Decision Making and Learning in Artificial Physical Systems – Long-Range Orientational Order in a Gibbsian Hard-Core Particle System, Value Learning of Spiking Neurons, and Exit Time Asymptotics of a Switching Process

by

Alexisz Tamás Gaál

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Mathematics

New York University

January, 2020

Charles M Newman

Daniel Stein

Abstract

In this dissertation, we present results in three distinct models of artificial physical systems concerning decision making and learning in an abstract sense. The first chapter discusses crystallization of physical fluids that is not rigorously understood mathematically. We study crystallization in a simplified model with hard-core interaction, and the absence of lattice defects. Point processes of hard spheres that are enumerated by and locally close to some three-dimensional rigid lattice are considered, and shown to exhibit long-range orientational order. This result extends an earlier work to the three-dimensional case which is – to our knowledge – the first rigorous result on long-range orientation order of hard spheres in three dimensions. Long-range orientational order is considered as a signature for crystallization. Relative orientation of neighboring atoms is learned from imposed, tight boundary conditions. In two dimensions, we define a Gibbsian point process with respect to a Poisson point process by a local, geometry dependent Hamiltonian on hard disks that informally speaking imposes a constraint that each point has exactly six neighbors in an annulus around them. Hence, in contrast to the earlier work, we do not require anymore that the process is globally enumerable by a triangular lattice. Existence of long-range orientational order carries over, and we obtain the existence of an infinite-volume Gibbs measure on two-dimensional point configurations that follow the orientation of a fixed triangular lattice arbitrary closely everywhere.

In the second chapter, we present a model for neural learning that is joint work with Johanni Brea, Robert Urbanczik, and Walter Senn. How predictions of future events of a

dynamically changing environment are learned on the neural level and what plasticity rule underlies their learning is not well understood. Dark clouds are a sign of upcoming rain, taking out the leash from the cabinet is a sign for our dog that the afternoon walk is about to begin. Animals can make such predictions. We demonstrate that a neuron which connects to a given number of presynaptic spiking neurons, which encode the environment, can learn such predictions following a biologically plausible learning rule. The learning rule allows the neuron to match its current firing rate to the expected future discounted reward determined by the environment. In a two-compartment neuron model, even though the plasticity window is on the scale of tens of milliseconds, predictions on the timescale of seconds can be made.

The third chapter is joint work with Yuri Bakhtin. We study stochastic processes with random switching, known as piecewise deterministic Markov processes. In general, these processes are defined by a family of vector fields and a collection of rates of switching between those vector fields. At each time the system is in some state where it evolves along one of the vector fields from the family. At random times, the system jumps between states, switching active vector fields from one to another according to the prescribed Markovian rates. In the limit of infinite switching rates, the evolution can be effectively described by the law of large numbers through the averaging of the vector fields involved. We show that a class of one-dimensional switching processes started at the origin with an unstable zero of the effective field at the origin escapes a given interval around the origin in a time that increases logarithmically in the switching rate. The location of escape (left or right edge) can be interpreted as a decision between two alternatives, and the time of exit as the decision time.

Contents

List of Figures

Chapter 1

Long-Range Orientational Order in Near-Lattice Gibbsian Hard-Core Particle Systems

1.1 Introduction

Random hard disk and hard sphere processes are one of the most easily defined, physically interesting point processes. Rigorous mathematical results about their behavior at high intensity are limited to two-dimensional systems. It remains an open question whether a phase transition with possibly orientational symmetry breaking occurs in either two or three dimensions. If breaking of rotational symmetry in either of these models could be shown, it would give rise to speculation whether such simple pair interaction could result in crystallization phenomena. In order to simplify the models, we exclude cavities and other crystal defects from the models and study random hard disk and sphere processes that are locally crystals. In our models, being locally crystal implies being a crystal on a long range. There is a lower bound on orientational correlation that is uniform in the distance and this bound can be made arbitrarily large by taking very "tight" boundary conditions. Theorem 1.2.1 is the first rigorous result about hard sphere long-range orientational order to our knowledge. The content of this chapter was first published in April 2018 on `https://math.nyu.edu/~gaal/`, then in [27].

In the previous works [29] and [28], we considered hard disk processes with disks of radius 1/2 that have the structure of a triangular lattice and neighboring disks have an upper bound on their distance. We showed the existence of natural "uniform" measures on these allowed configurations that exhibit uniform long-range orientational order. In the first half of this chapter, we show that the same arguments apply to some three-dimensional lattices. In the second half, we show that the result in the two-dimensional case can be formulated independently of an underlying triangular lattice structure that was explicitly present in the definition of the probability measures in [29]. Thus we show that being a crystal locally implies being a crystal on the long range in this particular model. We only require the local, geometry dependent condition that every point has exactly six points in an annulus with radii 1 and $1+\alpha$ around them. We will have the parameter α in both sections that gives the maximal distance of neighboring points. This α needs to be sufficiently small so that some local conditions are fulfilled, however it is on the macroscopic order of about 1/2, so not particularly small. Fluctuations from the orientation of a fixed lattice however can be made arbitrary small, in particular they can be made many orders smaller than α.

Similar but not hard-core models were considered in [42] without defects and in [34] and [4] with lattice defects. Models for dislocations were treated in [15] on the mesoscopic scale and in [32] for the Ariza-Ortiz model. Introducing bounded, separated missing regions as defects into our two-dimensional model is possible using similar techniques as in [34]. For three dimensions, we think it is possible but we have not carried it out. Also the techniques of Section 1.3 can possibly be carried out in three dimensions, but an analogue of Lemma 1.3.5 is required together with considering boundary conditions, since in three dimensions several close-packed lattices are possible analogues of the triangular lattice.

These simplified models with well-defined lattice structure and possible defects are motivated by more natural hard sphere models defined with respect to a Poisson point process at a given intensity $z > 0$. The set of Gibbs measures for these natural models is defined similarly to our definition of $\mathcal{G}^z$ in Section 1.3. They are basically sequential limits of Poisson point processes in bounded domains – as the domains tend to $\mathbb{R}^d$ – conditioned that no pair of points have distance smaller than one. In these natural models, instead of imposing complex geometry dependent interactions, merely hard-core repulsion is required. As a consequence, even at high intensity, all kinds of possible lattice defects emerge as soon as the domain gets large enough. It is believed that in dimensions two and greater there are multiple Gibbs measures in $\mathcal{G}^z$ for high enough intensity z. Their structure is believed to differ in the typical relative orientation of nearby points. It is shown in [43] that in dimension two any of these measures in $\mathcal{G}^z$ are translational invariant at any intensity $z > 0$, and in [44] a logarithmic lower bound is given on the mean square translational displacement of particles. These results prevent Gibbs measures from having long-range positional order.

One strategy of showing that $\mathcal{G}^z$ is not a singleton for $d \geq 2$ and $z > 0$ high enough, is to search for a measure in $\mathcal{G}^z$ that is not rotational invariant. Existence of such is called the breaking of rotational symmetry (of the energy function). Showing that such a measure is supported on a perturbed lattice structure with long-range orientational order would be an even stronger result which is connected to the widely studied crystallization problem, even though the crystallization problem is mostly studied for different interactions.

We would also like to mention the recent result [35] that at low intensity disagreement percolation results imply the uniqueness of the Gibbs state. While at high intensity it is shown in [3] that hard disks percolate with the percolation radius chosen sufficiently big which was generalized in [40] to arbitrary percolation radii. Percolation is necessary for crystallization, but to our knowledge breaking of rotational symmetry cannot be concluded from it.

1.2 The three-dimensional enumerated model

In this section we show that the arguments of [29] can be applied to some three-dimensional lattices to obtain similar results as in [29] about long-range orientational order for random perturbations of such lattices.

1.2.1 Configuration space

We consider three-dimensional lattices with well-defined distance between nearest neighbors (to be normalized to 1) that fulfill two conditions. Firstly, the nearest neighbor edges of the lattice have to define a tessellation of $\mathbb{R}^3$ by regular tetrahedra and octahedra. Secondly, the lattice has to be translational invariant in three linearly independent directions. We remark that regular tetrahedra and octahedra can be replaced by any rigid polyhedron (a polyhedron with all faces being triangles) that satisfies an analogue of the rigidity estimates in Lemmas 1.2.4 and 1.2.5, and their volume has positive partial derivatives with respect to their edge lengths. We note that by Cauchy's theorem, the volume is uniquely defined for rigid polyhedra when the edge lengths are given.

Examples of such lattices are the face-centered cubic lattice and the hexagonal close-packed lattice. For definitions see [39]. Note that being translational invariant does not mean that the lattice has to be a Bravais lattice, i.e. of the form $\mathbb{Z}n_1 + \mathbb{Z}n_2 + \mathbb{Z}n_3$ for some vectors $n_i \in \mathbb{R}^3$. Bravais lattices are translational invariant but a union of Bravais lattices might be still translational invariant, however not a Bravais lattice anymore for which the hexagonal close-packed lattice serves as an example.

Let the set $I \subset \mathbb{R}^3$ denote one of the lattices that fulfill both criteria. We assume $0 \in I$ and think of I as an index set which is going to be used to parametrize countable point configurations in $\mathbb{R}^3$. Let I have translational symmetry by the linearly independent vectors $t_1, t_2, t_3 \in \mathbb{R}^3$ and define the set $T = \mathbb{Z}t_1 + \mathbb{Z}t_2 + \mathbb{Z}t_3$. Define the quotient space $I_n := I/nT$ for $n \in \mathbb{N}$. We will think of I_n as a specific set of representatives in the half-open parallelepiped U_n spanned by nt_1, nt_2, nt_3, i.e. $U_n = n\{xt_1 + yt_2 + zt_3 \mid x, y, z \in [0, 1)\}$.

A *parametrized point configuration* in $\mathbb{R}^3$ is a map $\omega : I \to \mathbb{R}^2$, $x \mapsto \omega(x)$ that determines the point configuration $\{\omega(x) \mid x \in I\} \subset \mathbb{R}^3$. For the set of all parametrized point configurations we introduce the character $\Omega = \{\omega : I \to \mathbb{R}^2\}$. Note that a single point configuration

$\{\omega(x) \mid x \in I\}$ can be parametrized by many different $\omega \in \Omega$.

Let $\alpha \in (0, 1]$ be an arbitrary but fixed real to be fixed later. An *n-periodic parametrized point configuration* with edge length $l \in (1, 1+\alpha)$ is a parametrized configuration ω which satisfies the boundary conditions:

$$\omega(x + nt_i) = \omega(x) + lnt_i \quad \text{for all } x \in I \text{ and } i \in \{1, 2, 3\}. \tag{1.2.1}$$

The set of n-periodic parametrized configurations with edge length l is denoted by $\Omega^{per}_{n,l} \subset \Omega$. From now on we will omit the word parametrized because, in this section, we are going to work solely with *point configurations* which are parametrized by I. An n-periodic configuration is uniquely determined by its values on I_n. Therefore, we identify n-periodic configurations $\omega \in \Omega^{per}_{n,l}$ with functions $\omega : I_n \to \mathbb{R}^2$.

The bond set $E \subset I \times I$ contains index-pairs with Euclidean distance one; this is $E = \{(x, y) \in I \times I \mid |x - y| = 1\}$. We set $E_n = E/nT$, we can think of E_n as a bond set $E_n \subset I_n \times I_n$. Let $\mathcal{T}$ denote the set of convex polyhedra, as in the definition of I, whose edges are in E and provide a tessellation of $\mathbb{R}^3$, which is the Delaunay pre-triangulation, see [39]. Define $\mathcal{T}_n = \mathcal{T}/nT$. Each $\triangle \in \mathcal{T}$ can be triangulated into tetrahedra (not necessarily uniquely), let us fix such a T-periodic triangulation of $\mathcal{T}$. The set of all (necessarily not all regular) tetrahedra created this way define a tessellation of $\mathbb{R}^3$ and is denoted by $\text{triang}(\mathcal{T})$. We define $\text{triang}(\mathcal{T}_n) := \text{triang}(\mathcal{T})/nT$.

1.2.2 Probability space

By definitions of Ω and $\Omega_{n,l}^{per}$, we have $\Omega = (\mathbb{R}^3)^I$ and can identify $\Omega_{n,l}^{per} = (\mathbb{R}^3)^{I_n}$. Both sets are endowed with the corresponding product σ-algebras $\mathcal{F} = \bigotimes_{x\in I} \mathcal{B}(\mathbb{R}^3)$ and $\mathcal{F}_n = \bigotimes_{x\in I_n} \mathcal{B}(\mathbb{R}^3)$ where $\mathcal{B}(\mathbb{R}^3)$ denotes the Borel σ-algebra on each factor. The event of admissible n-periodic configurations $\Omega_{n,l} \subset \Omega_{n,l}^{per}$ is defined by the properties $(\Omega1) - (\Omega4)$:

$(\Omega1)$ $|\omega(x) - \omega(y)| \in (1, 1+\alpha)$ for all $(x, y) \in E$. For $\omega \in \Omega$ we define the extension $\hat{\omega} : \mathbb{R}^3 \to \mathbb{R}^3$ such that $\hat{\omega}(x) = \omega(x)$ if $x \in I$. On the closure of a tetrahedron $\triangle \in \text{triang}(\mathcal{T})$, the map $\hat{\omega}$ is defined to be the unique affine linear extension of the mapping defined on the corners of that tetrahedron.

$(\Omega2)$ The map $\hat{\omega} : \mathbb{R}^3 \to \mathbb{R}^3$ is injective (and thus bijective).

$(\Omega3)$ The map $\hat{\omega}$ is almost everywhere orientation preserving, this is to say that $\det(\nabla\hat{\omega}(x)) > 0$ for almost every $x \in \mathbb{R}^3$ with the Jacobian $\nabla\hat{\omega} : \mathbb{R}^3 \to \mathbb{R}^{3\times3}$.

$(\Omega4)$ The image $\hat{\omega}(\triangle)$ of a polyhedron $\triangle \in \mathcal{T}$ is a convex polyhedron.

The conditions $(\Omega3)$ and $(\Omega4)$ follow from conditions $(\Omega1)$ and $(\Omega2)$ up to the sign of the determinant in $(\Omega3)$ as it was also remarked in [34] on page 4. Since the proof is more analytic than stochastic, we also omit the proof and require them as technical conditions. Define the set of *admissible n-periodic configurations, with edge length l* as

$$\Omega_{n,l} = \{\omega \in \Omega_{n,l}^{per} \mid \omega \text{ satisfies } (\Omega1) - (\Omega4)\}.$$

The set $\Omega_{n,l}$ is open and nonempty subsets of $(\mathbb{R}^3)^{I_n}$. The scaled lattice $\omega_l(x) = lx$ for $x \in I$ and $1 < l < 1+\alpha$ is an element of $\Omega_{n,l}$. Any configuration $\omega \in \Omega_{n,l}$ is determined by a finite number of locations in $\mathbb{R}^3$. Each property $(\Omega 1) - (\Omega 4)$ is satisfied after small enough perturbation of these locations, therefore any $\omega \in \Omega_{n,l}$ has a neighborhood that is fully contained in $\Omega_{n,l}$, hence the openness of $\Omega_{n,l}$.

Clearly, $0 < \delta_0 \otimes \lambda^{I_n \setminus \{0\}}(\Omega_{n,l}) < \infty$ with the Lebesgue measure λ on $\mathbb{R}^3$ and the Dirac measure δ_0 in $0 \in \mathbb{R}^3$. The lower bound holds because $\Omega^0_{n,l}$ is nonempty and open in $(\mathbb{R}^3)^{I_n \setminus \{0\}}$ (similarly to the case of $\Omega_{n,l}$ above); the upper bound is a consequence of the parameter α in $(\Omega 1)$. Let the probability measure $P_{n,l}$ be

$$P_{n,l}(A) = \frac{\delta_0 \otimes \lambda^{I_n \setminus \{0\}}(\Omega_{n,l} \cap A)}{\delta_0 \otimes \lambda^{I_n \setminus \{0\}}(\Omega_{n,l})}$$

for any Borel measurable set $A \in \mathcal{F}_n$, thus $P_{n,l}$ is the uniform distribution on the set $\Omega_{n,l}$ with respect to the *reference measure* $\delta_0 \otimes \lambda^{I_n \setminus \{0\}}$. The first factor in this product refers to the component $\omega(0)$ of $\omega \in \Omega$.

1.2.3 Result

We have the following finite-volume result.

Theorem 1.2.1. *For α sufficiently small one has*

$$\lim_{l \downarrow 1} \sup_{n \in \mathbb{N}} \sup_{\triangle \in triang(\mathcal{T}_n)} \mathbb{E}_{P_{n,l}}[\ |\nabla \hat{\omega}(\triangle) - \mathrm{Id}|^2\ ] = 0 \tag{1.2.2}$$

with the constant value of the Jacobian $\nabla \hat{\omega}(\triangle)$ on the tetrahedron $\triangle$ from the triangulation of $\mathcal{T}_n$ and some norm $|\cdot|$ on $\mathbb{R}^{3\times 3}$.

The choice of α has to be such that the volume of any tetrahedron and octahedron with side lengths in $[1, 1+\alpha]$ is uniquely minimized by the regular tetrahedron and octahedron with side length 1 respectively (see proof of Theorem 1.2.1). The central argument is going to be the following rigidity theorem from [26, Theorem 3.1].

Theorem 1.2.2 (Friesecke, James and Müller). *Let U be a bounded Lipschitz domain in $\mathbb{R}^d$, $d \geq 2$. There exists a constant $C(U)$ with the following property: For each $v \in W^{1,2}(U, \mathbb{R}^d)$ there is an associated rotation $R \in \mathrm{SO}(d)$ such that*

$$||\nabla v - R||_{L^2(U)} \leq C(U)||\mathrm{dist}(\nabla v, \mathrm{SO}(d))||_{L^2(U)}.$$

This is a generalization of Liouville's theorem, which states that a map is necessarily a rotation whose Jacobian is a rotation in every point of its domain. We are going to set $v = \hat{\omega}|_{U_n}$ and $U = U_n$ which is a bounded Lipschitz domain. The function $\hat{\omega}|_{U_n}$ is linear on each triangle $\triangle \in \mathcal{T}_n$, thus piecewise affine linear on U_n. As a consequence, $\hat{\omega}|_{U_n}$ belongs to the class $W^{1,2}(U_n, \mathbb{R}^3)$. The following remark, which also appears in [26] at the end of Section 3, is essential to achieve uniformity in Theorem 1.2.2 in the parameter n.

Remark 1.2.3. *The constant $C(U)$ in* Theorem 1.2.2 *is invariant under scaling: $C(\gamma U) = C(U)$ for all $\gamma > 0$. Indeed, setting $v_\gamma(\gamma x) = \gamma v(x)$ for $x \in U$, we have $\nabla v_\gamma(\gamma x) = \nabla v(x)$ and hence $||\nabla v_\gamma - R||_{L^2(\gamma U)} = \gamma^{d/2}||\nabla v - R||_{L^2(U)}$ and $||\mathrm{dist}(\nabla v_\gamma, \mathrm{SO}(d))||_{L^2(\gamma U)} = \gamma^{d/2}||\mathrm{dist}(\nabla v, \mathrm{SO}(d))||_{L^2(U)}$. This implies that for the domains U_n $(n \geq 1)$, the corresponding* *constant $C(U_n)$ can be chosen independently of n.*

1.2.4 Proofs

We are going to show that the L^2-distance of the Jacobian $\nabla\hat{\omega}$ from the scaled identity matrix on U_n can be controlled by the difference of the areas of $\hat{\omega}(U_n)$ and U_n. Because of the periodic boundary conditions, $\lambda(\hat{\omega}(U_n))$ does not depend on configurations ω with $(\Omega 2)$, thus it provides a suitable uniform control on the set $\Omega_{n,l}$. Then we show that the expected square distance of $\nabla\hat{\omega}$ from the scaled identity matrix can be controlled by the expected square deviation of the polyhedra's edge lengths from one. The one should be associated with the lattice constant of the unscaled lattice.

The following two lemmas from [39] provide the desired rigidity estimate on tetrahedra and octahedra. They state that the distance from SO(3) of a piecewise affine linear map defined on the polyhedron can be controlled by terms that measure how the map deforms the edge lengths of the polyhedron. We conjecture that any convex, rigid polyhedron satisfies such rigidity estimates via Dehn's theorem and the Inverse function theorem. However in this dissertation, as its main concern is not rigidity theory, we will only consider tetrahedra and octahedra for which these estimates are already proven. Let $|M| = \sqrt{\mathrm{tr}(M^t M)}$ denote the Frobenius norm of a matrix $M \in \mathbb{R}^{3\times 3}$ and $|w|$ the Euclidean norm of $w \in \mathbb{R}^3$.

Lemma 1.2.4 ([39] Lemma 3.2.)**.** *There is a positive constant C_1 such that, for all linear maps $A : \mathbb{R}^3 \to \mathbb{R}^3$ with* $\det(A) > 0$ *and $w_1 = (1, 0, 0)$, $w_2 = (\frac{1}{2}, \frac{\sqrt{3}}{2}, 0)$, $w_3 = w_2 - w_1$,*

$w_4 = (\frac{1}{2}, \frac{\sqrt{3}}{6}, \frac{\sqrt{6}}{3})$, $w_5 = w_4 - w_2$, $w_6 = w_4 - w_1$ *and* $l \geq 1$, *the following inequality holds:*

$$\operatorname{dist}^2(A\ ,\ \mathrm{SO}(3)) := \inf_{R \in \mathrm{SO}(3)} |A - R|^2 \leq C_1 \sum_{i=1}^{6} (|Aw_i| - 1)^2. \tag{1.2.3}$$

A similar theorem holds for octahedra. Let $\mathcal{O}$ denote an octahedron with vertices P_i, $i \in \{1, \ldots, 6\}$, and edges P_iP_j for $i \neq j \pmod 3$.

Lemma 1.2.5 ([39] Lemma 3.4.)**.** *There is a constant* $C_2 > 0$ *such that*

$$\operatorname{dist}^2(\nabla u\ ,\ \mathrm{SO}(3)) \leq C_2 \sum_{i \neq j \ (mod\ 3)} (|u(P_iP_j)| - 1)^2 \quad \textit{almost everywhere in } \mathcal{O}, \tag{1.2.4}$$

for every $u \in \mathcal{C}^0(\mathcal{O}; \mathbb{R}^3)$ *such that* u *is piecewise affine with respect to the triangulation determined by cutting* $\mathcal{O}$ *along the diagonal* P_1P_4, $\det(\nabla u) > 0$ *a.e. in* $\mathcal{O}$, *and* $u(\mathcal{O})$ *is convex.*

Now, we prove the mentioned estimate, which provides control over the L^2-distance of $\nabla\hat{\omega}$ from the scaled identity matrix in terms of the edge length deviations.

Lemma 1.2.6. *For a polyhedron* $\triangle \in \mathcal{T}$, *let* $\mathcal{E}(\triangle)$ *denote the set of edges of* $\triangle$. *There is a constant* $c > 0$ *such that for all* $n \geq 1$ *and* $1 < l < 1 + \alpha$, *the inequality*

$$||\ \nabla\hat{\omega} - l\ \mathrm{Id}\ ||^2_{L^2(U_n)} \leq c \sum_{\triangle \in \mathcal{T}_n} \sum_{\{x,y\} \in \mathcal{E}(\triangle)} (|\omega(x) - \omega(y)| - 1)^2 \tag{1.2.5}$$

holds for all $\omega \in \Omega_{n,l}$, and hence

$$\mathbb{E}_{P_{n,l}}[\ ||\ \nabla\hat{\omega} - l\ \mathrm{Id}\ ||^2_{L^2(U_n)}\ ] \leq c \sum_{\triangle\in\mathcal{T}_n} \sum_{\{x,y\}\in\mathcal{E}(\triangle)} \mathbb{E}_{P_{n,l}}[\ (|\omega(x)-\omega(y)|-1)^2\ ] \tag{1.2.6}$$

where the L^2-norm is defined with respect to the scalar product on $\mathbb{R}^{3\times 3}$ that induces the Frobenius norm, and $|\cdot|$ denotes the Euclidean norm on $\mathbb{R}^3$.

