UNIVERSITY
OF OSLO

Computational Diversity in Neuroscience and Drug Design

Developing Novel Computational Tools Supporting The Study and Treatment of Severe Mental Disorders

Submitted for the Degree of Philosophiae Doctor

Series of dissertations submitted to the
Faculty of Mathematics and Natural Sciences, University of Oslo
No. 1234

ISSN 1234-5678

Print production: Reprosentralen, University of Oslo.

In loving memory of my father, Walter Buts.

We all live our lives dangerously, in a state of jeopardy, at the edge of calamity. You have discovered that the veil that separates your ordered life from disarray is wafer-thin. This is the ordinary truth of existence, from which none of us are exempt. In time we all find out we are not in control. We never were. We never will be.
– Nick Cave

Preface

The Brussels-born, French-Argentinian essayist Julio Cortázar once famously stated that *in quoting others, we cite ourselves*. In my case, this certainly rings true for the quote found at the beginning of this dissertation. The lesson that Nick Cave phrases so eloquently is one I had to learn time and time again during the course of my PhD studies: *We are not in control, we never were, and will never be.* When I set out for a long and cold commute on a dark October morning nearly six years ago, I had no idea of the turbulence my PhD journey would lead me through. On that first day as a PhD candidate, I could not have guessed that I would find myself, less than two years later, handing in my resignation and ending my career as theoretical physicist. Nor could I have guessed that afterwards I would leave Belgium to briefly join a biotech start-up in London and that I would eventually move to Norway to take up a PhD position in computational neuroscience.

My transition into Norway and neuroscience, which was planned for the spring of 2020, was delayed due to the outbreak of the coronavirus pandemic in Europe. While waiting for international travel to resume, I joined the open-science COVID Moonshot project, which aimed to crowdsource novel and patent-free SARS-CoV-2 main protease (Mpro) inhibitors. Several of the molecules I proposed, based on a prototype version of a molecular quality-diversity algorithm, were selected by the project for synthesis and potency tests against Mpro. The algorithm was completed and published, together with a discussion of the concepts behind it, in the first paper of this dissertation, *Illuminating Elite Patches of Chemical Space*. About a year and a half after this first paper, I wrote a single-author, follow-up paper, *Graph-Based Molecular Pareto Optimisation*, in response to questions from the community regarding the importance of chemical diversity in multi-objective molecular optimisation.

After the lifting of travel restrictions at the end of the first wave of the coronavirus pandemic in Europe, around July 2020, I managed to move to Norway. I finally started my PhD studies with Prof. Dr. Gaute Einevoll as part of the University of Oslo Life Science Convergence Environment 4MENT, which studies the molecular mechanisms underlying severe mental disorders such as schizophrenia and bipolar disorder. After many minor and somewhat diffuse contributions to the 4MENT collaboration throughout 2020 and 2021, we settled on focussing our efforts on accelerating complex and large-scale, biophysically-detailed neuron simulations in 2022. Specifically, we directed our attention to distilling the underlying differential equations into easier-to-evaluate artificial neural network models. The final results of this endeavour are currently still being collected and are presented here in a preliminary manuscript, *Multi-task Learning of Biophysically-Detailed Neuron Models*.

Jonas Verhellen
Oslo, April 2023

Abstracts

In English

The study of severe mental disorders, such as schizophrenia and bipolar disorder, is no longer only reserved for doctors or scientists in lab coats. Computational simulations are now commonplace in the (bio)medical sciences and are used to support the development of novel pharmacological interventions or to create detailed simulations of disease models. In this dissertation, we explain the current understanding of the molecular underpinnings of severe mental disorders and describe how computational simulations can play a role in the development of better treatment of these disorders. We specifically chronicle the development of two novel computational tools for the generation of drug-like molecules and one for the accelerated simulation of biophysically-detailed neuron models. The development of these tools relies on recent scientific advancements in the fields of quality-diversity algorithms and multi-task deep learning.

Schizophrenia and bipolar disorder are – in essence – disorders of the mind and hence they are inevitably linked to dysfunction of the brain and the central nervous system. The success of neuroscience in providing insights into the basic neurobiology of perception, cognition, memory, and to some extent emotions, paved the way for the scientific study of the biomolecular mechanisms of severe mental disorders. Unfortunately, for severe mental disorders neither the risk factors that contribute to disease development (known as the disease aetiology) nor the ways in which these risk factors lead to disease (known as the disease pathogenesis) are fully known. However, based on a series of modern and large-scale genetic screenings, dysfunction of glutamatergic neurotransmission, aberrant neuron morphology, faulty ubiquitination signalling and altered calcium ion dynamics can be implicated in these disorders.

In addition to our limited molecular understanding of schizophrenia and bipolar disorder, our current pharmacological treatment options are also limited and only effective for a fraction of the patients. To develop novel medication, the massive space of potentially pharmacologically active molecules has to be searched for the few compounds that have the right protein-interaction profile and physicochemical properties. A whole array of deep learning algorithms for generating drug molecules have been released in the past years. However, none of them managed to truly outperform the much more classical genetic algorithms. In this dissertation, we improved upon a well-known genetic algorithm for molecular optimisation by introducing chemical diversity through quality-diversity techniques borrowed from the field of soft robotics. In a follow-up project, we showed that in multi-objective optimisation, objective space diversity is more important than chemical diversity when trying to discover a Pareto front.

To facilitate better and larger simulations of networks of biophysically-detailed neuron models, recent academic interest has turned to deep learning. By distilling the complicated differential equations governing a biophysically-detailed neuron model into

an easier-to-evaluate artificial neural network, network simulations can be accelerated up to five orders of magnitude. Current distillation approaches only predict neuronal output (membrane potential, outgoing spikes,...) in a limited number of compartments. Ongoing work, presented in this dissertation, aims to explore state-of-the-art multi-task deep learning architectures to expand predictions of the membrane potential to all compartments of a biophysically-detailed neuron model. In addition to providing a stringent test of these architectures – membrane potential values are highly non-gaussian – this approach is also expected to be useful in the context of creating significantly faster local field potential simulations.

In addition to being a series of potentially useful tools for future research regarding the treatment and study of severe mental disorders, the three research projects presented here are connected through their use of computational diversity as a resource. Both research projects in computational drug design explicitly force candidate molecules to be diverse in either a chemical space (single objective optimisation) or in optimisation space (multiple objective optimisation). The research project focussing on the use of multi-task deep learning explores the use of neural architectures that enforce or encourage diversity in their internal representations. Throughout these projects, we explored many different ways to quantify diversity and noted a current lack of any unified framework to measure diversity of computational objects. With academic interest in diversity techniques increasing in many branches of the computational sciences, we look forward to seeing how the study of computational diversity evolves in the near future.

På Norsk

Studiet av alvorlige psykiske lidelsersom schizofreni og bipolar lidelse, er ikke lenger kun forbeholdt leger eller forskere i hvite laboratoriefrakker. Numeriske simuleringer er nå vanlig å ta i bruk i de (bio)medisinske vitenskapene og kan styrke utvikling av nye farmakologiske intervensjoner eller brukes til å lage detaljerte simuleringer av sykdomsmodeller. I denne avhandlingen forklarer vi den nåværende forståelsen av den molekylære bakgrunnen til alvorlige psykiske lidelser og vi forklarer hvordan numeriske simuleringer kan spille en rolle i utviklingen av bedre behandling av disse lidelsene. Vi går spesielt gjennom utviklingen av to nye simuleringsverktøy, et verktøy for syntetisering av medikamentlignende molekyler og et verktøy for akselerert simulering av biofysisk detaljerte nevronmodeller. Utviklingen av disse verktøyene er avhengig av nyere vitenskapelige fremskritt innen kvalitetsmangfold-algoritmer samt dyp læring med flere oppgaver.

Schizofreni og bipolar lidelse er, i hovedsak, sinnslidelser, og derfor er de uunngåelig knyttet til funksjonssvikt i hjernen og sentralnervesystemet. Nevrovitenskapens suksess med å gi innsikt i den grunnleggende nevrobiologien bak persepsjon, kognisjon, hukommelse og, til en viss grad, følelser, har banet vei for den vitenskapelige studien av biomolekylære mekanismer som ligger til grunn for alvorlige psykiske lidelser. For alvorlige psykiske lidelser er dessverre verken risikofaktorene som bidrar til sykdomsutvikling (kjent som etiologi) eller måtene disse faktorene fører til sykdom på (kjent som patogenesen) fullt ut forstått. Basert på moderne, storskala genetisk testing kan det imidlertid vise seg at dysfunksjon i glutamatergisk nevrotransmisjon,

avvikende nevronmorfologi, feil i ubiquitin-signalering og endret kalsiumion-dynamikk er involvert i disse lidelsene.

I tillegg til vår begrensede molekylære fo, rståelse av schizofreni og bipolar lidelse, er våre nåværende farmakologiske behandlingsalternativer også fåtallige og kun effektive for en brøkdel av pasientene. For å utvikle ny medisin, må det i det enorme utvalget av potensielle farmakologisk aktive stoffer søkes etter de få molekylene som har riktig protein-interaksjonsprofil og fysiokjemiske egenskaper. En hel rekke dyplæringsalgoritmer for å syntetisere nye legemidler har blitt utgitt de siste årene. Imidlertid har ingen klart å overgå de mer klassiske genetiske algoritmene. I denne avhandlingen forbedret vi en velkjent genetisk algoritme for molekylær optimalisering ved å introdusere kjemisk mangfold gjennom kvalitets- og mangfoldsteknikker lånt fra feltet myk robotikk. I et oppfølgingsprosjekt viste vi deretter at ved multi-objektiv optimalisering er diversitet i objektiv-rommet viktigere enn kjemisk mangfold når man prøver å oppdage en Pareto-front.

For å legge til rette for bedre og større simuleringer av nettverk av biofysisk detaljerte nevronmodeller, har akademisk interesse nylig vendt seg til dyp læring. Ved å destillere de kompliserte differensialligningene som ligger til grunn for en biofysisk detaljert nevronmodell til et kunstig nevralt nettverk som er lettere å evaluere,kan nettverkssimuleringene akselereres med opptil fem størrelsesordener. Nåværende tilnærminger kan bare forutsi nevralt output(membranpotensial, utgående aksjonspotensial, . . .) i et begrenset antall kompartementer Pågående arbeid, presentert i denne avhandlingen, tar sikte på å utforske toppmoderne fleroppgavers dyplæringsarkitekturer for å utvide prediksjon av membranpotensialet til alle kompartementer i en biofysisk detaljert nevronmodell. I tillegg til å gi en stringent test av disse arkitekturene - membranpotensialverdier er svært ikke-gaussiske - forventes denne tilnærmingen også å være nyttig for å lage betydelig raskere simuleringer av lokalfeltpotensialer.

I tillegg til å være en serie potensielt nyttige verktøy for fremtidig forskning på behandling og i studiet av alvorlige psykiske lidelser, så er de tre forskningsprosjektene som presenteres her knyttet sammen gjennom deres utnyttelse av beregningsmessig mangfold som ressurs. Begge forskningsprosjekter innen simuleringbasert legemiddeldesign tvinger kandidatmolekyler til å være diverse enten i det kjemisk rommet (singel-objektiv optimalisering) eller i et optimaliseringsrom (multi-objektiv optimalisering). Forskningsprosjektet som omhandler fleroppgavers dyplæring utforsker bruken av nevrale arkitekturer som enten tvinger eller oppfordrer til mangfold i deres interne representasjoner. Gjennom disse prosjektene undersøkte vi mange forskjellige måter å kvantifisere mangfold på, og bemerket en mangel på et enhetlig rammeverk for å måle mangfold av beregningsobjekter. Med en økende akademisk interesse for mangfoldsteknikker i mange grener av beregningsvitenskap, ser vi frem til å følge med på hvordan studiet av beregningsmessig mangfold utvikler seg i nær fremtid.

Translated by the Asbestkontoret hive mind.

Acknowledgements

During the years of my PhD studies, the necessity of friendship, kindness and having a sense of belonging became increasingly clear. I am filled with feelings of gratitude and admiration for the diverse group of friends and family that rallied around me and inspired me to go on, time and time again. While some friendships and family connections turned out be incredibly robust and nourishing, inevitably, others fell apart at the first signs of adversity or crumbled under the strain of distance and time. I am deeply saddened by (and in some cases thoroughly disappointed in) those friends and family members who, one way or another, forced me to continue on without them. I will not forget the role each of them played in the making of this dissertation, for good or for bad. I hope, that in time, I can look back to the period of my PhD studies as a positive and important formative experience. Difficult times come with important lessons and I want to thank everyone who was kind enough to help me figure those out.

In this category, first and foremost, my gratitude goes out to the collection of therapists who collectively helped me through difficult transitions, periods of grief and moments of doubt. Without them it is unlikely I would have ever returned to academia. Whenever professional psychotherapy was unavailable or insufficient, there luckily always was the invaluable support and the often very wise advice of my friends and family. I am deeply grateful to all of you! This dissertation is dedicated to my dad, who passed away a few weeks after I ended my first period of PhD studies. He prepared me for life in an idiosyncratic, inimitable and simply formidable fashion and not a single day passes by that I do not miss him. He lived a difficult life, but showed everyone that *hope springs eternal*, even amid adversity. Especially amid adversity. Dad, I wish you could have held on a bit longer, but I am grateful for every day we had.

In those difficult months between my first and second period of PhD studies, a small group of people went above and beyond to take care of me. Bruno Buts and Hilde Vanhaecke opened their doors for me and gave me a space to recover. Jeroen Buts and Isis Brun cheered me up in the early and disorienting weeks after my resignation and when I moved to London, Sarah Buts and Mario Berta provided me with a feeling of kinship in a large (and sometimes strange) city. In the months of the first European lockdown, Greta Van Der Kelen and Marc Buts very generously provided me with an apartment to stay and work in. My home town friends Arno Van Camp, Thierry Demol, and Evelyne Demol have been by my side throughout the high and lows of the past two decades and somehow each of them showed up, individually, at a crucial moment in my PhD studies to give me much needed encouragement. I will never forget. Thank you all!

My physics friends systematically supported and inspired me. I am so lucky to have met each and everyone of you, whether it was during my studies at the *Vrije Universiteit Brussel* or during my first period of PhD studies, which took place at the *Universiteit Gent*. Paul Coppin and Benoît Tuybens showed me the power of everyday kindness and the importance of consistent altruism. Through acts great and small,

including proofreading large parts of this dissertation, both of you have given me more, much more, than I will ever have the chance to reciprocate. You have also taught me that sometimes it is okay to receive more than you give. Timothy De Deyn, Pablo Correa Camiroaga, and Maurits van Roozendaal proofread the remaining parts of this dissertation. In the decade that we have known each other, all three of you of have lifted my spirits whenever I needed it and provided me with so much humour and companionship. I really enjoyed the many phone calls we have had over the years and cherish the few recent times we managed to meet up in person.

In the days before I left for London, Pieter Verhulst and Charlotte Verwimp invited me to a house warming party where they, together with Daphné Janssens and many of the friends mentioned above, provided me with a Brexit survival package. It was a wonderful and funny gesture and I often look back at our group picture from that evening with very fond feelings. Even though I have seen Jonathan De Ro less than I would have wished since leaving Belgium, I have experienced each and every time as an especially joyous occasion. Jonathan, I hope you are doing well! Twitter-based interactions with Stijn De Baerdemacker and Pieter Claeys were welcome reminders of the few but very real positive outcomes of my first period as a PhD student. I wish both of you many and very fulfilling years of leading vibrant research groups. Last but definitely not least, I have to express my gratitude to Ben Craps, who – to this day – is still, by far and away, the best person I have ever met in academia.

Sarah Kriebel Thon, Elise Reppe Olsen, Atle Eskeland Rimehaug and Maria Perona Fjeldstad made *Asbestkontoret* a significantly less toxic place to reside. I thank you all sincerely for the laughs and for all the times you allowed me to vent my frustrations. Sverre Grødem, Elise Holter Thompson, Gry Glenne, Torbjørn Vefferstad Ness, and Bente Iren Larsen gave me many useful insights into Norwegian society and made my time in Oslo a lot more fun. The same goes for Ali Telet, Sandra Stølen Bryne, Tina Kleven, Malin Benum Røe and Ivan Myhre Winje, without any of whom the *Centre for Integrative Neuroplasticity* (CINPLA) would be much less pleasant to hang out. Special recognition should go to Jeriek Van den Abeele and Kosio Beshkov. I feel lucky for having had the opportunity to work and write papers with both of you. In that same vein, I want to thank Sebastian Amundsen and Maria Lunde for choosing to complete their master thesis with me.

Twice during my PhD studies, once during each period, I was fortunate enough to join a summer school in a magnificent setting: the 2018 *Les Houches Summer School* on *Integrability in Atomic and Condensed Matter Physics* in the French alps and the 2021 *Summer School on Mathematical Methods in Computational Neuroscience* at the *Fred Kavli Science Center* in Eresfjord. The former was an indescribable, once-in-a-lifetime experience while the latter showed me a beautiful part of Norway and introduced me to some lovely people. Most noteably, I count my self lucky to have met Ula Włodkowska, Michael Tarlton, Abbas Karimi Rizi and Nicolai Waniek. All of you are uniquely passionate and interesting (neuro)scientists. Two other inspirational scientist whom I certainly have to mention and thank are Geoffrey Janssens and Klas Pettersen. Both of you systematically encouraged me to make the most of my capabilities and I am extremely grateful for that!

No acknowledgement section could be considered complete without the explicit mention of academic advisors and the funding sources that made the research discussed

in this dissertation possible. I want to thank Gaute Einevoll for being my main supervisor and for giving a PhD drop-out a second chance. I also want to thank Marianne Hafting Fyhn for accepting to be my secondary supervisor and for organising so many practical things in and around CINPLA. For legal and administrative reasons, it is important to note that my PhD fellowship was financed by the *4MENT Life Science Convergence Environment* at the *University of Oslo* and that all images in chapters one and two of this dissertation were created with *BioRender.com* through an academic license. Finally, I want to note – with gratitude – that large parts of this dissertation were written at the *Deichman Bjørvika*, the *Oslo Kaffebar* in Berlin, the law library at *Domus Juridica*, and in Paul Coppin's home office.

Contents

Preface **xi**

Abstracts **xiii**

Acknowledgements **xvii**

Contents **xxi**

List of Papers **xxiii**

List of Figures **xxv**

1 Introduction and Motivation **1**
Severe Mental Disorders: Schizophrenia and Bipolar Disorder 1
Difficulties in Drug Design: Cost, Attrition, and Complexity 13
Opportunities in Neuroscience: Accelerating Disease Models 21
References 26

2 Concepts and Methodology **53**
Quality-Diversity Methods in De Novo Drug Design 53
Distilling Biophysically-Detailed Neuron Models 59
References 64

3 Conclusion and Outlook **67**
Summary and Discussion 67
Contributions Made During the PhD 72
Concluding Remarks 75
References 75

Papers **80**

I Illuminating Elite Patches of Chemical Space **81**

II Graph-Based Molecular Pareto Optimisation **97**

III Multi-Task Learning of Biophysically-Detailed Neuron Models **117**

List of Papers

Included Papers

1. Verhellen, J. and Van den Abeele, J. “Illuminating Elite Patches of Chemical Space”. Published in *Chemical Science*, September 2020, Volume 11, Issue 42, pp. 11485–11491. DOI: 10.1039/D0SC03544K.

2. Verhellen, J. “Graph-Based Molecular Pareto Optimisation”. Published in *Chemical Science*, June 2022, Volume 13, Issue 25, pp. 7526–7535. DOI: 10.1039/D2SC00821A.

3. Verhellen, J. , Beshkov, K. , Amundsen, S. , Ness T. and Einevoll, G. “Multi-Task Learning of Biophysically-Detailed Neuron Models”. *Manuscript In Progress.*

Additional Papers

1. Beshkov, K. , Verhellen, J. and Lepperød, M. “Isometric Representations in Neural Networks Improve Robustness”. *Manuscript In Process of Resubmission.*

2. Verhellen, J. , Sacher, J. , MacKinnon, S. and Windemuth, A. “Morta: A Multi-Dimensional and Representative Archive of Molecules”. *Manuscript In Progress.*

List of Figures

1.1 Severe Mental Disorders Affect The Entire Body 2
1.2 Typical Onset and Progression of Schizophrenia 4
1.3 Frequency Spectrum of Genetic Schizophrenia Risk Factors 6
1.4 Four Aspects of Schizophrenia Pathophysiology 11
1.5 The Drug Development Pipeline 14
1.6 Evolution of Antipsychotic Compounds Through the Decades 16
1.7 Deep Generative Models for De Novo Drug Design 20
1.8 Basic Morphology and Properties of a Neuron 22
1.9 Neuroscience Platforms To Study Severe Mental Disorders 25

2.1 Graph-Based Genetic and Illumination Algorithms 54
2.2 Splitting Front Procedures in NSGA-II and NSGA-III 58
2.3 Hard Parameter Sharing for Prediction of Membrane Potentials 61
2.4 Soft Parameter Sharing for Prediction of Membrane Potentials 63

I.1 Pseudocode of MAP-Elites for Molecular Optimisation 84
I.2 Illumination of Elite Solutions for Rediscovery Benchmarks 84
I.3 Molecular Distributions for Rediscovery Benchmarks 88
I.4 Molecular Distributions for Median Molecule Benchmarks 89
I.5 Distribution of Function Evaluations and Cumulative Success Rate . . 91

II.1 Visualisation of a Pareto Front and Dominated Solutions 98
II.2 Examples of Mutations and Crossovers as Generated by GB-EPI 100
II.3 Pseudocode of NSGA for Molecular Optimisation 101
II.4 Splitting Front Procedure of NSGA-II and NSGA-III 102
II.5 Dominated Hypervolume, Geometric Mean, and Internal Similarity . . 106

III.1 Training and Validation Loss of the Distillation Process 120
III.2 Expert Diversity Throughout the Distillation Process 122
III.3 Schematic Representation of MTL Architectures 126

Chapter 1

Introduction and Motivation

Each of our minds is the complex biological outcome of every sensation that has ever run through our bodies, the product of all the ways we have ever been loved or disappointed, the aggregate of each time we have been hurt or celebrated. To adapt to the changing circumstances and various experiences of life, our brains reorganise themselves around underlying genetic and epigentic scaffolds. In health, this extraordinary mental flexibility – a hallmark of human resiliency – provides us with a wide playing field, but in disease we find ourselves traversing a narrow path: precariously balancing our own needs with those induced by a pathogen, dangerous genetic mutation or misplaced methyl tag. If that path becomes unbearably narrow and our internal balance impossible to maintain, we stray beyond the regular human experience and into the poorly understood realms of psychiatric disorder.

Depression, anxiety, and compulsive-obsessive behaviour have all been observed in the animal kingdom, but the severest mental disorders, such as schizophrenia and bipolar disorder, remain exclusive to humanity. These disorders have a unique potential for destruction, as they often affect personal relationships and social activities as well as the general health of those suffering from these diseases. Yet, at some recent point in humanity's 3.5 billion year evolutionary history, a predisposition for these disorders was instilled into our genes. Understanding the molecular origins of mental disorders is therefore not only of importance for the development of highly necessary novel clinical interventions, but it is also a unique opportunity to unravel a part of our shared and complex evolutionary backstory.

In the first section of this introductionary chapter, we provide the reader with an overview of the pathophysiology of schizophrenia and bipolar disorder and introduce them to the intricate and poorly understood network of molecular mechanisms involved. Having set the stage, we highlight the importance and difficulty of pharmacological intervention in severe mental disorders and then turn our attention to the specific scientific challenges addressed in this dissertation. We discuss the impact of computer simulations on drug design and the use of biophysically-detailed neuron simulations, and showcase the importance of both of these topics in the context of studying and treating severe mental disorders. Combined, these discussions provide the necessary context and motivation for the novel computational tools presented in the next chapter.

Severe Mental Disorders: Schizophrenia and Bipolar Disorder

Mental disorders are characterised [1] by a combination of significant disturbances in a person's behaviour, cognitive and daily functioning, or emotional regulation. Effective prevention and treatment options for some of the milder mental disorders exist [2, 3] but patient's access to effective and affordable care is unfortunately often limited [4]. For those patients that suffer from severe mental disorders, including schizophrenia and

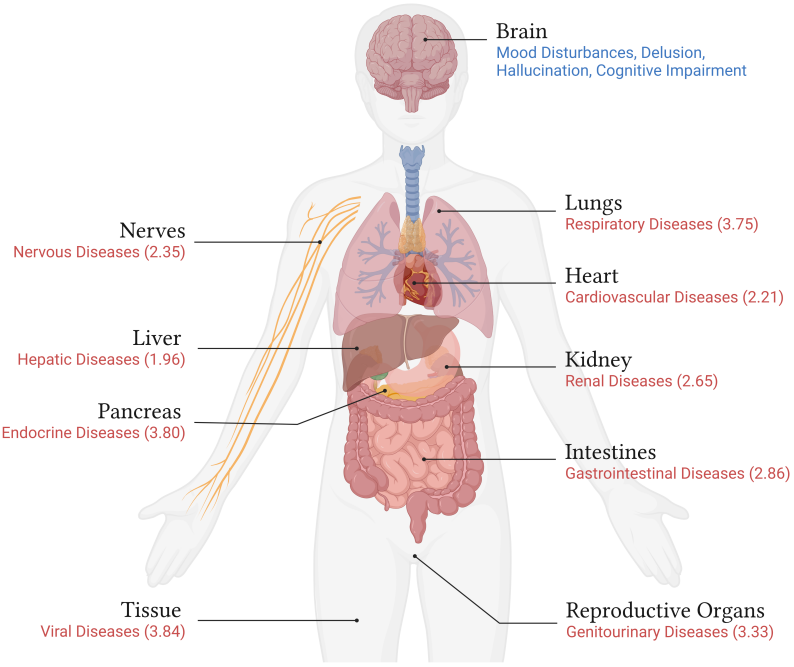

Figure 1.1: **Severe Mental Disorders Affect The Entire Body.** Severe mental disorders are characterised by their detrimental effect on brain functioning (shown in blue), but patients suffering from these diseases are also disproportionately affected by a range of other diseases (shown in red). For each cluster of diseases comorbid with schizophrenia or bipolar disorder, we highlight a representative organ and provide the 95% confidence interval mortality risk ratio [8] in incident and prevalent schizophrenia cases versus the general population.

bipolar disorder, effective curative or prophylactic treatment is generally non-existent or deeply insufficient [5]. Despite decades of research, people suffering from severe mental health disorders face an average 10 to 15 year reduction in life expectancy [6–8], in part due to current medical limitations and the severity of symptoms, see Figure 1.1. At the same time, stigma, discrimination, and human rights violations targeting patients with mental disorders, severe or otherwise, remain painfully common [9].

Schizophrenia [10] and bipolar disorder [11] manifest as a heterogeneous combination of positive (hallucinations, delusions, mania), negative (depression, anhedonia), and cognitive symptoms. Schizophrenia is mostly a psychotic disorder, characterised by hallucinations and psychosis, while bipolar disorder is primary diagnosed as a mood disorder, characterised by a recurring pattern of sequential periods of mania and depression. Combinations of psychosis and fluctuating mood disturbances in the same patient are grouped together under the diagnostic umbrella term schizoaffective disorder [12].

Unfortunately, due to the complexity of these diseases, patients diagnosed with the same mental disorder may not share any clinical symptoms [13] and overlapping symptoms may originate from different disorders [14], yet it remains clinically practical to make these diagnostic distinctions.

Whereas there is clear evidence that external factors play an important role in disease etiology, schizophrenia and bipolar disorder have been shown to have a heritability of around 70% [15]. This high level of heritability indicates that genetic variation creates an important biological foundation upon which mental disorders can develop. Modern genetic studies, making use of large patient cohorts and high-throughput screening techniques, have uncovered an extensive and partially overlapping genetic basis for schizophrenia and bipolar disorder [16], in line with their strong clinical overlap. Despite the impact on life expectancy and quality of life, severe mental disorders have managed to spread widely throughout the general population. The worldwide mean lifetime prevalence of schizophrenia and bipolar disorder has been estimated to be around 1% and 2%, respectively [17, 18], although large regional differences in prevalence rates have been noted.

Schizophrenia and Bipolar Disorder as Developmental Disorders

In contrast to most other neurological diseases, schizophrenia and bipolar disorder typically arise in early adolescence [19], and are presumed to be triggered by aberrations in brain development and maturation [20]. These aberrations are thought to be due to an interplay of genetic, epigenetic, environmental, and developmental factors, see Figure 1.2. A typical case of schizophrenia starts with a latent phase in childhood when sub-threshold psychotic symptoms and developmental delays might become apparent but often these go unnoticed [21]. The initial phase of schizophrenia is followed by a first episode of psychosis, typically in adolescence or young adulthood [22]. About 10% to 15% of patients recover after this first episode, a similar proportion of patients suffer from an unremitting and treatment-resistant form of schizophrenia, and the remainder develops a chronic and fluctuating relapse/remission disease pattern [23].

Each of the distinct phases in the progression of schizophrenia can be matched to a period in brain development [22, 23], underscoring the importance of neurodevelopment in psychiatric disorders. For instance, the initial, latent stage of schizophrenia mostly occurs during the years of early brain formation. Whereas the first episode of psychosis typically happens right after or in the midst of a period of intense brain reorganisation [24] and sexual maturation marked by adolescence, the chronic aspects of schizophrenia happen throughout life's later decades of brain maintenance. Brain formation and reorganisation are the result of an interplay of neurogenesis [25], synaptogenesis [26], gliogenesis [27], synaptic pruning [28], dendritic arborisation [29], cortical myelination [30], neuronal proliferation [31] and neuronal differentiation [32]. Dysfunction of any or all of these complex biological mechanisms is a potential suspect in the pathophysiology of schizophrenia and a plausible target for clinical intervention.

Early indications of the importance of neurodevelopment in schizophrenia arose from various epidemiological studies [33–35], linking prenatal adverse events such as maternal infection [36], malnutrition [37] or smoking [38] to later clinical manifestations of schizophrenia in offspring. Follow-up epidemiological studies also

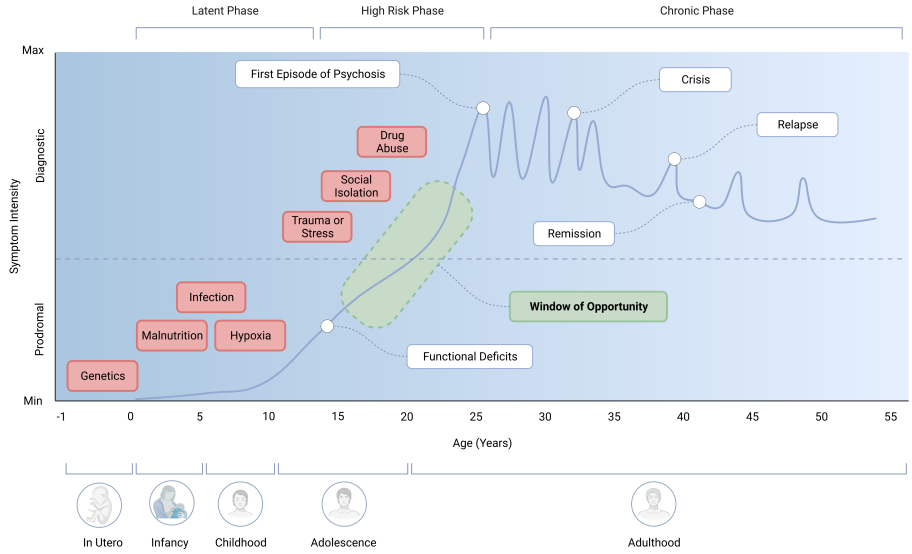

Figure 1.2: **Typical Onset and Progression of Schizophrenia.** Symptom intensity of schizophrenia (blue line), risk factors (red boxes), and disease progression milestones (white boxes) in function of age. A window of opportunity for altering disease course before the first episode of psychosis (shown in green) has gained significant academic interest. Figure based on Figures 1.a and 1.b in *Altering the Course of Schizophrenia: Progress and Perspectives* [23].

discovered statistical correlations to limited delays in the achievement of developmental milestones [39–41] regarding motor, speech, and cognitive function. The exact nature of the relationship between cognitive deficits and schizophrenia is a source of debate [42–44], and it remains unclear how minor developmental deficits in early life can transform to a severe psychotic illness later in life. The period of brain reorganisation during which the first episode of psychosis occurs is suspected to be involved in this transition. Although less strongly supported, viral infections in childhood have also been statistically linked to an increased risk of developing schizophrenia [45, 46].

A recent study of gene regulation in brain formation based on three-dimensional chromatin interaction profiles derived from human brain tissue, compared regulatory patterns for nine brain disorders: five psychiatric disorders and four neurodegenerative disorders [47]. This study indicates that genes involved in psychiatric disorders are preferentially expressed prenatally, and in neurons, in stark contrast to neurodegenerative disorders which show increasing genetic expression with age and affect a more diverse range of brain cell types. In addition, post-mortem DNA methylation analyses of fetal, post-natal, and adult human brain samples indicate that a significant fraction of the genes implicated in schizophrenia show preferential expression during fetal development [48,

49], further underscoring the importance of early development in schizophrenia.

Neurodevelopmental disturbances are less consistently implicated in bipolar disorders [50]. This difference is most likely due to the comparatively strong patient heterogeneity in bipolar disorder compared to schizophrenia [51, 52]. A substantial subset of patients with bipolar disorder exhibits sub-threshold manifestations of mood disorders before full disease onset, similar to the progression seen in schizophrenia patients, but it remains unclear if this observation extends to the entire patient population [53]. Generalisations from one subgroup of patients with bipolar disorder to another subgroup are notoriously difficult. Specifically, the subsets of early-onset patients or those diagnosed with schizoaffective disorder display clinical features and treatment outcomes that differ significantly from other diagnostic subgroups [54–56].

Given that the underlying neurodevelopmental aspects of schizophrenia and bipolar disorder typically start many years before diagnosis, identification of at-risk individuals in or before the latent stages of these diseases is increasingly seen as clinically important [57–59]. Large personal and societal costs could potentially be avoided through early detection and intervention. In the most ideal case, transition into psychosis could be halted by tailor-made clinical, pharmacological and psychosocial interventions [23]. To achieve this (or any other satisfactory) level of alteration in disease progression, better patient identification and stratification based on advanced multi-modal diagnostics needs to be combined with novel pharmacological options, all of which relies on a more precise molecular understanding of these disorders.

The Genetic Architecture of Schizophrenia and Bipolar Disorder

Determining the genetic architecture of any polygenic trait, such as schizophrenia and bipolar disorder, involves the testing of allele frequency differences in hundred of thousands of genetic variants across the human genome, in a process known as genome-wide association studies (GWASs) [60]. Whereas GWASs can be used to track down copy-number-variants (CNVs) or other sequence variations, they are most often employed to study common and small-scale genetic variations, known as single-nucleotide polymorphisms (SNPs), that affect a single base pair in a gene. The statistical association between genetic variants and diseases can be used for a range of downstream applications. For instance, the implication of pro-inflammatory pathways in a Crohn's disease GWAS have been used to support and justify clinical trials of drugs targeting these pathways [61, 62].

Because sequencing the whole genome of an individual is costly, only the most informative SNPs are considered in GWASs. Representative and pre-selected SNPs are used to discover genetic associations with a trait or disease [64]. Determining which genetic variants are causal to a trait remains challenging. Out of the more than 3500 GWASs that have been conducted so far, most have indicated that polygenic traits are influenced by many thousands of variants, each of them individually conferring little risk [65]. Similarly for schizophrenia and bipolar disorder: up to now 287 common risk loci have been identified for schizophrenia [66], and 64 for bipolar disorder [67]. To complicate matters, common genetic variants found in these GWASs only account for approximately 20% and 24% of the disease liability of bipolar disorder and schizophrenia, respectively.

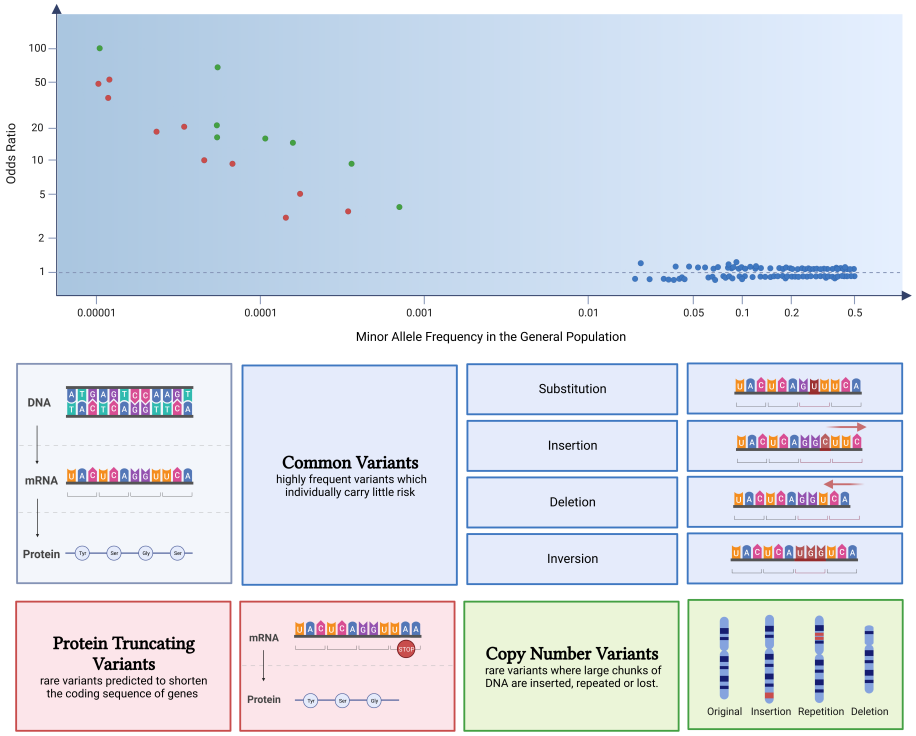

Figure 1.3: **Frequency Spectrum of Genetic Schizophrenia Risk Factors.** Significant genetic associations for schizophrenia from recent large scale genetic studies are shown (top), together with their impact on genetic material (bottom): common variants (blue), protein truncating variants (red), and copy number variations (green). Data points were manually reproduced from Figure 6.a in *Rare Coding Variants in Ten Genes Confer Substantial Risk for Schizophrenia* [63] for illustrative purposes.

Common variants associated with schizophrenia show a preferential expression in neurons and seem to implicate multiple pathways that target synaptic organisation and function. Indeed, several of the genes with common variants implicated in schizophrenia encode important neurotransmitter receptors or ion channels. These genes include voltage-gated calcium and chloride channels (*CACNA1C* and *CLCN3*) [68], glutamatergic (*GRM1*, *GRM3*, *GRIA1*, and *GRIN2A*) [69] and dopaminergic receptors (*DRD2*) [70] as well as the serine racemase enzyme (*SRR*) [71] which plays a role in the biosynthesis of an allosteric N-methyl-D-aspartate (NMDA) receptor ligand. Other GWAS implicated genes for schizophrenia include a gene that encodes a sarcoplasmic/endoplasmic reticulum calcium pump (*ATP2A2*) [72, 73], and genes involved in synaptic organisation and differentiation (*DLGAP2*, *LRRC4B*, *GPM6A*, and *PAK6*) [74].

The most significantly associated common variant for schizophrenia resides in the *C4A* gene [75, 76]. *C4A* is a component of the major histocompatibility complex (MHC) [77] which is an essential part of the adaptive immune system. Hence *C4A* has caused significant interest in the role of inflammation and the immune system in severe mental disorders [78–80]. Among the 64 genes implicated in bipolar disorder, seventeen are also implicated in schizophrenia. Surprisingly, unlike in schizophrenia, the MHC has been identified as a genome-wide significant risk factor for bipolar disorder through a common missense variant in the immune checkpoint gene *BTN2A1* [81] but not through the influence of the *C4A* gene. Other common variants associated with bipolar disorder are voltage-gated calcium and potassium channel encoding genes (*CACNB2* and *KCNB1*) [68] and a serotonin receptor gene (*HTR6*) [82].

Due to the limited risk contribution of common variants to schizophrenia and bipolar disorder, a substantial genetic contribution to disease risk might reasonably be expected to be found in other parts of the allele frequency spectrum [83], see Figure 1.3. Large-effect risk variants with low frequency in the population, often called rare variants, are suspected to play an important role. Several rare CNVs have been robustly associated with schizophrenia [84], most remarkably the 22q11.2 [85] and the 2p16.3 [86] deletions. The former is known to cause dramatically higher rates of schizophrenia, making it the single most important genetic risk factor for schizophrenia currently known, while the latter is notable because it only contains a single gene: the *NRXN1* gene which encodes the neurexin 1 protein. Neurexin 1 is abundant in GABAergic synapses, where it controls synaptic formation and transmitter release. CNVs associated with schizophrenia have a high pleiotropy [87], indicating a relationship with other disorders. Remarkably, bipolar disorder associated CNVs only seem to contribute to schizoaffective cases [88].

Analysis of rare coding variants in severe mental disorders is a powerful complementary approach to the study of common variants and CNVs. Exome studies [89], which only consider variations in the protein-coding regions of the genome, and proband studies [90], which search for de novo mutations in affected family members compared to their healthy ancestors or siblings, are typically conducted with higher sensitivity than GWASs. The Schizophrenia Exome Sequencing Meta-Analysis (SCHEMA) consortium [63] has aggregated and harmonised the exome sequences of nearly 25.000 schizophrenia patients and close to 100.000 healthy controls. This global meta-analysis effort has led to the identification of ten genes with rare, protein-truncating variants that carry a substantial risk with respect to schizophrenia (*SETDA1*, *CUL1*, *XPO7*, *TRIO*, *CACNA1G*, *SP4*, *GRIA3*, *GRIN2A*, *HERC1*, and *RB1CC1*).

The BipEx collaboration [91] performed a whole-exome meta-analysis of bipolar disorder, in a study comprising nearly 14.000 patients and a similar number of healthy controls. Despite the sizeable patient cohort, the BipEx study did not observe a single risk gene with exome-wide significance for bipolar disorder. However, several risk genes identified in the SCHEMA consortium analysis for schizophrenia were found to be significantly enriched for rare variants associated with bipolar disorder. Combining the data from the SCHEMA and BipEx studies revealed strong evidence that haploinsufficiency in the *AKAP11* gene confers risk for the development of severe mental disorders. The AKAP-11 protein [92] is suspected to interact with the GSK3B protein, which is assumed to be the main target of lithium therapy [93–95]. The overlap in risk for schizophrenia and bipolar disorder is hence clearly not limited to common

variants, further corroborating the clinical resemblances seen between the two disorders.

Molecular Mechanisms Underlying Schizophrenia and Bipolar Disorder

Uncovering the underlying biological mechanisms of the genetic variants associated with severe mental disorders, most of which only have a small effect, poses a serious challenge [96]. Despite the success in risk gene discovery through GWASs, the implicated common variants are often non-coding [97] or act through a regulatory function [98], making their interpretation difficult. CNVs on the other hand, each affect hundreds of kilobases of genome, making it hard to specify the mechanism by which they confer disease risk [99]. Most rare variants are hard to interpret correctly with the exception of rare coding or protein-truncating variants. Protein-truncating variants have been linked to loss-of-function for the affected protein, although in some limited cases gain-of-function has also been observed [100]. Because interpreting results from large-scale genetic screening is so challenging, there is no general consensus on the molecular mechanisms underlying severe mental disorders.

To create some sense of order in the chaos, we present an overview of a select amount of genetic sectors consistently implicated in schizophrenia through independent screening techniques, see Figure 1.4. In addition to highlighting several of the genes involved and discussing the presumed effects of mutations in these genes, we also discuss pre-clinical and clinical support for the presented sectors. Remarkably, as discussed later, dopamine dysregulation is not one of the sectors that is systematically implied in genetic screenings. Whereas this overview is not (and cannot possibly be) an exhaustive list of mechanisms involved in severe mental disorders, it aims to provide a solid and conservative foundation for further scientific discussions on the nature of these diseases. In addition, this overview could be used as a list of possible targets for pharmacological interventions. As a starting point for this overview, we used the ten genes with protein-truncating variants implicated in schizophrenia by the SCHEMA consortium.

Glutamatergic Neurotransmission

Excessive release of glutamate – the primary excitatory neurotransmitter of the human brain – is hypothesised to be at least partially responsible for psychotic symptoms and cognitive impairment in schizophrenia [101, 102]. Clinical administration of antagonists which bind to the phencyclidine pocket of NMDA receptors, an important type of glutamate receptors, have been shown to induce positive, negative, and cognitive symptoms of schizophrenia in healthy subjects [103, 104]. Crucially, the behavioural effects of these NMDA receptor antagonists persist even in the absence of aberrant dopamine activity [105]. Together, these results provide strong support for disruption of glutamatergic neurotransmission as a fundamental component of schizophrenia pathophysiology.

