



HAL
open science

The exploration-exploitation paradigm : a biophysical approach

Vito Dichio

► **To cite this version:**

Vito Dichio. The exploration-exploitation paradigm : a biophysical approach. Neurons and Cognition [q-bio.NC]. Sorbonne Université, 2023. English. NNT : 2023SORUS402 . tel-04406107

HAL Id: tel-04406107

<https://theses.hal.science/tel-04406107v1>

Submitted on 19 Jan 2024

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

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

Doctoral Thesis

The exploration-exploitation paradigm

A biophysical approach



Vito Dichio

A thesis submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

Jury chaired by prof. Guido Caldarelli

October 2023



Sorbonne université

21, Rue de l'école de médecine, 75006 Paris, France

Institut national de recherche en informatique et en automatique

2, Rue Simone IFF, 75012 Paris, France

Institut du cerveau

47, bd de l'hôpital 75013 Paris, France

Members of the Jury:

Fabrizio De Vico Fallani (Directeur de thèse)

Institut national de recherche en informatique et en automatique, Institut du cerveau

Guido Caldarelli (Président, Rapporteur)

Università Ca' Foscari

Marta Sales-Pardo (Rapporteuse)


Universitat Rovira i Virgili

Demian Battaglia (Examineur)

Université d'Aix-Marseille, Université de Strasbourg

Gašper Tkačik (Examineur)

Institute of Science and Technology Austria



Among Chuang-tzu's many skills, he was an expert draftsman. The king asked him to draw a crab. Chuang-tzu replied that he needed five years, a country house, and twelve servants. Five years later the drawing was still not begun. "I need another five years," said Chuang-tzu. The king granted them. At the end of these ten years, Chuang-tzu took up his brush and, in an instant, with a single stroke, he drew a crab, the most perfect crab ever seen.

– Italo Calvino, *Six memos for the next millennium*

富士山に一度も登らぬバカ、二度登るバカ
(Fools those who never climb the Mount Fuji, fools those who do it twice)

– Japanese proverb



Peak of the Mount Fuji (Japan), 3376m, 4.40 am (sunrise), 8 August 2023.

Contents

| | |
|------------------------------------------------------------|-------------|
| Contents | v |
| Acknowledgments | vii |
| Abstract | ix |
| Résumé en français | xi |
| A foreword | xiii |
| Overview of research contributions | xv |
| 1 Introduction | 1 |
| 2 Maxent inference for graphs | 9 |
| 2.1 Exponential random graph models | 9 |
| 2.1.1 A maximum entropy approach | 10 |
| 2.1.2 Model inference | 13 |
| 2.2 User handbook | 16 |
| 2.2.1 Interpretation of parameters | 16 |
| 2.2.2 The model specification problem | 18 |
| 3 Darwin and the others | 23 |
| 3.1 Subject matter | 23 |
| 3.2 Modelling evolution | 25 |
| 3.2.1 Statistical genetics | 25 |
| 3.2.2 On the fitness landscape | 28 |
| 3.3 Genetic algorithms | 30 |
| 4 The exploration-exploitation dynamics | 33 |
| 4.1 Fundamentals | 34 |
| 4.1.1 A graph dynamic | 35 |
| 4.1.2 Beyond Darwin, an interpretable GA | 36 |
| 4.2 The math of simple scenarios | 38 |
| 4.2.1 No exploitation | 38 |
| 4.2.2 Energy-like biological function | 39 |
| 4.2.3 Distance-like biological function | 40 |
| 4.3 Population-based simulations | 42 |
| 4.3.1 Gleaning dynamics from simulations | 43 |
| 5 Weaving the mind of a worm | 47 |
| 5.1 The brain wiring problem | 47 |
| 5.2 The mind of a worm | 50 |
| 5.2.1 <i>C. elegans</i> nervous system: a digest | 50 |
| 5.2.2 <i>C.elegans</i> brain maturation | 52 |
| 5.3 The EE development of a worm brain | 53 |
| 5.3.1 A minimal worm brain | 53 |
| 5.3.2 Topography of the functional landscape | 55 |
| 5.3.3 Worm brain maturation tracked down | 58 |

| | |
|--------------------------------------------------------|-----------|
| 5.4 Interpretation down to the synapse scale | 65 |
| 5.5 Paving ways | 67 |
| 6 Conclusions | 71 |
| | |
| APPENDIX | 75 |
| A Pen and paper EE dynamics | 77 |
| B Network measures | 79 |
| Glossary: biology for physicists | 83 |
| Bibliography | 85 |

Acknowledgments

Below are the words I read after being awarded the title of Doctor of Philosophy, 23 October 2023, Auditorium of the Institut du Cerveau, Paris, end of the ceremony. Surrounded by people I love, I tried in vain to hold back a some tears.

As you may know, if you are here, in recent times I have climbed a mountain. At the foot of the mountain, it was foggy, it was late evening, we could not see far. As we started, step by step, the path seemed smooth. I had never climbed a mountain; I am not familiar with such surroundings. Still, I thought, it is easy enough, the path is smooth. Then, it started to rain. A rainstorm, for hours, it was night time. We could not stop, we had to go. The stones were slippery, the rain paralysing. It is time to stop, I thought, one has to accept it, at some point, it's not a big deal, it is just a mountain. However, I had friends around me. Friends walking alongside me, probably with similar thoughts to the mines. Midway up a fucking mountain, night-time, under a heavy rain, I still thought that, among those people, that was to one of the most amazing paths I had taken in my life. At some point, climbing the mountain, we were so high that the clouds were left below, the rain with them. We were literally above the clouds. At that point I realised it was only a matter of time before we made it to the end. At 4 in the morning, on the top of mountain, I gazed out over the horizon, it was clear, it was immense. It was the sunrise, the most spectacular sunrise I have ever seen in my life.

In the first place, once more, I would like to express my most heartfelt thanks to the members of my Jury, for having accepted the invitation. At the cost of sounding rhetoric, it was an honour for me to stand on this stage and defend and discuss my research with you.

Then, what to say about Fabrizio. Fabrizio has been more than a supervisor, a friend – and (by the way) an excellent central defender on the football pitch. Thanks in particular for constant encouragement over the last three years. Also, for having granted me the possibility of spending some fruitless time in exploring the many wrong ideas I have had. More than the plots I have shown you here, more than the paper we have written, it is there that I have grown the most as a scientist.

There are so many people I should say thanks to, a PhD is a joint effort between a host of people. Before switching language, let me collectively thank all those people I love that have crowded my days, in my recent life. Those that have been fundamental for my scientific path, Mario, Erik, Hong-Li, Fulvio. Those that, this afternoon, decided to dedicate to my dissertation a few hours, here in this room – I know it's not just for the Italian cookies outside – and those that are listening to these words remotely, on Zoom, your presence means the world to me, in few selected moments of one's life, it is important to just be there. Those that have come by Rue Bonaparte, for a dinner, a drink, one night or few days. Those that I have met in some school or conference, in some corner of the world. And to this corner of the world, Paris, that I have come to call home, where I have most likely spent the most incredible years of my life. To all of you, thanks from the bottom of my heart, thanks for the times we had.

Voilà. Ça c'est le moment de partager quelques mots en français, j'y tiens beaucoup. Pour commencer, je voudrais remercier l'énorme famille de cet Institut, et d'ARAMIS. Ça fait trois ans que chaque fois qu'il y a une soutenance chez nous, il/elle dit que vous êtes incroyables, ambiance de ouf. Et vous savez quoi, c'est vrai, chaque fois c'est vrai. Donc, c'est à moi de le dire maintenant. Et je suis heureux que c'est à moi de le dire maintenant, mais j'en suis un peu triste aussi, je ne le cache pas. Je n'ai pas l'illusion de pouvoir avoir deux fois la chance que j'ai eu quand je vous ai rencontrés. Tellement de choses de cette thèse vont me manquer, et vous, bien sûr, plus que tout. Un énorme merci à Juliana, Charley, Remy, pour le parcours qu'on a fait ensemble. À Ravi et la team du foot du Dimanche. À ceux qui sont passés par la salle des stagiaires, les meilleurs postes de travail d'ARAMIS, by far. À Elisa, Domitille, Sophie et la team du quatrième.

Camille, Elise. Après le cinquième Spritz Campari, dans une place de Bologna, je vous ai dit un truc, et je voudrais que tout le monde sache. Vous êtes la joie de vivre en personne, j'adore ce que vous êtes, et ce que je

suis avec vous, grâce à vous. Merci aussi pour les slangs que vous m'avez appris. Aujourd'hui, je me nachav. No, je déconne, je déconne.

Et enfin, Tristan, Tristano. On a partagé bien plus qu'un bureau dans les dernières trois années, on a partagé un parcours de vie. Tu étais déjà là pour m'aider quand je ne comprenais rien de la procédure d'inscription au doctorat. Mais tu étais là aussi la dernière fois que je n'arrivais pas à parler, trop de larmes, je m'en souviens très bien. Au milieu, plein de choses. Il y a des stéréotypes sur les Parisiens et chaque fois que quelqu'un me demande je dis que non, ce n'est pas vrai, au contraire. Dans des cas comme ça, je pense surtout à toi, voilà.

Due ringraziamenti in Italiano. Ai tanti amici italiani qui a Parigi, a giudicare dalla mia personale e completamente unbiased statistica, una persona su due a Parigi è italiana. Grazie per tutte le chiacchiere, le serate, gli sprazzi di casa. Se non mi sono mai sentito solo in questi anni è perché casa era anche qui, casa eravate voi, #teamMattarella.

Per concludere, la casa con la C maiuscola. Nella vita un po' sconclusionata che mi trovo a vivere ci sono tuttavia alcuni punti fissi. Ad esempio, tornare giù nel mio paese e passare a prendere un caffè dalle zie. Oppure, chiamare in serata quasi ogni giorno il mio fratellino, che racconta i cazzi suoi e non fa domande. Ecco, ci tengo a ringraziare, per una volta in modo esplicito, i miei fratelli e sorelle, Cì, Roz e Aaaaangela. Non so come sia successo di preciso, ma siamo tutti arrivati adulti a volerci tutti un sacco bene. Forse non l'ho mai detto, ma sono davvero felice di essere vostro fratello, ancor di più di essere il fratello scherziero – eh sì sì è scherziero, lo è.

Gli ultimi due ringraziamenti li ho lasciati in fondo. A mia mamma, sora Anna. Mi dispiace per tutte le telefonate a cui non ho risposto negli anni, ma anche quando non l'ho fatto, mi è sempre stato caro il sapere che qualcuno, in qualunque situazione, in qualunque condizione, anche dall'altra parte del mondo, stesse pensando a me. Infine, a mio papà Raffaele, a l'ingegnere. A un certo punto tutto è sembrato andare un po' a rotoli, però stiamo ancora qua, io ho di nuovo una cravatta al collo e prendo un altro titolo di studio. Stasera andiamo a cena, e la prossima volta mi racconti se ha piovuto, se si può seminare, cosa dicevano gli antichi il giorno di San Vito.

Ah dimenticavo, un grazie anche a *cur' scem' d' Pelo*, il cagnolino più bello del mondo, che parla in dialetto e mi segue da Lemenzano.

Something I did not say that day but which was the only paragraph of a first version of these acknowledgments and which I would like to leave on the records. / My heartfelt gratitude to the many libraries around the world that have graciously hosted me during the countless hours spent in writing these pages.

- *Bibliothèque Sainte-Geneviève*, 10 place du Panthéon, Paris, France
- *Bibliothèque Mazarine*, 23 quai de Conti, Paris France
- *Bibliothèque Richelieu (BnF)*, 5 Rue Vivienne, Paris, France
- *Bibliothèque François-Mitterrand (BnF)*, Quai François Mauriac, Paris, France
- *Wienbibliothek im Rathaus*, Felderstraße 1, Wien, Austria
- *Universitätsbibliothek Wien*, Universitaetsring 1, Wien, Austria
- *Tokyo Metropolitan Central Library*, 5 Chome-7-13 Minamiazabu, Minato City, Tokyo, Japan
- *Katsushika City Library*, 6 Chome-7-13, Kanamachi, Katsushika City, Tokyo, Japan
- *Biblioteca Provinciale "T. Stigliani"*, Piazza Vittorio Veneto, Matera, Italia

Abstract

The study of living systems is notoriously challenging. The often-quoted daunting complexity of biological systems is primarily due to the intricacies of their interactions, their multiple organisation levels and their dynamic nature. In the quest to understand this complexity, parallels drawn with standard physics – in particular, statistical physics – are both useful and of limited use. On the one hand, they provide a rich set of theoretical and methodological building blocks for constructing theories and designing experiments. On the other hand, life also unfolds according to principles that are unparalleled in the physics of conventional matter.

A crucial difference lies in the notion of function: biological systems are shaped by the need to perform specific tasks. A general problem for living systems is to find and promote those configurations that yield improved or optimal functions, we call this the exploration-exploitation (EE) problem. One specific instance of the above is found in evolutionary biology. There, random genetic mutations sustain the exploration of the configuration space, with those leading to higher reproductive success being favoured by natural selection.

Inspired by the latter, we develop a novel formalism that encodes a general exploration-exploitation dynamics for biological networks. In particular, our EE dynamics is represented as an exploration of a functional landscape and consists of stochastic configuration changes combined with the state-dependent optimisation of an objective function (F metric). We begin by investigating its main features through the study of simple, analytically tractable functional landscapes. We deploy simulations for more general and complex applications.

We then turn to the brain wiring problem, i.e., the development of an individual's nervous system during its early life. We argue that this is another specific instance of the EE problem and therefore can be addressed by using our theoretical framework. In particular, we focus on brain maturation in the nematode *C. elegans*, the only organism for which a complete network of neurons and neuronal connections has been reconstructed, at multiple developmental time points (seven). We fix the network at birth and use the adult stage to infer (i) a parsimonious maxent (ERG) description of the F metric for the worm brain and (ii) the two parameters of our EE dynamics. According to the topography of its functional landscape, the adult brain is characterised by a tendency to form both triads and high degree nodes. We demonstrate that our EE dynamics in such landscape is capable of tracking down the entire developmental history. In particular, we show that the trajectory we obtain closely reproduces the other experimental time points that we did not use for inference. This is true both in the space of model statistics and for a number of other network properties. Additionally, we discuss a micro-level interpretation of the EE dynamics in terms of the underlying synapse formation process.

Our study is a first step towards the system-level understanding of the development of a natural brain and can be extended (i) to encompass more complex functional landscapes, (ii) to different organisms than the *C. elegans* and (iii) to several different problems than the brain wiring. Indeed, we posit that the exploration-exploitation paradigm is among those life-specific principles that we are just beginning to uncover.

Résumé en français

L'étude des systèmes vivants est notoirement difficile. La complexité déconcertante des systèmes biologiques, souvent citée, est principalement due à la complexité de leurs interactions, à leurs multiples niveaux d'organisation et à leur nature dynamique. Dans la quête de compréhension de cette complexité, les parallèles établis avec la physique standard - en particulier la physique statistique - sont à la fois utiles et d'une utilité limitée. D'une part, ils fournissent un riche ensemble d'éléments théoriques et méthodologiques pour construire des théories et concevoir des expériences. D'autre part, la vie biologique se déroule aussi selon des principes qui sont sans équivalent dans la physique de la matière conventionnelle.

Une différence cruciale réside dans la notion de fonction : les systèmes biologiques sont façonnés par la nécessité d'accomplir des tâches spécifiques. Un problème général pour les systèmes vivants est de trouver et de promouvoir les configurations qui produisent des fonctions améliorées ou optimales, ce que nous appelons le problème de l'exploration-exploitation (EE). Un exemple spécifique de ce problème se trouve dans la biologie évolutive. Dans ce cas, des mutations génétiques aléatoires soutiennent l'exploration de l'espace de configuration, celles qui correspondent à un succès reproductif plus élevé étant favorisées par la sélection naturelle.

Inspirés par ce dernier cas, nous développons un nouveau formalisme qui encode une dynamique générale d'exploration-exploitation pour les réseaux biologiques, représentée comme une exploration d'un paysage fonctionnel. En particulier, notre dynamique d'EE consiste en des changements de configuration stochastiques combinés à l'optimisation dépendante de l'état d'une fonction objective (métrique F). Nous commençons par étudier ses principales caractéristiques à travers l'étude de paysages fonctionnels simples et analytiquement traitables. Nous déployons des simulations pour des applications plus générales et plus complexes.

Nous nous penchons ensuite sur le problème du câblage du cerveau, c'est-à-dire le développement du système nerveux d'un individu tout au long de sa vie. Nous soutenons que ce dernier est un autre exemple spécifique du problème de l'EE et qu'il peut donc être traité à l'aide de notre cadre théorique. En particulier, nous nous concentrons sur la maturation du cerveau chez le nématode *C. elegans*, le seul organisme pour lequel un réseau complet de neurones et de connexions neuronales a été reconstruit, à plusieurs moments du développement. Nous fixons le réseau à la naissance et utilisons le stade adulte pour déduire (i) une description max.ent. parcimonieuse (ERG) de la métrique F pour le cerveau du ver et (ii) les deux paramètres de notre dynamique EE.

Selon la topographie de son paysage fonctionnel, le cerveau adulte est caractérisé par une tendance à former des triades et des nœuds de degré supérieur. Nous montrons que notre dynamique d'EE dans un tel paysage est capable de retracer toute l'histoire du développement. En particulier, nous montrons que la trajectoire que nous obtenons reproduit étroitement les autres points temporels expérimentaux que nous n'avons pas utilisés pour l'inférence. Ceci est vrai à la fois dans l'espace des statistiques du modèle et pour un certain nombre d'autres propriétés du réseau. En outre, nous discutons d'une interprétation micro-niveau de la dynamique de l'EE en termes de processus sous-jacent de formation des synapses.

Notre étude est un premier pas vers la compréhension au niveau du système du développement d'un cerveau naturel et peut être étendue (i) à des paysages fonctionnels plus complexes, (ii) à d'autres organismes que le *C. elegans* et (iii) à d'autres problèmes que le câblage du cerveau. En effet, nous pensons que le paradigme de l'exploration-exploitation fait partie de ces principes spécifiques à la vie que nous commençons à peine à découvrir.

A foreword

I have to say, I was somewhat reminded of Middle Age theological debates about how many angels can dance on the head of a needle – wrote an anonymous reviewer. He/she was reporting on the manuscript of a review I had submitted some time ago with my collaborators. A delightfully unfavourable assessment and, frankly, the reviewer was right. That first draft of our manuscript was indeed far too long, too mathsy and, ultimately, too obscure – a reader unfamiliar with those matters would never have made it to the end. The editor did not reject our paper, the report said ‘major revision’, it felt like a last chance. A famous quote from Blaise Pascal goes: *I have made this letter longer because I have not had time to make it shorter*. In the weeks that followed, I came to understand its meaning first hand. Months later, the review was accepted for publication.

There is a truth that everyone in academia knows, the students know it, the seniors know it, the blackboards know it, the orchids in the office know it and everyone repeats the same curse in chorus: *no one will ever read your thesis*. After all, why should anyone? Everything publishable has already been published or is at most under revision. With a few rare exceptions, PhD reports end up being more of a *souvenir* of your roaring 20s than a real scientific contribution. Your parents will keep a copy in plain sight on the brightest bookshelf, and you will pick it up from time to time, to read the acknowledgements page again and remember with nostalgia who was there, who not anymore, who not yet.

Done this way, I believe, it is a missed opportunity, besides a waste of time. As students, we spend a considerable amount of our PhD time crafting such a document, in the midst of (probably) the peak of our scientific creativity. How to rescue these pages from their doomed fate? This question has haunted me throughout the writing. I obviously have no illusions of success but I still believe it was worth trying. Concretely, this meant making some stylistic and editorial choices that, in my subjective and absolutely questionable judgement, resulted in a clearer and more useful manuscript.

First and foremost, the content had to be original. By this I do not mean the inclusion of new topics or results – all of those discussed here have already been made available to the scientific community by standard means. Rather, it meant minimising the overlap with the content of the papers: often adding to, sometimes subtracting from, in any case striving to offer a broader or complementary discussion.

Second, the content had to be coherent. This starkly contrasts with the typical developmental trajectory of a PhD project, which is often winding, riddled with branching points, wrong ideas, dead ends. Therefore, I deliberately omitted from this manuscript a minor but non trivial fraction of the work from these years. Instead, I preferred to direct the reader’s attention on a streamlined narrative of the core scientific idea of this project, in its formal, press release version.

Finally, the content had to be compact. Less is more, they say, and I tend to agree. To paraphrase a famous quote, everything should be as short as possible, but no shorter. My ambition was to write a self-contained document in which every single discussion had a purpose later on in the text. At the same time, I insisted in keeping the discussion focused on the essentials, avoiding long and tedious expositions of well-known subjects. Indeed, I suspect that the scientific community is not dying to learn about my view on the Markov processes, or my hot takes on the Bayes theorem.

An original, coherent and compact narrative: this was my North Star. Since much of the concrete meaning of these adjectives depends of their endpoint consumer, an essential question was: who am I talking to? The reader I had in mind was primarily a physicist, not necessarily conversant with the wonders of biology but with a basic knowledge of the fundamentals of statistical mechanics. However, I hope that a biologist who happens upon this work will find it digestible. As a physicist myself, I cannot be overly confident. Still, I think it is a reasonable hope.

A great deal of help in the writing came from the L^AT_EX class `kaobook v 0.9.8` by Federico Marotta – which I have come to love and recommend *urbi et orbi*. Thanks to its wide margins, I was able to establish a hierarchy of the information to be conveyed, and to separate the essentials from the frills. Indeed, I have employed

sidenotes for various purposes: to offer historical context, to expand on discussions, and to share more personal views. All of them can be ignored, and the narrative in the main body of the text should stand.

A final note, which I would like to label as the *AI statement* for this manuscript. There is not a single sentence on these pages whose content has been written alone by AI tools such as **Chat GPT**. I am far too jealous of my thoughts to blindly entrust them to algorithms that I neither understand (at present, nobody does) nor trust. I firmly claim the full intellectual authorship of these pages. On the other hand, there is not a single sentence on these pages that has not been grammatically reviewed and often improved by using both **Chat GPT** and **DeepL**. Ignoring the existence and usefulness of these tools or, worse, opposing them, would be irrational, useless and detrimental.

Tokyo, Japan
August 2023

Vito Dichio

Overview of research contributions

The source material for this manuscript has been shared with the community through the following scientific papers – in inverse chronological order.

- [1] **Vito Dichio** and Fabrizio De Vico Fallani. *The exploration-exploitation paradigm for networked biological systems*. In: arXiv e-prints 2306.17300 (2023)

Pocket abstract: The stochastic exploration of the configuration space and the exploitation of optimal functional states underlie many biological processes. The evolutionary dynamics stands out as a remarkable example. Here, we introduce a novel formalism that mimics evolution and encodes a general exploration-exploitation dynamics for biological networks. We apply it to the brain wiring problem, focusing on the maturation of the *C. elegans* connectome. We demonstrate that a parsimonious maxent description of the adult brain combined with our framework is able to track down the entire developmental trajectory.

- [2] **Vito Dichio** and Fabrizio De Vico Fallani. *Statistical models of complex brain networks*. In: Reports on Progress in Physics 86.10 (2023), p. 102601

Pocket abstract: Exponential random graph (ERG) models are a class of maximum entropy inference models for network data. Here, we provide a pedagogical introduction to the theory and methods of ERG models, including a critical discussion of the limitations of the approach. We review recent investigations based on their application to the study of functional, large-scale brain networks – fMRI, EEG data. ERG models can be used to characterise networks, discriminate between different states and predict the outcome of neurological diseases such as stroke.

- [3] **Vito Dichio**, Hong-Li Zeng, and Erik Aurell. *Statistical genetics in and out of quasi-linkage equilibrium*. In: Reports on Progress in Physics 86.5 (2023), p. 052601

Pocket abstract: Statistical genetics applies the principles of statistical physics to the study of evolution. In this review, we focus on quasi-linkage equilibrium (QLE) regime, which exists when mutations and recombinations happen at a sufficiently fast rate with respect to the strength of natural selection. Under QLE conditions, it is possible to infer epistatic interactions between genes. We also review recent investigations of the loss of a QLE state, e.g., the transition to a clonal competition state. We introduce and characterize a new transition from QLE to a non-random coexistence (NRC) regime.

The results of this thesis have been further disseminated as contributed talks at the following international scientific conferences – in inverse chronological order.

- *Statphys28*, Tokyo, Japan, August 2023
- *NetSci 2023*, Vienna, Austria, July 2023
- *APS March Meeting*, Las Vegas, USA, March 2023
- *CCS 2022*, Palma de Mallorca, Spain, October 2022



In addition to the principal subject of this thesis, I have been involved in the following two projects during my PhD. As they are outside the interest and scope of this manuscript, they will not be discussed further.

- [4] Hong-Li Zeng, Yue Liu, **Vito Dichio**, and Erik Aurell *Temporal epistasis inference from more than 3 500 000 SARS-CoV-2 genomic sequences*. In: Physical Review E 106.4 (2022), p. 044409

Pocket abstract: Building on our previous work, we use direct coupling analysis (DCA) to determine epistatic interactions between loci of the SARS-CoV-2 virus. Genomes are grouped by month of sampling (up to October 2021). We find that DCA terms are more stable over time than correlations, but still change over time as mutations disappear from the global population or reach fixation. We identify putative epistatic interaction mutations involving loci in the genomic region encoding the spike protein.

- [5] Hong-Li Zeng, Eugenio Mauri, **Vito Dichio**, Rémi Monasson, Simona Cocco and Erik Aurell *Inferring epistasis from genomic data with comparable mutation and outcrossing rate*. In: Journal of Statistical Mechanics: Theory and Experiment 2021.8 (2021), p. 08350

Pocket abstract: We address the problem of inferring the epistatic fitness from the evolutionary dynamics of a population, under quasi-linkage equilibrium conditions. We extend current state-of-the-art methods and build on a recently proposed technique that uses a Gaussian approximation for the genome probability distribution. We validate the results with *in silico* experiments.

Introduction

1

*Homo liber nulla de re minus quam de morte cogitat, et eius
sapientia non mortis sed vitae meditatio est*¹.

— Baruch Spinoza

1: Tr.: There is nothing over which a free man ponders less than death; his wisdom is to meditate not on death but on life.

At a typical physics conference, the parallel sessions dedicated to biological physics are often a sight to behold, as the most spectacular phenomena appear on the blackboard, *pardon*, on the screen. On one side of the room, the presenter discusses the collective behaviour of insect swarms, or the evolutionary dynamics in experimental populations of *E. coli*, or the electrical activity of individual neurons in the mouse hippocampus, or the layer formation in bacterial colonies [6–9] (...)

If we now turn our gaze to the other side of the room, we observe a diverse crowd of scientists. Some just passing by, perhaps taking a classical break from their quantum session, or taking a breath from a deadly series of talks on the latest theory of the universe in a few dozens of dimensions. However, most of the people in the room would probably describe their research field as biophysics, or physics of life, or biological physics². The questions are: what exactly is biological physics? What are these physicists looking for in biology?

Despite the long history of the subject [10–12], it is only recently that the community has organised itself, and only recently that biological physics has been recognised as a genuine, distinct sub-discipline of physics³ [13]. Perhaps because of this, it remains a somewhat nebulous or quirky field for many, and there are some misconceptions and misunderstandings about the subject that I will try to clear up briefly before delving into the contents of this manuscript⁴.

2: Identikit: their interest in biological problems arose very late in their university education, or even later. Their natural home is a physics department, but they are often found elsewhere. They cannot resist throwing in a reference to E. Schrodinger's *What is life?* whenever it is remotely possible, as I will do at the end of this chapter.

3: In 2022, the first decadal survey of biological physics – *Physics of Life* [13] – was published by the National Academies of the United States, a kind of historic moment. The survey is an extraordinarily rich and vivid portrait of the state of the art, including the (many) open challenges for the near future. If I may suggest: highly recommended.

4: What follows is a personal but of course not entirely original perspective on the subject. Therefore – full disclosure – let me acknowledge my main intellectual debts, the works that have molded the most my own views, as found in these first pages: [12–17]. The responsibility for any possible inaccuracy or fallacy is, of course, mine alone.

5: This very statement, as I understand it, is at the heart of what *thinking like a physicist* is supposed to mean.

Physics, biology, biological physics

There exist two broad ways in which academic disciplines define themselves: either by the object or by the style of their investigation.

As physicists, we undoubtedly belong to the second class. Physics spans the entire range of natural scales, from quarks to clusters of galaxies, and the frontiers of its exploration have been and are ever broadening. The *leitmotif* of our inquiry is the nature of the questions being formulated, and the nature of answers being pursued. In particular, we seek a parsimonious mathematical understanding of the phenomena, distilled into few general principles⁵.

(...) the physics community clings to the romantic notion that Physics is one subject. Not only is the book of Nature written in the language of mathematics, but also there is only one book, and we expect that if we really grasped its content, it could be summarized in very few pages.

W. Bialek, *Biophysics*, 2012

6: See footnote 4, chapter 1 in [14].

7: I am aware of the dangers of such a vague statement, as here I may have awakened the Cerberus of biology. Consider, for example, the level of experimental precision that geneticists can achieve with CRISPR gene editing versus the noisy and aggregated measurements of large-scale brain activity (e.g. fMRI), common in cognitive neuroscience [18, 19]. These techniques allow for equally interesting but fundamentally different questions. Indeed, the kind of theoretical questions one can meaningfully ask is constrained by the kind of experimental answers one can get.

8: There are many possible starting points. For example, staying alive involves solving a number of highly non-trivial physics problems (sensing the environment, navigating in physical space, converting energy...). Therefore, one possible line of investigation is to ask what are the physical problems that living systems have to solve. Others: how do living systems represent information? how do functions emerge from microscopic components? how do systems navigate parameter space? what are the physical limits of biological processes? how did life begin from a soup of molecules? [13]

9: An early view of biological physics regarded it as an *application* of the tools of physics to the problems of biology [10]. Today we find this view limiting, as it overlooks what is perhaps the most intellectually stimulating direction.

This, we strive to achieve by a tight dialogue between experiment and theory. Through the former, we question Nature, yearning for clues or for verdicts. Through the latter, we draw an understanding from what has been seen, and prescribe what ought to be seen⁶. If the approximate reasoning is granted – and often deemed necessary –, we still insist on the quantitative agreement between theory and experiment, between predictions and numerical facts about the world.

The character of the scientific enterprise is quite different in biology, which definitely belongs to the first class. In fact, biology is defined – Greek vocabulary in hands – as the study of living systems. Not only biology, but the also its many branches are strongly tied to the specific piece of the natural (biological) world they study. So that an ecologist, a geneticist and a cognitive neuroscientist may have very little in common, not only in terms of the system of interest, but also in terms of the nature of the questions that are formulated and the answers that are reasonably within reach⁷. Moreover, the vast majority of questions in biology are still investigated almost exclusively experimentally, so that theory is a much more unequal partner to experiment than in physics.

Where does biological physics fits within this picture? By now, the answer should be evident. Biological physics is the investigation *à la* physicist of the biological phenomena. The agenda is (i) identify the general principles that govern the phenomena of life, (ii) articulate them in a mathematical language and (iii) make quantitatively accurate predictions in agreement with experimental data.

By their very nature, the principles we seek should transcend the details of this or that particular system. Even more, they are expected to intersect with and manifest in several of the standard sub-disciplines of biology, and to cast a variety of seemingly disparate biological problems into a single, more fundamental physics problem⁸. In articulating principles, we borrow the formal and conceptual tools of statistical physics and information theory, but also mechanics and thermodynamics. Finally, for a genuine biological physics to deliver on its promises, we shall pursue in biology the same level of quantitative agreement with data that is standard in other physics domains.

It is important to emphasise that in biological physics, the symbiosis between physics and biology leads to an enrichment in both directions. However, one of them has only recently been fully appreciated⁹. As Stan Ulam said once, *ask not what physics can do for biology – ask what biology can do for physics* [20].

There is clearly something unique about the state of matter we call life, that has no equivalent in the physics of conventional matter. It is not a new force of nature that we are missing, the very carbon atoms and interactions that constitute the pencil I write with also form the neurons that guide my hand. What we are instead missing is a precise understanding of how evolution, adaptation and learning have shaped my own brain over very different time scales, so that I can now write about them. The enterprise of explaining these three processes exclusively in terms of standard condensed matter physics is doomed to failure. This because they are all related to the notion of *biological function*, which is essential for life but foreign to standard physics. Its centrality in biology cannot be overstated.

This to say, biology is not merely a playground for our physics tools. There is a new physics to be learnt from living systems, and this is the enterprise that we, as theorists, as biological physicists, are committed to.

From the principles to models

To carry out the programme of biological physics, we need to project general and abstract statements about the physics of life (principles) into models of real biological systems. In other words, we need to *make them work*. The path is anything but straightforward.

While a principle manifests across diverse problems, it is essential to tackle each of them individually, zeroing in on a specific problem, or *context*¹⁰. When we do so, the semantics of our statements translate into the language of mathematics and are rendered as a set of equations. This requires us to specify a formal representation for the biological systems under study¹¹ – which is in general context dependent. The equations have a number of parameters, whose biological interpretation is again context-bound. At this juncture, the details of the equations – including the values of the parameters – are unspecified¹².

To set them, the analysis must be further narrowed down to a particular system¹³. Only at this granular level is a model defined. Each and every (biological) physicist repeats the same mantra over and over again¹⁴

General principles, particular models. General principles, particular models. General principles, particular models (...)

Modelling may involve the formulation of additional theoretical assumptions or the setting of parameter values. Both of these endeavours benefit directly from the data we have on hand. They help us not only to trim our theoretical picture of the system, but also to fix (infer) the values of the parameters of the theory.

The task of modelling is full of nuances, especially when it comes to biological systems. Therefore, let me comment briefly on a few aspects.

Simple systems

The exploration of a physical principle usually begins from its simplest instances. For example, quite understandably the vast majority of physics students first encounter the principle of least action in classical – rather than quantum – mechanics. The first Lagrangian written on the blackboard is likely that of the simple – rather than Kapitza – pendulum.

Starting with simple systems is not just a pragmatic approach, it is a philosophical stance on the nature of understanding. By peeling back the layers of complexity, we do find an easier access to the underlying principles, that may remain otherwise obscured. The situation is no different when it comes to the biological matter, except that *the simplest biological systems already are of jaw-dropping complexity*.

This is true no matter where we look in biology. For instance, consider cyanobacteria: among the earliest life forms, they perform oxygenic

10: For instance, think about *homeostasis*, i.e., any self-regulating process by which a biological system maintains stable properties despite perturbations. In animal physiology, it may refer to the ability of maintaining a stable internal temperature (thermal homeostasis). In ecology, it may refer to the need of keeping a stable quantity of essential nutrients for the existing species – i.e., conserving the ecological stoichiometry. These are two different contexts for the same general principle [21–23].

11: As an example, throughout this manuscript we will represent systems as strings of zeros and ones, a binary, discrete representation. See below.

12: Here is where we can tread the well-trodden path of analysis of a typical problem in theoretical physics. This involves fixing the details in a convenient way – so to make calculations simpler, or at least possible – and starting to understand the resulting *toy models* (exact solutions, approximate solutions, asymptotic behaviours...). The hope is that we can learn from them about more complex cases. I am tempted to call this stage of analysis ‘preliminary’, but there are cases where people have been stuck in it for decades (and still are). Naturally, one can always resort to simulations. This is what we do after all in theoretical physics, we solve what we can, as much as we can, and we simulate the rest.

13: Consider again the environmental homeostasis in s.n. 10. A particular system could be a freshwater pond. A nutrient runoff from adjacent land causes an algae bloom in a pond. In turn, this causes a decrease of oxygen levels, affecting aquatic life. Yet, certain bacteria and plants can absorb these excess nutrients, curbing algae growth and restoring the pond’s balance [24].

14: Beware of the opposites: particular principles, general models.

15: Which we are very much grateful for, given that it is thought to have been responsible for the rise of atmospheric oxygen O_2 some 2.5 billion years ago.

16: A matter of semantics. I am aware of my loose use of the terms of *simplicity* — and *complexity* in this section. Indeed, it is tricky to define them formally and a significant debate exists around them — which I honestly find somewhat futile. My point here is to highlight that while it is fairly easy to tell what is *the simplest* or *the most complex*, it is much harder to say what is *simple* or *complex*. For instance, everyone agrees that the *C. elegans* has one of the simplest nervous systems, yet it is much more problematic (and probably pointless) to declare the *C. elegans* brain to be simple.

17: The landscape of possible answers is actually more multifaceted, for a more in-depth discussion of this topic, see [12].

18: This independence of large-scale properties from microscopic details should not sound new to those familiar with the theoretical apparatus of statistical mechanics.

19: Does this matter? After all, a model is what it is, regardless of my philosophical stance of the underlying biology. Well, in my opinion, it does matter, as it is tough to navigate the ocean without a star to steer by.

photosynthesis via an intricate molecular apparatus¹⁵. The roundworm *C. elegans*, despite the misfortune of having one of the smallest nervous systems known (302 neurons), is capable of locomotion, mating, chemosensation and more. *Archaea*, single-celled microorganisms, possess specialized membranes, enabling them to flourish in the most hostile environments [25–27]. My census of "simple" biological systems could continue: complexity is a ubiquitous and perhaps necessary feature of living systems.

This indubitably makes our job as physicists more challenging, and we should be more vigilant than ever before about the pitfalls that lie in wait. Yet, I do not intend to dishearten the reader, on the contrary. It is precisely the intrinsic complexity¹⁶ of living systems that makes the whole enterprise of biological physics so magnetic and, ultimately, rewarding.

Parameters

Our models have parameters. It is common sense that the more realistic we want our model to be, the more effects, therefore parameters, we shall include. Pushed to its limit, this reasoning would suggest that a *biological truth* is attainable only in the limit of infinite parameters. So says a influential book on the topic [17]:

We believe that "truth" (full reality) in the biological sciences has essentially infinite dimension (...) It is generally a mistake to believe that there is a simple "true model" in the biological sciences and that during data analysis this model can be uncovered and its parameters estimated.

*K. P. Burnham and D. R. Anderson,
Model selection and inference, 1998*

I do think that this point of view misses the focus of what we are trying to do. If by truth is meant the account of every possible fine-scale detail of a biological system then truth is unattainable, therefore uninteresting. On the contrary, it is very much interesting to ask: are all details *really* necessary?

Two contrasting pictures are the following¹⁷. One possibility is that (almost) all details are indeed necessary, and the multiplicity, or irreducibility, of parameters is hence an intrinsic characteristic of biological systems. This would spell doom for the physicists' dream of an understanding of the life phenomena in terms of a handful of principles (and parameters). An opposing view goes something like this. The vast majority of the microscopic details of a biological system are irrelevant, since the system functions are "robust" properties of the model, independent of the configuration of those details¹⁸. Instead, the relevant features of the system are controlled by (a few) parameters, that are fine-tuned on evolutionary scales by natural selection. It is from this tension between robustness of the functional outcomes and optimal tuning, that life unfolds.

There is no need for me to say where I stand between the two, the reader has already guessed¹⁹.

Data

In recent decades, most scientific fields have experienced an exponential surge in data volume²⁰, and particularly so in biology. Genomic sequencing now enables researchers to determine millions of DNA sequence reads in a single run, spanning from viruses to the entirety of the human genome. High-throughput mass-spectrometry churns out extensive datasets about protein composition and structure. Serial-section electron microscopy offers detailed three-dimensional reconstructions of an ever-growing number of natural brains [28–31] (...)

This has sparked a widespread data-centric enthusiasm. Some went so far as to say that data are all you need, the end of theory has come [32]²¹. Then, there is machine-learning. By summer 2023, there is no need to extol the impressive effectiveness of black-box artificial neural networks [36, 37], we are all astonished, all amazed. The paradigm seems to be: take the largest amount of data out there, feed them to your machine-learning architecture, *et voilà*, get the most accurate predictions. There seems to be no theory here, do we really need a theory?

