



LUND UNIVERSITY

Modelling and analysis of multi-scale networks

Ekström, Henrik

2024

Document Version:
Peer reviewed version (aka post-print)

[Link to publication](#)

Citation for published version (APA):
Ekström, H. (2024). *Modelling and analysis of multi-scale networks*. [Doctoral Thesis (compilation), Centre for Mathematical Sciences]. Lunds universitet, Media-Tryck .

Total number of authors:
1

General rights

Unless other specific re-use rights are stated the following general rights apply:
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Modelling and analysis of multi-scale networks

HENRIK EKSTRÖM

Lund University
Faculty of Science
Centre for Mathematical Sciences
Mathematical Statistics



Modelling and analysis of multi-scale networks

Modelling and analysis of multi-scale networks

Henrik Ekström



LUND
UNIVERSITY

Faculty Opponent

Senior Lecturer Jakob Björnberg
Chalmers University of Technology and University of Gothenburg
Gothenburg, Sweden

Evaluation Committee

Professor Maria Deijfen
Stockholm University
Stockholm, Sweden

Professor Jeffrey Steif
Chalmers University of Technology
Gothenburg, Sweden

Professor Maurice Duits
KTH Royal Institute of Technology
Stockholm, Sweden

Professor Anders Irbäck (substitute)
Lund University
Lund, Sweden

The thesis work is supported by Wallenberg AI, Autonomous Systems and Software Program (WASP), funded by the Knut and Alice Wallenberg Foundation.

© Henrik Ekström 2024

Faculty of Science , Department of Mathematics

ISBN: 978-91-8104-123-1 (print)

ISBN: 978-91-8104-124-8 (pdf)

ISSN: 1404-0034 2024:3

ISRN: LUNFMS-1031-2024

Printed in Sweden by Media-Tryck, Lund University, Lund 2024



Dedicated to all of you who made this happen.

Contents

List of publications	iii
Popular summary	v
Populärvetenskaplig sammanfattning	vii
Acknowledgements	ix
1 Introduction	1
1.1 Mathematically Modelling the Brain	1
1.2 Representation by graphs	2
1.3 Multiple levels of scale	3
1.4 The Hopfield Inverse Problem	5
2 An overview of the papers	7
2.1 Paper I, The Phases of a Discrete Flow of Particles on Graphs	7
2.2 Paper II, Non-Monotone Cellular Automata: Order Prevails Over Chaos	8
2.3 Paper III, A Central Limit Theorem and Decay of Correlations in the Coulomb Chain	10
2.4 Paper IV, Correlation Decay in the Coulomb Chain with Local Pairwise Interactions	13

List of publications

This thesis is based on the following publications:

- I **The Phases of a Discrete Flow of Particles on Graphs**
H. Ekström,
Markov Processes and Related Fields,
Special issue: Structure of Mathematical Physics, **26**(2), 343-364 (2020).

- II **Non-monotone cellular automata: Order prevails over chaos**
H. Ekström, T. Turova,
Biosystems, **120** (2022).

- III **A Central Limit Theorem and Exponential Correlation Decay
for the Coulomb Chain**
H. Ekström,
Pre-print, (2024).

- IV **Correlation decay in the Coulomb chain with local pairwise interactions**
H. Ekström,
Pre-print, (2024).

Popular summary

The term ‘artificial intelligence’, AI, is often used in popular culture whenever computers do things which were previously done by humans and require some form of reasoning beyond what we readily expect computers to be able to perform. The most common type of such a system falls under the umbrella term ‘Machine Learning’. These models use algorithms which we can describe and fine-tune but, at least currently, do not fully understand. Most machine learning algorithms are descended from early mathematical models which were inspired by last century’s progress in describing the components of a brain. A brain, be it humanoid, reptiloid or something else, consists of various types of neurons: small units which send electrical and chemical signals to each other. Each individual piece is rather well understood, but it is still a mystery how the combination of roughly, in the human case, several hundred billion such pieces connected by some quadrillion connections actually *works*. This thesis considers several disparate models which, while simple to define, exhibit interesting emergent behaviours.

A computer is a good example of something which despite consisting of simple units can perform astoundingly complex operations. At the ‘bottom’ level, you have a collection of ones and zeroes together with a few simple logical operations that combine them. At the ‘top’ level, you have a powerful machine capable of anything from making weather predictions to sharing cute photos of kittens wearing silly hats. In between, there are many, many layers of complexity which each expand your available set of operations using combinations of the previously available ones.

A brain also consists of simple units which we by now have quite a good understanding of. However, the levels of increasing complexity is still very mysterious. How is it that some neurons in your brain receiving stimuli from their ‘previous’ neighbours and sending further stimulus to their ‘next’ neighbours can result in anything from alerting you that the ground feels cold to imagining what a conver-

sation with a sentient box girder bridge would be like? We still lack intermediary ‘programming languages’ to take us from one point to the other. These trains of thought were partly what led to the investigations performed in this thesis. Owing to the complexity of the issue, the connections from these initial ideas to the presented results are not necessarily straightforward. However, an important step in trying to solve any big problem is to identify and examine the smaller problems you encounter on the way.

