SISSA
Scuola
Internazionale
Superiore di
Studi Avanzati

Physics Area - PhD course in
Statistical Physics

Numerical Investigations of Non-Reciprocal and Frustrated Systems

Candidate:
Gabriele Bandini

Advisors:
Prof. Andrea Gambassi
Dr. Asja Jelic

Academic Year 2025–2026



Cerca una maglia rotta nella rete

— Eugenio Montale, *In limine*

Abstract

This PhD thesis addresses two distinct problems in classical statistical physics: how vision-dependent interactions shape collective order, and how machine-learning techniques combined with renormalization-group ideas can improve the sampling of frustrated systems. The first and main part of the thesis focuses on non-reciprocal interactions in active matter, while the second part applies a wavelet-based approach to configuration sampling in statistical physics.

In particular, the first part investigates two-dimensional XY models with vision-dependent interactions. By introducing and comparing non-reciprocal and reciprocal realizations of these models, we show that the emergence of long-range orientational order is not caused by non-reciprocity alone. Instead, it primarily originates from the coupling between the spin orientations and the underlying lattice — a consequence of the vision-cone interaction. This coupling generates effective angular anisotropies, which selects preferred orientations. One of the reciprocal models also displays an unexpected order-by-disorder transition, in which thermal fluctuations strengthen orientational order. We then consider non-reciprocal models with smooth anisotropic interaction kernels and show that the selected pinning directions depend on both the kernel and the kind of terms retained in the microscopic dynamics. In particular, as the interaction strength of a spin with the neighbors depends on its orientation relative to the bonds connecting to them, the Langevin dynamics naturally gives rise to two distinct contributions, which we called reactive and proactive. The reactive term describes the response of a spin to the instantaneous orientations of the currently perceived neighbors, whereas the proactive term accounts for the gradient of the interaction strength. We clarify their different roles in producing pinning along specific lattice directions and investigate how this pinning appears in both local and global equations of motion. Their possible implications for the stability of ordered phases and for off-lattice active systems are also discussed.

The second part of the thesis applies the wavelet conditional renormalization group to the sampling of a frustrated soft-spin biaxial next-nearest-neighbor Ising model, for which conventional cluster algorithms cannot provide efficient sampling. The method decomposes microscopic configurations into coarse fields and wavelet coefficients and learns their conditional distributions independently at subsequent scales. At the Ising-like critical point, the conditional sampling dynamics shows no detectable critical slowing down. The reconstructed configurations reproduce local statistics and two-point correlation functions across most of the phase diagram, while discrepancies in the microscopic energy distribution reveal limitations associated with the finite-dimensional parametrization of the conditional energies. These limitations can, in principle, be reduced systematically by enriching the operator basis. This provides an interpretable route for improving the model, in contrast to more standard machine-learning approaches whose internal representations often remain difficult to understand.

Acknowledgments

I want to express my gratitude to everyone who accompanied me throughout my PhD journey.

I would first like to thank my supervisors, Andrea Gambassi and Asja Jelic, for guiding me with care and humanity through my first steps as a researcher; for encouraging me to trust my intuition and explore new ways of approaching problems; and for teaching me how to find the right balance between thinking freely and laterally and being pragmatic. I am happy and proud of the work we have done and of the results we have achieved together.

I would also like to thank my collaborators on my first project, Sarah Loos and Davide Venturelli, for the dedication and professionalism with which they contributed to its development. I thank Giulio Biroli and Patrick Charbonneau for their guidance throughout my second project, on wavelets. I am particularly grateful to Giulio for hosting me for three months at LPENS in Paris, in what became an incredibly enriching experience. I also thank my colleague Etienne Lempereur for his explanations and for the code used in the wavelet project.

I would then like to thank my colleagues at SISSA, especially my batchmates Giacomo, Konstantinos, Marcin, and Andrea, with whom I shared this four-year adventure most closely. I am also grateful to everyone in the StatPhys group, the students, postdocs, and professors, who made it such a lively, welcoming, and stimulating environment. In particular, I thank Cristiano Muzzi, with whom I began the wavelet project.

More broadly, I thank all my colleagues at SISSA with whom I shared both interdisciplinary exchanges and meaningful human connections, particularly those in the Condensed Matter, Neuroscience, and Data Science groups.

I would also like to thank the staff of the SISSA canteen, who know us by name and helped turn the many meals we shared at SISSA into genuine moments of conviviality.

I am deeply grateful to all the friends who stood by me throughout these years: those in Trieste, who filled the city with light during these four years; those in Pisa, who kept alive my love for the city where I spent my university years; those in Massa, my safe haven during the holidays; and those in Paris, who made my visiting period so special and whom I am going to see again very soon.

Finally, I thank my family, that loves and supports me every single day.

List of publications

The work presented in this thesis has led to the following publications and manuscripts:

[1] G. Bandini, D. Venturelli, S. A. M. Loos, A. Jelic, and A. Gambassi,
"The XY Model with Vision Cone: Non-Reciprocal vs. Reciprocal Interactions,"
J. Stat. Mech. **2025**, 053205 (2025). doi:10.1088/1742-5468/adc243.
This publication forms the basis of Chapters 2 and 3.

[2] G. Bandini, A. Jelic, and A. Gambassi,
"Proactivity and Pinning in the Non-Reciprocal XY Model with Vision Anisotropy,"
arXiv:2606.27310 (2026). doi:10.48550/arXiv.2606.27310.
This preprint forms the basis of Chapters 4 and 5.

[3] G. Bandini, G. Biroli, P. Charbonneau, and A. Gambassi,
"Overcoming Critical Slowing Down in Frustrated Spin Systems by Learned Multiscale Sampling,"
arXiv:2608.31114 (2026). doi:10.48550/arXiv.2608.31114 .
This manuscript forms the basis of Chapter 8.

Contents

Abstract iv

Acknowledgments v

List of publications vi

1 Introduction 1
- 1.1 Multiscale phenomena in statistical physics . . . 1
  - 1.1.1 From equilibrium to nonequilibrium and non-reciprocal active systems . . 1
  - 1.1.2 Multiscale structure as a computational opportunity . . . 9
- 1.2 Main contributions . . . 10
- 1.3 Outlook of the thesis . . . 12
- 1.4 Outline of the thesis . . . 14

I The non-reciprocal XY model with vision anisotropy 18

2 XY models with vision cone: non-reciprocal and reciprocal implementations 19
- 2.1 The XY model with vision cone: three variations . . . 20
  - 2.1.1 Non-reciprocal XY model — NRXY . . . 20
  - 2.1.2 Asymmetric reciprocal XY model — ARXY . . . 22
  - 2.1.3 Symmetric reciprocal XY model — SRXY . . . 22
  - 2.1.4 Energetically unfavorable range — EUR . . . 23
- 2.2 Methods . . . 24
  - 2.2.1 Observables . . . 25
  - 2.2.2 Phase diagrams of reference models . . . 27
- 2.3 Numerical results . . . 29
  - 2.3.1 The non-reciprocal XY model . . . 29
  - 2.3.2 The asymmetric reciprocal XY model . . . 30
  - 2.3.3 The symmetric reciprocal XY model . . . 31
- 2.4 Concluding remarks . . . 33

3 Spin–lattice coupling and emergent angular anisotropy 34
- 3.1 Spin–lattice and internal symmetry . . . 35
- 3.2 The asymmetric reciprocal model . . . 36
  - 3.2.1 Two different manifestations of $\mathbb{Z}_4$ symmetry: $\theta = 100^\circ$ and $\theta = 280^\circ$ . . . 36
  - 3.2.2 ARXY for $\theta = 180^\circ$: LRO without EUR . . . 39

3.3 The symmetric reciprocal model . . . 42
3.3.1 The role of redundant bonds for $\theta > 180°$ . . . 42
3.3.2 Order-by-disorder transition . . . 44
3.4 Concluding remarks . . . 46

4 Reactive and proactive dynamics in non-reciprocal XY models 48
4.1 From the selfish energy to the Langevin description . . . 49
4.1.1 Glauber vs. Langevin dynamics . . . 49
4.1.2 From active to proactive matter . . . 51
4.1.3 Interaction kernels . . . 52
4.2 Mean field: coupling strength at $T = 0$ . . . 54
4.3 Numerical results for the mean orientation . . . 55
4.3.1 Transient dynamics . . . 56
4.4 Concluding remarks . . . 58

5 Local and global lattice pinning in non-reciprocal XY models 60
5.1 Search for a pinning term . . . 61
5.1.1 Local and global equations of motion . . . 61
5.2 Kernel-dependent pinning mechanisms . . . 65
5.2.1 Pinning terms for the modulated coupling . . . 65
5.2.2 Pinning terms for the sinusoidal coupling . . . 66
5.2.3 Pinning terms for the von Mises coupling . . . 67
5.2.4 Proactivity and pinning: A summary . . . 69
5.2.5 Global pinning in the stationary state . . . 70
5.3 Comparison with correlation time and pinning from the structure factor . . . 74
5.3.1 Correlation time . . . 74
5.3.2 Pinning extracted from the structure factor . . . 76
5.4 Concluding remarks . . . 79

6 Future perspectives on non-reciprocal XY models 81
6.1 Metastability of the LRO phase in the NRXY model . . . 82
6.1.1 Defect-mediated decay of LRO phase . . . 82
6.1.2 A possible stabilizing role of the proactive term . . . 84
6.1.3 A preliminary switching experiment . . . 85
6.2 Proactive interactions in off-lattice active matter . . . 87
6.2.1 Overdamped intelligent active Brownian particles . . . 88
6.2.2 Discrete Vicsek-like implementation with reactive and proactive terms . . . 90
6.2.3 Emergent collective behavior . . . 92
6.3 Concluding remarks . . . 94

II Wavelet sampling for frustrated systems 96

7 Wavelet conditional renormalization group: efficient sampling at criticality 97
7.1 Wavelet decomposition and multiscale factorization . . . 98
7.1.1 RG and WCRG . . . 100
7.1.2 Coarse-grained and conditional energies . . . 101

7.2 Learning the multiscale model . . . 104
7.2.1 Direct coupling estimation by maximum likelihood . . . 104
7.2.2 Conditional coupling estimation by maximum likelihood . . . 107
7.2.3 Conditional coupling estimation by score matching . . . 109
7.3 Concluding remarks . . . 113

8 Wavelet conditional sampling of frustrated systems 114
8.1 Frustrated soft-spin models . . . 115
8.2 Wavelet conditional renormalization group . . . 118
8.3 Numerical results . . . 119
8.3.1 Wavelet sampling across the phase diagram . . . 119
8.3.2 Conditional sampling at the Ising-like critical point . . . 120
8.3.3 Configuration-wide observables . . . 124
8.4 Concluding remarks . . . 125

9 Conclusions 127

A Additional material for Chapters 2 and 3 131
A.1 Details of the Monte Carlo simulation . . . 131
A.2 Scaling of the second-moment correlation length . . . 131
A.3 NRXY model with different update protocols . . . 132
A.4 Analysis of the EPR in the NRXY model . . . 133
A.5 ARXY model for $\theta = 60^\circ$ . . . 136

B Additional material for Chapters 4 and 5 137
B.1 Crossover in the pinning direction for the reactive von Mises model on the triangular lattice . . . 137
B.2 Local equations for the spin . . . 138
B.2.1 Modulated coupling . . . 140
B.2.2 Sinusoidal coupling . . . 141
B.2.3 Von Mises coupling . . . 142
B.3 Global equations for the orientation . . . 144
B.3.1 Modulated coupling . . . 144
B.3.2 Sinusoidal coupling . . . 146
B.3.3 Von Mises coupling . . . 148

C Additional material for Chapters 7 and 8 153
C.1 Low-temperature analysis of the soft-spin BNNNI model . . . 153
C.2 Finite-size scaling of the Ising-like critical line . . . 156
C.3 Wavelet-transform conventions . . . 156
C.4 Learning the probability distributions . . . 158
C.4.1 Coarsest energy . . . 158
C.4.2 Conditional-energy ansatz . . . 159
C.4.3 Score-matching linear system . . . 159
C.4.4 Conditional free energy and reconstruction of microscopic energy . . . 160
C.5 Conditional sampling and reconstruction . . . 161

Bibliography **162**

Chapter 1

Introduction

1.1 Multiscale phenomena in statistical physics

Statistical physics aims to explain how the collective behavior of a macroscopic system emerges from the interactions among a large number of microscopic degrees of freedom. Even when the microscopic rules are simple, their consequences at the collective level may be highly non-trivial: these include phase transitions, spontaneous symmetry breaking, universal scaling behavior, topological defects, and complex macroscopic structures [4–6].

Collective phenomena are intrinsically multiscale: this is particularly evident near a continuous phase transition, where the correlation length of the system grows upon approaching the transition and eventually becomes the dominant macroscopic length scale. In this regime, fluctuations over a wide range of scales contribute simultaneously to the observable behavior of the system, making any description which assumes a separation between microscopic and macroscopic degrees of freedom inadequate.

The renormalization group provides a natural conceptual framework for addressing this problem [7–9]. Its basic idea is to construct a sequence of effective theories by progressively eliminating short-scale fluctuations and rescaling the remaining degrees of freedom. This procedure generates a flow in the space of possible interactions. Long-distance behavior close to critical points is controlled by the fixed points of this flow and by the operators that are relevant or marginal under repeated coarse graining. This perspective explains the concept of universality, i.e., why systems with very different microscopic details may eventually display the same large-scale critical behavior. Close to a fixed point, irrelevant microscopic features are progressively suppressed, while only a small number of relevant features of the systems, such as its internal symmetry and its dimensionality, determine the large-scale physics.

1.1.1 From equilibrium to nonequilibrium and non-reciprocal active systems

At thermal equilibrium, multiscale phenomena can be described within a coherent and well-established theoretical framework. The probability of each microscopic configuration of the system is governed by the Gibbs measure, while detailed balance ensures the absence of stationary probability currents in configuration space. Symmetries and conservation laws further constrain the form of the possible effective theories. The existence of such a framework, however, does not render equilibrium many-body systems trivial, but it allows the systematic development of

analytical and numerical methods. In particular, field theoretical tools based on the renormalization group have provided a detailed understanding of a broad class of continuous phase transitions.

Far from equilibrium, no comparably general framework is currently available [10, 11]. A stationary state is generically not described by a Gibbs distribution, detailed balance is in general violated, and probability currents may persist even when macroscopic observables appear stationary. The large-scale behavior of the statistical system may depend explicitly on the microscopic dynamics rather than only on a static energy function. Consequently, two systems characterized by similar symmetries and local interactions may display qualitatively different collective behavior when their dynamical rules differ.

Active matter provides a particularly rich setting in which to investigate these questions. Active systems are composed of units that consume energy and convert it into motion or mechanical forces. Examples range from molecular motors and cellular structures to bacterial suspensions, animal groups, synthetic swimmers, and self-propelled colloids [12–14].

Because energy is continuously injected and dissipated at the level of the individual constituents, active systems are intrinsically out of equilibrium. This sustained energy input generates collective phenomena with no direct equilibrium counterpart, including macroscopic polar order [15, 16], traveling density bands [17], clustering [18, 19], and motility-induced phase separation [20, 21]. A paradigmatic microscopic description of collective polar order is provided by the Vicsek model [15]. In this model, self-propelled point particles move at constant speed and repeatedly adjust their headings towards the average heading of their neighbors, subject to angular noise. This defines an *alignment interaction*, which tends to reduce the angular differences between the directions of motion of nearby agents, without necessarily producing attraction between their positions. The competition between alignment and noise produces a transition between disordered motion and collective polar order.

These collective states cannot be generally understood as equilibrium phases perturbed by a weak external drive. Activity modifies their fluctuations [22, 23], stability [24], the dynamics of defects [25–28], and their hydrodynamic modes [29] in a fundamental way. One quantitative signature of this departure from equilibrium is the entropy production rate. Microscopic field-theoretic formulations have provided explicit expressions for this quantity in specific instances of active-particle dynamics [30]. Its precise interpretation, however, depends on the microscopic variables and dissipative channels retained in the description. Particular care is thus required when moving from particle-based dynamics to coarse-grained field theories [31, 32].

At both the microscopic and coarse-grained levels, an important mechanism through which active behavior can emerge is provided by non-reciprocal interactions between particles, species, or fields.

Non-reciprocal interactions

Non-reciprocal interactions, in which the influence of one component on another is not matched by a corresponding reverse influence, are widespread in active and living systems and provide an important route to nonequilibrium collective behavior [33].

Focusing, for concreteness on microscopic degrees of freedom with pairwise interactions, let g_{ij} denote the influence exerted by degree of freedom j on the degree of freedom i. Reciprocity requires that this influence is balanced by the reverse one, whereas non-reciprocity arises when

$$g_{ij} \neq g_{ji}. \tag{1.1}$$

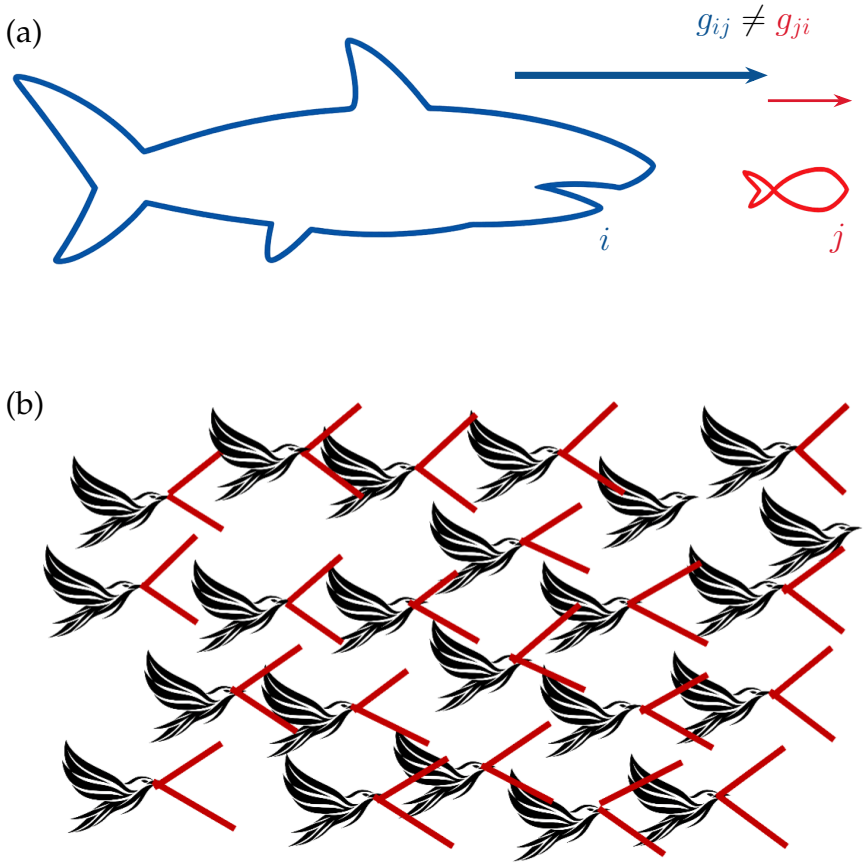


Figure 1.1: **Two routes to perception-induced non-reciprocity.** (a) In the presence of two species of agents, e.g., predators and preys, agents of different species respond to one another asymmetrically: predators pursue nearby prey, whereas prey tend to move away from approaching predators. (b) In the presence of a single species, non-reciprocity may arise from anisotropic perception. Because each agent responds only to neighbors falling within its field of view, one agent may influence another without being influenced in return. This second mechanism motivates the models studied in this thesis.

The quantities g_{ij} do not necessarily represent mechanical forces: they may describe alignment tendencies, chemical responses, or information-mediated influences. For mechanical interactions, the corresponding condition reads $\mathbf{F}_{j\to i} \neq -\mathbf{F}_{i\to j}$, representing an effective violation of the action–reaction symmetry embodied in Newton's third law. It is useful, however, to distinguish this force-level statement from the broader notion of a non-mutual response, which also applies when no mechanical forces are involved.

Origin and physical realizations of non-reciprocal interactions. A primary microscopic route to realizing non-reciprocity is provided by interactions mediated by a dynamical environment. In a closed mechanical system, equal and opposite internal forces ensure conservation of total momentum. An effective description, however, generally retains only a subset of the microscopic degrees of freedom. When a mediating field, substrate, or surrounding fluid is eliminated, for example, from a system in which they interact with some particles, the retained particles may

exchange momentum or energy through hidden channels. Accordingly, their effective interactions need not be symmetric or derivable from a potential which depends only on the retained variables. This mechanism underlies wake-mediated interactions in complex plasmas [34, 35], chemically and phoretically interacting particles [36–40], and droplets coupled through vapor [41], surfactant [42], or other environmental fields.

Within this broad class, the emergence of non-reciprocity becomes especially transparent in multicomponent systems, in which different species, internal states, or functional units respond differently to the same mediating field. The resulting dynamics may then acquire a pursuit–escape or predator–prey character: one component is effectively attracted toward the other, while the latter is repelled from the former, as sketched in Fig. 1.1(a). Such chase-and-run responses have been modeled or observed in cellular populations [43] and chemically communicating droplets [42].

A conceptually different situation arises when the interaction is not mediated by a mechanical or chemical force, but by the acquisition and processing of information. In biological and artificial active systems, agents may react selectively to visible, audible, or otherwise detected neighbors. The corresponding interaction rule may depend on the instantaneous orientation, internal state, past observations, or adaptive strategy of each agent, thereby continuously reshaping the directed network that determines which agents detect and respond to one another. Vision-based interactions, illustratively depicted in the two panels of Fig. 1.1 for the multi and single-species case, provide a particularly important example of this information-mediated non-reciprocity and will be discussed in detail below.

Non-reciprocal collective phenomena and dynamical self-assembly. Non-reciprocal interactions may generate a broad range of collective phenomena in spatially extended systems. Asymmetric chemical cross-interactions can drive multicomponent mixtures to separate into coexisting phases with different densities and compositions [44]. In the non-reciprocal Cahn-Hilliard model, for example, the different species are represented by distinct conserved scalar density fields. A static phase-separated state therefore consists of coexisting macroscopic domains in which these fields take different values, corresponding to phases with different compositions. As the asymmetry of the couplings increases, this static state can give way to traveling density waves with global polar order [45]. More generally, antagonistic couplings between scalar fields provide a route from stationary patterns to traveling or oscillatory states [46].

Similar effects arise in particle-based active mixtures. Non-reciprocal cross-species quorum-sensing interactions (i.e., a density-dependent response mediated by chemical signals) can produce chaotic chasing bands and anomalously slow coarsening [47], whereas non-reciprocal alignment can generate asymmetric, species-selective clusters in which a dense component chases a more dilute accumulation of the other species [48]. In mixtures of chiral active particles, asymmetric couplings can convert stationary instabilities into oscillatory ones and modify the relative orientation of the emerging flocks [49]. Recent work on non-reciprocal two-species extensions of the Vicsek model further indicates that globally coherent chiral motion is not a generic consequence of non-reciprocity, but survives only within a restricted regime of density, motility, inter-species coupling, and system size [50]. Beyond particle-based active mixtures, non-reciprocal coupling also enriches the collective dynamics of oscillator populations, reshaping chimera states and generating traveling chimeras and run-and-chase phases [51]. Non-reciprocity may also produce a collective property absent from either component in isolation, as illustrated by the emergence of long-range

polar order from the coupling of two nonpolar fields [52].

Non-reciprocity not only generates collective patterns, but can also be exploited as a mechanism for dynamical self-assembly. Already at the level of a few particles, chemically active colloids with unequal production and response coefficients experience asymmetric phoretic interactions. Particles that are not self-propelled in isolation can consequently assemble into composite structures with a net translational or angular velocity [38]. At larger scales, reciprocal interactions may encode several target structures, while additional non-reciprocal couplings drive persistent transitions between them [53]. The same departure from reciprocal relaxation can help a system leave metastable configurations and escape kinetic traps that would arrest conventional self-assembly [54]. Accordingly, non-reciprocity affects not only the structures that emerge in a system, but also the dynamics and pathways through which they are assembled.

Irreversibility, phase transitions, and defects in non-reciprocal systems. Non-reciprocal dynamics cannot, in general, be represented as a relaxation toward the minimum of a single free-energy functional. In stochastic systems, non-reciprocal couplings can sustain steady exchanges of energy and information between different degrees of freedom [55]. Their time irreversibility can be quantified through the entropy-production rate [56], which has also been studied in the non-reciprocal Cahn-Hilliard field theory [57].

The macroscopic consequences of a microscopic asymmetry depend on its fate under coarse graining. In active mixtures, non-reciprocity may effectively fade at large scales, leading to an effective equilibrium description, or survive and generate dynamical patterns [58]. More generally, the relevance of non-reciprocity for critical behavior can be assessed perturbatively through scaling arguments [59]. In single-species Ising models with state-dependent vision-cone couplings – where up and down spins assign larger weights to opposite subsets of bonds, so that the influence of one spin on another need not be mutual – , non-reciprocity modifies the phase diagram and produces anisotropic coarsening [60].

In systems containing two or more coupled species, or more generally several coupled order-parameter fields, some non-reciprocal phase transitions can be understood by linearizing the dynamical equations around a stationary state. Let $\boldsymbol{\phi} = (\phi_1, \ldots, \phi_n)$ collect these fields and denote a small perturbation by $\delta\boldsymbol{\phi}$. Its linearized evolution takes the form $\partial_t \delta\boldsymbol{\phi} = \mathcal{L}\delta\boldsymbol{\phi}$, where $\mathcal{L}$ is the stability operator. Non-reciprocal cross-couplings between species generically render $\mathcal{L}$ non-Hermitian. An exceptional point occurs when two or more eigenvalues of $\mathcal{L}$ and their corresponding eigenmodes coalesce, so that the stability operator no longer possesses a complete set of linearly independent eigenmodes. These spectral singularities can mark transitions from static order to oscillatory, traveling, or chiral phases and, in some cases, dynamically restore a spontaneously broken continuous symmetry. They may also be accompanied by a strong enhancement of fluctuations, entropy production, or chirality [61–63]. The notion of an exceptional point has also been generalized beyond linear stability. At bifurcations involving nonlinear attractors, generalized exceptional points can be identified through the coalescence of covariant Lyapunov vectors associated with vanishing Lyapunov exponents [64].

The behavior of topological defects provides another manifestation of non-reciprocity. Consider, for example, a two-dimensional lattice of planar spins. Vortices and antivortices (which carry opposite topological charges) correspond to configurations in which the spin orientation winds in opposite senses around their cores, i.e., plaquettes of a lattice. In the non-reciprocal XY model (that will be discussed in detail in this thesis), the coupling between neighboring spins is weighted

by an anisotropic kernel. The resulting dynamics of defects depends not only on the topological charge of each defect, but also on its shape, namely the detailed arrangement of the spins around its core. Non-reciprocity can therefore enhance or hinder the annihilation of vortex–antivortex pairs [65]. A continuum description of the same non-reciprocal $\mathrm{O}(2)$ model further shows [66] that non-reciprocity advects and reshapes patterns, controls their trajectories, and modifies the stability of configurations in which the orientation winds across the periodic system.

These examples show that non-reciprocity affects not only the stationary phases of a system, but also the dynamical pathways through which order is created, reorganized, or destroyed.

Perception-induced non-reciprocity

Among the mechanisms capable of generating asymmetric interactions, anisotropic perception is particularly relevant for active and living systems. In contrast to the physical realizations discussed above, in this case non-reciprocity does not necessarily originate from a force mediated by an environmental field that has been eliminated from the description. Instead, it arises at the level of the sensing and decision-making process of the individual agent. Each agent of, e.g., a flock of birds, constructs a partial representation of its surroundings and responds according to its own position, orientation, sensory apparatus, and possibly internal state. The resulting interaction is therefore mediated by information rather than by a direct mechanical force. Perception is generally limited and directional and the resulting interaction network need not be symmetric: agent i may perceive and respond to agent j, while j either does not perceive i or responds to it differently.

Experimental investigations of collectives such as insect swarms [67], bird flocks [68–71], schools of fish [72, 73], herds [74], and traffic-related systems [75], provide several indications that these social interactions are selective, anisotropic, and directed.

These observations have motivated the formulation and analysis of a broad family of theoretical models in which active agents respond differently to neighboring individuals according to how those individuals are perceived [76, 77]. In this thesis we always consider single-species systems, such as the one illustrated in Fig. 1.1(b). For the purposes of this introduction, we broadly classify these systems into two main categories.

The first class consists of *perception-dependent alignment models*, in which an agent aligns its heading with those of selected neighbors, while their relative positions determine which neighbors contribute to the alignment interaction and with what weight. The most direct examples are vision-cone variants of the Vicsek model. Relative to the standard Vicsek rule introduced above, the alignment average is restricted to neighbors lying within a finite angular sector centered on the instantaneous direction of motion of the agent under consideration. These finite fields of view, sometimes combined with stochastic or delayed communication, have been shown to modify the order–disorder transition and to generate condensation, milling, meandering, and spatially cohesive states [78–80]. More recent works show that narrowing the vision cone can drive a transition from globally coherent motion to locally ordered clusters with strong short-range correlations [81]. In a related single-species Vicsek model, spontaneous symmetry breaking can amplify microscopic vision-induced non-reciprocity at macroscopic scales, promoting traveling bands and a condensed traveling-line state [82].

A closely related variation of the Vicsek model replaces the sharp selection of the neighbors imposed by a vision cone with a fore-aft asymmetry. In this case, neighbors located ahead of and behind the agent under consideration both contribute in determining the alignment direction,

but with different weights. Even an arbitrarily weak preference for neighbors located in front of an agent can qualitatively modify the phase diagram of the Vicsek model and generate a regime characterized by multiple traveling bands and a banded liquid [83]. More generally, the antisymmetric component of a fore–aft exchange interaction acts as a non-mutual torque and can produce an advective transport of orientational information, the velocity of velocity differs from the material velocity of the flock [84].

The second class consists of *vision-based attraction and self-steering models*. Here, the headings of the perceived neighbors are not necessarily the main input of the orientational dynamics. Instead, an agent may respond to their spatial distribution, for example by turning towards the center of mass of the particles contained in its visual field. Vision-dependent attraction can generate aggregates, moving files, milling patterns, and nematic structures even in the absence of explicit velocity alignment [85].

This mechanism of steering towards the spatial distribution of visually detected neighbors has been developed systematically in models of so-called *intelligent active Brownian particles*. In these models, each particle reorients its propulsion axis — which determines the self-propulsion direction, — on the basis of visual information acquired from its surroundings. The finite rate at which this direction changes is quantified by a maneuverability parameter. Visual steering may also be combined with an independent polar-alignment interaction, through which a particle tends to orient its propulsion axis along the average heading of its neighbors. Depending on the relative importance of these two mechanisms, intelligent active particles can form worms, compact aggregates, milling states, and cohesive moving structures [86, 87]. Other vision-based steering models combine alignment with short-range collision avoidance, whereby an agent turns away from neighbors that are too close, and longer-ranged cohesion, whereby it turns towards regions occupied by other members of the group. The competition among these three responses can produce an order-disorder transition accompanied by large-scale vortical structures, while cohesion favors compact flock shapes within the ordered phase [88]. Accordingly, alignment and self-steering are not mutually exclusive: an agent may simultaneously attempt to reproduce the headings of its neighbors and steer according to their spatial bearings.

Most of these models consider self-propelled particles moving in continuous space: their orientational and translational dynamics are therefore intrinsically coupled.

The non-reciprocal XY model

A bridge towards lattice spin systems was presented in Ref. [84], where a model of spins fixed at lattice sites but otherwise able to rotate was constructed starting from a model of polar flocks with a fore–aft asymmetric non-mutual torque. Remarkably, the speed at which polarization information is advected remains nonzero even though the lattice spins themselves do not translate. Some phenomena conventionally associated with flocking can therefore originate from the non-reciprocal dynamics of the polarization field itself, rather than from the physical spatial displacement of the agents.

This idea was subsequently implemented in a particularly transparent form in Ref. [89], where the non-reciprocal XY model with vision-cone interactions was introduced. In this model, each spin is a two-component unit vector fixed at a site of a two-dimensional lattice and free to rotate within the lattice plane. It interacts only with those nearest neighbors that lie within an angular sector centered on its instantaneous orientation. Although the spins remain fixed on a regular

lattice, their orientations continuously redefine a directed network of active bonds. The dynamics is governed by Glauber-type update rules based on a local, spin-dependent quantity termed the *selfish energy* [90] that is the local energetic cost associated with rotating a spin, or, equivalently, its opposite represents the local work associated with the proposed move. In general, however, these local costs cannot be obtained as variations of a single global equilibrium Hamiltonian. Instead, the selfish energy defines the local update rule of a generally non-equilibrium dynamics.

The main interest of this model lies in its ability to support long-range orientational order in two spatial dimensions. In equilibrium conditions, the short-ranged interactions and the continuous rotational symmetry of the XY model preclude the existence of a finite-temperature long-range order, according to the Mermin-Wagner theorem [91]. In fact, only quasi-long-range order is observed, with spin-spin correlations which decay algebraically upon increasing the distance between the spins. The non-reciprocal XY model with vision-cone, however, falls outside the equilibrium assumptions of this theorem and, in fact, it was shown to develop a finite magnetization with exponential decay of spin-spin correlations at low temperature, to feature a directional propagation of topological defects and a nonzero entropy production [89]. Accordingly, it provides a minimal setting in which non-reciprocal orientational dynamics can qualitatively alter the ordering properties expected for two-dimensional XY-like models.

Subsequent studies have shown that this ordered state may be metastable and eventually decay through the nucleation and proliferation of defects [92, 93]. The dynamics of topological defects in the non-reciprocal XY model was investigated in detail in Ref. [65], where anisotropic interactions were shown to reshape the defects and to either enhance or hinder their annihilation. In formulating the corresponding Langevin description starting from the Glauber-type dynamics, that work also explicitly identified a contribution proportional to the derivative of the interaction kernel with respect to the orientation of the spin, which accounts for the variation of the interaction weights induced by the spin's own rotation.

Setting aside the question of the ultimate stability of the ordered state in the non-reciprocal XY model, one finds that such a state persists for remarkably long times. Within this long-lived ordered state, the mean orientation is not free to rotate continuously; rather, it tends to become pinned along specific directions selected by the underlying lattice [92]. A related form of lattice-dependent orientational pinning has also been reported in active crystals with turn-towards and turn-away interactions, including models relevant to metal-dielectric Janus colloids, where the direction of the emergent orientational order is controlled by the underlying crystalline lattice [94].

References [65, 89, 92, 93] motivated the investigation presented in the first part of this thesis. The lattice formulation provides a controlled setting in which the ingredients of perception-dependent interactions can be varied separately. In particular, it allows us to ask which properties originate from genuine non-reciprocity, and which are caused by the interplay between the angular anisotropy of the interaction kernel and the symmetry of the underlying lattice. It also enables a direct comparison between the Glauber and Langevin implementations of the dynamics, to separate the effects of different terms that emerge in the latter, and to determine whether pinning acts at the level of individual fluctuations, of the collective orientation, or both. The non-reciprocal XY model therefore offers a minimal framework in which interaction geometry, action-reaction asymmetry, defect dynamics, lattice pinning, and microscopic dynamical rules can be studied systematically in well-controlled conditions.

The original results of this research line, presented in Part I, of this thesis were published in Refs. [1, 2] and are summarized in Sec. 1.2.

1.1.2 Multiscale structure as a computational opportunity

The multiscale character of the collective behavior possibly emerging in statistical systems is not only a conceptual challenge. It is also one of the principal obstacles to the efficient numerical simulation of these systems and in particular in sampling their high-dimensional probability distributions with Markov chain Monte Carlo methods. The difficulty becomes particularly severe close to a continuous phase transition, where the correlation length grows with the linear system size L and local-update dynamics is affected by critical slowing down [95–98]. In this regime, the equilibration and autocorrelation times τ typically scale as

$$\tau \sim L^z, \tag{1.2}$$

where z is the so-called dynamical critical exponent which depends both on the model and on the gross features of the update rule. For algorithms which proceed via updates of local degrees of freedom, in two-dimensional systems z is commonly of order two or larger and may become substantially larger in the presence of disorder or complex free-energy landscapes [99, 100]. Consequently, the computational cost of producing statistically independent configurations (which have to be separated in time at least by an interval τ) increases rapidly as the thermodynamic limit $L \to \infty$ is approached.

Cluster algorithms constitute one of the most successful solutions to this problem. In several unfrustrated spin systems, the Swendsen-Wang and Wolff algorithms identify spatially correlated degrees of freedom and update them collectively, thereby strongly reducing critical slowing down [97, 101, 102]. Extending this strategy to frustrated systems is considerably more difficult. Although several generalizations of cluster methods have been proposed, improvements comparable to those achieved in non-frustrated systems have remained elusive even when such a frustration is weak [103, 104]. Recent studies have related this difficulty to the lack of a general constructive rule capable of identifying clusters underlying efficient sampling of the probability in configuration space [103, 105].

This limitation suggests a different perspective: rather than devising increasingly sophisticated update rules acting on the original microscopic variables, could one instead change the degrees of freedom that are being sampled altogether? Is it possible to find a representation in which the relevant variables possess only short-ranged correlations at every scale, allowing them to be sampled efficiently and the the complete microscopic configuration to be subsequently reconstructed on their bases?

Wavelet representations [106, 107] provide a natural framework for implementing this idea. An orthogonal wavelet transform, in fact, decomposes a microscopic configuration into coarse variables, which retain its long-wavelength structure, and wavelet coefficients, which encode the fluctuations introduced at increasingly finer scales. The wavelet conditional renormalization group (WCRG), introduced in Ref. [108], exploits this decomposition by factorizing the microscopic probability distribution into a coarse-scale distribution and a sequence of conditional distributions for the wavelet coefficients. At each scale, the conditional distribution of the wavelet coefficients is modeled and learned at fixed coarse variables. Here, "learning" means choosing a parameterized model for this conditional distribution and fitting its parameters using a set of training configurations, so that the fitted distribution approximates as closely as possible the conditional distribution from which the data were drawn. Once these conditional distributions have been fitted, new microscopic configurations can be generated recursively: one first samples the coarsest

field and then progressively reconstructs the shorter-wavelength degrees of freedom.

This approach belongs to the broader class of sampling methods based on machine learning, which includes normalizing-flow-augmented Monte Carlo schemes [109], autoregressive variational networks [110], diffusion-based samplers [111, 112], and machine-learning-assisted Monte Carlo algorithms [113]. Within this broader family, WCRG provides a comparatively interpretable alternative to more black-box architectures: the structure of the generative model is explicitly organized according to physical principles. The resulting procedure can indeed be viewed as a renormalization-group transformation carried out in reverse. Long-wavelength correlations are encoded directly in the coarse variables, while fluctuations are progressively restored through the conditional sampling of finer-scale wavelet degrees of freedom. Accordingly, correlations do not need to emerge through the slow propagation of local microscopic updates. Instead, the same hierarchy of scales that underlies critical slowing down is incorporated explicitly into the sampling strategy and turned into a computational resource.

Part II of this thesis investigates whether this strategy can provide an efficient sampling method for frustrated lattice field theories. In particular, we apply WCRG to a system of soft spins (i.e., spins which actually take continuous values) with competing interactions, for which conventional algorithms for the construction of clusters are not applicable. Our objectives are to assess the statistical accuracy of the generated configurations and to establish whether the corresponding multiscale algorithm suppresses the critical slowing down which affects standard local Monte Carlo sampling in the vicinity of critical points.

The original results of this research, presented in Part II, have been reported in Ref. [3] and are summarized in Sec. 1.2.

1.2 Main contributions

This thesis makes original contributions along two complementary research directions. The first concerns the emergence of collective order in lattice systems with orientational degrees of freedom and vision-induced interactions. The second, instead, concerns the development and critical assessment of a novel multiscale strategy for sampling frustrated statistical models.

The first research direction, developed in Part I of this thesis and focused primarily on the non-reciprocal XY model, led to the following main results:

1. Identifying the origin of long-range order in vision-cone models. A central result of this thesis is the demonstration that non-reciprocity is actually not necessary for the emergence of a bona-fide long-range orientational order in two-dimensional XY-like systems with vision-cone interactions. As discussed above, such an ordering was previously reported in Ref. [89] despite the fact that Mermin-Wagner theorem forbids it in the conventional XY model at equilibrium. This raised the question of whether the ordering mechanism originated from non-reciprocity itself or from the spatial anisotropy of the interactions within the system, due to the presence of vision cones.

To answer this question, we introduce two reciprocal equilibrium counterparts of the original non-reciprocal XY model, which also feature orientation-dependent interactions mediated by a vision cone, and which we term the asymmetric reciprocal XY model and the symmetric reciprocal XY model. Both are found to support true long-range order. This controlled comparison establishes

that the decisive ingredient for the emergence of such an order is the coupling between the spin orientation, the vision-cone geometry, and the fixed directions of the underlying lattice. Non-reciprocity determines the nonequilibrium character of the dynamics, but it is not the fundamental origin of the ordered phase.

2. Revealing an ordering mechanism beyond the energetically unfavorable range. We show that long-range order persists even at special vision-cone amplitudes for which the *energetically unfavorable range* (introduced in Ref. [89]) vanishes. According to earlier interpretations, the disappearance of this angular range was expected to restore an XY-like behavior, with quasi-long-range rather than true long-range order.

Our numerical results disprove this expectation. In fact, we demonstrate that the vision-cone interaction continues to couple the collective orientation to the lattice even at these special angles. Moreover, we identify an entropic mechanism that selects a discrete set of preferred orientations: states aligned with particular lattice directions possess a larger set of low-cost fluctuations and are therefore favored at finite temperature.

3. Discovering a thermally-induced increase of order. In the symmetric reciprocal XY model, we highlighted the existence of a reentrant order-by-disorder transition near a vision-cone amplitude of 180°. At very low temperatures, an accidental near-continuous degeneracy allows the system to display quasi-long-range order. Upon increasing the temperature, thermal fluctuations lift this degeneracy and generate an effective fourfold angular anisotropy, driving the system into a genuinely long-range-ordered phase.

Accordingly, orientational order is stabilized rather than destroyed upon increasing the temperature within a finite parameter range. This provides an unusual classical realization of order by disorder in which fluctuations strengthen the coupling between the collective polarization and the underlying lattice.

4. Identifying proactive sensing as a new mechanism in non-reciprocal active dynamics. A central contribution of this thesis is the identification and characterization of two distinct mechanisms, which we term *reactive* and *proactive*, that naturally emerge in the Langevin formulation of non-reciprocal orientational dynamics. The reactive contribution corresponds to the conventional alignment mechanism already considered in previous Langevin implementations of the (non-reciprocal) XY models: an agent (i.e., a spin) adjusts its orientation according to the neighbors that it currently perceives via its couplings.

On the other hand, the proactive term originates from the gradient of the perception kernel with respect to the agent own orientation. Accordingly, it embodies a form of "intelligent" active sensing: agents tend to reorient so as to optimize their sensory input, increasing the effective weight of visible neighbors that are aligned, or at least not strongly misaligned, with their own orientation.

We demonstrate that the reactive and proactive contributions play fundamentally different roles in determining the orientationally ordered state, providing a unified microscopic framework for understanding lattice-induced pinning in non-reciprocal systems. In particular, we show how the two mechanisms control both global and local pinning and how their interplay depends on the specific form of the anisotropic interaction kernel. By combining mean-field arguments, analytical derivations, and extensive numerical simulations, we provide a coherent description that

reconciles previous observations and resolves apparent discrepancies among different microscopic implementations of the dynamics.

The second research direction, mentioned at the beginning this Section and developed in Part II of this thesis, led to the following main results:

5. Overcoming critical slowing down in frustrated lattice systems. The second research line of this thesis addresses the numerical difficulty of sampling frustrated systems near continuous phase transitions. Standard local Monte Carlo algorithms suffer from critical slowing down, while cluster algorithms, that are highly successful solutions to this problem for unfrustrated ferromagnets, are not generally available in the presence of competing interactions.

We demonstrate that the wavelet conditional renormalization group (WCRG) provides a viable and effective alternative. In fact, within this approach the probability distribution is decomposed into a hierarchy of spatial scales, allowing long-wavelength structures and shorter-scale fluctuations to be sampled separately. Applied to the soft-spin two-dimensional biaxial next-nearest-neighbor Ising models (sBNNNI), this construction generates statistically independent configurations with a computational depth that grows only logarithmically with the linear system size.

6. Identifying the systematic limitations of multiscale generative sampling. We also assess systematically the accuracy of WCRG. Although the generated configurations reproduce the single-site field statistics and the two-point correlations measured by the structure factor with high accuracy, we observe a systematic discrepancy in the distribution of the microscopic energy.

We identify the source of this discrepancy in the finite-dimensional parametrization used to model the conditional energies. Exact coarse graining generates, at each scale, all interactions compatible with the symmetries of the microscopic model, whereas a practical implementation retains only a finite subset of them. The fitted conditional distributions therefore approximate the exact ones, and recursive WCRG sampling generates configurations from an approximate hierarchical measure rather than from the original microscopic Gibbs distribution. This approximation may have only small effects on single-site and two-point statistics while still changing the probabilities assigned to complete configurations. The difference becomes visible in the energy distribution because the original microscopic energy is evaluated on configurations drawn from this approximate measure.

1.3 Outlook of the thesis

Several questions remain open, while others have only tentative answers, supported by preliminary evidence. Below we comment on those which we consider the most relevant.

1. Metastability of the long-range ordered phase Recent studies based on Langevin dynamics have suggested that the long-range ordered phase observed in non-reciprocal XY models may be metastable [92, 93]. In the terminology developed in this thesis, the dynamics considered in these studies contains only the reactive contribution. After a possibly very long time, the ordered state may decay into a foam of asters, namely topological defects of charge $+1$ around which the spins

are radially arranged and point approximately toward the defect core. This decay is initiated by the rare nucleation of a first aster, which triggers a subsequent cascade of defect proliferation.

On the other hand, the numerical evidence reported in Ref. [65] shows that the inclusion of what this thesis identifies as the proactive term to the Langevin dynamics (for the von Mises interaction kernel they consider) actually enhances the defect annihilation of defects, suggesting a possible stabilization of the long-range ordered phase. Glauber-type Monte Carlo dynamics, effectively incorporates both reactive and proactive contributions. This observation motivates the conjecture that the same ordered state may be metastable under purely reactive Langevin dynamics but stable under the Glauber dynamics used in Monte Carlo simulations.

We investigate this possibility for the hard vision-cone kernel. Our numerical results offer preliminary support for this scenario. Starting from a configuration in which purely reactive Langevin dynamics rapidly nucleates asters, switching to Glauber dynamics suppresses their further proliferation and promotes the annihilation of the existing defects, allowing long-range order to be recovered. These observations suggest that the long-range ordered phases reported in the phase diagram of Chapter 2 are actually stable states of the Monte Carlo dynamics rather than long-lived metastable states. A systematic comparison of reactive and reactive-plus-proactive Langevin dynamics for smooth interaction kernels will nevertheless be necessary to establish this conjecture conclusively.

2. Reactive and proactive terms beyond lattice models A natural extension of the present work is to investigate the effects of the reactive and proactive terms in off-lattice active matter. This thesis explores in detail their role in the lattice model, but the physical implications in systems of self-propelled particles remain unexplored.

To address this question, we performed preliminary numerical investigations of a two-dimensional model of self-propelled particles moving at a constant speed and interacting via local alignment rules which change the direction of motion. In this setting, vision-cone aligment interactions have already been studied in the active-matter literature [78–82, 114]. However, we argue that the proactive contribution identified in this work on the lattice represents a qualitatively new interaction mechanism. Unlike standard interactions based on the presence of a vision cone, which are typically associated with alignment, avoidance, or following behaviors and which can often be interpreted as passive responses to the local environment, the proactive term represents a genuinely anticipatory strategy: the tendency to align becomes stronger precisely when a neighboring particle is about to leave (or just entered) the field of view of the agent.

Preliminary simulations suggest that the introduction of the proactive term in the dynamics alters qualitatively the collective evolution. In particular, the proactive interaction appears to generate an effective following or pursuit behavior, with particles spontaneously organizing into elongated queues or filamentary structures. This phenomenology resembles forms of steering interactions observed in biological and artificial active systems, but seems to originate from a fundamentally different microscopic mechanism.

Whether such a behavior can be captured within a coarse-grained hydrodynamic description remains an open question. We think that understanding the consequences of proactive interactions in off-lattice active matter constitutes a promising direction for future research.

3. Systematic improvement of the ansatz for wavelet sampling. The results presented in Part II of the thesis highlight the central trade-off of WCRG: the method provides fast multiscale sam-

pling, but the statistical accuracy of the generated configurations depends on how faithfully the parametrized conditional energies reproduce the exact conditional distributions. The interpretability of the construction turns this limitation into a natural route for systematic improvement. Since each conditional energy is associated with a definite reconstruction scale and is written in terms of an explicit operator basis, discrepancies can be traced back to specific levels of the hierarchy and the corresponding ansatz can be enlarged by including additional interactions allowed by symmetry. Future work should quantify how the errors in the generated probability distribution decrease as this operator basis is extended.

1.4 Outline of the thesis

This thesis is organized into two main parts, corresponding to two complementary research directions anticipated in Section 1.1. The first and principal part, comprising Chapters 2–6, investigates non-reciprocal XY systems with vision-induced anisotropic interactions. Its central aim is to understand how reciprocity, the actual network of vision-induced interactions, lattice structure, and microscopic implementation of the dynamics combine to determine possible collective behaviors and ordering. The second part, comprising Chapters 7 and 8, concerns multiscale sampling methods for frustrated statistical systems and, in particular, the wavelet conditional renormalization group. The first part is itself divided into three closely related blocks. Chapters 2 and 3 investigate the origin of long-range order in two-dimensional XY models with vision-cone interactions. Chapters 4 and 5 then examine how the observed ordered states depend on the microscopic dynamical prescription, proposing a quantitative description of their lattice-induced pinning.Chapter 6 finally discusses preliminary results and possible future extensions.

Part I of the thesis on non-reciprocal XY models is organized as follows.

- Chapter 2 introduces the family of XY models with vision-cone interactions considered in this thesis. We define one non-reciprocal (introduced in Ref. [89]) and two reciprocal implementations of the vision-cone interaction: the non-reciprocal XY model, the asymmetric reciprocal XY model, and the symmetric reciprocal XY model. Their phase diagrams are determined numerically by inspecting the finite-size behavior of observables such as the susceptibilities and second-moment correlation lengths. The main result is that bona-fide long-range orientational order occurs not only in the non-reciprocal model, but also in both reciprocal variants. The symmetric reciprocal model additionally displays an extended quasi-long-range-ordered phase and a reentrant ordered region in which a transition from quasi-long-range to long-range order occurs upon *increasing* the temperature.

- Chapter 3 develops a theoretical framework for interpreting these numerical results. In fact, vision-dependent interactions couple the orientation of the spins to the fixed spatial structure of the underlying lattice, thereby generating an effective angular anisotropy for the collective polarization of the spins. We analyze how this spin-lattice coupling depends on the vision-cone aperture and on the precise reciprocal implementation of the interaction. In the asymmetric reciprocal XY model, we compare different manifestations of the fourfold anisotropy and explain why long-range order persists even at special vision angles for which the energetically unfavorable angular range vanishes. In the symmetric reciprocal XY model, we show how redundant bonds generate an accidental continuous degeneracy of low-energy configurations and how thermal fluctuations lift this degeneracy, resulting into the observed

order-by-disorder transition. Together, Chapters 2 and 3 establish that non-reciprocity is not required for the emergence of long-range order: the decisive ingredient is the non-trivial coupling among spin orientation, interaction geometry, and lattice structure.

- Chapter 4 focused on the role of the microscopic dynamics in determining the emergent behavior of the non-reciprocal XY model. The Glauber-type Monte Carlo simulations of Chapter 2 are based on a local, configuration-dependent quantity, termed the *selfish energy*. This quantity, in the Monte Carlo simulation, plays the same role as the Hamiltonian in equilibrium. By considering smooth generalizations of the hard vision cone, we construct, for each spin, a Langevin dynamics from the local angular derivative of this quantity. The resulting drift naturally contains two contributions: (1) the reactive term, which describes the conventional alignment of each spin with the neighbors weighted by their instantaneous interaction strengths; (2) the proactive term, which accounts for the change of these weights caused by a rotation of the evolving spin. We compare the dynamics containing only the reactive term with the one containing both reactive and proactive terms, considering the modulated, sinusoidal, von Mises, and hard vision-cone kernels. With a combination of mean-field arguments and numerical simulations we show that the two different dynamics may lead to different pinning directions of the global polarization of the spins, relative to the square lattice on which the model is defined.

- Chapter 5 provides a quantitative characterization of the resulting pinning mechanisms. We distinguish the restoring force acting on the local orientational fluctuations from that one governing the collective orientation, i.e., the global polarization. By expanding the interaction kernels into angular harmonics, we derive local and global pinning coefficients for the reactive and proactive contributions. The analysis demonstrates that a finite pinning may emerge for the global polarization even when no explicit pinning term is present in the equation of motion of an individual spin, i.e., of the local polarization. We then compare the analytical expressions for the various pinning coefficients with the relaxation rate extracted from the autocorrelation of the global orientation and with the effective pinning inferred from the static structure factor of the local spins. When fluctuations of the local pinning remain sufficiently small, both estimates agree with the analytical global-pinning coefficient. The conditions under which this agreement holds are quantified in detail in the chapter; when local-pinning fluctuations become large, the two numerical estimates may instead depart substantially from the analytical prediction.

- Chapter 6 discusses two preliminary research directions emerging from the reactive-proactive framework. The first concerns the possible metastability of the phase with long-range order of the non-reciprocal XY model. Previous studies have shown that reactive implementations of the dynamics may lead to an eventual decay of such an order through the nucleation and proliferation of asters and shock lines. A preliminary numerical experiment presented here suggest that the proactive response may instead promote defect annihilation, and favor the recovery of orientational order. Although these observations are not yet sufficient to demonstrate the asymptotic stability of the ordered phase, they motivate a systematic comparison of defect nucleation and annihilation under reactive and reactive-plus-proactive dynamics. The second research direction concern understanding the role of active an proactive terms is active matter evolving off lattice. We formulate a Vicsek-like model of self-propelled particles

in which the reactive contribution aligns the headings of neighboring agents, while the proactive contribution generates a steering response associated with the gradient of the visual interaction kernel. Preliminary simulations indicate that this mechanism can suppress spatial dispersion of the particle by producing cohesive, elongated moving aggregates even in the absence of an explicit attractive force. These results suggest that proactivity may influence both the stability of ordered states on a lattice and the emergence of spatial cohesion in mobile active systems.

Part II of the thesis addresses the computational problem of sampling the equilibrium distribution of frustrated statistical systems close to criticality. In particular,

- Chapter 7 is intended as a review of the wavelet conditional renormalization group as a multiscale framework for learning and sampling critical probability distributions. An orthogonal wavelet transform decomposes a microscopic field into a coarse component, which contains its long-wavelength structure, and wavelet coefficients describing the fluctuations associated with successively finer scales. Iterating this transformation yields an exact factorization of the microscopic probability distribution into a coarsest-scale marginal and a hierarchy of conditional distributions. This factorization also defines a generative procedure in which a microscopic configuration is reconstructed recursively from coarse to fine scales. The chapter compares direct and conditional strategies for learning the corresponding energy models and shows that conditioning on the coarse variables improves the learning procedure under both maximum-likelihood estimation and score matching. In the former case, Monte Carlo sampling is restricted to the wavelet degrees of freedom, whose conditional distributions are generally easier to explore than the full critical measure. Finally, we present the score matching techinque applied to conditional distributions, which goes one step further by eliminating the need to sample (via Monte Carlo simulations) the parametrized model during training and thereby makes the learning procedure considerably faster. This computational advantage comes with a potential difficulty: score matching compares gradients of log-probabilities rather than the probabilities themselves, and in general a small score-matching error does not necessarily imply that the learned and target distributions are close. Previous work has shown, however, that an appropriate wavelet decomposition gives the conditional distributions of the wavelet fluctuations the property of strongly log-concavity. This property provides a bound on the discrepancy between the learned and target probabilities in terms of their score-matching error [115, 116]. Consequently, minimizing the score-matching loss guarantees that the learned conditional distribution is close to the original one. This last approach is the one adopted in this thesis: its performance is ultimately assessed numerically in the following chapter.

- Chapter 8 applies the framework developed in the previous chapter to the soft-spin biaxial next-nearest-neighbor Ising model (sBNNNI), a frustrated system for which conventional cluster algorithms do not provide a general remedy for critical slowing down. Using the conditional energy models learned from equilibrium Monte Carlo configurations, WCRG generates new microscopic configurations through the recursive coarse-to-fine procedure described above. Across representative regions of the phase diagram, these generated configurations reproduce the local field distribution and the two-point spin correlations, including along the Ising-like critical line and in the modulated phases. At the Ising-like

critical point, the conditional Monte Carlo chain used at each reconstruction scale decorrelates the corresponding wavelet variables in a number of local Metropolis sweeps that shows no detectable growth with the system size. Since reconstructing a system of linear size L requires $\log_2 L$ successive wavelet levels, the number of sequential sampling steps required to generate an independent configuration grows only logarithmically with L, rather than acquiring the additional power-law factor associated with critical slowing down in local microscopic sampling.

The chapter also identifies the main limitation of the method. Although WCRG-generated samples accurately reproduce the magnetization distribution and the two-point spin correlations, their microscopic-energy distribution differs systematically from that measured in the reference Monte Carlo configurations. Exact coarse graining generates an extended hierarchy of symmetry-allowed interactions in the conditional energies, whereas the practical implementation retains only a finite subset of them. The learned conditional distributions therefore approximate the exact ones, and the resulting errors can accumulate during the recursive reconstruction, producing a detectable bias in the energy distribution. WCRG therefore shifts the main difficulty from efficiently decorrelating critical configurations to representing their conditional probabilities with sufficient accuracy.

Finally, Chapter 9 summarizes the main results of the thesis and discusses and possible directions for future research.

Part I

The non-reciprocal XY model with vision anisotropy

Chapter 2

XY models with vision cone: non-reciprocal and reciprocal implementations

In this chapter, we introduce and numerically investigate a family of generalized two-dimensional XY models in which the orientation of each spin determines which of its nearest neighbors contribute to its dynamics. This perception-like interaction is implemented through a finite vision cone centered on the instantaneous direction of the spin. The resulting coupling between the internal orientational degree of freedom and the spatial structure of the underlying lattice introduces an anisotropy that is absent in the conventional XY model.

Our main objective is to disentangle the effects of this geometrical anisotropy given by the vision cone from those of non-reciprocity for what concerns the emergence of long-range order (LRO) at low temperature. Such order is forbidden in the conventional two-dimensional XY model by the Mermin–Wagner theorem [91]. Instead, below the Berezinskii–Kosterlitz–Thouless [117–119] transition, the model displays quasi-long-range order (QLRO) characterized by algebraically decaying correlations. To this end, we compare three models based on closely related vision-cone geometries but differing in the way the interaction rule is implemented: the non-reciprocal XY model introduced in Ref. [89], and two reciprocal equilibrium counterparts, the asymmetric reciprocal and the symmetric reciprocal XY models.

The numerical results presented in this chapter show that true LRO is not restricted to the non-reciprocal model. The non-reciprocal and asymmetric reciprocal models display remarkably similar phase diagrams. The symmetric reciprocal model also develops LRO, although its phase diagram is substantially different, with an extended QLRO phase and a reentrant transition from QLRO to LRO upon increasing temperature.

The purpose of this chapter is to establish these conclusion at the numerical level. The symmetry arguments that rationalize the observed phase structure and identify the mechanisms selecting the ordered directions are developed separately in Chapter 3. This separation allows the numerical evidence to be presented independently of its subsequent theoretical interpretation.

The rest of the chapter is organized as follows. Section 2.1 introduces the non-reciprocal, asymmetric reciprocal, and symmetric reciprocal vision-cone models, together with the geometrical notion of an energetically unfavorable range. Section 2.2 is dedicated to the numerical methods

used: Section 2.2.1 presents the numerical observables and finite-size diagnostics used to distinguish the different phases, while Sec. 2.2.2 illustrates their behavior in the 4-state and 6-state clock models and in the conventional XY model. This pedagogical review of the techniques used is propaedeutic for Section 2.3, that presents the numerical phase diagrams of the three vision-cone models. Finally, Sec. 2.4 summarizes the main numerical findings and prepares the symmetry-based analysis developed in the following chapter.

2.1 The XY model with vision cone: three variations

In this section, we introduce three variations of the classical XY model with short-range couplings and vision-cone (VC) interactions on the square lattice with periodic boundary conditions (PBCs).

To each lattice site i we associate a classical two-component spin variable $\mathbf{s}_i = (\cos\phi_i, \sin\phi_i)$, i.e., a unit-length vector parameterized by the angle ϕ_i that it forms with a reference direction on the plane. The spin is thought to belong to the same plane as the lattice (a crucial difference from models usually studied on the lattice, where the internal degrees of freedom are not coupled to the spatial structure), and the corresponding VC has an opening angle θ around the spatial direction of the spin. For convenience, the angle ϕ_i is determined with respect to one of the principal axes of the lattice.

The update protocol of the spins in the Monte Carlo (MC) dynamics is defined as follows: let $\Delta_i E$ denote the variation of some energy E (which we do not need to specify at this stage) caused by the spin update $\phi_i \to \phi_i'$ attempted at lattice site i. This proposed update is then accepted with probability

$$w_G\left(\phi_i \to \phi_i'\right) = \frac{1}{2}\left[1 - \tanh\left(\frac{\beta\,\Delta_i E}{2}\right)\right] \tag{2.1}$$

in the Glauber protocol [11]. The details of the implementation of the MC dynamics are reported in App. A.1, where we also discuss the strategy that we adopt in order to avoid that the MC simulations get trapped in long-lived metastable states at low temperatures. While for equilibrium systems this dynamics satisfies detailed balance — and thus leads to a distribution in which a configuration with energy E has probability $p \propto e^{-\beta E}$, where $\beta = 1/T$ is the inverse temperature — the situation is different for non-reciprocal models. In these systems, the dynamics leads to a non-equilibrium steady state, where β^{-1} does not correspond to the actual temperature of the system, and detailed balance is violated. Accordingly, the inherent arbitrariness in defining the update protocol leads to distinct steady states, as discussed in App. A.3.

The features of the actual update protocols for the three models, which we describe below, are summarized in Fig. 2.1.

2.1.1 Non-reciprocal XY model — NRXY

The non-reciprocal XY model (NRXY), represented in Fig. 2.1(a), was introduced in Ref. [89]. In this model, a spin at a lattice site i does not seek to minimize a global energy, but rather its own "selfish" local energy [90] E_i^{NR}, consisting of ferromagnetic interactions with the nearest neighbors within its vision cone. In particular,

$$E_i^{\mathrm{NR}} = -\sum_{j\in\mathcal{N}_i} J_{ij}\left(\phi_i\right)\cos\left(\phi_i - \phi_j\right), \tag{2.2}$$

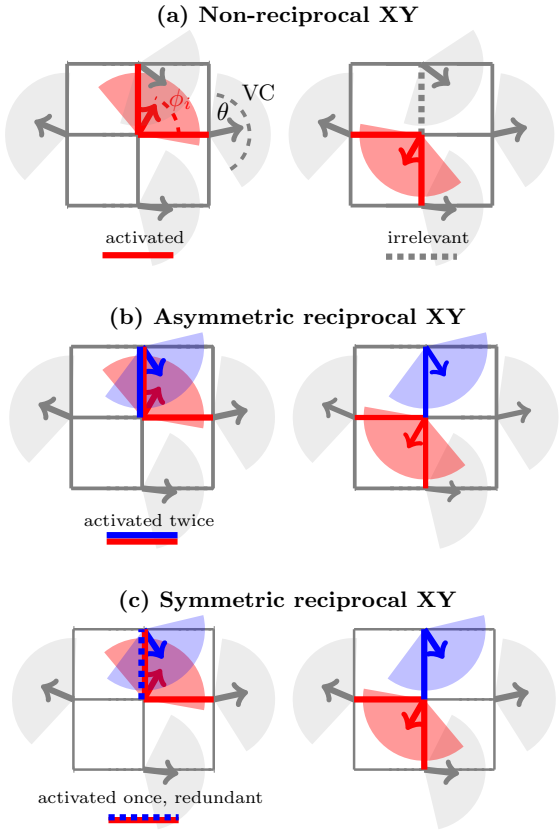


Figure 2.1: Graphical illustration of the three models analyzed in this work. (a) The NRXY model features a *selfish* energy (see Eq. (4.1)), in which the update of the (red) central spin accounts for the neighbors that are within its VC (and that are connected to it by the red bonds, which we term "activated"), while ignoring those outside that cone, even if they "see" the central spin (and to which they are connected by the dashed bond on the right, termed "irrelevant"). This choice results in a non-equilibrium dynamics. (b) In the ARXY model, the update rule depends on the variation of the total energy (see Eq. (2.4)), hence it takes into account both the neighbors of the spin that are seen by it (red bonds), and those who see it (blue bonds). Each bond can thus be not activated (grey), activated once (as in the right panel), or twice (as in the left panel). The resulting dynamics leads to an equilibrium state. (c) In the SRXY model, the update also takes into account the neighbors of the spin that are seen by it (red bonds), and those who see it (blue bonds) — see Eq. (2.5). However, a bond cannot be activated twice, and thus the virtual second activation is *redundant* (dashed blue and solid red bond). The resulting state is at equilibrium.

where $j \in \mathcal{N}_i$ indicates that the sum runs over all the nearest neighbors j of spin i, and the (non-symmetric) couplings $J_{ij}(\phi_i)$ take into account the vision cone as

$$J_{ij}(\phi_i) = \begin{cases} J, & \text{if } \min\left\{360^\circ - \Delta_{ij}, \Delta_{ij}\right\} \leq \theta/2, \\ 0, & \text{otherwise.} \end{cases} \tag{2.3}$$

Here $\Delta_{ij} = \left|\phi_i - \vartheta_{ij}\right|$ is the angle that the spin at site i forms with the lattice vector connecting site i with the nearest-neighbouring sites j, characterized by the angle $\vartheta_{ij} \in \{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$. In practice, $J_{ij}(\phi_i)$ equals J if the neighbor at site j is within the vision cone of the spin at site i, while it vanishes otherwise.

In what follows, the bonds connecting spin i to spin j for which $J_{ij}(\phi_i)$ does not vanish will be referred to as "activated". If, instead, $J_{ij}(\phi_i)$ vanishes but $J_{ji}(\phi_j)$ does not, the bond will be said to be "irrelevant" for the dynamics of the spin i. This occurs if the spin at site i does not see the spin

at site j, but the latter sees the former.

In the NRXY model, the selfish energy E_i^{NR} in Eq. (4.1) is then used to determine the probability to accept an attempted MC move, according to the update protocols in Eq. (2.1) with $\Delta_i E \mapsto \Delta_i E_i^{\mathrm{NR}}$. This implies that the decision to change the direction of spin i depends solely on interactions with the spins that i can see, without accounting for the energy cost associated with neighboring spins that see i but are not seen by it. As a consequence, this model operates out of equilibrium. Its characteristics will be thoroughly examined in Sec. 2.3.1, where we build upon the study conducted in Ref. [89].