The different genetic screening efforts discussed previously, have implicated components of glutamate receptors, such as *GRIA3* and *GRIN2A*, as well as

proteins indirectly involved in facilitating glutamatergic neurotransmission. For instance, the transcription factor *SP4* gene has been associated with schizophrenia through the SCHEMA meta-analysis, and experimentally observed to cause a dramatically decreased expression of *GRIN1*, a glutamate receptor component, in mice [106]. Similarly, *SETD1A* knockout in post-synaptic neurons has been shown to reduce glutamate release probability [107] and in-vitro studies of rodent neurons have shown that *TRIO* affects glutamatergic neurotransmission [108]. In addition, the NMDA receptor modulator enzyme serine racemase has been associated with schizophrenia susceptibility in humans and mice. [71].

Relevant Genes

GRIA3 (Glutamate Ionotropic Receptor AMPA Type Subunit 3), *GRIN2A* (Glutamate Ionotropic Receptor NMDA Type Subunit 2A), *SP4* (SP4 Transcription Factor), *GRIN1* (Glutamate Ionotropic Receptor NMDA Subunit zeta-1), *TRIO* (Trio Rho Guanine Nucleotide Exchange Factor), *SETD1A* (SET Domain Containing 1A, Histone Lysine Methyltransferase), *SRR* (Serine Racemase).

Neuron Morphology

As mentioned before, the immune system has long been suspected as a crucial player in severe mental disorders due to the presence of an important common variant in complement component gene *C4A* associated with schizophrenia. However, modern meta-analysis applied to existing large-scale genetic and transcriptomic datasets found that genes negatively co-expressed with *C4A* and associated with neuronal and synaptic pathways involved in synaptic pruning, exhibit a strong and specific enrichment for schizophrenia. The genes of the complement system, on the other hand, did not [109]. The previously discussed characteristic neurodevelopmental patterns of schizophrenia and bipolar disorder provide further clues to support the hypothesis that excessive synaptic pruning is involved in severe mental disorders.

Aberrant neuronal morphology is implicated in schizophrenia through versatile genes such as *SETD1A*, which is known to affect axonal branching [107] and *TRIO*, which is known to regulate neuronal migration, axonogenesis, axon guidance, and synaptogenesis through actin cytoskeleton remodelling [110–112]. These two implicated genes indicate that a range of possible variations in aberrant neuronal morphology could play a role in severe mental disorders. Finally, it is important to mention that staining techniques applied to post-mortem human brain tissue have indicated an association between a loss of dendritic spines density and the presence of schizophrenia [113, 114].

Relevant Genes

TRIO (Trio Rho Guanine Nucleotide Exchange Factor), *SETD1A* (SET Domain Containing 1A, Histone Lysine Methyltransferase), *C4A* (Complement Component 4A).

Ubiquitin Pathways

Genomic screenings efforts have associated common and rare mutations in several core components of ubiquitin ligases such as *CUL1*, *PJA1*, *CUL9*, and *HERC1*

with schizophrenia. Ubiquitin [115] is a small regulatory protein (76 amino acids) found almost everywhere in eukaryotic organisms, hence its name, and modification of proteins by ubiquitin is mediated by a set of ubiquitin ligases [116] and deubiquitinating enzymes [117]. The attachment of ubiquitin to a target protein is called ubiquitination. Ubiquitin ligases can also target histones and transcription factors, allowing them to play an important role in orchestrating gene expression, specifically in early mammalian development [118].

Ubiquitination is typically used by a cell to mark dysfunctional or misfolded proteins for degradation via the ubiquitin proteasome system (UPS) [119]. Dysregulation of the UPS has been linked to schizophrenia in ubiquitinated proteins in brain and blood samples of patients [120]. The UPS plays a crucial role in many basic cellular processes, including synaptic efficacy and cytokine production [121, 122]. This diversity of functions of ubiquitin and its proteasome system make interpreting their effect on schizophrenia and bipolar disorder challenging. Whereas there is mounting evidence that ubiquitin ligases are involved in severe mental disorders, their precise function or rather dysfunction also remains unclear.

Relevant Genes

CUL1 (E3 ubiquitin-protein ligase Cullin 1), *PJA1* (E3 ubiquitin-protein ligase Praja 1), *CUL9* (E3 ubiquitin-protein ligase Cullin-9), *HERC1* (Probable E3 ubiquitin-protein ligase HERC1).

Calcium Ion Dynamics

GWASs, exome sequencing, and proband studies have all indicated mutations in either $Ca_v3.1$ or $Ca_v3.3$ as risk factors for schizophrenia. While this is an indication for the importance of the $Ca_v3.x$ sector [123] in severe mental disorders, its interpretation is remarkably less straightforward. For instance, rare $Ca_v3.3$ mutations implicated in schizophrenia have been observed to affect an N-glycosylation site, limiting transport to the cell membrane, and indirectly reducing Ca^{2+} currents [124]. Similarly, protein truncating variants of $Ca_v3.1$, which are suspected to mark the protein for degradation and hence reduce Ca^{2+} currents, were associated with schizophrenia though the SCHEMA exome meta-analysis. On the other hand, a detailed analysis of an entire allelic series of $Ca_v3.3$ in a cohort of swedish patients has recently indicated that rare current reducing variants of $Ca_v3.3$ may actually protect against schizophrenia [125].

Loss of function mutations in the sarcoplasmic/endoplasmic reticulum calcium (SERCA) pump [126] encoded by *ATP2A2* are also implicated in schizophrenia and bipolar disorder [72]. The main function of SERCA pumps is to transport calcium from the cytosol into the sarcoplasmic reticulum. The number and efficiency of these SERCA pumps has been shown to be an important factor for the amount of Ca^{2+} that remain present in the sub-membrane area of a neuron after depolarisation [127]. Experiments on *ATP2A2* heterozygous neurons have observed that these neurons exhibit a slower decay of the cytosolic Ca^{2+} levels, which confirmed an essential role of the ATP2A2 pump for Ca^{2+} homeostasis in

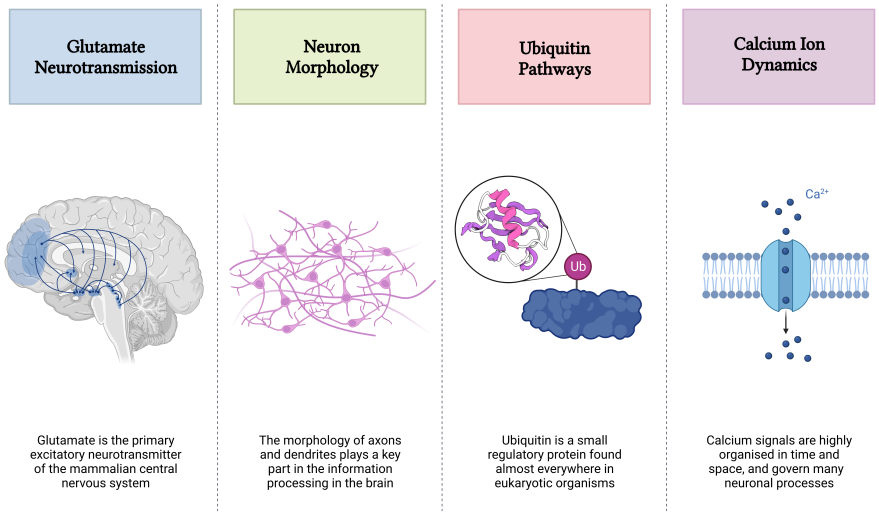

Figure 1.4: **Four Aspects of Schizophrenia Pathophysiology.** Based on a variety of different screening techniques for different sources of genetic risk associated with schizophrenia, we highlight four genetic sectors of importance: glutamatergic neurotransmission, neuron morphology, ubiquitin pathways, and calcium ion dynamics.

neurons [73]. Ca^{2+} currents play an important role in many neuronal processes and hence the downstream effect of changes in Ca^{2+}-levels are hard to predict.

Relevant Genes

CACNA1C ($Ca_v1.2$), *CACNA1G* ($Ca_v3.1$), *CACNA1I* ($Ca_v3.3$), *ATP2A2* (ATPase Sarcoplasmic/Endoplasmic Reticulum Ca^{2+} Transporting 2).

There are a few interesting observations to be made from this overview. First and foremost, the rather unsurprising importance of neuronal and synaptic dysfunction in severe mental disorders. Secondly, the consistent implication of the glutamatergic system (also supported by findings of lower post-mortem levels of glutamate receptors in schizophrenia patients [128]) and the absence of the dopaminergic system. Given the large body of post-mortem [129], preclinical [130], pharmacological [131] and neuroimaging studies [132] supporting dopamine dysregulation in severe mental disorders, genetic evidence is remarkably sparse. Only one dopamine receptor (D2) has been directly implicated in GWASs, indicating that aberrant dopamine signalling is likely due to complicated neuromodulatory effects [133]. This observation is even more remarkable in the light of the common use of dopamine antagonists as antipsychotics [134] (although only effective for a part of the patient population [135]) and the limited success of glutamatergic medication [136].

Thirdly, implicated genes can have multiple effects on different aspects of neuronal functioning. *TRIO* and *SETD1A* for example, affect both neuronal morphology and glutamatergic neurotransmission. It is not always clear which effects are relevant for the disorder, and whether or not separate effects are related. As a final observation, it is important to note the impact of post-translational protein modifications such as ubiquination, glycolisation or lipidation in severe mental disorders [137, 138]. In addition to marking proteins for degradation, post-translational protein modifications can alter the location of proteins in the cell, change their substrate activity, and promote or prevent protein-protein interactions [139, 140]. A growing body of literature reports glycosylation and lipidation abnormalities in many diseases [141, 142], including schizophrenia [143, 144] and bipolar disorder [145, 146].

From the genes with rare mutations identified in the SCHEMA meta-analysis, only *XPO7* (Exportin-7) and *RB1CC1* (RB1 Inducible Coiled-Coil 1) have remained undiscussed so far. The protein encoded by *XPO7* is thought to be the chaperone of other proteins, RNA, and smaller substrates through nuclear pores and into the cytoplasm [147]. The specific role of this protein in neuronal processes and by consequence severe mental disorders is currently unknown. Little is known about *RB1CC1* [148], except for the fact that its structure indicates it may be a regulator of the tumour suppressing *RB1* gene. Hence *RB1CC1* mutations are suspected to play a role in the occurrence of several types of cancer [149, 150]. The roles of *XPO7* and *RB1CC1* will hopefully become more clear as the molecular mechanisms involved in severe mental disorders are discovered.

The Importance of Mitochondrial DNA and Human Accelerated Regions

Novel sources of disease liability for schizophrenia and bipolar disorder are sought after intensely to explain the heritability of severe mental disorders not covered by the previously discussed genetic studies [151]. While human mitochondrial DNA is very small, it only encodes 37 genes, novel risk factors for schizophrenia and bipolar disorder could potentially be found in this genetic material [152]. A noteworthy decrease in mitochondrial DNA oxidation has been observed in patients with bipolar disorder and schizophrenia [153, 154]. At the same time, mitochondria are also involved in Ca^{2+} homeostasis, and might affect neuronal function through these pathways [155]. A more detailed understanding of mitochondrial pathophysiology is necessary to explain the role of mitochondrial oxidation in mood disorders and the overall importance of energy dysregulation and oxidative stress in severe mental disorders.

Another source of genetic risk for severe mental disorder might be found in recently discovered coding parts of DNA, called novel open reading frames (nORFs) [156], that were previously disregarded in screening studies. Evidence from massive proteomics and genome sequencing studies has revealed that eukaryotic genomes contain a substantial amount of small, previously uncharacterised open reading frames. These nORFs are typically less than 100 codons long, can be found in diverse regions of the genome, and fall outside the classical, and rather conservative, definition of a gene. Early studies have shown that RNA and transcripts encoded in these nORFs are involved in at least 150 rare diseases [157], and 22 different forms of cancer [158]. The study of nORFs in schizophrenia and bipolar disorders is still in its infancy, but initial results associated

56 and 40 differentially expressed nORFs to schizophrenia and bipolar disorder [159], respectively. For 21 of these nORFS, it was show that encoded transcripts can form tertiary structures, making them potential novel drug targets.

Remarkably, thirteen of the nORFs differentially expressed in sever mental disorder were found in genomic regions of rapid and human-specific evolution known as human accelerated regions (HARs) [160, 161]. HARs are short stretches of DNA, often found near telomeres or next to genes involved in transcription and DNA binding [162], that are strongly conserved across vertebrae species, including chimpanzees, but underwent high rates of mutations in early humans [163]. HARs, which are uniquely human, are suspected to play an important regulatory role in brain and limb development [164], and are known to be enriched source of risk variants for schizophrenia or altered cognitive behaviour [165]. Due to the involvement of HARs in severe mental disorders and the fact that psychotic traits appear to be limited to the human species, it has been hypothesised that schizophrenia and bipolar disorder are genetic by-products of human-specific brain evolution and associated energy consumption [166–168].

Difficulties in Drug Design: Cost, Attrition, and Complexity

Despite massive public and private investment in life sciences and the continual advances of molecular medicine, bringing new therapeutics to market has become an increasingly expensive and time-consuming endeavour. Over the past 60 years, the ratio of new FDA approved drugs per billion US dollars spent has steadily declined [169–171]. Currently, on average, more than twelve years pass between the first inception of a new experimental drug and that compound reaching the patient's bedside [172, 173]. Hotly debated estimates of research and development costs for novel drugs range from 161 million dollar to around 4.5 billion dollars [174]. The cost of developing drugs is largely due to the substantial proportion of failures of drug candidates throughout the development pipeline, see Figure 1.5, often due to unexpected toxicities or a lack of efficacy [175]. The lowest probability of success in drug development is seen for psychiatric disorders [176, 177], likely as a consequence of the complex and only partially understood molecular mechanism underlying these diseases.

During the first decades of the twentieth century, chemicals with medicinal potential were discovered haphazardly, synthesised by dye-producing companies, tried out in rudimentary cell and animal experiments, and clinically tested on whichever (often non-suspecting) group of patients was available [178]. Luckily, professionalism took over the field of pharmaceutical research and since the 1980's and 1990's drug discovery has relied strongly on lab-based target identification [179, 180] and high-throughput screening efforts [181]. High-throughput screening consists of automated experimental screenings of large, proprietary in-house chemical libraries of potential drug candidates. However, even the largest of these chemical libraries only represent a tiny fraction of the possibilities in chemical space and suffer from structural or (historical) biases [182–185]. Large-scale high-throughput screening efforts are regularly mired by false-positives, and drug candidates arising from them have a significant chance of failure in later toxicity screens, animal tests, or clinical trials [186–188].

Clinical trials are a series of rigorous studies of drug effects on human subjects [190].

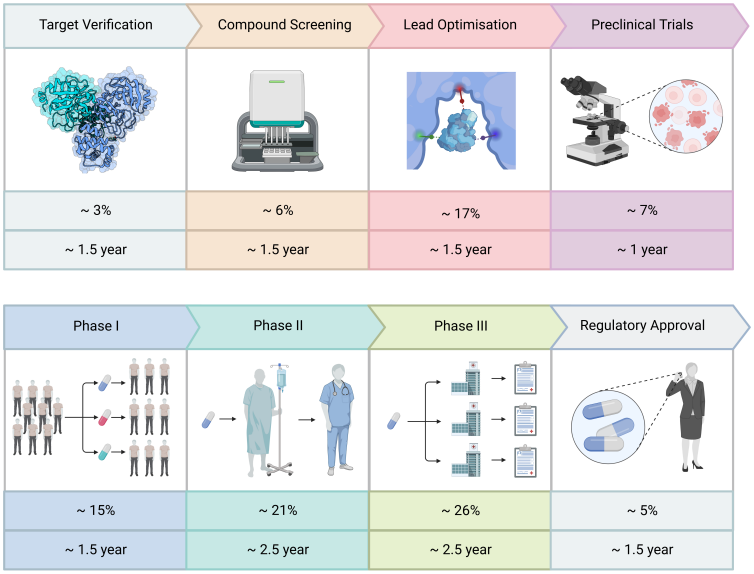

Figure 1.5: **The Drug Development Pipeline.** The process of drug discovery and development, step by step, and the percentage cost of developing a new molecular entity and corresponding cycle times [189] for each of these steps.

These trials are required to strictly follow guidelines from regulatory authorities in countries or regions where the drug candidate is intended to be registered and commercialised. Most commonly these authorities are the Food and Drug Administration (FDA, USA) [191], the European Medicines Agency (EMEA, EU) [192] or the Ministry of Health, Labour and Welfare (MHLW, Japan) [193]. A standard series of clinical trials starts with a small Phase I study and ends with a large-scale Phase III study. The primary purpose of a Phase I study is to evaluate basic safety of a drug candidate and discover side effects in healthy subjects. Using the dose and method found to be the safest in Phase I studies, a small group of patients are given the new treatment in the Phase II study. Phase III studies enrol a large number of patients in double-blind tests that compare the safety and effectiveness of the new treatment against the current standard treatment or a placebo treatment. If a compound successfully passes through these consecutive trials, it can be submitted for regulatory approval.

Unfortunately, up to 90% of drug candidates fail when they enter clinical studies. A review of clinical trial data [189] (during a seven year period starting in 2010) shows that late-stage clinical failures are most often due to a lack of clinical efficacy, unacceptable toxicity levels, poor absorption, distribution, metabolism, excretion, and toxicity (ADMET) characteristics, or a lack of commercial value or clinical need. Because finding small molecules that fit regulatory requirements and have sufficient therapeutic and commercial potential becomes harder and harder, novel drug modalities such as

antibodies [194], peptides [195], macrocyclic molecules [196], protein degraders [197–199] and gene editing therapies [200, 201] have attracted recent attention [202, 203]. However, in the light of a renewed appreciation of polypharmacology [204–206] and the novel opportunities provided by artificial intelligence breakthroughs [207], enthusiasm around small molecule drug design has (at least partially) rebounded. Because of their favourable properties for targeting the central nervous system [208, 209], we will limit our focus to the small molecule modality in this dissertation.

The Trials and Tribulations of Antipsychotic Medicine

A wide array of factors is involved in the current decline of efficiency in drug design, including decreased risk-tolerance of regulatory agencies [210], but for drugs related to severe mental disorders the situation is uniquely challenging [211]. The complex and poorly understood pathophysiology of schizophrenia and bipolar disorder, as discussed previously, complicates the identification of drug targets. In addition, it remains hard to predict the cumulative network effect of altering the molecular functioning of neuronal processes. On top of this, extra challenges with regard to penetrating the blood-brain barrier [212], toxicities [213, 214] and off-target effects [215] are present for molecules that act on the brain. Finally, antipsychotic medication deals with specific struggles regarding medication adherence with estimates of non-adherence for this class of drugs varying from 4% to 72% [216].

Before 1952, treatment options for severe mental disorders were typically limited to desperate interventions such as electroconvulsive therapy, prefrontal lobotomies, and insulin-induced coma therapy [217]. Early pharmacological attempts to create antipsychotic chemicals were based on the sedative and calming effects seen in animals, when exposed to compounds binding to histamine receptors. In 1947, Promethazine [218] (a first generation antihistamine) was found to enhance calmness in human subjects, but was not deemed sufficiently efficient to combat psychotic symptoms. As a follow-up, Chlorpromazine [219] – a Promethazine derivative – was tried in a military hospital in France with the aim to calm down patients before a surgery. In 1952, Chlorpromazine was successfully tried as an antipsychotic agent in 38 patients [220]. The drug significantly reduced positive psychotic symptoms such as hallucinations and delusional thinking [221].

After this initial breakthrough, Chlorpromazine was promoted as the first effective antipsychotic and prescribed around the world, which massively reduced the amount of chronically hospitalised patients with psychosis. This wide-spread success led to the development of a whole array of antipsychotic medication, now known as first generation antipsychotics [222], see Figure 1.6. Soon afterwards, seminal studies linked the effective doses of various antipsychotic drugs to corresponding antagonism levels of dopamine receptors [223–225]. Antipsychotic effects have consistently been shown to occur when the occupation of striatal D2 receptors is around 65% [226]. Further increases in antagonism levels have been associated with stronger side effects and not with improved antipsychotic efficacy [227]. Unfortunately, first generation antipsychotics fail to ameliorate core cognitive or negative symptoms (depression, anhedonia) in a large amount of patients [228].

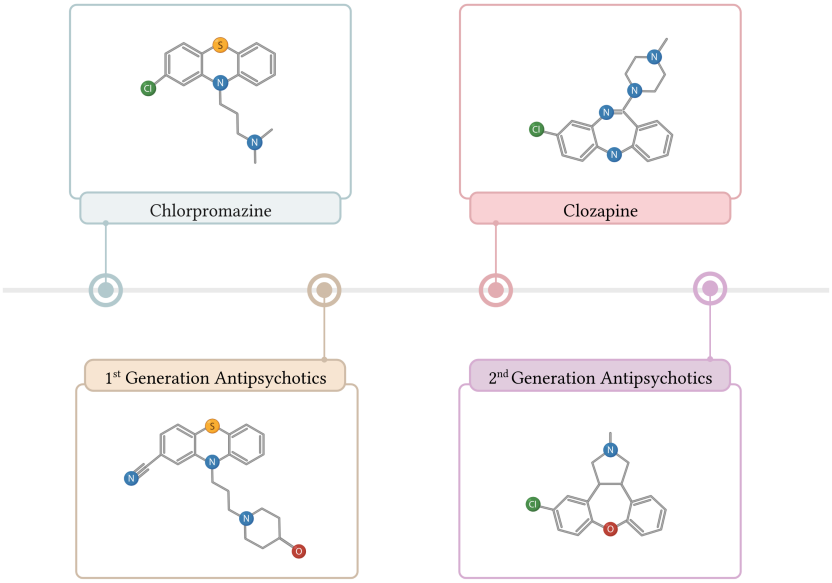

Figure 1.6: **Evolution of Antipsychotics Compounds Through the Decades.** After the unexpected discovery of the antipsychotic properties of Chlorpromazine [220], a group of first generation or *typical* antipsychotic, here represented by Periciazine [229], was developed until the early 1970's. From 1972 on, Clozapine [230] was established as the first second generation or *atypical* antipsychotic. The later compounds of the second generation of antipsychotics are represented here by Asenapine [231].

In 1961, a uniquely efficacious antipsychotic drug known as Clozapine was developed [230]. Initial pre-clinical optimism regarding this compound was limited, as Clozapine only inhibits 30% to 40% of stratial dopamine receptors [232, 233]. However, over time it has become the only antipsychotic medication that has proven its effectiveness against cases of otherwise treatment-resistant schizophrenia. Even at comparatively low levels of D2 antagonism, 60%–70% of patients with resistance to other antipsychotics were found to respond to Clozapine [234]. The precise pharmacological mechanisms facilitating Clozapine's effectiveness remain unknown [235] and debate about the compound's influence on glutamatergic neurotransmission continues actively [236–240]. Since its discovery, there have been many attempts to develop an equally effective compound with less side effects, but so far all these efforts have been in vain.

First generation antipsychotics typically cause a significant burden on patients through extrapyramidal side effects (EPSEs) [241] such as acute dystonias, akathisia, parkinsonism and tardive dyskinesia. Clozapine, while an effective antipsychotic, has the potential to cause serious cardiovascular conditions [242] including orthostatic hypotension, bradycardia, and syncope. In 1975, five Finnish patients on Clozapine

treatment died [243, 244] from secondary infections due to agranulocytosis. The severity of ESPEs in first generation antipsychotics led to the invention and introduction of the second generation antipsychotics in the 1990's [245]. While second generation antipsychotics have a similar effectiveness to first generation antipsychotics [246] and a lower propensity to cause EPSEs [247], they are often associated with metabolic abnormalities [248] such as weight gain, dyslipidaemia, or glucose dysregulation. Side effect of specific antipsychotics can typically be related to their binding profiles with regard to a wide variety of receptor types [249, 250].

Well over half a century after the invention of Chlorpromazine, the field of antipsychotic medication finds itself in dire need of innovative new drug targets and therapies [251]. Many of the major neuropsychiatric drug development programs at pharmaceutical companies have been shutdown or externalised, often due to high clinical failure rates and lack of success in standard research approaches [252]. For instance, determining a clinically optimal dose for novel compounds remains challenging due to a scarcity in precise biomarkers [253] and phase III clinical trials of a highly selective glycine reuptake inhibitor (Bitopertin) had to be discontinued due to a lack in efficacy [254]. To prevent a standstill in the development of novel antipsychotics, the European Union and the World Health Organisation have called for additional basic and applied research into these disorders [255]. In addition to the possibility of developing course-altering medication as mentioned earlier, approaches based on polypharmacology [256, 257] might also provide new and fruitful avenues for clinical progress.

Computational Drug Design in the Age of Artificial Intelligence

Computer-assisted drug design has long been seen as a potentially potent tool for improved compound selection and optimisation in challenging drug design projects, such as the development of novel antipsychotics. Recent advances in artificial intelligence (AI) and deep learning have been applied to drug design to improve the odds of selecting the right compounds for further development [207, 258, 259]. Machine learning models have been used to predict physicochemical properties [260], off-target biological activity [261], or synthetic accessibility [262–264] of candidate molecules. Often, these models are used in conjunction with expert medicinal chemists – colloquially known as drug hunters – to rank or prioritise certain compounds for further experimental studies. Deep learning approaches can also be used to predict biological activity of a proposed molecule against protein targets, by estimating its binding strength, but success in this specific (and crucially important) area has been limited [265, 266].

Traditionally, drug candidates used in large screenings efforts were required to conform to one or more rules of thumb to promote oral active in humans [267–269]. This also increases the probability of having compounds with good ADMET properties and easier or more convenient paths of synthesis. The most well-known of these heuristics is Lipinski's rule of five [270, 271]. According to this Lipinski's rule of five, a molecule with favourable drug-like properties has no more than five hydrogen bond donors, no more than ten hydrogen bond acceptors, a molecular mass below 500 atomic mass units, and an octanol-water partition coefficient that does not exceed five. The octanol-water partition coefficient [272, 273] is a measure of a compound's solubility

in fats and oils, also known as lipophilicity. Yet, even with these heuristic rules in place, in-vivo and in-vitro studies of toxicology [274], pharmacodynamics [275], and pharmacokinetics [276] properties in proposed drug candidates remain highly necessary to exclude compounds with a low chance of clinical success.

High-throughput screening has an estimated hit rate between 0.01% and 1% [277]. Consequently, there is a significant opportunity for alternative approaches in the drug discovery process. One such alternative approach, called rational drug design [278–280], reformulates drug discovery as an inverse design problem. Typically, expert medicinal chemists make use of experimentally determined protein structures and aim to develop small molecules tailor made to interact with certain binding sites on these targets. Most of the time, drugs are designed to inhibit the working of an enzyme's active site [281], due to comparatively easy target validation and binding options. However, by binding to regulatory (or *allosteric*) pockets, which are typically harder to find and bind to, the working of a target protein can be modulated in a more intricate fashion [282, 283]. To accelerate this process, active learning models [284] can be used which interactively indicate which molecules need to be synthesised and tested to gain a more accurate model which can then eventually help select the most promising candidates for further development.

Molecular Representations for Computational Drug Design

One of the main challenges of using deep or active learning to predict the properties of molecules, lies in molecular featurisation. Encoding a complex and dynamic, three-dimensional probability density of electrons and nuclei into a machine-readable format requires a significant level insight into the mixture of chemistry and computation that is known as cheminformatics, and can be done in a myriad of ways. The most common way to turn a molecule into a machine-readable object is through a string-of-characters encoding known as the simplified molecular input line-entry system (SMILES) [285, 286]. A SMILES string is typically obtained by depth-first traversing a hydrogen-reduced molecular graph and representing encountered atoms and bonds as symbols from a predefined set. This approach allows for the use of a powerful system of regular expressions tailored to SMILES strings, known as SMILES Arbitrary Target Specification (SMARTS) [287, 288].

A novel text encoding system for molecules, called self-referencing embedded strings [289] (SELFIES), was recently developed to have a one-on-one correspondence between all valid chemical structures and all allowed SELFIES encoding. Learned vector-encodings of text [290–292], sound [293], or images [294, 295] have become commonplace in deep learning and similarly learned representations of molecules, sometimes starting from SMILES or SELFIES, have also been invented in recent years [296]. Graph-based deep learning [297–300] is currently a very active field of research and a natural use-case for learning molecular embeddings [301–303]. Algorithmically encoding molecules in hand-crafted vector-encoding systems, however, has a long tradition that reaches back to the earliest days of the cheminformatics and these descriptors are still often used [304, 305].

The simplest hand-crafted molecular descriptors consist of a vector listing the individual physicochemical properties of a molecule. More intricate descriptors, such as

structural keys or hashed fingerprints, encode the presence of specific chemical groups and substructures into a bit-vector [306–308]. Structural keys check a molecule against a few hundred predefined structural patterns, whereas hashed fingerprints generate a bit-vector by iteratively hashing all present substructures under a certain size. One of the most well-known structural keys are the molecular access system (MACCS) keys [309–311]. Path-based fingerprints [312] on the other hand, break up a molecule along paths formed by its bonds. The substructures that are found along this path are then hashed into a bit-vector. Circular fingerprints [313], such as the extended connectivity fingerprints (ECFP) [314], follow a similar procedure but find and encode substructures by considering increasingly larger circular neighbourhoods around each atom.

Deep Learning for De Novo Drug Design

The success of neural network architectures [315] in generating images, text, or sound has sparked an interest in the cheminformatics community. Based on large datasets, such as ChEMBL [316] (2.3 million compounds), PubChem [317] (112 million compounds) or ZINC [318] (230 million compounds), deep learning models have been developed that aim to generate drug-like molecules [319, 320]. In practice these models learn underlying chemical patterns and reproduce and combine these to form novel molecules. A whole array of deep generative models with different architectures and acting on different molecular representations has been proposed for these purposes [321, 322], see Figure 1.7. To a limited extent, and depending on how the proposed molecules are evaluated, the design of novel compounds by these generative algorithms can be seen as a computational counterpart to high-throughput screening or the work done by medicinal chemists in rational drug design.

Most notable among the deep generative models for drug design are the recurrent neural networks (RNNs) [323] trained on SMILES. RNNs are commonly used as a generative model for data of a sequential nature such as natural language, or sound [324–330]. By applying RNNs to a SMILES representation of molecules, the model can learn both the syntax of this language as well as the statistical distribution of chemical patterns present in a database. Variational autoencoders (VAEs) [331, 332] trained on either SMILES, SELFIES, or directly on chemical graphs, consist of two parts: an encoder and a decoder. Both the encoder and the decoder are deep neural networks, which respectively compress a training molecule's representation to a continuous latent space, and aim to reconstruct it from that latent space. The latent space acts like an information bottle neck, and can be used to sample novel molecules which are expected to have a structure similar to those on which the VEA was trained.

Generative adversarial networks (GANs) [333, 334], and normalising flow models [335, 336] build on ideas developed around VAEs. Generative adversarial networks are formed by two deep neural networks, one which generates molecules from random Gaussian input, and a competing discriminatory network which aims to distinguishes candidates molecules produced by the generator from the molecules found in a training dataset. Through this zero-sum competition, both models are trained together until the generator produces molecules that are sufficiently indistinguishable from the training data. Normalising flow models are similar to variational autoencoders

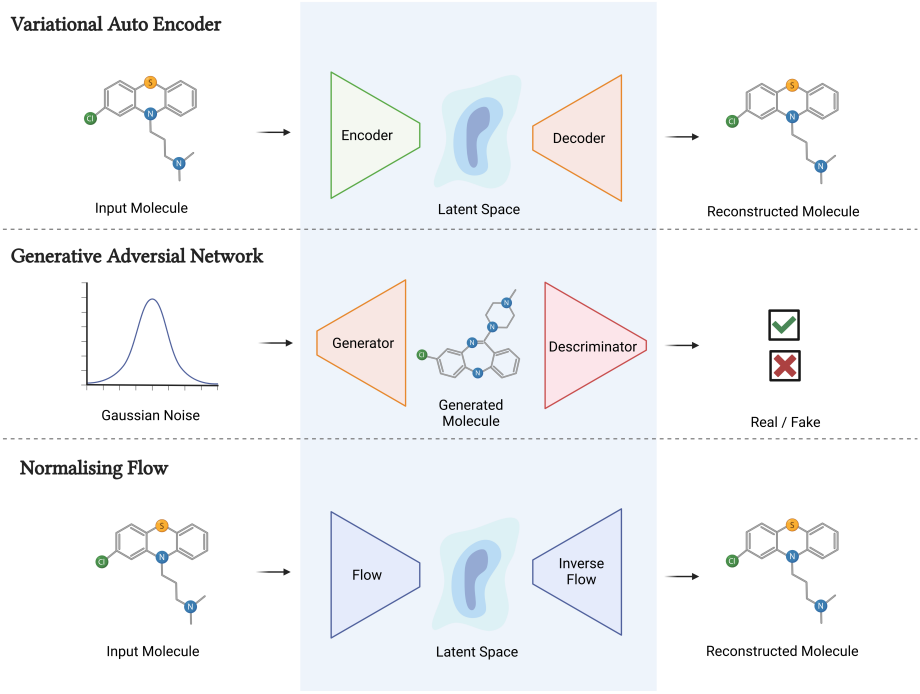

Figure 1.7: **Deep Generative Models for De Novo Drug Design.** Three of the major artificial neural network architectures commonly used in deep generative modelling approaches for de novo drug design are shown here: variational auto encoders (top), generative adversarial networks (middle), and normalizing flow (bottom). Based on Figure 1 of *Generative Models for Molecular Discovery: Recent Advances and Challenges* [322].

but make use of a single deep neural network with bijective layers. That single network is then simultaneously applied in one direction for encoding molecules into a latent space and inverted in the other direction for decoding, making them more stable to train than GANs or VAEs.

The methods discussed so far can be used to generate molecules similar to those in existing databases, but they do not optimise molecules for specific properties (such as a target lipophilicity or protein binding affinity). To obtain optimised molecules, reinforcement learning, bayesian optimisation methods, or genetic algorithms have to be used. Reinforcement learning [337] and bayesian optimisation methods [338] act on previously trained deep generative models which they then tune or use to generate molecules with the desired properties. Genetic algorithms, on the other hand, require nothing more than an initial population of molecules and a set of rules for molecular transformations which it then applies iteratively until it obtains sufficiently optimised molecules. Large benchmarking efforts [339, 340] on a variety of optimisation tasks have shown that a genetic algorithm applied directly to molecular graphs, known as the

graph-based genetic algorithm (GB-GA) [341], is generally more efficient and effective than deep learning based algorithms.

Unfortunately GB-GA is susceptible to stagnation around local optima. Stagnation is a common issue for genetic algorithms, and luckily a solution for the problem can be found in a remarkable diversification technique originally developed in the field of soft robotics. In the second chapter of this dissertation, we discuss the effect of applying this technique to GB-GA. Medicinal chemists often need to find a balance between many different aspects and properties of a single candidate molecule [342], such as potency and selectivity with regard to the chosen pharmacological targets, off-target activity, physicochemical properties, and synthetic accessibility. To solve these types of multi-objective optimisation (MOO) [342] problems, a series of genetic algorithm extensions called non-dominated sorting genetic algorithms (NSGAs) [343, 344] have been invented. In the second chapter we merge GB-GA with two different NSGA algorithms ans discuss the potential usefulness of the resulting algorithms in chemical optimisation and drug design.

Opportunities in Neuroscience: Accelerating Disease Models

Neuroscience is a research field devoted to the exploration and understanding of the central nervous system [345]. It is a multidisciplinary branch of science that combines theoretical debate and a diversity of experimental efforts with computational tools for brain simulation and neural data analysis. Most types of neural data recorded up to now are dividable into three categories: imaging data, spike data, and voltage trace data. Due to recent and fast-paced technological innovations, experimental data has become more reliably and abundantly available to the neuroscience community than ever before. Miniature two-photon miniscopes allow for the simultaneous calcium imaging recording of more than a thousand neurons in freely moving mice [346], miniaturised high-density probes enable stable and long-term recordings of spike dynamics of thousands of individual neurons [347], while recently developed patch-sequencing techniques combine single cell RNA-sequencing with patch-clamp electrophysiology recordings to facilitate multimodel characterisations of single neurons [348, 349].

In addition to novel tool-kits developed to analyse this recent influx of neural data, such as deep-learned latent spaces for neural dynamics [350] or topological data analysis [351] based on persistent homology, computational neuroscience has also seen the rapid development of increasingly larger simulations of networks of model neurons. While computationally expensive, large-scale computer models of networks of neurons form an important connection between single neuron and system-level analysis of the brain [352]. Being able to tune various parameters of neuron models, based on experimental inputs, is a way to use computational models as a counterpart to laboratory experiments. In these networks, model neurons can be represented at various levels of abstraction [353, 354]. The three most common levels of abstraction are biophysically-detailed models, point-neuron models, and population-level models.

In biophysically-detailed models, single neurons are represented by a compartmentalised version of their physical morphology. Tailor-made, phenomenological equations track the different ion currents that flow in, out, and through each of these compartments

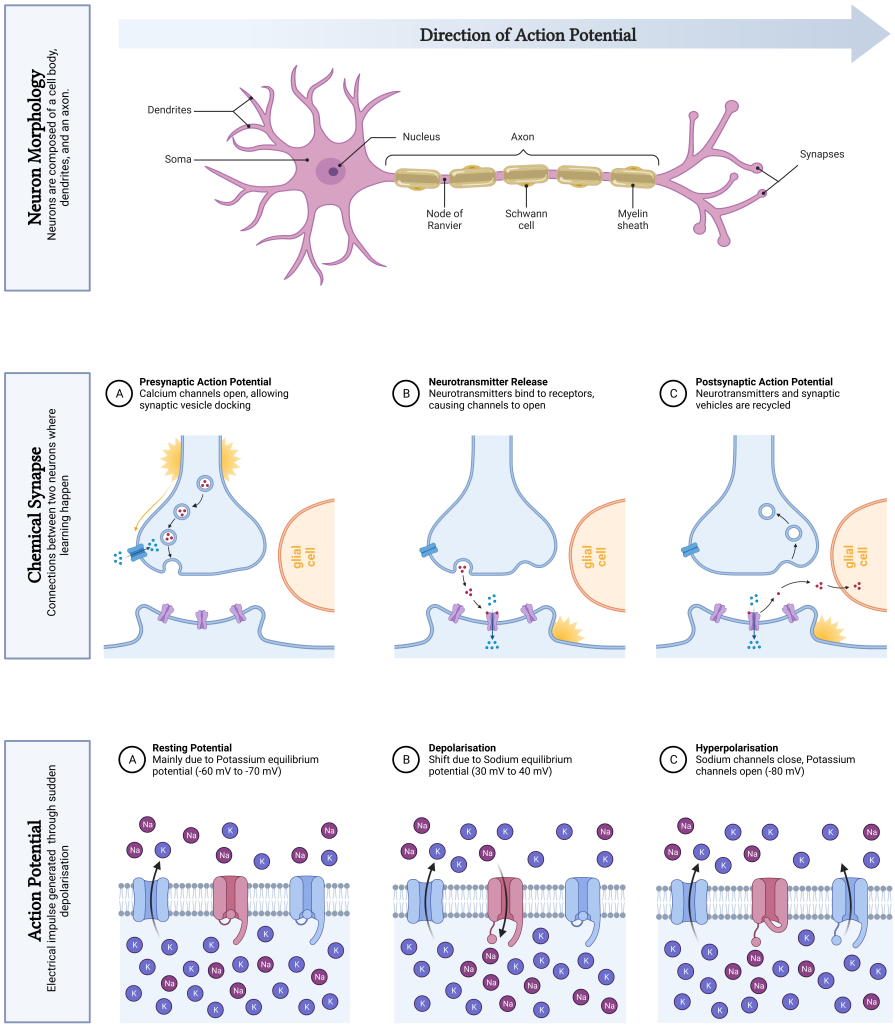

Figure 1.8: **Basic Morphology and Properties of a Neuron.** Neurons are electrically excitable cells that generate electrical pulses, known as action potentials, in their soma. The rest of the neuron conducts these pulses and passes on excitatory or inhibitory signals to other neurons via specialised (chemical) connections known as synapses. Schwann cells are a type of glial cell that cover axons in a myelin-sheath to improve action potential conduction. Microscopic gaps in the myelin-sheath, known as nodes of Ranvier, further facilitate efficient action potential conduction. In this figure the basic morphology of a neuron (top), the molecular mechanisms of a synapse (middle), and the different stages of an action potentials (bottom) are shown. In a biophysically-detailed neuron model, the neuron morphology is turned into a graph of connected compartments and both synapses and action potentials are described by a set of (differential) equations for each compartment.

to determine their membrane potential. Point-neuron models are simpler and do not have any explicitly modelled spatial extent. Point-neurons rely on integrate-and-fire dynamics to turn synaptic inputs into an output signal. Instead of directly simulating any form of electrophysiogical dynamics of an individual neuron at all, population-level models make use a single population-density equation for the entire voltage distribution of a population of neurons. In the past decade, biophysically-detailed models have been used to reconstruct the neuronal microcircuitry of the somatosensory cortex of the juvenile rat (Blue Brain Project) [355], and the mouse primary visual cortex (Allen Institute for Brain Science) [356] based on sparse biological data.

To obtain accurate biophysically-detailed network models [357, 358], simulations should not only be compared to experimental measures such as voltage traces and spiking statistics, but also to population-level measures such as local field potentials (LFPs), or electrocorticography (ECoG) and electroencephalography (EEG) measurements. In contrast to action potentials, LFPs capture the low-frequency part of electric potentials inside gray matter [359–361]. In addition to combining all these signals to pinpoint parameters and network geometries in brain simulations, biophysically-detailed models can also be used to predict ECoG and EEG signals based on underlying neural activity [362]. These models could, for instance, be used to predict the impact of different genetic mutations on EEG measurements, allowing us to match genetic information to clinical biomarkers of severe mental disorders [363–366].

Distilling Biophysically-Detailed Neuron Models

Unfortunately, simulating even a few seconds of biological activity of a large biophysically-detailed network model requires hours of supercomputer usage. The required computational resources are not always widely available outside large academic institutions, and the required runtime limits flexible use of the models. Faster biophysically-detailed network models would allow for better parameter optimisation and larger and more realistic model networks. The past years have seen the development of graphics processing units (GPU) based [328] and simplified neuron simulations [367] to fulfil some of these needs. To be able to run even larger and faster simulations, recent attention has turned to the idea of distilling computationally intensive biophysically-detailed neuron models into easier-to-evaluate artificial neural networks.

Several artificial neural network architectures that model the activity of biophysically-detailed cortical neuron models, including those with non-linear dendritic dynamics, have been devised and shown to be effective. Based on synaptic inputs, and in some cases the membrane potential of each compartment, at the previous time-step, these deep neural networks predict either the generation or absence of an outgoing action potential or the membrane potential in the soma of the modelled neuron. Current deep learning approaches that distil biophysically-detailed neuron models typically make use of convolutional neural networks (CNN) [368–370], a class of artificial neural network most often used to learn aspects of visual data because they efficiently learn spatial correlations. With some simple adaptations, CNNs can be turned into an effective tool for sequence modelling and forecasting known as a temporal convolutional network (TCN) [371].

A TCN is a causal, one-dimensional CNN with the same input and output sequence lengths for each layer, including the input and output layer. To ensure causality in each layer and to make sure that the output sequence has the same length as the input sequence, zero-padding and dilation are applied. This relatively simple architecture has been shown to be sufficient to predict action potentials [372]. Time-series forecasting of the soma membrane potential, i.e. a voltage trace, has been shown to work by adding a long short-term memory (LSTM) [373–376] network, a type of RNN, on top of a TCN architecture and by adding the membrane potential of each compartment at the previous time-step in the input data [377–379].

Instead of having to calculate large systems of coupled differential equations to simulate biophysically-detailed neurons, distilled neuron models can be evaluated through simple combinations of tensor operations. which can easily be accelerated on specialised equipment [380]. In addition to these speed-ups, deep learning approaches can evaluate inputs in batches so that simulation run-times can increase sub-linearly in function of the amount of simulated neurons. Simultaneous simulation of 5000 deep-learned neuron models have been shown to run faster than a classical simulations of a single biophysically-detailed neuron model [377]. In the following chapter, we show how to use multi-task learning (MTL) to optimizes a TCN model to perform simultaneous predictions of both action potentials and membrane potentials across all compartments of a biophysically-detailed neuron model.

In-Vivo, In-Vitro, and In-Silico Models of Severe Mental Disorders

Psychiatric disorders are difficult to study due to their complex genetic architecture, developmental nature, and the lack of direct access to the affected neural tissue of patients. Hence, to validate and interrogate the effect of genetic variations implicated in these diseases, neuroscientists make use of in-vivo, in-vitro, and in-silico models. These disease models (respectively making use of living animals, advanced laboratory techniques, or purely digital means) aim to gain novel insights into a disease or provide the pharmaceutical sector with a strongly supported targets for clinical intervention. On several occasions during the research period presented in this dissertation, there has been interaction and (informal) collaboration with experts, both at the University of Oslo (UiO) and the Oslo University Hospital (OUS), specialised in certain disease models for severe mental disorders. The scope and outcome of these interactions are briefly recapitulated here, as is the use and usefulness of biophysically-detailed neuron models for severe mental disorders.

