbor coordinate $\lambda_i(\mathbf{x})$ of \mathbf{x} in the power diagram $\overline{\text{Vor}}(\mathcal{P})$, $i = 1, \dots, n$.*

By construction, the \mathcal{T}_x -coordinates τ_i fulfill properties (ii) and (iii) of a system of coordinates over \mathcal{M} associated with \mathcal{P} as they are listed in Definition 3.1. The barycentric coordinate property 3.1 (i) is satisfied for the projected sample points $\overline{\mathbf{p}}_i$. With respect to \mathcal{P} , the barycentric coordinate property is only true if the manifold is locally planar so that all \mathcal{T}_x -neighbors of \mathbf{x} lie in the tangent space \mathcal{T}_x .

Locally bounded support

Let Δ_i denote the support of τ_i , i.e. the subset of the points $\mathbf{x} \in \mathcal{M}$ such that $\tau_i(\mathbf{x}) \neq 0$. In order to show the locality of Δ_i on \mathcal{M} , we apply the bound on the distance between a point $\mathbf{x} \in \mathcal{M}$ and its \mathcal{T}_x -neighbors.

Corollary 6.6 *The support Δ_i of τ_i is contained in a ball $B(\mathbf{p}_i, r)$ of radius $r = \frac{2\epsilon}{\sqrt{1-2\epsilon}-\epsilon} \text{lf s}(\mathbf{p}_i)$ centered at \mathbf{p}_i . For $\epsilon < 0.19$, $B(\mathbf{p}_i, r) \cap \mathcal{M}$ is a topological k -ball that contains Δ_i .*

Proof: Applying Lemma 6.3 and Lemma 5.2, we obtain that $\|\mathbf{x} - \mathbf{p}_i\| \leq \delta \text{lf s}(\mathbf{x}) \leq \frac{\delta}{1-\delta} \text{lf s}(\mathbf{p}_i)$ with $\delta = \frac{2\epsilon}{\sqrt{1-2\epsilon}}$ for any $\mathbf{x} \in \Delta_i$. Consequently, $B(\mathbf{p}_i, r)$ with $r = \frac{\delta}{1-\delta} \text{lf s}(\mathbf{p}_i) = \frac{2\epsilon}{\sqrt{1-2\epsilon}-2\epsilon} \text{lf s}(\mathbf{p}_i)$ contains Δ_i . With Lemma 5.4, $\mathcal{M} \cap B(\mathbf{p}_i, r)$ is a topological k -ball if $r < \text{lf s}(\mathbf{x})$, thus, for $\epsilon < \frac{\sqrt{17}-1}{16} \approx 0.195$. \square

6.3.2 Continuity properties of the coordinate function

In this section, we study the continuity of the function τ_i when \mathbf{x} moves on \mathcal{M} . Let us first state the lemma:

Lemma 6.7 *The \mathcal{T}_x -coordinate $\tau_i, i = 1, \dots, n$, associated with \mathcal{P} is continuous everywhere on \mathcal{M} .*

Proof: Assume that \mathcal{M} is parameterized by $\mathbf{u} = (u_1, \dots, u_k)$. The coordinate function $\tau_i(\mathbf{x}(\mathbf{u}))$ is, by definition, equal to the natural neighbor coordinate $\lambda_i((\mathbf{x}(\mathbf{u}), 0))$ in the power diagram $\overline{\text{Vor}}(\mathcal{P})$. When \mathbf{x} moves on \mathcal{M} , the projected sample points change their position and their weight in the power diagram continuously. Indeed, since \mathcal{M} is smooth, the projection onto the tangent space is a smooth mapping: $\mathbf{x}(\mathbf{u})$ is differentiable, and the normals $\overline{\mathbf{n}}_u^j \in \mathcal{N}_x(\mathbf{u})$ are continuous. The position of $\overline{\mathbf{p}}_i(\mathbf{u})$ is given by

$$\overline{\mathbf{p}}_i(\mathbf{u}) = \mathbf{p}_i - \sum_{\overline{\mathbf{n}}_u^j \in \mathcal{N}_x(\mathbf{u})} (\overline{\mathbf{x}(\mathbf{u})\mathbf{p}_i} \cdot \overline{\mathbf{n}}_u^j) \overline{\mathbf{n}}_u^j.$$

The weight of \mathbf{p}_i is

$$w_i(\mathbf{u}) = -\|\mathbf{p}_i - \overline{\mathbf{p}}_i(\mathbf{u})\|^2 = -\sum_{\overline{\mathbf{n}}_u^j \in \mathcal{N}_x(\mathbf{u})} (\mathbf{x}(\mathbf{u})\overline{\mathbf{p}}_i \cdot \overline{\mathbf{n}}_u^j)^2.$$

Recall also from Lemma 6.2, that \mathbf{x} is always in the convex hull of the projected sample points. Consequently, the \mathcal{T}_x -coordinates are continuous on all of \mathcal{M} because the natural neighbor coordinates are continuous as proven in Lemma 3.6. \square

Lemma 6.8 *The \mathcal{T}_x -coordinate $\tau_i, i = 1, \dots, n$, associated with \mathcal{P} is continuously differentiable everywhere on \mathcal{M} except at \mathcal{P} and at the points $\mathbf{x} \in \mathcal{M}$ such that the intersection of \mathcal{T}_x with a $(d-2)$ -dimensional Voronoi face of the Voronoi cell $V^+(\mathbf{x})$ in $\text{Vor}(\mathcal{P} \cup \{\mathbf{x}\})$ has dimension $k-1$.*

Proof: Consider the power diagram $\overline{\text{Vor}}(\mathcal{P})$. By Lemma 3.6, τ_i is C^1 continuous on $\mathcal{M} \setminus \mathcal{P}$ except at a finite number of other points. We want to characterize these points of C^1 discontinuity with respect to the d -dimensional Voronoi diagram $\text{Vor}(\mathcal{P})$.

The natural neighbor coordinate τ_i is not continuously differentiable at a point \mathbf{x} such that the bisector hyperplane of \mathbf{x} and $\overline{\mathbf{p}}_i$ contains a $(k-1)$ face of $\overline{\text{Vor}}(\mathcal{P})$ say $\overline{V}(\mathbf{p}_i, \mathbf{p}_j)$. (See also the proof of Lemma 3.6.) This means that a point v of the face has equal power with respect to \mathbf{x} , B_i , and B_j . But, by definition of $\overline{\text{Vor}}^+(\mathcal{P})$, this means also that v is at equal distance from \mathbf{p}_i , \mathbf{p}_j , and \mathbf{x} . Consequently, the $(k-1)$ dimensional power face of $\overline{V}(\mathbf{x})$ is part of the intersection of a $(d-2)$ dimensional face of $\text{Vor}^+(\mathcal{P})$ with \mathcal{T}_x . \square

6.3.3 Bounding the distance to the weighted barycenter

In this section, we bound the distance from $\mathbf{x} \in \mathcal{M}$ to its weighted barycenter $\mathbf{b}(\mathbf{x})$, that is $e(\mathbf{x}) = \mathbf{x} - \sum_i \tau_i(\mathbf{x})\mathbf{p}_i$. This is the error committed with respect to the barycentric coordinate property at $\mathbf{x} \in \mathcal{M}$.

Bounding $e(\mathbf{x})$ corresponds to bounding the distance of a \mathcal{T}_x -neighbor of \mathbf{x} to the tangent space \mathcal{T}_x , since

$$\mathbf{x} = \sum_i \tau_i(\mathbf{x})\overline{\mathbf{p}}_i = \sum_i \tau_i(\mathbf{x})(\mathbf{p}_i + \sum_{j=1}^{d-k} (\overline{\mathbf{x}}\overline{\mathbf{p}}_i \cdot \overline{\mathbf{n}}_x^j) \overline{\mathbf{n}}_x^j).$$

We show the following corollary which is more general.

Corollary 6.9 *Consider two points \mathbf{x} and \mathbf{y} on \mathcal{M} with $\|\mathbf{x} - \mathbf{y}\| \leq \rho \text{lfs}(\mathbf{x})$. If $\overline{\mathbf{y}}$ is the orthogonal projection of \mathbf{y} onto the tangent space \mathcal{T}_x of \mathcal{M} at \mathbf{x} , then*

$$\|\mathbf{y} - \overline{\mathbf{y}}\| \leq \frac{\|\mathbf{y} - \mathbf{x}\|^2}{2\text{lfs}(\mathbf{x})} = \frac{\rho^2}{2}\text{lfs}(\mathbf{x}).$$

Proof: Because the projection is orthogonal, $\|\mathbf{y} - \overline{\mathbf{y}}\| = \|\mathbf{y} - \mathbf{x}\| \sin \theta$ where $\theta = \angle \mathbf{y}\mathbf{x}\overline{\mathbf{y}}$. Since \mathbf{y} does not belong to any ball of radius $\text{lfs}(\mathbf{x})$ tangent to \mathcal{M} at \mathbf{x} , $\sin \theta \leq \frac{\|\mathbf{y} - \mathbf{x}\|}{2\text{lfs}(\mathbf{x})}$. \square

It is now straightforward to show the next proposition.

Proposition 6.10

$$\mathbf{x} = \sum_i \tau_i(\mathbf{x}) \mathbf{p}_i + \sum_{j=1}^{d-k} O(\epsilon^2) \text{lfs}(\mathbf{x}) \vec{\mathbf{n}}_x^j.$$

Proof: It is sufficient to prove that the distance between a \mathcal{T}_x -neighbor $\mathbf{p}_i \in \mathcal{P}$ of $\mathbf{x} \in \mathcal{M}$ and its orthogonal projection $\overline{\mathbf{p}}_i$ on the tangent space \mathcal{T}_x , is bounded by

$$\|\mathbf{p}_i - \overline{\mathbf{p}}_i\| \leq \frac{2\epsilon^2}{1-2\epsilon} \text{lfs}(\mathbf{x}) = O(\epsilon^2) \text{lfs}(\mathbf{x}).$$

With Lemma 6.3, we have $\|\mathbf{p}_i - \mathbf{x}\| \leq \frac{2\epsilon}{\sqrt{1-2\epsilon}} \text{lfs}(\mathbf{x})$ which implies, together with Corollary 6.9, the inequality above and the proposition. \square

Bounding $\|\mathbf{p}_i - \overline{\mathbf{p}}_i\|$ with respect to the curvature

Using differential geometry, we can express the distance of a \mathcal{T}_x -neighbor to the tangent plane in terms of the curvature of the manifold. We restrict our attention to surfaces in \mathbb{R}^3 .

The *normal curvature* κ_v at a point $\mathbf{x} \in \mathcal{M}$ in a given direction $\vec{\mathbf{v}}$ in the tangent plane \mathcal{T}_x is defined as the curvature of the intersection curve of the surface with the plane formed by $\vec{\mathbf{v}}$ and $\vec{\mathbf{n}}_x$. In the sequel, we relate the distance $\|\mathbf{p}_i - \overline{\mathbf{p}}_i\|$ to the normal curvature $\kappa_{\vec{\mathbf{v}}_i}$ at $\mathbf{x} \in \mathcal{M}$ where $\vec{\mathbf{v}}_i$ is the direction of $\overline{\mathbf{p}}_i$, i.e. $\vec{\mathbf{v}}_i = \overrightarrow{\mathbf{x}\overline{\mathbf{p}}_i}$. The minimum and the maximum normal curvature are called the *principal curvatures*. Its associated directions are the *principal directions*. See for example [48].

Define a local coordinate frame with origin \mathbf{x} , the tangent plane \mathcal{T}_x as plane $z = 0$, and the two principal directions of \mathcal{M} at \mathbf{x} as x - and y -axis. It is well known that, locally, any smooth surface is the graph of a differentiable function in such a coordinate system. It follows that, in a neighborhood of \mathbf{x} , S can be represented in the form $z = h(x, y)$. In our case, z is the distance of a point to the tangent plane \mathcal{T}_x , i.e. $z_i = \|\mathbf{p}_i - \overline{\mathbf{p}}_i\|$ for $\mathbf{p}_i = (x_i, y_i, z_i)$ in the barycentric coordinate system – see [48, pp. 163-165].

Corollary 6.11 *Let $\mathbf{p}_i \in \mathcal{P}$ be a \mathcal{T}_x -neighbor of $\mathbf{x} \in \mathcal{M}$ and $\overline{\mathbf{p}}_i = (x_i, y_i, 0)$ in the coordinate frame with origin \mathbf{x} , the two principal directions as x - and y -axis, and the normal $\vec{\mathbf{n}}_x$ aligned with the z axis. Then,*

$$\|\mathbf{p}_i - \overline{\mathbf{p}}_i\| = h(x_i, y_i) = \frac{\kappa_{v_i} r_i^2}{2} + R(x_i, y_i), \text{ with } \lim_{r_i \rightarrow 0} \frac{R(x_i, y_i)}{r_i^2} = 0,$$

where $r_i = \sqrt{x_i^2 + y_i^2}$ is the distance of $\overline{\mathbf{p}}_i$ to \mathbf{x} and κ_{v_i} is defined as above.

Proof: A Taylor's expansion about $(0, 0)$ shows that $h(x, y) = \frac{1}{2}(\kappa_1 x^2 + \kappa_2 y^2) + R(x, y)$ where $R(x, y)$ is the sum of the higher order terms. This is due to the fact that, in the tangent plane, $h(0, 0) = 0$, $h_x(0, 0) = 0$, and $h_y(0, 0) = 0$. $h_{xy}(0, 0) = 0$ because the coordinate system is aligned with the principal directions. It follows that $\kappa_1(\mathbf{x}) = h_{xx}(0, 0)$ and $\kappa_2(\mathbf{x}) = h_{yy}(0, 0)$. Using polar coordinates with respect to the same coordinate system, we get $h(x, y) = \frac{r^2(\kappa_1 \cos^2 \theta + \kappa_2 \sin^2 \theta)}{2} + R(x, y)$ where $r = \sqrt{x^2 + y^2}$ and θ is the angle of the vector (x, y) with the x -axis. With $\kappa_v = \kappa_1 \cos^2 \theta + \kappa_2 \sin^2 \theta$, we get the result. \square

6.3.4 Visualization of the coordinate function

In order to demonstrate the locality and the smoothness of the T -coordinate, we visualize the coordinate function τ_i with respect to a sample point \mathbf{p}_i on a parameterized surface. Figure 6.7 shows an example of the cylinder from two different viewpoints and of the torus. τ_i is computed on a (perturbed) regular grid of 40000 points with respect to a subset of 100 randomly chosen sample points. Each grid point is translated by the value of the coordinate τ_i in direction of the surface normal at that point. Of course, \mathbf{p}_i itself has the highest value for τ_i ($\tau_i(\mathbf{p}_i) = 1$). Note the locality of the coordinate function, as well as the C^1 discontinuity on the sample point \mathbf{p}_i itself.




Figure 6.7: The coordinate function τ_i on the cylinder and the torus.

6.4 Gradient of the \mathcal{T} -coordinate function defined on a surface in \mathbb{R}^3

In this section, we derive the formula for the gradient of the T -coordinate function τ_i for the case of a surface in \mathbb{R}^3 . Theoretically, this restriction is not necessary, but it simplifies the exposition. To make it clear, we change notations and call \mathcal{S} the 2-manifold \mathcal{M} . In fact, we compute the gradient of the volume function π_i which is the volume of the cell $\bar{V}_i(\mathbf{x}) = \bar{V}(\bar{\mathbf{p}}_i) \cap \bar{V}^+(\mathbf{x})$. It defines the coordinate function as $\tau_i(X) = \frac{\pi_i(X)}{\sum_j \pi_j(X)}$.

In order to compute the gradient of π_i as the point \mathbf{x} moves on the surface, we express the sample points in a canonical coordinate frame centered on \mathbf{x} , we translate the entire frame to the origin, and then we compute $\pi_i(\mathbf{0})$ as a function of the sample points which move in dependence of \mathbf{x} and as a function of the weights which also depend on \mathbf{x} . Without loss of generality, we suppose that \mathbf{x} moves on a geodesic curve γ parameterized by arc length s . We need to evaluate the differential of the position of the sample points as well as their weights in the power diagram as a function of the curve parameter s . Furthermore, we determine the differential of $\pi_i(\mathbf{0})$ as a function of the positions as well as the weights of the sample points. Let F be the function from \mathbb{R} to $\mathbb{R}^{3n} \times \mathbb{R}^n$ that maps the curve parameter s to the positions and the weights of the sample points in the canonical coordinate frame centered on \mathbf{x} .

$$\begin{aligned} \mathbb{R} &\xrightarrow{F} \mathbb{R}^{3n} \times \mathbb{R}^n && \xrightarrow{\pi_i} \mathbb{R} \\ s &\xrightarrow{F} \{(\bar{\mathbf{p}}_j(s), w_j(s))\} && \xrightarrow{\pi_i} \pi_i(\mathbf{0})(\{(\bar{\mathbf{p}}_j(s), w_j(s))\}). \end{aligned}$$

The position and the weight of the sample points as a function of s

Let $\mathbf{x} \in S$ and \vec{t}_0 be a unit vector tangent to S at \mathbf{x} . Let the geodesic γ be parameterized by the arc length s such that $\gamma(0) = \mathbf{x}$ and $\frac{d\gamma}{ds}(0) = \vec{t}_0$. For all s , we define a canonical frame $(\vec{e}_1(s), \vec{e}_2(s), \vec{e}_3(s))$ with origin $\gamma(s)$ such that $\vec{e}_1(s) = \vec{t}(s)$ where $\vec{t}(s)$ denotes the tangent vector to γ at $\gamma(s)$, $\vec{e}_3(s) = \vec{n}(s)$ where $\vec{n}(s)$ denotes the normal vector to S at $\gamma(s)$, and $\vec{e}_2(s) = \vec{u}(s)$ where $\vec{u}(s)$ denotes the vector tangent to S at $\gamma(s)$ such that $(\vec{t}(s), \vec{u}(s), \vec{n}(s))$ forms a direct orthonormal frame.

We can express any sample point \mathbf{p}_j in the frame $(\vec{t}(s), \vec{u}(s), \vec{n}(s))$:

$$\mathbf{p}_j = \alpha_j(s) \vec{t}(s) + \beta_j(s) \vec{u}(s) + \omega_j(s) \vec{n}(s).$$

In the sequel, we omit application (s) and write for example α instead of $\alpha(s)$ knowing that all quantities depend on s .

As s varies (\mathbf{x} moves along γ), the position of \mathbf{p}_j does not change. In particular,

$$\begin{aligned} \frac{d\mathbf{p}_j}{ds} &= 0 \\ &= \alpha'_j \vec{t} + \alpha_j \frac{d\vec{t}}{ds} + \beta'_j \vec{u} + \beta_j \frac{d\vec{u}}{ds} + \omega_j \vec{n} + \omega_j \frac{d\vec{n}}{ds} \\ &= \alpha'_j \vec{t} + \alpha_j h(\vec{t}, \vec{t}) \vec{n} + \beta'_j \vec{u} + \beta_j h(\vec{t}, \vec{u}) \vec{n} + \omega'_j \vec{n} \\ &\quad - \omega_j (h(\vec{t}, \vec{t}) \vec{t} + h(\vec{t}, \vec{u}) \vec{u}), \end{aligned}$$

where h is the second fundamental form of S at $\gamma(s)$.

Projecting \mathbf{p}_j onto \vec{t} , \vec{u} and \vec{n} , we obtain the following three equations which determine α'_j , β'_j and ω'_j :

$$\begin{aligned} \alpha'_j - \omega_j h(\vec{t}, \vec{t}) &= 0 \\ \beta'_j - \omega_j h(\vec{t}, \vec{u}) &= 0 \\ \alpha_j h(\vec{t}, \vec{t}) + \beta_j h(\vec{t}, \vec{u}) - \omega'_j &= 0. \end{aligned}$$

Since the projection of \mathbf{p}_j onto the tangent plane is $\bar{\mathbf{p}}_j = \alpha_j \vec{t} + \beta_j \vec{u}$, we obtain the tangential component $\frac{d\bar{\mathbf{p}}_j}{ds} (tan)$ of $\frac{d\mathbf{p}_j}{ds}$ as

$$\frac{d\bar{\mathbf{p}}_j}{ds} (tan) = \alpha'_j \vec{t} + \beta'_j \vec{u} = \omega_j h(\vec{t}, \vec{t}) \vec{t} + \omega_j h(\vec{t}, \vec{u}) \vec{u}.$$

Still we need to move this coordinate frame, so far centered on $\gamma(s)$, to the origin:

$$d(\bar{\mathbf{p}}_j - \gamma(s))_{(tan)} ds = (\omega_j h(\vec{t}, \vec{t}) - 1) \vec{t} + \omega_j h(\vec{t}, \vec{u}) \vec{u},$$

with $\frac{d\gamma}{ds}(s) = \vec{t}$. The weight of $\bar{\mathbf{p}}_j$ is, by definition, $w_j = -\omega_j^2$. It is invariant by translation parallel to the tangent plane. We get

$$\frac{dw_j}{ds} = -2 \omega_j (\alpha_j h(\vec{t}, \vec{t}) + \beta_j h(\vec{t}, \vec{u})).$$

The gradient of π_i as a function of s

In this section, we put together the ingredients developed in Section 3.4 and in the preceding paragraphs in order to obtain a formula for the differential of $\pi_i(\mathbf{x})$ in the direction \vec{t}_0 of the tangent plane. By varying \vec{t}_0 around \mathbf{x} , the differential can be obtained in all directions of the tangent plane. We use the notations as in the previous section but all functions on s will be evaluated on 0. In particular, h denotes the second fundamental form of S at \mathbf{x} , $\vec{\mathbf{t}} = \vec{t}_0$, $\vec{\mathbf{u}} = \vec{\mathbf{u}}(0)$, $\alpha_j = \alpha_j(0)$, $\beta_j = \beta_j(0)$, and $\omega_j = \omega_j(0)$. We obtain

$$\begin{aligned}
& D(\pi_i \circ F)(0) \\
&= D\pi_i(\mathbf{0})\left(\{(\overline{\mathbf{p}}_j - \gamma(0), w_j)\}\right) \left(\frac{d(\overline{\mathbf{p}}_j - \gamma(0))}{ds} \Big|_{(tan)}(0), \frac{dw_j}{ds}(0)\right) \quad (\text{Matrices of size } 3n \times 1 \text{ and } 1 \times 3n.) \\
&= \sum_j \frac{\partial}{\partial \overline{\mathbf{p}}_j} \pi_i(\mathbf{0})(\overline{\mathbf{p}}_j - \gamma(0)) \frac{d(\overline{\mathbf{p}}_j - \gamma(0))}{ds} \Big|_{(tan)}(0) + \frac{\partial \pi_i(\mathbf{0})}{\partial w_j} \frac{dw_j}{ds}(0) \\
&= \sum_{j \neq i} \left(\frac{\partial \pi_i(\mathbf{x})}{\partial \overline{\mathbf{p}}_j} (\overline{\mathbf{p}}_j) \frac{d\overline{\mathbf{p}}_j}{ds} \Big|_{(tan)}(0) + \frac{\partial \pi_i(\mathbf{x})}{\partial w_j} \frac{dw_j}{ds}(0) \right) + \frac{\partial \pi_i(\mathbf{x})}{\partial \overline{\mathbf{p}}_i} (\overline{\mathbf{p}}_i) \frac{d\overline{\mathbf{p}}_i}{ds} \Big|_{(tan)}(0) + \frac{\partial \pi_i(\mathbf{x})}{\partial w_i} \frac{dw_i}{ds}(0) \\
&= - \sum_{j \neq i} \left(\frac{v_{ij}(\mathbf{x})}{\|\overline{\mathbf{p}}_i - \overline{\mathbf{p}}_j\|^2} (\mathbf{c}_{ij}(\mathbf{x}) - \overline{\mathbf{p}}_j) ((\omega_j h(\vec{\mathbf{t}}, \vec{\mathbf{t}}) - 1) \vec{\mathbf{t}} + \omega_j h(\vec{\mathbf{t}}, \vec{\mathbf{u}}) \vec{\mathbf{u}}) \right) \\
&\quad - \sum_{j \neq i} \left(\frac{v_{ij}(\mathbf{x})}{\|\overline{\mathbf{p}}_i - \overline{\mathbf{p}}_j\|} \omega_j (\alpha_j h(\vec{\mathbf{t}}, \vec{\mathbf{t}}) + \beta_j h(\vec{\mathbf{t}}, \vec{\mathbf{u}})) \right) \\
&\quad + \left(\sum_{j \neq i} \left(\frac{v_{ij}(\mathbf{x})}{\|\overline{\mathbf{p}}_i - \overline{\mathbf{p}}_j\|^2} (\mathbf{c}_{ij}(\mathbf{x}) - \overline{\mathbf{p}}_i) - \frac{v_{ix}}{\|\overline{\mathbf{p}}_i - \mathbf{x}\|^2} (\mathbf{c}_{ix} \overline{\mathbf{p}}_i) \right) \left((\omega_i h(\vec{\mathbf{t}}, \vec{\mathbf{t}}) - 1) \vec{\mathbf{t}} + \omega_i h(\vec{\mathbf{t}}, \vec{\mathbf{u}}) \vec{\mathbf{u}} \right) \right) \\
&\quad + \left(\sum_{j \neq i} \left(\frac{v_{ij}(\mathbf{x})}{\|\overline{\mathbf{p}}_i - \overline{\mathbf{p}}_j\|} + \frac{v_{ix}}{\|\overline{\mathbf{p}}_i - \mathbf{x}\|} \right) \omega_i \left(\alpha_i h(\vec{\mathbf{t}}, \vec{\mathbf{t}}) + \beta_i h(\vec{\mathbf{t}}, \vec{\mathbf{u}}) \right) \right).
\end{aligned}$$

Notice, that the gradient is invariant by translation of the entire coordinate frame and, therefore, $\frac{\partial}{\partial \overline{\mathbf{p}}_j} \pi_i(\mathbf{0})(\overline{\mathbf{p}}_j - \mathbf{x}) = \frac{\partial}{\partial \overline{\mathbf{p}}_j} \pi_i(\mathbf{x})(\overline{\mathbf{p}}_j)$.

As expected, in the case $h(\vec{\mathbf{t}}, \vec{\mathbf{t}}) = h(\vec{\mathbf{t}}, \vec{\mathbf{u}}) = 0$ when the surface is locally flat, we obtain the same gradient as in the plane: From Minkowski's theorem applied to the cell $\overline{V}_i(\mathbf{x})$ (see Figure 6.8, $\overline{V}_i(\mathbf{x}) = \overline{V}(\overline{\mathbf{p}}_i) \cap \overline{V}^+(\mathbf{x})$), we know that

$$\sum_{i \neq j} \frac{v_{ij}(\mathbf{x})}{\|\overline{\mathbf{p}}_i - \overline{\mathbf{p}}_j\|} (\overline{\mathbf{p}}_j - \overline{\mathbf{p}}_i) - \frac{v_{ix}}{\|\overline{\mathbf{p}}_i - \mathbf{x}\|} (\mathbf{x} - \overline{\mathbf{p}}_i) = \mathbf{0}. \quad (6.1)$$




Figure 6.8: Minkowski's theorem applies to the cell $\bar{V}_i(\mathbf{x})$.

Because the weights are 0, the gradient expression simplifies, and with Equation 6.1, we get:

$$\begin{aligned}
 D(\pi_i \circ F)(0) &= \left(\sum_{j \neq i} \left(\frac{v_{ij}(\mathbf{x})}{\|\bar{\mathbf{p}}_i - \bar{\mathbf{p}}_j\|} ((\mathbf{c}_{ij}(\mathbf{x}) - \bar{\mathbf{p}}_j) - (\mathbf{c}_{ij}(\mathbf{x}) - \bar{\mathbf{p}}_i)) + \frac{v_{ix}}{\|\bar{\mathbf{p}}_i - \mathbf{x}\|} (\mathbf{c}_{ix} - \bar{\mathbf{p}}_i) \right) \vec{\mathbf{t}} \right. \\
 &= \left(\sum_{j \neq i} \frac{v_{ij}(\mathbf{x})}{\|\bar{\mathbf{p}}_i - \bar{\mathbf{p}}_j\|} (\bar{\mathbf{p}}_i - \bar{\mathbf{p}}_j) + \frac{v_{ix}}{\|\bar{\mathbf{p}}_i - \mathbf{x}\|} (\mathbf{c}_{ix} - \bar{\mathbf{p}}_i) \right) \vec{\mathbf{t}} \\
 &= \frac{v_{ix}}{\|\bar{\mathbf{p}}_i - \mathbf{x}\|} ((\bar{\mathbf{p}}_i - \mathbf{x}) + (\mathbf{c}_{ix} - \bar{\mathbf{p}}_i)) \vec{\mathbf{t}} \\
 &= \frac{v_{ix}}{\|\bar{\mathbf{p}}_i - \mathbf{x}\|} (\mathbf{c}_{ix} - \mathbf{x}) \vec{\mathbf{t}}.
 \end{aligned}$$

Summarizing, we put the lemma that follows from this section using the notations from above.