Note that the right side in equation (1.2.5) is strictly positive because of the boundary conditions (1.2.1) and because $l > 1$, whereas the left is zero for $\omega = \omega_l \in \Omega^{per}_{n,l}$. Since the measure $P_{n,l}$ is supported on the set $\Omega_{n,l}$, (1.2.6) follows from (1.2.5). Also note that c does not depend on n.

Proof. Let $\omega \in \Omega_{n,l}$ and $\mathcal{E}(\triangle)$ be the set of edges of a polyhedron $\triangle \in \mathcal{T}_n$. By Lemma 1.2.4 and Lemma 1.2.5 we conclude that on every polyhedron $\triangle \in \mathcal{T}_n$, we have

$$\mathrm{dist}^2\left(\nabla\hat{\omega}|_\triangle,\ \mathrm{SO}(3)\right) \leq \max\{C_1, C_2\} \sum_{\{x,y\}\in\mathcal{E}(\triangle)} (|\omega(x)-\omega(y)|-1)^2$$

where we used (Ω1), (Ω3) and (Ω4) to apply Lemmas 1.2.4 and 1.2.5 and with the constants C_1, C_2 from Lemmas 1.2.4 and 1.2.5. Orthogonality of functions which are nonzero only on disjoint polyhedra gives

$$||\ \mathrm{dist}(\nabla\hat{\omega}, \mathrm{SO}(3))\ ||^2_{L^2(U_n)} \leq C \sum_{\triangle\in\mathcal{T}_n} \sum_{\{x,y\}\in\mathcal{E}(\triangle)} (|\omega(x)-\omega(y)|-1)^2$$

with constant $C = \max\{C_1,\ C_2\} \max\{\sqrt{2}/12, \sqrt{2}/3\}$ where the second factor is the maximum of the volumes of a regular tetrahedron and octahedron. Applying Theorem 1.2.2 about geometric rigidity, we find an $R(\omega) \in \mathrm{SO}(3)$ such that

$$|| \nabla\hat{\omega} - R(\omega) ||^2_{L^2(U_n)} \leq K \, || \operatorname{dist}(\nabla\hat{\omega}, \mathrm{SO}(3)) ||^2_{L^2(U_n)},$$

with a constant $K > 0$ that does not depend on n by Remark 1.2.3. Due to the periodic boundary conditions (1.2.1), the function $\hat{\omega} - l \, \mathrm{Id}$ is n-periodic in the directions t_1, t_2, t_3, this is to say

$$\hat{\omega}(x + nt_i) - l(x + nt_i) = \hat{\omega}(x) - lx \quad \text{for all } x \in \mathbb{R}^3 \text{ and } \; i \in \{1, 2, 3\}. \tag{1.2.7}$$

By the fundamental theorem of calculus, the gradient of a periodic function is orthogonal to any constant function, and therefore

$$|| \nabla\hat{\omega} - l \, \mathrm{Id} ||^2_{L^2(U_n)} + || \, l \, \mathrm{Id} - R(\omega) ||^2_{L^2(U_n)} = || \nabla\hat{\omega} - R(\omega) ||^2_{L^2(U_n)}$$

by Pythagoras. Since $P_{n,l}$ is supported on the set $\Omega_{n,l}$, the lemma is established with $c = CK$. □

With Lemma 1.2.6 we can now prove Theorem 1.2.1.

Proof of Theorem 1.2.1. A generalization of Heron's formula for tetrahedra gives the volume $\lambda(\triangle)$ of the tetrahedron $\triangle$ with edge lengths u, v, w, U, V, W (opposite edges denoted with the same letter, lower case and capital)

$$\lambda(\triangle) = \frac{\sqrt{(-a+b+c+d)(a-b+c+d)(a+b-c+d)(a+b+c-d)}}{192\ uvw} \tag{1.2.8}$$

with

$$X = (w-U+v)(U+v+w) \qquad a = \sqrt{xYZ}$$

$$x = (U-v+w)(v-w+U) \qquad b = \sqrt{yZX}$$

$$Y = (u-V+w)(V+w+u) \qquad c = \sqrt{zXY}$$

$$y = (V-w+u)(w-u+V) \qquad d = \sqrt{xyz}$$

$$y = (V-w+u)(w-u+V)$$

$$Z = (v-W+u)(W+u+v)$$

$$z = (W-u+v)(u-v+W).$$

By first order Taylor approximation of (1.2.8) at the regular tetrahedron $\triangle_1$, denoting the edge lengths a_i, $i \in \{1, \ldots, 6\}$ we obtain

$$\lambda(\triangle) - \lambda(\triangle_1) = \frac{1}{12\sqrt{2}} \sum_{i=1}^{6} (a_i - 1) + o\left(\sum_{i=1}^{6} |a_i - 1|\right) \quad \text{as } a_i \to 1 \text{ for all } i.$$

For the octahedron, we obtain $\frac{1}{6\sqrt{2}}$ for the volume derivative in one edge b_1 at $b_1 = 1$ and the remaining 11 edges fixed at $b_i = 1$. This can be achieved by dividing the octahedron into

4 tetrahedrons that all have a common edge d that is a diagonal of the octahedron adjacent to x. Using the formula (1.2.8) and some elementary geometry of a regular trapezoid to see that $d = \sqrt{x+1}$, we obtain with the regular octahedron $\hexagon_1$ of edge length 1:

$$\lambda(\hexagon) - \lambda(\hexagon_1) = \frac{1}{6\sqrt{2}} \sum_{i=1}^{12} (b_i - 1) + o\left(\sum_{i=1}^{12} |b_i - 1| \right) \quad \text{as } b_i \to 1 \text{ for all } i.$$

We only need that the partial derivatives of the volume at $\triangle_1$ and $\hexagon_1$ are positive. By continuity, in a small neighborhood of the regular polyhedra, increasing one edge length increases the volume. Therefore we can choose $\alpha > 0$ from the definition of allowed configurations so small such that the polyhedra of the tessellation obtain minimal volume as the edge lengths go to 1. We choose $c_1 > 12\sqrt{2}$ and a corresponding $\alpha > 0$ so small that the inequalities

$$\begin{aligned} \sum_{i=1}^{6} (a_i - 1) &\leq c_1 (\lambda(\triangle) - \lambda(\triangle_1)) \\ \sum_{i=1}^{12} (b_i - 1) &\leq c_1 (\lambda(\hexagon) - \lambda(\hexagon_1)) \end{aligned} \tag{1.2.9}$$

are satisfied whenever $1 < a_i < 1 + \alpha$ and $1 < b_i < 1 + \alpha$. Let us fix such $c_1 > 0$ and $\alpha > 0$ and assume that $\Omega_{n,l}^{per}$ is defined by means of this α. Using (1.2.9) we can also estimate the squared edge length deviations:

$$\begin{aligned} \sum_{i=1}^{6} (a_i - 1)^2 &\leq c_1 \ \alpha \ (\lambda(\triangle) - \lambda(\triangle_1)) \\ \sum_{i=1}^{12} (b_i - 1)^2 &\leq c_1 \ \alpha \ (\lambda(\hexagon) - \lambda(\hexagon_1)) \end{aligned} \tag{1.2.10}$$

By equation (1.2.5) from Lemma 1.2.6 and (1.2.10), we get an upper bound on $||\nabla\hat{\omega} - l\ \mathrm{Id}||^2_{L^2(U_n)}$ in terms of the area differences. By summing up the contributions (1.2.10) of the polyhedra $\triangle \in \mathcal{T}_n$, we conclude for all $\omega \in \Omega_{n,l}$ that

$$|| \ \nabla\hat{\omega} - l\ \mathrm{Id}\ ||^2_{L^2(U_n)} \leq c_1\ \alpha\ c \sum_{\triangle\in\mathcal{T}_n} (\lambda(\hat{\omega}(\triangle)) - \lambda(\triangle)). \tag{1.2.11}$$

As a consequence of (Ω2) and the periodic boundary conditions (1.2.1), the right hand side in (1.2.11) does not depend on $\omega \in \Omega_{n,l}$. Hence, with $\omega_l \in \Omega_{n,l}$ we can compute

$$\sum_{\triangle\in\mathcal{T}_n} (\lambda(\hat{\omega}(\triangle)) - \lambda(\triangle)) = \sum_{\triangle\in\mathcal{T}_n} (\lambda(\hat{\omega}_l(\triangle)) - \lambda(\triangle)) = |U_n|(l^3 - 1). \tag{1.2.12}$$

The combination of the equations (1.2.11) and (1.2.12) gives

$$|| \ \nabla\hat{\omega} - l\ \mathrm{Id}\ ||^2_{L^2(U_n)} \leq c_1\ \alpha\ c\ |U_n|\ (l^3 - 1). \tag{1.2.13}$$

The reference measure $\delta_0 \otimes \lambda^{I_n\setminus\{0\}}$ and the set of allowed configurations $\Omega_{n,l}$ are invariant under the translations

$$\psi_b : \Omega^{\mathrm{per}}_{n,l} \to \Omega^{\mathrm{per}}_{n,l} \quad (\omega(x))_{x\in I} \mapsto (\omega(x+b) - \omega(b))_{x\in I}$$

for $b \in T$. As a consequence the matrix valued random variables $\nabla(\hat{\omega}(\triangle))$ are identically distributed for $\triangle, \widetilde{\triangle} \in \mathrm{triang}(\mathcal{T}_n)$ such that $\triangle = \widetilde{\triangle} \pmod T$. Thus for any $\triangle \in \mathrm{triang}(\mathcal{T}_1)$

the random variables $\nabla(\hat{\omega}(\triangle + t))_{t\in T}$ are identically distributed. Therefore

$$\mathbb{E}_{P_{n,l}}[\ ||\ \nabla\hat{\omega} - l\ \mathrm{Id}\ ||^2_{L^2(U_n)}\ ] = \sum_{\triangle\in\mathrm{triang}(\mathcal{T}_1)} |U_n(\triangle)|\ \mathbb{E}_{P_{n,l}}[\ |\nabla\hat{\omega}(\triangle) - l\ \mathrm{Id}|^2\ ]$$

with the regions $U_n(\triangle)$ of U_n taken up by T-translates of $\triangle$. Since the proportions $|U_n(\triangle)|/|U_n|$ are independent of n for any $\triangle \in \mathrm{triang}(\mathcal{T}_1)$, this equation together with (1.2.13), implies

$$\limsup_{l\downarrow 1}\ \sup_{n\in\mathbb{N}}\ \sup_{\triangle\in\mathrm{triang}(\mathcal{T}_n)} E_{P_{n,l}}[\ |\nabla\hat{\omega}(\triangle) - l\ \mathrm{Id}|^2\ ] = 0.$$

By means of the triangle inequality, we see that for all $\triangle \in \mathrm{triang}(\mathcal{T}_n)$ and $\omega \in \Omega_{n,l}$

$$|\nabla\hat{\omega}(\triangle) - \mathrm{Id}|^2 \le |\nabla\hat{\omega}(\triangle) - l\ \mathrm{Id}|^2 + c_2^2(l-1)^2 + 2c_2\ |l-1|\ |\nabla\hat{\omega}(\triangle) - l\ \mathrm{Id}|$$

with $c_2 = |\mathrm{Id}| > 0$. For $\omega \in \Omega_{n,l}$, the term $|\nabla\hat{\omega}(\triangle) - l\ \mathrm{Id}|$ is uniformly bounded for $l \in (1, \alpha)$ and $n \in \mathbb{N}$, which proves the theorem. $\square$

1.3 Two-dimensional model with local geometry dependent interactions

In this section, we extend the result of [29] about long-range orientational order in that we get rid of the a priori enumeration of two-dimensional hard disk configurations by an underlying triangular lattice and merely impose local geometry dependent conditions by means of a Hamiltonian H. The conditions impose that hard disks have exactly six neighbors that are not too far away. We show that long-range orientational order carries over to infinite-volume Gibbsian point processes defined by H.

1.3.1 Definitions

Let us cite some definitions from [22]. We equip the plane $\mathbb{R}^2$ with its Borel σ-algebra $\mathcal{B}(\mathbb{R}^2)$ and by λ we denote the Lebesgue measure on $(\mathbb{R}^2, \mathcal{B}(\mathbb{R}^2))$. The characters Λ and Δ will always denote measurable regions in $\mathbb{R}^2$ and the notation $\Delta \Subset \mathbb{R}^2$ means that in addition Δ is bounded. Consider the set $\mathcal{X} \subset 2^{(\mathbb{R}^2)}$ of locally finite point configurations in $\mathbb{R}^2$. That means $X \in \mathcal{X}$ is a subset $X \subset \mathbb{R}^2$ and for any $\Delta \Subset \mathbb{R}^2$, the intersection $X_\Delta := \mathrm{pr}_\Delta(X) := X \cap \Delta$ has finite cardinality $|X_\Delta| < \infty$. The counting variables $N_\Delta(X) := |X_\Delta|$ generate a σ-algebra $\mathcal{A} := \sigma(N_\Delta : \Delta \Subset \mathbb{R}^2)$ on $\mathcal{X}$. The union of $X, Y \in \mathcal{X}$ will be denoted by XY, this will be used when defining the configuration $X_\Lambda Y_{\Lambda^c}$ that agrees with X on Λ and with Y on the complement of Λ. In a sequence of set operation, unions XY are to evaluate first in order to reduce brackets. On the measurable space $(\mathcal{X}, \mathcal{A})$, we consider the Poisson point process Π^z with intensity $z > 0$. The measure Π^z is uniquely characterized by the properties that for all $\Delta \Subset \mathbb{R}^2$ under Π^z: (i) N_Δ is Poisson distributed with parameter $z\lambda(\Delta)$, and (ii) conditional on $N_\Delta = n$, the n points in Δ are independently and uniformly distributed on Δ for each integer $n \geq 1$. Similarly, configurations $\mathcal{X}_\Lambda = \{X_\Lambda : X \in \mathcal{X}\}$ in the set Λ carry the trace σ-algebra $\mathcal{A}'_\Lambda := \mathcal{A}|_{\mathcal{X}_\Lambda}$ and the reference measure Π^z_Λ which is the law of X_Λ if X is distributed according to Π^z. We will also need the pullback of $\mathcal{A}'_\Lambda$ to $\mathcal{X}$ defined by $\mathcal{A}_\Lambda := \mathrm{pr}_\Lambda^{-1} \mathcal{A}'_\Lambda \subset \mathcal{A}$. Finally, we define the shift group $\Theta = \{\theta_r : r \in \mathbb{R}^2\}$, where $\theta_r : \mathcal{X} \to \mathcal{X}$

is the translation by $-r \in \mathbb{R}^2$, consequently $N_\Delta(\theta_r X) = N_{\Delta+r}(X)$ for all $\Delta \Subset \mathbb{R}^2$.

We fix $\alpha > 0$ small enough, the size of α will be specified later. We change the notation of [29] from ϵ to α at this point to emphasize that α is fixed and not particularly small. Let $\Lambda^{1+\alpha} := \{x \in \mathbb{R}^2 : \ |x-y| < 1+\alpha \text{ for some } y \in \Lambda\}$ be the $(1+\alpha)$-enlargement of Λ. For $X \in \mathcal{X}$ we define the Hamiltonian $H_{\Lambda,Y}$ in Λ with boundary condition $Y \in \mathcal{X}$ by

$$H_{\Lambda,Y}(X) := \begin{cases} 0 & \text{if } |x-y| > 1 \text{ whenever } x \in X_\Lambda Y_{\Lambda^{1+\alpha}\setminus\Lambda} \text{ and } y \in X_\Lambda Y_{\Lambda^c} \text{ for } x \neq y, \\ & \text{and for all } x \in X_\Lambda Y_{\Lambda^{1+\alpha}\setminus\Lambda} : \ |X_\Lambda Y_{\Lambda^c} \cap A_{1,1+\alpha}(x)| = 6; \\ \infty & \text{otherwise.} \end{cases}$$

This is to say that $H_{\Lambda,Y}(X) \in \{0,\infty\}$ takes the value 0 if and only if every point of $X_{\Lambda^{1+\alpha}}$ has distance greater than one from points in $X_\Lambda Y_{\Lambda^c}$ and has exactly six $X_\Lambda Y_{\Lambda^c}$-neighbors in the annulus $A_{1,1+\alpha}(x) = \{y \in \mathbb{R}^2 : |y-x| \in (1,1+\alpha)\}$, otherwise H is defined to be infinity. Note that the only part of the boundary condition Y relevant for $H_{\Lambda,Y}(X)$ is in the region $\Lambda^{2(1+\alpha)} \setminus \Lambda$.

Definition 1.3.1. *We define the partition function* $Z^z_{\Lambda,Y}$ *by*

$$Z^z_{\Lambda,Y} := \Pi^z_\Lambda\{X_\Lambda : H_{\Lambda,Y}(X_\Lambda) = 0\} = \int e^{-H_{\Lambda,Y}(X)} \Pi^z_\Lambda(dX).$$

We call a boundary condition $Y \in \mathcal{X}$ *admissible for the region* $\Lambda \Subset \mathbb{R}^2$ *if* $0 < Z^z_{\Lambda,Y}$. *We write* $\mathcal{X}^{\Lambda,z}_*$ *for the set of all these* Y.

The set of admissible boundary conditions $\mathcal{X}^{\Lambda,z}_*$ is never empty as the $l \in (1,1+\alpha)$

multiply of a triangular lattice with lattice constant one is always in $\mathcal{X}_*^{\Lambda,z}$. We note that $H_{\Lambda,Y}(\emptyset) = 0$ for $Y_{\Lambda^{1+\alpha}} = \emptyset$ and also for specifically chosen Λ and possibly nonempty Y. The partition function $Z^z_{\Lambda,Y}$ is zero, if neither $Y_{\Lambda^{1+\alpha}\setminus\Lambda} = \emptyset$ nor the boundary condition $Y_{\Lambda^{1+\alpha}\setminus\Lambda}$ can be extended to a near triangular lattice configuration in $\Lambda^{1+\alpha}$.

Definition 1.3.2. *For $Y \in \mathcal{X}_*^{\Lambda,z}$, we define the Gibbs distribution in the region $\Lambda \Subset \mathbb{R}^2$ with boundary condition Y by the formula*

$$\gamma^z_\Lambda(F|Y) = \int_{\mathcal{X}_\Lambda} \mathbb{1}_F(XY_{\Lambda^c}) e^{-H_{\Lambda,Y}(X)} \Pi^z_\Lambda(\mathrm{d}X)/Z^z_{\Lambda,Y},$$

where $F \in \mathcal{A}$. Note that $\gamma^z_\Lambda(\cdot|Y)$ is a measure on the whole space $(\mathcal{X}, \mathcal{A})$.

In case of $Y_{\Lambda^\alpha\setminus\Lambda} \neq \emptyset$, the $\mathcal{X}_\Lambda$-marginal of the measure $\gamma^z_\Lambda(\cdot|Y)$ is uniform on the configurations in $\mathcal{X}_\Lambda$ that extended $Y_{\Lambda^\alpha\setminus\Lambda}$ to a near triangular lattice configuration in Λ^α. Otherwise if $Y_{\Lambda^\alpha\setminus\Lambda} = \emptyset$, then $\gamma^z_\Lambda(\cdot|Y) = \delta_{Y_{\Lambda^c}}$. Note that $(F, Y) \in (\mathcal{A}, \mathcal{X}) \mapsto \gamma^z_\Lambda(F|Y)$ is a probability kernel from $(\mathcal{X}, \mathcal{A}_{\Lambda^c})$ to $(\mathcal{X}, \mathcal{A})$, but the distribution $\gamma^z_\Lambda(\cdot|Y)$ has $\delta_{Y_{\Lambda^c}}$ as its marginal on $(\mathcal{X}_{\Lambda^c}, \mathcal{A}'_{\Lambda^c})$.

Definition 1.3.3 (infinite-volume Gibbs measure)**.** *A probability measure P on $(\mathcal{X}, \mathcal{A})$ is an infinite-volume Gibbs measure for $z > 0$ if $P(\mathcal{X}_*^{\Lambda,z}) = 1$ and*

$$\int f \mathrm{d}P = \int_{\mathcal{X}_*^{\Lambda,z}} \frac{1}{Z^z_{\Lambda,Y}} \int_{\mathcal{X}_\Lambda} f(XY_{\Lambda^c}) e^{-H_{\Lambda,Y}(X)} \Pi^z_\Lambda(\mathrm{d}X) P(\mathrm{d}Y)$$

for every $\Lambda \Subset \mathbb{R}^2$ and every measurable $f : \mathcal{X} \to [0, \infty)$. We denote the set of infinite-volume Gibbs measures by $\mathcal{G}^z$.

Note that the right hand side in the defining equality is equal to $\mathbb{E}_P[\gamma_\Lambda^z(f|\cdot)]$. Therefore, a measure P is infinite-volume Gibbs measure, if and only if $P\gamma_\Lambda^z = P$ for every $\Lambda \Subset \mathbb{R}^2$, where the product is understood as taking average with P in the second variable of γ_Λ^z. We can easily see a degenerated measure $\delta_\emptyset \in \mathcal{G}^z$, however we will be interested in more interesting Gibbs measures. In fact, as soon as $P(\emptyset) = 0$ for a measure $P \in \mathcal{G}^z$, we have that P is supported on hard disk configurations with infinitely many disks.

The Hamiltonian H implements an example of a k-nearest neighbor interaction as explained in [22, Chapter 4.2.1]. Therefore by [22, Lemma 5.1.], the kernels γ_Λ^z, γ_Δ^z for $\Lambda \subset \Delta \Subset \mathbb{R}^2$ and $Y \in \mathcal{X}_*^{\Lambda,z}$ satisfy the consistency conditions $\gamma_\Lambda^z(\mathcal{X}_*^{\Lambda,z}|Y) = 1$ and $\gamma_\Delta^z \gamma_\Lambda^z = \gamma_\Delta^z$, where the product is understood as product of probability kernels.

1.3.2 Results

We show the following generalization[1] of [29, Thm. 4.1].

Theorem 1.3.4. *Let $0 < \alpha$ be small enough (such that Lemma 1.3.5 and Theorem 1.3.6 hold true for the choice of this α). Then for every $2/(\sqrt{3}(1+\alpha)^2) < \rho < 2/\sqrt{3}$ (the density of centers in the densest packing of disks with diameter 1), there is a measure $P_\rho \in \cap_{z>0}\mathcal{G}^z$ such that*

(i) Density $= \rho$*: For any $\Lambda \Subset \mathbb{R}^2$, we have $\mathbb{E}_{P_\rho}[N_\Lambda] = \rho\lambda(\Lambda)$.*

(ii) Translational invariance*: The measure P_ρ is translational invariant in any direction in*

[1]The wording of Theorem 1.3.4 up to some minor modification in the definition of H was suggested by Franz Merkl in a talk at a conference (Trends in Mathematical Crystallization) held at Warwick University in May 2016

$\mathbb{R}^2$, *i.e.* $P_\rho \circ \theta_r^{-1} = P_\rho$ *for any* $r \in \mathbb{R}^2$.

(iii) Long-range orientational order*: Let* $x \in X$ *be the point with the smallest distance from the origin. It is a.s. unique. We have* $P_\rho(N_{A_{1,1+\alpha}(x)} = 6) = 1$*. Choose a random neighbor* $y \in X$ *of* x *(i.e.* $1 < |y - x| < 1 + \alpha$*) uniformly distributed among all six neighbors. Then as* $\rho \uparrow 2/\sqrt{3}$*, the law of* $y - x$ *w.r.t.* P_ρ *converges weakly to the uniform distribution on the 6th roots of unity in* $\mathbb{C} \hat{=} \mathbb{R}^2$*.*

Note that by translational invariance of P_ρ, property *(iii)* holds when initially picking the closest point x to any reference point $x_0 \in \mathbb{R}^2$ instead of the origin. Hence the long-range orientational order, as neighbors of x position themselves close to translates of the 6th roots of unity. The choice of α will be made somewhat explicit in the proof of Lemma 1.3.5. The set of Gibbs measures $\mathcal{G}^z$ is most likely independent of $z > 0$, however we will not pursue the proof of this statement as it leads to geometric considerations that are not in the center of our analysis.