Something which becomes quickly apparent when trying to probabilistically tackle the kind of models considered in this thesis is the level of dependency between the components. Most computer algorithms, while complex, are usually ‘linear’ in that every step builds on the result of the previous one. Furthermore, its physical makeup remains constant. A neuronal network does not work that way. In nature every neuron continuously interacts with its surroundings, and the surroundings change over time as a response to those interactions. A brain is *plastic*, meaning that the existence and strength of connections between neurons change over time. The plasticity itself is not treated in this work, and where relevant (specifically, in Paper II), the model components update synchronously. Even then, the behaviour is quite difficult to describe. In Paper II we derive results which describe the behaviour of a so called *cellular automaton* in certain special cases where this dependency is less impactful and discuss the behaviour observed via simulations in the remaining cases. Papers III and IV deal with a model which is less related to neuronal networks but where a similar dependency structure appears. There, a line of charged particles is considered where each particle only senses the charge of a given number of its nearest neighbours. Despite this, all particles are indirectly affecting each other. A range of methods (which are interesting in themselves) are used to prove that this dependency rapidly decreases when considering particles far apart. This allows a central limit theorem to be derived, at least in the case of only nearest-neighbour interactions.

It would be excessive to claim that this work takes us much closer to explaining the marvel that is the brain, but the pursuit has brought with it many interesting results that will hopefully be useful in following research, and it has raised some fascinating questions which can be explored further.

Populärvetenskaplig sammanfattning

Termen “artificiell intelligens”, AI, används ofta i populärkultur när datorer gör saker som tidigare gjorts av människor och kräver någon form av resonemang utöver vad vi först tänker att datorer kan utföra. De vanligaste typerna av sådana system faller under paraplybegreppet *maskininlärning*. Dessa modeller använder algoritmer som vi kan beskriva och finjustera men, åtminstone just nu, inte helt förstår. De flesta maskininlärningsalgoritmer är ättlingar till tidiga matematiska modeller som inspirerades av förra århundradets framsteg i att beskriva de biologiska komponenterna av en hjärna. En hjärna, så om den är humanoid, reptiloid eller något annat, består av olika typer av neuroner: små komponenter som skickar elektriska och kemiska signaler till varandra. Varje enskild komponent är ganska väl förstådd, men det är fortfarande ett mysterium hur kombinationen av ungefär, i det mänskliga fallet, flera hundra miljarder sådana, sammankopplade med några kvadriljoner förbindelser, faktiskt fungerar. Denna avhandling behandlar flera olika modeller som, även om de är enkla att definiera, besitter intressanta beteenden.

En dator är ett bra exempel på något som består av enkla komponenter men trots det kan utföra häpnadsväckande komplexa operationer. På “nedersta” nivå har man en samling ettor och nollor tillsammans med några enkla logiska operationer som kombinerar dem. På “toppnivån” har man en kraftfull maskin som kan göra allt från att göra väderprognoser till att dela foton på katter i lustiga hattar. Däremellan finns det många, många lager av komplexitet som var och en utökar ens tillgängliga uppsättning operationer med hjälp av kombinationer av de redan tillgängliga. En hjärna består också av enkla enheter som vi vid det här laget har en ganska god förståelse för. Nivåerna av ökande komplexitet är dock fortfarande väldigt mystiska. Hur kommer det sig att vissa neuroner i din hjärna

som tar emot stimuli från sina "tidigare" grannar och skickar ytterligare stimuli till sina "nästkommande" grannar kan resultera i allt från att varna dig om att marken känns kall till att föreställa dig hur en konversation med en talande lådbalksbro skulle vara? Vi saknar fortfarande mellanliggande "programmeringsspråk" som tar oss från de enkla komponenterna till det komplexa systemet. Dessa tankegångar var delvis vad som ledde till det som undersöks i denna avhandling. På grund av ämnets komplexitet är kopplingarna från de ursprungliga idéerna till de presenterade resultaten inte nödvändigtvis uppenbara. Ett viktigt steg i att försöka lösa ett stort problem är dock att identifiera och undersöka de mindre problem man stöter på på vägen.

Något som snabbt blir uppenbart när man försöker ta sig an den typ av modeller som behandlas i denna avhandling är graden av beroende mellan komponenterna. De flesta datoralgoritmer, även om de är komplexa, är vanligtvis "linjära" i det att varje steg bygger på resultatet av det föregående. Dessutom förblir datorns fysiska uppbyggnad konstant. Ett neuronalt nätverk fungerar inte så.

I naturen interagerar varje neuron kontinuerligt med sin omgivning, och omgivningen förändras över tid som ett svar på dessa interaktioner. En hjärna är *plastisk*, vilket innebär att förekomsten och styrkan av kopplingar mellan neuroner förändras över tid. Plasticiteten i sig behandlas inte här, och där det är relevant (specifikt i artikel II) uppdateras hela modellen samtidigt i varje steg. Trots detta är beteendet mycket svårt att förutspå. I artikel II härleder vi resultat som beskriver beteendet hos en så kallad *cellulär automat* i vissa speciella fall där detta beroende har mindre påverkan och diskuterar det beteende som observerats via simuleringar i de återstående fallen. Artiklarna III och IV behandlar en modell som är mindre relaterad till neuronala nätverk men där en liknande beroendestruktur framträder: där betraktas en linje av laddade partiklar där varje partikel bara känner laddningen från ett visst antal av sina närmaste grannar. Trots detta påverkar alla partiklar indirekt varandra. En rad metoder (som är intressanta i sig) används för att bevisa att detta beroende snabbt minskar när man betraktar partiklar långt ifrån varandra. Detta gör det möjligt att härleda en central gränsvärdessats, åtminstone för interaktioner mellan närmaste grannar.

Det vore övermaga att påstå att detta arbete tar oss mycket närmare en förklaring av hjärnans underverk, men jakten har fört med sig många intressanta resultat som förhoppningsvis kommer att vara användbara framöver, och den har väckt några fascinerande frågor som kan utforskas ytterligare.