Lemma 6.12 *The directional derivative of the volume function $\pi_i(\mathbf{x})$, $i = 1, \dots, n$, at a point $\mathbf{x} \in \mathcal{S}$ which is the volume of $\bar{V}(\mathbf{x}, \mathbf{p}_i)$ in the power diagram $\bar{\text{Vor}}(\mathbf{x})$ in direction $\vec{\mathbf{t}}_0$ is*

$$\nabla \pi_i(\mathbf{x}) \cdot \vec{\mathbf{t}}_0 = \sum_{j \neq i} \left(\frac{\partial}{\partial \bar{\mathbf{p}}_j} \pi_i(\mathbf{x})(\bar{\mathbf{p}}_j) \frac{d\bar{\mathbf{p}}_j}{ds} \Big|_{(tan)}(0) + \frac{d\pi_i(\mathbf{x})}{dw_j} \frac{dw_j}{ds} (0) \right) + \frac{\partial}{\partial \bar{\mathbf{p}}_i} \pi_i(\mathbf{x})(\bar{\mathbf{p}}_i) \frac{d\bar{\mathbf{p}}_i}{ds} \Big|_{(tan)}(0) + \frac{d\pi_i(\mathbf{x})}{dw_i} \frac{dw_i}{ds} (0).$$

6.5 Non-convex coordinate system for convex surfaces with barycentric property

In this section, we show how to construct a coordinate system on a convex $(d-1)$ -manifold that meets the barycentric property, i.e. we construct a coordinate function χ_j^i such that $\mathbf{x} = \sum_j \chi_j^i(\mathbf{x})\mathbf{p}_j$. \mathcal{P} is supposed to be an ϵ -sample with $\epsilon < 0.1$. Because of the restriction to hypersurfaces, we denote the manifold by \mathcal{S} (instead of \mathcal{M}).

As explained in Section 6.3.3, the \mathcal{T}_x -coordinate system does not necessarily verify the barycentric property unless all \mathcal{T}_x -neighbors of \mathbf{x} lie in the tangent (hyper-)plane \mathcal{T}_x of \mathcal{S} at \mathbf{x} . Of course, to obtain the barycentric property, we cannot avoid to give up the convexity condition and we must allow negative coordinates. Otherwise, either \mathbf{x} would have to lie in the convex hull of its neighbors or the partition of unity property of Definition 3.1 would be violated.

Inspired from Clarkson's idea to construct smoother coordinate functions (see Section 3.5), we combine the coordinate functions for \mathbf{x} with those for its T -neighbor \mathbf{p}_i in the Voronoi diagram $\text{Vor}^+(\mathcal{P} \setminus \{\mathbf{p}_i\})$ restricted to the tangent plane \mathcal{T}_x . It is denoted $\overline{\text{Vor}}^+(\mathcal{P} \setminus \{\mathbf{p}_i\})$. In the tangent plane, with respect to the projected data points $\overline{\mathbf{p}}_j$, the barycentric coordinate property is satisfied, i.e.

$$\mathbf{x} = \sum_i \tau_i(\mathbf{x})\overline{\mathbf{p}}_i. \quad (6.2)$$

Similarly, we can express a neighbor $\overline{\mathbf{p}}_i$ of \mathbf{x} as the weighted barycenter of \mathbf{x} and its other neighbors $\overline{\mathbf{p}}_j$ in $\overline{\text{Vor}}^+(\mathcal{P} \setminus \{\mathbf{p}_i\})$. As in Section 3.5, we arrange the equation for $\overline{\mathbf{p}}_i$ as to write \mathbf{x} as a convex combination of the points $\overline{\mathbf{p}}_j$ and $\overline{\mathbf{p}}_i$. Together with equation 6.2, we obtain two independent ways to write \mathbf{x} as a convex combination of the projected data points $\overline{\mathbf{p}}_j$. (We explain this in detail later.) Re-placing the projected data points $\overline{\mathbf{p}}_j$ by the original points \mathbf{p}_j yields in both cases an error in direction of the normal $\vec{\mathbf{n}}_x$ because $\overline{\mathbf{p}}_j$ is the projection of \mathbf{p}_j along $\vec{\mathbf{n}}_x$. Since \mathcal{S} is convex, we show later that the error is always in opposite directions, so a convex combination of both equations allows to cancel the error out and to establish the barycentric coordinate property.

The definition is restricted to hypersurfaces because in this case the normal space is one-dimensional and so is the error concerning the barycentric coordinate property. For manifolds with higher co-dimension, one has to be careful about the number of (linearly independent) equations needed to cancel out the error.

Definition of the non-convex \mathcal{T}_x -coordinate function χ_j^i

We define the new coordinate function χ_j^i . Let $\tau_x^+(\mathbf{p}_i)$ be the natural neighbor coordinate of $B_i = (\overline{\mathbf{p}}_i, -\|\mathbf{p}_i - \overline{\mathbf{p}}_i\|^2)$ with respect to \mathbf{x} in the power diagram $\overline{\text{Vor}}^+(\mathcal{P} \setminus \{\mathbf{p}_i\})$ defined in \mathcal{T}_x . Let $\tau_j^+(\mathbf{p}_i)$ be the natural neighbor coordinate of B_i with respect to $B_j = (\overline{\mathbf{p}}_j, -\|\mathbf{p}_j - \overline{\mathbf{p}}_j\|^2)$ in the same diagram.

We noticed that the barycentric coordinate property is satisfied in the tangent space, so that $\overline{\mathbf{p}}_i$ can be written as convex combination of its neighbors $\overline{\mathbf{p}}_j$ weighted by $\tau_j^+(\mathbf{p}_i)$. We denote by $I(\mathbf{x})$ the set

of indices of the \mathcal{T}_x -neighbors of \mathbf{x} and by $J_i(\mathbf{x})$ the set of indices of the natural neighbors of B_i in $\overline{\text{Vor}}^+(\mathcal{P} \setminus \{\mathbf{p}_i\})$ except \mathbf{x} .

$$\overline{\mathbf{p}}_i = \tau_x^+(\mathbf{p}_i) \mathbf{x} + \sum_{j \in J_i(\mathbf{x})} \tau_j^+(\mathbf{p}_i) \overline{\mathbf{p}}_j \iff \mathbf{x} = \frac{1}{\tau_x^+(\mathbf{p}_i)} \left(\overline{\mathbf{p}}_i - \sum_{j \in J_i(\mathbf{x})} \tau_j^+(\mathbf{p}_i) \overline{\mathbf{p}}_j \right). \quad (6.3)$$

Replacing the projected data points $\overline{\mathbf{p}}_i$ by the original points \mathbf{p}_i , we get

$$\mathbf{b}_i = \frac{1}{\tau_x^+(\mathbf{p}_i)} \left(\mathbf{p}_i - \sum_{j \in J_i(\mathbf{x})} \tau_j^+(\mathbf{p}_i) \mathbf{p}_j \right).$$

Let the signed distance of \mathbf{p}_i to the tangent hyperplane \mathcal{T}_x be denoted by $h_i = (\overline{\mathbf{p}}_i - \mathbf{p}_i) \cdot \overline{\mathbf{n}}_x$. Since \mathcal{S} is convex, we can suppose that $h_i > 0, i = 1, \dots, n$. (Otherwise, the orientation of $\overline{\mathbf{n}}_x$ can be reversed.) With Equation (6.3) and $\overline{\mathbf{p}}_i = \mathbf{p}_i + h_i \cdot \overline{\mathbf{n}}_x$, the (signed) distance between \mathbf{b}_i and \mathbf{x} in direction $\overline{\mathbf{n}}_x$ is

$$\begin{aligned} e_i &= (\mathbf{b}_i - \mathbf{x}) \cdot \overline{\mathbf{n}}_x = \left(\frac{1}{\tau_x^+(\mathbf{p}_i)} \left(\mathbf{p}_i - \sum_{j \in J_i(\mathbf{x})} \tau_j^+(\mathbf{p}_i) \mathbf{p}_j \right) - \frac{1}{\tau_x^+(\mathbf{p}_i)} \left(\overline{\mathbf{p}}_i - \sum_{j \in J_i(\mathbf{x})} \tau_j^+(\mathbf{p}_i) \overline{\mathbf{p}}_j \right) \right) \cdot \overline{\mathbf{n}}_x \\ &= \frac{1}{\tau_x^+(\mathbf{p}_i)} \left(-h_i + \sum_{j \in J_i(\mathbf{x})} \tau_j^+(\mathbf{p}_i) h_j \right). \end{aligned} \quad (6.4)$$

From Section 6.3.3, we re-call the error e_x concerning the barycentric coordinate property of the \mathcal{T}_x -coordinates:

$$e_x = (\mathbf{x} - \mathbf{b}(\mathbf{x})) \cdot \overline{\mathbf{n}}_x = \left(\sum_{i \in I(\mathbf{x})} \tau_i(\mathbf{x}) \overline{\mathbf{p}}_i - \sum_{i \in I(\mathbf{x})} \tau_i(\mathbf{x}) \mathbf{p}_i \right) \cdot \overline{\mathbf{n}}_x = \sum_i \tau_i(\mathbf{x}) h_i$$

We need to show that e_i and e_x have the same sign, thus, supposing $e_x > 0$, we need to show $e_i > 0$.

Lemma 6.13 *Let $h_i > 0, i = 1, \dots, n$. If \mathcal{S} is convex and $\epsilon < 0.1$, then $e_i > 0$.*

Proof: From Equation 6.4, it follows that

$$e_i > 0 \iff h_i < \sum_{j \in J_i(\mathbf{x})} \tau_j^+(\mathbf{p}_i) h_j.$$

Lemma 5.19 and Lemma 6.3 (about the locality of the \mathcal{T}_x -neighbors) allow to show that the angle between $\overline{\mathbf{n}}_x$ and the normal to \mathcal{S} at any point \mathbf{p} inside the ball B_x containing all \mathcal{T}_x -neighbors of \mathbf{x} is smaller than $\frac{\pi}{2}$ if $\epsilon < 0.1$. Thus, inside B_x , \mathcal{S} is locally the graph of the height function over \mathcal{T}_x . In the coordinate frame with \mathcal{T}_x as hyperplane $x_d = 0$, we get $\mathbf{p}_i = (\overline{\mathbf{p}}_i, h_i)$ and $\mathbf{x} = (\mathbf{x}, 0)$. The convexity of \mathcal{S} implies that \mathbf{p}_i lies outside $\text{conv}(\mathcal{P} \setminus \{\mathbf{p}_i\})$, thus, \mathbf{p}_i lies below the convex hull of its neighbors, i.e. it has a smaller x_{d-1}

coordinate in this coordinate frame than the convex hull points. Since $\bar{\mathbf{p}}_i$ is the convex combination of its neighbors, i. e.

$$\bar{\mathbf{p}}_i = \tau_x^+(\mathbf{p}_i) \mathbf{x} + \sum_{j \in J_i(\mathbf{x})} \tau_j^+(\mathbf{p}_i) \bar{\mathbf{p}}_j,$$

it follows that

$$h_i < \tau_x^+(\mathbf{p}_i) 0 + \sum_{j \in J_i(\mathbf{x})} \tau_j^+(\mathbf{p}_i) h_j.$$

□

Both errors e_i and e_x are in direction $\vec{\mathbf{n}}_x$ and we have

$$\mathbf{x} = \mathbf{b}_i - e_i \cdot \vec{\mathbf{n}}_x \quad \text{and} \quad \mathbf{x} = \mathbf{b}(\mathbf{x}) + e_x \cdot \vec{\mathbf{n}}_x.$$

The convex combination of both expressions yields

$$\begin{aligned} \mathbf{x} &= \frac{e_x}{e_i + e_x} \mathbf{b}_i + \frac{e_i}{e_i + e_x} \mathbf{b}(\mathbf{x}) \\ &= \frac{e_x}{e_i + e_x} \left(\frac{1}{\tau_x^+(\mathbf{p}_i)} (\mathbf{p}_i - \sum_{j \in J_i(\mathbf{x})} \tau_j^+(\mathbf{p}_i) \mathbf{p}_j) \right) + \frac{e_i}{e_i + e_x} \left(\sum_{i \in I(\mathbf{x})} \tau_i(\mathbf{x}) \mathbf{p}_i \right) \\ &=: \sum_{j \in I(\mathbf{x}) \cup J_i(\mathbf{x})} \chi_j^i(\mathbf{x}) \mathbf{p}_j. \end{aligned} \tag{6.5}$$

The coordinate function $\chi_j^i(\mathbf{x})$ is obtained by re-grouping all terms for \mathbf{p}_j .

$$\chi_i^i(\mathbf{x}) = \frac{e_i}{e_i + e_x} \tau_i(\mathbf{x}) + \frac{e_x}{e_i + e_x} \frac{1}{\tau_x^+(\mathbf{p}_i)} \tag{6.6}$$

$$\chi_j^i(\mathbf{x}) = \frac{e_i}{e_i + e_x} \tau_j(\mathbf{x}) - \frac{e_x}{e_i + e_x} \frac{\tau_j^+(\mathbf{p}_i)}{\tau_x^+(\mathbf{p}_i)}, \quad i \neq j. \tag{6.7}$$

Finally, for all $i \in I(\mathbf{x})$, we mix Equation (6.5) weighted by $\tau_i(\mathbf{x})$:

$$\mathbf{x} = \sum_{i \in I(\mathbf{x})} \sum_{j \in J_i(\mathbf{x})} \tau_i(\mathbf{x}) \chi_j^i(\mathbf{x}) \mathbf{p}_j.$$

If $e_x = 0$, we obtain $\mathbf{x} = \sum_{i \in I(\mathbf{x})} \tau_i(\mathbf{x}) \mathbf{p}_i$.

Locality of the non-convex \mathcal{T}_x -coordinate function χ_j^i

In the next lemma, we bound the maximum distance from \mathbf{x} to a sample point \mathbf{p}_j with $\chi_j^i(\mathbf{x}) \neq 0$. The normal $\vec{\mathbf{n}}_x$ to \mathcal{M} at \mathbf{x} is supposed to be known or estimated in a way that the angle between $\vec{\mathbf{n}}_x$ and the estimated normal is $O(\epsilon)$.

Lemma 6.14 *Let \mathcal{P} be an ϵ -sample of \mathcal{M} and $\mathbf{x} \in \mathcal{M}$. All sample points \mathbf{p}_j with $\chi_j^i(\mathbf{x}) \neq 0$ are contained in a ball of radius $6 \epsilon (1 + O(\epsilon)) \text{ lfs}(\mathbf{x})$ centered at \mathbf{x} .*

Proof: Let \mathbf{p}_i be a \mathcal{T}_x -neighbor of \mathbf{x} . If \mathbf{p}_j is a sample point with $\chi_j^i(\mathbf{x}) \neq 0$ then \mathbf{p}_i and \mathbf{p}_j have a common face in the Voronoi diagram $\overline{\text{Vor}}^+(\mathcal{P})$ that is restricted to \mathcal{T}_x . With Lemma 5.11 we can easily bound the distance between \mathbf{x} and an \mathcal{M} -neighbor \mathbf{p}_k of \mathbf{p}_i . We see this in detail later. However, the neighbors of \mathbf{p}_i in the Voronoi diagram restricted to \mathcal{T}_x must not necessarily be the same as the \mathcal{M} -neighbors (in the Voronoi diagram restricted to \mathcal{M}). As explained in Chapter 5, vertices of the d -dimensional Voronoi diagram might be close to \mathcal{M} so that the face $V(\mathbf{p}_i, \mathbf{p}_j)$ that intersects \mathcal{T}_x does not intersect \mathcal{M} .




Figure 6.9: The region \mathcal{I} includes $\overline{V}(\mathbf{p}_i)$ and therefore \mathbf{v} . (This situation is not realistic.)

Let \mathcal{H}_{ik}^i be the halfspace that is delimited by the bisector hyperplane \mathcal{H}_{ik} of \mathbf{p}_i and an \mathcal{M} -neighbor \mathbf{p}_k of \mathbf{p}_i and that contains \mathbf{p}_i . Consider the intersection \mathcal{I} of halfspaces \mathcal{H}_{ik}^i ,

$$\mathcal{I} = \bigcap \mathcal{H} \in \{\mathcal{H}_{ik}^i \mid \mathbf{p}_k \text{ is } \mathcal{M}\text{-neighbor of } \mathbf{p}_i\}.$$

Refer to the schematic drawing in Figure 6.9. By definition, the Voronoi cell $V(\mathbf{p}_i)$ is contained in \mathcal{I} because it is the intersection of the halfspaces \mathcal{H}_{ij}^i of all Delaunay neighbors of \mathbf{p}_i . A line through \mathbf{x} and a vertex $\mathbf{v} \in \overline{V}(\mathbf{p}_i)$ must intersect the boundary of \mathcal{I} in a point \mathbf{v}' that lies on a bisector hyperplane \mathcal{H}_{ik} of \mathbf{p}_i and an \mathcal{M} -neighbor \mathbf{p}_k of \mathbf{p}_i . Either \mathbf{v}' is further from \mathbf{x} or $\mathbf{v} = \mathbf{v}'$. In the sequel, we determine an upper bound for $\|\mathbf{x} - \mathbf{v}'\|$.

Because \mathbf{p}_i is a \mathcal{T}_x -neighbor of \mathbf{x} , Lemma 6.3 implies that the distance between \mathbf{x} and \mathbf{p}_i is

$$\|\mathbf{x} - \mathbf{p}_i\| \leq 2\epsilon (1 + O(\epsilon)) \text{lfs}(\mathbf{x}) \quad (6.8)$$

From Lemma 5.11 and the Lipschitz property of $\text{lfs}()$, we know that the distance from \mathbf{p}_i to \mathbf{p}_k is at most

$$\|\mathbf{p}_i - \mathbf{p}_k\| \leq \frac{2\epsilon}{1-\epsilon} \text{lfs}(\mathbf{p}_i) \leq 2\epsilon (1 + O(\epsilon)) \text{lfs}(\mathbf{x}). \quad (6.9)$$

We deduce that

$$\|\mathbf{x} - \mathbf{p}_k\| \leq \|\mathbf{x} - \mathbf{p}_i\| + \|\mathbf{p}_i - \mathbf{p}_k\| \leq 4\epsilon(1 + O(\epsilon))\text{fs}(\mathbf{x}). \quad (6.10)$$

Because \mathbf{p}_i and \mathbf{p}_k are \mathcal{M} -neighbors, there exists a point \mathbf{v}_m of the Voronoi face $V_{\mathcal{M}}(\mathbf{p}_i, \mathbf{p}_k)$ and \mathbf{p}_k and \mathbf{p}_i are the closest sample points to \mathbf{v}_m . Refer to Figure 6.10. Therefore,

$$\|\mathbf{p}_i - \mathbf{v}_m\| \leq \frac{\epsilon}{1 - \epsilon}\text{fs}(\mathbf{p}_i) \leq \epsilon(1 + O(\epsilon))\text{fs}(\mathbf{x}).$$




Figure 6.10: The point \mathbf{v}' cannot be far from \mathbf{x} .

To bound $\|\mathbf{x} - \mathbf{v}'\|$, we compute a bound on (1) $\|\mathbf{x} - \mathbf{v}_m\|$ and (2) $\|\mathbf{v}_m - \overline{\mathbf{v}_m}\|$ (where $\overline{\mathbf{v}_m}$ is the orthogonal projection of \mathbf{v}_m onto \mathcal{T}_x). Then, (3), we bound the distance from $\overline{\mathbf{v}_m}$ to \mathbf{v}' , in order to obtain

$$\|\mathbf{x} - \mathbf{v}'\| \leq \|\mathbf{x} - \overline{\mathbf{v}_m}\| + \|\overline{\mathbf{v}_m} - \mathbf{v}'\| \leq \|\mathbf{x} - \mathbf{v}_m\| + \|\mathbf{v}_m - \overline{\mathbf{v}_m}\| + \|\overline{\mathbf{v}_m} - \mathbf{v}'\|. \quad (6.11)$$

(1) By the triangle inequality, we know that

$$\|\mathbf{x} - \mathbf{v}_m\| \leq \|\mathbf{x} - \mathbf{p}_i\| + \|\mathbf{p}_i - \mathbf{v}_m\| \leq 3\epsilon(1 + O(\epsilon))\text{fs}(\mathbf{x}). \quad (6.12)$$

(2) With Equation 6.12 and Corollary 6.9, we obtain

$$\|\mathbf{v}_m - \overline{\mathbf{v}_m}\| = \frac{\|\mathbf{x} - \mathbf{v}_m\|^2}{2\text{fs}(\mathbf{x})} \leq \frac{9}{2}\epsilon^2(1 + O(\epsilon))^2\text{fs}(\mathbf{x}) = O(\epsilon^2)\text{fs}(\mathbf{x}). \quad (6.13)$$

(3) Now we can bound $\|\overline{\mathbf{v}_m} - \mathbf{v}'\|$. With Lemma 5.20, we know a lower bound on the angle α between the bisector hyperplane \mathcal{H}_{ik} and \mathcal{T}_x , and, consequently, an upper bound on the angle β at \mathbf{v}_m between the line $(\mathbf{v}_m, \mathbf{v}')$ and $(\mathbf{v}_m, \overline{\mathbf{v}_m})$. See also Figure 6.10. Notably,

$$\beta \leq \arcsin\left(\frac{\rho}{1 - \rho}\right) + \frac{\rho}{1 - 3\rho} \quad \text{with} \quad \rho = \frac{\|\mathbf{x} - \mathbf{p}_k\|}{\text{fs}(\mathbf{x})} \leq 4\epsilon(1 + O(\epsilon)), \quad \text{thus,} \quad \tan \beta = O(\epsilon).$$

$$\|\overline{\mathbf{v}_m} - \mathbf{v}'\| = \tan(\beta) \|\overline{\mathbf{v}_m} - \mathbf{v}_m\| = O(\epsilon^3) \text{lfs}(\mathbf{x}). \quad (6.14)$$

With Equations 6.12, 6.13 and 6.14, Equation 6.11 becomes

$$\begin{aligned} \|\mathbf{x} - \mathbf{v}\| &\leq \|\mathbf{x} - \mathbf{v}'\| \leq (3\epsilon(1 + O(\epsilon)) + O(\epsilon^2) + O(\epsilon^3)) \text{lfs}(\mathbf{x}) \\ &= \frac{3\epsilon}{1 - 3\epsilon} (1 + O(\epsilon)) \text{lfs}(\mathbf{x}). \end{aligned}$$

Suppose that \mathbf{v} is the furthest vertex $\mathbf{v} \in \overline{V}(\mathbf{p}_i)$ from \mathbf{x} . A sample point \mathbf{p}_j so that $\mathbf{v} \in \overline{V}(\mathbf{p}_i, \mathbf{p}_j)$ is at most twice as far from \mathbf{x} than \mathbf{v} since $\mathbf{v} \in V(\mathbf{p}_j)$ implies that $\|\mathbf{v} - \mathbf{p}_j\| \leq \|\mathbf{v} - \mathbf{x}\|$. With $\|\mathbf{x} - \mathbf{p}_j\| \leq 2\|\mathbf{x} - \mathbf{v}\|$, we obtain the result. \square

The support Λ_j of the all coordinate functions $\chi_j^i, i = 1, \dots, n$, is

$$\Lambda_j := \{\mathbf{x} \in \mathcal{U} \mid \chi_j^i(\mathbf{x}) \neq 0, i = 1, \dots, n\}.$$

It includes the support Δ_j of τ_j , i.e. $\Delta_j \subseteq \Lambda_j$. Additionally, it includes all points \mathbf{x} that have a \mathcal{T}_x -neighbor B_i which is a neighbor of B_j in the Voronoi diagram $\overline{\text{Vor}}^+(\mathcal{P} \setminus \{\mathbf{p}_i\})$ restricted to the tangent space of \mathbf{x} .

Corollary 6.15 *The support Λ_j of $\cup_i \chi_j^i$ is contained in a ball of radius $6\epsilon(1 + O(\epsilon)) \text{lfs}(\mathbf{p}_j)$ centered at \mathbf{p}_j .*

Proof: Applying Lemma 6.14 and Lemma 5.2, we obtain $\|\mathbf{x} - \mathbf{p}_j\| \leq \delta \text{lfs}(\mathbf{x}) \leq \frac{\delta}{1-\delta} \text{lfs}(\mathbf{p}_j) = 6\epsilon(1 + O(\epsilon))$ with $\delta = 6\epsilon(1 + O(\epsilon))$. \square

Continuity properties of the non-convex \mathcal{T}_x -coordinate function χ_j^i

The continuity properties of χ_j^i depend on the continuity properties of the \mathcal{T}_x -coordinate functions. We suppose without real loss of generality that the intersection of a tangent plane $\mathcal{T}_i, i = 1, \dots, n$ with a $(d-2)$ -dimensional Voronoi face of the Voronoi cell $V(\mathbf{p}_i)$ is either empty or it has dimension $k-2$. If it has dimension $k-1$, a small perturbation of \mathcal{T}_i removes this situation.

Lemma 6.16 *The coordinate function $\chi_j^i, i = 1, \dots, n, j \in I(\mathbf{x})$, is continuous and continuously differentiable everywhere on \mathcal{S} except at the points $\mathbf{x} \in \mathcal{S}$ such that the intersection of \mathcal{T}_x with a $(d-2)$ -dimensional Voronoi face of the Voronoi cell $V^+(\mathbf{x})$ in $\text{Vor}(\mathcal{P} \cup \{\mathbf{x}\})$ has dimension $k-1$.*

Proof: $\chi_j^i(\mathbf{x})$ is built from the functions $\tau_i(\mathbf{x}), \tau_j^+(\mathbf{p}_i)$ and $\tau_x^+(\mathbf{p}_i)$. Lemma 6.7 and Lemma 6.8 show the continuity properties of τ_i . With the same approach as in the proof of Lemma 3.7, we show that $\tau_j^+(\mathbf{p}_i)$ and $\tau_x^+(\mathbf{p}_i)$ have the same properties as $\tau_i(\mathbf{x})$. \square




Figure 6.11: The coordinate function $\sum_{i \in I_1(\mathbf{x})} \tau_i(\mathbf{x}) \chi_j^i(\mathbf{x})$ on the cylinder.

Visualization of the coordinate function

We demonstrate the locality and the smoothness of the non-convex T -coordinate as we did for the convex \mathcal{T}_x -coordinates. We visualize the coordinate function $\sum_{i \in I_1(\mathbf{x})} \tau_i(\mathbf{x}) \chi_j^i(\mathbf{x})$ with respect to a sample point \mathbf{p}_j on a parameterized surface. Figure 6.11 shows an example of the cylinder from two different view-points. $\sum_{i \in I_1(\mathbf{x})} \tau_i(\mathbf{x}) \chi_j^i(\mathbf{x})$ is computed on a (perturbed) regular grid of 40000 points with respect to a subset of 100 regularly placed sample points. Each grid point is translated by the value of the coordinate sum $\sum_j \tau_i(\mathbf{x}) \chi_j^i(\mathbf{x})$ in direction of the surface normal at that point. Of course, \mathbf{p}_j itself has the biggest coordinate value, $\sum_{i \in I_1(\mathbf{x})} \tau_i(\mathbf{p}_j) \chi_j^i(\mathbf{p}_j) = 1$. Notice the locality of the coordinate function and the zone where the function values are negative. The C^1 discontinuity at the sample point \mathbf{p}_j is less obvious than for the convex coordinate function τ_j .

Chapter 7

The Voronoi atlas of a point cloud on a k -manifold of \mathbb{R}^d

Suppose (again) that \mathcal{M} is a smooth and bounded k -manifold without boundary embedded in \mathbb{R}^d and $\mathcal{P} \subset \mathcal{M}$ is a set of sample point that fulfills the ϵ -sampling condition. In the preceding chapter, we examined the local restriction of the Euclidean Voronoi diagram of \mathcal{P} to the tangent space of a point $\mathbf{x} \in \mathcal{P}$. We showed that, in this restriction, the Voronoi cell of \mathbf{x} has good properties, namely, it is well defined everywhere on \mathcal{M} , it has a small, bounded diameter and it is homeomorphic to a neighborhood of \mathbf{x} in \mathcal{M} . The results apply equivalently if \mathbf{x} is a sample point $\mathbf{p}_i \in \mathcal{P}$ whose Voronoi cell $V(\mathbf{p}_i)$ is restricted to its tangent space \mathcal{T}_i . In this chapter, we define the (\mathcal{T} -restricted) Voronoi atlas of \mathcal{P} . This collection of




Figure 7.1: (a) The (\mathcal{T} -restricted) Voronoi atlas of a point set \mathcal{P} and (b) its dual.

cells contains the restriction of each Voronoi cell of the Euclidean Voronoi diagram to the tangent space of the point that defines the cell. Note that it is not a partition of a domain, as this is usually the case of a

Voronoi diagram because the so-called \mathcal{T} -restricted Voronoi cells have, in general, no common boundary. Refer to Figure 7.1(a). However, for each Voronoi cell on the manifold, there exist a bijection between this cell and the corresponding Voronoi cell in the tangent space, hence, the name “Voronoi atlas”.

Each \mathcal{T} -restricted Voronoi cell reveals the properties that were shown in the preceding chapter. In particular, it is homeomorphic to the corresponding cell in the Voronoi diagram restricted to \mathcal{M} – if the sampling condition is satisfied. Furthermore, both cells are close one to each other because they are contained in small balls centered on the sample point. The dimension of a \mathcal{T} -restricted Voronoi cell is the dimension of the manifold. We define also the dual of the \mathcal{T} -restricted Voronoi atlas. It is a collection of simplices that are part of the Delaunay triangulation of \mathcal{P} . For each sample point, a subset of the simplices incident to the point forms a k -manifold. In Chapter 8, we will show that this dual is an important tool to compute piecewise linear approximations of \mathcal{M} . Figure 7.1 depicts a Voronoi atlas as well as its dual.

For hypersurfaces ($k = d - 1$), the definition domain of the Voronoi atlas can be enlarged to connect the cells such that we obtain a Voronoi diagram that is, indeed, the covering of a hypersurface.¹ The nerve theorem can be applied to show that the dual Delaunay complex is homotopy equivalent to the hypersurface.²

In Section 7.3, we define a system of coordinates in the Voronoi atlas which has the advantage to be defined in the neighborhood of the manifold without exact knowledge of it.