1.3.3 Proofs

For a configuration $X \in \mathcal{X}$, we say that $H(X) = 0$ if for all $x, y \in X$, we have $|x - y| > 1$ and $|X \cap A_{1,1+\alpha}(x)| = 6$. This is the same as having $H_{\Lambda,X}(X) = 0$ for any $\Lambda \Subset \mathbb{R}^2$. For a configuration $\emptyset \neq X \in \mathcal{X}$ with $H(X) = 0$, we can define a simplicial complex $K(X)$ consisting of zero, one and two cells defined as follows. The set of zero cells $K_0(X)$ is $X \subset \mathbb{R}^2$. The set of one cells $K_1(X)$ are edges between zero cells of distance between 1 and $1 + \alpha$, and the two cells are triangles with sides in $K_1(X)$. We will see in the following lemma, that by definition of H and some geometric considerations, for α small enough, the graph defined by the one and two skeleton of this complex is locally, and therefore also globally isomorphic to the triangular lattice $I = \mathbb{Z} + \tau\mathbb{Z}$ with $\tau = e^{\frac{i\pi}{3}}$ with edge set $E = \{\{i, j\} \subset I : |i - j| = 1\}$. The set of triangles surrounded by three edges in E is denoted by $\mathcal{T}$, these are two cells if we regard I as a simplicial complex.

The most important lemma linking the theorem above to [29, Thm. 4.1] is the following.

Lemma 1.3.5. *With the choice of a small enough α, we have for any configuration $X \in \mathcal{X}$ with $H(X) = 0$, that the graph defined by the one and two skeletons of $K(X)$ is isomorphic to the triangular lattice I. In other words, there is a bijective map $\omega : I \to X$ such that for all $i, j \in I$: $|i - j| = 1$ if and only if $|\omega(i) - \omega(j)| \in (1, 1 + \alpha)$.*

Later on, we will choose α small enough such that Lemma 1.3.5 and Theorem 1.3.4 both work for that α. From the proof of the lemma it will be obvious that the choice of α does not need to be particularly small for it (and any smaller choice) to work.

Proof. We define for $i \in I$ its closest neighborhood $N(i) \subset I$ by $N(i) = \{j \in I : |i - j| \leq 1\}$. Let $X \in \mathcal{X}$ such that $H(X) = 0$. A map $\omega : N(i) \to X$ is called a local isomorphism at i if for all $j, k \in N(i)$, we have $|j - k| = 1$ if and only if $|\omega(j) - \omega(k)| \in (1, 1 + \alpha)$. By taking $\alpha > 0$ small enough, we can ensure that for all $i \in I$ and $x \in X$ there is a local isomorphism ω at i such that $\omega(i) = x$. To see this, observe that as $\alpha \to 0$, for every $y \in A_{1,1+\alpha}(x)$ there are exactly two points $y_1, y_2 \in A_{1,1+\alpha}(x) \setminus \{y\}$ such that $|y_i - y| \to 1$, for other $z \in A_{1,1+\alpha}(x) \setminus \{y\}$, we have $\liminf_{\alpha \to 0} |z - y| \geq \sqrt{3}$. Since we know that $|X \cap A_{1,1+\alpha}(y)| = 6$, a simple geometric consideration related to the kissing problem, gives that $y_1, y_2 \in A_{1,1+\alpha}(y)$, since if $y_i \notin A_{1,1+\alpha}(y)$ for $i \in \{1, 2\}$, for α small enough there was not enough space to place 6 points in $A_{1,1+\alpha}(y)$ having distance bigger than 1 from each other and from y_i. To be more precise, for all $i \in I$ and $x \in X$ there will be twelve such local isomorphisms taking rotations and reflection into account. We fix α small enough such that the local isomorphism property holds.

Let us construct a map $\omega : I \to X$ as follows. We fix an arbitrary $x_0 \in X$ and define $\omega|_{N(0)}$ to be one of the six orientational preserving local isomorphism at 0 with $\omega(0) = x_0$. Fix a spanning tree T of I. For each $i \in I$, there is a unique path on nearest neighbors in T connecting 0 to i. Since there are local isomorphism at each pair of points of I and X, we can successively, uniquely extend ω to vertices of T by choosing the unique of the six orientation preserving local isomorphisms that is consistent with T. This is to say that if for a neighbor i of j in T, we already assigned a point $\omega(i)$ then we already choose a local isomorphism at i with $i \mapsto \omega(i)$. Let us assign j to the point in X which is determined by this local isomorphism. Now, there is only one local isomorphism at j, which is consistent with the local isomorphism chosen at i in the sense that i has identical images under the two local isomorphisms. We use this local isomorphism to proceed with the construction and map all neighbors of j in T into X.

It remains to show that the map $\omega : I \to X$ is an isomorphism. To conclude ω is an isomorphism onto its image, we fix a loop γ starting and ending in $i \in I$ composed of a path in T and an edge between i and one of its neighbors in I to which it is not connected in T. We need to show that the map induced along γ with an initial orientational preserving local isomorphism $\omega|_{N(i)}$ at i, maps to a loop in $K(X)$ starting and ending in $\omega(i)$. To this end we can show a seemingly more general but equivalent statement. Take any loop $\gamma = (i_0, i_1, i_2, \ldots, i_n)$ at $0 \in I$ (i.e. $i_0 = i_n = 0$) and $x \in X$, fix a local isomorphism at 0 with $0 \mapsto x$ and show that the map induced along γ maps γ to a loop $\omega(\gamma)$ in X at x. Here ω is locally defined along the curve γ.

We can deform the loop γ to the boundary of a two cell that contains 0 by successively "removing" two cells that intersect γ and are inside of it. By removing a two cell, we mean one of the following. Two subsequent edges (i_{k-1}, i_k), (i_k, i_{k+1}) of γ, we can exchange for the unique edge (i_{k-1}, i_{k+1}) if $|i_{k-1} - i_{k+1}| = 1$, or we can exchange one edge (i_k, i_{k+1}) of γ for two edges (i_k, j) and (j, i_{k+1}) in I. For every such transformation of γ, we obtain a modified γ' and a map ω' that is uniquely determined by the local isomorphism at i_k and is the unique extension of the local isomorphism at 0 along γ'. Note that $\omega = \omega'$ on the domain that they are both defined and $\omega(\gamma)$ is closed if and only if $\omega'(\gamma')$ is. When after removing finitely many two cells, we arrive at $\gamma' = (0, i, j, 0)$ being the boundary of a two cell that contains the origin. Since $\omega'|_{\gamma'}$ should be the unique extension of the local isomorphism at 0 along γ, we see that $\omega'(\gamma')$ is closed and therefore so is $\omega(\gamma)$.

We showed that for neighbors i, j in I, also $\omega(i)$, $\omega(j)$ are neighbors in X. To obtain the converse statement and the injectivity of ω, we repeat the preceding procedure for the same map ω but with exchanged roles of I and X. This concludes the proof that ω is an isomorphism onto its image.

It remains to show that ω is surjective. Take now a curve $\hat{\gamma}$ in $K(X)$ from x_0 to some $y \in K(X)$. Note that $K(X)$ is a connected graph, as for small enough α and $x \neq y$ we can always find a neighbor z of x which is closer to y than x. The curve $\hat{\gamma}$ corresponds to a curve

γ in I from 0 to some $i \in I$. Applying the procedure from above to the concatenation of the path from 0 to i in T and the reverse of γ, we see that $\omega(i) = y$. $\square$

Lemma (1.3.5) can be also proved with the formalism of Čech cohomology using the de Rham isomorphism and can be generalized to configurations with point defects (missing points). The usefulness of the Čech cohomology and de Rham's theorem was pointed out to us by Franz Merkl. We decided to give another proof using less formalism.

To construct P_ρ, we use measures on periodic configurations. For $l > 1$ and $n \in \mathbb{N}$, let us define measures $P_{n,l}$ on n-periodic configurations as in [29]. A periodic, enumerated configuration $\omega \in \Omega_{n,l}^{per}$ is a map $I \to \mathbb{R}^2$ such that Theorem 1.3.6 holds true for this choice of α.

$$\omega(i + nj) = \omega(i) + lnj \quad \text{for all } i, j \in I. \tag{1.3.1}$$

It suffices to define an n-periodic, enumerated configuration on a set of n^2 representatives $I_n \subset I$ as equation (1.3.1) uniquely defines the configuration on the complement $(I_n)^c$. The event of admissible, n-periodic, enumerated configurations $\Omega_{n,l} \subset \Omega_{n,l}^{per}$ is defined by the properties $(\Omega 1) - (\Omega 3)$:

$(\Omega 1)$ $|\omega(i) - \omega(j)| \in (1, 1 + \alpha)$ for all $\{i, j\} \in E$.

For $\omega \in \Omega$ we define the extension $\hat{\omega} : \mathbb{R}^2 \to \mathbb{R}^2$ such that $\hat{\omega}(i) = \omega(i)$ if $i \in I$, and on the closure of any triangle $\triangle \in \mathcal{T}$, the map $\hat{\omega}$ is defined to be the unique affine linear extension of the mapping defined on the corners of $\triangle$.

$(\Omega 2)$ The map $\hat{\omega} : \mathbb{R}^2 \to \mathbb{R}^2$ is injective.

$(\Omega 3)$ The map $\hat{\omega}$ is orientation preserving, this is to say that $\det(\nabla \hat{\omega}(x)) > 0$ for all

$\triangle \in \mathcal{T}$ and $x \in \triangle$ with the Jacobian $\nabla\hat{\omega} : \cup\mathcal{T} \to \mathbb{R}^{2\times 2}$.

Define the set of *admissible, n-periodic, enumerated configurations* as

$$\Omega_{n,l} = \{\omega \in \Omega_{n,l}^{per} \mid \omega \text{ satisfies } (\Omega 1)\text{–}(\Omega 3)\}.$$

Let the probability measure $P_{n,l}$ be

$$P_{n,l}(A) = \frac{\delta_0 \otimes \lambda^{I_n \setminus \{0\}}(\Omega_{n,l} \cap A)}{\delta_0 \otimes \lambda^{I_n \setminus \{0\}}(\Omega_{n,l})}$$

for any Borel measurable set $A \in \mathcal{F}_n = \bigotimes_{i\in I_n} \mathcal{B}(\mathbb{R}^2)$, thus $P_{n,l}$ is the uniform distribution on the set $\Omega_{n,l}$ with respect to the *reference measure* $\delta_0 \otimes \lambda^{I_n\setminus\{0\}}$. The first factor in this product refers to the component $\omega(0)$. The parameter l in the definition of $\Omega_{n,l}$ and $P_{n,l}$ controls the density of periodic configurations such that $\rho = \frac{2}{l^2\sqrt{3}}$. We quote Theorem 4.1 from [29] which will be the major ingredient of the proof of Theorem 1.3.4.

Theorem 1.3.6. *For any $\alpha > 0$ small enough one has*

$$\lim_{l\downarrow 1} \sup_{n\in\mathbb{N}} \sup_{\triangle\in\mathcal{T}} \mathbb{E}_{P_{n,l}}[\ |\nabla\hat{\omega}(\triangle) - \mathrm{Id}|^2\ ] = 0 \qquad (1.3.2)$$

with the constant value of the Jacobian $\nabla\hat{\omega}(\triangle)$ on the set $\triangle \in \mathcal{T}$.

We note that the theorem holds for any $\alpha \in (0, \sqrt{3} - 1)$, however we omit the proof of this which is just a more careful consideration of arguments in the proof of [29, Theorem 4.1] and will refer to small enough α. The main observation needed for this explicit range of α where the theorem holds is, that the area of triangles with side lengths in the range $[1, \sqrt{3})$

is uniquely minimized by the regular triangle with side length 1. This observation is then utilized like in the similar proof of Theorem 1.2.1 in the 3D case. We note that Theorem 1.3.6 might work with $\alpha \geq \sqrt{3}-1$, however looking for the optimal upper bound is not the concern of this dissertation.

In the following we construct P_ρ as a limit of translational invariant versions of $P_{n,l}$ and show that this measure is a Gibbs measure in $\mathcal{G}^z$ for any $z > 0$. We follow ideas from [22] to construct a limiting measure. Fix $l > 1$ and define the measures G_n on $(\mathcal{X}, \mathcal{A})$ by specifying its marginal $(G_n)_{\Lambda_n}$ on $(\mathcal{X}_{\Lambda_n}, \mathcal{A}'_{\Lambda_n})$

$$(G_n)_{\Lambda_n} = \left(\frac{1}{\lambda(\Lambda_n)} \int_{\Lambda_n} \mathrm{Im}[P_{n,l}] \circ \theta_r \, \mathrm{d}r \right)_{\Lambda_n},$$

with the image measure $\mathrm{Im}[P_{n,l}]$ of $P_{n,l}$ under the map $\mathrm{Im} : \omega \mapsto \{\omega(x) : x \in I\}$ and the domain $\Lambda_n = l\{x + y\tau : x, y \in [-n/2, n/2)\}$. The averaging over $r \in \Lambda_n$ is necessary to obtain a translational invariant measure on the torus, since $\omega(0) = 0$ holds $P_{n,l}-$a.s.. The measure G_n is then defined by having i.i.d. projections on the sets $\{\Lambda_n + inl\}_{i \in I}$, which form a tiling of $\mathbb{R}^2$. In order to have translational invariant probability measures on $(\mathcal{X}, \mathcal{A})$, we consider the averaged measures

$$\hat{G}_n = \frac{1}{\lambda(\Lambda_n)} \int_{\Lambda_n} G_n \circ \theta_r \, \mathrm{d}r.$$

By definition and the periodicity of G_n, $\hat{G}_n$ are translational invariant. We will show that the sequence $(\hat{G}_n)_{n \in \mathbb{N}}$ is tight in the *topology of local convergence* on translational invariant probability measures on $\mathcal{X}$ generated by $P \to \int f \mathrm{d}P$ for functions f that are $\mathcal{A}_\Lambda$-measurable for some $\Lambda \Subset \mathbb{R}^2$. Such functions we call local and denote the set of local functions by $\mathcal{L}$.

The only difference to the definitions after Lemma 5.1. in [22] are in the nature of the measures $(G_n)_{\Lambda_n}$. In our case $(G_n)_{\Lambda_n}$ are measures that inherit geometric constraints from the structure of $P_{n,l}$ that are defined on toruses of different size. In [22] on the contrary, the authors use a measures $G^z_{\Lambda_n,\bar{\omega}}$ that have fixed boundary condition $\bar{\omega}$ on the complement of Λ_n.

For a shift invariant probability measure P on $(\mathcal{X}, \mathcal{A})$ and $\Lambda \Subset \mathbb{R}^2$ define the measure $P_\Lambda := P \circ \mathrm{pr}_\Lambda^{-1}$ and the *relative entropy* w.r.t. Π^z_Λ as

$$I(P_\Lambda | \Pi^z_\Lambda) := \begin{cases} \int f \ln f \mathrm{d}\Pi^z_\Lambda & \text{if } P_\Lambda << \Pi^z_\Lambda \text{ with density } f \\ \infty & \text{otherwise.} \end{cases}$$

The *specific entropy* of P w.r.t. Π^z is then defined by

$$I(P) := \lim_{n\to\infty} \frac{1}{\lambda(\Delta_n)} I(P_{\Delta_n} | \Pi^z_{\Delta_n}),$$

where $\Delta_n \Subset \mathbb{R}^2$ is a cofinite increasing sequence of sets. We refer to [30] and [31] for existence and properties of the specific entropy. We will set $z = 1$ and compute entropies relative to $\Pi^1_{\Delta_n}$. By [31, Proposition 2.6], the sublevel sets of I are sequentially compact in the topology of local convergence. Therefore, we only need to show that the specific entropies of the measures $\{\hat{G}_n\}_{n\in\mathbb{N}}$ are bounded by some constant. We start with citing and repeating the proof of a proposition from [28, Lemma 5.2] that provides lower bound on the partition sum.

Proposition 1.3.7. *For all $\alpha \in (0,1]$ and $l \in (1, 1+\alpha)$, there is an $r = r(\alpha, l) \in (0, 1/2)$ such that for $n \in \mathbb{N}$, we have*

$$\delta_0 \otimes \lambda^{I_n \setminus \{0\}}(\Omega_{n,l}) \geq (\pi r^2)^{|I_n|-1}. \tag{1.3.3}$$

Proof. For $r > 0$, we define, like in (3.2) in [34], the set of configurations which are close to the scaled, enumerated, standard configuration $\omega_l(i) = li$ for $i \in I$:

$$S_{n,l,r} = \{\omega \in \Omega_{n,l}^{per} \mid |\omega(i) - \omega_l(i)| < r \text{ for all } i \in I\}. \tag{1.3.4}$$

For sufficiently small $r > 0$, depending on α and l, we conclude, like in the proof of [34, Lemma 3.1], that $S_{n,l,r} \subset \Omega_{n,l}$. To prove this inclusion, we have to show the properties (Ω1)–(Ω3) for all $\omega \in S_{n,l,r}$. Let us compute for $(i,j) \in E$ and $\omega \in S_{n,l,r}$:

$$\begin{aligned} ||\omega(i) - \omega(j)| - l| &= ||\omega(i) - \omega(j)| - |\omega_l(i) - \omega_l(j)|| \\ &\leq |\omega(i) - \omega_l(i)| + |\omega(j) - \omega_l(j)| < 2r. \end{aligned}$$

If we choose $2r < \max\{l-1, 1+\alpha-l\} < 1$, then ω satisfies (Ω1). Condition (Ω2) is a consequence of the inequality $\langle v, \nabla\hat{\omega}(x)v\rangle > 0$ for all $v \in \mathbb{R} \setminus \{0\}$, and for all $x \in \mathbb{R}^2$ where $\hat{\omega}$ is differentiable. This inequality holds for small enough r since $\nabla\hat{\omega}$ is close to the identity uniformly on $\mathbb{R}^2$. Hence $\hat{\omega}$ is a bijection onto its image. Here we applied a theorem from analysis which states that a $\mathcal{C}^1$-map f from an open convex domain $U \subset \mathbb{R}^n$ into $\mathbb{R}^n$ with $\langle v, \nabla f(x) v\rangle > 0$ for all $v \in \mathbb{R}^n \setminus \{0\}$ and $x \in U$ is a diffeomorphism onto its image. However, $\nabla\hat{\omega}(x)$ is only piecewise differentiable, but on the straight line L connecting $x, y \in \mathbb{R}^2$ with $x \neq y$, there are only finitely many points $z \in \mathbb{R}^2 \cap L$ where the curve $(\hat{\omega}(ty + (1-t)x))_{t\in(0,1)}$ is not differentiable. Assume that $\langle v, \nabla\hat{\omega}(x)v\rangle > 0$ holds whenever $\hat{\omega}$ is differentiable in x. The curve is piecewise linear, and on each of these pieces, the derivative of the curve forms an acute angle with $y - x$, therefore the curve cannot be closed. Thus, the condition (Ω2) is satisfied in the case of a sufficiently small r. Furthermore, condition (Ω3) is satisfied by ω_l, therefore also by ω if r is sufficiently small. Hence $S_{n,l,r} \subset \Omega_{n,l}$ for some $r \in (0, 1/2)$, and we conclude

$$\delta_0 \otimes \lambda^{I_n\setminus\{0\}}(\Omega_{n,l}) \geq \delta_0 \otimes \lambda^{I_n\setminus\{0\}}(S_{n,l,r}) = (\pi r^2)^{|I_n|-1}$$

where the last equality is obtained by integrating over each $\omega(i)$ with $i \neq 0$ successively along a fixed spanning tree of I_n which gives a factor πr^2, and considering that $\omega_l(0) = 0$ and that the measure $\delta_0 \otimes \lambda^{I_n\setminus\{0\}}$ fixes $\omega(0) = 0$. □

Proposition 1.3.8. *The set $\{I(\hat{G}_n) : n \in \mathbb{N}\}$ is bounded, thus the set $\{\hat{G}_n : n \in \mathbb{N}\}$ is sequentially compact in the topology of local convergence. Therefore, there is a sequence $n_k \to \infty$ and a shift invariant measure P_ρ on $(\mathcal{X}, \mathcal{A})$ such that $\lim_{k\to\infty} \int f \mathrm{d}\hat{G}_{n_k} = \int f \mathrm{d}P_\rho$ for any $f \in \mathcal{L}$.*

Proof. As also noted in the proof of [22, Proposition 5.3], the definition of $\hat{G}_n$ implies that

$$I^z(\hat{G}_n) = \frac{1}{\lambda(\Lambda_n)} I\left((G_n)_{\Lambda_n} | \Pi^1_{\Lambda_n}\right).$$

The relative entropy $I\left((G_n)_{\Lambda_n} | \Pi^1_{\Lambda_n}\right)$ can be explicitly computed as follows. The measure $(G_n)_{\Lambda_n}$ is supported on configurations that have n^2 points in Λ_n and if Λ_n is folded into a torus, then each point x has exactly six neighbors in the annulus $A_{1,1+\alpha}(x)$ around it and no points closer than distance one. These configurations $\mathcal{X}_{n,l}$ are images of enumerated configurations $\mathcal{X}_{n,l} = (\text{Im } \Omega_{n,l})_{\Lambda_n}$. By Lemma 1.3.5, $(G_n)_{\Lambda_n}$ is the uniform distribution on these configurations with respect to $\Pi^1_{\Lambda_n}$. The density of $(G_n)_{\Lambda_n}$ w.r.t. $\Pi^1_{\Lambda_n}$ is given by $f = \mathbb{1}_{\mathcal{X}_{n,l}} / \Pi^1_{\Lambda_n}(\mathcal{X}_{n,l})$. To find the constant $\Pi^1_{\Lambda_n}(\mathcal{X}_{n,l})$ more explicitly, consider the expectation

$$\Pi^1_{\Lambda_n}[g] = e^{-\lambda(\Lambda_n)} \sum_{k=0}^{\infty} \int_{\Lambda_n^k} \frac{1}{k!} g(\{x_1, \ldots, x_k\})\ \lambda^k|_{\Lambda_n^k}(\mathrm{d}x_1, \ldots, \mathrm{d}x_k)$$

Consequently, we have

$$\Pi^1_{\Lambda_n}(\mathcal{X}_{n,l}) = \frac{e^{-\lambda(\Lambda_n)}}{n^2} \lambda(\Lambda_n)\ \delta_0 \otimes \lambda^{I_n \setminus \{0\}}(\Omega_{n,l}).$$

This follows since a factor $\frac{e^{-\lambda(\Lambda_n)}}{(n^2)!}$ comes from the density of $\Pi^1_{\Lambda_n}$ conditioned on n^2 points with respect to $\lambda^{(n^2)}|_{\Lambda_n^{(n^2)}}(\mathrm{d}x_1, \ldots, \mathrm{d}x_n^2)$. Then conditioned on the position of x_1, the volume of the allowed configurations by their shift invariance on the torus is $(n^2-1)!\ \delta_0 \otimes \lambda^{I_n \setminus \{0\}}(\Omega_{n,l})$, furthermore the first point can be distributed uniformly in Λ_n. The relative entropy is $I\left((G_n)_{\Lambda_n} | \Pi^1_{\Lambda_n}\right) = -\ln\left(\Pi^1_{\Lambda_n}(\mathcal{X}_{n,l})\right)$ and the specific entropy can be bounded using Proposition 1.3.7 and $\lambda(\Lambda_n) = n^2 l^2 \sqrt{3}/2$ for big enough n, we obtain

$$\begin{aligned} I\left((G_n)_{\Lambda_n}\right) = -\frac{\ln\left(\Pi^1_{\Lambda_n}(\mathcal{X}_{n,l})\right)}{\lambda(\Lambda_n)} &= 1 + \frac{n^2}{\lambda(\Lambda_n)} - \frac{\ln\left(\lambda(\Lambda_n)\right)}{\lambda(\Lambda_n)} - \frac{\ln\left(\delta_0 \otimes \lambda^{I_n \setminus \{0\}}(\Omega_{n,l})\right)}{\lambda(\Lambda_n)} \\ &\leq 1 + \frac{n^2}{\lambda(\Lambda_n)} - \frac{\ln\left(\lambda(\Lambda_n)\right)}{\lambda(\Lambda_n)} - \frac{|I_n - 1| \ln(\pi r^2)}{\lambda(\Lambda_n)} \\ &\leq 1 + \frac{2 - 2\ln(\pi r^2)}{l^2\sqrt{3}}. \end{aligned}$$

□

The next proposition shows that P_ρ is an infinite-volume Gibbs measure. Note that $\hat{G}_n$ and Λ_n depend on $l > 1$ which we fixed previously.