2.1.2 Asymmetric reciprocal XY model — ARXY

For the asymmetric reciprocal XY model (ARXY), the update probabilities in Eq. (2.1) are calculated on the basis of the difference $\Delta_i E \mapsto \Delta_i E^{\mathrm{AR}}$ of the global energy obtained by summing all the non-reciprocal ("selfish") energies in Eq. (4.1), i.e.,

$$\begin{aligned} E^{\mathrm{AR}} &= \frac{1}{2}\sum_i E_i^{\mathrm{NR}} \\ &= -\frac{1}{2}\sum_{\langle ij\rangle}[J_{ij}(\phi_i) + J_{ji}(\phi_j)]\cos(\phi_i - \phi_j), \end{aligned} \tag{2.4}$$

where the sum runs over all pairs $\langle ij\rangle$ of nearest-neighbor sites i and j on the lattice. The factor $1/2$ ensures that both the NRXY and ARXY models reduce to the standard XY model with the same coupling when the vision cone is removed, i.e., for $\theta = 360^\circ$.

Note that, while the individual couplings $J_{ij}(\phi_i)$ and $J_{ji}(\phi_j)$ are asymmetric (i.e., $J_{ij}(\phi_i) \neq J_{ji}(\phi_j)$), the total coupling between ϕ_i and ϕ_j — represented by the terms within the square brackets in Eq. (2.4) — is actually symmetric. Therefore, unlike the NRXY model, the dynamics of the spin at site i is influenced both by the spins it sees and by those that see it. This is depicted in Fig. 2.1(b), which illustrates how a bond (ij) can be activated zero, one, or two times, depending on $J_{ij}(\phi_i) + J_{ji}(\phi_j)$ being 0, J, or $2J$. We refer to this model as *asymmetric* because the lattice couplings are asymmetric, even though the eventual interaction between the neighboring spins in Eq. (2.4) (which enters the MC update) is symmetric. In this way, we distinguish it from the SRXY model introduced in the next section, where the couplings are symmetric. We emphasize in particular that the asymmetry of J_{ij} implies that $\Delta_i E^{\mathrm{AR}} \neq \Delta_i E_i^{\mathrm{NR}}$, as pointed out in Ref. [120]: this difference underscores the inherently non-equilibrium nature of NR systems, whose dynamics does not follow from a principle of global energy minimization. On the contrary, the ARXY is a proper statistical equilibrium model, in which a *bona-fide* temperature β^{-1} is defined. Its characteristics are examined in Sec. 2.3.2.

2.1.3 Symmetric reciprocal XY model — SRXY

Another reciprocal version of the XY model with VC is the symmetric reciprocal XY model (SRXY), whose couplings are devised to be symmetric while still retaining a VC. This is achieved by defining the total energy as

$$E^{\mathrm{SR}} = -\sum_{\langle ij\rangle} \tilde{J}_{ij}\left(\phi_i, \phi_j\right)\cos\left(\phi_i - \phi_j\right), \tag{2.5}$$

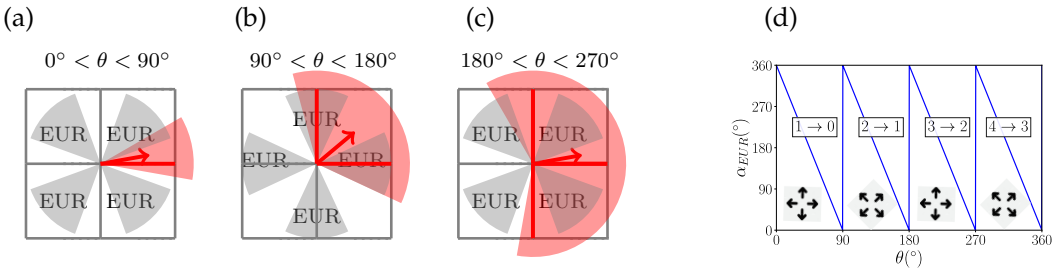


Figure 2.2: Illustration of the EUR. (a–c) When the direction of a spin enters the EUR (gray zone), the number of visible nearest neighbors decreases by one compared to its maximal possible value determined by the angular opening θ of the VC (red fan). Correspondingly, the number of activated links (red) changes (a) from 1 to 0 for $\theta \in [0°, 90°]$, (b) from 2 to 1 when $\theta \in [90°, 180°]$, (c) from 3 to 2 when $\theta \in [180°, 270°]$, and from 4 to 3 when $\theta \in [270°, 360°]$ (not shown). Considering as a reference a state in which all spins are aligned in the same direction, this results in a discontinuous increase in energy, making the configuration energetically unfavorable, whence the name or the region. The angular position of the EUR depends on the value of θ. (d) Size α_{EUR} of the EUR as a function of θ, exhibiting a periodicity of 90°. For each period, the black arrows indicate the spatial directions that are *favored* by a collective alignment of the spins on the lattice, resulting in the lowest possible energy. Upon varying θ, the corresponding loss of a nearest neighbor, from n to $n-1$, with $n = 1, 2, 3$, and 4 (indicated by $n \to n-1$ in the figure) has different effects depending on n (even for the same size α_{EUR} of the EUR), which are discussed in Sec. 3.2.1.

with

$$\begin{aligned}\frac{\tilde{J}_{ij}\left(\phi_i, \phi_j\right)}{J} = 1 + \operatorname{sgn}\big[&\Theta\left(\phi_i - \vartheta_{ij} + \theta/2\right) \\ &- \Theta\left(\phi_i - \vartheta_{ij} - \theta/2\right) + \Theta\left(\phi_j - \vartheta_{ji} + \theta/2\right) \\ &- \Theta\left(\phi_j - \vartheta_{ji} - \theta/2\right)\big],\end{aligned} \tag{2.6}$$

where Θ is the Heaviside function, while ϑ_{ij} are defined after Eq. (2.3). These couplings are symmetric, $\tilde{J}_{ij}\left(\phi_i, \phi_j\right) = \tilde{J}_{ji}\left(\phi_j, \phi_i\right)$. In practice, as shown in Fig. 2.1(c), this means that if *at least* one of the two nearest-neighbouring spins at sites i and j is looking at the other (i.e., one has the other within its vision cone), then the corresponding bond is activated in the sense that $\tilde{J}_{ij}$ takes the finite value J. If the spins are both looking at each other, the bond is still activated only once: this is the main difference with respect to the ARXY model. Accordingly, one of the two activations can be deemed *redundant*. If neither of the spins is looking at the other, then the bond is not activated and $\tilde{J}_{ij} = 0$. The resulting dynamics of this model is an equilibrium one. However, the intricate dependence of its couplings on the configuration of the system results in a behavior that is significantly distinct from those of the previously discussed NRXY and ARXY models, as we shall show in Sec. 2.3.3.

2.1.4 Energetically unfavorable range — EUR

Before proceeding to the analysis of the various models introduced above, here we recall the concept of energetically unfavorable range (EUR), introduced in Ref. [89], that will prove crucial for understanding our results. The coupling $J_{ij}(\phi_i)$ with VC given in Eq. (2.3), shared by both the NRXY and ARXY models, implies that the number of nearest neighbors that are visible to the spin i varies between the two consecutive integers $\lfloor \theta/90° \rfloor$ and $\lceil \theta/90° \rceil$ upon varying the spin orientation ϕ_i, where $\lfloor x \rfloor = \max(y \in \mathbb{Z}, y \leq x)$ and $\lceil x \rceil = \min(y \in \mathbb{Z}, y > x)$ are the floor and ceiling functions, respectively. These values follow from the choice of a square lattice: if the system

were defined, e.g., on a honeycomb (triangular) lattice, then the number of visible spins would vary between $\lfloor \theta/120^\circ \rfloor$ and $\lceil \theta/120^\circ \rceil$ ($\lfloor \theta/60^\circ \rfloor$ and $\lceil \theta/60^\circ \rceil$). The EUR corresponds to the set of values of ϕ_i for which the number of visible neighbors takes its minimal value. This is illustrated in Fig. 2.2(a)–(c) for three different values of the amplitude θ of the VC (red fan). The angular extension α_{EUR} of the EUR varies as a function of θ with a periodicity of 90°, as shown in Fig. 2.2(d); its functional form reads

$$\alpha_{\mathrm{EUR}} = 360^\circ - 4(\theta - 90^\circ \lfloor \theta/90^\circ \rfloor). \tag{2.7}$$

Moreover, both the extension and the centering of the EUR change discontinuously at multiples of 90° (on the square lattice). To see this, consider, for example, a vision cone as in Fig. 2.2(a), with $\theta \in [0^\circ, 90^\circ]$. Here, the central red spin can see one neighbor only if its arrow points approximately along one of the four lattice directions, i.e., if $|\phi_i - n \cdot 90^\circ| < (360^\circ - \alpha_{\mathrm{EUR}})/8$ for some value of $n \in \{0, 1, 2, 3\}$. Conversely, the spin sees no neighbors if it points away from the spatial directions in which the neighbors are located, i.e., if its orientation differs approximately by 45° from those of the lattice directions, with $|\phi_i - n \cdot 90^\circ - 45^\circ| < \alpha_{\mathrm{EUR}}/8$. Note that α_{EUR} quantifies the angular range within which the number of visible neighbors vanishes. When θ approaches 90° from below, the spin sees almost always one neighbor regardless of the spatial direction in which it points, so that $\alpha_{\mathrm{EUR}} \to 0^\circ$. Once θ exceeds 90°, as in Fig. 2.2(b), the spin can see two neighbors only if it points approximately 45° away from the lattice directions. This situation is actually opposite to that in the previous range. Consequently, the *centering* of the EUR changes discontinuously whenever θ crosses 90° and its integer multiples. In addition, α_{EUR} assumes its maximum value 360° immediately after $\theta = 90^\circ$, and then it decreases as θ increases — until θ reaches the value $\theta = 180^\circ$, where another analogous discontinuity occurs, as summarized in Fig. 2.2(d).

This implies that a perfectly aligned configuration of the spins maximizes the number of activated bonds and minimizes the energy when the system is oriented either along the lattice directions or along the lattice directions plus 45°, depending on the value of θ. These preferential orientations are indicated by the arrows close to the horizontal axis in Fig. 2.2(d). We anticipate here that the existence of this EUR actually shapes the phase diagrams of the NRXY and ARXY models, shown in, c.f., Figs. 2.4 and 2.5, respectively. Note that the impact on the system of losing sight of one of the n nearest neighbors by entering the $n \to n-1$ EUR depends not only on the size of this region, but also on how many other neighbors remain visible, i.e., on the value of θ. This results in significantly different outcomes, which will be discussed in Sec. 3.2.1.

Finally, we point out that the notion of EUR does not carry over straightforwardly to the SRXY model, because the corresponding coupling $\tilde{J}_{ij}(\phi_i, \phi_j)$ (see Eq. (2.6)) depends on the two angles $\phi_{i,j}$, and thus the activation of the four bonds around each spin depends on the value of its angle and of those of the 4 neighboring spins. In this case, the number of activated bonds equals an integer between $\lfloor \theta/90^\circ \rfloor$ and 4. This leads to significant differences in the phase diagram, as will be shown in Sec. 2.3.3.

2.2 Methods

The aim of this study is to characterize the phase diagram of the models introduced in the previous Section, focusing on identifying regions that exhibit LRO, QLRO, and DO, as well as understanding the transitions between these phases. This is achieved via MC simulations of the dynamics, which allow one to sample and analyze the observables we are interested in. Before conducting such

analysis, we recall the results of MC simulations of systems with well-known behavior, specifically the q-state clock model with $q = 4, 6$ and the XY model, which exhibit the three aforementioned phases, and transitions of second and infinite order separating them.

Details on the numerical implementation of the MC dynamics are provided in App. A.1.

2.2.1 Observables

The three phases mentioned above, i.e., LRO, QLRO, and DO, can be identified on the basis of the behavior of two key observables: the modulus of the magnetization and the correlation length. The (average) magnetization of the model is given by the vector

$$\mathbf{m} = \frac{1}{N}\sum_{i=1}^{N}\mathbf{s}_i = \left(\frac{1}{N}\sum_{i=1}^{N}\cos\phi_i, \frac{1}{N}\sum_{i=1}^{N}\sin\phi_i\right), \tag{2.8}$$

where the sum runs over all lattice sites of the model, which are $N = L^2$ in total on a square lattice with linear size L. Its modulus $m = |\mathbf{m}|$ is then given by

$$m = \frac{1}{N}\sqrt{\left(\sum_{i=1}^{N}\cos\phi_i\right)^2 + \left(\sum_{i=1}^{N}\sin\phi_i\right)^2}. \tag{2.9}$$

From the correlation function

$$G(r = |i-j|) = \langle \mathbf{s}_i \cdot \mathbf{s}_j \rangle = \langle \cos(\phi_i - \phi_j) \rangle, \tag{2.10}$$

we define its connected part $G_c(r) = G(r) - m^2$ and the exponential correlation length ξ as

$$\xi = \lim_{r\to\infty} \frac{r}{-\log G_c(r)}. \tag{2.11}$$

In an infinitely extended system, LRO phases are characterized by a non-zero magnetization m and an exponentially decaying correlation function, corresponding to a finite correlation length ξ. DO phases have zero magnetization $m = 0$ and also exhibit finite correlation lengths ξ. In contrast, QLRO phases, as well as critical points (CPs), feature a power-law decaying correlation function $G_c(r)$, implying an infinite correlation length ($\xi \sim L$ for finite system sizes), with a vanishing magnetization m. In the case of spatial dimension $d = 2$ on which we focus here, the above can be summarized as

$$G_c(r) \sim \begin{cases} e^{-r/\xi}, & \text{for LRO, DO,} \\ 1/r^{\eta}, & \text{for QLRO,} \end{cases} \tag{2.12}$$

$$m = \begin{cases} \text{const.}, & \text{for LRO,} \\ 0, & \text{for QLRO, DO.} \end{cases} \tag{2.13}$$

Notably, in the case of the XY model at low temperatures, the magnetization is subject to strong finite-size corrections due to slow spin-wave scaling [121, 122], described by

$$m \sim N^{-T/(8\pi J)}. \tag{2.14}$$

This scaling makes it difficult to remove finite-size effects in simulations. Furthermore, determining the exponential correlation length ξ from Eq. (2.11) is practically challenging for finite geometries. Instead, we will use the second-moment correlation length, defined as

$$\xi_2^2 = \frac{\int d^d r\, r^2\, G(r)}{\int d^d r\, G(r)} = -\frac{1}{\hat{G}(\mathbf{k})} \frac{\partial^2 \hat{G}(\mathbf{k})}{\partial k^2}\bigg|_{k=0}, \tag{2.15}$$

where $\hat{G}(\mathbf{k})$ is the Fourier transform of the correlation function:

$$\hat{G}(\mathbf{k}) = \frac{1}{N} \left\langle \sum_{\mu=x,y} \left| \sum_{i=1}^{N} s_{i\mu}\, \mathrm{e}^{i\mathbf{k}\cdot\mathbf{r}_i} \right|^2 \right\rangle, \tag{2.16}$$

with $s_{ix} = \cos\phi_i$ and $s_{iy} = \sin\phi_i$. On the lattice, ξ_2^2 in Eq. (2.15) is expressed as [123]

$$\xi_L^{(2)} = \frac{1}{2\sin(k_m/2)} \sqrt{\frac{\hat{G}(0)}{\hat{G}(\mathbf{k}_m)} - 1}, \tag{2.17}$$

where $k_m = 2\pi/L$, and $\mathbf{k}_m = (k_m, 0)$. Since in the LRO phase $\hat{G}(0) \sim L^d$ and $\hat{G}(\mathbf{k}_m) \sim L^0$, the behavior of $\xi_L^{(2)}/L$ in $d = 2$ is given by [124]

$$\xi_L^{(2)}/L \sim \begin{cases} L, & \text{for LRO,} \\ \text{const.}, & \text{for QLRO or CPs,} \\ 1/L, & \text{for DO.} \end{cases} \tag{2.18}$$

Accordingly, inspecting the behavior of $\xi_L^{(2)}/L$ upon increasing L turns out to be a reliable approach for identifying the three phases.

Another useful observable for characterizing phase transitions is the variance of the magnetization, defined as

$$\chi_m = \frac{N}{T} \left(\langle m^2 \rangle - \langle m \rangle^2 \right), \tag{2.19}$$

which is equal to the magnetic susceptibility under equilibrium conditions. Similarly, the variance of the internal energy

$$C_v = \frac{1}{NT^2} \left(\langle E^2 \rangle - \langle E \rangle^2 \right) \tag{2.20}$$

is equal to the specific heat at equilibrium.

In order to distinguish between the three phases, it often proves useful to consider also the vectorial magnetization in Eq. (2.8). This quantity is sampled along a MC run and visualized in a scatter plot, as an informative approximation of its actual distribution. In fact, the analysis of this distribution provides insight into the phase behavior of the system, which will be discussed in Sec. 2.2.2.

Finally, to quantify the irreversibility of the dynamics in the cases where the system is out of equilibrium, we consider the mean entropy production rate (EPR) per spin associated with a certain realization of the process $\{\phi_i(t)\}_{i,t}$, defined as [125–128]

$$\mathrm{EPR} = \left\langle \ln \frac{w\left(\phi_i(t) \to \phi_i(t+dt)\right)}{w\left(\phi_i(t+dt) \to \phi_i(t)\right)} \right\rangle / dt. \tag{2.21}$$

Here $w\left(\phi_i(t) \to \phi_i(t+dt)\right)$ is the transition rate for updating the angle ϕ_i from the value $\phi_i(t)$ at

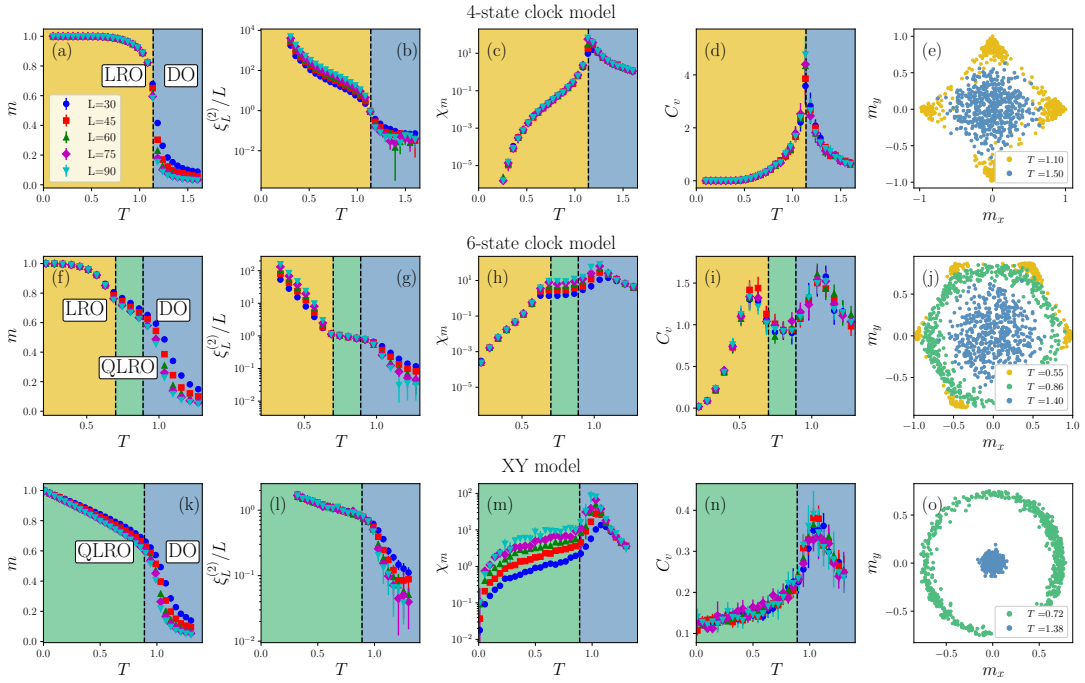


Figure 2.3: Temperature dependence of the observables defined in Sec. 2.2.1 for three benchmark models, namely: (a)–(e) the 4-state clock model, (f)–(j) the 6-state clock model, and (k)–(o) the XY model. The quantities m, $\xi_L^{(2)}/L$, χ_m, and C_v are shown for various system sizes L (indicated by the legend in panel (a)), while scatter plots of $\mathbf{m}$ are presented at selected temperatures for a fixed size $L = 90$. Note that the plots of $\xi_L^{(2)}/L$ (exhibiting the distinct behaviors outlined in Eq. (2.18)) and $\mathbf{m}$ are particularly effective for distinguishing the three phases: LRO (yellow), QLRO (green), and DO (blue). See Sec. 2.2.2 for a detailed description of the observed behaviors.

time t to the value $\phi_i(t + dt)$ at time $t + dt$ (where $dt = 1$ in the MC simulation), as described by the Glauber algorithm of Eq. (2.1). The EPR is computed by averaging the logarithm of the ratio of forward to reverse transition rates over the entire MC history. This measure provides valuable information about the deviation of the dynamics of the system from equilibrium.

2.2.2 Phase diagrams of reference models

Here we recall the behavior of the quantities defined in Sec. 2.2.1 for three paradigmatic models with well-known properties that we will use for comparison: the 4- and 6-state clock models and the XY model, all in spatial dimension $d = 2$. These models are all characterized by the same formal expression for the energy E, i.e.,

$$E = -J \sum_{\langle ij \rangle} \cos(\phi_i - \phi_j), \tag{2.22}$$

where the angular variable ϕ_i takes different values depending on the model. For the XY model, ϕ_i is continuous between 0 and 2π, while for the clock models, ϕ_i takes discrete values $\phi_i = 2\pi k/n$ for $k = 0, 1, \ldots, n-1$, with $n = 4$ or $n = 6$ for the 4-state and 6-state clock models, respectively.

The 4-state clock model. — The behavior of the 4-state clock model is summarized in Fig. 2.3(a)–(e), which shows the observables m, $\xi_L^{(2)}/L$, χ_m, and C_v for different system sizes L, and as a function of temperature, as well as $\mathbf{m} = (m_x, m_y)$ for two temperatures within the two phases. (Here-

after, $J = 1$ is assumed in all models.) This model has an internal $\mathbb{Z}_4$ symmetry and it undergoes a second-order phase transition from a LRO phase to a DO phase, belonging to the Ising universality class [129]. The magnetization m in panel (a) thus goes from a finite value at low temperatures to zero as the temperature is increased beyond the transition at $T_c = 1/\log(1+\sqrt{2}) \simeq 1.13$ [130], indicated by the dashed vertical line in all panels. The curves of the second-moment correlation length $\xi_L^{(2)}$ normalized by L, shown in panel (b), are size-dependent in the LRO and DO phases (scaling, respectively, as L and $1/L$ upon increasing L at fixed temperature, see Eq. (2.18)), and cross at the transition point.

The susceptibility in panel (c) is size-independent in both the LRO and DO phases but exhibits a size-dependent peak at the transition point. The same holds for the specific heat in panel (d).

Finally, the scatter plot of the vectorial magnetization $\mathbf{m}$ in panel (e) shows that, within the LRO phase, most of the spins align along one of the four allowed directions, with configuration changes occurring only due to finite-size effects. In contrast, in the DO phase, the magnetization no longer favors any of the four directions, forming an isotropic cloud around the origin.

The 6-state clock model. — Next, the behavior of the 6-state clock model is summarized in Fig. 2.3(f)–(j). With its internal $\mathbb{Z}_6$ symmetry, this model exhibits two infinite-order transitions: one from the LRO to the QLRO phase, and the other from the QLRO to the DO phase [131]. The corresponding transition temperatures, $T_c^{(1)} \simeq 0.7$ and $T_c^{(2)} \simeq 0.89$, indicated by the vertical dashed lines, are those numerically estimated in Ref. [132]. The magnetization m shown in panel (f) is finite in the LRO phase, remains finite in the QLRO phase due to finite-size effects — similarly to what happens in the XY model (see Eq. (2.14)) — while it vanishes in the DO phase. The curves of the second-moment correlation length $\xi_L^{(2)}$, normalized by L and shown in panel (g), scale as L in the LRO phase, collapse onto the same value within the finite range of temperatures corresponding to the QLRO phase, and then separate again, scaling as $1/L$, in the DO phase. These different scalings are shown, for selected temperatures, in Fig. A1(a) in App. A.2. The magnetic susceptibility in panel (h) does not scale with the system size in the LRO and DO phases, but it features a size dependence in the QLRO phase. The position of the scale-dependent peak at the transition from QLRO to DO tends logarithmically to the infinite-size transition temperature $T_{\mathrm{BKT}} \simeq 0.89$ [132].

The specific heat C_v in panel (i) shows size-independent peaks that do not correspond directly to the transition temperature, as it is also observed in other infinite-order transitions such as the one of the XY model [133]. The distribution of the vectorial magnetization $\mathbf{m}$ in panel (j) reveals that, in the LRO phase, $\mathbf{m}$ tends to align along one of the six available directions, with fluctuations due to finite-size effects. In the QLRO phase, the system no longer prefers any particular direction, and the magnetization remains finite due to finite-size effects, resulting in the observed ring-shaped distribution. In the DO phase, the distribution forms a cloud around the origin.

The XY model. — Finally, the behavior of the XY model is summarized in Fig. 2.3(k)–(o). With its internal $O(2)$ symmetry, this model exhibits a single infinite-order transition, from the QLRO phase to the DO phase [118, 119]. The transition temperature, indicated with a dashed line, is numerically estimated in Ref. [134]. The magnetization m in panel (k) remains finite in the QLRO phase due to significant finite-size effects, as described by Eq. (2.14), and then it decreases to zero in the DO phase. Here $\xi_L^{(2)}$, normalized by L and shown in panel (l), collapses onto the same value in the QLRO phase and then separates as $1/L$ in the DO phase. The magnetic susceptibility in panel (m) shows a size dependence in the QLRO phase while it is size-independent in the DO phase. The specific heat in panel (n) presents a size-independent peak that does not occur at the transition temperature. The distribution of the vector magnetization in panel (o) forms a ring in

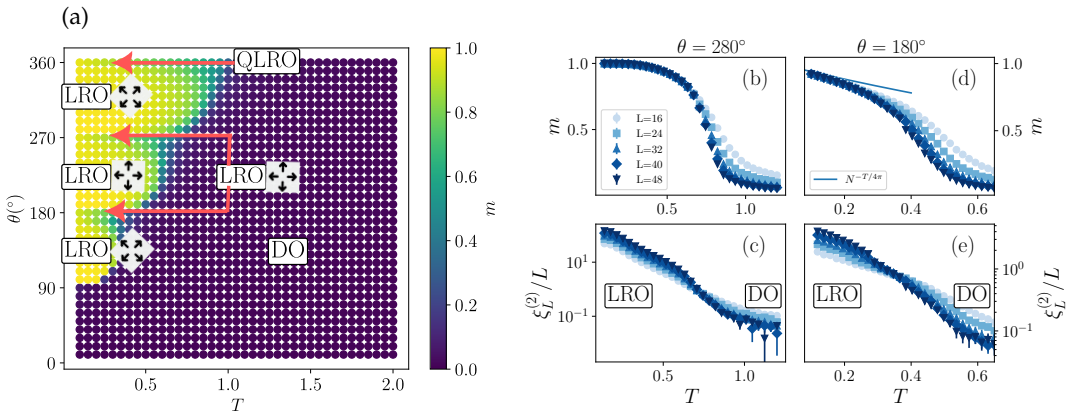


Figure 2.4: Results of the MC simulations for the NRXY model in the non-equilibrium steady state. (a) The magnetization m as a function of temperature T and VC angle θ reveals a phase diagram characterized by three distinct lobes. At low temperatures and for $\theta > 90°$, a LRO phase (in yellow) is observed, with the spins showing a collective alignment along one of the spatial directions indicated by the black arrows. At higher temperatures, the system undergoes a transition into a DO phase (dark blue). (b)–(c) Magnetization m and ratio $\xi_L^{(2)}/L$ as functions of T, for a vision cone angle $\theta = 280°$, showing a clear transition between LRO and DO phases. (d)–(e) Magnetization m and $\xi_L^{(2)}/L$ as functions of T, for $\theta = 180°$. In panel (d), the scaling of m described by Eq. (2.14) for the XY model is shown for $L = 48$ (solid blue line), exhibiting consistency with the numerical data at low temperature (symbols). However, the behavior of the data for $\xi_L^{(2)}/L$ in panel (e) — qualitatively similar to panel (c) — indicates that the model does not actually develop QLRO at low temperatures but that it undergoes a transition between a LRO and a DO phase, as in the previous case.

the QLRO phase, with a finite radius due to the significant finite-size effects. Finally, in the DO phase, the distribution becomes an isotropic cloud centered around the origin.

In the next section, the comparison with the behaviors described above will be used to analyze the phase diagram of the lattice models with VC introduced in Sec. 2.1.

2.3 Numerical results

Having established the finite-size diagnostics used to distinguish LRO, QLRO, and DO, we now apply them to the three vision-cone models introduced in Sec. 2.1.

2.3.1 The non-reciprocal XY model

We begin by analyzing the NRXY model of Sec. 2.1.1, previously studied in Ref. [89]. Our goal is to further investigate certain aspects of the model, highlighting that some of its features are equivalent to those of its reciprocal counterpart, that is, the asymmetric reciprocal model ARXY, whose analysis will be the subject of Sec. 2.3.2. In particular, we will show that the existence of a phase with LRO is due to the interaction being characterized by a vision cone, rather than its non-reciprocal nature.

The effects of the chosen update protocol on the stationary state are discussed in App. A.3. In general, the MC dynamics drives the system into a non-equilibrium stationary state with a non-vanishing EPR; we refer to App. A.4 for a detailed analysis of the EPR for different points in the phase diagram of the model.

In Fig. 2.4(a), we report the magnetization m in the steady state as a function of the temperature

T and the VC angle θ, for a system of size $L = 100$. In particular, this plot reveals that the phase diagram of the model features a structure with three lobes, characterized by an enhanced m, which start and end at angles that are multiples of 90°. This is a consequence of the periodicity of the occurrence of EUR, highlighted in Fig. 2.2. Within these lobes, a LRO phase is present, with the spins preferentially and collectively aligned along one of the four directions indicated by the black arrows in Fig. 2.4(a) for each lobe. An account of the origin of these lobes and why there is no LRO for $\theta < 90^\circ$ is given in the supplemental material of Ref. [89]. For $\theta = 360^\circ$, the standard XY model is recovered and we find a QLRO phase at low temperature.

In Figs. 2.4(b)–(e), we report m and $\xi_L^{(2)}/L$ as functions of the temperature T for various system sizes, and for two representative values of $\theta = 280^\circ$ and 180°. As already suggested in Ref. [89], for all values of the VC angle that have a finite EUR, such as $\theta = 280^\circ$, we expect a LRO phase. Indeed, the magnetization m in panel (b) approaches 1 for $T \to 0$, remaining flat at low temperatures, and goes to zero above the transition temperature. Similarly, the curves of $\xi_L^{(2)}/L$ for the various values of L, reported in panel (c), intersect at a single point, marking the transition from LRO to DO. In these two phases, the curves of $\xi_L^{(2)}/L$ scale as L and $1/L$, respectively, consistently with Eq. (2.18).

The case with $\theta = 180^\circ$ (similar to $\theta = 270^\circ$) — suggested to exhibit QLRO at low temperatures in Ref. [89] due to disappearance of EUR — is more nuanced. In panel (d) of Fig. 2.4, we observe that at low temperatures the dependence of the magnetization on T is not flat, but instead follows the scaling behavior of the XY model described by Eq. (2.14) (with the coupling J effectively halved because, for $\theta = 180^\circ$, only half of the bonds are activated). However, in panel (e), the curves of the temperature dependence of $\xi_L^{(2)}/L$ for various L do not collapse on a single curve at low temperatures, as expected in the presence of QLRO, but instead they intersect at a single point. For temperatures below and above the temperature of this intersection point, it actually turns out that $\xi_L^{(2)}/L$ scales as L and $1/L$, respectively, as shown in Fig. A1(b) of App. A.2. This behavior, combined with the observation that the system prefers alignment along the lattice directions at low temperatures (indicated by the black arrows in Fig. 2.4(a)), suggests that, for these values of θ, the phase has LRO at low temperatures despite a vanishing EUR. This phenomenon, also observed for the ARXY model, will be explained in Sec. 3.2.2.

2.3.2 The asymmetric reciprocal XY model

We next consider the ARXY model introduced in Sec. 2.1.2. Our purpose here is to determine its numerical phase diagram and compare it directly with that of the NRXY model. Although the two models share the same underlying directed VC couplings, the ARXY dynamics derives from a global energy and samples an equilibrium Boltzmann distribution. Their comparison therefore provides a direct test of whether the ordered phases observed in the NRXY model require non-reciprocal dynamics.

The behavior of the magnetization m of this model as a function of temperature T and the VC angle θ is reported in Fig. 2.5(a). In particular, we note that it is qualitatively very similar to that of the NRXY, presented in Fig. 2.4(a). This is rather remarkable, since in the ARXY discussed here the MC averages of the magnetization (and of the other observables defined in Sec. 2.2.1) provide accurate estimates of the averages sampled from the Boltzmann distribution, which characterizes equilibrium stationary states, whereas the NRXY is out of equilibrium. However, at a closer inspection, some quantitative differences emerge in the temperature dependence of the magnetization, as we show in Fig. A2(b) of App. A.3.

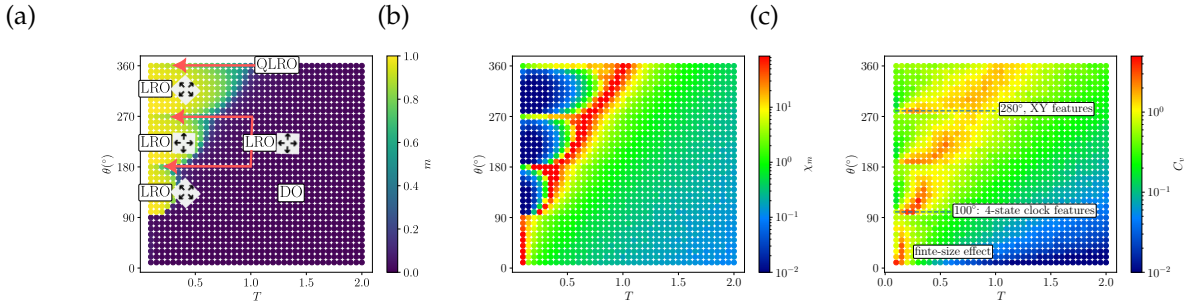


Figure 2.5: Results of the MC simulations for the equilibrium state of the ARXY model obtained for a system size $L = 100$. (a) The behavior of the magnetization m is characterized by the presence of three lobes, qualitatively indistinguishable from those of the NRXY model (see Fig. 2.4(a)), with the LRO phase present within the lobes and the DO phase outside them. (b) The low values of the magnetic susceptibility χ_m within the three lobes confirms the presence of the LRO phase (see the discussion in, c.f., Sec. 2.3.2). However, from this diagram it is not possible to conclude whether the system exhibits LRO or QLRO at low temperatures for $\theta = 180°$ and $270°$. The susceptibility χ_m as a function of temperature features a peak (in red) at the transition that separates the LRO phase from the DO phase for $\theta > 90°$, while it increases monotonically upon decreasing the temperature towards $T = 0$ for $\theta < 90°$. (c) The specific heat C_v shows peaks in correspondence of the transitions from LRO to DO when $\theta > 90°$, and at low temperatures for $\theta < 90°$. The origin of these low-temperature peaks is addressed in App. A.5.

The susceptibility χ_m as a function of T and θ, obtained for a system of size $L = 100$, is shown in Fig. 2.5(b). We observe that χ_m is small in the LRO phase, while it increases upon increasing the temperature T for a fixed value of θ, reaching its maximum along the boundaries of the three lobes identified in panel (a), which separate the LRO phase from the DO phase. The higher values of the susceptibility for $\theta = 180°$ and $270°$, compared to those observed at similar temperatures for angles immediately above and below, raise the question of whether QLRO is present here, as it is for $\theta = 360°$. However, as demonstrated for the NRXY model, this is not the case, as will be further elaborated in Sec. 3.2.2.

The behavior of the specific heat C_v for this model is shown in Fig. 2.5(c). Also in this case, we observe that, for a fixed value of θ, C_v as a function of T reaches its maxima along the boundaries of the three lobes separating the LRO phase from the DO phase. The values attained at these maxima actually depend on θ: in Sec. 3.2.1 we address the origin of this dependence, focusing in particular on $\theta = 100°$ and $\theta = 280°$. The peculiar behavior for $\theta < 90°$, where the system does not magnetize for the same reasons presented in the supplemental material of Ref. [89], is further investigated for $\theta = 60°$ in App. A.5.

Taken together, these results show that the ARXY model supports a LRO phase over essentially the same three-lobed region observed in the NRXY model. This similarity is especially significant because the two models differ in their thermodynamic character: the NRXY model reaches a non-equilibrium steady state, whereas the ARXY model is an equilibrium system. The numerical comparison therefore demonstrates that the existence of LRO and the basic three-lobed structure of the phase diagram cannot be attributed exclusively to non-reciprocity. The geometrical origin of the numerical results shown here will be analyzed in Chapter 3.

2.3.3 The symmetric reciprocal XY model

In this Section, we finally consider the SRXY model introduced in Sec. 2.1.3. As in the ARXY model, its dynamics satisfies detailed balance and samples an equilibrium distribution. Its interaction rule, however, differs from that of the ARXY model because a bond is activated only once whenever at

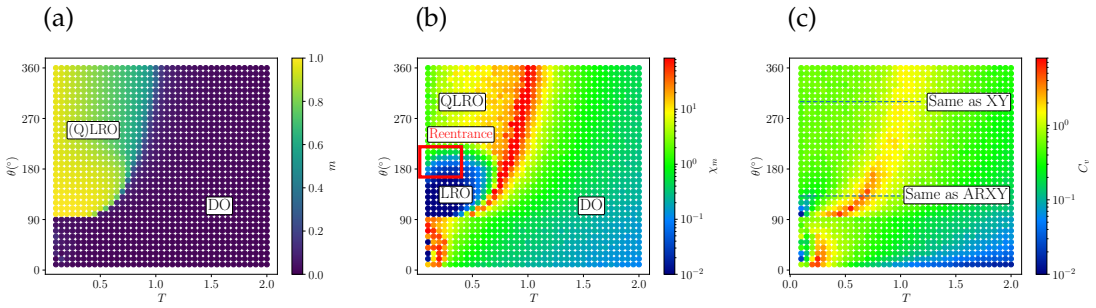


Figure 2.6: Results of the MC simulations for the equilibrium state of the SRXY model, obtained for a system of size $L = 100$. (a) The magnetization m as a function of the temperature T and the VC angle θ does not exhibit the three lobes seen in the previous two models, see Figs. 2.4(a) and 2.5(a). The behavior of m, however, is not sufficient to distinguish whether the low-temperature region with non-zero magnetization corresponds to a LRO or a QLRO phase. (b) The magnetic susceptibility χ_m, below the transition to the DO phase, is characterized by a region within which it is relatively small, indicating the presence of a LRO, and another one within which it takes larger values, corresponding to the QLRO phase. Moreover, for $\theta \gtrsim 180°$ we note the existence of a reentrance region (highlighted by a red box), where the system transitions from QLRO to LRO as the temperature T increases. (c) The specific heat C_v shows peaks corresponding to the transitions from the low-temperature region to DO, with higher peaks where the transition is from LRO to DO (as in the ARXY model), and lower peaks where the transition is from QLRO to DO (as in the XY model).

least one of the two neighboring spins sees the other. The purpose of this section is to determine how this alternative reciprocal implementation modifies the numerical phase diagram and, in particular, whether it also permits true LRO.

Figure 2.6 shows the behavior of m, χ_m and C_v as a function of the temperature T and the VC angle θ, obtained for a square system with $L = 100$. In particular, panel (a) presents the magnetization m, which, contrary to the previous two models, does *not* feature the three lobes which were visible in Figs. 2.4(a) and 2.5(a).

To determine whether the low-temperature region with finite magnetization corresponds to a LRO or QLRO phase, we analyze the magnetic susceptibility χ_m shown in Fig. 2.6(b). The plot reveals the presence of a lobe characterized by low values of the susceptibility, which suggests the possible presence of a LRO phase, and a region with higher susceptibility, associated with the QLRO phase. Rigorous proof of the presence of the LRO and QLRO phases is given in Sec. 3.3.1. Moreover, the plot in Fig. 2.6(c) shows that the specific heat C_v takes larger values along the peaks for $90° < \theta \lesssim 180°$ (corresponding to the transition from LRO to DO), and smaller values for $\theta \gtrsim 180°$ (corresponding to the transition from QLRO to DO). As we shall discuss in Sec. 3.3.1, these different behaviors observed for smaller or larger VC angles can be attributed to the presence of the redundant bonds introduced in Sec. 2.1.3 for $\theta > 180°$.

Interestingly enough, the lobe characterized by LRO extends beyond the range $\theta \in [90°, 180°]$ and, in particular, it exceeds $180°$, before showing a reentrant behavior at lower temperatures. The reentrance is highlighted by the red box in Fig. 2.6(b). This suggests a rather unusual phenomenon: for values of θ slightly above $180°$, increasing the temperature T drives the system from a QLRO phase into a LRO phase, resembling a type of order-by-disorder transition, which shall be discussed in Sec. 3.3.2.

The phase diagram of the SRXY model therefore differs qualitatively from those of both the NRXY and ARXY models. Nevertheless, it provides a second and independent reciprocal realization of a LRO phase. The presence of LRO in this model confirms that neither directed couplings

nor a non-vanishing entropy production rate are necessary for finite-temperature ordering in the class of systems considered here.

At the same time, the coexistence of LRO and QLRO regions, together with the reentrant transition from QLRO to LRO observed slightly above $\theta = 180^\circ$, shows that the precise implementation of the vision-dependent interaction remains essential. The numerical results thus separate two statements that might otherwise be conflated: non-reciprocity is not required for LRO, but changing how the vision-cone interaction is made reciprocal can substantially alter the topology of the phase diagram. A symmetry-based interpretation of these differences is developed in the following chapter.

2.4 Concluding remarks

In this chapter, we have introduced and numerically compared three two-dimensional XY models in which the interactions depend on a vision cone tied to the instantaneous spin orientation. The three models share the same basic perception-like geometry but differ in the way the corresponding interactions enter the dynamics. The NRXY model is intrinsically non-reciprocal and reaches a nonequilibrium stationary state, whereas the ARXY and SRXY models derive from global energy functions and define equilibrium statistical-mechanical systems.

The controlled comparison between these models leads to the conclusion that non-reciprocity is not required for the emergence of true LRO. Both reciprocal models support a finite-temperature LRO phase, and the NRXY and ARXY models display remarkably similar three-lobed phase diagrams despite their different thermodynamic character. Moreover, LRO persists in the NRXY and ARXY models at the special angles $\theta = 180^\circ$ and 270°, where the EUR vanishes. The presence of a finite EUR is therefore not a necessary condition for orientational ordering.

The second reciprocal variant, the SRXY, displays additional interesting phenomena: although it also supports LRO, its phase diagram differs qualitatively from those of the NRXY and ARXY models. It contains an extended QLRO region and, close to $\theta = 180^\circ$, exhibits a reentrant regime in which increasing temperature drives a transition from QLRO to LRO. The ordering behavior therefore depends highly on how the vision-dependent bonds are activated and counted.

Taken together, these results identify the orientation-dependent coupling between the internal spin degree of freedom and the spatial structure of the lattice as the central ingredient underlying the observed LRO. Non-reciprocity remains physically important because it determines the nonequilibrium character of the stationary state, including the violation of detailed balance and the presence of entropy production, but it is not the mechanism that uniquely accounts for orientational ordering.

The numerical evidence established here provides the basis for the analysis developed in the following chapter. There, we show how the vision-cone interaction generates an effective coupling between spin orientation and lattice directions, and use this mechanism to explain the different phase diagrams, the persistence of LRO at special vision angles, and the reentrant ordering behavior of the SRXY model.

Chapter 3

Spin–lattice coupling and emergent angular anisotropy

The numerical results presented in the previous chapter established that true LRO occurs not only in the non-reciprocal vision-cone model, but also in both of its reciprocal counterparts. The aim of the present chapter is to identify the microscopic mechanisms responsible for this ordering and to explain the qualitative differences between the phase diagrams of the asymmetric and symmetric reciprocal models.

In the conventional XY model, a global rotation of all spins leaves the interaction energy unchanged, independently of the geometry of the underlying lattice. Vision-cone interactions alter this separation between internal and spatial degrees of freedom: rotating the spins also changes which lattice bonds are active. As a result, the global orientation can acquire a non-trivial angular dependence tied to the directions of the lattice.

We probe this spin–lattice coupling by applying uniform rotations to selected configurations and measuring the resulting energy variation. For perfectly aligned states, this procedure characterizes the zero-temperature angular energy landscape. Applied to thermal configurations, it instead measures the effective anisotropy generated by fluctuations.

For the ARXY model, we show that the EUR provides an explicit source of fourfold anisotropy, but with different features for two different values of VC extension. We then examine the special case $\theta = 180^\circ$, where the EUR vanishes but fluctuations nevertheless select the lattice directions. Finally, for the SRXY model, we show that redundant bonds generate an accidental continuous degeneracy of nearly aligned configurations. Close to $\theta = 180^\circ$, thermal fluctuations lift this degeneracy and produce the observed reentrant transition from QLRO to LRO.

The rest of the chapter is organized as follows. Section 3.1 introduces global spin rotations as a diagnostic of the effective internal symmetry of the system. Section 3.2.1 compares the asymmetric reciprocal model at $\theta = 100^\circ$ and $\theta = 280^\circ$, emphasizing how two EUR zone of equal range can give rise to different thermodynamic behavior. Section 3.2.2 examines the special case $\theta = 180^\circ$ and develops a fluctuation-based explanation for the emergence of fourfold order in the absence of such a range. Section 3.3.1 discusses how redundant bonds produce an accidental continuous degeneracy in the symmetric reciprocal model, while Sec. 3.3.2 shows how thermal fluctuations lift this degeneracy and generate the observed order-by-disorder transition. Finally, Sec. 3.4 summarizes the resulting picture.

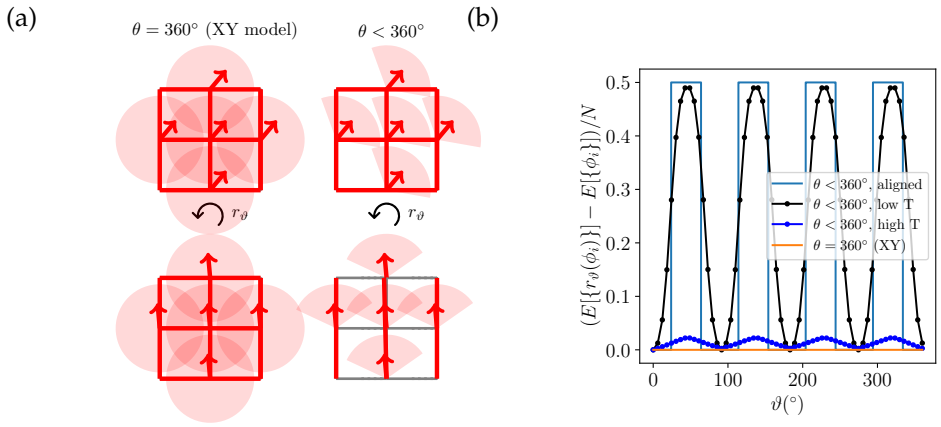


Figure 3.1: Illustration of the effects of a global rotation r_ϑ on the energy of a configuration, used as a tool to probe the symmetry of a system. (a) In the case of the XY model (i.e., of the ARXY with $\theta = 360°$), represented on the left, r_ϑ does not change the energy of the system, independently of the initial configuration. In fact, all spins continue to interact with their nearest neighbors, and the differences between angles (i.e., $\phi_i - \phi_j$) remain constant. For the ARXY with $\theta < 360°$, represented on the right, the angle differences between neighboring spins are again not altered by r_ϑ, but the structure of the activated bonds (highlighted in red) changes, resulting in a change of the energy and in the absence of an $O(2)$ symmetry in the system. (b) Dependence on the rotation angle ϑ of the energy difference per spin $(E[\{r_\vartheta(\phi_i)\}] - E[\{\phi_i\}])/N$ (displayed data correspond to $\theta = 315°$). For the XY model, the energy remains constant upon rotation, as expected. Conversely, for ARXY with $\theta < 360°$, the change in energy depends on the starting configuration. For a perfectly aligned configuration, the rotation results into a square wave with a periodicity of 90°. For a low-temperature thermal configuration (not perfectly aligned), the curve is smoothed but retains the 90° periodicity. At higher temperatures, corresponding to the DO phase, the curve is again almost independent of ϑ, similar to the XY model.

3.1 Spin–lattice and internal symmetry

In this section, we provide an interpretation of the results of the MC simulations presented in Chaper 2, which requires adapting the notion of internal symmetry to the types of models we are studying.

In the models typically examined in the context of statistical physics, the notion of internal symmetry plays a crucial role in characterizing phase transitions. In particular, the XY model has an $O(2)$ symmetry, which means that by applying a global rotation $r_\vartheta(\phi_i) = \phi_i + \vartheta \pmod{2\pi}$ of an angle ϑ to all spins (i.e., the internal degrees of freedom) of the system, the total energy does not change:

$$E_{XY}[\{\phi_i\}] = E_{XY}[\{r_\vartheta(\phi_i)\}]. \tag{3.1}$$

Crucially, in many models usually studied in statistical physics (including the XY model that is relevant here), the internal degrees of freedom are not coupled to the spatial structure of the system, characterized by the dimensionality and the lattice type. By contrast, in the models we are currently interested in, the vision cone acts as a bridge between the internal degrees of freedom and the lattice structure, rendering the actual couplings J_{ij} configuration dependent. In particular, when the vision cone is $\theta < 360°$, applying a rotation r_ϑ has also the effect of changing the structure of the activated bonds, as shown in Fig. 3.1(a), where the active bonds are highlighted in red.

Plotting the energy difference per spin, $(E[\{r_\vartheta(\phi_i)\}] - E[\{\phi_i\}])/N$, obtained upon rotating the system, provides information about the symmetry emerging from this interplay between the vision cone and the lattice structure. In fact, for the XY model (with vision cone $\theta = 360°$), the energy difference due to such a rotation is zero, as shown in Fig. 3.1(b), as expected due to the $O(2)$

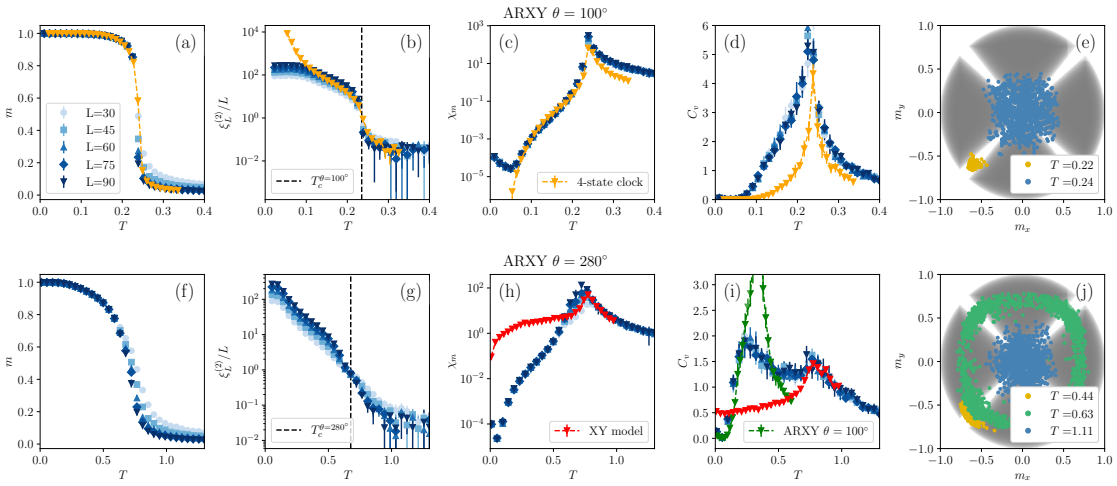


Figure 3.2: Behavior of the quantities introduced in Sec. 2.2.1 for the ARXY model with VC angles (a)–(e) $\theta = 100°$, and (f)–(j) $\theta = 280°$. Despite both VC angles sharing the same EUR size, the plots show marked differences. Notably, for $\theta = 100°$ (upper row of plots), the behavior is similar to that of the 4-state clock model (gold diamond symbols), as evidenced in panels (a)–(c) by the similar profile of m, $\xi_L^{(2)}/L$ and χ_m (upon rescaling of temperature), and by the comparable peak height of the specific heat C_v in panel (d). For $\theta = 280°$ (lower row of plots), we observe a hybrid behavior, as seen particularly in panel (i), where the specific heat initially grows similarly to the case $\theta = 100°$ (green triangles), then decreases, and later exhibits a second peak with the same profile as the XY model (red triangles in panels (h) and (i)).

symmetry of the model. Conversely, in the presence of a vision cone $\theta < 360°$ (here for the ARXY model), the energy difference obtained by rotating a perfectly aligned initial spin configuration takes the form of a square wave with peaks of value $1/2$ (see Eq. (2.4)). This corresponds to the energy gained when each spin stops interacting with a neighbor, and is due to the presence of the EUR discussed in Sec. 2.1.4. In the absence of symmetry, the difference between $E[\{r_\vartheta(\phi_i)\}]$ and $E[\{\phi_i\}]$ actually depends on the configuration $\{\phi_i\}$. For this reason, in Fig. 3.1(b) we apply the rotation r_ϑ also to the case of a thermal configuration of the system at a certain temperature T. In practice, at each step of the MC evolution, we rotate the configuration $\{\phi_i\}$ of the system by an angle ϑ, calculate the energy difference per spin $(E[\{r_\vartheta(\phi_i)\}] - E[\{\phi_i\}])/N$, and then we average over many MC realizations. The resulting curve, obtained at low temperatures in the LRO phase, is smoother compared to the one corresponding to an initial configuration with perfectly aligned spins, but it still exhibits a periodicity of 90°. We also note that, as the temperature T increases, the system is progressively less influenced by the presence of the EUR: the energy-difference curve becomes increasingly flatter, corresponding to the DO phase. This procedure will be used to rationalize the order-by-disorder transition observed for the SRXY model in, c.f., Sec. 3.3.2.

In the following, we apply these notions for elucidating different features of the transition from LRO to DO for two selected values of the VC angle.

3.2 The asymmetric reciprocal model

3.2.1 Two different manifestations of $\mathbb{Z}_4$ symmetry: $\theta = 100°$ and $\theta = 280°$

Numerical results. — Figure 3.2 reports the results of the MC simulations of the ARXY for the two values $\theta = 100°$ and $\theta = 280°$ of the VC angle. In both cases, the size α_{EUR} of the EUR is the same, i.e., $\alpha_{\text{EUR}} = 320°$. However, when a spin enters the EUR, the number of activated bonds changes

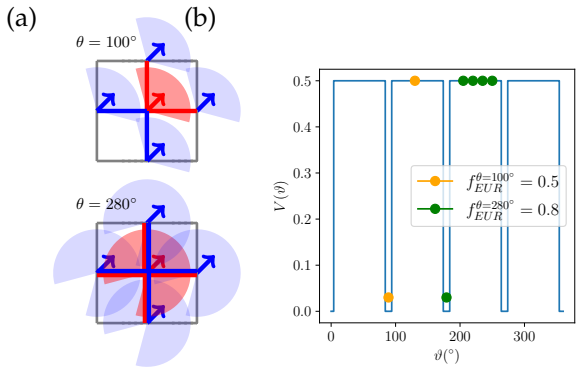


Figure 3.3: Illustration of the difference in the behavior of the ARXY model with VC angles $\theta = 100°$ and $\theta = 280°$ at the transition from the LRO to the DO phase. (a) Perfectly aligned configuration $\{\bar{\phi}_i = 45°\}$ — as an approximation for low-temperature configurations — for both VC angles. Even though the EUR has the same size, for $\theta = 280°$ the activated bonds (in red) are twice as many as for $\theta = 100°$, which doubles the energy scale of the model. (b) Profile of the energy variation per spin, obtained by applying a rotation r_ϑ to the aligned configurations shown in (a). The maxima correspond to the energy gain of $1/2$ when a spin enters the EUR and loses one nearest neighbor. Within our approximation, we consider this profile as a potential $V(\vartheta)$ (see Eq. (3.2)), to be used at finite temperatures to compute the fraction of spins f_{EUR} that are within the EUR (see Eq. (3.3)). At the transition temperature for $\theta = 100°$, only half of the spins are within the EUR, and the $\mathbb{Z}_4$ symmetry is more pronounced than for $\theta = 280°$, where 4 out of 5 spins are in the EUR, and the potential effectively seen by the system is flat, resembling that of the XY model.

from 2 to 1 for $\theta = 100°$, and from 4 to 3 for $\theta = 280°$. This entails a significantly different behavior of the observables for $\theta = 100°$ and $\theta = 280°$, shown in the upper and lower rows of Fig. 3.2, respectively.

For $\theta = 100°$, the behavior observed in Figs. 3.2(a)–(e) qualitatively resembles that of the 4-state clock model shown in Figs. 2.3(a)–(e), with an appropriate rescaling of the temperature. In particular, the magnetization m in panel (a), the curves of $\xi_L^{(2)}/L$ in panel (b) (which, for different system sizes L, do not collapse below the transition temperature, but instead intersect at the onset of the DO phase), and the magnetic susceptibility χ_m in panel (c) (that is size-independent below the transition temperature), indicate the presence of LRO at low temperature. Numerical data for the 4-state clock model with $L = 60$ are included for comparison, and they are represented by gold symbols and a dashed line in panels (a)–(d) of Fig. 3.2. By locating the crossing of the curves $\xi_L^{(2)}/L$ vs. T with different L, we estimate the transition temperature $T_c^{\theta=100°} \simeq 0.23$, which is indicated by a vertical dashed line in panel (b). This temperature is approximately four times smaller than that of the standard XY model. This fact can be rationalized by considering that, when the system is disordered, approximately one bond out of four is activated, and thus the coupling energy rescales accordingly.

The temperature dependence of the specific heat C_v in Fig. 3.2(d) is also qualitatively very similar to that of the 4-state clock model, and it features a peak of comparable magnitude. The vectorial magnetization (m_x, m_y) in Fig. 3.2(e) illustrates that, at low temperatures, the spins collectively align along one of the four preferred directions (complementary to the EUR), while at higher temperatures, in the DO phase, they form a rotationally symmetric cloud centered around $(0, 0)$.

For $\theta = 280°$ in Fig. 3.2, the size dependence of the magnetization m in panel (f), of $\xi_L^{(2)}/L$ in panel (g), and of χ_m in panel (h) are similar to those of the corresponding variables for the VC angle $\theta = 100°$, shown in the panels (a)–(c) above. This indicates the presence of a phase with LRO at low temperatures, and a DO phase at higher temperatures.

However, the specific heat C_v in panel (i) of Fig. 3.2 shows a rather peculiar dependence on temperature T: upon increasing T, there is a first rise in C_v that matches the increase of C_v for the case $\theta = 100^\circ$ of panel (d) (displayed again for convenience, with a temperature rescaling, with green symbols in panel (i)). Upon further increasing T, a second peak gradually appears, with the same profile as that of the XY model (red symbols) reported in Fig. 2.3(n). This suggests that, contrary to the case $\theta = 100^\circ$, for $\theta = 280^\circ$ the system displays also a behavior which is somehow reminiscent of the XY model. From the crossing of the $\xi_L^{(2)}/L$ curves, we estimate the transition temperature $T_c^{\theta=280^\circ} \simeq 0.68$, which is indicated by a vertical dashed line in panel (g). This temperature is approximately $3/4$ of that in the XY model. We rationalize this observation by noting that in an almost disordered phase, for $\theta = 280^\circ$, an average of three out of four bonds are activated, leading to a corresponding rescaling of the effective coupling. The scatter plot of the vectorial magnetization $\mathbf{m} = (m_x, m_y)$ reported in Fig. 3.2(j) also suggests that an intermediate case between disorder and the alignment in one of the four minima — such as those observed in panel (e) — appears. In this case, the magnetization m does not vanish, but $\mathbf{m}$ lacks a preferred direction and the system becomes isotropic, similarly to what happens for the XY model in the phase with QLRO (see Fig. 2.3(o)). The emergence of this QLRO-like behavior in a phase which actually admits a LRO can be explained as discussed below.

The role of internal symmetry. — To motivate the observations listed above, let us consider a zero-temperature configuration in which all spins are aligned, as illustrated in Fig. 3.3(a) for both values of the VC angle, e.g., forming an angle $\bar{\phi}_i = 45^\circ$ at site i. In the case of $\theta = 100^\circ$, each spin participates in the activation of two bonds (thus the energy per spin is $E^{\mathrm{AR}}/N = -1$, according to the normalization of Eq. (2.4)), while in the case of $\theta = 280^\circ$, each spin activates four bonds (and thus $E^{\mathrm{AR}}/N = -2$). This estimate of the energy magnitudes suggests that the transition temperature for the model (that is proportional to the effective coupling strength) with $\theta = 100^\circ$ is expected to be approximately half of the transition temperature of the model with $\theta = 280^\circ$. Indeed, as discussed above, we estimated them to be $T_c^{\theta=100^\circ} \simeq 0.23$ and $T_c^{\theta=280^\circ} \simeq 0.68$ from panels (b) and (g) of Fig. 3.2, respectively. By applying a rotation r_ϑ to the configuration $\{\bar{\phi}_i\}$, the energy variation per spin, i.e.,

$$V(\vartheta) = (E[\{r_\vartheta(\bar{\phi}_i)\}] - E[\{\bar{\phi}_i\}])/N, \tag{3.2}$$

shown in Fig. 3.3(b) as a function of ϑ, exhibits a periodicity of 90°. The upper plateaus of the profile correspond to the values of ϑ for which the rotated spins enter the EUR, losing one bond and thus increasing the energy per spin by $1/2$. This 90° periodicity implies that the interplay between the arrangement of the spins according to a square lattice and the vision cone actually reduces the symmetry of the model to $\mathbb{Z}_4$. This explains the features of the collective behavior of the model with $\theta = 100^\circ$, characterized by the fact that all observables reported in Fig. 3.2 are very similar to those of the 4-state clock model.

However, the same argument does not carry over to the case $\theta = 280^\circ$. In fact, while the EUR is not energetically favored, it is not strictly forbidden. Accordingly, at sufficiently high temperatures, a significant fraction of the spins actually occupies the EUR. In order to calculate the probability for this to occur at finite temperature, let us consider the profile shown in Fig. 3.3(b), that was obtained for an initial zero-temperature configuration. In a first approximation, we assume that most of the spins remain roughly aligned at finite temperature, with only a finite fraction f_{EUR} of them entering the EUR. This effectively reduces the problem to that of a single particle at inverse temperature

$\beta = 1/T$ in the one-dimensional potential provided by Fig. 3.3(b), with $V(\vartheta) \simeq 0$ or $\simeq 1/2$ (see Eq. (3.2)). In particular, $e^{-\beta V(\vartheta)}$ is approximately equal to 1 or $e^{-\beta/2}$, and is proportional to the probability for such particle to be, respectively, outside or inside the EUR. In terms of the total extension $\alpha_{\text{EUR}} \in [0^\circ, 360^\circ)$ of the latter, the fraction f_{EUR} of spins within the EUR is proportional to $\alpha_{\text{EUR}}\ e^{-\beta/2}$, which has to be normalized by the constant $Z = \alpha_{\text{EUR}} \cdot e^{-\beta/2} + (360^\circ - \alpha_{\text{EUR}}) \cdot 1$. Accordingly, one has

$$f_{\text{EUR}} = \frac{\alpha_{\text{EUR}}\ e^{-\beta/2}}{Z} = \frac{\alpha_{\text{EUR}}\ e^{-\beta/2}}{360^\circ - \alpha_{\text{EUR}}(1 - e^{-\beta/2})}. \tag{3.3}$$

This relation allows us to determine the value of f_{EUR} at the transition temperature for $\theta = 100^\circ$ and $\theta = 280^\circ$ on the basis of the knowledge of the transition temperatures, previously estimated from the crossing points in Figs. 3.2(b) and (g). Correspondingly, we obtain $f_{\text{EUR}}^{\theta=100^\circ} \simeq 0.5$ and $f_{\text{EUR}}^{\theta=280^\circ} \simeq 0.8$. This implies that the number of spins within the EUR for $\theta = 100^\circ$ equals that outside of it, meaning that the statistical weight of the four minima is actually relevant, and the behavior of the spins is influenced by their presence. On the contrary, for $\theta = 280^\circ$, the fraction of spins within the EUR is significantly larger, indicating that the statistical relevance of the four minima is reduced. As a result, the effective potential experienced by a significant fraction of the spins of the system, and resulting from that in Fig. 3.3(b), becomes essentially flat. Consequently, the system behaves like an XY model, for which a rotation r_ϑ leaves the energy unchanged. Note, however, that different quantities are influenced to a different extent by this approximate recovery of the $O(2)$ symmetry: the specific heat in Fig. 3.2(i), in fact, displays the same behavior observed for the XY model at the transition, while the ratio $\xi_L^{(2)}/L$ reported in Fig. 3.2(g) continues to feature the intersection typical of the transition from LRO to DO, indicating the presence of an underlying $\mathbb{Z}_4$ symmetry.

The above argument explains why the ARXY can behave very differently even for two values of θ for which the extension α_{EUR} of the EUR is the same: the change in the number of activated bonds per spin — going from 2 to 1 for $\theta = 100^\circ$ and from 4 to 3 for $\theta = 280^\circ$ — produces a significant difference. Specifically, it alters the energy scale of the systems relative to the energy lost upon deactivating a bond, which is always $1/2$.

Finally, for $\alpha_{\text{EUR}} = 0^\circ$ (which is equivalent to $\alpha_{\text{EUR}} = 360^\circ$, see Fig. 2.2(d)), the argument presented above would suggest that the model should behave like a standard XY model with $O(2)$ symmetry. In the next section, we will explain why this is not really the case for $\theta = 180^\circ$.

3.2.2 ARXY for $\theta = 180^\circ$: LRO without EUR

As discussed at the end of the previous section, the argument presented there indicates that, in the absence of an EUR, the model should display an $O(2)$ symmetry. This was also suggested in Ref. [89] for the NRXY model. Accordingly, one would expect that, in this case, the transition occurs between a QLRO phase and a DO phase. However, this is not what we find, as observed in Fig. 2.4(e) for the NRXY model, and as will be shown for the ARXY model in the present section.