7.1 Definition and basic properties

In this section, we define formally the \mathcal{T} -restricted Voronoi atlas associated with a set of points scattered on a k -manifold. For $k = d - 1$, there is an alternative definition of the Voronoi atlas which allows to enlarge the cells of the atlas such that it is actually a Voronoi diagram. This is presented in Section 7.2.

7.1.1 The Voronoi atlas

In the general case, let \mathcal{M} be a smooth k -manifold of \mathbb{R}^d , $k < d$, and $\mathcal{P} \subset \mathcal{M}$ be an ϵ -sample of \mathcal{M} . $\text{Vor}(\mathcal{P})$ is the Voronoi diagram of \mathcal{P} and $V(\mathbf{p}_i)$ is the Voronoi cell of \mathbf{p}_i . Let \mathcal{T}_i be the tangent space of \mathcal{M} at $\mathbf{p}_i \in \mathcal{P}$. The normal space \mathcal{N}_i is spanned by the normal vectors $\{\vec{\mathbf{n}}_i^l\}_{l=1,\dots,d-k}$. We suppose that the normal space is known. In the contrary case, it can be estimated with the techniques defined in Chapter 5.

Definition 7.1 (Voronoi atlas) *The (\mathcal{T} -restricted) Voronoi atlas of \mathcal{P} is*

$$\mathcal{A}(\mathcal{P}) := \bigcup_{i=1}^n (\mathcal{T}_i \cap V(\mathbf{p}_i)).$$

¹ Re-call from Chapter 5 that a covering of $\mathcal{F} \subseteq \mathbb{R}^d$ is a finite collection \mathcal{E} of subsets of \mathcal{F} such that $\mathcal{F} = \cup \mathcal{E}$.

² See Theorem 5.9.

In Observation 2.4, we showed that the intersection of $\text{Vor}(\mathcal{P})$ with a hypersurface of \mathbb{R}^d is a power diagram. The \mathcal{T} -restricted Voronoi atlas is therefore a collection of cells of different power diagrams because each cell is defined in the intersection of a tangent space with $\text{Vor}(\mathcal{P})$. The Voronoi cell of \mathbf{p}_i restricted to \mathcal{T}_i is

$$\overline{V}(\mathbf{p}_i) = V(\mathbf{p}_i) \cap \mathcal{T}_i.$$

The $(d-l)$ -dimensional face $V(\mathbf{p}_{j_0}, \dots, \mathbf{p}_{j_l})$, $l \leq k$, intersects \mathcal{T}_i in the $(k-l)$ -dimensional face

$$\overline{V}^i(\mathbf{p}_{j_0}, \dots, \mathbf{p}_{j_l}) = V(\mathbf{p}_{j_0}, \dots, \mathbf{p}_{j_l}) \cap \mathcal{T}_i.$$

The sample points $\mathbf{p}_{j_0}, \dots, \mathbf{p}_{j_l}$ are the *generators* of the face. Each vertex $\mathbf{v}_t = \overline{V}^i(\mathbf{p}_{j_0}, \dots, \mathbf{p}_{j_k})$ of $\mathcal{A}(\mathcal{P})$ is the center of an empty ball passing through the $k+1$ generators of \mathbf{v}_t . We call them the \mathcal{A} -Delaunay balls.

The orthogonal projection of a point \mathbf{x} onto \mathcal{T}_i is called $\overline{\mathbf{x}}^i$. Similarly, the sample point $\mathbf{p}_j \in \mathcal{P}$ projects orthogonally onto $\overline{\mathbf{p}}_j^i$. Let $\overline{\text{Vor}}^i(\mathcal{P})$ be the power diagram of the weighted points $B_j = (\overline{\mathbf{p}}_j^i, -\|\mathbf{p}_j - \overline{\mathbf{p}}_j^i\|^2)$ and $B_i = (\mathbf{p}_i, 0)$ defined in the tangent space \mathcal{T}_i .

A homeomorphism between the Voronoi cells $V_{\mathcal{M}}(\mathbf{p}_i)$ and $\overline{V}(\mathbf{p}_i)$

If \mathcal{P} is an ϵ -sample, $V(\mathbf{p}_i)$ intersects \mathcal{M} in a topological k -ball: The Voronoi cell $V_{\mathcal{M}}(\mathbf{p}_i)$ is contained in a small ball $B(\mathbf{p}_i, \rho \text{ lfs}(\mathbf{p}_i))$ centered on \mathbf{p}_i with radius $\rho \text{ lfs}(\mathbf{p}_i) \leq \frac{\epsilon}{1-\epsilon} \text{ lfs}(\mathbf{p}_i)$ (see also the proof of Lemma 5.11). With Lemma 5.4, $\mathcal{M} \cap B(\mathbf{p}_i, \rho \text{ lfs}(\mathbf{p}_i))$ is a topological k -ball if $\rho < 1$. In fact, $\frac{\epsilon}{1-\epsilon} < 1$ if $\epsilon < \frac{1}{2}$. Since $V_{\mathcal{M}}(\mathbf{p}_i)$ is the intersection of the d -dimensional convex polytope $V(\mathbf{p}_i)$ with \mathcal{M} and it is contained in $B(\mathbf{p}_i, \rho \text{ lfs}(\mathbf{p}_i))$, it is a topological k -ball.

Similarly, the cell $\overline{V}(\mathbf{p}_i)$ is homeomorphic to a topological k -ball because it is a (non-empty) cell in a k -dimensional power diagram. It is always non-empty and it contains \mathbf{p}_i because $\mathbf{p}_i \in V(\mathbf{p}_i)$ and $\mathbf{p}_i \in \mathcal{T}_i$, thus, $\mathbf{p}_i \in \overline{V}(\mathbf{p}_i) = V(\mathbf{p}_i) \cap \mathcal{T}_i$. Furthermore, $\overline{V}(\mathbf{p}_i)$ is close to $V_{\mathcal{M}}(\mathbf{p}_i)$ if the sampling is sufficiently dense. If the tangent space \mathcal{T}_i is known, the proof of Lemma 6.3 yields that the ball $B(\mathbf{p}_i, \rho \text{ lfs}(\mathbf{p}_i))$ contains $\overline{V}(\mathbf{p}_i)$ if $\rho = \frac{\epsilon}{\sqrt{1-2\epsilon}}$. With $\rho < 1 \iff \epsilon < \sqrt{2} - 1 < 0.41$, $B(\mathbf{p}_i, \rho \text{ lfs}(\mathbf{p}_i)) \cap \mathcal{M}$ is a topological k -ball. These remarks lead to the next lemma.

Lemma 7.2 *For $\epsilon < \sqrt{2} - 1 < 0.41$, the Voronoi cell $V_{\mathcal{M}}(\mathbf{p}_i)$ and the Voronoi cell $\overline{V}(\mathbf{p}_i)$ are homeomorphic.*

Proof: As argued before, in this case, both Voronoi cells, $V_{\mathcal{M}}(\mathbf{p}_i)$ and $\overline{V}(\mathbf{p}_i)$, are topological k -balls and therefore homeomorphic one to another. \square

We deduce that there exist a bijection between $V_{\mathcal{M}}(\mathbf{p}_i)$ and $\overline{V}(\mathbf{p}_i)$.

For curves ($k = 1$), the separate cells can be easily connected in order to obtain a global homeomorphism between the curve and a Voronoi atlas. Recall from Lemma 5.13 that, for $k = 1$, all Voronoi vertices

are far from the curve. Therefore, one can show that, for $\epsilon < 0.31$, a Voronoi edge $V(\mathbf{p}_i, \mathbf{p}_j)$ which is intersected by \mathcal{T}_i is also intersected by \mathcal{T}_j . This will be proven in Theorem 8.1.

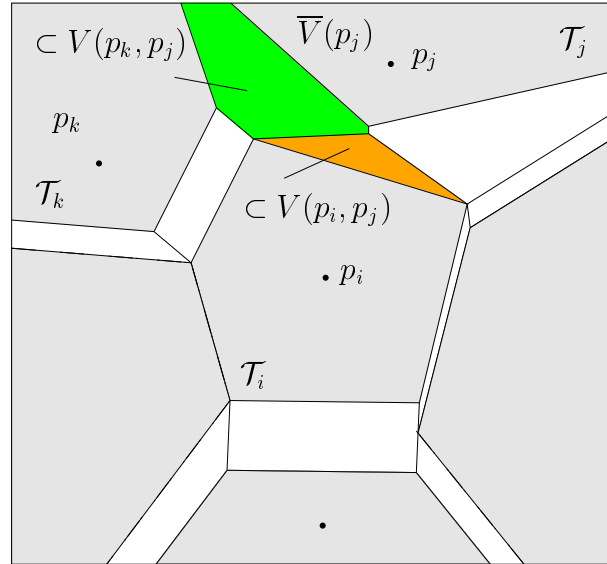


Figure 7.2: The Voronoi face $V(\mathbf{p}_i, \mathbf{p}_j)$ intersects \mathcal{T}_i but not \mathcal{T}_j . A top view on the cell $\overline{V}(\mathbf{p}_i)$ and some of its neighboring cells.

In general, it is not as easy to join a face $\overline{V}^i(\mathbf{p}_i, \mathbf{p}_j)$ to a corresponding face in \mathcal{T}_j because this face might not exist. Indeed, for $k > 1$, the d -dimensional Voronoi diagram might have Voronoi vertices that are close to \mathcal{M} (and, consequently, close to the tangent spaces \mathcal{T}_i and \mathcal{T}_j) so that a Voronoi face $V(\mathbf{p}_i, \mathbf{p}_j)$ might intersect \mathcal{T}_i but not \mathcal{T}_j . This is indicated in Figure 7.2.

7.1.2 The dual of the Voronoi atlas

The dual of $\mathcal{A}(\mathcal{P})$ is a collection of k -simplices $\mathcal{G}(\mathcal{P})$ that are part of the Delaunay triangulation $\text{Del}(\mathcal{P})$. It is defined in the sequel. At first, we define a subset of simplices of $\mathcal{G}(\mathcal{P})$ that is adjacent to a point $\mathbf{p}_i \in \mathcal{P}$ and that is homeomorphic to the neighborhood of \mathbf{p}_i in \mathcal{M} .

Definition 7.3 (Umbrella) *The (local) umbrella U_i of $\mathbf{p}_i \in \mathcal{P}$ is the set of k -faces of $\text{Del}(\mathcal{P})$ that are dual to the vertices of $\overline{V}^i(\mathbf{p}_i)$.*

It is called the umbrella of \mathbf{p}_i because it is a set of k -simplices incident to \mathbf{p}_i that forms a k -manifold.³ In three-dimensions, this resembles an umbrella because it is the set of triangles incident to one point that forms a surface around it. U_i forms a k -manifold because it is the dual of the k -dimensional Voronoi cell $\overline{V}(\mathbf{p}_i)$.

³ This name has been used for the three-dimensional case by Adamy et al. [1].

Definition 7.4 (\mathcal{T}_i -neighbors) The \mathcal{T}_i -neighbors of $\mathbf{p}_i \in \mathcal{P}$ are the vertices adjacent to \mathbf{p}_i in its umbrella U_i .

The \mathcal{T} -neighbors are all pairs of points \mathbf{p}_i and \mathbf{p}_j such that \mathbf{p}_j is a \mathcal{T}_i -neighbor of \mathbf{p}_i or \mathbf{p}_i is a \mathcal{T}_j -neighbor of \mathbf{p}_j . They will be connected by an edge in the dual of the \mathcal{T} -restricted Voronoi atlas $\mathcal{A}(\mathcal{P})$.




Figure 7.3: $\mathcal{G}(\mathcal{P})$ contains the entire sliver tetrahedra $(\mathbf{p}_i, \mathbf{p}_j, \mathbf{p}_l, \mathbf{p}_k)$ whose dual Voronoi vertex is v .

It is important to notice that this neighborhood relation is not necessarily symmetric. If \mathbf{p}_i has a \mathcal{T}_i -neighbor \mathbf{p}_j , this does not necessarily mean that \mathbf{p}_i is a \mathcal{T}_j -neighbor \mathbf{p}_j . This reflects the fact that a Voronoi face $V(\mathbf{p}_i, \mathbf{p}_j)$ might intersect \mathcal{T}_i but not \mathcal{T}_j if \mathbf{p}_i and \mathbf{p}_j are vertices of a sliver, i.e. a very flat simplex of the Delaunay triangulation whose vertices lie almost on a $(d - 1)$ -ball. The dual Voronoi vertex of this sliver lies between \mathcal{T}_i and \mathcal{T}_j . Figure 7.1 depicts a Voronoi atlas as well as its dual in the two-dimensional case. In Figure 7.3, we indicate the atlas as well as the dual in 3D.

Definition 7.5 (Dual of the Voronoi atlas) The dual of the \mathcal{T} -restricted Voronoi atlas $\mathcal{A}(\mathcal{P})$ is the set of simplices

$$\mathcal{G}(\mathcal{P}) := \bigcup_{i=1}^n U_i,$$

where U_i is the local umbrella of $\mathbf{p}_i \in \mathcal{P}$, i.e. the set of k -faces of $\text{Del}(\mathcal{P})$ that are dual to a vertex of $\bar{V}^i(\mathbf{p}_i)$.

In $\mathcal{G}(\mathcal{P})$, a $(k - 1)$ -dimensional simplex dual to $V(\mathbf{p}_i, \mathbf{p}_j)$ might be incident to only one k -dimensional simplex. Similarly, it might be incident to more than two k -dimensional simplices. In other words, $\mathcal{G}(\mathcal{P})$,

the collection of all umbrellas $U_i, i = 1, \dots, n$, is not combinatorially a k -dimensional triangulation nor homeomorphic to a k -manifold. A simplex that is generated by the \mathcal{T} -restricted Voronoi cell $\bar{V}(\mathbf{p}_i)$ with vertex \mathbf{p}_j is not necessarily part of the umbrella U_j . Figure 7.4 shows a reduced $\mathcal{G}(\mathcal{P})$ for two different data sets. Only those triangles are depicted that are part of at least l umbrellas with $l = 2$ in Figure 7.4(a) and $l = 3$ in Figure 7.4(b).




Figure 7.4: Reduced version of $\mathcal{G}(\mathcal{P})$ containing all triangles present in at least l umbrellas. (a) Model “Schale” with $|\mathcal{P}| = 2714, l = 2$, running time: 7,2s for $\text{Del}(\mathcal{P}) + 2\text{s}$ for $\mathcal{G}(\mathcal{P})$, (b) Model “Octopus” with $|\mathcal{P}| = 16,944, l = 3$, running time: 60s for $\text{Del}(\mathcal{P}) + 15\text{s}$ for $\mathcal{G}(\mathcal{P})$

7.2 Enlargement for hypersurfaces

For hypersurfaces, $k = d - 1$, there is an alternative definition of the \mathcal{T} -restricted Voronoi atlas which allows to enlarge the cells of the atlas such that the collection of cells is actually a Voronoi diagram. The resulting \mathcal{T} -restricted Voronoi diagram is defined by a subset of faces of a d -dimensional Voronoi diagram which form a $(d - 1)$ -manifold. The basic idea has been introduced by Nullans [83] for the reconstruction of geographic data.

In this section, we suppose that S is a closed orientable hypersurface and \mathcal{F} be the region enclosed by S . The point set \mathcal{P} is now an ϵ -sample of S and we suppose that the normal $\bar{\mathbf{n}}_i$ of S at a point $\mathbf{p}_i \in \mathcal{P}$ is known and oriented to the exterior of \mathcal{F} . In the contrary case, it can be estimated with the techniques defined in Chapter 5 and a consistent orientation can be obtained by propagation (see e.g. [8]). Let \mathcal{T}_i

be the tangent (hyper-)plane of \mathbf{p}_i . The positive halfspace defined by \mathcal{T}_i and $\vec{\mathbf{n}}_i$ is called \mathcal{T}_i^+ and the corresponding negative halfspace is called \mathcal{T}_i^- .

7.2.1 The \mathcal{T} -restricted Voronoi diagram

We construct a point set \mathcal{P}' in which each point $\mathbf{p}_i \in \mathcal{P}$ is replaced by two points \mathbf{p}_i^- and \mathbf{p}_i^+ that are slightly perturbed in each direction of the normal $\vec{\mathbf{n}}_i$. Let δ be an arbitrarily small positive real value.

$$\mathcal{P}' := \bigcup_{i=1}^n \{\mathbf{p}_i^+, \mathbf{p}_i^-\}, \text{ where } \mathbf{p}_i^- = \mathbf{p}_i - \delta \vec{\mathbf{n}}_i \text{ and } \mathbf{p}_i^+ = \mathbf{p}_i + \delta \vec{\mathbf{n}}_i.$$

The perturbation of \mathbf{p}_i is a *symbolic perturbation*. Conceptually, each geometric coordinate is replaced with a symbolically perturbed coordinate which means that δ is, in fact, a polynomial rather than a constant. When evaluating geometric tests, the substitution of the symbolically perturbed coordinates in the test results in a polynomial whose coefficients are determined by the original geometric coordinates.




Figure 7.5: (a) Duplicating the points separates all faces that are intersected by the tangent plane in two parts. (b) The face $V(\mathbf{p}_i, \mathbf{p}_j)$ is partitioned in $V(\mathbf{p}_i^-, \mathbf{p}_j^-)$, $V(\mathbf{p}_i^+, \mathbf{p}_j^-)$ and $V(\mathbf{p}_i^+, \mathbf{p}_j^+)$.

The bisector hyperplane of \mathbf{p}_i^- and \mathbf{p}_i^+ is the tangent hyperplane \mathcal{T}_i to \mathcal{S} at \mathbf{p}_i . Duplicating the point \mathbf{p}_i in this manner, splits the Voronoi cell $V(\mathbf{p}_i)$ in two parts, $V(\mathbf{p}_i^-) = V(\mathbf{p}_i) \cap \mathcal{T}_i^-$ and $V(\mathbf{p}_i^+) = V(\mathbf{p}_i) \cap \mathcal{T}_i^+$. Refer to figure 7.5(a). The Voronoi face $V(\mathbf{p}_i^-, \mathbf{p}_i^+)$ between both cells is exactly the cell $\bar{V}(\mathbf{p}_i)$ (with $\bar{V}(\mathbf{p}_i) = V(\mathbf{p}_i) \cap \mathcal{T}_i$). This is proven in the next lemma. The lower-dimensional faces of $V(\mathbf{p}_i)$ that are intersected by \mathcal{T}_i are likewise divided. See figure 7.5(b).

All faces of $\text{Vor}(\mathcal{P})$ that are not intersected by the bisector hyperplanes of their generators exist also in $\text{Vor}(\mathcal{P}')$. An empty sphere that is centered on a point of a face of $\text{Vor}(\mathcal{P})$ and that has the generators

of this face on its boundary, remains empty if each generator \mathbf{p}_i is replaced by \mathbf{p}_i^- and \mathbf{p}_i^+ because the latter points are at an infinitely small distance to \mathbf{p}_i . The Voronoi face has either the generator \mathbf{p}_i^- or \mathbf{p}_i^+ . It cannot have both \mathbf{p}_i^- and \mathbf{p}_i^+ as generators unless it is intersected by their bisector hyperplane, the tangent plane \mathcal{T}_i . This is indicated in Figure 7.6(a).

Lemma 7.6 *The Voronoi diagram $\text{Vor}(\mathcal{P}')$ contains all faces of $\mathcal{A}(\mathcal{P})$.*

Proof: Any point $\mathbf{v} \in \overline{V}(\mathbf{p}_i) = V(\mathbf{p}_i) \cap \mathcal{T}_i$ is at equal distance to \mathbf{p}_i^- and \mathbf{p}_i^+ and further away from any other point of \mathcal{P}' . Consequently, $\overline{V}(\mathbf{p}_i)$ is the $(d-1)$ -face $V(\mathbf{p}_i^+, \mathbf{p}_i^-)$ of $\text{Vor}(\mathcal{P}')$.

A $(d-l)$ -dimensional face F of $\overline{V}(\mathbf{p}_i)$, $0 < l \leq d$, is generated by the intersection of a $(d-l+1)$ -dimensional face of $\text{Vor}(\mathcal{P})$ with \mathcal{T}_i (see Section 2.4). A point $\mathbf{v} \in F$ is the center of an empty ball B with the l generators of F on its boundary. When \mathcal{P} is replaced by \mathcal{P}' , the same ball B is still empty. Since \mathbf{v} lies in \mathcal{T}_i , thus, on the bisector hyperplane of \mathbf{p}_i^- and \mathbf{p}_i^+ , it must contain both points, \mathbf{p}_i^- and \mathbf{p}_i^+ , on its boundary. \mathbf{v} is, therefore, the center of an empty ball with $l+1$ points of \mathcal{P}' on its boundary. Notice that B cannot contain both duplicates of any other generator $\mathbf{p}_j \neq \mathbf{p}_i$ of F unless the tangent spaces \mathcal{T}_j and \mathcal{T}_i intersect in \mathbf{v} . This is a degenerate situation that is removed by a small displacement of \mathbf{p}_i or \mathbf{p}_j . Generically, \mathbf{v} lies in a $(d-l)$ -face of $\text{Vor}(\mathcal{P}')$. See figure 7.6 (b). \square




Figure 7.6: (a) The balls B_1 and B_2 must be empty of sample points if \mathbf{p}_i and \mathbf{p}_j are duplicated. (b) $v_3 \in \mathcal{T}_i$ is the center of the empty ball B_3 with \mathbf{p}_i^- and \mathbf{p}_i^+ on its boundary. The tangent planes of \mathbf{p}_i and \mathbf{p}_k intersect in \mathbf{v}_4 . The empty ball B_4 passes through the duplicates of both sample points \mathbf{p}_i and \mathbf{p}_j .

If the sampling is sufficiently dense, except in the neighborhood of \mathcal{T}_i , $V(\mathbf{p}_i^-)$ lies inside \mathcal{F} and $V(\mathbf{p}_i^+)$ lies outside \mathcal{F} . This idea will be used to approximate \mathcal{F} by the union of Voronoi cells $V(\mathbf{p}_i^-)$. The boundary of this d -volume is a hypersurface which should approximate the surface S .

Definition 7.7 (\mathcal{T} -restricted Voronoi diagram) *Let S be a hypersurface and $\mathcal{P} \subset S$. The normal to S at $\mathbf{p}_i \in \mathcal{P}$ is $\vec{\mathbf{n}}_i$ and $\mathbf{p}_i^\pm = \mathbf{p}_i \pm \delta \vec{\mathbf{n}}_i$ for some arbitrary small real value δ . $\mathcal{P}' = \bigcup_{i=1}^n \{\mathbf{p}_i^+, \mathbf{p}_i^-\}$. A*

“bipolar” facet of $\text{Vor}(\mathcal{P}')$ is a facet $V(\mathbf{p}_i^-, \mathbf{p}_j^+)$. The \mathcal{T} -restricted Voronoi diagram of \mathcal{P} is the cell complex of bipolar facets of the Voronoi diagram $\text{Vor}(\mathcal{P}')$ and their sub-faces,

$$\text{Vor}_{\mathcal{T}}(\mathcal{P}) := \bigcup_{1 \leq i, j \leq n} V(\mathbf{p}_i^-, \mathbf{p}_j^+).$$




Figure 7.7: (a) $\text{Vor}_{\mathcal{T}}(\mathcal{P})$ contains all bipolar facets. (b) A zoom on the cell $V_{\mathcal{T}}(\mathbf{p}_i)$. Its annulus is $V(\mathbf{p}_i^-, \mathbf{p}_j^+)$ and $V(\mathbf{p}_i^-, \mathbf{p}_k^+)$.

We can show that the union of the faces of $\text{Vor}_{\mathcal{T}}(\mathcal{P})$ is a hypersurface: The facets that constitute $\text{Vor}_{\mathcal{T}}(\mathcal{P})$ are the boundary of the union of the Voronoi cells $V(\mathbf{p}_i^-)$, $\mathbf{p}_i \in \mathcal{P}$. The cells are convex solids which may share common $(d - 1)$ -dimensional faces. Any lower dimensional face shared by two cells of $\text{Vor}(\mathcal{P}')$ is necessarily part of their common $(d - 1)$ -dimensional face because, in the dual triangulation, the existence of a simplex with vertices p and q implies the existence of the edge between p and q which implies the existence of the dual $(d - 1)$ -dimensional Voronoi face. Thus, the boundary of the union of the cells is a hypersurface and, furthermore, $\text{Vor}_{\mathcal{T}}(\mathcal{P})$ is a covering of it (refer to Section 5.2).

Figure 7.7(a) depicts a two-dimensional \mathcal{T} -restricted Voronoi diagram $\text{Vor}_{\mathcal{T}}(\mathcal{P})$. The cell of \mathbf{p}_i in $\text{Vor}_{\mathcal{T}}(\mathcal{P})$ is the union of the facets of $\text{Vor}_{\mathcal{T}}(\mathcal{P})$ that are incident to $V(\mathbf{p}_i^-)$ as shown in figure 7.7(b).

Definition 7.8 The cell of $\mathbf{p}_i \in \mathcal{P}$ in $\text{Vor}_{\mathcal{T}}(\mathcal{P})$ is

$$V_{\mathcal{T}}(\mathbf{p}_i) := \bigcup_{1 \leq j \leq n} V(\mathbf{p}_i^-, \mathbf{p}_j^+).$$

The core of the cell $V_{\mathcal{T}}(\mathbf{p}_i)$ is $V(\mathbf{p}_i^-, \mathbf{p}_i^+)$. The annulus of $V_{\mathcal{T}}(\mathbf{p}_i)$ are the facets $V(\mathbf{p}_i^-, \mathbf{p}_j^+)$, $1 \leq j \leq n$, $i \neq j$, together with their faces and subfaces.

An annulus facet $V_{\mathcal{T}}(\mathbf{p}_i, \mathbf{p}_j)$ is either the facet $V(\mathbf{p}_i^-, \mathbf{p}_j^+)$ or $V(\mathbf{p}_i^+, \mathbf{p}_j^-)$ of $\text{Vor}(\mathcal{P}')$ depending on which one exists. Both exist only if $\mathcal{T}_i = \mathcal{T}_j$ because in this case, the facets $V(\mathbf{p}_i^-, \mathbf{p}_j^+)$ and $V(\mathbf{p}_i^+, \mathbf{p}_j^-)$ coincide and their dimension is $d - 2$. In the general case, $V_{\mathcal{T}}(\mathbf{p}_i, \mathbf{p}_j)$ is $(d - 1)$ -dimensional.

The definition of $\text{Vor}_{\mathcal{T}}(\mathcal{P})$ should be coupled with a condition on the Voronoi diagram $\text{Vor}(\mathcal{P}')$ in order to impose that its cells are connected and contractible. If the condition is fulfilled, the \mathcal{T} -restricted Voronoi diagram and, in particular, its dual triangulation (defined in the next section) have nice properties because we will show in the sequel that the nerve theorem applies.

Definition 7.9 *The \mathcal{T} -restricted Voronoi diagram $\text{Vor}_{\mathcal{T}}(\mathcal{P})$ is valid if the union of the bipolar facets of each cell $V(\mathbf{p}_i^-)$, $\mathbf{p}_i \in \mathcal{P}$, is a topological $(d-1)$ -ball.*

For small dimensions, $d \leq 3$, this condition is easily checked. It suffices to compute the bipolar facets of $\text{Vor}(\mathcal{P}')$ incident to \mathbf{p}_i^- while controlling that they build a connected component without holes.

The \mathcal{T} -restricted Voronoi diagram for an ϵ -sample \mathcal{P}

In this section, we consider the case that \mathcal{P} is an ϵ -sample of S . This is not imposed in Definition 7.7 but it allows to show some additional properties of $\text{Vor}_{\mathcal{T}}(\mathcal{P})$.

From the proof of Lemma 6.3, we know that the cell $\bar{V}(\mathbf{p}_i)$ is contained in a small ball centered at \mathbf{p}_i . The following lemma shows that any intersection of a Voronoi facet $V(\mathbf{p}_i, \mathbf{p}_j)$ with \mathcal{T}_j is also contained in a small ball centered on \mathbf{p}_i . Consequently, no tangent plane \mathcal{T}_j intersects a face $V(\mathbf{p}_i, \mathbf{p}_j)$ of $V(\mathbf{p}_i)$ far from \mathbf{p}_i .

Lemma 7.10 *A point \mathbf{v} of the intersection of the tangent plane \mathcal{T}_j of S at \mathbf{p}_j with the Voronoi facet $V(\mathbf{p}_i, \mathbf{p}_j)$ of $\text{Vor}(\mathcal{P})$, i.e. $\mathbf{v} \in V(\mathbf{p}_i, \mathbf{p}_j) \cap \mathcal{T}_j$, is at distance*

$$\|\mathbf{v} - \mathbf{p}_i\| \leq \frac{\epsilon}{\sqrt{1-2\epsilon}-2\epsilon} \text{lfs}(\mathbf{p}_i).$$

Proof: From the proof of Lemma 6.3, we know that

$$\|\mathbf{v} - \mathbf{p}_k\| = \|\mathbf{v} - \mathbf{p}_i\| \leq \frac{\epsilon}{\sqrt{1-2\epsilon}} \text{lfs}(\mathbf{p}_k).$$

Due to the Lipschitz property of lfs , we obtain

$$\begin{aligned} \|\mathbf{v} - \mathbf{p}_i\| &\leq \frac{\epsilon}{\sqrt{1-2\epsilon}} (\text{lfs}(\mathbf{p}_i) + \|\mathbf{p}_i - \mathbf{p}_k\|) \leq \frac{\epsilon}{\sqrt{1-2\epsilon}} (\text{lfs}(\mathbf{p}_i) + 2 \|\mathbf{v} - \mathbf{p}_i\|) \\ \Leftrightarrow (1 - \frac{2\epsilon}{\sqrt{1-2\epsilon}}) \|\mathbf{v} - \mathbf{p}_i\| &\leq \frac{\epsilon}{\sqrt{1-2\epsilon}} \text{lfs}(\mathbf{p}_i) \\ \Leftrightarrow \|\mathbf{v} - \mathbf{p}_i\| &\leq \frac{\epsilon}{\sqrt{1-2\epsilon}-2\epsilon} \text{lfs}(\mathbf{p}_i). \square \end{aligned}$$

Recall that \mathcal{F} is the region enclosed by S .