I think so. Making predictions about the facts of nature is an essential part of what we do as (biological) physicists. However, this is the means by which we achieve the goal of our scientific enterprise, not the goal itself. What we do is formulate hypotheses (theories) using transparent and interpretable mathematical models, based on empirical observations. Our theories make quantitative predictions, and if they accurately describe the data, then we claim to have achieved some form of understanding of the natural phenomenon. Understanding is the goal, prediction the means. Any finite amount of data will not suffice *alone* in this, in order for data to speak meaningfully, it must be meaningfully questioned²². A fundamental part of our job is and will remain to take a piece of paper, a pencil, sit down under a tree and wait for an apple to fall on our head.

This said, as theorists, we do share the data-enthusiasm. Even if data are not enough, data are definitely good. Data and theory should coexist and enrich each other: data informing theory building²³ and theory guiding data mining. It is a safe bet to predict that such a symbiotic relationship will become increasingly essential in the future of biological physics.

On the representation

There is one final point of importance that deserves further attention, the *representation* [38]. As the subject is vast and to make the discussion concrete, let me start by defining the mathematical representation of interest for most part of this manuscript.

A graph, or network²⁴, in its simplest form, is a collection of points (nodes) connected by lines (edges), fig. 1.1 [39]. A simple²⁵ graph G can be identified with a symmetric, binary matrix, with zero diagonal, i.e.,

$$G = \begin{bmatrix} 0 & & & \\ a_{21} & 0 & & \\ \vdots & \vdots & \ddots & \\ a_{N1} & a_{N2} & \dots & 0 \end{bmatrix} = G^T, \quad (1.1)$$

20: The sense of Wordsworthian sublime and dread at the same time of many scholars is conveyed by the apocalyptic vocabulary often employed: the data *deluge*, or *flood*, or *avalanche*, or *explosion* (...)

21: This is wrong in so many ways that it would be difficult to account for them all here. See [33–35].

22: Let me go further. There is a feeling that by piling on layers of artificial neurons we are not getting an inch closer to understanding what is going on. That is not necessarily good or bad, it depends on the question. I think that machine learning has finally freed theoretical physics from the anxiety of providing fast answers to (very hard) quantitative questions about biological systems. The process of our science may be slow, it may take time to disprove our wrong assumptions and identify the right ones. If you want a prompt prediction, machine-learn it. If you want to understand what is going on, ask the theorist – and be patient!

23: There is room here for all sorts of inference methods, white or at least grey boxes.

24: Here and everywhere in this text, I will use these two terms interchangeably.

25: A graph is said to be simple if its edges are undirected, unweighted (binary) and has no self-loops – the number of nodes is finite.

26: If N is the number of nodes, there are $L = N(N-1)/2$ possible edges, therefore an equivalent representation is that of a string of L binary values 00110... There are 2^L possible graphs.




Figure 1.1: A simple graph is a collection of points (nodes) and lines (edges).

27: Why should (ir)relevant variables even exist? It is very fundamental question, not exclusive of the studies of complex/biological systems. A classic and superb piece of literature on the topic is *The unreasonable effectiveness of mathematics in natural sciences* by E. Wigner [43].

28: A similar system of neurons and synapses – though not in the case of the human brain – will be a main focus of this manuscript. For this, a graph representation is certainly the simplest and most natural.

29: There is a number of unimportant details I am overlooking here. Of course, since the more fundamental representation of a brain is in terms of neurons and synapses, there should be a way of deducing the large-scale behaviours from the firing patterns of individual neurons. To my knowledge, however, we have no clue of how this could be done. Worse, a graph of the physical neural connections (which support the transfer of information) has never been drawn and is currently experimentally out of reach. For the human brain, the road is long.

30: Let me emphasise this point. In choosing a representation, we may of course make approximations that take us away from the "true" system. But when we do, we (hopefully) know what we are leaving out, what we are including, and we have strong theoretical control over what is going on in our model.

where $a_{ij} \in \{0, 1\}$ indicates the absence or presence of an edge²⁶ within the dyad (ij) – i.e., between the nodes i, j .

Representing a system as a graph has proved to be a valuable theoretical tool for the analysis of complex systems – including biological systems [39, 40]. Despite (and perhaps because of) its widespread adoption, there are periodic cries of alarm from the community warning of its potential misuse [41, 42].

The first (maybe trivial) point is to fully acknowledge the fundamental difference between the system, the data and their representation. The data are what we measure or observe about a system, an empirical fact. The representation is the way in which we represent the system, an abstraction. The construction or choice of a representation is therefore a genuinely theoretical act, involving assumptions about the class of systems that we are studying and the data we have on hand. A crucial one regards the choice of relevant variables, which is a foundational issue of any scientific approach²⁷. It is also important to say that there exist no intrinsically *correct* representation for any data or systems, as it depends on the answers being sought.

When it comes to networks, this translates to two fundamental questions: *what are the nodes? what are the edges?* Consider the human brain [44]. Down to the cellular scale, $\sim 10^{11}$ individual neurons (nodes) and their synapses (edges) form an intricate web of connections, which naturally lend themselves to a network representation²⁸. At a much larger scale, one of the most common experimental techniques, an electroencephalogram (EEG) allows to measure electrical activity in the brain using around a hundred of sensors (nodes) attached to the scalp. These signals can be then correlated, yielding a measure of functional relatedness (edges) between different brain regions. These two networks describe the same biological systems, but they differ substantially in terms of *what* they represent – neuron vs brain regions, direct and physical vs indirect and functional interactions²⁹.

The more artificial our definitions of *nodes* and *edges*, the greater the chance of introducing spurious effects over which we have no control. The approach we have taken in this work has been to focus on systems for which the path from the system to data and their representation is as short as possible and under theoretical control. For a genuinely theoretical approach, the latter is a *conditio sine qua non*³⁰. It is a practical matter, too. If the model works, we declare victory and rejoice of it. If it does not, whose fault is it? Is it a fallacy of the model? Or of the data we are using? Or of uncontrolled spurious biases in the representation?

As for me, if I fail, whenever I fail, every time I fail, I would like to take full responsibility for my own failure.

Overview of the manuscript

The body of this manuscript is structured into four chapters. My aim has been to prioritise a coherent and continuous flow of ideas from one to another and to reflect the historical development of the project. At the heart of this manuscript, the principle formulated in ch. 4 and

implemented in ch. 5. Yet, scientific ideas seldom emerge from nowhere, the one at the heart of this thesis was definitely not an *Eureka!* moment, but rather a gradual sedimentation of intuitions³¹. In ch. 2 and ch. 3 I therefore briefly lay out the methodological and theoretical ground from which later chapters have sprung. More specifically:

- **Chapter 2** introduces the task of inference from network data in the context of exponential random graph (ERG) models. The discussion is organised in two parts. The first deals with general theoretical aspects by placing the ERG approach in the larger context of maximum entropy inference. The second covers a range of methodological issues that arise in practical applications. The inference method introduced here find application in ch. 5.
- **Chapter 3** delves into evolutionary dynamics, starting with a concise overview of key concepts in evolutionary biology. Later, a recently proposed model of multilocus evolution is presented and its salient features discussed. The chapter concludes by examining extensions to genetic algorithms. The concepts introduced here constitute a theoretical background for ch. 4.
- **Chapter 4** introduces the exploration-exploitation (EE) paradigm. First, the rationale is thoroughly discussed and the basic formalism is established. As a first step, the theory is solved for a set of toy models, from which general conclusions can be drawn. Finally, a simulation framework for the theory is discussed and tested against the analytical solutions.
- **Chapter 5** deals with modelling the growth of the *C. elegans* brain using the EE framework. This system is naturally represented as a graph. The chapter begins with an illustration of the brain wiring problem and an overview of the *C. elegans* nervous system. A model of its development (from birth to adulthood) is formulated and made to work. The chapter concludes with a discussion of a biological mechanistic interpretation of the model, and an detailed outline of potential extensions.
- **Chapter 6** concludes this manuscript by summarising the lines of investigation followed in this project and and the main findings³².

This manuscript is essentially the story of a scientific idea, from what it has blossomed, how it has grown, what it might become. As E. Schrodinger wrote once, *I do not know whether my way of approach is really the best and simplest. But, in short, it was mine (...)* And I could not find any better or clearer way towards the goal than my own crooked one [45].

31: Many incorrect (not shown), but some, or at least one, worth pursuing.

32: An appendix is attached at the end of the manuscript. Appendix A presents mathematical details of the toy models in ch. 4. Appendix B discusses the set of network measures used in ch. 5. Finally, a Glossary provides brief definitions of key biological concepts discussed in the manuscript.

Maxent inference for graphs

Those who ignore statistics are condemned to reinvent it.

—Bradley Efron

2.1 ERG models 9

2.2 User handbook 16

The core problem in statistical inference is to recover the parameters $\chi \in \mathbb{R}^r$ of a statistical model

$$P(D; \chi) : \mathcal{D} \times \mathbb{R}^r \mapsto [0, 1] \quad (2.1)$$

from the data $D^* \in \mathcal{D}$, assuming that $D^* \sim P(D; \chi)$. A perfect reconstruction of χ from finite data is neither possible in theory nor in practice. Therefore, we rather seek a good approximation χ^* , close to the *true* values. The two fundamental tasks of statistical modelling thus consist in specifying (i) the model $P(D; \chi)$ and (ii) how to infer χ^* from D^* [46].

There is more than one interpretation for the notion of probability in (2.1) and, by consequence, for the statistical uncertainty that follows from it. One, phenomenological, interprets the uncertainty as the empirical variation in the data-generating process. Another, epistemological, interprets it as our uncertainty about the outcome, arising from the limited information at our disposal. In some fortunate cases, as for statistical mechanics, the two overlap. Regardless of the interpretation, however, we distinguish the notion of statistical uncertainty from that of measurement error: here, the data are assumed to be noiseless representations of the underlying system.

The present chapter focuses on the statistical inference based on *exponential random graph* (ERG) models. The data consist of unweighted, undirected graphs G (1.1). The starting point is an exponential, maxent distribution (sec. 2.1) [47–49]. The role of the modeller is to specify a model within the ERG framework and draw conclusions about the data from the inferred parameters (sec. 2.2).

A minimal version of the ERG modelling is considered here, for the sake of clarity. A number of generalisations of the present framework have been proposed, a survey of which is beyond scope of this manuscript, we refer to [50] for an entry point.

Main reference

📖 Vito Dichio, Fabrizio De Vico Fallani (2022). *Statistical models of complex brain networks: a maximum entropy approach*. In: Reports on Progress in Physics 86.10 (2023), p. 102601 [2].

2.1 Exponential random graph models

Where a streamlined theoretical and methodological minimum of ERG models is provided. Some emphasis is placed on the philosophy of the approach – often overlooked in the literature and the source of a number of misconceptions.

1: A very large family, indeed. The Bernoulli, Poisson, Gaussian, binomial, multinomial, Boltzmann, Rayleigh (...) distributions all belong to it. See [51, 52] for a taxonomy.

2: The discussion here will be general, focusing on methodological aspects. We refer to [2, 61, 62] for recent reviews of the applications of ERG models, including those to social sciences, economics, neuroscience (...)

3: It was formulated for the first time in the 1957 by Edwin Thompson Jaynes (1922-1998) [63, 64]. Jaynes dedicated much of his career to advocating for the principle of maximum entropy as a fundamental tool for statistical inference. He held a strong belief in Bayesian probability and often defended its interpretation as an extension of logic [65].

4: The maxent derivation here illustrated is nowadays standard. However, it is not the way the ERG models were originally introduced. Instead, the original formulation was based on the Hammersley-Clifford theorem for Markov graphs [59], and built on a previous work of J. Besag in the context of spatial models of lattice systems [67].

5: The argument here is deliberately qualitative, to avoid slowing down the discussion. It can be made more quantitative, though. See for instance [65, 66].

The interest in statistical exponential families dates back to the dawn of modern statistics¹ [53–55]. A number of mathematical properties makes them particularly apt for purposes of statistical inference [56, 57]. In the context of graph they burst onto the scene in the 1980s, mainly motivated by the study of social network interactions [58–60]. Later on, they attracted the attention of the physics community, encouraged by their formal similarity to the well-developed theory of classical statistical mechanics [47, 48].

When tailored to graphs, exponential distributions are referred to as *exponential random graph* (ERG) models and in this section we provide a minimalist theoretical introduction to them².

2.1.1 A maximum entropy approach

Let $G^* \in \mathcal{G}$ be an observed graph (data). Let us assume that all relevant information about the data can be reduced to a vector of *statistics* $\mathbf{x}(G^*) \in \mathbb{R}^r$. We postpone the discussion on the choice of the $\mathbf{x} : \mathcal{G} \mapsto \mathbb{R}^r$ (model selection problem) to sec. 2.2.2 and, until then, we consider it as given.

According to the *maximum entropy (maxent) principle*³ [63, 64, 66] the most unbiased model of the data, consistent with the current state of knowledge, is found by maximising the Shannon entropy

$$H(P) = - \sum_{G \in \mathcal{G}} P(G) \log P(G) \quad (2.2)$$

subject to the normalization $\sum_{G \in \mathcal{G}} P(G) = 1$ and to the (soft) constraints:

$$\sum_{G \in \mathcal{G}} \mathbf{x}(G)P(G) = \mathbf{x}(G^*) . \quad (2.3)$$

This constrained maximization problem is easily solved with the method of Lagrange multipliers [2, 48] and yields:

$$P(G|\boldsymbol{\theta}^*) = \frac{e^{\boldsymbol{\theta}^* \cdot \mathbf{x}(G)}}{\sum_{\tilde{G} \in \mathcal{G}} e^{\boldsymbol{\theta}^* \cdot \mathbf{x}(\tilde{G})}} , \quad (2.4)$$

where the parameters $\boldsymbol{\theta}^* \in \mathbb{R}^r$ are set so to satisfy (2.3). We refer to (2.4) as the ERG model of the data⁴. Before proceeding further, let us elucidate some fundamental underpinnings of the maxent approach.

On the rationale

Intuitively, the Shannon entropy is associated to the uncertainty of a random variable [68, 69].

An unconstrained maximisation of (2.2), subject only to the normalisation, would yield a flat distribution where each possible graph has probability 2^{-L} . This corresponds to the case where no information is encoded in the probability distribution. On the contrary, maximising (2.2) subject to the constraints (2.3) yields the distribution in which no information *other* than that contained in the constraints is taken into account. In this sense, the maxent distribution (2.4) is the most *unbiased*.⁵

The maxent simply prescribes the optimal approach to integrate any prior knowledge about the system – here, observed statistics – into the probability distribution. In the words of E. T. Jaynes [65]:

The information available defines constraints fixing some properties of the initial probability distribution, but not all of them. The ambiguity remaining is to be resolved by the policy of honesty; frankly acknowledging the full extent of its ignorance by taking into account all possibilities allowed by its knowledge.

E.T. Jaynes, Probability Theory - The Logic of Science, 2003

An important point is the following. In presenting the maxent method we have carefully avoided referring to $x(G^*)$ as *sufficient statistics*. They are, but the matter is more subtle.

By definition, a statistic $x(G)$ is said to be *sufficient* for the model $P(G|\theta)$ with parameters θ , if and only if the data reduction $x : \mathcal{G} \mapsto \mathbb{R}^r$ implies no information loss⁶. This is true in the case of the statistics $x(G)$ and the ERG model (2.4), since the data G appear in the distribution only through the statistics $x(G)$. But this is so by design, as a result the maxent construction: any choice of $x(G)$ would be sufficient for the resulting ERG model. In other words, the notion of sufficient statistics is determined a priori, as an hypothesis, rather assessed a posteriori, as a property of the distribution. We are therefore led to concede that different modellers with different amounts of information about the physical system will come up with different ERG distributions, leading to different predictions. The vast majority of these models, presumably, will be wrong.

The above is somehow bewildering if we embrace an *orthodox* school of thought, for which probabilities are long run frequencies of repeated experiments [70]. In this latter case in fact, we clearly would not want probabilities to depend on the state of knowledge of the modeller. This is indeed the major source of criticism to the maxent modelling approach [71, 72]: the nature out there remains indifferent to our knowledge or lack thereof. The argument is evidently true, but it misses the point. Again in the words of E.T. Jaynes [65]:

The principle of maximum entropy is not an oracle telling which predictions must be right; it is a rule for inductive reasoning that tells us which predictions *are most strongly indicated by our present information*.

E.T. Jaynes, Probability Theory - The Logic of Science, 2003

Here, differently from the orthodox view, probabilities are considered as epistemic statements, informed guesses on a phenomenon. It can be proved that if the information included in the maxent formulation encompasses all *relevant* constraints operating in a system, then the maxent distribution is the overwhelmingly most likely to be observed experimentally. What if, instead, the observations disagree with the predictions of the maxent model? For the maxent modeller, this is not a cause for embarrassment. It simply hints at the presence of additional or different constraints in the systems that have not yet been accounted for.

6: More precisely, let $P(D|\chi)$ be a parametric distribution for the data D , with parameters χ . Let $t(D)$ be any statistic of the data. According to the data processing inequality,

$$I(\chi; t(D)) \leq I(\chi; D),$$

where I is the mutual information [69]. In words, any manipulation of the data D can either reduce or preserve the information about the parameters χ . In this latter case, we call $t(D)$ sufficient statistics.

On the similarity with statistical mechanics

The reader conversant with statistical mechanics (SM) will readily recognise (2.4) as analogous to a Gibbs-Boltzmann distribution, i.e.,

$$P(G|\boldsymbol{\theta}) = \frac{e^{-\mathcal{H}(G,\boldsymbol{\theta})}}{\mathcal{Z}(\boldsymbol{\theta})}, \quad (2.5)$$

where $\mathcal{H}(G, \boldsymbol{\theta}) = -\boldsymbol{\theta} \cdot \mathbf{x}(G)$ is the (graph) Hamiltonian and

$$\mathcal{Z}(\boldsymbol{\theta}) = \sum_{\tilde{G} \in \mathcal{G}} e^{-\mathcal{H}(\tilde{G}, \boldsymbol{\theta})} \quad (2.6)$$

is the partition function [47, 48]. A Gibbs-Boltzmann distribution (canonical ensemble distribution) describes the statistical behavior of particles in a thermodynamic system at equilibrium. It is known from SM that all sorts of observables can be computed by differentiation from the free energy [73, 74]

$$\mathcal{F}(\boldsymbol{\theta}) = -\log \mathcal{Z}(\boldsymbol{\theta}). \quad (2.7)$$

7: By construction, the ensemble averages of the statistics match the experimental values $\langle x_\alpha \rangle = x_\alpha^*$, cf. (2.3).

For instance, the expected value of the α -th statistic⁷:

$$-\frac{\partial \mathcal{F}(\boldsymbol{\theta})}{\partial \theta_\alpha} = \frac{1}{\mathcal{Z}(\boldsymbol{\theta})} \sum_{\tilde{G}} x_\alpha(\tilde{G}) e^{-\mathcal{H}(\tilde{G}, \boldsymbol{\theta})} = \langle x_\alpha \rangle, \quad (2.8)$$

where we have used (2.4) and introduced the shorthand

$$\langle O \rangle = \sum_{\tilde{G} \in \mathcal{G}} O(\tilde{G}) P(\tilde{G}) \quad (2.9)$$

for a graph observable $O : \mathcal{G} \mapsto \mathbb{R}$. The formal analogy with (2.5) is powerful because it allows a number of results and methods from over a hundred years of SM to be translated directly into the ERG context [48, 75].

It is certainly not a stroke of luck. In two groundbreaking papers published in 1957, E. T. Jaynes demonstrated that, considering SM as a form of statistical inference, the Gibbs-Boltzmann distribution can be derived directly from the maxent principle [63, 64]. Indeed, when spoiled from its physical interpretation, the mathematical structure of the computations of SM turns out to be a general property the maxent formalism [65].⁸ There is, however, a crucial *caveat*.

8: This is the ultimate reason for the existence of so many journal articles in the literature with titles of the form "Statistical mechanics of (something else)."

SM is more than a statistical theory, it is a physical theory, it agrees with experiments, it *works*. As discussed in the previous section, the maxent argument is independent of any experimental verification. The very reason for the experimental success of SM, viewed as a maxent model, is that its choice of sufficient statistics – notably, the energy of a microstate – is *correct* for a thermodynamic system at equilibrium. The latter result is peculiar to the case of SM and does not generalise. In summary, when formulating a maxent model, we are allowed to borrow the formal structure of SM, but not (in general) its interpretation as a physical process.

On the purposes

We have recently argued that ERG models are certainly descriptive, might be predictive and they are not explicative [2].

The latter is a straightforward consequence of the discussion above. ERG models are agnostic about the data-generating process therefore, in general, they cannot answer to the question why does a phenomenon happen⁹ [76].

Prediction is here intended as feature generalisation, i.e., as the ability of a ERG model to predict the values of different statistics of the data $y_\alpha(G)$, $\alpha = 1, \dots, s$ than those used in the specification of the model. This is in line with similar maxent approaches – e.g., [78, 79]. Clearly, the potential predictive power of the ERG model is inherently linked to the selection of statistics, as they are the only means by which an ERG model is informed about the system. In the case where our hypothesis x was accurate, the resulting model would be capable of predicting *any other* test statistics y . In practice, this is very seldom the case. Nonetheless, ERG models constructed with incomplete information can still demonstrate strong predictive performance on specific subsets of test statistics.

Regardless of their predictive power, ERG models remain inherently descriptive. They provide a framework to characterise a system (read, compute observables), based on any hypothesis about the sufficient statistics. This makes them ideal for constructing null models [47]. In line with the view of the maxent principle as a *rule for inductive reasoning* (see above), ERG-based null models can always be used to lower bound the complexity of the *true* model. Furthermore, experimental deviations from the null predictions may contain useful information about the system, and suggest possible theoretical refinements.

2.1.2 Model inference

In deriving the ERG distribution (2.4) we have implicitly stated an inference (or inverse¹⁰) problem [81]. Let us highlight it.

Definition 2.1 (ERG inference) *Given a set of observations $x(G^*) \in \mathbb{R}^r$, the ERG inference consists in finding the parameters $\theta^* \in \mathbb{R}^r$ such that the constraints (2.3) are met.*

It is instructive to start by considering a simple solvable case.

Bernoulli random graphs

Let us consider an ERG model whose only statistic is the number of edges in the graph, i.e., $x(G) = \sum_{i<j} a_{ij}$ [48]. Given a graph G^* , with $x(G^*)$ edges, the goal of the ERG inference is to find the parameter θ^* such that $\langle x \rangle = x(G^*)$. The inference problem can be solved in two steps.

9: Here we adopt a classical (strict) definition of scientific explanation, articulated in a highly influential paper (1948) by Carl Hempel, and particularly close to the spirit of physics – (...) *the question "Why does the phenomenon happen?" is construed as meaning "according to what general laws, and by virtue of what antecedent conditions does the phenomenon occur?"* [76]. Nevertheless, this is not at all the only conceivable definition, as the subject has been extensively debated in the philosophy of science, see [77].

10: Given a model with known parameters, the *forward* or *direct* problem is to compute the values of the observables (data). This is the case, e.g., in statistical mechanics [74]. Conversely, given a set of observables (data), the *inverse* problem is to infer the unknown values of the model parameters. This is the case, e.g., for the inverse Ising problem [80].

First, we solve the forward problem, i.e., we express $\langle x \rangle$ as a function of θ . We start by evaluating (2.6):

$$\mathcal{L}(\theta) = \sum_{G \in \mathcal{G}} e^{\theta \sum_{i < j} a_{ij}} = \prod_{i < j} \sum_{a_{ij}=0,1} e^{\theta a_{ij}} = (1 + e^{\theta})^L, \quad (2.10)$$

where $L = N(N - 1)/2$ is the number of possible edges. We can obtain a simple analytical expression for (2.7) as well, which reads

$$\mathcal{F}(\theta) = -L \log(1 + e^{\theta}), \quad (2.11)$$

therefore, using (2.8):

$$\langle x \rangle = -\frac{\partial \mathcal{F}(\theta)}{\partial \theta} = \frac{L}{1 + e^{-\theta}}, \quad (2.12)$$

which solves the forward problem.¹¹

By a simple inversion of the latter formula, we can find the value θ^* for which the constraint (2.3) is met. Imposing $x(G^*) = \langle x \rangle$, we get

$$\theta^* = \log \left[\frac{d(G^*)}{1 - d(G^*)} \right], \quad (2.14)$$

where $d(G) = x(G)/L$ is the density of a graph. The latter expression solves the inverse problem. The ERG probability distribution can be finally written as

$$P(G|\theta^*) = \frac{e^{\theta^* x(G)}}{(1 + e^{\theta^*})^L} = \prod_{i < j} \frac{e^{\theta^* a_{ij}}}{(1 + e^{\theta^*})}. \quad (2.15)$$

A general framework: MLE

The ERG inference defined above can be placed in the broader context of the maximum likelihood estimation (MLE). Given an observed graph G^* and a model $P(G|\theta)$, according to the maximum likelihood principle, the best choice of the unknown parameters is given by

$$\theta^* = \arg \max_{\theta} \log P(G^*|\theta), \quad (2.16)$$

where $P(G^*|\theta)$ is the *likelihood* of the data, given the parameters θ ¹².

It is easy to show that the ERG inference problem can be derived from the maximum likelihood principle. Introducing the notation $\mathcal{L}(\theta) = \log P(G^*|\theta)$, the r equations (2.3) are found by setting to zero the derivatives with respect to the parameters:

$$0 = \frac{\partial \mathcal{L}(\theta)}{\partial \theta_{\alpha}} = \frac{\partial}{\partial \theta_{\alpha}} \left[\theta \cdot x(G^*) - \log \mathcal{L}(\theta) \right] \stackrel{(2.8)}{=} x_{\alpha}(G^*) - \langle x_{\alpha} \rangle. \quad (2.18)$$

There is of course an elephant in the room of this discussion, hidden in (2.16). In order to evaluate $P(G^*|\theta^*)$, we need to compute $\log \mathcal{L}(\theta)$, or $\mathcal{F}(\theta)$. In the case of the Bernoulli graphs, this could be done analytically (2.11), by virtue of the utmost simplicity of the model. In general, the

11: By defining $p = 1/(1 + e^{-\theta})$, we can rewrite the ERG probability distribution (2.4) as

$$P(G|\theta) = \frac{e^{\theta x(G)}}{(1 + e^{\theta})^L} \quad (2.13)$$

$$= p^{x(G)} (1 - p)^{L - x(G)},$$

which is the probability of a Bernoulli graph where each of the L possible edges appears independently with probability p , hence the name. Note also from (2.12) that $\langle x \rangle = Lp$, as it should be [82, 83].

12: The MLE estimator (2.16) can be derived from the Bayes theorem [81]. Accordingly, the posterior distribution $P(\theta|G^*)$ – which represents our knowledge after taking into account the information in the data – can be expressed as

$$P(\theta|G^*) = \frac{P(G^*|\theta)P(\theta)}{P(G^*)}, \quad (2.17)$$

where $P(\theta)$ is our prior information on the parameters. Our best choice for the parameters is the one that maximises the posterior distribution above. If $P(\theta)$ is a uniform distribution in the parameters space (no prior information available) this is the same as maximising the likelihood $P(G^*|\theta)$. The estimator (2.16) has a number of appealing properties, in particular it converges in probability to the true values (consistency) and reaches the Cramér-Rao lower bound (efficiency) [84].

computation of $\mathcal{F}(\boldsymbol{\theta})$ is an *extremely* difficult problem, well-known in statistical mechanics¹³.

In the vast majority of the cases, there is little choice but to resort to numerical approximations for (2.16), which we will now briefly describe.¹⁴

MCMC-MLE

The fundamental idea to circumvent the explicit evaluation of $\log \mathcal{L}(\boldsymbol{\theta})$ was introduced in the early 1990s [90]. Let us consider an arbitrary vector of parameters $\boldsymbol{\theta}_0 \in \mathbb{R}^r$, we can formally rewrite $\mathcal{L}(\boldsymbol{\theta})$ as

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}) &= \mathcal{L}(\boldsymbol{\theta}_0) \sum_{\tilde{G} \in \mathcal{G}} e^{(\boldsymbol{\theta} - \boldsymbol{\theta}_0) \cdot \mathbf{x}(\tilde{G})} \frac{1}{\mathcal{L}(\boldsymbol{\theta}_0)} e^{\boldsymbol{\theta}_0 \cdot \mathbf{x}(\tilde{G})} \\ &= \mathcal{L}(\boldsymbol{\theta}_0) \langle e^{(\boldsymbol{\theta} - \boldsymbol{\theta}_0) \cdot \mathbf{x}} \rangle_{\boldsymbol{\theta}_0} \end{aligned} \quad (2.19)$$

where the subscript $\langle \cdot \rangle_{\boldsymbol{\theta}_0}$ indicates the expectation value over the distribution $P(G|\boldsymbol{\theta}_0)$. The trick is now is to use the a Markov Chain Monte Carlo sampling¹⁵ to evaluate approximately the right hand side of (2.19). In particular, given a sample of m graphs G_1, \dots, G_m whose stationary distribution is $P(G|\boldsymbol{\theta}_0)$, we can approximate

$$\frac{\mathcal{L}(\boldsymbol{\theta})}{\mathcal{L}(\boldsymbol{\theta}_0)} = \langle e^{(\boldsymbol{\theta} - \boldsymbol{\theta}_0) \cdot \mathbf{x}} \rangle_{\boldsymbol{\theta}_0} \simeq \frac{1}{m} \sum_{i=1}^m e^{(\boldsymbol{\theta} - \boldsymbol{\theta}_0) \cdot \mathbf{x}(G_i)}. \quad (2.20)$$

We now consider the log-likelihood and note that the argument that maximises $\mathcal{L}(\boldsymbol{\theta})$, maximises the shifted log-likelihood $\tilde{\mathcal{L}}(\boldsymbol{\theta}) = \mathcal{L}(\boldsymbol{\theta}) - \mathcal{L}(\boldsymbol{\theta}_0)$, too. The latter, however, can be numerically approximated:

$$\begin{aligned} \tilde{\mathcal{L}}(\boldsymbol{\theta}) &= (\boldsymbol{\theta} - \boldsymbol{\theta}_0) \cdot \mathbf{x}(G^*) - \log \left\{ \frac{\mathcal{L}(\boldsymbol{\theta})}{\mathcal{L}(\boldsymbol{\theta}_0)} \right\} \\ &\stackrel{(2.20)}{\simeq} (\boldsymbol{\theta} - \boldsymbol{\theta}_0) \cdot \mathbf{x}(G^*) - \log \left\{ \frac{1}{m} \sum_{i=1}^m e^{(\boldsymbol{\theta} - \boldsymbol{\theta}_0) \cdot \mathbf{x}(G_i)} \right\}. \end{aligned} \quad (2.21)$$

A parameter estimation based on the maximisation of (2.21) converges to the same result as (2.16), in the limit $m \rightarrow \infty$ and it is used in practice as an approximation of the MLE.

Algorithm: MCMC-MLE (pseudocode)¹⁶.

```

τ ← 0
θτ ← θpl
while conv = F do
  τ += 1
  generate m graphs Gk ~ P(G|θτ-1) by MCMC
  θτ = arg maxθ [ (θ - θτ-1) · x(G*) - log [ 1/m ∑i=1m e(θ-θτ-1)·x(Gi) ] ]
  if (convergence criterion) then
    conv ← T

```

13: One might say – without fear of contradiction – that this is *the* problem of statistical mechanics. Just to give a flavor, a Nobel Prize for Physics has recently been assigned (G. Parisi, 2021) for the science that has blossomed from a *trick* in evaluating $\log \mathcal{L}$, under particular conditions [85, 86].

14: In some simple cases [87, 88], it is possible to work out the mean field theory of the model (exact in the limit of large network sizes) and possibly perform a diagrammatic expansion around the mean-field solution. However, the calculations, close to those of statistical field theory [89], soon become cumbersome and for all practical purposes, they are unworkable.

15: See also s.n. 22.

16: The algorithm starts from an initial guess of the parameters. Albeit arbitrary, an appropriate choice can aid in achieving rapid convergence. A commonly adopted approach is to use the parameters obtained through pseudo-likelihood (pl) maximization [49, 67]. The parameter space is explored iteratively by solving at each step a maximisation problem based on (2.21). The previous set of parameters serves each time as the starting point for the optimization process. Iterations continue until convergence is reached.

17: There are historical reasons for this. In fact, since the beginning [58], ERG models have been particularly popular within the domain of social sciences which, in turn, are particularly fond of R. The level of technical detail that packages such as `ergm` have reached is such that, over a typical 3-years PhD project, one would rather learn a new programming language than re-write everything from scratch. As the adage goes: *good programmers write good code; great programmers steal great code.*

18: Additional resources and examples can be found on the `statnet` website [50].

19: This is intuitive, and it is also elementary to show. First, since all dyads are independent, in the large graph limit $d(G) \sim p$, where p is the probability of having an edge between each possible dyads (connection probability). Let us focus on a single dyad. The *randomness* of the connection can be quantified by computing the Shannon entropy of the Bernoulli trial: $P_B(a_{ij} = 1) = 1$ and $P_B(a_{ij} = 0) = 1 - p$, i.e.,

$$H(P_B) = -p \log p - (1 - p) \log(1 - p),$$

which is maximum when $p = 1/2$ [69]. By consequence, we conclude that a Bernoulli graph with $d(G) = 1/2$ corresponds is the maximally random graph or, simply, random graph.

Software tools

In the last two decades, a number of open-source libraries have been developed that can be used to perform a ERG inference, as illustrated above. By far, the most popular implementation is the R-based `ergm` package, published in the 2008 by *Hunter et al.* [49] and recently updated [91]. It stands as a comprehensive tool, providing extensive functionality for fitting ERG models, including model specification, inference and diagnostics. It has served as a foundational library for numerous generalisations, an overview can be found in [50]. Unfortunately, no ERG implementation has attained an equivalent level of maturity outside the R programming language¹⁷.

The `ergm` library is the reference implementation of the ERG inference in this manuscript. In the GitHub folder `ergm_minimal` – originally presented in [2] – we have published the scripts for a minimal implementation of an ERG inference workflow¹⁸.

2.2 User handbook

Where two important and subtle aspects of the ERG method (at the beginning, at the end) are discussed in detail. The limitations of the ERG inference are highlighted, what they can(not) say, where they can(not) work.

Stripped down to the essentials, an ERG inference is a computational device that takes as input a real vector (graph statistics) and returns as output a real vector (estimated parameters), fig. 2.1.

Two key matters for practitioners are therefore (i) how to select the model statistics and (ii) how to interpret the estimated parameters. In this section, we tackle these two questions. For reasons that will be clear later, we do it in reverse order.

2.2.1 Interpretation of parameters

In sec. 2.1.1 we have argued that the ERG inference can be used for the purpose to characterise a system. Here, we quantify this qualitative statement.

Bernoulli strikes back

Let us start again by considering the ERG inference for the Bernoulli graphs (2.15), sec. 2.1.2.

As discussed above, we do have a mathematical understanding of how the data, through the ERG machinery, determine the inferred parameter, namely (2.14). In particular, if the original graph is maximally random – i.e., $d(G^*) = 1/2$ ¹⁹ – we obtain $\theta^* = 0$. If instead we start from a denser graph $d(G^*) > 1/2$, we get $\theta^* > 0$. Finally, $d^* < 1/2$ implies $\theta^* < 0$. There are two general lessons to be learnt from this: (i) the ERG model (2.15) automatically rules out model effects for which there is no evidence in the




Figure 2.1: ERG models workflow. a) A graph G^* is given, here, for instance, a brain network. b) The information of the graph G^* is condensed in a set of statistics $\mathbf{x}(G^*) \in \mathbb{R}^r$. Their choice is up to the modeller (model selection problem) and they represent the input of the ERG model inference. c) The output of an ERG inference is a vector of parameters $\theta^* \in \mathbb{R}^r$. d) The inferred parameters can be used to simulate n new networks G_1, \dots, G_n that statistically reproduce the observed one i.e. $\frac{1}{n} \sum_{i=1}^n \mathbf{x}(G_i) \sim \mathbf{x}(G^*)$.

data [65] and (ii) nonzero parameters are evidence of structure in the data, captured qualitatively by the sign of the parameter and quantitatively by its value.

However, the logic of an interpretation of the parameter θ is essentially the inverse of the one we have just illustrated. Namely, *given the result of the inference, we want to make a statement about the structure of the original data*. Two possible ways to interpret the ERG parameter θ^* are the following.

1. A *direct interpretation*, using (2.12) and (2.3):

$$d(G^*) = \frac{1}{1 + e^{-\theta^*}}. \quad (2.22)$$

2. A *dyadic (indirect) interpretation*, by computing the effect of the parameter θ^* on the probability of the existence of an edge between two nodes. Specifically, let us consider the dyad (ij) of the generic graph $G \in \mathcal{G}$, and let $P(a_{ij} = 1|\theta^*)$ be the probability of an edge between the nodes i, j ²⁰. Using (2.15) we have

$$\log \frac{P(a_{ij} = 1|\theta^*)}{P(a_{ij} = 0|\theta^*)} = \theta^*, \quad (2.23)$$

Suppose for instance $\theta^* > 0$. The probability of having an edge within the dyad ij is then higher than its complementary (absence of an edge) by a factor $\exp \theta^*$. This is a property of the ERG graph distribution. Nevertheless, we can leverage (2.3) to argue that, by ERG construction, it is also a property of the original graph, which therefore is denser-than-random.

Whenever a direct relation (2.22) is available, a dyadic interpretation (2.23) is clearly an unnecessary complication. When models become wilder than meek Bernoulli graphs, however, we soon loose the former, and the latter is all we have left.

20: In the case of a Bernoulli graph, we can treat each dyad independently. In fact, using (2.15) we have

$$\begin{aligned} P(a_{ij}|G_{\setminus ij}, \theta^*) &= \frac{P(G|\theta^*)}{P(G_{\setminus ij}|\theta^*)} \\ &= \frac{e^{\theta^* a_{ij}}}{1 + e^{\theta^*}} = P(a_{ij}|\theta^*), \end{aligned}$$

where $G_{\setminus ij}$ stands for all other dyads in the graph than ij .

The one and only general interpretation

In applications of ERG models of any interest, the vector of statistics includes several effects, that are in general mutually dependent. In such cases, as discussed in sec. 2.1.2, the inference problem cannot be solved analytically, no such equations as (2.14), (2.22) are available and no direct interpretation is possible.

The parameters θ^* are obtained by numerical approximation, are however still amenable to a dyadic interpretation. Let $P(a_{ij} = 1|G_{\setminus ij}, \theta^*)$ be the probability of an edge within the dyad (i, j) ²¹. Using (2.15) we have

$$\log \frac{P(a_{ij} = 1|G_{\setminus ij}, \theta^*)}{P(a_{ij} = 0|G_{\setminus ij}, \theta^*)} = \theta^* \cdot \Delta \mathbf{x}(G_{ij}), \quad (2.24)$$

where $\Delta \mathbf{x}(G_{ij})$ is the vector of *change statistics*. Introducing the shorthand $G_{+ij} = \{a_{ij} = 1, G_{\setminus ij}\}$ and $G_{-ij} = \{a_{ij} = 0, G_{\setminus ij}\}$, the α -th element of the change statistics is defined as:

$$\Delta x_\alpha(G_{ij}) = x_\alpha(G_{+ij}) - x_\alpha(G_{-ij}). \quad (2.25)$$

Whether or not the presence of an edge between the nodes ij is favoured is depends on the overall sign of the the right hand side of (2.24) and, therefore, on the combination of change statistics, weighted by the corresponding parameters²². We are in the position to state the following:

Definition 2.2 (ERG interpretation) *Given an ERG model (2.4), the dyadic interpretation (2.24) of the parameter θ_α is the change in the log probability of a graph, resulting from switching from G_{-ij} to G_{+ij} (i) per unit increase of the corresponding statistic $\Delta x_\alpha(G_{ij}) = 1$, and (ii) holding fixed the cumulative effect of the other statistics $\sum_{\beta \neq \alpha} \theta_\beta \Delta x_\beta(G_{ij})$.*

Once again, the ERG interpretation is based on a characterisation of the ensemble distribution $P(G|\theta^*)$, (2.4). This because, by construction (2.3), the average properties of the latter reflect those of the original graph.