Numerical results. — The numerical results for the ARXY model at $\theta = 180^\circ$ are shown in Fig. 3.4(a)–(e). By looking at the temperature dependence of $\xi_L^{(2)}/L$ in panel (b), one notices that the curves corresponding to different sizes L do not show the collapse expected for an extended critical phase (such as the one with QLRO) at low temperatures. On the contrary, the behavior is similar to that seen for the NRXY model in Fig. 2.4(e), indicating LRO. In addition, from the scatter

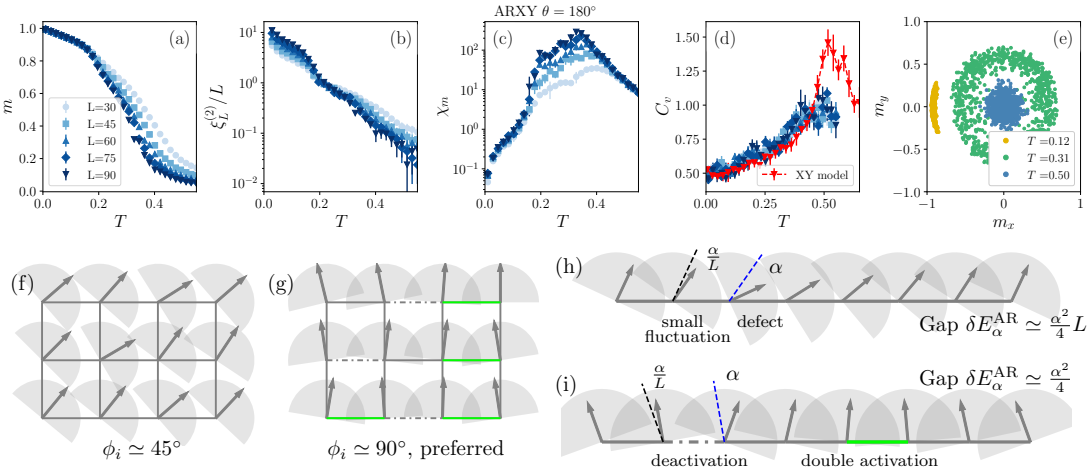


Figure 3.4: ARXY model for $\theta = 180°$. (a)–(e) Result of the MC simulations for the quantities introduced in Sec. 2.2.1. In panel (b), the behavior of the curves of $\xi_L^{(2)}/L$ upon varying L indicates the presence of LRO at low temperatures. The vectorial magnetization in panel (e) demonstrates that, at low temperatures, the spins preferentially align along one of the four (equivalent) lattice directions. (f)–(i) Illustration of the mechanism favoring configurations aligned along the lattice directions at low temperatures. (f) For a configuration aligned around 45°, all bonds remain activated (solid gray) even in the presence of small fluctuations. (g) For a configuration aligned around 90°, small fluctuations can deactivate one bond (dash-dotted gray) and doubly activate another one (in green). (h) For a gapless spin wave, the angular change between neighboring sites i and j scales as $\phi_i - \phi_j \sim 1/L$. A finite fluctuation α, indicated as a defect in the figure, increases the energy by $\alpha^2/2$, and the corresponding configuration has an energy that scales subextensively. (i) If the bond at the defect is deactivated, the energy increase is finite (1/2), but is compensated by the double activation of another bond further along the chain. The corresponding configuration (identical to that in panel (h), but rotated by 45°) is gapless in the thermodynamic limit.

plot of the vectorial magnetization **m** in panel (e) it seems that at low temperatures the system preferentially aligns along one of the four (actually equivalent) lattice directions. Upon increasing the temperature, a behavior similar to that observed for the ARXY model at $\theta = 280^\circ$ emerges, where a flavor of QLRO appears, and no spatial orientation seems to be preferred. This is evident from the magnetic susceptibility in panel (c) as well, which scales with the system size within a finite range of temperatures before the transition.

Heuristic explanation. — We aim to explain the fact that, at least at low temperatures, the various spatial orientations of the spins are not equivalent. This leads to a mechanism that reduces the symmetry from $O(2)$ to $\mathbb{Z}_4$, even in the absence of the EUR. To do so, we analyze the configurations with small fluctuations around a perfectly aligned state. In such a state, for $\theta = 180^\circ$, the number of activated bonds per spin is at most 2 (except for a perfect alignment with $\phi \in \{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$, which is physically unrealistic). As shown in Fig. 3.4(f), if the system predominantly aligns around 45° (or its odd multiples), such small oscillations do not alter the configuration of the bonds activated once by the interaction (these bonds are depicted in dark gray in the figure). In contrast, if the spins of the system align along one of the four lattice directions, as depicted in Fig. 3.4(g), then even small fluctuations may lead to the double activation of one bond (green line) at the expense of the complete deactivation of another bond (dash-dotted gray). Note that the PBCs imposed on the lattice ensure that the number of bonds that are activated twice is equal to that of the bonds that are deactivated, such that the energy of the configuration, up to fluctuations, is almost unaltered.

As we are about to show, the two reference configurations discussed above are actually characterized by different spectra of gapless excitations, and thus they are not equivalent in determining the thermodynamics of the system, therefore breaking the rotational symmetry.

Let us first consider the standard XY model, and its gapless spin-wave excitations: these are fluctuations where the angular change between neighboring sites i and j scales as $\phi_i - \phi_j \sim 1/L$, which lead to a local increase of energy $\sim 1/L^2$. Summing this increase over the entire system amounts to multiplying it by L^2, which renders an intensive energy increase $\propto L^0$, corresponding to an excitation that is gapless in the thermodynamic limit. These kind of excitations are also present for the ARXY model, irrespective of the direction of global alignment of the spins (i.e., for both 45° and 90°, and their multiples) — provided that we consider, in these cases, fluctuations that do not alter the global alignment of the configurations, such that it still possible to identify configurations for which $\phi_i \simeq 45^\circ$ or $\phi_i \simeq 90^\circ$, as shown in panels (f) and (g) of Fig. 3.4, respectively.

Still considering the standard XY model, we now focus on the energy increase δE_α that results from the formation of a line of defects, where the angle of a single spin along a row changes by a finite value α (potentially small, but not scaling as $1/L$) compared to the direction of collective alignment of the remaining spins. Referring to Fig. 3.4(h), assuming that one defect per lattice row is formed, the excitation energy is

$$\delta E_\alpha = L \left[\frac{\alpha^2}{2L^2}(L-1) + \frac{\alpha^2}{2} \right] \simeq \frac{\alpha^2}{2} L, \tag{3.4}$$

for large L. This results in a subextensive excitation.

Turning to the ARXY model, if the spins in the original configuration are aligned around 45°, and the fluctuations are not sufficiently large to alter the spatial arrangement of the activated lattice bonds, the excitation energy of the configuration in panel (h) is the same (up to a factor $1/2$ from Eq. (2.4)), given by $\delta E_\alpha^{\mathrm{AR}} \simeq (\alpha^2/4)L$.

On the other hand, if the configuration aligns around 90°, the same configuration actually becomes gapless. In panel (i), we show the same configuration as in panel (h) but rotated by 45°. At the defect point (along a row), one bond deactivates, gaining an energy of $1/2$ (independently of α, thus the functional dependence on the defect angle disappears). This gain is compensated by the double activation of another bond further along the row (because every deactivation must correspond to an activation of some other bond, due to PBCs). Consequently, the energy of this activation is, in this case,

$$\delta E_{\alpha}^{\mathrm{AR}} = L\left[\frac{\alpha^2}{4L^2}(L-2) + \frac{1}{2} - \frac{1}{2}\left(1 - \frac{\alpha^2}{2L^2}\right)\right] \simeq \frac{\alpha^2}{4}. \tag{3.5}$$

This results in an intensive excitation $\propto L^0$.

In other words, in the second case, the spectrum of strongly gapless excitations (where the global energy required for excitation does not scale with system size) is broader. By contrast, in the first case, the same configurations have an energy that scales subextensively. Our heuristic argument is then based on entropic reasoning: the system favors alignments around the spatial direction from which more microstates can be explored via fluctuations at no energy cost. We demonstrated that configurations near the lattice directions exhibit a higher number of strongly gapless states: consequently, the four lattice directions are preferred, breaking the $O(2)$ symmetry of the system and leaving a residual $\mathbb{Z}_4$ symmetry. Notably, this symmetry breaking occurs without the involvement of an EUR.

At present, we are unable to make a more direct comparison with the behavior of the 4-state clock model or the XY model, as was done in Sec. 3.2.1.

3.3 The symmetric reciprocal model

3.3.1 The role of redundant bonds for $\theta > 180°$

The quantities plotted in Fig. 2.6 reveal that, unlike the NRXY and ARXY models, the SRXY model features a single lobe of LRO for $90° < \theta \lesssim 180°$ and an extended QLRO phase for $\theta \gtrsim 180°$. In order to highlight the occurrence of these different behaviors within the ordered phase, we focus on the behavior of the SRXY model for $\theta = 300°$, which is reported in Fig. 3.5(a)–(e). The close match of the observables in panels (a)–(d) with the corresponding profiles for the XY model, that we report in red symbols for $L = 60$, allows one to conclude that the model exhibits only a QLRO and a DO phase. The behavior of the vectorial magnetization in panel (e) suggests a similar trend, forming a ring in the QLRO phase and an isotropic cloud centered around the origin in the DO phase. We refer to Figs. 2.3(k)–(o) for a comparison.

This restoration of the $O(2)$ symmetry in the SRXY model for $\theta > 180°$ can be understood through the redundant bond mechanism introduced in Sec. 2.1.3, as explained below.

We start by analyzing a typical low-temperature configuration of the model for the VC angle θ either with $90° < \theta < 180°$ or $\theta > 180°$, while applying the symmetry arguments already developed in the previous sections. For $90° < \theta < 180°$, Fig. 3.6(a) illustrates that, assuming a configuration with aligned spins (up to small fluctuations) at low temperatures, there is actually no difference between the asymmetric and the symmetric reciprocal models in terms of the number of activated bonds, and thus the corresponding energy is the same, up to a factor of 2 arising from the choice of normalization in Eq. (2.4) for the ARXY model. When a rotation r_{ϑ} is applied to

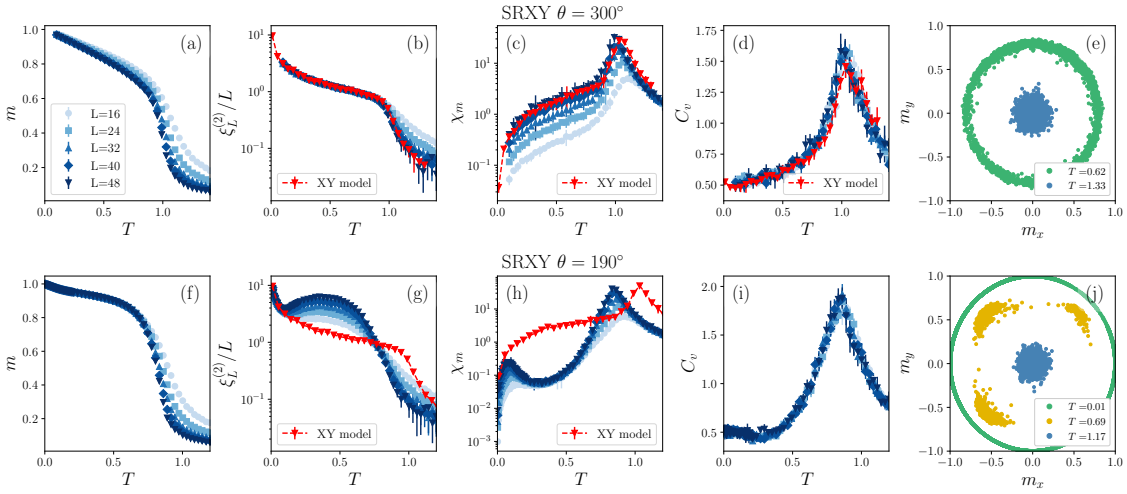


Figure 3.5: Behavior of the quantities introduced in Sec. 2.2.1 for the SRXY model with VC angles (a)–(e) $\theta = 300°$, and (f)–(j) $\theta = 190°$. For $\theta = 300°$, all the observables in panels (a)–(e) display behaviors that are similar to those observed in the XY model presented in Fig. 2.3(k)–(o) (and here indicated by red triangles, for a comparison). For $\theta = 190°$, evidence of an order-by-disorder transition is observed at low temperatures, specifically a transition from QLRO to LRO upon increasing the temperature. In panels (g) and (h), the behavior of $\xi_L^{(2)}/L$ and χ_m, respectively, are consistent at low temperatures with those of the XY model (i.e., with the presence of QLRO), but they deviate towards that expected in the presence of LRO as the temperature T increases. The vectorial magnetization $\mathbf{m}$ in panel (e) is also consistent with the fact that at low temperatures the system is isotropic, while it recognizes the presence of four preferred directions only at higher temperatures.

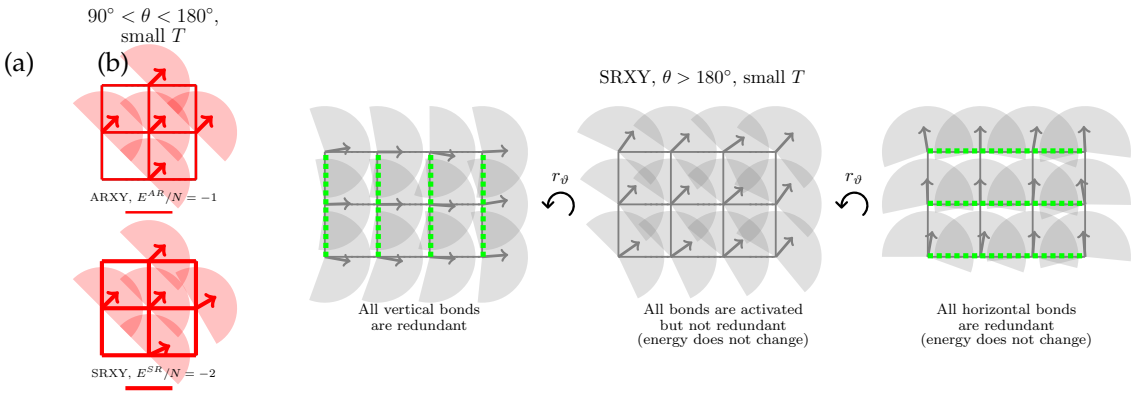


Figure 3.6: Illustration of the behavior of the SRXY model for $\theta \in [90°, 180°]$ and $\theta > 180°$. (a) For $\theta \in [90°, 180°]$, in a low-temperature configuration with all spins aligned up to small fluctuations, there is no difference between the ARXY and SRXY models (apart from a factor 2 in the corresponding values of the energy, due to the normalization chosen in Eq. (2.4)). In fact, all bonds are activated, an EUR exists, and the four directions that are preferred by the magnetization are the odd integer multiples of 45°. Consequently, as in the case of the ARXY model (see Fig. 2.5), there is a lobe with LRO at low temperatures in the phase diagrams in Fig. 2.6(a)–(b). (b) For $\theta > 180°$, the redundancy of bonds in the SRXY model makes it drastically different from the ARXY model. Notably, a nearly aligned low-temperature configuration does not have a privileged alignment direction in the SRXY model. In fact, starting for example from the configuration $\{\bar{\phi}_i \approx 0°\}_i$, all the vertical bonds are redundant; upon applying a global rotation r_ϑ of an angle $\vartheta \simeq 45°$, all the bonds remain activated, but leaving zero redundant bonds, so that the energy of the system does not change (see Eq. (2.5)). With a further rotation, the horizontal bonds become redundant, but the energy is still unchanged. This mechanism effectively restores the $O(2)$ symmetry.

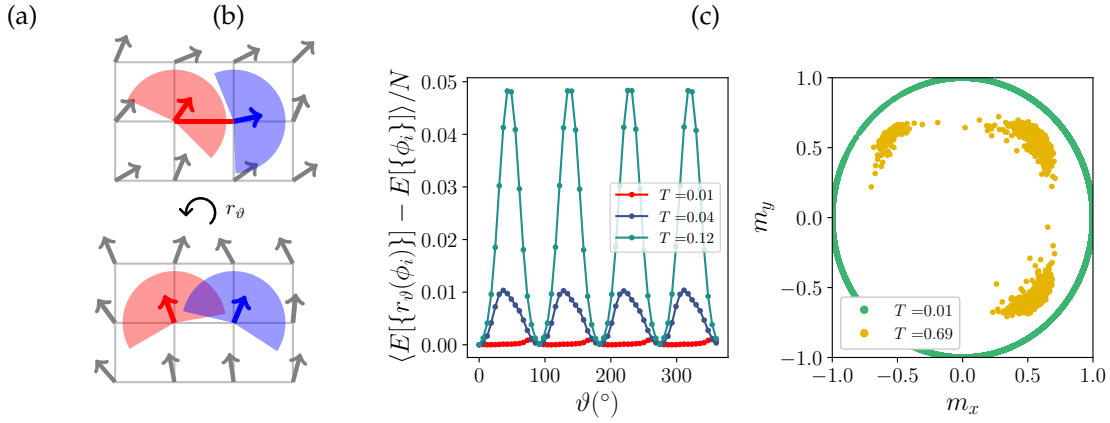


Figure 3.7: Illustration of the mechanism underlying the order-by-disorder transition observed in the SRXY model for $\theta = 190^\circ$. (a) As the temperature T increases, the system exhibits more pronounced fluctuations even at the local level, making the presence of the redundant bonds less likely. In the upper panel, the bond between the red spin and the blue spin is activated only through the interaction of the red spin with the blue one. Thus, unlike what shown in Fig. 3.6(b), the rotation r_ϑ of the spin configuration shown in the bottom panel can have the effect of deactivating a bond, thus increasing the energy. (b) Energy difference due to a rotation r_ϑ of the spins as a function of ϑ, for various temperatures T close to the order-by-disorder transition. While at very low temperatures the energy difference vanishes, as expected for a system with $O(2)$ symmetry, finite values with 90° periodicity emerge as T increases. (c) Scatter plot of the vectorial magnetization $\mathbf{m} = (m_x, m_y)$. While at low temperatures the angular distribution of $\mathbf{m}$ is isotropic, upon increasing the temperature T, $\mathbf{m}$ preferentially localizes along four directions: accordingly, disorder (in the form of fluctuations) reduces the symmetry from $O(2)$ to $\mathbb{Z}_4$.

this spin configuration, the resulting change of the energy as a function of the angle ϑ is like the one shown in Fig. 3.3(b), exhibiting the $\mathbb{Z}_4$ symmetry and therefore indicating a transition from LRO to DO. Indeed, within this range of VC angles, the transition from LRO to DO is observed also in the behavior of C_v reported in Fig. 2.6(c), in terms of the occurrence of higher values of C_v near the phase transition compared to the case with $\theta > 180^\circ$. In this respect, we remind (see also Figs. 2.3(d, n)) that the transition from LRO to DO is accompanied by a significantly more pronounced peak of C_v (that scales logarithmically with system size) compared to the case of the transition between QLRO and DO (that is size-independent).

The situation changes drastically for $\theta > 180^\circ$. In this case, applying the rotation r_ϑ to an aligned configuration does not change the energy of the SRXY model. The reason is the presence of the redundant bonds, as shown in Fig. 3.6(b): for a hypothetical almost aligned configuration with $\{\bar{\phi}_i \approx 0^\circ\}_i$ at low temperatures, all vertical bonds are redundant. However, the second "virtual" activation does not change the energy of the model (see Eq. (2.5)), contrary to what happens for the ARXY model. Rotating the system by an angle $\vartheta \simeq 45^\circ$ deactivates the redundancy of the vertical bonds, while the energy remains unchanged. This mechanism leads to a restoration of the $O(2)$ symmetry, which is reflected in the behavior of the susceptibility χ_m in Fig. 2.6(b): the two lobes with LRO that were present in the ARXY model for $\theta > 180^\circ$ are absent in the SRXY model, and they are replaced by a phase with QLRO.

3.3.2 Order-by-disorder transition

As we have already noted in Sec. 2.3.3, the behavior of the susceptibility χ_m in Fig. 2.6(b) exhibits a reentrance at low temperatures for VC angles $\theta \gtrsim 180^\circ$. Correspondingly, upon increasing the temperature T, the system undergoes a transition from a phase with QLRO into one with LRO.

This behavior is further confirmed by the dependence of the observables m, $\xi_L^{(2)}/L$, χ_m, C_v, and $\mathbf{m} = (m_x, m_y)$ shown in Figs. 3.5(f)–(j) for this model with $\theta = 190^\circ$. In fact, the susceptibility χ_m in Fig. 3.5(h) initially increases, upon increasing the temperature T, in a size-dependent manner, matching the values observed for the XY model, reported with red symbols for $L = 60$. At higher values of T, χ_m then decreases and becomes size independent, as expected in a phase with LRO, before forming a peak at the transition to the DO phase. Similarly, in Fig. 3.5(g) we observe that the curves for $\xi_L^{(2)}/L$ at various system sizes collapse at low temperatures, onto the same curve as for the XY model, which indicates the presence of a critical phase with QLRO. As the temperature increases, the system enters the phase with LRO, the curves corresponding to different values of L begin to separate (as expected within a phase with LRO), but they intersect again at a single point corresponding to the temperature of transition to the DO phase. Another manifestation of the occurrence of the transition from QLRO to LRO is shown in Fig. 3.5(j), where the vectorial magnetization $\mathbf{m} = (m_x, m_y)$ is finite and isotropic at very low temperatures. However, as the temperature increases, it becomes localized in four minima, before becoming isotropic again in the DO phase.

This transition from QLRO to LRO resembles an order-by-disorder transition, where, counter-intuitively, order increases upon increasing the temperature [135–137]. The underlying mechanism driving this phenomenon can be rationalized in terms of the existence of redundant bonds, and considering the rotation r_ϑ of the spin configuration as a tool to probe the internal symmetry of the system. As it has been already discussed in the previous section, an almost perfectly aligned configuration, such as the one in Fig. 3.6(b), can present redundant bonds. However, this conclusion hinges on the assumption that the spins are aligned at low temperatures, which is actually not the case for a system with $O(2)$ symmetry, because it cannot sustain an ordered state, which is destroyed by the spin-wave excitations at arbitrarily low temperatures. Nevertheless, the slow power-law decay of spatial correlations (see Eq. (2.12)) implies that, at least locally, the fluctuations of the spins are small, and thus they can be considered almost aligned across large regions of space. By applying a rigid rotation r_ϑ to the spin configuration, we expect to observe no change in energy (indeed, this is what happens when θ is *sufficiently* larger than 180°). However, as the temperature increases, local fluctuations become more pronounced, as illustrated in Fig. 3.7(a). As a consequence, as soon as θ is larger than 180°, some of the bonds that are not redundant might deactivate after a rotation r_ϑ. This is illustrated in Fig. 3.7 for the VC angle $\theta = 190^\circ$. In particular, panel (a) refers to the case of two neighboring spins that are connected by a non-redundant bond (activated only through the interaction of the red spin with the blue one): the rotation r_ϑ can lead to the deactivation of the bond, before it is activated again — this time by the interaction of the blue spin with the red one. Panel (b) of Fig. 3.7 shows the energy difference per spin after performing the rotation r_ϑ of the spin configuration with $\theta = 190^\circ$ at various temperatures close to the order-by-disorder transition. It can be seen that, as the temperature increases, regions emerge where the energy difference becomes increasingly more pronounced, with a dependence on ϑ showing a periodicity of 90°, thus again highlighting the $\mathbb{Z}_4$ symmetry. This is markedly different from the case shown in Fig. 3.1 for the ARXY model, where an increase in temperature led to the smoothing of the square wave profile. Here, instead, at low temperatures no direction is favored, while as the temperature increases, a sort of EUR forms dynamically: fluctuations make the various directions not equivalent, and therefore favor the selection of one of the preferred directions. Specifically, the four lattice directions are disfavored in this case. Referring to the lower panel of Fig. 3.7(a), we observe that alignment around one of the lattice directions makes the system more

susceptible to losing a bond following a small fluctuation. Consequently, the favored directions are the four odd integer multiples of 45°. This fact is also reflected in the vectorial magnetization shown in Fig. 3.7(c), which illustrates that at low temperatures there is no preferred direction, but one emerges as the temperature increases, demonstrating the transition from $O(2)$ to an effective $\mathbb{Z}_4$ symmetry.

3.4 Concluding remarks

In this chapter, we have investigated the microscopic mechanisms underlying the ordered phases identified numerically in the reciprocal vision-cone XY models (ARXY and SRXY). The central ingredient is the coupling between the internal spin orientation and the fixed spatial directions of the lattice. Because the spin orientations determine which bonds are active, a global rotation can modify the interaction energy even when all relative spin angles remain unchanged. Vision-cone interactions therefore generate a fourfold angular anisotropy on the square lattice that is absent from the conventional XY model.

For the ARXY model, the EUR provides the most direct manifestation of this anisotropy. The comparison between $\theta = 100^\circ$ and $\theta = 280^\circ$ shows, however, that its geometrical extent alone does not determine the observable behavior. The relevance of the fourfold modulation also depends on the overall number of active bonds and hence on the ratio between the angular barrier associated with losing a bond and the total interaction-energy scale. As a result, equal EUR extensions can produce very different behaviors: the $\mathbb{Z}_4$ anisotropy is already prominent for $\theta = 100^\circ$, whereas for $\theta = 280^\circ$ several observables display an extended XY-like regime.

The case $\theta = 180^\circ$ demonstrates that an explicit EUR is not required for orientational selection. Although perfectly aligned states possess an accidental continuous degeneracy, fluctuations around different global orientations are not equivalent. We identified a class of deformations whose total energy remains of order one for alignment along the lattice axes but grows as L for diagonal alignment. This difference suggests that axial configurations have access to a broader low-cost region of configuration space and therefore provides a plausible entropic mechanism for the observed pinning along lattice directions and LRO.

A different mechanism operates in the SRXY model. For $\theta > 180^\circ$, redundant bonds make the energy of nearly aligned configurations approximately independent of their global orientation, producing an accidental $O(2)$-like degeneracy of the low-energy configurations. This suppresses the effective fourfold anisotropy and accounts for the extended QLRO region. Close to $\theta = 180^\circ$, however, increasing temperature enhances local fluctuations and reduces bond redundancy. The degeneracy is then lifted and a $\mathbb{Z}_4$ anisotropy emerges, driving the reentrant transition from QLRO to LRO. The transition thus provides a classical realization of order by disorder in which thermal fluctuations strengthen the coupling between the collective orientation and the lattice.

These results explain why LRO can arise in reciprocal vision-cone models despite the absence of non-reciprocal dynamical couplings and stationary entropy production. Comparison with the NRXY model therefore separates two distinct roles: spin–lattice coupling controls the directional structure of orientational order, whereas non-reciprocity determines the nonequilibrium character of the stationary state.

The arguments developed here are based on geometrical, energetic, and entropic considerations and should be regarded as microscopic interpretations of the numerical phase diagrams rather than complete derivations of their long-wavelength theories. They nevertheless identify the

essential ingredients that such theories must retain: the orientation dependence of the active-bond network, the resulting effective lattice anisotropy and how it changes with temperature-dependent fluctuations.

Chapter 4

Reactive and proactive dynamics in non-reciprocal XY models

The results presented in Chapter 2 for the NRXY model were obtained using a Glauber-type dynamics based on the local quantity that we termed the *selfish energy*. In equilibrium systems, different stochastic dynamics satisfying detailed balance may be used to sample the same Boltzmann distribution. Their relaxation properties may differ, but their stationary thermodynamic behavior is determined by a common underlying energy functional. This equivalence does not hold generically in non-reciprocal systems, for which a global energy governing the many-body dynamics cannot, in general, be defined. The microscopic update rule then becomes part of the physical specification of the model and may affect not only its relaxation, but also its stationary properties.

In this chapter, we investigate this dynamical dependence in a broader class of NRXY models, including both the hard vision cone considered in the previous chapters and smooth generalizations of the corresponding interaction kernel. Starting from the same local selfish energy, we compare a Glauber-type implementation with a continuous-time Langevin dynamics obtained from the local angular derivative of the selfish energy. For a differentiable interaction kernel, the resulting deterministic drift naturally separates into two contributions.

The first is a conventional alignment torque, weighted by the instantaneous coupling to each neighbor. It describes a *reactive* response, in which the spin adjusts its orientation according to the neighboring orientations that it currently perceives. The second contribution originates from the derivative of the interaction kernel with respect to the spin orientation and it tends to increase the agent's effective number of perceived neighbors that are partially aligned with it. We refer to this kernel-gradient contribution as the *proactive* term.

We characterize the differences between a reactive-only Langevin dynamics and the complete dynamics containing both reactive and proactive contributions. To this end, we first introduce a mean-field argument relating the angular dependence of the total coupling to the expected pinning directions. We then numerically study the relaxation of the mean orientation for four different interaction kernels. The principal result is that the reactive-only and reactive-plus-proactive dynamics can select different directions of the collective polarization relative to the underlying lattice. Thus, in non-reciprocal systems, lattice pinning is controlled not only by the geometry of the interaction kernel, but also by the microscopic dynamical rule through which that kernel is

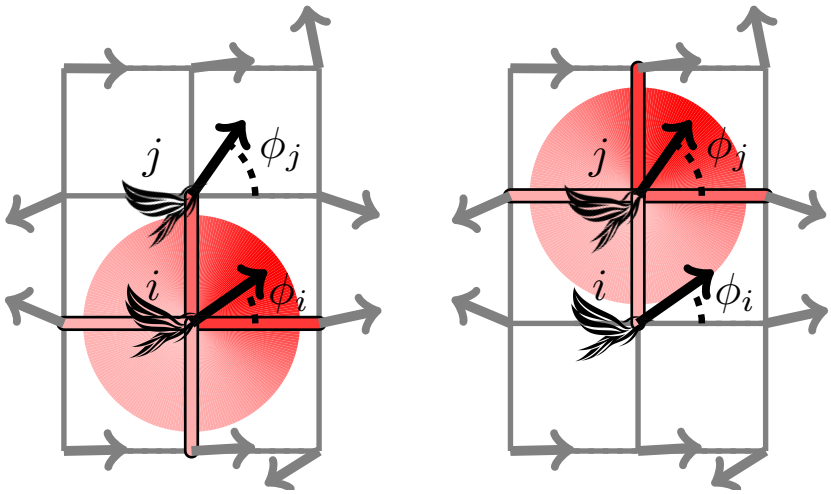


Figure 4.1: Schematic illustration of the non-reciprocal XY model with vision-induced *smooth* anisotropic interactions. The coupling between a spin (black arrow, representing for example a bird in a flock) and one of its nearest neighbors depends on the angle formed by the spin orientation and the bond connecting the two sites (e.g., $\langle i, j \rangle$), with stronger interactions occurring at smaller angles, i.e., with the neighbors situated in the direction indicated by the spin. The color intensity (red scale) represents the strength of the interaction as a function of the relative direction of the bond compared to that of the spin ϕ_i on the left panel and ϕ_j on the right one. In both panels, the resulting coupling strengths along each bond are also indicated. Because the interaction strength associated with a bond depends on the orientation of the spin that generates it, the same bond $\langle i, j \rangle$ generally contributes differently to the dynamics of spins i and j. This breaks reciprocity, as illustrated by the different color intensities assigned to the bond $\langle i, j \rangle$ in the two panels.

implemented.

The rest of the chapter is organized as follows. In Section 4.1 we introduce the models and the implementations of their dynamics. Section 4.1.1 introduces the generalized selfish energy and compares its Glauber-type and Langevin implementations. Section 4.1.2 separates the Langevin drift into reactive and proactive contributions and discusses their physical interpretation. Section 4.1.3 presents the four interaction kernels considered in the numerical analysis. Section 4.2 introduces the zero-temperature mean-field argument used to rationalize the expected pinning directions. Section 4.3 then compares the transient orientational selection produced by the different kernels and dynamical prescriptions. Finally, Sec. 4.4 summarizes the implications of these results.

4.1 From the selfish energy to the Langevin description

4.1.1 Glauber vs. Langevin dynamics

The non-reciprocal XY model with a generalized vision-mediated interaction is defined in a similar fashion as the sharp vision cone one introduced in Chapter 2: we are still in two spatial dimensions on a square lattice in which spin i, located at position $\mathbf{r}_i$, interacts with neighboring spins depending on its current orientation, represented by the angle ϕ_i it forms with a chosen lattice axis. Here we implement the dynamics of each spin in two closely related, yet physically distinct, ways, discussed further below. Both cases involve the "selfish" local energy functional, which for spin i reads

$$E_i = -J \sum_{j \in \mathcal{N}_i} g(\phi_i, \vartheta_{ij}) \cos(\phi_i - \phi_j), \tag{4.1}$$

where $\mathcal{N}_i$ denotes the set of the four nearest neighbors of spin i, and $\vartheta_{ij} = \arg(\mathbf{r}_j - \mathbf{r}_i)$ is the angle (integer multiple of $\pi/2$) of the lattice bond connecting sites i and j. The function $g(\phi_i, \vartheta_{ij})$ defines the *interaction kernel*, which modulates the coupling strength between spin i and j according to the field of view of spin i. Typically, spin i aligns more strongly with neighboring spins j whose relative positions $\mathbf{r}_j - \mathbf{r}_i$ lie along its current orientation ϕ_i, thereby introducing non-reciprocal interactions, as illustrated in Fig. 4.1.

Glauber dynamics

In the stochastic Glauber formulation of the (discrete-time) dynamics, each spin is updated according to the transition rates W that depend on the change of the local energy E_i associated with the attempted update. In practice, one runs a Monte Carlo simulation in which, at each step of the dynamics, a spin (say, i) is randomly selected for an attempted update of its orientation ϕ_i to ϕ_i', which is accepted with probability

$$W(\phi_i \to \phi_i') = \frac{1}{1 + \exp[(E_i' - E_i)/T]}. \tag{4.2}$$

Here T plays the role of an effective temperature while E_i and E_i' are the selfish energies of the spin configuration before and after the attempted change of orientation.

Langevin dynamics

An alternative continuous-time dynamics of the system is obtained on the basis of a Langevin equation which involves E_i and which drives the system towards its stationary points. In practice, the derivative of Eq. (4.1) with respect to ϕ_i provides the deterministic driving force of the evolution of ϕ_i to which a noise η_i is added. Accordingly, the resulting equation of motion is

$$\partial_t \phi_i = -J \sum_{j \in \mathcal{N}_i} g(\phi_i, \vartheta_{ij}) \, \sin(\phi_i - \phi_j) + J \sum_{j \in \mathcal{N}_i} \partial_{\phi_i} g(\phi_i, \vartheta_{ij}) \, \cos(\phi_i - \phi_j) + \sqrt{2T}\, \eta_i(t), \tag{4.3}$$

where η_i is a zero-mean Gaussian white noise with $\langle \eta_i(t) \eta_j(t') \rangle = \delta_{ij} \delta(t - t')$ and T plays the role of an effective temperature. Note that, at thermal equilibrium (which, inter alia, requires g to be independent of ϕ_i), both dynamics relax to the Boltzmann distribution $P \propto e^{-E/T}$, where $E = \sum_i E_i/2$. For non-reciprocal kernels, however, detailed balance is broken, and the two dynamics generally reach distinct stationary states.

The first term on the r.h.s. of Eq. (4.3) corresponds to the familiar torque of the XY model, which tends to align the spins, here weighted by the anisotropic kernel $g(\phi_i, \vartheta_{ij})$. The second term, proportional to the derivative $\partial_{\phi_i} g$, has no counterpart in equilibrium [65, 138] and is discussed in more detail in the following.

Although the Glauber and Langevin dynamics generally lead to different non-equilibrium stationary states, they originate from the same selfish energy and therefore share the same underlying microscopic interactions. One may thus expect several qualitative features of the stationary behavior to be preserved. On the other hand, removing the contribution associated with the derivative of the interaction kernel modifies the microscopic dynamics itself, and may therefore lead to qualitatively different behavior. As we shall see in Sec. 4.3, this is indeed the case for the pinning direction of the global magnetization.

4.1.2 From active to proactive matter

The second term on the r.h.s. of Eq. (4.3), proportional to $\partial_{\phi_i} g(\phi_i, \vartheta_{ij})$, is often omitted in the Langevin formulation of the non-reciprocal XY model, see, e.g., Refs. [84, 92, 93]. As we shall demonstrate below, this term actually introduces a qualitatively new ingredient in the microscopic dynamics, the physical interpretation of which can be understood by comparing the torques generated separately by the two contributions on the r.h.s. of Eq. (4.3):

- **Reactive alignment.** The first term, i.e.,

$$-J \sum_{j \in \mathcal{N}_i} g(\phi_i, \vartheta_{ij}) \sin(\phi_i - \phi_j), \tag{4.4}$$

  describes the usual interaction which favors spin alignment and which is the only present in the equilibrium XY model, in which g is constant. Specifically, spin i tends to rotate towards the orientations of its visible neighbors, with an effective coupling strength modulated by the kernel g as illustrated in Fig. 4.2(a). This mechanism is analogous to the alignment rule in Vicsek-type models of active matter, where each agent reacts by adjusting its orientation according to what it currently perceives in its local environment. Henceforth we refer to this as the *reactive terms*. Within this interaction term, the kernel g sets the relative weight of the contribution of each neighbor to the alignment dynamics.

- **Proactive adjustment.** The second term on the r.h.s. of Eq. (4.3), i.e.,

$$+J \sum_{j \in \mathcal{N}_i} \partial_{\phi_i} g(\phi_i, \vartheta_{ij}) \cos(\phi_i - \phi_j), \tag{4.5}$$

  has a qualitatively different effect. In fact, it causes spin i to rotate in the direction that increases (decreases) the value of the coupling kernel $g(\phi_i, \vartheta_{ij})$, when $\cos(\phi_i - \phi_j) > 0 \, (< 0)$. In practice, this term generates a torque on spin i which tends to maximize the effective number of visible neighbors that are at least partially aligned with that spin, as schematically represented in Fig. 4.2(b). Here, by partial alignment between two spins i and j we mean that one has a positive projection onto the other, i.e., $\cos(\phi_i - \phi_j) > 0$. We emphasize that this term is proportional to the gradient of the "perception" kernel g. The resulting feedback mechanism therefore induces a tendency to *maximise the perceived alignment* itself, rather than merely responding to it.

The behavior of the spin induced by the second term discussed above goes beyond conventional activity. For this reason, we refer to it as *proactivity*: each agent modifies its state (i.e., its orientation) in order to enhance its own sensory input (i.e., the effective number of perceived neighbors which are not misaligned), rather than merely aligning passively with it. At the collective level, proactivity can stabilize orientational order and suppress the nucleation of defects (the so-called asters [92, 93, 139]) in simulations of the Langevin dynamics in Eq. (4.3). In fact, the role of this additional proactive contribution in promoting defect annihilation during coarsening starting from a disordered initial state was previously highlighted in Ref. [65]. On this subject, we propose a conjecture and a preliminary result in Chapter 6.

It is worth mentioning that the Langevin dynamics in Eq. (4.3), which includes both the reactive and proactive terms, can be directly mapped onto the Glauber dynamics associated with the selfish

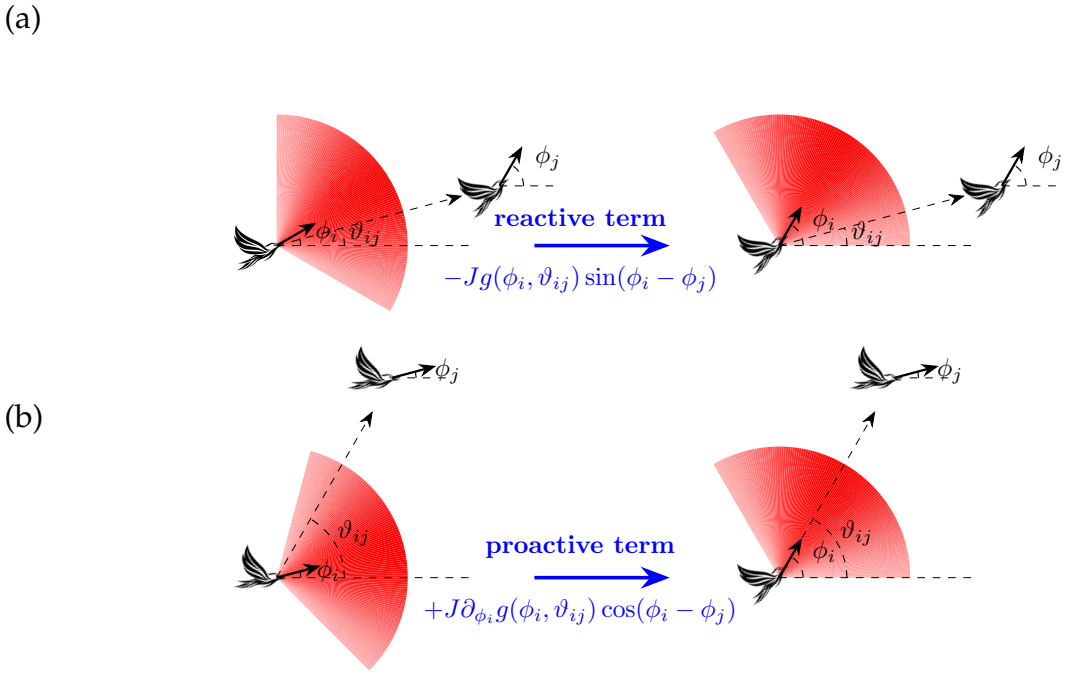


Figure 4.2: Comparison between the reactive and the proactive term in the Langevin dynamics of the non-reciprocal XY model. (a) The reactive term tends to align the orientation of an agent (e.g., a bird at site i) with the orientation of the neighbors it perceives (e.g., the one at site j), with a strength modulated by the kernel $g(\phi_i, \vartheta_{ij})$, such that $\phi_i \simeq \phi_j$. (Here ϑ_{ij} is the angle connecting the sites i and j, where $\vartheta_{ij} = n\pi/2$ with $n = 0, 1, 2, 3$ on the square lattice.) In contrast, (b) the proactive term drives the agent to reorient along the direction of the unit vector connecting it to its neighbors, such that $\phi_i \simeq \vartheta_{ij}$. This occurs through an interaction mechanism sensitive to the gradient of the vision coupling, effectively pushing the agent to maximize its sensory input, i.e., the effective number of visible neighbors which are not misaligned with the agent. As a consequence, the proactive contribution is sensitive to variations in the visual field: when a neighbor is about to exit the field of view, this term induces a reorientation that tends to keep it within sight.

energy in Eq. (4.1). By contrast, the Langevin dynamics with the sole reactive term can be simulated through Monte Carlo dynamics only by using the auxiliary Hamiltonian introduced in Ref. [138].

4.1.3 Interaction kernels

The physical behavior of the non-reciprocal XY model is dictated by the choice of the functional form of the interaction kernel $g(\phi_i, \vartheta_{ij})$. For concreteness, we consider here the following four representative cases, illustrated in Fig. 4.3:

1. **Modulated coupling (no vision bias)** [92]:

$$g(\phi_i, \vartheta_{ij}) = 1 + \epsilon \cos(4\phi_i), \tag{4.6}$$

which breaks rotational symmetry via an explicit fourfold modulation of the coupling which is independent of the spatial perception of spin i, i.e., $g(\phi_i, \vartheta_{ij})$ is actually independ of ϑ_{ij}. The parameter ϵ controls the degree of anisotropy. In the present work, we consider couplings which favor the alignment of neighboring spins an thus we assume $\epsilon \in [0, 1]$.

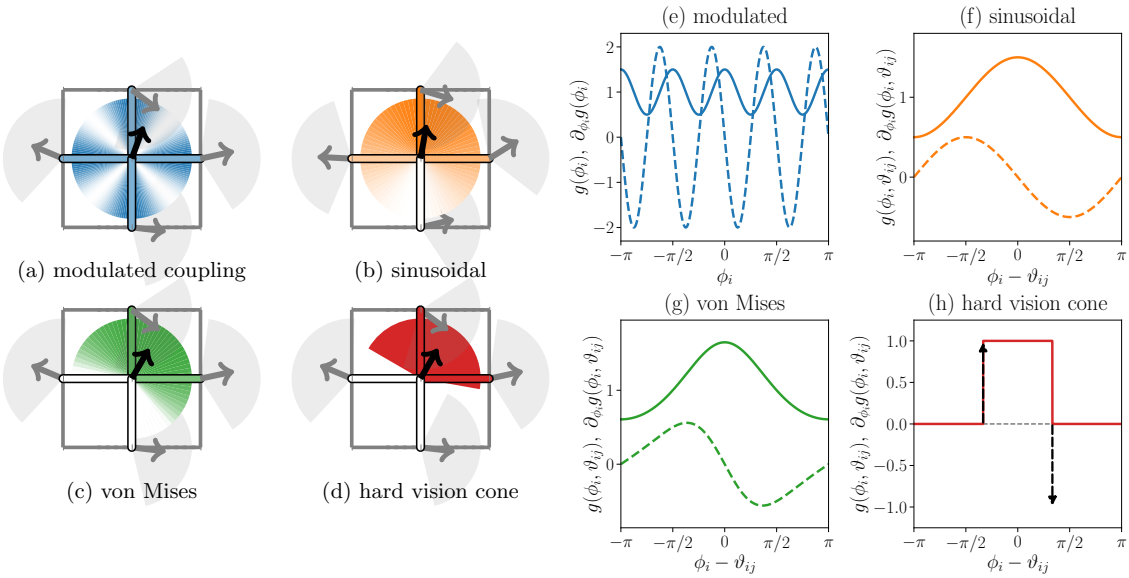


Figure 4.3: Representation of the interaction kernels considered in this work. Panels (a–d) illustrate the interaction strength, while panels (e–h) display the corresponding functional form g of the kernels (solid lines) and of their derivatives $\partial_{\phi_i} g$ (dashed lines). (a, e) Modulated coupling (Eq. (4.6)): all nearest-neighbors bonds have identical strength, with no vision bias or anisotropy (g is not function of ϑ_{ij}), although the interaction is modulated by the spin orientation. (b, f) Sinusoidal coupling (Eq. (4.7)): forward bonds are enhanced through a smooth sinusoidal angular modulation. (c, g) Von Mises coupling (Eq. (4.8)): similar to the sinusoidal case, but with a different angular modulation, characterized by a forward bias which becomes increasingly more pronounced upon increasing the parameter σ; (d, h) Hard vision cone (Eq. (4.9)): interactions are restricted only to neighbors that lie within a finite vision cone. Accordingly, the kernel derivative exhibits two Dirac delta contributions (dashed arrows) at the boundaries of the vision cone.

2. **Sinusoidal kernel** [92, 93, 140]:

$$g(\phi_i, \vartheta_{ij}) = 1 + \epsilon \cos(\phi_i - \vartheta_{ij}), \tag{4.7}$$

where the coupling is maximal when the neighboring spin j lies along the spatial direction indicated by spin i.

3. **von Mises kernel** [65]:

$$g(\phi_i, \vartheta_{ij}) = \exp[\sigma \cos(\phi_i - \vartheta_{ij})], \tag{4.8}$$

where σ is a parameter of the model, which is characterized by a smooth, positive-definite angular dependence of the coupling strenght on the angle formed by the direction of the neighbors and the one of the spin. For small values of σ, this expression reduces to the sinusoidal kernel above with $\epsilon = \sigma$.

4. **Hard vision cone (HVC)** [1, 89]:

$$g(\phi_i, \vartheta_{ij}) = \begin{cases} 1 & \text{if} \quad \cos(\phi_i - \vartheta_{ij}) \geq \cos(\theta/2), \\ 0 & \text{otherwise.} \end{cases} \tag{4.9}$$

In this case, only neighbors which are within an angular sector of width $\theta \in [0, 2\pi]$ centered around the direction of spin i contribute to the interaction. This kernel introduces discontinuities and sharp geometric constraints.

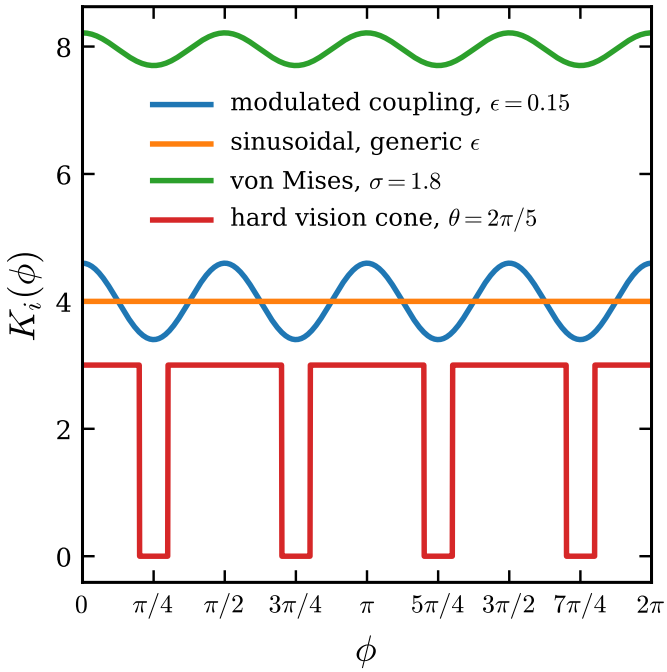


Figure 4.4: Total coupling $K_i(\phi)$ (see Eq. (4.10)) for a perfectly aligned spin configuration (corresponding to $T \to 0$) as a function of the global spin orientation ϕ, for the various interaction kernels considered in this work and listed in Sec. 4.1.3. The qualitative features of these curves are the same for generic choices of the parameters ϵ, σ, and θ.

Note that all these kernels, for either $\epsilon = \sigma = 0$ or $\theta = 2\pi$, render the equilibrium two-dimensional XY model with isotropic couplings.

4.2 Mean field: coupling strength at $T = 0$

We begin our analysis by considering the case in which fluctuations are negligible, i.e., in which all the spins are perfectly aligned. Formally, this corresponds to the limit $T \to 0$ of the dynamics in Eq. (4.3) with $\phi_i = \phi$ for all i, such that $\cos(\phi_i - \phi_j) = 1$ and $\sin(\phi_i - \phi_j) = 0$. Under this assumption, the *total coupling* $K_i(\phi)$ of spin i with *all* its nearest neighbors can be defined, on the basis of Eq. (4.1), as

$$K_i(\phi) = \sum_{j \in \mathcal{N}_i} g(\phi, \vartheta_{ij}); \tag{4.10}$$

this quantity depends only on the global orientation ϕ of the spins and is actually independent of the site index i. Figure 4.4 shows $K_i(\phi)$ as a function of ϕ for the different kernels considered in this work and summarized in Fig. 4.3.

In order to rationalize the emergence of ordering in the system, a key question to be addressed is whether the dynamics drives the system to some stationary point of the total coupling $K_i(\phi)$. In fact, if this is the case, the spins are expected to be pinned collectively with a value of ϕ corresponding to one of these stationary points.

In the presence of the reactive term alone, $K_i(\phi)$ is indeed *minimized*, as demonstrated for the modulated coupling kernel and further generalized in Ref. [92]. In fact, when the mean direction ϕ corresponds to a minimum of the coupling strength $K_i(\phi)$, a spin that fluctuates away from ϕ

experiences a stronger interaction with its neighbors than the spins that remain aligned along ϕ. Consequently, the fluctuating spin is pulled back toward the mean orientation more strongly than the rest of the system is pulled toward the fluctuating spin. This imbalance generates an effective restoring torque that stablizes the configuration around the direction of weakest coupling. Accordingly, configurations with weaker effective coupling are dynamically favored.

When the proactive term is added to the reactive one, the situation is reversed. As discussed in Sec. 4.1.2, the additional contribution in the Langevin equation generates a torque that tends to increase the "sensory input" of each spin, thereby *maximizing* its coupling $K_i(\phi)$ to the visible neighbors which are not misaligned. This can be understood more clearly by considering the Glauber dynamics associated with the selfish energy in Eq. (4.1), which naturally corresponds to the Langevin dynamics with both the reactive and proactive contributions. For an ordered state, the selfish energy in Eq. (4.1) becomes $E_i = -JK_i(\phi)$ and thus it is minimized when the coupling $K_i(\phi)$ is maximized, confirming the tendency towards configurations with stronger coupling.

Inspecting Fig. 4.4, we observe that the total coupling $K_i(\phi)$ is actually independent of ϕ for the sinusoidal kernel, while it exhibits a periodic modulation for the modulated coupling, the von Mises kernel, and the hard vision cone (see Sec. 4.1.3 for their definitions). This modulation reflects the $\mathbb{Z}_4$ symmetry of the underlying lattice (explicitly imposed in the case of the modulated coupling). Based on this plot, the argument presented above about the minimization/maximization of $K_i(\phi)$ by the reactive/total dynamics allows one to infer the preferred pinning directions for these three kernels, both for the sole reactive dynamics and for the dynamics that includes the proactive contribution. For example, for the modulated coupling and the von Mises kernel, the reactive dynamics promotes ordering with $\phi = \pi/4 + n\pi/2$, with $n = 0$, 1, 2, and 3, while the complete one with $\phi = n\pi/2$. These directions are summarized in Table 4.1, which will be illustrated and discussed in, c.f., Sec. 4.3 together with the numerical evidence.

An important observation is that, for the von Mises coupling kernel, the leading contribution that generates a modulation of the total coupling $K_i(\phi)$ arises only at order σ^4. This fact will be taken into account in Sec. 5.2.3 where the corresponding equations of motion are derived.

The arguments presented above, applicable in the limit of vanishing noise $T \to 0$, also indicate that the sinusoidal kernel alone is not capable of pinning the spins to any specific orientation ϕ, neither with the reactive dynamics nor with the one which includes the reactive and proactive terms since, in both cases, $K_i(\phi) = 0$ independently of ϕ. Nevertheless, as shown in Ref. [92], finite-temperature fluctuations can induce a weak but measurable pinning effect even for the sinusoidal kernel. Interestingly, fluctuations can also invalidate the mean-field prediction even when the coupling $K_i(\phi)$ has a finite but small modulation. We observe this behavior for the von Mises coupling on the triangular lattice, where the anisotropy of the coupling is weaker than on the square lattice (see the Supplementary Information of Ref. [65]). The corresponding analysis is presented in App. B.1.

In the next section, we systematically analyze how these pinning terms emerge in the presence of fluctuations, both for the sinusoidal and for the other interaction kernels considered in this work.

4.3 Numerical results for the mean orientation

In this section, we present the results of the numerical simulations of the non-reciprocal XY model, considering the different dynamical implementations and interaction kernels introduced in Sec. 2.1.

Table 4.1: Summary of the lattice-induced pinning directions for the interaction kernels and dynamical prescriptions considered. For the hard vision-cone kernel, $q = \lfloor \theta/(\pi/2) \rfloor \in \{0, 1, 2, 3\}$ labels the angular interval containing the aperture θ and determines the alternation between axial and diagonal pinning.

Kernel type	**Functional form**	**Vision bias**	**Pinning direction**	
			Reactive	**Reactive + proactive**
Modulated coupling	$1 + \epsilon \cos(4\phi_i)$	None	$\frac{\pi}{4} + n\frac{\pi}{2}$	$n\frac{\pi}{2}$
Sinusoidal	$1 + \epsilon \cos(\phi_i - \vartheta_{ij})$	Soft	$n\frac{\pi}{2}$	$n\frac{\pi}{2}$
von Mises	$\exp\left[\sigma \cos(\phi_i - \vartheta_{ij})\right]$	Soft	$\frac{\pi}{4} + n\frac{\pi}{2}$	$n\frac{\pi}{2}$
Hard vision cone	Step function with aperture θ	Hard	$\frac{\pi}{4}[1 - (q \bmod 2)] + n\frac{\pi}{2}$	$\frac{\pi}{4}(q \bmod 2) + n\frac{\pi}{2}$

In particular, we focus on the emerging mean orientation both during the transient dynamics and in the (quasi)stationary state, in order to compare the numerical results with the predictions formulated in the previous sections.

We performed Langevin simulations for the three smooth kernels (i.e., modulated, sinusoidal, and von Mises, see Sec. 4.1.3), considering either the reactive term alone (i.e., only the first term on the r.h.s. of Eq. (4.3)) or the complete dynamics which includes also the proactive term (i.e., both terms on the r.h.s. of Eq. (4.3)).

Note that, for the hard vision-cone kernel, only the reactive term can be straightforwardly implemented within the Langevin framework. In fact, the associated proactive contribution would involve Dirac delta functions in the corresponding Langevin equation, which are not amenable to numerical simulations. Accordingly, in the presence of the hard vision cone, instead of simulating the complete Langevin equation, we resort to Glauber dynamics (see Sec. 4.1.1) as a proxy for the former, which includes the proactive term. The Glauber update rule is constructed from the corresponding selfish energy, i.e., the energy functional whose variation generates the Dirac-delta terms appearing in the Langevin equations.

4.3.1 Transient dynamics

Starting from an initial configuration in which all the spins are aligned in the same direction, i.e., with the same value of ϕ_i, the system is allowed to evolve until all observables reach stationary values. In particular, the dynamics of the mean orientation $\phi(t)$ defined in Eq. (5.1) provides indication about the eventual pinning direction in phases exhibiting LRO. Even in those cases in which such a phase is known to be metastable due to eventual nucleation and proliferation of defects (for example the sinusoidal and hard vision cone with only the reactive term [92]), the parameters of the model are chosen such that, under Langevin dynamics, these processes occur on timescales much longer than the longest time reached in the present work [92, 139]. This separation of time scales allows us to gather sufficient statistics for a quantitative analysis of the effective pinning mechanisms.

Figure 4.5 reports the numerical data corresponding to the transient dynamics. In particular, starting from the initial condition $\phi_i(0) = \phi(0) = \pi/8$ for all i, the system always relaxes towards

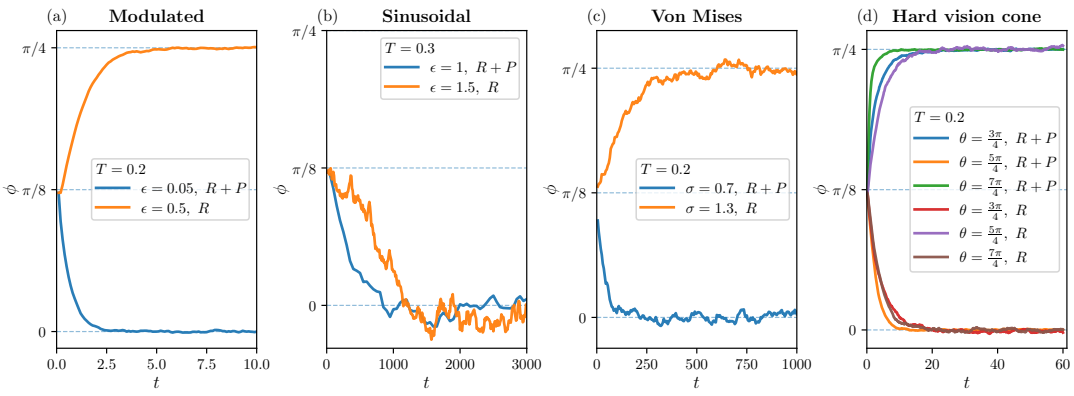


Figure 4.5: Results of the numerical simulation of the transient dynamics of the mean orientation $\phi(t)$ (see Eq. (5.1)) for the four interaction kernels considered ($L = 200$ for all simulations). Starting from an initial condition $\phi(0) = \pi/8$, we monitor whether the system relaxes toward $\Phi = 0$, corresponding to pinning along the lattice directions $n\pi/2$ with integer n, or toward $\Phi = \pi/4$, corresponding to pinning along the diagonal directions $\pi/4 + n\pi/2$. The numerical results clearly show that, whenever the mean-field argument presented in Sec. 4.2 — concerning the total coupling $K_i(\phi)$ — applies, the reactive dynamics (curves labeled by R in the figure) tends to minimize the coupling (argument of Ref. [92]), whereas the complete Langevin dynamics (curves labeled by R + P) or, equivalently, the Glauber dynamics tend to maximize it. For the sinusoidal kernel (panel (b)), the mean-field argument is inconclusive. In this case, both these simulations and the analytical considerations discussed in Sec. 5.2.2 show that pinning occurs along the lattice directions, consistently with previous results [92]. For the hard vision cone (panel (d)), the orientation that maximizes or minimizes the number of visible neighbors depends explicitly on the vision cone aperture θ, as discussed in Refs. [1, 89]: as a consequence, the pinning angle varies accordingly. When comparing the evolution under different dynamics (namely, the reactive Langevin, the complete Langevin or, alternatively, the Glauber dynamics) the parameters have been chosen such that the corresponding relaxation timescales are comparable, see, e.g., panels (a), (b), and (c). The conclusions presented here regarding the selection of the pinning direction, turn out to be independent of the specific values of the parameters used in the simulation, as we confirmed for a number of different choices.

one of the two possible states: $\Phi = 0$, corresponding to pinning along the lattice directions, or $\Phi = \pi/4$, corresponding to pinning along the diagonal directions. Which one is eventually selected depends on the choice of the kernel and of the dynamics, as shown by the various panels in the figure and as summarized in Table 4.1. For the modulated (see Fig. 4.5, panel (a)), von Mises (panel (c)), and hard vision cone (panel (d)) kernels, the observed pinning directions are consistent with the mean-field predictions discussed in Sec. 4.2. In fact, when only the reactive term is included in the dynamics (orange curves in Fig. 4.5), the selected direction corresponds to the minimum of the total coupling $K_i(\phi)$. This yields $\Phi = \pi/4$ for the modulated and von Mises kernels (panels (a) and (c) in Fig. 4.5) and

$$\Phi = \frac{\pi}{4}\left[1 - (q \bmod 2)\right] \quad \text{with} \quad q = \lfloor \theta/(\pi/2) \rfloor \in \{0, 1, 2, 3\} \tag{4.11}$$

for the hard vision cone kernel (panel (d)) with vision cone aperture θ. In practice, in the latter case, $\Phi = \pi/4$ for $0 < \theta < \pi/2$ or $\pi < \theta < 3\pi/2$, while $\Phi = 0$ otherwise. Upon including also the proactive term in the Langevin dynamics (or, equivalently, for the hard vision cone kernel, upon employing the corresponding Glauber dynamics), the selected direction (indicated by the blue curves in Fig. 4.5) corresponds, instead, to the maximum of the total coupling. Accordingly, one finds $\Phi = 0$ for the modulated and von Mises kernels (panels (a) and (c)), and

$$\Phi = \frac{\pi}{4}\,(q \bmod 2) \tag{4.12}$$

for the hard vision cone (panel (d)), i.e., $\Phi = \pi/4$ for $\pi/2 < \theta < \pi$ or $3\pi/2 < \theta < 2\pi$ while $\Phi = 0$ otherwise. These results are consistent with the findings reported in Ref. [92] where it was also noted, without further explanation, that the direction of the observed pinning changes when considering the Monte Carlo dynamics instead of the Langevin dynamics. The analysis presented here explains this observation. Notably, the sinusoidal kernel is the only case considered here in which the addition of the proactive term to the reactive dynamics does not alter the pinning direction, which remains aligned with the lattice axes ($\Phi = 0$), as shown by the curves in Fig. 4.5(b). Finally, we note that, in the limit $\sigma \ll 1$, the von Mises coupling becomes asymptotically equivalent to the sinusoidal coupling . It is therefore natural to ask whether there exists a threshold value σ^* (with $\sigma^* \ll 1$) at which, in the purely reactive dynamics, the preferred pinning direction changes abruptly from the lattice axes ($\Phi = 0$ for $\sigma < \sigma^*$) to the diagonal direction ($\Phi = \pi/4$ for $\sigma > \sigma^*$). At present, our numerical results at very small values of σ do not allow us to draw a definite conclusion. Resolving this issue would likely require simulations on significantly larger system sizes than those considered here. On the other hand, a crossover in the pinning direction as a function of σ turns out to occur on the triangular lattice, as detailed in App. B.1.

4.4 Concluding remarks

In this chapter, we have shown that the microscopic dynamical prescription is an essential part of the physical definition of a non-reciprocal XY model.

Starting from the local selfish energy, we identified two physically distinct contributions to the corresponding Langevin dynamics. The reactive term aligns each spin with the orientations of the neighbors that it currently perceives, with weights fixed by the instantaneous interaction kernel. The proactive term instead originates from the derivative of the kernel with respect to the

spin orientation. For vision-dependent kernels, this mechanism provides a microscopic realization of proactive sensing: agents do not merely respond to the information they are perceiving, but modify their orientation according to how their coupling with the environment is changing.

The angular dependence of the total coupling $K_i(\phi)$ provides a simple organizing principle for the resulting orientational selection. On a perfectly aligned configuration, the reactive-plus-proactive contribution generates a direct drift toward maxima of $K_i(\phi)$, which correspond to minima of the local selfish energy. By contrast, the reactive term alone favors orientations where the effective coupling is weaker. The two mechanisms can therefore select complementary directions even though they are constructed from the same interaction kernel.

The numerical results confirm this distinction. For the modulated and von Mises kernels, reactive-only dynamics selects the diagonal directions of the square lattice, whereas the complete reactive-plus-proactive dynamics selects the lattice axes. For the sinusoidal kernel, the aligned-state coupling is independent of the global orientation, and both dynamical prescriptions display a weaker, fluctuation-induced pinning along lattice directions. For the hard vision cone, both the reactive Langevin dynamics and the selfish-energy Glauber dynamics alternate between axial and diagonal preferred orientations whenever the aperture crosses an integer multiple of 90°. However, the two dynamics select opposite directions: the reactive Langevin dynamics favors orientations that minimize the number of active bonds, whereas the Glauber dynamics favors those that maximize it.

These results establish that the interaction kernel alone does not uniquely determine the large-scale behavior of a non-reciprocal orientational system. The same anisotropic coupling can produce different ordered states depending on whether agents react to their current sensory input or proactively modify that input through their own reorientation. This distinction provides a unified microscopic interpretation of previously observed differences between Monte Carlo and Langevin implementations and motivates the more detailed analysis of local and global pinning developed in the following chapter.

Chapter 5

Local and global lattice pinning in non-reciprocal XY models

The previous chapter showed that the orientation selected by the NRXY models depends both on the interaction kernel and on the microscopic implementation of the dynamics. In particular, reactive-only and reactive-plus-proactive dynamics may pin the collective polarization along different directions of the underlying square lattice. The aim of the present chapter is not to determine these directions again, but to characterize quantitatively the restoring mechanisms that stabilize them.

A natural question is whether the pinning of the collective orientation must be visible directly in the equation of motion of each individual spin. In equilibrium models with an explicit external anisotropy, the same restoring coefficient commonly appears in the local dynamics and in the equation governing the global orientation. This correspondence need not hold in a non-reciprocal system. Because directed interaction terms do not cancel pairwise when the microscopic equations are spatially averaged, an effective restoring torque may emerge for the collective orientation even when no explicit pinning term is present in the local equation of motion.

To investigate this distinction, we introduce separate measures of local and global pinning. The local coefficient $r_i^{\mathrm{loc}}(t)$ is defined from the term linear in the fluctuation of an individual spin around the selected orientation. Its average r^{loc} characterizes the restoring tendency visible at the single-spin level, while its fluctuations $\Delta_{r^{\mathrm{loc}}}$ quantify the spatial and temporal variability of this local restoring force. The global coefficient r^{glob}, instead, governs the relaxation of the collective orientation around its pinned direction. These quantities are conceptually distinct and, in general, need not coincide.

For the smooth interaction kernels introduced in the previous chapter, we derive the local and global equations of motion by expanding their reactive and proactive contributions into angular harmonics. This allows us to identify which harmonics generate an explicit local restoring force and which appear only after the microscopic dynamics is spatially averaged.

We subsequently compare the analytically defined coefficients with two independent numerical estimators. The first is the relaxation rate r^{OU} extracted from the autocorrelation of the mean orientation. The second estimator is the pinning $r^{S(\mathbf{k})}$ obtained from the structure factor of stationary configurations. It characterizes spatial fluctuations but does not necessarily coincide with the average local coefficient r^{loc}. We determine the conditions under which these comparisons

are meaningful and the regimes in which fluctuations prevent a direct comparison.