Lemma 7.11 *Consider a point $\mathbf{v} \in V(\mathbf{p}_i, \mathbf{p}_j)$ with*

$$\|\mathbf{v} - \mathbf{p}_i\| = \nu \text{lfs}(\mathbf{p}_i) \geq \frac{\epsilon(1-\epsilon)^2}{\sqrt{1-2\epsilon}} \text{lfs}(\mathbf{p}_i).$$

(a) If $\mathbf{v} \in \mathcal{F}$, then $\mathbf{v} \in V(\mathbf{p}_i^-, \mathbf{p}_j^-)$, i.e. $\mathbf{v} \in V(\mathbf{p}_i, \mathbf{p}_j) \cap \mathcal{T}_i^- \cap \mathcal{T}_j^-$.

(b) If $\mathbf{v} \notin \mathcal{F}$, then $\mathbf{v} \in V(\mathbf{p}_i^+, \mathbf{p}_j^+)$, i.e. $\mathbf{v} \in V(\mathbf{p}_i, \mathbf{p}_j) \cap \mathcal{T}_i^+ \cap \mathcal{T}_j^+$.

Proof: With Lemma 5.23, we know that the vector from \mathbf{p}_i to \mathbf{v} makes a small angle with the normal $\vec{\mathbf{n}}_i$ that is oriented in the same direction, $\angle(\vec{\mathbf{n}}_i, \vec{\mathbf{p}}_i\mathbf{v}) \leq \arcsin \frac{\epsilon}{\nu(1-\epsilon)} + \arcsin \frac{\epsilon}{(1-\epsilon)}$. Thus, if $\mathbf{v} \notin \mathcal{F}$ and the normal is oriented to the exterior of \mathcal{F} , we obtain (using maple) that

$$\angle(\vec{\mathbf{n}}_i, \vec{\mathbf{p}}_i\mathbf{v}) \leq \arcsin \frac{\epsilon}{\nu(1-\epsilon)} + \arcsin \frac{\epsilon}{(1-\epsilon)} \leq \frac{\pi}{2} \quad \text{if} \quad \nu \geq \frac{\epsilon(1-\epsilon)^2}{\sqrt{1-2\epsilon}}.$$

The same arguments works if $\mathbf{v} \in \mathcal{F}$ except that the normal is oriented to the opposite direction. \square

At last in this section, we state an open question which we believe to be true but which we could not yet prove.

Open question 7.12 *There exists a positive real constant c such that $\text{Vor}_{\mathcal{T}}(\mathcal{P})\overline{V}(\mathbf{p}_i)$ is valid in the sense of Definition 7.9 if \mathcal{P} is an ϵ -sample of \mathcal{S} with $\epsilon < c$.*

7.2.2 The \mathcal{T} -restricted Delaunay triangulation

In the preceding section, we defined the \mathcal{T} -restricted Voronoi diagram and we argued that it is a covering of a hypersurface. In this section, we define the dual Delaunay triangulation of this Voronoi diagram. We show that, if $\text{Vor}_{\mathcal{T}}(\mathcal{P})$ is valid, the nerve theorem (Theorem 5.9) can be applied which guarantees that the dual triangulation is homotopy equivalent to the domain covered by the Voronoi diagram, thus, to a hypersurface.

Definition 7.13 (\mathcal{T} -restricted Delaunay triangulation) *With the notations of Definition 7.7, the \mathcal{T} -restricted Delaunay triangulation $\text{Del}_{\mathcal{T}}(\mathcal{P})$ contains all faces of $\text{Del}(\mathcal{P})$ that are dual to faces in $\text{Vor}_{\mathcal{T}}(\mathcal{P})$. It consists of the edges that are dual to bipolar facets of $\text{Vor}(\mathcal{P}')$.*

Lemma 7.14 *If the \mathcal{T} -restricted Voronoi diagram $\text{Vor}_{\mathcal{T}}(\mathcal{P})$ is valid, then the \mathcal{T} -restricted Delaunay triangulation $\text{Del}_{\mathcal{T}}(\mathcal{P})$ is homotopy equivalent to the hypersurface defined by the union of cells of $\text{Vor}_{\mathcal{T}}(\mathcal{P})$.*

Proof: Let \mathcal{X} be the hypersurface that consists of the union of faces of $\text{Vor}_{\mathcal{T}}(\mathcal{P})$. It is indeed a hypersurface because it is the boundary of the union of Voronoi cells $V(\mathbf{p}_i^-)$, $i = 1, \dots, n$. We apply the nerve theorem, Theorem 5.9, to the restriction of $\text{Vor}(\mathcal{P})$ to \mathcal{X} . The restriction of $\text{Vor}(\mathcal{P})$ to a hypersurface is a finite closed covering of the latter (cf. [54]). In our case, it is $\text{Vor}_{\mathcal{T}}(\mathbf{p}_i)$.

We need to show that the conditions of the nerve theorem are satisfied, notably, the intersection of every set of Voronoi cells $V_{\mathcal{T}}(\mathbf{p}_i), V_{\mathcal{T}}(\mathbf{p}_j), \dots$ must be empty or contractible. Already, a cell $V_{\mathcal{T}}(\mathbf{p}_i)$ is contractible because $\text{Vor}_{\mathcal{T}}(\mathcal{P})$ is valid. The faces of $\text{Vor}_{\mathcal{T}}(\mathcal{P})$ are faces of $\text{Vor}(\mathcal{P}')$. By definition, they are convex and contractible. \square

7.3 Natural-neighbor coordinates defined in the Voronoi atlas

Let \mathcal{P} be an ϵ -sample of a k -manifold \mathcal{M} embedded in \mathbb{R}^d . The natural neighbor coordinates in the \mathcal{T} -restricted Voronoi atlas of \mathcal{P} can be defined in the usual manner. The difference between the \mathcal{T}_x -coordinate functions defined in the preceding chapter is quite subtle. While the \mathcal{T}_x -coordinate function $\tau_i(\mathbf{x})$ (of the preceding chapter) is defined for a point $\mathbf{x} \in \mathcal{M}$ in the tangent space \mathcal{T}_x of \mathbf{x} , the natural neighbor coordinates in the Voronoi atlas are defined in the tangent space \mathcal{T}_i of each sample point. Its definition domain Δ is restricted to the neighborhood of \mathcal{M} . Precisely, Δ is the union of \mathcal{A} -Delaunay balls, i.e. the union of (open) balls that are centered on vertices of the \mathcal{T} -restricted Voronoi atlas and empty of sample points. We show that Δ encloses the manifold \mathcal{M} if \mathcal{P} is an ϵ -sample of \mathcal{M} .

7.3.1 Definition of the coordinate function μ_i

Before we define the coordinate function more formally, we explain the intuition behind it. At the insertion of a point $\mathbf{x} \in \Delta$ into the d -dimensional Voronoi diagram $\text{Vor}(\mathcal{P})$, we consider $V^+(\mathbf{x})$, the cell of \mathbf{x} in $\text{Vor}(\mathcal{P} \cup \{\mathbf{x}\})$, and its subcells $V_i(\mathbf{x}) = V^+(\mathbf{x}) \cap V(\mathbf{p}_i)$. If the tangent space \mathcal{T}_i intersects $V_i(\mathbf{x})$, \mathbf{x} has a non-zero coordinate with respect to \mathbf{p}_i . This intersection is, in fact, the cell $\overline{V}_i(\mathbf{x})$ in the power diagram $\overline{\text{Vor}}^i(\mathcal{P})$ (defined by the intersection of $\text{Vor}(\mathcal{P})$ with \mathcal{T}_i). In this power diagram, \mathbf{x} corresponds to the weighted point $X_i = (\overline{\mathbf{x}}^i, -\|\overline{\mathbf{x}}^i - \mathbf{x}\|^2)$. The volume of the cell $\overline{V}_i(\mathbf{x})$ is denoted $\pi_i(X_i)$. (Of course, X_i is a function of \mathbf{x} .)

The natural coordinate function in the \mathcal{T} -restricted Voronoi atlas is defined as follows:

Definition 7.15 (Atlas-coordinate system associated with \mathcal{P}) *In the Atlas-coordinate system associated with a set of sample points $\mathcal{P} \subset \mathcal{M}$, the \mathcal{A} -coordinate function $\mu_i(\mathbf{x})$ of a point $\mathbf{x} \in \Delta$ is defined by*

$$\mu_i(\mathbf{x}) := \frac{\pi_i(X_i)}{\sum_j \pi_j(X_j)}.$$

The sample points with non-zero coordinate functions are called the neighbors of the point \mathbf{x} in the Voronoi atlas.

Definition 7.16 (Atlas-neighbors of \mathbf{x}) *Given a set of sample points $\mathcal{P} \subset \mathcal{M}$, the Atlas-neighbors of a point $\mathbf{x} \in \Delta$ are the sample points $\mathbf{p}_i \in \mathcal{P}$ such that $\mu_i(\mathbf{x}) \neq 0$.*

Verifying the conditions of Definition 3.1

In Chapter 3, we have analyzed the function $\pi_i(X_i)$ in detail. The coordinate function $\mu_i(\mathbf{x})$ inherits the continuity properties of $\pi_i(X_i)$ because it is simply the ratio of several functions $\pi_i(X_i)$ – even if they are not defined in the same power diagram. Since X_i changes smoothly with \mathbf{x} , the results hold. We can deduce that μ_i is C^0 continuous and that it is C^1 continuous except in a set of measure zero. We can

determine its gradient because Lemma 3.8 defines the gradient of $\pi_i(X_i)$. We establish these results in more detail in the next section.

The Atlas-coordinate system fulfills two properties of Definition 3.1: $\mu_i(\mathbf{x}) = \delta_{ij}$ because $\mu_j(\mathbf{p}_i) = 0$ for $j \neq i$ and $\mu_i(\mathbf{p}_i) > 0$. Also, $\sum_i \mu_i(\mathbf{x}) = 1$ by definition.

It does not satisfy the barycentric coordinate property of Definition 3.1. If the tangent spaces of all sample points \mathbf{p}_i with $\mu_i(\mathbf{x}) \neq 0$ coincide, i.e. the manifold is locally flat, the barycenter $\mathbf{b}(\mathbf{x}) = \sum_i \mu_i(\mathbf{x})\mathbf{p}_i$ is the projection of \mathbf{x} onto the (flat) manifold. Notice that \mathbf{x} is not necessarily on the manifold, so that even in this case $\mathbf{b}(\mathbf{x}) \neq \mathbf{x}$ is possible.

The distance between \mathbf{x} and the barycenter $\mathbf{b}(\mathbf{x})$ is bounded, nevertheless, by the simple fact that the distance of \mathbf{x} and its Atlas-neighbors is bounded. $\mathbf{b}(\mathbf{x})$ must lie in the convex hull of these neighbors.

Corollary 7.17 *The distance between \mathbf{x} and the barycenter $\mathbf{b}(\mathbf{x})$ defined by the \mathcal{A} -coordinate system is*

$$\|\mathbf{x} - \mathbf{b}(\mathbf{x})\| \leq \frac{2\epsilon}{\sqrt{1-2\epsilon}-2\epsilon} \text{ lfs}(\mathbf{x}).$$

Proof: Corollary 6.6 bounds the distance between $\mathbf{p}_i \in \mathcal{M}$ and a point $\mathbf{x} \in \mathcal{M}$ such that \mathbf{x} is a \mathcal{T}_i -neighbor of \mathbf{p}_i . This yields the result because $\mathbf{b}(\mathbf{x})$ cannot be further from \mathbf{x} than the furthest neighbor \mathbf{p}_i . \square

The support of the coordinate function

The support Δ_i of μ_i is the union of \mathcal{A} -Delaunay balls that are centered on vertices of $\overline{V}(\mathbf{p}_i)$. Inside Δ_i , the cell $V^+(\mathbf{x})$ has a non-empty intersection with $\overline{V}(\mathbf{p}_i)$. The radius r of such an \mathcal{A} -Delaunay ball is bounded by the upper distance between \mathbf{p}_i and a vertex of $\overline{V}(\mathbf{p}_i)$,

$$r \leq \frac{\epsilon}{\sqrt{1-2\epsilon}} \text{ lfs}(\mathbf{p}_i).$$

Lemma 7.18 *Let \mathcal{P} be an ϵ -sample of \mathcal{M} with $\epsilon < 0.31$. The support Δ_i contains the Voronoi cell $V_{\mathcal{M}}(\mathbf{p}_i)$ of \mathbf{p}_i restricted to \mathcal{M} . Consequently, the support Δ of the \mathcal{A} -coordinate system contains \mathcal{M} .*

Proof: A point $\mathbf{q} \in V_{\mathcal{M}}(\mathbf{p}_i)$ lies inside Δ_i if a part of $\overline{V}(\mathbf{p}_i)$ would be “stolen” by \mathbf{q} at its insertion into $\text{Vor}(\mathcal{P})$. This means that there exists a point $\mathbf{y} \in \overline{V}(\mathbf{p}_i)$ with $\|\mathbf{q} - \mathbf{y}\| \leq \|\mathbf{p}_i - \mathbf{y}\|$. We say that, in this case, \mathbf{q} is *in conflict* with \mathbf{y} . Assume, at first, that the orthogonal projection $\overline{\mathbf{q}}$ of \mathbf{q} onto \mathcal{T}_i lies in $\overline{V}(\mathbf{p}_i)$. We show that, in this case, \mathbf{q} is in conflict with $\overline{\mathbf{q}}$.

Let $\theta = \angle \mathbf{q}\mathbf{p}_i\overline{\mathbf{q}}$ and $\|\mathbf{p}_i - \mathbf{q}\| = \rho \text{ lfs}(\mathbf{p}_i)$. (Because $\mathbf{q} \in V_{\mathcal{M}}(\mathbf{p}_i)$, $\rho < \frac{\epsilon}{1-\epsilon}$ with Equation 5.1.) Thus,

$$\|\mathbf{p}_i - \overline{\mathbf{q}}\| = \|\mathbf{p}_i - \mathbf{q}\| \cos(\theta) \quad \text{and} \quad \|\mathbf{q} - \overline{\mathbf{q}}\| = \|\mathbf{p}_i - \mathbf{q}\| \sin(\theta).$$

(See Figure 7.8(a)) From Lemma 5.17, we know that $\theta \leq \arcsin(\frac{\rho}{2})$. It follows that $\|\mathbf{q} - \overline{\mathbf{q}}\| \leq \|\mathbf{p}_i - \overline{\mathbf{q}}\|$ if $\theta < \frac{\pi}{4} \iff \rho < \sqrt{2}$. With $\rho < \frac{\epsilon}{1-\epsilon}$, we deduce that \mathbf{q} is in conflict with $\overline{\mathbf{q}}$ if $\epsilon \leq \frac{1}{2}$.




Figure 7.8: $\mathbf{q} \in V_{\mathcal{M}}(\mathbf{p}_i)$ must be in conflict with the cell $\bar{V}(\mathbf{p}_i)$.

However, $\bar{\mathbf{q}}$ may not lie in $\bar{V}(\mathbf{p}_i)$. Let \mathbf{v} be the intersection point of the line $[\mathbf{p}_i, \bar{\mathbf{q}}]$ and a face $\bar{V}(\mathbf{p}_i, \mathbf{p}_j)$ of $\bar{V}(\mathbf{p}_i)$. We show that \mathbf{q} is in conflict with \mathbf{v} . From Lemma 5.17, we deduce that the angle α between any normal $\bar{\mathbf{n}}_i$ of \mathcal{M} at \mathbf{p}_i and the bisector hyperplane \mathcal{H}_{ij} that contains $\bar{V}(\mathbf{p}_i, \mathbf{p}_j)$ is small because $\|\mathbf{p}_i - \mathbf{p}_j\| \leq \frac{2\epsilon}{\sqrt{1-2\epsilon}} \text{lfs}(\mathbf{p}_i)$ by Lemma 6.3. In fact, $\alpha \leq \arcsin(\frac{\epsilon}{\sqrt{1-2\epsilon}})$. The point \mathbf{q} does not necessarily lie on \mathcal{H}_{ij} but it must lie inside the halfspace \mathcal{H}_{ij}^i induced by \mathcal{H}_{ij} that contains \mathbf{p}_i . Otherwise \mathbf{q} could not be in $V(\mathbf{p}_i)$ because it would be closer to \mathbf{p}_j than to \mathbf{p}_i .

In order to obtain $\|\mathbf{q} - \mathbf{v}\| \leq \|\mathbf{p}_i - \mathbf{v}\|$, the angle $\beta = \angle \mathbf{v}\mathbf{q}\mathbf{p}_i$ must be larger than θ . Refer to figure 7.8(b). We obtain

$$\begin{aligned} \beta \geq \frac{\pi}{2} - \theta - \alpha > \theta &\iff \frac{\pi}{2} > \alpha + 2\theta \iff \arcsin\left(\frac{\epsilon}{\sqrt{1-2\epsilon}}\right) + 2\arcsin\left(\frac{\epsilon}{2(1-\epsilon)}\right) \\ &\iff \epsilon < 0.31. \end{aligned}$$

□




Figure 7.9: The coordinate $\mu_i(\mathbf{x}) = 1$ because \mathbf{x} is only in conflict with $\mathbf{v} \in \bar{V}(\mathbf{p}_i)$ but with no other \mathcal{T} -restricted Voronoi cell.

Notice that the cell $\overline{V}(\mathbf{p}_i)$ can never be entirely contained in $V^+(\mathbf{x})$: As we noticed before, B_i must have a non-empty cell in the power diagram $\overline{\text{Vor}}^i(\mathcal{P} \cup \mathbf{x})$ because the weight of B_i is zero while the weight of all other sites, including X_i , is negative or zero. Also, it might occur that the coordinate function $\mu_i(\mathbf{x}) = 1$ while $\mathbf{x} \neq \mathbf{p}_i$. This is the case when \mathbf{x} lies in the support Δ_i but outside any support Δ_j . \mathbf{x} is in conflict with a vertex of $\overline{V}(\mathbf{p}_i)$ but not with the corresponding vertices in other tangent spaces (if they exist). Refer also to figure 7.9.

7.3.2 Continuity properties and gradient of the coordinate function μ_i

Lemma 7.19 *The \mathcal{A} -coordinate $\mu_i, i = 1, \dots, n$, associated with \mathcal{P} is continuous everywhere in Δ .*

Proof: The volume functions $\pi_i(X_i)$ are continuous everywhere in \mathcal{T}_i . The tangent space \mathcal{T}_i is fixed and $X_i = (\overline{\mathbf{x}}^i, \|\overline{\mathbf{x}}^i - \mathbf{x}\|)$ is a smooth function of \mathbf{x} . The ratio $\mu_i(\mathbf{x}) := \frac{\pi_i(X_i)}{\sum_i \pi_i(X_i)}$ is continuous if X_i and π_i is continuous, for all i . \square

Lemma 7.20 *The \mathcal{A} -coordinate $\mu_i, i = 1, \dots, n$, associated with \mathcal{P} is continuously differentiable everywhere in Δ except at \mathcal{P} and at the points $\mathbf{x} \in \Delta$ such that the intersection of $\mathcal{T}_i, i = 1, \dots, n$, with a $(d - 2)$ -dimensional Voronoi face of the Voronoi subcell $V_i(\mathbf{x})$ has dimension $k - 1$.*

Proof: The same arguments as for the proof of Lemma 6.8 apply. The coordinate function μ_i is differentiable iff the volume functions $\pi_i(X_i), i = 1, \dots, n$, are differentiable. They are defined in (different) power diagrams $\overline{\text{Vor}}^i(\mathcal{P})$. Lemma 3.6 about the differentiability of the volume function $\pi_i(X)$ shows that $\pi_i(X_i)$ is differentiable unless the bisector hyperplane of \mathbf{x} and \mathbf{p}_i contains a $(k - 1)$ -dimensional face of the power diagram dual to \mathbf{p}_i and another point \mathbf{p}_j . As we noticed in the proof of Lemma 6.8, this means that \mathbf{x}, \mathbf{p}_i and \mathbf{p}_j define a $(d - 2)$ -dimensional Voronoi face of $V^+(\mathbf{x})$ that intersects \mathcal{T}_i in a $(k - 1)$ -dimensional face. \square

The gradient of the coordinate function

We determine the gradient $\nabla \pi_i(X_i)$ as a function of \mathbf{x} . With $\mu_i(\mathbf{x}) = \frac{\pi_i(X_i)}{\sum_i \pi_i(X_i)}$, this yields the gradient $\nabla \mu_i$.

The volume function $\pi_i(X_i)$ is a composite function because X_i is a function of \mathbf{x} . All other elements defining $\pi_i(X_i)$ are constant, namely, the tangent space \mathcal{T}_i and the weighted points $B_j = (\overline{\mathbf{p}}_j^i, -\|\mathbf{p}_j - \overline{\mathbf{p}}_j^i\|)$ that define the power diagram.

$X_i = (\mathbf{f}_i(\mathbf{x}), w_i(\mathbf{x}))$ is defined by

$$\mathbf{f}_i(\mathbf{x}) = \overline{\mathbf{x}}^i = \mathbf{x} - \sum_{l=1}^{d-k} (\overline{\mathbf{p}}_i \overline{\mathbf{x}} \cdot \overline{\mathbf{n}}_i^l) \overline{\mathbf{n}}_i^l \quad (7.1)$$

and

$$w_i(\mathbf{x}) = -\|\bar{\mathbf{x}}^i - \mathbf{x}\|^2 = -\left\| \sum_{l=1}^{d-k} (\overrightarrow{\mathbf{p}_i \bar{\mathbf{x}}^i} \cdot \overrightarrow{\mathbf{n}_i^l}) \overrightarrow{\mathbf{n}_i^l} \right\|^2. \quad (7.2)$$

Since an orthogonal projection is always a linear transformation, the Jacobian matrix $\nabla \mathbf{f}_i$ and the gradient ∇w_i are constant. Denote by $\nabla X_i = (\nabla \mathbf{f}_i, \nabla w_i)$ the Jacobian of X_i .

The gradient of $\pi_i(X_i)$ is determined by Lemma 3.8. We need the following definitions:

1. \bar{v}_{ix} is the $(k-1)$ -dimensional volume of the face between X_i and B_i in $\bar{V}_i(\mathbf{x})$ if $k > 1$, and $\bar{v}_{ix}^i = 1$, if $k = 1$,
2. $\bar{\mathbf{c}}_{ix}$ is the centroid of this facet,

$$\nabla \pi_i(X_i) = \begin{cases} \frac{\bar{v}_{ix}}{\|\mathbf{p}_i - \bar{\mathbf{x}}^i\|} \begin{pmatrix} \bar{\mathbf{c}}_{ix} - \bar{\mathbf{x}}^i \\ \frac{1}{2} \end{pmatrix} & \text{if } \mathbf{x} \in \Delta_i, \\ \mathbf{0} & \text{otherwise,} \end{cases}$$

The composition of X_i and π_i is summarized as follows:

$$\begin{array}{ccc} \mathbb{R}^n & \xrightarrow{X_i} & \mathbb{R}^n \times \mathbb{R} & \xrightarrow{\pi_i} & \mathbb{R} \\ \mathbf{x} & \xrightarrow{X_i} & X_i = (\mathbf{f}_i(\mathbf{x}), w_i(\mathbf{x})) & \xrightarrow{\pi_i} & \pi_i(X_i). \end{array}$$

Notice that the gradient $\nabla \mu_i(\mathbf{x}) \neq \mathbf{0}$ only if \mathbf{x} lies in the support of at least one other coordinate function, $\mathbf{x} \in \bigcup_{j \neq i} (\Delta_i \cap \Delta_j)$. Otherwise, $\mu_i(\mathbf{x})$ is constant, $\mu_i(\mathbf{x}) = 1$.

By the chain rule, we obtain

$$\nabla \pi_i(X_i(\mathbf{x})) = \begin{cases} \nabla X_i(\mathbf{x}) \cdot \nabla \pi_i(X_i) & \text{if } \mathbf{x} \in \bigcup_{j \neq i} (\Delta_i \cap \Delta_j), \\ \mathbf{0} & \text{otherwise,} \end{cases}$$

Chapter 8

Applications

In this chapter, we present three different applications of the concepts defined in this thesis. In section 8.1, the \mathcal{T}_x -coordinates is applied to interpolation on surfaces. The following section deals with reconstruction of surfaces and curves from point clouds. Finally, in Section 8.3, we present a re-meshing technique that is based on centroidal Voronoi diagrams.

8.1 Interpolating a function on a surface S

In this section, the \mathcal{T}_x -coordinate system defined in Chapter 6 is used for scattered data interpolation on smooth surfaces. This means that we interpolate a function that is defined on the surface and that is known on discrete sample points. We will apply four interpolation methods described in Chapter 4 by replacing the natural neighbor coordinates by the \mathcal{T}_x -coordinates defined on the surface. The reader may recall that these methods depend only on the barycentric coordinate property (Definition 3.1 (i)) and on the continuity and differential properties of the coordinate function. Even though the barycentric coordinate property is not satisfied by the \mathcal{T}_x -coordinate, the interpolation methods can be applied. We restrict the experiments presented in this section to surfaces, i.e. 2-manifolds in \mathbb{R}^3 .

More precisely, the problem is stated as follows. We want to approximate $\Phi : S \rightarrow \mathbb{R}^l$ where $S \subset \mathbb{R}^3$ corresponds to a manifold \mathcal{M} defined as before and $l = \{1, 2, \dots\}$. We assume that an ϵ -sample \mathcal{P} of S with function values $\{(\mathbf{p}_i, z_i) : \mathbf{p}_i \in \mathcal{P}, z_i = \Phi(\mathbf{p}_i)\}$ is given. The interpolation is carried out for a point $\mathbf{x} \in S$ (and $\vec{\mathbf{n}}_x$ is its normal to S).

The section gives a short overview about previous work on this problem, then, it defines the different interpolation functions that are proposed and, finally, it presents some experimental results.

8.1.1 Introduction and related work

This problem, which is also called the 'surface on surface' problem or 'scattered data fitting on surfaces', arises in a variety of settings. In geodesy, geophysics, and meteorology, S is some model of the earth,

and the function to interpolate from a number of discrete measurements represents temperature, rainfall, pressure, etc. In other contexts, \mathcal{S} might be a complicated surface, e.g., the surface of some mechanical piece in *CAD*, a molecular surface or the wing of an airplane [17]. The interpolation of vector fields defined on surfaces has applications in fluid dynamics and in so-called moving boundary problems from partial differential equations. For example, the surface \mathcal{S} might represent the interface between solid and liquid phases of a material substance, e.g. ice and water, and it varies in time. The interpolated data consists of the normal vectors to \mathcal{S} , the mean curvature, or the velocity of the moving boundary. See [82] for an introduction to the “surface on surface” interpolation problem and notably to splines on surfaces.

Several methods exist to solve this problem. One of the most popular is to enlarge the definition of splines to treat the case of a non-planar parameter domain. This was first done for the spherical case in [3]. With this achievement, it suffices to partition a general surface \mathcal{S} into a collection of non-overlapping surface patches, for instance geodesic triangles, and to define a globally smooth interpolation function as piecewise polynomials on the patches that are carefully accorded at the boundaries. However, as Neamtu concludes in [82], the definition of spline spaces on general surfaces depends upon the existence of barycentric coordinates. For some special surfaces, smooth coordinate functions with good properties exists, but, to my knowledge, there exist no general way to define them.

Implicit Bernstein-Bézier patches are used by Bajaj et al. in [15] to reconstruct the surface \mathcal{S} from a sample set but also to approximate a scalar function defined at the sample points. A user-defined parameter bounds the approximation tolerance, i.e. the maximum distance between the approximated and the input data on the sample points. The surface is represented as the zero-contour of an implicit function which approximates the signed distance to \mathcal{S} . In a first step, the surface is approximated by so-called α -solids which are subsets of the Delaunay triangulation of \mathcal{P} . The distance function (defined by the distance to the α -solid) is approximated by piecewise polynomial functions obtained from least-square fitting with barycentric Bernstein-Bézier patches and incrementally refined until the approximation tolerance is respected. In the same way, a piecewise polynomial is constructed by least squares fitting that approximates the scalar function defined on the sample points. At last, smoothing is necessary to obtain globally C^1 smooth surfaces.

Radial basis functions (RBF) is another popular method used for interpolation. Lodha and Franke [77] give a recent introduction to the field (for interpolation in \mathbb{R}^2 and \mathbb{R}^3). For a long time, the practical use of RBF was prohibited by the large computation cost. According to the authors, new progress has made the RBF approach feasible even for very large data sets. In contrast to spline based methods, they have the advantage to be *meshless*, i.e. independent of a partition of the definition space. The main ingredient for RBFs is the radial function which is a function that is symmetric around a center point as it depends only on the distance to the center. Its support might be unbounded or compact. The solution to the interpolation problem is a linear combination of translates of radial functions plus, in some cases, a polynomial term. It is subject to a system of linear equations that impose the interpolation of the function data at the data

points. Only recently, methods have been developed that allow to compute the solution to this –often poor conditioned – linear system efficiently. For interpolation of functions on surfaces, *Hardy multiquadrics* were the first type of RBFs which have been used with applications in geophysics [69]. Applications of RBFs for scattered data interpolation on the sphere are summarized by Fasshauer and Schumaker in [58].