Large positive (negative) values of the parameter θ_α indicate an over-(under-) representation in the original graph of the corresponding x_α , with respect to the null expectation – i.e., $\langle x_\alpha \rangle$ of an ERG model with $\theta_\alpha = 0$ and unaltered $\theta_{\setminus \alpha}$ ²³. This has an important consequence: parameters associated with the same statistic within different ERG models cannot be directly compared. A comparison is possible only if the cumulative effect of the rest of the statistics of the models are held fixed.

2.2.2 The model specification problem

The core question of the *model specification problem* (or, *feature selection problem*) [92–94] in the context of ERG models is fairly simple to state: *for a given networked system, what is the best choice of the statistics $\mathbf{x}(G)$?*

Generally speaking, the optimal choice of statistics is one that most accurately embodies our hypotheses regarding the *relevant* characteristics of the system. When using an ERG as a null model, this choice represents

21: Note that dyads are in general dependent, by consequence the probabilities are to be conditioned on the rest of the graph $G_{\setminus ij}$.

22: The scenario here considered is also at the hearth of Markov chain Monte Carlo (MCMC) methods for ERG models. The goal is to construct a Markov chain that has (2.4) as its equilibrium (stationary) distribution. To this end, at each step of the Markov chain a random change to the current graph is proposed, its effect on the (log) probability of the graph is evaluated by (2.24), and accepted or rejected based on the Metropolis-Hastings rule [49].

23: Against sloppiness. A nasty habit in the ERG literature is to interpret the parameters exclusively in qualitative terms. When it comes to something as subtle as the interpretation of ERG parameters, def. 2.24, qualitative statements alone run the risk of being too vague, or even misleading, if not false. For instance, let us consider the following common statement: " $\theta > 0$ implies that the corresponding metric is higher than expected by chance". For ERG models with multiple statistics, this is only true in precise sense of def. 2.2. *Chance* here is *the rest of the ERG model* and not a maximally random graph. A common misinterpretation of the ERG parameters indeed arises from forgetting (ii) in def. 2.2.

the null hypothesis. When using an ERG as a model of the data, this choice represents (at least) our best theoretical approximation to the *true* state space of the system [17].

All sort of graph statistics $x : \mathcal{G} \rightarrow \mathbb{R}$ can be designed. The first and simplest class of ERG statistics is the one of *edge covariates*, i.e.,

$$x_{ec}(G) = \sum_{i < j} a_{ij} \gamma_{ij} \quad (2.26)$$

where $\gamma \in \mathbb{R}^{N \times N}$ is a real matrix, with the same algebraic properties as G . The latter assigns an attribute to all dyads in a graph, by performing the sum of the values γ_{ij} over all existing edges. A simple subcase of edge covariate is obviously the number of edges in a graph – i.e., (2.26) for $\gamma_{ij} = 1 \forall i, j$ –, which we have encountered when defining the Bernoulli random graphs, sec. 2.1.2. In the case of spatially embedded graphs, the matrix γ can be used to encode the physical distance between each any nodes. Alternatively, it can be used to quantify homophily or heterophily effects²⁴ on the edge formation based on a nodal (categorical) attribute η , i.e., $\gamma_{ij} = \delta_{\eta_i, \eta_j}$, where η_i, η_j are the nodal attributes of i, j and δ is the Kroenecker delta. ERG models based on statistics of the form (2.26) are still amenable to analytical treatment²⁵. However, they are of limited interest, since the interest is often in modelling the complex interactions between dyads.

Thus, we shall turn to consider graph statistics that encode dyadic dependencies. More specifically, we restrict our attention to *Markov graphs* [60, 67], i.e., we assume that any two dyads that do not share an endpoint are independent, conditional on the rest of the graph²⁶. A large family of graph statistics for modelling dyadic dependencies falls under the umbrella of *motifs counts*, i.e., the number of times a particular connection pattern occurs in the graph. However, early numerical investigations brought to light a significant hurdle to ERG inference based on simple motif counts: degeneracy.

Degeneracy, the trap of simplicity

Let us consider the simplest ERG model of dyadic dependence, which is based on the graph Hamiltonian:

$$-\mathcal{H}_{2\star}(G, \theta) = \theta_- \sum_{i < j} a_{ij} + \theta_\wedge \sum_{i < j, k} a_{ik} a_{kj}. \quad (2.27)$$

This is referred to as the *two star* model, and includes two terms. The first, of the kind (2.26), counts the number of edges ($-$). The second counts the number of two stars, i.e., a pair of edges that share a common node (\wedge)²⁷. In practice, it has soon be realised that ERG estimations based on (2.27) are not possible [60].

The point at which the approximate estimation described in sec. 2.1.2 fails is the MCMC sampling of (2.4). Almost everywhere in the ERG parameter space, graph Markov chains based on (2.24) get trapped in graph configurations that are either almost empty, or almost fully connected. This means that for the vast majority of the parameter combinations,

24: *Homophily (heterophily)* refers to the tendency of nodes to form connections with others that have similar (different) attributes or characteristics.

25: It is a straightforward generalisation of the discussion in sec. 2.1.2 for Bernoulli random graphs. See the section III.B, *Generalised random graphs*, in [87]

26: In other words, only when they share a node can any two dyads be statistically dependent, when fixing the rest of the graph. Note that the Markov dependence is a property of the graph probability distribution, and not of the individual graph.

27: For instance, a two star model could be used to characterise a sparse graph ($\theta_- < 0$) with many high degree nodes ($\theta_\wedge > 0$).

28: The phenomenology of the two star model has been analytically understood by Park & Newman [48, 87]. The degenerate behaviour is nothing but a symmetry breaking between high- and low-density phases. A continuous (second-order) phase transition exists is found for this system. An analogous degeneracy problem affects other simple ERG models, e.g., the Strauss clustering model [60, 88].

29: According to [97], $x(G)$ is stable if there exist $C > 0$ and $L_c > 0$ such that $U_x \equiv \max_{G \in \mathcal{G}} x(G) \leq CL \forall L > L_c$, where L is the number of dyads in the network. For instance, since $U_- = L$, the edges term is stable while $U_\wedge \sim NL$ implies that the two star term is unstable.

30: An exponential distribution is *curved* – in the sense of Efron [101, 102] – when its natural parameters satisfy constraints. Here, for instance, in order to model the information of the graph probability distribution, one should generally include in the graph Hamiltonian one statistic for each of the $N - 1$ possible degrees, each associated to an independent parameter $\theta^{(k)}$. In (2.28), we are imposing the following non-linear constraint on the parameter space:

$$\theta^{(k)} = \theta w_\lambda^{(k)},$$

where $w_\lambda^{(k)}$ are defined in (2.29). Thus, in this view, we are constraining a $N - 1$ dimensional parameter space associated to the degree distribution to a two-dimensional subspace, hence the name of "curved" model. However, we take a slightly different view on the roles of θ and $w_\lambda^{(k)}$, see later in the text.

31: More specifically, an ERG model with a curved statistic such as (2.28), (2.29) is stable as long as $\lambda > -\log 2$ [97]. Here this is always the case, since $\lambda > 0$.

(2.24) assigns a negligible probability to all but the two extreme, unrealistic graph configurations (*degeneracy*) [60]. The subset of the parameter space that yields non-degenerate probability distribution turns out to be negligible and in practice unattainable [95, 96]. The estimation task is therefore impracticable²⁸.

The root of the the degeneracy issue is therefore the geometry of the ERG distribution. Unstable graph statistics lead to degenerate model behaviours, this is the case for the two star model²⁹. In such cases, the nearest-neighbor log probabilities (2.24) are unbounded, and MCMC simulations are a waste of time and resources – see [97] for a detailed, general discussion. These theoretical investigation had a fairly simple conclusion: inference for ill-posed ERG models is not possible.

Therefore, let us state the following:

Definition 2.3 (ERG model specification) The task of selecting a set of graph statistics that (i) optimally represents a given hypothesis about the relevant features of data and (ii) yields a non-degenerate ERG model.

A general strategy to cure the degeneracy issue is to "add structure" to the ERG specification. For instance, this includes adding constraints on the block, multilayer, spatial structure of the input graph [98]. Another solution is to stabilise the ERG model by using more sophisticated and more realistic graph statistics, as opposed to simple motifs counts as the two stars in (2.27). As this is the case of interest for this report, we describe it in some detail.

Curving ERG models

A widely adopted choice to design non degenerate ERG models consists in using the so-called *curved statistics* [99, 100]. Let us start by considering the following *geometrically weighted degree* (*gwd*):

$$x_{gwd}(G|\lambda) = \sum_k w_\lambda^{(k)} x_d^{(k)}(G), \quad (2.28)$$

where $x_d^{(k)}(G)$ is the number of nodes in the graph G with degree k and

$$w_\lambda^{(k)} = e^\lambda \left\{ 1 - (1 - e^{-\lambda})^k \right\}, \quad (2.29)$$

$\lambda > 0$. The *gwd* statistic is a linear combination of the graph degree distribution. An ERG model containing (2.28) is curved³⁰ and stable³¹ [97, 98]. To interpret the role of (2.28), we can reason in an analogous way to sec. 2.2.1 [2, 103].

As a result of adding an edge to the dyad (ij), the degrees of both the extremal nodes increase by one unit. Let us call G_{-ij}, G_{+ij} the graphs before and after the edge addition. If k_i is the original degree of the node i , then

$$x_d^{(k_i)}(G_{+ij}) = x_d^{(k_i)}(G_{-ij}) - 1, \quad x_d^{(k_i+1)}(G_{+ij}) = x_d^{(k_i+1)}(G_{-ij}) + 1, \quad (2.30)$$

and analogously for the node j . For simplicity, let us assume that the edge addition does not produce any other change in the graph statistics




Figure 2.2: Parametric weights for curved ERG statistics [99, 100]. Solid lines plot (2.29) for three sample values of λ , as a function of k . Dashed lines plot the geometric decreasing factor as it appears in (2.31), (2.33). The parameters λ determine the geometry of the curved statistics, therefore they pertain to the model specification problem. In practice, they can be inferred from data together with the maxent parameters.

than (2.30). From (2.24), one finds:

$$\begin{aligned} \log \frac{P(a_{ij} = 1 | G_{\setminus ij}, \theta^*)}{P(a_{ij} = 0 | G_{\setminus ij}, \theta^*)} &= \theta_{gwd} \left[x_{gwd}(G_{+ij} | \lambda) - x_{gwd}(G_{-ij} | \lambda) \right] \\ &= \theta_{gwd} (1 - e^{-\lambda})^{k_i} + \theta_{gwd} (1 - e^{-\lambda})^{k_j}. \end{aligned} \quad (2.31)$$

Note that $(1 - e^{-\lambda})^k \rightarrow 0$ for $k \rightarrow \infty$. This means that effect of the *gwd* statistic decreases geometrically with the degree, with the rapidity of the decay controlled by λ – large values of λ correspond to slow decays, fig. 2.2. Let us consider for instance $\theta_{gwd} > 0$. In this case, (2.30) always implies an advantage for the edge addition. However, this effect wanes and eventually vanishes for increasing the original degrees of the nodes i, j ³². Intuitively, this is what prevents a statistic such as (2.28) from driving the system into a fully connected state.


Apart from yielding non-degenerate ERG models, the *gwd* term is also a more plausible statistic than the two stars in (2.27) for modelling, e.g., the presence of hubs in a graph. On the one hand, hubs would be reflected by a degree distribution skewed towards higher degrees. On the other hand, there must be a cut-off for this effect: due to the presence of physical and/or functional constraints, it is highly unrealistic in biological networks for nodes to be connected to large fractions of the graph. The *gwd* statistic – and in particular the λ parameter – is able to tune this trade-off.

The *gwd* is only one of the possible curved statistics. They are all defined as linear combinations of a distribution of graph counts, with weights (2.29). Another example of interest for this work is the *geometrically weighted edgewise shared partner* (*gwesp*):

$$x_{gwesp}(G | \lambda) = \sum_k w_\lambda^{(k)} x_{esp}^{(k)}(G), \quad (2.32)$$

where $w_\lambda^{(k)}$ are defined in (2.29) and $x_{esp}^{(k)}(G)$ is the number of dyads that are connected by an edge and that have exactly k neighbors in common. It can be used as a more nuanced proxy for a tendency towards triadic

³²: In practice, the presence of nodes with high enough degrees is indistinguishable from the random expectation, i.e., the one we would get in the absence of the *gwd* term.



closure [104], rather than a simple (and unstable) count of triangles. Similarly as above, a single edge addition to the dyad (il) that increases the number of common neighbors between the connected nodes i, j by one unit results in

$$\log \frac{P(a_{il} = 1 | G_{\setminus il}, \boldsymbol{\theta}^*)}{P(a_{il} = 0 | G_{\setminus il}, \boldsymbol{\theta}^*)} = \theta_{g_{wesp}} (1 - e^{-\lambda})^{k_{ij}}, \quad (2.33)$$

where k_{ij} is the initial number of common partners of the nodes i, j . Once again, in the case $\theta_{g_{wesp}} > 0$, the tendency to add shared partners is damped for increasing k_{ij} , thus sidestepping the pitfall of degeneracy.

33: Parameters estimation of curved ERG models was first discussed in [100] and implemented in the first release of the `ergm` package [49]. See [91] for an overview of the recent developments.

ERG inference for curved models builds upon and extends the general framework discussed in sec. 2.1.2³³. In practice, the parameters $\boldsymbol{\theta} \in \mathbb{R}^r$ which weight the terms in the graph Hamiltonian and those $\boldsymbol{\lambda} \in \mathbb{R}_+^q$ which govern the geometry of the curved statistics ($q \leq r$) are estimated simultaneously [100]. Nevertheless, we regard them as logically distinct. The former are the Lagrange multipliers derived from the maxent principle, sec. 2.1.1. The latter, pertain to the issue of model specification, in the sense of def. 2.3.

We set aside both graphs and inference methods for the moment. The former will resurface in chapter 4, the latter one chapter later. Instead, over the next few pages, our discussion will take a sharp turn in both style and subject.

Darwin and the others

3

The alternative to thinking in evolutionary terms
is not to think at all.

—Peter Medawar

| | | |
|-----|----------------------------|----|
| 3.1 | Subject matter | 23 |
| 3.2 | Modelling evolution . . . | 25 |
| 3.3 | Genetic algorithms | 30 |

24 November 1859. The publication of *On the Origin of Species* [105] by the English naturalist Charles R. Darwin (1809-1882) constitutes a pivotal juncture in the annals of the history of science. The core set of concepts and principles there articulated have survived almost unaltered to the present day, eliciting deep philosophical ramifications¹ [107]. The theory there presented – combined with Mendel’s genetics [108] – represents our current conceptual understanding of the emergence of life’s complexity [109, 110].

According to Darwin’s theory it is the natural selection, or the *survival of the fittest*, that drives the emergence of the biological complexity, which has no equivalent in the inanimate world [111]. It is arguably a remarkable achievement, for a single theory, to put forth a unifying explanation for the process that leads to intricately structured organisms, starting from a primordial soup of molecules. This process, which is the subject of the present chapter, is referred to as *evolution*.

In the eyes of a physicist, evolution is puzzling. If the general picture of Darwin’s evolution is accepted, the details are poorly understood. The evolutionary problem is a formidable testing ground for our style of scientific enquiry, trained on spin lattices and pairwise interactions, now facing genetic codes and mostly unknown interactions [15].

First, the quest for a theoretical understanding of the evolutionary process entails learning from the fellow biologists *what* happens and *how* it happens, i.e., the subject matter (sec. 3.1). A theoretical framework can be then be established, which in turn allows for hypotheses to be formulated and models to be mathematically stated (sec. 3.2). In so doing, connections with related problems – by analogy and/or by generalization – are sometimes unveiled, and fruitfully exploited (sec. 3.3).

1: *A spectre is haunting the modern world, Darwin’s spectre, Darwinism.* – begins a recent book by Michael R. Rose [106], paraphrasing the opening line of a well-known revolutionary book, a *Manifesto*. Too much emphasis? Apparently, not. Ernst W. Mayr (1904-2005), leading evolutionary biologist, mentions *On the Origin of Species* by Charles Darwin among the three most influential books ever written [107], together with *Das Kapital*, by Karl Marx, and the *Bible*, by many authors (or just one).

Main reference

📖 Vito Dichio, Hong-Li Zeng, and Erik Aurell. *Statistical genetics in and out of quasi-linkage equilibrium*. In: Reports on Progress in Physics 86.5 (2023), p. 052601 [3].

3.1 Subject matter

Where an attempt is made to summarise in a few pages the salient features of evolution, in the narrative initiated by Darwin and perpetuated by modern evolutionary biology.

2: Two classical, comprehensive references that extensively cover the biology outlined in this section are [112, 113].

The Darwinian theory, or *Darwinism*, is the foundational theory upon which the entire field of evolutionary biology is built². The subject matter of Darwinism is the evolutionary dynamics of a population - the latter defined as a group of individuals (organisms) of the same species that live in a specific geographical area and reproduce across successive generations.

With respect to an individual, a distinction is made between the genotype and phenotype. The genotype refers to the specific set of genes that an organism carries, which encode the instructions for building and operating an organism. These genes can exist in different versions, called alleles. The phenotype, on the other hand, represents the actual observable traits or characteristics of the organism. Phenotypes result from the expression of an organism's genes as well as the influence of environmental factors and the interactions between the two. The genotype-phenotype map (or GP map) is the term used to describe the relationship between an organism's genetic makeup (its genotype) and its observable traits (its phenotype). It translates the information stored in the genes into the physical, macroscopic traits. This translation process is complex and not fully understood, as it involves multiple steps of gene expression, regulation, and interaction, and it is influenced by environmental factors as well [114].

The key ingredient of evolution is *inheritance*: offspring inherit traits from their parents through genetic information passed down from generation to generation. At the population level, two opposing drivers define the evolutionary dynamics:

- (a) Genetic *variation*. It broadly refers to the increase of genetic heterogeneity within a population. It primarily stems from by stochastic events that introduce variability in the genetic makeup of individuals. The two most common sources of variability are:
 - i. Mutations, which involve random alterations of an individual's genotype. Beneficial, deleterious, and neutral mutations respectively enhance, impair, or do not noticeably affect an organism's ability to survive and reproduce in its environment.
 - ii. Recombinations, which involve the exchange and rearrangement of genetic material between genotypes. Importantly, they require a physical transfer of genetic material between two individuals during reproduction in sexual populations.
- (b) Natural *selection*. It acts upon the genetic heterogeneity within a population, by favoring individuals with traits that enhance their survival and reproductive success, while disadvantaging those with less favorable traits. It is ultimately due to the selective pressure exerted by the environment, which include a variety of factors such as resource availability, predation, and competition. Through natural selection, advantageous traits become more common in a population over time, leading to the adaptation of species to their ecological niches.

Inheritance, variation, and selection form the conceptual core of Darwinism, constituting the fundamental principles that underpin the theory. It is worth emphasizing that while variation happens at the level of the genotypes, selection operates on observable traits, hence on phenotypes.

As stated in the beginning, the subject of the evolutionary process is not the genotype nor the phenotype of isolated individuals, but rather the population as a whole. Evolution unfolds over multiple generations, by gradually changing the genetic composition of populations or, more precisely, the statistical properties of the genotype distribution. One can then take a step forward and frame the problem in terms of the dynamics of the distribution of allele frequencies. This latter change of perspective resonates with classical statistical mechanics, as asserted already in the early 1950s by the R. A. Fisher [115]³:

Now, the frequencies with which the different genotypes occur define the gene ratios characteristic of the population, so that it is often convenient to consider natural population not so much as an aggregate of living individuals as an aggregate of gene ratios. Such a change of viewpoint is similar to that familiar in the theory of gases, where the specification of the population of velocities is often more useful than that of a population of particles.

R. A. Fisher, Croonian Lecture, 1953

3.2 Modelling evolution

Where one theoretical approach to evolutionary dynamics is presented and stated in the mathematical language. Where the discussion also hinges on the contemplation of a popular metaphor.

The diversity observed in natural outcomes of the evolution implies the need for a probabilistic description of the evolutionary dynamics. Indeed, the notion of probability pops up everywhere in formulating the building blocks of Darwinism: genetic mutations happen by chance, recombination events randomly reshuffle the parental genetic material, natural selection enhances - but does not guarantee - the reproductive success of apt individuals (...) The Darwin's theory is inherently a statistical theory. A model of Darwin's evolution is not expected to predict what *must* happen, but to inform about what *could* happen ⁴.

A wide array of theoretical models have been proposed, encompassing various aspects of genetic variation and evolutionary processes. A review of these approaches is far beyond the scope of this section, the interested reader is referred to [113, 119, 120]. Here, our attention is rather directed towards a particular such approach, which is instructive, as it comprises the subject matter discussed earlier and formulates the essential theoretical tools employed in this work. It has been originally formulated by *Neher and Shraiman* [121], and recently reviewed in [3].

3.2.1 Statistical genetics

The interest of physicists to the problem of the evolutionary dynamics stems primarily from a conceptual similarity between thermodynamics and quantitative genetics⁵, as first noted by R. A. Fisher in the 1930 [122]. The lack of an energy-like conserved quantity for the evolutionary

3: Ronald Aylmer Fisher (1890-1962) was a renowned British statistician and biologist. He is regarded as the *founder of modern statistics* [116]. He was also among the founders, together with J. B. S. Haldane (1892-1964) and S Wright (1889-1988), of population genetics – the subfield of biology that studies the distributions and changes of allele frequencies in a population (bottom-up approach) [117].

4: The prominent role of chance in the Darwin's theory caused a sensation in the scientific community at the time, including some scorn reactions. One of his scientific mentors, John Herschel (1792-1871), privately dubbed his theory as *the Law of higgledy-piggledy* [118].

5: The study of continuous-varying phenotypic traits (height, weight...), and specifically on population-wide averages over several genetically diverse individuals (top-down approach).

dynamics hampers a straightforward translation of the equilibrium thermodynamics laws, while a non-equilibrium picture appear more appropriate [123, 124]. The comparison can still be pursued, by seeking a theory for quantitative genetics that parallels the role of statistical mechanics for thermodynamics. Specifically, a phenomenological theory that provides a link between the population-averaged phenotypic traits and the underlying genotype dynamics. Such a theory is referred to as statistical genetics.

In the framework proposed by *Neher and Shraiman* [3, 121], a genotype is represented as a string of L binary variables $g = \{\sigma_1, \dots, \sigma_L\}$, where $\sigma_i = \pm 1$. Each locus σ_i represents a spin-like biallelic gene⁶, the number of genes L in each genotype is fixed. Consequently, the genotype space \mathcal{G} is an L -dimensional hypercube. Each possible genotype is found in a population with probability $P(g, t)$, which depends on time. In particular, it changes under the effect of mutations, recombinations and selection.

6: Here, for simplicity, one gene = one locus (spin variable).

- (a) *Mutations*. In a time interval Δt , mutations change the genotype distribution as follows:

$$P(g, t + \Delta t) = P(g, t) + \Delta t \mu \sum_{i=1}^L [P(M_i g, t) - P(g, t)], \quad (3.1)$$

where $\mu > 0$ is the constant mutation rate, uniform across the genotype, and the operator M_i swaps the i -th locus i.e., it replaces $\sigma_i \rightarrow -\sigma_i$.

- (b) *Recombinations*. One such event consists in the exchange of genetic material between two individuals g^*, g^{**} , to form an offspring. The result is a novel genotype g , which randomly inherits parts of the parental genotypes (crossover). Formally, this can be described by a set of Boolean variables $\xi_i = \{0, 1\}$, defining $\sigma_i = \xi_i \sigma_i^* + (1 - \xi_i) \sigma_i^{**}$. In words, the locus σ_i of the offspring g is inherited from g^* if $\xi_i = 1$, from g^{**} if $\xi_i = 0$. Each crossover pattern $\xi = \{\xi_i\}$ comes with probability $C(\xi)$. The change of the genotype distribution induced by recombinations is:

$$P(g, t + \Delta t) = (1 - \Delta tr)P(g, t) + \Delta tr \sum_{\xi, \bar{g}} C(\xi) P(g^*, t) P(g^{**}, t), \quad (3.2)$$

where r is the recombination rate and the sum in the second terms runs over all possible crossover patterns ξ and the genetic material \bar{g} that is discarded during the recombination event, i.e., $\bar{\sigma}_i = (1 - \xi_i) \sigma_i^* + \xi_i \sigma_i^{**7}$.

7: Recombination acts as collision process in the theory of gases. An unspoken assumption in (3.2) is that all the two-genome distributions factorize $P_2(g^\alpha, g^\beta, t) \sim P(g^\alpha, t) P(g^\beta, t)$, which is equivalent to the assumption of molecular chaos (*Stosszahlansatz*). A critical discussion of this hypothesis can be found in [3].

- (c) *Natural selection*. A fitness-based scheme is formulated. The fitness of a genotype is proportional to the average number of offspring of an individual with genotype g [74]. Since it is a non-negative function, we choose an exponential representation $\text{fitness}(g) \sim \exp[\Delta t \varphi F(g)]$ where Δt is a time interval, $\varphi > 0$ is the selection rate and $F(g) \in \mathbb{R}$ is the fitness function. We postpone to sec. 3.2.2 a thorough discussion of the latter. Natural selection favors individuals with higher-than-the-average reproductive success as follows:

$$P(g, t + \Delta t) = \frac{e^{\Delta t \varphi F(g)}}{\langle e^{\Delta t \varphi F(g)} \rangle_t} P(g, t), \quad (3.3)$$

where we have used the notation $\langle e^{\Delta t \varphi F(g)} \rangle_t = \sum_g e^{\Delta t \varphi F(g)} P(g, t)$ for the population-average. It is important to stress that, according to (3.3), the notion of fitness is inherently comparative. It is not the intrinsic value of fitness that is important, but rather its value in comparison to average in the population, at a given time.

In the limit $\Delta t \rightarrow 0$, the combined action of mutations, recombinations and selection can be expressed in a unified phenomenological master equation [125]:

$$\begin{aligned} \frac{d}{dt} P(g, t) = & [F(g) - \langle F \rangle_t] P(g, t) + \mu \sum_{i=1}^L [P(M_i g, t) - P(g, t)] + \\ & + r \sum_{\xi, \bar{g}} C(\xi) [P(g^*, t) P(g^{**}, t) - P(g, t) P(\bar{g}, t)]. \end{aligned} \quad (3.4)$$

The master equation (3.4) describes the genotype dynamics in the limit of an infinite population i.e. $M \rightarrow \infty$, where M is the number of individuals in a population. This corresponds to the limit of a perfect sampling of the genotype distribution, which allows to neglect the sampling noise – referred to as genetic drift.

The dynamics of an asexual population can be studied by setting $r = 0$ in (3.4). In sexually reproducing populations, it is frequently observed that recombinations happen at much faster rate than mutations. In this case, the dynamics on the time scale r^{-1} is investigated by using (3.4) with $\mu \sim 0$ [121]. Note however that recombinations do not create nor eliminate alleles in the population. Therefore, for the long-term dynamics of the allele frequencies in a population must be governed by the influx of new mutations. Indeed, mutations play a more fundamental role in the evolutionary process compared to recombinations. Unlike the latter, which act on pre-existing variability, mutations actively spawn new variations⁸.

With (3.4) one can in principle to compute the dynamics of the average of any quantitative trait $O(g)$, i.e., of any function in the genotype space. In fact,

$$\frac{d}{dt} \langle O \rangle_t = \frac{d}{dt} \sum_g O(g) P(g, t) = \sum_g O(g) \left[\frac{d}{dt} P(g, t) \right], \quad (3.5)$$

using (3.4) to evaluate the boxed quantity. Moreover, it can be used to evaluate the dynamics of the distribution of any trait, which can be obtained from (3.4) by projection:


$$P(O, t) = \sum_g \delta[O - O(g)] P(g, t), \quad (3.6)$$

where $\delta(O)$ is the Dirac delta function. The equations (3.5), (3.6) fulfill the purpose of the theoretical framework, that is, to furnish a formal scaffolding that enables (i) hypotheses on the parameter space to be formulated and (ii) computations to be carried out. The results of such computations, in turn, are to be tested against experimental data⁹.

8: A simple example is that of a population of M identical individuals. Recombinations have no effect whatsoever in the genetic composition of the populations. For an evolutionary process to begin, the influx of new alleles (read, mutations) is needed.

9: A review of the results recently obtained starting from (3.5) falls outside the scope of this manuscript. A considerable interest has been devoted to the region of the parameter space corresponding to the *quasi-linkage equilibrium* phase, investigated both in theory [121, 126, 127] and numerically [3, 5, 128] using efficient simulation tools [129]. Notably, the framework outlined here has enabled compelling connections with experimental data, as explored and exploited in the case of the SARS-CoV-2 viral genomes [4, 130].

Figure 3.1: A beautiful, deceptive picture. A long-standing tradition, initiated by Wright [131], is to depict the notion of a fitness function a surface in an aesthetically pleasing 2D or 3D plot. Here, the xy plane is the state space (genotypes or phenotypes) and the z direction, or "height", stands for fitness. This representation makes intuitive the notion of evolution as an attempt to "climb the fitness hills". However, the true state space is multi-dimensional and intuitions developed in a lower dimensional representation can be highly misleading, both quantitatively and qualitatively (e.g., the notion of "neighbor"). For these reasons, it has been proposed to abandon at all the landscape picture [132, 133]. On our end, we recognize the efficacy of the abstract landscape illustration and semantics, we advise prudence when using it as a tool for intuitive understanding.



3.2.2 On the fitness landscape

10: Moreover, as old as population genetics itself. In fact, it was introduced by S. Wright in the 1932: "*The problem of evolution as I see it is that of a mechanism by which the species may continually find its way from lower to higher peaks in such a field.* (fitness landscape, ed.)" [131]

A key conceptual aspect of the framework illustrated above pivots around the question: what is the fitness function? It is a coarse-grained metaphor, at once useful and of limited use¹⁰.

In the previous section, the fitness function F has been assigned the role of providing a map from the genotype space to the reproductive success. It is the environment that exerts a selective pressure of a population (competition for limited resources), hence establishing the fitness of any individual. Selection, however, acts on phenotypic traits, hence it impacts genotypes in an indirect manner. This implies that the mapping from genotypes to fitness aggregates two distinct types of information: (i) the genotype-phenotype map and (ii) the phenotype-to-reproductive-success map [114].

If and when the theoretical picture of a fitness function holds valid, the evolutionary process can be thought as a dynamic on the *fitness landscape*, with the population probability distribution that, driven by natural selection, progressively climbs up the fitness hills, fig. 3.1. The shape of the fitness landscape is fixed once the parameters of F are specified. In particular, as a function on a L -dimensional hypercube, the fitness function can generally be decomposed as follows:

$$F(g) = \bar{F} + \sum_i F_i \sigma_i + \sum_{i<j} F_{ij} \sigma_i \sigma_j + \sum_{i<j<k} F_{ijk} \sigma_i \sigma_j \sigma_k + \dots, \quad (3.7)$$

where \bar{F} is a constant and the F_{i_1, \dots, i_k} are the parameters - there are 2^L of them, as expected for an exact representation. The higher the order of interactions that play a significant role in (3.7), the more rugged and multi-peaked will be resulting landscape. At the very least, (3.7) is useful, as it enables the explicit evaluation of the dynamics of phenotypic traits (3.5), (3.6).

11: What is the minimal genome size needed to support life? The record-holding organism, to date, is the bacterium *Candidatus Nasuia deltocephalinicola*, with 137 protein-coding genes [134]. A fitness function of the form in (3.7) would have $2^{137} \sim 10^{41}$ parameters (!).

The 2^L parameters of (3.7) are arguably a huge number¹¹. We therefore seek a theoretical argument to scale down the dimensionality of F , i.e., the number of degrees of freedom and, by consequence, that of parameters. Two possible approaches are the following.

- i. Landscapes of the form (3.7) have been thoroughly investigated within the framework of spin glass theory [85]. Notably, it is understood from these studies that the essential consequences of the complexity of such a landscape are already manifest in a model with only pairwise interactions [135]. Therefore, for the sake of simplicity, one starts exploring a simplified fitness function:

$$F(g) = \bar{F} + \sum_i F_i \sigma_i + \sum_{i < j} F_{ij} \sigma_i \sigma_j, \quad (3.8)$$

which has $\sim L^2$ parameters. As a matter of fact, this is the only (non-trivial) choice that has allowed to solve the evolutionary dynamics in a closed form [3, 127].

- ii. In order to mimic more closely the mechanism of natural selection, it is possible to define a set of phenotypic variables $f_{GP}(g) \in \mathbb{R}^r$, where $f_{GP} : \mathcal{G} \mapsto \mathbb{R}^r$ is the genotype-phenotype map. The cardinality of the phenotype space is necessarily lower than or equal to that of the genotype space, i.e., $r \leq 2^L$. A phenotype-fitness map $F : \mathbb{R}^r \mapsto \mathbb{R}$ can be defined on the phenotype space. The fitness function is then written as $F(g) = F(f_{GP}(g))$,

$$\mathcal{G} \xrightarrow{f_{GP}} \mathbb{R}^r \xrightarrow{F} \mathbb{R}. \quad (3.9)$$

For instance, if F is a simple linear combination of phenotypic traits, the resulting fitness function has a number r of parameters.

Historically, skepticism about the concept of a fitness landscape has largely stemmed from the lack of empirical data to outline its actual topography. In recent years, this critique has been somewhat offset with the advent of methods that allow the construction of *empirical fitness landscapes*¹² [114, 136]. Though still in their early stages, these experiments carry potential to drive a more comprehensive understanding of the shape and significance of the fitness function.

A different source of criticism has emerged in recent years, against the assumption of a constant environment, implicit in (3.7) [138, 139]. A time-dependent selection would rather be mediated by a *fitness seascape* $F(g, t)$. In fact, the hypothesis of a fixed environment is never exactly true, as even in the simplest lab environments a number of factors induce a temporal dynamic of the selective pressure, including modifications of the physical environment, frequency-dependent selection, co-evolution effects, interaction between species (ecology) [140]. The assumption of a fixed landscape holds approximately true when the timescale of the environmental changes, broadly speaking, is sufficiently long compared to that of the process under investigation. Unquestionably, it results in a remarkable simplification of the complexity of the phenomenon.

The picture of a fitness landscape, therefore, should be regarded as approximate in several senses. Yet, with the *caveat* above, it offers an instructive way of thinking about evolution.

12: These approaches consist in creating artificial mutants, each carrying one or more mutations with respect to the wild-type genotype, then measuring their fitness using a fitness proxy (e.g., antibiotic resistance) [136]. For instance, deep mutational scans [137] are able to test and assess the phenotypes of all single mutants, and several double- and triple-mutants of a wild-type genotype.

3.3 Genetic algorithms

Where evolution is recognised as solving a more abstract optimisation problem, the evolutionary algorithm is therefore isolated and used elsewhere.

Devoid of all details, the evolutionary problem described in sec. 3.2 is easily recognized as a particular instance of a general discrete optimization problem. The latter is defined as:

$$\max_y f(y) \quad \text{subject to} \quad y \in \mathcal{Y}, \quad (3.10)$$

where $f : \mathcal{Y} \mapsto \mathbb{R}$ is the *objective function*, \mathcal{Y} is the *feasible set* and the y_i , $i = 1, \dots, l$ are *binary decision variables* [141]. In sub-case of evolution, the binary decision variables are the biallelic genes σ_i , the feasible set is the genotype space \mathcal{G} , the objective function is the biological fitness F . In this (simplified) view, natural selection acts as an optimization process. Over time, evolution leads to the "optimization" of a population for survival and reproduction in their specific environment.

For all problems of the form (3.10), a straightforward algorithm exists. It involves generating a complete list of all possible $y \in \mathcal{Y}$, evaluating the objective function value for each solution, and identifying those y that yield the maximum value of f . Practically however, this becomes soon unfeasible, since it involves computing $\sim 2^l$ computations of f . The exponential explosion of the running time as a function of the dimension of the problem is often referred to – with a hint of desperation – as the *curse of dimensionality* [142]. In computer science, several approximate yet efficient methods have been devised to explore the feasible solution space of an optimization problem and discover the optimal solutions. These approaches, commonly known as *metaheuristics*¹³, encompass a range of techniques [143], often inspired by natural processes – simulated annealing being a celebrated example [146].

Genetic algorithms (GAs), in particular, leverage the parallelism demonstrated above with the evolutionary problem, to construct a metaheuristic for (3.10)¹⁴ [150]. In GAs, candidate solutions to a problem undergo successive generations of selection, recombination, and mutation to converge towards optimal or near-optimal solutions. GAs are population-based, i.e., they operate on a population of potential solutions, each represented as a binary strings. The following example presents a minimal version of the algorithm.

An illustrative case

Suppose an optimization problem

$$\max_y y_{[10]} |\sin y_{[10]}| \quad \text{subject to} \quad y \in \mathbb{Z}_2^5, \quad (3.11)$$

where the feasible set $\mathcal{Y} = \mathbb{Z}_2^5$ is the space of all possible binary bit-like strings¹⁵ of length 5, $f(y) = y_{[10]} |\sin y_{[10]}| \in \mathbb{R}^+$ is the objective function and $y_{[10]}$ is the decimal representation of the binary string y ¹⁶. Clearly, for (3.11) an exhaustive list of all possible $2^5 = 32$ strings in \mathcal{Y} could

13: Metaheuristics are problem-solving strategies that provide a general framework for solving optimization problems. They are not tailored to a specific problem (as it is the case of heuristics) but offer a set of guiding principles and strategies that can be applied to a wide range of problem domains. An overview of the existing approaches can be found in [143]. A substantial subset of such techniques is inspired by biological systems, as it is the case for the particle swarm optimization [144], ant colony optimization [145], and genetic algorithms (see below).

14: John H. Holland (1929-2015), father of GAs, opens an influential article in *Scientific American* by proclaiming: *Living organisms are consummate problem solvers. (...) Pragmatic researchers see evolution's remarkable power as something to be emulated rather than envied.* [147]. However, GAs were introduced by Holland and his students in the 1970s not with a specific application as (3.10) in mind, but with the general purpose of simulating and studying artificial adaptive systems in a computational environment [148, 149].

15: In the previous section, we used spin-like binary variables $\sigma_i \in \{-1, 1\}$, as it is common in statistical physics. Here, we rather use bit-like binary variables $y_i \in \{0, 1\}$, as it is common in computer science. The two sets are obviously equivalent.

16: In the language of sec. 3.2, we would say that y is a genotype, $y_{[10]}$ its phenotype and $f(y)$ its fitness.

be written, their f evaluated and the maximum determined. Here, for illustrative purposes, we ask (3.11) to be solved using a GA.