The rest of the chapter is organized as follows. Section 5.1.1 introduces the local and global orientational fluctuations and defines the corresponding pinning coefficients. Sections 5.2.1, 5.2.2, and 5.2.3 derive the pinning terms generated by the modulated, sinusoidal, and von Mises kernels, respectively. Section 5.2.4 summarizes the occurrence of explicit local pinning for the different dynamical contributions. Section 5.2.5 evaluates the global pinning coefficients in the long-lived pinned states and verifies their stability. Section 5.3 then compares these analytical coefficients with the relaxation rate of the collective orientation and with the effective pinning extracted from the static structure factor. Finally, Sec. 5.4 summarizes the main results and discusses their limitations and implications.

5.1 Search for a pinning term

When the system is in a (quasi)stationary long-range ordered (LRO) state at small but finite temperatures, the global orientation of the spins remains pinned along specific lattice directions. The selection of these directions depends on both the details of the dynamics and the form of the interaction kernel, as already discussed in Sec. 4.2 and as it will be proved numerically in, c.f., Sec. 4.3. In this section, we analyze how pinning emerges at the level of the equations of motion for the modulated, sinusoidal, and von Mises interaction kernels. The hard vision cone kernel, instead, is not amenable to an analytical treatment because of the discontinuities it generates in the corresponding Langevin equation. This is the reason why we will investigate it only numerically in, c.f., Sec. 4.3.

5.1.1 Local and global equations of motion

In order to write the equations of motion for both the local and the global spin orientation, we introduce the following notation. We denote by $\phi_i(t)$ the local spin at lattice site i, and by $\phi_{\mathbf{x}}(t)$ its counterpart on the continuum space, which will enter the evolution equations discussed further below. The global orientation of the model at time t is defined as the instantaneous spatial average,

$$\phi(t) = \langle \phi_i(t) \rangle = \frac{1}{N} \sum_{i=1}^{N} \phi_i(t), \tag{5.1}$$

where $N = L^2$ is the total number of spins in the (finite) system and $\langle \cdots \rangle$ denotes the spatial average over the lattice sites. When, in its (quasi)stationary state (i.e. stable or metastable), the system is pinned along a specific direction, the time average of the global orientation is

$$\Phi = \overline{\phi(t)} = \overline{\langle \phi_i(t) \rangle}. \tag{5.2}$$

Here, the overbar indicates a time average in the (quasi)stationary state. For the systems considered in this work, the long-time average orientation Φ is selected from a discrete set of angles determined by the symmetry of the square lattice. In particular,

$$\Phi \in \left\{ n\frac{\pi}{2}, \frac{\pi}{4} + n\frac{\pi}{2} \right\} \quad \text{with} \quad n = 0, 1, 2, 3, \tag{5.3}$$

corresponding to alignment along the principal lattice directions and along the lattice diagonals, respectively.

Fluctuations. We now consider the global and local fluctuations. In particular, the fluctuation $\delta\phi(t)$ of the global orientation $\phi(t)$ in the stationary state with respect to its average direction Φ is defined as

$$\delta\phi(t) = \phi(t) - \Phi. \tag{5.4}$$

At the local level, it is useful to distinguish between fluctuations with respect to the instantaneous global orientation, i.e.,

$$\delta\tilde{\phi}_i(t) = \phi_i(t) - \phi(t), \tag{5.5}$$

and fluctuations with respect to the long-time average orientation,

$$\delta\phi_i(t) = \phi_i(t) - \Phi. \tag{5.6}$$

Clearly, $\delta\phi_i(t) = \delta\tilde{\phi}_i(t) + \delta\phi(t)$. Analogous definitions hold for the fluctuations of the field $\phi_{\mathbf{x}}(t)$ on the continuum. We point out that in the thermodynamic limit ($N \to \infty$), the mean orientation $\phi(t)$ in the stationary state has vanishing fluctuation $\delta\phi(t)$, i.e., $\phi(t) \to \Phi$ and therefore $\delta\tilde{\phi}_i(t) \to \delta\phi_i(t)$.

Evolution of the local fluctuations. To determine the pinning terms acting on the local variable ϕ_i, we consider Eq. (4.3), substitute the explicit expressions of the interaction kernels, and expand the sums over nearest neighbors. The final result of this procedure can be conveniently cast in the form

$$\partial_t\phi_i = \sum_{\alpha\in\{\mathrm{R},\mathrm{P}\}}\sum_{n=0}^{\infty}\left[S_{\alpha,i}^{(n),\mathrm{loc}}(t)\,\sin(n\phi_i) + C_{\alpha,i}^{(n),\mathrm{loc}}(t)\,\cos(n\phi_i)\right] + \sqrt{2T}\eta_i(t), \tag{5.7}$$

where the analytic expressions of the coefficients $S_{\alpha,i}^{(n),\mathrm{loc}}(t)$ and $C_{\alpha,i}^{(n),\mathrm{loc}}(t)$ are derived in detail in Apps. B.2.1, B.2.2, and B.2.3. These coefficients depend on the instantaneous configuration of the neighbors of spin i and are functions of the local angular differences, i.e.,

$$S_{\alpha,i}^{(n),\mathrm{loc}}(t) = F_{\alpha}^{(n),\mathrm{loc}}\Big(\{\phi_i(t) - \phi_j(t)\}_{j\in\mathcal{N}_i}\Big), \tag{5.8}$$

and similarly for $C_{\alpha,i}^{(n),\mathrm{loc}}(t)$. In Eq. (5.7) the sum over $\alpha \in \{\mathrm{R},\mathrm{P}\}$ contains separately the contributions due to the reactive (R) and proactive (P) terms in the equation of motion (4.3).

When the spins are pinned around a given direction Φ, Eq. (5.7) can be expanded in terms of the fluctuations $\delta\phi_i$, defined in Eq. (5.6). The resulting equation generally takes the form

$$\partial_t\delta\phi_i = \kappa(\delta\phi_i)\,\nabla_{\mathrm{d}}^2\delta\phi_i - r_i^{\mathrm{loc}}(t)\,\delta\phi_i + \sqrt{2T}\eta_i(t) + \mathcal{A}_i, \tag{5.9}$$

where $\kappa(\delta\phi_i)$ is a modulated diffusion coefficient and ∇_{d}^2 denotes the discrete Laplacian

$$\nabla_{\mathrm{d}}^2\delta\phi_i = \frac{1}{a^2}\sum_{j\in\mathcal{N}_i}(\delta\phi_j - \delta\phi_i). \tag{5.10}$$

Here, a indicates the lattice spacing, which sets the unit of length (throughout this work we set $a = 1$). The effective diffusion term $\kappa(\delta\phi_i)\nabla_{\mathrm{d}}^2\delta\phi_i$ controls the propagation of local fluctuations, while the term linear in $\delta\phi_i$ acts as a local restoring force with the local pinning coefficient $r_i^{\mathrm{loc}}(t)$

that depends on the harmonic coefficients $S_{\alpha,i}^{(n),\mathrm{loc}}(t)$ and $C_{\alpha,i}^{(n),\mathrm{loc}}(t)$. Instead, $\mathcal{A}_i$ collects all the contributions arising from the original equation of motion in Eq. (5.7), which are not accounted for by the first two contributions.

The stability of the fluctuations $\delta\phi_i(t)$ is controlled by the average value of the local pinning coefficient

$$r^{\mathrm{loc}} = \overline{r_i^{\mathrm{loc}}(t)}, \tag{5.11}$$

which needs to be positive, i.e.,

$$r^{\mathrm{loc}} > 0. \tag{5.12}$$

Note that the average r^{loc} in Eq. (5.11) does not actually depend on i, due to the translational invariance of the system. In the simplest approximation, r^{loc} determines the long-time relaxation rate of $\delta\phi_i(t)$ following an initial fluctuation $\delta\phi_i(t=0)$. The coefficient r^{loc} might therefore be interpreted as the pinning strength controlling spatial fluctuations. In general, however, this identification is only approximate. Indeed, the effective pinning governing the long-wavelength dynamics is generally renormalized both by the fluctuations of $r_i^{\mathrm{loc}}(t)$ and by the additional contributions collected in $\mathcal{A}_i$. A useful measure of the magnitude of the first effect is provided by the standard deviation

$$\Delta_{r^{\mathrm{loc}}} = \sqrt{\overline{\left(r_i^{\mathrm{loc}}(t) - r^{\mathrm{loc}}\right)^2}}, \tag{5.13}$$

which quantifies the fluctuations of the local pinning coefficient around its mean value. When

$$\Delta_{r^{\mathrm{loc}}} \ll r^{\mathrm{loc}}, \tag{5.14}$$

the renormalization associated with the fluctuations of r_i^{loc} is expected to be weak, and r^{loc} provides a reliable approximation to the effective pinning governing the spatial dynamics. Conversely, when $\Delta_{r^{\mathrm{loc}}}$ becomes comparable or larger than r^{loc}, significant renormalization effects are expected and the effective pinning extracted from spatial correlations may differ substantially from r^{loc}. As discussed in Sec. 5.3.2, this criterion plays an important role when comparing analytical estimates of the pinning coefficient with those extracted from the static structure factor. More generally, although r^{loc} provides a natural measure of the local pinning strength, the effective pinning governing spatial fluctuations does not necessarily coincide with it.

Finally, note that the averages appearing in Eqs. (5.11) and (5.13), are taken over time. However, in the stationary state, the system is invariant both under time translations and, due to periodic boundary conditions, spatial translations. As a consequence, the quantities r^{loc} and $\Delta_{r^{\mathrm{loc}}}$ can be estimated equivalently from temporal or spatial averages. In practice, throughout this work we compute them from spatial averages over the lattice (see also App. B.2).

Evolution of the global fluctuation. To derive the evolution equation for the global fluctuation ϕ in Eq. (5.1), we again start from Eq. (4.3), substitute the explicit form of the interaction kernel, and sum both sides of the equation over all lattice sites in the system . This procedure leads to an exact expression of a form analogous to Eq. (5.7), i.e.,

$$\partial_t\phi = \sum_{\alpha\in\{\mathrm{R},\mathrm{P}\}}\sum_{n=0}^{\infty}\left[S_\alpha^{(n),\mathrm{glob}}(t)\,\sin(n\phi) + C_\alpha^{(n),\mathrm{glob}}(t)\,\cos(n\phi)\right] + \sqrt{\frac{2T}{N}}\eta(t). \tag{5.15}$$

The details of the derivation of the analytic expressions for the various interaction kernels considered in this work are reported in Apps. B.3.1, B.3.2, and B.3.3. The coefficients $S_\alpha^{(n),\mathrm{glob}}(t)$ and $C_\alpha^{(n),\mathrm{glob}}(t)$ in Eq. (5.15) involve contributions from all pairs of neighboring spins in the system and they can be expressed as functions of the fluctuations $\delta\tilde{\phi}_i(t)$ as

$$S_\alpha^{(n),\mathrm{glob}}(t) = \left\langle F_\alpha^{(n),\mathrm{glob}}\left(\delta\tilde{\phi}_i(t), \delta\tilde{\phi}_j(t)_{j\in\mathcal{N}_i}\right)\right\rangle, \tag{5.16}$$

and similarly for $C_\alpha^{(n),\mathrm{glob}}(t)$. Note that these coefficients depend on the fluctuations with respect to the instantaneous mean orientation $\phi(t)$, and not to its temporal average Φ.

Once the system is pinned, we can expand Eq. (5.15) in terms of the fluctuation $\delta\phi$ in Eq. (5.4), obtaining an equation that features the pinning term analogous to Eq. (5.9):

$$\partial_t\delta\phi = -r^{\mathrm{glob}}(t)\,\delta\phi(t) + \sqrt{\frac{2T}{N}}\eta(t) + \mathcal{A}(t). \tag{5.17}$$

This coefficient $r^{\mathrm{glob}}(t)$, determined by the harmonic coefficients $S_\alpha^{(n),\mathrm{glob}}(t)$ and $C_\alpha^{(n),\mathrm{glob}}(t)$ appearing in Eq. (5.15), plays the role of the global pinning coefficient and sets the strength of the restoring force acting on the mean orientation. In particular, in the stationary state, its time-dependence is due to the fluctuations of these coefficients as functions of time. In the thermodynamic limit, the amplitude of these fluctuations actually vanishes, as they refer to a global quantity, and thus $r^{\mathrm{glob}}(t)$ becomes time-independent and positive, i.e.,

$$r^{\mathrm{glob}}(t) = r^{\mathrm{glob}} > 0. \tag{5.18}$$

This latter fact ensures the stability of the pinned state against perturbations. Note that, by construction, $r^{\mathrm{glob}}(t)$ is a global quantity and thus it lacks any dependence on space which, instead, is pronounced in its "local counterpart" $r_i^{\mathrm{loc}}(t)$. In Eq. (5.17), the quantity $\mathcal{A}(t)$ (which results from a spatial average) collects all the contributions arising from the original equation of motion in Eq. (5.15), which are not captured by the linear global pinning term $-r^{\mathrm{glob}}(t)\delta\phi$: it has zero mean and its fluctuations vanish in the thermodynamics limit.

Remarks on the pinning terms. A key observation is that, in general,

$$r^{\mathrm{loc}} \neq r^{\mathrm{glob}}. \tag{5.19}$$

This means that the effective pinning coefficient that governs local fluctuations $\delta\phi_i$ does not necessarily coincide with the one that controls the fluctuations $\delta\phi$ of the global orientation, even after averaging over time.

As an illustrative example, consider first a case in which $r_i^{\mathrm{loc}} = r^{\mathrm{glob}}$, i.e., the reciprocal XY model in the presence of an external field which selects a preferred spin direction Φ. The corresponding dynamics is given by

$$\partial_t\phi_i = -J\sum_{j\in\mathcal{N}_i}\sin(\phi_i-\phi_j) - h\sin\left(\phi_i-\Phi\right). \tag{5.20}$$

Expanding this equation in terms of the local fluctuations $\delta\phi_i$ around the preferred direction, one finds that the local evolution equation for $\delta\phi_i$ contains a linear term $-h\,\delta\phi_i$ which provides a

restoring force toward $\delta\phi_i = 0$ and corresponds to $r_i^{\mathrm{loc}} = h$. Instead, summing Eq. (5.20) over all lattice sites, one obtains the following equation for the spatially averaged orientation $\phi(t)$:

$$\partial_t \phi(t) = -\frac{h}{N} \sum_i \sin(\phi_i - \Phi), \tag{5.21}$$

where the first term in Eq. (5.20) vanishes because the possible contribution of each ordered pair (i, j) is canceled by that of (j, i). Expanding this equation in terms of the global fluctuation $\delta\phi$ of $\phi(t)$ around the preferred direction Φ one finds $-h\,\delta\phi$, corresponding to $r^{\mathrm{glob}} = h$. Consequently, for this reciprocal XY model one obtains

$$r_i^{\mathrm{loc}} = r^{\mathrm{loc}} = r^{\mathrm{glob}} = h. \tag{5.22}$$

This equality no longer holds in the non-reciprocal XY model, where both local and global effective pinning terms depend explicitly on the instantaneous spin configuration.

The emergence of effective pinning terms in non-reciprocal XY models has been previously investigated in Refs. [92, 93]. In particular, Ref. [93] focuses on *local* equations of motion for $\delta\phi_i(t)$ in the presence of the reactive term, considering both sinusoidal and von Mises interaction kernels. Reference [92], instead, investigates the *global* dynamics of the system by deriving an effective equation for the collective variable $\delta\phi(t)$, again in the presence of the reactive term, for both sinusoidal and modulated coupling kernels. In the following, we extend these analyses by applying our framework to the modulated, sinusoidal, and von Mises coupling kernels, with particular emphasis on the role of the proactive term. We report the main findings of this work in the following three subsections, while the details of the numerical simulations are presented in Sec. 4.3.

5.2 Kernel-dependent pinning mechanisms

5.2.1 Pinning terms for the modulated coupling

The case of the modulated coupling was previously studied in Ref. [92], where only the reactive contribution to the dynamics was considered, with focus on the evolution equation for the global fluctuations $\delta\phi(t)$. Here, in order to determine the pinning terms, we expand Eq. (5.7) for the local orientation ϕ_i and Eq. (5.15) for the mean orientation $\phi(t)$ as detailed in Apps. B.2.1 and B.3.1, respectively. In the following, we discuss the resulting equations and their physical implications.

In the purely reactive case, the resulting time evolution equations at the local level *do not* exhibit any explicit pinning term. This can be seen from the expansion of $\partial_t \phi_i(t)$ reported in Eq. (B.5): when only the reactive contribution is retained, all coefficients have a vanishing mean value, resulting in a modulated diffusion term but no local restoring force. At first sight, this conclusion appears to contradict the argument presented in Sec. 4.2, based on which pinning is expected at the minima of $K_i(\phi)$, located at $\phi = \pi/4 + n\pi/2$, with $n \in \mathbb{Z}$. Moreover, the numerical simulations discussed in Sec. 4.3 clearly show pinning along the diagonal directions for the purely reactive dynamics. Our interpretation of this apparent discrepancy is the following. As discussed in Sec. 4.2 and in Ref. [92], the reactive dynamics tends to minimize the total coupling $K_i(\phi)$. This happens because the total coupling with the neighbors of a fluctuating spin which departs from the global alignment is stronger than the total coupling of its neighbors with it. As a consequence,

the fluctuating spin is pulled back toward the global alignment. However, this imbalance cannot be fully captured by the dynamic equations for ϕ_i alone: due to non-reciprocity, one must account not only for the bonds activated by ϕ_i, but also for those activated by its neighboring spins. For this reason, pinning does not emerge explicitly at the level of the local equations for ϕ_i, but it appears in the evolution equation for the fluctuations of the global orientation $\delta\phi(t)$. In particular, Ref. [92] expresses the quantity r^{glob} in terms of the fluctuations $\delta\tilde{\phi}_{\mathbf{x}}$ on the continuum as

$$-(16\epsilon/L^2)\int \mathrm{d}\mathbf{x}\,|\nabla\delta\tilde{\phi}_{\mathbf{x}}|^2, \tag{5.23}$$

which can then be estimated analytically. In our approach, we derive the corresponding expression on the lattice, i.e.,

$$r^{\text{glob}} = 4S_{\text{R}}^{(4),\text{glob}} \tag{5.24}$$

which can be evaluated numerically on the basis of the analytic expression in Eq. (B.41). Expanding this expression for small fluctuations and performing a gradient expansion (see App. B.3.1), one recovers Eq. (5.23). Note that, in the absence of an explicit local pinning term, this global pinning coefficient is a quantity that can be related to the pinning extracted from the static structure factor, as discussed in Sec. 5.3.2.

By contrast, when the proactive term is included, explicit pinning terms emerge in the evolution equations for both the local and the global orientations, although with different expressions (see Apps. B.2.1 and B.3.1 for details). The resulting complete dynamics eventually leads to alignment along the lattice directions.

5.2.2 Pinning terms for the sinusoidal coupling

The evolution equation for the local orientation ϕ_i in the presence of a sinusoidal kernel can be expanded as shown in App. B.2.2, while its continuum counterpart reproduces Eq. (5) of Ref. [93]. The corresponding equations for the global orientation $\delta\phi(t)$ are derived in App. B.3.2; upon expansion for small $\delta\tilde{\phi}_i$, they acquire the same structure as that of the equations reported in Ref. [92].

Since the system exhibits pinning along one of the four lattice directions, as confirmed by the numerical results presented in Sec. 4.3, we can set $\Phi = 0$ without loss of generality. We find that the evolution equations for both the local fluctuations $\delta\phi_i(t)$ and the global fluctuation $\delta\phi(t)$ contain effective pinning terms with the coefficients $r_i^{\text{loc}}(t)$ and r^{glob}, respectively. The reactive contribution of the sinusoidal coupling kernel, yields

$$r_i^{\text{loc}}(t) = -S_{\text{R},i}^{(1),\text{loc}}(t) \quad \text{and} \quad r^{\text{glob}} = -S_{\text{R}}^{(1),\text{glob}}, \tag{5.25}$$

where the explicit analytic expressions for $S_{\text{R},i}^{(1),\text{loc}}$ and $S_{\text{R}}^{(1),\text{glob}}$ are provided in Eqs. (B.14) and (B.49), respectively. This allows us to compare the average of the local pinning strength r_i^{loc}, with the global coefficient r^{glob}. For $\epsilon = 1$, $T = 0.2$, and $L = 500$, we obtain numerically

$$r^{\text{loc}} = 0.0015 \quad \text{with} \quad \Delta_{r^{\text{loc}}} = 0.4 \quad \text{and} \quad r^{\text{glob}} = 0.0022. \tag{5.26}$$

The positive value of r^{glob} indicates the presence of a global restoring force acting on the fluctuations

of the mean orientation $\delta\phi(t)$. Moreover, the values of the global and local pinning coefficients are quantitatively different, i.e., $r^{\mathrm{loc}} \neq r^{\mathrm{glob}}$.

An even more important observation concerns the fluctuations of the local pinning coefficient. Although the average value r^{loc} is positive, the fluctuations of r_i^{loc} around this mean, quantified by $\Delta_{r^{\mathrm{loc}}}$, are much larger than the mean itself (see Eq. (5.26)). Accordingly, the condition in Eq. (5.14) is strongly violated. In this regime, it is not obvious that the local dynamics can be described in terms of a single effective pinning obtained by renormalizing the average pinning coefficient r^{loc}. As discussed in Sec. 5.3.2, this makes any quantitative comparison between analytical pinning coefficients and those extracted from the stationary fluctuations considerably more subtle. Note that the numerical estimates reported in Eq. (5.26) were obtained for a specific set of simulation parameters. Nevertheless, the qualitative features $r^{\mathrm{loc}} \neq r^{\mathrm{glob}}$ and that $\Delta_{r^{\mathrm{loc}}} \gg r^{\mathrm{loc}}$ are expected to hold throughout the region of parameter space in which the (quasi)stationary LRO phase exists, as we verified numerically for various other choices.

As far as the proactive term is concerned, its inclusion in the dynamics also results in pinning along the lattice , as it is seen numerically in Sec. 4.3. In this case, however, the proactive contribution does not generate any local pinning term (see App. B.2.2); rather, it produces only a global pinning, which can be shown analytically to coincide with the one arising from the reactive contribution discussed above (see App. B.3.2).

5.2.3 Pinning terms for the von Mises coupling

The von Mises kernel (see Eq. (4.8)) represents a particularly instructive case. As discussed in Sec. 4.2, the mean-field analysis already reveals an explicit preference for alignment along specific directions: when both reactive and proactive terms are included in the Langevin dynamics, the system aligns along the lattice directions, whereas retaining only the reactive term favors alignment along the diagonal directions. This behavior is similar to that observed for the modulated coupling in Sec. 5.2.1 above, but differs from the case of the sinusoidal kernel in Sec. 5.2.2, for which no directional selection emerges at the mean-field level. At the same time, the von Mises kernel differs fundamentally from the modulated coupling because it intrinsically couples the spin degrees of freedom to the underlying lattice through its explicit dependence on the angles ϑ_{ij}. In the following, we show how these features manifest themselves at the level of the equations of motion, comparing them with what is reported in the literature.

A convenient starting point is to expand the von Mises kernel in Eq. (4.8) in powers of σ up to the fourth order, i.e.,

$$e^{\sigma\mathcal{C}} = 1 + \sigma\mathcal{C} + \frac{\sigma^2}{2}\mathcal{C}^2 + \frac{\sigma^3}{6}\mathcal{C}^3 + \frac{\sigma^4}{24}\mathcal{C}^4, \tag{5.27}$$

where $\mathcal{C} = \cos(\phi_i - \vartheta_{ij})$. We first focus on the purely reactive case, i.e., on the Langevin dynamics which includes only the reactive term. From the mean-field argument presented in Sec. 4.2 and as confirmed by the numerical results presented in Sec. 4.3, we know that the system features pinning along the diagonal directions, e.g., $\Phi = \pi/4$. We therefore seek to identify which term in the expansion in Eq. (5.27) is responsible for this behavior. The term of order σ^0 generates the reciprocal interaction of the usual XY model and thus it cannot explain the emergence of pinning. The contribution of order σ to $\partial_t\phi_i$ turns out to coincide, up to an inconsequential prefactor (see App. B.2.3), with the one due to a sinusoidal coupling and thus, as discussed in Sec. 5.2.2, it cannot by itself lead to diagonal pinning. The contribution of order σ^2 was analyzed in Ref. [93], where it

was shown that it corresponds to a term of the form

$$\langle \sin\phi_{\mathbf{x}} \, (\partial_x^2 \phi_{\mathbf{x}} - \partial_y^2 \phi_{\mathbf{x}}) \rangle \sin\phi_{\mathbf{x}} \tag{5.28}$$

when the system is described in terms of the variable $\phi_{\mathbf{x}}$ on the continuum. It was suggested there that this term might be responsible for pinning. However, since it is proportional to $\sin\phi_{\mathbf{x}}$, it would favor pinning at $\Phi = 0$, in contrast with the diagonal pinning observed here. Even including the contribution of order σ^3 of the von Mises kernel does not explain this observation, because the $\mathcal{O}(\sigma^3)$ contribution to the total coupling $K_i(\phi)$ of Eq. (4.10) vanishes for a globally aligned configuration and therefore it cannot induce any pinning at the mean-field level. This conclusion is further supported by our numerical simulations: truncating the kernel at order σ^3 results into a pinning along the lattice directions.

In fact, we identify the term of order σ^4 in Eq. (5.27) as the leading contribution responsible for diagonal pinning. Within the mean-field argument discussed in Sec. 4.2, this is indeed the first term that generates a total coupling $K_i(\phi)$ with the $\mathbb{Z}_4$ symmetry of the square lattice in a globally aligned state at angle ϕ. More specifically, $K_i(\phi)$ acquires a contribution of the form

$$\frac{\sigma^4}{12}(\cos^4\phi + \sin^4\phi). \tag{5.29}$$

These observations suggest that the pinning phenomenology of the von Mises coupling can be conveniently captured by a simplified model that retains only the contributions of order σ^0, σ, and σ^4 from the expansion in Eq. (5.27). The zeroth-order term is included to recover the reciprocal XY model as $\sigma \to 0$. The term of order σ is also retained because, as shown below, it generates a local pinning contribution. Interestingly, unlike the purely sinusoidal coupling, this contribution here favors the diagonal rather than the lattice directions. Within this truncation scheme, the model effectively reduces to the combination of a sinusoidal coupling and a modulated-coupling term.

The corresponding equations for the local orientation ϕ_i and for the mean orientation ϕ, obtained within this truncation, are derived in Apps. B.2.3 and B.3.3, respectively. In particular, the equation for ϕ_i, written in the form of Eq. (5.7), contains two leading harmonic contributions $S_{\mathrm{R}}^{(1),\mathrm{loc}}$ and $C_{\mathrm{R}}^{(1),\mathrm{loc}}$ (see Eqs. (B.25) and (B.26), respectively) which, by symmetry, satisfy $S_{\mathrm{R}}^{(1),\mathrm{loc}} = -C_{\mathrm{R}}^{(1),\mathrm{loc}}$ (see App. B.2.3), Moreover, numerical measurements show that $S_{\mathrm{R}}^{(1),\mathrm{loc}} < 0$. The corresponding contribution to $\partial_t \phi_i$ can therefore be rewritten as $\sqrt{2}\, S_{\mathrm{R}}^{(1),\mathrm{loc}} \sin(\phi_i - \pi/4)$, from which the pinning direction $\phi_i = \Phi = \pi/4$ naturally emerges. Upon linearization around $\Phi = \pi/4$, this term generates a local pinning term for $\delta\phi_i$ with coefficient

$$r_i^{\mathrm{loc}}(t) = -\sqrt{2} S_{\mathrm{R}}^{(1),\mathrm{loc}} > 0. \tag{5.30}$$

Similarly, the global pinning term governing the evolution of $\delta\phi$ takes the form of Eq. (5.17), with (see App. B.3.3)

$$r^{\mathrm{glob}} = -\sqrt{2} S_{\mathrm{R}}^{(1),\mathrm{glob}} + 4 S_{\mathrm{R}}^{(4),\mathrm{glob}}, \tag{5.31}$$

where the analytic expressions of the harmonic coefficients $S_{\mathrm{R}}^{(1),\mathrm{glob}}$ and $S_{\mathrm{R}}^{(4),\mathrm{glob}}$ are provided in Eqs. (B.60) and (B.64), respectively. The numerical estimates of these coefficients, reported in Sec. 5.2.5, yield $r^{\mathrm{glob}} > 0$, confirming the stability of the pinning.

It is worth emphasizing that the reactive coefficients $S_{\mathrm{R}}^{(1),\mathrm{loc}}$, $C_{\mathrm{R}}^{(1),\mathrm{loc}}$, $S_{\mathrm{R}}^{(1),\mathrm{glob}}$, and $C_{\mathrm{R}}^{(1),\mathrm{glob}}$ for the sinusoidal coupling have the same analytical expressions as for the von Mises kernel, but

with different prefactors. In particular, they are proportional to ϵ in the former case and to σ in the latter: compare, e.g., the expressions of $C_R^{(1),\mathrm{glob}}$ in Eqs. (B.50) and (B.61). (The same applies to the corresponding proactive coefficients, but we consider here the case of reactive dynamics.) Nevertheless, their physical effect is qualitatively different in the two models, which indeed feature different pinning directions. This is due to the fact that coefficients with the same analytical expression actually take different values depending on the pinning direction Φ selected by the dynamics.

In fact, for the sinusoidal kernel, pinning occurs always along the lattice directions, e.g., $\Phi = 0$. For this specific pinning, the symmetry argument discussed in Apps. B.2.2 and B.3.2 implies that $C_{\dots}^{(1),\cdots} = 0$, while the non-vanishing coefficients $S_{\dots}^{(1),\cdots}$ eventually yield pinning. By contrast, for the von Mises kernel with reactive dynamics, pinning occurs along the diagonal directions, e.g., $\Phi = \pi/4$. The relevant symmetry is then the diagonal reflection discussed in App. B.3.3, which implies $S_{\dots}^{(1),\cdots} = -C_{\dots}^{(1),\cdots}$. However, these first-harmonic terms alone are not sufficient to explain the observed diagonal pinning in the (reactive) von Mises case. Indeed, if only these $\mathcal{O}(\sigma)$ contributions were retained, the resulting dynamics would coincide with that of the sinusoidal kernel, thereby leading to pinning along the lattice directions. This means that the selection of pinning along the diagonal directions is actually driven by the fourth harmonic term, that however does not appear in the local pinning, but only in the global one with $S_{\mathrm{R}}^{(4),\mathrm{glob}}$ of $\mathcal{O}(\sigma^4)$, exactly as in the modulated coupling discussed in Sec. 5.2.1 above.

Summarizing, the fourth-harmonic contribution to the von Mises coupling is ultimately responsible for selecting the diagonal pinning directions. However, this contribution does not appear directly in the dynamics of individual spins. This local pinning, instead, originates from the $\mathcal{O}(\sigma)$ contribution, the analytical form of which is the same as in the sinusoidal kernel. The role of the fourth-harmonic term is to modify the stationary state, thereby altering the actual numerical values of the $\mathcal{O}(\sigma)$ coefficients, making them consistent with pinning along the diagonal directions.

In this context, a natural question is whether, for sufficiently small σ, the pinning direction eventually reverts to the lattice axes, or whether an arbitrarily small fourth-harmonic terms is already sufficient to stabilize the diagonal directions. Since the pinning effects become very weak in this regime, our numerical results are currently inconclusive on this point for what concerns the square lattice (data not shown here), but they are able to detect a crossover in the pinning direction as a function of σ on the triangular lattice (see App. B.1). Similarly, the effects on the transient dynamics of adding the terms of order σ^2 and σ^3 remains to be investigated numerically.

Finally, when the proactive term is included in the dynamics, the local pinning receives a contribution only from the $\mathcal{O}(\sigma^4)$ terms (see App. B.2.3), whereas the global pinning receives contributions from both the $\mathcal{O}(\sigma)$ and $\mathcal{O}(\sigma^4)$ terms (see App. B.3.3). In this case, the pinning occurs along the lattice directions, consistently with the mean-field argument presented in Sec. 4.2 and with the numerical results reported in Sec. 4.3.

5.2.4 Proactivity and pinning: A summary

Before illustrating, in the next Section, the results of our numerical simulations, we summarize the content of this Sec. 5.1. We have systematically investigated here the emergence of local and global pinning terms in the equations governing the dynamics of the individual spins and of the mean orientation for all the interaction kernels considered in this work.

A first important result is that pinning does not necessarily appear in the local equations of

Table 5.1: Summary of the presence of a non-vanishing local pinning coefficients r_i^{loc} in the effective equation for ϕ_i, coming from the reactive and proactive part of Eq. (5.7) for the various interaction kernels considered in this work.

Kernel type	**Reactive**	**Proactive**
Modulated	no	yes
Sinusoidal	yes	no
von Mises	yes at $\mathcal{O}(\sigma)$	yes at $\mathcal{O}(\sigma^4)$

motion; accordingly, it might be deceptive to focus exclusively on them. Moreover, the possible emergence of pinning and its eventual direction depend on the specific implementation of the dynamics (i.e., Langevin equation with the sole reactive term or with both reactive and proactive terms, this latter case being equivalent to Monte Carlo dynamics). In particular, for the modulated coupling, the reactive term does not generate any local pinning, whereas the proactive term does. For the sinusoidal kernel, the situation is reversed: the reactive term generates a local pinning contribution, while the proactive term does not. For the von Mises kernel (truncated such as to include terms of order σ^0, σ, and σ^4), instead, the reactive contribution to the local pinning originates from the term at $\mathcal{O}(\sigma)$, corresponding to the sinusoidal component of the interaction, whereas the proactive contribution to pinning appears only at $\mathcal{O}(\sigma^4)$, corresponding to the modulated kernel component. These results are summarized in Tab. 5.1.

By contrast, *all* the dynamical contributions considered here generate a global pinning term with coefficient r^{glob}, acting on the mean orientation $\phi(t)$. A more detailed analysis of this pinning mechanism is presented in Sec. 5.2.5, together with a discussion of its relation with the properties of the system in the stationary state.

5.2.5 Global pinning in the stationary state

In this section, we focus on the behavior of the system in the (quasi)stationary state. In particular, we verify that the pinning terms predicted from the analysis of the evolution equations for the global orientation $\phi(t)$ are consistent with the orientation effectively selected by the system and detected numerically.

After the system has reached the stationary state, we compute the harmonic coefficients $S_\alpha^{(n),\mathrm{glob}}(t)$ and $C_\alpha^{(n),\mathrm{glob}}(t)$ appearing in Eq. (5.15). This is done by numerically averaging over the fluctuations the corresponding expressions derived in the Apps. B.3.1, B.3.2, and B.3.3. for the different interaction kernels. We then perform additional time averages in order to estimate their mean values $\overline{S_\alpha^{(n),\mathrm{glob}}}$ and $\overline{C_\alpha^{(n),\mathrm{glob}}}$ together with the associated standard deviations. Since these coefficients are global quantities, their fluctuations are expected to vanish in the thermodynamic limit $L \to \infty$.

The resulting averages and standard deviations, along with the values of the parameters used in the simulations, are reported in Table 5.2. In the following, we discuss the implications of these numerical results for each interaction kernel separately. For notational simplicity, we henceforth omit the superscript “glob" and the indication of the time average from the harmonic coefficients.

- **Modulated coupling.** The analytical derivation of the dynamics of the mean orientation $\phi(t)$

Table 5.2: Numerical estimates of the harmonic coefficients $S_{\mathrm{R,P}}^{(n),\mathrm{glob}}$ and $C_{\mathrm{R,P}}^{(n),\mathrm{glob}}$ appearing in the equation of the mean orientation ϕ (see Eq. (5.15)) for the different interaction kernels and dynamical implementations. To simplify the notation, the superscript "glob" is omitted in the table. The reported values are time averages obtained for finite system sizes; uncertainties correspond to the associated standard deviations and are expected to vanish in the thermodynamic limit. Entries denoted by "–" indicate coefficients that are absent by construction. All the simulations are performed with $L = 200$ and $T = 0.2$. While the precise numerical values of the various coefficients depend on the chosen parameters, the eventual stability of the pinned state is actually determined only by the sign of their combinations which enter the global pinning coefficient r^{glob}. These signs are found to be robust upon changing the parameters as long as they give rise to a (quasi)stationary state with pinning.

Model	$S_{\mathrm{R}}^{(1)}$	$C_{\mathrm{R}}^{(1)}$	$S_{\mathrm{P}}^{(1)}$	$C_{\mathrm{P}}^{(1)}$	$S_{\mathrm{R}}^{(4)}$	$C_{\mathrm{R}}^{(4)}$	$S_{\mathrm{P}}^{(4)}$	$C_{\mathrm{P}}^{(4)}$
Modulated – reactive, $\epsilon = 0.5$	–	–	–	–	0.257(1)	0.0003(10)	–	–
Modulated – total, $\epsilon = 0.05$	–	–	–	–	0.0196(1)	0.00001(200)	-0.548(2)	0.0001(10)
Sinusoidal – reactive, $\epsilon = 1.5$	-0.0057(3)	-0.0001(3)	–	–	–	–	–	–
Sinusoidal – total , $\epsilon = 1.0$	-0.0065(3)	0.0001(2)	-0.0065(3)	0.0001(2)	–	–	–	–
von Mises – reactive , $\sigma = 1.5$	-0.0015(2)	0.0015(2)	–	–	0.00229(4)	0.000004(20)	–	–
von Mises – total , $\sigma = 0.7$	-0.0017(2)	0.00001(10)	-0.0017(2)	0.00001(10)	0.00120(3)	0.000001(10)	-0.0234(6)	-0.00006(20)

under the complete Langevin evolution is provided in App. B.3.1 and yields the equation:

$$\begin{aligned}\partial_t \phi = (S_{\mathrm{R}}^{(4)} + S_{\mathrm{P}}^{(4)}) \sin(4\phi) + \\ + (C_{\mathrm{R}}^{(4)} + C_{\mathrm{P}}^{(4)}) \cos(4\phi) + \sqrt{\frac{2T}{N}}\eta(t),\end{aligned} \tag{5.32}$$

where the analytic expressions of the harmonic coefficients $S_{\mathrm{R}}^{(4)}$, $S_{\mathrm{P}}^{(4)}$, $C_{\mathrm{R}}^{(4)}$, and $C_{\mathrm{P}}^{(4)}$ are provided by Eqs. (B.41), (B.43), (B.40), and (B.42), respectively. As implied by the symmetry arguments presented in App. B.3.1, and confirmed by the numerical averages reported in Table 5.2, the coefficients $C_{\mathrm{R}}^{(4)}$ and $C_{\mathrm{P}}^{(4)}$ vanish within the numerical precision. This holds both for the complete dynamics and for the one involving only the reactive term, in which case $C_{\mathrm{P}}^{(4)}$ is absent by construction.

When the dynamics is governed solely by the reactive term, Eq. (5.32) reduces to

$$\partial_t \phi = S_{\mathrm{R}}^{(4)} \sin(4\phi) + \sqrt{\frac{2T}{N}}\eta(t). \tag{5.33}$$

In our simulations, we let the system relax and reach its stationary state, which, as shown in Fig. 4.5(a), is pinned along one of the diagonal directions, $\Phi = \pi/4$, consistently with the mean-field arguments discussed in Sec. 4.2. Evaluating the coefficient $S_{\mathrm{R}}^{(4)}$ in this stationary state, we find $S_{\mathrm{R}}^{(4)} > 0$, as reported in Tab. 5.2.

Expanding around the pinned direction as $\phi = \Phi + \delta\phi$, with $\Phi = \pi/4$, one obtains, at linear order, an equation of the form of Eq. (5.17), where the relaxation rate is given by

$$r^{\mathrm{glob}} = 4S_{\mathrm{R}}^{(4)}. \tag{5.34}$$

Since $r^{\mathrm{glob}} > 0$, the stationary distribution of $\delta\phi$ is centered around zero with finite variance. This confirms that the pinned direction $\Phi = \pi/4$ is dynamically stable.

By contrast, in the presence of the additional proactive term, Eq. (5.32) becomes

$$\partial_t \phi = \left(S_{\mathrm{R}}^{(4)} + S_{\mathrm{P}}^{(4)} \right) \sin(4\phi) + \sqrt{\frac{2T}{N}} \eta(t). \tag{5.35}$$

In our simulations, the stationary state is found to be pinned along one of the lattice directions, e.g., $\Phi = 0$ (see Fig. 4.5(a)), consistently with the mean-field arguments of Sec. 4.2. In this state, we find numerically $S_{\mathrm{R}}^{(4)} + S_{\mathrm{P}}^{(4)} < 0$, as reported in Tab. 5.2.

Expanding around the pinned direction as $\phi = \Phi + \delta\phi$, with $\Phi = 0$, one obtains again Eq. (5.17) with relaxation rate

$$r^{\mathrm{glob}} = -4 \left(S_{\mathrm{R}}^{(4)} + S_{\mathrm{P}}^{(4)} \right). \tag{5.36}$$

Since $r^{\mathrm{glob}} > 0$, the pinned direction $\Phi = 0$ is indeed dynamically stable.

These stability arguments, established for the pinned directions $\Phi = \pi/4$ in the purely reactive dynamics and $\Phi = 0$ in the full dynamics, extend straightforwardly to all symmetry-equivalent diagonal and lattice directions.

- **Sinusoidal coupling.** The derivation of the evolution equation for the mean orientation is presented in App. B.3.2 and leads to

$$\begin{aligned} \partial_t \phi =& (S_{\mathrm{R}}^{(1)} + S_{\mathrm{P}}^{(1)}) \sin\phi \\ &+ (C_{\mathrm{R}}^{(1)} + C_{\mathrm{P}}^{(1)}) \cos\phi + \sqrt{\frac{2T}{N}} \eta(t), \end{aligned} \tag{5.37}$$

where the analytic expressions for the harmonic coefficients $S_{\mathrm{R}}^{(1)}$, $S_{\mathrm{P}}^{(1)}$, $C_{\mathrm{R}}^{(1)}$, and $C_{\mathrm{P}}^{(1)}$ are provided in Eqs. (B.49), (B.54), (B.50), and (B.55), respectively. However, under periodic boundary conditions, these coefficients are not independent but they satisfy the relations $S_{\mathrm{R}}^{(1)} = S_{\mathrm{P}}^{(1)}$ and $C_{\mathrm{R}}^{(1)} = C_{\mathrm{P}}^{(1)}$, as shown in Sec. B.3.2. In our simulations, the stationary state is found to be pinned along one of the lattice directions, e.g., $\Phi = 0$, for both the purely reactive and the complete dynamics (see Fig. 4.5(b)). Unlike the case of the modulated coupling, here we do not have a simple mean-field argument predicting the selected direction *a priori*.

Evaluating the coefficients in the stationary state, we find $S_{\mathrm{R}}^{(1)} < 0$ and $C_{\mathrm{R}}^{(1)} \simeq 0$ in the purely reactive case, while for the full dynamics we obtain $S_{\mathrm{R}}^{(1)} = S_{\mathrm{P}}^{(1)} < 0$ and $C_{\mathrm{R}}^{(1)} = C_{\mathrm{P}}^{(1)} \simeq 0$ (see Tab. 5.2; the vanishing of the coefficients $C_{\mathrm{R}}^{(1)}$ and $C_{\mathrm{P}}^{(1)}$ is theoretically expected by the symmetry argument reported in App. B.3.2). Expanding around the pinned direction as $\phi = \Phi + \delta\phi$, with $\Phi = 0$, one obtains Eq. (5.17) with relaxation rates $r^{\mathrm{glob}} = -S_{\mathrm{R}}^{(1)} > 0$ for the purely reactive dynamics, and $r^{\mathrm{glob}} = -\left(S_{\mathrm{R}}^{(1)} + S_{\mathrm{P}}^{(1)} \right) = -2S_{\mathrm{R}}^{(1)} > 0$ for the full dynamics. In both cases, we find that $r^{\mathrm{glob}} > 0$, confirming that the lattice direction $\Phi = 0$ is dynamically stable.

In App. B.3.2, we further show that Eq. (5.37) preserves the underlying $\mathbb{Z}_4$ symmetry of the lattice. The stability arguments presented above therefore straightforwardly extend to all equivalent lattice directions $\Phi = n\pi/2$, with $n \in \mathbb{Z}$.

- **von Mises coupling.** In this case, we simulated the dynamics using a simplified version of the von Mises kernel, obtained by truncating its expansion in σ and keeping only the terms

of order σ^0, σ^1, and σ^4 (see Sec. 5.2.3). The resulting dynamics is qualitatively equivalent to that generated by the complete von Mises kernel, while allowing for a considerably simpler analytical treatment.

The analytical derivation of the dynamics of the mean orientation is provided in App. B.3.3 and yields

$$\begin{aligned}\partial_t\phi = &\left(S_{\mathrm{R}}^{(1)} + S_{\mathrm{P}}^{(1)}\right)\sin\phi + \left(C_{\mathrm{R}}^{(1)} + C_{\mathrm{P}}^{(1)}\right)\cos\phi \\ &+ \left(S_{\mathrm{R}}^{(4)} + S_{\mathrm{P}}^{(4)}\right)\sin(4\phi) + \left(C_{\mathrm{R}}^{(4)} + C_{\mathrm{P}}^{(4)}\right)\cos(4\phi) + \sqrt{\frac{2T}{N}}\,\eta(t).\end{aligned} \tag{5.38}$$

where the analytic expressions of the first harmonic coefficients $S_{\mathrm{R}}^{(1)}$, $S_{\mathrm{P}}^{(1)}$, $C_{\mathrm{R}}^{(1)}$, and $C_{\mathrm{P}}^{(1)}$ are provided by Eqs. (B.60), (B.68), (B.61), and (B.69), respectively; those of the forth harmonic $S_{\mathrm{R}}^{(4)}$, $S_{\mathrm{P}}^{(4)}$, $C_{\mathrm{R}}^{(4)}$, and $C_{\mathrm{P}}^{(4)}$ are provided by Eqs. (B.64), (B.72), (B.65), and (B.73), respectively. In Eq. (5.38) we do not report the terms proportional to the second harmonics $\sin(2\phi)$ and $\cos(2\phi)$, which are discussed in detail in App. B.3.3. These terms are also of order $\mathcal{O}(\sigma^4)$, like the fourth-harmonic contributions. However, we find numerically that their contribution to the pinning is significantly smaller than that of the corresponding fourth-harmonic terms and therefore they do not affect the conclusions drawn here.

Symmetry arguments analogous to those used for the modulated coupling and discussed in App. B.3.3 imply that $C_{\mathrm{R}}^{(4)} = C_{\mathrm{P}}^{(4)} = 0$. In addition, by the same reasoning employed for the sinusoidal coupling, one finds that, under periodic boundary conditions and in the presence of both reactive and proactive contributions, $S_{\mathrm{R}}^{(1)} = S_{\mathrm{P}}^{(1)}$ and $C_{\mathrm{R}}^{(1)} = C_{\mathrm{P}}^{(1)}$. All these relations are confirmed by the numerical values reported in Tab. 5.2.

In the purely reactive case, the stationary state observed in simulations is pinned along one of the diagonal directions, $\Phi = \pi/4$ (see Fig. 4.5(c)), consistently with the mean-field arguments discussed in Sec. 4.2. Evaluating the coefficients in this stationary state, we find $S_{\mathrm{R}}^{(1)} = -C_{\mathrm{R}}^{(1)} < 0$, $S_{\mathrm{R}}^{(4)} > 0$, as reported in Tab. 5.2. Note that the relation $S_{\mathrm{R}}^{(1)} = -C_{\mathrm{R}}^{(1)}$ is theoretically expected based on the symmetries of the problem, as discussed in App. B.3.3. The resulting equation for the mean orientation therefore reduces to

$$\partial_t\phi = S_{\mathrm{R}}^{(1)}\big(\sin\phi - \cos\phi\big) + S_{\mathrm{R}}^{(4)}\sin(4\phi) + \sqrt{\frac{2T}{N}}\eta(t). \tag{5.39}$$

Expanding around the pinned direction as $\phi = \pi/4 + \delta\phi$, one obtains Eq. (5.17) with

$$r^{\mathrm{glob}} = -\sqrt{2}\,S_{\mathrm{R}}^{(1)} + 4S_{\mathrm{R}}^{(4)}. \tag{5.40}$$

Since the numerical values of the coefficients in Tab. 5.2 gives $r^{\mathrm{glob}} > 0$, the diagonally pinned direction is dynamically stable.

By contrast, when the proactive contribution is included, the stationary state observed in our simulation becomes pinned along one of the lattice directions, $\Phi = 0$ (see Fig. 4.5(c)), again in agreement with the mean-field arguments. In this case, we find numerically that $S_{\mathrm{R}}^{(1)} = S_{\mathrm{P}}^{(1)} < 0$, $C_{\mathrm{R}}^{(1)} = C_{\mathrm{P}}^{(1)} \simeq 0$ and $S_{\mathrm{R}}^{(4)} + S_{\mathrm{P}}^{(4)} < 0$ (with $S_{\mathrm{R}}^{(4)}$ and $S_{\mathrm{P}}^{(4)}$ having opposite signs), as reported in Tab. 5.2. The vanishing of the coefficients $C_{\mathrm{R}}^{(1)}$ and $C_{\mathrm{P}}^{(1)}$ observed numerically is also theoretically expected from the same symmetry arguments as in the

sinusoidal case. Expanding around the pinned direction $\Phi = 0$ as $\phi = \delta\phi$, we obtain Eq. (5.17) with

$$r^{\text{glob}} = -2S_{\text{R}}^{(1)} - 4(S_{\text{R}}^{(4)} + S_{\text{P}}^{(4)}) \tag{5.41}$$

Since the numerical values of the coefficients yield $r^{\text{glob}} > 0$, the pinned direction is dynamically stable.

The arguments above straightforwardly extend to all equivalent diagonal and lattice directions selected by the $\mathbb{Z}_4$ symmetry of the underlying lattice.

- **Hard vision cone.** This case does not admit a straightforward analytical treatment due to the singularities introduced by the discontinuous interaction kernel. We therefore restrict the discussion to the numerical results which show how the selected orientation depends on the aperture angle θ and of the type of dynamics, as already discussed in Sec. 4.3.1.

5.3 Comparison with correlation time and pinning from the structure factor

In this section, we consider two alternative approaches to determine quantities which may provide estimates of the strengths of the local and global pinning terms discussed above. In particular, we focus on (i) the correlation time associated with the Ornstein–Uhlenbeck (OU) process approximately obeyed by the mean orientation $\phi(t)$ and (ii) the pinning term extracted from the structure factor of the fluctuations in the stationary state, which was considered in Refs. [92, 93].

5.3.1 Correlation time

Once the mean orientation $\phi(t)$ is pinned along the direction Φ, its fluctuations $\delta\phi(t)$ around this value (see Eq. (5.4)) evolve according to Eq. (5.17). In the stationary state, $r^{\text{glob}}(t)$ becomes time-independent as discussed after Eq. (5.17). While $\mathcal{A}(t)$ has zero mean, i.e., $\overline{\mathcal{A}} = 0$, in systems of finite size it exhibits fluctuations that vanish in the thermodynamic limit, with standard deviation scaling as $\sigma_{\mathcal{A}} \sim N^{-1/2}$, where N is the total number of spins in the system. Likewise, the fluctuations of the mean orientation satisfy $\overline{\delta\phi} = 0$ and $\sigma_{\delta\phi} \sim N^{-1/2}$. Since both $\delta\phi$ and $\mathcal{A}$ vanish as $N^{-1/2}$ in the thermodynamic limit $N \to \infty$, we expect Eq. (5.17) to reduce to the linear evolution equation of a standard OU process whenever the component of $\mathcal{A}(t)$ correlated with the fluctuations $\delta\phi(t)$ is small compared to the deterministic pinning term $r^{\text{glob}}\delta\phi(t)$. Indeed, any component of $\mathcal{A}(t)$ that is uncorrelated with $\delta\phi(t)$ acts effectively as an additional noise source: it modifies the amplitude of the stationary fluctuations, but not the relaxation rate extracted from the autocorrelation function. This leads to the condition for the validity of the OU description:

$$\sigma_{\mathcal{A}} \, |\text{Corr}(\mathcal{A}, \delta\phi)| \ll r^{\text{glob}} \, \sigma_{\delta\phi}, \tag{5.42}$$

where $\text{Corr}(\mathcal{A}, \delta\phi)$ denotes the normalized correlation coefficient between $\mathcal{A}(t)$ and $\delta\phi(t)$ in the stationary state. Under this condition, the contribution of $\mathcal{A}$ in Eq. (5.17) becomes negligible, and the dynamics of $\delta\phi$ is effectively described by an OU process characterized by the relaxation rate $r^{\text{OU}} \equiv r^{\text{glob}}$ and noise amplitude $\sqrt{2T/N}$. The autocorrelation function of this process is given by

$$C_\phi(t) = \langle \delta\phi(0)\, \delta\phi(t) \rangle \sim \exp\left(-r^{\text{OU}}|t|\right), \tag{5.43}$$

which decays exponentially with increasing time, independently of the noise amplitude. The relaxation rate r^{OU} can therefore be extracted numerically by fitting $C_\phi(t)$ with an exponential law. This allows a direct comparison between the measured relaxation rate r^{OU} and the pinning coefficient predicted by the analytical theory r^{glob} and calculated from the simulations. As we show below, the agreement between the two quantities depends crucially on the validity of the condition in Eq. (5.42). In the following, we discuss two representative cases: the modulated coupling, for which the condition in Eq. (5.42) is satisfied and the OU description is accurate, and the sinusoidal coupling, for which it is not.

For the modulated coupling (see Sec. 4.1.3) the analytical expression for the dynamics of the mean orientation $\phi(t)$ in the purely reactive case is obtained from Eq. (5.32) by setting the proactive coefficients to zero. Expanding around one of the diagonal pinning directions, e.g., $\Phi = \pi/4$, and comparing the resulting expression with Eq. (5.17), one readily identifies

$$r^{\mathrm{glob}} = 4S_{\mathrm{R}}^{(4)} \quad \text{and} \quad \mathcal{A} = -C_{\mathrm{R}}^{(4)} \cos\left(4\delta\phi\right). \tag{5.44}$$

Numerical measurements performed at $T = 0.1$, $\epsilon = 0.3$, and for a system size $L = 200$ yield $r^{\mathrm{glob}} \simeq 0.40$, while $C_{\mathrm{R}}^{(4)}$ is compatible with zero within statistical fluctuations, as expected from the symmetry argument discussed in App. B.3.1. Moreover, we find $\sigma_{\delta\phi} \simeq 3 \times 10^{-3}$, $\sigma_{\mathcal{A}} \simeq 4 \times 10^{-4}$, and $\mathrm{Corr}(\mathcal{A}, \delta\phi) \simeq -0.17$ [1]. Accordingly, the component of $\mathcal{A}$ that is correlated with the fluctuations of the mean orientation — given by $\sigma_{\mathcal{A}}|\mathrm{Corr}(\mathcal{A}, \delta\phi)|$, — is much smaller than the restoring contribution $r^{\mathrm{glob}}\sigma_{\delta\phi}$, and the condition in Eq. (5.42) is satisfied. Thus, the dynamics of $\delta\phi$ is accurately described by an OU process with relaxation time $1/r^{\mathrm{glob}}$, so that the autocorrelation time $1/r^{\mathrm{OU}}$ determined numerically should coincide with the analytical prediction $1/r^{\mathrm{glob}}$. We test this relationship in Fig. 5.1(a) (with $T = 0.1$, $L = 200$, $dt = 0.05$) where, as a function of the non-reciprocity parameter ϵ, we compare the values of r^{glob}, obtained by numerically by averaging the corresponding analytical expressions (see Eqs. (5.44) and (B.49)), with the values of r^{OU}, extracted by fitting the autocorrelation function of the mean orientation with an exponential decay. The two estimates are found to be in excellent agreement over the entire range of non-reciprocity parameters ϵ considered. In Fig. 5.1(b) we compare r^{glob} and r^{OU} after including the proactive term in the dynamics ($T = 0.3$, $L = 200$, $dt = 0.05$), finding an equally good agreement; This confirms that the OU description remains valid.

The situation changes qualitatively for the sinusoidal coupling with the sole reactive Langevin dynamics. In this case, the analytical expression for the dynamics of the mean orientation $\phi(t)$ is obtained from Eq. (5.37) by setting the proactive coefficients to zero. Expanding around one of the pinned lattice directions, e.g., $\Phi = 0$, and comparing the resulting expression with Eq. (5.17), one identifies

$$r^{\mathrm{glob}} = -S_{\mathrm{R}}^{(1)} \qquad \text{and} \qquad \mathcal{A} = C_{\mathrm{R}}^{(1)} \cos\left(\delta\phi\right). \tag{5.45}$$

For the parameters $T = 0.2$, $L = 200$, and $\epsilon = 2.0$, we find numerically $r^{\mathrm{glob}} \simeq 10^{-2}$, $C_{\mathrm{R}}^{(1)}$ compatible with zero, $\sigma_{\mathcal{A}} \simeq 10^{-3}$, $\sigma_{\delta\phi} \simeq 6 \times 10^{-2}$, and $\mathrm{Corr}(\mathcal{A}, \delta\phi) \simeq 0.73$. In this case one finds

$$\sigma_{\mathcal{A}}\,|\mathrm{Corr}(\mathcal{A}, \delta\phi)| \sim r^{\mathrm{glob}}\sigma_{\delta\phi}, \tag{5.46}$$

[1] The correlation coefficient $\mathrm{Corr}(\mathcal{A}, \delta\phi)$ is computed from the stationary time series of $\mathcal{A}(t)$ and $\delta\phi(t)$ as the normalized covariance, $\mathrm{Corr}(\mathcal{A}, \delta\phi) = \overline{\left(\mathcal{A}(t) - \overline{\mathcal{A}}\right)\left(\delta\phi(t) - \overline{\delta\phi}\right)}/\sigma_{\mathcal{A}}\sigma_{\delta\phi}$, where the overline denotes a time average in the stationary state. By construction, $-1 \leq \mathrm{Corr}(\mathcal{A}, \delta\phi) \leq 1$.

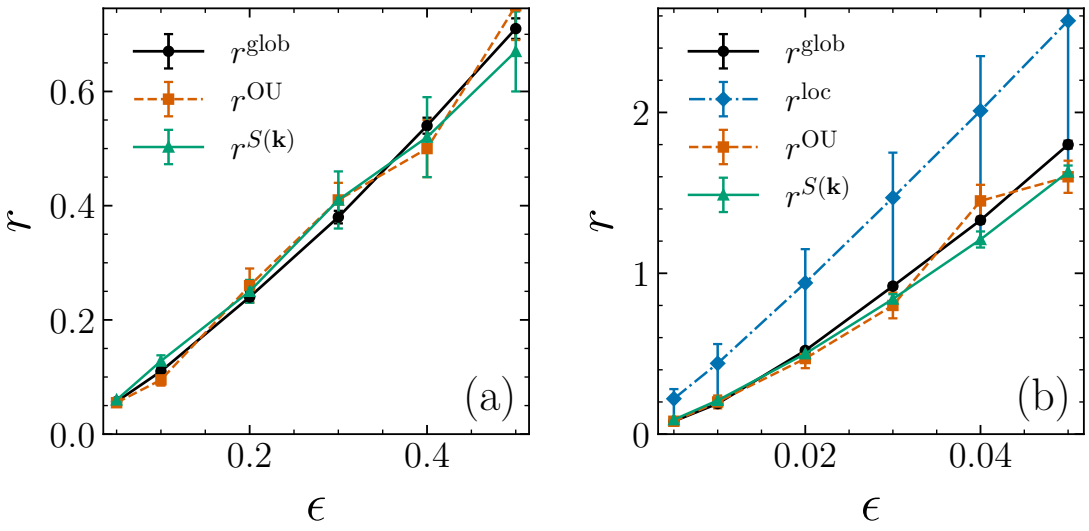


Figure 5.1: Pinning coefficients for the modulated-coupling kernel (see Sec. 4.1.3) obtained by using four different methods. Panel (a) corresponds to the Langevin dynamics with only the reactive term. In this case, the global pinning coefficient r^{glob}, the relaxation rate r^{OU} extracted from the OU autocorrelation time, and the pinning coefficient $r^{S(\mathbf{k})}$ obtained from the static structure factor are in excellent agreement. This indicates that, despite the absence of an explicit local pinning term in the continuum equations, the spatial fluctuations are accurately described by an effective pinning equal to r^{glob}. Panel (b) corresponds to the complete Langevin dynamics including both the reactive and proactive terms. Again, r^{glob}, r^{OU}, and $r^{S(\mathbf{k})}$ are found to be in very good agreement. By contrast, the mean local pinning coefficient r^{loc}, which is generated by the proactive term and which is absent in panel (a), differs systematically from these quantities. The ranges of the non-reciprocity parameter ϵ explored in panels (a) and (b) are different, but were chosen so as to produce pinning coefficients of comparable magnitude.

so that the condition in Eq. (5.42) is no longer satisfied. Accordingly, the contribution of $\mathcal{A}$ cannot be regarded as a small perturbation or as an additional uncorrelated source of noise. Instead, it contributes significantly to the dynamics of the mean orientation. Consistently with this picture, the relaxation rate extracted from the autocorrelation function is found to be $r^{\text{OU}} \simeq 10^{-3}$, approximately one order of magnitude smaller than the value $r^{\text{glob}} \simeq 10^{-2}$ predicted analytically. This discrepancy indicates that, in this regime, the dynamics of the mean orientation cannot be accurately described by a simple OU process with relaxation rate determined solely by the global pinning coefficient.

By systematically analyzing all the interaction kernels considered in this work (see Sec. 4.1.3), we find that the relaxation rate r^{OU} provides a reliable estimate of the global pinning coefficient r^{glob} only for the modulated coupling, both in the purely reactive dynamics (see Fig. 5.1(a)) and when the proactive term is included (see Fig. 5.1(b)). For the remaining kernels, the additional terms contained in $\mathcal{A}(t)$ contribute significantly to the dynamics of the mean orientation and do not allow a direct identification of r^{OU} and r^{glob}.

5.3.2 Pinning extracted from the structure factor

In this section, we compare the global pinning coefficient r^{glob} and the average local pinning coefficient r^{loc} with the effective pinning term $r^{S(\mathbf{k})}$ extracted from the static structure factor of

configurations in the stationary state.

We start from the evolution equation for the fluctuations $\delta\phi_i$ around a pinned direction Φ, given in Eq. (5.9). If the modulation of the diffusion coefficient is weak, i.e., if $\kappa(\delta\phi_i) \simeq \kappa$, the fluctuations of the local pinning coefficient satisfy $\Delta_{r^{\rm loc}} \ll r^{\rm loc}$, and the additional contributions $\mathcal{A}_i$ can be neglected, then Eq. (5.9) reduces to the simplest scalar theory characterized by a pinning coefficient $r^{\rm loc}$. In this regime, $r^{\rm loc}$ directly controls the long-wavelength spatial correlations. However, once the fluctuations of the local pinning coefficient and the neglected terms $\mathcal{A}_i$ become non-negligible, the effective pinning governing the spatial correlations may be renormalized by fluctuations. To investigate the magnitude of this effect, we compare the analytical estimates of the pinning coefficients with the pinning extracted from the static structure factor $S(\mathbf{k}) = \langle |\phi(\mathbf{k})|^2 \rangle$ of the fluctuations, where $\phi(\mathbf{k})$ denotes the Fourier transform of the local field ϕ_i.

For a scalar field obeying the equation

$$\partial_t \delta\phi_i = \kappa^{S(\mathbf{k})} \nabla_{\rm d}^2 \delta\phi_i - r^{S(\mathbf{k})} \delta\phi_i + \sqrt{2T}\,\eta_i, \tag{5.47}$$

the stationary structure factor is given by the lattice Ornstein–Zernike form

$$S(\mathbf{k}) = \frac{T}{r^{S(\mathbf{k})} + \kappa^{S(\mathbf{k})}\,\hat{k}^2}, \tag{5.48}$$

where $\hat{k}^2 = 4\sin^2(k_x/2) + 4\sin^2(k_y/2)$ is the squared lattice momentum. In the fits discussed below, the temperature T is fixed to the value used for the simulations, while both $r^{S(\mathbf{k})}$ and $\kappa^{S(\mathbf{k})}$ are treated as free parameters. The central question is whether the pinning coefficient $r^{S(\mathbf{k})}$ extracted from the spatial correlations can be identified with either the average local pinning coefficient $r^{\rm loc}$ or the global pinning coefficients $r^{\rm glob}$, and under which conditions such identification is valid.

We address this question for the modulated coupling kernel, considering both the purely reactive dynamics and the complete dynamics which includes also the proactive contribution. In the purely reactive case, the evolution equation (5.7) expressed in terms of the fluctuations $\delta\phi_i = \phi_i - \pi/4$ around the pinned direction $\Phi = \pi/4$ takes the form (see Eq. (B.5))

$$\partial_t \delta\phi_i = \left[1 - \epsilon\cos\left(4\delta\phi_i\right)\right] \nabla_{\rm d}^2 \delta\phi_i + \sqrt{2T}\,\eta_i(t). \tag{5.49}$$

Comparing with Eq. (5.9), one can identify $\kappa(\delta\phi_i) = 1 - \epsilon\cos\left(4\delta\phi_i\right)$, while no local pinning term of the form $r_i^{\rm loc}\delta\phi_i$ is present. This is consistent with the discussion in Sec. 5.2.1: the continuum equation does not contain a local pinning coefficient $r_i^{\rm loc}$ that can be directly compared with the effective pinning extracted from the structure factor. In the absence of the local pinning term, a natural approximation is to use the global pinning coefficient $r^{\rm glob}$ as a proxy for the pinning experienced by the local fluctuations and compare it with the effective pinning coefficient $r^{S(\mathbf{k})}$ extracted from the static structure factor. This is essentially the approach adopted in Ref. [92], where the local fluctuations are assumed to experience the same pinning strength as the global orientation. In the present case, $r^{\rm glob}$ is obtained by evaluating the analytical expression in Eq. (B.41) and averaging it over the stationary fluctuations.

In Fig. 5.1(a), we compare the values of $r^{S(\mathbf{k})}$, obtained by fitting the static structure factor with Eq. (5.48), with the corresponding values for $r^{\rm glob}$ for stationary configurations generated at $T = 0.1$ and $L = 800$, as a function of the non-reciprocity parameter ϵ. We find good agreement throughout the entire range of ϵ considered. This indicates that, for the modulated coupling, the

global pinning coefficient provides an accurate effective description of the spatial fluctuations despite the absence of an explicit local pinning term in the evolution equations of the fluctuations.