Proposition 1.3.9. *The measure P_ρ is an infinite-volume Gibbs measure $P_\rho \in \cap_{z>0}\mathcal{G}^z$.*

Proof. Fix $\Lambda \Subset \mathbb{R}^2$, $z > 0$ and $\rho < 2/\sqrt{3}$ large enough such that $2/(\sqrt{3}(1+\alpha)^2) < \rho$ where α is such that Lemma 1.3.5 holds with that α. Let $l > 1$ such that $\rho = 2/(l^2\sqrt{3})$. For $X \in \mathcal{X}$, let $\widetilde{X}_n$ be the periodic extension of X_{Λ_n} to $\mathcal{X}$, i.e. $\widetilde{X}_n = \cup_{i \in I} X_{\Lambda_n} + lni$. Let $\kappa > 0$ be so big such that $\Lambda^\kappa \setminus \Lambda$ contains a connected ring of triangles from $K_2(\widetilde{X}_n)$ for G_n-almost all X for all $n \in \mathbb{N}$. Consequently, for all $n \in \mathbb{N}$ large enough such that $\Lambda^\kappa \subset \Lambda_n$, the number of points in Λ conditioned on X_{Λ^c} is G_n-almost surely determined by the configuration in $\Lambda^\kappa \setminus \Lambda$. The measure $(G_n)_{\Lambda_n}$ is the uniform distribution of enumerable, allowed configurations with n^2 points on the torus. By Lemma 1.3.5, the conditional distribution of X_Λ given X_{Λ^c} under G_n is therefore the uniform distribution on configurations X_Λ such that $H_{\Lambda, X_{\Lambda^c}}(X_\Lambda) = 0$. Uniform distribution makes sense, as the number of points in Λ is almost surely constant with respect to the conditioned measure. Therefore, the factorized version of the conditional distribution of G_n given $\mathcal{A}_{\Lambda^c}$ is given by $\gamma_\Lambda(\cdot|\cdot)$, this is to say that

$$G_n(F) = \int_{\mathcal{X}} \gamma_\Lambda(F|Y) G_n(\mathrm{d}Y) \tag{1.3.5}$$

for any $F \in \mathcal{A}$ and $n \in \mathbb{N}$ big enough for $\Lambda^\kappa \subset \Lambda_n$. Since z is fixed, we can omit it as a superscript in γ^z.

The rest of the proof is as the proof of [22, Prop. 5.5.]. Define $\Lambda_n^\circ := \{r \in \mathbb{R}^2 : \Lambda^\kappa + r \subset \Lambda_n\}$ and the (subprobability) measures

$$\bar{G}_n := \frac{1}{|\Lambda_n|} \int_{\Lambda_n^\circ} G_n \circ \theta_r^{-1} \, \mathrm{d}r.$$

Then $\int f \mathrm{d}\hat{G}_n - \int f \mathrm{d}\bar{G}_n \to 0$ by the same argument as in [31, Lemma 5.7], therefore P_ρ can also be seen as an accumulation point of the sequence $(\bar{G}_n)$. Let $F \in \cup_{\Delta \Subset \mathbb{R}^2} \mathcal{A}_\Delta$ be a local set, using (1.3.5), we obtain for $r \in \Lambda_n^\circ$

$$G_n \circ \theta_r^{-1}(F) = \int_{\mathcal{X}} \gamma_\Lambda(F|Y) G_n \circ \theta_r^{-1}(\mathrm{d}Y).$$

Therefore averaging over $r \in \Lambda_n^\circ$ gives

$$\bar{G}_n(F) = \int_{\mathcal{X}} \gamma_\Lambda(F|Y) \bar{G}_n(\mathrm{d}Y). \tag{1.3.6}$$

Since the integrand on the right is a local function of Y, we can set $n = n_k$ and let $k \to \infty$, that gives (1.3.6) for P_ρ instead of $\bar{G}_n$. Since local sets generate the σ-algebra $\mathcal{A}$, (1.3.6) holds for P_ρ and $F \in \mathcal{A}$, which by monotone convergence shows that P_ρ is an infinite-volume Gibbs measure. □

Proof of Theorem 1.3.4. In Propositions 1.3.9 and 1.3.8, we showed the existence of a translational invariant measure $P_\rho \in \cap_{z>0}\mathcal{G}^z$ which is the local limit of the measures $(G_{n_k})_{k\geq 1}$, therefore P_ρ satisfies property (ii). Property (i) holds as it can be expressed by a local function and $\mathbb{E}_{G_{n_k}}[|X \cap B|] = \rho\lambda(B)$ for any $k \geq 1$ by the periodic boundary conditions. Similarly, property (iii) can be expressed by local functions depending on $\{x_0, x_1, ..., x_6\} \cap \Lambda_n$, where x_0 is the closest random point to the origin and x_i is the ith closest point to x_0. For n large enough we have $G_{n_k}(|\{x_0, x_1, ..., x_6\} \cap \Lambda_n| = 7) = 1$ for any $k \geq 1$ and therefore $P_\rho(|\{x_0, x_1, ..., x_6\} \cap \Lambda_n| = 7) = 1$. By Theorem 1.3.6 we have

$$\lim_{\rho\uparrow 2/\sqrt{3}} \sup_{k\geq 1} \mathbb{E}_{G_{n_k}} \left[\sum_{i=1}^{6} |\nabla\hat{\omega}(\triangle_i) - \mathrm{Id}|^2\right] = 0, \tag{1.3.7}$$

where $\{\triangle_i\}_{1\leq i\leq 6}$ are the random six triangles in $\mathcal{T}$ such that one of their vertices is mapped to x_0 under ω. Let $f : \mathbb{C}^6 \to \mathbb{R}$ be continuous, bounded and permutation invariant. We use the natural identification of topological spaces $\mathbb{C}\hat{=}\mathbb{R}^2$. Let $y_i = x_i - x_0$. By continuity of f, there is a constant $c > 0$ such that

$$\left|f(y_1, \ldots, y_6) - f(e^{i\pi/3}, e^{i2\pi/3}, \ldots, e^{i2\pi})\right| \leq c \sum_{i=1}^{6} |\nabla\hat{\omega}(\triangle_i) - \mathrm{Id}|^2 \tag{1.3.8}$$

G_{n_k}-a.s. for any $k \geq 1$. Combining equations (1.3.7) and (1.3.8), we obtain that

$$\lim_{\rho \uparrow 2/\sqrt{3}} \mathbb{E}_{P_\rho}\left[\left|f(y_1, \ldots, y_6) - f(e^{i\pi/3}, e^{i2\pi/3}, \ldots, e^{i2\pi})\right|\right] = 0$$

which concludes the proof of property (iii). □

Chapter 2

Value Learning of Spiking Neurons

2.1 Introduction

In the context of reinforcement learning, we consider an agent (in Section 2.3 represented by a two-compartment neuron) being exposed to a changing environment. The environment provides sensory stimulation (e.g. odors) in terms of presynaptic activity to a dendritic compartment, and reward information (e.g. encoded by a dopamine or octopamine firing rate) r_t^{M} at each discrete time step t to the soma. This reward information is encoded by nudging conductances g_t^{E} and g_t^{I} such that the matching potential U_t^{M}, which is elevated in case of a reward and decreased in case of a punishment (as defined in [46]), yields an instantaneous firing rate $r_t^{\mathrm{M}} = \phi(U_t^{\mathrm{M}})$.

We wish to adapt the synaptic strengths on the dendritic compartment such that the instantaneous firing rate, $\phi(U_t)$, in some form approximates the total discounted future reward, originally defined as $R_t^{\mathrm{M}} = \sum_{i=0}^{\infty} \gamma^i r_{t+i}^{\mathrm{M}}$, with some discount factor $\gamma \in [0, 1)$ (with the convention $0^0 = 1$). We propose a learning rule that achieves this goal. We analyze the rule's basic properties and demonstrate it in computer simulations. This part of the dissertation is joint work with Johanni Brea, Robert Urbanczik and Walter Senn [18].

2.2 The case of clamping

The case of clamping refers to a simplified model layout in which the reward r_t^{M} is intrinsically available to the synapses. In Section 2.3, we turn our attention to a two-compartment neuron model by R. Urbanczik and W. M. Senn, [46], where reward is driven by nudging. The two-compartment neuron model has the advantage that nudging, thus also the reward, might be provided by a neuronal network which is a biologically more reasonable model. It also enables predictions on the behavioral time scale of seconds, even though the plasticity window is tens of milliseconds.

The case of clamping captures the most important ideas of value learning in a simplified model. For this reason, we recognize the importance of this simplification on its own. Most strategies which are discussed in the present section can be also used later in Section 2.3.

2.2.1 The model

In this model, we study a single neuron with N presynaptic neurons. The synaptic input matrix, $\Theta \in \mathbb{R}^{N\times M}$, indicates M distinct input patterns consisting of the contribution of N presynaptic neurons. The entry $\Theta_{i,j}$ of Θ, for $i = 1, \ldots, N$ and $j = 1, \ldots, M$, is the contributions of the ith synapse to the jth pattern.

Let $y = (y_t)_{t\in\mathbb{Z}}$ be a stationary, irreducible Markov chain on some finite state space E with transition matrix π. The adapted process y is realized on a suitable filtered probability space $(\Omega, \mathcal{F}, P)$ with filtration $\mathcal{F}_t \subset \mathcal{F}$ for $t \in \mathbb{Z}$. Without loss of generality $E = \{e_1, ..., e_M\}$, where $e_j \in \mathbb{R}^M$ is the unit vector in the jth coordinate direction. Let $f : E \to E$ be a function, and set $m = |\text{range}(f)|$. We define the partially observable Markov chain $x_t = f(y_t)$. Note that x is not a Markov chain in general, unless f is injective. A state $e_j \in E$ represents the input pattern Θe_j, which is presented to the neuron whenever $x_t = e_j$. Hence, the stochastic process $\rho_t = \Theta x_t$ with values in $\mathbb{R}^N$ represents the unweighted, postsynaptic activity at time $t \in \mathbb{Z}$. Later on, ρ will represent the firing rate process of the presynaptic neurons. The dendritic potential $V_t(w)$ at time t is given by the scalar product

$$V_t(w) := w \cdot \text{PSP}_t := w \cdot \text{PSP}(\rho_t), \tag{2.2.1}$$

where $w \in \mathbb{R}^N$ denotes the synaptic weight vector and the function PSP : $\{\Theta e \mid e \in \text{range}(f)\} \to \mathbb{R}^N$ determines the unweighted postsynaptic potential. We can think of PSP as a function from the reals into the reals that is applied to each entry of ρ_t independently. In subsequent sections, we will add a random argument to PSP from a suitable probability space such that PSP will represent random potentials driven by Poisson processes with intensities given by the entries of ρ_t. We abbreviate $\text{PSP}_e := \text{PSP}(\Theta e)$, $V_e(w) := w \cdot \text{PSP}_e \in \mathbb{R}$, $V(w) = w^T\text{PSP} \in \mathbb{R}^m$ defined by the matrix $\text{PSP} := (\text{PSP}(\Theta e))_{e\in\text{range(f)}} \in \mathbb{R}^{N\times m}$, with a small abuse of notation, such that lower indices can be both vectors (denoting an input pattern $\text{PSP}_e = \text{PSP}(\Theta e)$) and integers (denoting a timepoint $\text{PSP}_t = \text{PSP}(\rho_t)$) for the sake

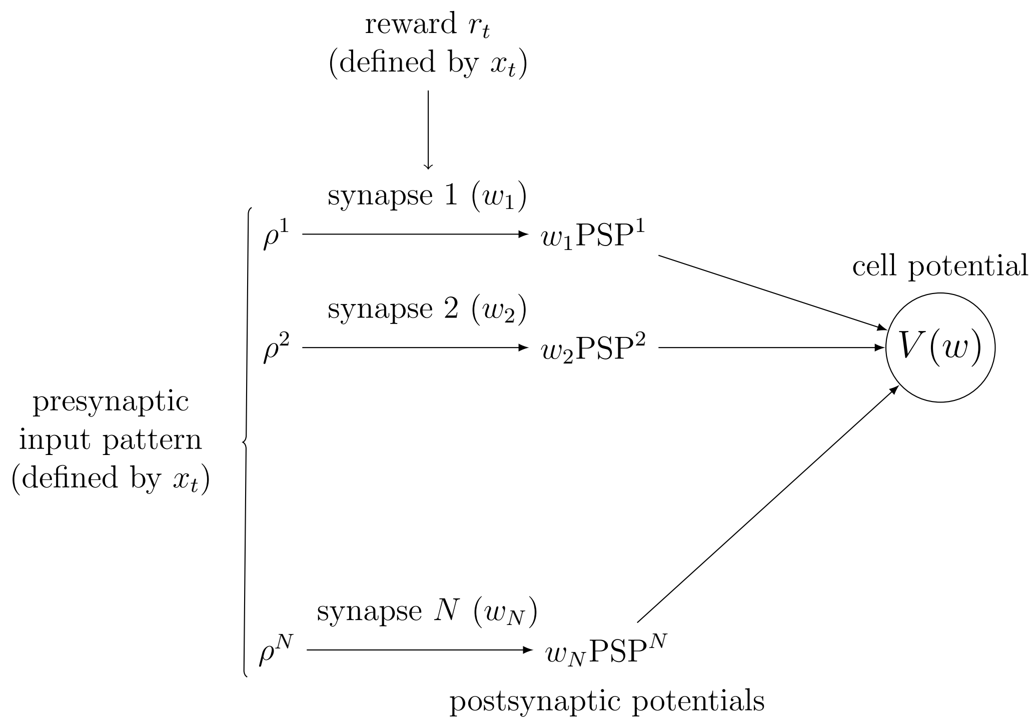

Figure 2.1: Design in the case of clamping

of notational simplicity. In Figure 2.1 we illustrated the schematic design of the system. The upper indices in the figure denote the entry of the vectors which notation we do not explicitly use, rather adhere to a matrix and vector notation. Since PSP deterministically depends on ρ in this simplified model, we can refer to PSP instead of ρ as the input pattern.

By stationarity, we have time independent probabilities $p_e := P(x_t = e)$ and $p_e \geq 0$ where equality holds if and only if $e \notin \text{range}(f)$. Let $\mathbb{E}[\cdot]$ denote expectation with respect to the probability measure P. For $e \in E$, let $\mathbb{E}_{x_t=e}[\cdot]$ denote the conditional expectation, conditioned on the event $\{x_t = e\}$, for the case of $t = 0$, we will write $\mathbb{E}_e[\cdot]$. For $e \notin \text{range}(f)$, conditional expectations are defined to be some fixed constant. Furthermore, let us define a reward function $r : E \to \mathbb{R}$. The real number $r_e := r(e)$ is called the *reward* due to the state

e. Throughout, we will use both lower index and argument notation for functions which are defined on E. For such a function h, we write h_e and $h(x_t)$, and we also note that we can identify h with $h \in \mathbb{R}^M$ via the linear isomorphism, induced by the mapping $e_i \mapsto i$ for $1 \leq i \leq M$.

Let us fix a discount factor $\gamma \in [0, 1)$ and a scale factor $\alpha > 0$. The *reinforcement value* of a state $e \in E$ is defined by

$$\mathbf{R}_e := \alpha \mathbb{E}_e \left[\sum_{t=0}^{\infty} \gamma^t r(x_t) \right] = \alpha \sum_{t=0}^{\infty} \gamma^t \mathbb{E}_e \left[r(x_t) \right], \tag{2.2.2}$$

the expected, discounted, future reward given the starting point $x_0 = e$. We also define the random variable $R_t = \alpha \sum_{n=0}^{\infty} \gamma^n r(x_{t+n})$. By stationarity, for any $t \in \mathbb{Z}$, we have $\mathbf{R}_e = \mathbb{E}_{x_t = e} [R_t]$. For the reason of normalization, we set $\alpha = 1 - \gamma$ in this section. Please note that since (x_t) is in general not a Markov chain, $\mathbf{R}_e$ is not the best estimate for the discounted future reward given the history at the time and before observing the state e. In case of a Markov chain (x_t), however, $\mathbb{E}[R_t | \mathcal{F}_n, n \leq t]$ is x_t-measurable.

Let ϕ be a continuously differentiable function of sublinear growth from the reals into the reals, i.e. there are positive constants a, b such that

$$|\phi(x)| \leq a|x| + b.$$

Furthermore, we assume that ϕ has strictly positive, bounded first derivative such that $r_e \in \text{range}(\phi)$ for all $e \in \text{range}(f)$. Since $\alpha = 1 - \gamma$, also $\mathbf{R}_e \in \text{range}(\phi)$ for all $e \in \text{range}(f)$, and ϕ is diffeomorphism onto range(ϕ) such that ϕ^{-1} exists. The goal of value learning is to adjust the weight vector, w, such that

$$\mathbf{R}(x_t) = \phi(V_t(w)) \tag{2.2.3}$$

at every time point $t \in \mathbb{Z}$. Equivalently, we can write

$$\phi^{-1}(\mathbf{R}) = w^T \mathrm{PSP}, \tag{2.2.4}$$

with the vector $\phi^{-1}(\mathbf{R})$, where we apply ϕ^{-1} entrywise to the vector $\mathbf{R} = (\mathbf{R}_e)_{e \in \mathrm{range}(f)}$. Equation (2.2.4) is a linear equation, which is to be solved for w. If PSP has rank m, then equation (2.2.4) has a solution, and the solutions form an affine linear subspace W^* of $\mathbb{R}^N$. This is the case, if patterns PSP_e for $e \in \mathrm{range}(f)$ are linearly independent in $\mathbb{R}^N$ which we will assume throughout the discussion.

2.2.2 Learning rule

In order to achieve equality in (2.2.3), we propose an online learning rule. For an arbitrary starting point $w_0 \in \mathbb{R}^N$, we define a stochastic sequence of synaptic weights w_t by the update rule

$$w_{t+1} = w_t + \eta_t \left(r(x_t)\ \widetilde{\mathrm{PSP}}_t - \phi(V_t(w_t))\ \mathrm{PSP}_t \right) \quad \text{for } t \in \mathbb{N}, \tag{2.2.5}$$

where $\eta_t > 0$ is a deterministic sequence of learning rates, and $\widetilde{\mathrm{PSP}}_t$ is called the *(discounted) eligibility trace*:

$$\widetilde{\mathrm{PSP}}_t(\Theta, x) := \alpha \sum_{i=0}^{\infty} \gamma^i \ \mathrm{PSP}_{t-i} \tag{2.2.6}$$

which is an exponential average of past postsynaptic potentials. The name eligibility trace comes from reinforcement learning and refers to the eligibility to update an entry of w. Note that if a given entry of PSP_t and the corresponding entry of $\widetilde{\mathrm{PSP}}_t$ are both zero, then the update rule (2.2.5) leaves the corresponding entry of w_t unchanged. Since PSP_{t-i} can only take finitely many values and $\gamma \in [0, 1)$, (2.2.6) is well-defined. In this simplified model, the synaptic update (2.2.5) directly depends on the current reward $r(x_t)$, which was indicated in Figure 2.1 as an arrow to the synapses.

Lemma 2.2.1. *The expectation of the update in rule (2.2.5) is equal to the expectation of* $\eta_t \ [\ R_t - \phi(V_t(w))\ ]\ \mathrm{PSP}_t$ *if we fix* w. *This is to say that for all* $t \in \mathbb{Z}$ *and* $w \in \mathbb{R}^N$ *we have*

$$\mathbb{E}\Big[r(x_t)\ \widetilde{\mathrm{PSP}}_t - \phi(V_t(w))\ \mathrm{PSP}_t\Big] = \ \mathbb{E}\Big[\,[\ R_t - \phi(V_t(w))\ ]\ \mathrm{PSP}_t\Big]. \tag{2.2.7}$$

Proof. We only need to show that for a stationary process $(x_t)_{t\in\mathbb{Z}}$ and two bounded, real valued functions h, f defined on the set where x takes its values, we have

$$\mathbb{E}[f_t \tilde{h}_t] = \mathbb{E}[\hat{f}_t h_t] \quad \text{for all } t \in \mathbb{Z},$$

with $f_t = f(x_t)$, $\hat{f}_t = \sum_{i=0}^{\infty} \gamma^i f(x_{t+i})$ and $\tilde{h}_t = \sum_{i=0}^{\infty} \gamma^i f(x_{t-i})$. However this is a direct consequence of stationarity, since $\mathbb{E}[f_t \; h_{t-i}] = \mathbb{E}[f_{t+i} \; h_t]$. Applying this observation to the bounded function r, and entrywise to the bounded components of the function PSP, we arrive at the statement of the lemma. □

Now, we turn our attention to the issue of convergence when following the update rule (2.2.5). We will show that (2.2.5) is a stochastic gradient descend algorithm. We follow the idea in [46, Supp. Information, pp 8.] and define the loss function

$$l(v, y) := -\int_0^v (y - \phi(s))\mathrm{d}s.$$

Assume that $y \in \mathrm{range}(\phi)$. Since ϕ is monotonically increasing, $l(\cdot, y)$ attains its unique global minimum at $v = \phi^{-1}(y)$. Let us define the cost function $J : \mathbb{R}^N \to \mathbb{R}$:

$$J(w) := \mathbb{E}[l(V_0(w), \mathbf{R}(x_0))] = \sum_{e \in E} l(V_e(w), \mathbf{R}_e) \; p_e$$

which has a global minimum at w if and only if $w \in W^*$. We remark that J is invariant under translations $w \to w + u$, $u \in TW^*$, where TW^* is the tangential space of W^*. For this purpose choose $w_1^*, w_2^* \in W^*$, and observe that V_e is linear and $V_e(w_1^* - w_2^*) = 0$ for all $e \in \mathrm{range}(f)$.

In the following lemma, we study the function J.

Lemma 2.2.2. *The ordinary differential equation (ODE)*

$$\dot{w}(t) = -\nabla J(w(t)) \tag{2.2.8}$$

is globally asymptotically stable and invariant under translations $w \to w + u$, $u \in TW^*$*. The set of equilibrium points is* W^**. For any* $w^* \in W^*$*, the domain of attraction* $D(w^*)$ *is the orthogonal affine linear subspace to* W^* *in* w^**, namely*

$$D(w^*) = \{w \in \mathbb{R}^N \mid w - w^* \perp u - w^* \text{ for all } u \in W^*\}. \tag{2.2.9}$$

This is to say that for all $w_0 \in D(w^*)$*, the solution* $w(t)$ *of (2.2.8) with initial condition* $w(0) = w_0$ *satisfies* $\lim_{t\to\infty} w(t) = w^*$*. Furthermore, if* $w(0) \in D(w^*)$ *for some* $w^* \in W^*$*, then* $w(t) \in D(w^*)$ *holds for all* $t \geq 0$*.*

Note that for $w_1^*, w_2^* \in W^*$ with $w_1^* \neq w_2^*$, we have $D(w_1^*) \cap D(w_2^*) = \emptyset$.

Proof. The gradient of J in w is given by

$$\nabla J(w) = -\sum_{e \in \text{range}(f)} (\mathbf{R}_e - \phi(V_e(w)))\, \text{PSP}_e\; p_e = -\mathbb{E}\left[(R_0 - \phi(V_0(w)))\, \text{PSP}_0\right]. \tag{2.2.10}$$

By the assumption that PSP has full rank, $\{\text{PSP}_e\}_{e\in\text{range}(f)}$ are linearly independent. Therefore $\nabla J(w) = 0$ implies $\mathbf{R}_e - \phi(V_e(w)) = 0$ because $p_e > 0$ for all $e \in \text{range}(f)$. We obtain that the set of equilibrium points of the ODE (2.2.8) is W^*.

The Hessian of J is given by

$$\text{Hess } J(w) = \sum_{e \in E} p_e \; \phi'(V_e(w)) \text{PSP}_e \; \text{PSP}_e^T \tag{2.2.11}$$

which is positive semidefinite, because $\phi' > 0$ and for $x \in \mathbb{R}^N$:

$$x^T \text{ Hess } J(w) \; x = \sum_{e \in E} p_e \; \phi'(V_e(w)) \; \left(x \cdot \text{PSP}_e\right)^2 .$$

The quadratic form x^T Hess $J(w)$ x is zero if and only if $\text{PSP}_e \perp x$ for all $e \in \text{range}(f)$. Let $TW^* = \{w \in \mathbb{R}^N \mid w = w_1^* - w_2^* \text{ for some } w_1^*, w_2^* \in W^*\}$ denote the tangential space of W^*. We have $TW^* = \text{Im}(\text{PSP})^\perp$. Consequently, J is strictly convex on each $D(w^*)$ and attains its unique minimum in w^*.