Acknowledgements

I am incredibly lucky to be surrounded by so many wonderful people who have made this journey possible. I thank my supervisor for the amount of time and effort she spent in spurring me onwards, for always being available despite being busy, and for believing in me. I thank WASP for funding the research. I thank my friends whose daily doses of silliness always heightens the mood. I thank Var Gladspexet and all its members for the periods of high-intensity chaos that consistently increased my energy long-term. Most importantly, I thank you!

Research context

Chapter 1

Introduction

1.1 Mathematically Modelling the Brain

The general subject matter of this thesis is to mathematically study one of the most astonishing of nature's creations: the brain. We are inspired by many fundamental questions: How does the brain function? Is there an advantage to the specific way a brain is made up or has evolution just resulted in something that works well enough but is actually some kind of local minima? Is there a link between the physiological structure of the neural pathways of a brain and the tasks which it can perform and, if so, can we describe it? Many of these questions are as old as our knowledge about the objects themselves, so we make no claim to have found a definitive answer to them. The core question concerns the (presumed) link between form and function. Is there a clear link between the structure of a neuronal network and its capacity to perform certain tasks? A lot of research is being done to find the neuronal connection structure of various species, a brain's *connectome*, e.g. for rodents [16], or to construct probabilistic models which produce structures similar to experimental findings [12]. However, it is still an active area of research exactly what the knowledge of neuronal structure of a brain can tell us about its function.

We seek a deeper understanding of the processes at work and the complexity resulting from them. A variety of mathematical approaches are used and we take a closer look at several models, some of which harking back to the roots of artificial intelligence research from the 1950s through 1980s [11, 13]. We approach the above questions analytically and probabilistically on a fundamental level by studying simple models and observing the complexities that naturally arise from

them, with the hope of providing a deep understanding of the processes involved. This can pave the way for future investigations of even more advanced models consisting of combinations of our simpler units.

Another feature that is closely related is that of encoding information. There are many open questions on what the set of signals fired by a brain at a given moment, a certain *state*, ‘means’. In relation to coding theory, one can think of the initial state of a model as a coded message and the resulting output as the decoded message. These could be thought to correspond to some sensory input and a decision about what action to take given the input, respectively. Due to the complexity of the dynamics in the considered models, it is not so straightforward to predict this outcome, making it a possible cryptographical tool. On the other hand, even with knowledge on the model parameters, the problem of finding an initial state that leads to a given output is highly non-trivial.

The field of machine learning is under a lot of development and recent years has seen an increase of activity in the quest for *explainable AI* [23]. We can think of that as a ‘top down’ approach: it investigates machine learning models used for decision-making (in e.g. medical context or self-driving cars) with the aim of making concrete the influence various factors have on the outcome. This could be of great use in model design and parameter choices. In the future it might lead to programs that not only identify the sought after features but also explain why the features were identified and/or estimate the likelihood of the findings in a way that is understandable to humans. Compared to this, our approach is more ‘bottom up’: our aim is also to understand complex models, but we limit the complexity in order to expand our understanding on a more fundamental level.

1.2 Representation by graphs

Several ways of examining various dynamics on different networks have been explored in this work, all rooted in (very simplified) models of neuronal networks. At the most basic level, the brain is made out of neurons which are connected to each other via synapses and can output chemical and electrical signals based on what signals it receives. In a human brain there are roughly $\mathcal{O}(10^{11})$ neurons and roughly $\mathcal{O}(10^{15})$ synaptic connections [18]. Thus, as a simplification (for details on the intricacies involved, see [5, 19]), the brain can be modelled by a graph, denote it G , with vertices representing neurons and directed edges representing synaptic connections. On this graph we define and study various processes, and there are many options on how to define them.

Firstly, the graph structure of a real brain is *plastic*, changing over time, in response to external stimuli and previous activity. A common analogy is the treading of a path: sending the same signals many times lowers the effort required to send them. In neuroscience, this is known as the *Hebbian Rule* see e.g. [19], after [11]. In this work, however, the structure will always be fixed in time. It will be more rewarding to study the effects of this plasticity after one understands the dynamics and its effects on a given structure.

Then there is the matter of the *activity* on the graph at any given time. Physiologically, we know that a neuron receives electrical and chemical impulses from its *neighbours* (neurons with a synaptic connection to it). The neuron itself *fires* similar signals to all of its neighbours if the total sum of the received signals exceeds some threshold value, the so-called *action potential*. Many models, e.g. so-called integrate-and-fire models [10], seek to capture the complex behaviour that emerges from the fact that a neuron's chemical levels and electrical potential differences over membranes need refractory time to reset before this neuron can fire again. We do not include such complexities in any of our models, focusing instead on a larger scale of information flow over the whole graph.

One complication we do consider is the existence of *inhibitory neurons*, see e.g. [18]. These are neurons such that a fired signal reduces the potential difference of the recipient, thus working against the other incoming signals which increase it. At first glance, from an evolutionary viewpoint, this might seem like a waste of energy. There are some beneficial consequences on a chemical and electrical level [10], but we find that there are also benefits in the resulting complexity. In mathematical terms, including inhibitory neurons breaks the monotonicity of the dynamics, which greatly complicates the analysis. Particularly, in [7] (paper II in this thesis), we provide an informative example on how inhibitory vertices can be computationally beneficial. The model in [7] is simple to define yet results in complex dynamics. Without inhibitory vertices, one would need to define a much more involved model to achieve the same behaviour. This highlights as possible role played by the inhibitory neurons in a brain.