In the case of the modulated-coupling kernel with the complete dynamics, the evolution equation (5.7) expressed in terms of the fluctuations $\delta\phi_i = \phi_i$ around the pinned direction $\Phi = 0$ takes the form (see Eqs. (B.5) and (B.7))

$$\begin{aligned}\partial_t \delta\phi_i = & \left[1 + \epsilon \cos\left(4\delta\phi_i\right)\right] \nabla_{\mathrm{d}}^2 \delta\phi_i \\ & + \frac{\sin\left(4\delta\phi_i\right)}{4\delta\phi_i} S_{\mathrm{P},i}^{(4),\mathrm{loc}}\, \delta\phi_i + \sqrt{2T}\, \eta_i(t).\end{aligned} \tag{5.50}$$

Comparing with Eq. (5.9), one readily identifies

$$\kappa(\delta\phi_i) = \left[1 + \epsilon \cos\left(4\delta\phi_i\right)\right] \quad \text{and} \quad r_i^{\mathrm{loc}} = \frac{\sin\left(4\delta\phi_i\right)}{4\delta\phi_i} S_{\mathrm{P},i}^{(4),\mathrm{loc}}. \tag{5.51}$$

Unlike the purely reactive case, an explicit pinning term is already present in the local equation of motion, since r_i^{loc} remains finite in the limit $\delta\phi_i \to 0$. This enables a direct comparison between the (spatial) average r^{loc} of r_i^{loc} and the effective pinning coefficient $r^{S(\mathbf{k})}$ extracted from the static structure factor.

In Fig. 5.1(b), we report the comparison between r^{loc} and $r^{S(\mathbf{k})}$ as a function of ϵ. We find that the two quantities are of the same order of magnitude, although they agree less than r^{glob} and r^{OU}. Interestingly, we find that the values of $r^{S(\mathbf{k})}$ are actually closer to the global pinning coefficient r^{glob} than to the local coefficient r^{loc}. This suggests that, in the regime where $\Delta_{r^{\mathrm{loc}}} \lesssim r^{\mathrm{loc}}$ (see the errobars of r^{loc} in Fig.5.1(b)[2]), the local pinning coefficient undergoes a non-negligible renormalization due to the remaining terms in the local equation of motion. As a result, the effective pinning governing the spatial correlations appears to be closer to the global pinning coefficient r^{glob} than to the local quantity r^{loc} itself.

The discussion above focused on the simple case of the modulated coupling. For interaction kernels with vision-induced anisotropy, additional complications arise. First, the form of the evolution equations of the fluctuations in the continuum, analogous to Eq. (5.9), become considerably more involved, containing several additional contributions that can renormalize the effective pinning. Second, the local pinning coefficient r_i^{loc} may exhibit large fluctuations. This is the case, for instance, of the sinusoidal coupling (see Sec. 4.1.3) for the parameters values considered in Sec. 5.2.2, which correspond to those studied in Ref. [92]. In fact, although the average value r^{loc} of r_i^{loc} remains positive, indicating an overall tendency toward pinning along lattice directions, strong local fluctuations ($\Delta_{r^{\mathrm{loc}}} \gg r^{\mathrm{loc}}$) can drive r_i^{loc} towards strongly negative values within finite spatial regions and finite time intervals. As a consequence, a quantitative comparison between r^{loc} and r^{glob} with the effective pinning coefficient $r^{S(\mathbf{k})}$ extracted from the static structure factor becomes considerably more subtle. This scenario likely contributes to the discrepancy, of approximately a factor of three, reported in Ref. [92] between the correlation length computed from the global pinning coefficient ($\xi = \sqrt{\kappa / r^{\mathrm{glob}}}$), and the value of the correlation length inferred directly from the structure factor.

To summarize, for the modulated-coupling kernel, the assumption that local fluctuations

[2]We emphasize that the error bars associated with r^{loc} are of a fundamentally different nature from those of r^{OU}, r^{glob}, and $r^{S(\mathbf{k})}$. While the fluctuations underlying the latter quantities vanish in the thermodynamic limit, the uncertainty associated with r_i^{loc} corresponds to the intrinsic standard deviation $\Delta_{r^{\mathrm{loc}}}$. It therefore reflects genuine spatial and temporal variability of the local pinning coefficient rather than finite-size statistical uncertainty.

experience the same effective pinning strength as the global orientation provides a remarkably accurate description. This is true both for the purely reactive dynamics, where no explicit local pinning term is present, and for the complete dynamics including the proactive contribution.

In the latter case, the results suggest that the local pinning coefficient is subject to significant renormalization due to the additional terms in the local dynamics, with r^{glob} playing the role of an effective renormalized pinning that controls the fluctuations of the spins.

The situation is considerably more subtle for the sinusoidal coupling. As already discussed in Sec. 5.3.1, even the comparison between r^{glob} and the relaxation rate r^{OU} fails in this case. Furthermore, the local pinning coefficient exhibits very large fluctuations, $\Delta_{r^{\mathrm{loc}}} \gg r^{\mathrm{loc}}$, making any effective mean-field description based on a single renormalized pinning coefficient highly questionable. We therefore expect the relation between the analytical pinning coefficients and the correlation length to be substantially more intricate in this case.

5.4 Concluding remarks

In this chapter, we have investigated how the lattice-induced pinning behavior identified in Chapter 4 emerges in the equations governing local and collective orientational degrees of freedom. Our main result is that local and global pinning are distinct properties. A restoring torque may appear in the equations for the collective orientation even when the equation of motion of an individual spin contains no explicit term proportional to its local angular fluctuation.

We formalized this distinction through the local coefficient $r_i^{\mathrm{loc}}(t)$, its stationary average r^{loc}, and the global coefficient r^{glob}. The former characterizes the restoring term in the local linearized dynamics, whereas the latter governs the relaxation of the zero-wavevector mode (the global orientation). In non-reciprocal systems, these coefficients originate from different combinations of microscopic terms and do not generally coincide.

The interaction kernels considered provide explicit examples of this distinction. For the modulated kernel with reactive-only dynamics, no local restoring term appears, although the global equation contains finite global pinning. The proactive contribution instead generates an explicit local term. For the sinusoidal kernel, the reactive contribution produces local pinning while the proactive one does not, although both contribute to the global restoring torque. Global pinning therefore cannot be inferred solely from the presence of a local pinning term.

For the von Mises coupling kernel with purely reactive dynamics, we show that local pinning originates already at order $\mathcal{O}(\sigma)$, corresponding to the sinusoidal component of the interaction. The global pinning term, however, receives contributions not only from the $\mathcal{O}(\sigma)$ term but also, crucially, from the $\mathcal{O}(\sigma^4)$ contribution, which is the first term to generate the $\mathbb{Z}_4$-symmetric modulation responsible for the selection of the diagonal directions. A low-order truncation of the kernel is therefore insufficient to reproduce the complete pinning mechanism.

We compared the analytical coefficients with the relaxation rate r^{OU} extracted from the autocorrelation of the mean orientation and with the effective pinning $r^{S(\mathbf{k})}$ obtained from the structure factor. The identification $r^{\mathrm{OU}} \simeq r^{\mathrm{glob}}$ is accurate when the residual terms in the collective equation are weakly correlated with the global fluctuation, as occurs for the modulated kernel. For the sinusoidal kernel, these terms are comparable with the linear restoring force, preventing a quantitative reduction to a simple Ornstein–Uhlenbeck process.

For the modulated-coupling kernel, we find that $r^{S(\mathbf{k})}$ is remarkably close to the global pinning coefficient r^{glob}, even when the explicit local pinning term is present. For more complex interaction

kernels, however, large fluctuations and additional renormalization effects can substantially modify the dynamics and therefore the relation between the analytical pinning coefficients and the effective pinning extracted from the static structure factor.

Chapter 6

Future perspectives on non-reciprocal XY models

In this chapter, we discuss several research directions that emerged during the course of this doctoral work but have not yet been explored sufficiently to develop the preliminary findings and conjectures into fully established results.

The first topic concerns the metastability of the long-range ordered (LRO) phase in the non-reciprocal XY model. The existence and nature of metastable states in this model have recently attracted considerable attention in the literature [92, 93, 139, 140]. Here, however, we focus on a different aspect, namely the role played by the proactive term introduced in Chapter 4. We shall argue that many of the features of the metastable dynamics change qualitatively once this additional contribution is taken into account. Since the proactive term has not been considered in previous studies which address metastability in the NRXY model, these observations open a number of interesting questions that deserve further investigation.

The second topic concerns the consequences of introducing the proactive interaction in systems off lattice. Rather than formulating a microscopic model in terms of coupled stochastic differential equations, our aim is to retain the spirit of the Vicsek approach and to construct a cellular-automaton-like dynamics in which the proactive interaction emerges naturally in the alignment rules. Preliminary numerical investigations suggest that this ingredient may generate novel collective behaviors, and in particular may induce effective steering interactions leading to the spontaneous formation of organized structures.

This chapter is organized as follows. In Sec. 6.1, we revisit the possible metastability of the LRO phase of the NRXY model. In Sec. 6.1.1 we review the defect-mediated decay mechanisms reported for constant-density flocks and non-reciprocal planar spin systems, emphasizing how non-reciprocity modifies the shape, mobility, and annihilation dynamics of topological defects. In Sec. 6.1.2 we discuss the possible role of the proactive term in stabilizing the dynamics, hindering the defect-mediated decay. In Sec. 6.1.3 we present a preliminary numerical experiment in which an aster-rich configuration generated under reactive dynamics progressively recovers orientational order after the proactive term is switched on in the dynamics.

In Sec. 6.2, we introduce the reactive and proactive terms in the dynamics of self-propelled particles moving in continuous space. In Sec. 6.2.1 we formulate the interaction in the language of overdamped intelligent active Brownian particles and discuss the difficulties associated with

a direct continuous-time numerical implementation in the absence of a fixed microscopic length scale. In Sec. 6.2.2 we construct a genuinely discrete Vicsek-like dynamics in which the reactive contribution aligns the headings of neighboring agents, while the proactive contribution generates a steering response determined by by the relative positions of neighboring agents and by the gradient of the vision kernel. Finally, in Sec. 6.2.3 we present preliminary numerical observations showing that the addition of proactivity is capable to suppress particle dispersion, promoting the spontaneous formation of cohesive, elongated moving structures.

6.1 Metastability of the LRO phase in the NRXY model

In Chapters 2, 4, and 5, we investigated the properties of the LRO phase emerging in the non-reciprocal XY model with anisotropic interaction due to the presence of a vision anisotropy in the interaction. In particular, we characterized its stationary fluctuations, its orientational pinning, and its dependence on the microscopic implementation of the dynamics. Throughout these studies, the numerical simulations were sufficiently long to reach an apparently stationary ordered regime, in which we collected the statistical data required for our analysis. Within the corresponding observation times, we did not detect any systematic breakdown of the LRO phase.

The absence of a detectable decay, however, should not be regarded as conclusive evidence of the asymptotic stability of the LRO phase. In fact, the relaxation time within the ordered phase and the time required for the nucleation of a rare topological excitation may be actually separated by several orders of magnitude. Consequently, a configuration may appear stationary over the numerically accessible times, which are relevant for the measurement of conventional observables, while it eventually decays via a rare, non-perturbative event. In this sense, the ordered states investigated in the previous chapters may be either genuinely stable or metastable with a nucleation time much longer than the duration of our simulations.

6.1.1 Defect-mediated decay of LRO phase

The metastability of ordered phases related to the one observed in the NRXY model was first discussed in Ref. [139] for constant-density flocks. In particular, it turned out that the homogeneous ordered phase predicted by the perturbative hydrodynamic theory is metastable with respect to the nucleation of a characteristic defect configuration. The first event in which such a nucleation occurs typically consists of a radially organized aster, the center of which carries topological charge $q = +1$ (see, c.f., Eq. (6.1)), surrounded by a semicircular shock line across which the orientation of the local field varies rapidly. This is illustrated in Figs. 6.1(a). Once the first aster has appeared, further asters may nucleate along the pre-existing shock lines and close to their intersections. The ordered phase is then progressively replaced by a dynamically evolving active foam consisting of asters separated by a network of shock lines. An example of this process, referred to as an "aster apocalypse" in Ref. [93], is shown in Fig. 6.1. The resulting mechanism is qualitatively different from the proliferation of point-like vortex-antivortex pairs in the equilibrium XY model. In the active case, the extended interfaces themselves constitute the regions in which fluctuations are amplified and from which new defects are preferentially generated. A close connection between this kind of evolution and the NRXY model is suggested by the continuum descriptions derived for non-reciprocal planar spins. At the mean-field hydrodynamic level, the leading contributions generated by vision-dependent non-reciprocal interactions have the same

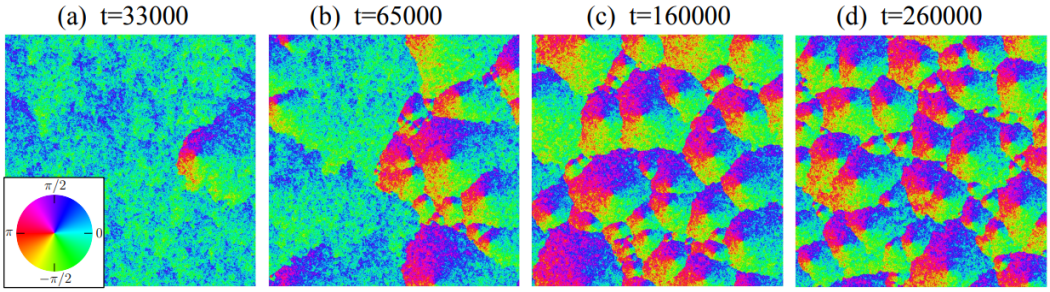


Figure 6.1: Nucleation of an aster from an initially ordered constant-density flock and subsequent proliferation of asters and shock lines. The local orientation is represented by the color wheel. The initially homogeneous ordered phase is eventually replaced by a dynamically evolving active foam. Adapted from Ref. [139].

structure as those entering constant-density flocking theories [92]. This correspondence does not imply that the two microscopic systems have identical defect dynamics, but it provides a strong motivation for expecting similar non-perturbative decay mechanisms.

Indeed, the nucleation of asters and the destruction of the polarized phase have subsequently been observed directly in lattice realizations of the NRXY model [92, 93]. In particular, Ref. [93] showed that defects of opposite charge can become effectively screened, or "cloaked", from one another. As a consequence of this screening, their annihilation is delayed, leaving sufficient time for additional defects to nucleate and to eventually wipe out the polarized state through a process analogous to the aster apocalypse mentioned above.

The point-like excitations of the reciprocal XY model are classified by the integer winding number

$$q = \frac{1}{2\pi} \oint_{\mathcal{C}} \nabla\phi \cdot d\boldsymbol{\ell}, \tag{6.1}$$

where $\mathcal{C}$ is a closed contour surrounding the defect core. In the reciprocal model, defects with charges $q = \pm 1$ have equivalent properties up to symmetry transformations. Vision-dependent non-reciprocal interactions substantially modify this picture, because the dynamics then depends not only on the topological charge but also on the detailed shape and orientation of the spin field surrounding the defect.

This effect was investigated systematically in Ref. [65]. For an isolated defect of charge $q = \pm 1$, the associated phase field $\Phi_q(r, \varphi)$ may be written on the continuum as

$$\Phi_q(r, \varphi) = q\varphi + \mu_q, \tag{6.2}$$

where φ is the polar angle around the defect core and μ_q is defined modulo 2π. The physical meaning of μ_q depends on the topological charge. For a $q = +1$ defect, μ_+ changes the intrinsic form of the texture: $\mu_+ = 0$ corresponds to a source, $\mu_+ = \pi$ to a sink, and $\mu_+ = \pm\pi/2$ to the two vortex-like configurations. At zero temperature, the source and the sink are both stationary configurations, but the source is an unstable fixed point, whereas the sink is stable. Generic initial configurations with $\mu_+ \neq 0, \pi \pmod{2\pi}$ are therefore progressively twisted toward the sink. In particular, Figs. 6.2(a)—(d) show the evolution from $\mu_+ = \pi/6$ to $\mu_+ = \pi$; the unstable source configuration is not shown. For a $q = -1$ defect, instead, μ_- specifies the global orientation of the texture: changing μ_- by $\Delta\mu_-$ rigidly rotates the defect by $\Delta\mu_-/2$. It therefore fixes the orientation of the defect and of its twofold symmetry axis, rather than measuring the degree of

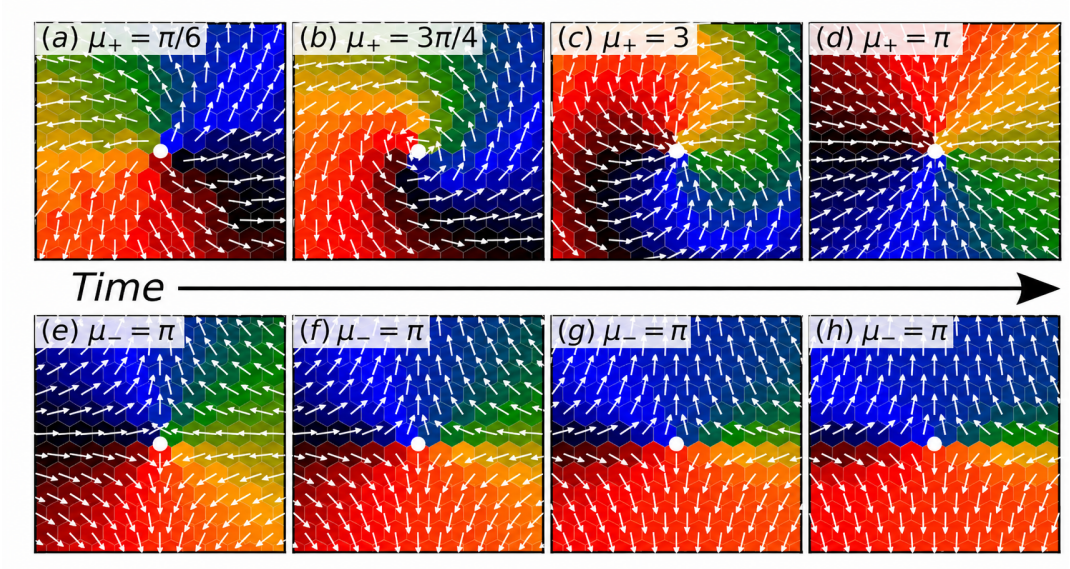


Figure 6.2: Evolution of topological defects induced by non-reciprocal vision-dependent interactions, in the case of von Mises kernel. White arrows indicate the directions of the spins at the various lattice sites, while the background colorcode represent the associated angle. (a)–(d) An initial configuration with a defect of charge $q = +1$ (see Eq. (6.1)), represented in panel (a), evolves toward the radially sink defect, also called aster, represented in panel (d), with the value of μ_+ progressively increasing from $\pi/6$ to π. (e)–(h) An initial configuration with a defect of charge $q = -1$, represented in panel (e), evolves toward a polarized defect characterized by a defined axis, shown in panel (h). In this case the value of μ_- remains constant. Adapted from Ref. [65].

polarization generated by the non-reciprocal dynamics. As shown in Figs. 6.2(e)–(h), an initially unpolarized $q = -1$ defect evolves toward a polarized configuration with a well-defined symmetry axis. Motion parallel to this axis is facilitated, whereas displacement in the perpendicular direction is hindered. Consequently, non-reciprocity may either accelerate or slow down the annihilation of a pair of oppositely charged defects, depending on their initial shapes and relative orientation.

6.1.2 A possible stabilizing role of the proactive term

The studies discussed above considered primarily the dynamics that, according to the terminology introduced in Chapter 4, contains only the reactive contribution. In this implementation, a spin responds to the instantaneous orientations of the neighbors currently lying inside its vision cone, without accounting for the change in the set of perceived neighbors caused by its own rotation.

However, as explained in Section 4.1.2, the derivative of the local, configuration-dependent selfish energy produces an additional contribution to the Langevin dynamics, which we referred to as the proactive term. This term accounts for the variation of the interaction kernel under a rotation of the spin. In the interpretation we developed, it describes a tendency of each agent to preserve, or possibly maximize, the number and strength of the interactions that it perceives. Accordingly, a spin experiences an additional response whenever its rotation is about to move one of its nearest neighbors across the boundary of its vision region.

Previous investigations of the dynamics of pairs of defects suggest that adding what we have identified as the proactive term may alter the stability of source-like configurations, favoring annihilation pathways which are suppressed or slowed down under purely reactive dynamics

[65]. This observation motivates the conjecture that proactivity influences not only the dynamics of individual defect pairs, but also the metastability of the macroscopic ordered phase. In particular, by enhancing the elimination of defects or preventing their long-lived cloaking, the proactive term might counteract the proliferation of asters, thereby stabilizing the LRO phase.

A complementary intuitive argument follows from the local interpretation of the proactive interaction. The formation of a shock line requires spins in neighboring regions to have sharply different orientations. Spins close to the interface between these regions (i.e., close to the shock line) continuously gain and lose perceived neighbors as their orientations fluctuate. A proactive response opposing such losses may favor the local reconstruction of a coherent neighborhood and suppress the persistence of strongly mismatched domains. Within this picture, the same microscopic mechanism responsible for the tendency of an agent to retain its perceived neighbors could also promote the healing of shock lines and the restoration of global order.

At this stage, these arguments remain conjectural. Establishing whether the proactive term makes the LRO phase asymptotically stable would require a systematic study of nucleation times, their dependence on system size, and the statistics of defect creation and annihilation. Here we restrict ourselves to presenting a simple numerical experiment that provides preliminary and qualitative support to this scenario.

6.1.3 A preliminary switching experiment

We consider the NRXY model on a square lattice with the hard vision-cone kernel introduced in Sec. 2.1.1 and 4.1.3. For this kernel, however, a direct implementation of the proactive contribution in a Langevin equation is problematic, because the derivative of the step-like interaction kernel consists of Dirac delta contributions localized at the boundary of the vision cone.

To overcome this difficulty, we compare the results of the evolution according to two different microscopic implementations of the dynamics: During the first part of the simulation, the system evolves according to the Langevin dynamics with the sole reactive term described in Sec. 4.1.2. At a prescribed switching time t_{sw}, this dynamics is replaced by the Glauber-type dynamics described in Sec. 4.1.1. Since the corresponding acceptance probabilities are computed from the variation of the configuration-dependent selfish energy associated with the proposed local update, this second implementation actually corresponds to a Langevin dynamics which accounts for both reactive and proactive terms.

Starting the dynamics from a perfectly ordered configuration would make the numerical observation of the first nucleation event prohibitively expensive. Accordingly, we exploit the observation done in Ref. [139] that shock lines and their intersections constitute preferential regions for the nucleation of new asters. The initial condition is therefore prepared by dividing the system into two ordered half-planes with different spin orientations ϕ_i, i.e.,

$$\phi_i(t=0) = \begin{cases} \phi_{\mathrm{lower}} & \text{for} \quad y_i < L/2, \\ \phi_{\mathrm{upper}} & \text{for} \quad y_i \geq L/2, \end{cases} \tag{6.3}$$

where $y_i = 1, 2, \ldots, L$ denotes the vertical coordinate of site i. The discontinuity of ϕ_i across the two domains provides a controlled initial interface from which defect nucleation may occur.

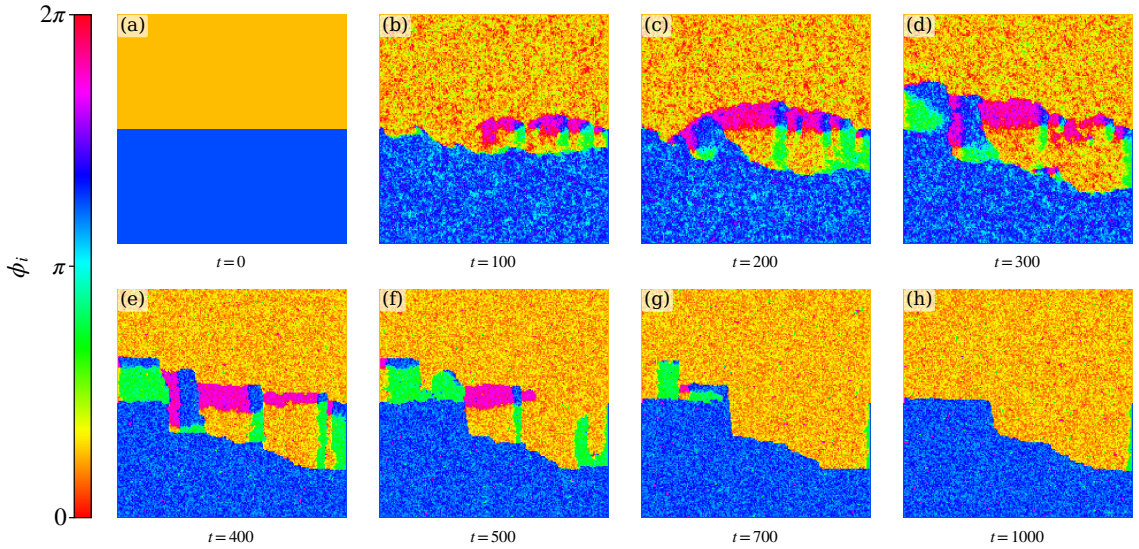


Figure 6.3: Preliminary results of the switching experiment in the NRXY model starting from the two-domain configuration in Eq. (6.3). Panels from (a) to (d) show the evolution of the model under the reactive Langevin dynamics. The initially flat interface becomes unstable and acts as a source for the nucleation and proliferation of asters. At time $t = t_{\rm sw} = 300$, the microscopic dynamics is changed to the Glauber implementation, which includes the proactive term. Panels from (e) to (h) show the subsequent evolution: the initial proliferation is reversed as the asters progressively disappear and a macroscopically ordered configuration with two bulk orientations is eventually recovered. The parameters used in the simulations are $L = 200$ for the lattice size, $\theta_{\rm R} = 4\pi/3$, $\theta_{\rm G} = 2\pi/3$, while the temperature of the thermal bath in both phases of the evolution was fixed to $T = T_{\rm R} = T_{\rm G} = 0.2$.

In the simulations shown in Fig. 6.3, we use

$$\phi_{\rm lower} = \frac{7}{4}\pi \quad \text{and} \quad \phi_{\rm upper} = \frac{1}{4}\pi. \tag{6.4}$$

The evolution with the reactive term for $t < t_{\rm sw}$ is performed with vision-cone opening $\theta_{\rm R} = 4\pi/3$ while, after the switch, we use a vision-cone angle $\theta_{\rm G} = 2\pi/3$ for the Glauber dynamics. These two values are chosen such that the orientations $\phi_{\rm lower}$ and $\phi_{\rm upper}$ used in the initial state correspond to the preferred zero-temperature bulk orientations for the respective dynamical implementations. Note that the choice $\theta_{\rm R} \neq \theta_{\rm G}$ is motivated by the fact discussed in Section 4.3 that, for a certain amplitude of the vision cone, the set of preferred directions of the reactive Langevin dynamics differs from that of the Glauber dynamics. The chosen values of the vision cone angles $\theta_{\rm R}$ and $\theta_{\rm G}$ avoid that the decay of a domain is trivially due to the fact that the orientation of their spins is no longer compatible with the pinning direction of the the dynamics applied after the switch.

Figure 6.3 shows the resulting evolution of the spin configuration, where the color of each pixel indicates the angle ϕ_i. In particular, the first line of subsequent snapshots corresponds to the evolution with the reactive Langevin dynamics, while the second line to the Glauber dynamics. Under the initial reactive dynamics, the initially imposed sharp interface (see panel (a)) rapidly becomes unstable as a number of asters nucleate along the interface and subsequently invade an increasingly larger fraction of the system, as in panel (d). This evolution confirms the role of shock lines as favorable locations for the nucleation of defects, which then trigger a defect-mediated decay of the initially ordered regions, as the one reported for both constant-density flocks and

reactive implementations of the NRXY model.

After the switch, the dynamics shows the reversed tendency, as clearly demonstrated by the subsequent panels (e)—(h) of Fig. 6.3. In fact, the pre-existing asters in panel (d) and (e) progressively shrink or annihilate, the network of shock lines progressively disappears, and the system returns towards an ordered configuration with a single, well-defined shock line separating two regions with different spin orientations, as illustrated moving from panel (e) to panel (h). We observed the same qualitative behavior in simulations performed for different values of ϕ_{lower} and ϕ_{upper} and of the corresponding θ_{R} and θ_{G}. Although these observations are based only on a limited set of runs, they consistently suggest that the microscopic response contained in the Glauber implementation suppresses the proliferation of defects generated by the reactive dynamics and ultimately promotes their annihilation.

This numerical results provide support to the findings of Chapter 2 and of Ref. [89] concerning the existence of a phase with LRO in the phase diagram of the NRXY model with hard vision cone and Glauber-type dynamics. In fact, the low-temperature phase found there displayed robust LRO over all the sizes and times investigated. The present switching experiment suggests that this apparent robustness might not be a consequence of having a very long nucleation time. Instead, it might be a consequence of the fact that the proactive response actively remove the defects that would otherwise destabilize the ordered state.

A hard vision-cone interaction combined with Glauber-type orientational updates has also been considered in a lattice-gas model in which spins are allowed to move within the lattice, where the formation of a single macroscopic aster was reported [141], in an apparent contradiction with what we have observed above. However, the comparison with the preliminary results presented here must be made cautiously. In that model, the agents possess translational mobility within the lattice, and the coupling between positional and orientational degrees of freedom can stabilize spatial structures which have no direct counterpart in the NRXY model on the lattice considered here.

Finally, the experiment in Fig. 6.3 should not be interpreted as a proof that the proactive term alone stabilizes the LRO phase. At the switching time, both the update rule and the vision-cone angle are modified. The observed reversal may therefore contain contributions associated with other differences between Langevin and Glauber dynamics. Nevertheless, the remarkably fast disappearance of defects after the switch, together with its robustness across the preliminary runs we have considered, provides a valuable indication that proactivity may favor the healing of shock lines and the annihilation of asters.

A more conclusive analysis could be performed using a differentiable vision kernel, for which reactive Langevin dynamics can be easily supplemented with the reactive term, while keeping all other microscopic parameters unchanged. It would then be possible to determine whether the proactive term merely increases the lifetime of the ordered phase or it suppresses defect proliferation strongly enough to render the LRO phase asymptotically stable. The extension of this investigation to different vision kernels, system sizes, and noise intensities remains an open direction for future work.

6.2 Proactive interactions in off-lattice active matter

In Chapters 4 and 5, we investigated the consequences of reactive and proactive interaction terms in a lattice model. While the presence of the lattice provided a controlled environment in which

the two terms could be isolated and characterized, an important open question concerns their role in systems of self-propelled particles moving in continuous space, i.e., off lattice.

Vision-based interactions are widespread in active-matter models and have been extensively studied in the context of flocking and collective motion. Most existing approaches, however, interpret visual information as a means to implement alignment, avoidance, or following behaviors [88]. In contrast, the proactive interaction introduced in this thesis originates from a different principle. Rather than being merely proportional to the coupling with the configuration of neighbouring agents, it is proportional to the gradient of that coupling. As a consequence, the interaction becomes strongest (in modulus) precisely when a neighbour is at the edge of the field of view (that is, when it has just entered it or is about to leave it).

Preliminary numerical investigations suggest that such a contribution can qualitatively modify the collective behaviour of self-propelled particles, generating effective pursuit dynamics and the spontaneous formation of elongated queues and filamentary structures. These observations motivate the extension of the reactive-proactive framework from lattice systems to off-lattice active matter.

The goal of this section is to proposed a minimal continuous-space model incorporating both reactive and proactive interactions. We begin from the framework of intelligent active Brownian particles (iABP), from which the vision-dependent angular dynamics emerges. We then derive a simplified overdamped description and finally introduce a discrete Vicsek-like implementation that preserves the physical interpretation of the reactive and proactive contributions while avoiding the introduction of instrumental numerical time and length scales for the numerical simulation.

6.2.1 Overdamped intelligent active Brownian particles

We first frame the model in the language of intelligent active Brownian particles [87, 88]. In this formulation, each particle (*agent*) is described by its position $\mathbf{x}_i$ and by an internal orientation vector

$$\hat{\mathbf{n}}_i = (\cos\phi_i, \sin\phi_i). \tag{6.5}$$

The translational underdamped dynamics is written as

$$m\ddot{\mathbf{x}}_i = f_p\hat{\mathbf{n}}_i - \gamma\dot{\mathbf{x}}_i, \tag{6.6}$$

where m is the mass, f_p the self-propulsion force and γ the translational friction. In the overdamped translational limit $m \to 0$, this reduces to

$$\dot{\mathbf{x}}_i = v_0\hat{\mathbf{n}}_i \quad \text{with} \quad v_0 = f_p/\gamma, \tag{6.7}$$

i..e, to the standard constant-speed active-particle kinematics.

The dynamics of the orientation $\hat{\mathbf{n}}_i$ can be written in the spirit of the inertial spin model [142], by introducing a spin variable $\mathbf{s}_i$, which acts as an angular momentum for the turning dynamics:

$$\dot{\hat{\mathbf{n}}}_i = \mathbf{s}_i \times \hat{\mathbf{n}}_i, \tag{6.8}$$

$$\chi\dot{\mathbf{s}}_i = \hat{\mathbf{n}}_i \times \left[-\eta\dot{\hat{\mathbf{n}}}_i + \boldsymbol{\zeta}_i + \mathbf{M}_i\right], \tag{6.9}$$

where χ is a rotational inertia, η is a rotational damping coefficient, $\boldsymbol{\zeta}_i$ is a stochastic torque and $\mathbf{M}_i$

is the generalized steering force generated by the interactions of the i-th agent with neighbouring ones. In Ref. [88], $\mathbf{M}_i$ was assumed to contain alignment, avoidance, and following contributions, all defined through vision-based neighbourhoods. Let us emphasize here that the proactive term introduced in this thesis does not belong to any of these categories of responses, though its effect is similar to the following contribution, as it promotes cohesion.

In the present work we do not aim at investigating inertial turning effects. Accordingly, we focus on the overdamped angular regime $\chi \to 0$ in which the spin variable can be adiabatically eliminated assuming $\dot{\mathbf{s}}_i \simeq 0$. The resulting alignment dynamics corresponds to the continuous-time counterpart of the Vicsek model, see Ref. [88]. This overdamped description can be written directly in terms of a first-order evolution equation for the orientation angle ϕ_i which characterizes $\hat{\mathbf{n}}_i$. In our case, this angular dynamics is driven by two distinct contributions, exactly as in the previously studied lattice models, i.e., a reactive and a proactive term, where the former corresponds to the standard alignment contribution of Ref. [88]:

$$\dot{\phi}_i = -J \sum_{j \neq i} g(\phi_i, \vartheta_{ij}, r_{ij}) \sin(\phi_i - \phi_j) + J \sum_{j \neq i} \partial_{\phi_i} g(\phi_i, \vartheta_{ij}, r_{ij}) \cos(\phi_i - \phi_j) + \sqrt{2T}\, \eta_i(t). \tag{6.10}$$

The interaction kernel g depends both on the relative distance r_{ij} between two particles and on their bearing angle ϑ_{ij}, defined as

$$r_{ij} = |\mathbf{r}_j - \mathbf{r}_i| \quad \text{and} \quad \vartheta_{ij} = \arg\left(\mathbf{r}_j - \mathbf{r}_i\right). \tag{6.11}$$

For simplicity, we assume that this dependence factorizes as

$$g(\phi_i, \vartheta_{ij}, r_{ij}) = g_r(r_{ij})\, g_\theta(\phi_i, \vartheta_{ij}), \tag{6.12}$$

in terms of the radial and the angular modulations g_r and g_θ, respectively. Moreover, we choose a radial modulation g_r that is approximately constant at short distances, while it decays exponentially upon increasing r beyond a characteristic range R_0:

$$g_r(r) = \begin{cases} 1, & \text{for} \quad r \leq R_0, \\ \exp\left(-\dfrac{r - R_0}{\ell_r}\right) & \text{for} \quad r > R_0, \end{cases} \tag{6.13}$$

where ℓ_r controls the decay length. The angular component g_θ of the interaction kernel is chosen from the same class of anisotropic functions as those introduced for the lattice model in Eqs. (4.7)—(4.9). Accordingly, all the angular dependences previously considered for the interactions, including von Mises kernels, sinusoidal kernels, and hard vision-cone interactions, can be directly implemented within the present off-lattice framework.

The problem of numerical integration Equations (6.7) and (6.10) provide a continuous-time description of the model. A direct numerical implementation of the prescribed evolution would require the introduction of a finite integration time step Δt. The corresponding displacement $ell_{\Delta t}$ of a particle during a single update is

$$\ell_{\Delta t} = v_0 \Delta t. \tag{6.14}$$

For the numerical integration to represent accurately the underlying continuous dynamics, this displacement should remain small compared to the characteristic spatial scales over which the local neighbourhood changes. In systems with excluded-volume interactions, such scales are naturally set by the interaction range or by the particle diameter. For example, one may supplement Eq. (6.7) with a repulsive force term,

$$\dot{\mathbf{r}}_i = v_0\hat{\mathbf{n}}_i + D_t\mathbf{F}_i, \tag{6.15}$$

where $\mathbf{F}_i$ derives from a hard-sphere, Lennard–Jones, or Weeks–Chandler–Andersen (WCA) (see, e.g., Ref. [87]). In such cases, the interaction introduces a microscopic length scale σ, and the integration step can be chosen such that the displacement during one update remains much smaller than this characteristic distance.

In the model we are considering here, however, no excluded-volume interaction is introduced. As a consequence, there is no microscopic length scale fixing the typical interparticle distance. The local spacing between neighbouring particles is, instead, an emergent quantity that depends on the collective state of the system and may vary significantly in space and time. In particular, dense clusters can form dynamically, reducing the local interparticle distance well below its average value.

Under these conditions, the quantity $\ell_{\Delta t} = v_0\Delta t$ acquires an artificial dynamical significance. If it becomes comparable to the local interparticle spacing, the neighbourhood of a particle may change substantially during a single numerical update. The resulting dynamics may then depend on the arbitrary choice of the integration step rather than solely on the interaction rules encoded in Eq. (6.10).

For this reason, we use the continuous overdamped equations only as a motivation for the structure of the interactions. The final model is formulated as a genuinely discrete Vicsek-like dynamics, in which one update corresponds to an elementary behavioural decision rather than to a small numerical approximation of an underlying continuous trajectory. The novelty of the following construction lies in the implementation of the proactive contribution of Eq. (6.10) within an otherwise standard discrete Vicsek-like dynamics.

6.2.2 Discrete Vicsek-like implementation with reactive and proactive terms

Having identified the active and proactive contributions at the level of the overdamped angular dynamics, we now construct a discrete-time model inspired by that of Vicsek. The goal is to preserve the physical interpretation of the two interaction mechanisms while avoiding the introduction of an arbitrary integration time step.

The reactive contribution, as in the Vicske model, is implemented as a weighted alignment with the headings of neighbouring particles. In particular, for the i-th particle we define the reactive vector

$$\mathbf{R}_i = \sum_{j\neq i} g_r(r_{ij})g_\theta(\phi_i, \vartheta_{ij})\hat{\mathbf{n}}_j. \tag{6.16}$$

in terms of the orientations $\hat{\mathbf{n}}_j$ in Eq. (6.5) of the particles around it. This definition of $\mathbf{R}_i$ generalizes the one entering in the standard Vicsek model, to the case in which each neighbor contributes with a weight determined both by its distance from the i-th particle and by its position within the visual field of the i-th particle.

The implementation of the proactive contribution requires a more careful interpretation. Con-

sider the contribution appearing in Eq. (6.10),

$$\partial_{\phi_i} g_\theta(\phi_i, \vartheta_{ij}) \cos(\phi_i - \phi_j), \tag{6.17}$$

where, for the moment, we omit the radial dependence.

To guide the construction of the discrete update rule, we make a specific assumption about the angular kernel. We assume that $g_\theta(\phi_i, \vartheta_{ij})$ is a differentiable and 2π-periodic function of the relative angle $\phi_i - \vartheta_{ij}$, symmetric about the viewing direction. We further assume that it has a unique maximum when particle i points towards particle j, a unique minimum in the opposite direction, and no other stationary points. These assumptions are satisfied by the smooth kernels considered in Sec. 4.1.3. They imply that

$$\partial_{\phi_i} g_\theta(\phi_i, \vartheta_{ij}) = 0 \tag{6.18}$$

if and only if

$$\phi_i = \vartheta_{ij}, \qquad \text{or} \qquad \phi_i = \vartheta_{ij} + \pi \pmod{2\pi}. \tag{6.19}$$

These two kernel-induced stationary orientations are respectively parallel and antiparallel to the bearing vector connecting particle i to particle j. Which of them is favored by the proactive interaction depends on the sign of the remaining factor $\cos(\phi_i - \phi_j)$. When the two particles i and j point in the same half-plane one has

$$\cos(\phi_i - \phi_j) > 0, \tag{6.20}$$

and thus the proactive interaction favours alignment with the bearing direction $\hat{\mathbf{r}}_{ij}$. Conversely, when the two particles are approximately oppositely oriented, i.e., they point in the two complementary half planes, one has

$$\cos(\phi_i - \phi_j) < 0, \tag{6.21}$$

and the direction promoted by the proactive term becomes $-\hat{\mathbf{r}}_{ij}$.

This construction differs from the standard following interaction of iABPs. In a pure following mechanism, a sufficiently strong following tendency may cause a particle to orient directly towards another agent even when the latter is approaching head-on. In contrast, the present rule distinguishes between neighbours moving in compatible and incompatible directions. Particles approaching from opposite directions do not generate a tendency to look towards each other; rather, they contribute with the opposite sign and therefore favour steering away from the corresponding bearing direction.

Motivated by the previous discussion, we construct the proactive interaction such that it fulfills the following requirements:

1. Unlike the reactive contribution, which is a weighted average of the heading vectors $\hat{\mathbf{n}}_j$ of neighbouring particles, the proactive contribution should depend on their spatial location. We therefore construct it as a weighted average of the bearing vectors $\hat{\mathbf{r}}_{ij}$.

2. The strength of the interaction should increase upon increasing the gradient of the visual kernel. Since the proactive term in Eq. (6.10) is proportional to $\partial_{\phi_i} g_\theta$, the magnitude of the weight is chosen proportional to $|\partial_{\phi_i} g_\theta(\phi_i, \vartheta_{ij})|$.

3. The interaction should be stronger for neighbours moving in a compatible direction. Accord-

ingly, we weight each contribution by $|\cos(\phi_i - \phi_j)|$.

4. The sign of the steering direction is determined by the sign of $\cos(\phi_i - \phi_j)$. If $\cos(\phi_i - \phi_j) > 0$, the preferred direction is $\hat{\mathbf{r}}_{ij}$, whereas if $\cos(\phi_i - \phi_j) < 0$ the preferred direction becomes $-\hat{\mathbf{r}}_{ij}$.

Combining these requirements leads naturally to the definition of the proactive steering vector $\mathbf{P}_i$ as

$$\mathbf{P}_i = \sum_{j \neq i} g_r(r_{ij}) \left|\partial_{\phi_i} g_\theta(\phi_i, \vartheta_{ij})\right| \cos(\phi_i - \phi_j) \hat{\mathbf{r}}_{ij}. \tag{6.22}$$

In fact, $\mathbf{P}_i$ points toward spatial regions occupied by neighbours moving in compatible directions that were on the edge of the field of view of the i-th agent.

In order to combine the effects of the reactive and the proactive terms, we first normalize the corresponding vectors,

$$\hat{\mathbf{R}}_i = \frac{\mathbf{R}_i}{|\mathbf{R}_i|}, \qquad \hat{\mathbf{P}}_i = \frac{\mathbf{P}_i}{|\mathbf{P}_i|}, \tag{6.23}$$

whenever their norm is non-zero. Then, we assume that the effective steering direction $\mathbf{A}_i$ is given by

$$\mathbf{A}_i = \hat{\mathbf{R}}_i + \lambda_p \hat{\mathbf{P}}_i, \tag{6.24}$$

where the parameter λ_p controls the strength of the proactive interaction relative to the reactive one.

The discrete-time dynamics of the orientation is finally given by

$$\phi_i(t+1) = \arg(\mathbf{A}_i) + \sigma \xi_i(t), \tag{6.25}$$

where $\xi_i(t)$ is a Gaussian random variable with zero mean and unit variance. This defines the new orientation vector $\hat{\mathbf{n}}_i(t+1)$ according to Eq. (6.5), i.e., $\hat{\mathbf{n}}_i(t+1) = (\cos\phi_i(t+1), \sin\phi_i(t+1))$. As in the standard Vicsek model, the particle position evolves according to

$$\mathbf{r}_i(t+1) = \mathbf{r}_i(t) + v_0 \hat{\mathbf{n}}_i(t+1). \tag{6.26}$$

In the limit $\lambda_p = 0$, the model reduces to an anisotropic Vicsek dynamics with distance-dependent interactions. Non-zero values of λ_p introduce the additional proactive steering mechanism. The following section is devoted to exploring how this additional interaction modifies the emergent collective behaviour.

6.2.3 Emergent collective behavior

The Vicsek model represents one of the simplest and most successful descriptions of collective motion. Despite its minimal microscopic rules, local alignment combined with self-propulsion is sufficient to produce a transition from a disordered phase to a macroscopically polarized state, together with propagating density structures and other characteristic features of flocking systems [15, 17].

Local alignment alone, however, does not guarantee spatial cohesion. In an unbounded domain, particles that move beyond the interaction range of the group cease to contribute to one another's orientational dynamics and may continue to separate. Consequently, the standard Vicsek model does not generically produce a self-confined flock. Periodic boundary conditions, external

confinement, or additional attractive interactions are therefore commonly introduced to maintain a finite average density and prevent the particles from dispersing indefinitely.

Several extensions of Vicsek-like dynamics address this limitation by introducing steering mechanisms that favor spatial cohesion [87, 88]. In their simplest form, these rules orient a particle towards the center of mass of its local neighborhood.

The proactive contribution introduced in Sec. 4.1.2 and implemented in the model described in Sec. 6.2.2 produces a qualitatively similar cohesive effect, although its microscopic interpretation is different. In particular, the proactive vector $\mathbf{P}_i$ does not simply point towards the region of highest local density. The contribution it receives from a given neighbor depends on three distinct ingredients: the spatial bearing of that neighbor, the compatibility between the two headings, and the angular gradient of the field-of-view kernel.

The last property is particularly important. Since the proactive weight is proportional to $\left|\partial_{\phi_i} g_\theta(\phi_i, \vartheta_{ij})\right|$, the interaction becomes appreciable when a neighbor lies in a region where a small rotation of the agent would substantially change its perceived interaction strength. For a localized field of view, this typically occurs close to its angular boundaries. The proactive rule therefore tends to rotate an agent so as to retain compatible neighbors that would otherwise leave its perception region.

Accordingly, the mechanism should not be interpreted as a conventional central attraction: no explicit positional force pulls two particles towards each other. Cohesion emerges indirectly from the modification of the heading dynamics: by steering towards suitably oriented neighbors located close to the edge of the visual field, particles reduce the probability of losing contact with the group.

Figure 6.4 illustrates the qualitative effect of this interaction. We consider the sinusoidal angular kernel of Eq. (4.7) and we compare configurations obtained in the absence (see panel (a)) and in the presence (panel (b)) of the proactive contribution. In both simulations, the particle positions and orientations are initially taken to be random, and periodic boundary conditions are imposed in both spatial directions. The two simulations are performed with the same values of the microscopic parameters and differ only in the value of the relative proactive strength λ_p.

For $\lambda_p = 0$, the orientational dynamics reduces to the anisotropic Vicsek-like alignment rule defined by Eq. (6.16). As shown in Fig. 6.4(a), the particles can develop locally coherent headings, but this orientational order is not accompanied by spatial self-confinement. Particles progressively leave one another's interaction range, and the population tends to spread over the available domain.

The behavior changes qualitatively when the proactive contribution is activated. In the example shown in Fig. 6.4(b), the particles spontaneously assemble into a cohesive moving structure. The aggregate is not isotropic but displays a markedly elongated shape.

This phenomenology resembles the cohesive behavior generated by density-seeking interactions, but it originates from a different microscopic rule. The proactive dynamics does not steer an agent indiscriminately towards all nearby particles. Neighbors moving in incompatible directions contribute with the opposite sign, while neighbors located in regions where the visual kernel is nearly constant produce only a weak proactive response. The resulting steering is therefore simultaneously selective in position, heading, and visual sensitivity.

The snapshots reported in Fig. 6.4 should be regarded as a preliminary demonstration of the mechanism rather than as a systematic characterization of the model. In particular, we have not yet determined whether the observed cohesive state remains stable in the thermodynamic limit,

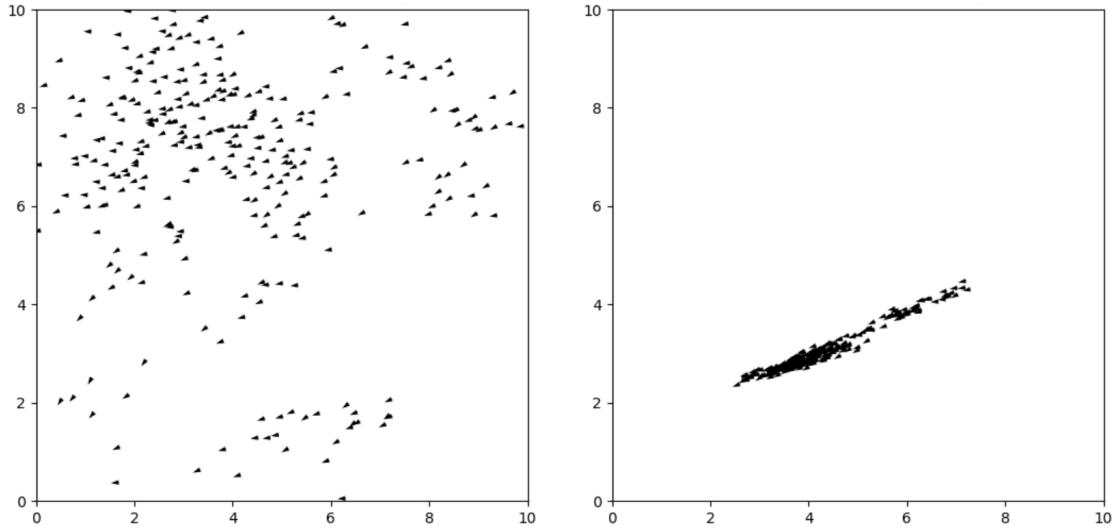


Figure 6.4: Representative snapshots of the discrete off-lattice dynamics introduced in Sec. 6.2.2, using the sinusoidal angular interaction kernel of Eq. (4.7). Each arrow represents an active particle and indicates its instantaneous heading.(a) For $\lambda_p = 0$, the model reduces to an anisotropic Vicsek-like dynamics. Although local polar order develops, no mechanism prevents particles from progressively dispersing throughout the available space. (b) When the proactive interaction is activated, here by setting $\lambda_p = 0.5$, particles tend to retain compatible neighbors close to the boundary of their field of view. The resulting steering suppresses the spatial dispersion of the population and leads to the spontaneous formation of a cohesive, elongated, collectively moving structure. The simulation involves $N = 300$ particles in box with side length $L = 10$, with velocity $v_0 = 0.5$, while the interaction is characterized by $\epsilon = 0.7$, $R_0 = 1.0$, $\ell_r = 1.0$, and the noise by the strength $\sigma = 0.6$.

whether its shape corresponds to a genuine stationary phase, or how it depends on noise, density, propulsion speed, and interaction range. Other collective states may arise in different regions of parameter space.

Nonetheless, these initial observations indicate that the proactive interaction introduced in this thesis may provide a minimal mechanism for combining alignment and cohesion without adding an explicit attractive force. A systematic exploration of the corresponding phase diagram, together with a quantitative characterization of the shape and stability of the emerging structures, constitutes a promising direction for future work.

6.3 Concluding remarks

In this chapter, we presented two promising research directions emerging from the reactive-proactive framework developed in this thesis. First, the switching experiment performed in the lattice NRXY model suggests that the proactive response may suppress the persistence and proliferation of asters, thereby favoring the healing of shock lines and the recovery of long-range orientational order. Although this observation does not yet demosntrate the asymptotic stability of the LRO phase, it motivates a systematic comparison between reactive and reactive-plus-proactive dynamics using differentiable vision kernels.

Second, we extended the notion of proactivity to the case of self-propelled particles in continuous space. The resulting discrete Vicsek-like model combines orientational alignment with a

selective steering mechanism that is enhanced near the boundaries of the visual field. Preliminary simulations indicate that this contribution can generate spatial cohesion and elongated moving aggregates without introducing an explicit attractive force. A characterization of the corresponding phase diagram, the stability of these structures, and their dependence on noise and interaction parameters remains an open problem.

Taken together, these observations suggest that proactivity may control both the long-time fate of ordered states and the emergence of cohesive collective motion, providing a promising direction for future investigations.

Part II

Wavelet sampling for frustrated systems

Chapter 7

Wavelet conditional renormalization group: efficient sampling at criticality

In the first part of this thesis, numerical simulations were primarily used as a tool for investigating the statistical properties of non-reciprocal many-body systems. Accordingly, the focus was on the physical phenomena emerging from the models under consideration, whereas the sampling algorithms themselves remained largely in the background.

The perspective adopted in this chapter is different. Here, the numerical method becomes the object of study. More specifically, we ask whether ideas originating from statistical physics, renormalization-group (RG) theory, and machine learning can be combined in order to construct efficient sampling strategies for critical systems.

As explained in Sec. 1.1.2, the central computational problem when simulating critical systems is critical slowing down: long-wavelength modes become increasingly expensive to decorrelate with local updates, while the construction of cluster for global updates do not extend generally to frustrated systems [97, 98, 103]. The wavelet conditional renormalization group (WCRG) [108] addresses this problem by changing variables. An orthogonal wavelet transform separates a field into a coarse component and localized fluctuations associated with finer scales [107, 143]. Rather than discarding the fluctuations as in an ordinary RG calculation, WCRG learns their conditional distributions and uses them to reconstruct configurations recursively.

This chapter provides a pedagogical introduction to the basic construction. Our original application of the method is presented in Chapter 8, where the WCRG is used to sample a frustrated two-dimensional statistical model for which no cluster algorithm is available. The purpose of the present chapter is, instead, to explain the multiscale factorization which the method is based on, the learning problems that arise at each scale, and the reason why the wavelet representation is able to avoid the critical slowing down affecting direct sampling and direct parameter estimation.

The chapter is organized as follows. In Sec. 7.1, we introduce the orthogonal wavelet decomposition and compare it with a standard Wilsonian RG transformation. We then derive the exact factorization of the microscopic probability distribution into a coarse distribution and a hierarchy of conditional distributions for the wavelet coefficients. In Sec. 7.2, we discuss how these distributions can be learned from data. We first introduce direct coupling estimation as a reference procedure and explain why it becomes inefficient at criticality. We then formulate the corresponding conditional problem in the wavelet basis. Finally, Sec. 7.2.3 introduces conditional

score matching, explains its computational advantage, and discusses the assumptions under which score accuracy controls the learned probability distribution.

7.1 Wavelet decomposition and multiscale factorization

One of the central problems in machine learning is the approximation of high-dimensional probability distributions. In the present context, we may think of a two-dimensional image defined on a square lattice of linear size L, where each pixel carries a real variable $\phi_0(x)$. In statistical physics, such an image could be the one generated by a field configuration of a statistical system. If the system is at equilibrium, its probability distribution $p_0(\phi_0)$ can be written in Boltzmann form,

$$p_0(\phi_0) = \frac{1}{Z_0} e^{-E_0(\phi_0)}, \tag{7.1}$$

where E_0 is the microscopic energy and Z_0 is the partition function.

We focus on models for which the microscopic energy contains a finite number of interaction terms. Even in this apparently favorable situation, learning and sampling the corresponding probability distribution can be difficult. The problem can be split into two steps. Given a dataset consisting of a collection of statistically independent configurations generated by some method, we would like first to determine, i.e., *learn* an accurate representation of p_0 (via a suitable parameterization). Once this representation is available, we would then like to use it in order to generate efficiently new statistically independent configurations.

Near criticality, both tasks become particularly challenging. The diverging correlation length makes the original microscopic variables strongly correlated over increasingly large distances and this leads to the critical slowing of sampling algorithms based on the local update of the microscopic degrees of freedom. The WCRG addresses this difficulty by decomposing the field into variables associated with different spatial scales.

We use the term *scale* to denote the level of resolution: the original field ϕ_0 defines the finest, microscopic resolution and is assumed to be defined on a lattice of linear size $L_0 = L$. Increasing j corresponds to progressively coarser resolutions and, thus, to increasingly large characteristic length scales.

At scale j, an orthogonal wavelet transform $(W_j, \bar{W}_j)$ maps the field ϕ_{j-1} at a finer scale into a coarse field ϕ_j and a set of wavelet coefficients $\bar{\phi}_j$,

$$\phi_j = W_j \phi_{j-1} \quad \text{and} \quad \bar{\phi}_j = \overline{W}_j \phi_{j-1}. \tag{7.2}$$

The field ϕ_j is defined on a lattice of linear size $L_j = L/2^j$ and contains the slowly varying component of ϕ_{j-1}, whereas the variables $\bar{\phi}_j$ contain the shorter-scale fluctuations separated from it at that coarse-graining step.

Along each spatial direction, the one-dimensional wavelet transform applies a low-pass filter, which retains the slowly varying component of the field, and a complementary high-pass filter, which extracts local variations associated predominantly with shorter wavelengths. In two spatial dimensions, a separable transform applies these filters independently along the two lattice directions, producing the four low–low, low–high, high–low, and high–high combinations. The low–low channel defines the coarse field ϕ_j, while the other three channels are collected in $\bar{\phi}_j$. Here, low and high refer to the relative frequency content selected by the filters, rather than to sharply

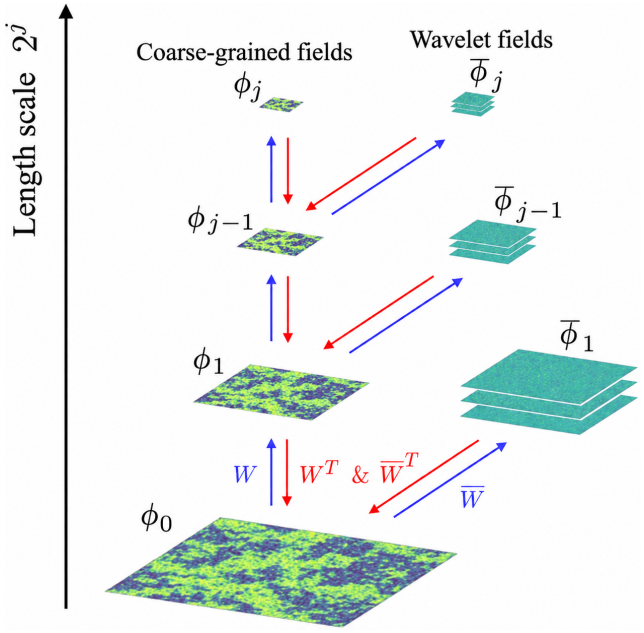


Figure 7.1: Illustration of the multiscale decomposition underlying the WCRG. At each step, the field ϕ_{j-1} is separated into slow degrees of freedom, represented by the coarse field ϕ_j, and fast degrees of freedom, represented by the wavelet coefficients $\bar{\phi}_j$. The procedure can be iterated recursively to obtain a hierarchy of scales. The inverse transform reconstructs the microscopic configuration from the coarsest field and the wavelet variables, providing a natural framework for generative sampling from coarse to fine scales. Adapted from Ref. [108].

separated Fourier bands. Their explicit definitions and normalization conventions are given in App. C.3.

For an orthogonal transform, the operators satisfy

$$W_j^{\mathrm{T}} W_j + \overline{W}_j^{\mathrm{T}} \overline{W}_j = I, \tag{7.3}$$

and the finer field is reconstructed as

$$\phi_{j-1} = W_j^{\mathrm{T}} \phi_j + \overline{W}_j^{\mathrm{T}} \bar{\phi}_j. \tag{7.4}$$

It is useful to collect the low- and high-pass maps into the orthogonal transform $\mathcal{W}_j$ such that

$$\mathcal{W}_j \phi_{j-1} = (\phi_j, \bar{\phi}_j) \quad \text{and} \quad \phi_{j-1} = \mathcal{W}_j^{-1}(\phi_j, \bar{\phi}_j). \tag{7.5}$$

Iterating the decomposition recursively, the microscopic configuration ϕ_0 can be reconstructed from the coarsest field ϕ_J and the collection of wavelet variables $\{\bar{\phi}_j\}_{j=1}^{J}$. At each reconstruction step,

$$(\phi_j, \bar{\phi}_j) \xrightarrow{\mathcal{W}_j^{-1}} \phi_{j-1}, \qquad j = J, J-1, \ldots, 1. \tag{7.6}$$

Starting from ϕ_J, the successive application of these inverse transforms incorporates the wavelet fluctuations at progressively finer scales until the microscopic field ϕ_0 is recovered. The construction is illustrated in Fig. 7.1.

7.1.1 RG and WCRG

The decomposition employed in the WCRG is conceptually similar to a Wilsonian RG transformation [6–8]. In fact, in the latter construction, a field $\phi(k)$ in Fourier space is first decomposed into slow and fast modes, i.e.,

$$\phi(k) = \begin{cases} \phi_<(k) & \text{for} \quad |k| < \Lambda/b, \\ \phi_>(k) & \text{for} \quad \Lambda/b < |k| < \Lambda, \end{cases} \tag{7.7}$$

where Λ is an ultraviolet cutoff and $b > 1$ is the coarse-graining factor. The effective theory at a larger length scale is then obtained by integrating over the fast modes,

$$e^{-E'[\phi_<]} = \int D\phi_> \, e^{-E[\phi_<,\phi_>]}. \tag{7.8}$$

Heuristically, in the WCRG, the coarse fields ϕ_j and the wavelet coefficients $\bar{\phi}_j$ play roles analogous to the slow and fast variables of a Wilsonian construction. There are, however, two important differences.

- The first concerns the role assigned to the fast variables. In a conventional RG transformation, the fast degrees of freedom $\phi_>$ are integrated out in order to define an effective energy for the remaining slow variables $\phi_<$. In the WCRG, this marginalization is only one part of the construction. At each scale j, one also retains the conditional probability distribution of the wavelet variables $\bar{\phi}_j$ at fixed coarse field ϕ_j. These conditional distributions are learned and subsequently sampled. The wavelet coefficients are therefore not merely auxiliary variables to be eliminated: they are the degrees of freedom from which finer-scale configurations are generated.

- The second difference concerns the representation used to separate spatial scales. A Wilsonian transformation based on a sharp momentum-shell decomposition is sharply localized in Fourier space and consequently delocalized in real space. A wavelet decomposition is not tied to this particular choice. Instead, the wavelet family determines the compromise between real-space and frequency localization.

The Haar transform [143] lies close to the real-space-localized end of this spectrum. Its filters have support on only two neighboring sites: the coarse variable is their normalized sum, while the wavelet coefficient is their normalized difference. In this sense, the Haar construction is reminiscent of a Kadanoff block-spin transformation [144], supplemented by the detail variables required to reconstruct the original field exactly. The coarse and wavelet variables, respectively, are obtained from neighboring sites according to

$$\phi_{j+1}(x) = \frac{\phi_j(2x) + \phi_j(2x+1)}{\sqrt{2}} \quad \text{and} \quad \bar{\phi}_{j+1}(x) = \frac{\phi_j(2x) - \phi_j(2x+1)}{\sqrt{2}}. \tag{7.9}$$

The inverse transform reconstructs the finer field through

$$\phi_j(2x) = \frac{\phi_{j+1}(x) + \bar{\phi}_{j+1}(x)}{\sqrt{2}} \quad \text{and} \quad \phi_j(2x+1) = \frac{\phi_{j+1}(x) - \bar{\phi}_{j+1}(x)}{\sqrt{2}}. \tag{7.10}$$

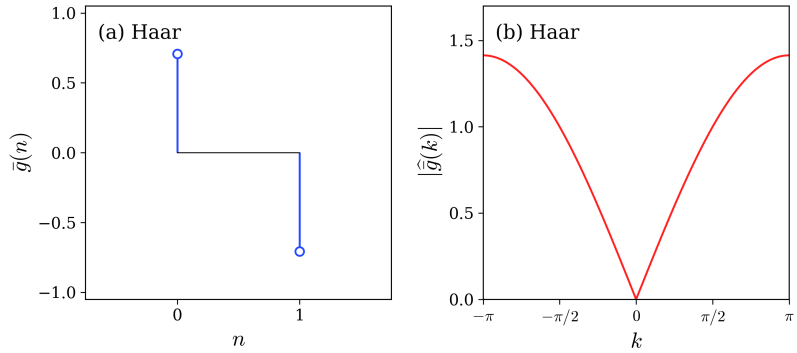


Figure 7.2: Haar high-pass filter in real and Fourier space. (a) The discrete filter $\bar{g}(n)$ as a function of n, which is supported on two neighboring sites. (b) Modulus of the discrete Fourier transform $\hat{\bar{g}}(k)$ of $\bar{g}(n)$, showing that it is non zero over all momentum space rather than being localization within a sharp momentum shell.

These relations are generated by the discrete low-pass and high-pass filters

$$g(n) = \frac{\delta_{n,0} + \delta_{n,1}}{\sqrt{2}} \quad \text{and} \quad \bar{g}(n) = \frac{\delta_{n,0} - \delta_{n,1}}{\sqrt{2}}, \tag{7.11}$$

respectively, in agreement with the general conventions introduced in App. C.3. The coarse variable is thus a normalized local sum, whereas the wavelet coefficient is a normalized local difference. The discrete Fourier transform of the Haar high-pass filter $\bar{g}(n)$ is

$$\hat{\bar{g}}(k) = \frac{1 - e^{-ik}}{\sqrt{2}} \quad \text{and} \quad |\hat{\bar{g}}(k)| = \sqrt{2} \left| \sin\left(\frac{k}{2}\right) \right|. \tag{7.12}$$

As shown in Fig. 7.2, the Fourier response of the high-pass filter vanishes at $k = 0$ and increases towards the edge of the Brillouin zone, corresponding to $k = \pm\pi$. The filter therefore preferentially selects short-wavelength fluctuations. Nevertheless, it is nonzero throughout almost the entire Brillouin zone and it is not concentrated around a single wavevector. Accordingly, although the Haar filter is indeed high-pass, it is only weakly localized in Fourier space and should not be identified with the sharp momentum shell of a Wilsonian transformation.

Higher-order wavelet families, such as Daubechies wavelets [107], provide a different compromise. Their filters extend over a wider real-space region, but their additional vanishing moments suppress low-frequency components more strongly and produce a more selective separation of frequency scales. The application in Chapter 8 uses Haar wavelets for homogeneous structures and Daubechies wavelets with four vanishing moments for modulated phases.

7.1.2 Coarse-grained and conditional energies

We now clarify the relation between the microscopic energy, the effective energies of the coarse fields, and the conditional energies of the wavelet variables. This relation is the central ingredient of the WCRG construction.

The probability distribution $p_j(\phi_j)$ of the field at scale j is written in terms of an effective energy E_j,

$$p_j(\phi_j) = \frac{1}{Z_j} e^{-E_j(\phi_j)}. \tag{7.13}$$

That energy is obtained recursively from the finer-scale distribution, starting from the microscopic energy E_0, by integrating out the wavelet variables,

$$e^{-E_j(\phi_j)} = \int \mathrm{d}\bar{\phi}_j \, e^{-E_{j-1}(\mathcal{W}_j^{-1}(\phi_j,\bar{\phi}_j))}. \tag{7.14}$$

Here and in the following, effective energies are defined up to additive constants, which are absorbed into the corresponding normalizations.