Other interpolation methods are variational methods, distance weighted methods, transfinite methods or multi-resolution methods. Most methods are defined for the spherical case. See [58] for a survey of the principal methods for scattered data fitting on the sphere. To our knowledge, there exist only few methods for interpolation on general surfaces. For instance, multiresolution techniques are widely used for constructing surfaces. Apart from a work by Sweldens [93] where wavelets are defined for arbitrary surfaces, there seems to be few activity wrt. the surface-on-surface problem.

Pottman [85] defines a transfinite method for arbitrary at least twice-differentiable surfaces. Transfinite methods are based on a partition of the function domain and the definition of smooth functions on the edges of the partition which are themselves interpolated over each cell.

In this context, our method has several advantages: It is meshless, i.e. it applies directly to the point samples without need of a prior subdivision or triangulation of the surface. The quality of the result depends on the curvature of the surface and on the sampling density. If the surface is locally planar, the interpolant has linear precision. If, additionally, the gradient $\nabla(\Phi(\mathbf{p}_i))$ of Φ at the sample points $\mathbf{p}_i, i = 1, \dots, n$, is known, we define an interpolant that reproduces exactly a quadratic function – again if the surface is locally planar. Our method generalizes easily to interpolation of vector-valued functions, i.e. functions $\Phi : \mathcal{S} \rightarrow \mathbb{R}^l, l = 2, 3 \dots$. Each coordinate of the result vector is interpolated independently as if it were a scalar function. The only draw-back lies in the limited differentiability on the data points themselves. With Sibson’s and Farin’s methods (see Sections 4.2 and 4.1), the interpolant is globally C^1 smooth but it seems difficult to increase the smoothness on the data points. The interpolants are tested in different settings. Examples are shown in Section 8.1.3.

8.1.2 Definition and properties of interpolants

In the sequel, the interpolation methods described in Chapter 4 are applied to interpolation on surfaces. The natural neighbor coordinate system defined in the entire Euclidean d -space is simply replaced by the \mathcal{T}_x -coordinate system whose definition domain is limited to the surface. For a point $\mathbf{x} \in \mathcal{S}$, we compute its \mathcal{T}_x -coordinates and we replace λ_i by the \mathcal{T}_x -coordinate function τ_i in the interpolation scheme. The result of the interpolation is the same whenever all \mathcal{T}_x -neighbors of \mathbf{x} lie in the tangent plane \mathcal{T}_x . If this is not the case, the \mathcal{T}_x -coordinates may not satisfy the barycentric coordinate property which causes an error in the interpolation scheme. However, this error can be bounded with respect to the local feature size $\text{lfs}(\mathbf{x})$ (we refer to Section 6.3.3). We tested the interpolation of scalar functions using Sibson’s interpolants (Z^0 and Z^1), Farin’s interpolant and the quadratic precision interpolant (I^0). We did not compute Clarkson’s coordinate function but we tested Sibson’s Z^0 interpolant with the non-convex coordinate system described

in Section 6.5. The experiments were run on a Bi-processor Pentium III computer, 1 GHz, 1024 MB RAM

8.1.3 Experimental results

Functions in 3D

We assume that we are given a real-valued function $f(\mathbf{x})$ that can be evaluated at all points $\mathbf{x} \in S$. For example, we interpolated the following three functions on the sphere:


- $f_1(x, y, z) = 4x^2 + y^4 + 6z^4$,
- $f_2(x, y, z) = e^{-2((x-1)^2+y^2+z^2)} + 0.5 e^{-4(x^2+(y-0.7)^2+(z-0.7)^2)} - 0.25 e^{-4(x^2+(y+0.7)^2+(z+0.7)^2)}$,
- $f_3(x, y, z) = 1 + x^8 + e^{2y^3} + e^{2z^2} + 10xyz$.

The functions f_2 and f_3 are test functions from [59] and [2].


To visualize the result, we deformed the sphere at each point along its normal by the amount of the corresponding interpolation result. Figures 8.2, 8.3, and 8.4 depict the result of the interpolation of f_1 , f_2 , and f_3 resp. knowing 50 and 250 pseudo-random sample points. In Figures 8.5, 8.6, and 8.7, the error statistics are depicted together with the exact result of the function application where the function is evaluated at the 6000 points of the sphere model and the sphere is deformed correspondingly. For each interpolant, mean and maximum errors are given with respect to the absolute difference between the actual and the interpolated function value on the 6000 evaluation points. The linear interpolant Z^0 with convex coordinates is the least accurate concerning mean as well as maximum error followed by the same interpolant using the non-convex coordinate system and the I^1 interpolant. Farin's interpolant and Sibson's Z^1 interpolant achieve comparable results, yet, the Z^1 interpolant yields to be slightly better than Farin's interpolant if few sample points are known. The running time for computing the interpolants on the 6000 points on the sphere is depicted in Table 8.1.

$ \mathcal{P} $	100	200	500	1000	2000
convex coordinates	15.14 s	16.84 s	18.25 s	18.48 s	19.19 s
non-convex coordinates	62.14 s	67.1 s	80.61 s	82.64 s	99.29 s

Figure 8.1: Running time (in seconds) for interpolation at 6000 points on the sphere.

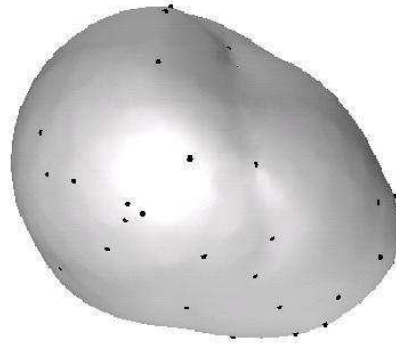


(a)




(b)

Figure 8.2: Interpolation of f_1 with (a) 50 and (b) 250 sample points.




(a)

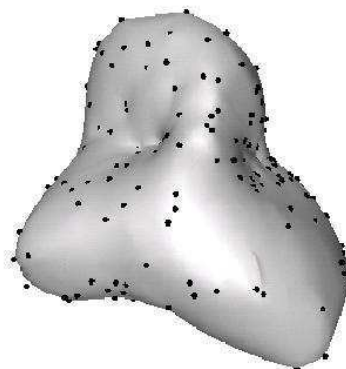


(b)

Figure 8.3: Interpolation of f_2 with (a) 50, (b) 250 sample points.



(a)



(b)

Figure 8.4: Interpolation of f_3 with (a) 50, (b) 250 sample points.





Figure 8.5: Error statistics for the interpolation of f_1 : (top) exact model and mean error statistic, (bottom) maximum error statistic. Z^0 interpolant: dotted curve with points, Z^0 interpolant with non-convex coordinates: dotted curve with circles, I^1 interpolant: dark solid curve with diamonds, Farin's interpolant: grey solid curve with crosses, Z^1 interpolant: slashed curve with boxes.





Figure 8.6: Error statistics for the interpolation of f_2 : (top) exact model and mean error statistic, (bottom) maximum error statistic. Interpretation of curves cf. Figure 8.5.





Figure 8.7: Error statistics for the interpolation of f_3 : (top) exact model and mean error statistic, (bottom) maximum error statistic. Interpretation of curves cf. Figure 8.5.

Functions on parameterized surfaces

In this section, we suppose that we are given a parameterized surface \mathcal{S} defined by

$$\mathbf{s} : \mathbb{R}^2 \rightarrow \mathbb{R}^3, \text{ such that } \mathbf{s}(u, v) = (x(u, v), y(u, v), z(u, v)) \in \mathcal{S}$$

and a function $h(u, v) : \mathcal{S} \rightarrow \mathbb{R}$ defined on \mathcal{S} . We evaluate \mathcal{S} on a regular grid of 40000 points (which is for numerical reasons slightly perturbed). The function h is evaluated on a random set of points uniformly distributed in parametric space. The gradient needs to be expressed with respect to the Cartesian coordinates

$$\mathbf{g}_i = \begin{pmatrix} \frac{\partial}{\partial x} \\ \frac{\partial}{\partial y} \\ \frac{\partial}{\partial z} \end{pmatrix} h \circ \mathbf{s}^{-1}(\mathbf{P}_i) = \begin{pmatrix} \frac{\partial u}{\partial x} \frac{\partial h}{\partial u}(u_i, v_i) + \frac{\partial v}{\partial x} \frac{\partial h}{\partial v}(u_i, v_i) \\ \frac{\partial u}{\partial y} \frac{\partial h}{\partial u}(u_i, v_i) + \frac{\partial v}{\partial y} \frac{\partial h}{\partial v}(u_i, v_i) \\ \frac{\partial u}{\partial z} \frac{\partial h}{\partial u}(u_i, v_i) + \frac{\partial v}{\partial z} \frac{\partial h}{\partial v}(u_i, v_i) \end{pmatrix}$$

We use Sibson's Z^0 and Z^1 interpolants, the quadratic precision interpolant I^1 and Farin's interpolant f^3 to estimate the function value on the grid points. The surface is deformed by the value of the interpolant in direction of the surface normal. An exact model is obtained by evaluating the function on the grid points and by deforming the surface accordingly. Figure 8.11 depicts an example of Farin's interpolant on the torus. The corresponding error statistics as well as the exact model are shown in Figures 8.10 and 8.9. The error statistics show that Sibson's Z^1 and Farin's f^3 interpolants yield the best results. Farin's f^3 interpolant has a smaller maximum error and a smaller mean error if the sampling of the function is dense. For sparser samples, Sibson's Z^1 interpolant is better with respect to the average error. The linear Z^0 interpolant with convex coordinates produces the largest errors, followed by Z^0 interpolant with non-convex coordinates and the I^1 interpolant. Notice that the non-convex coordinate system cannot be applied on the torus (since it is not a convex surface).

$ \mathcal{P} $	100	200	500	750	1000	2000
convex coordinates	155.44 s	179.26 s	219.71 s	241.47 s	260.13 s	302.34 s

Figure 8.8: Running time (in seconds) for interpolation at 40.000 points.




Figure 8.9: Exact model and mean error statistic for the interpolation of $h(u, v) = 0.6 \cos(6v)$ on the torus. Z^0 interpolant: dotted curve with points, I^1 interpolant: dark solid curve with diamonds, Farin's interpolant: grey solid curve with crosses, Z^1 interpolant: slashed curve with boxes.




Figure 8.10: Maximum error statistic for the interpolation of $h(u, v) = 0.6 \cos(6v)$ on the torus. Interpretation of curves cf. Figure 8.9.




Figure 8.11: Interpolation of $h(u, v) = 0.6 \cos(6v)$ on the torus with 100, 200, and 400 random sample points (from left to right). Top row: sample points. Bottom row: result of interpolation.

Vector fields on parameterized surfaces

The interpolation of vector fields, i.e. functions that are defined from the surface to a higher dimensional space, can be treated in the same way as scalar functions by interpolating each coordinate of the result separately with one of the interpolants presented in this chapter. For example, let $v : \mathcal{S} \rightarrow \mathbb{R}^3$ with $v(u, v) = (v_x(u, v), v_y(u, v), v_z(u, v))$, then, v_x , v_y , and v_z are interpolated independently. The error of the interpolation is measured by the squared distance between the vector obtained by applying the function on a point \mathbf{x} and the interpolation result at \mathbf{x} . Figure 8.12 shows the interpolation of $v(u, v)$ with $v_x(u, v) = -\cos(u) \cos(v)$, $v_y(u, v) = -\cos(u) \sin(v)$, and $v_z(u, v) = -\sin(u)$ on the cylinder. The function as well as the function gradient is known at a small number of random points, and the interpolants are evaluated on 40,000 grid points. For surfaces with boundary, all boundary grid points are part of the sample points. To visualize, we translate each point by the vector resulting from the interpolation or the function application. We tested Sibson's Z^0 and Z^1 interpolants and the I^1 interpolant. Concerning the error statistics, Figure 8.13, this time the mean error of the Z^0 interpolant is worse than the maximum error of I^1 and Z^1 . The Z^1 interpolant produces less error than the I^1 interpolant. The third and the fourth highest curves correspond to the maximum error in the I^1 and the Z^1 interpolant (in this order).

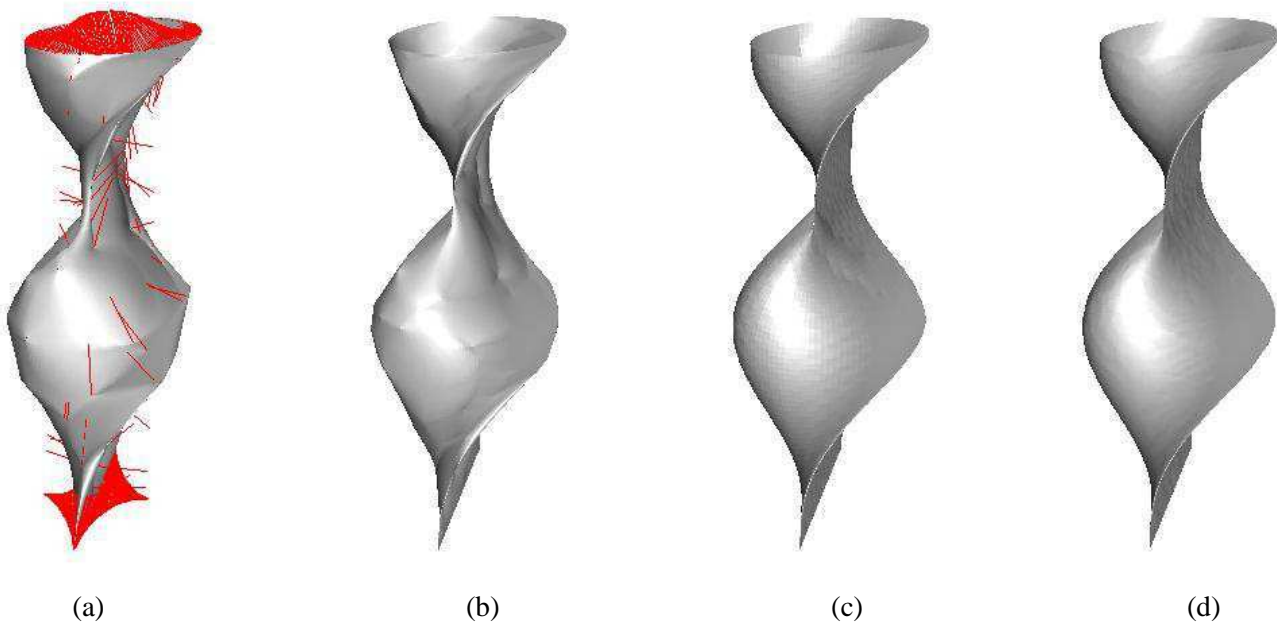


Figure 8.12: Interpolation of $v(u, v)$ on the cylinder with (a) 100 regularly spaced sample points (shown in red), (b) 500, (c) 1000 and (d) 4000 random sample points.




Figure 8.13: Exact model and error statistic for Figure 8.12. Interpretation of curves (in descending order). Z^0 interpolant maximum error: dotted curve with crosses, Z^0 mean error: dotted curve with points, I^1 interpolant maximum error: dark solid curve with diamonds, Z^1 interpolant maximum error: slashed curve with circles, I^1 mean error: dark solid curve with crosses, Z^1 mean error: slashed curve with boxes.

8.2 Delaunay-based surface and curve reconstruction

In this section, we describe how the concept of \mathcal{T} -neighbors can be applied to curve and surface reconstruction in \mathbb{R}^3 . The problem of reconstructing a surface or a curve from a set of scattered points has received a lot of attention in the last twenty years. One class of algorithms to solve this problem is based on the Delaunay triangulation of the input points. It has been shown that – under some sampling condition – the surface and the restriction of the Delaunay triangulation to this surface are homeomorphic (Lemma 5.12). Therefore, the approach chosen in this class of algorithms is to extract the restricted Delaunay triangulation (or a plausible estimation) from the Delaunay triangulation. In the sequel, we precisely define the problem and we overview several curve and surface reconstruction methods which are based on the Delaunay triangulation, and we situate our approach with respect to the existing methods. Then, we introduce the \mathcal{T} -neighbor based approach and discuss the relation to existing methods.

Problem statement

The reconstruction problem for a surface and resp. a curve is stated as follows:

Given a set of points \mathcal{P} sampled from a closed surface S (resp. from a closed curve \mathcal{C}) embedded in \mathbb{R}^3 as well as the corresponding normal vectors to S (resp. \mathcal{C}), construct a polygonal surface (resp. curve) \mathcal{R} such that the points of \mathcal{P} lie on \mathcal{R} and such that \mathcal{R} approximates S (resp. \mathcal{C}).

We suppose that it is known whether the points are issued from a curve or a surface. If this is not the case, the dimension of the manifold can be determined following the work of Dey et al. [46]. Similarly, if the normals are not part of the input, they can be estimated by one of the methods described in Chapter 5. Furthermore, we suppose that the sample set \mathcal{P} is an ϵ -sample of the manifold, i. e. it satisfies the ϵ -sampling condition for some $\epsilon < \frac{1}{2}$.

An algorithm for the reconstruction problem is *topologically correct* if the polygonal approximation \mathcal{R} is homeomorphic to the underlying manifold S (resp. \mathcal{C}). In addition, one would like that the Hausdorff distance between \mathcal{R} and S (resp. \mathcal{C}) is small.

8.2.1 Previous work

Even though there exist significant contributions to this problem that are not based on the Delaunay triangulation, we restrict our attention to Delaunay based methods in order to compare and classify our algorithm with respect to them.

Sculpture, wrapping, flow and convection

The first method for three-dimensional surface reconstruction based on the Delaunay triangulation is defined by Boissonnat [20] and is called *SCULPTURE*. Starting from the convex hull of the sample points (the exterior facets of the Delaunay triangulation), tetrahedra are removed one-by-one from the Delaunay triangulation following some priority criteria (based on local geometric properties) and ensuring topological correctness. During this process, the boundary of the remaining tetrahedra contains more and more sample points, and the algorithm stops when all sample points lie on it. Due to the topological rules, *SCULPTURE* works only for surfaces with genus zero. A major drawback is the sensibility to the order of removal: The algorithm might in some cases be blocked before finishing because the topological rules do not allow any further removal of simplices.

This “sculpturing” process can also be interpreted as the evolution of a surface to best fit the sample points. Evolution (or flow) based approaches have been recently published in [51], [62] and [31]. The new approaches are different in the sense that they compute a global ordering of the Delaunay simplices rather than depending on local criteria. The acyclic relation between the simplices is motivated from Morse theory in [51], from flow complexes in [62] or from a physical convection model in [31]. The resulting surface is some stable configuration of the evolution process. The result of these algorithms is manifold by construction.

Surface-based or advancing front methods

A second method by Boissonnat [20] constructs the surface directly from the point cloud. It is called *surface-based* approach in opposition to the volume-based nature of the *SCULPTURE* algorithm. Another common classification is *advancing front* method.

Starting from the shortest edge, a contour is propagated by adding triangles one by one. Given an edge of the contour, the sample points in the neighborhood are projected onto the (estimated) tangent plane of one of the endpoints. The next triangle is built from the contour edge and the point that maximizes the angle when its projected counterpart is connected to the edge. The process is repeated until the surface is closed, and all sample points are included. A priori, this is not a Delaunay based method, however, the author remarks that the candidate triangles may be chosen from the Delaunay triangulation [20].

A new method in which the surface is propagated or spread over the point cloud has been defined recently by Cohen-Steiner and Da [36]. The criterion to extend an edge is, this time, based on distance and angle measures. The surface is guaranteed to be manifold by the topological rules imposed during the propagation.

All algorithms presented so far ensure by construction that the reconstructed surface is manifold. There exist another category of algorithms that collect a subset of faces of the Delaunay triangulation without considering global topological criteria. For most curve reconstruction algorithms, the right topology is

guaranteed, nevertheless, by the sampling condition (refer to Section 5.3). In surface reconstruction, it can often be established in a post-treatment. Most algorithms presented in the sequel are of this type.

Alpha shapes

Alpha shapes by Edelsbrunner et al. were first defined for the two-dimensional curve reconstruction problem [52] and then enlarged to three-dimensions [53]. From the Delaunay triangulation of the input points, the simplices are removed whose radius of the smallest circumsphere is smaller than the parameter α . For the two dimensional case, it has been proven in [18] that there exist values of α for which the alpha shapes reconstructs the curve, provided that the sampling density is uniform. Several algorithms are based on α -shapes, notably, Teichmann et al. introduce anisotropy in order to improve the results, the Ball Pivoting Algorithm by Bernadini et al. [16] is very close to α -shape as it removes triangles with circumradius less than the radius of the pivoting ball.

Crust

The *CRUST* algorithm by Amenta et al. [10] is the first algorithm with theoretical guarantees for the three-dimensional surface reconstruction problem. The authors introduce a sampling condition based on the *local feature size*, i.e. the smallest distance of a surface point to the medial axis. In Definition 5.10, we call it ϵ -sampling condition.

First developed for the two-dimensional case only, the *CRUST* is based on the fact that the Voronoi vertices of the $2D$ Voronoi diagram of the sample points approximate the medial axis of the curve. Therefore, all edges of the dual Delaunay triangulation that traverse the medial axis are destroyed if the Voronoi vertices are inserted into the Delaunay triangulation of the sample points. Those Delaunay edges that connect two sample points in the enlarged Delaunay triangulation (sample points plus the Voronoi vertices) do not traverse the medial axis and are likely to belong to the Delaunay triangulation restricted to the curve. The collection of such edges is called the *CRUST*. For $\epsilon < 0.26$, the *CRUST* connects exactly all neighboring curve points and no others.

This idea does not extend directly to surface reconstruction in \mathbb{R}^3 , since in the three-dimensional case, Voronoi vertices might be far from the medial axis of the surface due to so-called sliver tetrahedra (refer to Chapter 5). We re-call that a sliver is generated by four roughly co-circular points so that its circumcenter may be arbitrary close to the surface independently from the sampling density.

It was the notion of poles that allowed to generalize the *CRUST*: the pole of a sample point (which is its furthest Voronoi vertex) is close to the medial axis, and the pole vector that is the vector from the sample point to the pole, is a good approximation of the normal to the surface at the sample point [9]. The second pole is the furthest Voronoi vertex such that the angle between the two pole vectors (the vectors from the sample point to the first and to the second pole) is bigger than $\frac{\pi}{2}$.

The 3D *CRUST* builds the Delaunay triangulation of the sample points plus the first and the second poles and extracts the triangles in which all three vertices are sample points. A second filtering step removes the triangles that exceed a bound on the angle between the triangle's normal and the vertices' pole vectors. Thirdly, poles and triangles are oriented, and the manifold is extracted from the collection of triangles. If the sampling is sufficiently dense (an ϵ -sample with $\epsilon < 0.1$), it is proven that the restricted Delaunay triangulation $\text{Del}_S(\mathcal{P})$ is homeomorphic to the surface (Lemma 5.12) and that the *CRUST* includes all triangles of the restricted Delaunay triangulation. However, it is not shown that the result of the manifold extraction step is homeomorphic to the surface.

Cocone

The *COCONE*-algorithm [7] presents a different Voronoi filtering criterion that has the same goal: to extract the Delaunay triangulation restricted $\text{Del}_S(\mathcal{P})$ to the surface. The result of the algorithm is (provably) homeomorphic to the surface given an ϵ -sample with $\epsilon < 0.006$. We present this algorithm in detail, since it is very close to the T -neighbor based approach that we will describe in the sequel of this chapter.

The global approach consists of computing a collection of candidate triangles and to extract a manifold from the triangle set. A candidate triangle is chosen from the Delaunay triangulation if its dual Voronoi edge intersects the *co-cones* of its three vertices. The co-cone of a sample point \mathbf{p}_i is the set of points \mathbf{x} that are close to the tangent plane at \mathbf{p}_i , in the sense that the vector $\overrightarrow{\mathbf{x}\mathbf{p}_i}$ forms an angle close to $\frac{\pi}{2}$ with the normal at \mathbf{p}_i (which can be estimated by the pole vector).

The collection of candidate triangles T selected by the *COCONE* algorithm meets the following three criteria [7]:

- I. T includes all triangles of the restricted Delaunay triangulation $\text{Del}_S(\mathcal{P})$,
- II. each triangle in T is small (the radius of its circumcircle is much smaller than the distance to the medial axis), and
- III. for each triangle, the angles between the triangle's normals and the surface normals at the vertices are small.

Under this conditions, a piecewise-linear manifold can be extracted from T that is homeomorphic to the surface. The manifold extraction algorithm is based on two principles: first, every two adjacent triangles should form an angle greater than $\frac{\pi}{2}$ at their common vertex, and, second, every sample point should be adjacent to at least one triangle. To meet the first property, so-called *sharp edges* are removed from the candidate set. An edge is sharp if the angle between two incident triangles is greater than $\frac{3\pi}{2}$ or if the edge has only one incident triangle. The authors show that no triangle of the restricted Delaunay triangulation $\text{Del}_S(\mathcal{P})$ is sharp if $\epsilon \leq 0.006$. After the removal of sharp edges, a manifold is constructed by a depth-first search along the outer boundary of each of its components. If the sampling condition is met, such a manifold must exist, and the success of the algorithm is proven.

This first version of the *COCONE* algorithm is very sensitive to the sampling condition, which is, in practice, met by very few point sets. Several papers revisit the *COCONE* algorithm to make it valuable in practice. We call them the *COCONE*-based algorithms since the basic idea stays the same.

Most of these ideas apply directly to the \mathcal{T} -neighbor based reconstruction method that we describe in the sequel of this section. Both methods are very close. Instead of restricting the Voronoi cells to the co-cone of the sample point, in \mathcal{T} -neighbor based reconstruction, they are restricted to the tangent space (which is a subspace of the co-cone).

The *COCONE* has been adopted to work for large sets of data in [95]. It has been used to detect undersampled zones and boundaries in [45] and [41]. The algorithm to detect oversampled zones as well as undersampled zones ([40]) served to describe a sample decimation algorithm in [94] and a decimation and re-sampling algorithm that deletes sample points in [96]. Robustness issues of the implementation are discussed in [97]. An algorithm using *COCONE* to find out the dimension of a manifold from the set of sample points is defined in [46].

Dey, Funke and Ramos [44] elaborated an algorithm that computes the *COCONE* triangles without constructing the entire Delaunay triangulation of the sample point and that is therefore more efficient. The authors suppose that the sampling is locally uniform. This implies that any sphere centered on a sample point and including all its neighbors in the restricted Delaunay triangulation $\text{Del}_{\mathcal{S}}(\mathcal{P})$ contains a constant number of points of \mathcal{P} . This assumption makes it efficient to collect all co-cone neighbors of a point from a data-structure for k -nearest neighbors search.

Notice that the normals to the surface at the sample points cannot be estimated from the poles without computing the three-dimensional Delaunay triangulation. The authors suggest to estimate the normals from the triangle formed by a sample point and its two nearest neighbors which form a sufficiently big angle. Refer to Lemma 5.18. In [60], Funke and Ramos describe a method to decimate sample points in order to meet the local uniformity. However, the implementation of this algorithm seems quite involved and to our knowledge, there exist no implementation so far which could prove the practical impact of the method.

We describe in more detail the ideas behind the boundary and dimension detection from [41] and [46]. In a second time, we give a sketch of the manifold extraction methods used for *COCONE*.

The special shape of the Voronoi cells, namely their long and skinny shape in the normal direction while being of small diameter in the tangent space, allows to detect boundaries as well as the dimension of the underlying manifold. With the help of a lemma equivalent to Lemma 6.3, the diameter of a Voronoi cell restricted to the *co-cone* of its generator sample point is upper bounded with respect to the local feature size. Also, the distance of the sample point to its pole is a lower bound for the local feature size. If the ratio of the diameter of the restricted Voronoi cell and the distance to the first pole as well as the second pole is greater than a certain constant, the sample point must be on the boundary of \mathcal{S} . This is shown in [41].

The dimension detection algorithm [46] exploits the same fact for manifolds with higher co-dimension. It detects the dimension of the normal space by checking if the Voronoi cell is long and skinny in the direction of a potential normal while having small diameter in its orthogonal space. The normals are approximated by pole vectors (see Definition 5.14). The Voronoi cell is successively restricted to the estimated tangent space until the ratio of the diameter of the restricted Voronoi cell and the distance to the poles is sufficiently small. The dimension of the manifold is the dimension of the first restricted Voronoi cell that meets the ratio bound. It is necessary for this algorithm that sample points cannot be arbitrary close to each other, so \mathcal{P} is supposed to be an (ϵ, δ) sample with $\delta > \frac{\epsilon}{2}$. The algorithm *SHAPE COCONE* computes a polygonal approximation of the manifold.

Two types of manifold extraction strategies are proposed for the *COCONE*-based algorithms: In [47], the *TIGHT COCONE* reconstructs closed surfaces by a volume-based approach that produces watertight polygonal reconstructions. In [41], a surface-based approach allows to reconstruct surfaces with boundaries. Starting from one triangle, the surface is developed like an advancing oriented front while choosing new triangles from the candidate set of the *COCONE*. Boundary edges were detected prior to the manifold extraction using the boundary detection algorithm. Both strategies can be applied almost directly to the candidate set computed by the \mathcal{T} -neighbor based approach. Details are given Section 8.2.2.

Power Crust

The *Power Crust* algorithm by Amenta et al. [8] is also based on the fact that the medial axis is approximated by the poles but it makes a very different use of this idea. Remember that the medial axis is the locus of the centers of medial balls, i.e. the largest empty spheres touching any point on the surface (Definition 5.10). Any surface point lies on the boundary of two medial balls: the outer medial ball that is outside the solid whose boundary is S and the inner ball that is included in the solid. The medial balls can be approximated by the Delaunay balls that are centered on a pole. This yields a good approximation because the poles are close to the medial axis and the Delaunay balls are the largest balls that are empty of sample points. They are called *polar balls*.