After the invention of recombinant DNA technologies in the early 1970's [381, 382], genetically altered rodent models started to be used intensely to study a variety of neurological diseases in-vivo. In the case of schizophrenia and bipolar disorder, a significant amount of attention has gone to *CACNA1C* gene knock-out rodent models and their aberrant behaviour in working-memory tasks [383, 384]. However, the specific molecular effects of common *CACNA1C* variants in schizophrenia and bipolar disorder are still unclear. In light of the more easily interpretable SCHEMA consortium exome sequencing results, we have proposed a *CACNA1G* gene knockout model based on the more clearly supported deleterious effect of protein truncating variants found for this gene. As a consequence, a CRISPR/Cas9 knockout protocol [382, 385–388] for

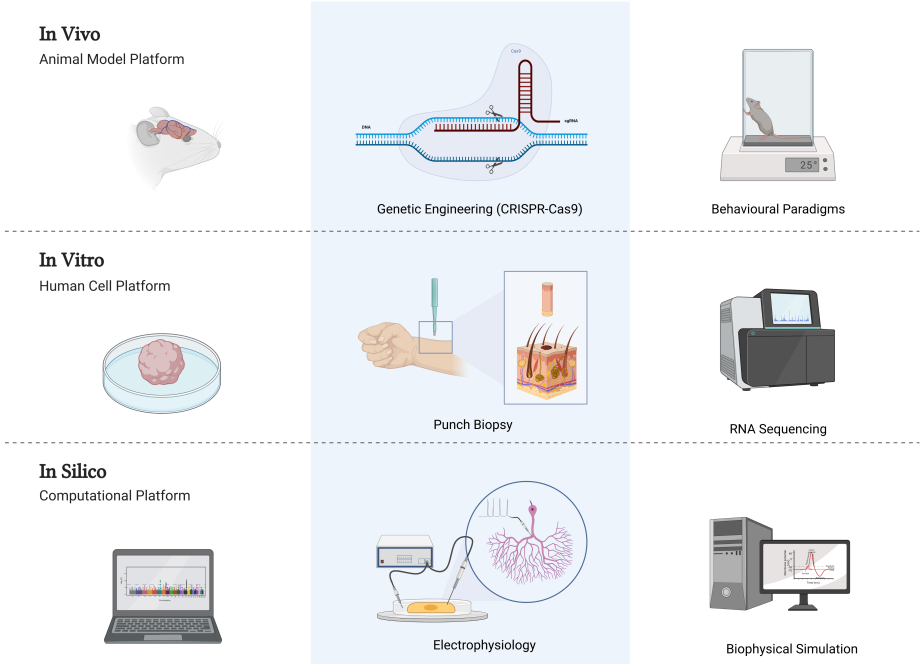

Figure 1.9: **Neuroscience Platforms To Study Severe Mental Disorders.** Different experimental platforms can be used to study severe mental disorders such as schizophrenia and bipolar disorder. Gene editing techniques (top) can be used to build animal models which can be tested on different behavioural tasks. Patient-derived brain organoids allow for the electrophysiological and genetic interrogation of human-like neural tissue (middle). Patch-clamp electrophysiology of single neurons allows for the construction of biophysically-detailed neuron model which can be tuned to emulate the effect of genetic mutations on small neural circuits (bottom).

CACNA1 genes in rodents and a repetition of working memory tasks with this novel phenotype has been included in a FODS application by members of the Hafting-Fyhn research group (UiO).

Blood samples (specifically the red blood cells they contain) have been used as a proxy to study mRNA expression levels in human brain cells [389]. An alternative approach, being applied by our collaborators from the Djurovic group (OUS), makes use of state-of-the-art cell cultures known as brain organoids, to create in-vitro access to proxy brain cells. These brain organoids are derived from patient-specific, pluripotent stem cells obtained through skin biopsies, and transformed into a self-organising and three-dimensional cell culture by means of an intricate and expensive chemical process [390–392]. We have supported their preliminary analysis, in which they compare brain organoids derived from fourteen schizophrenia patients with brain organoids derived from healthy controls, by providing a theoretical context regarding upregulation of an ion transporter and the effect on GABA neurotransmission as regulated through

intraneuronal chloride concentrations.

Attempts to evaluate the influence of genetic variants and differentially expressed protein levels in computational models have been made [393–396]. Due to the challenges in linking genetic mutations to their functional effects, ad-hoc parameter association and modification is often used to link these models to questions regarding the molecular mechanisms of severe mental disorders. Hence, these simulations (and their outcome) should be seen as prototypes rather than being accepted as precise disease models for severe mental disorders. Better knowledge of the functional effect of individual mutations, more accurate biophysically-detailed neuron models, and efficient large-scale biophysically-detailed simulations powered by deep learning approaches are necessary to turn the current methods into tools with real-world applications such as patient stratification or disease mechanism elucidation.

References

1. Edition, F. *et al.* Diagnostic and statistical manual of mental disorders. *Am Psychiatric Assoc* vol. 21, 591–643 (2013).
2. First, M. B. *et al.* DSM-5-TR: overview of what's new and what's changed. *World Psychiatry* vol. 21, 218 (2022).
3. Nussbaum, A. M. *et al. The Pocket Guide to the DSM-5-TR™ Diagnostic Exam* (American Psychiatric Pub, 2022).
4. Wainberg, M. L. *et al.* Challenges and opportunities in global mental health: a research-to-practice perspective. *Current psychiatry reports* vol. 19, 1–10 (2017).
5. Druss, B. G. Improving medical care for persons with serious mental illness: challenges and solutions. *Journal of Clinical Psychiatry* vol. 68, 40 (2007).
6. Liu, N. H. *et al.* Excess mortality in persons with severe mental disorders: a multilevel intervention framework and priorities for clinical practice, policy and research agendas. *World psychiatry* vol. 16, 30–40 (2017).
7. Hjorthøj, C., Stürup, A. E., McGrath, J. J. & Nordentoft, M. Years of potential life lost and life expectancy in schizophrenia: a systematic review and meta-analysis. *The Lancet Psychiatry* vol. 4, 295–301 (2017).
8. Correll, C. U. *et al.* Mortality in people with schizophrenia: a systematic review and meta-analysis of relative risk and aggravating or attenuating factors. *World psychiatry* vol. 21, 248–271 (2022).
9. Collins, P. Y. *et al.* Grand challenges in global mental health. *Nature* vol. 475, 27–30 (2011).
10. Kahn, R., Sommer, I. & Murray, R. Schizoprenia. *Nat Rev Dis Primers* vol. 1, 15067 (2015).
11. Vieta, E. *et al.* Bipolar Disorders. *Nat Rev Dis Primers* vol. 4, 18008 (2018).
12. Kantrowitz, J. T. & Citrome, L. Schizoaffective Disorder. *CNS drugs* vol. 25, 317–331 (2011).

13. Andreasen, N. C. A unitary model of schizophrenia: Bleuler's fragmented phrene as schizencephaly. *Archives of general psychiatry* vol. 56, 781–787 (1999).

14. Craddock, N. & Owen, M. J. The Kraepelinian dichotomy–going, going… but still not gone. *The British Journal of Psychiatry* vol. 196, 92–95 (2010).

15. Lichtenstein, P. *et al.* Common genetic determinants of schizophrenia and bipolar disorder in Swedish families: a population-based study. *The Lancet* vol. 373, 234–239 (2009).

16. Smeland, O. B. *et al.* Genome-wide analysis reveals extensive genetic overlap between schizophrenia, bipolar disorder, and intelligence. *Molecular psychiatry* vol. 25, 844–853 (2020).

17. Janoutová, J. *et al.* Epidemiology and risk factors of schizophrenia. *Neuroendocrinology Letters* vol. 37, 1–8 (2016).

18. Rowland, T. A. & Marwaha, S. Epidemiology and risk factors for bipolar disorder. *Therapeutic advances in psychopharmacology* vol. 8, 251–269 (2018).

19. Gore, F. M. *et al.* Global burden of disease in young people aged 10–24 years: a systematic analysis. *The Lancet* vol. 377, 2093–2102 (2011).

20. Birnbaum, R. & Weinberger, D. R. Genetic insights into the neurodevelopmental origins of schizophrenia. *Nature Reviews Neuroscience* vol. 18, 727–740 (2017).

21. Kahn, R. S. & Keefe, R. S. Schizophrenia is a cognitive illness: time for a change in focus. *JAMA psychiatry* vol. 70, 1107–1112 (2013).

22. Rapoport, J., Giedd, J. & Gogtay, N. Neurodevelopmental model of schizophrenia: update 2012. *Molecular psychiatry* vol. 17, 1228–1238 (2012).

23. Millan, M. J. *et al.* Altering the course of schizophrenia: progress and perspectives. *Nature Reviews Drug Discovery* vol. 15, 485–515 (2016).

24. Gonzalez-Escamilla, G., Muthuraman, M., Chirumamilla, V. C., Vogt, J. & Groppa, S. Brain networks reorganization during maturation and healthy aging-emphases for resilience. *Frontiers in psychiatry* vol. 9, 601 (2018).

25. Mira, H. & Morante, J. Neurogenesis from embryo to adult–lessons from flies and mice. *Frontiers in Cell and Developmental Biology* vol. 8, 533 (2020).

26. Garner, C. C., Zhai, R. G., Gundelfinger, E. D. & Ziv, N. E. Molecular mechanisms of CNS synaptogenesis. *Trends in neurosciences* vol. 25, 243–250 (2002).

27. Rusznák, Z., Henskens, W., Schofield, E., Kim, W. S. & Fu, Y. Adult neurogenesis and gliogenesis: possible mechanisms for neurorestoration. *Experimental neurobiology* vol. 25, 103 (2016).

28. Paolicelli, R. C. *et al.* Synaptic pruning by microglia is necessary for normal brain development. *science* vol. 333, 1456–1458 (2011).

29. Jan, Y.-N. & Jan, L. Y. Branching out: mechanisms of dendritic arborization. *Nature Reviews Neuroscience* vol. 11, 316–328 (2010).

30. Hill, R. A., Li, A. M. & Grutzendler, J. Lifelong cortical myelin plasticity and age-related degeneration in the live mammalian brain. *Nature neuroscience* vol. 21, 683–695 (2018).

31. Homem, C. C., Repic, M. & Knoblich, J. A. Proliferation control in neural stem and progenitor cells. *Nature Reviews Neuroscience* vol. 16, 647–659 (2015).

32. Agostini, M. *et al.* Metabolic reprogramming during neuronal differentiation. *Cell Death & Differentiation* vol. 23, 1502–1514 (2016).

33. Susser, E. S. & Lin, S. P. Schizophrenia after prenatal exposure to the Dutch Hunger Winter of 1944-1945. *Archives of general psychiatry* vol. 49, 983–988 (1992).

34. St Clair, D. *et al.* Rates of adult schizophrenia following prenatal exposure to the Chinese famine of 1959-1961. *Jama* vol. 294, 557–562 (2005).

35. Brown, A. S. & Lau, F. S. A review of the epidemiology of schizophrenia. *Handbook of Behavioral Neuroscience* vol. 23, 17–30 (2016).

36. Brown, A. S. & Patterson, P. H. Maternal infection and schizophrenia: implications for prevention. *Schizophrenia bulletin* vol. 37, 284–290 (2011).

37. Susser, E., Clair, D. S. & He, L. Latent effects of prenatal malnutrition on adult health: the example of schizophrenia. *Annals of the New York Academy of Sciences* vol. 1136, 185–192 (2008).

38. Hunter, A., Murray, R., Asher, L. & Leonardi-Bee, J. The effects of tobacco smoking, and prenatal tobacco smoke exposure, on risk of schizophrenia: a systematic review and meta-analysis. *Nicotine and Tobacco Research* vol. 22, 3–10 (2020).

39. Isohanni, M. *et al.* Early developmental milestones in adult schizophrenia and other psychoses. A 31-year follow-up of the Northern Finland 1966 Birth Cohort. *Schizophrenia research* vol. 52, 1–19 (2001).

40. Fish, B., Shapiro, T., Halpern, F. & Wile, R. The prediction of schizophrenia in infancy: III. A ten-year follow-up report of neurological and psychological development. *American Journal of Psychiatry* vol. 121, 768–775 (1965).

41. Sørensen, H. J. *et al.* Early developmental milestones and risk of schizophrenia: a 45-year follow-up of the Copenhagen Perinatal Cohort. *Schizophrenia research* vol. 118, 41–47 (2010).

42. Tripathi, A., Kar, S. K. & Shukla, R. Cognitive deficits in schizophrenia: understanding the biological correlates and remediation strategies. *Clinical Psychopharmacology and Neuroscience* vol. 16, 7 (2018).

43. Alkan, E., Davies, G. & Evans, S. L. Cognitive impairment in schizophrenia: relationships with cortical thickness in fronto-temporal regions, and dissociability from symptom severity. *npj Schizophrenia* vol. 7, 20 (2021).

44. Mascio, A. *et al.* Cognitive impairments in schizophrenia: a study in a large clinical sample using natural language processing. *Frontiers in digital health* vol. 3, 711941 (2021).

45. Nielsen, P. R., Benros, M. E. & Mortensen, P. B. Hospital contacts with infection and risk of schizophrenia: a population-based cohort study with linkage of Danish national registers. *Schizophrenia bulletin* vol. 40, 1526–1532 (2014).

46. Khandaker, G. M., Zimbron, J., Dalman, C., Lewis, G. & Jones, P. B. Childhood infection and adult schizophrenia: a meta-analysis of population-based studies. *Schizophrenia research* vol. 139, 161–168 (2012).

47. Won, H. *et al.* Chromosome conformation elucidates regulatory relationships in developing human brain. *Nature* vol. 538, 523–527 (2016).

48. Hannon, E. *et al.* Methylation QTLs in the developing brain and their enrichment in schizophrenia risk loci. *Nature neuroscience* vol. 19, 48–54 (2016).

49. Jaffe, A. E. *et al.* Developmental and genetic regulation of the human cortex transcriptome in schizophrenia. *BioRxiv,* 124321 (2017).

50. Kloiber, S. *et al.* Neurodevelopmental pathways in bipolar disorder. *Neuroscience & Biobehavioral Reviews* vol. 112, 213–226 (2020).

51. Bora, E. Differences in cognitive impairment between schizophrenia and bipolar disorder: Considering the role of heterogeneity. *Psychiatry and clinical neurosciences* vol. 70, 424–433 (2016).

52. Charney, A. *et al.* Evidence for genetic heterogeneity between clinical subtypes of bipolar disorder. *Translational psychiatry* vol. 7, e993–e993 (2017).

53. Fiedorowicz, J. G. *et al.* Subthreshold hypomanic symptoms in progression from unipolar major depression to bipolar disorder. *American Journal of Psychiatry* vol. 168, 40–48 (2011).

54. Schürhoff, F. *et al.* Early and late onset bipolar disorders: two different forms of manic-depressive illness? *Journal of affective disorders* vol. 58, 215–221 (2000).

55. Cirone, C. *et al.* What do we know about the long-term course of early onset bipolar disorder? A review of the current evidence. *Brain sciences* vol. 11, 341 (2021).

56. Miller, J. N. & Black, D. W. Schizoaffective disorder: A review. *Annals of clinical psychiatry: official journal of the American Academy of Clinical Psychiatrists* vol. 31, 47–53 (2019).

57. Perkins, D. O. *et al.* Towards a psychosis risk blood diagnostic for persons experiencing high-risk symptoms: preliminary results from the NAPLS project. *Schizophrenia bulletin* vol. 41, 419–428 (2015).

58. Guidotti, A. *et al.* Toward the identification of peripheral epigenetic biomarkers of schizophrenia. *Journal of neurogenetics* vol. 28, 41–52 (2014).

59. Day, F. L. *et al.* Blunted cortisol awakening response in people at ultra high risk of developing psychosis. *Schizophrenia research* vol. 158, 25–31 (2014).

60. Uffelmann, E. *et al.* Genome-wide association studies. *Nature Reviews Methods Primers* vol. 1, 1–21 (2021).

61. Currò, D., Pugliese, D. & Armuzzi, A. Frontiers in drug research and development for inflammatory bowel disease. *Frontiers in Pharmacology* vol. 8, 400 (2017).

62. Lee, J. C. & Parkes, M. Genome-wide association studies and Crohn's disease. *Briefings in functional genomics* vol. 10, 71–76 (2011).

63. Singh, T. *et al.* Rare coding variants in ten genes confer substantial risk for schizophrenia. *Nature* vol. 604, 509–516 (2022).

64. Ding, K. & Kullo, I. J. Methods for the selection of tagging SNPs: a comparison of tagging efficiency and performance. *European Journal of Human Genetics* vol. 15, 228–236 (2007).

65. Mills, M. C. & Rahal, C. A scientometric review of genome-wide association studies. *Communications biology* vol. 2, 1–11 (2019).

66. Trubetskoy, V. *et al.* Mapping genomic loci implicates genes and synaptic biology in schizophrenia. *Nature* vol. 604, 502–508 (2022).

67. Mullins, N. *et al.* Genome-wide association study of more than 40,000 bipolar disorder cases provides new insights into the underlying biology. *Nature genetics* vol. 53, 817–829 (2021).

68. Bloomfield, M. A. P., Howes, O. D. & Large, C. H. in *Psychotic Disorders: Comprehensive Conceptualization and Treatments* (Oxford University Press, Nov. 2020).

69. Egerton, A. *et al.* Glutamate in schizophrenia: neurodevelopmental perspectives and drug development. *Schizophrenia Research* vol. 223, 59–70 (2020).

70. Brisch, R. *et al.* The role of dopamine in schizophrenia from a neurobiological and evolutionary perspective: old fashioned, but still in vogue. *Frontiers in psychiatry* vol. 5, 47 (2014).

71. Labrie, V. *et al.* Serine racemase is associated with schizophrenia susceptibility in humans and in a mouse model. *Human molecular genetics* vol. 18, 3227–3243 (2009).

72. Nakamura, T. *et al.* Loss of function mutations in ATP2A2 and psychoses: A case report and literature survey. *Psychiatry and clinical neurosciences* vol. 70, 342–350 (2016).

73. Nakajima, K. *et al.* Brain-specific heterozygous loss-of-function of ATP2A2, endoplasmic reticulum Ca2+ pump responsible for Darier's disease, causes behavioral abnormalities and a hyper-dopaminergic state. *Human molecular genetics* vol. 30, 1762–1772 (2021).

74. Obi-Nagata, K., Temma, Y. & Hayashi-Takagi, A. Synaptic functions and their disruption in schizophrenia: From clinical evidence to synaptic optogenetics in an animal model. *Proceedings of the Japan Academy, Series B* vol. 95, 179–197 (2019).

75. Sekar, A. *et al.* Schizophrenia risk from complex variation of complement component 4. *Nature* vol. 530, 177–183 (2016).

76. Druart, M. *et al.* Elevated expression of complement C4 in the mouse prefrontal cortex causes schizophrenia-associated phenotypes. *Molecular psychiatry* vol. 26, 3489–3501 (2021).

77. Wieczorek, M. *et al.* Major histocompatibility complex (MHC) class I and MHC class II proteins: conformational plasticity in antigen presentation. *Frontiers in immunology* vol. 8, 292 (2017).

78. Figueiredo, T. C. & de Oliveira, J. R. M. Reconsidering the association between the major histocompatibility complex and bipolar disorder. *Journal of Molecular Neuroscience* vol. 47, 26–30 (2012).

79. Lehner, T. The genes in the major histocompatibility complex as risk factors for schizophrenia: de omnibus dubitandum. *Biological Psychiatry* vol. 72, 615–616 (2012).

80. Werner, M. C. F. *et al.* Immune marker levels in severe mental disorders: associations with polygenic risk scores of related mental phenotypes and psoriasis. *Translational psychiatry* vol. 12, 1–8 (2022).

81. Cano, C. E. *et al.* BTN2A1, an immune checkpoint targeting Vγ9Vδ2 T cell cytotoxicity against malignant cells. *Cell Reports* vol. 36, 109359 (2021).

82. Fukuo, Y. *et al.* Serotonin 6 receptor gene and mood disorders: case–control study and meta-analysis. *Neuroscience research* vol. 67, 250–255 (2010).

83. Mowry, B. J. & Gratten, J. The emerging spectrum of allelic variation in schizophrenia: current evidence and strategies for the identification and functional characterization of common and rare variants. *Molecular psychiatry* vol. 18, 38–52 (2013).

84. Bergen, S. E. *et al.* Joint contributions of rare copy number variants and common SNPs to risk for schizophrenia. *American Journal of Psychiatry* vol. 176, 29–35 (2019).

85. Owen, M. J. & Doherty, J. L. What can we learn from the high rates of schizophrenia in people with 22q11. 2 deletion syndrome? *World Psychiatry* vol. 15, 23 (2016).

86. Viñas-Jornet, M. *et al.* A common cognitive, psychiatric, and dysmorphic phenotype in carriers of NRXN1 deletion. *Molecular genetics & genomic medicine* vol. 2, 512–521 (2014).

87. Rees, E. & Kirov, G. Copy number variation and neuropsychiatric illness. *Current Opinion in Genetics & Development* vol. 68, 57–63 (2021).

88. Charney, A. W. *et al.* Contribution of rare copy number variants to bipolar disorder risk is limited to schizoaffective cases. *Biological psychiatry* vol. 86, 110–119 (2019).

89. Backman, J. D. *et al.* Exome sequencing and analysis of 454,787 UK Biobank participants. *Nature* vol. 599, 628–634 (2021).

90. Kim, S. H. *et al.* Proband-only clinical exome sequencing for neurodevelopmental disabilities. *Pediatric Neurology* vol. 99, 47–54 (2019).

91. Palmer, D. S. *et al.* Exome sequencing in bipolar disorder reveals shared risk gene AKAP11 with schizophrenia. *medRxiv* (2021).

92. Tanji, C. *et al.* A-kinase anchoring protein AKAP220 binds to glycogen synthase kinase-3β (GSK-3β) and mediates protein kinase A-dependent inhibition of GSK-3β. *Journal of Biological Chemistry* vol. 277, 36955–36961 (2002).

93. Jope, R. S. Lithium and GSK-3: one inhibitor, two inhibitory actions, multiple outcomes. *Trends in pharmacological sciences* vol. 24, 441–443 (2003).

94. Freland, L. & Beaulieu, J.-M. Inhibition of GSK3 by lithium, from single molecules to signaling networks. *Frontiers in molecular neuroscience* vol. 5, 14 (2012).

95. Kishore, B. K. & Ecelbarger, C. M. Lithium: a versatile tool for understanding renal physiology. *American Journal of Physiology-Renal Physiology* vol. 304, F1139–F1149 (2013).

96. Gibson, G. Rare and common variants: twenty arguments. *Nature Reviews Genetics* vol. 13, 135–145 (2012).

97. Ellingford, J. M. *et al.* Recommendations for clinical interpretation of variants found in non-coding regions of the genome. *Genome medicine* vol. 14, 1–19 (2022).

98. Võsa, U. *et al.* Large-scale cis-and trans-eQTL analyses identify thousands of genetic loci and polygenic scores that regulate blood gene expression. *Nature genetics* vol. 53, 1300–1310 (2021).

99. Nowakowska, B. Clinical interpretation of copy number variants in the human genome. *Journal of applied genetics* vol. 58, 449–457 (2017).

100. DeBoever, C. *et al.* Medical relevance of protein-truncating variants across 337,205 individuals in the UK Biobank study. *Nature communications* vol. 9, 1–10 (2018).

101. Moghaddam, B. & Javitt, D. From revolution to evolution: the glutamate hypothesis of schizophrenia and its implication for treatment. *Neuropsychopharmacology* vol. 37, 4–15 (2012).

102. Uno, Y. & Coyle, J. T. Glutamate hypothesis in schizophrenia. *Psychiatry and clinical neurosciences* vol. 73, 204–215 (2019).

103. Tsai, G. & Coyle, J. T. Glutamatergic mechanisms in schizophrenia. *Annual review of pharmacology and toxicology* vol. 42, 165 (2002).

104. Javitt, D. C. Glutamate and schizophrenia: phencyclidine, N-methyl-d-aspartate receptors, and dopamine–glutamate interactions. *International review of neurobiology* vol. 78, 69–108 (2007).

105. Carlsson, M. & Carlsson, A. The NMDA antagonist MK-801 causes marked locomotor stimulation in monoamine-depleted mice. *Journal of neural transmission* vol. 75, 221–226 (1989).

106. Zhou, X. *et al.* Reduced NMDAR1 expression in the Sp4 hypomorphic mouse may contribute to endophenotypes of human psychiatric disorders. *Human molecular genetics* vol. 19, 3797–3805 (2010).

107. Mukai, J. *et al.* Recapitulation and reversal of schizophrenia-related phenotypes in Setd1a-deficient mice. *Neuron* vol. 104, 471–487 (2019).

108. Sadybekov, A., Tian, C., Arnesano, C., Katritch, V. & Herring, B. E. An autism spectrum disorder-related de novo mutation hotspot discovered in the GEF1 domain of Trio. *Nature communications* vol. 8, 1–13 (2017).

109. Kim, M. *et al.* Brain gene co-expression networks link complement signaling with convergent synaptic pathology in schizophrenia. *Nature neuroscience* vol. 24, 799–809 (2021).

110. Dent, E. W., Gupton, S. L. & Gertler, F. B. The growth cone cytoskeleton in axon outgrowth and guidance. *Cold Spring Harbor perspectives in biology* vol. 3, a001800 (2011).

111. Liebl, E. C. *et al.* Dosage-sensitive, reciprocal genetic interactions between the Abl tyrosine kinase and the putative GEF trio reveal trio's role in axon pathfinding. *Neuron* vol. 26, 107–118 (2000).

112. Van Haren, J. *et al.* Dynamic microtubules catalyze formation of navigator-TRIO complexes to regulate neurite extension. *Current Biology* vol. 24, 1778–1785 (2014).

113. Moyer, C. E., Shelton, M. A. & Sweet, R. A. Dendritic spine alterations in schizophrenia. *Neuroscience letters* vol. 601, 46–53 (2015).

114. Glausier, J. R. & Lewis, D. A. Dendritic spine pathology in schizophrenia. *Neuroscience* vol. 251, 90–107 (2013).

115. Swatek, K. N. & Komander, D. Ubiquitin modifications. *Cell research* vol. 26, 399–422 (2016).

116. Zheng, N. & Shabek, N. Ubiquitin ligases: structure, function, and regulation. *Annual review of biochemistry* vol. 86, 129–157 (2017).

117. He, M. *et al.* The emerging role of deubiquitinating enzymes in genomic integrity, diseases, and therapeutics. *Cell & bioscience* vol. 6, 1–15 (2016).

118. Cruz Walma, D. A., Chen, Z., Bullock, A. N. & Yamada, K. M. Ubiquitin ligases: guardians of mammalian development. *Nature Reviews Molecular Cell Biology* vol. 23, 350–367 (2022).

119. Schwartz, A. L. & Ciechanover, A. Targeting proteins for destruction by the ubiquitin system: implications for human pathobiology. *Annual review of pharmacology and toxicology* vol. 49, 73–96 (2009).

120. Bousman, C. A. *et al.* Elevated ubiquitinated proteins in brain and blood of individuals with schizophrenia. *Scientific reports* vol. 9, 1–8 (2019).

121. Speese, S. D., Trotta, N., Rodesch, C. K., Aravamudan, B. & Broadie, K. The ubiquitin proteasome system acutely regulates presynaptic protein turnover and synaptic efficacy. *Current biology* vol. 13, 899–910 (2003).

122. Çetin, G., Klafack, S., Studencka-Turski, M., Krüger, E. & Ebstein, F. The Ubiquitin–Proteasome System in Immune Cells. *Biomolecules* vol. 11, 60 (2021).

123. Andrade, A. *et al.* Genetic associations between voltage-gated calcium channels and psychiatric disorders. *International journal of molecular sciences* vol. 20, 3537 (2019).

124. Andrade, A. *et al.* A rare schizophrenia risk variant of CACNA1I disrupts CaV3. 3 channel activity. *Scientific reports* vol. 6, 1–13 (2016).

125. Baez-Nieto, D. *et al.* Analysing an allelic series of rare missense variants of CACNA1I in a Swedish schizophrenia cohort. *Brain* vol. 145, 1839–1853 (2022).

126. Xu, H. & Van Remmen, H. The SarcoEndoplasmic Reticulum Calcium ATPase (SERCA) pump: A potential target for intervention in aging and skeletal muscle pathologies. *Skeletal Muscle* vol. 11, 1–9 (2021).

127. Singh, N., Bartol, T., Levine, H., Sejnowski, T. & Nadkarni, S. Presynaptic endoplasmic reticulum regulates short-term plasticity in hippocampal synapses. *Communications biology* vol. 4, 1–13 (2021).

128. Lum, J. S., Millard, S. J., Huang, X.-F., Ooi, L. & Newell, K. A. A postmortem analysis of NMDA ionotropic and group 1 metabotropic glutamate receptors in the nucleus accumbens in schizophrenia. *Journal of Psychiatry and Neuroscience* vol. 43, 102–110 (2018).

129. Powchik, P. *et al.* Postmortem studies in schizophrenia. *Schizophrenia bulletin* vol. 24, 325–341 (1998).

130. Conn, K.-A., Burne, T. H. & Kesby, J. P. Subcortical dopamine and cognition in schizophrenia: Looking beyond psychosis in preclinical models. *Frontiers in Neuroscience* vol. 14, 542 (2020).

131. Martel, J. C. & Gatti McArthur, S. Dopamine receptor subtypes, physiology and pharmacology: new ligands and concepts in schizophrenia. *Frontiers in pharmacology* vol. 11, 1003 (2020).

132. Kraguljac, N. V. *et al.* Neuroimaging biomarkers in schizophrenia. *American Journal of Psychiatry* vol. 178, 509–521 (2021).

133. Laruelle, M., Frankle, W. G., Narendran, R., Kegeles, L. S. & Abi-Dargham, A. Mechanism of action of antipsychotic drugs: from dopamine D2 receptor antagonism to glutamate NMDA facilitation. *Clinical therapeutics* vol. 27, S16–S24 (2005).

134. Lally, J. & MacCabe, J. H. Antipsychotic medication in schizophrenia: a review. *British medical bulletin* vol. 114, 169–179 (2015).

135. Huhn, M. *et al.* Comparative efficacy and tolerability of 32 oral antipsychotics for the acute treatment of adults with multi-episode schizophrenia: a systematic review and network meta-analysis. *The Lancet* vol. 394, 939–951 (2019).

136. Kantrowitz, J. T. *et al.* Proof of mechanism and target engagement of glutamatergic drugs for the treatment of schizophrenia: RCTs of pomaglumetad and TS-134 on ketamine-induced psychotic symptoms and pharmacoBOLD in healthy volunteers. *Neuropsychopharmacology* vol. 45, 1842–1850 (2020).

137. Mueller, T. M. & Meador-Woodruff, J. H. Post-translational protein modifications in schizophrenia. *npj Schizophrenia* vol. 6, 1–16 (2020).

138. Kim, H. K. & Andreazza, A. C. The relationship between oxidative stress and post-translational modification of the dopamine transporter in bipolar disorder. *Expert Review of Neurotherapeutics* vol. 12, 849–859 (2012).

139. Conibear, A. C. Deciphering protein post-translational modifications using chemical biology tools. *Nature Reviews Chemistry* vol. 4, 674–695 (2020).

140. Ramazi, S. & Zahiri, J. Post-translational modifications in proteins: resources, tools and prediction methods. *Database* vol. 2021 (2021).

141. Reily, C., Stewart, T. J., Renfrow, M. B. & Novak, J. Glycosylation in health and disease. *Nature Reviews Nephrology* vol. 15, 346–366 (2019).

142. Chen, B., Sun, Y., Niu, J., Jarugumilli, G. K. & Wu, X. Protein lipidation in cell signaling and diseases: function, regulation, and therapeutic opportunities. *Cell chemical biology* vol. 25, 817–831 (2018).

143. Williams, S. E., Mealer, R. G., Scolnick, E. M., Smoller, J. W. & Cummings, R. D. Aberrant glycosylation in schizophrenia: a review of 25 years of post-mortem brain studies. *Molecular psychiatry* vol. 25, 3198–3207 (2020).

144. Vila-Rodriguez, F., Honer, W. G., Innis, S. M., Wellington, C. L. & Beasley, C. L. ApoE and cholesterol in schizophrenia and bipolar disorder: comparison of grey and white matter and relation with APOE genotype. *Journal of Psychiatry and Neuroscience* vol. 36, 47–55 (2011).

145. Huang, Y.-J. *et al.* State-dependent alterations of lipid profiles in patients with bipolar disorder. *The International Journal of Psychiatry in Medicine* vol. 53, 273–281 (2018).

146. Carter, C. Multiple genes and factors associated with bipolar disorder converge on growth factor and stress activated kinase pathways controlling translation initiation: implications for oligodendrocyte viability. *Neurochemistry international* vol. 50, 461–490 (2007).

147. Aksu, M. *et al.* Xpo7 is a broad-spectrum exportin and a nuclear import receptor. *Journal of Cell Biology* vol. 217, 2329–2340 (2018).

148. Chano, T. *et al.* Identification of RB1CC1, a novel human gene that can induce RB1 in various human cells. *Oncogene* vol. 21, 1295–1298 (2002).

149. Chano, T., Kontani, K., Teramoto, K., Okabe, H. & Ikegawa, S. Truncating mutations of RB1CC1 in human breast cancer. *Nature genetics* vol. 31, 285–288 (2002).

150. Chano, T. *et al.* RB1CC1 activates RB1 pathway and inhibits proliferation and cologenic survival in human cancer. *PLoS One* vol. 5, e11404 (2010).

151. Owen, M. J. & Williams, N. M. Explaining the missing heritability of psychiatric disorders. *World Psychiatry* vol. 20, 294 (2021).

152. Stewart, J. B. & Chinnery, P. F. Extreme heterogeneity of human mitochondrial DNA from organelles to populations. *Nature Reviews Genetics* vol. 22, 106–118 (2021).

153. Bodenstein, D. *et al.* Mitochondrial DNA content and oxidation in bipolar disorder and its role across brain regions. *npj Schizophrenia* vol. 5, 1–8 (2019).

154. Chestkov, I. *et al.* ROS-induced DNA damage associates with abundance of mitochondrial DNA in white blood cells of the untreated schizophrenic patients. *Oxidative Medicine and Cellular Longevity* vol. 2018 (2018).

155. Daniels, T. E., Olsen, E. M. & Tyrka, A. R. Stress and psychiatric disorders: The role of mitochondria. *Annual review of clinical psychology* vol. 16, 165–186 (2020).

156. Puntambekar, S. *et al.* Evolutionary divergence of novel open reading frames in cichlids speciation. *Scientific reports* vol. 10, 1–18 (2020).

157. Neville, M. D. *et al.* A platform for curated products from novel open reading frames prompts reinterpretation of disease variants. *Genome research* vol. 31, 327–336 (2021).

158. Erady, C. *et al.* Pan-cancer analysis of transcripts encoding novel open-reading frames (nORFs) and their potential biological functions. *NPJ genomic medicine* vol. 6, 1–17 (2021).

159. Erady, C. *et al.* Novel open reading frames in human accelerated regions and transposable elements reveal new leads to understand schizophrenia and bipolar disorder. *Molecular psychiatry* vol. 27, 1455–1468 (2022).

160. Pollard, K. S. *et al.* An RNA gene expressed during cortical development evolved rapidly in humans. *Nature* vol. 443, 167–172 (2006).

161. Pollard, K. S. *et al.* Forces shaping the fastest evolving regions in the human genome. *PLoS genetics* vol. 2, e168 (2006).

162. Levchenko, A., Kanapin, A., Samsonova, A. & Gainetdinov, R. R. Human accelerated regions and other human-specific sequence variations in the context of evolution and their relevance for brain development. *Genome biology and evolution* vol. 10, 166–188 (2018).

163. Ryu, H. *et al.* Massively parallel dissection of human accelerated regions in human and chimpanzee neural progenitors. *BioRxiv,* 256313 (2018).

164. Capra, J. A., Erwin, G. D., McKinsey, G., Rubenstein, J. L. & Pollard, K. S. Many human accelerated regions are developmental enhancers. *Philosophical Transactions of the Royal Society B: Biological Sciences* vol. 368, 20130025 (2013).

165. Cheung, J. P., Tubbs, J. D. & Sham, P. C. Extended gene set analysis of human neuro-psychiatric traits shows enrichment in brain-expressed human accelerated regions across development. *Schizophrenia Research* vol. 246, 148–155 (2022).

166. Crespi, B., Summers, K. & Dorus, S. Adaptive evolution of genes underlying schizophrenia. *Proceedings of the Royal Society B: Biological Sciences* vol. 274, 2801–2810 (2007).

167. Preuss, T. M. Human brain evolution: from gene discovery to phenotype discovery. *Proceedings of the National Academy of Sciences* vol. 109, 10709–10716 (2012).

168. Doan, R. N. *et al.* Mutations in human accelerated regions disrupt cognition and social behavior. *Cell* vol. 167, 341–354 (2016).

169. Madsen, E. S. & Wu, Y. Low R&D efficiency in large pharmaceutical companies. *American Journal of Medical Research* vol. 3, 141–141 (2016).

170. Schuhmacher, A. *et al.* R&D efficiency of leading pharmaceutical companies–A 20-year analysis. *Drug discovery today* vol. 26, 1784–1789 (2021).

171. Scannell, J. W., Blanckley, A., Boldon, H. & Warrington, B. Diagnosing the decline in pharmaceutical R&D efficiency. *Nature reviews Drug discovery* vol. 11, 191–200 (2012).

172. Freeman, M. W. & Dervan, A. P. The path from bench to bedside: considerations before starting the journey. *Journal of Investigative Medicine* vol. 59, 746–751 (2011).

173. Seyhan, A. A. Lost in translation: the valley of death across preclinical and clinical divide–identification of problems and overcoming obstacles. *Translational Medicine Communications* vol. 4, 1–19 (2019).

174. Schlander, M., Hernandez-Villafuerte, K., Cheng, C.-Y., Mestre-Ferrandiz, J. & Baumann, M. How much does it cost to research and develop a new drug? A systematic review and assessment. *PharmacoEconomics* vol. 39, 1243–1269 (2021).

175. Boyer, S., Brealey, C. & Davis, A. M. Attrition in drug discovery and development. *Attrition in the Pharmaceutical Industry: Reasons, Implications, and Pathways Forward,* 5–45 (2015).

176. Pangalos, M. N., Schechter, L. E. & Hurko, O. Drug development for CNS disorders: strategies for balancing risk and reducing attrition. *Nature Reviews Drug Discovery* vol. 6, 521–532 (2007).

177. Thomas, D. W. *et al.* Clinical development success rates 2006–2015. *BIO Industry Analysis* vol. 1, 25 (2016).

178. Drews, J. Drug discovery: a historical perspective. *science* vol. 287, 1960–1964 (2000).

179. Wang, S., Sim, T. B., Kim, Y.-S. & Chang, Y.-T. Tools for target identification and validation. *Current opinion in chemical biology* vol. 8, 371–377 (2004).

180. Schenone, M., Dancik, V., Wagner, B. K. & Clemons, P. A. Target identification and mechanism of action in chemical biology and drug discovery. *Nature chemical biology* vol. 9, 232–240 (2013).

181. Macarron, R. *et al.* Impact of high-throughput screening in biomedical research. *Nature reviews Drug discovery* vol. 10, 188–195 (2011).

182. Caraus, I., Alsuwailem, A. A., Nadon, R. & Makarenkov, V. Detecting and overcoming systematic bias in high-throughput screening technologies: a comprehensive review of practical issues and methodological solutions. *Briefings in bioinformatics* vol. 16, 974–986 (2015).

183. Iskandar, S. E., Haberman, V. A. & Bowers, A. A. Expanding the chemical diversity of genetically encoded libraries. *ACS combinatorial science* vol. 22, 712–733 (2020).

184. Akella, L. B. & DeCaprio, D. Cheminformatics approaches to analyze diversity in compound screening libraries. *Current Opinion in Chemical Biology* vol. 14, 325–330 (2010).

185. Wawer, M. J. *et al.* Toward performance-diverse small-molecule libraries for cell-based phenotypic screening using multiplexed high-dimensional profiling. *Proceedings of the National Academy of Sciences* vol. 111, 10911–10916 (2014).

186. Rishton, G. M. Reactive compounds and in vitro false positives in HTS. *Drug discovery today* vol. 2, 382–384 (1997).

187. Bibette, J. Gaining confidence in high-throughput screening. *Proceedings of the National Academy of Sciences* vol. 109, 649–650 (2012).

188. David, L. *et al.* Identification of Compounds That Interfere with High-Throughput Screening Assay Technologies. *ChemMedChem* vol. 14, 1795–1802 (2019).

189. Sun, D., Gao, W., Hu, H. & Zhou, S. Why 90% of clinical drug development fails and how to improve it? *Acta Pharmaceutica Sinica B* (2022).

190. Friedman, L. M., Furberg, C. D., DeMets, D. L., Reboussin, D. M. & Granger, C. B. *Fundamentals of clinical trials* (Springer, 2015).

191. Ciociola, A. A. *et al.* How drugs are developed and approved by the FDA: current process and future directions. *Official journal of the American College of Gastroenterology| ACG* vol. 109, 620–623 (2014).

192. Motola, D., De Ponti, F., Rossi, P., Martini, N. & Montanaro, N. Therapeutic innovation in the European Union: analysis of the drugs approved by the EMEA between 1995 and 2003. *British journal of clinical pharmacology* vol. 59, 475–478 (2005).

193. Ono, S. Ministry of Health, Labour and Welfare (MHLW, Japan). *Wiley encyclopedia of clinical trials,* 1–8 (2007).

194. Kaplon, H., Muralidharan, M., Schneider, Z. & Reichert, J. M. *Antibodies to watch in 2020* in *MAbs* **12** (2020), 1703531.

195. Drucker, D. J. Advances in oral peptide therapeutics. *Nature reviews Drug discovery* vol. 19, 277–289 (2020).

196. Vinogradov, A. A., Yin, Y. & Suga, H. Macrocyclic peptides as drug candidates: recent progress and remaining challenges. *Journal of the American Chemical Society* vol. 141, 4167–4181 (2019).

197. Schapira, M., Calabrese, M. F., Bullock, A. N. & Crews, C. M. Targeted protein degradation: expanding the toolbox. *Nature reviews Drug discovery* vol. 18, 949–963 (2019).

198. Chamberlain, P. P. & Hamann, L. G. Development of targeted protein degradation therapeutics. *Nature chemical biology* vol. 15, 937–944 (2019).

199. Paiva, S.-L. & Crews, C. M. Targeted protein degradation: elements of PROTAC design. *Current opinion in chemical biology* vol. 50, 111–119 (2019).

200. Li, H. *et al.* Applications of genome editing technology in the targeted therapy of human diseases: mechanisms, advances and prospects. *Signal transduction and targeted therapy* vol. 5, 1–23 (2020).

201. Raguram, A., Banskota, S. & Liu, D. R. Therapeutic in vivo delivery of gene editing agents. *Cell* (2022).

202. Blanco, M.-J. & Gardinier, K. M. *New chemical modalities and strategic thinking in early drug discovery* 2020.

203. Müller, C. E., Hansen, F. K., Gütschow, M., Lindsley, C. W. & Liotta, D. *New Drug Modalities in Medicinal Chemistry, Pharmacology, and Translational Science: Joint Virtual Special Issue by Journal of Medicinal Chemistry, ACS Medicinal Chemistry Letters, and ACS Pharmacology & Translational Science* 2021.

204. Ravikumar, B. & Aittokallio, T. Improving the efficacy-safety balance of polypharmacology in multi-target drug discovery. *Expert opinion on drug discovery* vol. 13, 179–192 (2018).

205. Bolognesi, M. L. *Harnessing polypharmacology with medicinal chemistry* 2019.

206. Chaudhari, R., Fong, L. W., Tan, Z., Huang, B. & Zhang, S. An up-to-date overview of computational polypharmacology in modern drug discovery. *Expert opinion on drug discovery* vol. 15, 1025–1044 (2020).

207. Schneider, P. *et al.* Rethinking drug design in the artificial intelligence era. *Nature Reviews Drug Discovery* vol. 19, 353–364 (2020).

208. Mikitsh, J. L. & Chacko, A.-M. Pathways for small molecule delivery to the central nervous system across the blood-brain barrier. *Perspectives in medicinal chemistry* vol. 6, PMC–S13384 (2014).

209. Xiong, B. *et al.* Strategies for structural modification of small molecules to improve blood–brain barrier penetration: A recent perspective. *Journal of Medicinal Chemistry* vol. 64, 13152–13173 (2021).

210. Eichler, H.-G. *et al.* The risks of risk aversion in drug regulation. *Nature reviews Drug discovery* vol. 12, 907–916 (2013).