The lemma follows from Lyapunov's global asymptotic stability theorem ([45, § 13. VII. Theorem]) applied to the Lyapunov function $V(w) = J(w) - J(w^*)$ on $D(w^*)$. Since $\nabla J(w) \in \text{Im}(\text{PSP})$, we conclude that $\nabla J(w) \perp TW^*$, and therefore $w(t) \in D(w^*)$ for all $t \geq 0$, whenever $w(0) \in D(w^*)$. This concludes the proof. □

Now, we show with the ODE method for adaptive algorithms driven by Markov chains [17, Part I. 2, page 40] that (2.2.5) is a stochastic gradient descend method which converges almost surely to w^* whenever $w(0) \in D(w^*)$.

Define the map $F(y, z) = (f(y), z)$ for $(y, z) \in E \times \mathbb{R}^N$. We set $X_t := \left(x_t, \widetilde{\text{PSP}}_t\right)$ which is a stationary, partially observable Markov chain with state space $E \times \mathbb{R}^N$, since $X_t = F(Y_t)$ where $Y_t = (y_t, \widetilde{\text{PSP}}_t)$ is a Markov chain on $E \times \mathbb{R}^N$ and with respect to $(\mathcal{F}_t)_{t \in \mathbb{Z}}$. For $w \in \mathbb{R}^N$, let us define the map H_1:

$$H_1(w, Y_t) := r(x_t)\ \widetilde{\mathrm{PSP}}_t - \phi(V_t(w))\ \mathrm{PSP}_t. \tag{2.2.12}$$

Distinguishing between H_1 and H_2 in Section 2.3, we use different indices. In the present section, we only discuss H_1, so let us write H instead of H_1 in the current section. The learning rule (2.2.5) then can be written as

$$w_{t+1} = w_t + \eta_t H(w_t, Y_t). \tag{2.2.13}$$

Furthermore, we set $h(w) := \mathbb{E}[H(w, Y_0)]$, thus $h \equiv -\nabla J$ by Lemma 2.2.1 and equation (2.2.10). Therefore, the expectation of (2.2.5) is a discretized version of the ODE (2.2.8) with step size η_t. Similar to Lemma 2.2.2, we wish to conclude for (2.2.13) that $\lim_{t\to\infty} w_t = w^*$ a.s. given that $w_0 \in D(w^*)$.

The proof of the following lemma is straightforward.

Lemma 2.2.3. *We have the following growth estimate with the Euclidean norm* $|\cdot|$

$$\sup_{w\in\mathbb{R}^N}\ \sup_{Y=\left(e,\widetilde{\mathrm{PSP}}\right)\in E\times\mathbb{R}^N} \frac{|H(w,Y)|}{(1+|w|)\left(1+\left|\widetilde{\mathrm{PSP}}\right|\right)} < \infty. \tag{2.2.14}$$

Let Π denote the transition kernel of (Y_t). This is to say that for $Y = \left(y, \widetilde{\mathrm{PSP}}\right) \in E\times\mathbb{R}^N$,

$S \subseteq E$, and any Borel set $B \subseteq \mathbb{R}^N$, we have

$$\Pi(Y, S \times B) = \sum_{s \in S} \pi_{y,s} \mathbb{1}_{\{\gamma \widetilde{\mathrm{PSP}} + \mathrm{PSP}_{f(s)} \in B\}}. \tag{2.2.15}$$

Since the irreducible Markov chain y on the finite state E space has an invariant measure, the assumptions in Theorem 5 of [17, Part II., 2.2.1] are fulfilled by Π. In our case, Π does not depend on the weight w. As a consequence of this observation and Lemma 2.2.3, all assumptions of [17, Part II., 1.9.1] are fulfilled by Π and H. This holds especially for the Poisson equation (A4):

$$(\mathbb{1} - \Pi)\nu_w = H(w, \cdot) - h(w)$$

by setting $\nu_w(Y) = H(w, Y) - h(w)$. By boundedness of ϕ', we obtain

$$|H(w, Y) - H(w', Y)| \leq \sup \phi' \; |w - w'| \; |\mathrm{PSP}| \leq C|w - w'|,$$

with the constant $C = |\mathrm{PSP}| \; \sup \phi' > 0$ using the Frobenius norm $|\mathrm{PSP}|$ of the PSP matrix. Therefore, we obtain the following theorem from [17, Part II., 1.9.2, Theorem 17] and Lemma 2.2.2 with the choice $U(w) = (J(w) - J(w^*))^2$.

Theorem 2.2.4. *Let $(\eta_t)_{t \in \mathbb{N}}$ be a sequence of learning rates such that $\sum_{t \in \mathbb{N}} \eta_t = \infty$ and $\sum_{t \in \mathbb{N}} \eta_t^2 < \infty$. Then the learning rule (2.2.13) converges with probability one for any initial parameters $w_0 \in D(w^*)$ and $Y_0 \in E \times \mathbb{R}^n$ to $w^* \in W^*$.*

In particular, equation (2.2.3) holds in the limit $t \to \infty$, achieving the predictions

$$\lim_{t\to\infty} |\mathbf{R}(x_t) - \phi(V_t(w_t))| = 0 \quad \textit{almost surely.}$$

The proof of this theorem is based on the observation that the expectation of the update rule (2.2.5) is a discretized version of the ODE (2.2.8) with step size η_t which tends to zero fast enough such that the paths of the ODE solution and the learning trajectory converge in fashion of Euler's method, but slow enough for the learning not to stop and build time averages that coincide with probability ensemble averages by the ergodicity of the underlying Markov chain (x_t). For the details we refer to [17, Part II., 1.9.2, Theorem 17].

In simulations, we will choose η small, but fixed. We will see that in such case w_t in the limit fluctuates closely around the target w^*. In the neuroscientific application of Theorem 2.2.4, it is not reasonable to have $\eta_t \to 0$, nevertheless η_t could be modulated by biological processes.

2.3 Two-compartment neuron model

In this section we apply the value learning algorithm from the previous section to a two-compartment neuron model from [46]. The result is a biologically plausible model for spike time dependent synaptic plasticity that enables the neuron to learn predictions of future events on behavioral timescales. Even though synapses are only eligible to update up to some tens of milliseconds after a presynaptic spike, the two-compartment neuron model allows to predict an event seconds before its occurrence.

In this section, we are going to distinguish between dendritic potential $V_t(w)$ and somatic potential U_t as in [46], however, in contrast to [46], we work in a time discrete model $t \in \mathbb{Z}$. The evolution of U_t is random just as the evolution of $V_t(w)$. Let us define y_t, $x_t = f(y_t)$, Θ, ρ_t, ϕ, $V_t(w)$, PSP_t, PSP and $\widetilde{\mathrm{PSP}}_t$ as in the previous section.

The continuous, two-compartment neuron model from [46] assumes that the evolution of the somatic potentials U_t is described by an ODE that is driven by the dendritic potential $V_t(w)$ and an input current I_t which nudges the soma:

$$C\ \dot{U}_t = -g^{\mathrm{L}}\ U_t + g^{\mathrm{D}}\ (V_t(w) - U_t) + I_t(U_t), \tag{2.3.1}$$

with constant capacitance C to be set to unitless 1 and constant conductances g^{L} and g^{D}. In contrast to the single compartment neuron, for the two-compartment neuron, the reward signal r_t from the previous section will be derived from the input current I_t. This gives the theoretic possibility that the reward is generated by a distinct network of neurons, however, for our analysis we will treat I_t as a parameter of the model that we can define arbitrarily and indirectly through some conductance parameters.

The input current I_t is defined by means of time varying conductances g_t^{E} and g_t^{I} that connect the soma to constant reservoirs of excitatory E^{E} and inhibitory E^{I} potentials. The meaning of excitatory and inhibitory in this context is that $E^{\mathrm{E}} > 0$ is above, while $E^{\mathrm{I}} < 0$ is below the resting potential 0 of the somatic membrane potential. Typical values of cell

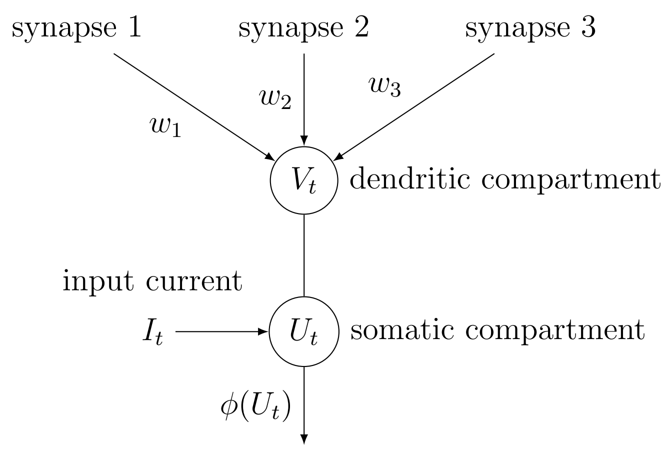

Figure 2.2: Design of the two-compartment neuron

membrane potential range from -40 mV to -80 mV, and the typical value of resting potential in an excitable neuron is around –60 mV to -70 mV. We set the resting potential to be unitless and equal to zero, and define the input current by

$$I_t(U_t) = g_t^{\mathrm{E}}(E^{\mathrm{E}} - U_t) + g_t^{\mathrm{I}}(E^{\mathrm{I}} - U_t). \tag{2.3.2}$$

The physical meaning of the equations (2.3.1) and (2.3.2) is illustrated in Figure 2.2, where arrows indicate the only direction of influence between pairs of objects. The outgoing arrow below the soma indicates the neuron's axon. The ODE (2.3.1) means that the somatic potential is the consequences of three current flows, one due to the soma's coupling to the dendrite (with conductance g^{D}), another leak current due to ion channels (with leak conductance g^{L}) that try to reset the membrane to its resting potential, and an input current I_t that drives the somatic potential towards $E^{\mathrm{E}} > 0$ in rewarding events and towards $E^{\mathrm{I}} < 0$ in case of punishment (negative reward).

We define the *matching potential* U_t^{M} as the solution of 2.3.2 with $I_t(U^{\mathrm{M}}) \equiv 0$:

$$U_t^{\mathrm{M}} = \frac{g_t^{\mathrm{E}}\ E^{\mathrm{E}} + g_t^{\mathrm{I}}\ E^{\mathrm{I}}}{g_t^{\mathrm{E}} + g_t^{\mathrm{I}}}.$$

The matching rate $\phi(U_t^{\mathrm{M}})$ takes the role of the reward signal from the previous section. The matching potential is the potential of the soma for which the input current is zero. We interpret it as the potential at which the neuron fully obtained the hypothetical reward provided. Instead of working with the solution of (2.3.1), we define the somatic potential U_t by the steady state ($\dot{U} \equiv 0$) of (2.3.1):

$$U_t(w) = \lambda_t V_t^*(w) + (1 - \lambda_t) U_t^{\mathrm{M}}, \tag{2.3.3}$$

where $\lambda_t = \frac{g^L + g^D}{g^L + g^D + g_t^{\mathrm{E}} + g_t^{\mathrm{I}}}$ is called *nudging factor* and $V_t^*(w) = \frac{g^D}{g^L + g^D} V_t(w)$ is called the *attenuated dendritic potential*. For $e \in E$, we define $V_e^*(w) = \frac{g^D}{g^L + g^D} V_e(w)$ with the small abuse of notation similarly as in the previous section, and the vector $V^*(w) = (V_e^*(w))_{e \in \mathrm{range}(f) \subset E} = \frac{g_D}{g_L + g_D} w^T \mathrm{PSP}$. The case $\lambda \equiv 0$ corresponds to the case of clamping from the previous section. The approximation (2.3.3) is motivated by the strong coupling limit case of large total conductance $g^L + g^D + g_t^{\mathrm{E}} + g_t^{\mathrm{I}} >> 1$ and slowly changing inhomogeneities $g_t^{\mathrm{E}}, g_t^{\mathrm{I}}, V_t$ in [46] and [18].

We define the *matching value* and the *somatic value* by

$$\mathbf{R}_e^{\mathrm{M}} = \mathbb{E}_e\left[\sum_{t=0}^{\infty} \gamma^t \phi(U_t^{\mathrm{M}})\right] \quad \text{and} \quad \mathbf{R}_e(w) = \mathbb{E}_e\left[\sum_{t=0}^{\infty} \gamma^t \phi(U_t)\right]$$

respectively. The goal of learning is to adjust the weight vector, w, such that $\phi(V_t^*(w))$ predicts the matching value $\mathbf{R}^{\mathrm{M}}(x_t)$. Setting $r_t = \phi(U_t^{\mathrm{M}})$ in rule (2.2.5) would result in predicting $\mathbf{R}^{\mathrm{M}}(x_t)$ as shown in Theorem 2.2.4. However, for purposes of biological modeling, the learning rule cannot depend on directly $\phi(U_t^{\mathrm{M}})$ but on $\phi(U_t)$ (or more precisely, on the somatic spiking with rate $\phi(U_t)$). We study the question, whether setting $r_t = \phi(U_t)$ in (2.2.5) also leads to the same, or at least a similar result. The somatic firing rate $\phi(U_t)$ respectively the reward vector

$$\mathbf{R}(w) = (\mathbf{R}_e(w))_{e \in \mathrm{range}(f)}$$

serves as a moving target since U_t depends on w, see (2.3.3). In the next subsection, we will see that setting $\phi(U_t)$ as reward instead of $\phi(U_t^{\mathrm{M}})$ in the learning rule, has biologically interesting consequences for the time window of the predictions.

2.3.1 Fixed point of the learning

In the following theorem, we look at the solution space of equation $\phi(V_e^*(w)) = \alpha\, \mathbf{R}_e(w)$ for $e \in E$ for linear rate function ϕ.

Theorem 2.3.1. *If the rate function ϕ is linear and $\alpha\, \lambda_{\max} < 1 - \gamma$, then the system of equations*

$$\phi(V_e^*(w)) = \alpha\, \mathbf{R}_e(w) \quad \textit{for } e \in \mathrm{range}(f) \tag{2.3.4}$$

has a nonempty affine linear solution space W^*. If, in addition, $\lambda_t = \lambda$ is constant, then for

$w^* \in W^*$ *we obtain*

$$\phi(V_e^*(w^*)) = \frac{\alpha(1-\lambda)}{1-\alpha\lambda}\mathbb{E}_{x_t=e}\left[\sum_{i=0}^{\infty}\left(\frac{\gamma}{1-\alpha\lambda}\right)^i \phi(U_{t+i}^{\mathrm{M}})\right] \quad \textit{for all } e \in \mathrm{range}(f). \tag{2.3.5}$$

We see from (2.3.5) that setting the reward $r_t = \phi(U_t)$ in rule (2.2.5), in case of convergence to this fixed point, would result in predicting future behavior of the matching potential through the matching reward, but with a different discount factor γ and a different normalizing constant α. We will call

$$\gamma_{\text{eff}} = \frac{\gamma}{1-\alpha\lambda}$$

the *effective discount factor* and $\frac{\alpha(1-\lambda)}{1-\alpha\lambda}$ the *effective normalizing constant*. We can see that the effective discount factor converges to ∞ as $\alpha\lambda \to 1$. This allows the two-compartment neuron to learn predictions on a behavioral timescale by choosing $\alpha\lambda < 1$ close enough to 1 for $\gamma_{\text{eff}} < 1$ to be sufficiently large.

Proof. We have $\phi(U_t) = \lambda_t\phi(V_t^*(w)) + (1-\lambda_t)\phi(U_t^{\mathrm{M}})$, because ϕ is linear. Since equation (2.3.4) is an inhomogeneous system of linear equations, the solution space W^* is affine linear or empty. For $e \in \mathrm{range}(f)$ and any random variable R, we have

$$\mathbb{E}_{x_0=e}[R] = \frac{\sum_{s\in f^{-1}(\{e\})} p_s \; \mathbb{E}_{y_0=s}[R]}{\sum_{s\in f^{-1}(\{e\})} p_s}.$$

By the assumption rank(PSP) $= m$, there exist linearly independent $v_1, \ldots, v_M$ in $\mathbb{R}^N$ such that for all $e \in \text{range}(f)$

$$\text{PSP}_e = \frac{\sum_{s \in f^{-1}(\{e\})} p_s v_s}{\sum_{s \in f^{-1}(\{e\})} p_s}.$$

Thus by linearity of (2.3.4), we obtain (2.3.5) for nonbijective f, provided we have proved the theorem for bijective f. Without loss of generality, we assume that $f : E \to E$ is bijective, hence (x_t) is a stationary Markov chain. By stationarity, we can write (2.3.4) as

$$\begin{aligned} \phi(V_t^*) &= \alpha \; \mathbb{E} \left[\sum_{i=0}^{\infty} \gamma^i (\lambda_{t+i} \phi(V_{t+i}^*) + (1 - \lambda_{t+i}) \phi(U_{t+i}^{\mathrm{M}})) \middle| \, x_t \right] \\ &= \alpha \sum_{i=0}^{\infty} \gamma^i (v_i(x_t) + u_i(x_t)) \end{aligned} \tag{2.3.6}$$

with deterministic functions v_i and u_i:

$$v_i(e) = \mathbb{E}_e \left[\lambda_i \; \phi(V_i^*) \right] \quad \text{and} \quad u_i(e) = \mathbb{E}_e \left[(1 - \lambda_i) \; \phi(U_i^{\mathrm{M}}) \right] \quad \text{for } e \in E \text{ and } i \in \mathbb{N}_0.$$

We would like to express v_i in terms of $(u_j)_{j \in \mathbb{N}}$. We can also express v_j and u_j by the following iteration:

$$\begin{aligned} v_0(e) &= \mathbb{E} \left[\lambda_0 \; \phi(V_0^*) \middle| \, x_0 = e \right] = \lambda(s) \; \phi(V^*(e)) \\ v_j(e) &= \sum_{h \in E} \pi^j(e, h) v_0(h), \end{aligned} \tag{2.3.7}$$

and for u_j similarly

$$u_0(e) = \mathbb{E}_e\left[(1-\lambda_0)\ \phi(U_0^{\mathrm{M}})\right] = (1-\lambda(e))\phi(U^{\mathrm{M}}(e))$$

$$u_j(e) = \sum_{h\in E} \pi^j(e,h)u_0(h). \tag{2.3.8}$$

In matrix notation, $v_j = \pi^j v_0$ and $u_j = \pi^j u_0$. The condition (2.3.6) for $i = 0$ is sufficient, since it implies the conditions for $i \in \mathbb{N}$. In the vector notation, we obtain from (2.3.6) with the diagonal matrix $(\mathrm{diag}(\lambda))_{ee} = \lambda_e$:

$$v_0 = \alpha \sum_{j=0}^{\infty} \gamma^j \mathrm{diag}(\lambda)\left(\pi^j v_0 + \pi^j u_0\right) \tag{2.3.9}$$

If we find the function v_0, which satisfies the above condition, then the fixed point equation (2.3.4) becomes a linear equation for w with inhomogeneity $\phi^{-1}(\mathrm{diag}(1/\lambda)v_0)$. In fact,

$$V^*(w) = \phi^{-1}(\mathrm{diag}(1/\lambda)v_0) \tag{2.3.10}$$

is a linear equation with unknown w. In the steady state, $\phi(V^*(w)) = \mathrm{diag}(1/\lambda)v_0$ gives the dendritic prediction. Define the matrix $A = \sum_{j=0}^{\infty} \gamma^j \pi^j$, which is the expected, discounted, future matrix. Let us write $v = v_0$ and $u = u_0$. Then (2.3.9) becomes

$$v = \alpha\ \mathrm{diag}(\lambda)(Av + Au)$$

where λ, v and u are functions on E. This is equivalent to $(\mathbb{1} - \alpha\mathrm{diag}(\lambda)A)v = \alpha\mathrm{diag}(\lambda)Au$.

Thus, if $\alpha|\text{diag}(\lambda)A|_\infty < 1$, then

$$v = \sum_{j=0}^{\infty} (\alpha \ \text{diag}(\lambda)A)^j \ \alpha \ \text{diag}(\lambda)Au.$$

Therefore, in the steady state, we have

$$\phi(V^*(w)) = \sum_{j=0}^{\infty} \text{diag}(\lambda)^{-1}(\alpha \ \text{diag}(\lambda)A)^{j+1}u. \tag{2.3.11}$$

The condition $\alpha|\text{diag}(\lambda)A|_\infty < 1$ is satisfied if $\alpha \ \lambda_{\max} < 1-\gamma$, because $|\text{diag}(\lambda)A|_\infty \leq \lambda_{\max} \sum_{j=0}^{\infty} \gamma^j |\pi^j|_\infty = \frac{\lambda_{\max}}{1-\gamma}$. Therefore the solutions form a nonempty affine linear space W^*. Note that the linear equation (2.3.10) has a solution by the full-rank assumption on the matrix PSP. Also note that (2.3.11) is the analogue of (2.3.5) in the general case of not constant nudging factor λ.

To obtain the second part of the theorem, assume that λ is constant, and simplify (2.3.11) for $w^* \in W^*$

$$\phi(V^*(w^*)) = \sum_{j=0}^{\infty} \lambda^j (\alpha A)^{j+1} u = (1-\lambda) \sum_{j=0}^{\infty} \lambda^j (\alpha A)^{j+1} \phi(U^{\mathrm{M}}).$$

For $n \in \mathbb{N}_0$ multiple application of A gives

$$A^n \phi(U^{\mathrm{M}})(e) = \mathbb{E}_e \left[\sum_{j_1=0}^{\infty} \cdots \sum_{j_n=0}^{\infty} \gamma^{j_1+\cdots+j_n} \phi(U^{\mathrm{M}}_{j_1+\cdots+j_n}) \right].$$

Now, we collect all terms $\phi(U^{\mathrm{M}}_i)$ such that $A^n \phi(U^{\mathrm{M}})(e) = \mathbb{E}_e \left[\sum_{i=0}^{\infty} \alpha_{n,i} \gamma^i \phi(U^{\mathrm{M}}_i) \right]$, where

$\alpha_{n,i} = \binom{n+i-1}{i}$, the number of combinations with repetition. Thus, we obtain

$$\begin{aligned}\phi(V_t^*(w^*)) &= (1-\lambda)\sum_{j=0}^{\infty}\lambda^j\alpha^{j+1}\mathbb{E}\left[\sum_{i=0}^{\infty}\alpha_{j+1,i}\gamma^i\phi(U_{t+i}^{\mathrm{M}})\middle| x_t\right]\\ &= (1-\lambda)\mathbb{E}\left[\sum_{i=0}^{\infty}\tilde{\gamma}_i\phi(U_{t+i}^{\mathrm{M}})\middle| x_t\right],\end{aligned}$$

where $\tilde{\gamma}_i = \gamma^i\sum_{j=0}^{\infty}\lambda^j\alpha^{j+1}\alpha_{j+1,i} = \gamma^i\sum_{j=0}^{\infty}\lambda^j\alpha^{j+1}\binom{j+i}{i} = \alpha\frac{\gamma^i}{(1-\alpha\lambda)^{i+1}}$. This also shows that the series is almost surely absolutely convergent and that the summation rearrangements we made are valid. Therefore,

$$\phi(V_t^*(w^*)) = \alpha\frac{1-\lambda}{1-\alpha\lambda}\;\mathbb{E}\left[\sum_{i=0}^{\infty}\left(\frac{\gamma}{1-\alpha\lambda}\right)^i\phi(U_{t+i}^{\mathrm{M}})\middle| x_t\right].$$

□

2.3.2 Convergence of learning

Let us assume that ϕ is linear and $\alpha\ \lambda_{\max} < 1-\gamma$. In particular W^* is nonempty and affine linear. We define (Y_t) as we did before Lemma 2.2.3. Similarly to Theorem 2.2.4, we would like to prove that the sequence of weights (w_t), which is defined by the modified learning rule

$$\begin{aligned}w_{t+1} &= w_t + \eta_t\ \left(\phi(U(w_t))\ \widetilde{\mathrm{PSP}}_t - \phi(V_t^*(w_t))\ \mathrm{PSP}_t\right)\\ &= w_t + \eta_t\ H_2(w_t, Y_t)\quad \text{for } t\in\mathbb{N},\end{aligned}\tag{2.3.12}$$

converges to some $w^* \in W^*$ with probability one. First, we need to prove that for any starting point $w(0)$, the solution of the ODE

$$\dot{w}(t) = h_2(w(t)) \tag{2.3.13}$$

converges almost surely to some $w^* \in W^*$, where $h_2(w) = \mathbb{E}[H_2(w, Y_0)]$. The following lemma will be an analogue of Lemma 2.2.2. For notational simplicity, we omit the lower index 2 from the functions H_2 and h_2.