Amenta et al. develop some heuristics to distinguish outer and inner polar balls. Of course this resumes to decide whether a pole lies inside or outside the solid. Then, the *power crust* computes the power diagram of the polar balls, and the reconstructed surface are the facets of the power diagram that are dual to edges between inner and outer polar balls. The result is guaranteed to constitute a manifold.

In contrary to the other Delaunay based methods, the *power crust* is quite robust against noise. In [8], the authors conjecture that the power crust is the three-dimensional equivalent of the following method by Attali.

r -regular shapes

Attali [11] introduces the notion of *normalized mesh* which corresponds to the restricted Delaunay triangulation (Section 5). The algorithm, valid only in $2D$, relies again on the observation that the circumspheres to Delaunay triangles approximate the medial ball. The filter criterion is related to another property of medial balls: since the medial balls are tangent to the curve, the angle at a curve point between the center of the inner and the outer medial ball is π . Consequently, if an edge is part of the reconstruction, the angle formed at one of the sample points by the vectors to the two Voronoi vertices dual to the two triangles adjacent to the edge should also be close to π . Attali filters the Delaunay triangulation with respect to this angle criterion. An r -regular shape is a shape whose curvature is bounded by r . If the underlying curve is r -regular, then a minimum sampling density exists such that the algorithm by Attali produces a polygonal curve with the same topology as the original curve. This idea was difficult to extend to $3D$ due to the presence of slivers.

Localized Delaunay approach

This approach by Gopi et al. [65] is very similar to the \mathcal{T} -neighbor based approach in the sense that the surface is approximated locally around each sample point by a two-dimensional local triangulation. In contrast to most other methods presented so far, it does not construct the $3D$ Delaunay triangulation. The surface normals are approximated from the k -nearest neighbors of each point. For each sample point \mathbf{p}_i , potential neighbors are collected in some sphere of influence around \mathbf{p}_i whose radius is a multiple of the distance from \mathbf{p}_i to its nearest neighbor in \mathcal{P} . This candidate point set is further pruned with respect to the distance of the candidate points to the (estimated) tangent plane. At last, the candidate points are mapped (or projected) onto the tangent plane in a way that the distance between \mathbf{p}_i and its potential neighbor is not changed by the mapping. The local Delaunay triangulation of \mathbf{p}_i and the projected candidate points is computed and the “umbrella” of triangles around \mathbf{p}_i with their original $3D$ geometry are potential triangles of the polygonal reconstruction. The authors categorize their approach within the advancing front approaches since the manifold is extracted in this way, by propagating a contour. Of course, the local triangulations are not always coherent with one another and simple heuristics are proposed to disambiguate these cases.

Natural neighbor interpolation

Boissonnat and Cazals [22] define an implicit function whose zero-set approximates the surface S . Following an idea of Hoppe et al. [72], they associate, to each sample point, the signed distance function representing the distance to its tangent plane. The implicit function is defined by natural neighbor interpolation of the signed distance functions. They show that the zero-level of the implicit function cannot be far from the underlying surface and that, on a surface point, the implicit function has a small value. As

$\epsilon \rightarrow 0$, the zero-contour of the implicit function approaches the surface and the Delaunay triangulation of the point set restricted to the implicit function is a polygonal approximation homeomorphic to \mathcal{S} . Furthermore, the quality of the polygonal approximation can be evaluated by looking at the distance of its facets to the implicit surface. It can be refined by adding points on the zero-level of the implicit surface wherever the quality of the polygonal reconstruction is poor. Moreover, the implicit function can be used to simplify the polygonal surface by deleting points while guaranteeing that the reconstruction stays correct [21].

Umbrella Check

Last but not least, we present an original method to re-enforce the manifold property with the help of a system of linear constraints. The *Gabriel* complex is a subset of the Delaunay triangulation that contains all faces whose smallest circumsphere is empty of sample points. Adamy et al. define in [1] an order on the Gabriel simplices which is related to the largest empty ball passing through the vertices of the simplex and being empty of sample points. The simplex has higher priority and it is more likely to belong to the reconstruction if there is at least one side on which all other sample points are far away.

Secondly, they formulate a topological condition for the result, called *umbrella condition*: Because on a closed surface, the neighborhood of each sample point is homeomorphic to a disk, this should also be the case in the polygonal reconstruction. For each sample point, there must exist a local “umbrella” of triangles, i.e. a set of triangles that forms a disk around it. The algorithm chooses such an umbrella for each sample point from the ordered set of Gabriel simplices. The existence of local patches that are homeomorphic to the surface does not yet guarantee global coherence. Therefore, it is necessary to do a topological clean-up. First, all triangles are deleted that are either incident to a vertex but none of its incident triangles are or two umbrellas are in conflict at a triangle’s vertex. This way, some holes might have been created which must be filled while respecting the topological constraints. The idea is to formulate topological surface conditions as a system of linear constraints such that an integer value solution specifies a topological correct solution to fill the hole.

Curve reconstruction methods

There exists several methods that are designed specially for curve reconstruction. An experimental study by Althaus et al. described in [6] compares five approaches for curve reconstruction in $2D$ with respect to the reconstruction quality as well as the running time. Amenta and Bern’s *Crust*, a Delaunay approach by Gold [64] and the *Conservative Crust* by Dey, Mehlhorn and Ramos [43] were reported to be quite sensible concerning the sampling density. The two remaining methods, notably the TSP-algorithm by Giesen [61] (later improved by [5]) and the *NN-Crust* by Dey and Kumar [42], seem to work best. The TSP-algorithm computes the Traveling Salesmen Path of \mathcal{P} . Giesen showed that there exists a sampling density for which the TSP is a correct reconstruction of the curve. It is notable that this algorithm does also work for non-smooth, open curves and in arbitrary dimensions. The extension in [5] is written for the

two-dimensional case only.

The *NN-Crust* [42] is easily generalized to curves in higher dimensions. The method is very simple: Each each sample point is connected to its nearest neighbors. Then, vertices that are not incident to two edges are further connected to the closest point that makes an angle greater than $\frac{\pi}{2}$ with the first edge. The authors show that the *NN-Crust* yields a correct polygonal reconstruction of the curve if \mathcal{P} is an ϵ -sample with $\epsilon < \frac{1}{3}$.

8.2.2 Reconstruction with \mathcal{T} -neighbors: the basic idea

In this section, we explain the basic idea behind reconstruction with \mathcal{T} -neighbors. The main ingredients are explained in Section 7.1.2. Roughly speaking, our approach consists of approximating the manifold, in a first time, by the dual $\mathcal{G}(\mathcal{P})$ of the \mathcal{T} -restricted Voronoi atlas (Definition 7.5) or by a subset of it. We obtain a local approximation of the manifold around each sample point. For $k > 1$, the local approximations do not match consistently. Thus, the reconstruction algorithm needs a post-processing step that establishes a closed manifold by adding faces of $\text{Del}(\mathcal{P})$ to $\mathcal{G}(\mathcal{P})$ and by deleting some faces. Other algorithms which are based on a similar approach suffer from the same problem and propose manifold extraction methods. See for example the co-cone algorithm [7] and the λ -complex [1].

We refer to the computation of $\mathcal{G}(\mathcal{P})$ as the \mathcal{T} -neighbor computation because $\mathcal{G}(\mathcal{P})$ connects all pairs of \mathcal{T} -neighbors by an edge.

Outline: the \mathcal{T} -neighbor based reconstruction algorithm performs the following steps:

- (S1) Normal estimation (e.g. via poles)
- (S2) \mathcal{T} -neighbor computation
- (S3) Manifold composition

The \mathcal{T} -neighbor based reconstruction algorithm is very close to the *COCONE*-based algorithms, *CO-CONE* [7] for surfaces in 3D and *COCONE-SHAPE* algorithm for general manifolds embedded in \mathbb{R}^d [46]. The *COCONE*-based algorithms collect all Delaunay k -simplices such that the dual Voronoi face intersects the co-cones of its vertices. They are called the *candidate simplices* T . Our algorithm is more strict and collects all Delaunay faces such that the dual Voronoi face intersect the tangent spaces of the vertices. Of course, since the tangent planes are included in the co-cones, a Voronoi edge intersects the tangent planes only if it intersects the co-cones. We deduce that $\mathcal{G}(\mathcal{P}) \subseteq T$.

This observation makes it possible to use several approaches that have been proposed for *COCONE*-based algorithms. Notably, the shape dimension algorithm and the boundary detection algorithm apply directly to our approach.

One important difference between both methods lies in the fact that, provided the sampling condition is satisfied, the set of candidate simplices T computed by a *COCONE*-based algorithm must contain the simplices of the Delaunay triangulation $\text{Del}_{\mathcal{M}}(\mathcal{P})$ restricted to the manifold \mathcal{M} . This is not always the case for the Delaunay complex $\mathcal{G}(\mathcal{P})$. Consequently, the manifold extraction algorithms proposed by Dey et al. may not be applied directly to our approach because $\mathcal{G}(\mathcal{P})$ may not contain a subset of simplices homeomorphic to a k -manifold.

For closed surface reconstruction in $3D$, the *TIGHT-COCONE* algorithm can be adapted to our algorithm. Also other methods could be applied to establish a closed manifold from $\mathcal{G}(\mathcal{P})$. For example, the LP-based approach by Adamy et al. [1] seems particularly adequate because it is specially designed to fix local holes by extending the surface through them. Its disadvantage is the running time which is very dependent on the particular model and the size and nature of the holes. Advancing-front approaches such as the algorithm of Cohen-Steiner and Da [36] could be run locally in order to propagate the surface through the holes.

In the sequel, we consider in more detail the case of curves in $3D$. For curves, we show the correctness of the algorithm independently of the dimension of the embedding space.

8.2.3 \mathcal{T} -neighbor based curve reconstruction

In this section, we prove that the collection of edges $\mathcal{G}(\mathcal{P})$ reconstructs the smooth curve \mathcal{C} if it is well sampled. Figure 8.14 shows an example of a reconstructed curve.




Figure 8.14: The example of a reconstructed curve given 60 sample points (two different viewpoints).

Topological correctness

Let \mathcal{C}_R be the polygonal approximation of \mathcal{C} obtained by connecting each pair of sample points that is adjacent on \mathcal{C} . If \mathcal{P} is an ϵ -sample ($\epsilon < 1$), the curve \mathcal{C} is homeomorphic to its polygonal approximation \mathcal{C}_R because any arc between two adjacent sample point on \mathcal{C} can be straightened out to a segment without traversing the medial axis. Consequently, the segment cannot intersect any other component of \mathcal{C} . We show under which sampling condition $\mathcal{G}(\mathcal{P}) = \mathcal{C}_R$ which proves that $\mathcal{G}(\mathcal{P})$ is homeomorphic to \mathcal{C} . We state this result formally.

Theorem 8.1 *The collection of edges $\mathcal{G}(\mathcal{P})$ is homeomorphic to the curve \mathcal{C}*

- (a) *if $\epsilon < 0.31$ and the normals and bi-normals to \mathcal{C} are known at each sample point,*
- (b) *if $\epsilon < 0.135$, if the normals are estimated by pole vectors.*
- (c) *if $\epsilon < 0.26$, if the normals are estimated from the nearest neighbor.*

Proof: Suppose that the distance from a sample point $\mathbf{p}_i \in \mathcal{P}$ to a Voronoi vertex \mathbf{v} of $\overline{V}(\mathbf{p}_i)$ (the Voronoi cell restricted to the tangent line \mathcal{T}_i) is smaller than $\text{lfs}(\mathbf{p}_i)$. We show that under this condition $\mathcal{G}(\mathcal{P}) = \mathcal{C}_R$, i.e. (1) $\mathcal{G}(\mathcal{P}) \subseteq \mathcal{C}_R$ and (2) $\mathcal{C}_R \subseteq \mathcal{G}(\mathcal{P})$. The same arguments apply if the tangent lines are estimated (by either method) and if the distance from a sample point $\mathbf{p}_i \in \mathcal{P}$ to a Voronoi vertex $\tilde{\mathbf{v}}$ of $\tilde{V}(\mathbf{p}_i)$ (the Voronoi cell restricted to the estimated tangent line $\tilde{\mathcal{T}}_i$) is smaller than $\text{lfs}(\mathbf{p}_i)$.

Part 1: Let $e = (\mathbf{p}_i, \mathbf{p}_j)$ be an edge of \mathcal{C}_R . We show that e is part of the restricted Delaunay triangulation if $\epsilon < 1$. Then, we show under which sampling condition e is part of $\mathcal{G}(\mathcal{P})$. Let \mathbf{q} be the intersection point of the bisector plane \mathcal{H}_{ij} with the arc of \mathcal{C} formed by \mathbf{p}_i and \mathbf{p}_j that is empty of other sample points. Refer to Figure 8.15.




Figure 8.15: (a) If the sphere centered on \mathbf{q} is empty of sample points, $e \in \text{Del}_{\mathcal{C}}(\mathcal{P})$. If the sphere centered on \mathbf{v} is empty of sample points, $e \in \text{Del}_{\mathcal{T}_i}(\mathcal{P})$.

By Lemma 5.4, we know that the ball centered on \mathbf{q} that passes through the sample point closest to \mathbf{q} intersects \mathcal{C} in a topological disk. Otherwise its radius would be greater than $\text{lfs}(\mathbf{q})$ which contradicts

the sampling condition. Since \mathbf{q} lies in the arc of \mathcal{C} covered by \mathbf{p}_i and \mathbf{p}_j and since this arc is empty of other sample points, \mathbf{p}_i and \mathbf{p}_j must be the sample points closest to \mathbf{q} (and, by definition of \mathbf{q} , both are at the same distance to \mathbf{q}). Consequently, \mathbf{q} lies on the Voronoi face dual to e and e is part of the restricted Delaunay triangulation $\text{Del}_{\mathcal{C}}(\mathcal{P})$. We need to show under which sampling condition the intersection of \mathcal{H}_{ij} with the tangent line \mathcal{T}_i lies also on the Voronoi face dual to e . If this is the case, e is an edge of $\text{Del}_{\mathcal{T}_i}(\mathcal{P})$ and, therefore, of $\mathcal{G}(\mathcal{P})$.

Let \mathbf{v} be this intersection point, $\mathbf{v} := \mathcal{H}_{ij} \cap \mathcal{T}_i$. If $\|\mathbf{p}_i - \mathbf{v}\| < \frac{1}{2}\text{lfs}(\mathbf{p}_i)$, the ball with center \mathbf{v} passing through \mathbf{p}_i intersects \mathcal{C} in a topological disk (to be exact, its intersection with \mathcal{C} is the arc covered by \mathbf{p}_i and \mathbf{p}_j). Thus, by the same argument as above, it contains no other sample points, and we found a necessary and sufficient condition for e to be in $\mathcal{G}(\mathcal{P})$.

(a) From the proof of Lemma 6.3(a), we know that the distance between the Voronoi vertex \mathbf{v} of the restricted Voronoi cell $\bar{V}(\mathbf{p}_i)$ and \mathbf{p}_i is at most $\|\mathbf{p}_i - \mathbf{v}\| \leq \frac{\epsilon}{\sqrt{1-2\epsilon}} \text{lfs}(\mathbf{p}_i)$. In consequence, if $\frac{\epsilon}{\sqrt{1-2\epsilon}} < 1 \iff \epsilon < \sqrt{2} - 1 \approx 0.31$, e is an edge of $\mathcal{G}(\mathcal{P})$.

(b) The same argumentation holds if the normals are estimated via poles: from the proof of Lemma 6.3(b), we know that the distance between the Voronoi vertex $\tilde{\mathbf{v}}$ of the restricted Voronoi cell $\tilde{V}(\mathbf{p}_i)$ and \mathbf{p}_i at most is

$$\|\mathbf{p}_i - \tilde{\mathbf{v}}\| \leq \frac{1}{2} \frac{\epsilon}{1-\epsilon} \sec(\arcsin(\frac{\epsilon}{1-\epsilon}) + \beta) \text{lfs}(\mathbf{p}_i),$$

where β is the error bound between the normals to \mathcal{C} at \mathbf{p}_i and the estimated normals (in the sense of Lemma 5.15). Thus,

$$\|\mathbf{p}_i - \tilde{\mathbf{v}}\| \leq \text{lfs}(\mathbf{p}_i) \iff \epsilon \leq \frac{-1 + \sin(\beta) + \sqrt{2 - 2\sin(\beta)}}{1 + \sin \beta}.$$

In the general case, $\beta \leq 4\sqrt{d-k} \arcsin(\frac{\epsilon}{1-\epsilon})$ and for $d = 3$ and $k = 1$, the bound evaluates to $\epsilon < 0.135$.

(c) If the normal space is estimated to be orthogonal to the edge with the nearest neighbor of \mathbf{x} , we know from Lemma 5.17 that the angle $\beta \leq \arcsin \frac{\epsilon}{1-\epsilon}$. We obtain

$$\|\mathbf{p}_i - \tilde{\mathbf{v}}\| \leq \frac{1}{2}\text{lfs}(\mathbf{p}_i) \iff \epsilon \leq 0.26.$$

Part 2: Now, let $e = (\mathbf{p}_i, \mathbf{p}_j)$ be an edge of $\mathcal{G}(\mathcal{P})$ and \mathbf{v} be the vertex dual to e in $\bar{V}(\mathbf{p}_i)$ (resp. in $\tilde{V}(\mathbf{p}_i)$). We show that $e \in \mathcal{C}_R$ if $\|\mathbf{p}_i - \mathbf{v}\| \leq \text{lfs}(\mathbf{p}_i)$. (In the first part of the proof, we showed under which condition on ϵ , this bound must be true.) In this case, the ball centered on \mathbf{v} passing through \mathbf{p}_i and \mathbf{p}_j intersects \mathcal{C} in a topological disk (again by Lemma 5.4). Since \mathbf{v} is a Voronoi vertex, it is empty of sample points, and in particular the arc of \mathcal{C} inside the ball contains no sample points. Yet, \mathbf{p}_i and \mathbf{p}_j bound the arc on the each side, so that \mathbf{p}_i and \mathbf{p}_j are adjacent on \mathcal{C} . Consequently, e is an edge of \mathcal{C}_R if the sampling condition holds. \square

Notice that, for the cases (a) and (c), this proof is independent of the dimension of the embedding space since it is based on Lemma 6.3(a) and Lemma 5.4 which are independent of the dimension. However, the

normal estimation via poles is less exact if the dimension of the embedding space is higher than 3. In the two-dimensional space, the result is $\epsilon < 0.4$.

8.3 Centroidal Voronoi diagrams on a surface

Surfaces represented by a set of unordered sample points or by surface meshes are encountered in many application areas such as computer graphics, computer aided design (CAD) and reverse engineering, image processing, and scientific computation. Rarely, these point clouds are evenly distributed over the surface, nor do they follow any adaptive distribution depending on local curvatures or similar criteria that allow to optimize between the number of sample points and the information they carry. Rather, the point clouds are issued from laser scanners or similar acquisition devices and the sampling density is far from optimal. Also, it is often suitable to have regular meshes where the triangles' angles are maximized. This is, for example, very crucial for finite element computations on the mesh. In this section, we restrict ourselves to the three-dimensional case which has the most practical impact. However, the idea should apply to any dimension.

8.3.1 Introduction and previous work

In this section, we define a centroidal Voronoi diagram in general. Then, we introduce previous work on the enlargement of this concept to sample points issued from a surface and we present a definition of the centroidal Voronoi diagram on surfaces in detail. In the next section, we present our solution that is an approximation of the centroidal Voronoi diagrams on surfaces. Experimental results are shown in the last section.

Centroidal Voronoi diagrams

The mass centroid \mathbf{z}_i of the Voronoi cell $V(\mathbf{p}_i)$ is defined by

$$\mathbf{z}_i = \frac{\int_{V(\mathbf{p}_i)} \mathbf{x} d\mathbf{x}}{\int_{V(\mathbf{p}_i)} d\mathbf{x}} \quad \text{for } i = 1, \dots, n. \quad (8.1)$$

A *centroidal Voronoi tessellation (CVT)* is a Voronoi diagram such that $\mathbf{z}_i = \mathbf{p}_i, i = 1, \dots, n$. See [49] for a recent overview about centroidal Voronoi diagram, including a proof of existence, applications, and algorithms to compute CVTs.

Constrained centroidal Voronoi diagrams and related work

Recently, Du, Gunzburger and Ju [50] define a centroidal Voronoi diagram for the case that the point set is constrained to lie on a (hyper-)surface of \mathbb{R}^d . This problem has been motivated in the beginning of this section. The so-called *constrained centroidal Voronoi tessellation (CCVT)* is defined with respect to

the (Euclidean) Voronoi diagram Vor_S restricted to the surface. The definition of the centroidal Voronoi diagram needs to be modified because, in general, the mass centroid of a Voronoi cell in Vor_S does not lie on the surface. The generalized definition proposed in [50] is as follows.

The *constrained mass centroid* \mathbf{z}_i^c of a Voronoi cell $V_S(\mathbf{p}_i)$ is the solution to the following minimization problem:

$$\mathbf{z}_i^c = \arg \min_{\mathbf{z} \in S} F_i(\mathbf{z}), \quad \text{with} \quad F_i(\mathbf{z}) = \int_{V_S(\mathbf{p}_i)} \|\mathbf{x} - \mathbf{z}\|^2 d\mathbf{x}. \quad (8.2)$$

The authors show that there exists always a solution for this minimization problem and, furthermore, that this solution is the projection onto S of the (unconstrained) centroid of $V_S(\mathbf{p}_i)$ defined by Equation 8.1 (with $V(\mathbf{p}_i) = V_S(\mathbf{p}_i)$) along the normal to S at \mathbf{z}_i^c . This yields the following algorithm to compute the CCVT by a Lloyd's relaxation [76].

1. Build the Voronoi diagram $V_S(\mathcal{P})$.
2. Compute the centroids of the n Voronoi cells and project them onto the surface S .
3. Replace \mathcal{P} by the set of projected centroids and repeat Steps 1 and 2 until satisfactory convergence is achieved.

In the traditional setting without restriction to a surface, Du, Faber and Gunzburger [49] report a significantly faster convergence of Lloyd's relaxation algorithm in comparison with other methods. In the surface setting, the major drawback is the computation of the Voronoi diagram restricted to the surface. According to the authors, there is no software available, except for Voronoi diagrams on the sphere [86]. Therefore, probabilistic methods have been developed for its computation that do not need the computation of $\text{Vor}_S(\mathcal{P})$. The reader may refer to [50] for further reading.

For isotropic surface remeshing, Alliez et al. [4] employ centroidal Voronoi diagrams that are defined in the parameter space of the surface. The results are very convincing as to the regularity of the mesh, etc. Nevertheless, this approach has several disadvantages. In particular, in order to compute the parameterization, it is necessary to flatten the surface which means to cut it in the case of closed surfaces. This might lead to artifacts along the cut edges. Also, it is necessary to add a density function when computing the centroidal Voronoi diagram in order to compensate the flattening of the surface.

In this thesis, we propose a method to approximate the constrained centroidal Voronoi diagram defined by Du et al. Instead of computing the CCVT in the surface restricted Voronoi diagram, we rely on the restriction of each Voronoi cell to the tangent plane of its generator. In Chapter 7, we showed that, if the sampling of the surface is sufficiently dense, the cell of a sample point in the so-called Voronoi atlas is very close to the corresponding cell in the Voronoi diagram restricted to S . This encouraged us to test Lloyd's method in the Voronoi atlas. Our method is very efficient since it is based on the computation of very small power cells which determine the Voronoi cells in the atlas. It can be applied on closed smooth

surfaces, as well as on triangle meshes. No restriction is made about the genus, or the number of connected components of the surface.


8.3.2 T -restricted centroidal Voronoi diagrams

To avoid the computation of $\text{Vor}_S(\mathcal{P})$, we propose to approximate it by the Voronoi atlas $\mathcal{A}(\mathcal{P})$ that is defined in Chapter 7. Because every cell in $\mathcal{A}(\mathcal{P})$ is very close to the corresponding cell in $\text{Vor}_S(\mathcal{P})$ (by Lemma 6.3, both are contained in a small ball centered on the generator), it seems reasonable that both solutions converge –up to the approximation error– to a comparable solution. We tested this method experimentally but we cannot yet present formal proofs.

The experimental results showed that the energy of each iteration, in terms of distance between the centroid and the generator of each cell, decreases rapidly and continuously. The dual Delaunay triangulation becomes very regular with respect to the distribution of edge lengths and angles. Almost all angles in the triangulation are sixty degrees. Some charts are shown in Section 8.3.3.

8.3.3 Experimental results of re-meshing with Centroidal Voronoi diagrams

The presented examples are computed starting from a triangulated surface. For each vertex \mathbf{p}_i of the triangular mesh, we compute its Voronoi cell $\bar{V}(\mathbf{p}_i)$ and the centroid of this cell. The latter is projected onto the triangular surface. In the next Lloyd iteration, \mathbf{p}_i is replaced by the projected centroid of its cell and the process is repeated.






Figure 8.16: A surface representing the Connolly surface of an alanine molecule and its re-meshed version (after 40 iterations). Running time: ~ 80 s per iteration. (The original is courtesy from Steve Oudot)




Figure 8.17: A zoom on the Connolly surface (a) the original (b) after 40 iterations.

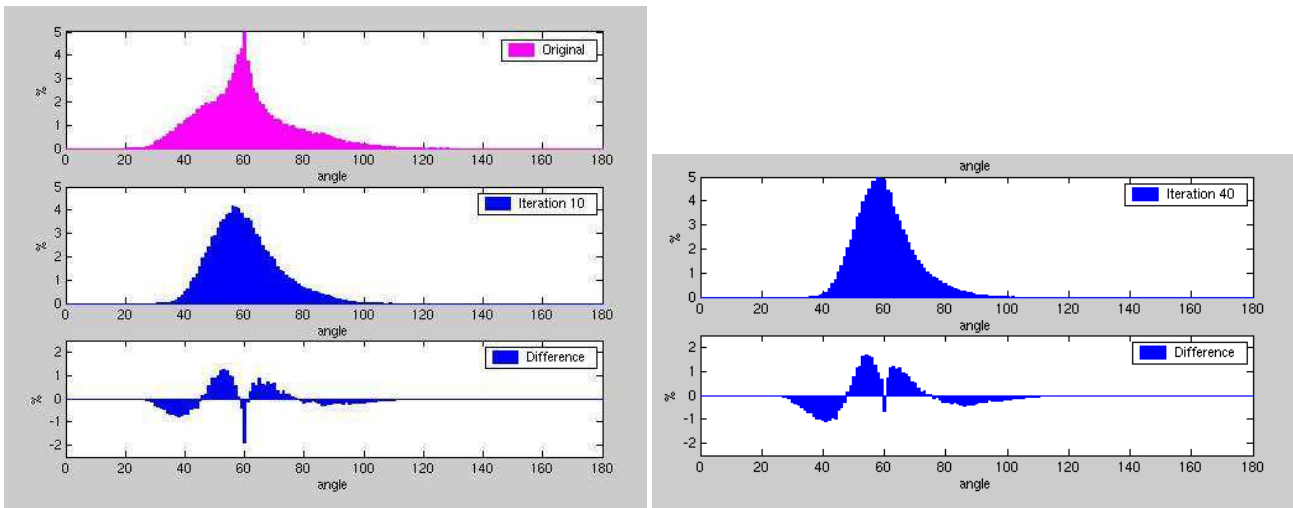


Figure 8.18: Histograms of angles depicting the difference between the original and each re-meshed version (a) after 10 iterations, (b) after 40 iterations.

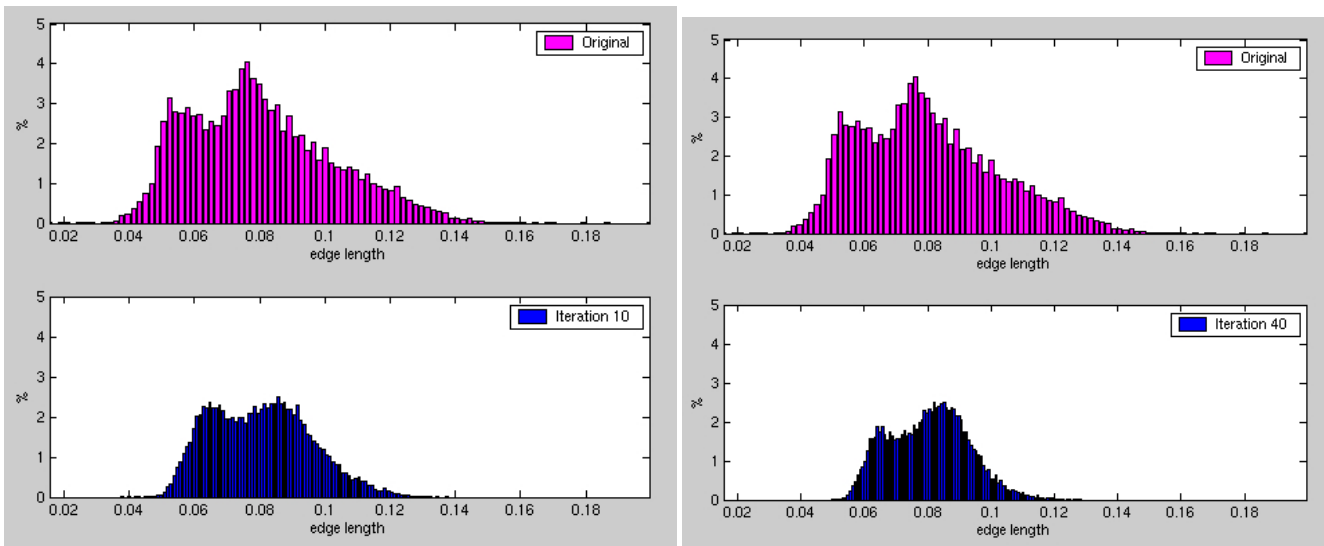


Figure 8.19: Histograms of edge lengths depicting the difference between the original and each re-meshed version (a) after 10 iterations, (b) after 40 iterations.