1.3 Multiple levels of scale

One interesting feature prevalent in nature but not often explicitly thought about is the seemingly separate, but definitely linked, levels of scale. On one level, we have atoms and molecules interacting to create and react to electrical and chemical processes. At a higher level, the behaviour of those atoms is interpreted as signals

traversing (and also altering) our neural network. Then, somehow, the behaviour of those signals e.g. help us decide what to have for lunch. The levels in between are quite mysterious, but it is easy to see how ‘zooming out’ from simple components opens up increasingly complicated phenomena.

A *Cellular automaton* is a model in which small components, cells, each have one of a set of *states* at any given time. When updated, a cell follows some simple dynamical process which depends on the current states of its neighbouring cells. The exact update function, number of states and which cells are considered to be neighbours can be defined in many ways, but the interesting question is how these kinds of cellular automata can behave over time. Perhaps the most known model is Conway’s game of Life, treated in detail in [1].

The models under consideration in this work all, to some extent, display different behaviours when considered at different scales. In [6] (paper I in this thesis) the flow of incoming particles on a graph with several outgoing directions and different speeds along the edges is considered. This gives rise to limiting cycles describing the (eventually cyclical) sequence of chosen directions. On a larger scale, this simplifies to the expected behaviour: the incoming particle flow gets separated into several particle flows which proportionally agree with the edge speeds.

The first level of complexity in [7] (paper II) is the activation of vertices on a cellular automaton. The second level of complexity is the activation patterns of the whole network over time. A third level can be the *interpretation* of that activation. We show that for the considered model the activation will become cyclical whatever the initial activation is. It is natural to classify the states by which cycle it will eventually reach. Interestingly, this defines some kind of metric on the state space: states which are pre-periodic to a cycle can be thought of as being infinitely far from all other cycles and the number of time-steps required to enter the cycle can be defined as its ‘distance’ to that cycle. Further, if one were to introduce noise to the model (alternatively, some external input affecting certain vertices), this metric would change: some of the hitherto impossible to reach cycles could become within reach. One way one could define the distance from a given state to a given cycle is the amount of noise required to make the target cycle a possible outcome. It is an open question whether it is possible to reach any cycle from any state, if one also requires the probability of that outcome to be above some threshold level.

In [8] (paper III in this thesis) a close-range variant of a typical model from statistical physics, the *Coulomb chain* is considered. The model, originally introduced by Malyshev [15], considers a density function consisting of a product of

Coulomb interactions between particles within a constant range. Note that even though the position of each particle is only present in factors containing its neighbours, all particle positions become dependent. In this setting, with the range reduced to each particle's nearest and next-to-nearest neighbours, a central limit theorem is derived which describes the large-scale behaviour of the inter-particle spacings. In [9], the correlation decay, which was required for the proof of the central limit theorem, is proved for the extended case of each particle interacting with its K nearest neighbours, for any constant K .

1.4 The Hopfield Inverse Problem

The original Hopfield model, introduced in [13] and further studied in e.g. [3, 14], is one where for $N \in \mathbb{N}$ one has a graph $G(V, E)$ with N two-state vertices $V \in \{-1, 1\}^N$ and edges with weights $\omega_{ij} \in [0, 1]$ ($i, j \in \{1, \dots, N\}$ and $\omega_{ii} = 0$). A set $\tilde{\xi}$ of $K \in \mathbb{N}$ given (or randomly chosen) *patterns* $\xi_k \in \{-1, 1\}^N$ are ‘learned’ by tuning the weights (often assumed to be symmetric). This is done so that if the model is given an input $\tilde{\xi} \in \{-1, 1\}^N$ as the initial state, $\sigma(0) = \tilde{\xi}$, the imposed dynamics given by

$$\sigma_i(t+1) = \operatorname{sgn} \left\{ \sum_{j=1}^N \xi_{ij} \sigma_j(t) \right\}, \quad (1.1)$$

terminates in one of the patterns ξ_k . The interpretation is that the final pattern is the one that is ‘closest’ to the initial one. The greatest number of patterns where the weights can be tuned to achieve this is called the model’s *storage capacity*. It is shown in [13] that the storage capacity is proportional to N , and later the storage capacity was rigorously bounded from below by $0.055N$ [17]. The set of states which lead to a learned pattern is called that pattern’s *basin of attraction*. It has been shown [2] that a surprisingly large number of weights can be randomly removed whilst the model still terminates at a pattern close (in a precise sense) to one of the learned ones. Specifically, for i.i.d. random patterns it is shown that for N neurons where each pair are connected at random with a (small) probability p , where for some constant c we have

$$p > \sqrt{\frac{c \ln(N)}{N}}, \quad (1.2)$$

there exists a constant α_c such that this *diluted* Hopfield model has a storage capacity of $\alpha_c p N$ patterns. For $c \approx 7$, $\alpha_c \approx 0.027$ and as c decreases, α_c decreases

‘roughly proportionally’ by $-1/\ln(c)$. It is also pointed out that $p = \ln(N)/N$ is the lowest probability where the diluted Hopfield could possibly work, as it is known from random graph theory that the network would be disconnected for lower values of p .

This inspires the following lines of investigations and conjectures. Say that instead of a random selection of weights are removed and the model is studied for random patterns, the *patterns* are (randomly) selected first and the *smallest* weights are removed after tuning them. The model can in this case possibly be pruned even further and still function as intended. This could provide results for the dilute Hopfield model for the region where $p \in \{\ln(N)/N, \sqrt{c \ln(N)/N}\}$.

In relation to our work, another question arises: is it possible to predict what pruned graph a given set of patterns will result in? Inversely, given a pruned graph, is there a relation to a set of patterns that the corresponding dilute Hopfield model is able to learn?