The algorithm starts by randomly instantiating an initial population of M strings, say $M = 4$. Three evolutionary operators are then iteratively applied at each generation:

1. *Selection.* The objective function f is evaluated for all $y^{(i)}$, $i = 1, \dots, M$ in the last generation. Each string $y^{(i)}$ has a probability to reproduce proportional to its value of the objective function. In the simplest scheme (biased roulette wheel), the probability corresponds to

$$f(y^{(i)}) / \sum_j f(y^{(j)}) . \quad (3.12)$$

See tab. 3.1 for an example. As a result of selection, the population average of the objective function $\langle f \rangle = \sum_{j=1}^M f(y^{(j)})$ increases.

| i | $y^{(i)}$ | $y_{[10]}^{(i)}$ | $f(y^{(i)})$ | (3.12) | new count |
|-----|-----------|------------------|--------------|--------|-----------|
| 1 | 01101 | 13 | 5.46 | 0.10 | 0 |
| 2 | 10101 | 21 | 17.57 | 0.33 | 1 |
| 3 | 01111 | 15 | 5.44 | 0.19 | 1 |
| 4 | 11010 | 26 | 3.30 | 0.38 | 2 |

Table 3.1: Example of selection. Each existing string $y^{(i)}$ as a probability (3.12) to be selected proportional to its value of the objective function f . In the last column, the result of a sampling of $M = 4$ individuals. The population average $\langle f \rangle$ increases from 13.15 to 16.74.

2. *Recombinations.* Pairs of strings undergo recombination with rate r . A popular such scheme (single-point crossover) entails swapping a portion of the parental strings, e.g.

$$\begin{array}{c} 01111 \\ 11010 \end{array} \rightarrow \begin{array}{c} 01 | 111 \\ 11 | 010 \end{array} \rightarrow \begin{array}{c} 01010 \\ 11111 \end{array} , \quad (3.13)$$

where the position of the cut $|$ is randomly chosen.

3. *Mutations.* Each existing bit in the population is flipped with uniform rate μ , e.g.

$$01010 \rightarrow 01111 . \quad (3.14)$$

By repeatedly applying the three aforementioned steps, the population gradually converges towards the optimal solution (3.11). However, the parameters governing the strength of evolutionary operators must be tuned to achieve a balance between the exploitation of those strings that better approximate the optimal solution (selection) and the exploration of the solution set (recombinations, mutations)¹⁷.

Remarks

A number of successful applications of GAs to real-world optimization problems demonstrate the interest of the scientific community in the approach – see [151] and references therein.¹⁸ Indeed, genetic algorithms exhibit two compelling features that make them highly attractive.

- GAs are population-based.* Realistic objective function are largely multimodal, meaning that they exhibit multiple peaks and local optima. Any search algorithm in such a landscape should avoid getting ensnared within a local optima and rather pursuit the global optimum. GAs rise to the occasion by unleashing multiple

17: In the GAs literature, however, mutations have a secondary role with respect to recombinations. They are mostly regarded as a mere insurance policy against the loss of diversity (s.n. 8) and a premature convergence of the search algorithm [147, 150]. By consequence, mutation rates are typically set to low values.

18: Curiously, there is no consensus on the exact reasons *why* GAs work. An overview of the debate can be found in [151, 152].

walkers that venture through the solution space simultaneously. In the case where a walker became trapped, the others would aid in circumventing the trap in the subsequent generation. This mitigates – even though, does not eliminate – the peril of stagnation around sub-optimal solutions.

- ii. *GAs have minimal assumptions about f .* As it should be clear from the example above, GAs solely rely on the payoff (objective function) values assigned to individual strings. For instance, they do not require the computation of derivatives, as gradient-based methods do. Even more, GAs do not even necessitate a mathematical expression for f . Consistently with the logic of GAs, an objective function can be considered a black box that takes an input (string) and produces a real number as output. This implies the formulation of GAs is *problem independent*, since it does not rely of problem-specific information about the topography of the search landscape.

Clearly, enhancing robustness and generality does not come without drawbacks, and GAs are not an exception, notably:

- iii. *GAs are computationally demanding.* The major computational bottleneck of GAs is the large number of objective function evaluations, which scales linearly with the population size. Therefore, GAs simulations are highly sensitive to the complexity of the search landscape. The latter, in turn, depends on the problem at hand, therefore there is no universal approach to guarantee timely convergence. Approximate methods and/or additional assumptions should be tailored to the specific phenomenon under investigation.

The exploration-exploitation dynamics

4

No one trusts a model except the man who wrote it;
everyone trusts an observation, except the man who made it.

—Harlow Shapley

| | | |
|-----|----------------------------|----|
| 4.1 | Fundamentals | 34 |
| 4.2 | Simple scenarios | 38 |
| 4.3 | Simulations | 42 |

To navigate its microscopic world, the bacterium *Escherichia coli* uses chemotaxis to find nutrients and avoid harm. When a desirable sugar molecule is in the vicinity, its receptors alert the bacterium's tiny flagella, which spin counterclockwise to gently glide the *E. coli* towards the source. Conversely, if a harmful substance is detected, the flagella whirl clockwise, causing the *E. coli* to tumble and change direction, scurrying away from the danger [153, 154].



When a foraging honeybee (*Apis mellifera*) comes across a rich source of nectar, it makes its way back to the hive to perform the iconic 'waggle dance'. This intricate dance communicates the location of the food source to her hive mates, inviting them to take advantage of the newfound resources. Some, however, deliberately ignore the information in the dance and choose to continue exploring the environment for new nectar sources. This provides the colony with a safety net against the whims of flower nectar production [155–157].



The slime mold *Physarum polycephalum* utilizes a form of reactive navigation to explore its environment. It is composed of several oscillating units, with the oscillation frequency varying according to environmental cues. In the presence of attractants like food, the oscillation frequency increases, inducing cytoplasmic flow towards the food source, while repellents like light or salts reduce the frequency. Moreover, as the slime mold forages, it leaves behind a trail of nonliving extracellular slime, which it later avoids. This simple, noneuronal organism is able to solve complex navigation tasks, such as the U-shaped trap problem [158].



When trying to attract a mate, male fruit flies (*Drosophila melanogaster*) serenade females with a courtship song, produced by wing vibration [159, 160]. Fruit flies display two major modes of song (sine and pulse). Mode transitions, and variable mode durations result in individual courtship sessions that are highly diverse. Recent evidence shows that simple flies use sensory information – in particular, the female's position and velocity – to pattern his song sequences over short timescales [161].



(...)

Flicking through the catalogue of biological dynamics, the most extraordinary behaviours stand out. The examples provided above represent a concise selection of noteworthy cases, chosen among numerous other possible instances¹. There is a clear leitmotif that connects them all, and with the evolutionary dynamics of ch. 3, and with another case that we will discuss in ch. 5. Indeed, they can all be brought under the same conceptual umbrella: *seeking a goal under uncertainty*. Indeed, this seems

1: See also the ant foraging behaviour [162], the fish predatory strategies [163], the root growth dynamics [164], the immune system response to pathogens [165].

to be a ubiquitous requirement of life, from the simplest organisms to complex human social behaviours [166, 167]. There are two possible reasons for a concept to be ubiquitous in a scientific domain: either it is trivial, or it is fundamental. Here, we will argue for the latter.

We start by formalising what we call the *exploration-exploitation* dynamics and providing a mathematical formulation for the case where a biological system can be represented as a graph (sec. 4.1). We start by studying simple, solvable models, which are instructive to showcase the essential features of the EE dynamics (sec. 4.2). Finally, we briefly describe and test the simulations we have designed to cope with complex scenarios (sec. 4.3).

Main reference

📖 Vito Dichio & Fabrizio De Vico Fallani. *The exploration-exploitation paradigm for networked biological systems*. In: arXiv e-prints 2306.17300 (2023) [1].

4.1 Fundamentals

Where the core idea of this dissertation is presented in its final, press-ready version. In particular, where the exploration-exploitation problem for (networked) biological systems is defined and formalised.

The line of thought is fairly straightforward. Biological systems inherently and universally exhibit randomness. Nonetheless, their dynamics are shaped by functional constraints, therefore randomness and function must coexist. Even more, biological systems achieve high-level functions not only *in spite of* randomness but also *through* randomness² [173, 174].

The space in which the randomness unfolds – typically, the systems' configuration space – is heavily constrained. In fact, in biology, systems have functions and we can reasonably expect the overwhelming majority of possible configurations to be *non functional* or *poorly functional* [14, 175]. Broadly speaking, therefore, the relation between randomness and biological function has a dual nature. On one side, randomness frequently triggers changes that are detrimental. On the other hand, it serves as an essential mechanism for exploring the range of possible configurations, and identifying those that enhance the system function³. If the details of this interplay are context-dependent, it is meaningful to look for general principles.

In the following, we develop a formalism that embeds the discussion above into two fundamental concepts: *exploration*, *exploitation*. Exploration indicates the stochastic search of the configuration space. Exploitation refers to the use of the configurations that have been found to optimise the system function. The optimisation problem implicit in these definitions, in turn, is formalised by stating (i) how the optimal states are encoded and (ii) how the system approaches them.

To proceed further, we need to define the characteristics of the system configuration space, i.e., we need to select a representation. Henceforth,

2: In the last two decades, with the advent of quantitative biology, there has been a paradigm shift in how we look at the role of chance in living systems. There is growing evidence that randomness is not always a hindrance to biological function, and it can be more than just *noise*, it can be a potential asset in the workings of life. Long established in the context of evolution, it is now recognised, e.g., within the domains of molecular biology [168, 169], cell biology [170], neuroscience [171, 172], to name a few.

3: In this paragraph, we have used the words *randomness* and *biological function* rather loosely. On the one hand, it was necessary to keep the discussion general. On the other hand, the exact definition of these two concepts is more problematic than it may seem at first sight. The semantic discussion about the meaning of both is instructive, since goes deep into the foundations of biology. Excellent starting points are, e.g. [175, 176].

we will narrow our discussion to consider a graph representation and develop an exploration-exploitation (EE) graph dynamics.

4.1.1 A graph dynamic

Consider a biological system that can be represented as an unweighted, undirected graph $G \in \mathcal{G}$ over N nodes, as defined in (1.1). Let $P(G, t)$ be the probability of observing the graph G at time t . Exploration and exploitation affect the graph probability distribution as follows.

Exploration

The most elementary configuration change involves the reversal of one dyadic state, which means that if there was no edge present, one is created $0 \rightarrow 1$, or if an edge was already in place, it gets removed $1 \rightarrow 0$. Let $M_{ij} : \mathcal{G} \mapsto \mathcal{G}$ be the operator which, acting on the graph G , inverts the state of its dyad (ij) . Again, by simplicity, we assume that dyadic reversals occur throughout the graph at a uniform, time-independent rate μ (exploration rate)⁴.

In the time interval Δt , the graph probability distribution changes as a result of the exploration as follows

$$P(G, t + \Delta t) = P(G, t) + \Delta t \mu \sum_{i < j} [P(M_{ij}G, t) - P(G, t)]. \quad (4.1)$$

Exploitation

Let us assume that the notion of *biological function* for a given system can be mathematically represented as a function on the graph configuration space $F : \mathcal{G} \mapsto \mathbb{R}$, which we refer to as *functional metric*, or simply *F metric*. As the argument goes, biological systems attempt to maximise their function, hence to increase the value of their *F metric*⁵. Optimal states correspond to the maxima of *F* or, pictorially, to the peaks in the *F landscape*.

In a time interval Δt , exploitation changes the graph probability distribution according to:

$$P(G, t + \Delta t) = \frac{e^{\Delta t \varphi F(G)}}{\langle e^{\Delta t \varphi F} \rangle_t} P(G, t), \quad (4.2)$$

where $\langle \cdot \rangle_t$ stands for the ensemble average at time t , i.e., $\langle e^{\Delta t \varphi F} \rangle_t = \sum_G e^{\Delta t \varphi F(G)} P(G, t)$ and $\varphi > 0$ (exploitation rate) is an overall scaling. In words, the way in which exploitation approaches the most functional configurations is by amplifying (reducing) the likelihood of those graphs that exhibit higher (lower) *F* values than the ensemble average at time t , fig. 4.1. The topography of the *F* landscape plays the delicate role of simultaneously encoding (i) the mechanisms that yield high functional states and (ii) the constraints that restrict the dynamics in the graph configuration space. It generalises the notion of *environment* of evolutionary dynamics to what we might designate as the *functional context*. Clearly, *F* is problem-dependent.

4: This is the simplest, non-trivial, exploration scheme. In fact, (i) it is defined for each microscopic degree of freedom (dyad), (ii) it happens at a constant rate and (iii) it does not depend on the dyads (iv) nor on their history. Each of these four assumptions can be relaxed to design more complex exploration algorithms.

5: In the following, we will mostly use "F metric" instead of "function", to avoid confusion with the homonymous general mathematical notion.




Figure 4.1: Dynamic of the one-dimensional *F* distribution (illustration). At each point in time, the likelihood of those graphs that have higher *F* values than the ensemble average (dotted lines) increases (upwards arrows), the likelihood of those that have lower *F* values decreases (downwards arrows). As a result, the *F* distribution as a whole shifts towards higher *F* values.

In the following, we further assume the existence of a lower dimensional *sufficient* representation of a graph G , in terms of $r < 2^L$ statistics $\mathbf{x}(G) \in \mathbb{R}^r$. Therefore, in this latter space the F metric is mathematically defined $F(G) = F(\mathbf{x}(G))$,

$$\mathcal{G} \xrightarrow{\mathbf{x}} \mathbb{R}^r \xrightarrow{F} \mathbb{R}. \quad (4.3)$$

We can collect (4.1) and (4.2) in a single expression, which defines our EE graph dynamics⁶:

6: This dynamics is correctly normalised over the graph ensemble $\langle \cdot \rangle_t$ at each time t . Note that the two terms in $\Delta P_{\Delta t}(G, t)$ are normalised independently, as it should be.

$$P(G, t + \Delta t) = P(G, t) + \underbrace{\Delta t \mu \sum_{i < j} [P(M_{ij}G, t) - P(G, t)] + \left[\frac{e^{\Delta t \varphi F(G)}}{\langle e^{\Delta t \varphi F} \rangle_t} - 1 \right] P(G, t)}_{\Delta P_{\Delta t}(G, t)}. \quad (4.4)$$

It is convenient to define an adimensional parameter to weight the relative strengths of exploration and exploitation. We call it *functional pressure*,

$$\rho = \varphi / \mu. \quad (4.5)$$

For mild functional pressures, $\rho \sim 0$ the dynamic is dominated by random dyadic mutations, and it is similar to a random walk in the graph space \mathcal{G} . On the contrary, $\rho \rightarrow \infty$ corresponds to the limit of a perfectly exploitative dynamics, where only the most functional graph configurations carry significant probability.

4.1.2 Beyond Darwin, an interpretable GA

It will not have escaped the attention of the reader that the theoretical structure assembled in the previous section bears resemblance to the evolutionary dynamics we formulated in ch. 3. Of course, this is no mere coincidence.

The EE graph dynamics *is* an evolutionary dynamics – in the sense of sec. 3.2.1 – for genotypes of length L , based on mutations and natural selection. It is worth stressing the semantic: the EE dynamics defined in (4.1) - (4.2) is not *analogous* to the mutation-selection dynamics (3.1) - (3.3): mathematically speaking, they are the *identical*⁷. Naturally, what changes from one case to the other is the interpretation we give to the same mathematical objects.


7: This is possible in the first place because we have used the same representation for genotypes and graphs, both of which live in the space of L -dimensional binary strings – though, technically, the first g is a vector of length L , the second G is a matrix with L degrees of freedom. This is also the reason why we have used the same notation \mathcal{G} for the space of genotypes and graphs, the same F for the fitness function and biological function and so on.

In the previous section, we have formulated the EE dynamics on a very general basis. The key logical step is to recognise that the evolutionary dynamics (without recombinations) is a particular instance of the more general EE dynamics where (i) exploration is interpreted as genetic mutations and (ii) exploitation is interpreted as natural selection – the "biological function" to be maximised is the fitness, or reproductive success. See tab. 4.1 for a complete list of correspondences. In this sense, the EE dynamics is a generalisation of an evolutionary process without recombinations.

| EE graph dyn. | Notation | Evolutionary dyn. |
|-----------------------------|----------------------|----------------------|
| graph space | \mathcal{G} | genotype space |
| number of dyads | L | genome length |
| graph (unw., und.) | G/g | genotype |
| graph statistics | x/f_{GP} | phenotype (traits) |
| dyad (bit-like) | a_{ij}/σ_{ij} | locus (spin-like) |
| time window | T | # generation |
| # samples* | M | population size |
| bio. function (F metric) | F | fitness function |
| exploitation r. | φ | natural selection r. |
| exploration r. | μ | mutation r. |
| \times | r | recombination r. |

In other words, in writing (4.4) we borrow the *algorithm* coded by Nature itself to solve the exploration-exploitation problem in the specific context of evolution. The crucial motivation for this is that evolutionary dynamics exhibits a ubiquitous, general and emergent property of the dynamics of living systems: it is *self-referential* [15]. This means that the time-evolution operator that governs the dynamic depends on the state of the system, therefore on its history, as it is manifest in the term $\langle \exp(\Delta t \varphi F) \rangle_t$ in (3.4) or (4.4). This is *the key, defining feature of biological dynamics*⁸.

In sec. 3.3 we have already discussed a very similar generalisation of the evolutionary dynamics, that of genetic algorithms (GAs). Indeed, our EE dynamics can be regarded as a GA where mutations are the crucial exploration mechanism and, once again, recombinations are left behind. In fact, unlike mutation and selection, it is not clear what recombination should correspond to, outside the context of the evolutionary process⁹. This is not an embarrassment for GAs, since their generalisation of the evolutionary process is purely *algorithmic*, i.e., they treat evolution as a computational strategy for solving optimization problems, no less, no more.



Here, however, we have a more ambitious plan. We want to leave open the possibility of interpreting the EE dynamic as a real-world process, when used outside in the evolutionary context. We do want our EE dynamics – lifted from evolution – to be meaningfully instantiated whenever the exploration-exploitation mechanisms cooperate in shaping a biological process. For this to be possible, we must handle the lesson we learn from

Table 4.1: Translation rules from the vocabulary of EE graph dynamics to that of evolutionary dynamics. We group by colour those terms that refer to the structure of the configuration and state spaces (top), to the structure of the dynamics (middle) and to the dynamic parameters (bottom). *See sec. 4.3.

8: Self-referentiality has no equivalent in the dynamics of conventional matter. The bewilderment and despair of a physicist – accustomed to the educated, gentle phenomenology of the dynamics of condensed matter – is well summarised in a recent review on the subject: *To a physicist, this sounds strange and mysterious: What is the origin of this feature that sets biological systems apart from physical ones? Aren't biological systems ultimately physical ones anyway; thus, why is self-reference an exclusive feature of biological systems (whatever they are!)?* [15]

9: In evolution, in fact, the probability distribution $P(g, t)$ is interpreted as resulting from individuals existing at the same time in a physical space, in the infinite population limit. Because of this, individuals can physically exchange genetic material. In general, we will not require such a strict physical interpretation of the EE probability distribution, and will rather use it as an abstract probabilistic tool, see later in ch. 5.

Figure 4.2: Relationship between exploration - exploitation (EE) dynamics, evolutionary dynamics (evol) and genetic algorithms (GAs). On the left, exploration mechanisms (mutation μ , light brown and recombinations r , brown); on the right, exploitation (selection φ , gray). The EE dynamics (area within thick black line) generalise the evolutionary process (circle) without recombinations. Alternatively, they can be regarded as GAs (rectangle) without recombinations. Dotted area indicates the "interpretability" as a real-world process. Recombinations are only interpretable within the evolutionary context.

evolution with caution, keeping what is essential and leaving back the rest, fig. 4.2.

4.2 The math of simple scenarios

Where from the investigation of simple case studies, general characteristics of the EE dynamics are unveiled. Pen and paper, let the theorist theorizing.

The graph dynamic (4.4) is not a model for anything in particular, it is a framework, or a theory for a class of phenomena [12]. In order to specify a model we must specify the context-specific representation of the notion of biological function. Having set the theoretical stage, much of our role as modellers boils down to one simple, crucial question: *what is F?*

As a first step, it is important to disentangle the intrinsic characteristics of the EE dynamic from those originating from the complexity of F . To do this, we start by studying (4.4) with models of F of minimal complexity, amenable to analytical investigation¹⁰.

10: The reader who is not interested in the more mathematical aspects can skip to the end of this section, read the remark 4.2.1 and go on. You are welcome!

Preliminaries

The following calculations simplify if, instead of the dyadic bit-like variables $a_{ij} = \{0, 1\}$ we use the equivalent representation with spin-like variables $\sigma_{ij} = \{-1, 1\}$. The two are related by the following¹¹:

$$\sigma_{ij} = 2a_{ij} - 1, \quad a_{ij} = \frac{1 + \sigma_{ij}}{2}. \quad (4.6)$$

11: Note that the analogy between graphs and spin systems is such that the dyads and not the nodes are equivalent to spins in classical statistical mechanics.

Analogous relations hold between the *average graph density* $d \in [0, 1]$ and the *average magnetisation* $m \in [-1, 1]$:

$$m = \frac{1}{L} \sum_{i < j} \langle \sigma_{ij} \rangle = 2d - 1, \quad d = \frac{1}{L} \sum_{i < j} \langle a_{ij} \rangle = \frac{1 + m}{2}. \quad (4.7)$$

In this section, we will consider (4.4) in the continuous time limit $\Delta t \rightarrow 0$, implying that $\Delta t \varphi F \ll 1$. The graph dynamics can be then described by the following master equation:

$$\frac{d}{dt} P(G, t) \stackrel{(a)}{=} \mu \sum_{i < j} [P(M_{ij}G, t) - P(G, t)] + \varphi [F(G) - \langle F \rangle_t] P(G, t). \quad (4.8)$$

In (a) we have used $e^{\pm x} \sim 1 \pm x$ for $x \sim 0$.

The dynamic of the expected value (ensemble average) of any graph observable $O : \mathcal{G} \mapsto \mathbb{R}$ and its time-dependent probability distribution can be calculated by (3.5) and (3.6), respectively.

4.2.1 No exploitation

A trivial scenario is the one in which exploitation is turned off, $F = \text{const}$, let us discuss it briefly. Due to the influx of random dyadic inversions $M_{ij}\sigma_{ij} \rightarrow -\sigma_{ij}$, any initial graph structure is eventually corrupted and the system slides towards randomness¹². The rapidity of this process is

12: On the meaning of *randomness* for graphs, see also s.n. 19

tuned by μ . Formally,

$$\begin{aligned} \frac{d}{dt} \langle \sigma_{ij} \rangle_t &\stackrel{(a)}{=} \sum_G \sigma_{ij} \mu \sum_{k<l} [P(M_{kl}G, t) - P(G, t)] \\ &= \mu \left[\sum_G \sigma_{ij} P(M_{ij}G, t) - \sum_G \sigma_{ij} P(G, t) \right] \\ &\stackrel{(b)}{=} -2\mu \langle \sigma_{ij} \rangle_t . \end{aligned} \quad (4.9)$$

In (a) we have used (4.8). In (b) we took advantage of the symmetry of the spin-like representation: $\sum_G \sigma_{ij} P(M_{ij}G, t) = \sum_G -\sigma_{ij} P(G, t)$.

The latter is the differential equation of an exponential decay with characteristic time $(2\mu)^{-1}$. The solution is straightforward, $\langle \sigma_{ij} \rangle_t = e^{-2\mu t} \langle \sigma_{ij} \rangle_{t_0}$. Since the dynamics of each σ_{ij} are independent, the magnetization also exhibits the same exponential decay behavior,

$$m_t = e^{-2\mu t} m_{t_0} . \quad (4.10)$$

In terms of graph density, this implies that under the action of exploration alone, the average state of the system melts down in an Erdős-Rényi random graph with connection probability $p = 1/2$.

4.2.2 Energy-like biological function

By (4.3), much of the complexity of the F metric arises from that of the graph state space \mathbb{R}^r . A simple, non-trivial state space is the one that represents each graph by its number of edges, i.e., $x(G) = \sum_{i<j} a_{ij} \in \mathbb{N}$. Therefore, let us consider the following

$$F(G) = -\frac{1}{L} \sum_{i<j} a_{ij} , \quad (4.11)$$

where each existing edge implies a fixed penalty¹³. The scenario considered is one in which the existence of any possible edge in the graph representation of the system is detrimental. In app. A, we plug (4.11) in (4.8) to derive the following dynamic for the ensemble average of the ij spin variable:

$$\frac{d}{dt} \langle \sigma_{ij} \rangle_t = -2\mu \langle \sigma_{ij} \rangle_t - \frac{\varphi}{2L} \sum_{k<l} [\langle \sigma_{ij} \sigma_{kl} \rangle_t - \langle \sigma_{ij} \rangle_t \langle \sigma_{kl} \rangle_t] . \quad (4.12)$$

The dynamics of the average spin variables are now coupled¹⁴. The coupling term is the sum of a row of the spin covariance matrix C_t where $(C_t)_{\sigma_{ij}, \sigma_{kl}} = C_0 v_t(\sigma_{ij}, \sigma_{kl}) = \langle \sigma_{ij} \sigma_{kl} \rangle_t - \langle \sigma_{ij} \rangle_t \langle \sigma_{kl} \rangle_t$.

To proceed, we restrict ourselves to the case where the covariance matrix has an approximately diagonal form, i.e.,

$$(C_t)_{\sigma_{ij}, \sigma_{kl}} \sim \mathcal{O}(\epsilon) \quad \text{for } \sigma_{ij} \neq \sigma_{kl} . \quad (4.13)$$

Discarding all terms $\mathcal{O}(\epsilon)$ in (4.12), we get:


$$\frac{d}{dt} \langle \sigma_{ij} \rangle_t \sim -2\mu \langle \sigma_{ij} \rangle_t - \frac{\varphi}{2L} [1 - \langle \sigma_{ij} \rangle_t^2] . \quad (4.14)$$

The latter is valid for $L\epsilon \ll 1$, which means either small graph sizes (small L) or mild functional pressures ρ , for which the dynamics are close

13: The case of a fixed benefit is equivalent, modulo a minus sign in (4.11) and those that follow from it.

14: This may seem odd at first, since F contains only single dyadic variables. For instance, a Gibbs-Boltzmann distribution with an Hamiltonian of the form (4.11), factorises in the spin variables and gives decoupled expected values. Accordingly, an MCMC dynamic based on such probability distribution does not introduce correlations between the spin variables. Where do the coupling come from? While the F metric (4.11) is "energy-like", the dynamic (4.8) is radically different from an MCMC dynamic in an energy landscape. It is the term $\langle F \rangle$ in (4.8) that is responsible for the coupling, since it contains an information about all dyadic variables, cf. (4.11), app. A.

Figure 4.3: EE dynamic with an energy-like F metric (4.11), under decoupling approximation (4.13). Here, $N = 10$ ($L = 45$), $\bar{\mu} = 1/L$, $\bar{\rho} = 2 \times 10^2$, $d_0 = 0.9$. Inset: plot of $\dot{m} = f(m)$, where f is as in (4.14). There are two fixed points (4.14), i.e., solutions of $f(m) = 0$: m_1 (red, unstable) and m_2 (green, stable). Shaded area corresponds to the inaccessible regions, $m \in [-1, 1]$. The initial condition m_0 (white diamond) and the one-dimensional vector field (green arrows) are indicated. Main: dynamic of the average graph density d_t , obtained from (4.15) and (4.7). The value of the asymptotic state d_2 solely depends on L and $\bar{\rho}$ (held fixed) while the rapidity of the approach to it is tuned by the exploration rate μ .



15: We stress that the independence of the dyad dynamics results here from restricting to a specific regime. In the absence of information about the structure of the covariance matrix, the decoupling approximation can be verified *a posteriori*, as we will do in sec. 4.3.1.

16: (4.15) describes exhaustively the dynamics in the graph state space, this is what we call a *solution* of the EE dynamic. The time course of any other graph feature, by construction, must be a function of (4.15).

17: A straightforward generalisation of this discussion is the one for the case of a F metric written as a general edge covariate (2.26), i.e.,

$$F(G) = \frac{1}{L} \sum_{i < j} \gamma_{ij} a_{ij}. \quad (4.17)$$

Following the same steps, we get an expression which is analogous to (4.14) but with the substitution $\varphi \rightarrow \varphi \gamma_{ij}$. In this case, however, even under decoupling approximation and the same initial conditions, the dyadic dynamics are different, since each is subject to a dyad-specific exploitation rate $\gamma_{ij} \varphi$. It is not possible to trade the dynamic equation of any $\langle \sigma_{ij} \rangle$ for that of m_t .

18: Worse, exponentially challenging. This is one of the many manifestations of the same, well-studied problem in the physics of disordered systems, namely the (NP-hard) problem of finding the ground state in a spin-glass landscape. [85].

to those in sec. 4.2.1 (small ϵ). By (4.14), the dynamics of the dyads are decoupled and this is the reason why we refer to (4.13) as the *decoupling approximation*¹⁵. If we use the same initial conditions $\langle \sigma_{ij} \rangle_{t_0} = \sigma_0 \forall i, j$, the same differential equation (4.14) holds for the magnetisation, since $m_t = \langle \sigma_{ij} \rangle$. It can be explicitly integrated by partial fractions, the solution being:

$$m_t = m_2 \left[1 + \frac{m_1/m_2 - 1}{1 + \frac{m_1 - m_0}{m_0 - m_2} e^{2\mu t \sqrt{1 + (2L/\rho)^2}}} \right], \quad (4.15)$$

where $m_0 = m_{t_0}$ and

$$m_1 = 2L/\rho + \sqrt{1 + (2L/\rho)^2}, \quad m_2 = 2L/\rho - \sqrt{1 + (2L/\rho)^2} \quad (4.16)$$

are the fixed points of the dynamic, fig. 4.3. The one described by (4.15) is a relaxation dynamic to the stable fixed point m_2 ¹⁶. Regarded as a function of ρ , we find consistently $m_2 \rightarrow 0$ for mild functional pressures ($\rho \rightarrow 0$), which corresponds to the case of sec. 4.2.1. In the opposite limit of perfect exploitation ($\rho \rightarrow \infty$), we have $m_2 \rightarrow -1$, corresponding to empty graphs¹⁷.

As for the general characteristics of the dynamic (4.15), we find that the value of the stable fixed point m_2 depends only on the ratio between the number of degrees of freedom (dyads) L and the functional pressure ρ ; the rapidity of the approach to the stable fixed point depends on the exploration rate μ , with higher values yielding faster approaches; for any $\rho < \infty$ the asymptotic value is not attainable, as m_2 balances the strengths of the exploration and exploitation drivers.

4.2.3 Distance-like biological function

The problem of determining the global maximum of an F metric can be challenging¹⁸. As theorists, however, we have the right to flip the script, with pen and paper. We can choose an optimal state and then shape our F metric to suit it. A simple way to do this is to formulate F as a




Figure 4.4: EE dynamic with a distance-like F metric (4.19), under decoupling approximation (4.20). Here, $N = 10$ ($L = 45$), $\bar{\mu} = 1/L$, $\bar{\rho} = 5 \times 10^2$, $d_0 = 0.9$, $E^* = 3$. Inset: plot of $\dot{m} = f(m)$, where f is as in (4.14). There are three fixed points, i.e., solutions of $f(m) = 0$: m_1 , m_3 (red, unstable) and m_2 (green, stable). Shaded area corresponds to the inaccessible regions, $m \in [-1, 1]$. The initial condition m_0 (white diamond) and the one-dimensional vector field (green arrows) are indicated. Main: dynamic of the average graph density d_t , obtained from m_t using (4.7). The value of the asymptotic state d_2 solely depends on L and $\bar{\rho}$ (held fixed) while the rapidity of the approach to it is tuned by the exploration rate μ .

distance function in the graph state space. If $x : \mathcal{G} \mapsto \mathbb{R}^r$ defines the state space and $x^* \in \mathbb{R}^r$ is (by definition) the optimal state, we can use, e.g., a squared distance¹⁹

$$F(G) \propto -[x(G) - x^*]^2. \quad (4.18)$$

Regardless of the complexity of the state space, the above F metric has a simple structure: there is a single, global maximum corresponding to those graphs with the same statistics as x^* . However, as in the previous section, we consider the simple, one-dimensional state space of the number of edges $x(G) = \sum_{i < j} a_{ij} \in \mathbb{N}$, which is amenable to analytical investigation. In particular, we consider

$$F(G) = -\frac{1}{L^2} \left(\sum_{i < j} a_{ij} - E^* \right)^2, \quad (4.19)$$

where $0 < E^* < L$ is by construction the optimal number of edges in a graph. In app. A, we follow similar steps as in sec. 4.2.2 to derive an approximate dynamic equation for the magnetisation, under decoupling approximation²⁰. As a result, we get

$$\dot{m}_t = -2\mu m_t - \frac{\varphi}{L^2} (1 - m_t^2) \left[\frac{L-1}{2} m_t + \frac{L}{2} - E^* \right]. \quad (4.21)$$

The latter can be integrated numerically, fixing the boundary conditions $m_{t_0} = \sigma_0$. Qualitatively, we find the same behaviour as in sec. 4.2.2, fig. 4.4. It is worth highlighting it:

Remark 4.2.1 (Characteristics of the EE dynamics) The EE dynamics (4.8) is such that (i) the asymptotic state ($t \rightarrow \infty$) depends on the functional pressure ρ and approaches $\max F$ in the limit of perfect exploitation²¹ $\rho \rightarrow \infty$; (ii) for fixed ρ , higher exploration rates μ correspond to faster approaches to the asymptotic state.

While the above findings have the robustness of an analytical result, they

19: Simple alternatives are the absolute value or the distance (l^2 -norm). The choice of a squared distance, however, simplifies the math below.

20: In particular, we generalise here the definition of decoupling approximation to the regime where

$$\langle \sigma^{(1)} \dots \sigma^{(k)} \rangle \sim \prod_{i=1}^k \langle \sigma^{(i)} \rangle, \quad (4.20)$$

where the left-hand side contains no repeated spin variables. The approximate dynamic (4.21) is expected to hold for $L\epsilon \ll 1$ where ϵ is the order of magnitude of spin correlations. The validity of the hypothesis will be assessed *a posteriori*, see sec. 4.3.1.

21: Note that, (a) for any nonzero exploration rate $\mu > 0$, the $\max F$ cannot be reached exactly and (b) for $\mu = 0$ there would be no dynamics at all. The limit of perfect exploitation should be intended as the limit of $\varphi \rightarrow \infty$ for $\mu \neq 0$, cf. (4.8).

were obtained under the rather limiting assumption of the decoupling approximation. Therefore, our conclusions need to be assessed numerically, which will be done in sec. 4.3.1.

4.3 Population-based simulations

Where a simulation scheme for the EE dynamics – inspired, once again, by evolution – is deployed and its main features presented. It is to be used there where the pen cannot get.

22: Here, we discuss the general features of the simulations. A detailed description of the design of the code can be found as documentation in the GitHub folder.

23: In general, one has two possible strategies for simulating forward population dynamics: (i) tracking the number of individuals (or frequency) associated with each possible graph, or (ii) tracking the graphs associated with the individuals present in the population. We opt for the latter, since the former requires listing and tracking all 2^L possible graph configurations, which quickly becomes infeasible as L increases.

24: In the case of an edge toggle, one has $\sigma_{ij} \rightarrow -\sigma_{ij}$. Later in ch. 5, we will use growth-only dyadic mutations, i.e., $\sigma_{ij} \rightarrow |\sigma_{ij}|$.

25: In practice, the simulation step can always be defined as $\Delta t = 1$ by rescaling accordingly:

$$T \rightarrow vT \quad \mu \rightarrow \mu/v.$$

The computational problem of setting up simulations for equations of the form (4.4) has recently been addressed, e.g., by [127, 129], in the context of evolutionary models. We develop a similar computational framework, coded in Python 3.9.7 and freely available on the GitHub folder `EE-graph-dyn`²².

Once again, the core idea of the simulations is to mimic the evolutionary process by simultaneously tracking the dynamics of an entire population of individuals. Each individual is associated to a graph, i.e., to a binary strings with L bit-like entries $a_{ij} = 0/1$ (dyads). Our population-based simulations keep track of all the individuals *existing* within the population, at each time t ²³. To speed up the simulations, we group similar individuals into a *clone*, that is a pair (G, n) , where n is the number of individuals associated with the same graph G . At time t , the population is thus defined as the set of existing clones $\mathcal{P}(t) = (\mathbf{G}(t), \mathbf{n}(t))$. The population size (total number of individuals) $\sum_{\alpha} n_{\alpha}(t) = M$ is held fixed while the total number of clones $M_c(t) \leq M$ fluctuates.

At each simulation step, the population is updated, new clones are created by dyadic mutations (exploration), their size updated by functional selection (exploitation), fig. 4.5. In particular:

- ▶ *Exploration*. Each dyad of each individual in the population mutates²⁴ with probability $1 - e^{-\Delta t \mu} \sim \Delta t \mu$. The exploration rate μ is uniform across dyads.
- ▶ *Exploitation*. The F values of the graphs associated to each clone are computed. The clone sizes are then updated by extracting M independent samples from a multinomial distribution where each graph G_{α} is selected with probability

$$p_{\alpha} = n_{\alpha} e^{\Delta t \varphi F(G_{\alpha})} / \sum_{\beta} n_{\beta} e^{\Delta t \varphi F(G_{\beta})}, \quad \alpha \in 1, \dots, M_c(t). \quad (4.22)$$

Our simulations have six parameters, summarised in tab. 4.2. The structural parameters N, T set the geometry of the simulations. The former is the (fixed) number of nodes of each graph; the latter is the size of the time window to be simulated. There are two internal degrees of freedom: the population size M and the time step Δt – for technical convenience, it is often preferable to set the inverse time step $\nu = \Delta t^{-1}$. Finally, two parameters control the dynamics of the system, the exploration rate μ and the relative strength of exploitation ρ ²⁵.

Algorithm: EE graph dynamics, forward simulations (pseudocode).

$\mathcal{P}(0) = (\mathbf{G}_0, \mathbf{n}_0)$

$t = 0$

while $t < T$ **do**

Exploration: $\sigma_{ij} \rightarrow -\sigma_{ij}$ with probability $\Delta t \mu \forall (i, j), \forall G_\alpha$

update $\mathcal{P}^* = (\mathbf{G}^*, \mathbf{n}^*)$

compute $F(G_\alpha^*) \forall G_\alpha^*$

Exploitation: M draws from a multinomial distribution with

$p_\alpha = n_\alpha^* e^{\Delta t \varphi F(G_\alpha^*)} / \sum_\beta n_\beta^* e^{\Delta t \varphi F(G_\beta^*)} \Rightarrow$ compute new counts \mathbf{n}^{**}

set $\mathcal{P}(t) = (\mathbf{G}^*, \mathbf{n}^{**})$

$t += \Delta t$

| Parameter | Description |
|-----------|-----------------------------------|
| N | number of nodes |
| T | time window span |
| M | population size |
| ν | inverse time step Δt^{-1} |
| μ | exploration rate |
| ρ | functional pressure φ/μ |

Table 4.2: Parameters of simulations for EE graph dynamics. Our computational framework has six degrees of freedom, which we group by color: structural parameters (top), internal degrees of freedom (middle) and parameters of the dynamics (bottom).

The running time of a single simulation has an obvious linear scaling with the inverse time step ν , since the same operations described above are repeated a number νT of times. A linear scaling is also observed with the population size M , which is reasonable since both the evaluation of dyadic mutations and of the F metric have to be performed independently for each clone – in the worst case, $M_c(t) \sim M$, fig. 4.6 (left). The dependence on the number of nodes N (or equivalently, on L) is trickier, since it depends strongly on the complexity of the graph operations involved in the evaluation of F . For a F metric as simple as (4.11), and for large M , we observe an approximately linear scaling in L , fig. 4.6 (right).