Figure 8.20: A zoom on another triangulated surface and its re-meshed version (100 iterations). (The original is courtesy from Steve Oudot)

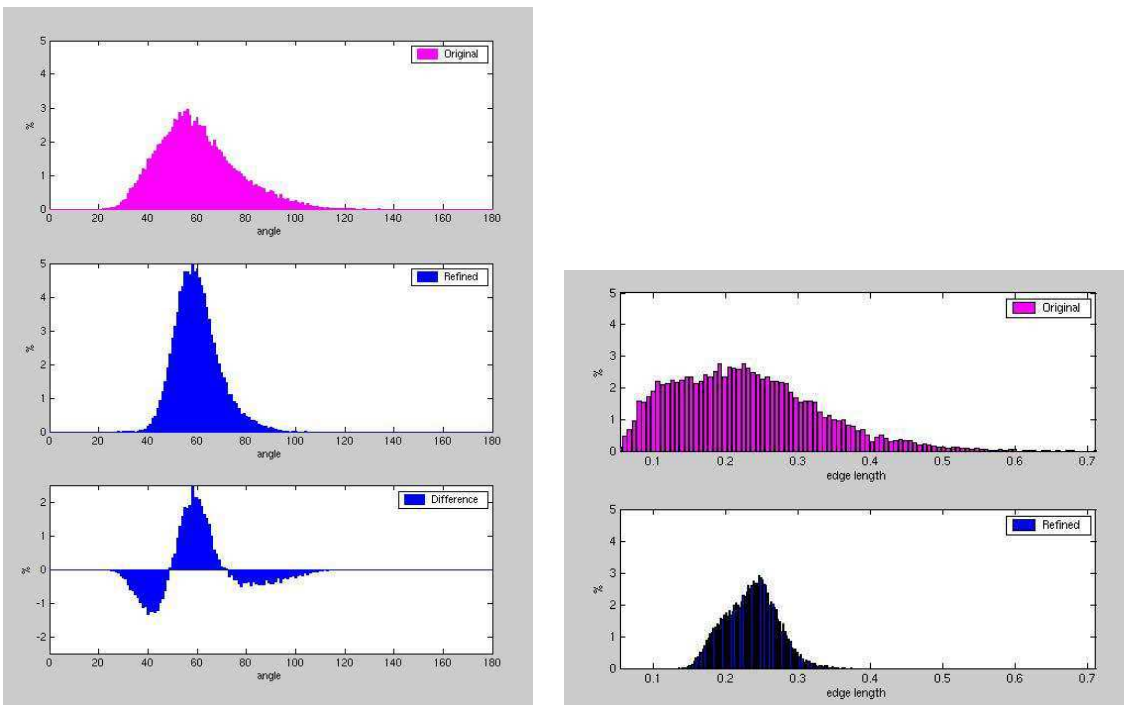


Figure 8.21: Histograms of angles and edge lengths depicting the difference between the original and the re-meshed version of Figure 8.20.

Chapter 9

Implementation of \mathcal{T} -coordinates in CGAL

9.1 Introduction

This chapter describes the implementation of the coordinate computation and the interpolation application for three-dimensional data taken from curves or surfaces. The implementation is based on CGAL (Computational Geometric Algorithms Library)[30], a C++ library of geometric algorithms that is developed by a consortium of eight research teams in Europe and Israel. CGAL started as a research program in the framework of two successive European projects. From now on, CGAL is commercialized by a private company but it remains free for academic use. The goal of CGAL is to provide robust, efficient, flexible and easy to use implementations of geometric algorithms and data-structures. The benefits are particularly interesting as geometric algorithms are known to be hard to implement because they are highly sensitive to arithmetic errors and because they include a lot of special case treatment. The programs described in the sequel make use of the $3D$ Delaunay triangulation as well as the $2D$ regular triangulation of CGAL. See e.g. [23] for an introduction to triangulations in CGAL and [70] for the general concepts. The manual pages of CGAL are available at <http://www.cgal.org>.

Design overview and specifications

The implementation of the coordinate computation as well as its application to interpolation on surfaces is designed in order to achieve a maximum of robustness and flexibility. All predicates are based on the computation of low degree determinants which allows fast filtering to avoid robustness problems. At a higher level, the template mechanism is used to assure flexibility for example in the nature of the function values to interpolate, e.g. the same software interpolates real value and vector functions.

The classes can be divided in several levels of dependency. The lowest level contains the computation of the $1D$ and $2D$ power diagrams that describe the intersection of a $3D$ Voronoi diagram with a line or a plane. For the $2D$ case, this means to compute a regular triangulation by using specially designed geo-

metric tests applied on the the $3D$ points. This module is called *Voronoi Sections*. A thorough description is given in Section 9.2.1.

The second level models the point cloud itself and contains the functions to compute the T -neighborhood and T -coordinates. It is a Delaunay triangulation of the point cloud that contains additional data, added to the vertices of the Delaunay triangulation, like the normals of the surface at the points. The additional functionalities are the computation of the Voronoi cell $\bar{V}(\mathbf{x})$ of a given point $\mathbf{x} \in \mathcal{S}$, the computation of the T -coordinates of \mathbf{x} , and the computation of the pole of \mathbf{x} in order to approximate the normal of \mathcal{S} at \mathbf{x} . This class is called *Surface neighbors*. It is described in Section 9.2.2.

At the highest level, the concept of T -coordinates is applied to function interpolation on a surface. There are two different settings: either the surface is defined by a parameterized surface (the function of the surface is known) or it is just known from the sample points. In either case, the sample points are associated to a function value and, if known, to the gradient of this function at the sample point. The function might be a real or a vector value function. There are four different types of interpolants: the linear interpolant, a quadratic interpolant, Farin's interpolant and Sibson's Z^1 interpolant. For the last three interpolants, it is necessary to know the gradient of the function. See Section 9.2.3 for the details.

9.2 Main classes and their functionalities

The design of CGAL is based on a strict separation between the combinatorial part of an algorithm and the geometric primitives necessary for the algorithm to deduce the combinatorial information. The geometric primitives and the arithmetic used for the computations are defined in the so-called *traits* class. For a given algorithm, like the computation of a regular triangulation, the traits class can be easily replaced using the template mechanism of C++. As long as the combinatorial part of the algorithm remains the same, the implementation of the algorithm is unchanged – apart from the geometric tests. The software described in the sequel uses this technique at a large extend.

9.2.1 Voronoi sections

In this section, we describe the details of the implementation of Voronoi sections in $3D$. In other words, we compute the power diagram that describes the intersection of a $3D$ Voronoi diagram with a plane \mathcal{H} ($2D$ Voronoi intersection), or the partition of a line l into intervals that correspond to the $3D$ Voronoi cells intersected by l ($1D$ Voronoi intersection). As usual, we proceed by computing the dual of the Voronoi intersection diagram, that is a regular triangulation for the $2D$ intersection and a sequence of points for $1D$.

The plane \mathcal{H} (resp. the line l) is a member of the traits class that is used in the regular triangulation. (The traits class is instantiated each time a Voronoi intersection diagram is computed.) The points of \mathcal{P} are inserted incrementally, and the natural neighbors as well as the natural coordinates of a point with respect

to the Voronoi intersection diagram might then be computed.

2D Voronoi intersection

In Chapter 6, we showed that the Voronoi intersection diagram is the power diagram of the projected sample points weighted by the negated squared distance between the projected and the original point. Consequently, a straightforward implementation constructs this power diagram via its dual regular triangulation using CGAL. However, it should be avoided to explicitly construct the projected points and to run the algorithm on the derived data because this increases numerical rounding errors. Also, CGAL provides efficient filtering methods which allow to do certified exact computation. These techniques cannot be used on derived data.

We explain how the classical incremental algorithm works, in order to show that a simple change in the traits class of the regular triangulation is sufficient in order to compute the dual of the Voronoi intersection diagram. The vertex v exists in the power diagram iff there is an empty sphere centered on v that is orthogonal to three sites. In the dual regular triangulation, this means that the power distance between the triangle's orthosphere and all sites of \mathcal{B} is positive or zero. During the incremental construction of a regular triangulation, a triangle (B_j, B_k, B_l) of the current triangulation is destroyed at the insertion of a new site B_i iff the power distance between B_i and the triangle's orthosphere is negative. In this case, the vertex v that is dual to the triangle does no longer exist. After collecting all triangles that are in conflict with B_i , the second part of the algorithm, namely, the creation of the new triangles incident to B_i , is purely combinatoric.

We call *power test* the predicate that computes the sign of the power distance between a site B_i and the orthosphere of three sites B_j, B_k and B_l , and say that B_i is *in conflict* with the triangle (B_j, B_k, B_l) if the power test is negative. The incremental algorithm to construct the regular triangulation depends only on two predicates: the predicate *power test* and the predicate *orientation* that determines the orientation of a triangle in the plane (i.e. whether the triangle's vertices are given in clockwise or counterclockwise order).

The following observation makes it possible to compute the Voronoi intersection diagram without constructing the projected points: When computing the intersection of a 3D Voronoi diagram with a plane, a Voronoi vertex is generated by the intersection of a Voronoi edge of the 3D-Voronoi diagram with the plane. The existence of a point \mathbf{v} on the Voronoi edge \mathbf{e} is witnessed by an empty sphere centered on \mathbf{v} having the three generating points of \mathbf{e} on its boundary (namely, $\mathbf{p}_j, \mathbf{p}_k, \mathbf{p}_l$ if $\mathbf{e} = V(\mathbf{p}_j) \cap V(\mathbf{p}_k) \cap V(\mathbf{p}_l)$). In particular, this must be the case for $\mathbf{v} = \mathbf{e} \cap \mathcal{H}$, and a vertex \mathbf{v} exists in the 2D Voronoi intersection diagram iff there exists a sphere centered on $\mathbf{v} \in \mathcal{H}$ that has three sample points on its boundary and is empty of sample points. Concerning the dual, this means that a triangle (B_j, B_k, B_l) exists in the regular triangulation iff there is an empty sphere centered on \mathcal{H} having $\mathbf{p}_k, \mathbf{p}_l, \mathbf{p}_m$ on its boundary.

It follows that the dual of the Voronoi intersection diagram for \mathcal{H} can be computed by simply replacing the *power test* predicate in the 2D incremental construction algorithm by a `plane_centered_in_`

sphere-test that tests whether the point p_i to insert is inside the sphere centered on \mathcal{H} and passing through three points forming a triangle. In fact, the `plane_centered_in_sphere`-test resumes to computing the sign of a 4×4 and a 3×3 determinant as described in detail below. The *orientation* predicate is replaced by a three-dimensional orientation test that determines the orientation of the triangle $(p_j p_k p_l)$ seen from a point $p_k + \vec{n}$ where \vec{n} is the normal to \mathcal{H} . With help of Figure 9.1, it is quite easy to see that the orientation of $(\overline{p_j p_k p_l})$ seen from the positive halfspace determined by \mathcal{H} (e.g. from the point $p_k + \infty \vec{n}$) must be the same as the orientation of $(p_j p_k p_l)$ seen from $p_k + \vec{n}$.




Figure 9.1: The orientation of $(\overline{p_j p_k p_l})$.

CGAL provides a very convenient way to replace predicates: the regular triangulation class is templated by the traits classes which defines all predicates. Therefore, we only need to define our proper traits class that defines the weighted points to be 3D points and that calls *power test* what is, in fact, the `plane_centered_in_sphere`-test. The plane \mathcal{H} for which we compute the Voronoi intersection diagram is a member of the traits class. It is defined by a point and the normal vector \vec{n} to \mathcal{H} and needs to be known at the creation of an instance of the traits class.

The 2D Voronoi intersection diagram has the same functionalities as CGAL's `Regular_neighbors_2.h` class including the incremental construction of the regular triangulation, the dual functionalities and the computation of T -coordinates. In addition, there exist access functions to obtain the plane \mathcal{H} defining the intersection.

1D Voronoi intersections

The 1D case is not available in its full functionality. In fact, we compute only the cell $\overline{V}(\mathbf{x})$ of a point \mathbf{x} . We assume, furthermore, that l is defined by \mathbf{x} and by two orthonormal vectors to \mathcal{M} at \mathbf{x} , the normal and the `bi_normal`. (This corresponds to our setting.) We compute $V^+(\mathbf{x}) \cap l$ as well as the decomposition of $V^+(\mathbf{x}) \cap l$ into subcells $\overline{V}(\mathbf{x}, p_i)$, $i = 1, \dots, n$, (see Section 3). The T -neighbors of \mathbf{x} are the sample

points \mathbf{p}_i with $V(\mathbf{x}, \mathbf{p}_i) \neq \emptyset$, and the T -coordinates are defined by the ratio of the lengths of $V(\mathbf{x}, \mathbf{p}_i)$ and $V^+(\mathbf{x})$. The reasoning that allows to avoid the explicit projection of the sample points onto the line l is the same as in the $2D$ case: a vertex $v = V(\mathbf{p}_j) \cap V(\mathbf{p}_k) \cap l$ exists in the Voronoi intersection diagram iff the sphere centered on l with \mathbf{p}_j and \mathbf{p}_k on its boundary is empty. The predicate that tests for a point \mathbf{p}_i if it is inside the sphere centered on l with \mathbf{p}_j and \mathbf{p}_k on its boundary is called `line_centered_in_sphere` test. We say that \mathbf{p}_i is in conflict with the edge $(\mathbf{p}_k, \mathbf{p}_l)$ if the `line_centered_in_sphere` test is positive. The algorithm is based on a second predicate, `smaller_along_line_3`, that compares the position of the projections of two points on the oriented line l . The incremental algorithm is very simple:

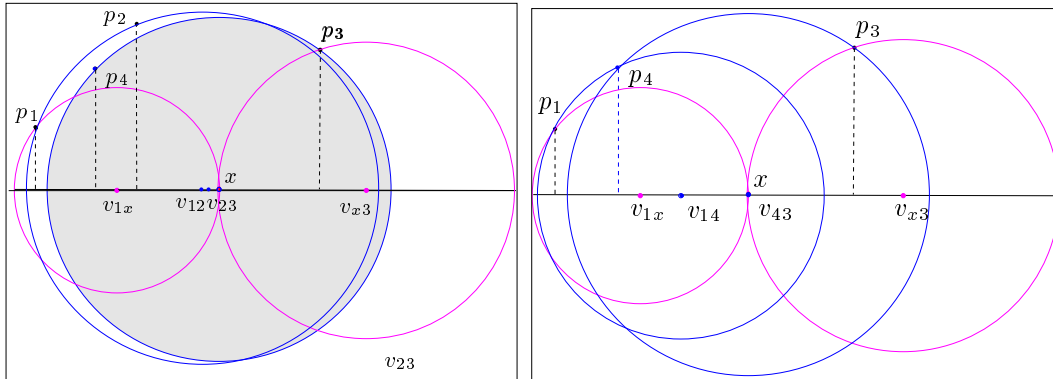


Figure 9.2: Insertion of \mathbf{p}_4 : (a) \mathbf{p}_4 is in conflict with $(\mathbf{p}_1, \mathbf{p}_2)$. \mathbf{p}_2 is not in conflict with \mathbf{p}_4 and \mathbf{p}_3 . (b) $\bar{V}(\mathbf{x})$ after the insertion.

we maintain the list of T -neighbors of \mathbf{x} ordered by their projections along l (`smaller_along_line_3`). The extremal points, \mathbf{p}_{\min} and \mathbf{p}_{\max} are the leftmost and the rightmost points of this list. A new point \mathbf{p}_i is located in the T -neighbor list (i.e. its projection is located on the line l), and we test if it is in conflict with its neighbors in the list. If the test is negative, \mathbf{p}_i is not a T -neighbor of \mathbf{x} . Otherwise, \mathbf{p}_i is inserted into the T -neighbor list. (If its projection is left of $\bar{\mathbf{p}}_{\min}$ on l , we test if \mathbf{p}_i is in conflict with $(\bar{\mathbf{p}}_{\min}, \mathbf{x})$. If the test is positive, \mathbf{p}_i is inserted into the list, and all points with projection left of \mathbf{p}_i are removed. If the projection is right of $\bar{\mathbf{p}}_{\max}$, we proceed similarly.) The `line_centered_in_sphere` test determines recursively whether a neighbor \mathbf{p}_j of \mathbf{p}_i is in conflict with \mathbf{p}_i and the second neighbor on the other side of \mathbf{p}_j . If not, \mathbf{p}_j is removed because $V(\mathbf{x}, \mathbf{p}_j)$ no longer intersects l . See also Figure 9.2.

Predicates for Voronoi intersection diagrams

In this section, we give the details of the `plane_centered_in_sphere` predicate. The parameters of the predicate are \mathcal{H} , determined by a point \mathbf{a} and the normal $\vec{\mathbf{n}}$, as well as four points $\mathbf{p}, \mathbf{q}, \mathbf{r}$ and \mathbf{t} . The predicate decides whether \mathbf{t} lies inside the sphere centered on \mathcal{H} with $\mathbf{p}, \mathbf{q}, \mathbf{r}$ on its boundary.

First, all points are translated such that \mathbf{p} becomes the origin ($\mathbf{p} := \mathbf{0}, \mathbf{a} := \mathbf{a} - \mathbf{p}, \mathbf{q} := \mathbf{q} - \mathbf{p}$,

$\mathbf{r} := \mathbf{r} - \mathbf{p}, \mathbf{t} := \mathbf{t} - \mathbf{p}$). Whether \mathbf{t} lies inside the sphere centered on \mathcal{H} with \mathbf{p}, \mathbf{q} , and \mathbf{r} on its boundary is determined by the sign of the following expression:

$$\begin{vmatrix} r_x & r_y & r_z & \mathbf{r}^2 \\ q_x & q_y & q_z & \mathbf{q}^2 \\ n_x & n_y & n_z & 2 \mathbf{n} \cdot \mathbf{a} \\ t_x & t_y & t_z & \mathbf{t}^2 \end{vmatrix} / \begin{vmatrix} n_x & n_y & n_z \\ q_x & q_y & q_z \\ r_x & r_y & r_z \end{vmatrix} \quad (9.1)$$

The `line_centered_in_sphere` predicate with parameters l , determined by a point \mathbf{a} and the normals $\vec{\mathbf{n}}$ and $\vec{\mathbf{b}}$, as well as three points \mathbf{p}, \mathbf{q} , and \mathbf{r} is determined by the sign of the following expression:

$$\begin{vmatrix} r_x & r_y & r_z & \mathbf{r}^2 \\ q_x & q_y & q_z & \mathbf{q}^2 \\ n_x & n_y & n_z & 2 \mathbf{n} \cdot \mathbf{a} \\ b_x & b_y & b_z & 2 \mathbf{b} \cdot \mathbf{a} \end{vmatrix} / \begin{vmatrix} q_x & q_y & q_z \\ n_x & n_y & n_z \\ b_x & b_y & b_z \end{vmatrix} \quad (9.2)$$

9.2.2 Surface and curve neighbors

The next level of the software models the point cloud \mathcal{P} and provides a way to compute the Voronoi cell $\bar{V}(\mathbf{x})$ of a point $\mathbf{x} \in \mathcal{M}$ as well as its T -coordinates. Of course, this class could simply be a container for the points of \mathcal{P} , and, given a manifold point \mathbf{x} with its normal(s), the Voronoi intersection diagram for \mathbf{x} can be computed using all points of \mathcal{P} . However, from the preceding chapters, it follows that only few points of \mathcal{P} are T -neighbors of \mathbf{x} , and it is, in fact, too costly to consider all of \mathcal{P} when computing the Voronoi cell $\bar{V}(\mathbf{x})$. So, we do not compute the entire Voronoi intersection diagram with the plane $T_{\mathbf{x}}$ but we insert only the sample points relevant to compute $\bar{V}(\mathbf{x})$ and its partition into subcells $\bar{V}(\mathbf{x}, \mathbf{p}_j)$. There are different ways to filter \mathcal{P} such that we get all T -neighbors of \mathbf{x} but as few points of \mathcal{P} as possible. We chose the following: we know that the T -neighbors of \mathbf{x} are a subset of the Delaunay neighbors of \mathbf{x} in $\text{Vor}^+(\mathcal{P})$. By maintaining the three-dimensional Delaunay triangulation of \mathcal{P} , we can efficiently determine the neighbors of \mathbf{x} and insert only these into the Voronoi intersection diagram. It is sometimes necessary to compute $\text{Del}(\mathcal{P})$, anyway, if the normal of \mathbf{x} needs to be determined by its pole vector.

The classes `Surface_neighbors_3.h` and `Curve_neighbors_3.h` inherit from the 3D Delaunay triangulation. In addition to the insert, remove, and dual functionalities inherited from `Delaunay_triangulation_3.h`, they allow to compute the T -neighbors of a surface (curve) point supplied with its normal(s). The traits class of `Surface_neighbors_3.h` (`Curve_neighbors_3.h`) has no special requirements other than to follow the model for the Delaunay triangulation traits class. However, if one wants to compute the T -coordinates for a sample point itself or if the reconstruction functionalities described in Section 8.2.2 and 8.2.3 are used, a special `Triangulation_vertex_base_with_id_and_normal.h` is needed in order to associate each sample point with a normal and an id. The vertex id is only necessary for visualization purposes and might be omitted.

***T*-coordinate computation**

The T -coordinates can be computed for a point $\mathbf{x} \in \mathcal{M}$ that is not a sample point or for a sample point $\mathbf{p}_i \in \mathcal{P}$ with respect to the Voronoi diagram of $\mathcal{P} \setminus \{\mathbf{p}_i\}$. In both cases, the normal(s) to the manifold needs to be supplied. The result of the coordinate computation is an instance of the class `Voronoi_intersection_coordinates.h` that contains a vector of pairs associating a sample point (via a `vertex_handle`) to the corresponding coordinate function value.

In the case $\mathbf{x} \notin \mathcal{P}$, the `find_conflicts` function of the Delaunay triangulation allows to simulate the insertion of \mathbf{x} into the Delaunay triangulation. The sample points in conflict with \mathbf{x} are the Delaunay neighbors of \mathbf{x} in $\text{Vor}^+(\mathcal{P})$. Similarly, for $\mathbf{p}_i \in \mathcal{P}$, we insert only the Delaunay neighbors of \mathbf{p}_i into the Voronoi intersection diagram with the plane T_i which is sufficient to compute the cell $V(\mathbf{p}_i) \cap T_i$ and its partition into subcells $\bar{V}(\mathbf{p}_i, \mathbf{p}_j)$.

Reconstruction

The reconstruction function allows to dump an unordered set of triangles that approximate the surface. It is the dual $\mathcal{G}(\mathcal{P})$ of the Voronoi atlas of \mathcal{P} defined in Section 7.1.2 or its subsets $\mathcal{G}_i(\mathcal{P}), i \leq 3$. For this, we need to compute the local umbrella U_i for each sample point $\mathbf{p}_i \in \mathcal{P}$. The theoretical properties of $\mathcal{G}(\mathcal{P})$ and $\mathcal{G}_3(\mathcal{P})$ are described in Section 7.1.2 and 8.2.2. The manifold extraction procedure is not implemented, thus, the collection of triangles does not constitute a manifold. In the class, there exists also a function to compute the umbrella of a sample point separately. This is of interest if a local approximation of the surface around a specific sample point is needed. The next section describes a fast method to compute the local umbrella U_i of $\mathbf{p}_i \in \mathcal{P}$.

Local umbrella computation

The following method allows to compute all triangles incident to a sample \mathbf{p}_i in a two-dimensional triangulation – without computing the entire triangulation. It is inspired from [65] who proposes a similar approach to compute a local Delaunay triangulation.

Let $n_i = |\mathcal{N}_i|$ denote the number of sample points that remain after the filtering step as potential neighbors of \mathbf{p}_i . In our case, \mathcal{N}_i is the set of neighbors of \mathbf{p}_i in the Delaunay triangulation. To compute the umbrella U_i of $\mathbf{p}_i \in \mathcal{P}$, i.e. the triangles of $\text{Del}_{T_i}(\mathcal{P})$ that are incident to \mathbf{p}_i , we propose a local method with $O(n_i \log n_i)$ running time. This complexity is theoretically the same as for the computation of the regular triangulation of \mathcal{N}_i , yet, in practice, it is a lot faster. The output is a sorted list of the T_i -neighbors of \mathbf{p}_i . If the filtering is efficient, so that n_i is small compared to the total number of points, this method is very efficient. If the surface is well sampled, the number neighbors of \mathbf{p}_i in the three-dimensional Delaunay triangulation is constant in average. A local umbrella can then be computed in constant time.

The major steps are as follows: starting with the edge to the closest neighbor \mathbf{p}_{\min} of \mathbf{x} in \mathcal{P} , the points

$p_j \in \mathcal{N}_i$ are sorted with respect to the angle formed by \overline{p}_{\min}, x and \overline{p}_j , the projection of p_j onto T_i . The sorting takes $O(n_i \log n_i)$ time. The predicate necessary for the sorting consists of at most three 3D orientation tests.

Then, the T_i -neighbors are computed in the following loop inspired from the Graham scan algorithm for convex hull. We iterate over the ordered set of points (contained in an array called `adjacent`) and eliminate points until all triplets of consecutive points fulfill the empty-sphere property (refer to Section 2.2).

```
//compute local umbrella of p:
//neighbors: the array containing the result
//adjacent:  the array containing the candidate points,
//           sorted by angle around p

//the outer loop:
for(int i=0; i < adjacent.size(); ++i){
  //the inner loop:
  while (k > 0 &&
         traits.power_test(neighbors[k-1]->point(), neighbors[k]->point(),
                           p, adjacent[i]->point()) == ON_POSITIVE_SIDE){
    --k;
    neighbors.pop_back();
  }
  k++;
  neighbors.push_back(adjacent[i]);
}
```

The entire loop takes linear time which we can see by the following argument: each time the test in the inner loop evaluates true, a point is erased from \mathcal{N}_i , thus, there is at most a linear number of negative tests. For each sample point passed in the outer loop, there is at most one positive evaluation of the test.

Normal estimation

If the normal is not supplied by the user, it is estimated by the pole vector which is the vector from the data point to its furthest Voronoi vertex. The function `set_poles()` computes the pole vector for all sample points and updates the appropriate field in the `vertex_base` instance. If this functionality is used, it is recommended to use the `Triangulation_cell_base_with_lazy_info.h` such that the dual of

each cell can be kept in memory and repeated computations of the same Voronoi vertex are avoided. In the case of curves, we compute a second pole vector as described in Section 5.3 and in [46].

An alternative method is implemented that estimates the normals from the nearest neighbors of a point. The theoretical justification is described in Section 5.3. For curves, this method works very well. The normals are orthogonal to the line segment that connects the sample point \mathbf{p}_i to its nearest neighbor. For surfaces, two nearest neighbors are chosen such that the triangle of \mathbf{p}_i and its two neighbors has a large angle at \mathbf{p}_i . This method appears to be very unstable in practice.

9.2.3 Function interpolation on a surface

The application that we developed based on `Surface neighbors_3` is function interpolation on a surface. To each sample point, a function value of any suitable user-defined type is given. The value type is specified by a template parameter of the class `Interpolation`. In addition, the gradient value, a vector containing the directional derivatives of the function at the sample point in x , y , and z direction might be given for each sample point. The four different interpolants I^0 , I^1 , f^3 and Z^1 described in Chapter 4 are implemented. The result of the interpolation is of the user-defined value type. The interpolation functions are called with a surface point $\mathbf{x} \in S \notin \mathcal{P}$ together with the surface normal $\vec{n}_{\mathbf{x}}$ as parameter. For testing, it is possible to pass the interpolated function to the class `Interpolation` (via a functor) and to compute some error statistics concerning the mean and the maximum error between the interpolated function value and the real one.

Visualization is possible by moving each surface point by the interpolated function value in direction of the normal at the point. The reconstruction functions of `Surface neighbors_3` allow to visualize the result in order to evaluate the smoothness and exactness of the method. Another possibility for visualization is to color the points according to the function value. This has not been implemented.

Interpolation over parametric surfaces

The knowledge of the surface, e.g. in form of a parameterization, allows to interpolate the function on a grid of surface points. The sample points may either be distributed at random or on a grid. We implemented an extension to the `Interpolation` class which takes, in addition to sample points and function values, the parameter function in form of a function. It evaluates the interpolant over a grid. The size of the grid is specified by the number of grid points in u and in v direction. In this setting, the surface is easily approximated by a quadrilateral grid over the grid points. Visualization is done as before by the deformation approach.



Conclusion



For a point cloud issued from a manifold, this thesis presents an efficient and practical approach to approximate the so-called manifold restricted Voronoi diagram. This is the intersection of the d -dimensional Voronoi diagram of a point set $\mathcal{P} \subset \mathcal{M}$ with the manifold \mathcal{M} . It represents an important way to define a Voronoi diagram on a manifold. However, without knowing the manifold, this diagram cannot be computed and even in the contrary case, it is very costly to compute. By linearly approximating the manifold in each Voronoi cell with the tangent space to \mathcal{M} at the generator of the cell, we present an approximation of this Voronoi diagram that is provably close to the exact solution if the sample points are sufficiently dense. This Voronoi diagram defines in a natural way a neighborhood relation between points that are close on the manifold as well as a local coordinate system. We described, at first, the local neighborhood and the coordinate system for one point that lies on the manifold. Then, the Voronoi atlas is considered which is the collection of local neighborhoods for each sample point. It is shown to be useful in the context of manifold reconstruction and re-meshing.