Lemma 2.3.2. *Assume that ϕ is linear. Then the ordinary differential equation (2.3.13) is globally asymptotically stable. The set of equilibrium points is W^*. For any $w^* \in W^*$, the domain of attraction is the orthogonal affine linear subspace to W^* in w^*, for which we write $D(w^*)$. This is to say that for all $w_0 \in D(w^*)$, the solution $w(t)$ of (2.3.13) with initial condition $w(0) = w_0$ satisfies $\lim_{t\to\infty} w(t) = w^*$. Furthermore, if $w(0) \in D(w^*)$ for some $w^* \in W^*$, then $w(t) \in D(w^*)$ holds for all $t \geq 0$.*

Proof. By the proof of Lemma 2.2.1, we have

$$h(w) = \mathbb{E}\left[\left(\mathbf{R}_{x_0}(w) - \phi(V_0^*(w))\right) \mathrm{PSP}_0\right],$$

thus, by the full rank property of PSP, the set of equilibrium points of (2.3.13) is W^*. By linearity of (2.3.4), the statement $w \in W^*$ can be characterized by $w^T\mathrm{PSP} = \beta$ for some fixed vector $\beta \in \mathbb{R}^m$, therefore $TW^* = \mathrm{Im}(\mathrm{PSP})^{\perp}$. Therefore $w(t) \in D(w^*)$, whenever $w(0) \in D(w^*)$. It remains to show that $D(w^*)$ is the domain of attraction of $w^* \in W^*$. For this purpose, we look for a suitable Lyapunov function. Let us define the obvious guess of a Lyapunov function $L(w) = \mathrm{dist}(w, W^*)^2/2$, which is zero exactly on W^*. The gradient is given by $\nabla L(w) = w - w^*$, where w^* denotes the orthogonal projection of w onto W^*. If we show that $h(w) \cdot \nabla L(w) \leq 0$, and equality holds if and only if $w \in W^*$, then the lemma follows from Lyapunov's global stability theorem ([45, § 13. VII. Theorem]). For $w, u \in \mathbb{R}^N$, we have

$$h(w + u) = h(w) + \mathbb{E}[(R_0^*(u) - \phi(V_0^*(u)))\ \mathrm{PSP}_0]$$

where $R_t^*(u) = \alpha \sum_{i=0}^{\infty} \gamma^i \phi(\lambda_{t+i} V_{t+i}^*(u))$. Let us define $\Delta w = w - w^*$ and compute

$$\begin{aligned}
& h(w) \cdot \nabla L(w) = \\
& h(w^* + \Delta w) \cdot \Delta w = \\
& h(w^*) \cdot \Delta w + \mathbb{E}\left[(R_0^*(\Delta w) - \phi(V_0^*(\Delta w)))\ \mathrm{PSP}_0\right] \cdot \Delta w = \\
& \mathbb{E}\left[(R_0^*(\Delta w) - \phi(V_0^*(\Delta w)))\ \mathrm{PSP}_0\right] \cdot \Delta w = \\
& \phi' \ \frac{g_D}{g_L + g_D}\ \mathbb{E}\left[\left(\alpha \sum_{t=0}^{\infty} \gamma^t \lambda_t \mathrm{PSP}_t \cdot \Delta w - \mathrm{PSP}_0 \cdot \Delta w\right) \mathrm{PSP}_0 \cdot \Delta w\right] \leq \\
& \phi' \ \frac{g_D}{g_L + g_D}\ \mathbb{E}\left[\alpha\ \lambda_{max} \sum_{t=0}^{\infty} \gamma^t\ |\mathrm{PSP}_t \cdot \Delta w|\ |\mathrm{PSP}_0 \cdot \Delta w| - (\mathrm{PSP}_0 \cdot \Delta w)^2\right].
\end{aligned}$$

By the Cauchy-Schwarz inequality applied to the minuend of the difference, we have

$$\begin{aligned}
& h(w) \cdot \nabla L(w) \leq \\
& \phi' \ \frac{g_D}{g_L + g_D} \left(\alpha\ \lambda_{max} \sum_{t=0}^{\infty} \gamma^t\ - 1\right)\ \mathbb{E}\left[(\mathrm{PSP}_0 \cdot \Delta w)^2\right].
\end{aligned}$$

If $\alpha\ \lambda_{max} \leq 1-\gamma$, then the upper bound is nonpositive, and since $TW^* = \mathrm{Im}(\mathrm{PSP})^\perp$, it is zero if and only if $\Delta w = 0$. Thus $h(w) \cdot \nabla L(w) \leq 0$, and equality holds if and only if $w \in W^*$. □

From [17, 1.9.2 Theorem 17] and the properties of Π from Section 2.2.2, we conclude almost sure convergence in the two-compartment model under the assumption of a linear rate function ϕ.

Theorem 2.3.3. *Assume that ϕ is linear. Let (η_t) be a sequence of learning rates such that $\sum_{t\in\mathbb{N}} \eta_t = \infty$ and $\sum_{t\in\mathbb{N}} \eta_t^2 < \infty$. Then the learning rule (2.3.12) converges with probability one for any initial parameters $w_0 \in D(w^*)$ and $Y_0 \in E \times \mathbb{R}^n$ to some $w^* \in W^*$. In particular, equations (2.3.4) and (2.3.5) hold in the limit $t \to \infty$, achieving the desired predictions*

$$\lim_{t\to\infty} |\phi(V_t^*(w_t)) - \alpha\ \mathbf{R}_{x_t}(w_t)| = 0 \quad \textit{almost surely.}$$

2.3.3 Time-continuous, spiking model

In order to approximate the time-continuous model, we introduce the time step $0 < \Delta t \leq 1$. Instead of the time independent transition matrix π, we introduce the time step dependent transition matrix $\tilde{\pi} = \Delta t \pi + (1-\Delta t)\mathrm{Id}_M$. In the limit $\Delta t \to 0$, we obtain a time-continuous Markov chain with generator $\pi - \mathrm{Id}_M$. Fix $\tau > 0$, we scale the parameters α, γ, η by replacing them according to $\alpha \to \frac{\alpha}{\tau}\Delta t$, $\gamma \to e^{-\Delta t/\tau}$, and $\eta \to \eta \Delta t$.

We define the postsynaptic potential process PSP_t as follows. Assume that the entries of Θ are nonnegative, thus the process ρ has nonnegative entries. Define a $\mathbb{R}^N$ valued Cox process $(C_t)_{t\in\mathbb{R}}$ (inhomogenous Poisson process with random rate process) with rate process $(\rho_t)_{t\in\mathbb{R}}$. Jumps of C_t correspond to spikes of the corresponding presynaptic neuron. Here, $\rho_t = \sum_{n=-\infty}^{\infty} \rho_n \mathbb{1}_{\Delta t[n,n+1)}(t)$ for $t \in \mathbb{R}$. With $\tau_l > \tau_s > 0$, we define *spike response kernel* $\kappa(t) = c_\kappa \theta(t)(e^{-t/\tau_l} - e^{-t/\tau_s})$, with constant $c_\kappa > 0$ and the Heaviside function θ. The post synaptic potential and the eligibility trace are given by:

$$\mathrm{PSP}_t = C * \kappa(t) \qquad \widetilde{\mathrm{PSP}}_t = \frac{\alpha}{\tau} \int_0^\infty e^{-\frac{s}{\tau}} \mathrm{PSP}_{t-s} \, \mathrm{d}s, \tag{2.3.14}$$

where we think of each entry C^i as a measure on $[0, \infty)$ placing a Dirac mass at the locations of spikes (jumps) of $(C_t^i)_{t\geq 0}$ for $i = 1, \ldots, N$. Thus PSP_t^i is the sum of response kernels, each started at one of the jumps of C_t^i. Again, we set $V_t(w) = w \cdot \mathrm{PSP}_t$. The update rule in the time-continuous case becomes:

$$\frac{\mathrm{d}w}{\mathrm{d}t}(t) = \eta \, \left[\phi(U_t) \, \widetilde{\mathrm{PSP}}_t - \phi(V_t^*(w_t)) \, \mathrm{PSP}_t \, \right] \quad \text{for } t \in \mathbb{R}, \tag{2.3.15}$$

with a learning rate $\eta > 0$. Taking the limit $\Delta t \to 0$, we see from Theorem 2.3.1 that the rule (2.3.15) achieves the predictions

$$\phi(V_t^*(w^*)) = \frac{\alpha}{\tau}(1-\lambda)\mathbb{E}\left[\int_0^\infty e^{-s/\tau_{\mathrm{eff}}} \phi(U_{t+s}^{\mathrm{M}}) \middle| \, x_t \right], \tag{2.3.16}$$

for constant λ and linear ϕ, where $\tau_{\mathrm{eff}} = \frac{\tau}{1-\alpha\lambda}$. For the sake of computer simulations, we will approximate the ODE (2.3.15) by the process

$$w_{t+\Delta t} = w_t + \Delta t\ \eta\ \Big[\phi(U_t)\ \widetilde{\mathrm{PSP}}_t - \phi(V_t^*(w_t))\ \mathrm{PSP}_t\ \Big] \quad \text{for } t \in \Delta t\ \mathbb{N}, \tag{2.3.17}$$

where $\widetilde{\mathrm{PSP}}_t = \frac{\alpha \Delta t}{\tau} \sum_{i=0}^{\infty} e^{-\frac{i\Delta t}{\tau}} \mathrm{PSP}_{t-i\Delta t}$.

2.4 Simulations

We have tested the learning rules (2.3.12) and (2.3.17) in computer simulations. We look at an experiment with five states of which only one is rewarding in the case of a deterministic Markov chain. Simulations are carried out using a linear and a sigmoid rate function with different discount factors γ. Then we turn our attention to the setting of an arbitrary transition matrix, ten states with one rewarding state and a linear rate function. In the cases where we use a biologically artificial linear rate function, we can compare the outcome with our theoretical results (2.3.5) and (2.3.11). We only perform experiments with constant nudging factor λ. In the spiking model with rule (2.3.17), we simulated an environment of 2 seconds where only the last 200 ms is rewarding and plotted simulated predictions against the theoretic values (2.3.5).

2.4.1 Two-compartment model without spiking

In this subsection we set the reversal potential for excitation $E^{\mathrm{E}} = 0$ mV and for inhibition $E^{\mathrm{I}} = -75$ mV, and define the rate functions

$$\phi_{\mathrm{sigmoid}}(U) = \frac{\phi_{max}}{1 + \exp(-\frac{U-\theta}{U_{73\%}-\theta})}, \qquad \phi_{\mathrm{linear}}(U) = \frac{60}{75}(U + 75),$$

where ϕ_{sigmoid} has a soft spiking threshold at $\theta = -40$ mV and reaches approximately

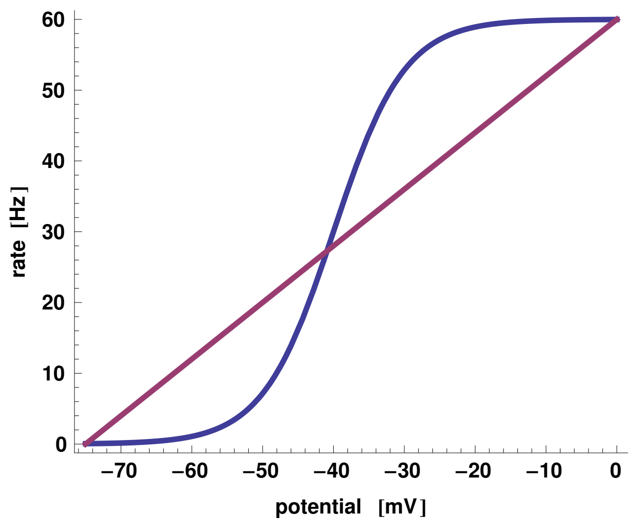

Figure 2.3: Linear and sigmoid rate functions ϕ

73% of its maximum $\phi_{max} = 60$ Hz at $U_{73\%} = -35$ mV (see Figure 2.3). Thus we have $\phi_{\mathrm{sigmoid}}(\theta) = 30$ Hz and $\phi_{\mathrm{sigmoid}}(U_{73\%}) \approx 43.86$ Hz. At the resting potential of -70 mV, we have spontaneous spiking at less than 0.15 Hz, and 4 Hz respectively. The choice of rate functions is arbitrary, we could have chosen any continuously differentiable, increasing function of sublinear growth.

We define a deterministic chain on five states $E = \{e_1, e_2, \cdots, e_5\}$. All states are nonrewarding except the last one e_5 (Figure 2.4, **Left**). In the rewarding state, we set $U^{\mathrm{M}} = 0$ mV, in other states $U^{\mathrm{M}} = -75$ mV. Note that the experiment only depends on U^{M}, but not the specific choice of the conductances g^{E}, g^{I} or the reversal potentials. We keep λ constant to demonstrate the validity of (2.3.5). Furthermore, we define $g_L = 0.1$ μS and $g_D = 2$ μS to get $V_t^*(w) = (2/2.1)\ V_t(w)$.

In the simulation, we compute the weight evolution of a single neuron which connects to 50 presynaptic neurons. We start the learning at w_0 which is normal with mean zero and standard deviation 5. We set the learning rate $\eta = 0.08$ and run the simulation for 500 000 transitions of the Markov chain. Much less rounds are sufficient for learning, but we want to show the accuracy of the theoretic finding. Note that in these simulations there is no randomness yet. After learning, we look at the dendritic predictions $\phi(V^*(e_j))$ for $j = 1, \cdots, 5$ (Figure 2.4, **Middle**).

After doing so for different values of $\gamma \in \{0.1, 0.4, 0.7\}$ and $\lambda \in \{0, 0.05, ..., 0.95\}$, we plot the dependence of an effective discount factor γ_{eff} on λ for $\gamma \in \{0.1, 0.4, 0.7\}$ (Figure 2.4, **Right**). In the case of a linear rate function, we know from equation 2.3.5 that $\gamma_{\mathrm{eff}} = \frac{\gamma}{1-\alpha\lambda}$.

We also performed simulations for random transitions. We generate a random transition matrix on 10 states, where only the last state is rewarding. We set the learning rate to $\eta = 1$ and perform 100 000 transitions of the Markov chain. Using the linear rate function and setting $\gamma = 0.4$ and $\lambda = 0.7$, we obtain Figure 2.5 of the theoretic (according to 2.3.5) and learned values of the dendritic predictions $\phi(V^*)$.

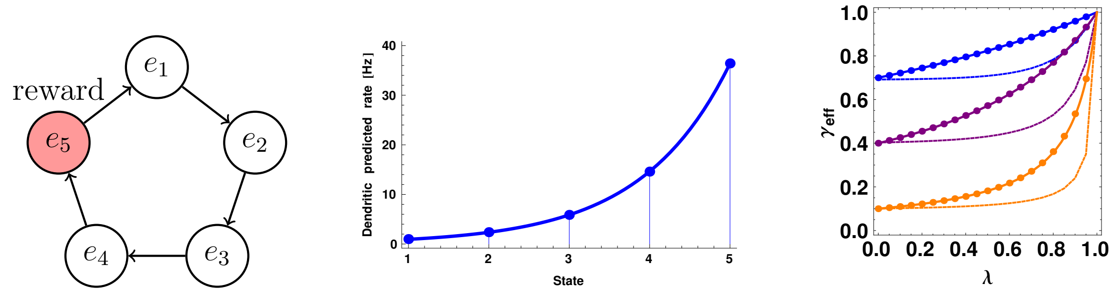

Figure 2.4: **The effective discount factor γ_{eff} depends through λ on the coupling conductances in the two-compartment model. Left:** We look at a deterministic Markov chain, with positive reward in state e_5 and zero reward otherwise (negligible reward for nonlinear ϕ). **Middle:** For nonlinear rate functions ϕ, the learned, future discounted values $\phi(V^*)$ are fitted with an exponential function to find γ_{eff}. **Right:** The effective discount factor γ_{eff} increases with λ for $\gamma \in \{0.1, 0.4, 0.7\}$; Continuous lines: relationship for linear rate (2.3.5); Dashed lines: linear interpolation for learned values in case of sigmoid rate.

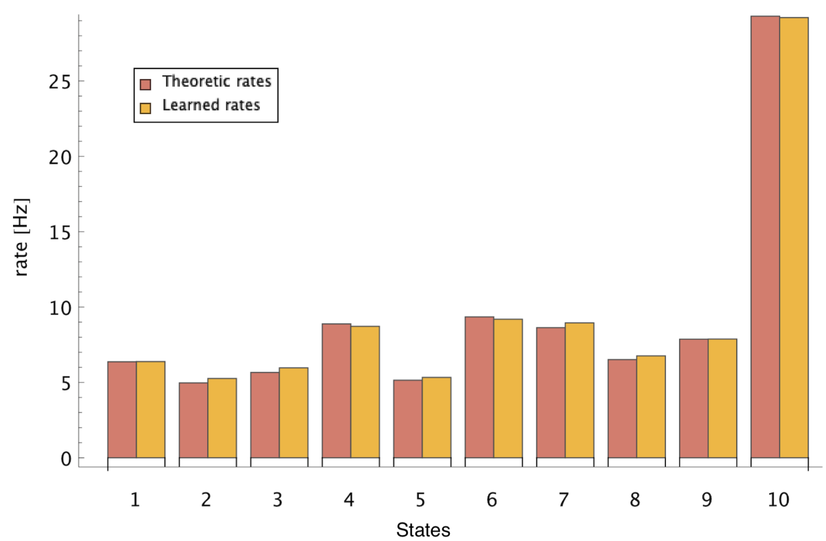

Figure 2.5: Theoretical and learned rates $\phi(V_t^*)$ for each of the 10 states for a nondeterministic Markov chain with state 10 being rewarding.

2.4.2 Two-compartment model with spiking

In another simulation, we assumed a two-compartment spiking model with presynaptic spiking. In this subsection, we use unitless potentials and resting potential 0 as described in Section 2.3 and the rate function $\phi(U) = U \times 0.03$ [kHz] for $U \geq 0$, and 0 otherwise. We introduced a time step $\Delta t = 0.1$ ms as described in Subsection 2.3.3. The rates ρ_t are discrete Ornstein Uhlenbeck processes with a length of 2 s independent for each neuron, which is kept constant for 2 ms then decays with a factor $e^{-1/50}$ after each of those 2 ms, reaching an effective decay of $1/e$ after 100 ms. The Gaussian noise standard deviation of the discrete Ornstein Uhlenbeck processes is 0.2 which is added after every 2 ms. Postsynaptic spikes are Poisson distributed with rates ρ_t and a refractory period of 1 ms. This choice makes patterns vary sufficiently continuously but still easily distinguishable after long enough time. The rates ρ_t are being kept constant for 2 ms each time after they are generated. This setting corresponds to 1000 states of a partially observable, deterministic, cyclic Markov chain, each state being presented for the duration of $20 = 2$ ms$/\Delta t$ steps to the dendritic compartment. The idea of introducing Ornstein Uhlenbeck processes into the simulations was developed by J. Brean and W. Senn in [18].

Reward is provided to the soma during the last 200 ms in each round. Initial weights are set to be zero to be able to present the bootstrapping effect graphically. The learning rate is set $\eta = 0.001$, which is multiplied by $\Delta t = 0.1$ as shown in (2.3.17). Other parameters are displayed below Figure 2.6. Note that we use theoretic values for comparison according to Theorem 2.3.1 instead of the continuous time limit (2.3.16), even though we replaced α, γ, η by the scaled parameters from Section 2.3.3. Results are plotted in Figure 2.6.

It is visible from Figure 2.6 that the dendritic potential is successively elevated after each round of 2 s of learning. The bootstrapping effect of the dendritic compartment predicting its own predictions results (after long enough learning) in an effective discount of time window 600 ms (meaning $\gamma_{\text{eff}} = e^{-\Delta t/600ms}$), while the plasticity window is less than 20 ms

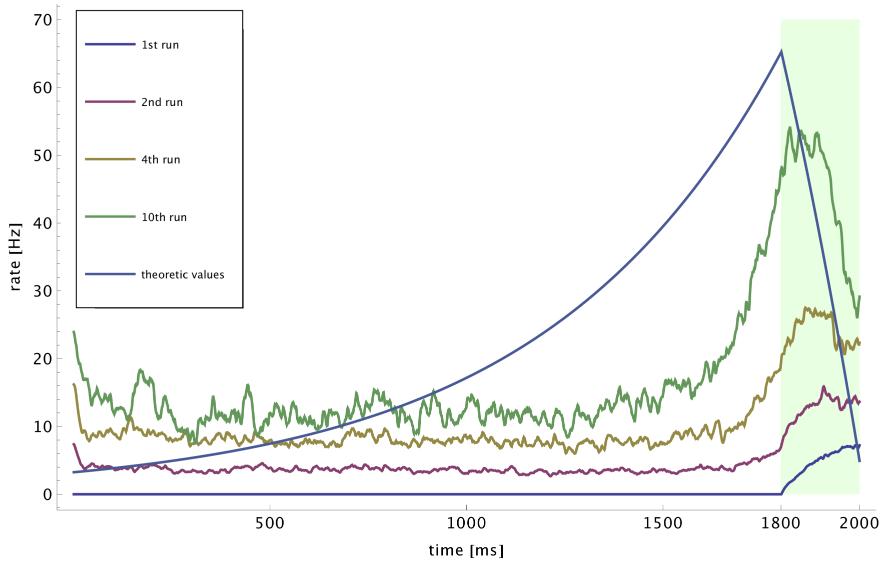

Figure 2.6: **A two-compartment neuron learns to represent the future, expected, discounted reward based on presynaptic spiking. The rate $\phi(V_t^*)$ is plotted against time for each of the training sessions.** The Poisson firing rate of 2000 presynaptic neurons depend on the "state of the environment", changing after every 20 ms. Reward is only delivered during the last 200 ms. Parameters: $\tau_{\text{eff}} = 600$ ms, $\gamma_{\text{eff}} = e^{-\Delta t/\tau_{\text{eff}}}, \tau \approx 8.17$ ms, $\gamma = e^{-\Delta t/\tau}$, $\lambda = .8$, $\alpha = 1.23$, $\Delta t = 0.1$, $\eta = 0.001$, $U_t^{\text{M}} \in \{0, 0.4\}$, $\tau_l = e^{-\Delta t/10 \text{ ms}}$, $\tau_s = e^{-\Delta t/(10/3) \text{ ms}}$, $c_\kappa = 5$, $V_t^*(w) = 0.96\ V_t(w)$.

for depression (window of PSP) and less than 35 ms for potentiation (window of $\widetilde{\text{PSP}}$) after a presynaptic spike (Figure 2.7). We can see that in the first run, the dendritic prediction is not elevated until the onset of the reward. During the first rewarding phase, weights with active synapses within 35 ms (plasticity window) before the onset of the reward are learning to predict the elevated predictions of the first run. We can see that in the second phase the baseline is elevated. The reason is that some neurons that are involved during and immediately before the rewarding phase are also spiking at different times during the entire 2 s of the training. In the second run, somatic prediction rate elevates even 35 ms before the onset of the actual reward. Our theory does not imply an elevation above the baseline more than 35 ms before the onset of the reward in the second run, though choosing the

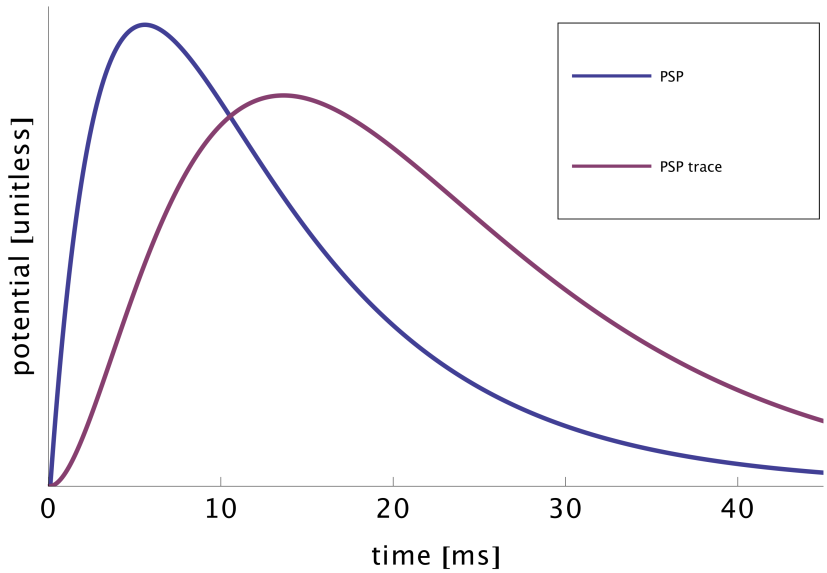

Figure 2.7: Postsynaptic potential (PSP) and its trace ($\widetilde{\text{PSP}}$) after a presynaptic spike

postsynaptic spiking rates to be Ornstein Uhlenbeck processes with a decay time window of 100 ms, we can expect elevation above the baseline even 100 ms before presenting the reward signal. First the elevation is learnt successively in every round, then during the 20th training session also depression of the baseline at the beginning of the training session is visible (Figure 2.8, **Upper**).