4.3.1 Gleaning dynamics from simulations

At each time t , the raw information provided by the simulations comes in the form of a snapshot of the population $\mathcal{P}(t) = (\mathbf{G}(t), \mathbf{n}(t))$. Assuming that it is a representative sample of the entire probability distribution, we can compute from it the distribution of any graph observable $O : \mathcal{G} \mapsto \mathbb{R}$ by adapting (3.6) to

$$P(O, t) = \frac{1}{M} \sum_{\alpha=1}^{M_c(t)} n_\alpha(t) \delta[O - O(G_\alpha(t))], \quad (4.23)$$

where δ is the Dirac-delta ($\int dO \delta(O) = 1$). By consequence, the expected value of O at time t is

$$\langle O \rangle_t \sim \frac{1}{M} \sum_{\alpha=1}^{M_c(t)} n_\alpha(t) O(G_\alpha(t)). \quad (4.24)$$

To conclude, let us showcase an example of EE graph dynamics simulations. Let us reconsider the case discussed in sec. 4.2.3, i.e., an EE dynamics with distance-like F metric (4.19), fig. 4.7.

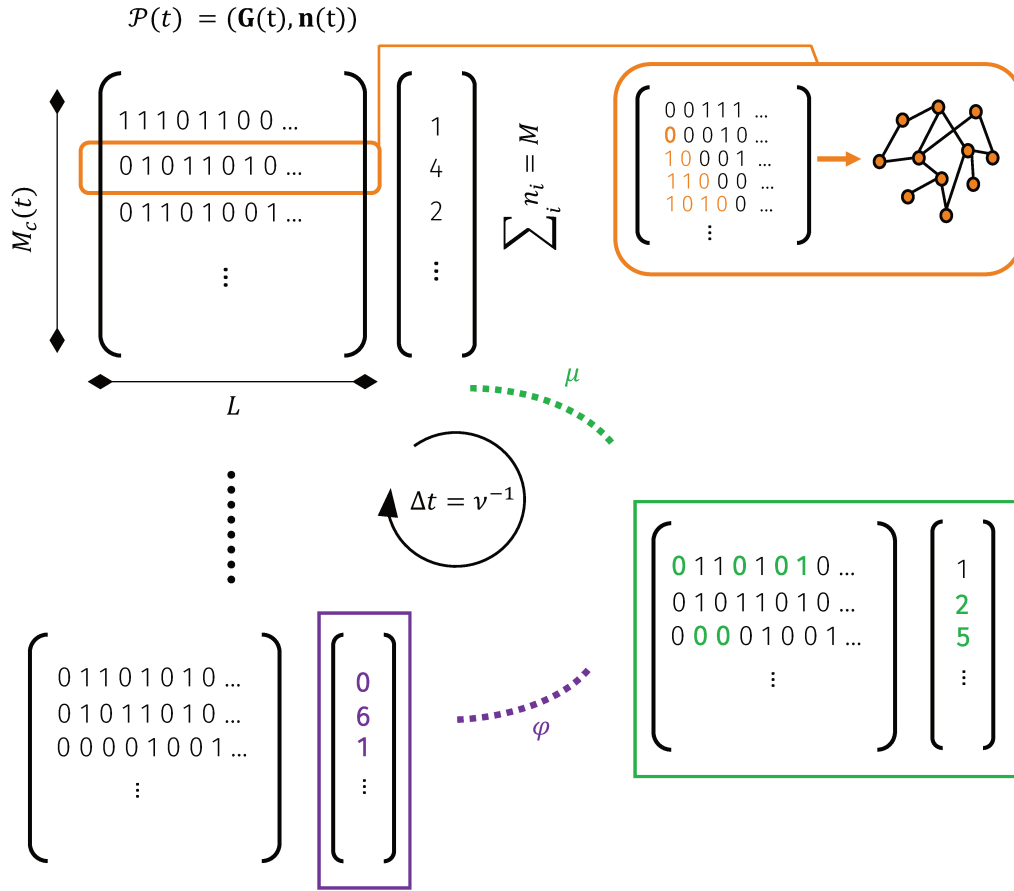
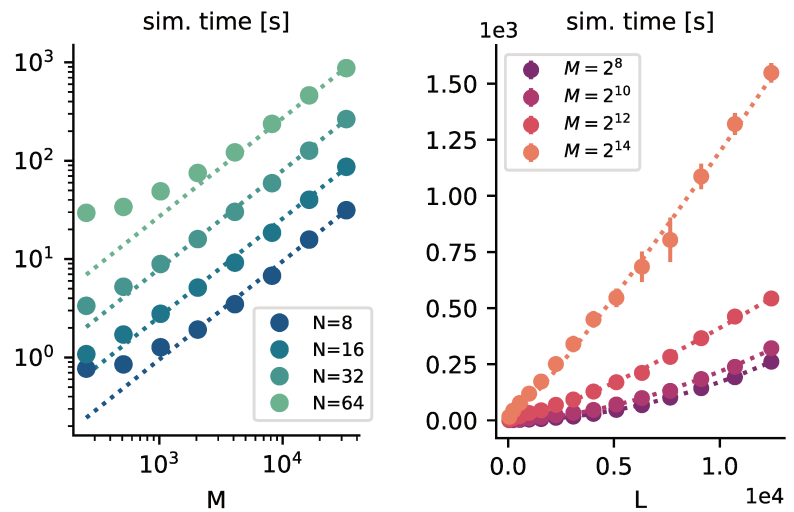


Figure 4.5: Population-based simulations for eq.(4.4), simulation step Δt . A population $\mathcal{P}(t)$ at time t is made up of a matrix $G(t)$ with dimensions $M_c(t) \times L$ along with a vector $\mathbf{n}(t)$ of length $M_c(t)$ (top left). Each row of $G(t)$ corresponds to a unique graph configuration, specifically, to the vectorised lower triangular matrix of an adjacency matrix (top right). First, the configuration space is explored by introducing random dyadic mutations (edge toggles) at the rate μ in each individual of the population (green, bottom right). As a consequence, both the matrix G and the counts \mathbf{n} change. As new configurations are found, the number of clones M_c increases. Second, the more functional graphs proliferate within the population. The exploitation alters only the count vector \mathbf{n} and is regulated by the parameter ϕ . Those clones that end up with zero-size are removed from G before the next time-step. The total number of individuals M and time interval Δt are free internal parameters of the simulations.

Figure 4.6: Simulation (running) time as a function of the population size M and number of dyads L . Dots are average running times obtained from 10 simulation runs, error bars are indicated when visible. Here $T = 100$, $v = 1$, $\mu = 1.0 \times 10^{-4}$, $\rho = 5$, an F metric (4.11). (Left) The sim. time scales linearly with M (log-log scale). Dotted lines are the a linear fit with the curve $y = aM$, where a is a parameter. (Right) The sim. time has a mild exponential dependence on L . Dotted lines are a linear fit with the curve $y = aL^b$, where a, b are parameters. We find $b = 1.95, 1.71, 1.29, 1.12$ for increasing M . In the large M limit, we expect $b \rightarrow 1$, indicating an approximately linear dependence on L . The dependence on L however, strongly depends on the complexity of the F metric.






Figure 4.7: Exploration-exploitation (EE) dynamic (4.4) with distance-like F metric (4.19). Parameters: $N = 10$, $\bar{\mu} = 1/L$, $\bar{\rho} = 5 \times 10^3$, $T = 50$, $\Delta t = 1/\nu = 10^{-2}$, $E^* = 3$, $\Delta t = 1$, $M = 4096$. (a) Distribution of F values at three different time points t , the corresponding $\langle F \rangle_t$ are indicated (white bars), cf. fig. 4.1. Exploitation amplifies the probability of those graphs exhibiting higher F values than the ensemble average. Consequently, the distribution shifts towards higher F values, until the max is reached (red dotted line). (b) The average graph density $d_t = \sum_{i < j} \langle a_{ij} \rangle / L$ exhibits a relaxation dynamic towards the steady-state value d_∞ , in proximity to the target density E^*/L . Shaded area indicates the 95% confidence interval, black dashed line represents the analytical solution from (4.19). (c) The value d_∞ depends on the functional pressure ρ , the target is approached for $\rho \rightarrow \infty$. The rapidity of the approach to the steady state value depends on ρ and, for fixed ρ , on the exploration rate μ (not shown). (d) The distance from the target diminishes for increasing exploration rate μ and the exploitation rate $\varphi = \rho\mu$. (e) The value d_∞^{dec} is invariably nearer to the target than d_∞ , with the difference between the two tapering off for increasing φ – where both converge towards E^*/L .

We use (4.23) to evaluate the instantaneous distribution of the F values, at three illustrative time points, fig. 4.7(a). As a result of the exploitation, these distributions are skewed towards the high- F direction, the more so the higher the functional pressure. The dynamics in the one-dimensional F corroborate the pattern we have depicted in fig. 4.1. We further use (4.24) to evaluate the time dynamics of the average graph density d_t , as defined in (4.7). We find the same qualitative behaviour as for the analytical solution discussed in sec. 4.2.3. Specifically, after an initial transient phase, a steady-state value is attained, striking a balance between the strengths of exploration and exploitation, fig. 4.7(b). This steady-state value is moderated by the functional pressure ρ – the higher the pressure, the closer the density approaches the target value, fig. 4.7(c-d).

Finally, we perform a comparison between our simulations and the numer-

ical solution we obtained from (4.21), under decoupling approximation. We present evidence that the asymptotic state of the latter consistently lies intermediate to the target density and the simulated dynamics, the difference between the three vanishing for increasing functional pressure, fig. 4.7(e). We conclude that the decoupling approximation universally exhibits qualitative agreement with simulations across the parameter space and quantitative agreement for large values of ρ .

This result, in turn, allows us to regard at the characteristics of the EE dynamics derived under rather specific conditions (Remark 4.2.1) as universal attributes of the EE dynamics. The next natural step is to direct our attention towards more realistic systems with more complex functional landscapes.

Weaving the mind of a worm

5

With four parameters I can fit an elephant,
and with five I can make him wiggle his trunk.

— John von Neumann (attributed)

The exploration-exploitation paradigm is general and theoretically sound for biological dynamics. Therefore it should apply elsewhere than in the evolutionary context. The aim of this chapter is to show that it does apply elsewhere. Between the *should* and the *does* there are a number of theoretical and methodological details to be worked out, approximations to be made along the way and, most importantly, plenty of biology to be learned.

The beginning of wisdom – as they say – is the definition of terms¹. The process we analyse in this chapter is the development of a natural nervous system, from birth to adulthood – or, as it has been called, the *brain wiring dynamics*. To do so, we shall focus on a specific organism, for which a natural (and almost obliged) choice is a tiny, transparent worm, the nematode *Caenorhabditis elegans*, or simply *C. elegans*.

We found it convenient to organise this chapter as a long-form scientific paper². We begin with a general discussion of the problem (sec. 5.1) and a description of the essential features of the organisation of the worm nervous system, including the data we use (sec. 5.2). We then turn to describe in detail our EE model of the *C. elegans* brain maturation, including our main results (sec. 5.3). We provide a plausible interpretation of how the exploration-exploitation dynamics could be implemented at a fine-scale (sec. 5.4) and conclude with an overview of the many possible generalisations of the model presented (sec. 5.5).

Main reference

📖 Vito Dichio & Fabrizio De Vico Fallani. *The exploration-exploitation paradigm for networked biological systems*. In: arXiv e-prints 2306.17300 (2023) [1].

5.1 The brain wiring problem

Where the (genetically encoded) growth of a brain – furiously debated among neuroscientist – is recognised as a specific instance of the EE dynamics. Where also the fundamental facts about the brain are: developmental variability, functional robustness.

The first occurrence of the wording appears in a recent perspective article by Hassan and Hiensinger [177]. The first lines read:

| | | |
|-----|--------------------------|----|
| 5.1 | The brain wiring problem | 47 |
| 5.2 | The mind of a worm . . . | 50 |
| 5.3 | The EE brain maturation | 53 |
| 5.4 | Interpretation | 65 |
| 5.5 | Perspectives | 67 |

1: This is often attributed to Socrates, but – if I understand correctly – there is no such quotation in the writings of Plato, Socrates' press office. To err on the side of caution, and to avoid offending any Greek philosophers who might pass through here, we refer to a generic "they".

2: In *Physical Review Letters* – a flagship publication for physicists – typical articles span just four or five pages. This concise format assumes that readers have a considerable amount of knowledge, as details are often distilled for brevity. Here, I do not. Our journey to the results might be a longer read, but it is supposed to be pedagogical. The main reference of this chapter [1], on the other hand, is written in a short format, for the already expert, for the impatient, or simply for the lazy.

3: *Wires, circuits, junctions* (...) The language of neuroscience is steeped in the vocabulary of telecommunications. This is a long-standing metaphor that goes back to the days of the telegraph [178]. As early as 1875, the German physicist Hermann von Helmholtz wrote: "Nerves have often and not unsuitably been compared to telegraph wires. Such a wire conducts one kind of electric current and no other; it may be stronger, it may be weaker, it may move in either direction; it has no other qualitative differences. Nevertheless (...) we can send telegraphic dispatches, ring bells, explode mines, decompose water, move magnets, magnetise iron, develop light, and so on. So with the nerves." [179]

The brain, as we neuroscientists like to say, is *really* complex. A good deal of our efforts are therefore dedicated to figuring out just how this apparent complexity is generated: where does the information to build a brain come from, and how is such information turned into synapse-specific wiring? We call this the "brain wiring problem".

In the rest of this manuscript, we will work alongside our fellow neuroscientists and confront the same compelling challenge. The brain wiring³ dynamic is a developmental dynamic that unfolds during a lifespan and involves the formation, growth and establishment of an individual's nervous system – primarily, neurons and synapses, fig. 5.1.

The crucial empirical observation is that, although the functional outcomes are highly reproducible and almost invariable, the nervous system is not *hardwired*. From worms to humans, neuroanatomical differences are observed between the nervous systems of any two individuals, even when they are genetically identical and even when environmental factors are controlled [180]. For instance, the branching patterns of neuronal connections vary in lower isogenic animals such as worms, grasshoppers and locusts [181–183] but also in mammals, including monozygotic human twins [184]. This form of stochasticity, which (i) is not due to genetic differences, (ii) is not induced by the environment, and (iii) nevertheless leads to equally functional outcomes, has been referred to as *genetically encoded* [177] or *intrinsic chance* [185]. Much of the solution to the brain wiring puzzle lies in answering the question: what is the origin of genetically encoded stochasticity?




Figure 5.1: Four common type of neuron: (a) unipolar, (b) bipolar, (c-d) multipolar, in particular: motor neuron (c) and Purkinje cell (d). The *cell body* (soma) is the neuron's core and houses the nucleus. This is where the neuron's fundamental metabolic activities occur. Extending from the cell body are *dendrites*, which are tree-like projections that receive signals from other neurons and transmit them to the cell body. The *axon*, on the contrary, is a long slender projection, that serves to transmit electrical impulses away from the cell body. This allows the neuron to communicate with other neurons, muscles or glands. The contact zone between dendrites and axons where two cells exchange chemical and electrical signals are called *synapses*. Together, these components form the communication network within the nervous system. The examples here proposed showcase the great variety of neuronal geometries found in Nature. Adapted from Kandel et al., *Principles of neural science* [186].

One possibility is that the observed variability is noise of a molecular code. In a nutshell, genes encode molecules, molecular mechanisms

drive the growth of the nervous system⁴, neuron by neuron, synapse by synapse. An appropriate spatiotemporal regulation of the latter, in turn, results in the synapse-specific wiring of the brain [187–189]. This specification process is noisy and occasionally results in inaccurate outcomes. Whenever such misspecifications do not impair the system's functionality, they reveal themselves as variability in the observed systems. A computer scientist faced with such a brain wiring algorithm would probably be appalled. Indeed, programming for each input/output is a highly inefficient coding strategy. Moreover, it seems implausible, at the very least, that a system as complex as the human brain ($\sim 10^{15}$ synapses) can be exhaustively specified by a single genome down to the finest spatiotemporal scale⁵.

A contrasting view has recently emerged: it is not the precise result, but the wiring algorithm that is genetically encoded [177, 190, 191]. Accordingly, neural circuits grow based on simple, genetically encoded, pattern formation rules⁶. The variability of the outcomes is not due to misspecified molecular instructions, but rather is an intrinsic and essential feature of the dynamics. This because the brain wiring is a stochastic process that generates patterns, and patterns can be realised in a variety of different ways. Our computer scientist would be relieved: encoding a finite set of (possibly simple) rules is certainly a less daunting programming task than fine-coding a nervous system. Indeed, from an algorithmic point of view, a stochastic process based on a few algorithmic constraints and otherwise random appears to be an efficient, flexible – and maybe ideal – way to explore an unknown environment.

It further follows from this view that the *functionality* is an attribute of the algorithm, rather than of the outcome: a functional rule-set is the one that leads to functional configurations of the nervous system. If the set of such configurations is large enough, the whole brain wiring process turns out to be robust, since small configuration changes do not affect the system's functionality [190]. In this sense, allowing for variability of the outcomes is an insurance policy against failure in the case of perturbations⁷.

Down to the neuronal scale, an experimental evidence consistent with this view is the fact that the synapse formation process is largely non-specific. The growth of each branch of a dendritic tree happens through a series of stochastic local decisions in an unknown molecular environment⁸. This allows the process to cope with unforeseen environmental conditions. For example, neurons that innervate incorrect target regions will form synapses wherever they land, regardless of how inappropriate the targets may be [186]. In the absence of other potential partners, they can even form perfectly functional synapses with themselves, known as autapses [194].

At this point, we cannot resist the temptation to draw a parallel with our exploration-exploitation framework. It is straightforward to rewrite the above paragraphs in the language of sec. 4.1: the brain wiring process is a self-referential (state-dependent) biological process, that unfolds as a random exploration of the configuration space under the action of a set of functional drivers and constraints. The dynamics (4.4) endowed with a choice for the F metric is precisely the way in which we specify a brain wiring algorithm, genetically encoded by assumption. The observed


4: The main way this happens is through so-called *guidance cues*, biochemical signals (molecules) that guide the growing axons of neurons.

5: This has been called the blueprint problem: a deterministic molecular code, accurate at every spatiotemporal scale, would be at least as complicated as the resulting wiring diagram [177].

6: These include, among others, spacing between axons, self-avoidance, lateral inhibition. For two recent examples in the *Drosophila* brain, see [192, 193].

7: The amount of variability is regulated by the algorithm itself: strict and/or complex functional requirements will yield a narrow distribution of outcomes, while simple wiring rules will allow for a broad outcome variability. This *degree of variability* is likely to be subject to evolutionary pressure and optimised by natural selection [190].

8: One might think at a branch of a dendritic tree as a navigator in a maze, who ignores both the maze map and the position of the other navigators. It only has algorithmic information of the sort: "at each crossroad, choose the wider path" or "if you see a lemon tree, turn around" or "if possible, avoid passing by the owl's nest" and so on. At each new point in the maze, the navigator makes decisions based on its rules and trying to accommodate constraints in the best possible way.



9: The running title comes from a bit of a lab inside joke, we all have some. The anecdote has been recently re-evoked: “We sent the text and figures to the publishers in a number of ring binders (no PDFs in those days). We had labelled these notebooks “the mind of a worm” in order to identify them on a shelf among all the other similar-looking notebooks. We were amused the printers had picked this up and used it as a running title, so we let the name stick.” [196]

10: At the time of the writing of this manuscript – summer 2023 – a preprint has appeared on bioRxiv with the reconstruction of a new, complete nervous system – that of the *Drosophila melanogaster*, ~ 500× larger than the one of the *C. elegans* [203]. Similar neural maps at the synapse level exist for other model systems at various stages of completion, including the mouse [204–206], the larval zebrafish [207], the tadpole larva [208]. The experimental technique employed is for the reconstructions is the serial section electron microscopy [209–211].

11: It is nothing short of impressive that such a tiny nervous system is able to support the wide behavioral array of a *C. elegans*. Beyond the basics of locomotion, foraging, and feeding, the worm can discern and navigate towards or away from various chemicals, odors, temperature gradients, and food sources. Furthermore, it demonstrates social awareness, detecting the presence, density, and even sex of neighboring nematodes [198, 212].

12: Each neuron in the worm’s nervous system is identified by a code. The nomenclature system consists in two or three letters (or, occasionally, numbers), followed by the position in worm’s body D/V (dorsal/ventral), R/L (right/left) [195]. For instance, the code AFDL stands for Amphid Finger-like Endings Dorsal Left.

variability of the nervous system is nothing but the statistical uncertainty associated with the *ensemble* probability distribution.

Below, we will follow this line of thought and tackle the problem of brain wiring with our EE dynamics. To work out the details, we need to focus on a specific context. The time has come to talk about worms.

5.2 The mind of a worm

Where a biological minimum is provided of the neuroanatomy of a tiny little worm. Where also the unique and recently published experimental data are presented, colorful brain networks are flashed.

It took thirteen years, the first time. But finally, in the 1986, Sydney Brenner and his collaborators published *The structure of the nervous system of the nematode Caenorhabditis elegans*, the first complete anatomical reconstruction of a natural nervous system, a major milestone in the history of neuroscience [195]. The report (over 300 pages) sported a rather solemn running title: *the mind of a worm*⁹.

The *Caenorhabditis elegans*, or *C. elegans*, is a small (~ 1 mm long), transparent nematode, or roundworm, that has been widely used as a model organism in biological research, from genetics to behaviour [197–202]. It is especially valuable due to its relative simplicity. This organism has held a unique position in the field of neuroscience since, until very recently¹⁰, it has remained the only one with a fully reconstructed connectome – i.e., a detailed charting of all its neural connections. Below, we provide a brief overview of the essential features of the *C. elegans* neural network. There are excellent resources available on the web for all aspects of the *C. elegans* biology, including the nervous system, in particular [Wormatlas](#) and [WormBook](#).

5.2.1 *C. elegans* nervous system: a digest

The nervous system of an adult hermaphrodite *C. elegans* consists of just 302 neurons, organised in two independent nervous system: a large somatic nervous system (282 neurons) and a small pharyngeal nervous system (20 neurons)¹¹, uniquely identifiable¹². The vast majority of neurons have simple unipolar or bipolar morphology, fig. 5.1 (a-b).

With few exceptions, neuron cell bodies are found in clusters, called *ganglia*. Several head ganglia – including the retrovesicular ganglion and ventral ganglion –, located around the nerve ring, host the largest collection of cell bodies. The second largest collection of cell bodies is found in the tail ganglia. Any projection from the cell bodies is generically called a *process* – predominantly, axons and dendrites. Neuronal processes extend from the ganglia and travel in longitudinal nerve bundles to different regions of the nervous system. Out of the total 302 neurons in the adult hermaphrodite, 180 project axons/dendrites into the nerve ring. The most prominent nerve bundles are the nerve ring, ventral nerve cord and dorsal nerve cord, Fig. 5.2.




Figure 5.2: The majority of neuron cell bodies are found in clusters, named ganglia, located in the head – VG, ventral ganglion; RVG, retrovesicular ganglion – and in the tail – PAG, pre-anal ganglion; DRG, dorsorectal ganglion; LG, umbar ganglion. Neuronal processes run in nerve bundles, the major is the nerve ring (head), the ventral and dorsal cord rung, that along the whole worm body. Adapted from Cook et al., *Whole-animal connectomes of both *Caenorhabditis elegans* sexes* [213].

The neurons in the *C. elegans* nervous system are classified into different categories based on their morphology, function, and connectivity.

- ▶ *Sensory* neurons are the primary receptors of environmental stimuli, ranging from temperature changes to chemical signals.
- ▶ *Interneurons* are information processors, responsible for transmitting signals between other classes of neurons.
- ▶ *Motor* neurons control the contraction and relaxation of muscles, on which they primarily form synapses.
- ▶ *Modulatory* neurons release neuromodulators, molecules that alter the activity of other neurons or influence the strength of the signals they send.

Neurons exchange information mainly through ~ 1500 electrical and ~ 5000 chemical synapses. The former, also called *gap junctions*, are specialized channels that directly connect the cytoplasm of adjacent cells, allowing various molecules, ions, and electrical impulses to pass between the cells. The latter, *chemical synapses*, function as specialised junctions that facilitate the one-way relay of chemical signals, or neurotransmitters, from a presynaptic to one or more postsynaptic cells. Chemical synapses are found between neighbouring processes. Therefore, it is the neighbourhood of the processes that predominantly determines the connectivity between neurons. The nerve ring hosts the highest density of these synapses, followed by the ventral and dorsal cord.

Recently, a number of studies have revised the original annotations of the hermaphrodite, adult *C. elegans* nervous system, updated its connectome, and measured it in with the tools of network science [213–216]. However, it is important to understand that, strictly speaking, *the C. elegans connectome does not exist*. This because (i) due to technical limitations, the whole animal connectome is constructed by patching together regions of the nervous system obtained from different animals. Some connections are not even observed, but inferred from the similarity patterns of certain regions [213]. Furthermore, (ii) although the overall structure is highly stereotyped, individual connectomes differ in detail due to natural developmental variability, sec. 5.1.

Given these limitations, and the availability of the adult connectome only, it is unsurprising that fewer studies have examined the growth of the *C.*

C. elegans neural network throughout development (from embryo to adult) [217–220]. This circumstance took a sharp turn in 2021, as we are about to discuss.

5.2.2 *C. elegans* brain maturation

13: After the cellular differentiation, the neuron body cells migrate across the worm's body to take up their final positions in adulthood. This happens in a highly stereotyped way. Either simultaneously or after migration, neurons extend their processes to find partners for synaptic connections. Note, however, that the two processes (neuron establishment and synapse formation) are distinct [222].

The lifecycle of *C. elegans* unfolds in a series of distinct stages [221]. The initial embryonic stage of *C. elegans* concludes with hatching (birth). The post-embryonic development consists of four larval stages ($L_1 - L_4$) and adulthood. Along with other biological structures, the nervous system of the worm matures during development. Neurodevelopment in *C. elegans* is characterised by a stereotyped cell lineage, resulting in consistent neural placement across individuals¹³. The formation of synapses, or synaptogenesis, and process outgrowth, underlie the specific connectivity patterns of the mature nervous system, which are enhanced through substantial remodeling during post-embryonic growth [220, 222].




Figure 5.3: *C. elegans* brain networks at different developmental ages. The embryonic stage (E, dotted line) terminates with the hatching (birth, 0 h, white circle). The post-embryonic stages include four larval ages ($L_1 - L_4$, dashed line) and ends with the onset of adulthood (red circle). The adult stage (A, solid line) lasts about 2 – 3 days. The dataset consists of eight snapshots (microscope icons), including one at birth 0 h, three $L_1 \sim 5$ h, 8 h and 16 h, one $L_2 \sim 23$ h, one $L_3 \sim 27$ h and two adults, both ~ 45 h. The brain consists in the nerve ring and ventral ganglion, see fig. 5.2. The nodes are neurons belonging to four different categories: inter- (red), modulatory (yellow), motor (blue) and sensory (pink) neurons, see tab. 5.1. Edges are here defined as synaptic connections, i.e., an edge exists between two nodes if at least one chemical synapse exists between them.

Table 5.1: List of the 180 neurons of the adult *C.elegans* brain (hermaphrodite, N2), as reported in [223]. Interneurons in red, modulatory in yellow, motor in blue, sensory in pink. We have marked with an asterisk* those neurons that were not present at birth. Each neuron in the worm nervous system is uniquely identified by a code, which consists in two or three letters (or, occasionally, numbers), followed by the position in worm's body D/V (dorsal/ventral), R/L (right/left) [195]. The left-right symmetry increases over time and reaches the ~ 90% in the adult brain.

| | | | | | | | | | | | | | |
|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|
| ADAL | ADAR | AIAL | AIAR | AIBL | AIBR | AINL | AINR | AIYL | AIYR | AIZL | AIZR | AVAL | AVAR |
| AVBL | AVBR | AVDL | AVDR | AVEL | AVER | AVJL | AVJR | BDUL | BDUR | PVCL | PVCR | PVPL | PVPR |
| PVR | PVT | RIAL | RIAR | RIBL | RIBR | RIFL | RIFR | RIGL | RIGR | RIH | RIML | RIMR | RIPL |
| RIPR | RIR | ADEL | ADER | AIML | AIMR | ALA | AVFL* | AVFR* | AVHL | AVHR | AVKL | AVKR | AVL* |
| CEPDL | CEPDR | CEPVL | CEPVR | DVC | HSNL* | HSNR* | PVNL* | PVNR* | PVQL | PVQR | RICL | RICR | RID |
| RIS | RMGL | RMGR | IL1DL | IL1DR | IL1L | IL1R | IL1VL | IL1VR | RIVL | RIVR | RMDDL | RMDDR | RMDL |
| RMDR | RMDVL | RMDVR | RMED | RMEL | RMER | RMEV | RMFL* | RMFR* | RMHL* | RMHR* | SIADL | SIADR | SIAVL |
| SIADR | SIBDL | SIBDR | SIBVL | SIBVR | SMBDL | SMBDR | SMBVL | SMBVR | SMDDL | SMDDR | SMDVL | SMDVR | URADL |
| URADR | URAVL | URAVR | ADFL | ADFR | ADLL | ADLR | AFDL | AFDR | ALML | ALMR | ALNL* | ALNR* | AQR* |
| ASEL | ASER | ASGL | ASGR | ASHL | ASHR | ASIL | ASIR | ASJL | ASJR | ASKL | ASKR | AUAL | AUAR |
| AVM* | AWAL | AWAR | AWBL | AWBR | AWCL | AWCR | BAGL | BAGR | DVA | FLPL | FLPR | IL2DL | IL2DR |
| IL2L | IL2R | IL2VL | IL2VR | OLLL | OLLR | OLQDL | OLQDR | OLQVL | OLQVR | PLNL* | PLNR* | SAADL | SAADR |
| SAAVL | SAAVR | SDQL* | SDQR* | URBL | URBR | URXL | URXR | URYDL | URYDR | URYVL | URYVR | | |

In 2021, the research into neural development has been boosted by the release of an unprecedented dataset. *Witoli et al.* have published the electron microscopy reconstruction of the *C. elegans* brain across different stages of the worm development [223].

More specifically, eight *C. elegans* – wild-type N2, hermaphrodite, isogenic, reared in the same environment – were selected for imaging at different post-embryonic stages¹⁴. These comprise one at birth 0h, three $L_1 \sim 5$ h, 8h and 16h, one $L_2 \sim 23$ h, one $L_3 \sim 27$ h and two adults, both ~ 45 h, fig. 5.3. The brain – i.e., nerve ring and ventral ganglion¹⁵ – of each specimen was entirely imaged by serial section electron microscopy. Each cell was identified based on its unique morphology and position, tab. 5.1. The totality of its chemical synapses and a subset of its gap junctions were manually annotated¹⁶. From birth to adult, the number of nodes increased from 161 to 180, that of chemical synapses from ~ 1300 to ~ 8000 . In contrast to mammals [224], synaptic pruning does not occur, and the removal of synaptic connections is rarely observed.

5.3 The EE development of a worm brain

Where the core results of this PhD are illustrated, a parsimonious white-box model of a worm brain development is formulated. Where some understanding is reached of a (seemingly complex, certainly fascinating) biological process.

Here, we will detail our exploration-exploitation model for the maturation of the *C. elegans* brain. Along the way, we mark any hypothesis and/or assumption with the icon (\triangle)_i¹⁷, and discuss them in sec. 5.5.

5.3.1 A minimal worm brain

Methods

To start with, we exclude gap junctions from our analysis (\triangle)₁ because they were only partially annotated in [223]. As for the chemical synapses,

14: The developmental age of each specimen is estimated using the known and stereotypical cell division pattern [221]. Thus, precise temporal annotation is not available.

15: As mentioned above, a large proportion of *C. elegans* neurons are located close to the brain or extend their processes and form synapses within it. However, it is worth stressing that, as in humans, the *brain not* the entire nervous system, fig. 5.2.

16: By focusing solely on the brain, one gains the ability to reconstruct it completely for an individual, without the need to stitch together segments from different specimens. This is important for the purposes of our investigation, because to capture individual variability, we prefer individual reconstructions to collages.

17: The ambition of this chapter is to speak out clearly every single assumption. This is why, every time we make one, we ring a bell (\triangle)_i (*i* counts the assumptions). We also try to resist the temptation to summon *common practices* or *established methods* to justify their use or, worse, to conceal them.

18: The contrary, however, is not true: a new synapse does not necessarily mean a new synaptic connection, in the case where it is added to and strengthens an existing one. This further implies that the biology of the connection removal process is not the same of that of the synapse elimination process. In our case, as we will see in sec. 5.3.3, this is no cause for concern.

a number of them can be found between each pair of neurons. The networks of chemical synapses are therefore weighted and directed. In this work, we transform them however into the form (1.1), i.e., we consider the unweighted $(\Delta)_2$ and undirected $(\Delta)_3$ networks of chemical synapses.

Casting weighted to unweighted networks entails casting the number of synapses between a pair of neurons in a binary state, which we call connection/non-connection. In particular, a (directed) connection exists from a presynaptic to a postsynaptic neuron if at least one synapse exist between the two, fig. 5.4. The dynamics of synaptic connections is a low-dimensional projection of the dynamics of synapses. Importantly however, the formation of a new connection implies (by definition) that of a new synapse. Therefore, the biology of the connection formation process is, loosely speaking, the same as that of the synaptogenesis process¹⁸. It has also been shown that both synapses and synaptic connections exhibit a qualitatively similar developmental dynamic, namely a near-linear rate of addition [223].





Figure 5.4: Multiple synapses in both direction can exist between two neurons. The network of chemical synapses is therefore weighted (W) (by the synapse number) and directed (D). A directed synaptic connection (UW,D) exists if at least one synapse is observed in the same direction. An undirected connection (UW, UD) exists if at least one directed connection is observed.

Figure 5.5: Triad census. Left: There exist sixteen directed graph patterns between three nodes. Each is denoted by a three-digit code, representing (i) the count of mutual links, (ii) of single links, (iii) and non-existent links, respectively. Additionally, (iv) a letter can be appended to indicate if the pattern has a cycle (C), a transitive (T), an upward (U), or a downward (D) connection structure. In bold red, the codes of those motifs that are over-represented in the adult *C. elegans* network of (directed) synaptic connections. With the exception of 030C, all motifs involving connections between each pair of nodes (green) are over-represented. Those involving one empty dyad (violet) tend to be under-represented [213, 214]. Right: There exist four unique connectivity patterns are possible among three nodes in an undirected graph. Color code highlights corresponding patterns.



19: In general, it can introduce spurious connection reciprocities and distort the network information flow – e.g., by obscuring causal relationships between different connections.

The further reduction to an undirected network means that an undirected connection is placed between two neurons if there is at least one directed connection (thus, a synapse) between them, regardless of its direction. This is a more delicate assumption¹⁹, which we can motivate as follows. Previous analyses [213, 214] of the (adult) *C. elegans* network of chemical

synapses have examined its triad census, i.e., the counts of all possible directed connection patterns between triples of nodes – there exist sixteen patterns. These studies have demonstrated that those connection patterns involving the same number of empty dyads exhibit consistent statistical characteristics, in the sense that they are all (with few exceptions) over- or under-represented relative to a randomised null model²⁰, fig. 5.5. The simple patterns listed in the triad census are the building blocks of more intricate network motifs. In turn, this means that, when it comes to the analysis of graph patterns, the use of undirected connections results in a minor distortion of the original directed network.

As a coarse-graining procedure, the projection to an unweighted, undirected graph implies a loss of information²¹. This sacrifice, however, is not worthless. What we gain is the noteworthy possibility of writing a simple model for the *C. elegans* brain maturation in terms of a handful of graph motifs²², as we will explain in more detail in the following.

Results

We obtain eight unweighted, undirected graphs²³. Unless otherwise specified, we will refer to the undirected synaptic connections simply as edges. A preliminary step in our modelling approach is to measure the network properties of interest and how they change during development. In tab. 5.2, we report the computation of a representative subset of standard graph metrics, app. B.

The number of neurons (nodes) increases from 161 at birth to 180 in the adult stage, with a burst of neuronal births at the turn of the larval stages L_1 and L_2 , consistent with what was previously reported²⁴ [217, 218]. In parallel, we observe a 2.7-fold increase in neuronal connectivity, from 617 edges at birth to ~ 1650 for the adult. This further results in an increase of both the number of two-stars (or, connected triples) and triangles. Also, the average geodesic distance between any two nodes decreases from ~ 3 to ~ 2.2 . In summary, the adult *C. elegans* network develops to become more closely interconnected.

Arguably, this leads to enhanced functionality of the *C. elegans* brain network. We can gain insights into this by computing the average local efficiency and clustering coefficient, which both increase during development. The first suggests an deployment of a biological strategy for increasing the system’s redundancy and robustness. The second is compatible with an increasingly modular organisation throughout development, and suggests an improvement in local information processing.

Crucially, all the network metrics here considered indicate a consistent and monotonous trend throughout the worm’s brain maturation.

5.3.2 Topography of the functional landscape

The crucial methodological step in EE modelling is the specification of an F metric and the inference of the topography of the resulting functional landscape. In this work, we do so by ERG inference, ch. 2. In particular, without loss of generality, we can express the F metric as a linear combination $F(G) = \theta \cdot x(G)$ of graph statistics $x \in \mathbb{R}^r$ with

20: The randomisation procedure used in both [213, 214] preserves in-degree and out-degree and the numbers of bidirectional and unidirectional connections for each neuron. See Figure 7 in [214] and an updated version of the same exact plot, Extended Data Figure 7 in [213].

21: In fact, we can derive the edge dynamics of the unweighted undirected graph from that of the original weighted directed graph, but not vice versa. The arguments we have presented above are not intended to prove that our coarse graining is a mere rephrasing of the network’s original information. Rather, they are intended to convince the reader that it is not meaningless, in the sense that it does not change the nature of the problem.

22: This is typically the point of a manuscript at which one summons the spirit of the 14th-century Franciscan friar, named William of Ockham, and its renowned razor. The relationship between the the *Ockham’s razor* – also, the *principle of parsimony* – and the scientific truth is riddled with nuances. A discussion of the use of the Ockham’s razor as an abductive heuristic can be found in [225, 226].

23: The scripts for both the preprocessing step and preliminary data analyses have been written using RStudio with R v4.0.4. They can be found in the Github folder: [EE-graph-dyn](#).

24: The *C. elegans* neurons are born in two separate bursts of cell differentiation. The first, major one happens during the embryonic stage (before hatching) and lasts approximately four hours. The second, minor, happens over seventeen hours during the post-embryonic stage, as here observed [217].

Table 5.2: Properties of the *C. elegans* networks of undirected synaptic connections. Each row corresponds to a graph, the first (birth) and the last two (adulthood) are highlighted. We compute the number of **nodes**, **edges**, **two-stars** – or connected triples –, **triangles**, the average shortest path (**av.sh.path**) – or average geodesic distance –, the average local efficiency (**loc.eff.**) and the average clustering coefficient (**clust.coeff.**). See app. B for the definitions.

| t[h] | nodes | edges | two-star | triangles | av.sh.path | glob.eff. | clust.coeff. |
|-----------|-------|-------|----------|-----------|------------|-----------|--------------|
| 0 | 161 | 617 | 5976 | 346 | 2.993 | 0.380 | 0.208 |
| 5 | 162 | 782 | 9273 | 601 | 2.712 | 0.416 | 0.232 |
| 8 | 162 | 788 | 9299 | 614 | 2.712 | 0.416 | 0.245 |
| 16 | 168 | 907 | 11838 | 830 | 2.617 | 0.428 | 0.246 |
| 23 | 173 | 1166 | 18449 | 1406 | 2.430 | 0.459 | 0.262 |
| 27 | 174 | 1175 | 18866 | 1433 | 2.429 | 0.458 | 0.274 |
| 45 | 180 | 1633 | 34124 | 2889 | 2.217 | 0.498 | 0.286 |
| 45 | 180 | 1669 | 35677 | 3003 | 2.206 | 0.501 | 0.292 |

linear coefficients $\theta \in \mathbb{R}^r$. With loss of generality instead, we propose the following parsimonious, coarse-grained model for the *C. elegans* brain maturation $(\hat{\Delta})_4$:

$$F(G) = \theta_{gwd} x_{gwd}(G|\lambda_{gwd}) + \theta_{gwesp} x_{gwesp}(G|\lambda_{gwesp}), \quad (5.1)$$

where the model statistics

$$\mathbf{x}(G) = \begin{bmatrix} x_{gwd}(G|\lambda_{gwd}) \\ x_{gwesp}(G|\lambda_{gwesp}) \end{bmatrix} \in \mathbb{R}^2 \quad (5.2)$$

25: Note that the two graph statistics used here are defined in the case of undirected, unweighted graphs. Simple representations allow for simple models. Furthermore, in defining this model, we have not included a term to represent the graph's edge count (or density). In fact, the latter is a degree of freedom which is controlled by the exploration rate in an EE dynamics, therefore it is not at the disposal of the F metric specification. See later in sec. 5.3.3.

have been defined in (2.28) and (2.32), respectively²⁵.