In research of the Hopfield model, one studies an *energy function* which has $\bar{\xi}$, the patterns to be learned, as minima. However, this energy function will typically have additional, so called ‘spurious’, minima in addition to the desired ones.

Chapter 2

An overview of the papers

2.1 Paper I, The Phases of a Discrete Flow of Particles on Graphs

Flows on graphs are typically studied in the continuous setting. This paper takes a close look at the intricacies that arise when considering a discrete flow on a graph with one vertex, one incoming edge and k outgoing edges. On these edges we assign different speeds. We let a steady stream of equidistant particles enter the system on the incoming edge. When they arrive at the ‘crossroad’ vertex, they choose the edge where the distance to the next particle is maximal and proceed to travel on that edge in its associated speed. Considering the system as a model of very thin electrical wires, the dynamics can be thought to result from the particles having a repellent charge to each other, and the various speeds can be thought of as the outgoing wires having different resistance.

For a given set of outgoing speeds, we consider the input of the system to be the speeds of the outgoing edges as well as the initial position of the nearest particle along them. It is shown that, after some time has elapsed, the imposed dynamics result in a repeating cycle of chosen edges. We define the output of the model to be the list of edges which each incoming particle chooses during the course of one such cycle.

We derive, for $k \in \{2, 3\}$, a complete description of what the possible limiting cycles are, given the configuration of speeds. This is done in a combinatorial way, treating a general configuration of speeds and going through the possible outcomes case by case, until a repetition is reached. Crucially, we show that in the case $k \geq 3$ the limiting cycles are not unique given the outgoing speeds. They are

in fact dependent on the initial placement on each outgoing edge of the particle nearest the vertex. In terms of encoding, the model (for $k \geq 3$) provides a highly non-trivial map from a set of k outgoing speeds and initial particle placements to a limit cycle.

Future works could consider an extension of the above model where the speed of the incoming stream of varies over time. This introduces even more complexity: due to the dependence on the initial configuration, this opens up the possibility of switching between limiting cycles by a temporary speed alteration. If the speed changes in a periodical way, the output would also be a limit cycle. This is a natural consideration, as each outgoing edge in the original model will (after some time has elapsed) have a repeating pattern of distances between the particles traversing it. Taking one (or several) outgoing edge(s) from the above model and treating it as an incoming stream for a new vertex which itself has several outgoing edges. This second vertex is locally equivalent to the extension discussed above, where the time between each incoming particle (i.e. distance between them while traversing it) is determined by the first model. This opens up the possibility of considering discrete particle flows on *trees*. Even for one additional ‘layer’ of such a tree, it is an open question for this extended model to what extent the initial particle placements affect the outcome.

2.2 Paper II, Non-Monotone Cellular Automata: Order Prevails Over Chaos

Paper II in this thesis studies the effect that adding inhibitory influences of certain vertices in dynamical model has on the dynamics. We consider a graph $G(V, E)$ on which, for each time t , a subset $A_t \subseteq V$ of the vertices are considered to be *active*. Typically, one initialises A_0 , the activity at time $t = 0$ and studies the behaviour of the model under a given rule for how this activity spreads.

In what is known as *bootstrap percolation*, each vertex becomes active (or ‘infected’) at time $t + 1$ if the number of neighbouring vertices that are active at time t exceeds some threshold value r . The activation is strictly increasing and each activation set A_0 will lead to a final state A_∞ which is stable under updates. A question one can ask is this: if each vertex is in A_0 (initially active) with some probability p , what is the probability of there being an infinite connected component of neighbouring active vertices in A_∞ ? With the bootstrap threshold r chosen correctly (so as not to invoke trivial dynamics where all or no vertices are in A_∞), there can be a phase transition at some critical p_c where the existence of

a finite connected component becomes likely.

Inspired by the existence of inhibitory neurons discussed in the introduction, we extend the bootstrap percolation model in the following way. Each vertex is assigned a type, either excitatory or inhibitory. For each time step, synchronously, each vertex becomes active if the difference of the number of active excitatory and inhibitory neighbours exceeds the threshold value r , as opposed to bootstrap percolation where only the sum of active neighbours is considered. It quickly becomes apparent that this change has many complicated consequences. Most results in the field of percolation theory become inapplicable, since the monotonicity of the evolution of activity is a key assumption. Instead, the model becomes closer to a *cellular automaton*. These are models where identical units update their state according to a function of their neighbours. These simple units form a complex system. For our model, one can for each time step t let vertex i have the state $\sigma_v(t) = (a_v(t), b_v)$ where $a \in \{0, 1\}$ describes its (time-dependent) activity and $b \in \{-1, 1\}$ describes its type. Let $n(v)$ denote the neighbours of the vertex v . The dynamics is then given by

$$a_v(t+1) = \begin{cases} 1, & \text{if } \sum_{w \in n(v)} a_w(t) \cdot b_w \geq r, \\ 0, & \text{otherwise.} \end{cases} \quad (2.1)$$

We study in depth the graph of a two-dimensional lattice where each vertex is connected to its four nearest neighbours. A vertex is defined to be inhibitory if both its indices are even and excitatory otherwise. Since this is a periodic graph, an equivalent model can be formulated as a classical one-state cellular automaton by mapping clusters of four vertices to one cell, but doing so makes the update rule very involved and unintuitive. This shows the complexity benefits of introducing inhibitory vertices.