Furthermore, this thesis develops the basic properties of natural neighbor coordinates in power diagrams. Even if this concept was defined by Aurenhammer in [12], the properties were not yet elaborated in detail. A survey over interpolation methods based on natural neighbor coordinates is as well presented. This survey completes some missing details of the original papers and presents a method by Clarkson [35] that has been, so far, unpublished.

A few open questions remain to be solved. First, we could not yet show whether Clarkson's coordinates are C^1 continuous or not and second, question 7.12 is not yet proven.

There are several important directions for future research. I will present some of the ideas in more detail.

Manifolds embedded in high dimensions

It seems straightforward to apply the presented concepts to point clouds issued from low dimensional manifolds that are embedded in higher dimensional Euclidean space. This topic attracts the attention of several research groups as well in computational geometry, e.g. Giesen and Wagner [63], as in other fields

such as cognitive science and image processing, see for instance [88], [98] and [78].

In the theoretical part of this thesis, the generalization has already been taken into account. However, further research is necessary to elaborate practical solutions. The approximation of tangent spaces seems possible based on the work of Har-Peled and Varadarajan [68]. They propose a method to compute a linear subspace of a fixed dimension that best fits a given point set. The complexity is linear with respect to the dimension of the embedding space. If the dimension of the manifold is small ($k < 3$), we presented in this thesis an efficient algorithm to compute, locally, the Voronoi cell restricted to the tangent space \mathcal{T}_i (Chapter 9). For higher dimensions, similar local methods should be possible but this question needs further investigation.

Estimation of geodesic quantities and enlargement of the Voronoi atlas on general k -manifolds


If the sampling condition holds, it should be possible to bound the maximum distance between the length of a geodesic on the manifold and the corresponding shortest paths in the Voronoi atlas. In the same context, one can determine the maximum error when approximating the area of an ϵ -sampled manifold by the area of the Voronoi cells in the atlas. Both achievements should enable the theoretical analysis of the approximation of the centroidal Voronoi diagram presented in Section 8.3 because they allow to bound the distance between a centroid in the surface restricted Voronoi diagram and its corresponding centroid in the Voronoi atlas.

In addition, it would be nice to define a natural enlargement of the Voronoi atlas in order to connect the Voronoi cells. For hypersurfaces, this is achieved in Section 7.2. Some ideas for general manifolds came up but they remain to be formalized in detail.

Noisy sampling

Another important enlargement is the treatment of noisy sample sets. In practice, data is always noisy – at least due to the fact that it is discretized. It seems quite possible that the methods presented in this thesis adapt to noisy samples. In [80], Mitra and Nguyen describe how normals can be estimated from noisy data. Also, it should be possible to detect outliers in a pre-treatment of the data. Of course, the Voronoi cells do not have the same properties for noisy data. In particular, it is no longer guaranteed that each Voronoi cell is long and skinny which is a major argument of most Voronoi-based surface reconstruction algorithms. However, it should nevertheless be possible to show that the restriction of a Voronoi cell to the tangent space of the manifold at the generator has a small diameter. Also, there is a subset of points for which the Voronoi cell is long and skinny – if not in every direction of the normal space, so at least in a subset of normal directions. A formalization of the properties should be worthwhile. The use of higher-order Voronoi diagrams in the context of noise also deserves further investigation.

Conclusion



Nous avons introduit dans cette thèse une approche pratique des diagrammes de Voronoï restreints à une variété \mathcal{M} qui consiste, à partir d'un ensemble de points \mathcal{P} échantillonnant \mathcal{M} , à approximer la restriction d'une cellule de Voronoï à la variété par sa restriction à l'espace tangent à la variété au point définissant la cellule. Une telle cellule de Voronoï restreinte à l'espace tangent permet de définir une relation de voisinage entre le point de la variété définissant la cellule et les points de l'échantillon qui possède une bonne propriété de localité ainsi qu'un système de coordonnées associés à cette notion de voisinage dont nous avons démontré les propriétés fondamentales: localité du support, continuité et différentiabilité des fonctions de coordonnées, etc. Ensuite, nous avons présenté la notion d'atlas de Voronoï et montré son utilité dans le contexte de reconstruction de variété et du remaillage. L'effectivité de l'ensemble des méthodes élaborées nous a permis de proposer des solutions efficaces pour les problèmes d'interpolation de fonctions, de reconstruction de surfaces et de remaillage.

Deux questions ouvertes restent pourtant à aborder: la question ouverte 7.12 et la preuve de la différentiabilité ou non-différentiabilité des fonctions de coordonnées de Clarkson définit dans le paragraphe 3.5.

Nous présentons dans la suite quelques directions de recherche intéressantes qui se sont présentées durant nos travaux.

Reconstruction de variétés plongées dans des grandes dimensions

Le thème de la reconstruction d'une variété (de petite dimension) plongée dans un espace euclidien de grande dimension attire l'attention des chercheurs aussi bien dans le domaine de la géométrie algorithmique (cf. Giesen et Wagner [63]), que dans celui de la science cognitive et du traitement d'images (voir [88], [98] et [78]).

La plupart des résultats théoriques de cette thèse sont indépendants de la dimension de la variété et de l'espace ambiant et s'applique donc directement à ce problème. Des recherches restent cependant nécessaires afin d'obtenir des solutions pratiques. Har-Peled et Varadarajan [68] ont proposé une méthode permettant le calcul d'un sous-espace linéaire de dimension fixée qui s'adapte le mieux à un ensemble de points donné et qui semble rendre possible l'approximation d'espaces tangents. La complexité de leur

méthode est linéaire en la dimension de l'espace dans lequel la variété est plongée. Cette thèse propose un algorithme efficace pour le calcul local de cellules de Voronoï restreintes à l'espace tangent T_i lorsque la dimension de la variété est petite ($k < 3$) (cf. chapitre 9), il serait intéressant d'étudier la généralisation de cette algorithme à de plus grandes dimensions.

Estimation de quantités géométriques et généralisation de l'atlas de Voronoï à des k -variétés

Si la condition d'échantillonnage est satisfaite, le calcul d'une borne supérieure sur la distance maximale entre une géodésique sur la variété et le plus court chemin correspondant dans l'atlas de Voronoï ainsi que le calcul majoration de l'erreur d'approximation commise en approchant l'aire d'une variété ϵ -échantillonnée par l'aire des cellules de l'atlas de Voronoï semble possible. Ce type de résultat permettrait l'obtention de garanties théoriques sur la convergence de l'approximation du diagramme de Voronoï centroïdal présentée à la section 8.3.

Il serait de plus intéressant de définir un élargissement de l'atlas de Voronoï qui connecterait les cellules de Voronoï. Cette thèse propose une solution pour le cas des hypersurfaces (cf. section 7.2), la généralisation au cas des variétés générales semble envisageable et reste à traiter.

Données bruitées

En pratique les points de données sont toujours bruités du fait de leur discrétisation, le traitement de ce type de données s'avère donc être particulièrement crucial. Les méthodes présentées dans cette thèse semblent pouvoir s'adapter à ce problème. L'estimation des normales à partir d'un échantillon bruité est possible par des méthodes telles que celle proposée par Mitra et Nguyen [80]. De plus, une grande parties des *outliers* peuvent être détecter lors d'un phase de pré-traitement des points de l'échantillon. Malheureusement, certaines propriétés du diagramme de Voronoï ne sont plus vérifiées dans le cas de données bruitées: on ne garantit désormais plus, par exemple, que chaque cellule est longue et fine, propriété utilisée par la plupart des algorithmes de reconstruction basée sur le diagramme de Voronoï. Il semble cependant envisageable de montrer que la restriction d'une cellule de Voronoï à l'espace tangent de la variété au point définissant la cellule a alors un petit diamètre et qu'il existe un sous-ensemble de points de données pour lequel les cellules de Voronoï sont longues et fines au moins pour un sous-ensemble des directions de l'espace normal. Une formalisation de ces propriétés serait digne d'intérêt. L'utilisation de diagrammes de Voronoï d'ordre supérieur semble aussi constituer une piste prometteuse pour ce type de problème.

Bibliography

- [1] U. Adamy, J. Giesen, and M. John. New techniques for topologically correct surface reconstruction. In *Proceedings of IEEE Visualization 2000*, pages 373–380, 2000.
- [2] P. Alfeld, M. Neamtu, and L. Schumaker. Fitting scattered data on sphere-like surfaces using spherical splines. *J. Comput. Appl. Math.*, 73:5–43, 1996.
- [3] Peter Alfeld, Marian Neamtu, and Larry L. Schumaker. Bernstein-bzier polynomials on spheres and sphere-like surfaces. *Comput. Aided Geom. Design.*, 13:333–349, 1996.
- [4] Pierre Alliez, Éric Coline de Verdière, Olivier Devillers, and Martin Isenburg. Isotropic surface remeshing. In *International Conference on Shape Modeling and applications*, pages 49–58, 2003.
- [5] E. Althaus and K. Mehlhorn. Polynomial time TSP-based curve reconstruction. In *Proc. 11th ACM-SIAM Sympos. Discrete Algorithms*, pages 686–695, January 2000.
- [6] E. Althaus, K. Mehlhorn, S. Näher, and S. Schirra. Experiments on curve reconstruction. In *Proc. 2nd Workshop Algorithm Eng. Exper.*, pages 103–114, January 2000.
- [7] N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. In *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, pages 213–222, 2000.
- [8] N. Amenta, S. Choi, and R. K. Kolluri. The power crust, unions of balls, and the medial axis transform. *Comput. Geom. Theory Appl.*, 19:127–153, 2001.
- [9] Nina Amenta and Marshall Bern. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.*, 22(4):481–504, 1999.
- [10] Nina Amenta, Marshall Bern, and David Eppstein. The crust and the β -skeleton: Combinatorial curve reconstruction. *Graphical Models and Image Processing*, 60:125–135, 1998.
- [11] D. Attali. r -regular shape reconstruction from unorganized points. *Comput. Geom. Theory Appl.*, 10:239–247, 1998.
- [12] F. Aurenhammer. Linear combinations from power domains. *Geom. Dedicata*, 28:45–52, 1988.

- [13] F. Aurenhammer. A relationship between Gale transforms and Voronoi diagrams. *Discrete Appl. Math.*, 28:83–91, 1990.
- [14] F. Aurenhammer. Voronoi diagrams: A survey of a fundamental geometric data structure. *ACM Comput. Surv.*, 23(3):345–405, September 1991.
- [15] C. Bajaj, F. Bernardini, and G. Xu. Reconstructing surfaces and functions on surfaces from unorganized 3D data. *Algorithmica*, 19:243–261, 1997.
- [16] C. L. Bajaj, F. Bernadini, and G. Xu. Reconstructing surfaces and functions on surfaces from unorganized three-dimensional data. *Algorithmica*, 19(1–2):243–261, September 1997.
- [17] R. Barnhill, B. Piper, and K. Rescorla. Interpolation to arbitrary data on a surface. In G. Farin, editor, *Geometric Modeling: Algorithms and New Trends*, pages 281–289. SIAM, Philadelphia, 1987.
- [18] Fausto Bernardini and Chrajit L. Bajaj. Sampling and reconstructing manifolds using alpha-shapes. In *Proceedings of the 9th Canadian Conference on Computational Geometry (CCCG'97)*, 1997.
- [19] J-D. Boissonnat and F. Cazals. Natural neighbor coordinates of points on a surface. *Computational Geometry: Theory and Applications*, 19(2-3):155–173, 2001.
- [20] Jean-Daniel Boissonnat. Geometric structures for three-dimensional shape representation. *ACM Trans. Graph.*, 3(4):266–286, 1984.
- [21] Jean-Daniel Boissonnat and Frédéric Cazals. Coarse-to-fine surface simplification with geometric guarantees. In A. Chalmers and T.-M. Rhyne, editors, *Eurographics'01*, Manchester, 2001. Blackwell.
- [22] Jean-Daniel Boissonnat and Frédéric Cazals. Smooth surface reconstruction via natural neighbour interpolation of distance functions. *Comp. Geometry Theory and Applications*, pages 185–203, 2002.
- [23] Jean-Daniel Boissonnat, Olivier Devillers, Monique Teillaud, and Mariette Yvinec. Triangulations in CGAL. In *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, pages 11–18, 2000.
- [24] Jean-Daniel Boissonnat and Julia Flötotto. A local coordinate system on a surface. In *Proc. 7th ACM Symposium on Solid Modeling and Applications*, 2002.
- [25] Jean-Daniel Boissonnat and Julia Flötotto. A coordinate system associated with points scattered on a surface. *Computer-Aided Design*, 36:161–174, 2004.
- [26] Jean-Daniel Boissonnat and Mariette Yvinec. *Algorithmic Geometry*. Cambridge University Press, UK, 1998. Translated by Hervé Brönnimann.

- [27] J. Braun and M. Sambridge. A numerical method for solving partial differential equations on highly irregular evolving grids. *Nature*, 376:655–660, 1995.
- [28] J. L. Brown. Natural neighbor interpolation on the sphere. In P.-J. Laurent, A. Le Méhauté, and L. L. Schumaker, editors, *Wavelets, Images, and Surface Fitting*, pages 67–74. A. K. Peters, Wellesley MA, 1994.
- [29] J. L. Brown. Systems of coordinates associated with points scattered in the plane. *Comput. Aided Design*, 14:547–559, 1997.
- [30] *The CGAL Manual*, 2002. Release 2.4.
- [31] Raphaëlle Chaine. A geometric convection approach of 3-d reconstruction. In L. Kobbelt, P. Schröder, and H. Hoppe, editors, *Proc. of Eurographics and ACM SIGGRAPH Symp. on Geometry Processing*, pages 233–244, 2003.
- [32] F. Chazal and R. Soufflet. Stability and finiteness properties of medial axis and skeleton. *Journal of Control and Dynamical Systems*, to appear.
- [33] Frédéric Chazal. Personal communication.
- [34] L. P. Chew. Guaranteed-quality mesh generation for curved surfaces. In *Proc. 9th Annu. ACM Sympos. Comput. Geom.*, pages 274–280, 1993.
- [35] K. L. Clarkson. <http://cm.bell-labs.com/who/clarkson/hulltalk.ps.gz>.
- [36] David Cohen-Steiner and Frank Da. A greedy delaunay based surface reconstruction algorithm. Rapport de recherche 4564, INRIA, 2002.
- [37] de Boor. B-form basics. In G. Farin, editor, *Geometric Modeling: Algorithms and New Trends*, pages 131–148. SIAM (Philadelphia), 1987.
- [38] C. de Boor. *A practical guide to splines*. Springer-Verlag, New York, NY, 1978.
- [39] Olivier Devillers, Stefan Meiser, and Monique Teillaud. The space of spheres, a geometric tool to unify duality results on Voronoi diagrams. Report 1620, INRIA Sophia-Antipolis, Valbonne, France, 1992.
- [40] T. K. Dey, J. Giesen, S. Goswami, J. Hudson, R. Wenger, and W. Zhao. Undersampling and oversampling in sample based shape modeling. In *Proc. IEEE Visualization 2001*, pages 83–90, 2001.
- [41] T. K. Dey, J. Giesen, N. Leekha, R. Wenger, and W. Zhao. Detecting boundaries for surface reconstruction using co-cones. *Intl. J. Comput. Graphics and CAD/CAM, Invited paper*, to appear.

- [42] T. K. Dey and P. Kumar. A simple provable algorithm for curve reconstruction. In *Proc. 10th ACM-SIAM Sympos. Discrete Algorithms*, pages 893–894, January 1999.
- [43] T. K. Dey, K. Mehlhorn, and E. A. Ramos. Curve reconstruction: Connecting dots with good reason. *Comput. Geom. Theory Appl.*, 15:229–244, 2000.
- [44] Tamal K. Dey, Stefan Funke, and Edgar A. Ramos. Surface reconstruction in almost linear time under locally uniform sampling. In *Abstracts 17th European Workshop Comput. Geom.*, pages 129–132. Freie Universität Berlin, 2001.
- [45] Tamal K. Dey and Joachim Giesen. Detecting undersampling in surface reconstruction. In *Proc. 17th Annu. ACM Sympos. Comput. Geom.*, pages 257–263, 2001.
- [46] Tamal K. Dey, Joachim Giesen, Samrat Goswami, and Wulue Zhao. Shape dimension and approximation from samples. In *Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 772–780. ACM Press, 2002.
- [47] Tamal K. Dey and Samrat Goswami. Tight cocone: A water-tight surface reconstructor. In *Proc. 8th ACM Symposium on Solid Modeling and Applications*. ACM Press, 2003.
- [48] Manfredo P. do Carmo. *Differential geometry of curves and surfaces*. Prentice-Hall, 1976.
- [49] Qiang Du, Vance Faber, and Max Gunzburger. Centroidal voronoi tessellations: Applications and algorithms. *SIAM Review*, 41(4):637–676, 1999.
- [50] Qiang Du, Max Gunzburger, and L. Ju. Constrained centroidal voronoi tessellations for surfaces. *SIAM Journal of Scientific Computing*, 24(5):1488–1506, 2003.
- [51] H. Edelsbrunner. Surface reconstruction by wrapping finite point sets in space. *Discrete and Comput. Geom.—The Goodman-Pollack Festschrift*, to appear 2003.
- [52] H. Edelsbrunner, D. G. Kirkpatrick, and R. Seidel. On the shape of a set of points in the plane. *IEEE Trans. Inform. Theory*, IT-29:551–559, 1983.
- [53] H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graph.*, 13(1):43–72, January 1994.
- [54] H. Edelsbrunner and N. R. Shah. Triangulating topological spaces. In *Proc. 10th Annu. ACM Sympos. Comput. Geom.*, pages 285–292, 1994.
- [55] G. Farin. Smooth interpolation to scattered 3d data. In R.E. Barnhill and Boehm W., editors, *Surfaces in CAGD*, pages 43–63. North Holland Publishing Company, 1983.
- [56] G. Farin. Surfaces over Dirichlet tessellations. *Comput. Aided Geom. Design*, 7:281–292, 1990.

- [57] G. Farin. *Curves and Surfaces for Computer Aided Geometric Design: A Practical Guide*. Academic Press, 1993.
- [58] G. Fasshauer and L. Schumaker. Scattered data fitting on the sphere. In T. Lyche M. Daehlen and L. L. Schumaker, editors, *Methods for Curves and Surfaces II*, pages 117–166. Vanderbilt University Press, Nashville TN, 1998.
- [59] Thomas A. Foley, David A. Lane, Gregory M. Nielson, Richard Franke, and Hans Hagen. Interpolation of scattered data on closed surfaces. *Computer Aided Geometric Design*, 7(1-4):303–312, 1990.
- [60] Stefan Funke and Edgar A. Ramos. Smooth-surface reconstruction in near-linear time. In *Proc. 13th ACM-SIAM Symp. on Discrete Algorithms, San Francisco, California*, pages 781 – 790, 2002.
- [61] J. Giesen. Curve reconstruction, the TSP, and Menger’s theorem on length. In *Proc. 15th Annu. ACM Sympos. Comput. Geom.*, pages 207–216, 1999.
- [62] J. Giesen and M. John. Surface reconstruction based on a dynamical system. In *Proc. 23rd Annual Conference of the European Association for Computer Graphics (Eurographics)*, pages 363–371, 2002.
- [63] Joachim Giesen and Uli Wagner. Shape dimension and intrinsic metric from samples of manifolds with high co-dimension. In *Proc. 19th Annu. ACM Sympos. Comput. Geom.*, pages 329–337, 2003.
- [64] C. Gold. Crust and anti-crust: A one-step boundary and skeleton extraction algorithm. In *Proc. 15th Annu. ACM Sympos. Comput. Geom.*, pages 189–196, 1999.
- [65] M. Gopi, S. Krishnan, and C. T. Silva. Surface reconstruction based on lower dimensional localized Delaunay triangulation. In *Eurographics*, 2000.
- [66] Branko Grünbaum. *Convex polytopes*. Wiley-Interscience, 1967.
- [67] Dianne Hansford. Bézier techniques. In G. Farin, J. Hoschek, and M.-S. Kim, editors, *Handbook of computer aided geometric design*. Elsevier Science B.V., 2002.
- [68] S. Har-Peled and K.R. Varadarajan. High-dimensional shape fitting in linear time. In *Proc. 19th Annu. ACM Sympos. Comput. Geom.*, pages 39–47, 2003.
- [69] R.L. Hardy. Multiquadric equations of topography and other irregular surfaces. *J. Geophysical Res.*, 76:1905–1915, 1971.
- [70] Susan Hert, Michael Hoffmann, Lutz Kettner, Sylvain Pion, and Michael Seel. An adaptable and extensible geometry kernel. Rapport de recherche 4270, INRIA, 2001.

- [71] Hisamoto Hiyoshi and Kokichi Sugihara. Voronoi-based interpolation with higher continuity. In *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, pages 242–250, 2000.
- [72] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. *Comput. Graphics*, 26(2):71–78, 1992. Proc. SIGGRAPH '92.
- [73] Richard Kunze, Franz-Erich Wolter, and Thomas Rausch. Geodesic Voronoi diagrams on parametric surfaces. Technical Report 2, Welfen Laboratory, Universität Hannover, July 1997.
- [74] Greg Leibon and David Letscher. Delaunay triangulations and Voronoi diagrams for Riemannian manifolds. In *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, pages 341–349, 2000.
- [75] J. Leray. Sur la forme des espaces topologiques et sur les points fixes des représentations. *J. Math. Pure Appl.*, 24:95–167, 1945.
- [76] Stuart P. Lloyd. Least squares quantization in pcm. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982.
- [77] S.K. Lodha and R. Franke. Scattered data interpolation: Radial basis and other methods. In G. Farin, J. Hoschek, and M.-S. Kim, editors, *Handbook of computer aided geometric design*. Elsevier Science B.V., 2002.
- [78] Facundo Memoli and Guillermo Sapiro. Distance functions and geodesics on point clouds. Preprint, May 2003.
- [79] John W. Milnor. *Topology from the Differential Viewpoint*. University of Virginia Press, 1965.
- [80] Niloy J. Mitra and An Nguyen. Estimating surface normals of a noisy point cloud data. In *Proc. 19th Annu. ACM Sympos. Comput. Geom.*, pages 322–328, 2003.
- [81] L. Moccozet and N. Magnenat Thalmann. Dirichlet free-form deformations and their application to hand simulation. In *Proceedings of Computer Animation '97, Geneva, Switzerland*, 1997.
- [82] Marian Neamtu. Splines on surfaces. In G. Farin, J. Hoschek, and M.-S. Kim, editors, *Handbook of computer aided geometric design*. Elsevier Science B.V., 2002.
- [83] Stéphane Nullans. *Reconstruction tridimensionnelle de structures géologiques à partir de données hétérogènes*. Thèse de doctorat en sciences, université de Nice-Sophia Antipolis, France, 1998.
- [84] B. Piper. Properties of local coordinates based on dirichlet tessellations. *Computing Suppl.*, 8:227–239, 1993.
- [85] H. Pottmann. Interpolation on surfaces using minimum norm networks. *Comput. Aided Geom. Design.*, 7:313–321, 1992.

- [86] R. J. Renka. Algorithm 772: STRIPACK: Delaunay triangulation and voronoi diagrams on the surface of a sphere. *ACM Trans. Math. Softw.*, 23:416–434, 1997.
- [87] Günter Rote. Personal communication.
- [88] Sam Roweis and Lawrence Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, Dec. 22 2000.
- [89] M. Sambridge, J. Braun, and H. McQueen. Geophysical parameterization and interpolation of irregular data using natural neighbours. *Geophys. J. Int.*, 122:837–857, 1995.
- [90] R. Sibson. A vector identity for the Dirichlet tessellation. *Math. Proc. Camb. Phil. Soc.*, 87:151–155, 1980.
- [91] R. Sibson. A brief description of natural neighbour interpolation. In Vic Barnett, editor, *Interpreting Multivariate Data*, pages 21–36. John Wiley & Sons, Chichester, 1981.
- [92] N. Sukumar and B. Moran. C1 natural neighbor interpolant for partial differential equations. *Numerical Methods for Partial Differential Equations*, 15(4):417–447, 1999.
- [93] W. Sweldens. The lifting scheme: a construction of second generation wavelets. *SIAM J. Math. Anal.*, 29:511–546, 1998.
- [94] J. Giesen T. K. Dey and J. Hudson. Decimating samples for mesh simplification. In *Proc. 13th Canadian Conf. on Computational Geometry*, pages 85–88, 2001.
- [95] J. Giesen T. K. Dey and J. Hudson. Delaunay based shape reconstruction from large data. In *Proc. IEEE Symposium in Parallel and Large Data Visualization and Graphics*, pages 19–27, 2001.
- [96] J. Giesen T. K. Dey and J. Hudson. Sample shuffling for quality hierarchic surface meshing. In *Proc. 10th Int. Meshing Roundtable, Sandia Nat. Lab.*, pages 143–154, 2001.
- [97] J. Giesen T. K. Dey and W. Zhao. Robustness issues in surface reconstruction. In *Proc. Intl. Conf. Comput. Sci., San Francisco, California*, pages 28–30, 2001.
- [98] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, Dec. 22 2000.
- [99] D. F. Watson. *nngriidr: An implementation of natural neighbor interpolation*. David Watson, 1994.
- [100] David F. Watson. *Contouring: A Guide to the Analysis and Display of Spatial Data*. Pergamon, 1992.
- [101] F. E. Wolter. Cut locus and medial axis in global shape interrogation and representation. Technical Report 92-2, MIT, Dept. Ocean Engg., Design Lab, Cambridge, MA 02139, USA, January 1992.

AVIS DU JURY SUR LA REPRODUCTION DE LA THESE SOUTENUE

1. TITRE DE LA THESE : Un système de coordonnées associé à un échantillon de points d'une surface :
définition, propriétés et applications
2. NOM ET PRENOM DE L'AUTEUR : Madame FLOTOTTO Julia
3. MEMBRES DU JURY : Pierre Bernhard, Bernard Lacolle, André Liecker,
Jean Marie Morvan, Günter Rote, Jean-Daniel Boissanaud.
4. PRESIDENT DU JURY : Pierre BERNHARD
5. DATE DE LA SOUTENANCE : 22 septembre 2003.
6. REPRODUCTION DE LA THESE SOUTENUE :
 - Thèse pouvant être reproduite en l'état
 - Thèse ne pouvant pas être reproduite
 - Thèse pouvant être reproduite après correction suggérées au cours de la soutenance

Signature du Président du jury

P. B - L I

UNIVERSITE DE NICE-SOPHIA ANTIPOLIS

ECOLE DOCTORALE DES SCIENCES ET TECHNOLOGIES DE L'INFORMATION ET DE LA COMMUNICATION

Rapport de soutenance

De doctorat de Madame FLOTOTTO Julia

Titre du mémoire : « Un système de coordonnées associé à un échantillon de points d'une surface :
définition, propriétés et applications »

Membres du jury : Pierre BERNHARD, Jean-Daniel BOISSONNAT, Bernard LACOLLE, André
LIEUTER, Jean-Marie MORVAN, Gunter ROTE.

Mademoiselle Julia Flötotto présente un travail qui utilise des outils mathématiques difficiles pour aborder des questions extrêmement pratiques et d'une grande importance dans de nombreuses applications. De ce fait, l'érudition nécessaire pour effectuer un tel travail est assez considérable : géométrie classique, géométrie différentielle, géométrie numérique, informatique proprement dite... À ce titre, les chapitres 1 à 5 du mémoire, loin d'être une simple compilation, présentent une véritable synthèse de la littérature existante. Et les applications présentées dans le chapitre 8, ainsi que l'implémentation pratique dans CGAL, présentée au chapitre 9, montrent la fécondité et l'effectivité des résultats originaux obtenus, qui font l'objet des chapitres 6 et 7.

Le mémoire est bien rédigé et d'une lecture agréable. Ajoutons qu'il constitue une référence utile sur la notion de "voisins naturels", clarifiant bien la littérature et apportant quelques démonstrations qui y faisaient défaut.

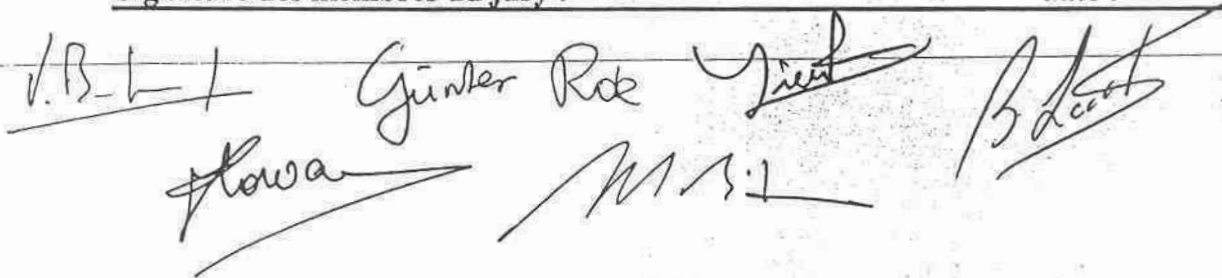
L'exposé a été très clair et didactique. Julia Flötotto a su choisir dans son travail les parties les plus significatives et les exposer avec le juste degré de technicité qui convient à cet exercice.

Enfin, les réponses apportées par Julia Flötotto aux questions posées par le jury ont montré qu'elle avait une bonne culture du domaine, et savait prendre du recul par rapport à son travail.

En conséquence, le jury décide d'accorder à Mademoiselle Julia Flötotto le grade de docteur en informatique de l'Université de Nice-Sophia Antipolis avec la mention très honorable.

Signature des membres du jury :

date :



Ecole Doctorale STIC de l'Université de Nice-Sophia Antipolis