According to (5.1), the biological function of a worm brain network can be characterised in terms of two complementary graph statistics $\mathbf{x} \in \mathbb{R}^2$. The first one, x_{gwd} , based on the graph degree distribution, highlights node connectivity. The second one, x_{gwesp} , based on the distribution of edgewise shared partners, captures relational patterns. Together, they provide a comprehensive view of both node attributes and network configurations.

Methods

Modulo a minus sign, the F metric (5.1) corresponds to the graph Hamiltonian \mathcal{H} obtained when constructing an ERG model with graph statistics x_{gwd}, x_{gwesp} . Therefore, we can use the ERG inference to estimate the four parameters - two linear coefficients θ and two decay parameters λ - of (5.1). In other words, we employ the ERG methods to infer the topography of the functional landscape. To ensure that the correct (functional) balance of model statistics can be achieved at the end of the developmental process $(\hat{\Delta})_5$, we use the two adult *C. elegans* brain snapshots G_T^* as input for the inference, $T = 45$ h.

Listing 5.1: ERG inference based on (5.1), library `ergm` v4.3.2 for R v4.0.4, code available in the Github folder `EE-graph-dyn`. G represents the input graph for the inference, either $G_{T,1}^*$ or $G_{T,2}^*$. The decay parameters of the curved statistics are estimated as well (`fixed=F`). The model is constrained to those graphs that have the same number of edges as the G . As initial guess of the four parameters, we use $(1, 1, 1, 1)$.

```
# ERGM formula
ergm(formula = G ~
      gwdegree(fixed=F)+gwesp(fixed=F),
      constraints = ~ edges,
      control=snctrl(init = c(1,1,1,1))
    )
```

As the ERG inference is defined for one single graph, an output procedure is therefore required. A simple choice is that to use the so-called *mean-ERG* ($\hat{\mathcal{A}}_6$), originally proposed in [227]. Accordingly, the inference is performed independently for each network, resulting in multiple estimates of each parameter. The final estimate is determined by averaging the corresponding values across all networks²⁶. In doing so, we implicitly assume that each network is a different realisation of the same (bio)physical system.

Results

The ERG inference based on (5.1) for the two *C. elegans* brain networks is performed as discussed in ch. 2 and yields the estimates summarised in tab. 5.3.

| | θ_{gwd}^* | λ_{gwd}^* | θ_{gwesp}^* | λ_{gwesp}^* |
|-------------|------------------|-------------------|--------------------|---------------------|
| $G_{T,1}^*$ | 0.45 ± 0.20 | 1.91 ± 0.46 | 0.626 ± 0.056 | 1.432 ± 0.067 |
| $G_{T,2}^*$ | 0.43 ± 0.20 | 1.97 ± 0.48 | 0.529 ± 0.048 | 1.542 ± 0.075 |

26: This is an instance of the more general problem of constructing a *group representative network* (GRN). There are several other methods, ranging from the cruder to the more sophisticated alternatives [1]. We consider the mean-ERG to be a lower limit of methodological complexity. We are essentially limited here by the availability of only two adult networks.

Table 5.3: ERG estimation based on (5.1) for the two adult worms G_T^* . The maxent parameters θ^* are both significant and positive for all networks. The parameters λ^* controlling for the geometric decays of the model statistics are significant – and positive by construction.

The emerging picture is of an adult *C. elegans* brain network characterised by a propensity for (i) the presence of highly connected nodes ($\theta_{gwd} > 0$) and (ii) triadic closure ($\theta_{gwesp} > 0$). The former is consistent with the presence of medium and large hub nodes. The latter can reflect an underlying graph modular structure and is compatible with a common neighbor rule – i.e., neuron pairs with more shared neighbors have a higher likelihood of connection – for the worm’s neuronal wiring. These characteristics of the worm’s adult brain network have been extensively documented in recent years [215, 220, 228–230].

The results of the mean-ERG construction for the parameters of the *C. elegans* functional landscape are shown in tab. 5.4.

| θ_{gwd}^* | λ_{gwd}^* | θ_{gwesp}^* | λ_{gwesp}^* |
|------------------|-------------------|--------------------|---------------------|
| 0.44 | 1.94 | 0.578 | 1.487 |

Table 5.4: Parameters of the *C. elegans* functional landscape, as defined by the F metric (5.1). Mean-ERG based on the estimation obtained from two adult worm’s brain, tab. 5.3.

We observe that an EE dynamics based on the F metric (5.1) with the above parameters would favour the emergence of hub-like structures and the strengthening of triadic closure throughout development. This is consistent with two of the developmental principles highlighted in [223], namely (i) that well-connected neurons receive more inputs and (ii) that network modularity increases with time. The latter, in turn, align with the trends described in tab. 5.2, i.e., with a developmental process that progressively weaves a more tightly connected, robust and efficient network topology.

Our minimal model of the *C. elegans* neurofunctional landscape is based on two graph statistics and has only four parameters that are amenable to biological interpretation and are inferred from the data. All that remains for us is to unleash an EE dynamic on it.

5.3.3 Worm brain maturation tracked down

Using an EE dynamics to model the *C. elegans* brain network development means representing it as a stochastic dynamics of a probability distribution on a functional landscape. The rationale for this has been discussed in sec. 5.1 in the more general context of the brain wiring problem and originates from the need to accommodate two fundamental observations: the existence of developmental variability²⁷ and the robustness of the functional outcome.

In using this picture, we implicitly make an important assumption, which we might call the *hypothesis of functional homogeneity* (Δ)₇. This means that the same definition of biological function holds true throughout the whole developmental process²⁸. We will treat this hypothesis self-consistently and verify it a posteriori.

To our best effort, no analytical treatment is possible for an EE dynamic (4.4) with the F score in (5.1), therefore, we resort to simulations, described in sec. 4.3. In the context of the EE dynamic, three elements need to be set: the boundary conditions and the two EE parameters, i.e., the exploration rate μ and the exploitation rate φ – or, equivalently, the functional pressure ρ . The collections of the methods we employ is summarised in fig. 5.6.

Methods

The boundary conditions are fixed by setting the birth connectome as the starting point for the EE dynamic (Δ)₈, i.e.,

$$P(G = G_0^*, t = 0) = 1. \quad (5.3)$$

In fact, as reported by *Nicosia et al.* [218] the embryonic and post-embryonic stages represent two distinct phases in the maturation of the *C. elegans* brain, with the hatching (birth) serving as a watershed. Qualitatively, these two phases are likely driven by the same developmental principles. Yet, in quantitative terms, they differ fundamentally²⁹. All of the data we use were collected at different post-embryonic stages. Therefore, in line with the hypothesis of functional homogeneity, we restrict our modeling to the post-embryonic developmental phase, setting the birth connectome as the starting point.

The removal of existing synaptic connection is rarely observed during the worm's brain maturation. Therefore, we adopt an exploration scheme where only the formation of new connections is permitted³⁰ (Δ)₉. As previously reported, the number of synapses increases approximately linearly with time [223]. For simplicity, we further assume that the formation of new connections occurs uniformly across all neuron pairs (Δ)₁₀. The exploration rate μ – i.e., the number of edges added to the graph per dyad and per unit time – is then computed as

$$\mu^* = \frac{1}{TL} \sum_{i < j} [\bar{a}_{ij}(G_T^*) - a_{ij}(G_0^*)], \quad (5.4)$$

27: In [223], the 43% of the (directed) synaptic connections were found not to be conserved between isogenic individuals, contrary to the common assumption that the *C. elegans* brain is hardwired. Interestingly, not all of these connections consist of only a few synapses.

28: Cf. with the discussion on the fitness landscape/seascape, sec. 3.2.2.

29: For instance, there is a prominent difference in the rate at which the connections appear during development, which is accelerated in the pre-embryonic phase ($\sim N^2$, where N is the number of neurons) and linear after birth ($\sim N$).

30: Luckily! As discussed in s.n. 18, the biology of the process of synapse formation coincides with that of the process of connection formation, but the same is not true of synapse elimination. In practice, for the *C. elegans* the latter does not occur. Using the unweighted representation of synaptic connections is therefore less harmful in this case.




Figure 5.6: Exploration-exploitation dynamics for the *C. elegans* brain maturation, illustration. The process is depicted as the dynamics of a probability distribution on a functional landscape. Our starting point is the dataset presented by Witoliet et al. [223], which consists of eight reconstructions of the worm’s brain (nerve ring and ventral ganglion), at different developmental ages $t = 0$ h, 5 h, 8 h, 16 h, 23 h, 27 h and 45 h ($\times 2$) after hatching (birth). In particular, we consider the unweighted, undirected networks of chemical synapses – some are omitted for visual clarity. The configuration of the brain at birth G_0^* (white graph, red background) is set as the starting point of the dynamics. The two adult brain networks G_T^* are used for the inference of the model parameters. In particular, we use the framework of exponential random graphs (ERG) to infer the topography of the functional landscape $F(G)$ (5.1), as encoded in the parameters $\theta \in \mathbb{R}^2$ (Lagrange multipliers) and $\lambda \in \mathbb{R}_+^2$ (decay parameters of the curved statistics). The final estimations are obtained by mean-ERG construction. As for the parameters of the dynamics, the exploration rate $\mu \in \mathbb{R}_+$ is assumed to be uniform across node pairs and constant in time. It is evaluated by computing the average increase of connections per dyad and per unit time, from birth to adulthood. Finally, the functional pressure $\rho \in \mathbb{R}_+$ is used to inform the EE dynamics about the age of the adult worm and is fixed by minimising the Mahalanobis distance between the experimental model statistics at adulthood (average) and their simulated distribution based on (4.4). Our model of the *C. elegans* brain maturation has six parameters, all of which lend themselves to biological interpretation.

where $T = 45$ h, $\sum_{i<j} a_{ij}(G_0^*)$ is the number of edges of the birth connectome and $\sum_{i<j} \bar{a}_{ij}(G_T^*)$ is the mean³¹ number of edges between the two adult worms.

This leaves us with only one degree of freedom, the functional pressure ρ . The value of this parameter is subject to biological tuning and is peculiar to both the process and the system under consideration – in our case, the brain wiring dynamics and the *C. elegans* nervous system. The regulation of the functional pressure ensures that, as development progresses, specialised functional circuits emerge and mature appropriately before the organism reaches its adult stage. For this reason, we use it to inform the EE dynamics about the age of the adult worms³². Intuitively, we intend to harness the degree of freedom of the functional pressure to minimise the distance between our simulations and the experimental values, at the adult age $T = 45$ h.

To do so, we resort to the Mahalanobis distance $(\hat{\Delta})_{11}$ [231], fig. 5.7(a).

31: We use here the notation of the bar for the mean value, in this case:

$$\bar{a}_{ij}(G_T^*) = \frac{1}{2}[a_{ij}(G_{T,1}^*) + a_{ij}(G_{T,2}^*)]$$

32: It is indeed a matter of time. For every $\rho > 0$, the distribution of F values shifts towards higher values and, waiting long enough, it reaches any desired target. Thus, it is not a question of *if* the EE dynamics reach a predetermined functional value, but *when* they do so. High values of ρ rapidly advance the F distribution while low values instead imply a slow progress.

33: It does so by computing the Euclidean distance of the *whitened* (standardized) data. More precisely, consider the *whitening* transformation [232]

$$\mathbf{y} \rightarrow \Sigma_Q^{-\frac{1}{2}} \mathbf{y} \quad \forall \mathbf{y} \in \mathbb{R}^r,$$

where $\Sigma_Q^{-\frac{1}{2}}$ is the inverse principal square root of the covariance matrix Σ_Q . It is possible to show that (a) the transformed variables have unit diagonal (white) covariance matrix and (b) the distance (5.5) corresponds to the Euclidean distance of the transformed variables.

34: This means that it is invariant under affine transformations of the form

$$\mathbf{y} \rightarrow A\mathbf{y} + \mathbf{b} \quad \forall \mathbf{y} \in \mathbb{R}^r,$$

where A is an $r \times r$ matrix and $\mathbf{b} \in \mathbb{R}^r$.

35: The same would not be true if instead we considered a distance function based on the F metric.

Figure 5.7: Mahalanobis distance in the space of model statistics (5.2). (a) Illustration of the Mahalanobis distance (5.5). It is a natural generalisation of the Euclidean distance to that between a point $\mathbf{y}^* \in \mathbb{R}^r$ (here, $r = 2$) and a multivariate distribution Q in the same space, with mean $\langle \mathbf{y} \rangle_Q$ and covariance matrix Σ_Q^{-1} . (b) Example of the correlation structure between the model statistics (5.2) from an EE simulation of the *C. elegans* brain maturation. The two statistics show a clear anti-correlation.

Consider a a multivariate distribution Q on \mathbb{R}^r and point $\mathbf{y}^* \in \mathbb{R}^r$. The Mahalanobis distance δ_Q^{mah} between the distribution Q and the point \mathbf{y} is defined as

$$\delta_Q^{mah} = \sqrt{(\langle \mathbf{y} \rangle_Q - \mathbf{y}^*)^\top \Sigma_Q^{-1} (\langle \mathbf{y} \rangle_Q - \mathbf{y}^*)} \quad (5.5)$$

where $\langle \mathbf{y} \rangle_Q$ and Σ_Q are the mean and covariance matrix of Q . It is a natural generalisation of the Euclidean distance and has two intriguing, i.e., (i) it accounts for the covariance structure of the distribution³³ and (ii) it is scale invariant³⁴.

Back to the EE dynamics, we compute (5.5) where:

- Q is the simulated two-dimensional distribution of model statistics $\mathbf{x} = (x_{gwd}, x_{gwesp})$, computed as (4.23), at the adult age $T = 45$ h – we use T to label the distribution.
- $\langle \mathbf{y} \rangle_Q = \langle \mathbf{x} \rangle_T$ and $\Sigma_Q = \Sigma_T$ are the ensemble average and covariance matrix, at $T = 45$ h, evaluated with (4.24).
- $\mathbf{y}^* = \bar{\mathbf{x}}(\mathbf{G}_T^*)$ are the experimental values of the model statistics, averaged over the two adult connectomes, i.e., $y_i^* = [x_i(\mathbf{G}_{T,1}^*) + x_i(\mathbf{G}_{T,2}^*)]/2$.

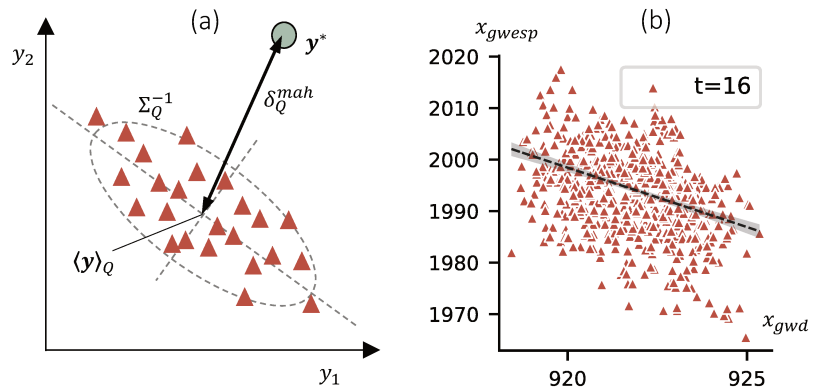
The optimal functional pressure ρ^* is then defined as:

$$\rho^* = \min_{\rho} \delta_T^{mah} \quad (5.6)$$

where

$$\delta_T^{mah} = \sqrt{(\langle \mathbf{x} \rangle_T - \bar{\mathbf{x}}(\mathbf{G}_T^*))^\top \Sigma_T^{-1} (\langle \mathbf{x} \rangle_T - \bar{\mathbf{x}}(\mathbf{G}_T^*))}. \quad (5.7)$$

The two properties of the Mahalanobis distance mentioned above turn out to be particularly useful in the present case. In fact (i) the two model statistics (5.2) exhibit a pronounced anti-correlation, fig. 5.7(b). Moreover (ii) the distance (5.7), hence the estimation (5.6), would not change if we scaled the statistics by their corresponding ERG parameter θ . Therefore, (5.7) offers a common ground for comparing models distinguished by varying sets of θ ³⁵.



Our model of the *C. elegans* brain maturation is exhaustively specified by the three expressions (5.3), (5.4), (5.6). It has six parameters, all of which inferred by using the adult worms exclusively, all of which lend themselves to biological interpretation.

If the theory is fully defined, our simulations have two more degrees of freedom that need to be set, i.e., the population size (number of samples) M and the inverse time interval ν , tab. 4.2³⁶. The first is set to $M = 2^{10}$. The second is fixed by requiring that a single edge addition occurs in each graph during in a simulation step. This means choosing $\Delta t = (L\mu^*)^{-1}$.


As a final note, our simulations are currently not able to model the appearance of neurons (nodes)³⁷. Consequently, we take the adult connectome as reference and embed the experimental networks at earlier developmental stages within larger networks that matched the adult node count of 180 (Δ)₁₂.

Results

The exploration rate and the functional pressure of the EE dynamics for the *C. elegans* brain maturation – evaluated by (5.4) and (5.6) – turn out to be

$$\mu^* = 1.426 \times 10^{-3} h^{-1}, \quad (5.8)$$

$$\rho^* = 9.017 \times 10^2. \quad (5.9)$$



36: These are chosen in order to optimise the trade-off between computational time and reliability of the simulations. Small M result in noisy empirical statistics (4.24), therefore M should be as large as possible – here, we have checked that the results do not change when using $M = 2^{11}, 2^{12}$. As for ν , setting it as described in the text helps in designing efficient simulations. The running time scales linearly with both M, ν . With the settings described here, a single simulation run currently takes ~ 1.5 h. Scripts available in the Github folder `EE-graph-dyn`.


37: This affects 19 of the 180 adult brain neurons, see tab. 5.1. The majority of them appears during the post-embryonic burst of cell differentiation between the L_1 and L_2 stages, tab. 5.2 and s.n. 24.

Figure 5.8: Exploration rate for the EE dynamics. The number of connections at birth (red circle) and average of two adults (white circles) determine the constant exploration rate μ^* (red line) as per (5.4). Thus, the number of connections grows linearly (red line). The shaded region represents a linear fit across all time points, including intermediate ages (black circles). We fit the line $y = \mu t + y_0$ and shade between lines with slopes $\mu^{**} \pm \delta\mu^{**}$ (same intercept y_0^{**}) – double asterisks indicate the estimated values. The line $\sim \mu^* t$ falls within the shaded area.

The exploration rate determines the frequency of structural changes, fig. 5.8. On average, $L\mu^* \sim 23$ new connections appear per hour. As for the exploitation strength, in fig. 5.9 we show that the average Mahalanobis distance (5.7) is a convex function of the functional pressure ρ . If ρ is too low, the dynamics resemble a random formation of new connections, and are therefore unlikely to result in functional configurations. Less intuitively, if ρ is too large, the dynamics also deviate from the experimental values, as they overestimated the strength of the functional selection. The *true* biological process is therefore not an upper bound of the corresponding EE dynamics, in the $\rho \rightarrow \infty$ limit. Rather, it corresponds to a finite value of functional pressure that has been plausibly calibrated by its evolutionary history.

By design, the estimates (5.8) and (5.9) were obtained using only the adult stage of worm development. Nevertheless, we can legitimately ask how the estimates would change if instead we used the whole information available, including the configurations of the worm brain at intermediate developmental stages, i.e., those at 5 h, 8 h, 16 h, 23 h and 27 h. For instance, (i) a linear fit of the rate at which new connections appear, using the entire time series, results in $\mu^{**} = (1.389 \pm 0.079) \times 10^{-3} h^{-1}$, which

Figure 5.9: Functional pressure for the EE dynamics. We run $S = 100$ simulations $\forall \rho \in \{200 + 20i, 0 \leq i \leq 60\}$. For each ρ , we compute the mean and standard deviation of δ_T^{mah} (5.7) over the S simulations (green line and shaded area, respectively). We fit the mean values by a quadratic curve $\tilde{\delta}_T^{mah}(\rho) = a\rho^2 + b\rho + c$ (dash dotted line) and take its minimum $\rho^* = -b^*/2a^*$ as estimate of the functional pressure for the EE dynamics. The same procedure, defining instead the Mahalanobis distance over the whole dataset (5.10), gives ρ^{**} . The values $\tilde{\delta}_T^{mah}(\rho^*)$ and $\tilde{\delta}_T^{mah}(\rho^{**})$ are highlighted (red and white diamonds, respectively). The two overlap within one standard deviation.



encompasses the value in (5.8), fig. 5.8. Similarly, we can define for the functional pressure an equivalent minimisation problem to (5.6) over the whole dataset as:

$$\rho^{**} = \min_{\rho} \sum_{t \in \mathbf{t}^*, t > 0} \delta_t^{mah}, \quad (5.10)$$

where each δ_t^{mah} is defined as in (5.7), using the graph configuration(s) at time $t \in \mathbf{t}^* = (0, 5, 8, 16, 23, 27, 45) h$. This procedure results in $\rho^{**} = 7.001 \times 10^2$. The distance from the experimental values of the model statistics in the adult stage, using ρ^{**} , is compatible within one standard deviation with the one obtained using ρ^* , fig. 5.9.

In essence, we find that using the full time series for inference, rather than just the adult stage, does not significantly alter the estimates of the parameters of the EE dynamics. This suggests that the adult stage captures most of the salient information about the brain wiring dynamics of the *C. elegans*, and other stages might have redundant or less significant impacts.

To delve deeper into this observation, we can look at a single simulation of the *C. elegans* brain maturation process, obtained using the exploration rate and functional pressure μ^*, ρ^* . In fig. 5.10, we show the dynamic in the space of model statistics (5.2). For comparison, we also plot a null model of the same process in which $\rho = 0$, i.e., a random formation of new connections at rate μ^* . We find that the statistics at the adult stage are accurately reproduced by our simulation – as it should be, in view of (5.6). Remarkably, however, even though we did not use the other observed developmental ages for parameter inference, our simulations closely approximate them as well. This holds true for the statistics (5.2) and, by consequence, for the F metric (5.1), which drives the dynamics.

This hints that our EE graph dynamics, informed about birth and adulthood, is able to capture the entire developmental trajectory³⁸. Upon reflecting on the model's design, this provides *a posteriori* validation for our hypotheses (i) that the model (5.1) encapsulates the fundamental drivers of the *C. elegans* brain's growth process, and that (ii) that the functional homogeneity discussed above holds true for the process considered. Importantly, our simulations are defined and open for inspection at all times $t \in [0, T]$, not just at the ages reconstructed in our source dataset. This opens up the possibility of using them to predict and

38: It is understood that this statement is contingent upon the representation chosen for both the system and the process. For example, by construction our model cannot reproduce any fine-scale detail of the worm brain networks related to node-specific effects – since nodes are indistinguishable in (5.1) – or to the directed nature of the connections – since we use undirected graphs. Similarly, a constant mutation rate and functional pressure cannot reproduce transient dynamic patterns, which are apparent in fig. 5.10.




Figure 5.10: Optimal reconstruction of the *C. elegans* brain maturation. We fix μ^*, ρ^* as in (5.8), (5.9) and run 500 simulations. The one that minimises (5.7) is here shown. Main: dynamics in the space of model statistics (5.2). Circles represent the eight experimental points for $t^* = 0\text{ h}, 5\text{ h}, 8\text{ h}, 16\text{ h}, 23\text{ h}, 27\text{ h}$ and 45 h ($\times 2$). Diamonds indicate the expected values $\langle x \rangle_{t^*}$ based on the EE dynamics. Finally, squares indicate the expected values of the statistics using a null model in which the exploitation is turned off, $\rho = 0$. Dashed and dotted lines shows the entire developmental trajectories, as obtained from the simulations. Inset: experimental values of the F metric (5.1) (white circles), the corresponding expected values based on the EE dynamics (red diamonds and dashed line). Our optimal EE graph dynamic, informed by the birth and adult configurations of the worm brain network, also closely captures the developmental ages not included in the inference process. By construction however, our model cannot capture fine-scale details of the dynamics, e.g., the slow down of the experimental progression in the space of model statistics between $t = 23\text{ h}$ and $t = 27\text{ h}$.

generate reliable estimates for those stages of worm brain maturation for which there is currently no data.

A meaningful follow-up question to ask is whether our model is able to reproduce other features of the data – not included in the formulation –, and which ones. We refer to this characteristic as *feature generalisation*. In fig. 5.11, 5.12, we demonstrate the ability of our model to generalise its predictions across different aspects of the data on which it was not explicitly trained. Once again, this is true not only at adulthood, but across the entire developmental process.

Collectively, these findings demonstrate that already a simple, low-dimensional, and coarse-grained model of the *C. elegans* brain maturation, properly informed about the final steady state of the process, can quantitatively replicate several features of the experimental data across development. Importantly, the model proposed has only six interpretable parameters – four for the functional landscape, two for the dynamics³⁹. To the best of our knowledge, this is a first for neurodevelopmental

39: Back to the quote that opened this chapter: who knows what an elephant can do with six parameters! But for this thesis, this was our best effort. The whole anecdote related to that quote, as told by its main character, F. Dyson, can be found in [233].

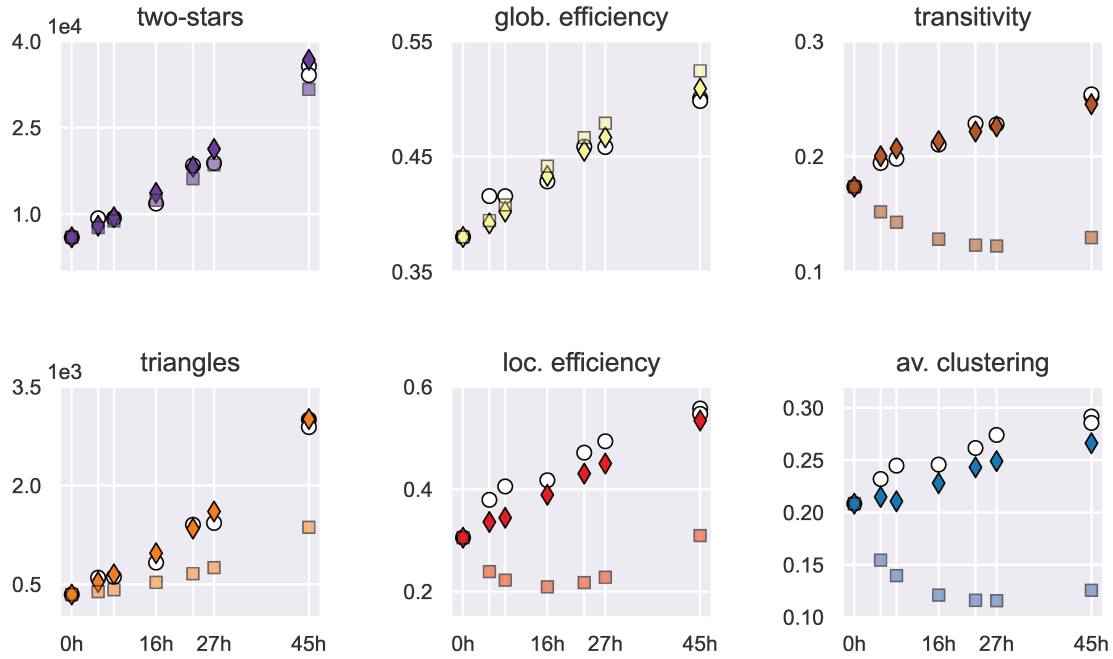



Figure 5.11: Feature generalisation for the EE model of the *C. elegans* brain network maturation. We compare the experimental values (circles) of a set of graph statistics with those obtained from simulations (diamonds) and those corresponding to a null model with no exploitation (squares). The same simulation illustrated in fig. 5.10 is used here. The graph statistics include (i-ii) two motif counts (two-stars, triangles), which represent the fundamental building blocks of higher order graph patterns; (iii-iv) two efficiency measures (global and local efficiency), which quantify the efficiency of information transmission within the network at a global and local scale; (v-vi) two measures of clustering (transitivity and average clustering coefficient), which indicate the presence of tightly knit communities or groups. See app. B for a more detailed description of these (standard) graph measures. The null model (random formation of new connections) produces similar values to the experimental ones for the number of two stars and global efficiency, but fails elsewhere. On the contrary, our optimal EE model closely reproduces the experimental values for the entire set of graph statistics presented here.

Figure 5.12: Complementary cumulative degree distribution at adulthood. The distributions associated to the two adult worms are shown as red solid lines. Blue dashed line for the curve obtained from simulations as $P_{deg}^{(c)}(k) = \frac{1}{N} \sum_{j \geq k} \langle x_d^{(j)} \rangle$. Gray dashed line for the null model (no exploitation). Inset: zoom in the high-degree tail (loglog plot). Green dashed line corresponds to a power law behaviour. Our EE model is able to reproduce closely the cumulative degree distribution of the adult worms, including the high-degree tail behaviour. Similar plots can be obtained for each developmental age.



40: Of course, this is mainly due to the fact that data such as that used here has only been available for a few years, but still.

biology studies⁴⁰. Returning to the general discussion of the brain wiring problem in sec. 5.1, our EE dynamics provide an appropriate framework within which to specify a (genetically encoded) developmental algorithm. In the most conservative interpretation, the one described in this chapter is precisely an algorithm, consistent with the general features of the biological dynamics under investigation (in particular, the self-referentiality) and whose validity is demonstrated *a posteriori* by the results discussed above.

Nevertheless, we intend to put forward a bolder interpretation of the EE

dynamics, at the local scale, consistent with our current understanding of the synapse formation process.

5.4 Interpretation down to the synapse scale

Where a biological interpretation of EE dynamics is speculated. Where in particular individual neurons in a single developing system make local decisions based only on their limited knowledge.

An interpretation of our model requires the specification of how the EE dynamics is implemented by the biological process of synapse formation, which happens at the scale of the single neuron. Here, we speculate on a plausible biological interpretation of the EE dynamics, within a graph representation of the system.

Most *C. elegans* neurons have only one or two processes that extend in parallel bundles along the worm's body⁴¹. These processes grow during development, guided by molecular cues. Presynaptic sites appear as *en passant* swellings on the shaft of the axon. The postsynaptic processes are dendrites or as spine-like protrusions [223]. Occasionally, these dendrites or protrusions form new synaptic connections, fig. 5.13(a).

To illustrate the essential idea, let us consider for simplicity the scenario illustrated in fig. 5.13(b). At time t , there exist synaptic connections between neurons AB , DE , and BC . In the time interval Δt , postsynaptic processes from both neurons C and D grow sufficiently close to the axon shaft of A , and hold the potential to develop into new synaptic connections. Conversely, neuron E has no such process, so it cannot form a connection with A . If Δt is sufficiently small, only one of the two possible synaptic connections AC or AD is likely to be formed. Which of the potential connections materialises first is determined stochastically. However, the connection with a greater functional advantage will plausibly have a higher probability of forming. Let us suppose – again, for simplicity – that the notion of biological function for this system is simply represented by the number of triangles in the undirected graph representation of the system, fig. 5.13(c). Consequently, we expect the AC connection to form preferentially, given it leads to the formation of the ABC triangle.

The example above illustrates how the EE dynamics could be implemented for a single developing system⁴². From the standpoint of the individual neuron, the process of synaptogenesis consists of a series of stochastic decisions about which other neuron to connect with⁴³. These decisions are biased towards those connections that lead to higher functional gains, which in turn are evaluated based on the information available to the neuron at any given time.

According to this interpretation, the exploration consists in the formation (or extension) of neuronal processes that *could* lead to a new synaptic connection and thus do not themselves correspond to the formation of physical connections⁴⁴.

On the other hand, the exploitation consists in assigning higher probabilities of formation to those potential connections that would lead to higher functional gains. In particular, suppose that G is the graph configuration

41: See sec. 5.2 for an introduction to the *C. elegans* neurobiology.

42: The argument straightforwardly generalises to more complex notions of biological functions.

43: It is worth stressing that this is a pictorial way of understanding the process and, clearly, only an effective description, which is based on (and assumes) a graph representation of the system. At the molecular level, the synaptogenesis is regulated by a complex forest of biochemical mechanisms. Strictly speaking, therefore, it is not necessary for neurons to "compute" any notion of biological function and "make decisions". Once again, the developmental rules are ultimately genetically encoded and biochemically implemented.

44: Note that this marks a sharp difference with the context of the evolutionary dynamics, where different exploration events are interpreted as a set of different genetic mutations, each associated with a distinct individual in a population. Here there is one and only one copy of the system, and the exploration events correspond to the formation of potential, not physical, connections.




Figure 5.13: Interpretation of the EE dynamics. (a) Synapse formation, schematics. Presynaptic sites, appear as swellings (black circles) on the axon shaft (thick black line) of a presynaptic neuron (gray circle). Postsynaptic neuronal processes - dendrites (black lines) and spine-like protrusions (thin black lines) - sprout from a postsynaptic neuron (white circle). Occasionally, they form synaptic connections with physically proximal presynaptic sites (red dashed lines). The presynaptic processes of the postsynaptic neuron and the postsynaptic processes of the presynaptic neuron are not shown here. (b) A simple scenario. We represent in a cartoonish physical space one presynaptic neuron A and four postsynaptic neurons B, C, D, E . The axon shaft extending from A is represented by the thick black line. At time t there is a synaptic connection between the nodes AB (black line connecting B to the axon shaft). Additional connections exist between the neurons BC and DE (not shown in the physical space, indicated by the gray dashed line). After a time interval Δt , postsynaptic neuronal processes extend from the neurons C, D towards the axon, potentially leading to new connections (red dotted lines). On the contrary, no such postsynaptic process exists for the neuron E . (c) Representation of the scenario in (b) in the corresponding graph space (undirected connections). We assume that the biological function (F metric) simply consists in the count of triangles. The two potential connections between AD and AC at time $t + \Delta t$ can be represented as two different graph configurations, G_1, G_2 , associated to different F values. G_2 , by virtue of its higher F , will be realised with higher probability. (d) Decision tree for two time steps of the EE dynamic (example). Here, $\Delta t = 1$. Each square represents a graph. In blue, we indicate the F values. In black, the unconditioned probabilities computed at each time as $\exp[F(G_i)] / \sum_j \exp[F(G_j)]$ where the sum runs over all graphs at that time (column). In brown, the probabilities conditioned on the previous time-point. They can be computed either as above, restraining the sum to those graphs that come from the same parent graph at the previous time, as in eq.(5.11). Alternatively, they can be evaluated starting from the unconditioned EE probabilities and using $P(G_i, t + 1 | G_j, t) = P(G_i, t + 1 \cap G_j, t) / P(G_j, t)$, where $P(G_j, t) = \sum_k P(G_k, t + 1 \cap G_j, t)$. In bold-red we highlight the most likely developmental trajectory.

at time t and \tilde{G} are the potential graph configurations at time $t + \Delta t$, then

$$P(G_i, t + \Delta t | G, t) = e^{\Delta t F(G_i)} / \sum_{G_j \in \tilde{G}} e^{\Delta t F(G_j)}, \quad (5.11)$$

where we have taken $\varphi = 1$ for simplicity. By definition, only one of them is eventually realised in Δt . The exploration-exploitation cycle is iterated throughout the process and defines the EE stochastic trajectory of a single, developing system in the configuration space.

We stress, however, that the information contained in (4.4) is much more general, as it allows to compute the (unconditioned) probability of all possible configurations that could have appeared by the time t , including those that result from very unlikely developmental paths, fig. 5.13(d). This provides the justification for using the EE dynamics to capture the intersubject variability of the brain wiring, that results from slightly different developmental trajectories.

5.5 Paving ways

Where an agenda of the next steps can be found. Some easy to take, some more ambitious, some optimistic, none unfeasible. Enough for another PhD project, a bold postdoc or – question by question, answer by answer – an entire early career.

The model presented in this chapter for the brain wiring problem, like any other model, is shaped by the assumptions made during its formulation. Some relate to the choice of the representation, some to the model, some to the theory and some others to the biological system itself – we list them all in tab. 5.5. So far in this chapter, we have moved from broad themes to granular insights. In this final section, we do the opposite, rewinding the tape of our discussion from the specific to the general and pointing out a (small) subset of the (many) possible extensions and generalisations of the concepts and methods exposed above⁴⁵.

| (A) | Description | Step beyond |
|-----|--------------------------------|-------------|
| 1 | connections are synapses | ••• |
| 2 | graphs are unweighted | •• |
| 3 | graphs are undirected | • |
| 4 | F metric is ERG-like | •• |
| 5 | inference from the adult stage | •• |
| 6 | mean ERG (F landscape) | ••• |
| 7 | functional homogeneity | •• |
| 8 | start from brain at birth | •• |
| 9 | connections are only formed | • |
| 10 | uniform exploration rate | • |
| 11 | min. Mahalanobis distance | •• |
| 12 | no node dynamics | • |

The model

First and foremost, our model of biological function (5.1) for the *C. elegans* brain. A number of factors (drivers and constraints) have the potential to play a role in molding the fine-scale details of an adult worm wiring, which have not been included in our coarse-grained model [220, 223, 234, 235].

For instance, the graphs considered in this research are embedded in the physical space. One possible way of taking this into account would be to calculate the total cost of the *C. elegans* brain, i.e. the sum of the physical soma-soma distances between connected neurons. The latter is not strictly minimised [236, 237], yet is likely to play a role in shaping the adult connectome⁴⁶ [218, 220, 238].

Homophily effects based on various cellular attributes influence the neuronal (synaptic) connectivity. Neurons in the adult worm brain are more likely to be connected if they differentiate close together in time [217] and if they belong to a bilaterally symmetric pair [239], tab. 5.1. In fact, homophily effects in the adult brain based on birth cohort (pre- or after-hatching) and symmetric pairing have been demonstrated [220].

45: Of all the sections typically found in a scientific paper, the Discussion is surely the most literary, the most replete with metaphors. Some papers claim to *shed light* or *open doors*. Others boldly *navigate uncharted waters*, *open horizons*, or *break new ground*. Yet, some prefer the humility of merely *scratching the surface*, being *the tip of an iceberg*, or *a drop in the ocean*. In our case, we do choose to *pave ways*, a challenging task – as most of the carpentry work –, but enduring and steadfast.

Table 5.5: A complete list of assumptions enforced in our *C. elegans* brain maturation model. They are listed in order of appearance in the text (first column). In the second column, we give a synthetic description. In the third column, we qualitatively indicate – according to our how difficult is the task of removing and/or generalising the assumptions: if easy (•), challenging (••) or arduous (•••). Those related to a lack of data are marked with three bullets, as there is not much we can do, as humble theoretical physicists.