Because the wavelet transform has unit Jacobian, the joint probability density of $(\phi_j, \bar{\phi}_j)$ is simply

$$p_{j-1}\Big(\mathcal{W}_j^{-1}(\phi_j,\bar{\phi}_j)\Big). \tag{7.15}$$

Accordingly, the conditional distribution of the wavelet variables at fixed coarse field is

$$\bar{p}_j(\bar{\phi}_j \mid \phi_j) = \frac{p_{j-1}\Big(\mathcal{W}_j^{-1}(\phi_j,\bar{\phi}_j)\Big)}{p_j(\phi_j)}. \tag{7.16}$$

Equivalently, this conditional probability can be written as

$$\bar{p}_j(\bar{\phi}_j \mid \phi_j) = \exp\left[-\bar{E}_j(\bar{\phi}_j;\phi_j) + \bar{F}_j(\phi_j)\right], \tag{7.17}$$

where the conditional energy $\bar{E}_j$ is defined from the finer-scale effective energy $\bar{E}_{j-1}$ as,

$$\bar{E}_j(\bar{\phi}_j;\phi_j) = E_{j-1}\Big(\mathcal{W}_j^{-1}(\phi_j,\bar{\phi}_j)\Big). \tag{7.18}$$

The corresponding free-energy contribution is fixed by normalization, i.e.,

$$e^{-\bar{F}_j(\phi_j)} = \int \mathrm{d}\bar{\phi}_j \, e^{-\bar{E}_j(\bar{\phi}_j;\phi_j)}. \tag{7.19}$$

The exact microscopic distribution p_0 can be factorized into a product of conditional probabilities associated with the successive wavelet decompositions, i.e.,

$$p_0(\phi_0) = p_J(\phi_J)\prod_{j=1}^{J}\bar{p}_j(\bar{\phi}_j \mid \phi_j). \tag{7.20}$$

Equivalently, using Eqs. (7.13) and (7.17), the microscopic energy can be written in a hierarchical form,

$$E_0(\phi_0) = E_J(\phi_J) + \sum_{j=1}^{J}\left[\bar{E}_j(\bar{\phi}_j;\phi_j) - \bar{F}_j(\phi_j)\right], \tag{7.21}$$

up to an irrelevant additive constant. Note that all the relationships written so far in this Section are actually exact. The idea behind WCRG is that expressing the probability distribution as a cascade of conditional probability distributions allows one to circumvent critical slowing down if the distributions $\bar{p}_j(\bar{\phi}_j \mid \phi_j)$ are easy to sample.

The practical challenge of this approach is to learn the conditional distributions at all scales from the available data. The coarsest energy $E_J(\phi_J)$ is approximated by a parametric energy $E_{\theta_J}(\phi_J)$. Each exact conditional energy in Eq. (7.18) is, instead, approximated in terms of a finite set of possible functions of the field (i.e., operators). If the finer-scale operators are collected in a

vector $\Psi_{j-1}(\phi_{j-1})$, we define

$$\bar{\Psi}_j(\bar{\phi}_j;\phi_j) = \Psi_{j-1}\Big(\mathcal{W}_j^{-1}(\phi_j,\bar{\phi}_j)\Big)\,, \tag{7.22}$$

and parametrize the conditional energy as

$$\bar{E}_{\bar{\theta}_j}(\bar{\phi}_j;\phi_j) = \bar{\theta}_j^{\mathrm{T}}\bar{\Psi}_j(\bar{\phi}_j;\phi_j). \tag{7.23}$$

In practice, the conditional energy is assumed to be a linear combination, with coefficients set by the parameter vector $\bar{\theta}_j$, of the operators which are element of the vector $\Psi_{j-1}(\phi_{j-1})$. The conditional model is fitted independently at every scale. Its parameters $\bar{\theta}_j$ therefore defines the best approximation within the selected operator family.

After learning (the procedure is detailed in Sec. 7.2), the normalized conditional distribution is

$$\bar{p}_{\bar{\theta}_j}(\bar{\phi}_j \mid \phi_j) = \exp\left[-\bar{E}_{\bar{\theta}_j}(\bar{\phi}_j;\phi_j) + \bar{F}^{\mathrm{ind}}_{\bar{\theta}_j}(\phi_j)\right], \tag{7.24}$$

where the conditional free energy induced by the learned energy is

$$\bar{F}^{\mathrm{ind}}_{\bar{\theta}_j}(\phi_j) = -\log\int \mathrm{d}\bar{\phi}_j\, \exp\left[-\bar{E}_{\bar{\theta}_j}(\bar{\phi}_j;\phi_j)\right]. \tag{7.25}$$

Once the learning is done, one can start sampling new configurations, that we call synthetic: they are generated by first sampling the coarsest field from $p_{\theta_J}(\phi_J)$, and then sampling the wavelet variables recursively from Eq. (7.24). The resulting microscopic measure is

$$p_0^{\mathrm{synth}}(\phi_0) = p_{\theta_J}(\phi_J)\prod_{j=1}^{J}\bar{p}_{\bar{\theta}_j}(\bar{\phi}_j \mid \phi_j), \tag{7.26}$$

which corresponds to having the following implicit dimensionless energy:

$$E_0^{\mathrm{synth}}(\phi_0) = E_{\theta_J}(\phi_J) + \sum_{j=1}^{J}\left[\bar{E}_{\bar{\theta}_j}(\bar{\phi}_j;\phi_j) - \bar{F}^{\mathrm{ind}}_{\bar{\theta}_j}(\phi_j)\right]. \tag{7.27}$$

In general, this energy differs from the original microscopic energy E_0, because the conditional energies are represented within a finite ansatz class and the parameters are learned on the basis of a finite set of configurations. The distinction between the original energy in Eq. (7.21) and the implicit synthetic energy in Eq. (7.27) will be important in Sec. 8.3.3.

Note that the induced free energies in Eq. (7.25) are not actually necessary for doing conditional sampling. If one wishes to construct an explicit microscopic energy functional associated with the learned WCRG hierarchy, they can be approximated by the fitted functions $\bar{F}_{\alpha_j}(\phi_j)$; see App. C.4.4 for details. This gives the explicit WCRG energy

$$E_0^{\mathrm{WCRG}}(\phi_0) = E_{\theta_J}(\phi_J) + \sum_{j=1}^{J}\left[\bar{E}_{\bar{\theta}_j}(\bar{\phi}_j;\phi_j) - \bar{F}_{\alpha_j}(\phi_j)\right]. \tag{7.28}$$

7.2 Learning the multiscale model

We now discuss how the probability distributions introduced above can be learned from a dataset and why the wavelet representation is useful near criticality.

Critical slowing down is not only a problem for the final generation of statistically independent configurations. It also affects the learning process. When the target distribution describes a system close to a continuous phase transition, its long-range correlations make the inference of the corresponding probability distribution increasingly difficult as the system size grows.

Below we discuss three related learning strategies. First, we consider direct coupling estimation, in which the distribution of a coarse field is learned by maximum likelihood. This provides a conceptual reference, although this approach is practically hindered by the fact that, during learning, one has to sample from the parametrized distribution, that is affected by critical slowing down if the systems is critical. Second, we use the same maximum likelihood method for the conditional distribution of the wavelet variables. Fixing the slow coarse variables, we are left with wavelet variables with short-range correlations that are fast to sample. Finally, we introduce conditional score matching, which does not learn the full conditional probability but only the derivatives of the conditional log probabilities and is constructed in a way such to avoid the need of Monte Carlo sampling the parametrized model during the training procedure.

7.2.1 Direct coupling estimation by maximum likelihood

Before introducing the learning strategy based on the conditional procedure, for the purpose of illustration it is useful to consider the conceptually simpler case in which the probability distribution of a coarse field is learned directly.

Suppose that the target distribution at scale j is

$$p_j(\phi_j) = \frac{1}{Z_j} e^{-E_j(\phi_j)}. \tag{7.29}$$

The exact energy generated by coarse graining the field generally contains an infinite set of operators which are compatible with the symmetries of the model. We restrict this set to contain only quadratic interaction kernels and a local potential, i.e.,

$$E_j(\phi_j) \simeq \frac{1}{2} \sum_{x,y \in \Lambda_j} K_j(x-y)\, \phi_j(x) \phi_j(y) + \sum_{x \in \Lambda_j} V_j(\phi_j(x)), \tag{7.30}$$

where $\Lambda_j \subset \mathbb{Z}^2$ denotes the lattice at scale j. The ansatz is not expected to reproduce every operator generated by coarse graining: neglecting them can produce systematic errors as demonstrated in, c.f., Sec. 8.3.3.

The local potential in Eq. (7.30) is, in turn, expanded in a finite set of functions which form a sort fo a basis for the expansion, i.e.,

$$V_j(\phi) = \sum_{\alpha} \lambda_{j,\alpha} s_\alpha(\phi), \tag{7.31}$$

where s_α may be chosen, for example, as hat or sigmoid functions (See Eq. C.32 and the relative Appendix). Collecting into the vector $\Psi_j(\phi_j)$ all the operators which appear in Eq. (7.30), the

complete energy can be written as

$$E_{\theta_j}(\phi_j) = \theta_j^{\mathrm{T}} \Psi_j(\phi_j). \tag{7.32}$$

This model is linear in the parameters θ_j, although the components of Ψ_j may be nonlinear functions of the field. The operator basis contains quadratic terms,

$$\Psi^{(2)}_{j,r}(\phi_j) = \sum_{x \in \Lambda_j} \phi_j(x)\phi_j(x+r), \tag{7.33}$$

for the finite set of lattice displacements r kept in the ansatz, and local terms,

$$\Psi^{(1)}_{j,\alpha}(\phi_j) = \sum_{x \in \Lambda_j} s_\alpha\big(\phi_j(x)\big)\,. \tag{7.34}$$

The corresponding model distribution is

$$p_{\theta_j}(\phi_j) = \frac{1}{Z_{\theta_j}} e^{-E_{\theta_j}(\phi_j)}. \tag{7.35}$$

In order the determine the values of the parameter θ_j in Eq. (7.32) such that $p_{\theta_j}(\phi_j)$ is the best estimate of the (unknown) target distribution $p_j(\phi_j)$ in Eq. (7.29), we minimize the Kullback–Leibler (KL) divergence

$$D_{\mathrm{KL}}(p_j \| p_{\theta_j}) = \int \mathrm{d}\phi_j \, p_j(\phi_j) \log \frac{p_j(\phi_j)}{p_{\theta_j}(\phi_j)}. \tag{7.36}$$

Minimizing this divergence is equivalent to maximizing the likelihood of the training data under the model in Eq. (7.35) [145]. Up to a contribution which is independent of θ_j, the KL divergence can be written as

$$D_{\mathrm{KL}}(p_j \| p_{\theta_j}) = \Big\langle E_{\theta_j} \Big\rangle_{p_j} + \log Z_{\theta_j} + \text{const.} \tag{7.37}$$

The derivative of the partition function is

$$\frac{\partial}{\partial \theta_{j,l}} \log Z_{\theta_j} = \frac{1}{Z_{\theta_j}} \frac{\partial Z_{\theta_j}}{\partial \theta_{j,l}} = -\Big\langle \Psi_{j,l} \Big\rangle_{p_{\theta_j}}. \tag{7.38}$$

Consequently, deriving $D_{\mathrm{KL}}(p_j \| p_{\theta_j})$ in Eq. (7.37) in order to determine its minima, and using Eq. (7.38), one finds

$$\frac{\partial}{\partial \theta_{j,l}} D_{\mathrm{KL}}(p_j \| p_{\theta_j}) = \Big\langle \Psi_{j,l} \Big\rangle_{p_j} - \Big\langle \Psi_{j,l} \Big\rangle_{p_{\theta_j}}. \tag{7.39}$$

The optimum is therefore reached when the operator averages predicted by the model coincide with those measured in the data. In order to determine the value of the parameters which satisfy this condition, one can conveniently perform a gradient descent, iteratively updating the values of the parameters as

$$\theta^{(t+1)}_{j,l} = \theta^{(t)}_{j,l} - \varepsilon_{\mathrm{lr}} \frac{\partial D_{\mathrm{KL}}}{\partial \theta_{j,l}}, \tag{7.40}$$

where the derivative of the KL divergence is calculated as in Eq. (7.39), while $\varepsilon_{\mathrm{lr}}$ is the learning rate.

Although this procedure is conceptually simple, it becomes inefficient near criticality for the two different reasons described below.

Sampling the model distribution. The first expectation value in Eq. (7.39) can be estimated directly from the training configurations, since these are distributed according to the target p_j. The second expectation value is, instead, taken over the current model distribution $p_{\theta_j^{(t)}}$. A Monte Carlo simulation of this model is therefore required at every optimization step. If the current model is close to a critical distribution, this auxiliary simulation suffers from critical slowing down. The evaluation of the learning gradient then becomes increasingly expensive as the correlation length and the system size increase.

Conditioning of the optimization problem. A second source of slowing down is associated with the shape of the optimization landscape. Close to an optimum, the convergence of the gradient descent is controlled by the Hessian

$$H_{j,lm} = \frac{\partial^2 D_{\mathrm{KL}}}{\partial\theta_{j,l}\partial\theta_{j,m}}. \tag{7.41}$$

For the family models of Eq. 7.32 considered here, the Hessian is the covariance matrix of the operators, i.e.,

$$H_{j,lm} = \left\langle \Psi_{j,l}\Psi_{j,m} \right\rangle_{p_{\theta_j}} - \left\langle \Psi_{j,l} \right\rangle_{p_{\theta_j}} \left\langle \Psi_{j,m} \right\rangle_{p_{\theta_j}}. \tag{7.42}$$

Let $\lambda_{\min}$ and $\lambda_{\max}$ denote the smallest and largest eigenvalues of the Hessian. Within the quadratic approximation around the optimum, the gradient descent in Eq. (7.40) is stable for

$$0 < \varepsilon_{\mathrm{lr}} < \frac{2}{\lambda_{\max}}, \tag{7.43}$$

while its convergence rate is controlled by the spectrum of the Hessian and, in particular, by the ratio

$$\mathcal{C}_j = \frac{\lambda_{\max}}{\lambda_{\min}} \tag{7.44}$$

which is the *condition number* of the problem [146, 147]. When $\mathcal{C}_j$ is large, the objective function (D_{KL} in our case) is very steep in some parameter directions and very flat in others. A learning rate $\varepsilon_{\mathrm{lr}}$ which is small enough to remain stable along the steep directions produces a slow evolution along the flat ones, and the optimization problem is said to be *poorly conditioned*.

A simple illustration of the increase of the condition number is provided by the two-point function of the critical field. In fact, at criticality,

$$G(r) = \langle \phi_j(x)\phi_j(x+r) \rangle \sim \frac{1}{r^{d-2+\eta}}, \tag{7.45}$$

where η is the anomalous dimension. In Fourier space,

$$G(k) \sim k^{-2+\eta}. \tag{7.46}$$

The longest-wavelength mode in a finite system has $k_{\min} \sim 2\pi/L_j$, whereas the largest lattice momentum is of order one. The condition number of the field covariance therefore scales as

$$\mathcal{C}_{\phi_j} \sim L_j^{2-\eta}. \tag{7.47}$$

Accordingly, at criticality, the problem is increasingly ill-conditioned as the system size increases.

Note, however, that the Hessian in Eq. (7.42) is not the field covariance itself, but the covariance of the operators used in the ansatz. Nevertheless, these operators generically inherit the long-range fluctuations from the critical field, possibly with a different value of the corresponding exponent η. The argument above is therefore to be understood as heuristic rather than as a rigorous derivation of the Hessian spectrum.

Direct estimation thus suffers from two related problems: the model distribution is expensive to sample, and the parameter optimization becomes increasingly poorly conditioned. The WCRG addresses both problems by learning conditional distributions of the wavelet variables instead.

7.2.2 Conditional coupling estimation by maximum likelihood

At scale j, the target joint distribution of the coarse and wavelet variables is

$$p_{j-1}(\bar{\phi}_j, \phi_j) = p_j(\phi_j)\bar{p}_j(\bar{\phi}_j \mid \phi_j). \tag{7.48}$$

The coarse-field distribution p_j is fixed by the data. The learning problem consists in approximating the conditional distribution $\bar{p}_j$ by the conditional model in Eq. (7.24). The corresponding model joint distribution is

$$p_{\bar{\theta}_j}(\bar{\phi}_j, \phi_j) = p_j(\phi_j)\bar{p}_{\bar{\theta}_j}(\bar{\phi}_j \mid \phi_j). \tag{7.49}$$

Because the target and the model distributions share the same coarse marginal, their KL divergence can be written as an average of conditional divergence, i.e.,

$$D_{\mathrm{KL}}(p_{j-1} \| p_{\bar{\theta}_j}) = \int \mathrm{d}\phi_j \, p_j(\phi_j) D_{\mathrm{KL}}(\bar{p}_j(\bar{\phi}_j \mid \phi_j) \| \bar{p}_{\bar{\theta}_j}(\bar{\phi}_j \mid \phi_j)). \tag{7.50}$$

The learning procedure therefore adjusts the conditional model for coarse configurations distributed as in the data.

The conditional operator vector and the energy ansatz which is linear in the parameters were defined in Eqs. (7.22) and (7.23). Proceeding as in the direct problem, the gradient of the joint KL divergence in Eq. (7.50) turns out to be

$$\frac{\partial D_{\mathrm{KL}}}{\partial \bar{\theta}_{j,l}} = \left\langle \Psi_{j,l} \right\rangle_{p_{j-1}} - \left\langle \Psi_{j,l} \right\rangle_{p_{\bar{\theta}_j}}. \tag{7.51}$$

In turn, this can be used in the gradient descent procedure analogous to Eq. (7.40), which now takes the form

$$\bar{\theta}_{j,l}^{(t+1)} = \bar{\theta}_{j,l}^{(t)} - \varepsilon_{\mathrm{lr}} \frac{\partial D_{\mathrm{KL}}}{\partial \bar{\theta}_{j,l}}. \tag{7.52}$$

Let the training dataset contain N_d decomposed configurations,

$$\left\{ \bar{\phi}_{j,n}^{(d)}, \phi_{j,n}^{(d)} \right\}_{n=1}^{N_d}, \tag{7.53}$$

where the superscript $^{(d)}$ indicates configurations that belong to the traninig dataset. The first expectation value in Eq. (7.51) is estimated by

$$\widehat{\left\langle \bar{\Psi}_{j,l} \right\rangle}_{\mathrm{data}} = \frac{1}{N_d} \sum_{n=1}^{N_d} \bar{\Psi}_{j,l} \left(\bar{\phi}_{j,n}^{(d)}; \phi_{j,n}^{(d)} \right). \tag{7.54}$$

The second, instead, must be evaluated under the current conditional model. At each optimization step, the coarse fields are drawn from the dataset and kept fixed. For each $\phi_{j,n}^{(d)}$, the wavelet variables are sampled via a monte Carlo simulation according to

$$\bar{\phi}_{j,n}^{(\mathrm{MC})} \sim \bar{p}_{\bar{\theta}_j^{(t)}}\left(\bar{\phi}_j \mid \phi_{j,n}^{(d)}\right). \tag{7.55}$$

The model average is then estimated as

$$\widehat{\left\langle \bar{\Psi}_{j,l} \right\rangle}_{\mathrm{model}} = \frac{1}{N_d} \sum_{n=1}^{N_d} \bar{\Psi}_{j,l}\left(\bar{\phi}_{j,n}^{(\mathrm{MC})}; \phi_{j,n}^{(d)}\right). \tag{7.56}$$

At first sight, this procedure appears similar to direct coupling estimation because it still requires a Monte Carlo simulation at every learning step. The crucial difference, however, is that the Monte Carlo dynamics acts only on the wavelet coefficients, while the slowly varying coarse field is held fixed. This changes both the sampling problem and the conditioning of the optimization, as explained below.

Conditional sampling. Sampling the full field close to criticality requires the relaxation of long-wavelength modes whose characteristic relaxation times increase with system size. In the conditional problem, these modes are fixed by ϕ_j. The Monte Carlo dynamics explores only the fluctuations encoded by $\bar{\phi}_j$. In the critical models investigated in Ref. [108], the thermalization and autocorrelation times of this conditional dynamics were found to be independent of the total system size.Accordingly, the conditional formulation strongly reduces the sampling bottleneck affecting direct estimation.

Conditional optimization. For the case of the conditional coupling estimation, the Hessian of the conditional KL divergence is again a covariance matrix, given by

$$\bar{H}_{j,lm} = \left\langle \left\langle \bar{\Psi}_{j,l}\bar{\Psi}_{j,m} \right\rangle_{\bar{p}_{\bar{\theta}_j}(\bar{\phi}_j|\phi_j)} - \left\langle \bar{\Psi}_{j,l} \right\rangle_{\bar{p}_{\bar{\theta}_j}(\bar{\phi}_j|\phi_j)} \left\langle \bar{\Psi}_{j,m} \right\rangle_{\bar{p}_{\bar{\theta}_j}(\bar{\phi}_j|\phi_j)} \right\rangle_{p_j(\phi_j)}. \tag{7.57}$$

Accordingly, for each fixed coarse field ϕ_j, the inner averages measure correlations generated by the fluctuations of the wavelet coefficients $\bar{\phi}_j$. The resulting conditional covariance is then averaged over coarse fields sampled from $p_j(\phi_j)$. This is the main difference from the direct coupling estimation discussed in Section 7.2.1, where the Hessian directly probes correlations of the operators under the full coarse-field distribution.

A central result of Ref. [108] is that the covariance eigenvalues of the wavelet fields remain bounded above and below independently of the coarse-graining scale. Schematically, if the original field has long-wavelength covariance

$$\lambda_\phi(k) \sim k^{-\zeta}, \tag{7.58}$$

then the corresponding wavelet covariance satisfies

$$A \leq \lambda_{\bar{\phi}_j} \leq B, \tag{7.59}$$

where the constants A and B do not increase upon increasing the system size. This gives a precise meaning to the statement that the wavelet variables have short-range correlations after the slow

field has been fixed.

The Hessian in Eq. (7.57) is the covariance of the conditional operators rather than that of the wavelet field itself. Using a basis of local operators, it is natural to expect that the improved conditioning of the wavelet variables also prevents a strong broadening of the operator covariance spectrum. Although this final step is not supported by a rigorous theorem, it is consistent with the numerical learning dynamics reported in Ref. [108]. The available results therefore show that both the auxiliary sampling time and the optimization time are strongly reduced in the conditional formulation, with no observable critical slowing down in the systems studied there.

7.2.3 Conditional coupling estimation by score matching

Conditional maximum-likelihood estimation avoids the slow modes of the original field, but it still requires sampling from the current conditional model at every gradient-descent iteration. Score matching provides a further simplification by replacing the comparison of probability densities with a comparison of their derivatives [148].

For a conditional probability $p_j(\bar{\phi}_j \mid \phi_j)$, the score $s_j(\bar{\phi}_j \mid \phi_j)$ is the gradient of its logarithm with respect to the wavelet variables, i.e.,

$$s_j(\bar{\phi}_j \mid \phi_j) = \nabla_{\bar{\phi}_j} \log \bar{p}_j(\bar{\phi}_j \mid \phi_j), \tag{7.60}$$

and similarly for the learned conditional distribution , i.e.,

$$s_{\bar{\theta}_j}(\bar{\phi}_j \mid \phi_j) = \nabla_{\bar{\phi}_j} \log \bar{p}_{\bar{\theta}_j}(\bar{\phi}_j \mid \phi_j). \tag{7.61}$$

For the exponential model of Eq. (7.24),

$$s_{\bar{\theta}_j} = -\nabla_{\bar{\phi}_j} \bar{E}_{\bar{\theta}_j}(\bar{\phi}_j; \phi_j), \tag{7.62}$$

because the conditional free energy depends only on ϕ_j and thus it drops out upon differentiation with respect to $\bar{\phi}_j$.

The discrepancy between target and model scores is measured by the conditional Fisher divergence,

$$D_F(p_{j-1} \| p_{\bar{\theta}_j}) = \frac{1}{2} \int \mathrm{d}\phi_j \, p_j(\phi_j) \int \mathrm{d}\bar{\phi}_j \, \bar{p}_j(\bar{\phi}_j \mid \phi_j) \left\| s_j(\bar{\phi}_j \mid \phi_j) - s_{\bar{\theta}_j}(\bar{\phi}_j \mid \phi_j) \right\|^2. \tag{7.63}$$

The target score is not known explicitly. The essential observation of score matching is that integration by parts removes it from the parameter-dependent part of the Fisher divergence. One obtains

$$D_F(p_{j-1} \| p_{\bar{\theta}_j}) = \text{const} + \left\langle \frac{1}{2} \left\| \nabla_{\bar{\phi}_j} \bar{E}_{\bar{\theta}_j} \right\|^2 - \Delta_{\bar{\phi}_j} \bar{E}_{\bar{\theta}_j} \right\rangle_{p_{j-1}}, \tag{7.64}$$

where $\Delta_{\bar{\phi}_j}$ is the Laplacian with respect to the wavelet variables. We then replace the expectation value over the target distribution with an average over the training dataset of Eq. (7.53): we thus

obtain the empirical Fisher divergence

$$D_F^{\text{emp}} = \frac{1}{N_d} \sum_{n=1}^{N_d} \left[\frac{1}{2} \left\| \nabla_{\bar{\phi}_j} \bar{E}_{\bar{\theta}_j} \left(\bar{\phi}_{j,n}^{(d)}; \phi_{j,n}^{(d)} \right) \right\|^2 - \Delta_{\bar{\phi}_j} \bar{E}_{\bar{\theta}_j} \left(\bar{\phi}_{j,n}^{(d)}; \phi_{j,n}^{(d)} \right) \right]. \tag{7.65}$$

For the linear conditional ansatz of Eq. (7.23), the empirical loss D_F^{emp} is a quadratic function of the parameters, i.e.,

$$D_F^{\text{emp}} = \frac{1}{2} \bar{\theta}_j^{\mathrm{T}} \bar{M}_j \bar{\theta}_j - \bar{b}_j^{\mathrm{T}} \bar{\theta}_j + \text{const.}, \tag{7.66}$$

where

$$(\bar{M}_j)_{lm} = \frac{1}{N_d} \sum_{n=1}^{N_d} \nabla_{\bar{\phi}_j} \bar{\Psi}_{j,l} \cdot \nabla_{\bar{\phi}_j} \bar{\Psi}_{j,m}, \tag{7.67}$$

$$(\bar{b}_j)_l = \frac{1}{N_d} \sum_{n=1}^{N_d} \Delta_{\bar{\phi}_j} \bar{\Psi}_{j,l}. \tag{7.68}$$

All operator components and derivatives in these equations are evaluated on the n-th training configuration. If $\bar{M}_j$ is invertible, the optimal values $\bar{\theta}_j^\star$ of the parameters $\bar{\theta}_j$ is obtained directly as

$$\bar{\theta}_j^\star = \bar{M}_j^{-1} \bar{b}_j. \tag{7.69}$$

If the operator basis is redundant the matrix $\bar{M}_j$ might be non-invertible: in that case one may use a pseudoinverse or a weak regularization. In either case, the important point is that the parameters are determined directly from the data without the need of running an iterative Monte Carlo simulation of the conditional model. More details on the practical implementation of score matching used for the model studied in Chapter 8 are given in App. C.4.

From similar scores to similar distributions

Matching the score is formally a weaker condition than matching the probability densities themselves. Accordingly, it is necessary to understand under which conditions a small Fisher divergence implies a small KL divergence.

Let $x = (x_1, \ldots, x_d) \in \mathbb{R}^d$, and consider a probability density

$$q(x) \propto e^{-E(x)}. \tag{7.70}$$

The Hessian of the energy E is the $d \times d$ symmetric matrix whose entries are

$$\left[\nabla_x^2 E(x) \right]_{ab} = \frac{\partial^2 E(x)}{\partial x_a \, \partial x_b}. \tag{7.71}$$

The density q is said to be α-strongly log-concave if there exists a constant $\alpha > 0$ such that

$$\alpha I_d \preceq \nabla_x^2 E(x) = -\nabla_x^2 \log q(x), \qquad \forall x \in \mathbb{R}^d, \tag{7.72}$$

where I_d is the d-dimensional identity matrix and $\preceq$ denotes the Löwner order on symmetric matrices. Thus, $A \preceq B$ means that $B - A$ is positive semidefinite. If the model conditional distribution is uniformly α-strongly log-concave for every coarse configuration, a logarithmic

Sobolev inequality implies that the conditional KL divergence is bounded from above by the Fisher divergence [116]. With the normalization used in Eq. (7.63), this means

$$D_{\mathrm{KL}}(p_{j-1}\|p_{\bar{\theta}_j}) \leq \frac{1}{\alpha} D_F(p_{j-1}\|p_{\bar{\theta}_j}). \tag{7.73}$$

Thus, under the condition of a uniform convexity, a small score-matching loss D_F guarantees that the learned conditional distribution is also close to the target in terms of the KL divergence. The use of conditional strong log-concavity in multiscale generative models is developed in Ref. [115].

To understand why conditioning on the coarse field can produce a convex problem, consider an effective energy at scale $j-1$ of the form

$$E_{j-1}(\phi_{j-1}) = \frac{1}{2}\phi_{j-1}^{\mathrm{T}} K_{j-1}\phi_{j-1} + \sum_x V_{j-1}(\phi_{j-1}(x)). \tag{7.74}$$

Its Hessian is

$$\nabla^2_{\phi} E_{j-1}(\phi_{j-1}) = K_{j-1} + \mathrm{diag}\left[V''_{j-1}(\phi_{j-1}(x))\right]. \tag{7.75}$$

Suppose that the negative curvature of the local potential is bounded from below by a constant $-\gamma < 0$, i.e., that

$$-\gamma \leq V''_{j-1}(t), \qquad \forall t \in \mathbb{R}. \tag{7.76}$$

The potential may therefore have a non-convex region, as in a double-well potential, but its negative curvature is bounded.

We now want to understand if we can obtain, from a non-convex energy E_{j-1}, a conditional energy uniformly convex with respect to the wavelet variables. The Hessian of the conditional energy is obtained by taking Eq. (7.75) and projecting it on the wavelet space, as

$$\nabla^2_{\bar{\phi}_j} \bar{E}_j = \overline{W}_j \left[K_{j-1} + \mathrm{diag}\left(V''_{j-1}(\phi)\right)\right] \overline{W}_j^{\mathrm{T}}. \tag{7.77}$$

Let $\mu_j^{\min}$ denote the smallest eigenvalue of the projected quadratic kernel $\overline{W}_j K_{j-1} \overline{W}_j^{\mathrm{T}}$. Equation (7.76) then gives

$$(\mu_j^{\min} - \gamma) I \preceq \nabla^2_{\bar{\phi}_j} \bar{E}_j. \tag{7.78}$$

A sufficient condition for conditional strong log-concavity is therefore

$$\mu_j^{\min} > \gamma. \tag{7.79}$$

The physical meaning of this condition is simple. The local potential can produce a negative curvature, but the short-wavelength modes are stiffer than the long-wavelength ones. If the wavelet variables isolate a sufficiently stiff range of fluctuations, the positive curvature of the quadratic term dominates the non-convexity of the local potential.

In order to illustrate this fact in a simple case, let us assume that the quadratic kernel in Eq. (7.74) has the continuum-like spectrum

$$\widehat{K}_{j-1}(k) \simeq c_K |k|^{\nu} \quad \text{with} \quad \nu > 0. \tag{7.80}$$

Accordingly, a projector $\overline{W}_j$ whose Fourier support contains only $|k| \geq k_0$ has

$$\mu_j^{\min} \geq c_K k_0^{\nu}. \tag{7.81}$$

The sufficient condition in Eq. (7.79) then becomes

$$k_0 > \left(\frac{\gamma}{c_K}\right)^{1/\nu}. \tag{7.82}$$

This argument explains why separating out the long-wavelength components can turn a non-convex full distribution into a well-behaved conditional distribution for the remaining fluctuations.

Computational efficiency and statistical reliability

The previous argument provides a sufficient condition rather than a statement that applies automatically to every wavelet transform. In particular, the idealized frequency-space argument assumes a projector whose Fourier support is sharply restricted to a sufficiently high-frequency sector.

Wavelets with compact support in real space, such as Haar wavelets, are not strictly band limited in frequency, as illustrated in Fig. 7.2. The Haar high-pass filter vanishes at zero frequency and increases towards the edge of the Brillouin zone. It therefore suppresses slowly varying components and assigns greater weight to short-wavelength fluctuations. Nevertheless, its Fourier response is broad: low-frequency components are suppressed rather than removed completely.

Consequently, a Haar wavelet subspace does not satisfy exactly the strict support assumption used to derive Eq. (7.82). That equation therefore cannot be invoked as a rigorous guarantee of conditional strong log-concavity for a generic Haar decomposition. It nevertheless identifies the mechanism by which a wavelet decomposition can improve the conditions for score matching. When the quadratic stiffness increases with the wavevector, suppressing the soft long-wavelength components and assigning more weight to the stiffer short-wavelength ones tends to increase the curvature of the energy within the wavelet subspace. Conditioning on the coarse field therefore leaves a set of fluctuations whose conditional energy is more likely to be convex than the energy of the complete field.

An important distinction must be made at this point. The computational advantage of score matching does not depend on the log-concavity. For the linear energy parameterization considered here, the empirical score-matching loss is quadratic in the parameters and its minimum is obtained by solving a single linear system. Accordingly,the parameter estimation is performed in one step, without the iterative alternation between parameter updates and Monte Carlo sampling required by maximum-likelihood estimation. This remains true whether or not the distribution being learned is log-concave.

Log-concavity determines, instead, how the score-matching objective should be interpreted. Score matching controls directly the difference between the gradients of the target and parametrized log probabilities. For a strongly log-concave conditional distribution, matching accurately these scores provides also a quantitative control over the difference between the corresponding probability distributions through Eq. (7.73).

This distinction motivates the practical strategy adopted in the following chapter. At each scale, we use score matching to learn the conditional distribution of the wavelet coefficients through a

single parameter-estimation step. The performance of the method must nevertheless be assessed numerically, since the Haar decomposition satisfies the ideal frequency-localization assumptions only approximately. A second approximation arises from truncating the exact conditional energies to a finite operator vector $\bar{\Psi}_j$. The next chapter therefore tests separately whether the learned conditional chains avoid critical slowing down and whether their recursive composition reproduces the relevant statistics of the target distribution.

7.3 Concluding remarks

In this chapter, we have introduced the main concepts underlying the Wavelet Conditional Renormalization Group. The central idea is to replace the direct modeling of a critical field, whose long-wavelength modes are strongly correlated, with a hierarchy of conditional models for fluctuations associated with subsequent spatial scales.

An orthogonal wavelet transform decomposes each field ϕ_{j-1} into a coarse component ϕ_j and a set of wavelet coefficients $\bar{\phi}_j$. Iterating this decomposition leads to the exact probability factorization of Eq. 7.20.

This identity provides both a multiscale representation of the statistical measure and a generative procedure. A microscopic configuration can be constructed recursively by first sampling the coarsest field and then adding the wavelet fluctuations scale by scale.

We then compared three strategies for learning the distributions which appears in this factorization. Direct coupling estimation based on maximum-likelihood fits the coarse-field energy but requires repeated sampling of a potentially critical model and thus it inherits the poor conditioning produced by long-range fluctuations. Conditional coupling estimation fixes the slow coarse field and samples only the wavelet variables. In the critical models studied in previous work, this substantially reduces the learning time.

Conditional score matching takes a further step by eliminating the need to sample the parametrized distribution during training. Integration by parts rewrites the score-matching loss entirely in terms of derivatives of the model energy evaluated on the training configurations. For an energy linear in its parameters, the resulting optimization is a quadratic problem that can be solved by linear algebra.

The theoretical justification for this procedure relies on the conditional distributions being sufficiently regular. If the quadratic stiffness of the wavelet modes dominates the possible negative curvature of the local potential, the conditional energy is strongly convex and a small Fisher divergence guarantees a small KL divergence. Ideal frequency-localized projectors make this mechanism transparent. Practical compactly supported wavelets do not necessarily satisfy the same assumptions exactly, and their performance must therefore be checked numerically.

The discussion developed here establishes the reference framework for Chapter 8. There, the WCRG will be applied to a frustrated lattice model in order to address two separate issues: whether the conditional sampling procedure avoids the critical slowing down which affects conventional algorithms, and whether the finite-dimensional conditional-energy ansatz reproduces the complete statistical measure with sufficient accuracy.

Chapter 8

Wavelet conditional sampling of frustrated systems

Building on the WCRG framework developed in the previous chapter, we now test its ability to learn and sample a frustrated system close to criticality. Such systems provide a demanding benchmark: cluster methods strongly reduce critical slowing down in simple unfrustrated spin models [97], but comparably effective constructions are not generally available in the presence of competing interactions, even when the frustration is weak [103–105]. WCRG approaches the problem through a multiscale representation rather than by constructing correlated microscopic clusters.

We consider soft-spin ϕ^4 models with competing interactions on a square lattice. Their coupling stencil encompasses the axial next-nearest-neighbor Ising (ANNNI), biaxial next-nearest-neighbor Ising (BNNNI), and third-nearest-neighbor Ising (3NNNI) geometries. Our numerical analysis focuses on the soft-spin BNNNI (sBNNNI) model, whose phase diagram contains homogeneous ferromagnetic and paramagnetic phases separated by an Ising-like critical line, together with modulated disordered and ordered phases.

The WCRG construction proceeds in two distinct stages. During learning, the conditional energy of the wavelet coefficients is inferred at each scale by score matching, using equilibrium Monte Carlo configurations as training data. As discussed in Sec. 7.2.3, score matching fits an unnormalized probability model by comparing gradients of log probabilities. For the linear-in-parameter ansatz used here, the optimization reduces to a single regularized linear system and does not require sampling (with Monte Carlo) the parametrized distribution. During generation, by contrast, the coarsest field is sampled first and the microscopic configuration is reconstructed recursively through conditional Monte Carlo sampling of the learned wavelet distributions.

The analysis leads to three main results. First, the generated configurations reproduce the local field distributions and the two-point correlations, as measured by the structure factor, across representative regions of the sBNNNI phase diagram. The method captures both homogeneous and finite-wavevector structures, although the strongly constrained antiphase remains a challenging case. Second, at the Ising-like critical point, the autocorrelation time of conventional local Monte Carlo dynamics grows algebraically with the system size, whereas the conditional wavelet updates remain decorrelated in $\mathcal{O}(1)$ sweeps at every scale. Since the number of reconstruction levels grows as $\log_2 L$, the sequential number of conditional sampling stages required for a complete

reconstruction grows as $\mathcal{O}(\log_2 L)$. Third, the magnetization distribution is accurately reproduced, whereas the histogram of the microscopic energy displays a residual mismatch. As we discuss below, this discrepancy is related to the finite-dimensional ansatz used for the learned conditional energies and to the difference between the original Gibbs measure and the approximate hierarchical measure sampled by the reconstruction procedure.

The chapter is organized as follows. In Sec. 8.1, we introduce the frustrated soft-spin models and discuss the phase diagram of the sBNNNI model. In Sec. 8.2, we specify the WCRG implementation used for the present application. Section 8.3.1 compares original Monte Carlo configurations with WCRG-generated samples across representative points of the phase diagram. In Sec. 8.3.2, we analyze the autocorrelation times at criticality, while Sec. 8.3.3 considers configuration-wide observables, with particular emphasis on the distributions of the magnetization and of the microscopic energy.

8.1 Frustrated soft-spin models

In order to leverage wavelet decompositions, which naturally act on real-valued signals [106, 107], we consider continuous spin variables $\phi_0(i) \in \mathbb{R}$, where i denotes the lattice site at position $\mathbf{r}_i = (x_i, y_i)$. The dimensionless microscopic energy on a square lattice is

$$E_0(\phi_0; T) = \frac{1}{2T} \sum_{i,j} \phi_0(i) K_{ij} \phi_0(j) + \sum_i \left[\phi_0^2(i) + \lambda \left(\phi_0^2(i) - 1 \right)^2 \right]. \tag{8.1}$$

It enters the microscopic probability measure as

$$p_0(\phi_0; T) = \frac{1}{Z_0(T)} \exp[-E_0(\phi_0; T)]. \tag{8.2}$$

Following Refs. [108, 149], this definition is not of the conventional form $H(\phi_0)/T$ with a temperature-independent Hamiltonian: the factor $1/T$ multiplies only the quadratic interaction kernel, while the local potential is not rescaled by temperature. The resulting low-temperature limit is discussed in App. C.1.

The frustrated quadratic interactions in Eq. (8.1) are encoded by the translationally invariant 5×5 stencil

$$\mathbf{K} = \begin{pmatrix} 0 & 0 & \kappa_y & 0 & 0 \\ 0 & \kappa_d & -J & \kappa_d & 0 \\ \kappa_x & -J & 0 & -J & \kappa_x \\ 0 & \kappa_d & -J & \kappa_d & 0 \\ 0 & 0 & \kappa_y & 0 & 0 \end{pmatrix}, \tag{8.3}$$

whose rows and columns encode the relative displacements $\Delta x, \Delta y \in \{+2, +1, 0, -1, -2\}$. More precisely, $K_{ij} = [\mathbf{K}]_{\Delta x, \Delta y}$, where $\mathbf{r}_i - \mathbf{r}_j = (\Delta x, \Delta y)$. Nearest-neighbor interactions have ferromagnetic strength $J > 0$, axial next-nearest-neighbor interactions along $\hat{x}$ and $\hat{y}$ have antiferromagnetic strengths $\kappa_x \geq 0$ and $\kappa_y \geq 0$, respectively, and diagonal nearest-neighbor interactions have strength $\kappa_d \geq 0$. All other interactions vanish.

The quartic term in Eq. (8.1) controls the stiffness of the local spin amplitude. In the Gaussian limit $\lambda \to 0$, the model reduces to a free quadratic theory. In the hard-spin limit $\lambda \to \infty$, the local

field is restricted to $\phi_0(i) = \pm 1$, recovering the corresponding frustrated Ising model. For finite λ, the model is a soft-spin ϕ^4 theory and differs from its hard-spin counterpart through non-universal properties and corrections to scaling, while retaining the same equilibrium universality class along the Ising-like transition line considered below [149].

The interaction stencil in Eq. (8.3) contains three standard frustrated geometries: the ANNNI model, with $\kappa_x > 0$, $\kappa_y = \kappa_d = 0$; the BNNNI model, with $\kappa_x = \kappa_y > 0$, $\kappa_d = 0$; and the 3NNNI model, with both axial and diagonal frustrating interactions. In the following, we focus on the sBNNNI model and set

$$J = 1, \qquad \kappa_x = \kappa_y = \kappa, \qquad \kappa_d = 0, \qquad \lambda = 1. \tag{8.4}$$

We vary the frustration parameter κ and the temperature T.

The resulting equilibrium phase diagram, determined through local Metropolis Monte Carlo simulations, is shown in Fig. 8.1. It contains five distinct phases, which are identified from the structure factor

$$S(q_x, q_y) = \frac{1}{L^2} \left\langle \left| \sum_i \phi_0(i) \exp\left[\mathrm{i}(q_x x_i + q_y y_i)\right] \right|^2 \right\rangle, \tag{8.5}$$

and, in particular, from its diagonal cut

$$S_{\mathrm{diag}}(q) = S(q, q). \tag{8.6}$$

The position of the dominant peak identifies the ordering wavevector, while the finite-size scaling of its amplitude distinguishes ordered and disordered phases [124].

The ferromagnetic (FM) phase is characterized by a peak at $q = 0$, whose amplitude scales extensively as $S_{\mathrm{diag}}(0) \sim L^2$. The homogeneous paramagnetic (PM) phase also exhibits a peak at $q = 0$, but its amplitude remains finite as $L \to \infty$. The modulated paramagnetic phase (MPM) has a peak at a finite wavevector $q > 0$, again with an amplitude that remains finite in the thermodynamic limit. In the incommensurate (IC) phase, the structure factor develops a peak at $0 < q^\star < \pi/2$ whose amplitude scales as $S_{\mathrm{diag}}(q^\star) \sim L^{2-\eta(T,\kappa)}$. Finally, the antiphase (AF) is a commensurate ordered state characterized by a peak at $q = \pi/2$ with amplitude proportional to L^2. In real space, it corresponds to a checkerboard modulation with domains of width two lattice spacings.

For the soft-spin model considered here, the IC phase extends over a broader region than in the corresponding discrete-spin BNNNI model studied in Ref. [150]. A zero-temperature variational analysis illustrating how a multiharmonic IC state can become energetically favorable at finite λ is presented in App. C.1. Its purpose is to explain qualitatively the enlarged IC region rather than to determine the two-dimensional phase boundaries quantitatively.

The phase boundaries in Fig. 8.1 are intended as guides to the eye, except for the four marked transition points between the FM and homogeneous PM phases. These points are consistent with the two-dimensional Ising universality class and were determined from the finite-size scaling of the magnetic susceptibility and second-moment correlation length, using the known exponents $\nu = 1$ and $\gamma = 7/4$; see App. C.2. The critical-sampling analysis below focuses on

$$\kappa = 0.2, \qquad T = T_{\mathrm{pc}}(L), \tag{8.7}$$

where $T_{\mathrm{pc}}(L)$ is the size-dependent pseudocritical temperature. The bulk extrapolation gives

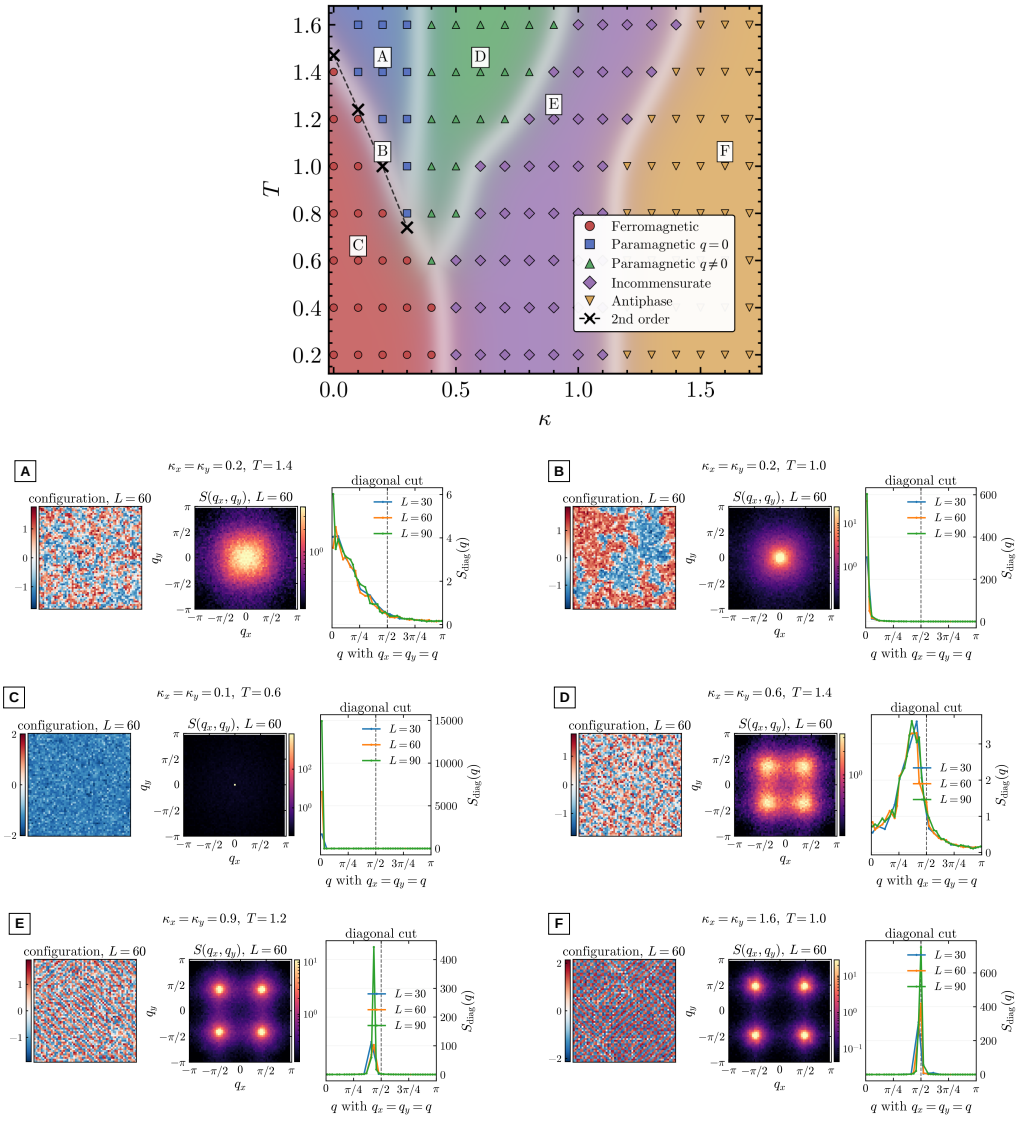


Figure 8.1: Phase diagram of the sBNNNI model with $\lambda = 1$ and representative behaviors in the different phases. **Top:** Phase diagram in the (κ, T) plane obtained from Monte Carlo simulations. The five phases are identified from the behavior and finite-size scaling of the structure factor $S(q_x, q_y)$, sampled at the points indicated by symbols. The faint boundaries are guides to the eye, whereas the black crosses denote the critical temperatures along the Ising-like ferromagnetic-to-paramagnetic transition, determined through finite-size scaling of the magnetic susceptibility and of the second-moment correlation length; see App. C.2. **Bottom:** For each representative point A–F, a typical equilibrium configuration in real space, the averaged structure factor $S(q_x, q_y)$, and its diagonal cut $S_{\mathrm{diag}}(q)$ from Eq. (8.6) for three system sizes. The peak position and the scaling of its amplitude distinguish homogeneous and modulated phases, as well as ordered and disordered phases.

$T_c = 0.987(1)$, while for $L = 64$ the pseudocritical temperature used below is $T_{\mathrm{pc}}(64) = 1.000(2)$.

In the following, we apply the WCRG to representative points spanning all five phases of the sBNNNI model and to the Ising-like critical line, with particular emphasis on the latter.

8.2 Wavelet conditional renormalization group

The general WCRG construction was developed in Chapter 7. We therefore do not repeat here the wavelet decomposition, the exact multiscale factorization, or the derivation of conditional score matching. We only specify the choices made for the sBNNNI model and recall the distinction between the original and learned probability measures that will be used in the numerical analysis.

At scale j, the microscopic field is decomposed into the coarse field ϕ_j and the three wavelet channels collected in $\bar{\phi}_j$, as in Eqs. (7.2) and (7.5). The explicit two-dimensional construction, including the low-pass map, the three high-pass channels, and the normalization conventions, is given in App. C.3. In the present application, the decomposition is iterated down to $L_J = 1$.

The original coarse-grained and conditional energies are defined in Eqs. (7.14) and (7.18). The resulting exact factorization of the microscopic probability distribution and the equivalent hierarchical representation of the dimensionless energy are given in Eqs. (7.20) and (7.21), respectively. For the present model, their microscopic starting point is Eq. (8.1).

The coarsest energy $E_J(\phi_J)$ is approximated by a parametric energy $E_{\theta_J}(\phi_J)$. Since $L_J = 1$, this is a one-dimensional potential expanded on a fixed family of sigmoid functions. At every finer scale, the exact conditional energy is approximated by the linear-in-parameter ansatz of Eq. (7.23). In the present implementation, the operator vector contains translationally invariant quadratic interactions and local potential terms of the reconstructed finer-scale field. The explicit parametrization is given in App. C.4.

The parameters $\bar{\theta}_j$ are learned independently at each scale by conditional score matching. As derived in Sec. 7.2.3, the empirical score-matching loss is quadratic in the parameters and its minimum is obtained by solving a regularized linear system. No Monte Carlo sampling of the current parametrized distribution is required during this learning step. The scale-dependent training pairs used to evaluate the empirical averages entering the linear system are defined in Eq. (7.53).

Synthetic microscopic configurations are generated recursively. One first samples the coarsest field from $p_{\theta_J}(\phi_J)$. Then, for $j = J, J-1, \ldots, 1$, the wavelet variables are sampled from Eq. (7.24) at fixed ϕ_j, and the next finer field is reconstructed through Eq. (7.5). We use local Metropolis updates of the individual wavelet coefficients. At fixed ϕ_j, the induced free energy in Eq. (7.25) is constant with respect to the sampled wavelet variables and cancels from the Metropolis ratio. Thus, the induced free energy does not need to be evaluated explicitly to generate configurations. Further details of the conditional sampler are given in App. C.5.

The recursive procedure samples the synthetic measure in Eq. (7.26), which is associated with the implicit dimensionless energy in Eq. (7.27). These contain the learned approximations and need not coincide with the original measure in Eq. (7.20), even for an arbitrarily large training set, when the finite operator basis cannot represent the original conditional energies. This distinction will be important when we discuss the microscopic energy distribution in Sec. 8.3.3.

The induced free energies in Eq. (7.25) are not computed during conditional sampling. If one wishes to construct an explicit microscopic energy functional associated with the learned hierarchy, they can be approximated by fitted functions $\bar{F}_{\alpha_j}(\phi_j)$, giving the explicit WCRG energy in Eq. (7.28). The fitted free energy $\bar{F}_{\alpha_j}$ is an explicit approximation of $\bar{F}^{\mathrm{ind}}_{\bar{\theta}_j}$: its parametrization and estimation are described in App. C.4.4.

The choice of wavelet family is adapted to the structure of each phase. We use Haar wavelets [143] in the homogeneous FM and PM phases and along the Ising-like critical line. For the MPM,

IC, and AF phases, we use Daubechies wavelets with four vanishing moments (DB4) [107]. Their smoother and wider filters are better suited to finite-wavevector modulations.

8.3 Numerical results

In this section, we test the WCRG construction as a sampling method for the sBNNNI model. We first compare Monte Carlo and WCRG-generated configurations across representative regions of the phase diagram. We then analyze the decorrelation of the conditional wavelet chains at criticality, where local Monte Carlo dynamics suffers from critical slowing down. Finally, we study configuration-wide observables, focusing on the magnetization and the microscopic energy as diagnostics of the generated measure.

8.3.1 Wavelet sampling across the phase diagram

We apply the WCRG construction to the representative points of the phase diagram shown in Fig. 8.1. Unless otherwise stated, we use systems of linear size $L = 64$. For each condition (T, κ), we first generate a reference ensemble of equilibrium configurations of the microscopic model in Eq. (8.1) using local single-site Metropolis updates. These configurations are used both as the original samples for comparison and as the training data for the conditional WCRG models.

For the ordered FM, IC, and AF phases, each run is initialized from an independent random configuration. We equilibrate the system until the magnetization and the spatial second moment, defined below in Eqs. (8.11) and (8.12), reach stationary values, and then discard an additional time window of at least five autocorrelation times of these observables. In the PM and MPM phases, and at the Ising-like critical point, we start from a disordered configuration and save configurations only after the same observables have thermalized. Consecutive saved configurations are separated by at least twice the measured autocorrelation time of these observables. At the critical point, this procedure explicitly reflects the critical slowing down of the underlying local Metropolis dynamics; its scaling with L is analyzed in Sec. 8.3.2.

The Monte Carlo configurations are transformed into the scale-dependent training pairs of Eq. (7.53), from which the conditional WCRG models are learned. After training, synthetic configurations are generated recursively from coarse to fine scales according to the procedure described in Sec. 8.2.

Figure 8.2 compares original and synthetic configurations at the representative points A–F. For each phase, we show a typical microscopic configuration, the distribution of local field values $P(\phi_0(i))$, and the diagonal structure factor $S_{\mathrm{diag}}(q)$. The local field distribution tests whether the reconstruction captures the single-site statistics of the soft-spin variable, whereas the structure factor probes spatial correlations and, in particular, the dominant ordering wavevector.

The agreement is generally good across the phase diagram. In the homogeneous phases and at the critical point, the reconstructed samples reproduce both the local field statistics and the peak structure at small wavevector. In the MPM and IC regimes, the method similarly captures the finite-wavevector peak of $S_{\mathrm{diag}}(q)$, showing that the synthetic configurations retain the modulated structure of the original ones.

The AF regime presents the most pronounced discrepancy. This case presents a particularly stringent test for the present reconstruction scheme, because the ordered state is a rigid commensurate modulation whose wavelength is tied directly to the microscopic lattice scale. In this

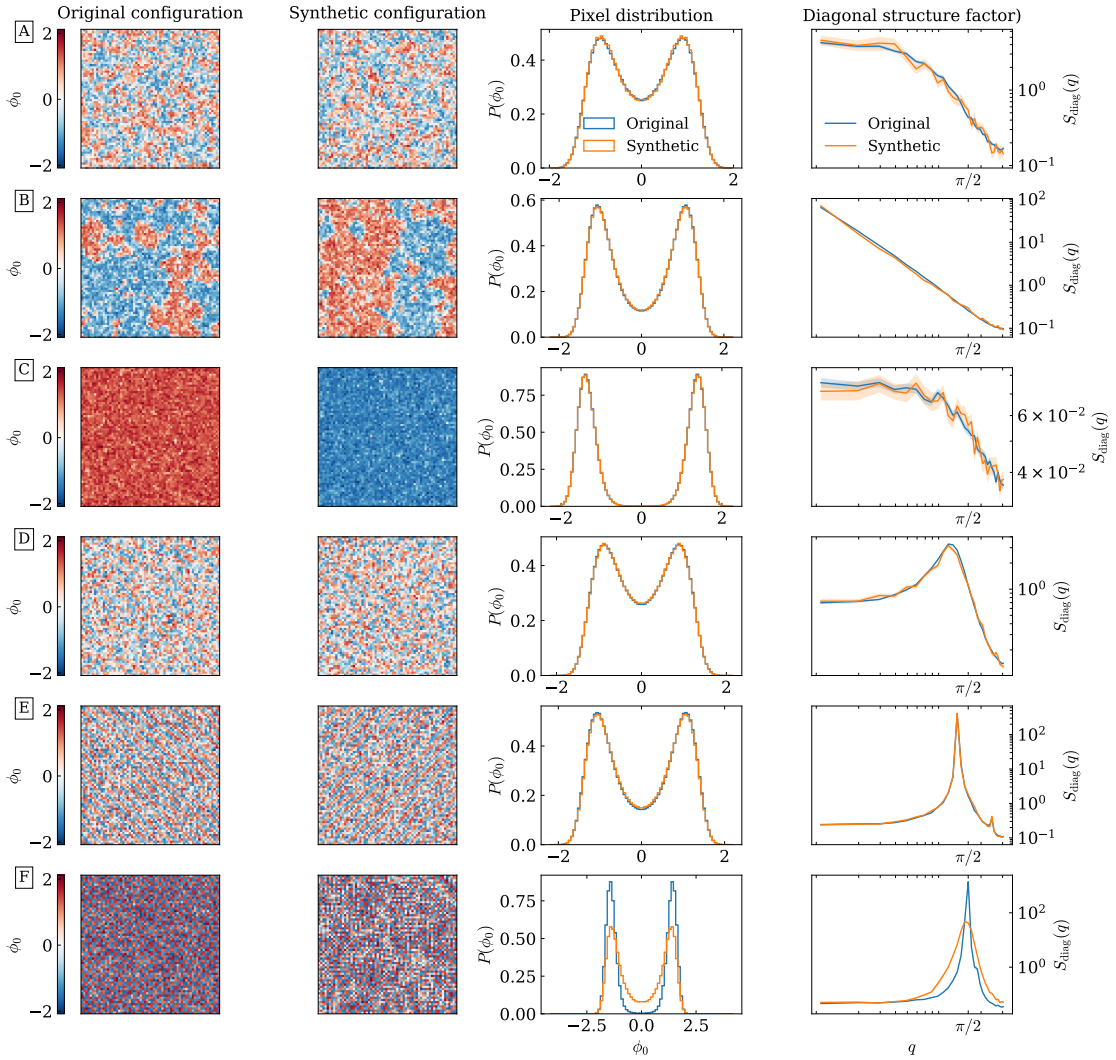


Figure 8.2: Comparison between original Monte Carlo configurations and samples generated through WCRG for representative points of the sBNNNI phase diagram A–F in Fig. 8.1. From top to bottom, the rows correspond to a point in the PM phase, along the Ising-like critical line, in the FM phase, in the MPM phase, in the IC phase, and in the AF phase. For each point, the panels show an original Monte Carlo configuration, a synthetic configuration reconstructed through recursive conditional sampling, and the comparison between original and reconstructed distributions of the local field $P(\phi_0(i))$ and of the diagonal structure factor $S_{\mathrm{diag}}(q)$ from Eq. (8.6). Haar wavelets are used for the FM, critical, and PM points, whereas DB4 wavelets are used for the MPM, IC, and AF phases. Good agreement is obtained for both local statistics and two-point correlations in all cases except the AF phase, for which the reconstruction is hindered for the reasons explained in the text.

situation, the low-pass fields generated at intermediate scales carry little information about the alternating pattern. After coarse graining, the checkerboard modulation is largely averaged out; the relevant information is concentrated only in the finest wavelet degrees of freedom. As a result, the conditional reconstruction exploits the multiscale hierarchy much less effectively than in phases where the dominant structures are also visible at coarse and intermediate scales.

8.3.2 Conditional sampling at the Ising-like critical point

At a continuous phase transition, critical slowing down originates from the growth of the correlation length ξ. In a finite system at criticality, ξ becomes comparable with L. For Markov

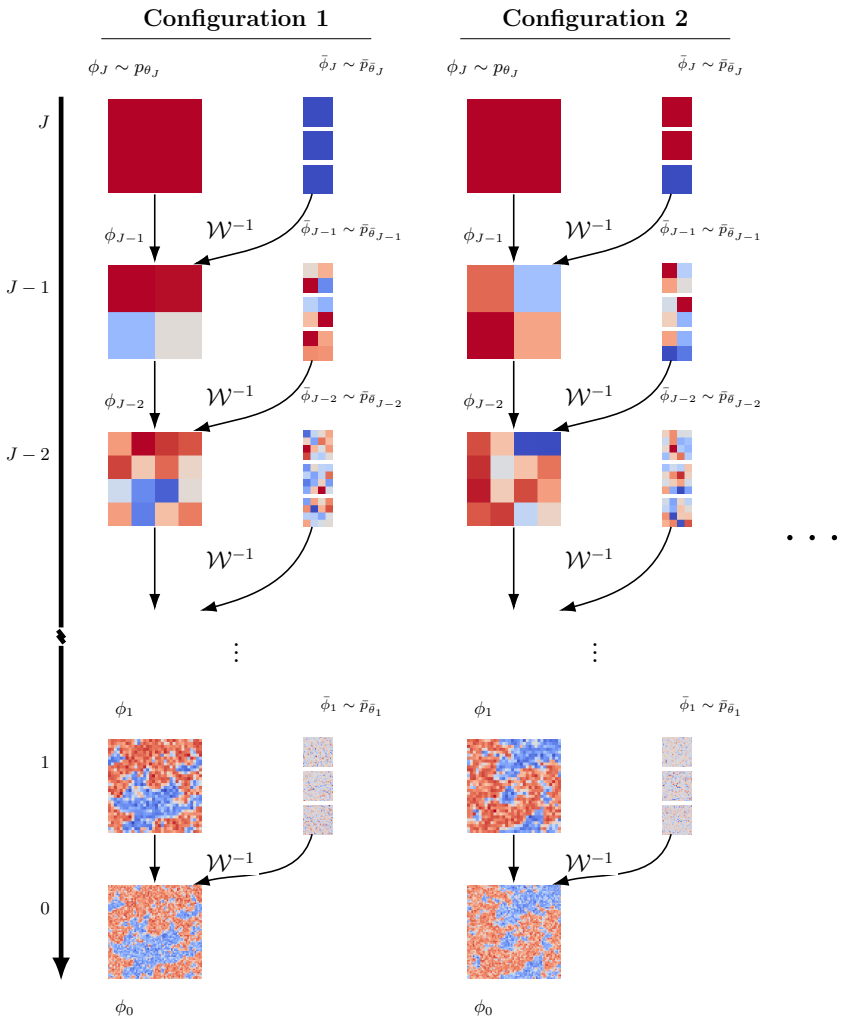


Figure 8.3: Multiscale sampling and reconstruction of two representative configurations. Each realization is generated starting from the coarsest scale (top configurations, ϕ_J) and moving towards the finer ones (bottom configurations, ϕ_0) by accounting recursively for the details encoded in three sampled wavelets at each scale (denoted $\bar{\phi}_J, \ldots, \bar{\phi}_1$ and represented by smaller squares). Colors encode ϕ_j and $\bar{\phi}_j$ at the various levels $j = J, J-1, \ldots, 0$, as in the previous figures. In practice, the coarsest variable ϕ_J is sampled first. Then, at each scale $j = J, J-1, \ldots, 1$, the wavelet variables $\bar{\phi}_j$ are sampled from the learned conditional distribution at fixed ϕ_j, and the next finer field ϕ_{j-1} is reconstructed through the inverse wavelet transform $\mathcal{W}^{-1}$. The large-scale structure of the eventual finer configuration ϕ_0 is therefore determined already at the coarse levels and it is progressively dressed by shorter-wavelength fluctuations, rather than assembled through local propagation on the microscopic lattice. Note that at the coarsest scale configurations 1 and 2 are rather similar and so are two of the three initial wavelets $\bar{\phi}_J$. Minor differences on that scale suffice to give rise to completely different clusters at the finest scale.

chain Monte Carlo algorithms based on local microscopic updates, reorganizing correlated structures then requires information to propagate over distances of order L, leading to a growth of equilibration and decorrelation times with system size [95, 96].

To quantify this effect, we consider an observable $O(t)$ measured along a trajectory, with time t measured in Monte Carlo sweeps, for a step size chosen such that the acceptance rate is about 50%.

Its normalized autocorrelation function is

$$C_O(t) = \frac{\langle O(t_0+t)O(t_0)\rangle - \langle O\rangle^2}{\langle O^2\rangle - \langle O\rangle^2}, \tag{8.8}$$

where the average is taken over the reference time t_0 along the stationary trajectory. At long times, we fit

$$C_O(t) \sim \exp(-t/\tau_O), \tag{8.9}$$

which defines the exponential autocorrelation time τ_O. For conventional local dynamics at criticality, one expects

$$\tau_O \sim L^{z_O}, \tag{8.10}$$

where the dynamical exponent may depend on both the update rule and the observable.

As a microscopic reference dynamics, we use local single-site Metropolis updates at $\kappa = 0.2$ and $T = T_{\rm pc}(L)$. We monitor the magnetization

$$m_0 = \frac{1}{L^2}\sum_i \phi_0(i), \tag{8.11}$$

and the spatial second moment

$$\langle \phi_0^2\rangle = \frac{1}{L^2}\sum_i \phi_0^2(i). \tag{8.12}$$

As shown in Fig. 8.4, both observables display the expected algebraic growth. A power-law fit gives

$$\tau_{m_0},\, \tau_{\langle\phi_0^2\rangle} \sim L^{z_{\rm loc}}, \qquad z_{\rm loc} = 2.11(4), \tag{8.13}$$

which is compatible, within the accuracy of the simulations, with the known dynamical critical exponent of local single-spin dynamics in the two-dimensional Ising universality class [151, 152].