211. Zhu, T. Challenges of Psychiatry Drug Development and the Role of Human Pharmacology Models in Early Development—A Drug Developer's Perspective. *Frontiers in Psychiatry* vol. 11, 562660 (2021).

212. He, Q. *et al.* Towards improvements for penetrating the blood–brain barrier—recent progress from a material and pharmaceutical perspective. *Cells* vol. 7, 24 (2018).

213. Kearney, M. A. C. T. E. & Olson, K. R. Comparing acute toxicity of first-and second-generation antipsychotic drugs: a 10-year, retrospective cohort study. *The Journal of Clinical Psychiatry* vol. 70, 3439 (2009).

214. Minns, A. B. & Clark, R. F. Toxicology and overdose of atypical antipsychotics. *The Journal of emergency medicine* vol. 43, 906–913 (2012).

215. Kaar, S. J., Natesan, S., Mccutcheon, R. & Howes, O. D. Antipsychotics: mechanisms underlying clinical response and side-effects and novel treatment approaches based on pathophysiology. *Neuropharmacology* vol. 172, 107704 (2020).

216. Acosta, F. J., Hernandez, J., Pereira, J., Herrera, J. & Rodriguez, C. J. Medication adherence in schizophrenia. *World journal of psychiatry* vol. 2, 74 (2012).

217. Lavretsky, H. History of schizophrenia as a psychiatric disorder. *Clinical handbook of schizophrenia* vol. 1 (2008).

218. Shen, W. W. A history of antipsychotic drug development. *Comprehensive psychiatry* vol. 40, 407–414 (1999).

219. Ban, T. A. Fifty years chlorpromazine: a historical perspective. *Neuropsychiatric disease and treatment* vol. 3, 495 (2007).

220. López-Muñoz, F. *et al.* History of the discovery and clinical introduction of chlorpromazine. *Annals of Clinical Psychiatry* vol. 17, 113–135 (2005).

221. Rosenbloom, M. Chlorpromazine and the psychopharmacologic revolution. *Jama* vol. 287, 1860–1861 (2002).

222. Ramachandraiah, C. T., Subramaniam, N. & Tancer, M. The story of antipsychotics: Past and present. *Indian journal of psychiatry* vol. 51, 324 (2009).

223. Creese, I., Burt, D. R. & Snyder, S. H. Dopamine receptor binding predicts clinical and pharmacological potencies of antischizophrenic drugs. *Science* vol. 192, 481–483 (1976).

224. Seeman, P., Lee, T., Chau-Wong, M. & Wong, K. Antipsychotic drug doses and neuroleptic/dopamine receptors. *Nature* vol. 261, 717–719 (1976).

225. Janowsky, D. S. & Risch, C. Amphetamine psychosis and psychotic symptoms. *Psychopharmacology* vol. 65, 73–77 (1979).

226. Nord, M. & Farde, L. Antipsychotic occupancy of dopamine receptors in schizophrenia. *CNS neuroscience & therapeutics* vol. 17, 97–103 (2011).

227. Farde, L. *et al.* Positron emission tomographic analysis of central D1 and D2 dopamine receptor occupancy in patients treated with classical neuroleptics and clozapine: relation to extrapyramidal side effects. *Archives of general psychiatry* vol. 49, 538–544 (1992).

228. Spark, D. L., Fornito, A., Langmead, C. J. & Stewart, G. D. Beyond antipsychotics: a twenty-first century update for preclinical development of schizophrenia therapeutics. *Translational Psychiatry* vol. 12, 1–11 (2022).

229. Matar, H. E., Almerie, M. Q., Makhoul, S., Xia, J. & Humphreys, P. Pericyazine for schizophrenia. *Cochrane Database of Systematic Reviews* (2014).

230. Fitton, A. & Heel, R. C. Clozapine. *Drugs* vol. 40, 722–747 (1990).

231. Minassian, A. & Young, J. W. Evaluation of the clinical efficacy of asenapine in schizophrenia. *Expert opinion on pharmacotherapy* vol. 11, 2107–2115 (2010).

232. Meltzer, H. Y. An overview of the mechanism of action of clozapine. *The journal of clinical psychiatry* (1994).

233. Hippius, H. The history of clozapine. *Psychopharmacology* (1989).

234. Lally, J., Gaughran, F., Timms, P. & Curran, S. R. Treatment-resistant schizophrenia: current insights on the pharmacogenomics of antipsychotics. *Pharmacogenomics and personalized medicine* vol. 9, 117 (2016).

235. Gammon, D., Cheng, C., Volkovinskaia, A., Baker, G. B. & Dursun, S. M. Clozapine: Why is it so uniquely effective in the treatment of a range of neuropsychiatric disorders? *Biomolecules* vol. 11, 1030 (2021).

236. Goldstein, M. E., Anderson, V. M., Pillai, A., Kydd, R. R. & Russell, B. R. Glutamatergic neurometabolites in clozapine-responsive and-resistant schizophrenia. *International Journal of Neuropsychopharmacology* vol. 18 (2015).

237. Sacchi, S. *et al.* Olanzapine, but not clozapine, increases glutamate release in the prefrontal cortex of freely moving mice by inhibiting D-aspartate oxidase activity. *Scientific reports* vol. 7, 1–13 (2017).

238. Fukuyama, K., Kato, R., Murata, M., Shiroyama, T. & Okada, M. Clozapine normalizes a glutamatergic transmission abnormality induced by an impaired NMDA receptor in the thalamocortical pathway via the activation of a group III metabotropic glutamate receptor. *Biomolecules* vol. 9, 234 (2019).

239. McQueen, G. *et al.* Changes in brain glutamate on switching to clozapine in treatment-resistant schizophrenia. *Schizophrenia bulletin* vol. 47, 662–671 (2021).

240. Hribkova, H. *et al.* Clozapine Reverses Dysfunction of Glutamatergic Neurons Derived From Clozapine-Responsive Schizophrenia Patients. *Frontiers in cellular neuroscience* vol. 16 (2022).

241. Blair, D. & Dauner, A. Extrapyramidal symptoms are serious side-effects of antipsychotic and other drugs. *The Nurse Practitioner* vol. 17, 56–62 (1992).

242. Yuen, J. W. *et al.* Clozapine-induced cardiovascular side effects and autonomic dysfunction: a systematic review. *Frontiers in neuroscience* vol. 12, 203 (2018).

243. DE, L. *et al.* Clozapine-induced agranulocytosis. A genetic and epidemiologic study. (1977).

244. Lahdelma, L. & Appelberg, B. Clozapine-induced agranulocytosis in Finland, 1982-2007: long-term monitoring of patients is still warranted. *The Journal of clinical psychiatry* vol. 73, 3029 (2012).

245. Aichhorn, W., Whitworth, A. B., Weiss, E. M. & Marksteiner, J. Second-generation antipsychotics. *Drug safety* vol. 29, 587–598 (2006).

246. Davis, J. M., Chen, N. & Glick, I. D. A meta-analysis of the efficacy of second-generation antipsychotics. *Archives of general psychiatry* vol. 60, 553–564 (2003).

247. Divac, N., Prostran, M., Jakovcevski, I. & Cerovac, N. Second-generation antipsychotics and extrapyramidal adverse effects. *BioMed research international* vol. 2014 (2014).

248. Hirsch, L. *et al.* Second-generation antipsychotics and metabolic side effects: a systematic review of population-based studies. *Drug safety* vol. 40, 771–781 (2017).

249. Sykes, D. A. *et al.* Extrapyramidal side effects of antipsychotics are linked to their association kinetics at dopamine D2 receptors. *Nature communications* vol. 8, 1–11 (2017).

250. Siafis, S., Tzachanis, D., Samara, M. & Papazisis, G. Antipsychotic drugs: from receptor-binding profiles to metabolic side effects. *Current neuropharmacology* vol. 16, 1210–1223 (2018).

251. Cunningham Owens, D. & Johnstone, E. C. The development of antipsychotic drugs. *Brain and Neuroscience Advances* vol. 2, 2398212818817498 (2018).

252. Brady, L. S., Potter, W. Z. & Gordon, J. A. Redirecting the revolution: new developments in drug development for psychiatry. *Expert opinion on drug discovery* vol. 14, 1213–1219 (2019).

253. Tricklebank, M. D., Robbins, T. W., Simmons, C. & Wong, E. H. Time to re-engage psychiatric drug discovery by strengthening confidence in preclinical psychopharmacology. *Psychopharmacology* vol. 238, 1417–1436 (2021).

254. Bugarski-Kirola, D., Wang, A., Abi-Saab, D. & Blättler, T. A phase II/III trial of bitopertin monotherapy compared with placebo in patients with an acute exacerbation of schizophrenia–results from the CandleLyte study. *European Neuropsychopharmacology* vol. 24, 1024–1036 (2014).

255. Mäki-Marttunen, T. *et al.* Biophysical psychiatry—how computational neuroscience can help to understand the complex mechanisms of mental disorders. *Frontiers in psychiatry* vol. 10, 534 (2019).

256. Lin, S.-K. Antipsychotic polypharmacy: a dirty little secret or a fashion? *International Journal of Neuropsychopharmacology* vol. 23, 125–131 (2020).

257. Azorin, J.-M. & Simon, N. Antipsychotic polypharmacy in schizophrenia: evolving evidence and rationale. *Expert Opinion on Drug Metabolism & Toxicology* vol. 16, 1175–1186 (2020).

258. Smalley, E. AI-powered drug discovery captures pharma interest. *Nature Biotechnology* vol. 35, 604–605 (July 2017).

259. Probst, D. Growing pains: Reacting to negative impacts of deep learning on machine learning for chemistry (2022).

260. Huang, D. *et al.* A Unified System for Molecular Property Predictions: Oloren ChemEngine and its Applications (2022).

261. Naga, D., Muster, W., Musvasva, E. & Ecker, G. F. Off-targetP ML: an open source machine learning framework for off-target panel safety assessment of small molecules. *Journal of cheminformatics* vol. 14, 1–19 (2022).

262. Gao, W. & Coley, C. W. The synthesizability of molecules proposed by generative models. *Journal of chemical information and modeling* vol. 60, 5714–5723 (2020).

263. Liu, C.-H. *et al.* RetroGNN: Approximating retrosynthesis by graph neural networks for de novo drug design. *arXiv preprint arXiv:2011.13042* (2020).

264. Thakkar, A., Chadimová, V., Bjerrum, E. J., Engkvist, O. & Reymond, J.-L. Retrosynthetic accessibility score (RAscore)–rapid machine learned synthesizability classification from AI driven retrosynthetic planning. *Chemical science* vol. 12, 3339–3349 (2021).

265. Ghasemi, F., Mehridehnavi, A., Pérez-Garrido, A. & Pérez-Sánchez, H. Neural network and deep-learning algorithms used in QSAR studies: merits and drawbacks. *Drug discovery today* vol. 23, 1784–1790 (2018).

266. Zakharov, A. V. *et al.* Novel consensus architecture to improve performance of large-scale multitask deep learning QSAR models. *Journal of chemical information and modeling* vol. 59, 4613–4624 (2019).

267. Doak, B. C., Over, B., Giordanetto, F. & Kihlberg, J. Oral druggable space beyond the rule of 5: insights from drugs and clinical candidates. *Chemistry & biology* vol. 21, 1115–1142 (2014).

268. Benet, L. Z., Hosey, C. M., Ursu, O. & Oprea, T. I. BDDCS, the Rule of 5 and drugability. *Advanced drug delivery reviews* vol. 101, 89–98 (2016).

269. Bergström, C. A. & Porter, C. J. Understanding the challenge of beyond-rule-of-5 compounds. *Adv Drug Deliv Rev* vol. 101, 1–5 (2016).

270. Lipinski, C. A., Lombardo, F., Dominy, B. W. & Feeney, P. J. Experimental and computational approaches to estimate solubility and permeability in drug discovery and development settings. *Advanced drug delivery reviews* vol. 23, 3–25 (1997).

271. Pollastri, M. P. Overview on the Rule of Five. *Current protocols in pharmacology* vol. 49, 9–12 (2010).

272. Sangster, J. Octanol-water partition coefficients of simple organic compounds. *Journal of Physical and Chemical Reference Data* vol. 18, 1111–1229 (1989).

273. Sangster, J. M. *Octanol-water partition coefficients: fundamentals and physical chemistry* (John Wiley & Sons, 1997).

274. Blomme, E. A. & Will, Y. Toxicology strategies for drug discovery: present and future. *Chemical research in toxicology* vol. 29, 473–504 (2016).

275. Buxton, I. L. Pharmacokinetics and pharmacodynamics. *Goodman & Gilman's The Pharmacological Basis of Therapeutics, 11th ed. New York: McGraw-Hill,* 1–39 (2006).

276. Jang, G. R., Harris, R. Z. & Lau, D. T. Pharmacokinetics and its role in small molecule drug discovery research. *Medicinal research reviews* vol. 21, 382–396 (2001).

277. Woodward, P. W., Williams, C., Sewing, A. & Benson, N. Improving the design and analysis of high-throughput screening technology comparison experiments using statistical modeling. *SLAS Discovery* vol. 11, 5–12 (2006).

278. Mandal, S., Mandal, S. K., *et al.* Rational drug design. *European journal of pharmacology* vol. 625, 90–100 (2009).

279. Mavromoustakos, T. *et al.* Strategies in the rational drug design. *Current medicinal chemistry* vol. 18, 2517–2530 (2011).

280. Zhu, H. Big data and artificial intelligence modeling for drug discovery. *Annual review of pharmacology and toxicology* vol. 60, 573 (2020).

281. Copeland, R. A., Harpel, M. R. & Tummino, P. J. Targeting enzyme inhibitors in drug discovery. *Expert opinion on therapeutic targets* vol. 11, 967–978 (2007).

282. Wang, Y., Yu, Z., Xiao, W., Lu, S. & Zhang, J. Allosteric binding sites at the receptor–lipid bilayer interface: novel targets for GPCR drug discovery. *Drug discovery today* vol. 26, 690–703 (2021).

283. Lu, S., Shen, Q. & Zhang, J. Allosteric methods and their applications: facilitating the discovery of allosteric drugs and the investigation of allosteric mechanisms. *Accounts of Chemical Research* vol. 52, 492–500 (2019).

284. Eisenstein, M. Active machine learning helps drug hunters tackle biology. *Nature biotechnology* vol. 38, 512–515 (2020).

285. Anderson, E. *et al. SMILES, a line notation and computerized interpreter for chemical structures* (US Environmental Protection Agency, Environmental Research Laboratory, 1987).

286. Weininger, D. SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. *Journal of chemical information and computer sciences* vol. 28, 31–36 (1988).

287. Sayle, R. *1st-class SMARTS patterns* in *EuroMUG 97* (1997).

288. Ehmki, E. S., Schmidt, R., Ohm, F. & Rarey, M. Comparing molecular patterns using the example of SMARTS: Applications and filter collection analysis. *Journal of Chemical Information and Modeling* vol. 59, 2572–2586 (2019).

289. Krenn, M., Hase, F., Nigam, A., Friederich, P. & Aspuru-Guzik, A. Self-referencing embedded strings (SELFIES): A 100% robust molecular string representation. *Machine Learning: Science and Technology* vol. 1, 045024 (2020).

290. Church, K. W. Word2Vec. *Natural Language Engineering* vol. 23, 155–162 (2017).

291. Goldberg, Y. & Levy, O. word2vec Explained: deriving Mikolov et al.'s negative-sampling word-embedding method. *arXiv preprint arXiv:1402.3722* (2014).

292. Rong, X. word2vec parameter learning explained. *arXiv preprint arXiv:1411.2738* (2014).

293. Li, J., Dai, W., Metze, F., Qu, S. & Das, S. *A comparison of deep learning methods for environmental sound detection* in *2017 IEEE International conference on acoustics, speech and signal processing (ICASSP)* (2017), 126–130.

294. He, X. & Deng, L. Deep learning for image-to-text generation: A technical overview. *IEEE Signal Processing Magazine* vol. 34, 109–116 (2017).

295. Xu, Y. *et al. Deep learning of feature representation with multiple instance learning for medical image analysis* in *2014 IEEE international conference on acoustics, speech and signal processing (ICASSP)* (2014), 1626–1630.

296. David, L., Thakkar, A., Mercado, R. & Engkvist, O. Molecular representations in AI-driven drug discovery: a review and practical guide. *Journal of Cheminformatics* vol. 12, 1–22 (2020).

297. Zhou, J. *et al.* Graph neural networks: A review of methods and applications. *AI Open* vol. 1, 57–81 (2020).

298. Wu, Z. *et al.* A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems* vol. 32, 4–24 (2020).

299. Asif, N. A. *et al.* Graph neural network: A comprehensive review on non-euclidean space. *IEEE Access* vol. 9, 60588–60606 (2021).

300. Zhou, Y. *et al.* Graph Neural Networks: Taxonomy, Advances, and Trends. *ACM Transactions on Intelligent Systems and Technology (TIST)* vol. 13, 1–54 (2022).

301. Wieder, O. *et al.* A compact review of molecular property prediction with graph neural networks. *Drug Discovery Today: Technologies* vol. 37, 1–12 (2020).

302. Jiang, D. *et al.* Could graph neural networks learn better molecular representation for drug discovery? A comparison study of descriptor-based and graph-based models. *Journal of cheminformatics* vol. 13, 1–23 (2021).

303. Wang, Y., Wang, J., Cao, Z. & Farimani, A. B. Molclr: Molecular contrastive learning of representations via graph neural networks. *arXiv preprint arXiv:2102.10056* (2021).

304. Todeschini, R. & Consonni, V. *Handbook of molecular descriptors* (John Wiley & Sons, 2008).

305. Todeschini, R. & Consonni, V. Molecular descriptors. *Recent Advances in QSAR Studies,* 29–102 (2010).

306. Mauri, A. in *Ecotoxicological QSARs* 801–820 (Springer, 2020).

307. Cereto-Massagué, A. *et al.* Molecular fingerprint similarity search in virtual screening. *Methods* vol. 71, 58–63 (2015).

308. Duan, J., Dixon, S. L., Lowrie, J. F. & Sherman, W. Analysis and comparison of 2D fingerprints: insights into database screening performance using eight fingerprint methods. *Journal of Molecular Graphics and Modelling* vol. 29, 157–170 (2010).

309. Xue, L. & Bajorath, J. Molecular descriptors in chemoinformatics, computational combinatorial chemistry, and virtual screening. *Combinatorial chemistry & high throughput screening* vol. 3, 363–372 (2000).

310. Durant, J. L., Leland, B. A., Henry, D. R. & Nourse, J. G. Reoptimization of MDL keys for use in drug discovery. *Journal of chemical information and computer sciences* vol. 42, 1273–1280 (2002).

311. Janela, T., Takeuchi, K. & Bajorath, J. Introducing a Chemically Intuitive Core-Substituent Fingerprint Designed to Explore Structural Requirements for Effective Similarity Searching and Machine Learning. *Molecules* vol. 27, 2331 (2022).

312. Bajusz, D., Rácz, A. & Héberger, K. Chemical Data Formats, Fingerprints, and Other Molecular Descriptions for Database Analysis and Searching. Compr. *Med. Chem. III* vol. 3, 8 (2017).

313. Glen, R. C. *et al.* Circular fingerprints: flexible molecular descriptors with applications from physical chemistry to ADME. *IDrugs* vol. 9, 199 (2006).

314. Rogers, D. & Hahn, M. Extended-Connectivity Fingerprints. *Journal of Chemical Information and Modeling* vol. 50, 742–754 (May 2010).

315. Foster, D. *Generative deep learning: teaching machines to paint, write, compose, and play* (O'Reilly Media, 2019).

316. Mendez, D. *et al.* ChEMBL: towards direct deposition of bioassay data. *Nucleic Acids Research* vol. 47, D930–D940 (Nov. 2018).

317. Kim, S. *et al.* PubChem in 2021: new data content and improved web interfaces. *Nucleic acids research* vol. 49, D1388–D1395 (2021).

318. Sterling, T. & Irwin, J. J. ZINC 15 – Ligand Discovery for Everyone. *Journal of Chemical Information and Modeling* vol. 55, 2324–2337 (Nov. 2015).

319. Bian, Y. & Xie, X.-Q. Generative chemistry: drug discovery with deep learning generative models. *Journal of Molecular Modeling* vol. 27, 1–18 (2021).

320. Sousa, T., Correia, J., Pereira, V. & Rocha, M. Generative deep learning for targeted compound design. *Journal of Chemical Information and Modeling* vol. 61, 5343–5361 (2021).

321. Meyers, J., Fabian, B. & Brown, N. De novo molecular design and generative models. *Drug Discovery Today* vol. 26, 2707–2715 (2021).

322. Bilodeau, C., Jin, W., Jaakkola, T., Barzilay, R. & Jensen, K. F. Generative models for molecular discovery: Recent advances and challenges. *Wiley Interdisciplinary Reviews: Computational Molecular Science,* e1608 (2022).

323. Blaschke, T. *et al.* REINVENT 2.0: an AI tool for de novo drug design. *Journal of chemical information and modeling* vol. 60, 5918–5922 (2020).

324. Hyötyniemi, H. Turing machines are recurrent neural networks. *Proceedings of step* vol. 96 (1996).

325. Rumelhart, D. E., Hinton, G. E. & Williams, R. J. Learning representations by back-propagating errors. *nature* vol. 323, 533–536 (1986).

326. Fernández, S., Graves, A. & Schmidhuber, J. *An application of recurrent neural networks to discriminative keyword spotting* in *International Conference on Artificial Neural Networks* (2007), 220–229.

327. Hopfield, J. J. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences* vol. 79, 2554–2558 (1982).

328. Amsalem, O. *et al.* An efficient analytical reduction of detailed nonlinear neuron models. *Nature communications* vol. 11, 1–13 (2020).

329. Yang, Y., Dan, X., Qiu, X. & Gao, Z. FGGAN: Feature-guiding generative adversarial networks for text generation. *IEEE Access* vol. 8, 105217–105225 (2020).

330. Kaliakatsos-Papakostas, M., Gkiokas, A. & Katsouros, V. in *Proceedings of the Audio Mostly 2018 on Sound in Immersion and Emotion* 1–7 (2018).

331. Kwon, Y. *et al.* Efficient learning of non-autoregressive graph variational autoencoders for molecular graph generation. *Journal of Cheminformatics* vol. 11, 1–10 (2019).

332. Alperstein, Z., Cherkasov, A. & Rolfe, J. T. All smiles variational autoencoder. *arXiv preprint arXiv:1905.13343* (2019).

333. Guimaraes, G. L., Sanchez-Lengeling, B., Outeiral, C., Farias, P. L. C. & Aspuru-Guzik, A. Objective-reinforced generative adversarial networks (ORGAN) for sequence generation models. *arXiv preprint arXiv:1705.10843* (2017).

334. De Cao, N. & Kipf, T. MolGAN: An implicit generative model for small molecular graphs. *arXiv preprint arXiv:1805.11973* (2018).

335. Shi, C. *et al.* Graphaf: a flow-based autoregressive model for molecular graph generation. *arXiv preprint arXiv:2001.09382* (2020).

336. Madhawa, K., Ishiguro, K., Nakago, K. & Abe, M. Graphnvp: An invertible flow model for generating molecular graphs. *arXiv preprint arXiv:1905.11600* (2019).

337. Zhou, Z., Kearnes, S., Li, L., Zare, R. N. & Riley, P. Optimization of molecules via deep reinforcement learning. *Scientific reports* vol. 9, 1–10 (2019).

338. Griffiths, R.-R. & Hernández-Lobato, J. M. Constrained Bayesian optimization for automatic chemical design using variational autoencoders. *Chemical science* vol. 11, 577–586 (2020).

339. Du, Y., Fu, T., Sun, J. & Liu, S. MolGenSurvey: A Systematic Survey in Machine Learning Models for Molecule Design. *arXiv preprint arXiv:2203.14500* (2022).

340. Gao, W., Fu, T., Sun, J. & Coley, C. W. Sample efficiency matters: a benchmark for practical molecular optimization. *arXiv preprint arXiv:2206.12411* (2022).

341. Jensen, J. H. A graph-based genetic algorithm and generative model/Monte Carlo tree search for the exploration of chemical space. *Chemical Science* vol. 10, 3567–3572 (2019).

342. Nicolaou, C. A. & Brown, N. Multi-objective optimization methods in drug design. *Drug Discovery Today: Technologies* vol. 10, e427–e435 (2013).

343. Bagchi, T. P. in *Multiobjective Scheduling by Genetic Algorithms* 171–202 (Springer US, Boston, MA, 1999).

344. Deb, K., Pratap, A., Agarwal, S. & Meyarivan, T. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation* vol. 6, 182–197 (2002).

345. Kandel, E. R. *et al. Principles of neural science* (McGraw-hill New York, 2000).

346. Zong, W. *et al.* Miniature two-photon microscopy for enlarged field-of-view, multi-plane and long-term brain imaging. *Nature methods* vol. 18, 46–49 (2021).

347. Steinmetz, N. A. *et al.* Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science* vol. 372, eabf4588 (2021).

348. Mahfooz, K. & Ellender, T. J. in *Patch clamp electrophysiology* 179–189 (Springer, 2021).

349. Lipovsek, M. *et al.* Patch-seq: Past, present, and future. *Journal of Neuroscience* vol. 41, 937–946 (2021).

350. Schneider, S., Lee, J. H. & Mathis, M. W. Learnable latent embeddings for joint behavioral and neural analysis. *arXiv preprint arXiv:2204.00673* (2022).

351. Gardner, R. J. *et al.* Toroidal topology of population activity in grid cells. *Nature* vol. 602, 123–128 (2022).

352. Einevoll, G. T. *et al.* The scientific case for brain simulations. *Neuron* vol. 102, 735–744 (2019).

353. Sterratt, D., Graham, B., Gillies, A. & Willshaw, D. *Principles of computational modelling in neuroscience* (Cambridge University Press, 2011).

354. Dai, K. *et al.* Brain Modeling ToolKit: An open source software suite for multiscale modeling of brain circuits. *PLOS Computational Biology* vol. 16, e1008386 (2020).

355. Markram, H. *et al.* Reconstruction and simulation of neocortical microcircuitry. *Cell* vol. 163, 456–492 (2015).

356. Billeh, Y. N. *et al.* Systematic integration of structural and functional data into multi-scale models of mouse primary visual cortex. *Neuron* vol. 106, 388–403 (2020).

357. Hagen, E., Næss, S., Ness, T. V. & Einevoll, G. T. Multimodal modeling of neural network activity: computing LFP, ECoG, EEG, and MEG signals with LFPy 2.0. *Frontiers in neuroinformatics* vol. 12, 92 (2018).

358. Hagen, E., Næss, S., Ness, T. V. & Einevoll, G. T. LFPy–multimodal modeling of extracellular neuronal recordings in Python. *bioRxiv,* 620286 (2019).

359. Pettersen, K. H. & Einevoll, G. T. Amplitude variability and extracellular low-pass filtering of neuronal spikes. *Biophysical journal* vol. 94, 784–802 (2008).

360. Lindén, H., Pettersen, K. H. & Einevoll, G. T. Intrinsic dendritic filtering gives low-pass power spectra of local field potentials. *Journal of computational neuroscience* vol. 29, 423–444 (2010).

361. Pettersen, K. H., Lindén, H., Dale, A. M. & Einevoll, G. T. Extracellular spikes and CSD. *Handbook of neural activity measurement* vol. 1, 92–135 (2012).

362. Næss, S. *et al.* Biophysically detailed forward modeling of the neural origin of EEG and MEG signals. *NeuroImage* vol. 225, 117467 (2021).

363. Featherstone, R. E. *et al.* EEG biomarkers of target engagement, therapeutic effect, and disease process. *Annals of the New York Academy of Sciences* vol. 1344, 12–26 (2015).

364. De Bock, R. *et al.* EEG microstates as biomarker for psychosis in ultra-high-risk patients. *Translational Psychiatry* vol. 10, 1–9 (2020).

365. Kim, H. K., Blumberger, D. M. & Daskalakis, Z. J. Neurophysiological biomarkers in schizophrenia—P50, mismatch negativity, and TMS-EMG and TMS-EEG. *Frontiers in Psychiatry* vol. 11, 795 (2020).

366. Barros, C., Silva, C. A. & Pinheiro, A. P. Advanced EEG-based learning approaches to predict schizophrenia: Promises and pitfalls. *Artificial intelligence in medicine* vol. 114, 102039 (2021).

367. Ben-Shalom, R. *et al.* NeuroGPU: Accelerating multi-compartment, biophysically detailed neuron simulations on GPUs. *Journal of Neuroscience Methods* vol. 366, 109400 (2022).

368. Fukushima, K. & Miyake, S. in *Competition and cooperation in neural nets* 267–285 (Springer, 1982).

369. Gu, J. *et al.* Recent advances in convolutional neural networks. *Pattern recognition* vol. 77, 354–377 (2018).

370. LeCun, Y. *et al.* Backpropagation applied to handwritten zip code recognition. *Neural computation* vol. 1, 541–551 (1989).

371. Bai, S., Kolter, J. Z. & Koltun, V. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271* (2018).

372. Beniaguev, D., Segev, I. & London, M. Single cortical neurons as deep artificial neural networks. *Neuron* vol. 109, 2727–2739 (2021).

373. Hochreiter, S. & Schmidhuber, J. LSTM can solve hard long time lag problems. *Advances in neural information processing systems* vol. 9 (1996).

374. Hochreiter, S. & Schmidhuber, J. Long short-term memory. *Neural computation* vol. 9, 1735–1780 (1997).

375. Gers, F. A., Schmidhuber, J. & Cummins, F. Learning to forget: Continual prediction with LSTM. *Neural computation* vol. 12, 2451–2471 (2000).

376. Graves, A. & Schmidhuber, J. Framewise phoneme classification with bidirectional LSTM and other neural network architectures. *Neural networks* vol. 18, 602–610 (2005).

377. Olah, V. J., Pedersen, N. P. & Rowan, M. J. Ultrafast simulation of large-scale neocortical microcircuitry with biophysically realistic neurons. *Elife* vol. 11, e79535 (2022).

378. Shi, X. *et al.* Convolutional LSTM network: A machine learning approach for precipitation nowcasting. *Advances in neural information processing systems* vol. 28 (2015).

379. Davison, A. P. & Appukuttan, S. Computational Neuroscience: A faster way to model neuronal circuitry. *Elife* vol. 11, e84463 (2022).

380. Justus, D., Brennan, J., Bonner, S. & McGough, A. S. *Predicting the computational cost of deep learning models* in *2018 IEEE international conference on big data (Big Data)* (2018), 3873–3882.

381. Wright, S. Recombinant DNA technology and its social transformation, 1972-1982. *Osiris* vol. 2, 303–360 (1986).

382. Hsu, P. D., Lander, E. S. & Zhang, F. Development and applications of CRISPR-Cas9 for genome engineering. *Cell* vol. 157, 1262–1278 (2014).

383. Moon, A. L., Haan, N., Wilkinson, L. S., Thomas, K. L. & Hall, J. CACNA1C: association with psychiatric disorders, behavior, and neurogenesis. *Schizophrenia bulletin* vol. 44, 958–965 (2018).

384. Temme, S. J., Bell, R. Z., Fisher, G. L. & Murphy, G. G. Deletion of the mouse homolog of CACNA1C disrupts discrete forms of hippocampal dependent memory and neurogenesis with in the dentate gyrus. *Eneuro* (2016).

385. Adli, M. The CRISPR tool kit for genome editing and beyond. *Nature communications* vol. 9, 1–13 (2018).

386. Jinek, M. *et al.* A programmable dual-RNA–guided DNA endonuclease in adaptive bacterial immunity. *science* vol. 337, 816–821 (2012).

387. Doudna, J. A. & Charpentier, E. The new frontier of genome engineering with CRISPR-Cas9. *Science* vol. 346, 1258096 (2014).

388. Qi, L. S. *et al.* Repurposing CRISPR as an RNA-guided platform for sequence-specific control of gene expression. *Cell* vol. 152, 1173–1183 (2013).

389. Sullivan, P. F., Fan, C. & Perou, C. M. Evaluating the comparability of gene expression in blood and brain. *American Journal of Medical Genetics Part B: Neuropsychiatric Genetics* vol. 141, 261–268 (2006).

390. Qian, X., Song, H. & Ming, G.-l. Brain organoids: advances, applications and challenges. *Development* vol. 146, dev166074 (2019).

391. Di Lullo, E. & Kriegstein, A. R. The use of brain organoids to investigate neural development and disease. *Nature Reviews Neuroscience* vol. 18, 573–584 (2017).

392. Quadrato, G., Brown, J. & Arlotta, P. The promises and challenges of human brain organoids as models of neuropsychiatric disease. *Nature medicine* vol. 22, 1220–1228 (2016).

393. Mäki-Marttunen, T. *et al.* Alterations in schizophrenia-associated genes can lead to increased power in delta oscillations. *Cerebral Cortex* vol. 29, 875–891 (2019).

394. Mäki-Marttunen, T. *et al.* Computational modeling of genetic contributions to excitability and neural coding in layer V pyramidal cells: applications to schizophrenia pathology. *Frontiers in computational neuroscience* vol. 13, 66 (2019).

395. Devor, A. *et al.* Genetic evidence for role of integration of fast and slow neurotransmission in schizophrenia. *Molecular psychiatry* vol. 22, 792–801 (2017).

396. Mäki-Marttunen, T. *et al.* Pleiotropic effects of schizophrenia-associated genetic variants in neuron firing and cardiac pacemaking revealed by computational modeling. *Translational psychiatry* vol. 7, 1–13 (2017).

Chapter 2

Concepts and Methodology

In the previous chapter, we introduced the pathophysiology and underlying genomics of severe mental disorders, and showed how computational tools in drug design and neuroscience can be useful in studying and potentially treating these diseases. We highlighted challenges in de novo drug design regarding optimisation stagnation and noted opportunities to accelerate computational neuroscience models through deep learning. In this chapter, we delve deeper into the concepts and ideas necessary to understand and evaluate the three research projects presented in this dissertation. Two of these projects focus on topics regarding quality-diversity methods in computational drug design, and one research project explores the premise of using multi-task deep learning with expert diversity for the distillation of biophysically-detailed neuron models.

A common theme in these three different research projects, aside from their connection to the study of severe mental disorders, is their consistent use of some form of computational diversity as a resource. In the first project, we make use of a quality-diversity method, which leverages chemical diversity to avoid being trapped in local minima, to solve the stagnation issues of GB-GA. The follow-up project shows that NSGAs do not rely on chemical diversity to deal with complex chemical trade-offs in optimisation problems but rather make use of diversity enforced in optimisation space. Finally, in the project that distils biophysically-detailed neuron models into MTL Architectures, we tested deep learning models that encourage diversity in internal data representations by feeding the data through multiple copies of the same input architecture.

Implementation details, exact results and complete benchmarks for each project can be found in the corresponding manuscripts, which are presented at the end of this dissertation. Open-source software for each of the research projects is (or will be) available online and can be found in the Github repositories cited in the manuscripts. These repositories are living documents, and have in some cases undergone updates and improvements since the publication of the manuscripts, and might undergo further updates after the publication of this dissertation. A history of all changes to the software can be found in the track-changes of the respective Github repositories. The majority of the data generated for the plots and results, and benchmarks presented in the manuscripts has been made available at the time of publication, all other (often more cumbersome) data can be obtained from the authors upon simple request.

Quality-Diversity Methods in De Novo Drug Design

In optimisation problems where there is no analytical gradient to compute, as is typically the case in molecular optimisation, genetic algorithms are known to be powerful optimisation tools [1–3]. These algorithms, such as GB-GA [4], are inspired by the biological phenomena of evolution by natural selection. In nature, we observe variation

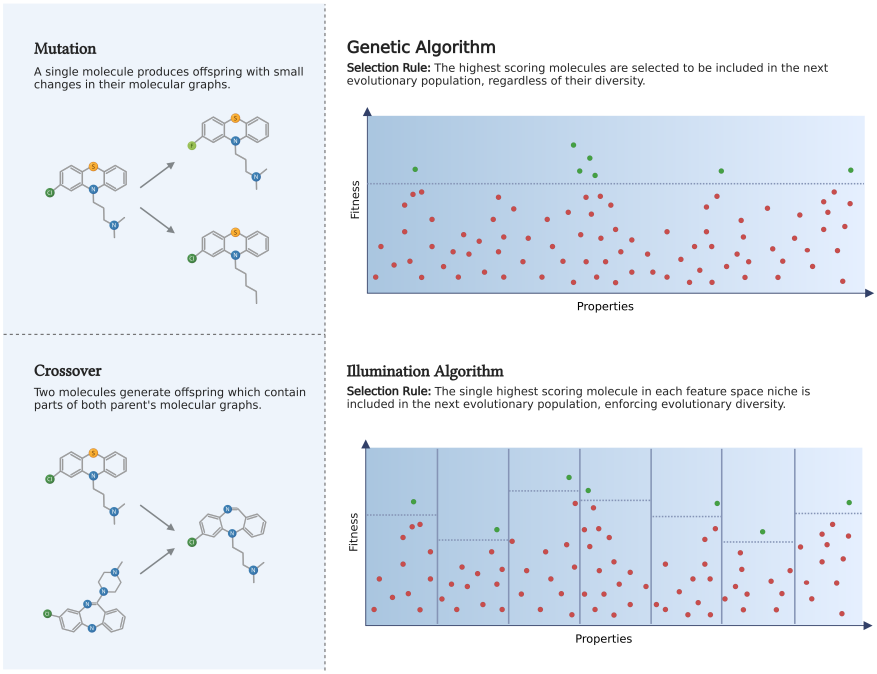

Figure 2.1: **Graph-Based Genetic and Illumination Algorithms.** Both GB-GA and GB-EPI make use of graph-based mutations and crossovers (left), but differ in how they select a population of molecules for inclusion in the next generation (right). A non-exhaustive list of examples of mutations and crossovers on molecular graphs is shown here to illustrate the concepts behind these operators rather than the specific technicalities of their implementation in GB-GA and GB-EPI. The two different selection procedures are illustrated by drawing a demarcation line (dotted) between the selected (green) and discarded (red) individuals of the evolutionary population.

within all types of life-forms due to an interplay of random genetic mutations and genetic inheritance. Natural selection acts on this variation though differential survival, based on how successful a life-form is at reproducing. Genetic algorithms aim to emulate this mechanism by subjecting a population of candidate solutions of an optimisation problem to iterative cycles of variation and differential selection based on their optimisation performance. In this way, the potential for a candidate solution to participate in the next cycle of the algorithm is tied to their ability to solve a given optimisation problem.

Iteration after iteration, the repeated application of simple natural selection rules, as such described above, forces the population of candidate solutions to move towards a better optimisation performance. In the genetic algorithm community, a candidate solution is called an individual and its optimisation performance is known as its *fitness*. The initial population of individuals (i.e. candidate solutions) is typically randomly generated or chosen from an existing database of possible solutions to the optimisation problem. Starting from this initial population, new individuals are created by two

variation operators, known as *mutations* and *crossovers*, that mimic natural variation. Mutations randomly change small aspects of existing individuals whereas crossovers combine parts of existing individuals into a novel candidate solution.

Many slight variations on the basic genetic algorithmic scheme, for specific use-cases, have been devised over the years. In GB-GA, candidate molecules are represented by their molecular graphs, and mutations and crossovers act on these graphs by randomly changing atom or bond types or by merging molecular fragments into a new molecule. Schematically, GB-GA can be represented by the following three recurring steps:

1. **Evaluation:** Evaluate the fitness of each molecule in the current evolutionary population, with respect to a given optimisation function (also known as a scoring or objective function).

2. **Selection:** Rank all the molecules according to their fitness and retain only a subsection of the population based on that ranking, either directly or through a weighted stochastic process.

3. **Variation:** Apply the variation operators, i.e. mutations and crossovers, to the underlying molecular graphs of the selected subpopulation to create new molecules to add to the population.

To obtain relevant and realistic molecules, molecular graphs that have been generated with incorrect valences (as determined by RDKit), or those that contain macrocycles or allene centers in rings are systematically discarded from the population. The initial set of molecules used in GB-GA is typically randomly chosen from the first 1000 molecules in the ZINC dataset [5].

While far more efficient than deep learning methods for de novo drug design, GB-GA sporadically suffers from stagnation issues [6, 7]. Stagnation occurs when the algorithm cannot find its way out of a population of sub-optimal molecules due to a lack of (molecular) diversity. Molecules in or near a local fitness optimum are hard to remove from the evolutionary population because almost all mutations and crossovers produce molecules with a lower fitness. As a consequence, molecules from the local fitness optimum start to dominate the evolutionary population by filling it with highly similar molecules which also lie in the local optimum. This causes stagnation of the genetic algorithm. When stagnation is encountered in GB-GA, the algorithm needs to be either restarted completely or the evolutionary population has to be supplemented with an ad-hoc variety of other molecules.

To encourage broad exploration of the search space, and hence avoid stagnation, a new class of optimisation algorithms called quality-diversity (QD) algorithms [8] was developed. These algorithms are exemplified by the multi-dimensional archive of phenotypic elites (MAP-Elites) algorithm [9–12]. MAP-Elites is a simple and efficacious QD approach that enforces *population diversity* explicitly by splitting the evolutionary population into feature space niches and retaining only the single most fit individual within each niche. In MAP-Elites, each candidate solution is associated with a vector describing a handful of user-defined properties. Based on this descriptor, the candidate solution is assigned to a niche in feature space. MAP-Elites and GB-GA serve together as the core architectures of the Graph-Based Elite Patch Illumination (GB-EPI) algorithm, presented in Paper I: *Illuminating Elite Patches of Chemical Space*.

The Graph-Based Elite Patch Illumination Algorithm

The GB-EPI algorithm builds on the inner mechanisms of GB-GA, in the sense that it continues to employ molecular graphs to which mutations and crossovers are applied. GB-EPI implements the necessary feature vector as a list of physicochemical properties of a molecule, effectively turning the archive of niches into a map of chemical space. Dividing chemical space into feature-based niches and explicitly enforcing population diversity is a paradigm shift away from GB-GA [13–15]. GB-EPI allows users to choose their own features of interest and the value ranges in which to explore these features. This enables chemists to demarcate the subset of chemical space in which the algorithm can search for locally optimal molecules. In the manuscript, pseudo-code for the MAP-Elites algorithm as applied to de novo molecule design in GB-EPI is provided as well as a link to the Github repository containing a full and open-source implementation of GB-EPI.

The molecules retained at the end of a GB-EPI run form a patchwork of locally optimal (or *elite*) solutions which clarify design trade-offs, where their fitness scores are an indication of how optimisation performance varies over chemical space. We hope that access to this information will both encourage interaction between the algorithm and medicinal chemists, as well as aid the development of more realistic scoring functions for molecular optimisation. On top of these additional use-cases, GB-EPI also provides the community with a straightforward and interpretable optimisation algorithm that avoids stagnation and is more efficient than GB-GA. At every generation, GB-EPI contains solutions spread out over the feature space. Solutions in far-away niches are an ever-present resource of diversity that can be accessed to accelerate optimisation or to escape stagnation.

A multitude of small technical improvements on the original GB-GA and MAP-Elites algorithms have been included in GB-EPI. Among others, these improvements include the decoupling of mutations and crossovers for improved optimisation, positional analogue scanning for broader exploration [16], memoisation to reduce unnecessary function calls [17], and the use of niches organised along a centroidal Voronoi tessellation (CVT) [18]. Further motivation and implementation details for all of these technical improvements can be found in the manuscript and accompanying code. Finally, it is worth noting that GB-EPI has a broad set of options for filtering out molecules based on ADMET property calculations and structural filters, and has built-in concurrency for both fitness and feature vector calculations, which significantly decreases the overall runtime of the algorithm.

Benchmarking Molecular Optimisation Through GuacaMol

To assess the efficiency of different generative models for de novo molecular design, the London-based bioinformatics company BenevolentAI has released an open-source benchmarking suite called GuacaMol [19]. This benchmarking suite provides researchers with a variety of computationally affordable molecular optimisation tasks. In the manuscript where we introduce GB-EPI, we made use of GuacaMol to compare the performance of a representative SMILES RNN model with GB-GA and GB-EPI on three explicit rediscovery tasks and two median molecule tasks. The Guacamol

benchmarking suite automatically removes molecules highly similar to the targets from the database of initial molecules usable by the generative algorithms. In addition, to make the benchmark more informative, we also recorded the results for all three algorithms on both median molecules benchmarks for a more randomised set of starting molecules.

Explicit rediscovery tasks require the generative algorithms being benchmarked to rediscover an existing FDA-approved molecule through a computationally affordable similarity measure as a fitness function. Typically, this measure is the Tanimoto similarity applied to some version of extended-connectivity fingerprints [20]. The median molecules benchmarks on the other hand, require the generative models to find a molecule that maximises the geometric mean of similarities to two different existing molecules. The median molecules benchmarks starting from highly randomised molecules are far more strenuous than the explicit rediscovery benchmarks. On these more difficult benchmarks, GB-EPI was shown to outperform both the SMILES RNN model and GB-GA. Notably GB-EPI was observed to cover a significantly larger part of the chemical space. More details, and exact results for the benchmarks can be found in the manuscript.