46: However, this is a simplistic way of thinking at role of distance in the worm nervous system. As discussed, neuronal process mostly run in parallel bundles along the worm body. Therefore it is more the physical neighborhood of the neurons' processes that matters and in particular the contact area between adjacent processes. Such data is increasingly available [234]. For instance, the contact area at birth between pairs of neurons correlates with the probability of a forming a new connections [223].




Figure 5.14: Information flow by cell type. Black, gray, and dashed black arrows indicate feedforward, feedback and recurrent connections, as defined in [223]. The number (and strength) of feedforward connections during the *C. elegans* development.

47: A *functional circuit* is a group of interconnected neurons that work together to perform a specific function or behavior [240, 241]. Examples are the navigation, touch sensitivity, chemosensation [242–244].

48: The reader who has taken the trouble to read ch. 3 will recognise the parallel of this discussion (and others in the present section) with the one in sec. 3.2.2. For instance, we are here proposing to move from a functional landscape to a *functional seascape*, which might possibly co-develop with the system itself.

49: For instance, a constant, uniform and scalar exploration rate would become a matrix to account for non-uniform rates, a vector to account for time variability, or a vector of matrices in the case of both.

In addition, the symmetry of the wiring increases over the course of development [223].

Synaptic connections are inherently directed, which is essential for detailing the information flow, fig. 5.4. For example, hub neurons receive disproportionately more input connections while the number of outputs remains stable. In addition, synaptogenesis preferentially creates new connections in the direction from sensory to motor neurons, increasing the "feedforward bias", fig. 5.14 [223].

The ERG-like F metric in our framework provides a straightforward means of incorporating all these effects into a model by formulating them as edge covariates (2.26), with an appropriate choice of matrix γ , see sec. 2.2.2. The ERG parameters estimation then has the role of weighting the relative importance of each factor.

Finally, F can be endowed with terms that can account for the presence of stable neuronal circuits in the wiring⁴⁷. A more realistic picture of the wiring variability is that of a core circuit, that is conserved across individuals and development, embedded in a background of variable connectivity [223, 234]. This (soft) constraint can be accounted for by a term in F that penalises those graphs that do not contain a given functional circuit.

The framework

A different set of possible extensions of the model here presented concerns the EE framework illustrated in ch. 4.

A natural initial step is to integrate node formation dynamics alongside the EE edge dynamics. Indeed, 19 out of the 180 neurons in the adult worm brain differentiate post-hatching. The time of their emergence is documented in the literature [217].

It is more challenging to relax the assumption of time homogeneity of the process and to introduce a time dependence for the parameters of the dynamics, e.g., $\mu(t)$, $\rho(t)$. The latter case is the one we would generally expect, since the worm brain maturation is influenced by several internal (and external) factors, that might manifest as a slow-down or speed-up of the developmental process⁴⁸ – e.g., fig. 5.10. Here, however, we are essentially limited by the data – seven developmental ages –, which do not justify the use of more refined time dependencies.

As we have discussed in sec. 5.3.3, the developmental processes during the embryonic and post-embryonic stages operate in distinct regimes and the latter has been studied here. We argue that the general EE structure of the problem should not change, while the parameters and possibly the F metric should. It would then be interesting to design an EE model of the embryonic phase of the *C. elegans* brain maturation and connected it to the post-embryonic model here described, at the moment of hatching.

Our EE simulation framework can accommodate (with varying degrees of effort) the ensemble of generalisations outlined here, mainly by adapting the mathematical structure of the parameters⁴⁹ and/or of the dynamic entities.

The system

Time to get ambitious. If we insist on a model for the brain wiring problem grounded on the experimental evidence, data are vital. The kind of data relevant to the research question under consideration here - in particular, electron microscopic reconstructions of nervous systems - is becoming increasingly available for ever larger systems. Let us therefore imagine for a paragraph (hopefully near) future in which we have the data we are looking for.

There is no fundamental reason to restrict our analysis to the connection defined as chemical synapses. A complete map of gap junctions (undirected connections) for the *C. elegans* across development would be easily incorporated into our analysis.

There is also no fundamental reason to restrict to the nematode *C. elegans* as the general principles that drive the wiring development do not depend on the specific system by which they are implemented. An EE model of the brain wiring dynamics for the fruit fly, mouse, zebrafish, tadpole and more⁵⁰ [203–208] might be within reach. It would be then interesting then to use the EE framework as a common ground to compare equivalent models across different natural nervous systems⁵¹.

Finally, it would be interesting to go beyond connectomics and look at functional connectivity⁵² through the same lens [245]. In fact, a more direct relation with the notion of biological function exists for the signal propagation atlas, as the one very recently reconstructed for the *C. elegans* by Randi *et al.* [246]. Here, individual neurons are excited by optogenetic stimulation and the activity induced in other neurons is recorded, thus defining a graph of directed, weighted functional connections.

50: We will resist the temptation to include the human brain in this list, since, to the present day, it would be ludicrous to even think of having such data. We leave here this side-note, in the hope of returning to it one day and being amazed at how quickly it has been disproved.

51: For example, the dream plot we have in mind is a scatterplot, where on the two axes we have μ, ρ and each point represents an EE model of a different brain system.

52: Connectomics assumes that the anatomical map of connections, like synapses and gap junctions, is key to understanding neural functions. While essential, this overlooks aspects such as inhibitory or excitatory nature of synapses, extrasynaptic communication via signaling molecules [216], and the timescales of signals propagation.

*Addo' arrivo' chiant' u zipp'*¹.

– Popular Lucanian wisdom

1: Tr.: Wherever you get, plant a stick.

The discussion presented up to this point has chronicled the birth, development and implementation of what Schrodinger would have called a *naive physicist's idea about organisms* [45]. As we stand at the threshold of the conclusion of both this manuscript and this academic project, it is an opportune moment to step back, reflect upon the journey undertaken, and discern the patterns that have emerged from the collective body of work presented, fig. 6.1.

At the heart of our scientific discourse was the exploration-exploitation (EE) paradigm, which was posited as a general dynamic principle for biological systems. It applies whenever the dynamics of a system arise from the interplay of (i) the variability introduced by stochastic state changes and (ii) a state-dependent optimisation of a biological function.

One obvious context in which the EE paradigm manifests itself is the evolutionary dynamics. In the simplest scenario, the latter results from the combined effect of random genetic mutations (exploration) and natural selection (exploitation). Crucially, this example showcases the self-referential nature of biological dynamics. For these reasons, we have devoted (ch. 3) to discussing the core concepts of evolution, ranging from biological foundations to modelling efforts, to algorithms inspired by the evolutionary processes.

Our foray into the realm of evolutionary biology had an underlying purpose, which was made explicit in (ch. 4). From the specific case of evolution, we learned the formal structure of a general exploration-exploitation dynamics (4.4)². We chiseled away the context-dependent details of the evolutionary dynamics to unveil and discuss the underlying context-free EE algorithm. The resulting theoretical picture is that of a stochastic evolution of a probability distribution on a functional landscape. The study of analytically tractable toy models allowed us to elucidate the main characteristics of the EE dynamics.

2: This, holding fixed the mathematical representation of the system – essentially a string of zeros and ones.

Then, we took the leap. We started (ch. 5) by arguing that the brain wiring dynamics – i.e., the development from birth to adulthood of a nervous system, here the graph of neurons and connections between them – is another manifestation of the EE paradigm. To work out the details, it is necessary to focus on a particular system, in our case it was the brain of the nematode *C. elegans*. Within the EE framework, we were able to specify a model of worm brain maturation with only six parameters, all inferred from data and all amenable to biological interpretation. We offered a putative, biologically realistic interpretation of the EE dynamics in terms of the synapse formation process.

Our main result is that a parsimonious characterisation of the adult *C. elegans* brain combined with our EE dynamics is able to quantitatively reproduce the entire developmental trajectory, as reconstructed




Figure 6.1: Visual summary of a PhD project. We started from the context of evolutionary dynamics to formulate a context-free exploration-exploitation (EE) problem, which is claimed to be general for a class of biological dynamics. We developed a theoretical framework, studied it analytically and developed simulations. We have then used it to tackle another type of biological dynamics, namely the developmental dynamics of a nervous system. In this context, we studied a specific system, the *C. elegans* brain. We have developed an EE white-box model of brain growth that has been validated throughout development.

experimentally by serial section electron microscopy. To the best of our knowledge, this stands as the first theoretical model of system-wide neurodevelopmental dynamics for a living system, that is (i) firmly anchored in experimental data across development and (ii) wholly interpretable. More generally, our results support the recently proposed view of brain wiring dynamics as driven by a set of simple and genetically encoded wiring rules.

There is one thorny issue that is the main theoretical bottleneck of the EE approach. What is the *biological function*? More specifically, what is the mathematical expression of the F metric? How to identify the relevant features? How to learn the topography of an F landscape from the data? These problems, of course, are much more general. These questions echo a foundational inquiry in physics: what is the energy function of a physical system? When it comes to complex systems – which includes, but is not limited to biological systems –, there is no Delphi’s oracle, the answer is nuanced, intricate³.

3: The lexicon varies as a function of the context, so that energy $f.$, cost $f.$, utility $f.$, fitness $f.$, our biological $f.$ (...) all essentially refer to the same mathematical entity.

To build a functional landscape, we have proposed here the use of an inferential approach that is both principled and data-driven. In particular, we used the maxent inference scheme associated with the exponential random graph (ERG) models. This approach was apt for studying the brain wiring dynamics because (i) the data were naturally represented as a graph, and (ii) inference needed to be drawn from a singular realisation of the system. In (ch. 2), we dwelt on the theoretical underpinnings of the ERG models and offered a pedagogical guide to assist the interested users in their application.

A fruitful theory is a theory that leaves the theorist with more and more precise questions about the subject matter than at the beginning. Several

have been sketched in sec. 5.5 for the brain wiring problem, and represent the agenda for the upcoming developments. Yet, I claimed generality⁴ for the EE paradigm and for such a claim demonstrating its applicability in at least two distinct contexts – evolution, brain wiring – was only the bare minimum. This manuscript should therefore serve as a guide to unifying and approaching new problems in biology along the same lines – some mentioned in the text, many more probably unforeseen by the writer himself. As is often the case at the beginnings, our progress is but a grain in the granary of what remains.

It is no job for the hasty, though. One profound lesson from physics resonates especially when it comes to the study of living systems. That is, general principles do not emerge unless we look for them. This is the central message that we hope will be a legacy of this work.

Here we stand. We shall resist the temptation to dismiss the phenomena of life as too messy for the physics-style of scientific inquiry. If we do so, then a theoretical physics of biological systems – built on solid, compelling principles and grounded in experimental data – becomes not only possible but also one of the most fascinating frontiers⁵ of modern physics.

4: Naturally, generality does not equate to explaining *everything*. There are numerous dynamical processes in biology that cannot be adequately described or explained within the exploration-exploitation framework. Just as gravity offers little insight into the bizarre nature of strong nuclear interactions. It is trivial to say, but better to say it.

5: I may not be completely impartial.

APPENDIX

We here provide a step-by-step derivation of the results discussed in sec. 4.2.2, 4.2.3, for EE dynamics under simple scenarios. The three fundamental ingredients are (i) the EE dynamics in the continuous time limit (4.8) (ii) a formal specification of the F metric as in (4.11), (4.19) and (iii) the dynamic of any graph observable $O : \mathcal{G} \mapsto \mathbb{R}$, i.e.,

$$\frac{d}{dt} \langle O \rangle_t = \frac{d}{dt} \sum_G O(G) P(G, t) = \sum_G O(G) \frac{d}{dt} P(G, t). \quad (\text{A.1})$$

Energy-like biological function

Consider the F metric (4.11). The exploitation term in the EE dynamics (4.8) can be written as

$$[F(G) - \langle F \rangle_t] P(G, t) = -\frac{1}{L} \sum_{i < j} [a_{ij} - \langle a_{ij} \rangle_t] P(G, t), \quad (\text{A.2})$$

$$\stackrel{(a)}{=} -\frac{1}{L} \sum_{i < j} \left[\frac{1 + \sigma_{ij}}{2} - \left\langle \frac{1 + \sigma_{ij}}{2} \right\rangle_t \right] P(G, t), \quad (\text{A.3})$$

$$= -\frac{1}{2L} \sum_{i < j} [\sigma_{ij} - \langle \sigma_{ij} \rangle_t] P(G, t), \quad (\text{A.4})$$

where in (a) we have used (4.6). Note that, when switching from the bit-wise to the spin-wise representation of the dyadic variables, it is also implied that:

$$\sum_G = \sum_{a_{11}=0,1} \cdots \sum_{a_{LL}=0,1} = \sum_{\sigma_{11}=\pm 1} \cdots \sum_{\sigma_{LL}=\pm 1}. \quad (\text{A.5})$$

The dynamics of the expected value $\langle \sigma_{ij} \rangle$ can be evaluated using (A.1).

$$\begin{aligned} \frac{d}{dt} \langle \sigma_{ij} \rangle_t &= \sum_G \sigma_{ij} \left\{ \mu \sum_{k < l} [P(M_{kl}G, t) - P(G, t)] - \frac{\varphi}{2L} \sum_{k < l} [\sigma_{kl} - \langle \sigma_{kl} \rangle_t] P(G, t) \right\} \\ &\stackrel{(a)}{=} -2\mu \langle \sigma_{ij} \rangle_t - \frac{\varphi}{2L} \sum_{k < l} \sum_G [\sigma_{ij} \sigma_{kl} - \langle \sigma_{ij} \rangle_t \langle \sigma_{kl} \rangle_t] P(G, t) \\ &= -2\mu \langle \sigma_{ij} \rangle_t - \frac{\varphi}{2L} \sum_{k < l} [\langle \sigma_{ij} \sigma_{kl} \rangle_t - \langle \sigma_{ij} \rangle_t \langle \sigma_{kl} \rangle_t], \end{aligned} \quad (\text{A.6})$$

where in (a) we have used (4.9). The latter corresponds to (4.12). Under the hypothesis of decoupling approximation (4.13), we discard all terms in the last sum except $\langle \sigma_{ij} \sigma_{ij} \rangle_t - \langle \sigma_{ij} \rangle_t \langle \sigma_{ij} \rangle_t = 1 - \langle \sigma_{ij} \rangle_t^2$ and obtain (4.14). The same differential equation holds for the magnetisation,

$$\dot{m}_t = -2\mu m_t - \frac{\varphi}{2L} [1 - m_t^2]. \quad (\text{A.7})$$

Solving (A.7) is a simple calculus exercise – separable variables, partial fraction decomposition. The general solution is

$$m_t = m_2 \left[1 + \frac{m_1/m_2 - 1}{1 + c e^{2\mu t} \sqrt{1 + (\varphi/2L)^2}} \right]. \quad (\text{A.8})$$

Fixing the constant c by requiring $m_{t_0} = m_0$ results in (4.15).

Distance-like biological function

We follow the exact same steps as in the previous section. Consider the F metric (4.19), we can rewrite it as

$$\begin{aligned}
 F(G) &\stackrel{(a)}{=} -\frac{1}{L^2} \left[\sum_{i<j} \frac{1+\sigma_{ij}}{2} - E^* \right]^2 \\
 &= -\frac{1}{L^2} \left[\left(\frac{L}{2} - E^* \right)^2 + \left(\frac{L}{2} - E^* \right) \sum_{i<j} \sigma_{ij} + \frac{1}{4} \sum_{i<j, k<l} \sigma_{ij} \sigma_{kl} \right] \\
 &= -\frac{1}{L^2} \left[\left(\frac{L}{2} - E^* \right)^2 + \left(\frac{L}{2} - E^* \right) \sum_{i<j} \sigma_{ij} + \frac{L}{4} + \frac{1}{4} \sum_{\substack{i<j, k<l \\ (i,j) \neq (k,l)}} \sigma_{ij} \sigma_{kl} \right] \tag{A.9}
 \end{aligned}$$

where in (a) we used (4.6). The exploitation term in the EE dynamics (4.8) can be written as

$$[F(G) - \langle F \rangle_t] P(G, t) = -\frac{1}{L^2} \left[\left(\frac{L}{2} - E^* \right) \sum_{i<j} [\sigma_{ij} - \langle \sigma_{ij} \rangle_t] + \frac{1}{4} \sum_{\substack{i<j, k<l \\ (i,j) \neq (k,l)}} [\sigma_{ij} \sigma_{kl} - \langle \sigma_{ij} \sigma_{kl} \rangle_t] \right] P(G, t), \tag{A.10}$$

where all the terms that are constant in (A.9) cancel out. The dynamics of the expected value $\langle \sigma_{ij} \rangle$ can be evaluated using (A.1).

$$\begin{aligned}
 \frac{d}{dt} \langle \sigma_{ij} \rangle_t &= -2\mu \langle \sigma_{ij} \rangle_t - \frac{\varphi}{L^2} \left[\left(\frac{L}{2} - E^* \right) \sum_{k<l} [\langle \sigma_{ij} \sigma_{kl} \rangle_t - \langle \sigma_{ij} \rangle_t \langle \sigma_{kl} \rangle_t] + \right. \\
 &\quad \left. + \frac{1}{4} \sum_{\substack{k<l, m<n \\ (k,l) \neq (m,n)}} [\langle \sigma_{ij} \sigma_{kl} \sigma_{mn} \rangle_t - \langle \sigma_{ij} \rangle_t \langle \sigma_{kl} \sigma_{mn} \rangle_t] \right] P(G, t), \\
 &= -2\mu \langle \sigma_{ij} \rangle_t - \frac{\varphi}{L^2} \left[\left(\frac{L}{2} - E^* \right) \left[1 - \langle \sigma_{ij} \rangle_t^2 + \sum_{\substack{i<j, k<l \\ (i,j) \neq (k,l)}} [\langle \sigma_{ij} \sigma_{kl} \rangle_t - \langle \sigma_{ij} \rangle_t \langle \sigma_{kl} \rangle_t] \right] + \right. \\
 &\quad \left. + \frac{1}{2} \sum_{\substack{k<l \\ (i,j) \neq (k,l)}} [\langle \sigma_{kl} \rangle_t - \langle \sigma_{ij} \rangle_t \langle \sigma_{ij} \sigma_{kl} \rangle_t] + \right. \\
 &\quad \left. + \frac{1}{4} \sum_{\substack{k<l, m<n \\ (i,j) \neq (k,l) \neq (m,n)}} [\langle \sigma_{ij} \sigma_{kl} \sigma_{mn} \rangle_t - \langle \sigma_{ij} \rangle_t \langle \sigma_{kl} \sigma_{mn} \rangle_t] \right] P(G, t). \tag{A.11}
 \end{aligned}$$

Enforcing now the decoupling approximation (4.20) we get:

$$\frac{d}{dt} \langle \sigma_{ij} \rangle_t = -2\mu \langle \sigma_{ij} \rangle_t - \frac{\varphi}{L^2} \left[\left(\frac{L}{2} - E^* \right) \left[1 - \langle \sigma_{ij} \rangle_t^2 \right] + \frac{1}{2} \left[1 - \langle \sigma_{ij} \rangle_t^2 \right] \sum_{\substack{k<l \\ (i,j) \neq (k,l)}} \langle \sigma_{kl} \rangle_t \right] \tag{A.12}$$

Using the same initial conditions for all dyads, the last sum can be approximated as $\sim (L-1) \langle \sigma_{ij} \rangle$. The same differential equation can then be written for the magnetisation m_t , the result is precisely (4.21).

Network measures

B

The general purpose of network measures is *descriptive*. They do not add information, quite the opposite. Computing measures on a network entails excluding all information except that which relates to the specific attribute of the network we seek to illuminate. We here provide a synthetic summary¹ of the network measures used in sec. 5.3.3. Popular libraries such as NetworkX for Python or igraph for R/Python can be used for computation.

Let us consider, as usual, an undirected, unweighted graph (or network) G (1.1). The degree of the node i , we recall, is defined as:

$$k_i = \sum_j a_{ij}. \quad (\text{B.1})$$

We can categorise our measures into three groups, based on the specific network feature they examine: clustering, efficiency and degrees.

Network motifs

Much of the content of this manuscript is based on the enumeration of network *motifs*. Essentially, these are identifiable patterns or subgraphs within a network that occur more frequently than would be statistically expected in a random network [247]. Identifying a motif count entails (i) counting the occurrences of a given subgraph and (ii) evaluating its statistical significance.

In a few fortunate cases², the count of network motifs can be expressed in terms of powers of the adjacency matrix. The simplest motif count is the number of edges, trivially. The number of connected triples $\#_{\wedge}$ and that of triangles $\#_{\Delta}$ can be expressed as:

$$\#_{\wedge}(G) = \sum_{i < j, q} a_{iq} a_{jq} = \frac{1}{2} \sum_{i, j} (G^2)_{ij} - \text{Tr}(G^2). \quad (\text{B.2})$$

$$\#_{\Delta}(G) = \sum_{i < j < q} a_{ij} a_{jq} a_{iq} = \frac{1}{6} \text{Tr}(G^3), \quad (\text{B.3})$$

where Tr is the trace operator, fig. B.1. The number $x_d^{(k)}$ of nodes with degree k is simply

$$x_d^{(k)}(G) = \sum_i \delta_{k, k_i}, \quad (\text{B.4})$$

where δ is the Kronecker delta, a k_i as in (B.1). Finally, the number $x_{esp}^{(k)}$ of connected dyads whose extremal nodes share exactly k partners is found as:

$$x_{esp}^{(k)}(G) = \sum_{i < j} \delta_{k, B_{ij}} \text{ with } B = G^2 \odot G \quad (\text{B.5})$$

where \odot is the Hadamard (element-wise) product, $B_{ij} = \sum_q a_{iq} a_{qj} a_{ij}$. In ch. 5, the statistical significance of the motif counts was assessed either

1: Detailed discussions and long catalogues of other measures can be found in any monograph on network science, our main references are [39, 40]. We will not here enter into competition with the thousands of items in the academic literature that have covered these topics.

2: These happen to span the totality of the cases discussed in ch. 5. Our EE simulations gain considerable advantage – in terms of computation time – from explicit formulae for the computation of an F metric building blocks.




Figure B.1: A triangle (top) is a triple of nodes i, j, q with $a_{ij} = a_{iq} = a_{ja} = 1$. A connected triple (below) is a pair of edges $a_{ij} = a_{iq} = 1$. Each triangle contains three connected triples.

within the ERG framework (non-zero inferred parameter) or by direct comparison with a null model.

Clustering

A common property of a number of real-world networks – in particular, social networks – is the presence of tightly knit communities or groups [248, 249]. A straightforward manifestation of such a *clustering behaviour* is a higher-than-random connection probability for two nodes that share a common partner. In other words, if nodes i and j are both connected to node q , they are more likely to be directly connected to each other as well. Such a behaviour can be quantified by the following two metrics:

- *Transitivity, T* . It is the ratio between the number of existing triangles $\#_{\Delta}$ (B.3) and the number of connected triples $\#_{\wedge}$ (B.2). Formally,

$$T = 3 \frac{\#_{\Delta}(G)}{\#_{\wedge}(G)}, \quad T \in [0, 1]. \quad (\text{B.6})$$

- *Average clustering coefficient, C* . The local clustering coefficient for a node i is defined as

$$C_i = \frac{\# \text{ connected pairs of neighbors of } i}{\# \text{ pairs of neighbors of } i} = \frac{\sum_{j < l} a_{ij} a_{jl} a_{il}}{\frac{1}{2} k_i (k_i - 1)}. \quad (\text{B.7})$$




Figure B.2: Subnetwork induced by the node i and its neighbors (orange). Here, i has four neighbors. There exist six possible pairs of neighbors, two of which are connected, therefore $C_i = 1/3$.

It is computed by considering the subnetwork induced by the node i and its first neighbors and quantifies the relative number of neighbors of i that are also themselves neighbors, fig. B.2. The average clustering coefficient then simply takes the average value over the node set:

$$C = \frac{1}{N} \sum_{i=1}^N C_i, \quad C \in [0, 1]. \quad (\text{B.8})$$

Although both T and C approach a value of 1 in the limit of perfect transitivity, they do not convey identical information. The clustering coefficient is relatively more sensitive to nodes with low degrees because it averages across all nodes. On the contrary, transitivity is more affected by nodes with high degrees, as they are the ones that influence the number of triangles more. As a result, while transitivity provides a more comprehensive picture of the overall structure of the network, the clustering coefficient sheds more light on local structures or subnetworks embedded in the larger network.

Efficiency

Another class of metrics in network science is designed to quantify the *efficiency* of information or resource transmission within the network [250]. This notion is based on the fundamental premise that the proximity of two nodes in a network graph strongly correlates with the efficiency of their information exchange. The distance d_{ij} between two given nodes i, j is defined as the length of the geodesic between them, i.e., the number

of edges that form the shortest path from one to the other, fig. B.3. By definition, if i, j are disconnected, $d_{ij} = \infty$.

- *Global efficiency*. It is defined as the harmonic mean of geodesic lengths. The global efficiency E_g is defined as

$$E_g = \frac{1}{L} \sum_{i < j} \frac{1}{d_{ij}}. \quad (\text{B.9})$$

A network exhibiting high global efficiency typically features brief paths connecting any two nodes, implying swift information distribution throughout the network. This is observed in random and ‘small-world’ networks.

- *Local efficiency*. Once again, it is possible to measure the efficiency at a local level and then average this across all nodes. This metric, known as the local efficiency or E_l , is defined as:

$$E_l = \frac{1}{N} \sum_i E_g^{(i)}, \quad (\text{B.10})$$

where $E_g^{(i)}$ is the global efficiency of the subgraph induced by the node i and its neighbors. High local efficiency in a network means that removing a node would not significantly disrupt communication between its neighbours. For this reason, it is often considered a proxy for the robustness or resilience of the network to attacks. Much like global efficiency, local efficiency is a characteristic trait of small-world networks. However, unlike global efficiency, high local efficiency is not typically found in random networks.

Degrees

The degree distribution encapsulates fundamental information about the network’s structure, robustness to failures, and information spreading dynamics [83]. The *degree distribution* of a given network is simply given by

$$P_{deg}(k) = x_d^{(k)} / N, \quad (\text{B.11})$$

where $x_d^{(k)}$ is as in (B.4). Alternatively, one can look at the *cumulative degree distribution*

$$P_{deg}^{(c)}(k) = P_{deg}(j \geq k) = \frac{1}{N} \sum_{j \geq k} x_d^{(j)} \quad (\text{B.12})$$

Cumulative distributions provide a more effective visualization of the degree structure, smoothing out fluctuations and making it easier to identify long-tail behaviors indicative of scale-free networks³.




Figure B.3: The geodesic between the nodes i, j (orange) has length $d_{ij} = 2$. For comparison, another, longer path connecting the same two nodes is highlighted (cyan).

3: In particular, if the underlying distribution has a power-law behaviour $\sim k^{-\gamma}$, then the cumulative distribution goes as $\sim k^{-\gamma+1}$ [251]

Glossary: biology for physicists

caenorhabditis elegans Abbreviated as *C. elegans*. A free-living transparent nematode (roundworm) about 1 mm in length. It is often used as a model organism in biological research, due to its simplicity, well-defined anatomy, and short life cycle. Despite this, it possesses a wide behavioral array. Beyond the basics of locomotion, foraging, and feeding, the worm can discern and navigate towards or away from various chemicals, odors, temperature gradients, and food sources. Furthermore, it demonstrates social awareness, detecting the presence, density, and even sex of neighboring nematodes. 47

Darwinism Charles Darwin's theory explaining the mechanism of evolution by natural selection. According to this theory, evolution results from the interaction of three principles: heredity, variation and survival (natural selection). It describes how species adapt to their environment over time. 24

evolution The gradual change in inherited characteristics of biological populations over successive generations. It results from several different processes: (i) mutations, changes in the genetic sequence (ii) recombinations, exchanges of genetic material between individuals (iii) genetic drift, random changes in gene frequency and (iv) natural selection (see entry: *selection*). 23

fitness Organism's ability to survive and reproduce in a given environment. It is a measure of the relative reproductive success of individuals with specific traits. Often defined as proportional to the average number of offspring of an individual. 26

fitness landscape A metaphorical representation used in evolutionary biology to illustrate the relationship between genotypes and their associated fitness in a given environment. In this landscape, each point represents a unique genotype, and the elevation at that point represents the corresponding fitness. Peaks on the landscape represent optimal genotypes with high fitness. The structure of the fitness landscape influences the paths evolution may take. 28

function The specific role or task performed by a component (e.g., molecule, cell, organ) within a living organism. They are essential for the organism's survival, growth, and reproduction, therefore subject to evolutionary pressure. 34

genotype The genetic makeup of an organism, representing the specific combination of genes present in its DNA. It serves as the blueprint for the organism's traits and characteristics. 24

inheritance The process by which genetic information is passed from one generation to the next. This transfer of genetic material occurs during reproduction, ensuring the continuity and maintenance of traits within a species. 24

nervous system A network of specialized cells (neurons) that coordinate and regulate the activities of an organism. In *C. elegans*, the nervous system of an adult hermaphrodite consists of 302 neurons, uniquely identifiable. Most of them are found in clusters, called *ganglia*. Neuronal processes extend from the ganglia and travel in longitudinal nerve bundles to different regions of the nervous system. The most prominent are the nerve ring, ventral nerve cord and dorsal nerve cord. 48

neuron Also, or nerve cell. It is the basic structural and functional unit of the nervous system. Neurons transmit information using electrical and chemical signals. They consist of a cell body, dendrites (receiving inputs), and an axon (transmitting outputs). Neurons play a crucial role in processing and transmitting information in the nervous system. The morphology of a neuron can vary substantially. In *C. elegans*, they are mostly unipolar or bipolar.. 48

phenotype The observable characteristics and traits of an organism. In general, the phenotype is determined by both the genomic makeup (genotype) and environmental factors. It includes features like appearance, behavior, and physiological functions. 24

- selection** The mechanism through which certain heritable traits confer advantages to individuals, increasing their likelihood of survival and reproduction. It is driven by the interplay between organisms and their environment, favoring traits that enhance an organism's fitness for its ecological niche. 24
- self-referential** Said of biological dynamics, where the update rules change during the time evolution of the system, in a manner that depends on the state and thus on the history of the system.. 37
- statistical genetics** In the sense of Neher-Shraiman, a statistical multilocus theory that explains how the laws of quantitative genetics – i.e., the study of phenotypic variation among individuals – emerge from the stochastic evolutionary dynamics in the space of genotypes. 26
- synapse** A specialized junction between two neurons, where information is transferred from one cell to another. There can be electrical or chemical. The former, also called *gap junctions*, are specialized channels that directly connect the cytoplasm of adjacent cells, allowing various molecules, ions, and electrical impulses to pass between the cells. The latter, *chemical synapses*, function as specialised junctions that facilitate the one-way relay of chemical signals, or neurotransmitters, from a presynaptic to one or more postsynaptic cells. 48
- trait** A specific characteristic or feature of an organism that can be inherited or influenced by environmental factors. Traits contribute to the overall phenotype and are subject to evolutionary pressure. 24
- variation** Diversity observed in the traits and characteristics among individuals within a population, arising from genetic mutations, recombination, and other sources. It introduces differences in physical and behavioral attributes, providing the raw material upon which natural selection acts. It is crucial for a population's adaptability to changing environments. 24

Bibliography

Here are the references in citation order.

- [1] Vito Dichio and Fabrizio De Vico Fallani. 'The exploration-exploitation paradigm for networked biological systems'. In: *arXiv e-prints* (June 2023) (cited on pages xv, 34, 47, 57).
- [2] Vito Dichio and Fabrizio De Vico Fallani. 'Statistical models of complex brain networks'. In: *Reports on Progress in Physics* 86.10 (2023), p. 102601 (cited on pages xv, 9, 10, 13, 16, 20).
- [3] Vito Dichio, Hong-Li Zeng, and Erik Aurell. 'Statistical genetics in and out of quasi-linkage equilibrium'. In: *Reports on Progress in Physics* 86.5 (2023), p. 052601 (cited on pages xv, 23, 25–27, 29).
- [4] Hong-Li Zeng et al. 'Temporal epistasis inference from more than 3 500 000 SARS-CoV-2 genomic sequences'. In: *Physical Review E* 106.4 (2022), p. 044409 (cited on pages xvi, 27).
- [5] Hong-Li Zeng et al. 'Inferring epistasis from genomic data with comparable mutation and outcrossing rate'. In: *Journal of Statistical Mechanics: Theory and Experiment* 2021.8 (2021), p. 083501 (cited on pages xvi, 27).
- [6] Andrea Cavagna et al. 'Natural swarms in 3.99 dimensions'. In: *Nature Physics* (2023), pp. 1–7 (cited on page 1).
- [7] Benjamin H Good et al. 'The dynamics of molecular evolution over 60,000 generations'. In: *Nature* 551.7678 (2017), pp. 45–50 (cited on page 1).
- [8] Leenoy Meshulam et al. 'Coarse graining, fixed points, and scaling in a large population of neurons'. In: *Physical review letters* 123.17 (2019), p. 178103 (cited on page 1).
- [9] Katherine Copenhagen et al. 'Topological defects promote layer formation in *Myxococcus xanthus* colonies'. In: *Nature Physics* 17.2 (2021), pp. 211–215 (cited on page 1).
- [10] John R Loofbourow. 'Borderland problems in biology and physics'. In: *Reviews of modern physics* 12.4 (1940), p. 267 (cited on pages 1, 2).
- [11] Hans Frauenfelder, Peter G Wolynes, and Robert H Austin. 'Biological physics'. In: *Reviews of Modern Physics* 71.2 (1999), S419 (cited on page 1).
- [12] William Bialek. 'Perspectives on theory at the interface of physics and biology'. In: *Reports on Progress in Physics* 81.1 (2017), p. 012601 (cited on pages 1, 4, 38).
- [13] National Academies of Sciences Engineering and Medicine. *Physics of life*. The National Academies Press, 2022 (cited on pages 1, 2).
- [14] William Bialek. *Biophysics: searching for principles*. Princeton University Press, 2012 (cited on pages 1, 2, 34).
- [15] Nigel Goldenfeld and Carl Woese. 'Life is physics: evolution as a collective phenomenon far from equilibrium'. In: *Annu. Rev. Condens. Matter Phys.* 2.1 (2011), pp. 375–399 (cited on pages 1, 23, 37).
- [16] Philip Charles Nelson et al. *Biological physics: energy, information, life*. WH Freeman New York, 2008 (cited on page 1).
- [17] Kenneth P Burnham et al. *Practical use of the information-theoretic approach*. Springer, 1998 (cited on pages 1, 4, 19).
- [18] Rodolphe Barrangou and Jennifer A Doudna. 'Applications of CRISPR technologies in research and beyond'. In: *Nature biotechnology* 34.9 (2016), pp. 933–941 (cited on page 2).
- [19] Nikos K Logothetis. 'What we can do and what we cannot do with fMRI'. In: *Nature* 453.7197 (2008), pp. 869–878 (cited on page 2).
- [20] Hans Frauenfelder. 'Ask not what physics can do for biology—ask what biology can do for physics'. In: *Physical biology* 11.5 (2014), p. 053004 (cited on page 2).

- [21] Walter Bradford Cannon. 'The wisdom of the body'. In: (1939) (cited on page 3).
- [22] KP Ivanov. 'The development of the concepts of homeothermy and thermoregulation'. In: *Journal of Thermal Biology* 31.1-2 (2006), pp. 24–29 (cited on page 3).
- [23] Dag O Hessen et al. 'Ecological stoichiometry: an elementary approach using basic principles'. In: *Limnology and Oceanography* 58.6 (2013), pp. 2219–2236 (cited on page 3).
- [24] Val H Smith. 'Eutrophication of freshwater and coastal marine ecosystems a global problem'. In: *Environmental Science and Pollution Research* 10 (2003), pp. 126–139 (cited on page 3).
- [25] Patricia Sánchez-Baracaldo and Tanai Cardona. 'On the origin of oxygenic photosynthesis and Cyanobacteria'. In: *New Phytologist* 225.4 (2020), pp. 1440–1446 (cited on page 4).
- [26] Lisa R Girard et al. 'WormBook: the online review of *Caenorhabditis elegans* biology'. In: *Nucleic acids research* 35.suppl_1 (2007), pp. D472–D475 (cited on page 4).
- [27] Yosuke Koga. 'Thermal adaptation of the archaeal and bacterial lipid membranes'. In: *Archaea* 2012 (2012) (cited on page 4).
- [28] Francis S Collins, Michael Morgan, and Aristides Patrinos. 'The Human Genome Project: lessons from large-scale biology'. In: *Science* 300.5617 (2003), pp. 286–290 (cited on page 5).
- [29] Elaine R Mardis. 'The impact of next-generation sequencing technology on genetics'. In: *Trends in genetics* 24.3 (2008), pp. 133–141 (cited on page 5).
- [30] Ruedi Aebersold and Matthias Mann. 'Mass-spectrometric exploration of proteome structure and function'. In: *Nature* 537.7620 (2016), pp. 347–355 (cited on page 5).
- [31] Kevin L Briggman and Davi D Bock. 'Volume electron microscopy for neuronal circuit reconstruction'. In: *Current opinion in neurobiology* 22.1 (2012), pp. 154–161 (cited on page 5).
- [32] Chris Anderson. 'The end of theory: The data deluge makes the scientific method obsolete'. In: *Wired magazine* 16.7 (2008), pp. 16–07 (cited on page 5).
- [33] Sabina Leonelli. 'What difference does quantity make? On the epistemology of Big Data in biology'. In: *Big data & society* 1.1 (2014), p. 2053951714534395 (cited on page 5).
- [34] Hykel Hosni and Angelo Vulpiani. 'Data science and the art of modelling'. In: *Lettera Matematica* 6 (2018), pp. 121–129 (cited on page 5).
- [35] Lisa Gitelman. *Raw data is an oxymoron*. MIT press, 2013 (cited on page 5).
- [36] Alon Halevy, Peter Norvig, and Fernando Pereira. 'The unreasonable effectiveness of data'. In: *IEEE intelligent systems* 24.2 (2009), pp. 8–12 (cited on page 5).
- [37] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 'Deep learning'. In: *nature* 521.7553 (2015), pp. 436–444 (cited on page 5).
- [38] Paul Cilliers. *Complexity and postmodernism: Understanding complex systems*. routledge, 2002 (cited on page 5).
- [39] Mark Newman. *Networks*. Oxford university press, 2018 (cited on pages 5, 6, 79).
- [40] Vito Latora, Vincenzo Nicosia, and Giovanni Russo. *Complex networks: principles, methods and applications*. Cambridge University Press, 2017 (cited on pages 6, 79).
- [41] Carter T Butts. 'Revisiting the foundations of network analysis'. In: *science* 325.5939 (2009), pp. 414–416 (cited on page 6).
- [42] Leto Peel, Tiago P Peixoto, and Manlio De Domenico. 'Statistical inference links data and theory in network science'. In: *Nature Communications* 13.1 (2022), p. 6794 (cited on page 6).
- [43] Eugene P Wigner. 'The unreasonable effectiveness of mathematics in the natural sciences'. In: *Mathematics and science*. World Scientific, 1990, pp. 291–306 (cited on page 6).
- [44] Olaf Sporns. *Networks of the Brain*. MIT press, 2016 (cited on page 6).
- [45] Erwin Schrödinger. *What is Life?* Cambridge University Press, Cambridge, 1944 (cited on pages 7, 71).
- [46] David Roxbee Cox. *Principles of statistical inference*. Cambridge university press, 2006 (cited on page 9).