In seeking to describe the behaviour of the studied cellular automaton, additional models are proposed which can describe the amount of activation for a certain regime of initial activation. We show that with one single initially active excitatory vertex, the activation will spread as strips of activation in either the horizontal, vertical or all four cardinal directions. Further, we list all possible outcomes of having two initially active vertices, by far the most common of which is a strip of activity from one vertex growing into a strip originating from the other and stopping its growth in that direction. This allows for a simplified model of sparse initial activation of the system. The (four time-steps cyclical) strips of activation are mapped to a continuous model of lines similarly growing along the cardinal directions from random points until they reach another line. The total length of

lines resulting from such growth is proportional to the total activation in the original model. This model was recently studied in [4] where it was shown that the correlation of the length of the lines emanating from two particles exponentially decays with the distance between them.

The lines from the latter model partition the space into rectangular areas which seem, from simulations, to result in regions with long and narrow horizontal/vertical strips. Can we define a measure that gives the correct probabilities for these? If one labels each point by the ratio of height and width, it would be interesting to study the correlation of two points having the same direction. Additionally, what would be the distribution of width-to-height ratio? Returning to the original model, such a description could be useful in describing the final activation patterns.

Another open question is whether there exists an energy function that has the cyclical states as minima, as is the case in the Hopfield model [13]. The two models might seem very disparate, but they share the feature of ‘retrieving patterns’, and basins of attraction for certain (in our case cyclical) final states.

2.3 Paper III, A Central Limit Theorem and Decay of Correlations in the Coulomb Chain

Paper III in this thesis considers a model of the Coulomb chain introduced by Malyshev in [15]. The model consists of $N \in \mathbb{N}$ identical particles that are confined to one dimension. Each particle experiences three-dimensional Coulomb interactions with its nearest and next-to-nearest neighbours. This work proves exponential decay of correlations between sets of consecutive inter-particle distances and uses this result to prove a central limit theorem.

First, we define the model. Let $\bar{P} \in \mathbb{R}^N$ denote the consecutive positions of the particles, so that we have

$$0 \leq P_1 \leq \cdots \leq P_{N-1} \leq P_N, \quad (2.2)$$

and we let the model be circular by treating the first and last particles as neighbours. Let $\beta, \gamma \in \mathbb{R}^+$ be parameters governing the strength of nearest neighbour and next-to-nearest neighbour interactions, respectively. The Gibbs energy function

of the considered model is given by

$$\begin{aligned} \tilde{H}(\bar{P}) = & \frac{\beta}{P_1} + \sum_{i=2}^N \left(\frac{\beta}{P_i - P_{i-1}} \right) \\ & + \frac{\gamma}{P_1 + (P_N - P_{N-1})} + \frac{\gamma}{P_2} + \sum_{i=3}^N \left(\frac{\gamma}{P_i - P_{i-2}} \right), \end{aligned} \quad (2.3)$$

where the terms outside the sums account for the circular nature of the model. We let $\bar{Y} \in [0, 1]^N$ denote the distances between each pair of consecutive particles by defining

$$\begin{aligned} Y_1 &:= P_1, \\ Y_i &:= P_i - P_{i-1}, \quad i \in \{2, \dots, N\}. \end{aligned} \quad (2.4)$$

This we assume to be bounded, and without loss of generality bounded by one. To simplify the notation, all indices are henceforth to be taken modulo N . We may now rewrite the energy of the system, Equation (2.3), into

$$H(\bar{Y}) = \sum_{i=1}^N \left(\frac{\beta}{Y_i} + \frac{\gamma}{Y_i + Y_{i+1}} \right), \quad (2.5)$$

and define a probability distribution for \bar{Y} by

$$\begin{aligned} f_{\bar{Y}}(\bar{y}) &= \frac{1}{Z_N} e^{-H(\bar{y})}, \\ Z_N &= \int_{[0,1]^N} e^{-H(\bar{y})} d\bar{y}. \end{aligned} \quad (2.6)$$

The first theorem in this work proves exponential decay of the correlation between two sets of consecutive variables. We use the notation that for a vector $\bar{Y} \in [0, 1]^N$ and an index set $I \subset \{1, \dots, N\}$, we denote by \bar{Y}_I the vector containing the elements of \bar{Y} whose index are in I .

Theorem 2.3.1 *Let $\beta > 0$ and $\gamma \geq 0$ be arbitrary parameters and let $\bar{Y} \in [0, 1]^N$ be a random vector with the density function $f_{\bar{Y}}(\bar{y})$ given by Equations (2.5) and (2.6). There exist positive constants $C = C(\beta, \gamma)$ and $\alpha = \alpha(\gamma)$ such that for any disjoint sets of (in the circular sense) consecutive indices $I, J \subset \{1, \dots, N\}$ with $r(I, J)$ denoting the distance between them, we have*

$$\left| f_{\bar{Y}_I | \bar{Y}_J = \bar{y}_J}(\bar{y}_I) - f_{\bar{Y}_I}(\bar{y}_I) \right| \leq C e^{-\alpha r(I, J)} f_{\bar{Y}_I}(\bar{y}_I), \quad (2.7)$$

for any $\bar{y}_I \in [0, 1]^{|I|}$ and $\bar{y}_J \in [0, 1]^{|J|}$. Furthermore, we have that

$$\alpha(\gamma) \xrightarrow{\gamma \rightarrow 0} \infty. \quad (2.8)$$

Theorem 2.3.1 is an extension of Theorem 2.1 in [21], which proved the case when $|I| = |J| = 1$. Note that only the distance between the sets I and J is present in Equation (2.7) and not their sizes.