Running the experiment for 100 cycles with the same parameters as in Figure 2.6, we obtain the plot in Figure 2.8, **Lower**. The signal approximately fits the theoretic values after 100 cycles of learning. The dendritic compartment signals the expected future reward at 1200 ms with a predicted somatic firing rate of about 20 Hz already 600 ms prior to its onset. Since there is no reward or punishment signal yet at 1200 ms (meaning $U^{\mathrm{M}} = 0$), we have $\lambda = .8$, therefore $\phi(U_t) = \phi(\lambda_t V_t^*) = \lambda_t\ \phi(V_t^*) = .8 \times 20$ Hz $= 16$ Hz by equation (2.3.3), so the soma fires at 16 Hz.

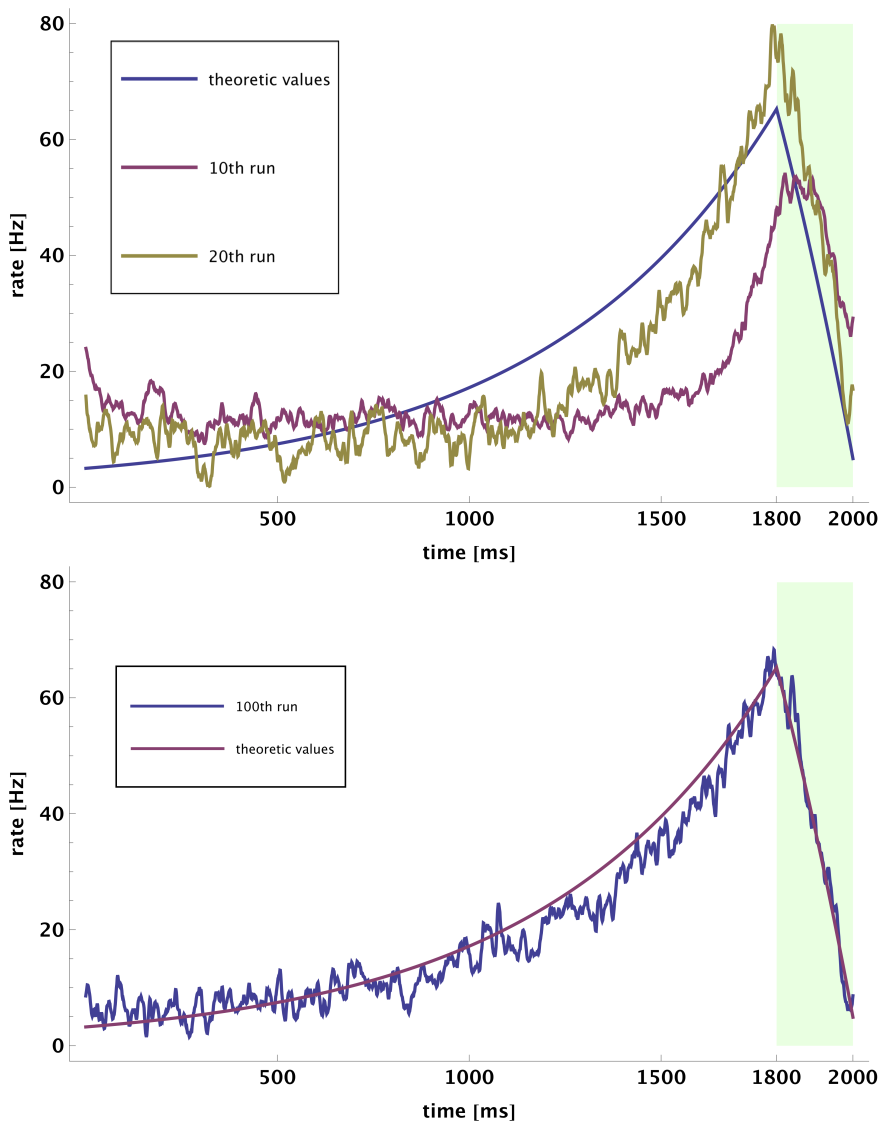

Figure 2.8: Long training sessions of a two-compartment neuron with presynaptic spiking. The rate $\phi(V_t^*)$ is plotted against time for each of the training sessions.

Chapter 3

Exit Time Asymptotics of a Switching Process

3.1 Introduction and Main result

In [37], [23], [20], [5], [6], [7], [1], [9], [8], [14], [13], [12], exit problems for processes in neighborhoods of unstable equilibria under the influence of white noise of small magnitude ϵ were studied. Among other results, it was obtained (under various additional sets of technical assumptions in [20], [5], [7], [1]) that if the origin is a hyperbolic critical point of a smooth vector field b, with simple leading eigenvalue $\lambda > 0$ of the linearization, then for any initial condition belonging to the stable manifold of the origin, the time τ_ϵ when the solution of the Itô SDE

$$\mathrm{d}X_\epsilon(t) = b(X_\epsilon(t))\mathrm{d}t + \epsilon\sigma(X_\epsilon(t))\mathrm{d}W(t), \tag{3.1.1}$$

with nondegenerate smooth diffusion matrix σ exits a small neighborhood U of the origin satisfies the following limit theorem: there are numbers $D_{\pm 1}$ such that

$$\left(X_\epsilon(\tau_\epsilon),\ \tau_\epsilon - \frac{1}{\lambda}\ln\frac{1}{\epsilon}\right) \xrightarrow{d} \left(q_B,\ D_B + \frac{1}{\lambda}\ln|N|\right), \quad \epsilon \to 0, \tag{3.1.2}$$

where $q_{\pm 1}$ are the points where the invariant curve associated with the eigenvalue λ intersects the boundary of U, B and N are independent random variables, B is 1/2-Bernoulli, N is standard Gaussian.

It is clear that such an asymptotic result for solutions of an SDE must describe the asymptotic behavior for a whole class of systems that are well approximated by this kind of SDE or its exemplar one-dimensional additive noise linear version

$$\mathrm{d}X_\epsilon(t) = \lambda X_\epsilon(t)\mathrm{d}t + \epsilon \mathrm{d}W.$$

In fact, it was shown in [8] that the exit times for Glauber dynamics for the Curie–Weiss mean field model belong to the universality class associated with (3.1.2), i.e., they satisfy a similar limit theorem under the infinite system size limit, with limiting distribution being $\ln|N|$ up to scaling and translation. These results were used in [9] in the context of decision/reaction times in psychology.

It is interesting to explore this universality class further and study processes of a totally different nature that exhibit similar behavior. In this dissertation, we study a family of processes with random switching, also known under the names of hybrid systems, piecewise deterministic Markov processes (PDMP), and random evolutions. The bibliography on these processes is growing. Interestingly, they were introduced and rediscovered many times by different groups of researchers. Here we give just a few references to works of some of these groups: [36], [19], [33], [2], [25], [47], [16], [11], [38], [41].

In general, these processes are defined by a family of vector fields and a collection of rates of switching between those vector fields. At each time the system is in the state where it evolves along one of the vector fields from the family. At random times, the system jumps between states switching active vector fields from one to another according to the prescribed Markovian rates.

In the limit of infinite switching rates, the evolution can be effectively described by the law of large numbers through the averaging of the vector fields involved. One can also state a central limit theorem and, moreover, a functional central limit theorem for such systems on a finite time interval, see, e.g., [2, Chapter 4].

More interesting questions involve the behavior of such systems over unbounded time intervals. In this dissertation, we study a class of switching processes on time scales logarithmic in the switching rate μ and show that it belongs to the universality class associated with (3.1.2). We consider processes driven intermittently by two vector fields in one dimension. The main condition on these vector fields is that their average defines an unstable critical point. The exit from a neighborhood of that unstable equilibrium takes a logarithmically long time in the switching rate μ. Thus, similarly to the situation in [8], the usual techniques of weak convergence on a finite time horizon are not sufficient for obtaining the desired universality result and have to be supplemented with additional arguments. This is joint work with Yuri Bakhtin and was published in [10].

Let us now describe the system we are interested in and our main result more formally. Let $f_1, f_{-1} \in \mathcal{C}^2(\mathbb{R})$ with $a_1 = f_1'(0)$ and $a_{-1} = f_{-1}'(0)$. We also define

$$F(x) = \frac{1}{2}(f_1(x) + f_{-1}(x)), \quad x \in \mathbb{R}, \tag{3.1.3}$$

and require that $a = F'(0) = (a_1 + a_{-1})/2 > 0$. Furthermore, we require that $|f_1(0)| = |f_{-1}(0)| > 0$ and $\mathrm{sgn}(F(x)) = \mathrm{sgn}(x)$, where $\mathrm{sgn}(z) = z/|z|$ for $z \neq 0$ and $\mathrm{sgn}(0) = 0$. We will study the dynamics driven by these functions on a finite segment $[-R, R]$ for some $R > 0$, so without loss of generality we will assume that the functions $f_{\pm 1}$ and their first two derivatives are bounded.

An example of such pair of functions is given by $f_1(x) = e^{ax}$, $f_{-1}(x) = -e^{-ax}$ for $x \in [-R, R]$ and $a > 0$.

Let $(\sigma_t^\mu)_{t\geq 0}$ be a homogeneous, rate $\mu > 0$, right continuous Markov process on the state space $\{+1, -1\}$ with an arbitrary initial distribution on $\{+1, -1\}$ at time 0. The realizations of this process almost surely make finitely many switches between 1 and -1 on any finite time interval, so omitting the exceptional set, we can work on a probability space $(\Omega, \mathcal{F}, P)$ that guarantees that the number of switches is locally finite for *all* realizations.

We will study the random trajectories $(x_t^\mu)_{t\geq 0}$ defined by

$$\frac{\mathrm{d}x_t^\mu}{\mathrm{d}t} = f_{\sigma_t^\mu}(x_t^\mu), \quad t \geq 0,$$
$$x_0^\mu = 0.$$

The paths of the stochastic process $(x^\mu_t)_{t\geq 0}$ are continuous. They switch between the dynamics governed by f_1 and f_{-1} intermittently, being controlled by the Markov chain $(\sigma^\mu_t)_{t\geq 0}$. For any $r > 0$, we define the exit time from the interval $[-r, r]$:

$$\tau^\mu(r) = \inf\{t : |x^\mu_t| \geq r\}.$$

Our main result describes the joint asymptotic behavior of the exit time $\tau^\mu(r)$ and exit location $x^\mu_{\tau^\mu(r)}$ as $\mu \to \infty$. The sign of $x^\mu_{\tau^\mu(r)}$ can be interpreted as a decision between two alternative directions of exit made by the system by the time $\tau^\mu(r)$. Let us define

$$K(r) = \int_0^r \left(\frac{1}{F(x)} - \frac{1}{ax}\right) \mathrm{d}x, \quad r \neq 0. \tag{3.1.4}$$

Let us also define $D(0)$ to be arbitrary and

$$D(r) = K(r) + \frac{\log|r|}{a} + \frac{\log(\sqrt{2a}/|f_1(0)|)}{a}, \quad r \neq 0.$$

Theorem 3.1.1. *For any $r > 0$, as $\mu \to \infty$,*

$$\left(x^\mu_{\tau^\mu(r)},\ \tau^\mu(r) - \frac{1}{2a}\log\mu\right) \xrightarrow{d} \left(r \cdot \operatorname{sgn} N,\ -\frac{1}{a}\log|N| + D(r\cdot \operatorname{sgn} N)\right),$$

where N is a standard Gaussian random variable.

Our strategy for the proof of Theorem 3.1.1 consists of studying the asymptotic exit time of $(x^\mu_t)_{t\geq 0}$ from the interval $[-\mu^\gamma, \mu^\gamma]$ for $\gamma \in (1/4, 1/2)$, then the remaining time to hit the boundary $\{r, -r\}$. In Lemma 3.2.4, we show that up to an additive constant depending on a and $|f_1(0)|$ (the last term in the definition of D), the time needed to exit $[-\mu^\gamma, \mu^\gamma]$ is $-\frac{1}{a}\log|N| + \frac{1/2-\gamma}{a}\log\mu$, after which the process becomes deterministic in the limit $\mu \to \infty$, and is driven by F, see Lemma 3.2.6. By Lemma 3.2.5, the process driven by F and started from $\{\mu^\gamma, -\mu^\gamma\}$ requires $K(r \cdot \operatorname{sgn} N) + \frac{\log|r|}{a} + \frac{\gamma}{a}\log\mu$ time to hit $\{r, -r\}$, depending on the sign of exit direction $\operatorname{sgn} N$. Summing up the two contributions to $\tau^\mu(r)$, the terms $\pm\frac{\gamma}{a}\log\mu$ cancel, and we obtain Theorem 3.1.1.

3.2 Proof

For brevity, throughout this section, we will often omit μ in the notation and use $\sigma_t = \sigma_t^\mu$, $x_t = x_t^\mu$, $\tau(r) = \tau^\mu(r)$, etc.

By Taylor's theorem, there are functions $R_1, R_{-1} : \mathbb{R} \to \mathbb{R}$ such that

$$f_\sigma(x) = \sigma f_1(0) + a_\sigma x + R_\sigma(x), \quad x \in \mathbb{R},\ \sigma \in \{1, -1\}, \tag{3.2.1}$$

and $|R_\sigma(x)| \le c\, 2^{-1}x^2$, where

$$c = \max_{\sigma\in\{-1,1\}} \sup_{x\in\mathbb{R}} |f''_\sigma(x)|. \tag{3.2.2}$$

The generator of the Markov process (σ_t, x_t) on any bounded, smooth function $g : \{-1, 1\} \times \mathbb{R} \to \mathbb{R}$ is given by

$$\mathcal{L}g(\sigma, x) = f_\sigma(x)\partial_x g(\sigma, x) + \mu(g(-\sigma, x) - g(\sigma, x)). \tag{3.2.3}$$

Applying (3.2.3) to the functions $g(\sigma,x)=\sigma$ and $\tilde{g}(\sigma,x)=\sigma x$ we obtain by Proposition 1.7 in [24, Chapter 4] that the processes

$$\begin{aligned} Z_t &= Z_t^{\mu} = \sigma_t + 2\mu\int_0^t \sigma_s \mathrm{d}s, \\ \tilde{Z}_t &= \tilde{Z}_t^{\mu} = \sigma_t x_t - \int_0^t f_{\sigma_s}(x_s)\sigma_s \mathrm{d}s + 2\mu\int_0^t \sigma_s x_s \mathrm{d}s \end{aligned} \tag{3.2.4}$$

are local martingales with quadratic variations $[Z]_t = 4B(t)$, where $B(t)=|\{s\in[0,t]:\sigma_s\neq\sigma_{-s}\}|$ denotes the number of jumps of σ up to time $t\geq 0$, and

$$[\tilde{Z}]_t = 4\sum_{s\in[0,t]:\ \sigma_s\neq\sigma_{-s}} |x_s|^2.$$

Moreover, the true martingale property also follows since there is a constant $C>0$ such that $|x_t|\leq Ct$ for all $t>0$ due to our assumptions on $f_{\pm 1}$. Note that $(B(t))_{t\geq 0}$ is a rate μ Poisson process.

Integration of (3.2.1) and substitution of (3.2.4) gives for any $t\in[0,\infty)$

$$\begin{aligned} x_t &= \int_0^t f_{\sigma_s}(x_s)\mathrm{d}s = \int_0^t \left(\sigma_s f_1(0) + a x_s + (a_{\sigma_s}-a)x_s + R_{\sigma_s}(x_s)\right)\mathrm{d}s \\ &= a\int_0^t x_s \mathrm{d}s + \frac{f_1(0)}{2\mu}(Z_t-\sigma_t) + \int_0^t (a_{\sigma_s}-a)x_s \mathrm{d}s + \int_0^t R_{\sigma_s}(x_s)\mathrm{d}s. \end{aligned}$$

Noting that $a_{\sigma_s}-a=(\Delta a/2)\sigma_s$ with $\Delta a = a_1 - a_{-1}$, we can use (3.2.4) to rewrite this as

$$x_t = a\int_0^t x_s \mathrm{d}s + \frac{f_1(0)}{2\mu}(Z_t-\sigma_t) + \frac{\Delta a}{4\mu}\int_0^t f_{\sigma_s}(x_s)\sigma_s \mathrm{d}s + \frac{\Delta a}{4\mu}(\tilde{Z}_t - \sigma_t x_t) + \int_0^t R_{\sigma_s}(x_s)\mathrm{d}s.$$

The variation of constants formula gives for any $t \in [0, \infty)$

$$x_t = e^{at}\left(\frac{f_1(0)}{2\mu}\int_0^t e^{-as}\mathrm{d}(Z-\sigma)_s + \frac{\Delta a}{4\mu}\int_0^t e^{-as} f_{\sigma_s}(x_s)\sigma_s \mathrm{d}s \right.$$
$$\left. + \frac{\Delta a}{4\mu}\int_0^t e^{-as}\mathrm{d}(\tilde{Z}-\sigma x)_s + \int_0^t e^{-as} R_{\sigma_s}(x_s)\mathrm{d}s\right). \quad (3.2.5)$$

Lemma 3.2.1. *Suppose there is a sequence of stopping times θ^μ with respect to the natural filtration of σ satisfying*

$$\theta^\mu \xrightarrow{P} \infty, \qquad \text{as } \mu \to \infty.$$

Then as $\mu \to \infty$, the random variable

$$I^\mu = \frac{1}{\sqrt{\mu}}\int_0^{\theta^\mu} e^{-as}\mathrm{d}(Z^\mu - \sigma)_s \quad (3.2.6)$$

converges in distribution to $\mathcal{N}(0, 2a^{-1})$.

Proof. By the alternating series test, $\int_0^{\theta^\mu} e^{-as}\mathrm{d}\sigma_s$ exists and belongs to the interval $(-2, 2)$. Consequently, $\mu^{-1/2}\int_0^{\theta^\mu} e^{-as}\mathrm{d}\sigma_s \to 0$ as $\mu \to \infty$. Therefore, we only need to study convergence in distribution of the part with the martingale integrator.

For the rest of the proof, we follow the ideas in the proof of [8, Lemmas 3.1, 3.2]. First, we define a martingale

$$V_t = V_t^\mu = \int_0^{t\wedge\theta^\mu} e^{-as}\mathrm{d}Z_s \quad (3.2.7)$$

with quadratic variation $[V]_t = \int_0^{t\wedge\theta^\mu} e^{-2as}\mathrm{d}[Z]_s$. Then we define a time-changed martingale

$U_t = U_t^\mu = V_{g(t)}$ for $t \in [0, 2a^{-1}]$, where

$$g(s) = -\frac{\log(1 - as/2)}{2a}, \qquad s \in [0, 2a^{-1}),$$

and $g(2a^{-1}) = \infty$. We will prove that for any $t \in [0, 2a^{-1}]$, the quadratic variation of U satisfies

$$\mu^{-1}[U]_t \xrightarrow{P} t, \quad \mu \to \infty. \tag{3.2.8}$$

Then, by [8, Theorem 3.1], which is just a specific case of Theorem 1.4 in [24, Chapter 7], $\mu^{-1/2} U_t$ converges to $\mathcal{N}(0, t)$ in distribution for any $t \in [0, 2a^{-1}]$ as $\mu \to \infty$. Therefore, $V_\infty^\mu = U_{2a^{-1}}^\mu \xrightarrow{d} \mathcal{N}(0, 2a^{-1})$ as $\mu \to \infty$.

It remains to prove (3.2.8). For all $t \in [0, 2a^{-1}]$,

$$[U]_t = \sum_{\substack{s: s \le g(t) \wedge \theta^\mu \\ \sigma(s) \ne \sigma(s-)}} H(s) \tag{3.2.9}$$

for $H(s) = 4e^{-2s}$. We claim that for any nonincreasing function $H(\cdot)$,

$$\mu^{-1} \sum_{\substack{s: s \le g(t) \wedge \theta^\mu \\ \sigma(s) \ne \sigma(s-)}} H(s) \xrightarrow{P} \int_0^{g(t)} H(s) \mathrm{d}s, \quad \mu \to \infty. \tag{3.2.10}$$

To prove this relation, we first note that it holds for $H(s) = \mathbb{1}_{[0,h]}(s)$, for any $h > 0$, since the Law of Large Numbers implies

$$\mu^{-1} \sum_{\substack{s: s \le g(t) \wedge \theta^\mu \\ \sigma(s) \neq \sigma(s-)}} \mathbb{1}_{[0,h]}(s) = \mu^{-1} B(g(t) \wedge h \wedge \theta^\mu) \xrightarrow{P} g(t) \wedge h = \int_0^{g(t)} \mathbb{1}_{[0,h]}(s) \mathrm{d}s.$$

Using this and approximating monotone functions with sums of indicator functions, we obtain (3.2.10) which, combined with (3.2.9), gives (3.2.8) and completes the proof of the lemma. □

For any $\gamma > 0$ we can define $\theta^\mu = \inf\{t : |x_t| \ge \mu^{-\gamma}\}$. Note that for μ large enough to ensure $\mu^{-\gamma} \le r$, we have $\theta^\mu \le \tau^\mu(r)$.

Lemma 3.2.2. *The random variables*

$$J^\mu = \int_0^{\theta^\mu} e^{-as} \mathrm{d}(\tilde{Z} - \sigma x)_s \tag{3.2.11}$$

satisfy

$$P(|J^\mu| > \mu^{-\delta + 1/2}) \to 0 \tag{3.2.12}$$

for any $\delta < \gamma$ as $\mu \to \infty$. Consequently, the sequence $(\mu^{-1/2} J^\mu)$ converges to zero in probability as $\mu \to \infty$.