- [47] Giulio Cimini et al. 'The statistical physics of real-world networks'. In: *Nature Reviews Physics* 1.1 (2019), pp. 58–71 (cited on pages 9, 10, 12, 13).
- [48] Juyong Park and Mark EJ Newman. 'Statistical mechanics of networks'. In: *Physical Review E* 70.6 (2004), p. 066117 (cited on pages 9, 10, 12, 13, 20).
- [49] David R Hunter et al. 'ergm: A package to fit, simulate and diagnose exponential-family models for networks'. In: *Journal of statistical software* 24.3 (2008), nihpa54860 (cited on pages 9, 15, 16, 18, 22).
- [50] Pavel N. Krivitsky et al. *Statnet: Tools for the Statistical Modeling of Network Data*. <https://statnet.org>. Statnet Development Team, 2003–2022 (cited on pages 9, 16).
- [51] Rolf Sundberg. *Statistical modelling by exponential families*. Vol. 12. Cambridge University Press, 2019 (cited on page 10).
- [52] Frank Nielsen and Vincent Garcia. 'Statistical exponential families: A digest with flash cards'. In: *arXiv e-prints* (Nov. 2009), arXiv:0911.4863 (cited on page 10).
- [53] Georges Darmois. 'Sur les lois de probabilitée estimation exhaustive'. In: *CR Acad. Sci. Paris* 260.1265 (1935), p. 85 (cited on page 10).
- [54] Bernard O Koopman. 'On distributions admitting a sufficient statistic'. In: *Transactions of the American Mathematical society* 39.3 (1936), pp. 399–409 (cited on page 10).
- [55] Edwin James George Pitman. 'Sufficient statistics and intrinsic accuracy'. In: *Mathematical Proceedings of the cambridge Philosophical society*. Vol. 32. Cambridge University Press. 1936, pp. 567–579 (cited on page 10).
- [56] Ole Barndorff-Nielsen. *Information and exponential families: in statistical theory*. John Wiley & Sons, 2014 (cited on page 10).
- [57] Lawrence D Brown. *Fundamentals of statistical exponential families: with applications in statistical decision theory*. Ims, 1978 (cited on page 10).
- [58] Paul W Holland and Samuel Leinhardt. 'An exponential family of probability distributions for directed graphs'. In: *Journal of the american Statistical association* 76.373 (1981), pp. 33–50 (cited on pages 10, 16).
- [59] Ove Frank and David Strauss. 'Markov Graphs'. In: *Journal of the American Statistical Association* 81.395 (1986), pp. 832–842 (cited on page 10).
- [60] David Strauss. 'On a general class of models for interaction'. In: *SIAM review* 28.4 (1986), pp. 513–527 (cited on pages 10, 19, 20).
- [61] Dean Lusher, Johan Koskinen, and Garry Robins. *Exponential random graph models for social networks: Theory, methods, and applications*. Cambridge University Press, 2013 (cited on page 10).
- [62] Saeid Ghafouri and Seyed Hossein Khasteh. 'A survey on exponential random graph models: an application perspective'. In: *PeerJ Computer Science* 6 (2020), e269 (cited on page 10).
- [63] Edwin T. Jaynes. 'Information theory and statistical mechanics'. In: *Physical review* 106.4 (1957), p. 620 (cited on pages 10, 12).
- [64] Edwin T. Jaynes. 'Information theory and statistical mechanics. II'. In: *Physical review* 108.2 (1957), p. 171 (cited on pages 10, 12).
- [65] Edwin T Jaynes. *Probability theory: The logic of science*. Cambridge university press, 2003 (cited on pages 10–12, 17).
- [66] Edwin T. Jaynes. 'On the rationale of maximum-entropy methods'. In: *Proceedings of the IEEE* 70.9 (1982), pp. 939–952 (cited on page 10).
- [67] Julian Besag. 'Spatial Interaction and the Statistical Analysis of Lattice Systems'. In: *Journal of the Royal Statistical Society. Series B (Methodological)* 36.2 (1974), pp. 192–236 (cited on pages 10, 15, 19).
- [68] Claude E. Shannon. 'A mathematical theory of communication'. In: *The Bell system technical journal* 27.3 (1948), pp. 379–423 (cited on page 10).
- [69] Thomas M Cover and Joy A. Thomas. *Elements of information theory*. John Wiley & Sons, 1999 (cited on pages 10, 11, 16).

- [70] Zoltan Dienes. 'Bayesian versus orthodox statistics: Which side are you on?' In: *Perspectives on Psychological Science* 6.3 (2011), pp. 274–290 (cited on page 11).
- [71] Erik Aurell. 'The maximum entropy fallacy redux?' In: *PLoS computational biology* 12.5 (2016), e1004777 (cited on page 11).
- [72] Gennaro Auletta, Lamberto Rondoni, and Angelo Vulpiani. 'On the relevance of the maximum entropy principle in non-equilibrium statistical mechanics'. In: *The European Physical Journal Special Topics* 226.10 (2017), pp. 2327–2343 (cited on page 11).
- [73] Kerson Huang. *Statistical Mechanics*. 2nd ed. John Wiley & Sons, 1987 (cited on page 12).
- [74] Luca Peliti. 'Introduction to the statistical theory of Darwinian evolution'. In: *arXiv e-prints* (Dec. 1997), cond-mat/9712027 (cited on pages 12, 13, 26).
- [75] Ginestra Bianconi. 'Entropy of network ensembles'. In: *Physical Review E* 79.3 (2009), p. 036114 (cited on page 12).
- [76] Carl G Hempel and Paul Oppenheim. 'Studies in the Logic of Explanation'. In: *Philosophy of science* 15.2 (1948), pp. 135–175 (cited on page 13).
- [77] Wesley C Salmon. *Four decades of scientific explanation*. University of Pittsburgh Press, 2006 (cited on page 13).
- [78] Thierry Mora et al. 'Maximum entropy models for antibody diversity'. In: *Proceedings of the National Academy of Sciences* 107.12 (2010), pp. 5405–5410 (cited on page 13).
- [79] Xiaowen Chen et al. 'Searching for collective behavior in a small brain'. In: *Physical Review E* 99.5 (2019), p. 052418 (cited on page 13).
- [80] H Chau Nguyen, Riccardo Zecchina, and Johannes Berg. 'Inverse statistical problems: from the inverse Ising problem to data science'. In: *Advances in Physics* 66.3 (2017), pp. 197–261 (cited on page 13).
- [81] David JC MacKay. *Information theory, inference and learning algorithms*. Cambridge university press, 2003 (cited on pages 13, 14).
- [82] Paul Erdős, Alfréd Rényi, et al. 'On the evolution of random graphs'. In: *Publ. math. inst. hung. acad. sci* 5.1 (1960), pp. 17–60 (cited on page 14).
- [83] Réka Albert and Albert-László Barabási. 'Statistical mechanics of complex networks'. In: *Reviews of modern physics* 74.1 (2002), p. 47 (cited on pages 14, 81).
- [84] Richard J Rossi. *Mathematical statistics: an introduction to likelihood based inference*. John Wiley & Sons, 2018 (cited on page 14).
- [85] Marc Mézard, Giorgio Parisi, and Miguel Angel Virasoro. *Spin glass theory and beyond: An Introduction to the Replica Method and Its Applications*. Vol. 9. World Scientific Publishing Company, 1987 (cited on pages 15, 29, 40).
- [86] Tommaso Castellani and Andrea Cavagna. 'Spin-glass theory for pedestrians'. In: *Journal of Statistical Mechanics: Theory and Experiment* 2005.05 (2005), P05012 (cited on page 15).
- [87] Juyong Park and Mark EJ Newman. 'Solution of the two-star model of a network'. In: *Physical Review E* 70.6 (2004), p. 066146 (cited on pages 15, 19, 20).
- [88] Juyong Park and Mark EJ Newman. 'Solution for the properties of a clustered network'. In: *Physical Review E* 72.2 (2005), p. 026136 (cited on pages 15, 20).
- [89] Giorgio Parisi. *Statistical Field Theory*. Avalon Publishing, 1998 (cited on page 15).
- [90] Charles J Geyer. 'Markov chain Monte Carlo maximum likelihood'. In: (1991) (cited on page 15).
- [91] Pavel N. Krivitsky et al. 'ergm 4: New features for analyzing exponential-family random graph models'. In: *Journal of Statistical Software* 105 (2023), pp. 1–44 (cited on pages 16, 22).
- [92] Ronald A Fisher. 'On the mathematical foundations of theoretical statistics'. In: *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character* 222.594-604 (1922), pp. 309–368 (cited on page 18).

- [93] Kenji Kira and Larry A Rendell. 'The feature selection problem: Traditional methods and a new algorithm'. In: *Proceedings of the tenth national conference on Artificial intelligence*. 1992, pp. 129–134 (cited on page 18).
- [94] Isabelle Guyon and André Elisseeff. 'An introduction to variable and feature selection'. In: *Journal of machine learning research* 3.Mar (2003), pp. 1157–1182 (cited on page 18).
- [95] Tom A. B. Snijders. 'Markov chain Monte Carlo estimation of exponential random graph models'. In: *Journal of Social Structure* 3.2 (2002), pp. 1–40 (cited on page 20).
- [96] Alessandro Rinaldo, Stephen E Fienberg, and Yi Zhou. 'On the geometry of discrete exponential families with application to exponential random graph models'. In: *Electronic Journal of Statistics* 3 (2008), pp. 446–484 (cited on page 20).
- [97] Michael Schweinberger. 'Instability, sensitivity, and degeneracy of discrete exponential families'. In: *Journal of the American Statistical Association* 106.496 (2011), pp. 1361–1370 (cited on page 20).
- [98] Michael Schweinberger et al. 'Exponential-Family Models of Random Graphs: Inference in Finite, Super and Infinite Population Scenarios'. In: *Statistical Science* 35.4 (2020), pp. 627–662 (cited on page 20).
- [99] Tom AB Snijders et al. 'New specifications for exponential random graph models'. In: *Sociological methodology* 36.1 (2006), pp. 99–153 (cited on pages 20, 21).
- [100] David R Hunter and Mark S Handcock. 'Inference in curved exponential family models for networks'. In: *Journal of Computational and Graphical Statistics* 15.3 (2006), pp. 565–583 (cited on pages 20–22).
- [101] Bradley Efron. 'Defining the curvature of a statistical problem (with applications to second order efficiency)'. In: *The Annals of Statistics* (1975), pp. 1189–1242 (cited on page 20).
- [102] Bradley Efron. 'The geometry of exponential families'. In: *The Annals of Statistics* (1978), pp. 362–376 (cited on page 20).
- [103] David R Hunter. 'Curved exponential family models for social networks'. In: *Social networks* 29.2 (2007), pp. 216–230 (cited on page 20).
- [104] Ginestra Bianconi et al. 'Triadic closure as a basic generating mechanism of communities in complex networks'. In: *Physical Review E* 90.4 (2014), p. 042806 (cited on page 22).
- [105] Charles R. Darwin. *On the Origin of Species by Means of Natural Selection*. Murray, 1859 (cited on page 23).
- [106] Michael R Rose. *Darwin's spectre: evolutionary biology in the modern world*. Princeton University Press, 2000 (cited on page 23).
- [107] Ernst Mayr. 'The philosophical foundations of Darwinism'. In: *Proceedings of the American Philosophical Society* 145.4 (2001), pp. 488–495 (cited on page 23).
- [108] Julian Huxley. 'Evolution. The modern synthesis.' In: *Evolution. The Modern Synthesis*. (1942) (cited on page 23).
- [109] Christoph Adami. 'What is complexity?' In: *BioEssays* 24.12 (2002), pp. 1085–1094 (cited on page 23).
- [110] Christoph Adami. 'The use of information theory in evolutionary biology'. In: *Annals of the New York Academy of Sciences* 1256.1 (2012), pp. 49–65 (cited on page 23).
- [111] Yuri I Wolf, Mikhail I Katsnelson, and Eugene V Koonin. 'Physical foundations of biological complexity'. In: *Proceedings of the National Academy of Sciences* 115.37 (2018), E8678–E8687 (cited on page 23).
- [112] William S. Klug et al. *Concepts of Genetics*. 12th ed. Always Learning. Pearson, 2019 (cited on page 24).
- [113] Daniel L. Hartl and Andrew G. Clark. *Principles of population genetics*. 4th ed. Vol. 116. Sinauer associates, 1997 (cited on pages 24, 25).
- [114] Susanna Manrubia et al. 'From genotypes to organisms: State-of-the-art and perspectives of a cornerstone in evolutionary dynamics'. In: *Physics of Life Reviews* 38 (2021), pp. 55–106 (cited on pages 24, 28, 29).

- [115] Ronald A. Fisher. 'Croonian Lecture-Population genetics'. In: *Proceedings of the Royal Society of London. Series B-Biological Sciences* 141.905 (1953), pp. 510–523 (cited on page 25).
- [116] C. Radhakrishna Rao. 'Ronald A. Fisher: The founder of modern statistics'. In: *Statistical science* (1992), pp. 34–48 (cited on page 25).
- [117] Matthew B. Hamilton. *Population Genetics*. 2nd ed. John Wiley & Sons, 2021 (cited on page 25).
- [118] Charles H. Pence. 'Sir John FW Herschel and Charles Darwin: Nineteenth-century science and its methodology'. In: *Hopos: the Journal of the International Society for the History of Philosophy of Science* 8.1 (2018), pp. 108–140 (cited on page 25).
- [119] Brian Charlesworth and Deborah Charlesworth. 'Population genetics from 1966 to 2016'. In: *Heredity* 118.1 (2017), pp. 2–9 (cited on page 25).
- [120] James F. Crow. *An introduction to population genetics theory*. Scientific Publishers, 2017 (cited on page 25).
- [121] Richard A Neher and Boris I Shraiman. 'Statistical genetics and evolution of quantitative traits'. In: *Reviews of Modern Physics* 83.4 (2011), p. 1283 (cited on pages 25–27).
- [122] Ronald A. Fisher. *The genetical theory of natural selection*. Clarendon Press, 1930 (cited on page 25).
- [123] Guy Sella and Aaron E Hirsh. 'The application of statistical physics to evolutionary biology'. In: *Proceedings of the National Academy of Sciences* 102.27 (2005), pp. 9541–9546 (cited on page 26).
- [124] Riccardo Rao and Stanislas Leibler. 'Evolutionary dynamics, evolutionary forces, and robustness: A nonequilibrium statistical mechanics perspective'. In: *Proceedings of the National Academy of Sciences* 119.13 (2022), e2112083119 (cited on page 26).
- [125] Crispin W. Gardiner. *Handbook of stochastic methods for physics, chemistry, and the natural sciences*. Springer, 1985 (cited on page 27).
- [126] Richard A Neher and Boris I Shraiman. 'Competition between recombination and epistasis can cause a transition from allele to genotype selection'. In: *Proceedings of the National Academy of Sciences* 106.16 (2009), pp. 6866–6871 (cited on page 27).
- [127] Eugenio Mauri, Simona Cocco, and Rémi Monasson. 'Gaussian closure scheme in the quasi-linkage equilibrium regime of evolving genome populations'. In: *Europhysics Letters* 132.5 (2021), p. 56001 (cited on pages 27, 29, 42).
- [128] Hong-Li Zeng and Erik Aurell. 'Inferring genetic fitness from genomic data'. In: *Physical Review E* 101.5 (2020), p. 052409 (cited on page 27).
- [129] Fabio Zanini and Richard A Neher. 'FFPopSim: an efficient forward simulation package for the evolution of large populations'. In: *Bioinformatics* 28.24 (2012), pp. 3332–3333 (cited on pages 27, 42).
- [130] Hong-Li Zeng et al. 'Global analysis of more than 50,000 SARS-CoV-2 genomes reveals epistasis between eight viral genes'. In: *Proceedings of the National Academy of Sciences* 117.49 (2020), pp. 31519–31526 (cited on page 27).
- [131] Sewall Wright. 'The roles of mutation, inbreeding, crossbreeding and selection in evolution.' In: *Proceedings of the Sixth International Congress on Genetics*. 1932, pp. 356–366 (cited on page 28).
- [132] Jonathan Kaplan. 'The end of the adaptive landscape metaphor?' In: *Biology & Philosophy* 23 (2008), pp. 625–638 (cited on page 28).
- [133] David M McCandlish. 'Visualizing fitness landscapes'. In: *Evolution* 65.6 (2011), pp. 1544–1558 (cited on page 28).
- [134] Nancy A Moran and Gordon M Bennett. 'The tiniest tiny genomes'. In: *Annual review of microbiology* 68 (2014), pp. 195–215 (cited on page 28).
- [135] David Sherrington and Scott Kirkpatrick. 'Solvable model of a spin-glass'. In: *Physical review letters* 35.26 (1975), p. 1792 (cited on page 29).
- [136] J. Arjan G.M. De Visser and Joachim Krug. 'Empirical fitness landscapes and the predictability of evolution'. In: *Nature Reviews Genetics* 15.7 (2014), pp. 480–490 (cited on page 29).

- [137] Douglas M. Fowler and Stanley Fields. 'Deep mutational scanning: a new style of protein science'. In: *Nature methods* 11.8 (2014), pp. 801–807 (cited on page 29).
- [138] Ville Mustonen and Michael Lässig. 'From fitness landscapes to seascapes: non-equilibrium dynamics of selection and adaptation'. In: *Trends in genetics* 25.3 (2009), pp. 111–119 (cited on page 29).
- [139] Michael Lässig, Ville Mustonen, and Aleksandra M Walczak. 'Predicting evolution'. In: *Nature ecology & evolution* 1.3 (2017), p. 0077 (cited on page 29).
- [140] Richard A. Neher and Aleksandra M. Walczak. 'Progress and open problems in evolutionary dynamics'. In: *arXiv e-prints* (Apr. 2018), arXiv:1804.07720 (cited on page 29).
- [141] Michel Bierlaire. *Optimization: Principles and Algorithms*. 2nd. Lausanne: EPFL Press, 2018 (cited on page 30).
- [142] Richard E Bellman. *Adaptive Control Processes: A Guided Tour*. Princeton Univ. Press, 1961 (cited on page 30).
- [143] Michel Gendreau and Jean-Yves Potvin. *Handbook of metaheuristics*. Vol. 2. Springer, 2010 (cited on page 30).
- [144] James Kennedy and Russell Eberhart. 'Particle swarm optimization'. In: *Proceedings of ICNN'95-international conference on neural networks*. Vol. 4. IEEE. 1995, pp. 1942–1948 (cited on page 30).
- [145] Marco Dorigo, Mauro Birattari, and Thomas Stutzle. 'Ant colony optimization'. In: *IEEE computational intelligence magazine* 1.4 (2006), pp. 28–39 (cited on page 30).
- [146] Scott Kirkpatrick, C Daniel Gelatt Jr, and Mario P Vecchi. 'Optimization by simulated annealing'. In: *science* 220.4598 (1983), pp. 671–680 (cited on page 30).
- [147] John H Holland. 'Genetic algorithms'. In: *Scientific american* 267.1 (1992), pp. 66–73 (cited on pages 30, 31).
- [148] John H. Holland. *Adaptation in Natural and Artificial Systems*. MIT Press, 1992 (cited on page 30).
- [149] Kenneth A. De Jong. 'Genetic algorithms are NOT function optimizers'. In: *Foundations of genetic algorithms*. Vol. 2. Elsevier, 1993, pp. 5–17 (cited on page 30).
- [150] David E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, 1989 (cited on pages 30, 31).
- [151] Colin R. Reeves. 'Genetic algorithms'. In: *Handbook of metaheuristics* (2010), pp. 109–139 (cited on page 31).
- [152] Colin Reeves and Jonathan E Rowe. *Genetic algorithms: principles and perspectives: a guide to GA theory*. Vol. 20. Springer Science & Business Media, 2002 (cited on page 31).
- [153] Daniel J Webre, Peter M Wolanin, and Jeffery B Stock. 'Bacterial chemotaxis'. In: *Current Biology* 13.2 (2003), R47–R49 (cited on page 33).
- [154] Victor Sourjik and Ned S Wingreen. 'Responding to chemical gradients: bacterial chemotaxis'. In: *Current opinion in cell biology* 24.2 (2012), pp. 262–268 (cited on page 33).
- [155] Thomas D Seeley, Scott Camazine, and James Sneyd. 'Collective decision-making in honey bees: how colonies choose among nectar sources'. In: *Behavioral Ecology and Sociobiology* 28 (1991), pp. 277–290 (cited on page 33).
- [156] Fred C Dyer. 'The biology of the dance language'. In: *Annual review of entomology* 47.1 (2002), pp. 917–949 (cited on page 33).
- [157] Christoph Grüter and Walter M Farina. 'The honeybee waggle dance: can we follow the steps?' In: *Trends in ecology & evolution* 24.5 (2009), pp. 242–247 (cited on page 33).
- [158] Chris R Reid et al. 'Slime mold uses an externalized spatial "memory" to navigate in complex environments'. In: *Proceedings of the National Academy of Sciences* 109.43 (2012), pp. 17490–17494 (cited on page 33).
- [159] Barry J Dickson. 'Wired for sex: the neurobiology of *Drosophila* mating decisions'. In: *Science* 322.5903 (2008), pp. 904–909 (cited on page 33).

- [160] Reuven Dukas and Andrew Scott. 'Fruit fly courtship: The female perspective'. In: *Current Zoology* 61.6 (2015), pp. 1008–1014 (cited on page 33).
- [161] Philip Coen et al. 'Dynamic sensory cues shape song structure in *Drosophila*'. In: *Nature* 507.7491 (2014), pp. 233–237 (cited on page 33).
- [162] Tomer J Czaczkes, Christoph Grüter, and Francis LW Ratnieks. 'Trail pheromones: an integrative view of their role in social insect colony organization'. In: *Annual review of entomology* 60 (2015), pp. 581–599 (cited on page 33).
- [163] David W Sims et al. 'Scaling laws of marine predator search behaviour'. In: *Nature* 451.7182 (2008), pp. 1098–1102 (cited on page 33).
- [164] A Glyn Bengough et al. 'Root responses to soil physical conditions; growth dynamics from field to cell'. In: *Journal of experimental botany* 57.2 (2006), pp. 437–447 (cited on page 33).
- [165] Grégoire Altan-Bonnet, Thierry Mora, and Aleksandra M Walczak. 'Quantitative immunology for physicists'. In: *Physics Reports* 849 (2020), pp. 1–83 (cited on page 33).
- [166] Thomas T Hills et al. 'Exploration versus exploitation in space, mind, and society'. In: *Trends in cognitive sciences* 19.1 (2015), pp. 46–54 (cited on page 34).
- [167] Robert C Wilson et al. 'Balancing exploration and exploitation with information and randomization'. In: *Current opinion in behavioral sciences* 38 (2021), pp. 49–56 (cited on page 34).
- [168] Arjun Raj and Alexander Van Oudenaarden. 'Nature, nurture, or chance: stochastic gene expression and its consequences'. In: *Cell* 135.2 (2008), pp. 216–226 (cited on page 34).
- [169] Mads Kaern et al. 'Stochasticity in gene expression: from theories to phenotypes'. In: *Nature Reviews Genetics* 6.6 (2005), pp. 451–464 (cited on page 34).
- [170] Gábor Balázsi, Alexander Van Oudenaarden, and James J Collins. 'Cellular decision making and biological noise: from microbes to mammals'. In: *Cell* 144.6 (2011), pp. 910–925 (cited on page 34).
- [171] A Aldo Faisal, Luc PJ Selen, and Daniel M Wolpert. 'Noise in the nervous system'. In: *Nature reviews neuroscience* 9.4 (2008), pp. 292–303 (cited on page 34).
- [172] Daqing Guo et al. 'Functional importance of noise in neuronal information processing'. In: *Europhysics Letters* 124.5 (2018), p. 50001 (cited on page 34).
- [173] Lev S Tsimring. 'Noise in biology'. In: *Reports on Progress in Physics* 77.2 (2014), p. 026601 (cited on page 34).
- [174] Kunihiro Kaneko. *Life: an introduction to complex systems biology*. Springer, 2006 (cited on page 34).
- [175] Justin Garson. *What biological functions are and why they matter*. Cambridge University Press, 2019 (cited on page 34).
- [176] Thomas Heams. 'Randomness in biology'. In: *Mathematical Structures in Computer Science* 24.3 (2014), e240308 (cited on page 34).
- [177] Bassem A Hassan and P Robin Hiesinger. 'Beyond molecular codes: simple rules to wire complex brains'. In: *Cell* 163.2 (2015), pp. 285–291 (cited on pages 47–49).
- [178] Matthew Cobb. 'A brief history of wires in the brain'. In: *Frontiers in Ecology and Evolution* 9 (2021), p. 760269 (cited on page 48).
- [179] Hermann L. F. Helmholtz. *On the Sensations of Tone as a Physiological Basis for the Theory of Music*. Ed. by Alexander J. Translator Ellis. 3rd ed. Cambridge University Press, 2009 (cited on page 48).
- [180] Peter GH Clarke. 'The limits of brain determinacy'. In: *Proceedings of the Royal Society B: Biological Sciences* 279.1734 (2012), pp. 1665–1674 (cited on page 48).
- [181] Samuel Ward et al. 'Electron microscopical reconstruction of the anterior sensory anatomy of the nematode *Caenorhabditis elegans*'. In: *Journal of Comparative Neurology* 160.3 (1975), pp. 313–337 (cited on page 48).
- [182] Corey S Goodman. 'Isogenic grasshoppers: genetic variability in the morphology of identified neurons'. In: *Journal of Comparative Neurology* 182.4 (1978), pp. 681–705 (cited on page 48).

- [183] John D Steeves and Keir G Pearson. 'Variability in the structure of an identified interneurone in isogenic clones of locusts'. In: *Journal of Experimental Biology* 103.1 (1983), pp. 47–54 (cited on page 48).
- [184] J Eric Schmitt et al. 'Review of twin and family studies on neuroanatomic phenotypes and typical neurodevelopment'. In: *Twin Research and Human Genetics* 10.5 (2007), pp. 683–694 (cited on page 48).
- [185] Caleb Ellicott Finch and Thomas BL Kirkwood. *Chance, development, and aging*. Oxford University Press, USA, 2000 (cited on page 48).
- [186] Eric R Kandel et al. *Principles of neural science*. Vol. 4. McGraw-hill, 2000 (cited on pages 48, 49).
- [187] Alex L Kolodkin and Marc Tessier-Lavigne. 'Mechanisms and molecules of neuronal wiring: a primer'. In: *Cold Spring Harbor perspectives in biology* 3.6 (2011), a001727 (cited on page 49).
- [188] Shaul Yogev and Kang Shen. 'Cellular and molecular mechanisms of synaptic specificity'. In: *Annual review of cell and developmental biology* 30 (2014), pp. 417–437 (cited on page 49).
- [189] Joshua R Sanes and S Lawrence Zipursky. 'Synaptic specificity, recognition molecules, and assembly of neural circuits'. In: *Cell* 181.3 (2020), pp. 536–556 (cited on page 49).
- [190] P Robin Hiesinger and Bassem A Hassan. 'The evolution of variability and robustness in neural development'. In: *Trends in Neurosciences* 41.9 (2018), pp. 577–586 (cited on page 49).
- [191] P Robin Hiesinger. 'Brain wiring with composite instructions'. In: *BioEssays* 43.1 (2021), p. 2000166 (cited on page 49).
- [192] Marion Langen et al. 'Mutual inhibition among postmitotic neurons regulates robustness of brain wiring in *Drosophila*'. In: *Elife* 2 (2013), e00337 (cited on page 49).
- [193] Marion Langen et al. 'The developmental rules of neural superposition in *Drosophila*'. In: *Cell* 162.1 (2015), pp. 120–133 (cited on page 49).
- [194] JM Bekkers and CF Stevens. 'Excitatory and inhibitory autaptic currents in isolated hippocampal neurons maintained in cell culture.' In: *Proceedings of the National Academy of Sciences* 88.17 (1991), pp. 7834–7838 (cited on page 49).
- [195] John G White et al. 'The structure of the nervous system of the nematode *Caenorhabditis elegans*'. In: *Philos Trans R Soc Lond B Biol Sci* 314.1165 (1986), pp. 1–340 (cited on pages 50, 53).
- [196] John G White. 'Getting into the mind of a worm—a personal view'. In: *WormBook: The Online Review of C. elegans Biology [Internet]*. WormBook, 2018 (cited on page 50).
- [197] LaDeana W Hillier et al. 'Genomics in *C. elegans*: so many genes, such a little worm'. In: *Genome research* 15.12 (2005), pp. 1651–1660 (cited on page 50).
- [198] Mario de Bono and Andres Villu Maricq. 'Neuronal substrates of complex behaviors in *C. elegans*'. In: *Annu. Rev. Neurosci.* 28 (2005), pp. 451–501 (cited on page 50).
- [199] Titus Kaletta and Michael O Hengartner. 'Finding function in novel targets: *C. elegans* as a model organism'. In: *Nature reviews Drug discovery* 5.5 (2006), pp. 387–399 (cited on page 50).
- [200] Ann K Corsi, Bruce Wightman, and Martin Chalfie. 'A transparent window into biology: a primer on *Caenorhabditis elegans*'. In: *Genetics* 200.2 (2015), pp. 387–407 (cited on page 50).
- [201] Eric S Haag, David HA Fitch, and Marie Delattre. 'From "the worm" to "the worms" and back again: the evolutionary developmental biology of nematodes'. In: *Genetics* 210.2 (2018), pp. 397–433 (cited on page 50).
- [202] Tosif Ahamed, Antonio C Costa, and Greg J Stephens. 'Capturing the continuous complexity of behaviour in *Caenorhabditis elegans*'. In: *Nature Physics* 17.2 (2021), pp. 275–283 (cited on page 50).
- [203] Sven Dorkenwald et al. 'Neuronal wiring diagram of an adult brain'. In: *bioRxiv* (2023), pp. 2023–06 (cited on pages 50, 69).
- [204] Moritz Helmstaedter et al. 'Connectomic reconstruction of the inner plexiform layer in the mouse retina'. In: *Nature* 500.7461 (2013), pp. 168–174 (cited on pages 50, 69).
- [205] Wei-Chung Allen Lee et al. 'Anatomy and function of an excitatory network in the visual cortex'. In: *Nature* 532.7599 (2016), pp. 370–374 (cited on pages 50, 69).

- [206] Larry F Abbott et al. 'The mind of a mouse'. In: *Cell* 182.6 (2020), pp. 1372–1376 (cited on pages 50, 69).
- [207] David Grant Colburn Hildebrand et al. 'Whole-brain serial-section electron microscopy in larval zebrafish'. In: *Nature* 545.7654 (2017), pp. 345–349 (cited on pages 50, 69).
- [208] Kerriane Ryan, Zhiyuan Lu, and Ian A Meinertzhagen. 'The CNS connectome of a tadpole larva of *Ciona intestinalis* (L.) highlights sidedness in the brain of a chordate sibling'. In: *Elife* 5 (2016), e16962 (cited on pages 50, 69).
- [209] Moritz Helmstaedter, Kevin L Briggman, and Winfried Denk. '3D structural imaging of the brain with photons and electrons'. In: *Current opinion in neurobiology* 18.6 (2008), pp. 633–641 (cited on page 50).
- [210] Jeff W Lichtman, Hanspeter Pfister, and Nir Shavit. 'The big data challenges of connectomics'. In: *Nature neuroscience* 17.11 (2014), pp. 1448–1454 (cited on page 50).
- [211] Ben Mulcahy et al. 'A pipeline for volume electron microscopy of the *Caenorhabditis elegans* nervous system'. In: *Frontiers in neural circuits* 12 (2018), p. 94 (cited on page 50).
- [212] Zeynep F. Altun and David H. Hall. 'Nervous System - General Description'. In: *WormAtlas Hermaphrodite Handbook*. 2005 (cited on page 50).
- [213] Steven J Cook et al. 'Whole-animal connectomes of both *Caenorhabditis elegans* sexes'. In: *Nature* 571.7763 (2019), pp. 63–71 (cited on pages 51, 54, 55).
- [214] Lav R Varshney et al. 'Structural properties of the *Caenorhabditis elegans* neuronal network'. In: *PLoS computational biology* 7.2 (2011), e1001066 (cited on pages 51, 54, 55).
- [215] Emma K Towilson et al. 'The rich club of the *C. elegans* neuronal connectome'. In: *Journal of Neuroscience* 33.15 (2013), pp. 6380–6387 (cited on pages 51, 57).
- [216] Barry Bentley et al. 'The multilayer connectome of *Caenorhabditis elegans*'. In: *PLoS computational biology* 12.12 (2016), e1005283 (cited on pages 51, 69).
- [217] Sreedevi Varier and Marcus Kaiser. 'Neural development features: spatio-temporal development of the *Caenorhabditis elegans* neuronal network'. In: *PLoS computational biology* 7.1 (2011), e1001044 (cited on pages 52, 55, 67, 68).
- [218] Vincenzo Nicosia et al. 'Phase transition in the economically modeled growth of a cellular nervous system'. In: *Proceedings of the National Academy of Sciences* 110.19 (2013), pp. 7880–7885 (cited on pages 52, 55, 58, 67).
- [219] Bradly Alicea. 'The emergent connectome in *Caenorhabditis elegans* embryogenesis'. In: *BioSystems* 173 (2018), pp. 247–255 (cited on page 52).
- [220] Anand Pathak, Nivedita Chatterjee, and Sitabhra Sinha. 'Developmental trajectory of *Caenorhabditis elegans* nervous system governs its structural organization'. In: *PLoS computational biology* 16.1 (2020), e1007602 (cited on pages 52, 57, 67).
- [221] John E Sulston and H Robert Horvitz. 'Post-embryonic cell lineages of the nematode, *Caenorhabditis elegans*'. In: *Developmental biology* 56.1 (1977), pp. 110–156 (cited on pages 52, 53).
- [222] Georgia Rapti. 'A perspective on *C. elegans* neurodevelopment: from early visionaries to a booming neuroscience research'. In: *Journal of Neurogenetics* 34.3-4 (2020), pp. 259–272 (cited on page 52).
- [223] Daniel Witvliet et al. 'Connectomes across development reveal principles of brain maturation'. In: *Nature* 596.7871 (2021), pp. 257–261 (cited on pages 53, 54, 57–59, 65, 67, 68).
- [224] Anna Y Klintsova and William T Greenough. 'Synaptic plasticity in cortical systems'. In: *Current opinion in neurobiology* 9.2 (1999), pp. 203–208 (cited on page 53).
- [225] Hugh G Gauch. *Scientific method in practice*. Cambridge University Press, 2003 (cited on page 55).
- [226] Elliott Sober. *Ockham's razors*. Cambridge University Press, 2015 (cited on page 55).
- [227] Sean L Simpson, Malaak N Moussa, and Paul J Laurienti. 'An exponential random graph modeling approach to creating group-based representative whole-brain connectivity networks'. In: *Neuroimage* 60.2 (2012), pp. 1117–1126 (cited on page 57).

- [228] Aurina Arnatkeviciute et al. 'Hub connectivity, neuronal diversity, and gene expression in the *Caenorhabditis elegans* connectome'. In: *PLoS computational biology* 14.2 (2018), e1005989 (cited on page 57).
- [229] Raj Kumar Pan, Nivedita Chatterjee, and Sitabhra Sinha. 'Mesoscopic organization reveals the constraints governing *Caenorhabditis elegans* nervous system'. In: *PloS one* 5.2 (2010), e9240 (cited on page 57).
- [230] Aharon Azulay, Eyal Itskovits, and Alon Zaslaver. 'The *C. elegans* connectome consists of homogenous circuits with defined functional roles'. In: *PLoS computational biology* 12.9 (2016), e1005021 (cited on page 57).
- [231] Prasanta Chandra Mahalanobis. 'On the generalized distance in statistics'. In: *Proceedings of the National Institute of Sciences (Calcutta)* 2 (1936), pp. 49–55 (cited on page 59).
- [232] Agnan Kessy, Alex Lewin, and Korbinian Strimmer. 'Optimal whitening and decorrelation'. In: *The American Statistician* 72.4 (2018), pp. 309–314 (cited on page 60).
- [233] Freeman Dyson. 'A meeting with Enrico Fermi'. In: *Nature* 427.6972 (2004), pp. 297–297 (cited on page 63).
- [234] Christopher A Brittin et al. 'A multi-scale brain map derived from whole-brain volumetric reconstructions'. In: *Nature* 591.7848 (2021), pp. 105–110 (cited on pages 67, 68).
- [235] Mark W Moyle et al. 'Structural and developmental principles of neuropil assembly in *C. elegans*'. In: *Nature* 591.7848 (2021), pp. 99–104 (cited on page 67).
- [236] Alfonso Pérez-Escudero and Gonzalo G de Polavieja. 'Optimally wired subnetwork determines neuroanatomy of *Caenorhabditis elegans*'. In: *Proceedings of the National Academy of Sciences* 104.43 (2007), pp. 17180–17185 (cited on page 67).
- [237] Alfonso Pérez-Escudero, Marta Rivera-Alba, and Gonzalo G de Polavieja. 'Structure of deviations from optimality in biological systems'. In: *Proceedings of the National Academy of Sciences* 106.48 (2009), pp. 20544–20549 (cited on page 67).
- [238] Beth L Chen, David H Hall, and Dmitri B Chklovskii. 'Wiring optimization can relate neuronal structure and function'. In: *Proceedings of the National Academy of Sciences* 103.12 (2006), pp. 4723–4728 (cited on page 67).
- [239] Oliver Hobert, Robert J Johnston Jr, and Sarah Chang. 'Left–right asymmetry in the nervous system: the *Caenorhabditis elegans* model'. In: *Nature Reviews Neuroscience* 3.8 (2002), pp. 629–640 (cited on page 67).
- [240] Oliver Hobert. 'Behavioral plasticity in *C. elegans*: paradigms, circuits, genes'. In: *Journal of neurobiology* 54.1 (2003), pp. 203–223 (cited on page 68).
- [241] Athanasios Metaxakis, Dionysia Petratou, and Nektarios Tavernarakis. 'Multimodal sensory processing in *Caenorhabditis elegans*'. In: *Open biology* 8.6 (2018), p. 180049 (cited on page 68).
- [242] Jesse M Gray, Joseph J Hill, and Cornelia I Bargmann. 'A circuit for navigation in *Caenorhabditis elegans*'. In: *Proceedings of the National Academy of Sciences* 102.9 (2005), pp. 3184–3191 (cited on page 68).
- [243] Martin Chalfie et al. 'The neural circuit for touch sensitivity in *Caenorhabditis elegans*'. In: *Journal of Neuroscience* 5.4 (1985), pp. 956–964 (cited on page 68).
- [244] Emily R Troemel, Bruce E Kimmel, and Cornelia I Bargmann. 'Reprogramming chemotaxis responses: sensory neurons define olfactory preferences in *C. elegans*'. In: *Cell* 91.2 (1997), pp. 161–169 (cited on page 68).
- [245] H Sebastian Seung. 'Towards functional connectomics'. In: *Nature* 471.7337 (2011), pp. 171–172 (cited on page 69).
- [246] Francesco Randi et al. 'Neural signal propagation atlas of *C. elegans*'. In: *arXiv e-prints*, arXiv:2208.04790 (Aug. 2022), arXiv:2208.04790. doi: [10.48550/arXiv.2208.04790](https://doi.org/10.48550/arXiv.2208.04790) (cited on page 69).
- [247] Ron Milo et al. 'Network motifs: simple building blocks of complex networks'. In: *Science* 298.5594 (2002), pp. 824–827 (cited on page 79).

- [248] Duncan J Watts and Steven H Strogatz. 'Collective dynamics of 'small-world' networks'. In: *nature* 393.6684 (1998), pp. 440–442 (cited on page 80).
- [249] Steven H Strogatz. 'Exploring complex networks'. In: *nature* 410.6825 (2001), pp. 268–276 (cited on page 80).
- [250] Vito Latora and Massimo Marchiori. 'Efficient behavior of small-world networks'. In: *Physical review letters* 87.19 (2001), p. 198701 (cited on page 80).
- [251] Guido Caldarelli. *Scale-Free Networks - Complex Webs in Nature and Technology*. 2007 (cited on page 81).