Pareto Optimisation and De Novo Drug Design

Due to the success of GB-EPI in the GuacaMol median molecules benchmarks, and the fact that GB-EPI covers wide parts of chemical space by design, questions about the algorithms usefulness in MOO problems arose quickly after its publication. Instead of testing GB-EPI's efficiency on benchmarks that aggregate performance across different objective functions into a single value, we turned to Pareto optimisation problems to truly probe the algorithms efficiency. In Pareto optimisation, algorithms are tasked to find a set of solutions that covers all the optimal trade-offs between objectives. None of these solutions can be improved in any one objective without lowering the performance with regard to at least any one other objective. Solutions fulfilling these requirements are called *Pareto dominant* and they form a structure in optimisation space known as the *Pareto front* [21].

Making conclusions about the efficiency and potency of GB-EPI in this setting was made significantly more difficult by the lack of comparable, open-source implementation of baseline genetic algorithm methods for Pareto optimisation in de novo drug design. Therefore we implemented open-source versions of NSGA-II and NSGA-III for molecular MOO [22–24]. In Paper II: *Graph-Based Molecular Pareto Optimisation* we presented these algorithms and compared their optimisation efficiency with with GB-EPI on seven different newly introduced MOO benchmarks. As a novelty, we tracked the chemical diversity of the evolutionary populations of these algorithms during optimisation and discovered that NSGA-II and NSGA-III obtain high levels of optimisation efficiency without relying on chemical diversity.

Instead of applying differential selection within niches, NSGA algorithms divide an evolutionary population based on a ranking of Pareto dominance ranking. NSGA algorithms only truly differ in how they use objective space diversity to differentiate between candidate solutions with a similar rank, see Figure 2.2. NSGA-II, for instance, promotes the selection of molecules with a larger optimisation space distance from

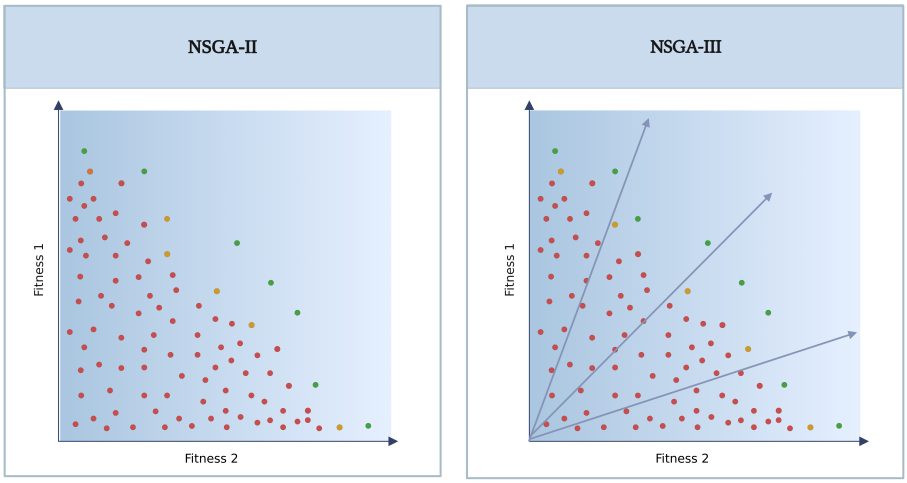

Figure 2.2: **Splitting Front Procedures in NSGA-II and NSGA-III.** NSGA-II and NSGA-III both select molecules based on their Pareto dominance ranking, but differ in how they select molecules with an identical ranking. For illustrative purposes, we here choose to apply the splitting procedures on the second Pareto front. The members of the evolutionary population selected for the next generation are either Pareto dominant (green) or more diverse in the optimisation space (orange) than their discarded counterparts (red). In NSGA-II, diversity is measured through means of a crowding measure which is used to select the most isolated individuals. NSGA-III, on the other hand, makes use of orthogonal distances to pre-defined reference directions to enforce diversity by selecting the individuals with the smallest orthogonal distance to each of the directions and the fitness directions.

other molecules. Whereas NSGA-III enforces diversity by selecting molecules close to pre-defined (equally spaced) reference directions in optimisation space. Note that the implementations of NSGA-II and NSGA-III presented here make use of the previously discussed improvements on GB-GA to efficiently calculate fitnesses and modify molecules. Further implementation details for these algorithms, and comments on the incorporation of recent technical improvements for NSGAs can be found in the manuscript.

Dominated Hypervolume and Internal Structural Similarity

The dominated hypervolume measure calculates the size of the region Pareto dominated by a set of points in objective space and is an indicator of how close the points are to the ideal Pareto front and how spread out these points are over objective space [25]. Using the dominated hypervolume as an evaluation metric, both NSGA-II and NSGA-III were

shown to outperform GB-EPI (which optimised the aggregated version of the same benchmarks). In general, the effectiveness of NSGA-II and NSGA-III is known to dependent on the specific Pareto optimisation problems they are applied to. In our benchmarks no consistent difference in term of efficiency was found between NSGA-II and NSGA-III. Markedly, NSGA-II and NSGA-III manage to produce solutions with a higher aggregated objective score than GB-EPI while not directly optimising for this metric.

To gain deeper insight into the differences between the optimisation algorithms, we also calculated the internal structural similarity of their respective evolutionary populations, as a measure of chemical diversity. No unified nor generally accepted framework for measuring chemical similarity of two or more compounds currently exists. The previously mentioned Tanimoto similarity (applied to ECFPs) is often used to compare two molecules. These Tanimoto similarities can also be used to measure similarity across a population of molecules by, for instance, calculating the determinant of the matrix of all pair-wise similarities. While this approach is at least minimally informative, it suffers from a computational cost that scales quadratically with the amount of molecules in the population. In an attempt to address this issue and to standardise similarity computations over a large group of molecules, extended similarity metrics have recently been proposed in the literature [26, 27].

Extended similarity metrics were designed to compare a set of equal-length bit-vectors, (like ECFPs) based on the number of coincidences of ones and zeros in each position along the vectors. Extended similarity metrics can be constructed so that, when they are applied to a pair of bit-vectors, they naturally reduce to existing binary measures such as Tanimoto similarities. In paper II, one such extended similarity index is applied to each generation of the evolutionary populations of NSGA-II, NSGA-III, and GB-EPI to study how internal similarity evolves throughout optimisation. This measurement demonstrated that NSGA-II and NSGA-III achieve their efficient Pareto optimisation while increasing internal structural similarity, in contrast to GB-EPI. Whether other, more involved, similarity indices can incorporate diversity in objective space and structural diversity into one single measurement remains an open and interesting question.

Distilling Biophysically-Detailed Neuron Models

Computational modelling has become a cornerstone of modern neuroscience [28]. As discussed in the first chapter, electrophysiological properties of neurons can be simulated at many different levels of detail. Multi-compartmental, biophysically-detailed neuron models are among some of the most involved and detailed neuron models available, but require significant computational resources [29, 30]. By distilling these costly models into ANN architectures, which by design are computationally cheaper and easier to evaluate, large-scale, biophysically-detailed network models can be democratised and will hopefully become a regular occurrence in neuroscience publications. A simple CNN architecture adjusted for causality, known as a TCN has been shown to be able to accurately predict action potentials based on an input consisting of a brief history (17 ms - 205 ms) of synaptic inputs [31, 32]. To predict voltage traces, especially in multiple

compartments simultaneously, more involved architectures seem to be necessary.

Recent work has claimed that both feed-forward networks and TCNs are capable of predicting sub-threshold membrane potentials but unable to predict voltage traces that include action potentials [33]. It is claimed, that in the presence of action potentials these models fail to deal with the non-normal distribution of voltage data and converge to the mean membrane potential of the training dataset. The use of a combined TCN-LSTM architecture, applied to a combination of synaptic inputs and voltage traces, on the other hand, was shown to be able to predict both sub-threshold activity and action potentials. Encouragingly, the interpretability and generalisation capability of the TCN-LSTM architecture was shown to be comparatively high. Exploration of TCN-LSTM parameters after training showed that the model gives equal importance to distal and proximal synaptic inputs and the TCN-LSTM model is efficiently re-trainable to account for specific changes in the underlying biophysically-detailed neuron. In addition, it is suggested that the TCN-LSTM model is a good starting point for predicting optical readouts (i.e. fluorescent calcium indicators).

The TCN-LSTM architecture was initially trained on a biophysically-detailed neuron model of a layer V pyramidal neuron, but later (in the same paper) the architecture was also used to establish distilled neuronal models for the other major cortical pyramidal neuron types (layer II/III, layer IV and layer VI) [34]. For a single neuron simulation, the TCN-LSTM architecture was shown to give speed-ups of up to three orders of magnitude. For network simulations, an acceleration of close to five orders of magnitude has been recorded. With the availability of these efficient tools for running biophysically-detailed neuron models, parameter scans of large network simulations become feasible. As an example, the authors that introduced TCN-LSTM distillation to biophysically-detailed neuron models studied the influence of recurrent connections and excitatory drive in a cortical network of layer V pyramidal neurons to understand network instability in the context of Rett syndrome pathophysiology [35, 36].

Two of the most obvious applications of efficient and accurate distillation of biophysically-detailed neuron models are accelerating large-scale networks, such as the previously mentioned rat sensory somatic cortex and the Allen V1 cortical models, and the creation of in-silico disease models for a range of neurological diseases including severe mental disorders. A biophysically-detailed simulation of the mouse V1 cortical area could, for instance, be accelerated using no more than 114 different distilled neuron models, each of which corresponds to one type of neuron in the network model. However, as explained in the previous chapter, for both existing large-scale networks of brain areas and novel in-silico disease models, it would be highly beneficial to be able to simulate LFPs. The current implementations of distilled neuron models only predict the membrane potential in a single or limited amount of compartments. To efficiently extend the membrane potential predictions to every compartment, we turned to ideas from MTL which were previously have been applied to other biological and medical datasets [37, 38].

Heterogeneous Multi-task Learning with Expert Diversity

Predicting multiple heterogeneous targets is a challenge for traditional deep learning architectures. The simplest approach to predict a heterogeneous set of targets is to make

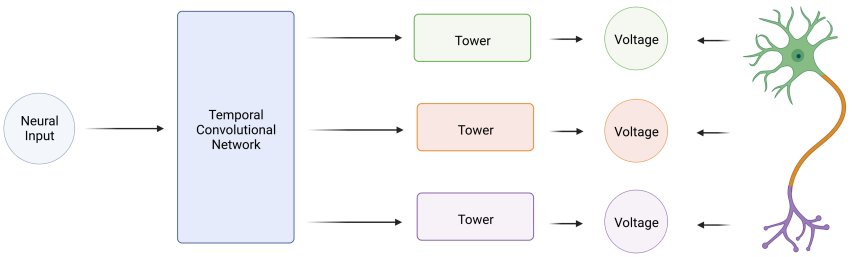

Figure 2.3: **Hard Parameter Sharing for Prediction of Membrane Potentials.** Neural network architecture of hard parameter sharing in multi-task learning for the simultaneous prediction of membrane potentials. This network is composed of a comparatively large, shared bottom and smaller task-specific branches known as towers. Each of the towers predicts the membrane voltage for one compartment of the biophysically-detailed neuron model. Optionally, an extra tower can be added for a binary prediction of the presence or absence of an outgoing spike. The input data is either a history of synaptic inputs or a combined history of synaptic inputs and membrane voltages.

a separate deep learning model for each target. Unfortunately, this approach incurs a computational cost that increases linearly with the amount of tasks that need to be learned. Whereas separate deep learning models can be used in situations when the different tasks are not strongly correlated with each other, it is often more efficient (and sometimes more effective) to leverage associations between tasks to train a single model that makes simultaneous predictions for all tasks [39, 40]. Architectures of this sort can be divided into two main groups: the hard parameter sharing models and the soft parameter sharing models. In our research, we have applied one type of hard parameter sharing model and two types of soft parameter sharing models to predict the membrane potentials across all compartments of a cortical layer V biophysically-detailed neuron model [41].

In hard parameter sharing models, a single neural network (often called *the bottom*) is used to learn a representation which is shared by a series of smaller task-specific networks (known as *the towers*) that each produce the output for an individual task [42]. The main advantage of this approach is that the amount of parameters in the model scales more slowly with respect to the amount of tasks. Soft parameter sharing models on the other hand, make use of multiple encoders (also known as *the experts*) that learn different representations of the input data which are then mixed-and-matched to the experts. The main advantage of soft parameter sharing is that it provides flexible

feature sharing which can encourage the network to learn both task-specific and globally shared representations. For our purposes, we will focus on two architectures that use a data-driven gating mechanism to facilitate soft parameter sharing. These two approaches are known as the multi-gate mixture-of-experts (MMoE) [43] and multi-gate mixture-of-experts with exclusivity (MMoEEx) [44].

MMoE and MMoEEx mix-and-match feature representations learned by the experts through adaptive gate functions which are trained on the input data to give different weights to different representations through matrix multiplication. The fact that these adaptive gate functions are exposed to the input data allows for dynamic parameter allocation between task-specific and global representations. However, the only initial source of diversity among the experts lies in the random initialisation of the network architecture. As a consequence, there is no strict guarantee that different experts specialise in learning either task-specific or shared representations. To encourage further specialisation among the experts, MMoEEx adds exclusion and exclusivity conditions to the MMoE architecture. The MMoEEx approach forces some experts to only contribute representations to some towers while other experts can learn representations available to all towers which has caused improved learning in some MTL tasks.

Multi-task Learning for Biophysically Detailed Neuron Models

In our exploratory study, we use a TCN architecture for both the bottom and all the experts. These TCNs are used to learn representations of a short history of synaptic inputs and voltages traces. In the case of our hard parameter sharing model, we pass this learned representation to experts consisting of a simple feed-forwards neural network for each compartment, see Figure 2.3. As a form of guidance for the algorithm, we also predict the absence or presence of an action potential in the soma. In the case of MMoE or MMoEEx, we use a homogenous set of TCN architectures as experts. However, due to the particularities of our input data, and more specifically the large amount of input features, we add a single extra expert to compress the input data before presenting it to the gating functions, see Figure 2.4. Otherwise, according to our experience exploring this, the size of the gating functions (in terms of learnable parameters) becomes prohibitive large with respect to the use of specialised acceleration hardware.

Implementation details, training results, and a study regarding the importance of expert diversity in our models is provided in a preliminary version of our manuscript *Multi-stask Learning of Biophysically-Detailed Neuron Models*, which can be found at the end of this dissertation. In short, our initial conclusions and insights indicate that using a TCN architecture as the bottom or experts in our hard or soft parameter sharing schemes is sufficient to predict both the action potentials and the sub-threshold dynamics of the voltage traces across compartments. Without giving any extra weight to the binary prediction task, we observe no to little improvement on this single task during training. One important aspect in training these multi-task architectures for predicting neural behaviour is the need to include a history of standardised voltage traces of each compartment in the training or test data. Future work solidifying and clarifying these results, and a link to an open-source implementation of these architectures will be provided in the final version of the manuscript.

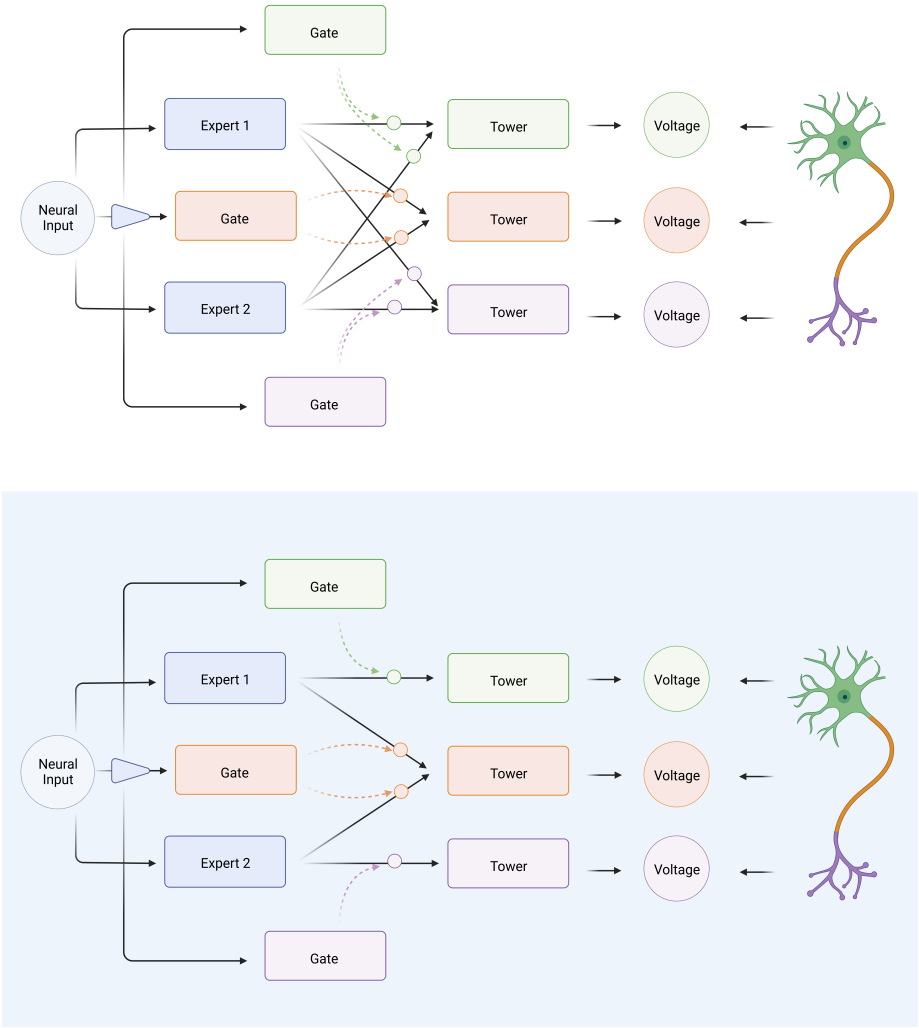

Figure 2.4: **Soft Parameter Sharing for Prediction of Membrane Potentials.** Network architectures of the soft parameter sharing approaches multi-gate mixture-of-experts (top) and multi-gate mixture-of-experts with exclusivity (bottom) for the simultaneous prediction of membrane potentials. These networks are composed of a (partially) shared set of experts and task-specific branches known as towers. The connections between the outputs of the experts and the towers are mixed and matched by gating functions trained on the (compressed) input data.

References

1. Maier, H. *et al.* Introductory overview: Optimization using evolutionary algorithms and other metaheuristics. *Environmental Modelling & Software* vol. 114, 195–213 (2019).
2. Holland, J. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence* (MIT Press, 1992).
3. Goldberg, D. E. & Holland, J. H. Genetic Algorithms and Machine Learning. *Machine Learning* vol. 3, 95–99 (1988).
4. Jensen, J. H. A graph-based genetic algorithm and generative model/Monte Carlo tree search for the exploration of chemical space. *Chemical Science* vol. 10, 3567–3572 (2019).
5. Sterling, T. & Irwin, J. J. ZINC 15 – Ligand Discovery for Everyone. *Journal of Chemical Information and Modeling* vol. 55, 2324–2337 (Nov. 2015).
6. Nigam, A. *et al.* Augmenting Genetic Algorithms with Deep Neural Networks for Exploring the Chemical Space. *International Conference on Learning Representations* (2020).
7. Zhou, Z. & Harris, K. D. M. Counteracting stagnation in genetic algorithm calculations by implementation of a micro genetic algorithm strategy. *Physical Chemistry Chemical Physics* vol. 10, 7262–7269 (48 2008).
8. Chatzilygeroudis, K., Cully, A., Vassiliades, V. & Mouret, J.-B. in *Black Box Optimization, Machine Learning, and No-Free Lunch Theorems* 109–135 (Springer, 2021).
9. Vassiliades, V. *et al. Comparing Multimodal Optimization and Illumination* in *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (Association for Computing Machinery, Berlin, Germany, 2017), 97–98.
10. Mouret, J.-B. & Clune, J. Illuminating search spaces by mapping elites. *arXiv e-prints,* arXiv:1504.04909 (Apr. 2015).
11. Mouret, J.-B. & Clune, J. *An algorithm to create phenotype-fitness maps* in *Proceedings of the Artificial Life Conference* (2012), 593–594.
12. Lehman, J. & Stanley, K. O. *Evolving a Diversity of Virtual Creatures through Novelty Search and Local Competition* in *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation* (Association for Computing Machinery, Dublin, Ireland, 2011), 211–218.
13. Bleicher, S. S. *et al.* Divergent behavior amid convergent evolution: A case of four desert rodents learning to respond to known and novel vipers. *PLOS ONE* vol. 13, 1–17 (Aug. 2018).
14. Rundle, H. *et al.* Divergent Selection and the Evolution of Signal Traits and Mating Preferences. *PLOS Biology* vol. 3, e368 (Oct. 2005).

15. Gulick, J. Divergent Evolution. *Nature* vol. 39, 54–55 (Nov. 1888).

16. Pennington, L. D. *et al.* Positional Analogue Scanning: An Effective Strategy for Multiparameter Optimization in Drug Design. *Journal of Medicinal Chemistry* (2020).

17. Hughes, J. *Lazy memo-functions* in *Conference on Functional Programming Languages and Computer Architecture* (1985), 129–146.

18. Vassiliades, V. *et al.* Using Centroidal Voronoi Tessellations to Scale Up the Multi-dimensional Archive of Phenotypic Elites Algorithm. *IEEE Transactions on Evolutionary Computation* (Aug. 2017).

19. Brown, N. *et al.* GuacaMol: Benchmarking Models for de Novo Molecular Design. *Journal of Chemical Information and Modeling* vol. 59, 1096–1108 (2019).

20. Rogers, D. & Hahn, M. Extended-Connectivity Fingerprints. *Journal of Chemical Information and Modeling* vol. 50, 742–754 (May 2010).

21. Van Veldhuizen, D. A., Lamont, G. B., *et al. Evolutionary computation and convergence to a pareto front* in *Late breaking papers at the genetic programming 1998 conference* (1998), 221–228.

22. Deb, K., Pratap, A., Agarwal, S. & Meyarivan, T. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation* vol. 6, 182–197 (2002).

23. Deb, K. & Jain, H. An Evolutionary Many-Objective Optimization Algorithm Using Reference-Point-Based Nondominated Sorting Approach, Part I: Solving Problems With Box Constraints. *IEEE Transactions on Evolutionary Computation* vol. 18, 577–601 (2014).

24. Jain, H. & Deb, K. An Evolutionary Many-Objective Optimization Algorithm Using Reference-Point Based Nondominated Sorting Approach, Part II: Handling Constraints and Extending to an Adaptive Approach. *IEEE Transactions on Evolutionary Computation* vol. 18, 602–622 (2014).

25. Guerreiro, A. P., Fonseca, C. M. & Paquete, L. The hypervolume indicator: Problems and algorithms. *arXiv preprint arXiv:2005.00515* (2020).

26. Miranda-Quintana, R. A., Bajusz, D., Rácz, A. & Héberger, K. Extended similarity indices: the benefits of comparing more than two objects simultaneously. Part 1: Theory and characteristics. *Journal of cheminformatics* vol. 13, 1–18 (2021).

27. Miranda-Quintana, R. A., Rácz, A., Bajusz, D. & Héberger, K. Extended similarity indices: the benefits of comparing more than two objects simultaneously. Part 2: speed, consistency, diversity selection. *Journal of cheminformatics* vol. 13, 1–16 (2021).

28. Einevoll, G. T. *et al.* The scientific case for brain simulations. *Neuron* vol. 102, 735–744 (2019).

29. Markram, H. *et al.* Reconstruction and simulation of neocortical microcircuitry. *Cell* vol. 163, 456–492 (2015).

30. Billeh, Y. N. *et al.* Systematic integration of structural and functional data into multi-scale models of mouse primary visual cortex. *Neuron* vol. 106, 388–403 (2020).

31. Bai, S., Kolter, J. Z. & Koltun, V. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271* (2018).

32. Beniaguev, D., Segev, I. & London, M. Single cortical neurons as deep artificial neural networks. *Neuron* vol. 109, 2727–2739 (2021).

33. Olah, V. J., Pedersen, N. P. & Rowan, M. J. Ultrafast simulation of large-scale neocortical microcircuitry with biophysically realistic neurons. *Elife* vol. 11, e79535 (2022).

34. Gouwens, N. W. *et al.* Systematic generation of biophysically detailed models for diverse cortical neuron types. *Nature communications* vol. 9, 1–13 (2018).

35. Roche, K. J. *et al.* Electroencephalographic spectral power as a marker of cortical function and disease severity in girls with Rett syndrome. *Journal of neurodevelopmental disorders* vol. 11, 1–14 (2019).

36. Steffenburg, U., Hagberg, G. & Hagberg, B. Epilepsy in a representative series of Rett syndrome. *Acta paediatrica* vol. 90, 34–39 (2001).

37. Johnson, A. E. *et al.* MIMIC-III, a freely accessible critical care database. *Scientific data* vol. 3, 1–9 (2016).

38. Wang, Y. *et al.* Pubchem bioassay: 2017 update. *Nucleic acids research* vol. 45, D955–D963 (2017).

39. Zhang, Y. & Yang, Q. A survey on multi-task learning. *IEEE Transactions on Knowledge and Data Engineering* (2021).

40. Standley, T. *et al. Which tasks should be learned together in multi-task learning?* in *International Conference on Machine Learning* (2020), 9120–9132.

41. Hay, E., Hill, S., Schürmann, F., Markram, H. & Segev, I. Models of neocortical layer 5b pyramidal cells capturing a wide range of dendritic and perisomatic active properties. *PLoS computational biology* vol. 7, e1002107 (2011).

42. Caruana, R. *Multitask learning: A knowledge-based source of inductive bias1* in *Proceedings of the Tenth International Conference on Machine Learning* (1993), 41–48.

43. Ma, J. *et al. Modeling task relationships in multi-task learning with multi-gate mixture-of-experts* in *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining* (2018), 1930–1939.

44. Aoki, R., Tung, F. & Oliveira, G. L. Heterogeneous multi-task learning with expert diversity. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* vol. 19, 3093–3102 (2022).

Chapter 3

Conclusion and Outlook

Severe mental disorders such as schizophrenia and bipolar disorder impose a heavy cost on society and form a significant challenge for the scientific, psychotherapeutic, and medical communities. The search for effective cures or treatments of these diseases is an ongoing endeavour, hampered by shortcomings in our current biomolecular understanding of mental functioning. In this dissertation, we have aimed to give a sense of the complexity involved in studying severe mental disorders, and tried to highlight the diversity of tools – computational or otherwise – involved therein. In this final chapter, we recapitulate the motivation and methods discussed in the preceding chapters, and discuss opportunities for future research. In addition, we provide an overview of the research and teaching contributions made by the author during the period of these PhD studies. As an endpoint to this dissertation, we provide the reader with brief concluding remarks.

Summary and Discussion

This dissertation covers a wide range of research topics that are linked together (conceptually and technically) through the study of severe mental disorders and the use of computational diversity as a resource. Current treatments options for schizophrenia and bipolar disorder, in terms of pharmacological and psychosocial intervention, have a limited rate of success. The development of novel compounds to treat severe mental disorders is curbed by large caveats in our current understanding of the biomolecular mechanisms underlying these diseases, the general challenges in high-throughput screening or de novo drug design, and the lack of biophysically-detailed computer models that can combine and integrate the effects of many different genetic variations. For the latter two of these challenges, we introduced novel computational tools and performed analyses of the importance of diversity in these tools.

Biomolecular Mechanisms of Severe Mental Disorders

Schizophrenia and bipolar disorder are heterogeneous diseases characterised by a wide range of clinical manifestations regarding cognitive, perceptive, and emotional functioning. While external factors such as pre-natal malnutrition, childhood infection or substance abuse are known to contribute to the development of these disorders, there is no doubt that genetic variations form the main biological foundation upon which which mental disorders can develop. Unfortunately, much of the spectrum of genetic variations that cause or facilitate the pathophysiology of schizophrenia and bipolar disorder remains unknown. In this dissertation, we have attempted to make a synthesis of current genetic screening efforts and argue that at least four distinct biomolecular sectors (glutamatergic neurotransmission, neuron morphology, ubiquitin pathways, and

calcium ion dynamics) are strongly implicated in severe mental disorders. Recent research also seems to indicate that neurodevelopmental pathways, novel open reading frames in human accelerated regions, and plausibly mitochondrial dysfunction are involved in these disorders.

As mentioned before, the exact biomolecular mechanisms underlying severe mental disorders remain unclear. Even for known genetic risk factors, extensive in-vitro, in-vivo, and in-silico studies are necessary to elucidate the impact of genetic variants on neuronal functioning. Development of disease models for severe mental disorders is ongoing at both the University of Oslo and the Oslo University Hospital. It is important to note that each of these techniques comes with scientific and technical advantages and disadvantages. Brain organoids have become a potent model of early human brain development but face technological challenges and are hard to link to specific behavioural phenotypes. Animal models, on the other hand, can be used to study behaviour but are hard to match to clinical biomarkers. Biophysically-detailed neuron models have the potential to bridge the gap between these experimental modalities but require significant improvements before they can be considered to accurately represent neuronal biology.

De Novo Drug Design in The Age of Deep Learning

Novel antipsychotic compounds for use in patients with severe mental disorders are highly sought after. Drug design is unfortunately a challenging endeavour because of the sheer size of the drug-like chemical space that has to be explored and the combination of different property requirements that needs to be fulfilled. Faced with such a challenging and important chemical design problem, and in light of the recent success of deep generative models for others modalities such as image, sound, or text, computational chemists turned to the use of deep learning architectures. These artificial neural network architectures are trained on large databases of drug-like molecules and learn to reproduce underlying patterns and chemical motives. Current machine learning efforts to coax these algorithms into producing molecules with optimised properties remain plagued by a significant lack of efficiency.

In this dissertation, we discussed the introduction of a quality-diversity technique called graph-based elite patch illumination that largely resolves efficiency issues in single objective optimisation tasks for molecular design. This algorithm relies on ideas from genetic algorithms but avoids the typical stagnation issues by enforcing diversity of the evolutionary population explicitly. In addition to providing a reliable source of diversity to facilitate efficient optimisation, graph-based elite patch illumination also generates an overview of the potential for optimisation in feature space. To a limited extent, graph-based elite patch illumination has already seen practical use. An early prototype of the algorithm was used to propose several molecules that were synthesised and assayed as potential antiviral compounds against COVID-19 by the COVID Moonshot effort.

In reaction to questions from the cheminformatics community, we tested the capabilities of graph-based elite patch illumination in multi-objective Pareto optimisation problems in a follow-up paper. To have access to baselines for comparison, we implemented open-source and graph-based versions of two non-dominated sorting

genetic algorithms: NSGA-II and NSGA-III. Our benchmarking efforts showed that these algorithms remain state-of-the-art in terms of efficient multi-objective optimisation for molecular design and outperform the graph-based elite patch illumination algorithm. Remarkably, these non-dominated sorting genetic algorithms were also shown not to rely on explicit chemical diversity to obtain efficient optimisation. To calculate the chemical diversity of the evolutionary populations in these algorithms, we made use of extended similarity indices.

Accelerating Biophysically-Detailed Neuron Models

Biophysically-detailed neuron models contain biologically interpretable parameters and allow for the calculation of local field potentials, making them highly relevant for the study of complex neurological disorders. The challenge of using these models mostly lies in their high computational cost: tuning parameters or running simulations of large networks of model neurons requires computational resources on a scale that is often unavailable outside specialised research centers. This problem has largely been solved, in a series of recent papers, by distilling the input-output relationship of a biophysically-detailed neuron model into a comparatively simple artificial neural network architecture. For small network simulations this approach has already been shown to produce a decrease in runtime of near to five orders of magnitude.

In this dissertation, we discussed how to apply the current forefront in artificial neural network architectures for multi-task learning to predict the membrane potential in each compartment of a biophysically-detailed neuron model simultaneously. Voltage traces of multiple compartments within the same neuron are strongly correlated to each other due to the morphology of the underlying neuron model, making this distillation problem an excellent candidate for efficient multi-task learning. On the other hand, the values of these voltages traces follow a non-normal distribution because of the presence of action potentials. This task can, hence, be seen as a stringent test of multi-task learning architectures and specialised optimisation procedures of these architectures. In addition, distillation of a full biophysically-detailed neuron will also allow for the calculation of local field potentials as a post-processing step.

Future Work and Possible Technical Improvements

During the three years discussed in this dissertation, many different research avenues were explored. Due to the usual practical limitations of time, energy and computational resources, only a fraction of the considered research ideas were turned into actual projects. At the same time, it is imperative to be aware of the current shortcomings and opportunities for improvement of the research projects that have been presented in this dissertation. Therefore, we list a limited selection of suggestions for future work below. For clarity, we group these comments by subject as discussed above: disease models, molecular optimisation, and distilling biophysically-detailed neurons.

Disease Models

To obtain a coherent and accurate picture of the biomolecular mechanisms of severe mental disorders, deliberate and thoughtful use of combined in-vitro, in-vivo, and in-silico disease models will be necessary. At the same time, each of these modalities is struggling with its own technical issues, and has its own unique opportunities for improvement.

- The important effects of electrophysiological activity on gene expression [1–3] are often not taken into account when analysing cell cultures or brain organoids, and hence the wide adaptation of patch-sequencing techniques might allow for the discovery of novel mechanisms involved in severe mental disorders.
- Brain organoids currently suffer from issues with regard to cell maturation, reproducibility, perfusable vasculature and development of a blood-brain barrier [4, 5]. Yet, these cell cultures can still serve as a unique platform to study the impact of common variants implicated in neurological diseases and their synergetic effects through the application of CRISPR/Cas9 techniques [6–8].
- Creating animal models of disorders that are both strongly polygenic and reliant on external factors is notoriously difficult [9, 10]. However, the effects of genetic variations in the genes implicated through the SCHEMA consortium [11] should be reflected well by CRISPR/Cas9 knock-out models. Alternatively, cross-species conserved biological functions could be targeted to make novel animal models for mental disorders [12].
- Several obstacles still stand in the way of the common and accurate simulation of large-scale networks of biophysically-detailed neuron models. Whereas challenges regarding speed and computational cost are being resolved, problems with biological accuracy remain. More systematic and standardised electrophysiological measurements will be necessary to make the construction of detailed and accurate networks representative of large brain areas possible [13–15].
- Most current biophysically-detailed neuron models group different ion channels of a similar type together into one ad-hoc ion current model for simplicity. The construction of an accurate ion current model for each ion channel in Channelpedia [16], and the common use of patch-sequencing [17, 18], could facilitate the development of biophysically-detailed neuron models with significantly improved biological detail and accuracy.

Molecular Optimisation

The past years have seen the revival of traditional computational techniques in molecular optimisation and design. Genetic algorithms remain the driving force behind the current state-of-the-art molecular optimisation algorithms, despite the large investments of academia and industry in their deep learning counterparts.

Questions about how to integrate both methodologies to obtain the best of both worlds remain open.

- Genetic algorithms apply mutations and crossovers randomly, without taking rewards or past experience into account. To rectify this waste of information, the choice of variation operations applied to any given molecule could be decided upon through reinforcement learning or a contextual non-linear bandit algorithm. Applications of these ideas have been shown to be effective in quality-diversity algorithms developed for latent space illumination and evolutionary robotics [19, 20].
- In our study of molecular Pareto optimisation, we made use of the dominated hypervolume as performance indicator for the potency of the studied algorithms. However, the dominated hypervolume can also be used as a resource within the optimization process. Algorithms that make use of this approach exists, and could – maybe even should – be tested in the setting of molecular multi-objective optimisation [21, 22].

Distilling Neurons

Even though distilled biophysically-detailed neuron models have been shown to be efficient and relatively accurate, they have not yet been incorporated in simulations containing more than a few thousand neurons. Examples of larger neuron networks include the neuronal microcircuitry of the somatosensory cortex of the juvenile rat (Blue Brain Project) [23] and the mouse primary visual cortex (Allen Institute for Brain Science) models [24]. For the current situation to change, large-scale data acquisition and training efforts will need to be combined with better integration of local field potential calculating tools. In addition, the discovery of memory-efficient task-balancing techniques for multi-task learning will be necessary in complement to such efforts.

- We have taken initial steps to expand deep learning models to be able to simultaneously predict membrane potentials in each compartment of a biophysically-detailed neuron model. In theory, this should suffice to calculate local field potentials. In practice, integration of distilled neuron models with tools such as LFPy still needs to be completed and optimised [25, 26].
- Task balancing can be necessary in multi-task learning to avoid a limited number of tasks from dominating the loss function. A two-step optimization technique from transfer-learning called model agnostic meta-learning has shown great promise in this regard [27, 28]. Unfortunately, this technique is currently limited to models with a small amount of trainable parameters.
- Questions regarding the importance of diversity among experts during and after training remain open. For instance, it remains unclear whether or not diversity measures can be used as effective regularisers for multi-task

learning. As seen in the other projects of this dissertation, there is also an unmet need for a unified framework for measuring and interpreting computational diversity.

Contributions Made During the PhD

In addition to the research discussed above, the three year period covered by these PhD studies also included several other academic contributions such as teaching, supervision and outreach. In addition, 30 credits of coursework were completed as part of the PhD studies. The research discussed in this dissertation was disseminated through several poster presentations, one oral presentation and two peer reviewed publications. An overview of these contributions is provided in this section. For convenience, these contribution are grouped into three categories: (1) teaching, outreach and supervision, (2) research results and publications, and (3) courses and reports.

Teaching, Outreach and Supervision

No formal teaching requirements were part of the PhD contract, but informal contacts led to the participation of the author of this dissertation in minor teaching duties. Every year, the University of California, San Diego, the University of Oslo, and Simula Research Laboratory Simula organise a summer school in computational physiology. In the 2021 edition, the author of this dissertation gave a lecture at this summer school titled *Integrating Omic Data and Computational Simulations in Neurological Models of Schizophrenia*. Similarly, for the Norwegian Artificial Intelligence Research Consortium (NORA), a Norwegian collaboration between eight universities, three university colleges and five research institutes, the author contributed to a workshop on AlphaFold v2.0 and RoseTTAFold with a session on the practical use of these tools. This workshop was hosted by NORA in collaboration with the University of Oslo, dScience, Elixir Norway, Oslo Cancer Cluster and the Centre for Digital Life Norway, and recordings of the sessions can be found online.

Pint of Science is an international organisation that aims to bring researchers from cities across the world into local pubs, bars or public spaces to share and discuss their research findings with a general audience. For Pi-day, celebrated on March 14th 2022, Pint of Science Norway organised a podcast special on the topic of applied maths and biology. Together with Dr. Vegard Vinje (Simula research laboratory) and Assistant Prof. Leiv Øyehaug (OsloMet), the author of this dissertation was invited to discuss how mathematical modelling can be used to better understand biology. The episode is available on most well-known podcast platforms. The Centre for Digital Life Norway is a national centre for biotechnology research, education and innovation focussed on transdisciplinary collaboration. At their annual conference Digital Life Norway organises a scientific image exhibition and competition. In 2022, the author participated in this competition with a 3D rendered image of neuroreceptors.

Since the spring semester of 2022, the author has been involved in the supervision of several master students with regards to their master thesis. At first, this consisted of support for prof. Gaute Einevoll's *spesialpensum* regarding biophysically-detailed

neuron models and specifically the Allen mouse primary visual cortex model. In the academic year 2022-2023, the author took on the role of main supervisor for two masters students (Sebastian Amundsen and Maria Lunde) and the role of secondary supervisor for one masters student (Marcus Berget). Sebastian's research work has focussed on the distillation of biophysically-detailed neuron models as discussed in this dissertation and Maria's work focusses on the distillation of the Victor–Purpura spike synchrony metric into a Siamese neural network. Details of their research will be made available in their respective master theses, which are expected to be submitted in the spring semester of 2023. The topic of Marcus' project is related to extending inception loops to a variety of neural data types.

Research Results and Publications

The research outlined in this dissertation, which was conducting as part of a PhD fellowship awarded by the University of Oslo, resulted in two peer reviewed publications, several poster presentations, and one oral presentation at a conference. Multiple other research projects are in the final stages of completion and are expected to result in publications in the near future. Reproductions of the manuscripts included in the dissertation can be found in the pages directly following the current chapter. A survey of those research publications and presentations is provided here, together with an overview of publications expected in the near future. Where applicable, the author's specific contributions are provided as corresponding to the official co-author declarations.

Illuminating Elite Patches of Chemical Space

This paper introduced quality-diversity techniques to single objective optimisation for small molecules and was published in Chemical Science in September 2020. The paper was co-authored with Dr. Jeriek Van den Abeele. The author of this dissertation was responsible for the initial GB-EPI concept and implementation of the main code, as well as the refactoring of the code into a public Python repository. The author also contributed to discussions regarding the performance studies and led the writing of the paper, especially in sections two and four.

Graph-Based Molecular Pareto Optimisation

This paper introduced the use of chemical diversity and the dominated hypervolume indicator to the study of non-dominated sorting genetic algorithms for Pareto optimisation of small molecules. The paper was published in Chemical Science in June 2022 and is the result of a single-author project. Hence, the author of this dissertation was responsible for all aspects of this paper.

Multi-task Learning of Biophysically-Detailed Neuron Models

In this research project we aim to distil biophysically-detailed neuron models into multi-task learning architectures. The initial concept and implementation have been shared between the author of this dissertation and Kosio Bechkov. Implementation of the hard parameter sharing model was partially carried out by Sebastian Amundsen. Useful references to relevant literature have been provided

by Dr. Torbjørn Vefferstad Ness and Prof. Gaute Einevoll. The project is ongoing and expected to result in a journal publication in the near future.

Isometric Representations in Neural Networks Improve Robustness

In this paper, which is in the process of resubmission, we enforced the conservation of metric structure during the training of artificial neural networks to obtain isometric and robust within-class representations. Kosio Bechkov is the main author of this paper. The paper is co-authored by Dr. Mikkel Elle Lepperød and the author of this dissertation. The latter was responsible for refactoring of parts of the code and supported the writing and original submission of the paper.

Morta: A Multi-Dimensional and Representative Archive of Molecules

For this project, the author of this dissertation and (former) members of the Canadian biotech company Cyclica have come together to create representative down-sampled archives of large molecular databases through the use of quality-diversity techniques and chemical diversity measures. Details regarding the distribution of responsibilities in this public-private partnership and a manuscript describing the obtained results will be available in the near future, as will open-source results and the code generated for this project.

Oral Presentation at the Nordic AI Meet

In November 2022, the author of this dissertation presented an oral presentation titled *Distilling Computational Neuroscience* at the second Nordic Conference for Young AI researchers (Nordic AI Meet) in Oslo. In this talk, the previously discussed work on distilling biophysically-detailed neuron models and the Victor–Purpura spike synchrony metric were presented. This presentation paper was co-authored by Kosio Bechkov and Prof. Gaute Einevoll, both of whom supported abstract submission and design of the slides for the presentation.

Poster Presentations at the FENS Forum and the Bernstein Conferences

The author of this dissertation had abstracts for poster presentations accepted at the Federation of European Neuroscience Societies (FENS) Forum in Paris and the Bernstein Conferences in Berlin, which took place in July and September of 2022 respectively. Both poster presentations covered topics related to biophysically-detailed neuron models, disease models for severe mental disorders, and the 4MENT collaboration in general.

Courses and Reports

As part of a PhD Degree at the Faculty of Mathematics and Natural Sciences at the University of Oslo, 30 credits of coursework and a scientific dissertation have to be successfully completed. The written part of the dissertation is presented in this document, and the oral part will be presented in a trial lecture and public defence. The 30 credits of coursework were obtained through the successful completion of the following courses: Science, Ethics and Society (MNSES9100, five credits), Communicating Scientific Research (Simula, five credits), Molecular Medicine (MF9120BTS, ten credits), and

Computational Neuroscience (FYS388, ten credits). As part of these courses, a number of reports were written on topics ranging from vaccine equity to personalised medicine.

Concluding Remarks

As discussed throughout this dissertation, schizophrenia and bipolar disorder are among the most human of diseases. No other species of animal is known to suffer from either of these disorders and important genetic risk factors are lodged deep inside the most human parts of our DNA. Yet, for much of recorded history, severe mental disorders and the behaviour they cause were seen as a sign of the supernatural. Even though these beliefs gradually faded, it took until the turn of the 20th century before moral judgements were sufficiently contained to earnestly apply the scientific method to mental disorders. Initial scientific progress was fast – only 44 years passed between the coining of the term *schizophrenia* by the the Swiss psychiatrist Paul Eugen Bleuler and the discovery of the first antispychotic medication – but in the past decades, we have witnessed painful stagnation in the study and treatment of psychosis and mood disorders.

Current antipsychotic treatments can only help a small fraction of patients affected by severe mental disorders and simultaneously carry a risk of intense and potentially deadly side-effects. On average, patients diagnosed with severe mental disorders suffer from a reduction in life expectancy of more than a decade and the gap in life expectancy with the general population keeps growing. A student once famously asked the anthropologist Margaret Mead what she considered to be the first sign of civilization. Her answer was simple: *A healed femur.* Mead explained that the first sign of civilization is neither a clay pot nor an early weapon but the emergence of compassion, exemplified in the protection and companionship necessary to mend a complex bone fracture. Whether or not history will judge the first decades of the 21st century as a turning point or a failure in humanity's quest to extend this care to its most vulnerable members is yet to be decided.