Proof. For any $t \ge 0$, we use integration by parts to write

$$\int_0^t e^{-as} \mathrm{d}(\sigma x)_s = e^{-at} \sigma_t x_t + a \int_0^t e^{-as} \sigma_s x_s \mathrm{d}s,$$

$$\left|\int_0^{\theta^\mu} e^{-as}\mathrm{d}(\sigma x)_s\right| \le e^{-a\theta^\mu}|\sigma_{\theta^\mu}x_{\theta^\mu}| + a\int_0^\infty e^{-as}|\sigma_s x_s|\mathrm{d}s \le 2\mu^{-\gamma}.$$

For the other term of J including the martingale integrator, we can apply Chebyshev inequality followed by Proposition 6.1 in [24, Chapter 2]

$$P\left(\left|\int_0^{\theta^\mu} e^{-as}\mathrm{d}\tilde{Z}_s\right| > \mu^{-\delta+1/2}\right) \le \mu^{2\delta-1}\mathbb{E}\left[\left(\int_0^{\theta^\mu} e^{-as}\mathrm{d}\tilde{Z}_s\right)^2\right] = \mu^{2\delta-1}\mathbb{E}\left[\int_0^{\theta^\mu} e^{-2as}\mathrm{d}[\tilde{Z}]_s\right].$$

Since the quadratic variation $[\tilde{Z}]_{s\wedge\theta^\mu}$ is stochastically dominated by $4\mu^{-2\gamma}B(s)$, by the Campbell formula in the second step

$$\mathbb{E}\left[\int_0^{\theta^\mu} e^{-2as}\mathrm{d}[\tilde{Z}]_s\right] \le \mathbb{E}\left[\int_0^\infty e^{-2as}4\mu^{-2\gamma}\mathrm{d}B(s)\right] = \frac{4\mu^{1-2\gamma}}{2a},$$

which multiplied by $\mu^{2\delta-1}$ converges to zero. □

Lemma 3.2.3. *If $0 < \gamma < \frac{1}{2}$, then $\theta^\mu \xrightarrow{P} \infty$ as $\mu \to \infty$.*

Proof. We need to check that for an arbitrary $T > 0$, $P(\theta^\mu < T) \to 0$ as $\mu \to \infty$. Recalling the definition of V in (3.2.7) and estimating

$$\mathbb{E}[V]_{\theta^\mu} \le \frac{2\mu}{a}$$

by Campbell's formula, then using the Chebyshev inequality and the martingale property of $V^2 - [V]$ (see Proposition 6.1 in [24, Chapter 2]), we obtain

$$P\left(\frac{1}{2\mu}|V_{\theta^\mu}| > \mu^{-\gamma/2-1/4}\right) \le \frac{\mathbb{E}[V]_{\theta^\mu}}{4\mu^2\mu^{-\gamma-1/2}} \le \frac{2\mu a^{-1}}{4\mu^2\mu^{-\gamma-1/2}} \to 0 \qquad (3.2.13)$$

as $\gamma < \frac{1}{2}$. By (3.2.12), we also have

$$P\left(\frac{1}{\mu}|J^\mu| > \mu^{-\gamma/2-1/4}\right) \to 0. \tag{3.2.14}$$

On the set $\{\theta^\mu < T\}$, we can use (3.2.5) and (3.2.2) to see that if μ is large enough to guarantee $\sup_{\sigma\in\{-1,1\},x\in[-\mu^{-\gamma},\mu^{-\gamma}]}|f_\sigma(x)| \le 2|f_1(0)|$, then

$$\begin{aligned}\mu^{-\gamma} = |x_{\theta^\mu}| = e^{a\theta^\mu}\bigg|\frac{f_1(0)}{2\mu}\int_0^{\theta^\mu} e^{-as}\mathrm{d}(Z-\sigma)_s + \int_0^{\theta^\mu} e^{-as}R_{\sigma_s}(x_s)\mathrm{d}s \\ + \frac{\Delta a}{4\mu}\int_0^{\theta^\mu} e^{-as}f_{\sigma_s}(x_s)\sigma_s\mathrm{d}s + \frac{\Delta a}{4\mu}J^\mu\bigg| \\ \le e^{aT}\left(\frac{|f_1(0)|}{2\mu}|V_{\theta^\mu}| + \frac{|f_1(0)|}{\mu} + \frac{c\mu^{-2\gamma}}{2a} + \frac{|f_1(0)\Delta a|}{2\mu a} + \frac{|\Delta a|}{4\mu}|J^\mu|\right).\end{aligned}$$

Therefore, on the set $\{\theta^\mu < T\} \cap \left\{\frac{1}{2\mu}|V^\mu_{\theta^\mu}| \le \mu^{-\gamma/2-1/4}, \frac{1}{\mu}|J^\mu| \le \mu^{-\gamma/2-1/4}\right\}$,

$$\mu^{-\gamma} \le e^{aT}\left(|f_1(0)|\mu^{-\gamma/2-1/4} + \frac{|f_1(0)|}{\mu} + \frac{c\mu^{-2\gamma}}{2a} + \frac{|f_1(0)\Delta a|}{2\mu a} + \frac{|\Delta a|}{4}\mu^{-\gamma/2-1/4}\right)$$

that is impossible for large μ and $\gamma < \frac{1}{2}$. Combining this with (3.2.13) and (3.2.14), we complete the proof. □

Lemma 3.2.4. *If $\frac{1}{4} < \gamma < \frac{1}{2}$, then*

$$\left(\mathrm{sgn}(x_{\theta^\mu}),\ \theta^\mu - \frac{1/2-\gamma}{a}\log\mu\right) \xrightarrow{d} \left(\mathrm{sgn}(H),\ -\frac{1}{a}\log|H|\right), \quad \mu\to\infty,$$

where H is a random variable with $\mathrm{Law}(H) = \mathcal{N}(0, f_1(0)^2(2a)^{-1})$.

Proof. Using (3.2.5) at time θ^μ, and introducing

$$H^\mu = \frac{f_1(0)}{2} I^\mu + \sqrt{\mu} \int_0^{\theta^\mu} e^{-as} R_{\sigma_s}(x_s) \mathrm{d}s + \frac{\Delta a}{4\sqrt{\mu}} \int_0^{\theta^\mu} e^{-as} f_{\sigma_s}(x_s) \sigma_s \mathrm{d}s + \frac{\Delta a}{4\sqrt{\mu}} J^\mu, \quad (3.2.15)$$

where I^μ was defined in (3.2.6) and J^μ in (3.2.11), we obtain

$$\theta^\mu = \frac{1/2 - \gamma}{a} \log \mu - \frac{1}{a} \log |H^\mu| \quad (3.2.16)$$

and

$$\operatorname{sgn}(x_{\theta^\mu}) = \operatorname{sgn} H^\mu. \quad (3.2.17)$$

The first term on the right-hand side of (3.2.15) converges in distribution to H by Lemma 3.2.1. The second term converges to 0 almost surely due to

$$\sqrt{\mu} \left| \int_0^{\theta^\mu} e^{-as} R_{\sigma_s}(x_s) \mathrm{d}s \right| \leq \frac{\sqrt{\mu}}{a} c \; 2^{-1} \mu^{-2\gamma} \to 0.$$

Since $f_{\sigma_s}(x_s)$ and σ_s are bounded, the third term on the right-hand side of (3.2.15) converges to 0 almost surely, and the fourth term converges to zero in probability by Lemma 3.2.2. Since the distribution of H has no atom at 0, the only point of discontinuity of functions $x \mapsto \log |x|$ and $x \mapsto \operatorname{sgn} x$, the lemma follows now from (3.2.16) and (3.2.17). □

Now, we consider the exit time from the fixed interval $[-r, r]$. We define $\eta_t = x_{\theta^\mu + t}$,

$\sigma'_t = \sigma_{\theta^\mu + t}$ and the exit time

$$\nu(r) = \inf\{t \geq 0 : |\eta_t| \geq r\}.$$

Recalling (3.1.3), we define $(S^t)_{t\geq 0}$ as the flow associated with the ODE

$$\dot{z} = F(z).$$

Note that (S^t) preserves the sign of the initial condition. For $\delta \neq 0$ we introduce $t(\delta, r)$ to be the only solution of $|S^t\delta| = r$.

Lemma 3.2.5. *For any $r \neq 0$,*

$$\lim_{\delta\to 0}\left(t(\delta, r) - \frac{1}{a}\log\frac{r}{\delta}\right) = K(r),$$

where δ approaches zero from the right if $r > 0$, and from the left if $r < 0$, and $K(r)$ is given in (3.1.4)*.*

Proof. By separation of variables,

$$\begin{aligned} t(\delta, r) = \int_\delta^r \frac{\mathrm{d}x}{F(x)} &= \int_\delta^r \left(\frac{1}{F(x)} - \frac{1}{ax}\right)\mathrm{d}x + \int_\delta^r \frac{1}{ax}\mathrm{d}x \\ &= \int_\delta^r \left(\frac{1}{F(x)} - \frac{1}{ax}\right)\mathrm{d}x + \frac{1}{a}\log\frac{r}{\delta}. \end{aligned}$$

Letting $\delta \to 0$ and using $F'(0) = a$, we complete the proof. □

Lemma 3.2.6. *There is $r_0 > 0$ such that for any $r \in (0, r_0)$*

$$\sup_{0\le t\le t(\eta_0, r\cdot \mathrm{sgn}(\eta_0))} |\eta_t - S^t\eta_0| \xrightarrow{P} 0 \quad \text{as } \mu \to \infty.$$

Proof. Choosing $g(\sigma, x) = f_{-\sigma}(x)$, we obtain for $|x| < R$

$$\mathcal{L}g(\sigma, x) = f_\sigma(x) f'_{-\sigma}(x) + 2\mu\sigma G(x),$$

where $G(x) = \frac{1}{2}(f_1(x) - f_{-1}(x)) > 0$. By Proposition 1.7 in [24, Chapter 4], the process $Z'(t) = -g(\sigma'_t, \eta_t) + \int_0^t \mathcal{L}g(\sigma'_s, \eta_s)\mathrm{d}s$ is a martingale.

We will only prove

$$\left(\sup_{0\le t\le t(\eta_0, r\cdot \mathrm{sgn}(\eta_0))} |\eta_t - S^t\eta_0|\right) \mathbb{1}_{\{\mathrm{sgn}(\eta_0)=\alpha\}} \xrightarrow{P} 0, \quad \mu \to 0, \tag{3.2.18}$$

for $\alpha = 1$. The case of $\alpha = -1$ is treated similarly.

The rest of the proof follows closely the proof of [8, Lemma 3.6]. We define $\Delta(t) = \eta_t - S^t\eta_0$. Since for $t \ge 0$,

$$\eta_t = \eta_0 + \int_0^t F(\eta_s)\mathrm{d}s + \int_0^t \sigma'_s G(\eta_s)\mathrm{d}s,$$
$$S^t\eta_0 = \eta_0 + \int_0^t F(S^t\eta_0)\mathrm{d}s,$$

we have

$$\Delta(t) = \int_0^t (F(\eta_s) - F(S^t\eta_0))\mathrm{d}s + \int_0^t \sigma'_s G(\eta_s)\mathrm{d}s.$$

Using the Lipschitz constant $L(r)$ of F on $[-r, r]$, we obtain on the event $\{\mathrm{sgn}(\eta_0) = 1\}$ for any $t \le t(\eta_0, r) = t(\mu^{-\gamma}, r)$

$$|\Delta(t \wedge \nu(r))| \le L(r) \int_0^{t\wedge\nu(r)} |\Delta(s)|\mathrm{d}s + \left| \int_0^{t\wedge\nu(r)} \sigma'_s G(\eta_s)\mathrm{d}s \right|. \tag{3.2.19}$$

Note that by the definition of Z', we have $2\mu \int_0^t \sigma'_s G(\eta_s)\mathrm{d}s = Z'(t) + f_{-\sigma'_t}(\eta_t) - \int_0^t f_{\sigma'_s}(\eta_s) f'_{-\sigma'_s}(\eta_s)\mathrm{d}s$. Defining $A_1(r) = \sup_{\sigma\in\{1,-1\}, z\in[-r,r]} |f_\sigma(z)|$ and $A_2(r) = \sup_{\sigma\in\{1,-1\}, z\in[-r,r]} |f_\sigma(z) f'_{-\sigma}(z)|$, we can bound

$$2\mu \left| \int_0^{t\wedge\nu(r)} \sigma'_s G(\eta_s)\mathrm{d}s \right| \le \sup_{s \le t(|\eta_0|, r)\wedge\nu(r)} |Z'(s)| + A_1(r) + A_2(r)(t(\eta_0, r) \wedge \nu(r))$$

for any $t \in [0, t(\eta_0, r)]$ on the event $\{\mathrm{sgn}(\eta_0) = 1\}$. We claim that if $\delta < 1/2$, then

$$P\left((2\mu)^{-1} \sup_{s\le t(\eta_0, r)\wedge\nu(r)} |Z'(s)| > \mu^{-\delta}, \mathrm{sgn}(\eta_0) = 1 \right) \to 0, \quad \mu \to \infty.$$

The quadratic variation process $[Z']_t \le 4A_1^2(r)(B(t+\tau) - B(\tau))$ is stochastically dominated by $4A_1^2(r)$ times a rate μ Poisson process. Using the Chebyshev inequality followed by

Burkholder–Davis–Gundy inequality [21, Thm. 92, Chap. 7], we have with some $\tilde{C} > 0$

$$P\left(\sup_{s \le t(\eta_0, r)\wedge\nu(r)} |Z'(s)| > 2\mu\mu^{-\delta}, \operatorname{sgn}(\eta_0) = 1\right) \le \tilde{C}\mu^{2(\delta-1)}\mathbb{E}\left[[Z']_{t(\mu^{-\gamma},r)}\right]$$

$$\le \tilde{C}\mu^{2(\delta-1)}4A_1^2(r)\,\mathbb{E}[B(t(\mu^{-\gamma},r)+\tau) - B(\tau)]$$

$$= \tilde{C}\mu^{2(\delta-1)}\ 4A_1^2(r)\,\mu\, t(\mu^{-\gamma},r)$$

which converges to zero by Lemma 3.2.5 as $\mu \to \infty$, whenever $\delta < 1/2$. We conclude that for any $r > 0$

$$P\left(\left|\int_0^{t(\eta_0,r)\wedge\nu(r)} \sigma_s' G(\eta_s)\mathrm{d}s\right| > \mu^{-\delta}, \operatorname{sgn}(\eta_0) = 1\right) \to 0.$$

For $r_0 > 0$, on the event

$$\left\{\left|\int_0^{t(\eta_0,r_0)\wedge\nu(r_0)} \sigma_s' G(\eta_s)\mathrm{d}s\right| \le \mu^{-\delta}, \operatorname{sgn}(\eta_0) = 1\right\},$$

Gronwall's inequality applied to (3.2.19) implies that there is $C > 0$ such that for $t \le t(\eta_0, r_0)$,

$$|\Delta(t \wedge \nu(r_0))| \le e^{L(r_0)t(\eta_0,r_0)}\mu^{-\delta} \le C\mu^{\gamma L(r_0)/a}\mu^{-\delta}. \tag{3.2.20}$$

Since $F'(0) = a$, we can choose r_0 so small such that $\gamma L(r_0)/a < 1/2$. Consequently, for some $\delta < 1/2$, we have $\rho = \delta - \gamma L(r_0)/a > 0$ so that the right hand side in the bound above converges to 0. For any $r \in (0, r_0)$, we conclude using (3.2.20) at $t = \nu(r_0)$ that

$P(\nu(r_0) < t(\eta_0, r), \mathrm{sgn}(\eta_0) = 1) \to 0$. Now, using (3.2.20) at $s \le t(\eta_0, r) < t(\mu^{-\gamma}, r_0)$, we have

$$P\left(\nu(r_0) \ge t(\eta_0, r), \sup_{s \le t(\eta_0, r)} |\Delta(s)| > C\mu^{-\rho}, \mathrm{sgn}(\eta_0) = 1\right) \to 0,$$

which implies (3.2.18) for $\alpha = 1$. □

Combining Lemmas 3.2.4, 3.2.5 and 3.2.6, we conclude

Lemma 3.2.7. *For any* $r \in (0, r_0)$

$$\left(x_{\tau(r)}, \tau(r) - \frac{1}{2a}\log\mu\right) \xrightarrow{d} \left(r \cdot \mathrm{sgn}(H), -\frac{1}{a}\log|H| + \frac{\log r}{a} + K(r \cdot \mathrm{sgn}\, H)\right) \quad \text{as } \mu \to \infty.$$

Now let us consider the exit time from an arbitrary interval.

Lemma 3.2.8. *For* $r \in (0, r_0)$*, define* $\eta_t = x_{\tau(r)+t}$*. Then, for any* $T > 0$*,*

$$\sup_{0 \le t \le T} |\eta_t - S^t \eta_0| \xrightarrow{P} 0 \quad \text{as } \mu \to \infty.$$

The proof follows the same steps as the proof of Lemma 3.2.6, however easier as the time horizon is fixed this time. Our main result follows from Lemmas 3.2.7 and 3.2.8.

Bibliography

[1] S. Almada and Y. Bakhtin. Normal forms approach to diffusion near hyperbolic equilibria. *Nonlinearity*, 24:1883–1907, 2011.

[2] V. V. Anisimov. *Switching processes in queueing models*. Applied Stochastic Methods Series. ISTE, London; John Wiley & Sons, Inc., Hoboken, NJ, 2008.

[3] D. Aristoff. Percolation of hard disks. *J. Appl. Probab.*, 51(1):235–246, 2014.

[4] S. Aumann. Spontaneous breaking of rotational symmetry with arbitrary defects and a rigidity estimate. *J. Stat. Phys.*, 160:168–208, 2015.

[5] Y. Bakhtin. Exit asymptotics for small diffusion about an unstable equilibrium. *Stochastic Process. Appl.*, 118(5):839–851, 2008.

[6] Y. Bakhtin. Small noise limit for diffusions near heteroclinic networks. *Dynamical Systems*, 25(3):413–431, 2010.

[7] Y. Bakhtin. Noisy heteroclinic networks. *Probab. Theory Relat. Fields*, 150:1–42, 2011.

[8] Y. Bakhtin. Decision making times in mean-field dynamic Ising model. *Ann. Henri Poincaré*, 13(5):1291–1303, 2012.

[9] Y. Bakhtin and J. Correll. A neural computation model for decision-making times. *J. Math. Psych.*, 56(5):333–340, 2012.

[10] Y. Bakhtin and A. T. Gaál. Exit time asymptotics for dynamical systems with fast random switching near an unstable equilibrium. *Stochastics and Dynamics*, Vol. 20, No. 1, 12 pp., 2020.

[11] Y. Bakhtin and T. Hurth. Invariant densities for dynamical systems with random switching. *Nonlinearity*, 25(10):2937–2952, 2012.

[12] Y. Bakhtin and Zs. Pajor-Gyulai. Tails of exit times from unstable equilibria on the line. *arXiv e-prints*, page arXiv:1810.05341, Oct. 2018. Accepted at *J. Appl. Probab.*

[13] Y. Bakhtin and Zs. Pajor-Gyulai. Malliavin calculus approach to long exit times from an unstable equilibrium. *Ann. Appl. Probab.*, 29(2):827–850, 2019.

[14] Y. Bakhtin and Zs. Pajor-Gyulai. Scaling limit for escapes from unstable equilibria in the vanishing noise limit: Nontrivial jordan block case. *Stochastics and Dynamics*, Vol. 19, No. 3, 21 pp., 2019.

[15] R. Bauerschmidt, D. Conache, M. Heydenreich, F. Merkl, and S. W. W. Rolles. Dislocation lines in three-dimensional solids at low temperature. *Annales Henri Poincaré*, 20(9):3019–3057, Jul 2019.

[16] M. Benaïm, S. Le Borgne, F. Malrieu, and P.-A. Zitt. Quantitative ergodicity for some switched dynamical systems. *Electron. Commun. Probab.*, 17, no. 56:1–14., 2012.

[17] A. Benveniste, M. Metivier, and P. Priouret. *Adaptive Algorithms and Stochastic Approximations*. Springer-Verlag, 1990.

[18] J. Brea, A. T. Gaál, R. Urbanczik, and W. Senn. Prospective coding by spiking neurons. *PLoS Comput. Biol.*, 12(6), 2016.

[19] M. H. A. Davis. *Markov models and optimization*, volume 49 of *Monographs on Statistics and Applied Probability*. Chapman & Hall, London, 1993.

[20] M. V. Day. On the exit law from saddle points. *Stochastic Process. Appl.*, 60(2):287–311, 1995.

[21] C. Dellacherie and P.-A. Meyer. *Probabilities and potential. B*, volume 72 of *North-Holland Mathematics Studies*. North-Holland Publishing Co., Amsterdam, 1982. Theory of martingales, Translated from the French by J. P. Wilson.

[22] D. Dereudre, R. Drouilhet, and H.-O. Georgii. Existence of Gibbsian point processes with geometry-dependent interactions. *Probab. Theory Relat. Fields*, Volume 153, Issue 3–4:643–670, 2012.

[23] A. Eizenberg. The exit distributions for small random perturbations of dynamical systems with a repulsive type stationary point. *Stochastics*, 12(3-4):251–275, 1984.

[24] S. N. Ethier and Th. G. Kurtz. *Markov processes. Characterization and convergence.* Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. John Wiley & Sons Inc., New York, 1986.

[25] A. Faggionato, D. Gabrielli, and M. Ribezzi Crivellari. Non-equilibrium thermodynamics of piecewise deterministic Markov processes. *J. Stat. Phys.*, 137(2):259–304, 2009.

[26] G. Friesecke, R. James, and S. Müller. A theorem on geometric rigidity and the derivation of nonlinear plate theory from three-dimensional elasticity. *Comm. Pure Appl. Math.*, Volume 55, Issue 11:1461–1506, 2002.

[27] A. T. Gaál. Long-range orientational order of a random near lattice hard sphere and hard disk process. arXiv:1911.02167 [math.PR], 2019, accepted for publication by the Applied Probability Trust (`http://www.appliedprobability.org`) in *J. Appl. Probab.* 57(2), 2020.

[28] A. T. Gaál. Spontaneous breaking of rotational symmetry in a probabilistic hard disk model in Statistical Mechanics (Master's thesis). Retrieved from `https://math.nyu.edu/~gaal/thesis/GaalMASoSe2013.pdf` on Nov 25 2019, September 2013.

[29] A. T. Gaál. Long-range order in a hard disk model in statistical mechanics. *Electron. Commun. Probab.*, 19(9):1–9., 2014.

[30] H.-O. Georgii. *Gibbs Measures and Phase Transitions.* de Gruyter, Berlin, 1988.

[31] H.-O. Georgii and H. Zessin. Large deviations and the maximum entropy principle for marked point random fields. *Probab. Theory Relat. Fields*, 96:177–204, 1993.

[32] A. Giulian and F. Theil. Long range order in atomistic models for solids. arXiv:1907.07923 [math.AP], 2019.

[33] R. Hersh. The birth of random evolutions. *Math. Intelligencer*, 25(1):53–60, 2003.

[34] M. Heydenreich, F. Merkl, and S. Rolles. Spontaneous breaking of rotational symmetry in the presence of defects. *Electron. J. Probab.*, 19(111):1–17, 2014.

[35] Ch. Hofer-Temmel. Disagreement percolation for the hard-sphere model. *Electron. J. Probab.*, Volume 24, paper no. 91, 22 pp., 2019.

[36] M. Kac. A stochastic model related to the telegrapher's equation. *Rocky Mountain J. Math.*, 4:497–509, 1974. Reprinting of an article published in 1956, Papers arising from a Conference on Stochastic Differential Equations (Univ. Alberta, Edmonton, Alta., 1972).

[37] Y. Kifer. The exit problem for small random perturbations of dynamical systems with a hyperbolic fixed point. *Israel J. Math.*, 40(1):74–96, 1981.

[38] S. D. Lawley, J. C. Mattingly, and M. C. Reed. Stochastic switching in infinite dimensions with applications to random parabolic PDE. *SIAM J. Math. Anal.*, 47(4):3035–3063, 2015.

[39] G. Lazzaroni, M. Palombaro, and A. Schlömerkemper. Rigidity of three-dimensional lattices and dimension reduction in heterogeneous nanowires. *Discrete Contin. Dyn. Syst. S.*, 10:119–139, 2017.

[40] A. Magazinov. On percolation of two-dimensional hard disks. *Commun. Math. Phys.*, 364(1):1–43, Nov 2018.

[41] F. Malrieu. Some simple but challenging Markov processes. *Ann. Fac. Sci. Toulouse Math. (6)*, 24(4):857–883, 2015.

[42] F. Merkl and S. Rolles. Spontaneous breaking of continuous rotational symmetry in two dimensions. *Electron. J. Probab.*, Volume 14, paper no. 57:1705–1726, 2009.

[43] T. Richthammer. Translation-invariance of two-dimensional Gibbsian point processes. *Commun. Math. Phys.*, 274:81–122, 2007.

[44] T. Richthammer. Lower bound on the mean square displacement of particles in the hard disk model. *Commun. Math. Phys.*, 345:1077–1099, 2016.

[45] J. L. Salle and S. Lefschetz. *Stability by Liapunov's Direct Method With Appliations.* Academic Press, United Kingdom Edition, 1961.

[46] R. Urbanczik and W. Senn. Learning by the dendritic prediction of somatic spiking. *Neuron*, 81(3):521–528, 2014.

[47] G. G. Yin and Ch. Zhu. *Hybrid switching diffusions (Properties and applications)*, volume 63 of *Stochastic Modelling and Applied Probability.* Springer, New York, 2010.