The WCRG strategy is different. It does not attempt to reorganize a complete microscopic critical configuration through successive local moves. The coarse field ϕ_j already contains the large-scale morphology, and the conditional chain samples only the wavelet variables associated with the refinement from scale j to $j-1$. The reconstruction protocol is illustrated in Fig. 8.3.

The magnetization is *not* a useful observable for monitoring a conditional wavelet chain. With the orthogonal convention used here, wavelet modes have zero spatial average. At fixed ϕ_j, updating $\bar\phi_j$ therefore does not change the uniform component of the reconstructed field, as discussed in App. C.3. We instead monitor the second moment of the field reconstructed at that level,

$$O_{j,L}(t) = \frac{1}{L_{j-1}^2}\sum_{i\in\Lambda_{j-1}} \phi_{j-1}^2(i,t), \tag{8.14}$$

where $L \equiv L_0$ is the final microscopic size and $L_j = L/2^j$ is the size of the coarse field at reconstruction level j, and extract the corresponding conditional autocorrelation time $\bar\tau_j(L)$ from Eq. (8.9). The reported values are obtained by fitting the autocorrelation function averaged over 50 independent chains, with error bars indicating the corresponding statistical uncertainty.

The results in Fig. 8.4 show a qualitative difference between local and conditional dynamics. Within the range of sizes and scales investigated,

$$\bar\tau_j(L) = \mathcal{O}(1), \tag{8.15}$$

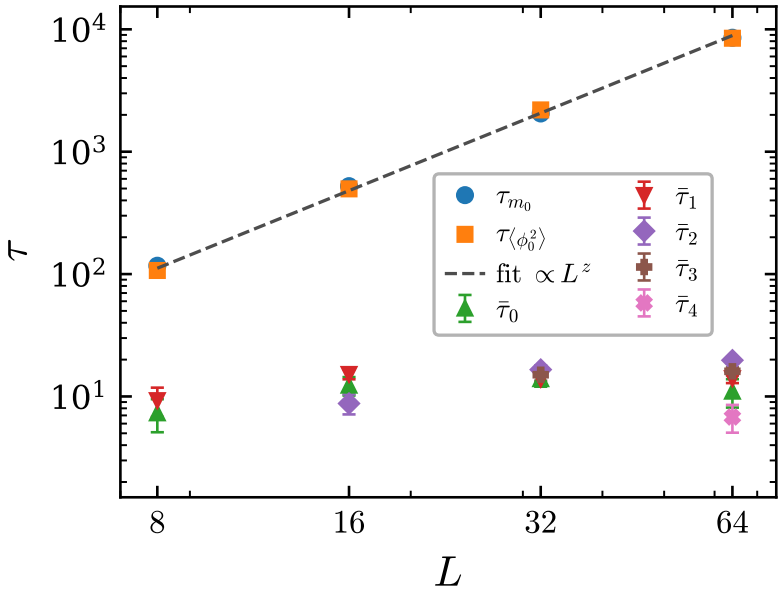


Figure 8.4: Autocorrelation times at the Ising-like critical point of the sBNNNI model with $\kappa = 0.2$ and $T = T_{\rm pc}(L)$. For local microscopic Metropolis dynamics, the autocorrelation times of the magnetization and of the second moment grow algebraically, with the fit in Eq. (8.13). By contrast, the conditional autocorrelation times $\bar{\tau}_j(L)$, measured for the observable in Eq. (8.14) during wavelet sampling at fixed coarse field, remain of order unity. Within the investigated sizes, no systematic growth is detected with either the final microscopic size L or the scale-dependent size $L_j = L/2^j$.

with no detectable systematic increase as either L or L_j grows. A fixed number of conditional sweeps is therefore sufficient to decorrelate the wavelet variables at every reconstruction level.

A complete microscopic realization requires one conditional sampling stage for each wavelet level. Since the hierarchy is iterated down to $L_J = 1$, the number of levels is

$$J = \log_2 L. \tag{8.16}$$

The sequential decorrelation depth of the full reconstruction, measured in conditional-sweep units, consequently scales as

$$\tau^{\rm depth}_{\rm WCRG} \sim \sum_{j=1}^{J} \bar{\tau}_j(L) = \mathcal{O}(J) = \mathcal{O}(\log_2 L). \tag{8.17}$$

This logarithmic scaling should not be confused with the total number of elementary local proposals. A conditional sweep at scale j updates a number of wavelet degrees of freedom proportional to L_j^2, and therefore

$$\mathcal{N}_{\rm WCRG} \propto \sum_{j=1}^{J} L_j^2 = \mathcal{O}(L^2), \tag{8.18}$$

up to scale-independent factors. The relevant point is that the cost is linear in the number of microscopic degrees of freedom and does not acquire the additional critical factor L^{z_O} characteristic of local microscopic sampling.

The discussion above assumes that the conditional energies have already been learned from the reference configurations. In the present implementation, the training configurations are themselves produced by conventional Monte Carlo and therefore inherit its critical cost. Moreover,

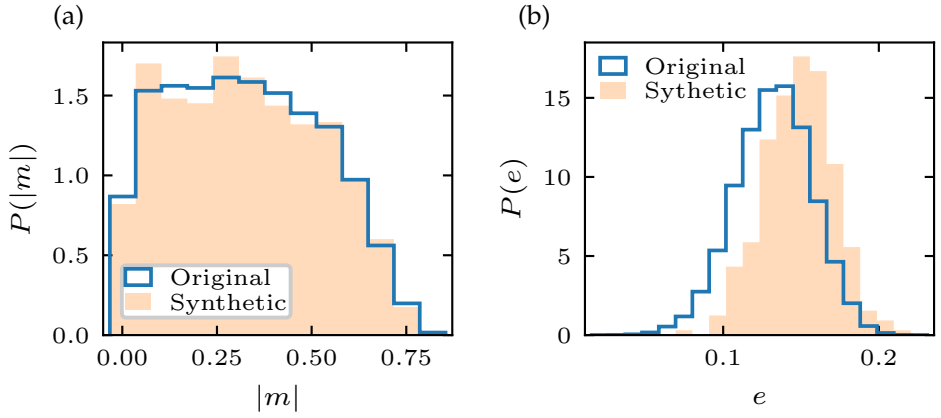


Figure 8.5: Comparison between configuration-wide observables for original and synthetic configurations at the Ising-like critical point of the sBNNNI model, with $\kappa = 0.2$, $L = 64$, and $T = T_{\mathrm{pc}}(64) = 1.0$. Distribution of (a) the absolute magnetization $|m_0|$ and (b) the microscopic energy density e from Eq. (8.19). The magnetization shows excellent agreement, whereas the energy presents a visible mismatch.

maintaining a fixed reconstruction accuracy as L increases may require an increasing number of training configurations. At present we are not aware of any studies of this issue in the context of WCRG, which is left for future investigation, although we expect this dependence to be rather weak.

8.3.3 Configuration-wide observables

The comparisons in Sec. 8.3.1 assess the local field distribution and the structure factor, which probe single-site statistics and two-point correlations. We now consider observables defined over an entire configuration. The simplest examples are the magnetization from Eq. (8.11) and the dimensionless microscopic energy per site

$$e = \frac{E_0(\phi_0; T)}{L^2}. \tag{8.19}$$

Figure 8.5 compares these two quantities for original and synthetic configurations at $\kappa = 0.2$, $L = 64$, and the corresponding pseudocritical temperature $T_{\mathrm{pc}}(64) = 1.0$.

Magnetization. Figure 8.5(a) shows excellent agreement between the absolute-magnetization distributions of original and synthetic configurations. This agreement should nevertheless be interpreted with caution. The distribution $P(m_0)$ encodes all moments $\langle m_0^n \rangle$, which are spatial integrals of microscopic multipoint correlation functions. Reproducing the complete hierarchy of correlations therefore implies the correct magnetization distribution, but the converse is not true: reproducing $P(m_0)$ alone does not guarantee that all microscopic correlations are correct.

In the present construction, the success in reproducing the magnetization is largely built into the multiscale reconstruction scheme of the WCRG. . At the coarsest scale $L_J = 1$, the field ϕ_J consists of a single degree of freedom proportional to the spatial average of ϕ_0. Learning its distribution amounts to fitting a one-dimensional effective potential. Since the wavelet modes have vanishing spatial averages (see App. C.3), the subsequent recursive reconstruction preserves the uniform component, up to the normalization convention of the transform. The magnetization distribution is therefore fixed at the coarsest scale and preserved throughout reconstruction.

Energy. Figure 8.5(b) shows a visible mismatch between the distributions of the microscopic energy evaluated on original and synthetic configurations. This observable provides a more sensitive test of the generated measure.

The origin of the discrepancy can be understood by comparing the original and learned hierarchies. The original microscopic energy in Eq. (7.21) involves the original conditional energies $\bar{E}_j$ from Eq. (7.18) and their original free energies $\bar{F}_j$ from Eq. (7.19). These objects are determined by the original coarse-grained energies generated through Eq. (7.14), which generally contain the complete hierarchy of interactions compatible with the symmetries of the model.

By contrast, conditional sampling uses the finite-dimensional energy ansatz in Eq. (7.23). The corresponding normalized distribution is Eq. (7.24), in which the learned energy induces the free energy $\bar{F}^{\mathrm{ind}}_{\bar{\theta}_j}$ defined in Eq. (7.25). Since $\bar{E}_{\bar{\theta}_j}$ is generally only an approximation of $\bar{E}_j$, its induced normalization is also generally different from $\bar{F}_j$. The original and learned conditional distributions therefore need not coincide,

$$\bar{p}_{\bar{\theta}_j}(\bar{\phi}_j \mid \phi_j) \neq \bar{p}_j(\bar{\phi}_j \mid \phi_j), \tag{8.20}$$

even for an infinite number of training samples if the retained operator basis is not expressive enough.

The reconstruction consequently samples the synthetic measure in Eq. (7.26), with implicit energy Eq. (7.27), rather than the original Gibbs measure in Eq. (8.2). In general,

$$p_0^{\mathrm{synth}}(\phi_0) \neq p_0(\phi_0), \qquad E_0^{\mathrm{synth}}(\phi_0) \neq E_0(\phi_0). \tag{8.21}$$

The histogram is obtained by evaluating the original microscopic energy E_0 on configurations generated from the synthetic measure. Because these configurations are not distributed exactly according to p_0, the distribution of $E_0[\phi_0^{\mathrm{synth}}]$ need not coincide with the distribution of $E_0[\phi_0^{\mathrm{orig}}]$. The mismatch in Fig. 8.5(b) is therefore to be expected, and it is an indication that a more refined finite dimensional representation of $\bar{E}_j$ is required to capture the energy distribution. This research direction is left for future work.

8.4 Concluding remarks

In this chapter, we applied the WCRG to the sBNNNI model as a benchmark for learned multiscale sampling in a frustrated system. The method decomposes microscopic configurations into coarse fields and wavelet degrees of freedom, learns the conditional distribution of the wavelet variables independently at each scale, and generates new configurations recursively from coarse to fine resolutions.

We showed that WCRG-generated configurations reproduce the main statistical features of Monte Carlo samples across representative points of the sBNNNI phase diagram. In particular, the method captures both local field distributions and diagonal structure factors in the homogeneous phases, along the Ising-like critical line, and in the modulated MPM and IC regimes. The AF regime, however, represents a pathological case.

At the Ising-like critical point, conventional local Monte Carlo dynamics displays the expected critical slowing down, with $z_{\mathrm{loc}} = 2.11(4)$. In contrast, the conditional wavelet chains remain decorrelated in $\mathcal{O}(1)$ sweeps at each reconstruction scale, within the range of sizes investigated. Since the number of wavelet levels grows as $\log_2 L$, the time required to generate independent

configurations grows as logarithmically in L. The total number of elementary updates remains proportional to the number of microscopic degrees of freedom, but it does not acquire the additional critical factor L^z characteristic of local microscopic sampling.

The analysis of configuration-wide observables clarifies the main limitation of the method. The magnetization distribution is accurately reproduced but is already fixed at the coarsest scale. The microscopic-energy histogram instead retains a visible mismatch. This difference reflects the tradeoff of the present construction: local Monte Carlo samples the original Gibbs measure but slows down at criticality, whereas WCRG provides efficient generation from an approximate hierarchical measure whose accuracy is controlled by the expressiveness of the learned conditional energies, that here are limited to quadratic-plus-local ansatz.

From this perspective, WCRG is best viewed as a complementary strategy to Monte Carlo. It trades exact sampling from the original Gibbs measure for a fast generative procedure that can, in principle, be systematically improved. The present results suggest that this tradeoff can be particularly useful in frustrated systems, where conventional cluster constructions are ineffective.

Chapter 9

Conclusions

In this concluding chapter, we summarize the main results of the two parts of this thesis and discuss the open questions and directions for future research that emerge from them.

Part I

Chapters 2 and 3 of this thesis focused on the numerical investigation of three two-dimensional XY models with vision-dependent interactions: one non-reciprocal version introduced in Ref. [89] and two reciprocal versions that we defined for the purpose of disentangling the ingredient of non-reciprocity from the one of the vision cone for what concerned the effects on the collective behavior. Although the non-reciprocal XY model, the asymmetric reciprocal model, and the symmetric reciprocal model differ in their microscopic dynamics, all three display true long-range orientational order. In particular, the close similarity between the phase diagrams of the non-reciprocal and of the asymmetric reciprocal models demonstrates that long-range order is not a consequence of non-reciprocity alone. Rather, the essential ingredient is the vision-dependent geometry of the interactions. By coupling the internal orientation of each spin to the fixed spatial directions of the lattice, the vision cone makes a global solid rotation of all the spins physically different from a symmetry transformation of the conventional XY model.

The reciprocal models made it possible to identify two distinct microscopic mechanisms through which the spin-lattice coupling generates an effective fourfold anisotropy. For some vision-cone apertures, particular global orientations deactivate lattice bonds and are therefore energetically unfavorable. In this case, the Z_4 anisotropy is already present at the level of the energy of uniformly aligned configurations. The asymmetric reciprocal model at $\theta = 180^\circ$ constitutes a different case: in fact, all perfectly aligned states are energetically degenerate, but their spectra of low-energy fluctuations depend on their orientation relative to the lattice. When the orientation occurs along the lattice axes, we identified a class deformations that has an energy cost of order one, whereas the corresponding deformations occurring when the orientation is along the lattice diagonals have a cost that scales as the linear size L. As a consequence, states aligned along lattice axes possess a broader set of low-energy fluctuations and this leads to an entropic selection mechanism that favors alignment along lattice directions, despite the exact energetic degeneracy of all uniformly aligned configurations.

The symmetric reciprocal model provides another particularly clear example of the spin-lattice interplay. For vision-cone apertures larger than 180°, in a nearly aligned low-temperature

configuration, a bond can be activated simultaneously by both of its endpoint spins, although this second activation does not change its energetic contribution, and we called it *redundant*. This bond redundancy makes the energy almost independent of the global orientation and produces an accidental continuous $O(2)$ degeneracy, consistent with the observation of quasi-long-range order. Close to 180°, thermal fluctuations make simultaneous activation from both endpoints less likely. Once the redundancy is reduced, rotating the configuration can change the set of active bonds and generate an orientation-dependent energy cost with Z_4 periodicity. Accordingly, temperature reveals the underlying lattice anisotropy and produces the observed reentrant transition from quasi-long-range to long-range order.

The subsequent analysis of Chapter 4 showed that the functional form of the anisotropic interaction kernel does not by itself determine the selected ordered state. Starting from a local configuration-dependent selfish energy associated with the interactions of a single spin with its neighbors and depending on the orientation of that spin, we derived, for different interaction kernels, a stochastic Langevin dynamics containing two distinct terms, that we called reactive and proactive. The reactive contribution describes the response to the orientations of the currently perceived neighbors of a spin, whereas the proactive contribution accounts for the gradient of the interaction coupling. This second contribution effectively pushes the agent to maximize its sensory input, i.e., the effective number of visible neighbors which are not misaligned with it. Retaining only the first or both terms may result into different pinning directions even when they originate from the same interaction kernel. For example, for the smooth kernels we studied, i.e., the von Mises kernel, the reactive dynamics favors diagonal alignment, while the dynamics which includes also the proactive term selects the lattice axes. Accordingly, the presence of the proactive term can qualitatively alter the macroscopic state by favoring orientations that preserve or increase the total perceived coupling.

The more detailed analysis of Chapter 5 further distinguished between local and global pinning. For some interaction kernels, the microscopic evolution equations contain an explicit local restoring term, whereas in other cases, no such term is present and orientational pinning emerges only through the collective dynamics of the mean orientation. When the remaining fluctuating contributions are sufficiently weak, the global orientation can be accurately described by an effective Ornstein-Uhlenbeck process, whose relaxation time provides a direct measure of the global pinning strength. Comparison with the static structure factor further showed that the effective pinning governing spatial correlations can be closer to the global restoring coefficient than to the spatial average of the local one.

This description, however, breaks down for more complex interaction kernels, where large fluctuations and additional renormalization effects can substantially modify the dynamics and weaken the correspondence between the analytical pinning coefficients and the effective pinning extracted from the static structure factor. We identified the conditions under which the different estimators provide a quantitatively consistent characterization of the pinning mechanism. A natural direction for future work is therefore to investigate regimes in which these conditions no longer hold (primarily due to fluctuations) by developing a coarse-grained, renormalization-group (RG) description of non-reciprocal XY models with anisotropic interactions. Such an approach could clarify how the reactive and proactive contributions, together with the lattice-induced pinning mechanisms identified in this work, evolve under coarse graining and ultimately determine the asymptotic ordering properties of the system. Recent RG and scaling approaches to non-reciprocal phase transitions provide useful guidance in this direction [52, 59, 153–156].

This reactive-proactive framework also suggests broader consequences for non-reciprocal many-body systems, that are presented as future perspectives in Chapter 6. Preliminary observations reported indicate that the proactive response may promote defect annihilation and suppress the proliferation of asters and shock lines that destabilize ordered states under purely reactive dynamics. Whether this mechanism is sufficient to render the long-range ordered phase asymptotically stable remains an open question. Extending the same reactive-proactive distinction to off-lattice active matter provides another promising application. In a Vicsek-like implementation, the proactive contribution acts as a novel steering mechanism that has not been previously considered in the context of intelligent active particles. Preliminary simulations suggest that this mechanism can suppress dispersion and promote the formation of cohesive queues or elongated moving structures. More broadly, this points towards a class of intelligent interaction rules in which agents respond not only to their present environment but also according to how their coupling with the environment is changing.

More broadly, one might wonder whether mechanisms analogous to those identified here play a role in biological collectives. Visual sensory networks can shape coordinated motion and the propagation of behavioral changes in animal groups [157]. Separately, non-reciprocal active torques have been predicted to enhance the speed and efficiency of information transfer during collective turning events [158]. While anisotropic interactions are known to be relevant in systems such as bird flocks, the underlying anisotropy may differ qualitatively from the forward-oriented vision kernels considered here. For instance, inference of the effective interactions from the experimental data of starling flocks suggests stronger interactions with lateral neighbors rather than with those located in front [71]. Recent models have shown that partial or structured visual fields can strongly reshape clustering, traveling bands, and density fluctuations [114]. Investigating how different forms of perceptual anisotropy shape collective ordering and whether proactive-like contributions play a role in biological groups remains an open question.

Part II

The second part of the thesis addressed the problem of the efficient sampling of critical systems. To overcome the long-range correlations that make direct sampling difficult near criticality, we adopted the wavelet conditional renormalization group (WCRG), that we reviewed in Chapter 7, which rewrites a microscopic probability distribution as an exact hierarchy composed of a coarse-field distribution and conditional distributions of wavelet variables at successive scales. This representation separates long-wavelength structures from the shorter-scale fluctuations added during reconstruction. Conditional score matching then provides a one-step estimation of the parameters at each scale, without requiring the repeated Monte Carlo simulations of the parametrized model that appear in maximum-likelihood learning. Under suitable convexity conditions, a small score-matching error also controls the difference between the learned and target conditional probabilities. Practical wavelets satisfy these ideal assumptions only approximately, so statistical accuracy must ultimately be tested numerically.

This test was performed in Chapter 8 on the soft-spin biaxial next-nearest-neighbor Ising model model, a frustrated spin system on the lattice containing homogeneous, modulated, and Ising-like critical regimes, for which standard cluster strategies for efficient sampling are ineffective. Across most of its phase diagram, the reconstructed configurations reproduce the local field distribution and the two-point correlations measured by the structure factor. At the Ising-like critical point,

local microscopic Monte Carlo dynamics displays the expected algebraic critical slowing down, while the conditional wavelet chains remain decorrelated in a number of sweeps that shows no detectable growth with system size. The sequential depth of the reconstruction therefore increases only as $\log_2 L$, without the additional critical factor that affects local microscopic sampling.

The accuracy of the generated measure is nevertheless limited by the finite parameterization of the conditional energies. While the magnetization distribution is accurately reproduced, the microscopic energy histogram are affected by a systematic mismatch. The learned conditional energies, together with their induced free energies, define an approximate hierarchical measure that may differ from the original Gibbs distribution. From this perspective, WCRG is best viewed as a complementary strategy to Monte Carlo. It trades exact sampling at fixed microscopic energy for a fast generative procedure that can, in principle, be systematically improved. The present results suggest that this tradeoff can be particularly useful in frustrated systems, where conventional cluster constructions are ineffective.

This places WCRG within the broader line of machine-learning-assisted sampling approaches in statistical physics, that include normalizing-flow-augmented Monte Carlo schemes, autoregressive variational networks, diffusion-based samplers, and machine-learning-assisted Monte Carlo algorithms [109–113]. Its specific contribution is to combine learning with an explicit multiscale RG structure: the learned objects are conditional energies associated with definite length scales, rather than a single black-box generator. This organization makes it possible to associate approximation errors with specific reconstruction levels and to improve the model in a controlled way.

Natural directions for further investigations include adding symmetry-allowed operators beyond the quadratic-plus-local ansatz used in the present work. Developing better ansätze would also be expected to resolve the poor sampling of the antiphase. Another promising direction would be to trade some of the interpretability of the present ansatz for greater expressive power, for example by parametrizing the conditional models with convolutional neural networks having local receptive fields, following Ref. [159]. A systematic study of how the approximation errors in the learned conditional distributions decrease with increasing model expressiveness, together with the development of practical convergence diagnostics, would help quantify this trade-off.

Appendix A

Additional material for Chapters 2 and 3

A.1 Details of the Monte Carlo simulation

We report here the details of the MC simulations. Once the model, system size, and temperature are set, the MC dynamics is initiated in two possible ways (which turn out to lead to identical results): (i) In a disordered configuration (with vanishing magnetization), but with a magnetic field selecting one of the four directions in which the system might order. As the simulation proceeds, the strength of the magnetic field is made to vanish. (ii) In an ordered configuration, with all spins aligned along one of the directions of possible order. This is to prevent the system from getting trapped in a metastable state with two or more domain walls (at low temperatures, when we expect a LRO phase with spins aligned except for fluctuations). After allowing for a thermalization time, estimated by observing when the quantities of interest at temperatures close to the phase transition have reached a steady value, the sampling begins. Between each sampling, we perform a number of sweeps proportional to L^2, due to the scaling of the autocorrelation time τ, which scales as $\tau \sim \xi^2$ for local update algorithms. With the sampled variables, such as the magnetization and the energy, we calculate the observables described in Sec. 2.2.1, and estimate the corresponding statistical error by using a bootstrap method that accounts for the residual correlation between data points via a blocking technique [160].

A.2 Scaling of the second-moment correlation length

Here we present the scaling behavior of the second-moment correlation length $\xi_L^{(2)}$ defined in Eq. (2.17) as a function of the system size L, for the two significant cases of the 6-state clock model and and the NRXY model with $\theta = 180^\circ$. As discussed in Sec. 2.2.1, $\xi_L^{(2)}$ is expected to scale as Eq. (2.18) within the three possible phases encountered in this work, i.e., LRO, QLRO, and DO.

In particular, the 6-state clock model exhibits all these three phases within different temperature ranges, as illustrated in Fig. 2.3(f)–(j). In Fig. A1(a), we report $\xi_L^{(2)}/L$ as a function of the system size L for various temperatures. Specifically, upon increasing L, $\xi_L^{(2)}/L$ grows as L (dashed line) at low temperatures (LRO phase), remains constant in the intermediate temperature regime (QLRO phase), while it decreases as $1/L$ (dash-dotted line) at high temperatures (DO phase).

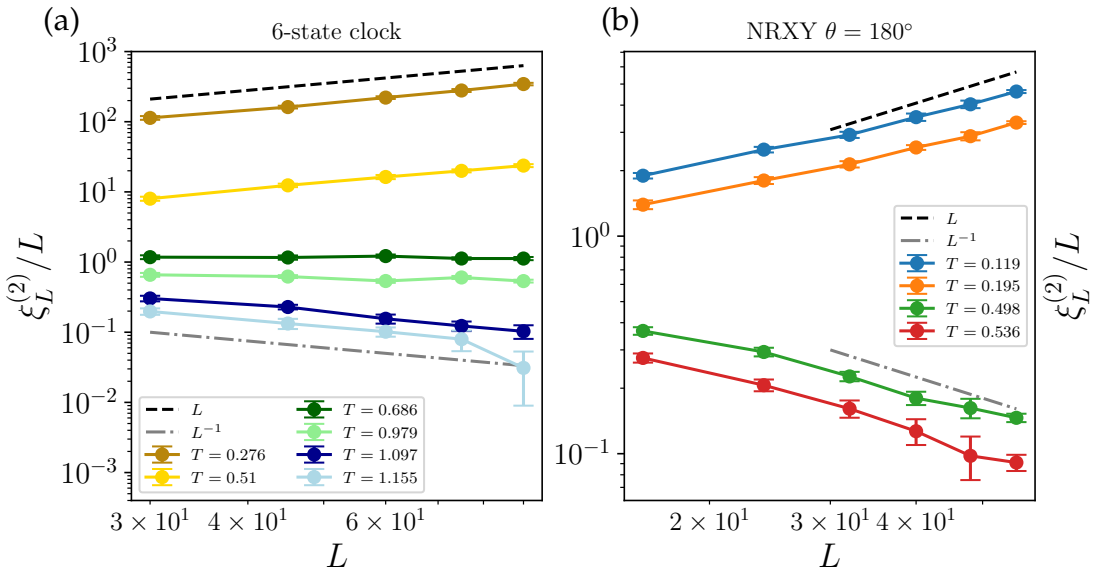


Figure A1: Dependence of the second-moment correlation length $\xi_L^{(2)}$ on the system size L for selected values of the temperature. (a) In the 6-state clock model we observe, depending on the temperature T, three different scalings of $\xi_L^{(2)}/L$ as a function of L, namely as L (dashed line, to guide the eyes), constant, and $1/L$ (dash-dotted line), corresponding to phases with LRO, QLRO and DO, respectively — see Eq. (2.18). In Fig. 2.3(g) the same quantity is plotted as a function of the temperature, for different system sizes. (b) In the NRXY model with $\theta = 180°$ we observe that, depending on the temperature, $\xi_L^{(2)}/L$ scales as L (dashed line) and $1/L$ (dash-dotted line) upon increasing L, in the LRO and DO phase, respectively. In Fig. 2.4(e) the same quantity is plotted as a function of the temperature, for different system sizes.

In Fig. A1(b), instead, we present the corresponding plot for the NRXY model at $\theta = 180°$. Here, we highlight the two observed scaling behaviors: $\xi_L^{(2)}/L \sim L$ (dashed line) in the LRO phase, and $\xi_L^{(2)}/L \sim 1/L$ (dash-dotted line) in the DO phase.

A.3 NRXY model with different update protocols

In a non-equilibrium model such as the NRXY model, the choice of the MC update protocol is generically expected to affect the stationary state reached by the dynamics and, consequently, the quantities measured in such a state. In principle, there are infinitely many ways to implement MC dynamics — for instance, one could choose the Glauber transition rate or the Metropolis one [161], and updates can be performed either sequentially across the lattice or by randomly selecting sites for updates. In general, these choices can lead to different entropy production rates, see Eq. (2.21). To illustrate this point, in Fig. A2(a) we present simulations of the NRXY model at various temperatures for $\theta = 225°$, and for a few selected choices of update protocols. Here we observe that the EPR increases when employing Glauber-like dynamics, which is characterized by lower acceptance rates, especially at high temperatures, compared to Metropolis dynamics. The EPR increases even further when updates are performed sequentially — this behavior can be attributed to the faster propagation of domain wall defects (see the video included in the Supplementary Material), which enhances the speed of irreversible dynamics.

However, beyond the quantitative differences, Fig. A2(a) shows that the qualitative trend of the EPR is quite robust upon changing the update protocols. (In passing, we recall that for the SRXY and ARXY models investigated in this work, the EPR vanishes identically due to their equilibrium

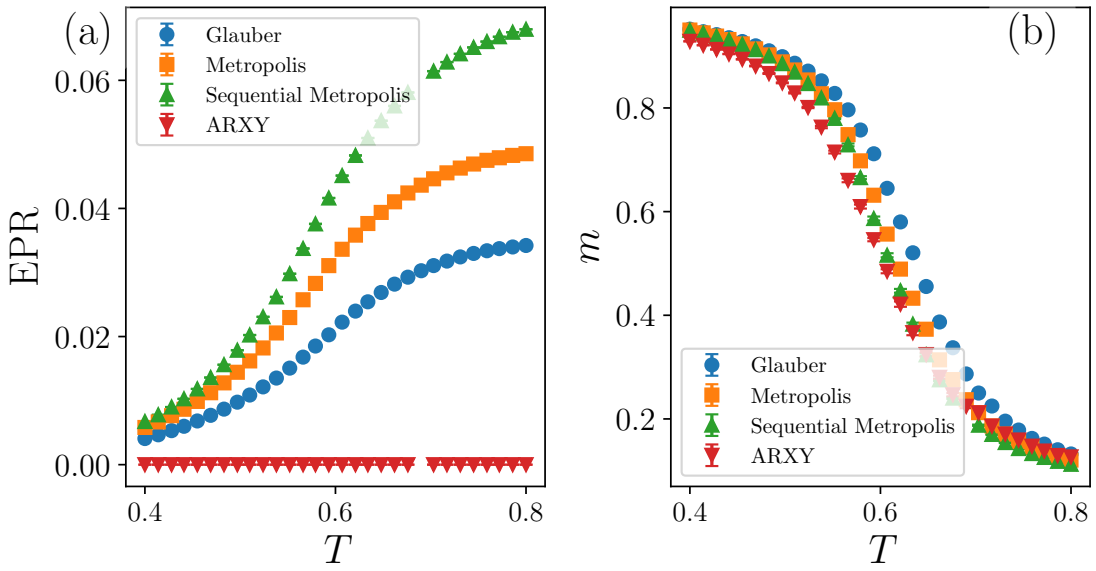


Figure A2: Temperature dependence of (a) the entropy production rate and (b) the magnetization m for various update protocols in the MC dynamics (see the legend) of the NRXY model with $\theta = 225°$. The dependence of these quantities on the protocols demonstrates that the latter lead to distinct steady states. In particular, it turns out that the EPR in panel (a) is the smallest for the Glauber dynamics (which has the lowest acceptance ratio probability), and the largest when implemented with sequential sweeps, while it takes intermediate values for the Metropolis dynamics. Clearly, the EPR vanishes for the equilibrium ARXY model. The magnetization m in panel (b) (as well as other observables) varies quantitatively, depending on the update protocol. However, these differences do not affect qualitatively the resulting phase diagram of the model.

nature.) Similarly, most other static (i.e., equal-time) observables that we measured in this work for the NRXY model turn out to exhibit only minor quantitative (but not qualitative) differences depending on the update rule. For example, in Fig. A2(b) we report the magnetization m obtained using various update protocols for simulating the NRXY model, showing their overall qualitative consistency. In particular, these values of m are also very similar to those of the ARXY model (discussed in Sec. 2.3.2) for the same value of θ, which (being an equilibrium model) reaches the same equilibrium state regardless of the chosen update protocol.

A.4 Analysis of the EPR in the NRXY model

As we discussed in Sec. 2.1.1, the stationary state reached by the NRXY model is a non-equilibrium one, characterized by a non-vanishing EPR. As shown in App. A.3, the actual value of the EPR depends on the chosen update protocol, while its overall behavior does not. Here, we investigate the dependence of the EPR (see Eq. (2.21)) on the value of the VC amplitude θ.

Figure A3(a) shows the EPR as a function of the temperature T and of θ, determined from the MC simulation of the model using Glauber dynamics with random updates and with $L = 100$. Interestingly enough, we observe that, generically, the EPR is rather small and it either vanishes or approaches zero within the lobes in the phase diagram (see Fig. 2.4(a)) corresponding to LRO, particularly in their lower parts. (Note that, as expected, the EPR vanishes identically for $\theta = 360°$ because, correspondingly, the NRXY model becomes the equilibrium XY model.) In order to explain this fact, we observe that the spin configurations within these lobes are characterized by

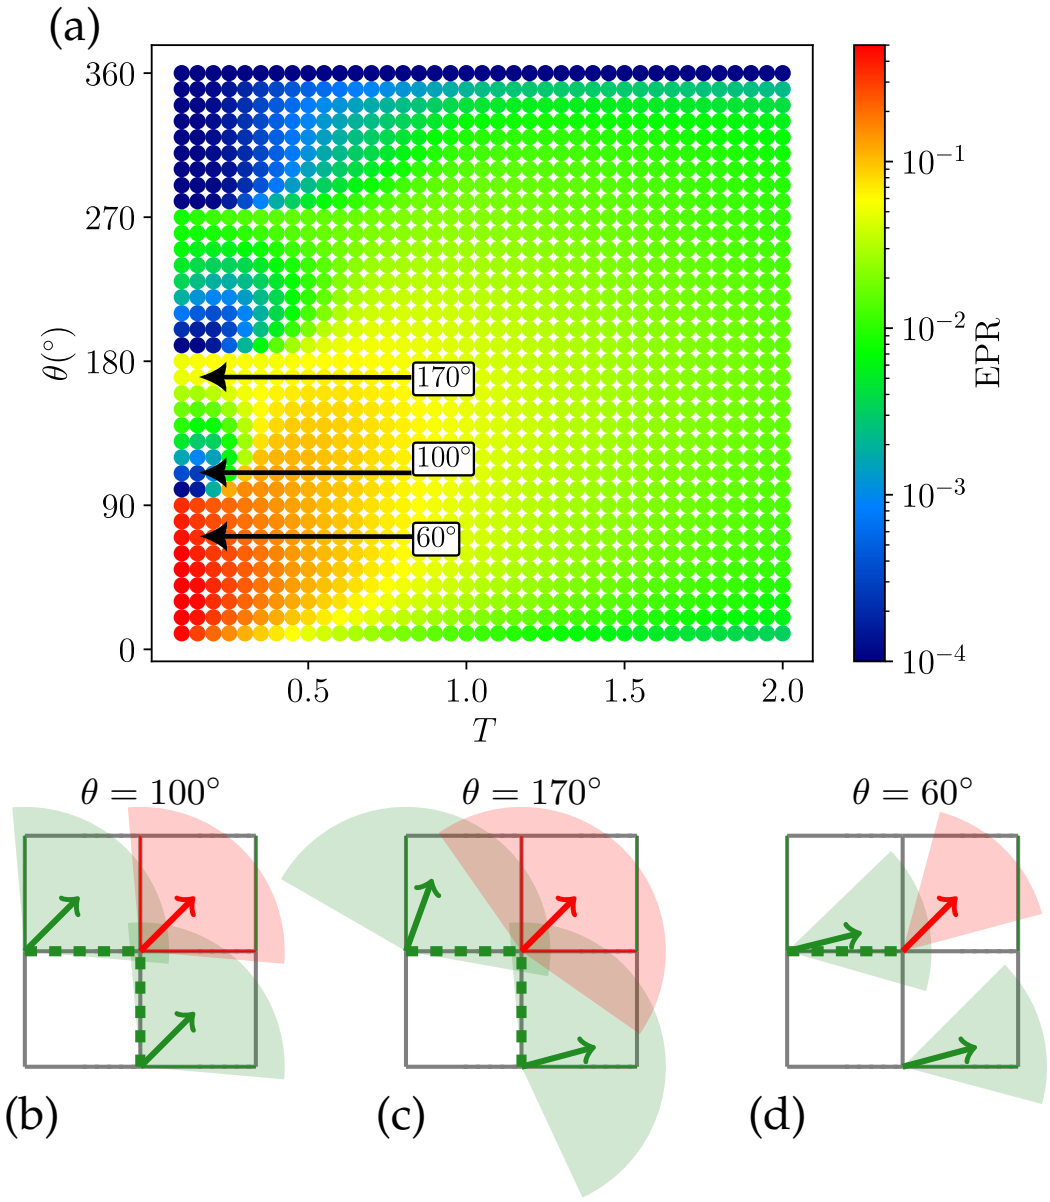


Figure A3: (a) Phase diagram of the EPR for the non-reciprocal XY model, obtained using Glauber updates with random site selection for a system of size $L = 100$. The EPR vanishes at $\theta = 360°$ (corresponding to the standard equilibrium XY model) and within the lobes with LRO, particularly in their lower parts, where the EUR is almost 360°. (b) Typical low-temperature configuration for $\theta = 100°$. The green spins, which make, for the update of the red spin, the difference between the NRXY model and the ARXY model, are almost aligned. As a result, they approximately act as a constant magnetic field, which allows one to write the variation of the selfish energy during an update as the variation of a global energy, see Eq. (A.1). This renders the model effectively an equilibrium one. (c) Typical low-temperature configuration for $\theta = 170°$. Since the smaller EUR allows for a greater range of motion for the spins even within the LRO phase, the constant magnetic field approximation is no longer applicable, and the system begins to exhibit its non-reciprocal nature. (d) Typical configuration for $\theta = 60°$. Since, for this value of the vision cone θ, there is no LRO even at low temperatures, the approximation in terms of the presence of a magnetic field is no longer viable and, as a consequence, the EPR takes its largest values within this region of the phase diagram.

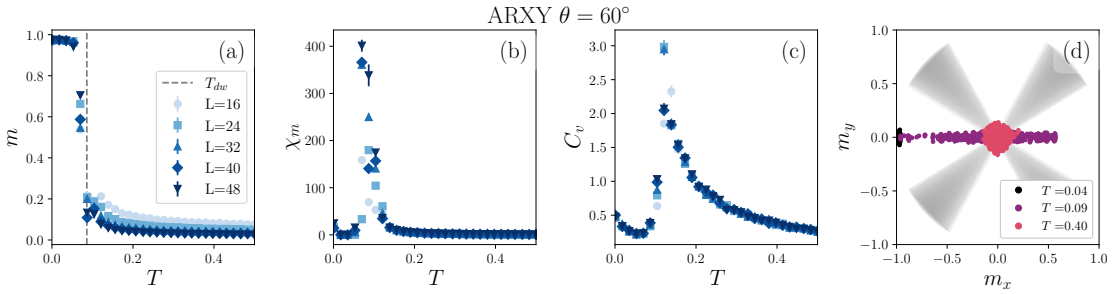


Figure A4: MC results for the ARXY model with VC angle $\theta = 60°$. The magnetization in panel (a) becomes nonzero at a finite temperature, despite the system being effectively one-dimensional. We estimate the temperature (vertical dashed line) at which this occurs in App. A.5. Correspondingly, panels (b) and (c) show that the magnetic susceptibility and specific heat exhibit non-analytic behavior at this transition. In panel (d), the vectorial magnetization illustrates alignment along the horizontal direction at low temperatures, transitioning to alignment with no preferred orientation along the horizontal directions, and eventually becoming disordered as the temperature increases. For a detailed discussion, we refer to App. A.5.

almost perfectly aligned spins (indeed, the corresponding large EUR implies limited freedom of rotation at low temperatures). As shown in Fig. A3(b) for the choice $\theta = 100°$, the bonds activated by the green spins — which, in the NRXY model, are irrelevant for updating the red spin (hence the dashed green lines) — mark the difference between the selfish (E_i^{NR}, see Eq. (4.1)) and asymmetric reciprocal (E^{AR}, see Eq. (2.4)) energy functionals. Accordingly, these green spins act like a magnetic field $\mathbf{h}$ aligned with the overall magnetization of the model ($\mathbf{h} = h\hat{\mathbf{m}}$, with $\hat{\mathbf{m}} = \mathbf{m}/|\mathbf{m}|$), whose fluctuations diminish as the EUR extends (i.e., approaching the lower parts of the lobes). As a result, the difference $\Delta_i E_i^{\mathrm{NR}}$ in the selfish energy E_i^{NR} due to an update in which spin i is updated from being $\mathbf{s}_i$ to $\mathbf{s}_i'$ can be written as the variation of a global functional:

$$\Delta_i E_i^{\mathrm{NR}} = \Delta_i E^{\mathrm{AR}} + \mathbf{h} \cdot (\mathbf{s}_i' - \mathbf{s}_i) = \Delta_i (E^{\mathrm{AR}} - \mathcal{E}^{\mathrm{mag}}), \tag{A.1}$$

where $\mathcal{E}^{\mathrm{mag}} = -\mathbf{h} \cdot \mathbf{m}$. Correspondingly, the model behaves effectively as an equilibrium one, with vanishing EPR. This behavior is fundamentally different from other non-reciprocal two-species lattice models [90], where even a minimal non-reciprocity at low temperatures leads to phenomena that are markedly different from those observed in the corresponding reciprocal versions.

Upon increasing the value of θ within each lobe, i.e., upon moving from its bottom part towards its upper part, the freedom of movement of the green spins in Fig. A3(b) increases due to a decrease of the EUR, as shown in Fig. A3(c). Correspondingly, the effective magnetic field h fluctuates more, and its very notion is no longer useful. Consequently, the difference in the selfish energy can no longer be described by a global functional, as in Eq. (A.1). At this point, the model begins to feel its non-reciprocal nature, and the EPR is no longer zero. This effect is even more pronounced for $\theta = 60°$, as shown in Fig. A3(d). In this case, no LRO phase exists even at low temperature (as we discuss in App. A.5), making the description in terms of an almost constant magnetic field no longer appropriate; this is indeed the region with the largest EPR (see Fig. A3(a)).

A.5 ARXY model for $\theta = 60^\circ$

Here we comment on the numerical results for the ARXY model with $\theta = 60^\circ$, which are summarized in Fig. A4. We first note that, as shown by the behavior of the magnetization m for the NRXY reported in Fig. 2.4(a) and for the ARXY in Fig. 2.5(a), for $\theta \in (0^\circ, 90^\circ)$ there is no lobe with LRO at low temperatures. This was rationalized in Ref. [89] in terms of an XY model with randomly diluted bonds, in which each bond can be activated with probability p [162, 163]. In this model, the QLRO phase at low temperature can only form if the number of activated bonds exceeds the percolation threshold p_c (equal to $1/2$ for the model on a square lattice).

In the ARXY model for $\theta < 90^\circ$, the fraction of activated bonds is equal to $1/2$ only for configurations in which the spins are aligned along the lattice direction, where the system effectively behaves as a collection of L almost independent one-dimensional chains of length L: in particular, in the limit $\theta \to 0$, this is approximated by one-dimensional Ising chains. As it is well-known, one-dimensional models do not exhibit LRO at finite temperature and, similarly, no LRO phase emerges here. On the other hand, they have a critical point at $T = 0$ characterized by exponential singularities of the correlation length [164].

However, the LRO, which is absent in the thermodynamic limit, might appear as a finite-size effect for large but finite systems. This can be understood in terms of the presence of domain walls which destroy the ordered state if their free energy F_{dw} is negative. Indeed, considering a single line of the ARXY model, a domain-wall defect has an energy $E_{\mathrm{dw}} = J/2 = 1/2$ and an entropy $S_{\mathrm{dw}} = \log L$. Accordingly, the corresponding free energy is given by $F_{\mathrm{dw}} = 1/2 - T \log L$ and it is negative for $T_{\mathrm{dw}}(L) > 1/(2 \log L)$: as shown in Fig. A4(a), the system, in fact, acquires a finite magnetization m for $T < T_{\mathrm{dw}}(L = 56)$. Correspondingly, the susceptibility in Fig. A4(b) and the specific heat in Fig. A4(c) exhibit non-analytical behavior at the corresponding temperature. Finally, Fig. A4(d) shows the vectorial magnetization $\mathbf{m}$, revealing the aforementioned organization into horizontal stripes at low temperatures (the equivalent ordering in vertical stripes can be obtained by changing the initial conditions in the MC dynamics). In this regime, all rows align in the same direction. As the temperature increases, different rows adopt distinct magnetization orientations until the system reaches a fully disordered configuration.

Appendix B

Additional material for Chapters 4 and 5

B.1 Crossover in the pinning direction for the reactive von Mises model on the triangular lattice

Throughout this work, we have focused on the non-reciprocal XY model on the square lattice, where the lattice geometry enhances the anisotropic character of the interactions compared, e.g., to the triangular lattice, as discussed also in the Supplementary Information of Ref. [65]. Here, instead, we consider the triangular lattice, following Refs. [65, 93], and show that the weaker lattice anisotropy due to such a lattice gives rise to a crossover in the selected pinning direction as the interaction strength is varied.

We start again from the Langevin dynamics in Eq. (4.3), but we consider only the purely reactive contribution, i.e., the first term on the r.h.s. of Eq. (4.3). On the triangular lattice, each spin has six nearest neighbors. Their angular positions ϑ_{ij} are given by

$$\vartheta_{ij} \in \left\{ n\frac{\pi}{3} \right\} \quad \text{with} \quad n = 0, \dots, 5, \tag{B.1}$$

We consider the von Mises interaction kernel defined in Eq. (4.8), and investigate which global orientation is selected by the reactive dynamics.

As in Sec. 4.2, we first focus on the mean-field description in which all spins are perfectly aligned with $\phi_i = \phi$ for all i. In this case, the total coupling of a spin to its nearest neighbors is

$$K_i(\phi) = \sum_{j \in \mathcal{N}_i} \exp\left[\sigma \cos(\phi - \vartheta_{ij})\right]. \tag{B.2}$$

This quantity is independent of the lattice site i and inherits the $\mathbb{Z}_6$ symmetry of the triangular lattice. Figure A1(a) shows the shifted coupling $K_i(\phi) - K_{i,\min}$ for two different values of the von Mises non-reciprocal parameter σ. The subtraction of the minimum $K_{i,\min} = \min_\phi K_i(\phi)$ from $K_i(\phi)$ does not alter the position of its stationary points, but makes the amplitude of the modulation for different values of σ easier to compare. The modulation is substantially weaker for $\sigma = 3$ than for $\sigma = 6$, indicating that thermal fluctuations are expected to compete more effectively with the deterministic pinning mean-field mechanism at smaller σ.

For the reactive dynamics, the argument of Sec. 4.2 predicts that the selected orientations should correspond to the minima of the total coupling $K_i(\phi)$. With the definition of ϕ and ϑ_{ij} used in Fig. A1(a), these mean-field minima are located at

$$\Phi \in \left\{ n\frac{\pi}{3} + \frac{\pi}{6} \right\} \quad \text{with} \quad n = 0, \dots, 5, \tag{B.3}$$

which are the six directions reciprocal to the principal directions of the triangular lattice. However, this zero-temperature argument does not provide a quantitative criterion for the stability of the selected direction at finite temperature. At finite T, spatial fluctuations dress the ordered state and generate additional pinning contributions that are not captured by the simple minimization of $K_i(\phi)$, leading to competition between deterministic and fluctuation-induced pinning mechanisms. A quantitative prediction of the value of T beyond which fluctuations become sufficiently strong to overcome the mean-field pinning would require evaluating the corresponding coefficients $S_\alpha^{(n),\mathrm{glob}}(t)$ and $C_\alpha^{(n),\mathrm{glob}}(t)$ introduced in Eq. (5.15) for the triangular lattice. We do not pursue this calculation here. However, at a qualitative level, one can expect that the mean-field prediction should be increasingly accurate as σ increases and temperature decreases, i.e., when the modulation of $K_i(\phi)$ is more pronounced and thermal fluctuations are less effective. In fact, Fig. A1(b) confirms this expectation. We report the time series of the mean orientation ϕ for two values of σ. Simulations were performed at $T = 0.2$, $L = 100$, with the integration time step dt chosen such that $K_{i,\max}(\phi)dt \sim 0.01$. For $\sigma = 3$, the system eventually pins along the principal lattice direction $\Phi = 0$, rather than along the direction predicted by the zero-temperature minimization of $K_i(\phi)$. This indicates that, in this case, fluctuation-induced pinning dominates over the mean-field contribution. For the larger value $\sigma = 6$, instead, the system eventually pins at $\phi = \pi/6$, corresponding to one of the reciprocal directions selected by the minima of $K_i(\phi)$. In this regime the mean-field mechanism becomes dominant again.

These results reveal a crossover between two competing pinning mechanisms. For weak anisotropy, finite-temperature fluctuation-induced pinning dominates and selects the principal lattice directions, as in the sinusoidal reactive kernel, whose zero-temperature total coupling is actually independent of ϕ. As the anisotropy increases, the deterministic mean-field mechanism prevails, and the system is pinned at the minima of $K_i(\phi)$.

A systematic characterization of this crossover in the parameter space (T, σ) is left for future work. We expect that an analogous crossover occurs also on the square lattice. In this case, however, the stronger anisotropy of the total coupling $K_i(\phi)$ (see Fig. A1(a)) implies that the fluctuation-dominated pinning should occur for smaller values of σ. Resolving the corresponding pinning direction in this weak-anisotropy regime would therefore require substantially larger system sizes and longer simulation times than those considered here.

B.2 Local equations for the spin

In order to identify the local pinning terms, we start from the Langevin equation reported in Eq. (4.3), substitute the different interaction kernels discussed in Sec. 4.1.3, and rearrange the resulting expressions so as to highlight the various contributions to the expansion in Eq. (5.7) introduced in Sec. 5.1.1. The coefficients $S_{\alpha,i}^{(n),\mathrm{loc}}(t)$ and $C_{\alpha,i}^{(n),\mathrm{loc}}(t)$ of the expansion depend on the instantaneous configuration of the neighboring spins of spin i and can be expressed as functions of the local angular differences, as in Eq. (5.8). The mean values and fluctuations of these coefficients

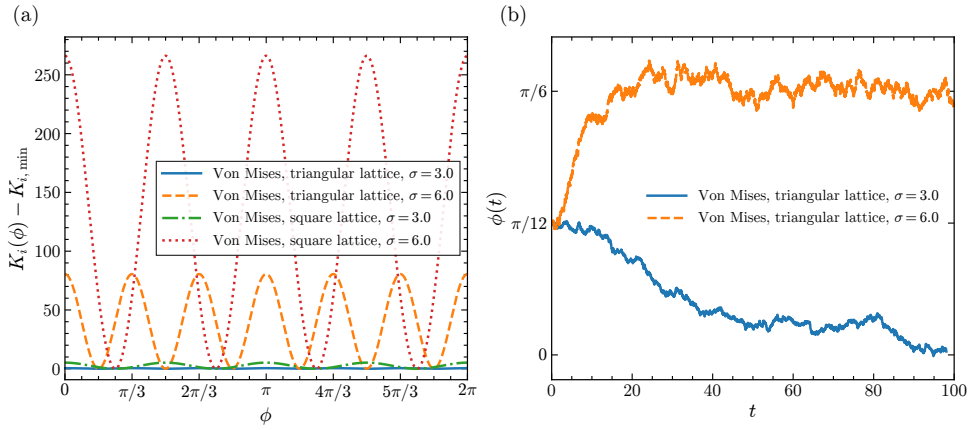


Figure A1: Crossover of the pinning direction for the reactive von Mises model on the triangular lattice. (a) Zero-temperature total coupling $K_i(\phi) - K_{i,\min}$ as a function of the global spin orientation ϕ, for two values of the parameter σ, i.e., $\sigma = 3$ and $\sigma = 6$. In addition to the sixfold modulation associated with the $\mathbb{Z}_6$ symmetry of the triangular lattice, we show the corresponding coupling for the same von Mises kernel on the square lattice, which displays the expected fourfold modulation. For a certain value of σ, the lattice-induced anisotropy is more pronounced on the square than on the triangular lattice, as indicated by its larger amplitude. Subtracting $K_{i,\min}$ does not affect the position of the stationary points, but makes the amplitudes of the different modulations easier to compare. (b) Time evolution of the mean orientation $\phi(t)$ for the same two values of σ as in panel (a) and starting from $\phi(0) = \pi/12$. For $\sigma = 3$, the system eventually pins along a principal lattice direction, given here by $\Phi = 0$. This fact indicates that the fluctuation-induced pinning dominates over the zero-temperature mean-field prediction. For $\sigma = 6$, instead, the system eventually pins at $\phi = \pi/6$, corresponding to one of the reciprocal directions selected by the minima of $K_i(\phi)$. The dynamics corresponds to $T = 0.2$ for both values of σ

can be computed numerically in the stationary state, once the system is pinned along a given direction Φ. Since, in the stationary state, the system is both time-translation invariant and, due to the periodic boundary conditions, space-translation invariant, averages and standard deviations can be computed either over time or over space. Here, we choose to perform spatial averages, as this makes it more transparent which terms vanish on average as a consequence of the periodic boundary conditions. Expanding around the pinned configuration in terms of the fluctuations $\delta\phi_i = \phi_i - \Phi$, one can investigate whether a local restoring term of the form $-r_i^{\mathrm{loc}}\delta\phi_i$ emerges in the (overdamped) evolution equation of ϕ_i, with r_i^{loc} depending on the coefficients $S_{\alpha,i}^{(n),\mathrm{loc}}$ and $C_{\alpha,i}^{(n),\mathrm{loc}}$. In general, r_i^{loc} is itself a fluctuating quantity, characterized by a mean value r^{loc} and a standard deviation $\Delta_{r^{\mathrm{loc}}}$ (see Eq. (5.13)), as discussed in Sec. 5.1.1.

In order to formulate a description on the continuum, we formally replace $\delta\phi_i \to \delta\phi_{\mathbf{x}}$, $S_{\alpha,i}^{\cdots} \to S_{\alpha,\mathbf{x}}^{\cdots}$, $C_{\alpha,i}^{\cdots} \to C_{\alpha,\mathbf{x}}^{\cdots}$, and perform a gradient expansion, e.g.,

$$\delta\phi_{i\pm\hat{\mu}} = \delta\phi_{\mathbf{x}} \pm a\,\partial_\mu\delta\phi_{\mathbf{x}} + \frac{1}{2}a^2\,\partial_\mu^2\delta\phi_{\mathbf{x}} + \mathcal{O}(a^3), \tag{B.4}$$

where $\hat{\mu}$ is the unit vector of the μ-th spatial direction, while we eventually set the lattice spacing a to unity. We emphasize that this construction on the continuum is mainly intended to facilitate comparison with coarse-grained descriptions which are reported in the literature, in particular that one of Refs. [93]. For notational simplicity, throughout this appendix we consider only the deterministic part of the dynamics and omit the Gaussian noise term.

B.2.1 Modulated coupling

We consider Eq. (4.3) for ϕ_i with the modulated coupling kernel (see Eq. (4.6)) and rewrite it in the form of Eq. (5.7), in order to identify the effective pinning terms. We then introduce the coarse-grained field $\phi_{\mathbf{x}}$ and perform the continuum expansion described in Eq. (B.4).

Reactive term. The reactive contribution (indicated by the subscript R) reads

$$\begin{aligned}(\partial_t \phi_i)_{\mathrm{R}} &= -\sum_{j\in\mathcal{N}_i} [1+\epsilon\cos(4\phi_i)]\sin(\phi_i-\phi_j) \\ &= C_{\mathrm{R},i}^{(0),\mathrm{loc}} + C_{\mathrm{R},i}^{(4),\mathrm{loc}}\cos(4\phi_i),\end{aligned} \tag{B.5}$$

where $\epsilon C_{\mathrm{R},i}^{(0),\mathrm{loc}} = C_{\mathrm{R},i}^{(4),\mathrm{loc}} = \epsilon\sum_{j\in\mathcal{N}_i}\sin(\phi_i-\phi_j)$. The coefficients $C_{\mathrm{R},i}^{(0),\mathrm{loc}}$ and $C_{\mathrm{R},i}^{(4),\mathrm{loc}}$ vanish upon averaging over the entire lattice. Consequently, the reactive contribution does not generate any local pinning term, but only a modulation of the diffusion. This conclusion is also confirmed in the continuum limit, in which the previous equation becomes

$$(\partial_t \phi_{\mathbf{x}})_{\mathrm{R}} \simeq [1+\epsilon\cos(4\phi_{\mathbf{x}})]\,\nabla^2\phi_{\mathbf{x}}. \tag{B.6}$$

This expression explicitly shows the presence of a modulated diffusion term and the absence of a local restoring force.

Proactive term. We now consider the proactive contribution (indicated by the subscript P),

$$(\partial_t \phi_i)_{\mathrm{P}} = -4\epsilon\sum_{j\in\mathcal{N}_i}\sin(4\phi_i)\cos(\phi_i-\phi_j) = S_{\mathrm{P},i}^{(4),\mathrm{loc}}\sin(4\phi_i) \quad \text{with} \quad S_{\mathrm{P},i}^{(4),\mathrm{loc}} = -4\epsilon\sum_{j\in\mathcal{N}_i}\cos(\phi_i-\phi_j). \tag{B.7}$$

To identify the local pinning coefficient (see Eq. (5.9)), we expand this equation in terms of the fluctuations $\delta\phi_i$ around one of the pinned directions, e.g., $\Phi = 0$. Isolating the term which is linear in $\delta\phi_i$ by writing $\sin(4\delta\phi_i) = \frac{\sin(4\delta\phi_i)}{4\delta\phi_i}4\delta\phi_i$,one obtains the following expression for the coefficient r_i^{loc} (introduced in Eq. (5.9)) of such a linear term:

$$r_i^{\mathrm{loc}} = -4\frac{\sin 4\delta\phi_i}{4\delta\phi_i}S_{\mathrm{P},i}^{(4),\mathrm{loc}}. \tag{B.8}$$

Passing to the continuum limit, one finds

$$\sum_{j\in\mathcal{N}_i}\cos(\phi_i-\phi_j) \simeq 4 - |\nabla\delta\phi_{\mathbf{x}}|^2, \tag{B.9}$$

which yields

$$(\partial_t\delta\phi_{\mathbf{x}})_{\mathrm{P}} \simeq -r_{\mathbf{x}}^{\mathrm{loc}}\delta\phi_{\mathbf{x}} \simeq -16\epsilon\frac{\sin(4\delta\phi_{\mathbf{x}})}{4\delta\phi_{\mathbf{x}}}\left(4-|\nabla\delta\phi_{\mathbf{x}}|^2\right)\delta\phi_{\mathbf{x}}. \tag{B.10}$$

Accordingly, the proactive interaction generates a genuinely local pinning mechanism. By contrast, in the purely reactive dynamics, the selection of preferred orientation emerges only at the collective level, as it is shown furtgher below in App. B.3.1, where the corresponding equations for the global orientation are derived.

B.2.2 Sinusoidal coupling

We rewrite Eq. (4.3) for ϕ_i in the case of the sinusoidal coupling (see Eq. (4.7)) in the form of Eq. (5.7), in order to identify the pinning terms. We then take the continuum limit by introducing the coarse-grained field $\phi_{\mathbf{x}}$ and perform the continuum expansion as in Eq. (B.4).

Reactive term. We consider the reactive term

$$(\partial_t \phi_i)_{\mathrm{R}} = -\sum_{j\in\mathcal{N}_i} \Big[1 + \epsilon \cos(\phi_i - \vartheta_{ij})\Big] \sin(\phi_i - \phi_j), \tag{B.11}$$

which can be rewritten in the form

$$\begin{aligned}(\partial_t \phi_i)_{\mathrm{R}} = &-\sum_{j\in\mathcal{N}_i} \sin(\phi_i - \phi_j) \\ &- \epsilon \cos\phi_i \Big[\sin(\phi_i - \phi_{i+\hat{x}}) - \sin(\phi_i - \phi_{i-\hat{x}})\Big] \\ &- \epsilon \sin\phi_i \Big[\sin(\phi_i - \phi_{i+\hat{y}}) - \sin(\phi_i - \phi_{i-\hat{y}})\Big],\end{aligned} \tag{B.12}$$

where $\hat{x}$ and $\hat{y}$ are unit vectors along the x and y axis of the lattice. Comparing this expression with Eq. (5.7), we identify the coefficients

$$C^{(1),\mathrm{loc}}_{\mathrm{R},i} = -\epsilon \Big[\sin(\phi_i - \phi_{i+\hat{x}}) - \sin(\phi_i - \phi_{i-\hat{x}})\Big], \tag{B.13}$$

$$S^{(1),\mathrm{loc}}_{\mathrm{R},i} = -\epsilon \Big[\sin(\phi_i - \phi_{i+\hat{y}}) - \sin(\phi_i - \phi_{i-\hat{y}})\Big]. \tag{B.14}$$

The symmetry properties of these coefficients follow from the same symmetry arguments discussed for the corresponding global coefficients in App. B.3.2. We therefore refer the reader to the symmetry analysis presented there. In particular, when the system is pinned along $\Phi = 0$, one has $C^{(1),\mathrm{loc}}_{\mathrm{R},i} = 0$ and $S^{(1),\mathrm{loc}}_{\mathrm{R},i} \neq 0$, whereas for a state pinned along $\Phi = \pi/2$ the roles of the two coefficients are exchanged: $S^{(1),\mathrm{loc}}_{\mathrm{R},i} = 0$ and $C^{(1),\mathrm{loc}}_{\mathrm{R},i} = -S^{(1),\mathrm{loc}}_{\mathrm{R},i}|_{\Phi=0}$. These relations are the local counterparts of those derived in App. B.3.2 for the global coefficients.

Expanding Eq. (B.12) in spatial gradients, we find the following expression on the continuum:

$$(\partial_t \phi_{\mathbf{x}})_{\mathrm{R}} = \nabla^2 \phi_{\mathbf{x}} + 2\epsilon(\sin\phi_{\mathbf{x}} \partial_y \phi_{\mathbf{x}} + \cos\phi_{\mathbf{x}} \partial_x \phi_{\mathbf{x}}) + \frac{\epsilon}{3}\left\{\sin\phi_{\mathbf{x}}\left[\partial_y^3 \phi_{\mathbf{x}} - (\partial_y \phi_{\mathbf{x}})^3\right] + \cos\phi_{\mathbf{x}}\left[\partial_x^3 \phi_{\mathbf{x}} - (\partial_x \phi_{\mathbf{x}})^3\right]\right\}, \tag{B.15}$$

which corresponds to Eq. (5) in Ref. [93]. When we search for the first potentially non-vanishing contributions to the local pinning — i.e., linear in the fluctuations of $\phi_{\mathbf{x}}$ — we see that the terms $\partial_x \phi_{\mathbf{x}}$, $\partial_y \phi_{\mathbf{x}}$, $\partial_x^3 \phi_{\mathbf{x}}$, $\partial_y^3 \phi_{\mathbf{x}}$ are total derivatives and therefore vanish upon spatial averaging under periodic boundary conditions. By contrast, the terms $(\partial_y \phi_{\mathbf{x}})^3$ and $(\partial_x \phi_{\mathbf{x}})^3$ are not total derivatives and thus provide the candidate for the leading (non-vanishing) local pinning contributions.

Proactive term. We now consider the proactive contribution

$$(\partial_t \phi_i)_{\mathrm{P}} = -\epsilon \sum_{j\in\mathcal{N}_i} \sin(\phi_i - \vartheta_{ij})\, \cos(\phi_i - \phi_j), \tag{B.16}$$

which can be written in the form

$$\begin{aligned}(\partial_t\phi_i)_{\mathrm{P}} =& -\epsilon\Big[\cos(\phi_i-\phi_{i+\hat{x}})-\cos(\phi_i-\phi_{i-\hat{x}})\Big]\sin\phi_i \\ &+\epsilon\Big[\cos(\phi_i-\phi_{i+\hat{y}})-\cos(\phi_i-\phi_{i-\hat{y}})\Big]\cos\phi_i \qquad \text{(B.17)}\\ =&C^{(1),\mathrm{loc}}_{\mathrm{P},i}\cos\phi_i + S^{(1),\mathrm{loc}}_{\mathrm{P},i}\sin\phi_i. \qquad \text{(B.18)}\end{aligned}$$

The coefficients $C^{(1),\mathrm{loc}}_{\mathrm{P},i}$ and $S^{(1),\mathrm{loc}}_{\mathrm{P},i}$ (the expressions of which are readily obtained by comparing the last equality with the previous two lines) vanish upon averaging over a system with periodic boundary conditions. This can be seen also on the continuum by expanding the equation in gradients, as one obtains

$$(\partial_t\phi_{\mathbf{x}})_{\mathrm{P}} = \frac{\epsilon}{2}\Big\{\sin\phi_{\mathbf{x}}\,\partial_x\big[(\partial_x\phi_{\mathbf{x}})^2\big]-\cos\phi_{\mathbf{x}}\,\partial_y\big[(\partial_y\phi_{\mathbf{x}})^2\big]\Big\}. \qquad \text{(B.19)}$$

Note that the terms $\partial_x\big[(\partial_x\phi_{\mathbf{x}})^2\big]$ and $\partial_y\big[(\partial_y\phi_{\mathbf{x}})^2\big]$ corresponding to $S^{(1),\mathrm{loc}}_{\mathrm{P},i}$ and $C^{(1),\mathrm{loc}}_{\mathrm{P},i}$, which could be responsible for pinning, are actually total derivatives and thus they have a vanishing spatial average in the presence of periodic boundary conditions. Accordingly, for the sinusoidal coupling, we conclude that the proactive term does not induce local pinning, and thus a possible local pinning must originate from the reactive contribution.

B.2.3 Von Mises coupling

We now consider Eq. (4.3) for the evolution of ϕ_i in the case of the von Mises coupling (see Eq. (4.8)), casting in the form of Eq. (5.7) in order to identify the pinning terms. As discussed in Sec. 5.2.3, we restrict the analysis to the contributions of order $\mathcal{O}(\sigma^0)$, $\mathcal{O}(\sigma)$, and $\mathcal{O}(\sigma^4)$, which are sufficient to understand the emergence of pinning in the model. Within this approximation, the resulting dynamics turns out to contain terms which also appear in the evolution with the sinusoidal and the modulated coupling, as discussed further below.