Will we find the courage and humility to accept that severe mental disorders can affect any of us – directly or indirectly – and invest in accessible and life-long psychosocial support for society at large? Will we fund dedicated interdisciplinary research groups and create extensive public-private partnerships to bring novel treatment plans to patients? Will we outgrow our childhood fears and learn to treat severe mental disorders without stigma or distrust? Only time will tell... but we have to rely on that old adage, *scientia vincere tenebras*, and keep trying to cure or prevent these diseases. Now more than ever, science and society have the tools and knowledge at their disposal for great improvement in prophylactic and curative treatments of schizophrenia and bipolar disorder. This much I know to be true.

References

1. Lee, P. R. & Fields, R. D. Activity-dependent gene expression in neurons. *The Neuroscientist* vol. 27, 355–366 (2021).

2. Iacobas, D. A., Iacobas, S., Lee, P. R., Cohen, J. E. & Fields, R. D. Coordinated activity of transcriptional networks responding to the pattern of action potential firing in neurons. *Genes* vol. 10, 754 (2019).

3. Holt, C. E., Martin, K. C. & Schuman, E. M. Local translation in neurons: visualization and function. *Nature structural & molecular biology* vol. 26, 557–566 (2019).

4. Dixon, T. A. & Muotri, A. R. Advancing preclinical models of psychiatric disorders with human brain organoid cultures. *Molecular Psychiatry,* 1–13 (2022).

5. Koo, B., Choi, B., Park, H. & Yoon, K.-J. Past, present, and future of brain organoid technology. *Molecules and cells* vol. 42, 617 (2019).

6. Kurishev, A. O., Karpov, D. S., Nadolinskaia, N. I., Goncharenko, A. V. & Golimbet, V. E. CRISPR/Cas-Based Approaches to Study Schizophrenia and Other Neurodevelopmental Disorders. *International Journal of Molecular Sciences* vol. 24, 241 (2022).

7. Rajarajan, P., Flaherty, E., Akbarian, S. & Brennand, K. J. CRISPR-based functional evaluation of schizophrenia risk variants. *Schizophrenia research* vol. 217, 26–36 (2020).

8. Schrode, N. *et al.* Synergistic effects of common schizophrenia risk variants. *Nature genetics* vol. 51, 1475–1485 (2019).

9. Richtand, N. M., Harvey, B. H. & Hoffman, K. L. Animal Models in Psychiatry: Translating Animal Behavior to an Improved Understanding and Treatment of Psychiatric Disorders. *Frontiers in Psychiatry* vol. 13, 876155–876155 (2022).

10. Nestler, E. J. & Hyman, S. E. Animal models of neuropsychiatric disorders. *Nature neuroscience* vol. 13, 1161–1169 (2010).

11. Singh, T. *et al.* Rare coding variants in ten genes confer substantial risk for schizophrenia. *Nature* vol. 604, 509–516 (2022).

12. Sall, S., Thompson, W., Santos, A. & Dwyer, D. S. Analysis of major depression risk genes reveals evolutionary conservation, shared phenotypes and extensive genetic interactions. *Frontiers in psychiatry* vol. 12, 1180 (2021).

13. Hay, E., Hill, S., Schürmann, F., Markram, H. & Segev, I. Models of neocortical layer 5b pyramidal cells capturing a wide range of dendritic and perisomatic active properties. *PLoS computational biology* vol. 7, e1002107 (2011).

14. Gouwens, N. W. *et al.* Systematic generation of biophysically detailed models for diverse cortical neuron types. *Nature communications* vol. 9, 1–13 (2018).

15. Masoli, S. *et al.* Single neuron optimization as a basis for accurate biophysical modeling: the case of cerebellar granule cells. *Frontiers in cellular neuroscience* vol. 11, 71 (2017).

16. Ranjan, R. *et al.* Channelpedia: an integrative and interactive database for ion channels. *Frontiers in neuroinformatics* vol. 5, 36 (2011).

17. Mahfooz, K. & Ellender, T. J. in *Patch clamp electrophysiology* 179–189 (Springer, 2021).

18. Lipovsek, M. *et al.* Patch-seq: Past, present, and future. *Journal of Neuroscience* vol. 41, 937–946 (2021).

19. Fontaine, M. & Nikolaidis, S. Differentiable quality diversity. *Advances in Neural Information Processing Systems* vol. 34, 10040–10052 (2021).

20. Tjanaka, B., Fontaine, M. C., Togelius, J. & Nikolaidis, S. *Differentiable Quality Diversity for Reinforcement Learning by Approximating Gradients* in *ICLR Workshop on Agent Learning in Open-Endedness* (2022).

21. Beume, N., Naujoks, B. & Emmerich, M. SMS-EMOA: Multiobjective selection based on dominated hypervolume. *European Journal of Operational Research* vol. 181, 1653–1669 (2007).

22. Bader, J. & Zitzler, E. HypE: An algorithm for fast hypervolume-based many-objective optimization. *Evolutionary computation* vol. 19, 45–76 (2011).

23. Markram, H. *et al.* Reconstruction and simulation of neocortical microcircuitry. *Cell* vol. 163, 456–492 (2015).

24. Billeh, Y. N. *et al.* Systematic integration of structural and functional data into multi-scale models of mouse primary visual cortex. *Neuron* vol. 106, 388–403 (2020).

25. Hagen, E., Næss, S., Ness, T. V. & Einevoll, G. T. LFPy–multimodal modeling of extracellular neuronal recordings in Python. *bioRxiv,* 620286 (2019).

26. Hagen, E., Næss, S., Ness, T. V. & Einevoll, G. T. Multimodal modeling of neural network activity: computing LFP, ECoG, EEG, and MEG signals with LFPy 2.0. *Frontiers in neuroinformatics* vol. 12, 92 (2018).

27. Finn, C., Abbeel, P. & Levine, S. *Model-agnostic meta-learning for fast adaptation of deep networks* in *International conference on machine learning* (2017), 1126–1135.

28. Aoki, R., Tung, F. & Oliveira, G. L. Heterogeneous multi-task learning with expert diversity. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* vol. 19, 3093–3102 (2022).

Papers

Paper I

Illuminating Elite Patches of Chemical Space

I

Jonas Verhellen, Jeriek Van den Abeele

Published in *Chemical Science*, September 2020, Volume 11, Issue 42, pp. 11485–11491. DOI: 10.1039/D0SC03544K.

Abstract

In the past few years, there has been considerable activity in both academic and industrial research to develop innovative machine learning approaches to locate novel, high-performing molecules in chemical space. Here we describe a new and fundamentally different type of approach that provides a holistic overview of how high-performing molecules are distributed throughout a search space. Based on an open-source, graph-based implementation [J. H. Jensen, Chem. Sci., 2019, 10, 3567–3572] of a traditional genetic algorithm for molecular optimisation, and influenced by state-of-the-art concepts from soft robot design [J. B. Mouret and J. Clune, Proceedings of the Artificial Life Conference, 2012, pp. 593–594], we provide an algorithm that (i) produces a large diversity of high-performing, yet qualitatively different molecules, (ii) illuminates the distribution of optimal solutions, and (iii) improves search efficiency compared to both machine learning and traditional genetic algorithm approaches.

I.1 Introduction

Recent years have seen a surge [1–10] of machine learning (ML) papers focused on generating *de novo* molecules optimised for performance with regard to a chosen objective function, e.g. melting point [11] or binding affinity to a target protein [12]. These ML models aim to generate chemical compounds which exhibit desired behaviour, without reverting to explicit chemical rules, patterns or transformations. Instead, ML models learn from experimental data, and attempt to extrapolate the relevant aspects of the underlying chemistry. In terms of performance, however, ML models for molecular optimisation are rivalled by more traditional and often simpler, rule-based approaches [13, 14] such as genetic algorithms (GA).

In this paper, we introduce a novel rule-based algorithm which we call graph-based elite patch illumination (GB-EPI). This algorithm enforces diversity among a set of high-performing molecules and leverages [15–18] them to obtain efficient optimisation. In addition, GB-EPI provides the user with a map relating the performance of generated molecules to chosen physicochemical properties. The algorithmic methodology of GB-EPI is discussed in the section, followed by results of standard benchmarks and an in-depth comparative efficiency analysis between a graph-based genetic algorithm (GB-GA) and GB-EPI.

I.2 Algorithmic Methodology

The goal of a classical optimisation algorithm is to obtain the highest performing solution in a search space. If the exact mathematical form of the evaluation function is inaccessible, as is typically the case in molecular optimisation, heuristic search methods [19] become a necessity. Many of these heuristic methods are inspired by biological phenomena. Genetic algorithms [20, 21] are based on the theory of evolution and aim to optimise with regard to an evaluation function, incrementally improving on existing solutions. Specifically, novel solutions are generated by randomly changing or stochastically combining solutions from the existing population. In the genetic algorithm community, these two operations are respectively known as *mutations* and *crossovers*. Solutions found by genetic algorithms are called *phenotypes* and each solution is described by an underlying *genome*. The performance of a solution with respect to the chosen evaluation function is known as the *fitness* of a phenotype.

Genetic algorithms can be highly effective in straightforward optimisation problems, but are known to struggle [22, 23] when trying to cross low-performing valleys or to break out of local optima, and both of these occurrences can lead to evolutionary stagnation. We have based GB-EPI on an existing genetic algorithm for molecular optimisation, but evade evolutionary stagnation by enforcing molecular diversity. Moreover, GB-EPI speeds up the optimisation process by decoupling mutations from crossovers, and introduces the concept of positional analogue scanning to genetic algorithms. These and other technical aspects[1] of GB-EPI are discussed in the upcoming paragraphs.

[1] A lightweight, open-source version of the GB-EPI algorithm is available for download at https://github.com/Jonas-Verhellen/argenomic.

I.2.1 Graph-Based Genetic Algorithm

The current leading rule-based model for molecular optimisation is the graph-based genetic algorithm [14] (GB-GA). In GB-GA, genomes of molecules are encoded by their molecular graphs. Novel molecules are generated by mutating or combining the graphs of molecules in the existing population. The initial population of candidate molecules is typically obtained from freely accessible molecular data-sets like ZINC [24] or ChEMBL [25]. Every generation, as a form of selection pressure, only the most fit molecules (with respect to the evaluation function) present in the population are retained.

This paper, and hence our algorithm, builds on the conceptual developments made in GB-GA by continuing to work with molecular graphs as genomes. We maintain the graph-based aspect of the crossover and mutation operators, but apply crossovers and mutations in parallel instead of sequentially. Our motivation for decoupling these two operators lies in the fact that crossovers customarily only support efficient exploration of chemical space in the early generations of a genetic algorithm. Later on, the nearly-converged solutions are typically only improved by the comparatively smaller effects of mutations.

I.2.2 Multi-Dimensional Archive of Phenotypic Elites

The multi-dimensional archive of phenotypic elites algorithm [16] (MAP-Elites) is a simple, efficacious and surprisingly powerful tool developed in the context of soft robot design, and serves as the core architecture of our GB-EPI algorithm. MAP-Elites mimics diversity in biological evolution and explores the search space by introducing the concept of niches [26–28] to genetic algorithms. In MAP-Elites, candidate solutions are generated by a genetic algorithm but are assigned to different niches depending on their characterising features. Each generation, the best performing solution in each of the individual niches – with respect to a global evaluation function – is retained.

Dividing the search space into feature-based niches and explicitly enforcing population diversity stands in stark contrast with classical genetic algorithms which typically only retain the top high-scoring solutions regardless of their diversity, or lack thereof. The enforced variation between niches makes crossovers more diverse, and by mutating existing solutions, potent scaffolds can spread into other niches. Most importantly, because at every generation MAP-Elites contains solutions spread out over feature space, diverse solutions in far-away niches can be used a resource to escape stagnation. In Figure I.1 we provide pseudocode of the MAP-Elites algorithm for *de novo* molecule design as applied in GB-EPI.

In practical terms, users of GB-EPI can choose their own features of interest, and define relevant ranges of variation to construct a feature space. If, for instance, a user wants to find medicinally relevant molecules in chemical space, they could construct a feature space based on physicochemical properties like lipophilicity and molecular mass, and practical concerns like synthetic accessibility. The chosen ranges in which to explore these features can be used to specify a desired subset of chemical space in which to generate new molecules.

The fitness score obtained by the molecule occupying a niche at the end of a GB-EPI run represents the capability of the corresponding part of feature space to contain high-

Algorithm: MAP-Elites for molecule generation in GB-EPI

Input: N – number of generations, $\mathscr{M}_0$ – initial population
$\mathscr{P}_0 \leftarrow \mathsf{fitness}(\mathscr{M}_0)$;
for $i = 1 \rightarrow N$ **do**
  $\mathscr{M}_i \leftarrow \mathscr{M}_{i-1}, \mathscr{P}_i \leftarrow \mathscr{P}_{i-1}$;
  $\mathscr{M}' \leftarrow \mathsf{mutation}(\mathscr{M}_i) + \mathsf{crossover}(\mathscr{M}_i)$;
  for *molecule* **in** $\mathscr{M}'$ **do**
    niche $\leftarrow \mathsf{features}$(*molecule*);
    performance $\leftarrow \mathsf{fitness}$(*molecule*);
    if *performance* > $\mathscr{P}_{i-1}$[*niche*] **then**
      $\mathscr{M}_i$[*niche*] $\leftarrow$ *molecule*;
      $\mathscr{P}_i$[*niche*] $\leftarrow$ *performance*;
  end
end
Result: $\mathscr{M}_N$ – molecules, $\mathscr{P}_N$ – performances

Figure I.1: Pseudocode description of the MAP-Elites algorithm adapted to the setting of molecular optimisation.

performance molecules. In this way, GB-EPI *illuminates* the relationship between the chosen features of interest and how varying them affects performance, either positively or, equally relevant, negatively. As can be seen in Figure I.2, the molecules at the end of a GB-EPI run form a patchwork of locally elite solutions (with respect to the chosen evaluation function) in a part of chemical space.

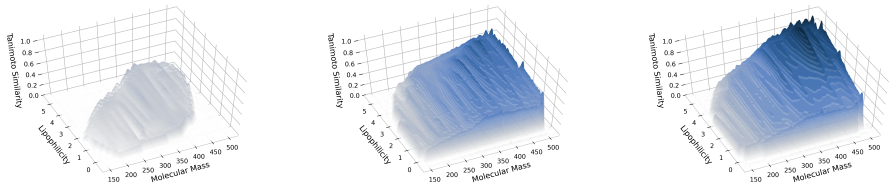

(a) (b) (c)

Figure I.2: Illumination of a patch of elite solutions for the rediscovery of Troglitazone after (a) 1 generation, (b) 200 generations, and (c) 400 generations. For this visualisation, the feature space of GB-EPI was spanned by molecular mass and lipophilicity, and divided into 200 niches. The starting population consisted of 100 random compounds from a standardised subset of the ChEMBL database, further described in Section II.3. The surface was obtained by interpolating, refining and triangulating the results. Darker shading indicates higher Tanimoto similarities with respect to Troglitazone.

I.2.3 Centroidal Voronoi Tesselations

In a regular grid partition, the number of niches grows exponentially with the dimensionality of feature space. To effectively partition high-dimensional feature spaces into niches, we can rely on a technique from computational geometry called the centroidal Voronoi tessellation [29–31] (CVT). The CVT can be used to create a pre-defined number of niches, irrespective of the dimensionality of the feature space. Because the number of niches is fixed, the use of a CVT partition in MAP-Elites maintains selection pressure for performance, even in higher-dimensional feature spaces [32].

A CVT is constructed by forming the lattice reciprocal to the cluster centroids of a uniform distribution over feature space. Each of the lattice cells outlines the space contained in a single niche. Computationally, the centroidal Voronoi tessellation can be constructed by Lloyd's clustering algorithm [33]. Efficient look-up of the nearest centroid to a given point in feature space is necessary to determine the niche to which a new solution belongs. Fortunately, this is made possible through fast multi-dimensional tree algorithms [34].

I.2.4 Positional Analogue Scanning and Memoisation

Changes in molecular interactions and physicochemical properties resulting from small molecular structure modifications are used in *in vitro* medicinal chemistry to optimise lead compounds [35]. To minimise the number of experimental design cycles in lead optimisation, medicinal chemists apply small structure modifications in systematic batches, in a procedure known as positional analogue scanning [36]. During this procedure, series of molecular analogues of a lead compound are generated by the systematic exchange of heteroatoms or functional groups, and rapidly evaluated.

Similar to the small structure modifications used in the lab, GB-GA uses molecular mutations to work towards compounds with desired properties. Inspired by the success of positional analogue scanning, we repurpose the mutation operator in GB-EPI to systematically return not just a single mutated molecule, but all of its positional analogues. This approach accelerates convergence by allowing a potent design to spread out to several niches in a single generation. To speed up convergence even further, we extend the mutation operator to allow for the addition and removal of user-specified functional groups.

Memoisation [37] is a computational technique that ensures that a program does not unnecessarily repeat calculations, by keeping an on-the-fly record of obtained results. To balance memory and efficiency, the set of remembered results is typically limited to a fixed size and controlled by a first-in-first-out replacement algorithm. In this paper, memoisation was applied to fitness calculation, as this often carries the prohibitive computational cost, but memoisation can be readily extended to the other calculations in the algorithm. We note that memoisation can be used to reduce or even fully resolve the computational overhead introduced by positional analogue scanning.

I.2.5 Filters and Parallelism

To rule out unwanted and potentially toxic molecules, we use functional group knowledge from the ChEMBL database [25] and a combination of ADME property calculations [38–40]. We remove undesirable compounds before they enter the evaluation step of the algorithm. Removing these compounds at an early stage makes the algorithm more efficient, increases the predictive value of the final outcome, and significantly decreases overall processing time.

To reduce clock time, we implemented a concurrent version of GB-EPI. The program distributes function evaluations, mutations and crossovers over a CPU/GPU architecture and receives performance scores, new molecules and behavioural descriptors from the individual nodes. Concurrency has no effect on the overall results obtained by the algorithm. All of the experiments in this paper can be reproduced either with or without concurrency.

I.3 Results and Benchmarks

To standardise the assessment of models for *de novo* molecular design, the bioinformatics company BenevolentAI released a benchmarking suite named GuacaMol [13]. The suite is open source and is meant to provide researchers with a variety of molecular optimisation tasks, related to the basic needs of computational and medicinal chemists. In this paper, we use GuacaMol as a starting point to quantify the performance of GB-EPI. We present and compare the results on the selected benchmarks for a deep-learning algorithm (SMILES LSTM), a rule-based algorithm (GB-GA), and the illumination algorithm GB-EPI presented in this paper.

SMILES LSTM [41] is a deep learning model for *de novo* molecule generation, based on natural language processing and reinforcement learning. SMILES LSTM uses a simple text representation of molecules known as Simplified Molecular-Input Line-Entry System [42] (SMILES) strings and trains a recurrent neural network (RNN) as a statistical language model for these textual descriptors of molecular structures. To obtain numerical stability in training through back-propagation, the RNN is enhanced with long short-term memory [43] (LSTM) cells, making it capable of learning dependencies from larger collections of information.

After the SMILES LSTM model is sufficiently trained to produce chemically feasible SMILES strings, reinforcement learning [44] is applied to bias the generation of new chemical structures towards molecules with the desired chemical properties. Reinforcement learning is powerful, yet brittle; initialisation of the underlying LSTM network and the hyperparameters of the reinforcement learning algorithm must be done carefully. If successful, however, SMILES LSTM is able to cover and explore a large portion of chemical space [45].

In this paper, we run both SMILES LSTM and GB-GA in their standard GuacaMol baseline implementations [13]. In particular, for each rediscovery target, the GB-GA algorithm was run with a mating pool of 200 molecules for a total of 1,000 generations, unless there was no improvement for 5 consecutive generations. The SMILES LSTM baseline is a pre-trained recurrent neural network model, further optimised for each specific benchmark over 20 epochs by means of a hill-climbing algorithm. Each epoch

the model generates 8192 molecules, of which the best 1024 are used to steer the reinforcement learning algorithm for further tuning.

I.3.1 Rediscovery of Small Molecule Drugs

Rediscovery benchmarks, which require the explicit rediscovery of a target molecule on top of scoring for similarity, are a common potency test for *de novo* molecule generating models. By requiring explicit rediscovery, these benchmarks are more robust against exploitation [46] of metric deficiencies by generative models than - for instance - similarity metrics with a thresholded linear score modifier [13]. The similarity between a generated molecule and the target compound is determined by the Tanimoto similarity of their extended-connectivity fingerprints [47] (ECFPs).

ECFPs are circular topological fingerprints, meaning that they encode molecular structures in terms of concentric atomic neighbourhoods. These fingerprints were originally [48] designed for similarity searching in high-throughput screening, but have also found applications in chemical clustering and compound library analysis. The main advantage of ECFPs, compared to more involved similarity measures, is that they can be rapidly calculated and inherently represent the presence or absence of molecular substructures.

In GuacaMol, three marketed and FDA-approved drugs are proposed as targets for rediscovery: Celecoxib (an anti-inflammatory), Troglitazone (an antidiabetic), and Thiothixene (an antipsychotic). Together, these three ligands cover a wide range of physicochemical properties and pharmacological applications. To increase the effectiveness of the benchmarks, molecules highly similar to the targets (bit-vector Tanimoto similarity above 0.323) were removed by GuacaMol from the database of initial molecules. That initial database is derived from ChEMBL, which exclusively consists of molecules that have both been synthesised in a lab and tested against biological targets.

To set up GB-EPI for the rediscovery benchmarks, we chose the feature space to be spanned by molecular mass, 140 u. to 520 u., and lipophilicity, $\log P = -0.4$ to $\log P = 5.6$. The ranges were chosen to roughly correspond to properties of orally active drugs, and the space was feature subdivided into 150 niches. More complex, higher-dimensional feature spaces are possible and often advisable, but here we limit the algorithm to its simplest form. The number of generations for GB-EPI was limited to a maximum of 400.

GB-EPI is successful in rediscovering these three drug-like molecules, just as SMILES LSTM and GB-GA. Whereas the power to differentiate between models through these GuacaMol rediscovery tasks can hence be debated, these simple tasks do give insight in the properties of the algorithms. The letter-value plots [49] in Figure I.3 show that the three distributions obtained by the algorithms at the end of each of the GuacaMol rediscovery benchmarks are highly distinct from each other. Whereas the GB-GA population provides a concentrated group of high-scoring molecules, SMILES LSTM generates a broad distribution of molecules with a few high-scoring outliers.

GB-EPI combines diversity with local selection pressure, and the obtained population distributions reflect this by having median scores above those of SMILES LSTM, and a more balanced spread than the distributions of GB-GA. While GB-GA

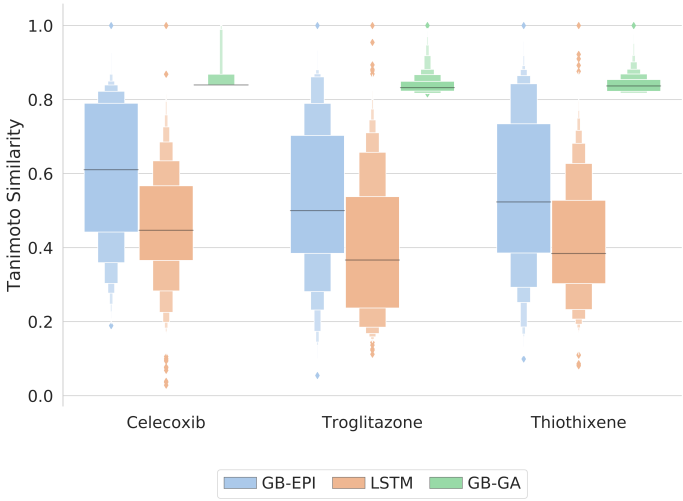

Figure I.3: Letter-value plots [49] of the final molecule distributions obtained by GB-EPI, SMILES LSTM, and GB-GA for the GuacaMol rediscovery benchmarks in terms of Tanimoto similarity to the target. The length of the innermost box represents the interquartile range, whereas the protruding boxes represent subsequent interquantiles (i.e. interoctiles, intersedecimiles, ...). The horizontal line marks the median, while outliers (conventionally assumed to be the outer 0.7% of the population) are shown as individual diamonds beyond the largest interquantile displayed.

only retains the highest-scoring molecules in its population, GB-EPI deliberately keeps lower-scoring molecules that are the best in their niche. In fact, the GB-GA median lies near the bottom of the narrow interquartile range because most of the molecules proposed by GB-GA have high internal similarity [22] and hence nearly identical scores.

I.3.2 Simultaneous Similarity for Conflicting Compounds

In a median molecules benchmark, the goal is to maximise similarity to several small-drug molecules simultaneously. The standard GuacaMol benchmark starts from the highest scoring molecules in the ChEMBL subset described in Section I.3.1. These benchmarks are explicitly designed to be conflicting and can be regarded as challenging tasks. The GuacaMol benchmarking suite provides two of these tasks: Camphor vs. Menthol (two topical antitussives) and Tadalafil vs. Sildenafil (two drugs used to treat erectile dysfunction and pulmonary hypertension).

To increase the real-world relevance of these benchmarks, we filter out molecules that contain macrocycles, fail at Veber's rule [40], or raise structural alerts from ChEMBL. The feature space of GB-EPI was again chosen to be spanned by lipophilicity and molecular mass. For both benchmarks, the feature space of GB-EPI was divided into 200 niches and the algorithm ran for 600 iterations. Furthermore, the GB-GA algorithm was only halted after 50 consecutive iterations without progress.

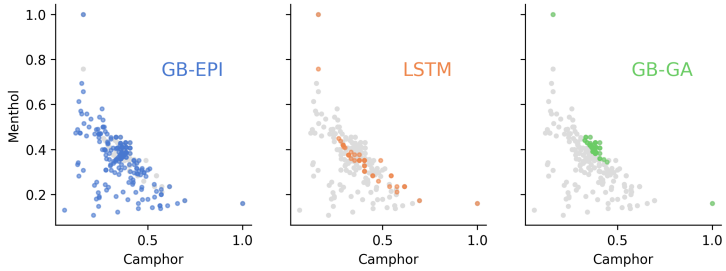

(a) Camphor vs. Menthol

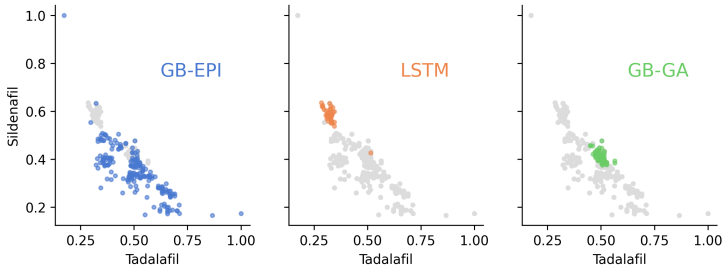

(b) Tadalafil vs. Sildenafil

Figure I.4: Distribution of proposed median molecules – coloured and highlighted by algorithm type – for the conflicting targets in the GuacaMol benchmarks, after filtering out structurally problematic molecules from the 100 highest-scoring ones. For Camphor vs. Menthol, the ranges of feature space for GB-EPI were chosen to be $\log P = -0.4$ to 5.6, and 100 u. to 350 u. For Tadalafil vs. Sildenafil, the ranges were $\log P = -0.4$ to 5.6, and 350 u.to 600 u. GB-EPI's inherent strategy to explore broader swaths of chemical space in an optimisation problem is clear in both figures. In contrast, the molecules proposed by GB-GA are focused around small regions of high-scoring median molecules.

As shown in Table I.1 and Figure I.4, these median molecules benchmarks are far more strenuous than the rediscovery benchmarks and can differentiate between the different models more accurately. Here, SMILES LSTM scores lower than the rule-based algorithms GB-GA and GB-EPI. To ensure an accurate comparison between the three generative models, two of which are pure optimisation algorithms (SMILES LSTM, GB-GA) and one of which (GB-EPI) balances quality and diversity, we only recorded the single highest score obtained by each algorithm.

To make the benchmark more informative, we also recorded the results for all algorithms on both benchmarks for a completely random subset of the standardised dataset. In the randomised subset benchmarks, GB-GA and GB-EPI begin with 100 arbitrary compounds, whereas the SMILES-LSTM model is pre-trained on a larger set of molecules from the same collection but not hyper-tuned by top scoring molecules from the dataset. Both SMILES LSTM and GB-GA have trouble crossing the larger distance in chemical space to the median molecules and score significantly lower than

Table I.1: Results for the Maximum Median Molecule

Benchmark	GB-EPI	SMILES LSTM	GB-GA
Standard			
Camphor vs. Menthol	0.419	0.415	0.419
Tadalafil vs. Sildenafil	0.453	0.422	0.453
Randomised			
Camphor vs. Menthol	0.419	0.400	0.345
Tadalafil vs. Sildenafil	0.370	0.368	0.313

GB-EPI.

I.3.3 Comparing Efficiency of GB-EPI and GB-GA

To study the difference in efficiency of GB-EPI and GB-GA, we make a statistical analysis of a representative rediscovery task (Troglitazone). In line with earlier work [14, 50] on the efficiency of GB-GA, we calculate the average number of fitness function evaluations and CPU time needed for rediscovery, and the rediscovery success rate of both algorithms. As we learned from the median molecule task, starting from a randomised set of molecules elucidates the exploratory power of the algorithms more.

Therefore, we start this rediscovery task with the 100 top-scoring molecules from 10,000 molecules randomly chosen from a 1.6 million ChEMBL subset, as constructed by Henault et al. [50]. In this subset all molecules with a bit-vector Tanimoto similarity to the target above 0.323 are removed [13]. Table I.2 shows the results for 100 runs of GB-EPI and GB-GA (with settings taken from Henault et al. [50]), both with a maximum of 1,000 generations per run.

Table I.2: Efficiency of GB-EPI and GB-GA in the rediscovery of Troglitazone, in terms of the average number of required score evaluations and CPU time in the case of a successful run, and the overall success ratio over 100 independent, randomly seeded runs of both algorithms.

Algorithm	Evaluations	CPU Time	Success Ratio
GB-EPI	14,258	3 min. 5 sec.	100 %
GB-GA	24,216	11 min. 37 sec.	81 %

While chemical space consists of an estimated 10^{60} molecules, it has been argued [50] that the perfect, omnipotent search algorithm would be able to find small drug-like molecules (i.e. excluding peptides, antibodies, ...) in a few hundred transformation operations (crossovers and mutations) and corresponding fitness evaluations. With this idealised benchmark in mind, it can be observed from Table I.2 and Fig. I.5 that GB-EPI makes a sizeable improvement (approx. 41%) to the average number of function evaluations needed for rediscovery. Similarly, we note that the average CPU time needed for rediscovery decreased starkly (approx. 73%) in GB-EPI compared to GB-GA.

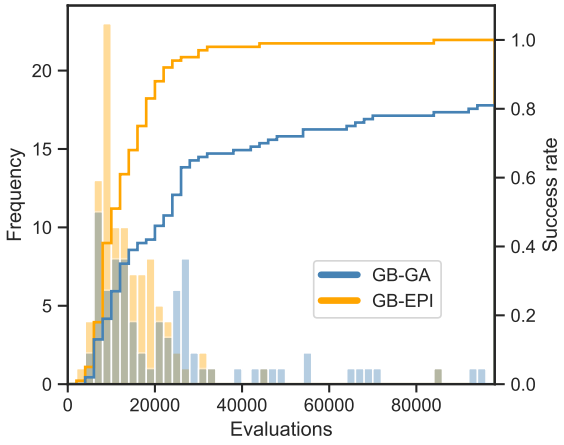

Figure I.5: Distribution of the number of score function evaluations necessary for the rediscovery of Troglitazone and corresponding cumulative success rate, for 100 independent runs of GB-GA (blue) and GB-EPI (orange). Both distributions are shown on the same scale.

In addition, the success ratios affirm that GB-GA suffers from stagnation issues, whereas GB-EPI can leverage molecular diversity to escape local optima of the scoring metric. The success rate of GB-GA for this rediscovery is 81%, meaning that at least 3 GB-GA searches are needed for the rediscovery to succeed with at least 99% certainty. Taking this into account would further increase the number of score evaluations to about 70.000 before an expected successful rediscovery. Similarly the expected CPU time before rediscovery by GB-GA will be of the order of 35 minutes.

I.4 Conclusion and Outlook

This paper introduces the concept of illumination to *de novo* molecule generating algorithms through an algorithm called GB-EPI. Previous molecular optimisation algorithms, like SMILES LSTM and GB-GA, aim to obtain the highest performing solution in chemical space. In contrast, our novel algorithm constructs a whole patch of high-performing solutions spread out over niches covering a selected part of chemical space. By exploring what is chemically possible, in addition to leveraging diversity to efficiently discover what is purely optimal, GB-EPI illuminates design trade-offs and encourages synergy between design algorithms and human chemists.

For instance, researchers wishing to understand how the binding affinity with a target protein changes with physicochemical properties of an inhibitor could use GB-EPI to scan a feature space spanned by the lipophilicity, molar refractivity, and mass of the candidate molecules. In contrast, an industrial chemist could find more use in a feature space spanned by estimated production costs and synthetic accessibility. In both cases,

molecules that are predicted to have a desired combination of properties can easily be selected for further examination.

Future extensions of GB-EPI could include adaptive meshing of the centroidal Voronoi tessellations [51] to increase the number of niches in the most suitable regions of feature space, surrogate modelling techniques [52, 53] to reduce the number of necessary fitness function evaluations, or crossovers based on intermolecular correlations [54]. In addition, deep learning models could be trained to predict which mutations are most beneficially applied to which molecules. Combined, these extensions have the potential to significantly speed up the current GB-EPI algorithm.

Some attention should also be drawn to the exciting prospect of steering GB-EPI by direct experimental feedback. Through active learning [55] – a small-data alternative to deep learning – and graph-based retrosynthesis [56, 57], molecules proposed by GB-EPI could be selected for *in vitro* synthesis and analysis.[2] The experimental results could then be used to update the fitness model. The practical aspects of this iterative loop could perhaps even be executed autonomously by a robotics platform, creating a self-driving laboratory [59] for molecular design.

Conflicts of interest

There are no conflicts to declare.

Acknowledgements

The authors wish to acknowledge useful feedback on this manuscript by P. Coppin. Jonas Verhellen was supported by the UiO:Life Science convergence environment 4MENT.

References

1. Moret, M. *et al.* Generative molecular design in low data regimes. *Nature Machine Intelligence* vol. 2, 171–180 (Mar. 2020).
2. Schneider, P. *et al.* Rethinking drug design in the artificial intelligence era. *Nature Reviews Drug Discovery* (Dec. 2019).
3. Zhavoronkov, A. *et al.* Deep learning enables rapid identification of potent DDR1 kinase inhibitors. *Nature Biotechnology* vol. 37, 1546–1696 (2019).
4. Popova, M., Isayev, O. & Tropsha, A. Deep reinforcement learning for de novo drug design. *Science Advances* vol. 4, eaap7885 (2018).
5. Segler, M. H. S. *et al.* Generating Focused Molecule Libraries for Drug Discovery with Recurrent Neural Networks. eng. *ACS Central Science* vol. 4, 120–131 (Jan. 2018).

[2] A preliminary version of the GB-EPI algorithm was used to propose *de novo* molecules for inhibiting the SARS-CoV-2 main protease, which were selected by the COVID Moonshot initiative [58] to be synthesised and analysed in activity assays.

6. Gómez-Bombarelli, R. *et al.* Automatic Chemical Design Using a Data-Driven Continuous Representation of Molecules. *ACS Central Science* vol. 4. PMID: 29532027, 268–276 (2018).
7. Kusner, M. J. *et al.* Grammar Variational Autoencoder. *International Conference on Machine Learning* (2017).
8. Smalley, E. AI-powered drug discovery captures pharma interest. *Nature Biotechnology* vol. 35, 604–605 (July 2017).
9. Manglik, A. *et al.* Structure-based discovery of opioid analgesics with reduced side effects. *Nature* vol. 537, 185–190 (Sept. 2016).
10. Reymond, J.-L. The Chemical Space Project. *Accounts of Chemical Research* vol. 48. PMID: 25687211, 722–730 (2015).
11. Popova, M. *et al.* MolecularRNN: Generating realistic molecular graphs with optimized properties. *arXiv e-prints,* arXiv:1905.13372 (2019).
12. Elton, D. C. *et al.* Deep learning for molecular design—a review of the state of the art. *Molecular Systems Design & Engineering* vol. 4, 828–849 (2019).
13. Brown, N. *et al.* GuacaMol: Benchmarking Models for de Novo Molecular Design. *Journal of Chemical Information and Modeling* vol. 59, 1096–1108 (2019).
14. Jensen, J. H. A graph-based genetic algorithm and generative model/Monte Carlo tree search for the exploration of chemical space. *Chemical Science* vol. 10, 3567–3572 (2019).
15. Vassiliades, V. *et al. Comparing Multimodal Optimization and Illumination* in *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (Association for Computing Machinery, Berlin, Germany, 2017), 97–98.
16. Mouret, J.-B. & Clune, J. Illuminating search spaces by mapping elites. *arXiv e-prints,* arXiv:1504.04909 (Apr. 2015).
17. Mouret, J.-B. & Clune, J. *An algorithm to create phenotype-fitness maps* in *Proceedings of the Artificial Life Conference* (2012), 593–594.
18. Lehman, J. & Stanley, K. O. *Evolving a Diversity of Virtual Creatures through Novelty Search and Local Competition* in *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation* (Association for Computing Machinery, Dublin, Ireland, 2011), 211–218.
19. Maier, H. *et al.* Introductory overview: Optimization using evolutionary algorithms and other metaheuristics. *Environmental Modelling & Software* vol. 114, 195–213 (2019).
20. Holland, J. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence* (MIT Press, 1992).
21. Goldberg, D. E. & Holland, J. H. Genetic Algorithms and Machine Learning. *Machine Learning* vol. 3, 95–99 (1988).

22. Nigam, A. *et al.* Augmenting Genetic Algorithms with Deep Neural Networks for Exploring the Chemical Space. *International Conference on Learning Representations* (2020).

23. Zhou, Z. & Harris, K. D. M. Counteracting stagnation in genetic algorithm calculations by implementation of a micro genetic algorithm strategy. *Physical Chemistry Chemical Physics* vol. 10, 7262–7269 (48 2008).

24. Sterling, T. & Irwin, J. J. ZINC 15 – Ligand Discovery for Everyone. *Journal of Chemical Information and Modeling* vol. 55, 2324–2337 (Nov. 2015).

25. Mendez, D. *et al.* ChEMBL: towards direct deposition of bioassay data. *Nucleic Acids Research* vol. 47, D930–D940 (Nov. 2018).

26. Bleicher, S. S. *et al.* Divergent behavior amid convergent evolution: A case of four desert rodents learning to respond to known and novel vipers. *PLOS ONE* vol. 13, 1–17 (Aug. 2018).

27. Rundle, H. *et al.* Divergent Selection and the Evolution of Signal Traits and Mating Preferences. *PLOS Biology* vol. 3, e368 (Oct. 2005).

28. Gulick, J. Divergent Evolution. *Nature* vol. 39, 54–55 (Nov. 1888).

29. Du, Q., Faber, V. & Gunzburger, M. Centroidal Voronoi tessellations: Applications and algorithms. *SIAM Review* vol. 41, 637–676 (1999).

30. Du, Q. & Gunzburger, M. Grid generation and optimization based on centroidal Voronoi tessellations. *Applied mathematics and computation* vol. 133, 591–607 (2002).

31. Liu, Y. *et al.* On centroidal voronoi tessellation—energy smoothness and fast computation. *ACM Transactions on Graphics (ToG)* vol. 28, 1–17 (2009).

32. Vassiliades, V. *et al.* Using Centroidal Voronoi Tessellations to Scale Up the Multi-dimensional Archive of Phenotypic Elites Algorithm. *IEEE Transactions on Evolutionary Computation* (Aug. 2017).

33. Du, Q., Emelianenko, M. & Ju, L. Convergence of the Lloyd algorithm for computing centroidal Voronoi tessellations. *SIAM Journal on Numerical Analysis* vol. 44, 102–119 (2006).

34. Ramasubramanian, V. & Paliwal, K. K. Fast k-dimensional tree algorithms for nearest neighbor search with application to vector quantization encoding. *IEEE Transactions on Signal Processing* vol. 40, 518–531 (1992).

35. Meanwell, N. A. Improving Drug Design: An Update on Recent Applications of Efficiency Metrics, Strategies for Replacing Problematic Elements, and Compounds in Nontraditional Drug Space. *Chemical Research in Toxicology* vol. 29, 564–616 (2016).

36. Pennington, L. D. *et al.* Positional Analogue Scanning: An Effective Strategy for Multiparameter Optimization in Drug Design. *Journal of Medicinal Chemistry* (2020).

37. Hughes, J. *Lazy memo-functions* in *Conference on Functional Programming Languages and Computer Architecture* (1985), 129–146.

38. Lipinski, C. A. *et al.* Experimental and computational approaches to estimate solubility and permeability in drug discovery and development settings. *Advanced Drug Delivery Reviews* vol. 23. In Vitro Models for Selection of Development Candidates, 3–25 (1997).

39. Egan, W. J. *et al.* Prediction of Drug Absorption Using Multivariate Statistics. *Journal of Medicinal Chemistry* vol. 43, 3867–3877 (2000).

40. Veber, D. F. *et al.* Molecular Properties That Influence the Oral Bioavailability of Drug Candidates. *Journal of Medicinal Chemistry* vol. 45, 2615–2623 (2002).

41. Segler, M. H. *et al.* Generating focused molecule libraries for drug discovery with recurrent neural networks. *ACS Central Science* vol. 4, 120–131 (2018).

42. Anderson, E. *et al. SMILES, a line notation and computerized interpreter for chemical structures* (US Environmental Protection Agency, Environmental Research Laboratory, 1987).

43. Hochreiter, S. Long Short-Term Memory. *Neural Computation* vol. 9, 1735–1780 (1997).

44. Hessel, M. *et al. Rainbow: Combining improvements in deep reinforcement learning* in *Thirty-Second AAAI Conference on Artificial Intelligence* (2018).

45. Arús-Pous, J. *et al.* Exploring the GDB-13 chemical space using deep generative models. *Journal of Cheminformatics* vol. 11, 1–14 (2019).

46. Renz, P. *et al.* On failure modes of molecule generators and optimizers. *ChemRxiv e-prints* (Apr. 2020).

47. Rogers, D. & Hahn, M. Extended-Connectivity Fingerprints. *Journal of Chemical Information and Modeling* vol. 50, 742–754 (May 2010).

48. Morgan, H. L. The Generation of a Unique Machine Description for Chemical Structures-A Technique Developed at Chemical Abstracts Service. *Journal of Chemical Documentation* vol. 5, 107–113 (May 1965).

49. Hofmann, H., Wickham, H. & Kafadar, K. Letter-Value Plots: Boxplots for Large Data. *Journal of Computational and Graphical Statistics* vol. 26, 469–477 (2017).

50. Henault, E. S., Rasmussen, M. H. & Jensen, J. H. Chemical Space Exploration: How Genetic Algorithms Find the Needle in the Haystack. *PeerJ Physical Chemistry* vol. 2, e11 (Apr. 2020).

51. Hu, K. & Zhang, Y. J. Centroidal Voronoi tessellation based polycube construction for adaptive all-hexahedral mesh generation. *Computer Methods in Applied Mechanics and Engineering* vol. 305, 405–421 (2016).

52. Gaier, A., Asteroth, A. & Mouret, J.-B. *Data-Efficient Exploration, Optimization, and Modeling of Diverse Designs through Surrogate-Assisted Illumination* in *Proceedings of the Genetic and Evolutionary Computation Conference* (Association for Computing Machinery, Berlin, Germany, 2017), 99–106.

53. Kim, S. H. & Boukouvala, F. Machine learning-based surrogate modeling for data-driven optimization: a comparison oKim2020f subset selection for regression techniques. *Optimization Letters* vol. 14, 989–1010 (June 2020).

54. Vassiliades, V. & Mouret, J.-B. *Discovering the Elite Hypervolume by Leveraging Interspecies Correlation* in *Proceedings of the Genetic and Evolutionary Computation Conference* (Association for Computing Machinery, Kyoto, Japan, 2018), 149–156.

55. Eisenstein, M. Active machine learning helps drug hunters tackle biology. *Nature Biotechnology* vol. 38, 512–514 (May 2020).

56. Schwaller, P. *et al.* Predicting retrosynthetic pathways using transformer-based models and a hyper-graph exploration strategy. *Chemical Science* vol. 11, 3316–3325 (12 2020).

57. Somnath, V. R. *et al.* Learning Graph Models for Template-Free Retrosynthesis. *arXiv e-prints,* arXiv:2006.07038 (2020).