Although the energy function in (2.5) consists of terms with at most two variables, all variables in \bar{Y} are indeed dependent. The proof of Theorem 2.3.1 involves rewriting the density functions on the left hand side of Equation (2.7) as ‘chunks’ of integrals over the sets I, J and the two remaining sets of consecutive variables in $\{1, \dots, N\} \setminus (I \cup J)$. One then gets an expression which is a difference of functions inside a large amount of integrals, divided by another large amount of integrals. The difference itself can be bounded and then a number of approximations are made which split the integrals into their respective ‘chunks’ of consecutive variables, breaking the overlapping dependencies. It is shown that all (possibly small) factors in the denominator can be cancelled out with corresponding factors in the numerator.

The result of Theorem 2.3.1 was a necessity in proving the central limit theorem for the considered model. For simplicity, the random variables are centralised.

Theorem 2.3.2 *Let $\bar{X} \in [-1, 1]^N$ be a zero-mean random vector with density $f_{\bar{X}}(\bar{x})$ satisfying the condition that there exists constants $C, \alpha \in \mathbb{R}^+$ such that for any disjoint sets of consecutive indices $I, J \subset \{1, \dots, N\}$, Equation (2.7) is fulfilled. Furthermore, let $f_{\bar{X}}(\bar{x})$ satisfy the condition that there for any $\delta \in [0, 1]$ exists a constant $c > 0$ such that for large N we have*

$$\text{Var}\left(\sum_{i=1}^{\lfloor N^\delta \rfloor} X_i\right) \geq c \lfloor N^\delta \rfloor. \quad (2.9)$$

Denote

$$\sigma_N^2 := \frac{1}{N} \text{Var}\left(\sum_{i=1}^N X_i\right), \quad (2.10)$$

and set

$$\zeta_N := \frac{1}{\sqrt{N\sigma_N^2}} \sum_{i=1}^N X_i. \quad (2.11)$$

Then for any $\varepsilon \in (0, 1/4)$, we have

$$\sup_z |\mathbb{P}(\zeta_N \leq z) - \mathbb{P}(Z \leq z)| = \mathcal{O}(N^{-\frac{1}{4}+\varepsilon}), \quad (2.12)$$

where Z is a standard normal random variable.

The proof of Theorem 2.3.1 was inspired by the work of Schmuland and Sun in [20], but several technical adjustments were needed to account for the difference in setting compared to that work, which considers a random process on an infinite d -dimensional lattice with a different correlation decay condition than the one in our case.

The proof of Theorem 2.3.2 centres around splitting the (circular) string of random variables into subsets of p and q consecutive variables, where $q \ll p$ and both are (the integer parts of) exponents of N . There are a total of $k = \lfloor N/(p+q) \rfloor$ subsets of each type. The sum ζ_N is then approximated in a number of steps. Firstly, the contribution of terms from the smaller sets are neglected. Secondly, the sum of terms from each of the larger sets are approximated as independent of each other. Thirdly, the difference between the sum of the (now independent) random variables and the standard normal is bound by the Berry-Esseen theorem.

The error bounds of all the above approximations contain the quantities p and q (as well as k , itself a function of p and q), and one can tune the exponents defining p and q to achieve the desired final rate of $\mathcal{O}(N^{-1/4+\varepsilon})$.

One noteworthy change from the corresponding proof in [20] is the definition of p and q , which in that work is defined immediately instead of being fine-tuned towards the end. The result in [20] is not directly comparable to Theorem 2.3.2, but it is encouraging that the changed exponents, together with the altered setting, allowed us to improve the rate from $\mathcal{O}(N^{-1/9})$ to arbitrarily close to $\mathcal{O}(N^{-1/4})$.

The last result in Paper III is confirming that the condition in Equation (2.9) holds for the Coulomb chain model. This was shown using a result from [22] and was the final step needed to confirm the conjecture in [21] of a central limit theorem for the inter-particle distances in the nearest and next-to-nearest Coulomb chain model.

2.4 Paper IV, Correlation Decay in the Coulomb Chain with Local Pairwise Interactions

In Paper IV, the result on correlation decay in Paper III, [8], is extended to apply for the circular Coulomb chain model with pairwise interactions between each particle

and its K nearest neighbours in each direction, where $K \in \mathbb{N}$ is an arbitrarily fixed constant. This alteration brought additional technical difficulties compared to the previous work and it was even a challenge to find a notation that made the work digestible. The model under consideration is the same as in Paper III but now contains K parameters β_1, \dots, β_K . Letting \bar{Y} still denote the distance between consecutive particles, the Gibbs energy function is now of the form

$$H_K(\bar{Y}) = - \sum_{k=1}^K \sum_{i=1}^N \frac{\beta_k}{Y_i + \dots + Y_{i+k-1}}. \quad (2.13)$$

The main difference compared to Paper III is that instead of needing to split integrals by altering a factor containing two variables, one needs to alter a much more involved factor containing $K - 1$ variables. An additional notational difficulty arises from this fact: the factors involved, let us denote them $\mathcal{T}(\bar{x}, \bar{y})$ for $\bar{x}, \bar{y} \in [0, 1]^{K-1}$, are not symmetric in the sense that $\mathcal{T}(\bar{x}, \bar{y}) \neq \mathcal{T}(\bar{y}, \bar{x})$, but rather in the sense that

$$\begin{aligned} \mathcal{T}((x_1, x_2, \dots, x_{K-1}), (y_1, y_2, \dots, y_{K-1})) \\ = \mathcal{T}((y_{K-1}, \dots, y_2, y_1), (x_{K-1}, \dots, x_2, x_1)). \end{aligned} \quad (2.14)$$

Nevertheless, we show that Theorem 2.3.1 indeed also holds for the Coulomb chain model for arbitrary constant $K > 2$. A promising line of inquiry for future research is to prove that this model also satisfies the central limit theorem 2.3.2 for arbitrary K . The result needed for this is that there exists a constant c such the lower bound of the variance,

$$\text{Var} \left(\sum_{i=1}^{\lfloor N^\delta \rfloor} Y_i - \mathbb{E}(Y_i) \right) \geq c \lfloor N^\delta \rfloor, \quad (2.15)$$

holds for any $\delta \in [0, 1]$. In fact, for Theorem 2.3.2 to hold for $\varepsilon \in (0, 1/4)$, it is sufficient to prove that there exists such a constant so that Equation (2.15) holds for $\delta = 1 - 2\varepsilon$.