58. Chodera, J., Lee, A. A., London, N. & von Delft, F. Crowdsourcing drug discovery for pandemics. *Nature Chemistry* vol. 12, 581–581 (2020).

59. MacLeod, B. P. *et al.* Self-driving laboratory for accelerated discovery of thin-film materials. *Science Advances* vol. 6, eaaz8867 (2020).

Authors' addresses

Jonas Verhellen University of Oslo, Postboks 1337 Blindern, 0316 Oslo, Norway, jonasver@uio.no

Jeriek Van den Abeele University of Oslo, Postboks 1337 Blindern, 0316 Oslo, Norway, jonasver@uio.no

Paper II

Graph-Based Molecular Pareto Optimisation

Jonas Verhellen

Published in *Chemical Science*, June 2022, Volume 13, Issue 25, pp. 7526–7535.
DOI: 10.1039/D2SC00821A.

Abstract

Computer-assisted design of small molecules has experienced a resurgence in academic and industrial interest due to the widespread use of data-driven techniques such as deep generative models. While the ability to generate molecules that fulfil required chemical properties is encouraging, the use of deep learning models requires significant, if not prohibitive, amounts of data and computational power. At the same time, open-sourcing of more traditional techniques such as graph-based genetic algorithms for molecular optimisation [Jensen, *Chem. Sci.*, 2019, **12**, 3567-3572] has shown that simple and training-free algorithms can be efficient and robust alternatives. Further research alleviated the common genetic algorithm issue of evolutionary stagnation by enforcing molecular diversity during optimisation [Van den Abeele, *Chem. Sci.*, 2020, **42**, 11485-11491]. The crucial lesson distilled from the simultaneous development of deep generative models and advanced genetic algorithms has been the importance of chemical space exploration [Aspuru-Guzik, *Chem. Sci.*, 2021, **12**, 7079-7090]. For single-objective optimisation problems, chemical space exploration had to be discovered as a usable resource but in multi-objective optimisation problems, an exploration of trade-offs between conflicting objectives is inherently present. In this paper we provide state-of-the-art and open-source implementations of two generations of graph-based non-dominated sorting genetic algorithms (NSGA-II, NSGA-III) for molecular multi-objective optimisation. We provide the results of a series of benchmarks for the inverse design of small molecule drugs for both the NSGA-II and NSGA-III algorithms. In addition, we introduce the dominated hypervolume and extended fingerprint based internal similarity as novel metrics for these benchmarks. By design, NSGA-II, and NSGA-III outperform a single optimisation method baseline in terms of dominated hypervolume, but remarkably our results show they do so without relying on a greater internal chemical diversity.

II.1 Introduction

Machine learning has recently assumed a prominent role [1] in chemistry: predicting ADMET properties [2], supporting molecular dynamics simulations [3], and assisting in

the design of small molecules without reverting to explicit rules or expert knowledge [4–12]. However, training-free optimisation algorithms that comprehensively traverse and explore chemical space have been shown to be more efficient [13, 14] than their machine learning counterparts in discovering high-performing *de novo* molecules. Sometimes this search in chemical space reduces to an optimisation for a single property like melting point [15] or protein binding affinity [16], but often there are additional requirements that make it necessary to optimise for additional properties such as low toxicity [17], high synthesizability [18] or off-target activity. In the case that multi-objective optimisation is necessary, a trade-off between different (and possibly competing) optimisation objectives has to be defined.

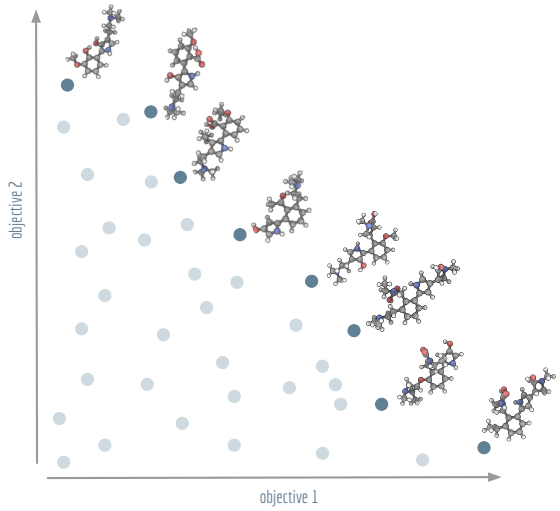

Figure II.1: Visualisation of a Pareto front (dark blue) and dominated solutions (light blue). Example molecules shown at the Pareto front were generated by NSGA-II for Tanimoto similarities with regard to lysergic acid diethylamide (objective 1) and psilocybin (objective 2).

In current molecular generative model benchmarks [13], typically either the median or the geometric mean of the objective is chosen as a stand-in aggregate fitness function. To give relative importance to the different objectives, domain experts can assign weights to them or combine appropriate modifying functions to obtain a single, fine-tuned objective function. However, many fields of science and engineering make use of an alternative approach to multi-objective optimisation by searching for a set of so-called Pareto optimal solutions [19]. All solutions in a Pareto optimal set are characterised by the fact that there are no other individual solutions that have a higher (or equal) fitness in all objective functions. Together, the set of Pareto optimal solutions form an optimal envelope in objective space known as the Pareto front, see Figure II.1.

The Pareto front provides a family of solutions, all equivalent in principle, aiding domain experts to make choices when trade-offs between objectives are not known

beforehand. Over the past two decades, a family of algorithms known as the non-dominated sorting genetic algorithms [20] (NSGA) has been developed for finding Pareto fronts. In a complex process, such as drug design, having access to a technique complementary to single objective optimisation, can yield deeper insights and improve efficiency. Therefore, in this paper, we provide the community with state-of-the-art and open-source implementations of the NSGA-II and NSGA-III algorithms [21–23] based on a popular graph-based genetic algorithm [24] (GB-GA) for molecular optimisation.

A newer generation of NSGA algorithm, NSGA-III, which uses a more complex means of ensuring coverage of the entire Pareto front, was originally reported to be an improvement over NSGA-II. However, later analyses [25, 26] have shown that for a wide range of computational experiments NSGA-III does not consistently outperform NSGA-II in every use-case. Therefor we compare the performance of NSGA-III and NSGA-II on a set of small molecule multi-objective optimisation benchmarks, making use of the dominated hypervolume as a novel measure of the effectiveness in these type of problems. As a baseline, we make use of a state-of-the-art single-objective optimisation algorithm that employs the geometric mean as a surrogate aggregate fitness function. Whereas proprietary applications of NSGA-II to molecular design have been reported [27, 28], there is a lack of open-source implementations of both NSGA-II and NSGA-III for the inverse design of small molecules. We anticipate that our results and the availability of the code will encourage the development of more powerful Pareto optimisation algorithms for chemistry as well as their widespread adoption in computer-assisted chemical design.

II.2 Algorithmic Methodology

NSGA-II and NSGA-III are genetic algorithms tailored to finding Pareto fronts. In this section, we introduce the fundamentals of genetic algorithms in the context of small molecule design and discuss the importance of balancing quality with diversity. We then describe the general framework of non-dominated sorting genetic algorithms and elaborate upon the NSGA-II and NSGA-III algorithms and their differences. In the remainder of the section, we discuss technical aspects such as structural alert based chemical filters, memoisation, the construction of reference directions (only used in NSGA-III), positional analogue scanning, and parallelism.

II.2.1 Genetic Algorithms

A genetic algorithm is, as the name suggests, a heuristic search method [29] inspired by the process of natural evolution. Genetic algorithms [30, 31] can achieve highly effective single-objective optimisation by consistently and incrementally improving a selection of trial solutions. The current set of the solutions used by the algorithm is known as the (evolutionary) population. In each iteration of the algorithm – known as a generation – novel solutions are generated by stochastically changing or combining the current solutions. In the genetic algorithm community, these two operations for generating new solutions are known as *mutations* and *crossovers*, respectively. At the end of each generation, the population is reduced to its original size by selecting only

the highest performing molecules for survival. Eventually, the selection pressure in this procedure forces the population of solutions towards an optimum.

For small molecule optimisation, these ideas can be implemented by representing solutions (i.e. molecules) by either their molecular graphs, or by text representation such as the simplified molecular-input line-entry system [32] (SMILES) or self-referencing embedded strings [33] (SELFIES). The graph representation has been used in the graph-based genetic algorithm (GB-GA) which was shown to outperform machine learning approaches [24]. In Figure II.2, we show examples of mutations and crossovers on molecular graphs. To rule out graphs that represent impossible chemical configurations, only those that can be correctly translated to and from SMILES are retained. The initial population of candidate molecules is typically obtained from public databases like ZINC [34] or ChEMBL [35].

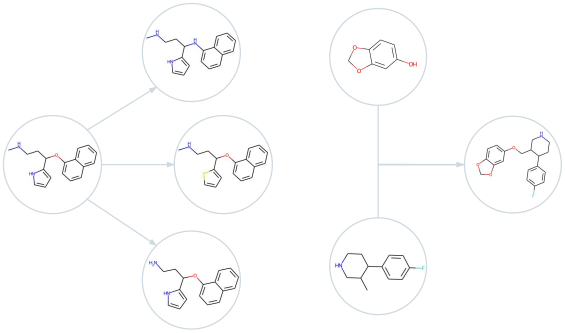

Figure II.2: Examples of mutations (left) and a crossover (right) as generated by GB-EPI. Note that minor changes to chemical structure can be used to efficiently achieve optimisation even for challenging objectives.

II.2.2 Quality-Diversity Algorithms

Unfortunately, genetic algorithms are known to be vulnerable to evolutionary stagnation when encountering low-performing valleys or local optima [36]. Enforcing diversity [37] in the population of molecules a genetic algorithm uses can alleviate these issues. Quality-diversity algorithms [38], such as the graph-based elite patch illumination algorithm [39] (GB-EPI), obtain this diversity by splitting the population into niches based on their physicochemical properties. In each generation, the best performing molecule in each of the individual niches is retained, rather than selecting the highest-scoring solutions regardless of their diversity.

Alternatively, the superfast traversal, optimisation, novelty, exploration and discovery algorithm [40] (STONED) leverages molecular diversity through the use of SELFIES. In contrast to the more traditionally used SMILES, SELFIES can be mutated arbitrarily at any position in the string to produce new strings that represent valid molecular structures. The STONED algorithm uses this property of SELFIES to preserve diversity in its population. By varying the position of modification within

the string, the algorithm balances exploration and exploitation to avoid stagnation in low-performing valleys or local optima.

II.2.3 Non-dominated Sorting Genetic Algorithms

In contrast to single-objective optimisation problems, in which diversity had to be discovered as a usable resource, diversity is inherently present in multi-objective optimisation problems. The presence of diversity is most obvious when considering a Pareto front, in which solutions to multi-objective optimisation problems must involve trade-offs to satisfy the conflicting demands of different objective functions. Several algorithms with different properties and varying levels of complexity have been proposed for finding Pareto optimal fronts. The main class of algorithms used for this task are the non-dominated sorting genetic algorithms, NSGA-II and NSGA-III.

Algorithm: Non-dominated Sorting Genetic Algorithms

Input: N – number of generations, $\mathscr{M}_0$ – initial population
for $i = 0 \rightarrow N$ **do**
  $\mathscr{F}_i \leftarrow \mathsf{fitnesses}(\mathscr{M}_i)$;
  $\mathscr{M}' \leftarrow \mathsf{mutation}(\mathscr{M}_i) + \mathsf{crossover}(\mathscr{M}_i)$;
  $\mathscr{F}' \leftarrow \mathsf{fitnesses}(\mathscr{M}')$;
  fronts $\leftarrow \mathsf{sorting}(\mathscr{M}' + \mathscr{M}_i, \mathscr{F}' + \mathscr{F}_i)$;
  for *front* **in** *fronts* **do**
    if *splitting_front(front)* **then**
      $\mathscr{M}_{i+1} \leftarrow \mathsf{splitting\_procedure}(\mathit{front})$;
    else
      $\mathscr{M}_{i+1} \leftarrow \mathit{front}$;
  end
end
Result: $\mathscr{M}_N$ – molecules, $\mathscr{F}_N$ – fitnesses

Figure II.3: Pseudocode description of a generic non-dominated sorting genetic algorithm adapted to the setting of molecular optimisation.

Non-dominated sorting genetic algorithms [20] are, in essence, genetic algorithms that evaluate and select on the Pareto dominating status of each solution in the evolutionary population as shown in Figure II.3. Instead of selecting molecules based on a fitness function, these algorithms sort all solutions into a series of fronts, see Figure II.4 (a), each front dominated by the previous fronts. The first front (dark blue) is the set of completely non-dominated individuals in the current population, the second front (light blue) is the set of individuals dominated only by the individuals in the first front, and so on for all other fronts formed by the remaining individuals in the population (white). The algorithm accepts the fronts, with all of its individuals, into the evolutionary population in ascending order, until the maximum size of the evolutionary population has been reached.

The final front accepted by a non-dominated sorting genetic algorithm might, and often will, contain more individuals than can be added to the surviving evolutionary population without exceeding its size limit. This set of individuals is known in the multi-objective optimisation community as the *splitting front* [20]. Because there is no difference between the individuals in the splitting front in terms of Pareto dominance, further criteria are used to select which individuals are retained and which are discarded. In the splitting front selection procedure for non-dominated sorting genetic algorithms, this criteria is typically a measure of diversity. The NSGA-II and NSGA-III algorithms both rely on a diversity criteria, but differ significantly in how they enforce this diversity, see Figure II.4 (b) and (c).

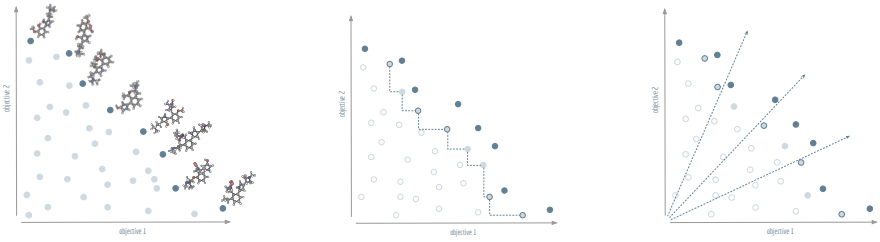

(a) Example Setup (b) NSGA-II (c) NSGA-III

Figure II.4: Visualisation of the splitting front procedure of non-dominated sorting genetic algorithms: (a) The Pareto dominant front is shown in dark blue, the splitting front is light blue, and the remaining solutions are white. For this example, the second front is chosen as the splitting front, and it is assumed that five more solutions need to be picked to complete the population. These solutions will be indicated with a dark blue circumference. (b) The selection procedure of NSGA-II calculates a distance in objective space to the nearest neighbours in the front. The outermost solutions are picked by default, the remaining solutions are chosen according to the furthest distance from neighbours. (c) The selection procedure of NSGA-III calculates the orthogonal distance to predefined reference directions in objective space and selects the closest solution for each axis. Note that the two objective axes are also used as reference directions so that the outermost solutions are picked by default.

II.2.4 NSGA-II

NSGA-II [21] makes use of a *crowding distance* to differentiate within the splitting front. The crowding distance is calculated for each individual, and indicates how closely the individual is surrounded by the other members of the splitting front. For NSGA-II, the crowding distance used is the Manhattan distance [41] in objective space. A larger crowding distance indicates a less crowded individual. Within a splitting front, NSGA-II orders all individuals by their crowding distances, and subsequently accepts the molecules with the largest crowding distance into the evolutionary population until the maximum size is reached. The outer solutions in the splitting front are assigned an infinite crowding distance to ensure that they are retained in each generation.

II.2.5 NSGA-III

In contrast to NSGA-II, the NSGA-III algorithm [22, 23], uses reference directions [42, 43] instead of a crowding distance to enforce diversity in the selection of solutions within the splitting front. Reference directions are determined by a predefined set of points on the unit simplex in fitness space. Each reference direction is defined as a ray originating from the origin and passing through exactly one of these points. NSGA-III assigns a reference direction to each solution in the population based on the nearest perpendicular distance (in normalised fitness space) to the corresponding direction. In the splitting front selection procedure, the NSGA-III algorithm prioritises reference directions that are underrepresented in the current surviving evolutionary population.

If a reference direction does not have any solution assigned to it after reaching the splitting front, then the molecule in the splitting front with the smallest perpendicular distance to this direction is selected for survival. If all underrepresented reference directions have been assigned one surviving solution, and the maximum size of the surviving population has not been reached, the remaining solutions are selected by a stochastic procedure. Note that NSGA-III selects the solutions in the fronts before the splitting front in its entirety, like in NSGA-II. However, contrary to NSGA-II's crowding distance which is calculated within the splitting front, the reference directions used in NSGA-III take into account the diversity of the entire surviving population.

II.2.6 Reference Directions

The reference directions determine the diversity in the selection of solutions from the splitting front, so these directions are typically chosen to be well distributed over the unit simplex. Traditionally the reference direction generation method of Das and Dennis has been used for NSGA-III. Unfortunately, due to the highly structured (combinatorial) nature of the Das-Dennis reference direction generating procedure [42], the method cannot produce an arbitrary number of directions. In addition, it has been shown that most of the reference directions generated by the Das-Dennis method cross through the boundaries of the unit simplex rather than the interior [44], inducing a bias in the selection of solutions from the splitting front.

To alleviate the issues of the Das-Dennis method, an energy-based approach has recently been proposed [43] in the multi-objective optimisation literature. Inspired by methods in physics, a generalisation of the potential energy called the *Riesz s-energy* [45] is calculated for a given number of reference points on the unit simplex. The Riesz s-energy U_s is defined between two points p_1, p_2 in s-dimensional Euclidean space as,

$$U_s(p_1, p_2) = \frac{1}{\|p_1 - p_2\|^s}\,. \tag{II.1}$$

The location of the points along of the unit simplex are then optimised to minimise the combined Riesz s-energy of all the reference points. This allows for the construction of an arbitrary number of well-spaced reference directions. The results in this paper were obtained using the Riesz s-energy method to generate the reference directions for NSGA-III, with s equal to the square root of the number of objective functions as suggested in the original paper [43].

II.2.7 Shared Technical Properties

We follow the example of GB-EPI [39] and include a series of minor but important technical features to our NSGA-II and NSGA-III implementations, focused on improved chemical optimisation or higher relevance and better quality of the generated molecules. For instance, our NSGA-II and NSGA-III implementations make use of decoupled crossovers and mutations. As shown in GB-EPI, early on in an evolutionary algorithm, crossovers support the efficient exploration of chemical space, while later on local mutations are beneficial in improving the nearly-converged solutions. Therefore it is beneficial to apply both operators separately rather than in sequence.

Similarly, we follow the example of GB-EPI to apply the computational equivalent of *in vitro* positional analogue scanning [46] by repurposing the mutation operator to systematically return not just a single mutation of a molecule, but all of its positional analogues. To offset the the computational overhead introduced by positional analogue scanning and to improve efficiency in general, we store a record of obtained fitness calculations. This approach is known as memoisation [47] and ensures that an algorithm does not unnecessarily repeat calculations. To further reduce clock time, we also implemented concurrency for the objective function evaluations and remove undesirable compounds based on structural ADMET filters [48–50] before they enter the evaluation step of the algorithm.

II.3 Benchmarks

To the test the potency of our open-source implementations of NSGA-II and NSGA-III for multi-objective optimisation in drug design, we extend the use of tasks devised in the GuacaMol benchmarking suite [13] by the bioinformatics company BenevolentAI. From the suite we selected multi-parameter optimisation (MPO) tasks with three or more objectives that aim to fine-tune the structural or physicochemical properties of five FDA-approved drugs: Cobimetinib (a mitogen-activated kinase inhibitor), Fexofenadine (a second-generation antihistamine), Osimertinib (a Tyrosine kinase inhibitor), Perindopril (a long acting ACE inhibitor), and Ranolazine (an anti-anginal drug). We search for a set of molecules that span the entirety of the Pareto front instead of trying to optimise a single value like the geometric mean.

The objectives in these benchmarks, as shown in Table II.2, are either similarity metrics that measure the distance to the corresponding drug molecule, or specific properties such as the amount of rotatable bonds in a molecule, the topological polar surface area [52] (TPSA) or the lipophilicity partition coefficient [53] (log(P)). The similarity metrics are calculated using the Tanimoto similarity [54, 55], of the fingerprints of the target and the generated candidate molecule. The fingerprints used here are either extended-connectivity fingerprints [56, 57] (ECFP/FCFP) which encode molecular structures in terms of concentric atomic neighbourhoods, or atom-pair fingerprints [58] (AP) which encode molecules based on their atom pairs and their bond distance. The main advantage of fingerprint-based similarities compared to more involved similarity measures is that they can be rapidly calculated and inherently represent the presence or absence of molecular substructures or atom pairs.

Table II.1: Overview of the multi-objective optimisation benchmarks used in this paper, the first five benchmarks are adapted from the Guacamol suite while the latter two benchmarks were constructed to emulate the demands of poly-pharmacology projects. The upper row of each task represents the values calculated for each objective. The lower rows show the modifiers applied to each of these values. The fingerprints used to calculate the similarities are denoted as arguments of the Tanimoto function, the parameters used for the modifiers are displayed as arguments of the corresponding functions. For the poly-pharmacology benchmarks, the genes targeted for activity are indicated. The CNS function calculates the central nervous system desirability score (high blood-brain-barrier permeability and low toxicity potential) as proposed by Pfizer [51].

Task \ Objective	I	II	III	IV	V
Cobimetinib					
	Tanimoto(FCFP4)	Tanimoto(ECFP6)	Rotatable Bonds	Aromatic Rings	CNS(0.5)
	Clipped(0.7)	MinGaussian(0.75, 0.1)	MinGaussian(3, 1)	MaxGaussian(3, 1)	-
Fexofenadine					
	Tanimoto(AP)	TPSA	log(P)	-	-
	Clipped(0.8)	MaxGaussian(90, 10)	MinGaussian(4, 1)	-	-
Osimertinib					
	Tanimoto(FCFP4)	Tanimoto(ECFP6)	TPSA	log(P)	-
	Clipped(0.8)	MinGaussian(0.85, 0.1)	MaxGaussian(95, 20)	MinGaussian(1, 1)	-
Pioglitazone					
	Tanimoto(ECFP4)	Molecular Weight	Rotatable Bonds	-	-
	Gaussian(0, 0.1)	Gaussian(356, 10)	Gaussian(2, 0.5)	-	-
Ranolazine					
	Tanimoto(AP)	log(P)	TPSA	Fluorine Count	-
	Clipped(0.7)	MaxGaussian(7, 1)	MaxGaussian(95, 20)	Gaussian(1, 1)	-
DAP Kinases					
	hERG	SCN2A	DAPk1	DRP1	ZIPk
	Gaussian(0, 0.1)	Gaussian(0, 0.1)	Clipped(0.8)	Clipped(0.8)	Clipped(0.8)
Antipsychotics					
	hERG	5-HT2A	5-HT2B	DRD2	CNS(0.5)
	Gaussian(0, 1.0)	Clipped(0.8)	Clipped(0.8)	Clipped(0.8)	-

The raw scores obtained from similarity or property measurements are post-processed by modifier functions that map the scores to the [0, 1] interval and allow the objective to be fine-tuned. The modifier functions used in this paper are *Clipped(value)*, *Gaussian(mean, variance)*, *MinGaussian(mean, variance)*, and *MaxGaussian(mean, variance)*. The *Clipped* modifier is a thresholded modifier in which values above a given threshold are mapped to one, while values below threshold decrease linearly to zero. The *Gaussian* modifiers target a specific value, returning high scores when the underlying value is near the target. The *Min* and *Max* versions of this modifier map the input value to one if it is lower or higher than the target value, respectively. For example, in the Fexofenadine benchmark a molecule with a Tanimoto similarity higher than 0.8, a TPSA above 90.0 and a log(P) below 4.0 would score perfectly on each objective. More information on the modifiers can be found in the supplementary information accompanying the Guacamol paper [13].

Precise evaluation of generative models in terms of their value to pharmaceutical drug design programs can be challenging. To increase relevance, with respect to real-life drug design projects, while maintaining the efficient benchmark evaluations necessary for iterative design and statistical analysis, we integrate an existing data-driven surrogate

model for target activity into the Guacamol benchmarking suite [13]. We make use of a previously proposed surrogate model [59], minding the separation of concerns [60], that has been used to study failure modes in molecule generation. This model ranks molecules based on the ratio of trees in a random forest classifier, trained on ChEMBL activity data [35], predicting that the molecule is active. In the model, binary ECFP fingerprints [57] of size 1024 and radius 2 are used as features.

In this paper, we provide two novel benchmarks for Pareto optimisation making use of this model. Inspired by the demands of a multi-target drug discovery project [61], we have constructed a multi-kinase inhibitor task and a multi-neuroreceptor binding antipsychotics task. In the kinase inhibitor task, we aim for molecules that inhibit three DAP kinases [62] (DAPk1, DRP1, and ZIPk) often implicated in cancer while trying to avoid activity against common off-target ion channels [63, 64] (hERG, and SCN2A). In the ongoing search for novel anti-psychotic medication, focus has shifted [65] to combined binders of serotonergic receptors (5-HT2A, and 5-HT2B) and a more classical target: the dopaminergic DRD2 receptor. In the multi-receptor antipsychotica task, we target these three receptors, and aim to avoid an off-target ion channel (hERG) while fulfilling the Pfizer central nervous system desirability requirements.

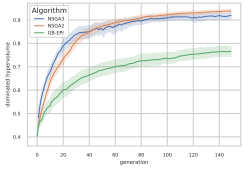

(a) Cobimetinib - Dominated Hypervolume

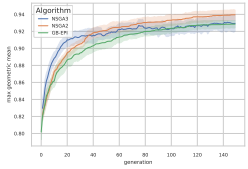

(b) Cobimetinib - Geometric Mean

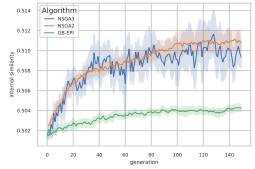

(c) Cobimetinib - Internal Similarity

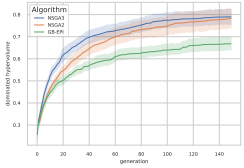

(d) Fexofenadine - Dominated Hypervolume

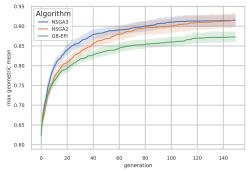

(e) Fexofenadine - Geometric Mean

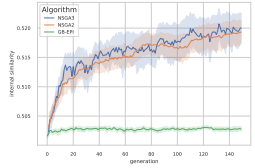

(f) Fexofenadine - Internal Similarity

Figure II.5: Timeseries plots with variance bands of the dominated hypervolume, the maximum geometric mean, and internal similarity for the Cobimetinib (a,b,c) and Fexofenadine (d,e,f) tasks as a function of generations of the evolutionary populations. The mean value (solid line) and the 95% confidence interval (variance bands) over twenty runs of NSGA-II (orange), NSGA-III (blue), and GB-EPI (green, optimising the geometric mean) are shown. Details of the experimental setup for these results, including hyperparameters, initial population and chemical filters are discussed in subsection II.4.

II.3.1 Dominated Hypervolume

In multi-objective problems, tracking the evolution of an algorithm or measuring the quality of a Pareto front with respect to a single parameter can be challenging. In previous benchmarking efforts for optimisation algorithms of small molecules, the geometric mean of the objectives has traditionally been used as both an aggregate objective and as a metric. From a technical point of view, the geometric mean is the exponential of the arithmetic mean of the log-transformed set of objective scores. As a consequence, the geometric mean for strictly positive values is sensitive to severe underperformance in any single objective, making it a relevant measure for many multi-objective optimisation problems. However, other indicators of the quality of Pareto fronts have been developed by the multi-objective optimisation community. One such metric is the dominated hypervolume [66], which we introduce to the domain of chemical optimisation as an alternative measure for multi-objective optimisation benchmarks.

The dominated hypervolume (also known as Lebesgue measure [67] or S-metric [68]) maps a set of points in objective space to the size of of the region Pareto dominated by that set. The hypervolume has to be bounded from below by a reference point, which for the purposes of this paper will systematically be chosen to be the origin of objective space. The dominated hypervolume simultaneously takes into account the proximity of the points to the ideal Pareto front and their spread over the objective space. For problems with less than five objectives, the dominated hypervolume can be calculated exactly. However, for higher-dimensional multi-objective optimisation problems, calculating the dominated hypervolume precisely can be computationally expensive and hence a smorgasbord of efficient approximation methods [69, 70] for the dominated hypervolume has been developed.

II.3.2 Internal Similarity

In comparing the performance of the different algorithms discussed in this paper, it is useful to differentiate whether algorithms encourage a significantly different amount of chemical diversity in their evolutionary populations. In cheminformatics, similarity between two molecules is usually quantified based on metrics applied to binary fingerprints that featurise chemical substructures. To calculate the diversity of molecules, the pairwise similarity of each combination of molecules in a set has been traditionally calculated using a binary similarity index, like the Tanimoto similarity [54, 55], and summarised in an aggregate metric. However, the recent development of extended similarity metrics [71, 72] enables the simultaneous and straightforward comparison of an arbitrary number of bitvectors such as molecular fingerprints.

In this paper we make use of extended similarity indices to calculate and track the internal similarity of evolutionary populations. Extended similarity metrics, which compare a stack of bitvectors, have the advantage [71] that they do not require the full similarity matrix of the compound pool or aggregate metric. In addition to being more efficient, extended similarity metrics reduce to the traditional binary similarity metrics if applied to a set of two molecules. According to computational experiments, two newly proposed extended similarity metrics [72] are highly advantageous compared to

Table II.2: The dominated hypervolume, maximum geometric mean, internal similarity, and cumulative fitness calls after 150 generations, for seven multi-objective optimisation tasks averaged over 20 runs of the GB-EPI, NSGA-II, and NSGA-III algorithms. Details of the experimental setup for these results, including hyperparameters, construction of the initial population, and chemical filters are discussed in subsection II.4. Mean average values for each of the measures are given with standard deviations.

Algorithm	Task	Dominated Hypervolume	Geometric Mean	Internal Similarity	Fitness Calls (Cumulative)
GB-EPI					
	Cobimetinib	0.77 ± 0.05	0.93 ± 0.01	0.50 ± 0.00	13577 ± 1224
	Fexofenadine	0.67 ± 0.07	0.87 ± 0.03	0.50 ± 0.00	17985 ± 1398
	Osimertinib	0.54 ± 0.04	0.85 ± 0.01	0.50 ± 0.00	12982 ± 1351
	Pioglitazone	0.98 ± 0.04	0.99 ± 0.01	0.50 ± 0.00	13160 ± 3104
	Ranolazine	0.46 ± 0.04	0.81 ± 0.02	0.50 ± 0.00	16859 ± 1537
	DAP Kinases	0.03 ± 0.05	0.46 ± 0.06	0.51 ± 0.00	23545 ± 3150
	Antipsychotics	0.09 ± 0.02	0.57 ± 0.06	0.51 ± 0.00	21905 ± 3073
NSGA-II					
	Cobimetinib	0.94 ± 0.02	0.94 ± 0.01	0.51 ± 0.00	17784 ± 1753
	Fexofenadine	0.78 ± 0.10	0.92 ± 0.04	0.52 ± 0.00	20268 ± 2909
	Osimertinib	0.66 ± 0.03	0.89 ± 0.01	0.52 ± 0.00	16848 ± 2655
	Pioglitazone	1.00 ± 0.00	1.00 ± 0.00	0.51 ± 0.00	19944 ± 4765
	Ranolazine	0.68 ± 0.06	0.87 ± 0.02	0.51 ± 0.00	21259 ± 2181
	DAP Kinases	0.05 ± 0.03	0.50 ± 0.07	0.52 ± 0.00	24350 ± 3826
	Antipsychotics	0.08 ± 0.03	0.50 ± 0.05	0.51 ± 0.00	21246 ± 1909
NSGA-III					
	Cobimetinib	0.92 ± 0.03	0.93 ± 0.02	0.51 ± 0.00	14224 ± 1807
	Fexofenadine	0.79 ± 0.00	0.91 ± 0.03	0.52 ± 0.01	12950 ± 2326
	Osimertinib	0.66 ± 0.03	0.89 ± 0.01	0.52 ± 0.00	11052 ± 2337
	Pioglitazone	1.00 ± 0.00	1.00 ± 0.00	0.51 ± 0.01	10639 ± 2736
	Ranolazine	0.63 ± 0.06	0.85 ± 0.02	0.51 ± 0.00	17949 ± 2732
	DAP Kinases	0.04 ± 0.02	0.48 ± 0.07	0.51 ± 0.01	22454 ± 3440
	Antipsychotics	0.05 ± 0.03	0.49 ± 0.04	0.52 ± 0.01	32991 ± 3473

the extended Tanimoto similarity: the extended Baroni-Urbani-Buser similarity index and the extended Faith similarity index. Throughout this paper will make use of the extended Faith similarity index.

II.4 Results

To increase the real-life relevance of the benchmarks used here, we run each algorithm 20 times for 150 generations per benchmark. We also reject molecules that either trigger the structural alerts from GSK [73], or those that contain ring allenes, macrocycles, an abundance of hologenicity (#F > 6, #Br > 3, #Cl > 3), rotatable bonds (>10) or hydrogen acceptors/donors (>10). In addition, the initial populations used in this paper consist of a hundred molecules randomly sampled from the Guacamol [13] subset of ChEMBL [35]. All these molecules are neutral, do not contain salts and have Tanimoto similarities below 0.323 to any of ten FDA approved drugs (Celecoxib, Aripiprazole, Cobimetinib, Osimertinib, Troglitazone, Ranolazine, Thiothixene, Albuterol, Fexofenadine, Mestranol).

Based on previous work comparing single objective optimisation methods, we choose GB-EPI (with geometric mean as surrogate fitness function) as a representative baseline to compare against NSGA-II and NSGA-III. For GB-EPI, we choose four medicinally relevant features of interest to span the archive: molecular weight (ranged

from 140 to 555), log(P) (0.0 to 7.0), TPSA (0 to 140), and molar refractivity (40 to 130). For fair comparison, molecules exceeding these ranges are excluded from the evolutionary populations of NSGA-II and NSGA-III during the benchmarks. Based on previous experience with GB-EPI, the archive size for was set to 150 and the batch size to 20. The archive size in quality-diversity algorithms, such as GB-EPI, is the counterpart of the population size in traditional genetic algorithms. In general, the batch size refers to the amount of molecules submitted to mutation and crossover per generation. For NSGA-II, we used a population size of 100 (corresponding to the initial population) and a batch size of 20. For NSGA-III, we used the same batch size but experimentation guided us towards a smaller total evolutionary population: we settled on the use 25 reference directions, and a population size of 35 molecules. These hyperparameters were chosen to support global performance of each individual algorithm without disrupting splitting procedures, as a consequence the amount of fitness calls varies across algorithms and generations.

In Figure II.5 the evolution of the dominated hypervolume, maximum geometric mean and internal similarity of the NSGA-II, NSGA-III, and GB-EPI algorithms is shown for two representative benchmarks (Cobimentib and Fexofenadine). Throughout the computational experiments GB-EPI, which optimises directly for the geometric mean, is used as a baseline comparison method. As expected, NSGA-II and NSGA-III successfully out-compete the GB-EPI baseline in terms of dominated hypervolume for both benchmarks. In contrast to GB-EPI, the NSGA algorithms are designed specifically to optimise the Pareto front, the quality of which is measured by the dominated hypervolume. The geometric mean follows trends similar to the dominated hypervolume in the benchmarks. However, the values of the maximal geometric mean lie close to each other and the 95% confidence interval of GB-EPI overlaps with NSGA-II and NSGA-III during the latter stages of the Cobimentib task.

An overview of the results for the multi-objective benchmarks is shown in Table II.2 in terms of averages and standard deviations. NSGA-II and NSGA-III perform better than the baseline on each of the benchmarks for both dominated hypervolume and maximum geometric mean with the exception of the antipsychotics task. In that task, similarity between the three receptor targets disadvantages NSGA-III due to its rigid reference directions. For the Fexofenadine and Pioglitazone benchmarks, GB-EPI lies within one standard deviation of either NSGA-II or NSGA-III for both metrics. Note that to obtain the global maximum geometric mean of these benchmarks or the global optimum of one of the objectives, direct optimisation should be used. In principle, Pareto optimisation algorithms should reach these types of global optima, but significantly less efficiently as the evolutionary population is spread out over objective space. Conversely, when using a single aggregation function, the solutions tend to lie close to each other in objective space, and don't cover the entirety of the Pareto front.

To study the comparative efficiency of each algorithm, we track the cumulative number of function calls over the full 150 generations for the twenty individual runs of each algorithm. This has the advantage that it does not interrupt the splitting front procedure, as might be the case when working with a fixed and limited function call budget. An overview of the mean and standard deviation of the cumulative fitness calls of each algorithm is shown in Table II.2. NSGA-III consistently outperforms NSGA-II in terms of efficiency, and is more efficient than GB-EPI in all benchmarks where they

have similar performance for dominated hypervolume and geometric mean. In contrast to single objective optimisation problems, where a lower internal similarity has been regarded as beneficial, for multi-objective optimisation the algorithms which encourage greater internal similarity are better performing.

II.5 Conclusion and Outlook

This paper introduces two novel open-source and graph-based implementations of non-dominated sorting genetic algorithms, NSGA-II and NSGA-III, for small molecule multi-objective optimisation. The performance of these algorithms is compared to a single objective quality-diversity algorithm (GB-EPI) on four metrics: dominated hypervolume, maximal geometric mean, internal similarity and efficiency. Previous benchmarks for generative models of small molecules focused on the maximal geometric mean as a sole aggregate indicator of success in multi-objective optimisation. However, the Pareto front – the collection of optimal points in objective space – is not solely characterised by the geometric mean of a single molecule. In this paper we show that the size of the hypervolume dominated in objective space (with respect to the origin) is a useful, often more discriminative, alternative metric in generative model benchmarks.

The performance of NSGA-II and NSGA-III for graph-based optimisation of molecules is encouraging. Both algorithms specialise in finding the optimal Pareto front and our benchmarks show that this approach is superior compared to GB-EPI (which optimises the geometric mean directly). In line with analyses of purely numerical benchmarks found in the literature, NSGA-III does not always outperform NSGA-II in our chemical benchmarks, indicating that the two algorithm produce similar results according to this metric. Throughout all the benchmarks presented in this paper however, NSGA-III seems to be the most efficient in its use of function calls. Notably, and in contrast to single objective optimisation, the higher performing algorithms NSGA-II and NSGA-III have a higher and faster increasing internal similarity in their evolutionary populations than the baseline.

The above discussed efficiency, performance, and flexibility of the graph-based implementations of NSGA-II and NSGA-III for small molecule multi-objective optimisation as provided with this paper, allows the community to use these algorithms for practical use. In addition, these implementations can be used as future baselines and as starting points for future developments in this field. One such possible development would be to further reduce the amount of function calls through the use of contextual multi-armed bandits [74], or Gaussian processes [75] to prune the amount of molecules presented to the evaluation step of the algorithms. Finally, the algorithms presented here can be integrated into the workflow for multi-objective tasks given to self-driving laboratories [76] or other set-ups making use of active learning [77].

Data Availability

Full code for the implementations of NSGA-II and NAGA-III is available at: https://github.com/Jonas-Verhellen/MolecularGraphPareto.

Author contributions

The author confirms sole responsibility for the following: study conception and design, data collection, analysis and interpretation of results, and manuscript preparation.

Conflicts of interest

There are no conflicts to declare.

Acknowledgements

The author wishes to acknowledge useful feedback on this manuscript by K. Beshkov, P. Coppin and E. Herbert. This article was funded by UiO:Life Science through the 4MENT convergence environment.

References

1. Schneider, P. *et al.* Rethinking drug design in the artificial intelligence era. *Nature Reviews Drug Discovery* (Dec. 2019).
2. Cáceres, E. L., Tudor, M. & Cheng, A. C. Deep learning approaches in predicting ADMET properties. *Future Medicinal Chemistry* vol. 12, 1995–1999 (2020).
3. Doerr, S. *et al.* TorchMD: A Deep Learning Framework for Molecular Simulations. *J. Chem. Theory Comput.* vol. 17, 2355–2363 (Apr. 2021).
4. Moret, M. *et al.* Generative molecular design in low data regimes. *Nature Machine Intelligence* vol. 2, 171–180 (Mar. 2020).
5. Zhavoronkov, A. *et al.* Deep learning enables rapid identification of potent DDR1 kinase inhibitors. *Nature Biotechnology* vol. 37, 1546–1696 (2019).
6. Popova, M., Isayev, O. & Tropsha, A. Deep reinforcement learning for de novo drug design. *Science Advances* vol. 4, eaap7885 (2018).
7. Segler, M. H. S. *et al.* Generating Focused Molecule Libraries for Drug Discovery with Recurrent Neural Networks. eng. *ACS Central Science* vol. 4, 120–131 (Jan. 2018).
8. Gómez-Bombarelli, R. *et al.* Automatic Chemical Design Using a Data-Driven Continuous Representation of Molecules. *ACS Central Science* vol. 4. PMID: 29532027, 268–276 (2018).
9. Kusner, M. J. *et al.* Grammar Variational Autoencoder. *International Conference on Machine Learning* (2017).
10. Smalley, E. AI-powered drug discovery captures pharma interest. *Nature Biotechnology* vol. 35, 604–605 (July 2017).
11. Manglik, A. *et al.* Structure-based discovery of opioid analgesics with reduced side effects. *Nature* vol. 537, 185–190 (Sept. 2016).

12. Reymond, J.-L. The Chemical Space Project. *Accounts of Chemical Research* vol. 48. PMID: 25687211, 722–730 (2015).

13. Brown, N. *et al.* GuacaMol: Benchmarking Models for de Novo Molecular Design. *Journal of Chemical Information and Modeling* vol. 59, 1096–1108 (2019).

14. Henault, E. S., Rasmussen, M. H. & Jensen, J. H. Chemical Space Exploration: How Genetic Algorithms Find the Needle in the Haystack. *PeerJ Physical Chemistry* vol. 2, e11 (Apr. 2020).

15. Popova, M. *et al.* MolecularRNN: Generating realistic molecular graphs with optimized properties. *arXiv e-prints,* arXiv:1905.13372 (2019).

16. Elton, D. C. *et al.* Deep learning for molecular design—a review of the state of the art. *Molecular Systems Design & Engineering* vol. 4, 828–849 (2019).

17. Smith, D. A., Harrison, A. & Morgan, P. Multiple Factors Govern the Association between Pharmacology and Toxicity in a Class of Drugs: Toward a Unification of Class Effect Terminology. *Chemical Research in Toxicology* vol. 24. PMID: 21391650, 463–474 (2011).

18. Gao, W. & Coley, C. W. The synthesizability of molecules proposed by generative models. *Journal of chemical information and modeling* vol. 60, 5714–5723 (2020).

19. Van Veldhuizen, D. A., Lamont, G. B., *et al. Evolutionary computation and convergence to a pareto front* in *Late breaking papers at the genetic programming 1998 conference* (1998), 221–228.

20. Bagchi, T. P. in *Multiobjective Scheduling by Genetic Algorithms* 171–202 (Springer US, Boston, MA, 1999).

21. Deb, K., Pratap, A., Agarwal, S. & Meyarivan, T. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation* vol. 6, 182–197 (2002).

22. Deb, K. & Jain, H. An Evolutionary Many-Objective Optimization Algorithm Using Reference-Point-Based Nondominated Sorting Approach, Part I: Solving Problems With Box Constraints. *IEEE Transactions on Evolutionary Computation* vol. 18, 577–601 (2014).

23. Jain, H. & Deb, K. An Evolutionary Many-Objective Optimization Algorithm Using Reference-Point Based Nondominated Sorting Approach, Part II: Handling Constraints and Extending to an Adaptive Approach. *IEEE Transactions on Evolutionary Computation* vol. 18, 602–622 (2014).

24. Jensen, J. H. A graph-based genetic algorithm and generative model/Monte Carlo tree search for the exploration of chemical space. *Chemical Science* vol. 10, 3567–3572 (2019).

25. Ishibuchi, H., Imada, R., Setoguchi, Y. & Nojima, Y. *Performance comparison of NSGA-II and NSGA-III on various many-objective test problems* in *2016 IEEE Congress on Evolutionary Computation (CEC)* (2016), 3045–3052.

26. Campos Ciro, G., Dugardin, F., Yalaoui, F. & Kelly, R. A NSGA-II and NSGA-III comparison for solving an open shop scheduling problem with resource constraints. *IFAC-PapersOnLine* vol. 49. 8th IFAC Conference on Manufacturing Modelling, Management and Control MIM 2016, 1272–1277 (2016).
27. Pophale, R., Daeyaert, F. & Deem, M. W. Computational prediction of chemically synthesizable organic structure directing agents for zeolites. *J. Mater. Chem. A* vol. 1, 6750–6760 (23 2013).
28. Besnard, J. *et al.* Automated design of ligands to polypharmacological profiles. *Nature* vol. 492, 215–220 (Dec. 2012).
29. Maier, H. *et al.* Introductory overview: Optimization using evolutionary algorithms and other metaheuristics. *Environmental Modelling & Software* vol. 114, 195–213 (2019).
30. Holland, J. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence* (MIT Press, 1992).
31. Goldberg, D. E. & Holland, J. H. Genetic Algorithms and Machine Learning. *Machine Learning* vol. 3, 95–99 (1988).
32. Anderson, E. *et al. SMILES, a line notation and computerized interpreter for chemical structures* (US Environmental Protection Agency, Environmental Research Laboratory, 1987).
33. Krenn, M., Häse, F., Nigam, A., Friederich, P. & Aspuru-Guzik, A. Self-referencing embedded strings (SELFIES): A 100% robust molecular string representation. *Machine Learning: Science and Technology* vol. 1, 045024 (Oct. 2020).
34. Sterling, T. & Irwin, J. J. ZINC 15 – Ligand Discovery for Everyone. *Journal of Chemical Information and Modeling* vol. 55, 2324–2337 (Nov. 2015).
35. Mendez, D. *et al.* ChEMBL: towards direct deposition of bioassay data. *Nucleic Acids Research* vol. 47, D930–D940 (Nov. 2018).
36. Zhou, Z. & Harris, K. D. M. Counteracting stagnation in genetic algorithm calculations by implementation of a micro genetic algorithm strategy. *Physical Chemistry Chemical Physics* vol. 10, 7262–7269 (48 2008).
37. Nigam, A. *et al.* Augmenting Genetic Algorithms with Deep Neural Networks for Exploring the Chemical Space. *International Conference on Learning Representations* (2020).
38. Mouret, J.-B. & Clune, J. Illuminating search spaces by mapping elites. *arXiv e-prints,* arXiv:1504.04909 (Apr. 2015).
39. Verhellen, J. & Van den Abeele, J. Illuminating elite patches of chemical space. *Chem. Sci.* vol. 11, 11485–11491 (42 2020).
40. Nigam, A., Pollice, R., Krenn, M., Gomes, G. d. P. & Aspuru-Guzik, A. Beyond generative models: superfast traversal, optimization, novelty, exploration and discovery (STONED) algorithm for molecules using SELFIES. *Chem. Sci.* vol. 12, 7079–7090 (20 2021).

41. Krause, E. F. *Taxicab geometry: An adventure in non-Euclidean geometry* (Courier Corporation, 1986).

42. Das, I. & Dennis, J. E. Normal-Boundary Intersection: A New Method for Generating the Pareto Surface in Nonlinear Multicriteria Optimization Problems. *SIAM Journal on Optimization* vol. 8, 631–657 (1998).

43. Blank, J., Deb, K., Dhebar, Y., Bandaru, S. & Seada, H. Generating Well-Spaced Points on a Unit Simplex for Evolutionary Many-Objective Optimization. *IEEE Transactions on Evolutionary Computation* vol. 25, 48–60 (2021).

44. Deb, K., Bandaru, S. & Seada, H. *Generating uniformly distributed points on a unit simplex for evolutionary many-objective optimization* in *International Conference on Evolutionary Multi-Criterion Optimization* (2019), 179–190.

45. Hardin, D. & Saff, E. Minimal Riesz energy point configurations for rectifiable d-dimensional manifolds. *Advances in Mathematics* vol. 193, 174–204 (2005).

46. Pennington, L. D. *et al.* Positional Analogue Scanning: An Effective Strategy for Multiparameter Optimization in Drug Design. *Journal of Medicinal Chemistry* (2020).

47. Hughes, J. *Lazy memo-functions* in *Conference on Functional Programming Languages and Computer Architecture* (1985), 129–146.

48. Lipinski, C. A. *et al.* Experimental and computational approaches to estimate solubility and permeability in drug discovery and development settings. *Advanced Drug Delivery Reviews* vol. 23. In Vitro Models for Selection of Development Candidates, 3–25 (1997).

49. Egan, W. J. *et al.* Prediction of Drug Absorption Using Multivariate Statistics. *Journal of Medicinal Chemistry* vol. 43, 3867–3877 (2000).

50. Veber, D. F. *et al.* Molecular Properties That Influence the Oral Bioavailability of Drug Candidates. *Journal of Medicinal Chemistry* vol. 45, 2615–2623 (2002).

51. Wager, T. T., Hou, X., Verhoest, P. R. & Villalobos, A. Central nervous system multiparameter optimization desirability: application in drug discovery. *ACS chemical neuroscience* vol. 7, 767–775 (2016).

52. Prasanna, S. & Doerksen, R. Topological polar surface area: a useful descriptor in 2D-QSAR. *Current medicinal chemistry* vol. 16, 21–41 (2009).

53. Wildman, S. A. & Crippen, G. M. Prediction of Physicochemical Parameters by Atomic Contributions. *Journal of Chemical Information and Computer Sciences* vol. 39, 868–873 (1999).

54. Jaccard, P. The distribution of the flora in the alpine zone. 1. *New phytologist* vol. 11, 37–50 (1912).

55. Tanimoto, T. T. Elementary mathematical theory of classification and prediction (1958).

56. Morgan, H. L. The Generation of a Unique Machine Description for Chemical Structures-A Technique Developed at Chemical Abstracts Service. *Journal of Chemical Documentation* vol. 5, 107–113 (May 1965).

57. Rogers, D. & Hahn, M. Extended-Connectivity Fingerprints. *Journal of Chemical Information and Modeling* vol. 50, 742–754 (May 2010).

58. Cahart, R., Smith, D. & Venkataraghavan, R. Atom pairs as molecular features in structure–activity studies: definition and application. *J. Chem. Inf. Comput. Sci* vol. 25, 64–73 (1985).

59. Renz, P., Van Rompaey, D., Wegner, J. K., Hochreiter, S. & Klambauer, G. On failure modes in molecule generation and optimization. *Drug Discovery Today: Technologies* vol. 32-33. Artificial Intelligence, 55–63 (2019).

60. Reeves, S. *et al.* Assessing methods and obstacles in chemical space exploration. *Applied AI Letters* vol. 1, e17 (2020).

61. Proschak, E., Stark, H. & Merk, D. Polypharmacology by Design: A Medicinal Chemist's Perspective on Multitargeting Compounds. *Journal of Medicinal Chemistry* vol. 62, 420–444 (2019).

62. Shohat, G., Shani, G., Eisenstein, M. & Kimchi, A. The DAP-kinase family of proteins: study of a novel group of calcium-regulated death-promoting kinases. *Biochimica et Biophysica Acta (BBA) - Proteins and Proteomics* vol. 1600. 7th European Symposium on Calcium-Binding Proteins in Normal and Transformed Cells, 45–50 (2002).

63. Garrido, A., Lepailleur, A., Mignani, S. M., Dallemagne, P. & Rochais, C. hERG toxicity assessment: Useful guidelines for drug design. *European Journal of Medicinal Chemistry* vol. 195, 112290 (2020).

64. Lin, A. *et al.* Off-target toxicity is a common mechanism of action of cancer drugs undergoing clinical trials. *Science translational medicine* vol. 11, eaaw8412 (2019).

65. Kondej, M., Stępnicki, P. & Kaczor, A. A. Multi-target approach for drug discovery against schizophrenia. *International journal of molecular sciences* vol. 19, 3105 (2018).

66. Guerreiro, A. P., Fonseca, C. M. & Paquete, L. The hypervolume indicator: Problems and algorithms. *arXiv preprint arXiv:2005.00515* (2020).

67. Zitzler, E., Brockhoff, D. & Thiele, L. *The hypervolume indicator revisited: On the design of Pareto-compliant indicators via weighted integration* in *International Conference on Evolutionary Multi-Criterion Optimization* (2007), 862–876.

68. Beume, N. S-metric calculation by considering dominated hypervolume as Klee's measure problem. *Evolutionary Computation* vol. 17, 477–492 (2009).

69. Nowak, K., Märtens, M. & Izzo, D. *Empirical performance of the approximation of the least hypervolume contributor* in *International Conference on Parallel Problem Solving From Nature* (2014), 662–671.

70. Biscani, F. & Izzo, D. A parallel global multiobjective framework for optimization: pagmo. *Journal of Open Source Software* vol. 5, 2338 (2020).

71. Miranda-Quintana, R. A., Bajusz, D., Rácz, A. & Héberger, K. Extended similarity indices: the benefits of comparing more than two objects simultaneously. Part 1: Theory and characteristics. *Journal of cheminformatics* vol. 13, 1–18 (2021).

72. Miranda-Quintana, R. A., Rácz, A., Bajusz, D. & Héberger, K. Extended similarity indices: the benefits of comparing more than two objects simultaneously. Part 2: speed, consistency, diversity selection. *Journal of cheminformatics* vol. 13, 1–16 (2021).

73. Chakravorty, S. J. *et al.* Nuisance compounds, PAINS filters, and dark chemical matter in the GSK HTS collection. *SLAS DISCOVERY: Advancing Life Sciences R&D* vol. 23, 532–545 (2018).

74. Lu, T., Pál, D. & Pál, M. *Contextual multi-armed bandits* in *Proceedings of the Thirteenth international conference on Artificial Intelligence and Statistics* (2010), 485–492.

75. Bradford, E., Schweidtmann, A. M. & Lapkin, A. Efficient multiobjective optimization employing Gaussian processes, spectral sampling and a genetic algorithm. *Journal of global optimization* vol. 71, 407–438 (2018).

76. MacLeod, B. P. *et al.* Self-driving laboratory for accelerated discovery of thin-film materials. *Science Advances* vol. 6, eaaz8867 (2020).

77. Reker, D. & Schneider, G. Active-learning strategies in computer-assisted drug discovery. *Drug Discovery Today* vol. 20, 458–465 (2015).

Authors' addresses

Jonas Verhellen University of Oslo, Postboks 1337 Blindern, 0316 Oslo, Norway, jonasver@uio.no

Paper III

Multi-Task Learning of Biophysically-Detailed Neuron Models

Jonas Verhellen, Kosio Beshkov, Sebastian Amundsen, Gaute Einvoll

This manuscript is in progress and contains preliminary results.

Abstract

Understanding the operating principles underlying the electrical activity of the human brain requires integrated research efforts at the molecular, cell, circuit, and systems levels. To study local neural circuits, computationally expensive simulations of biophysically-detailed neuron models can be employed. Recent innovations have shown that it is possible to accurately predict the biophysical behaviour of detailed neuron models with artificial neural networks (ANN) in terms of spikes, electrical potentials and optical readouts. While these methods have the potential to accelerate large network simulations with several orders of magnitude compared to conventional differential equation-based modelling, they currently only predict outputs for the soma or a few chosen neuron compartments. Based on state-of-the-art ANN architectures for multi-task learning (MTL), we present a novel approach allowing for the simultaneous prediction of membrane potentials in each compartment of a neuron model. In addition to providing the necessary information to calculate local field potentials, learning all membrane potentials simultaneously also serves as a challenging benchmark for MTL architectures due to the presence of correlations between membrane potentials in neighbouring compartments and the non-Gaussian distribution of membrane potential values in general.

III.1 Introduction

In the seven decades since Hodgkin and Huxley's original description of the action potential in terms of ion channel gating [1–7], the scientific community has developed a rather comprehensive understanding of how individual neurons process information. Neurons receive thousands of synaptic inputs in dendritic branches where inputs interact with a plethora of local non-linear processes to culminate in the soma which, if sufficiently depolarised, triggers the generation of an action potential. The behaviour of large networks of neurons, however, remains poorly understood. Qualitative insights can be gleaned from statistical correlations between recorded neural activity and sensory stimulation or animal behaviour as seen in experimental studies [8–11]. Unfortunately, these correlations offer little information on how networks of neurons perform neural computation or how these networks give rise to neural representations. Mechanistic modelling, in which detailed neuron models or networks of detailed neuron models are simulated on a computer, offers an alternative approach to study the network dynamics of neural circuits [12, 13].

For a long time, the computational exploration of networks of neurons mimicking entire cortical areas was deemed practically unfeasible due to the extraordinary associated computational cost. A series of relatively recent pioneering efforts, facilitated by large supercomputers, paved the way for this type of neuroscience by constructing simulations containing tens of thousands of model neurons to mimic specific cortical columns in mammalian sensory cortices [14–18]. Even more recent advances have reduced the need of supercomputers for large simulations of detailed neuron networks, by distilling the output of biophysically-detailed neuron models into easier-to-evaluate artificial neural networks (ANN). Previous work focused on predicting outgoing action potentials (i.e. spikes) or other experimental variables (membrane potentials, specific ion currents or optical readouts) in a limited number of compartments [19, 20]. In this paper, we expand on these approaches by simultaneously predicting the presence of an action potential and the membrane potentials for all compartments.

By predicting outgoing action potentials and membrane potentials across the biophysically-detailed neuron model simultaneously, we have made the transition from single-task learning to heterogeneous multi-task learning (MTL). In contrast to single-task learning, heterogeneous MTL does not train a separate model for each target but optimises a single neural architecture to predict multiple heterogeneous targets simultaneously. In our work, task heterogeneity manifests itself as the difference between regression tasks on one hand, i.e. membrane potential predictions, and the binary classification task for the prediction of action potentials on the other hand. MTL approaches aim to leverage statistical relationships between multiple targets to improve generalisation and efficiency across tasks [21–25]. To capture shared patterns that could be missed in single-task learning, multi-task learning generally relies on either one of two categories of neural architectures, respectively known as hard parameter [26] and soft parameter sharing models [27].

Hard parameter sharing models for MTL consist of a shared bottom and task-specific output branches whereas soft parameter sharing models make use of dedicated sets of learning parameters and feature mixing mechanisms. In this paper, we apply a single type of hard parameter sharing model and two types of soft parameter sharing models

(MMoE [28] and MMoEEx [29]) to the challenge of distilling the full electrophysiology of a biophysically-detailed layer 5 pyramidal neuron model [30]. Specific details of the MTL architectures used in this study can be found in the methods and materials section. Importantly, we found that the soft parameter sharing models substantially outperform the hard parameter sharing model on this task. In addition, we tracked the development of diversity in the internal representations of the soft parameter sharing models during training. In the future, we will investigate the importance of different aspects of the training regime with regard to task-balancing and the generalisation properties of several optimisation schemes.

III.2 Results

To test the capability of different MTL architectures in accurately representing the full dynamic membrane potential of each compartment of a large biophysically-detailed neuron model, we trained the architectures on a balanced dataset of simulated neural activity. For each target (a collection of membrane potentials and the presence or absence of an outgoing action potential) in the dataset, we provide a 100 ms history of neural activity and synaptic inputs from all compartments of the biophysically-detailed neuron model to the MTL models. Further details regarding the dataset and pre-processing steps can be found in the methods and materials section. To facilitate a fair comparison between MTL methods, we constructed a hard parameter sharing model (14 million trainable parameters, 27.930 MB) that is close in size to the soft parameter sharing models (12 million trainable parameters, 23.918 MB). Similarly, we used the same training procedures for each of the models (Adam [31] without task-balancing) and evaluated the inference speed in a single session on publicly available hardware.

III.2.1 Multi-Task Prediction of Membrane Potential Dynamics in Biophysically-Detailed Neuron Models

After training all three MTL models, we observed that the soft parameter sharing models strongly outperformed the hard parameter sharing model both in terms of training and generalisation loss, see Figure III.1. During training, the hard parameter model reached a training loss of 0.439 and a validation loss of 0.680, the MMoE model reached a training loss of 0.241 and a validation loss of 0.533, and the MMoEEx model reached a training loss of 0.262 and a validation loss of 0.551. The MMoE model, saved at its best validation performance, predicts the membrane potential of a compartment up to an average accuracy of 3.77 mV. Unfortunately, with the current set of hyperparameters in the loss function ($\gamma = 1$, $w_i = w_{spike} = 1$), all of the explored models prioritise learning the membrane potentials of the compartments of the neuron model over learning an accurate prediction of spike generation in the soma.

Figure III.1: Training loss (left) and validation loss (right) of the hard parameter model, the MMoE model, and the MMoEEx model.

III.2.2 Measuring Experts Diversity in MMoE and MMoEEx Trained on Neural Data

In previous work on soft parameter sharing models, it has been argued that a higher diversity among experts could, in some cases, lead to an improvement in training and generalisation results for heterogeneous MTL [28, 29]. To analyse the influence of expert diversity on the soft parameter sharing models presented in this paper, we computed several diversity metrics during and after the training of the MMoE and MMoEEx models. In addition to the diversity score, which was previously introduced as a measure of expert diversity, we also tracked the determinant and permanent of the standardised distance matrix between experts. While MMoEEx was originally introduced to the literature as a method to induce more diversity among experts, for the task at hand however, we observed that the MMoEEx expert diversity converged to a lower value than the final MMoE expert diversity in all three of our measures. Curiously, for both models the diversity measures initially increase and afterwards either stabilise (in the case of MMoE) or reduce to down below their initial values (in the case of MMoEEx), see the middle panel of Figure III.2.

The distances between the experts in MMoE and MMoEEx, at the end of the training procedure, are presented in the heatmaps at the top of Figure III.2. It is important to note that, the MMoEEx model has an additional hyperparameter α, which controls the number of tasks each expert participates in. For the present task, α was set to 0.1, which means that each expert of MMoEEx learns a representation that can contribute to 90% of the tasks. It is possible that after a thorough hyperparameter scan, a value of α can be found for which the MMoEEx model outperforms MMoE. Finally, given the large number of partially correlated tasks in the distillation of a biophysically-detailed neuron model, we were in an unique position to explore how individual experts in MMoE and MMoEEx contribute to correlated prediction tasks. In both models, we observe

correlations between the mean weights projected from each expert to each compartment for neighbouring compartments, see bottom panel of Figure III.2. Finally, it is also worth noting that despite the difficulties in successfully predicting somatic spikes, this task receives high average weights from some of the experts in MMoEEx.

III.2.3 Ultra-fast Simulation of Full Membrane Potential Dynamics of Multiple Cells

Traditional simulation environments such as NEURON rely on numerical integration of compartment-specific differential equations that represent the active and passive biophysical mechanisms of the modeled neuron. This approach requires a significant amount of computational resources and one of the main attractive features of distilling biophysically-detailed neuron models into ANNs is the significant reduction in simulation runtimes deep learning architectures provide. Previous work from the literature, where only the output of a limited number of compartments was predicted, has already shown that NEURON simulations of biophysically-detailed neuron models are significantly slower than their ANN counterparts. Especially in simulations with multiple neurons where accelerators like GPU's can be used to parallel process independent timesteps of the model. In practice, the batch-size of the MTL models during inference equals the amount of timesteps individually evaluated in parallel and can be most easily interpreted as the amount of model neurons being evaluated at the same time.

In Table III.1, we recorded the mean and standard deviation runtimes of seven independent simulations for each of the three MTL models. Note that each simulated timestep is independent of the others and that the simulation time presented in the table is cumulative. So, according to the previously mentioned interpretation, we have simulated 100 ms of single neuron model, simultaneous 1 ms simulations for a 100 neurons and simultaneous 10 ms simulations for 100 neurons. Clearly, being able to evaluate the MTL models in parallel on a GPU has a significant effect on runtimes: a 1000 ms of simulation time (spread over a 100 neurons) can be predicted in near real time by the hard parameter sharing model. MMoE and MMoEEx have similar inference runtimes, but are significantly slower than the hard parameter sharing model. This is likely due to the lower amount of convolutions filters in the hard parameter sharing model and the presence of large matrix multiplications in the gates of the soft parameter sharing models.

III.3 Discussion

III.3.1 Advantages and Limitations of Multi-task Learning in Distilling Biophysically-Detailed Neuron Models

We explored the capacity of three state-of-the-art MTL neural network architectures [26, 28, 29] in the distillation of the complex membrane potential dynamics of a multi-compartment, biophysically-detailed model of a layer 5 pyramidal neuron [30]. Given the fact that the biophysical model under study has 639 compartments, this is an unusually challenging MTL problem. The current results, as reported here, are

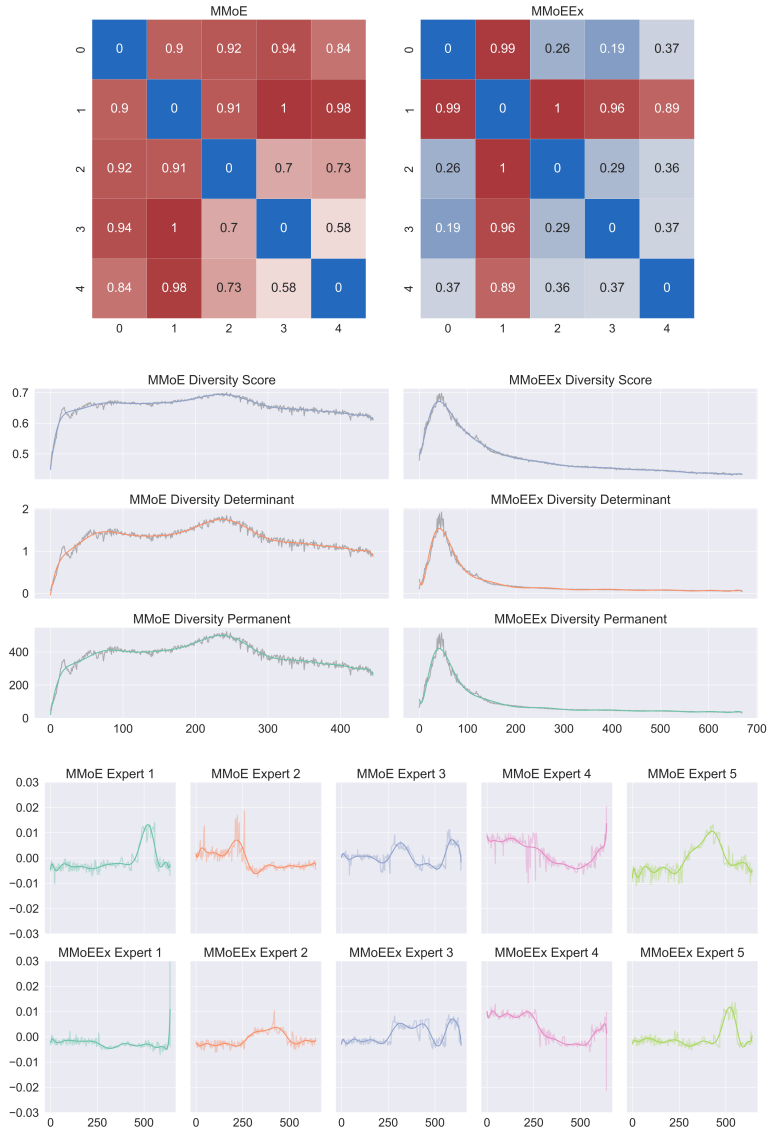

Figure III.2: **Top** - Heatmaps of the distances between the experts in the MMoE ($\bar{d}_1 = 0.68$, $\bar{d}_2 = 1.57$, and $\bar{d}_3 = 451.63$) and MMoEEx ($\bar{d}_1 = 0.45$, $\bar{d}_2 = 0.09$, and $\bar{d}_3 = 47.69$) models. **Middle** - Smoothed diversity measures (colours) for MMoE (left) and MMoEEx (right) shown together with the corresponding raw values (grey). **Bottom** - Smoothed curves (dark colours) of the average weights projecting from the 5 experts to the towers for MMoE (top) and MMoEEx (bottom) shown together with the corresponding raw values (light colours). The final average weight value represents the tower responsible for the binary prediction of a spike in the soma.

Table III.1: Inference Speed of the Three MTL Architectures after Training on Neural Data Generated by a Biophysically-Detailed Model of a Layer 5 Pyramidal Neuron.

Runtime (s)	Std. Dev. (ms)	Model	Hardware	Datapoints	Batch Size
21.2	287	MMoE	CPU	100	1
21.4	361	MMoEEx	CPU	100	1
7.50	284	Hard Parameter	CPU	100	1
20.4	120	MMoE	GPU	100	1
20.3	217	MMoEEx	GPU	100	1
7.83	286	Hard Parameter	GPU	100	1
0.30	3.25	MMoE	GPU	100	100
0.31	2.24	MMoEEx	GPU	100	100
0.19	12.4	Hard Parameter	GPU	100	100
2.10	89.7	MMoE	GPU	1000	100
2.18	164	MMoEEx	GPU	1000	100
0.98	205	Hard Parameter	GPU	1000	100

encouraging. The MMoE model with the lowest validation loss has a mean average error of 3.77 mV in membrane potential prediction across all compartments. This mean average error is well below the standard deviation of experimentally measured action potential peak membrane voltages in layer 5 pyramidal neurons during perisomatic step current firing (4.97 mV - 6.93 mV) or back-propagating action potential Ca^{2+} firing (5 mV) [30, 32]. One point worth noting is that, up to now, all the trained models failed to learn the binary prediction of somatic spikes. This issue can potentially be addressed by increasing the value of the hyperparameter γ and hence giving more weight to this specific task during learning.

Further exploration of the soft parameter sharing models through the use of task balancing methods, as discussed in the methods and materials section, might improve upon the results obtained so far. Hard parameter sharing models typically do not benefit from task balancing approaches. Experimental recordings of the after-hyperpolarization depth of the membrane potential in the soma of layer 5 pyramidal neurons have a slightly lower standard deviation (3.58 mV - 5.82 mV) than electrophysiological measurements of action potentials [30, 33, 34], indicating that a more sensitive evaluation of distilled neuron models could be based on their performance in specific neuronal scenario's. It is also important to note that one expected use-case of these MTL models is in the acceleration of LFP calculations where it is well known that the low frequency aspects of membrane potential dynamics matter most [35–37]. In conclusion, it seems that distilled biophysically-detailed neuron models, while already effective, could be improved based on bespoke neuroscience-based metrics rather than those regularly used in deep learning.

III.3.2 Measuring Experts Diversity in MMoE and MMoEEx Trained on Neural Data

To test the conjecture that in soft parameter sharing MTL models high diversity between experts can be beneficial in training and generalisation, we computed several diversity

measures. Indeed, the model with highest diversity (MMoE) performed the best, despite the fact that unlike its counterpart (MMoEEx), it has no explicit inductive bias towards expert diversity. Further exploration of the exclusivity hyperparameter α might lead to more desirable results for the MMoEEx model. Additionally, increasing the number of experts should only lead to improvements in training and generalisation if a substantial amount of independence between the tasks had not yet been incorporated. Given the low dimensionality of the data used in this project, as described in the methods and materials section, it is unlikely that our current MTL problem would benefit appreciably from a higher number of experts. Future work can make these statements more precise by studying the contribution of individual experts to prediction of membrane potentials in morphologically distinct neuronal compartments.

III.3.3 Hardware Acceleration of Distilled Neuron Models and Simulation Runtimes

Biophysically-detailed neuron models distilled into ANNs can be evaluated at significantly higher speeds than their classical counterparts. Because a single instance of a deep learning model can be used to predict outputs for multiple instances of the same neuron model in parallel, on accelerators such as GPUs, these models are deemed to be particularly suited to accelerating large networks of model neurons. The biophysically-detailed model of the mouse V1 cortical area developed by the Allen Institute [38], for instance, could be accelerated by 114 deep learning models each representing one model neuron type. The current biophysically-detailed Allen V1 model makes use of neuron models with passive dendrites which should be significantly easier to distil into an MTL architecture than the layer 5 pyramidal neuron model represented here. Further acceleration of the MTL models discussed in this paper could be achieved by running model inference on multiple GPUs or more advanced accelerators such as TPUs and IPUs.

III.4 Methods

III.4.1 Multi-compartmental NEURON Simulations and Data Balancing

As a baseline for training and testing, we used an existing dataset of electrophysiological data generated in NEURON [19] based on a well-known biophysical-detailed and multi-compartment model of cortical layer V pyramidal cells [30]. This model contains a wide range of dendritic (Ca^{2+}-driven) and perisomatic (Na^{+}-driven) active properties which are represented by ten key active ionic currents that are unevenly distributed over different dendritic compartments. The data was generated in response to presynaptic spike trains sampled from a Poisson process with additional temporal variety due to resampling. For the purposes of this paper, it is important to note that the biophysical-detailed model contains 639 compartments and 1278 synapses and that 128 simulations of the complete model for 6 seconds of biological time each were included before data balancing.

The subthreshold dynamics of the membrane potential in a compartment have weak variance and as a result one would expect them to be easy to predict. However, suprathreshold deviations generated by dendritic spiking events can be more problematic to predict. To address this issue, we implemented a form of data balancing. We first identified the time points at which somatic spikes occurred and afterwards we used them to create a dataset in which one third of the targets included a spiking event. Additionally, we standardised the membrane potential through z-scoring. Principal components analysis on the target data showed that the largest 5 principal components explain 88% of the variance in the data. We used input data consisting of membrane potentials and incoming synaptic events across all compartments during a 100 ms time window, and target data consisting of 1 ms of membrane potentials and a binary value for the presence or absence of a somatic spike.

III.4.2 Hard and Soft Parameter Sharing Architectures for Multi-task Learning

To learn spatiotemporal relationships between neural inputs and outputs (synaptic events and voltage traces), we make use of a generic temporal convolutional network (TCN) architecture [39]. Whereas learning on sequence tasks is often associated with recurrent architectures [40, 41] such as GRU's and LSTM's, previous research has shown that certain convolutional architectures can outperform canonical recurrent networks on tasks such as audio synthesis, language modelling, and machine translation without suffering from the vanishing gradient problem or a lack in memory retention. A TCN is a standard one-dimensional convolutional network endowed with a causal structure that has been used previously to learn representations of neural data. The causality in a TCN is guaranteed, layer by layer, by solely applying causal convolutions and adding appropriated zero-padding and dilation. In this work, we use a TCN architecture to learn representations to feed into hard and soft parameter sharing architectures for multi-task learning.

One of the initial works discussing modern multi-task learning makes use of hard parameter sharing through shared initial layers of the neural network architecture ("the bottom") and top layers ("the towers") which are task-specific. This approach has the advantage of scaling well with the number of tasks but the disadvantage that the shared representation learned by the bottom can become biased towards the tasks with a strong loss signal. To improve on hard parameter sharing, soft parameter sharing architectures make use of dedicated representations for each task. In this paper, we use the multi-gate mixture-of-experts (MMoE) [28] and multi-gate mixture-of-experts with exclusivity (MMoEEx) [29] architectures as soft parameter sharing models. These models combine representations learned by several shared bottoms ("the experts") using gating functions, which apply linear combinations based on learnable and data-dependent weights. MMoEEx expands on MMoE by forcing a random subsection of these weights to be zero to encourage diversity among the representations provided by the experts.

We use a three layered TCN (32, 16, and 8 channels respectively) with a kernel size of 10 and a dropout of 0.2 for the bottom architecture and for each of the expert architectures. For each of the towers we use three-layered feed-forward networks with

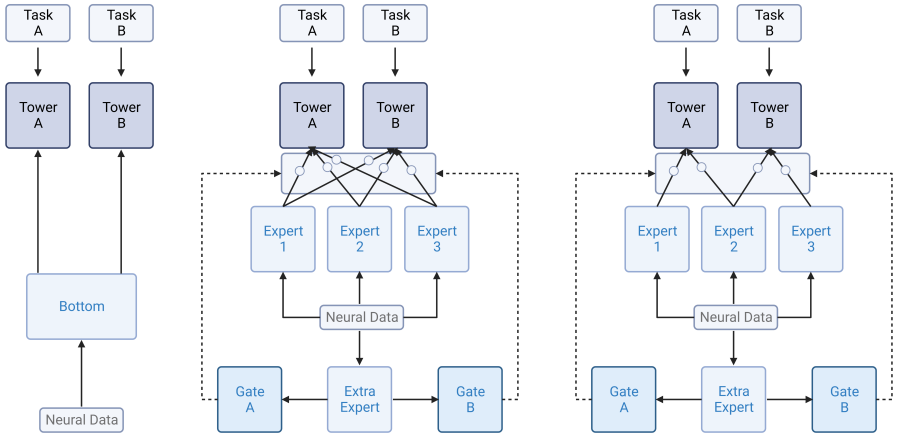

Figure III.3: Schematic representation of the architectures of the hard parameter model (left), the MMoE model (middle), and the MMoEEx model (right).

ELU activation functions and 10 (soft parameter sharing) or 25 (hard parameter sharing) hidden nodes each. We replace the final RELU activation function in the original TCN implementation with a sigmoid activation function for increased stability. For the soft parameter sharing models, we added in an extra TCN expert, see Figure III.3, that compresses the data before feeding it into the gating functions to avoid quadratic growth of the trainable parameters with respect to the size of the input data (time-window length times the amount of features). Based on the previously discussed principal component analysis of the target data, we chose to use 5 experts for both MMoE and MMoEEx in the numerical experiments presented in this paper. To facilitate efficient training, we implemented mixed precision and multi-GPU training for all three models. All models were trained on (up to 6) parallel RTX2080Ti GPUs, while inference experiments were performed on Tesla P100 GPUs, which are publicly available through Kaggle.

III.4.3 Task Balancing and Expert Diversity

Effective multi-task learning typically requires some form of task balancing to reduce negative transfer or to prevent one or more tasks from dominating the optimisation procedure. To avoid these issues we make use of loss-balanced task weighting (LBTW) [42] which dynamically updates task weights in the loss function during training. For each batch, LBTW calculates loss weights based on the ratio between the current loss and the initial loss for each task, and a hyperparameter ξ. As ξ goes to 0, LBTW approaches standard multitask learning training. All taken together, the loss function can be summarised as

$$L(\hat{y}, y) = L_{spike}(\hat{y}_{spike}, y_{spike}) + \sum_{i=0}^{N} L_i(\hat{y}_i, y_i)\,, \tag{III.1}$$

or equivalently,

$$L(\hat{y},y) = w_{spike} \cdot \gamma \cdot \mathrm{BCE}_{logits}(\hat{y}_{spike}, y_{spike}) + \sum_{i=0}^{N} w_i \cdot \mathrm{MSE}(\hat{y}_i, y_i)\,, \tag{III.2}$$

where $\hat{y}$ and y are the target and the predicted data respectively, w_{spike} and w_i are the task weight, γ is a hyperparameter that determines the importance of spike prediction, BCE_{logits} is the binary cross entropy loss combined with a sigmoid activation function, and MSE is the mean square error. The summation index i runs over all N compartments of the biophysically-detailed neuron model, and the LBTW task weights are recalculated every epoch E, for each batch B, according to

$$w_{spike} = \left(\frac{L_{spike,E,B}(\hat{y}_{spike}, y_{spike})}{L_{spike,E,0}(\hat{y}_{spike}, y_{spike})}\right)^{\xi} \quad \text{and} \quad w_i = \left(\frac{L_{i,E,B}(\hat{y}_i, y_i)}{L_{i,E,0}(\hat{y}_i, y_i)}\right)^{\xi}. \tag{III.3}$$

During training of the different multi-task methods, we consistently made use of a batch-size of 32 data points. To explore different optimisation methods we tested both Adam [31] with a starting learning rate of 0.001 and a weight decay (i.e. L^2 penalty) of 0.0001, and SGD with warm restarts using a cosine annealing schedule [43] defined by

$$\eta_t = \eta_{min} + \frac{1}{2}(\eta_{max} - \eta_{min})\left(1 + \cos\left(\frac{T_{cur}}{T_i}\pi\right)\right) \tag{III.4}$$

where η_{max} is the initial learning rate, T_{cur} is the number of epochs since the latest restart and T_i is the number of epochs between two restarts. When $T_{cur} = T_i$ in this annealing schedule, we set η_t equal to η_{min} and when T_{cur} is set to 0 after a restart, we also set η_t equal to η_{max}. For the SGD optimiser itself, we use a starting learning rate of 0.05 and a weight decay of 0.0005. We set the dampening to 0, and the Nesterov momentum [44] to 0.9.

To measure the diversity in expert representations in MMoE and MMoEEx, we made use of diversity measurement proposed in the original MMoEEx paper. In that paper, the diversity between two experts n and m is calculated as a (real-valued) distance $d(n,m)$ between the learned representations f_n and f_m, as defined by

$$d(n,m) = \sqrt{\sum_{i=0}^{N} (f_n(x_i) - f_m(x_i))^2}\,, \tag{III.5}$$

where N is the number of samples x_i in the dataset (often the validation dataset) which is used to probe the diversity of the experts. The diversity matrix D of a trained MMoE or MMoEEx model is defined by calculating all pairwise distances between expert representations as described above and normalising the matrix. In this normalised matrix, a pair of experts with distance close to 0 are considered near-identical, and experts with a distance close to 1 are considered to be highly diverse. To compare two different models in terms of overall diversity, we respectively define the first, second, and third diversity scores of a model as the mean entry ($\bar{d}_1$) and the determinant ($\bar{d}_2$), and the permanent ($\bar{d}_3$) of its diversity matrix D.

III.5 Data and Code Availability

Pre-processed neural data, numerical results, complete computational workflows, and the weights of the trained MTL networks will be made publicly available together with the completed version of this manuscript.

III.6 Acknowledgments

We wish to thank Torbjørn Vefferstad Ness for his insightful comments on the topic of learning computational neuron models with artificial neural networks and for referring the authors to relevant literature. This research was funded by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement Nº 945371 and by UiO:Life Science through the 4MENT convergence environment. Figure III.3 has been made with Biorender.com through an academic license.

References

1. Hodgkin, A. L., Huxley, A. F. & Katz, B. Measurement of current-voltage relations in the membrane of the giant axon of Loligo. *The Journal of physiology* vol. 116, 424 (1952).
2. Hodgkin, A. L. & Huxley, A. F. Currents carried by sodium and potassium ions through the membrane of the giant axon of Loligo. *The Journal of physiology* vol. 116, 449 (1952).
3. Hodgkin, A. L. & Huxley, A. F. The components of membrane conductance in the giant axon of Loligo. *The Journal of physiology* vol. 116, 473 (1952).
4. Hodgkin, A. L. & Huxley, A. F. The dual effect of membrane potential on sodium conductance in the giant axon of Loligo. *The Journal of physiology* vol. 116, 497 (1952).
5. Hodgkin, A. L. & Huxley, A. F. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of physiology* vol. 117, 500 (1952).
6. Häusser, M. The Hodgkin-Huxley theory of the action potential. *Nature neuroscience* vol. 3, 1165–1165 (2000).
7. Brown, A. *The Hodgkin and Huxley papers: still inspiring after all these years* 2022.
8. Hubel, D. H. & Wiesel, T. N. Receptive fields of single neurones in the cat's striate cortex. *The Journal of physiology* vol. 148, 574 (1959).
9. Hubel, D. H. & Wiesel, T. N. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of physiology* vol. 160, 106 (1962).

10. Hubel, D. H. & Wiesel, T. N. Ferrier lecture-Functional architecture of macaque monkey visual cortex. *Proceedings of the Royal Society of London. Series B. Biological Sciences* vol. 198, 1–59 (1977).

11. Hubel, D. H. & Wiesel, T. N. Receptive fields and functional architecture of monkey striate cortex. *The Journal of physiology* vol. 195, 215–243 (1968).

12. Sterratt, D., Graham, B., Gillies, A. & Willshaw, D. *Principles of computational modelling in neuroscience* (Cambridge University Press, 2011).

13. Einevoll, G. T. *et al.* The scientific case for brain simulations. *Neuron* vol. 102, 735–744 (2019).

14. Traub, R. D. *et al.* Single-column thalamocortical network model exhibiting gamma oscillations, sleep spindles, and epileptogenic bursts. *Journal of neurophysiology* vol. 93, 2194–2232 (2005).

15. Potjans, T. C. & Diesmann, M. The cell-type specific cortical microcircuit: relating structure and activity in a full-scale spiking network model. *Cerebral cortex* vol. 24, 785–806 (2014).

16. Markram, H. *et al.* Reconstruction and simulation of neocortical microcircuitry. *Cell* vol. 163, 456–492 (2015).

17. Schmidt, M., Bakker, R., Hilgetag, C. C., Diesmann, M. & van Albada, S. J. Multi-scale account of the network structure of macaque visual cortex. *Brain Structure and Function* vol. 223, 1409–1435 (2018).

18. Arkhipov, A. *et al.* Visual physiology of the layer 4 cortical circuit in silico. *PLoS computational biology* vol. 14, e1006535 (2018).

19. Beniaguev, D., Segev, I. & London, M. Single cortical neurons as deep artificial neural networks. *Neuron* vol. 109, 2727–2739 (2021).

20. Olah, V. J., Pedersen, N. P. & Rowan, M. J. Ultrafast simulation of large-scale neocortical microcircuitry with biophysically realistic neurons. *Elife* vol. 11, e79535 (2022).

21. Finn, C., Abbeel, P. & Levine, S. *Model-agnostic meta-learning for fast adaptation of deep networks* in *International conference on machine learning* (2017), 1126–1135.

22. Zhang, Y. & Yang, Q. An overview of multi-task learning. *National Science Review* vol. 5, 30–43 (2018).

23. Ruder, S. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098* (2017).

24. Lee, S. & Son, Y. Multitask learning with single gradient step update for task balancing. *Neurocomputing* vol. 467, 442–453 (2022).

25. Mallya, A., Davis, D. & Lazebnik, S. *Piggyback: Adapting a single network to multiple tasks by learning to mask weights* in *Proceedings of the European Conference on Computer Vision (ECCV)* (2018), 67–82.

26. Caruana, R. *Multitask learning: A knowledge-based source of inductive bias1* in *Proceedings of the Tenth International Conference on Machine Learning* (1993), 41–48.

27. Duong, L., Cohn, T., Bird, S. & Cook, P. *Low resource dependency parsing: Cross-lingual parameter sharing in a neural network parser* in *Proceedings of the 53rd annual meeting of the Association for Computational Linguistics and the 7th international joint conference on natural language processing (volume 2: short papers)* (2015), 845–850.

28. Ma, J. *et al. Modeling task relationships in multi-task learning with multi-gate mixture-of-experts* in *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining* (2018), 1930–1939.

29. Aoki, R., Tung, F. & Oliveira, G. L. Heterogeneous multi-task learning with expert diversity. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* vol. 19, 3093–3102 (2022).

30. Hay, E., Hill, S., Schürmann, F., Markram, H. & Segev, I. Models of neocortical layer 5b pyramidal cells capturing a wide range of dendritic and perisomatic active properties. *PLoS computational biology* vol. 7, e1002107 (2011).

31. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

32. Le Bé, J.-V., Silberberg, G., Wang, Y. & Markram, H. Morphological, electrophysiological, and synaptic properties of corticocallosal pyramidal cells in the neonatal rat neocortex. *Cerebral cortex* vol. 17, 2204–2213 (2007).

33. Larkum, M. E., Zhu, J. J. & Sakmann, B. A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature* vol. 398, 338–341 (1999).

34. Larkum, M. E., Zhu, J. J. & Sakmann, B. Dendritic mechanisms underlying the coupling of the dendritic with the axonal action potential initiation zone of adult rat layer 5 pyramidal neurons. *The Journal of physiology* vol. 533, 447–466 (2001).

35. Pettersen, K. H. & Einevoll, G. T. Amplitude variability and extracellular low-pass filtering of neuronal spikes. *Biophysical journal* vol. 94, 784–802 (2008).

36. Pettersen, K. H., Lindén, H., Dale, A. M. & Einevoll, G. T. Extracellular spikes and CSD. *Handbook of neural activity measurement* vol. 1, 92–135 (2012).

37. Lindén, H., Pettersen, K. H. & Einevoll, G. T. Intrinsic dendritic filtering gives low-pass power spectra of local field potentials. *Journal of computational neuroscience* vol. 29, 423–444 (2010).

38. Billeh, Y. N. *et al.* Systematic integration of structural and functional data into multi-scale models of mouse primary visual cortex. *Neuron* vol. 106, 388–403 (2020).

39. Bai, S., Kolter, J. Z. & Koltun, V. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271* (2018).

40. Hochreiter, S. & Schmidhuber, J. Long short-term memory. *Neural computation* vol. 9, 1735–1780 (1997).

41. Chung, J., Gulcehre, C., Cho, K. & Bengio, Y. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555* (2014).

42. Liu, S., Liang, Y. & Gitter, A. *Loss-balanced task weighting to reduce negative transfer in multi-task learning* in *Proceedings of the AAAI conference on artificial intelligence* **33** (2019), 9977–9978.

43. Loshchilov, I. & Hutter, F. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983* (2016).

44. Sutskever, I., Martens, J., Dahl, G. & Hinton, G. *On the importance of initialization and momentum in deep learning* in *International conference on machine learning* (2013), 1139–1147.

Authors' addresses

Jonas Verhellen University of Oslo, Postboks 1337 Blindern, 0316 Oslo, Norway, jonasver@uio.no

Kosio Beshkov University of Oslo, Postboks 1337 Blindern, 0316 Oslo, Norway, kosio.neuro@gmail.com

Sebastian Amundsen University of Oslo, Postboks 1337 Blindern, 0316 Oslo, Norway, sebastian.amundsen@fys.uio.no

Gaute Einvoll University of Oslo, Postboks 1337 Blindern, 0316 Oslo, Norway, gaute.einevoll@fys.uio.no