Bibliography

- [1] Adamatzky, A., *Game of Life Cellular Automata*, Springer, London, 2010.
- [2] Bovier, A., Gayrard, V., *Rigorous Bounds on the Storage Capacity of the Dilute Hopfield Model*, *Journal of Statistical Physics* **69**, 597–627 (1992).
DOI 10.1007/BF01050427.
- [3] Bovier, G., Gayrard, V., *Hopfield Models as Generalized Random Mean Field Models*, in: Bovier, A., Picco, P. (eds), *Mathematical Aspects of Spin Glasses and Neural Networks*, *Progress in Probability*, **41**, Birkhäuser Boston, 1998.
DOI 10.1007/978-1-4612-4102-7_1.
- [4] Ewers, E. (2024) *Process of a Random Planar Quadrangulation*, Master's Theses in Mathematical Sciences, Lund University (2024).
- [5] Fornito, A., Zalesky, A., Breakspear, M., *Graph Analysis of the Human Connectome: Promise, progress, and pitfalls*, *NeuroImage* **80**, (2013).
10.1016/j.neuroimage.2013.04.087.
- [6] Ekström, H., *The Phases of a Discrete Flow of Particles on Graphs*, *Markov Processes and Related Fields*, **26**(2), 343–364 (2020).
<https://dlib.eastview.com/browse/doc/58666123>.
- [7] Ekström, H., Turova, T., *Non-Monotone Cellular Automata: Order Prevails over Chaos*, *Biosystems*, **120** (2022). 10.1016/j.biosystems.2022.104756.
- [8] Ekström, H., *A Central Limit Theorem and Exponential Correlation Decay for the Coulomb Chain*, 10.48550/arXiv.2404.14884, 2024.
- [9] Ekström, H., *Correlation Decay in the Coulomb Chain with Local Pairwise Interactions*, 2024.

- [10] Gerstner W., Kistler W.M., Naud R., and Paninski, L. *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition*, Cambridge University Press, 2014.
- [11] Hebb, D.O., *The Organization of Behavior*, Wiley, United Kingdom, 1949.
- [12] Hill, S.L., Wang, Y., Riachi, I., Schürmann, F., Markram, H., *Statistical Connectivity Provides a Sufficient Foundation for Specific Functional Connectivity in Neocortical Neural Microcircuits*, DOI 10.1073/pnas.1202128109.
- [13] Hopfield, J.J., *Neural Networks and Physical Systems with Emergent Collective Computational Abilities*, Proceedings of the National Academy of Sciences of the United States of America, **79**(8), 2554-2558 (1982). DOI 10.1073/pnas.79.8.2554.
- [14] Löwe, M., Vermet, F., *The Storage Capacity of the Hopfield Model and Moderate Deviations*, Statistics & Probability Letters, **75**(4), 237-248 (2005). DOI 10.1016/j.spl.2005.06.001.
- [15] Malyshev, V.A., *Phase Transitions in the One-Dimensional Coulomb Medium*, Problems of Information Transmission, **51**, 31-36 (2014). DOI 10.1134/S0032946015010044.
- [16] Mäki-Marttunen, T, *Modelling Structure and Dynamics of Complex Systems: Applications to Neuronal Networks*, PhD Thesis, Tampere University of Technology, **1183**, (2013). <http://urn.fi/URN:ISBN:978-952-15-3203-0>.
- [17] Newman, C.M., *Memory Capacity in Neural Network Models: Rigorous Lower Bounds*, Neural Networks, **1**(3), 223-238 (1988). DOI 10.1016/0893-6080(88)90028-7.
- [18] Rolls, E.T., *Cerebral Cortex: Principles of Operation*, Oxford University Press, Oxford, 2016. DOI 10.1093/acprof:oso/9780198784852.001.0001.
- [19] Rolls, E.T., Treves, A., *Neural Networks and Brain Function*, Oxford University Press, 1998. DOI 10.1093/acprof:oso/9780198524328.001.0001.
- [20] Schmuland, B., Sun, W., *A Central Limit Theorem and Law of the Iterated Logarithm for a Random Field with Exponential Decay of Correlations*, Canadian Journal of Mathematics, **56**(1), 209-224 (2004).

BIBLIOGRAPHY

DOI 10.4153/CJM-2004-010-6.

- [21] Turova, T.S., *Exponential Decay of Correlations in the One-Dimensional Coulomb Gas Ensembles*, Journal of Mathematical Physics, **63**(5), (2022). DOI 10.1063/5.0089803.
- [22] Turova, T.S., *Antiferromagnetic Covariance Structure of Coulomb Chain*, 10.48550/arXiv.2304.06500 (updated 12 Feb, 2024).
- [23] Yang, W., Wei, Y., Wei, H. et al., *Survey on Explainable AI: From Approaches, Limitations and Applications Aspects*, Human-Centric Intelligent Systems, **3**, 161–188 (2023). DOI 10.1007/s44230-023-00038-y.



LUND
UNIVERSITY

Doctoral Theses in Mathematical Sciences 2024:3

ISBN 978-91-8104-123-1

ISSN 1404-0034