Reactive term. We write the local equation for the reactive term

$$(\partial_t\phi_i)_{\mathrm{R}} = -\sum_{j\in\mathcal{N}_i}\left[1+\sigma\cos(\phi_i-\vartheta_{ij})+\frac{\sigma^4}{24}\cos^4(\phi_i-\vartheta_{ij})\right]\sin(\phi_i-\phi_j), \qquad \text{(B.20)}$$

as

$$\begin{aligned}(\partial_t\phi_i)_{\mathrm{R}} =& -\sum_{j\in\mathcal{N}_i}\sin(\phi_i-\phi_j)\\ &-\sigma\cos\phi_i\Big[\sin(\phi_i-\phi_{i+\hat{x}})-\sin(\phi_i-\phi_{i-\hat{x}})\Big]\\ &-\sigma\sin\phi_i\Big[\sin(\phi_i-\phi_{i+\hat{y}})-\sin(\phi_i-\phi_{i-\hat{y}})\Big]\\ &-\frac{\sigma^4}{24}\cos^4\phi_i\Big[\sin(\phi_i-\phi_{i+\hat{x}})+\sin(\phi_i-\phi_{i-\hat{x}})\Big]\\ &-\frac{\sigma^4}{24}\sin^4\phi_i\Big[\sin(\phi_i-\phi_{i+\hat{y}})+\sin(\phi_i-\phi_{i-\hat{y}})\Big], \qquad \text{(B.21)}\end{aligned}$$

which, after using the following trigonometric identities,

$$\cos^4 x = \frac{3}{8} + \frac{1}{2}\cos(2x) + \frac{1}{8}\cos(4x) \quad \text{and} \quad \sin^4 x = \frac{3}{8} - \frac{1}{2}\cos(2x) + \frac{1}{8}\cos(4x), \tag{B.22}$$

can be rewritten as in Eq. (5.7), i.e.,

$$(\partial_t \phi_i)_{\mathrm{R}} = C_{\mathrm{R},i}^{(0),\mathrm{loc}} + S_{\mathrm{R},i}^{(1),\mathrm{loc}} \sin\phi_i + C_{\mathrm{R},i}^{(1),\mathrm{loc}} \cos\phi_i + C_{\mathrm{R},i}^{(2),\mathrm{loc}} \cos(2\phi_i) + C_{\mathrm{R},i}^{(4),\mathrm{loc}} \cos(4\phi_i) \tag{B.23}$$

with

$$C_{\mathrm{R},i}^{(0),\mathrm{loc}} = -\left(1 + \frac{\sigma^4}{64}\right) \sum_{j\in\mathcal{N}_i} \sin(\phi_i - \phi_j), \tag{B.24}$$

$$S_{\mathrm{R},i}^{(1),\mathrm{loc}} = -\sigma\left[\sin(\phi_i - \phi_{i+\hat{y}}) - \sin(\phi_i - \phi_{i-\hat{y}})\right], \tag{B.25}$$

$$C_{\mathrm{R},i}^{(1),\mathrm{loc}} = -\sigma\left[\sin(\phi_i - \phi_{i+\hat{x}}) - \sin(\phi_i - \phi_{i-\hat{x}})\right], \tag{B.26}$$

$$C_{\mathrm{R},i}^{(2),\mathrm{loc}} = \frac{\sigma^4}{48}\left[\sin(\phi_i - \phi_{i+\hat{y}}) + \sin(\phi_i - \phi_{i-\hat{y}}) - \sin(\phi_i - \phi_{i+\hat{x}}) - \sin(\phi_i - \phi_{i-\hat{x}})\right], \tag{B.27}$$

$$C_{\mathrm{R},i}^{(4),\mathrm{loc}} = -\frac{\sigma^4}{192} \sum_{j\in\mathcal{N}_i} \sin(\phi_i - \phi_j). \tag{B.28}$$

Note that apart from a prefactor, $C_{\mathrm{R},i}^{(1),\mathrm{loc}}$ and $S_{\mathrm{R},i}^{(1),\mathrm{loc}}$ have the same expression as for the sinusoidal coupling, see Eqs. (B.13) and (B.14)). Moreover, upon averaging over the system with periodic boundary conditions, only the coefficients $C_{\mathrm{R},i}^{(1),\mathrm{loc}}$ and $S_{\mathrm{R},i}^{(1),\mathrm{loc}}$ may have non-vanishing mean values and therefore they may contribute to the local pinning. By contrast, the coefficients $C_{\mathrm{R},i}^{(0),\mathrm{loc}}$, $C_{\mathrm{R},i}^{(2),\mathrm{loc}}$, and $C_{\mathrm{R},i}^{(4),\mathrm{loc}}$ have zero mean and thus they only contribute to the diffusive part of the dynamics. Moreover, symmetry properties imply that $S_{\mathrm{R},i}^{(1),\mathrm{loc}} = -C_{\mathrm{R},i}^{(1),\mathrm{loc}}$ when the system is pinned along diagonal directions, as discussed further below in App. B.3.3.

Proactive term. We now focus on the contributions of the proactive terms to the equation of motion, keeping only the terms of order $\mathcal{O}(\sigma)$ and $\mathcal{O}(\sigma^4)$ in $\partial_{\phi_i} g(\phi, \vartheta_{ij})$ (the term of order $\mathcal{O}(\sigma^0)$ does not contribute to $\partial_{\phi_i} g(\phi, \vartheta_{ij})$). Starting from Eq. (4.3) and (4.8), one finds

$$\begin{aligned}(\partial_t \phi_i)_{\mathrm{P}} &= -\sum_{j\in\mathcal{N}_i}\left[\sigma\sin(\phi_i - \vartheta_{ij}) + \frac{\sigma^4}{6}\sin(\phi_i - \vartheta_{ij})\cos^3(\phi_i - \vartheta_{ij})\right]\cos(\phi_i - \phi_j) \\ &= -\sum_{j\in\mathcal{N}_i}\left[\sigma\sin(\phi_i - \vartheta_{ij}) + \frac{\sigma^4}{24}\sin\big(2(\phi_i - \vartheta_{ij})\big) + \frac{\sigma^4}{48}\sin\big(4(\phi_i - \vartheta_{ij})\big)\right]\cos(\phi_i - \phi_j).\end{aligned} \tag{B.29}$$

This expression can eventually be rewritten as in Eq. (5.7), i.e.,

$$(\partial_t \phi_i)_{\mathrm{P}} = S_{\mathrm{P},i}^{(1),\mathrm{loc}} \sin\phi_i + C_{\mathrm{P},i}^{(1),\mathrm{loc}} \cos\phi_i + S_{\mathrm{P},i}^{(2),\mathrm{loc}} \sin(2\phi_i) + S_{\mathrm{P},i}^{(4),\mathrm{loc}} \sin(4\phi_i), \tag{B.30}$$

where

$$S_{\mathrm{P},i}^{(1),\mathrm{loc}} = -\sigma\Big[\cos(\phi_i - \phi_{i+\hat{x}}) - \cos(\phi_i - \phi_{i-\hat{x}})\Big], \tag{B.31}$$

$$C_{\mathrm{P},i}^{(1),\mathrm{loc}} = \sigma\Big[\cos(\phi_i - \phi_{i+\hat{y}}) - \cos(\phi_i - \phi_{i-\hat{y}})\Big], \tag{B.32}$$

$$S_{\mathrm{P},i}^{(2),\mathrm{loc}} = -\frac{\sigma^4}{24}\Big[\cos(\phi_i - \phi_{i+\hat{x}}) + \cos(\phi_i - \phi_{i-\hat{x}}) - \cos(\phi_i - \phi_{i+\hat{y}}) - \cos(\phi_i - \phi_{i-\hat{y}})\Big], \tag{B.33}$$

$$S_{\mathrm{P},i}^{(4),\mathrm{loc}} = -\frac{\sigma^4}{48}\Big[\cos(\phi_i - \phi_{i+\hat{x}}) + \cos(\phi_i - \phi_{i-\hat{x}}) + \cos(\phi_i - \phi_{i+\hat{y}}) + \cos(\phi_i - \phi_{i-\hat{y}})\Big]. \tag{B.34}$$

In particular, the coefficients $S_{\mathrm{P},i}^{(1),\mathrm{loc}}$ and $C_{\mathrm{P},i}^{(1),\mathrm{loc}}$, up to an overall prefactor, have the same expression as for the sinusoidal coupling, see Eq. (B.18). As discussed in App. B.2.2, these terms vanish upon averaging over the whole system with periodic boundary conditions. They do not contribute to the average local pinning.

By contrast, the coefficients $S_{\mathrm{P},i}^{(2),\mathrm{loc}}$ and $S_{\mathrm{P},i}^{(4),\mathrm{loc}}$ are not constrained to vanish by the same argument and can contribute to pinning. We verified numerically that they both contribute to pinning along lattice directions. However, the magnitude of the $S_{\mathrm{P},i}^{(2),\mathrm{loc}}$ term is smaller than that of $S_{\mathrm{P},i}^{(4),\mathrm{loc}}$ for the parameters considered.

This difference in magnitude is already apparent from their continuum expansions,

$$S_{\mathrm{P},\mathbf{x}}^{(2),\mathrm{loc}} \simeq -\frac{\sigma^4}{24}\Big[(\partial_y \phi_{\mathbf{x}})^2 - (\partial_x \phi_{\mathbf{x}})^2\Big], \tag{B.35}$$

$$S_{\mathrm{P},\mathbf{x}}^{(4),\mathrm{loc}} \simeq -\frac{\sigma^4}{48}\Big[4 - |\nabla \phi_{\mathbf{x}}|^2\Big]. \tag{B.36}$$

In fact, the coefficient $S_{\mathrm{P},\mathbf{x}}^{(2),\mathrm{loc}}$ starts only at second order in spatial gradients and is therefore subleading. Conversely, $S_{\mathrm{P},\mathbf{x}}^{(4),\mathrm{loc}}$ contains a zeroth-order contribution in the gradient expansion, exactly as the proactive contribution of the modulated-coupling kernel discussed in App. B.2.1.

B.3 Global equations for the orientation

In this appendix, we derive the evolution equation for the mean (global) orientation $\phi(t)$, defined in Eq. (5.1). To this end, we rewrite $\partial_t \phi(t) = N^{-1} \sum_i \partial_t \phi_i$ as a sum over lattice bonds, obtaining an expression of the general form reported in Eq. (5.15). For notational simplicity, throughout this appendix we consider only the deterministic part of the dynamics and omit the Gaussian noise term.

B.3.1 Modulated coupling

We consider the reactive (first line) and the proactive (second line) contributions together, and obtain

$$\begin{aligned}\partial_t \phi = &-\frac{\epsilon}{N}\sum_{\langle i,j\rangle}\big[\cos(4(\phi + \delta\tilde{\phi}_i)) - \cos(4(\phi + \delta\tilde{\phi}_j))\big]\sin(\delta\tilde{\phi}_i - \delta\tilde{\phi}_j)\\ &-\frac{4\epsilon}{N}\sum_{\langle i,j\rangle}\big[\sin(4(\phi + \delta\tilde{\phi}_i)) + \sin(4(\phi + \delta\tilde{\phi}_j))\big]\cos(\delta\tilde{\phi}_i - \delta\tilde{\phi}_j),\end{aligned} \tag{B.37}$$

where we expressed the local value ϕ_i of the phase in terms of its deviation $\delta\tilde{\phi}_i$ from the global orientation ϕ, according to Eq. (5.5). After using standard trigonometric identities, Eq. (B.37) can

be rearranged into the form

$$\begin{aligned}\partial_t\phi = &-2\epsilon\cos(4\phi)\Big\langle\,\big[\cos(4\delta\tilde{\phi}_i)-\cos(4\delta\tilde{\phi}_j)\big]\sin(\delta\tilde{\phi}_i-\delta\tilde{\phi}_j)\Big\rangle\\ &+2\epsilon\sin(4\phi)\Big\langle\,\big[\sin(4\delta\tilde{\phi}_i)-\sin(4\delta\tilde{\phi}_j)\big]\sin(\delta\tilde{\phi}_i-\delta\tilde{\phi}_j)\Big\rangle\\ &-8\epsilon\cos(4\phi)\Big\langle\,\big[\sin(4\delta\tilde{\phi}_i)+\sin(4\delta\tilde{\phi}_j)\big]\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_j)\Big\rangle\\ &-8\epsilon\sin(4\phi)\Big\langle\,\big[\cos(4\delta\tilde{\phi}_i)+\cos(4\delta\tilde{\phi}_j)\big]\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_j)\Big\rangle,\end{aligned} \tag{B.38}$$

where the averages $\langle\cdots\rangle$ are taken over all the lattice bonds at a given time, i.e., $\langle\dots\rangle = \frac{1}{2N}\sum_{\langle ij\rangle}(\dots)$. Note that the first two lines of the expression above originate from the *reactive* part of the dynamics, while the last two stem from the *proactive* term.

By collecting the coefficients of $\sin(4\phi)$ and $\cos(4\phi)$, Eq. (B.38) can be written in the form of Eq. (5.15) as

$$\partial_t\phi = \big(S_{\mathrm{R}}^{(4),\mathrm{glob}} + S_{\mathrm{P}}^{(4),\mathrm{glob}}\big)\sin(4\phi) + \big(C_{\mathrm{R}}^{(4),\mathrm{glob}} + C_{\mathrm{P}}^{(4),\mathrm{glob}}\big)\cos(4\phi) \tag{B.39}$$

where the coefficients $S_{\mathrm{R,P}}^{(4),\mathrm{glob}}$ and $C_{\mathrm{R,P}}^{(4),\mathrm{glob}}$ are defined as

$$C_{\mathrm{R}}^{(4),\mathrm{glob}} = -2\epsilon\Big\langle\,\big[\cos(4\delta\tilde{\phi}_i)-\cos(4\delta\tilde{\phi}_j)\big]\sin(\delta\tilde{\phi}_i-\delta\tilde{\phi}_j)\Big\rangle, \tag{B.40}$$

$$S_{\mathrm{R}}^{(4),\mathrm{glob}} = +2\epsilon\Big\langle\,\big[\sin(4\delta\tilde{\phi}_i)-\sin(4\delta\tilde{\phi}_j)\big]\sin(\delta\tilde{\phi}_i-\delta\tilde{\phi}_j)\Big\rangle, \tag{B.41}$$

$$C_{\mathrm{P}}^{(4),\mathrm{glob}} = -8\epsilon\Big\langle\,\big[\sin(4\delta\tilde{\phi}_i)+\sin(4\delta\tilde{\phi}_j)\big]\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_j)\Big\rangle, \tag{B.42}$$

$$S_{\mathrm{P}}^{(4),\mathrm{glob}} = -8\epsilon\Big\langle\,\big[\cos(4\delta\tilde{\phi}_i)+\cos(4\delta\tilde{\phi}_j)\big]\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_j)\Big\rangle. \tag{B.43}$$

Symmetry considerations. The dynamics of the system is generically invariant under a global $\mathbb{Z}_4$, i.e., under the simultaneous transformation $\phi_i \mapsto \phi_i + n\pi/4$ at all lattice sites i, with $n \in \{0,1,2,3\}$. Once the pinning direction Φ has been selected spontaneously by the system as discussed in Sec. 5.2.1, the dynamics remains invariant under the global reflection symmetry $\mathcal{R} : \delta\phi \to -\delta\phi, \delta\tilde{\phi}_i \to -\delta\tilde{\phi}_i$ done simultaneously at all lattice sites i. As a consequence, the fluctuation-induced coefficients can be classified according to their parity under $\mathcal{R}$. In particular, $C_{\mathrm{R}}^{(4),\mathrm{glob}}$ in Eq. (B.40) and $C_{\mathrm{P}}^{(4),\mathrm{glob}}$ in Eq. (B.42) are odd under $\mathcal{R}$ and therefore they vanish identically, whereas $S_{\mathrm{R}}^{(4),\mathrm{glob}}$ in Eq. (B.41) and $S_{\mathrm{P}}^{(4),\mathrm{glob}}$ in Eq. (B.43) are even and thus they might be finite. Accordingly, the only term in Eq. (B.39) which survives is the one proportional to $\sin(4\phi)$. The overall sign of the corresponding prefactor (determined via numerical simulations, see Sec. 4.3) selects the pinned orientations of the global direction and locks the system either to $\Phi = 0 + n\pi/2$ or to $\Phi = \pi/4 + n\pi/2$, with $n \in \mathbb{Z}$.

Small-fluctuation expansion. We now expand the coefficients $S_{\mathrm{R}}^{(4),\mathrm{glob}}$ and $S_{\mathrm{P}}^{(4),\mathrm{glob}}$ in Eqs. (B.41) and (B.43), respectively, assuming small fluctuations $\delta\tilde{\phi}_i$ and then consider formally to the continuum limit by introducing the field $\delta\tilde{\phi}_{\mathbf{x}}(t)$ instead of $\delta\tilde{\phi}_i$. For the reactive contribution, we obtain

$$(\partial_t\phi)_{\mathrm{R}} = 4\epsilon\,\big\langle|\nabla\delta\tilde{\phi}_{\mathbf{x}}|^2\big\rangle\,\sin(4\phi), \tag{B.44}$$

which is actually the expression found in Ref. [92]. For the proactive contribution, one finds, instead,

$$(\partial_t\phi)_{\rm P} = -4\epsilon\sin(4\phi)\Big[4 - 32\,\langle\delta\tilde{\phi}_{\mathbf{x}}^2\rangle - \langle|\nabla\delta\tilde{\phi}_{\mathbf{x}}|^2\rangle\Big]. \tag{B.45}$$

Note that, differently from the dynamics of the local phase fluctuation $\delta\phi_{\mathbf{x}}$ on the continuum reported in Eq. (B.10), the spatially averaged onsite fluctuations $\langle\delta\tilde{\phi}_{\mathbf{x}}^2\rangle$ of the field explicitly appear in the evolution equation for the global phase.

B.3.2 Sinusoidal coupling

For the case of the sinusoidal coupling in Eq. (4.7), we find convenient to treat separately the reactive and proactive contributions to the equation of motion in Eq. (4.3), which we then use to determine the equation of motion for the average orientation $\phi(t)$, as done above.

Reactive term. The deterministic contribution to the Langevin dynamics originating from the reactive term is

$$\begin{aligned}(\partial_t\phi)_{\rm R} = -\frac{\epsilon}{N}\sum_i\Big\{&\big[\cos(\phi+\delta\tilde{\phi}_i)+\cos(\phi+\delta\tilde{\phi}_{i+\hat{x}})\big]\sin(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{x}})\\ &+\big[\sin(\phi+\delta\tilde{\phi}_i)+\sin(\phi+\delta\tilde{\phi}_{i+\hat{y}})\big]\sin(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{y}})\Big\},\end{aligned} \tag{B.46}$$

where, as done in the previous sections, the local value ϕ_i of the phase is expressed in terms of its deviation $\delta\tilde{\phi}_i$ from the global orientation ϕ, according to Eq. (5.5). Applying trigonometric identities, we obtain

$$\begin{aligned}(\partial_t\phi)_{\rm R} = -\frac{\epsilon}{N}\sum_i\Big\{&\big[(\cos\delta\tilde{\phi}_i+\cos\delta\tilde{\phi}_{i+\hat{x}})\cos\phi-(\sin\delta\tilde{\phi}_i+\sin\delta\tilde{\phi}_{i+\hat{x}})\sin\phi\big]\sin(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{x}})\\ &+\big[(\cos\delta\tilde{\phi}_i+\cos\delta\tilde{\phi}_{i+\hat{y}})\sin\phi+(\sin\delta\tilde{\phi}_i+\sin\delta\tilde{\phi}_{i+\hat{y}})\cos\phi\big]\sin(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{y}})\Big\}.\end{aligned} \tag{B.47}$$

Collecting the coefficients of $\sin\phi$ and $\cos\phi$ and indicating with $\langle\cdots\rangle$ the average over the lattice sites, the result can be written as

$$\begin{aligned}(\partial_t\phi)_{\rm R} = &-\epsilon\Big[\big\langle\big(\cos\delta\tilde{\phi}_i+\cos\delta\tilde{\phi}_{i+\hat{y}}\big)\sin\big(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{y}}\big)\big\rangle\\ &\qquad-\big\langle(\sin\delta\tilde{\phi}_i+\sin\delta\tilde{\phi}_{i+\hat{x}})\sin(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{x}})\big\rangle\Big]\sin\phi\\ &-\epsilon\Big[\big\langle\big(\sin\delta\tilde{\phi}_i+\sin\delta\tilde{\phi}_{i+\hat{y}}\big)\sin\big(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{y}}\big)\big\rangle\\ &\qquad+\big\langle(\cos\delta\tilde{\phi}_i+\cos\delta\tilde{\phi}_{i+\hat{x}})\sin(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{x}})\big\rangle\Big]\cos\phi.\end{aligned} \tag{B.48}$$

Defining the coefficients

$$S_{\rm R}^{(1),\rm glob}(t) = -\epsilon\left[\Big\langle(\cos\delta\tilde{\phi}_i+\cos\delta\tilde{\phi}_{i+\hat{y}})\sin(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{y}})\Big\rangle-\Big\langle(\sin\delta\tilde{\phi}_i+\sin\delta\tilde{\phi}_{i+\hat{x}})\sin(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{x}})\Big\rangle\right], \tag{B.49}$$

$$C_{\rm R}^{(1),\rm glob}(t) = -\epsilon\left[\Big\langle(\sin\delta\tilde{\phi}_i+\sin\delta\tilde{\phi}_{i+\hat{y}})\sin(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{y}})\Big\rangle+\Big\langle(\cos\delta\tilde{\phi}_i+\cos\delta\tilde{\phi}_{i+\hat{x}})\sin(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{x}})\Big\rangle\right], \tag{B.50}$$

the previous equation eventually can be written in the form of Eq. (5.15) as

$$(\partial_t\phi)_{\mathrm{R}} = S_{\mathrm{R}}^{(1),\mathrm{glob}}\sin\phi + C_{\mathrm{R}}^{(1),\mathrm{glob}}\cos\phi. \tag{B.51}$$

Proactive term. The proactive contribution to the deterministic part of the Langevin dynamics, summed over all the lattice, reads

$$\begin{aligned}(\partial_t\phi)_{\mathrm{P}} = -\frac{\epsilon}{N}\sum_i\Big\{&\big[\sin(\phi+\delta\tilde{\phi}_i)-\sin(\phi+\delta\tilde{\phi}_{i+\hat{x}})\big]\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{x}})\\ &+\big[-\cos(\phi+\delta\tilde{\phi}_i)+\cos(\phi+\delta\tilde{\phi}_{i+\hat{y}})\big]\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{y}})\Big\}.\end{aligned} \tag{B.52}$$

By using trigonometric identities this equation can be cast in the form

$$\begin{aligned}(\partial_t\phi)_{\mathrm{P}} = -\frac{\epsilon}{N}\sum_i\Big\{&\big[\sin\phi\,(\cos\delta\tilde{\phi}_i-\cos\delta\tilde{\phi}_{i+\hat{x}})+\cos\phi\,(\sin\delta\tilde{\phi}_i-\sin\delta\tilde{\phi}_{i+\hat{x}})\big]\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{x}})\\ &+\big[-\cos\phi\,(\cos\delta\tilde{\phi}_i-\cos\delta\tilde{\phi}_{i+\hat{y}})+\sin\phi\,(\sin\delta\tilde{\phi}_i-\sin\delta\tilde{\phi}_{i+\hat{y}})\big]\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{y}})\Big\}.\end{aligned} \tag{B.53}$$

Introducing

$$S_{\mathrm{P}}^{(1),\mathrm{glob}}(t) = -\epsilon\left[\left\langle(\cos\delta\tilde{\phi}_i-\cos\delta\tilde{\phi}_{i+\hat{x}})\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{x}})\right\rangle+\left\langle(\sin\delta\tilde{\phi}_i-\sin\delta\tilde{\phi}_{i+\hat{y}})\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{y}})\right\rangle\right], \tag{B.54}$$

$$C_{\mathrm{P}}^{(1),\mathrm{glob}}(t) = -\epsilon\left[\left\langle(\sin\delta\tilde{\phi}_i-\sin\delta\tilde{\phi}_{i+\hat{x}})\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{x}})\right\rangle-\left\langle(\cos\delta\tilde{\phi}_i-\cos\delta\tilde{\phi}_{i+\hat{y}})\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{y}})\right\rangle\right], \tag{B.55}$$

the proactive contribution in Eq. (B.53) can eventually be written as

$$(\partial_t\phi)_{\mathrm{P}} = S_{\mathrm{P}}^{(1),\mathrm{glob}}\sin\phi + C_{\mathrm{P}}^{(1),\mathrm{glob}}\cos\phi. \tag{B.56}$$

Finally, combining this proactive contribution with the reactive one in Eq. (B.51) one finds the complete equation for the mean orientation of the sinusoidal coupling in the general form of Eq. (5.15):

$$\partial_t\phi = (\partial_t\phi)_{\mathrm{R}} + (\partial_t\phi)_{\mathrm{P}} = (S_{\mathrm{R}}^{(1),\mathrm{glob}} + S_{\mathrm{P}}^{(1),\mathrm{glob}})\sin\phi + (C_{\mathrm{R}}^{(1),\mathrm{glob}} + C_{\mathrm{P}}^{(1),\mathrm{glob}})\cos\phi. \tag{B.57}$$

Symmetry considerations. We now analyze the coefficients obtained above from the perspective of the symmetries of the system, as done above for the modulated coupling, in order to show that the $\mathbb{Z}_4$ symmetry of the square lattice is preserved by the equations.

If the ordered state is pinned along $\Phi = 0$, one has $\phi = \delta\phi$ (see Eq. (5.4)), and the symmetry transformation is given by the global reflection $\mathcal{R}_y : \delta\phi \to -\delta\phi$ and $\delta\tilde{\phi}_i \to -\delta\tilde{\phi}_i$, combined with a spatial reflection $y \to -y$. Under this transformation, the coefficients $S_R^{(1),\mathrm{glob}}$ and $S_P^{(1),\mathrm{glob}}$ in Eqs. (B.49) and (B.54), respectively, are even, and therefore their averages might be different from zero. By contrast, $C_R^{(1),\mathrm{glob}}$ and $C_P^{(1),\mathrm{glob}}$ in Eqs. (B.50) and (B.55), respectively, are odd and thus their spatial averages vanish. As a result, the evolution of the mean orientation prescribed by Eq. (B.57) involves only the pinning term of the form $-r^{\mathrm{glob}}\sin\delta\phi$, with $r^{\mathrm{glob}} = S_R^{(1),\mathrm{glob}}$ in the purely reactive case, and $r^{\mathrm{glob}} = S_R^{(1),\mathrm{glob}} + S_P^{(1),\mathrm{glob}}$ for the dynamics which involves both the

reactive and the proactive terms. The numerical simulations discussed in Sec. 5.2.5 indicate that $r^{\mathrm{glob}} > 0$ in both cases.

If, instead, the system is pinned along $\Phi = \pi/2$, one has $\delta\phi = \phi - \pi/2$ (see Eq. (5.4)). The symmetry transformation is given by the global reflection $\mathcal{R}_x : \delta\phi \to -\delta\phi$ and $\delta\tilde{\phi}_i \to -\delta\tilde{\phi}_i$, now combined with $x \to -x$. Under $\mathcal{R}_x$, the non-vanishing contributions come from $C_R^{(1),\mathrm{glob}}$ and $C_P^{(1),\mathrm{glob}}$, because it is easy to check that $S_R^{(1),\mathrm{glob}}$ and $S_P^{(1),\mathrm{glob}}$ in Eqs. (B.49) and (B.54) are now odd and therefore vanish. The pinning term for the total dynamics then takes the form $-r^{\mathrm{glob}} \sin\delta\phi$, using $\cos(\pi/2 + \delta\phi) = -\sin\delta\phi$, with $r^{\mathrm{glob}} = -(C_R^{(1),\mathrm{glob}} + C_P^{(1),\mathrm{glob}})$. Moreover, a rotation of the lattice by $\pi/2$ (counterclockwise) maps $x \to y$ and $y \to -x$. Under this transformation, one finds that $S_R^{(1),\mathrm{glob}} \leftrightarrow -C_R^{(1),\mathrm{glob}}$ (compare Eqs. (B.49) and (B.50)) and similarly $S_P^{(1),\mathrm{glob}} \leftrightarrow -C_P^{(1),\mathrm{glob}}$ (compare Eqs. (B.54) and (B.55)). Accordingly, the pinning strength r^{glob} at $\Phi = \pi/2$ has the same sign as r^{glob} at $\Phi = 0$. Repeating the same argument for $\Phi = \pi$ and $\Phi = 3\pi/2$, one concludes that the pinning coefficient r^{glob} has the $\mathbb{Z}_4$ symmetry of the square lattice, and it leads always to a pinning along the lattice directions. The same symmetry argument hold for the local counterparts for these coefficients, as anticipated in App. B.2.2.

Equality of $S_{\mathrm{R}}^{(1),\mathrm{glob}}$ and $S_{\mathrm{P}}^{(1),\mathrm{glob}}$ in the presence of periodic boundary conditions. We show here that $S_{\mathrm{R}}^{(1),\mathrm{glob}}(t)$ in Eq. (B.49) and $S_{\mathrm{P}}^{(1),\mathrm{glob}}(t)$ in Eq. (B.54) are actually equal when periodic boundary conditions are assumed on the lattice. We first consider the vertical contribution, i.e., the one which involves the sites i and $i + \hat{y}$, corresponding to the first average in square brackets in Eq. (B.49) for $S_{\mathrm{R}}^{(1),\mathrm{glob}}(t)$ and to the second one in Eq. (B.54) for $S_{\mathrm{P}}^{(1),\mathrm{glob}}(t)$. Using elementary trigonometric identities, one finds that the difference $(\cos\delta\tilde{\phi}_i + \cos\delta\tilde{\phi}_{i+\hat{y}})\sin(\delta\tilde{\phi}_i - \delta\tilde{\phi}_{i+\hat{y}}) - (\sin\delta\tilde{\phi}_i - \sin\delta\tilde{\phi}_{i+\hat{y}})\cos(\delta\tilde{\phi}_i - \delta\tilde{\phi}_{i+\hat{y}})$ between the quantities which are then averaged in the two contributions is equal to $\sin\delta\tilde{\phi}_i - \sin\delta\tilde{\phi}_{i+\hat{y}}$. Similarly, for the remaining horizontal contribution, i.e., the one which involves the sites i and $i + \hat{x}$ (the second average in square brackets in Eq. (B.49) for $S_{\mathrm{R}}^{(1),\mathrm{glob}}(t)$ and to the first one in Eq. (B.54) for $S_{\mathrm{P}}^{(1),\mathrm{glob}}(t)$), one finds that the difference $-(\sin\delta\tilde{\phi}_i + \sin\delta\tilde{\phi}_{i+\hat{x}})\sin(\delta\tilde{\phi}_i - \delta\tilde{\phi}_{i+\hat{x}}) - (\cos\delta\tilde{\phi}_i - \cos\delta\tilde{\phi}_{i+\hat{x}})\cos(\delta\tilde{\phi}_i - \delta\tilde{\phi}_{i+\hat{x}})$ is actually equal to $\cos\delta\tilde{\phi}_i - \cos\delta\tilde{\phi}_{i+\hat{x}}$. Accordingly, the difference $S_R^{(1),\mathrm{glob}}(t) - S_P^{(1),\mathrm{glob}}(t)$ between the two coefficients can be written as $S_R^{(1),\mathrm{glob}}(t) - S_P^{(1),\mathrm{glob}}(t) = -\epsilon\big[\langle\sin\delta\tilde{\phi}_i - \sin\delta\tilde{\phi}_{i+\hat{y}}\rangle + \langle\cos\delta\tilde{\phi}_i - \cos\delta\tilde{\phi}_{i+\hat{x}}\rangle\big]$. Under periodic boundary conditions, the spatial average $\langle\cdots\rangle$ of a discrete gradient — such as those on the r.h.s. of the previous equation — vanishes identically, i.e., $\sum_i \big[f_i - f_{i+\hat{\mu}}\big] = 0$, for $\mu \in \{x, y\}$. Accordingly, we conclude that $S_R^{(1),\mathrm{glob}}(t) = S_P^{(1),\mathrm{glob}}(t)$ for periodic boundary conditions. Similar equality holds for $C_R^{(1),\mathrm{glob}}(t)$ and $C_P^{(1),\mathrm{glob}}(t)$.

B.3.3 Von Mises coupling

As done above, also for the case of the von Mises coupling in Eq. (4.8), we find convenient to treat separately the reactive and proactive contributions to the equation of motion in Eq. (4.3), which we then use to determine the equation of motion for the average orientation $\phi(t)$. In particular, we consider the orders $\mathcal{O}(\sigma)$ and $\mathcal{O}(\sigma^4)$ of the expansion of the interaction kernel in increasing powers of σ. The order $\mathcal{O}(\sigma^0)$ does not appear in the equations for the mean orientation, because it is reciprocal.

Reactive term. The deterministic contribution to the Langevin dynamics in Eq. (4.3) originating

from the reactive part alone is

$$\begin{aligned}(\partial_t\phi)_{\mathrm{R}} = &-\frac{1}{N}\sum_i \Bigg\{\sigma\left[\cos(\phi+\delta\tilde{\phi}_i)+\cos(\phi+\delta\tilde{\phi}_{i+\hat{x}})\right] \\ &\qquad +\frac{\sigma^4}{24}\left[\cos^4(\phi+\delta\tilde{\phi}_i)-\cos^4(\phi+\delta\tilde{\phi}_{i+\hat{x}})\right]\Bigg\}\sin\left(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{x}}\right) \\ &-\frac{1}{N}\sum_i \Bigg\{\sigma\left[\sin(\phi+\delta\tilde{\phi}_i)+\sin(\phi+\delta\tilde{\phi}_{i+\hat{y}})\right] \\ &\qquad +\frac{\sigma^4}{24}\left[\sin^4(\phi+\delta\tilde{\phi}_i)-\sin^4(\phi+\delta\tilde{\phi}_{i+\hat{y}})\right]\Bigg\}\sin\left(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{y}}\right).\end{aligned} \tag{B.58}$$

where, as done in the previous sections, the local value ϕ_i of the phase is expressed in terms of its deviation $\delta\tilde{\phi}_i$ from the global orientation ϕ, according to Eq. (5.5). Using the trigonometric identities in Eq. (B.22), we can rewrite this equation in the form of Eq. (5.15), i.e.,

$$\begin{aligned}(\partial_t\phi)_{\mathrm{R}} = S_{\mathrm{R}}^{(1),\mathrm{glob}}\sin\phi + C_{\mathrm{R}}^{(1),\mathrm{glob}}\cos\phi + S_{\mathrm{R}}^{(2),\mathrm{glob}}\sin(2\phi) \\ + C_{\mathrm{R}}^{(2),\mathrm{glob}}\cos(2\phi) + S_{\mathrm{R}}^{(4),\mathrm{glob}}\sin(4\phi) + C_{\mathrm{R}}^{(4),\mathrm{glob}}\cos(4\phi).\end{aligned} \tag{B.59}$$

where $S_{\mathrm{R}}^{(1),\mathrm{glob}}$ and $C_{\mathrm{R}}^{(1),\mathrm{glob}}$ have, up to a factor, the same expressions as those for the sinusoidal interaction (see Eqs. (B.49) and (B.50), respectively), i.e.,

$$S_{\mathrm{R}}^{(1),\mathrm{glob}} = -\sigma\left[\left\langle(\cos\delta\tilde{\phi}_i-\cos\delta\tilde{\phi}_{i+\hat{x}})\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{x}})\right\rangle+\left\langle(\sin\delta\tilde{\phi}_i-\sin\delta\tilde{\phi}_{i+\hat{y}})\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{y}})\right\rangle\right], \tag{B.60}$$

$$C_{\mathrm{R}}^{(1),\mathrm{glob}} = -\sigma\left[\left\langle(\sin\delta\tilde{\phi}_i-\sin\delta\tilde{\phi}_{i+\hat{x}})\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{x}})\right\rangle-\left\langle(\cos\delta\tilde{\phi}_i-\cos\delta\tilde{\phi}_{i+\hat{y}})\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{y}})\right\rangle\right], \tag{B.61}$$

while

$$\begin{aligned}S_{\mathrm{R}}^{(2),\mathrm{glob}} = \frac{\sigma^4}{48}\Big\langle&\left[\sin(2\delta\tilde{\phi}_i)-\sin(2\delta\tilde{\phi}_{i+\hat{x}})\right]\sin\left(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{x}}\right) \\ &-\left[\sin(2\delta\tilde{\phi}_i)-\sin(2\delta\tilde{\phi}_{i+\hat{y}})\right]\sin\left(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{y}}\right)\Big\rangle,\end{aligned} \tag{B.62}$$

$$\begin{aligned}C_{\mathrm{R}}^{(2),\mathrm{glob}} = \frac{\sigma^4}{48}\Big\langle&-\left[\cos(2\delta\tilde{\phi}_i)-\cos(2\delta\tilde{\phi}_{i+\hat{x}})\right]\sin\left(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{x}}\right) \\ &+\left[\cos(2\delta\tilde{\phi}_i)-\cos(2\delta\tilde{\phi}_{i+\hat{y}})\right]\sin\left(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{y}}\right)\Big\rangle,\end{aligned} \tag{B.63}$$

$$S_{\mathrm{R}}^{(4),\mathrm{glob}} = \frac{\sigma^4}{96}\Big\langle\left[\sin(4\delta\tilde{\phi}_i)-\sin(4\delta\tilde{\phi}_j)\right]\sin\left(\delta\tilde{\phi}_i-\delta\tilde{\phi}_j\right)\Big\rangle, \tag{B.64}$$

$$C_{\mathrm{R}}^{(4),\mathrm{glob}} = -\frac{\sigma^4}{96}\Big\langle\left[\cos(4\delta\tilde{\phi}_i)-\cos(4\delta\tilde{\phi}_j)\right]\sin\left(\delta\tilde{\phi}_i-\delta\tilde{\phi}_j\right)\Big\rangle. \tag{B.65}$$

the expression of $S_{\mathrm{R}}^{(4),\mathrm{glob}}$ and $C_{\mathrm{P}}^{(R),\mathrm{glob}}$ reported here are the same, up to a factor, as the corresponding ones for the modulated coupling, reported in Eqs. (B.41) and (B.40), respectively.

Symmetry considerations. If the system is pinned along $\Phi = \pi/4$ (the reactive dynamics leads to pinning to diagonals as discussed in Sec. 5.2.3), the stationary state is invariant under the diagonal reflection $\mathcal{R}_d$: $\delta\phi \to -\delta\phi$, $\delta\tilde{\phi}_i \to -\delta\tilde{\phi}_i$, combined with $x \leftrightarrow y$. Under this symmetry, one finds that $S_{\mathrm{R}}^{(1),\mathrm{glob}} \leftrightarrow -C_{\mathrm{R}}^{(1),\mathrm{glob}}$, (compare Eqs. (B.60) and (B.61)) from which it follows that $S_{\mathrm{R}}^{(1),\mathrm{glob}} = -C_{\mathrm{R}}^{(1),\mathrm{glob}}$. Numerically, we further find that $S_{\mathrm{R}}^{(1),\mathrm{glob}} < 0$ (see Tab. 5.2) Accordingly, the first two terms on the r.h.s. of Eq. (B.59) can be written as $S_{\mathrm{R}}^{(1),\mathrm{glob}}(\sin\phi - \cos\phi) = \sqrt{2}\, S_{\mathrm{R}}^{(1),\mathrm{glob}} \sin(\phi - \pi/4)$, and expanding in $\delta\phi$ around $\phi = \pi/4 + \delta\phi$, one obtains $\sqrt{2}\, S_{\mathrm{R}}^{(1),\mathrm{glob}}\, \delta\phi$, which contributes to a restoring force toward $\delta\phi = 0$, consistently with the stability of the diagonal pinned direction. A similar symmetry argument applies for the corresponding local coefficients $S_{\mathrm{R},i}^{(1),\mathrm{loc}}$ and $C_{\mathrm{R},i}^{(1),\mathrm{loc}}$ discussed in App. B.2.3, which yields the relation $S_{\mathrm{R},i}^{(1),\mathrm{loc}} = -C_{\mathrm{R},i}^{(1),\mathrm{loc}}$ anticipated in the main text and thus to Eq. (5.30).

For the second harmonic, the same symmetry implies that the term $S_{\mathrm{R}}^{(2),\mathrm{glob}}$ in Eq. (B.59) is odd under $\mathcal{R}_d$ and therefore vanishes in the diagonal pinned state. By contrast, $C_{\mathrm{R}}^{(2),\mathrm{glob}}$ is even under $\mathcal{R}_d$ and is not constrained to vanish by symmetry. Expanding the corresponding contribution around $\phi = \pi/4 + \delta\phi$, one finds $C_{\mathrm{R}}^{(2),\mathrm{glob}} \cos(2\phi) = C_{\mathrm{R}}^{(2),\mathrm{glob}} \cos(\pi/2 + 2\delta\phi) \simeq -2C_{\mathrm{R}}^{(2),\mathrm{glob}}\delta\phi$, which therefore contributes to the global relaxation rate. We find numerically that $C_{\mathrm{R}}^{(2),\mathrm{glob}} > 0$ but we not report this contribution in Tab. 5.2, since we find numerically that it is much smaller than the one arising from the fourth harmonic.

We now turn to the fourth-harmonic terms. In this case, the diagonal reflection $\mathcal{R}_d$ reduces to the same symmetry argument used for the modulated-coupling kernel in App. B.3.1, since the two spatial directions enter on the same footing. As a consequence, $C_{\mathrm{R}}^{(4),\mathrm{glob}}$ is odd under the symmetry and vanishes, whereas $S_{\mathrm{R}}^{(4),\mathrm{glob}}$ is even and can be nonzero. Expanding around $\phi = \pi/4 + \delta\phi$, one has $S_{\mathrm{R}}^{(4),\mathrm{glob}} \sin(4\phi) = S_{\mathrm{R}}^{(4),\mathrm{glob}} \sin(\pi + 4\delta\phi) \simeq -4S_{\mathrm{R}}^{(4),\mathrm{glob}}\delta\phi$, corresponding to a contribution $4S_{\mathrm{R}}^{(4),\mathrm{glob}}$ to the global relaxation rate when $S_{\mathrm{R}}^{(4),\mathrm{glob}} > 0$.

Altogether, the reactive von Mises dynamics pinned along the diagonal direction is therefore characterized, to linear order in $\delta\phi$, by a restoring term with relaxation rate $r^{\mathrm{glob}} = -\sqrt{2}\, S_{\mathrm{R}}^{(1),\mathrm{glob}} + 2C_{\mathrm{R}}^{(2),\mathrm{glob}} + 4S_{\mathrm{R}}^{(4),\mathrm{glob}}$, where the second-harmonic contribution is allowed by symmetry but is numerically subleading for the parameters considered.

We also find that, for the parameters reported in Tab. 5.2, the contribution proportional to $S_{\mathrm{R}}^{(1),\mathrm{glob}}$ is smaller than the one proportional to $S_{\mathrm{R}}^{(4),\mathrm{glob}}$. However, this hierarchy is parameter dependent: $S_{\mathrm{R}}^{(1),\mathrm{glob}}$ originates from the $\mathcal{O}(\sigma)$ part of the kernel, whereas $S_{\mathrm{R}}^{(4),\mathrm{glob}}$ originates from the $\mathcal{O}(\sigma^4)$ contribution. Therefore, in the small-σ regime, the first-harmonic term is expected to become the leading contribution to the global pinning. We did not explore this regime systematically, since in the purely reactive case the pinning becomes too weak to be clearly resolved with the system sizes accessible in our simulations.

Proactive term. The deterministic contribution of the proactive term to the evolution equation

for $\phi(t)$ is given by

$$
\begin{aligned}
(\partial_t\phi)_{\rm P} = &-\frac{1}{N}\sum_i \Bigg\{ \sigma\Big[\sin(\phi+\delta\tilde{\phi}_i)-\sin(\phi+\delta\tilde{\phi}_{i+\hat{x}})\Big] \\
&\quad -\frac{\sigma^4}{6}\Big[\sin(\phi+\delta\tilde{\phi}_i)\cos^3(\phi+\delta\tilde{\phi}_i)+\sin(\phi+\delta\tilde{\phi}_{i+\hat{x}})\cos^3(\phi+\delta\tilde{\phi}_{i+\hat{x}})\Big]\Bigg\}\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{x}}) \\
&+\frac{1}{N}\sum_i \Bigg\{ \sigma\Big[\cos(\phi+\delta\tilde{\phi}_i)-\cos(\phi+\delta\tilde{\phi}_{i+\hat{y}})\Big] \\
&\quad +\frac{\sigma^4}{6}\Big[\cos(\phi+\delta\tilde{\phi}_i)\sin^3(\phi+\delta\tilde{\phi}_i)+\cos(\phi+\delta\tilde{\phi}_{i+\hat{y}})\sin^3(\phi+\delta\tilde{\phi}_{i+\hat{y}})\Big]\Bigg\}\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{y}}),
\end{aligned}
\tag{B.66}
$$

where, as done above, ϕ_i which enters in the definition of $\phi(t)$ is expressed in terms of its deviation $\delta\tilde{\phi}_i$ from the global orientation ϕ, according to Eq. (5.5). Using the trigonometric identities $\sin x\cos^3 x = \frac{1}{4}\sin(2x)+\frac{1}{8}\sin(4x)$ and $\cos x\sin^3 x = \frac{1}{4}\sin(2x)-\frac{1}{8}\sin(4x)$ this equation can be cast in the form of Eq. (5.15), with

$$
\begin{aligned}
(\partial_t\phi)_{\rm P} = &\, S_{\rm P}^{(1),\rm glob}\sin\phi + C_{\rm P}^{(1),\rm glob}\cos\phi + S_{\rm P}^{(2),\rm glob}\sin(2\phi) \\
&+ C_{\rm P}^{(2),\rm glob}\cos(2\phi) + S_{\rm P}^{(4),\rm glob}\sin(4\phi) + C_{\rm P}^{(4),\rm glob}\cos(4\phi).
\end{aligned}
\tag{B.67}
$$

where the first harmonics $S_{\rm P}^{(1),\rm glob}$ and $C_{\rm P}^{(1),\rm glob}$ have, up to a factor, the same expression as those for the sinusoidal interaction reported in Eqs. (B.54) and (B.55), respectively, i.e.,

$$
S_{\rm P}^{(1),\rm glob}(t) = -\sigma\left[\Big\langle(\cos\delta\tilde{\phi}_i-\cos\delta\tilde{\phi}_{i+\hat{x}})\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{x}})\Big\rangle+\Big\langle(\sin\delta\tilde{\phi}_i-\sin\delta\tilde{\phi}_{i+\hat{y}})\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{y}})\Big\rangle\right],
\tag{B.68}
$$

$$
C_{\rm P}^{(1),\rm glob}(t) = -\sigma\left[\Big\langle(\sin\delta\tilde{\phi}_i-\sin\delta\tilde{\phi}_{i+\hat{x}})\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{x}})\Big\rangle-\Big\langle(\cos\delta\tilde{\phi}_i-\cos\delta\tilde{\phi}_{i+\hat{y}})\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{y}})\Big\rangle\right],
\tag{B.69}
$$

while the second harmonics are given by

$$
\begin{aligned}
S_{\rm P}^{(2),\rm glob} = &\frac{\sigma^4}{24}\Big\langle\big[\cos(2\delta\tilde{\phi}_i)+\cos(2\delta\tilde{\phi}_{i+\hat{y}})\big]\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{y}})\Big\rangle \\
&-\frac{\sigma^4}{24}\Big\langle\big[\cos(2\delta\tilde{\phi}_i)+\cos(2\delta\tilde{\phi}_{i+\hat{x}})\big]\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{x}})\Big\rangle,
\end{aligned}
\tag{B.70}
$$

$$
\begin{aligned}
C_{\rm P}^{(2),\rm glob} = &\frac{\sigma^4}{24}\Big\langle\big[\sin(2\delta\tilde{\phi}_i)+\sin(2\delta\tilde{\phi}_{i+\hat{y}})\big]\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{y}})\Big\rangle \\
&-\frac{\sigma^4}{24}\Big\langle\big[\sin(2\delta\tilde{\phi}_i)+\sin(2\delta\tilde{\phi}_{i+\hat{x}})\big]\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_{i+\hat{x}})\Big\rangle.
\end{aligned}
\tag{B.71}
$$

Finally, for the fourth harmonic,

$$
S_{\rm P}^{(4),\rm glob} = -\frac{\sigma^4}{24}\Big\langle\big[\cos(4\delta\tilde{\phi}_i)+\cos(4\delta\tilde{\phi}_j)\big]\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_j)\Big\rangle,
\tag{B.72}
$$

$$
C_{\rm P}^{(4),\rm glob} = -\frac{\sigma^4}{24}\Big\langle\big[\sin(4\delta\tilde{\phi}_i)+\sin(4\delta\tilde{\phi}_j)\big]\cos(\delta\tilde{\phi}_i-\delta\tilde{\phi}_j)\Big\rangle.
\tag{B.73}
$$

As expected, the expression of $S_{\rm P}^{(4),\rm glob}$ and $C_{\rm P}^{(4),\rm glob}$ reported here are the same, up to a factor, as

the corresponding ones for the modulated coupling, reported in Eqs. (B.43) and (B.42), respectively.

Symmetry considerations. If the system is pinned along the lattice direction $\Phi = 0$ (the case realized by the full von Mises dynamics, see Sec. 5.2.3), the relevant symmetry is the reflection $\mathcal{R}_y$ introduced in App. B.3.2. Under this transformation, the proactive coefficients $S_{\mathrm{P}}^{(1),\mathrm{glob}}$, $S_{\mathrm{P}}^{(2),\mathrm{glob}}$, and $S_{\mathrm{P}}^{(4),\mathrm{glob}}$ are even and therefore they may be nonzero. For the fourth harmonic, the argument reduces to the same one discussed in App. B.3.1, since the two spatial directions enter symmetrically in the corresponding expressions.

For the reactive contribution, the coefficients $S_{\mathrm{R}}^{(1),\mathrm{glob}}$, $S_{\mathrm{R}}^{(2),\mathrm{glob}}$, and $S_{\mathrm{R}}^{(4),\mathrm{glob}}$ are likewise even under $\mathcal{R}_y$. Furthermore, as shown in App. B.3.2, one has $S_{\mathrm{R}}^{(1),\mathrm{glob}} = S_{\mathrm{P}}^{(1),\mathrm{glob}}$. Therefore, after expanding the sum of Eqs. (B.59) and (B.67) around $\Phi = 0$ with $\phi = \delta\phi$, the global relaxation rate takes the form $r^{\mathrm{glob}} = -2S_{\mathrm{R}}^{(1),\mathrm{glob}} - 2S_{\mathrm{R}}^{(2),\mathrm{glob}} - 2S_{\mathrm{P}}^{(2),\mathrm{glob}} - 4S_{\mathrm{R}}^{(4),\mathrm{glob}} - 4S_{\mathrm{P}}^{(4),\mathrm{glob}}$. In Eq. (5.41) of the main text, we reported only the first- and fourth-harmonic contributions. The first harmonic is of order $\mathcal{O}(\sigma)$ and is therefore expected to dominate in the small-σ regime, even though, as already discussed in this section, we did not explore this regime numerically. The fourth harmonic is of order $\mathcal{O}(\sigma^4)$ and provides the leading higher-order correction. The second harmonic is also of order $\mathcal{O}(\sigma^4)$; however, we find numerically that its contribution to the pinning is significantly smaller than that of the fourth harmonic for the parameter values considered in this work. For this reason, it has been omitted from the discussion in the main text.

Finally, we comment on the equivalence among all lattice directions. Under a rotation by $\pi/2$, one has $x \to y$ and $y \to -x$. As discussed in App. B.3.2, this implies $S_{\mathrm{R}}^{(1),\mathrm{glob}} \leftrightarrow -C_{\mathrm{R}}^{(1),\mathrm{glob}}$ (compare Eqs. (B.60) and (B.61)) and, similarly, $S_{\mathrm{P}}^{(1),\mathrm{glob}} \leftrightarrow -C_{\mathrm{P}}^{(1),\mathrm{glob}}$ (compare Eqs. (B.68) and (B.69)). The fourth-harmonic terms (Eqs. (B.64), (B.65), (B.72), and (B.73)) are unchanged by this rotation, since their analytical expression does not distinguish between the lattice directions. The second-harmonic coefficients (Eqs. (B.62), (B.63), (B.70), and (B.71)) instead, change sign, i.e., $S_{\mathrm{R}}^{(2),\mathrm{glob}} \to -S_{\mathrm{R}}^{(2),\mathrm{glob}}$, $C_{\mathrm{R}}^{(2),\mathrm{glob}} \to -C_{\mathrm{R}}^{(2),\mathrm{glob}}$, $S_{\mathrm{P}}^{(2),\mathrm{glob}} \to -S_{\mathrm{P}}^{(2),\mathrm{glob}}$, and $C_{\mathrm{P}}^{(2),\mathrm{glob}} \to -C_{\mathrm{P}}^{(2),\mathrm{glob}}$. This sign change is consistent with the equivalence of the pinned lattice directions. For instance, a second-harmonic contribution of the form $(S_{\mathrm{R}}^{(2),\mathrm{glob}} + S_{\mathrm{P}}^{(2),\mathrm{glob}})\sin(2\phi)$ becomes $-(S_{\mathrm{R}}^{(2),\mathrm{glob}}|_{\Phi=0} + S_{\mathrm{P}}^{(2),\mathrm{glob}}|_{\Phi=0})\sin(2\phi)$ after a rotation by $\pi/2$. Expanding around $\Phi = \pi/2$, i.e., $\phi = \pi/2 + \delta\phi$, one has $\sin(2\phi) = \sin(\pi + 2\delta\phi) \simeq -2\delta\phi$. Accordingly, the additional minus sign in the coefficient is exactly compensated by the expansion of $\sin(2\phi)$ around the rotated pinned direction, yielding the same restoring contribution as the one around $\Phi = 0$. This implies that the pinning strength is the same along all lattice directions, as required by the underlying $\mathbb{Z}_4$ symmetry.

Appendix C

Additional material for Chapters 7 and 8

C.1 Low-temperature analysis of the soft-spin BNNNI model

We distinguish between the discrete-spin BNNNI model (dBNNNI), whose fields take the values $\phi_i = \pm 1$, and the soft-spin sBNNNI model considered in this work, for which $\phi_i \in \mathbb{R}$. The dBNNNI model is formally obtained by taking $\lambda \to \infty$ at fixed temperature. The relation between this hard-spin limit and the low-temperature analysis below requires some care, because the latter is performed at fixed finite λ.

An important qualitative difference between the phase diagram in Fig. 8.1 and that of the discrete model studied in Ref. [150] is the extent of the incommensurate (IC) phase. In the soft-spin model, the IC regime remains broad and appears to extend down to comparatively low temperatures, whereas in the discrete model it disappears as $T \to 0$, where the ferromagnetic phase gives way directly to the antiphase. The purpose of the analysis below is not to determine the two-dimensional phase boundaries quantitatively, but to identify a low-temperature mechanism by which continuous field amplitudes at finite λ can stabilize modulated profiles over an interval of frustration strengths.

We consider the auxiliary one-dimensional analogue whose dimensionless energy, consistently with Eq. (8.1), is

$$E_T[\phi] = -\frac{J}{T}\sum_i \phi_i\phi_{i+1} + \frac{\kappa J}{T}\sum_i \phi_i\phi_{i+2} + \sum_i \left[\phi_i^2 + \lambda\left(\phi_i^2 - 1\right)^2\right]. \tag{C.1}$$

Because $1/T$ multiplies only the quadratic interaction terms, the low-temperature behavior is controlled by a balance between the quadratic interaction energy and the quartic confinement. At the level of scaling,

$$\frac{\phi^2}{T} \sim \lambda\phi^4, \tag{C.2}$$

so that

$$\phi^2 \sim \frac{1}{\lambda T}. \tag{C.3}$$

This suggests introducing directly the rescaled field

$$\psi_i = \sqrt{\lambda T}\,\phi_i, \tag{C.4}$$

in terms of which Eq. (C.1) becomes

$$\begin{aligned} E_T[\psi] = \frac{1}{\lambda T^2}\left[-J\sum_i \psi_i\psi_{i+1} + \kappa J\sum_i \psi_i\psi_{i+2} + \sum_i \psi_i^4\right] \\ + \frac{1-2\lambda}{\lambda T}\sum_i \psi_i^2 + \lambda N. \end{aligned} \tag{C.5}$$

For fixed finite λ, the term in square brackets dominates as $T \to 0$. Its prefactor is positive and therefore does not affect the profile selected by minimization. The leading low-temperature variational problem thus reduces to minimizing the reduced functional

$$\mathcal{F}[\psi] = -J\sum_i \psi_i\psi_{i+1} + \kappa J\sum_i \psi_i\psi_{i+2} + \sum_i \psi_i^4. \tag{C.6}$$

The functional $\mathcal{F}$ depends neither on T nor on λ. These parameters determine only the relation between the original and rescaled amplitudes, $\phi_i = \psi_i/\sqrt{\lambda T}$, and the overall scale of the leading energy. In particular, at fixed finite λ, the original field diverges as $T^{-1/2}$, whereas the rescaled field remains finite.

We compare three variational families. The FM state is uniform,

$$\psi_i = m. \tag{C.7}$$

The one-dimensional antiphase is

$$\psi_i = A(-1)^{\lfloor i/2 \rfloor}, \tag{C.8}$$

corresponding to the repeating pattern $++--$. Finally, we consider the multiharmonic IC profile

$$\psi_i = A\sum_{n=0}^{N_h-1} c_{2n+1}\cos[(2n+1)q_\star(\kappa)i], \qquad c_1 = 1, \tag{C.9}$$

with $N_h = 5$. The wavevector is chosen to minimize the quadratic part of Eq. (C.1),

$$q_\star(\kappa) = \arccos\left(\frac{1}{4\kappa}\right), \qquad \kappa \geq \frac{1}{4}. \tag{C.10}$$

The overall amplitude and the coefficients c_3, c_5, c_7, c_9 are optimized variationally. Fixing $q = q_\star$ is an approximation, because nonlinear terms can shift the optimal modulation wavevector.

For a generic profile $\psi_i = Af_i$, the reduced functional per site is

$$\frac{\mathcal{F}}{N} = a_2[f]A^2 + a_4[f]A^4, \tag{C.11}$$

with

$$a_2[f] = -J\langle f_i f_{i+1}\rangle + \kappa J\langle f_i f_{i+2}\rangle, \tag{C.12}$$

$$a_4[f] = \langle f_i^4\rangle. \tag{C.13}$$

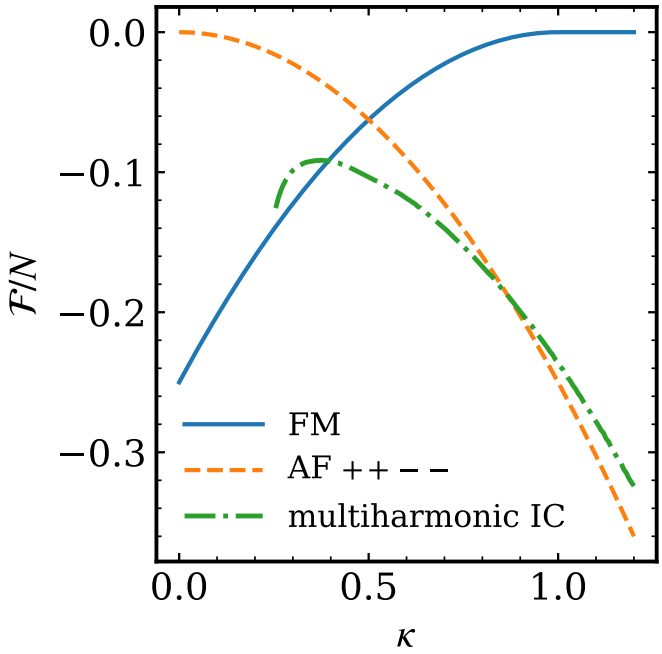


Figure A1: Reduced functional density $\mathcal{F}/N$ of the FM state (solid line), period-four antiphase (dashed line), and multiharmonic IC profile (dash-dotted line). The IC branch is shown for $\kappa \geq 1/4$, where $q_\star$ in Eq. (C.10) is real. Within the variational families considered, it has the lowest leading energy approximately for $0.393 \lesssim \kappa \lesssim 0.869$. Since $\mathcal{F}$ is independent of T and λ, the profiles and crossing points are universal within the leading low-temperature scaling problem.

When $a_2 < 0$, minimization over the amplitude gives

$$A^2 = -\frac{a_2}{2a_4}, \qquad \frac{\mathcal{F}^{\min}}{N} = -\frac{a_2^2}{4a_4}, \tag{C.14}$$

while $A = 0$ otherwise. The FM and antiphase branches are therefore

$$\frac{\mathcal{F}^{\mathrm{FM}}}{N} = -\frac{J^2(1-\kappa)^2}{4}, \qquad \kappa < 1, \tag{C.15}$$

$$\frac{\mathcal{F}^{\mathrm{AF}}}{N} = -\frac{J^2\kappa^2}{4}. \tag{C.16}$$

Figure A1 compares these branches with the optimized multiharmonic one. Within the variational families considered, the IC profile has the lowest leading energy approximately for

$$0.393 \lesssim \kappa \lesssim 0.869. \tag{C.17}$$

This statement concerns the restricted ansatz space and does not establish the global minimizer of Eq. (C.6). Because $\mathcal{F}$ contains neither T nor λ, the three variational profiles and their crossing points are independent of both parameters within this leading scaling problem. The λ dependence enters only through the common prefactor $1/\lambda$ of the leading contribution to the original energy in Eq. (C.5).

Within its limitations, this calculation provides qualitative support for the extended IC region in Fig. 8.1. Although it concerns an auxiliary one-dimensional model and a restricted multiharmonic ansatz, it shows that continuous field amplitudes at finite λ can favor an IC profile over a finite

interval of frustration strengths. This contrasts with the direct low-temperature FM–AF transition reported for the discrete model in Ref. [150]. The discrete result cannot be recovered by simply taking $\lambda \to \infty$ in the reduced problem, which was obtained by considering $T \to 0$ at fixed finite λ. Relating the two limits therefore requires some care, but a detailed analysis lies beyond the scope of this analysis.

C.2 Finite-size scaling of the Ising-like critical line

We determine the Ising-like critical point used in the dynamical analysis for the sBNNNI parameters fixed in Eq. (8.4), with $\kappa = 0.2$. The magnetic susceptibility is computed from the uniform magnetization as

$$\chi(T, L) = \frac{L^2}{T} \left(\langle m_0^2 \rangle - \langle m_0 \rangle^2 \right). \tag{C.18}$$

For every size, the pseudocritical temperature $T_{\mathrm{pc}}(L)$ is obtained from the maximum of $\chi(T, L)$, determined by fitting the data around the peak with a quadratic polynomial. The bulk critical temperature is extracted from

$$T_{\mathrm{pc}}(L) = T_c + aL^{-1/\nu}, \tag{C.19}$$

with the Ising value $\nu = 1$ [165, 166].

As an independent check, we calculate the second-moment correlation length

$$\xi_2 = \frac{1}{2\sin(\pi/L)} \sqrt{\frac{S(\mathbf{0})}{S(\mathbf{q}_{\mathrm{min}})} - 1}, \tag{C.20}$$

where $\mathbf{q}_{\mathrm{min}} = (2\pi/L, 0)$. Close to criticality, the expected finite-size scaling forms are

$$\chi(T, L) = L^{\gamma/\nu} f_\chi \left[(T - T_c) L^{1/\nu} \right], \tag{C.21}$$

and

$$\frac{\xi_2(T, L)}{L} = f_\xi \left[(T - T_c) L^{1/\nu} \right], \tag{C.22}$$

with $\gamma = 7/4$ and $\nu = 1$.

The same value

$$T_c = 0.987(1) \tag{C.23}$$

produces the susceptibility collapse, the correlation-length collapse, and the finite-size extrapolation shown in Fig. A2. The quality of these three tests is consistent with the two-dimensional Ising universality class and fixes the critical benchmark used in Sec. 8.3.2.

C.3 Wavelet-transform conventions

The conceptual role of the wavelet decomposition was discussed in Chapter 7. Here we specify the conventions used in the numerical implementation. We employ orthogonal discrete wavelet transforms with periodic boundary conditions.

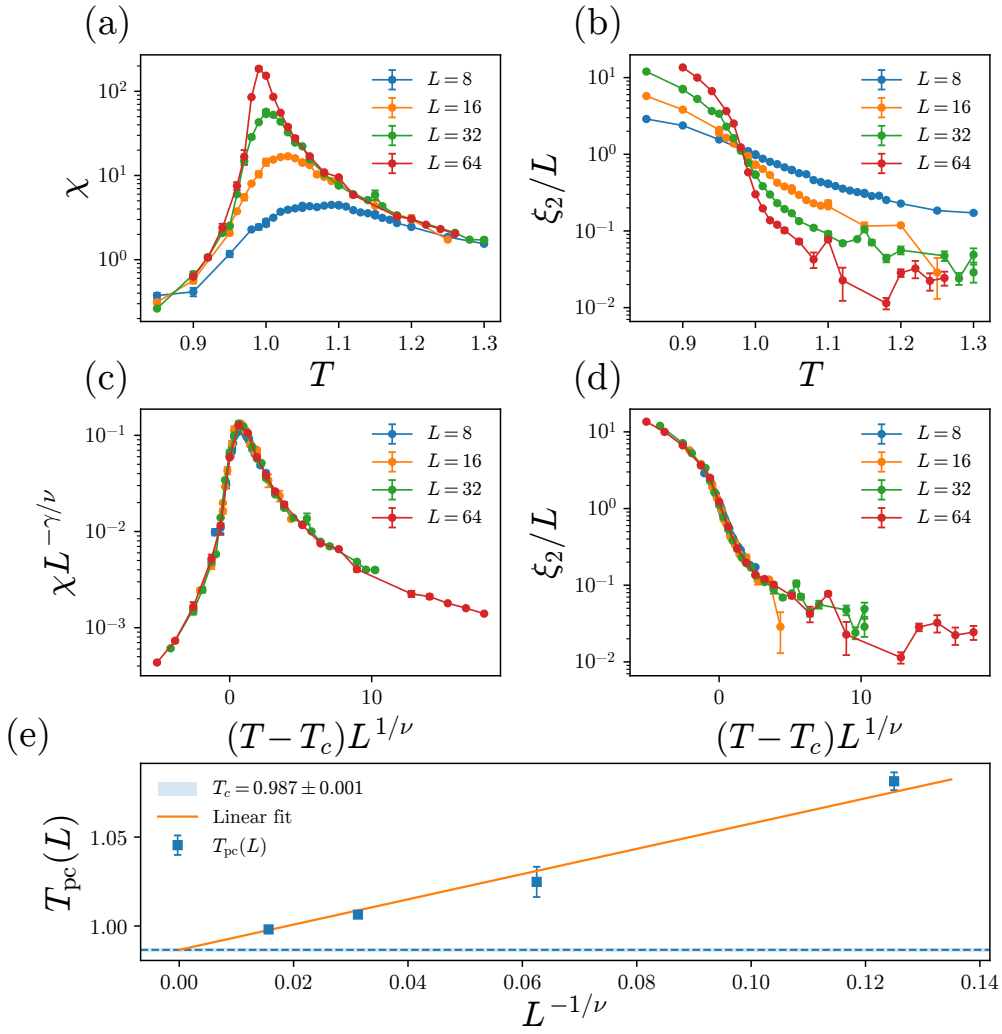


Figure A2: Finite-size scaling analysis for $\kappa_x = \kappa_y = 0.2$, $\kappa_d = 0$, $\lambda = 1$, and $J = 1$. Panels (a) and (b) show the susceptibility and ξ_2/L as functions of temperature. Panels (c) and (d) show the corresponding collapses using the two-dimensional Ising exponents. Panel (e) reports the pseudocritical temperatures and the fit to Eq. (C.19).

In one spatial dimension, the low-pass and high-pass maps act on a field ϕ_{j-1} of size L_{j-1} as

$$\phi_j(x) = \sum_{n=0}^{L_{j-1}-1} g_j(n)\, \phi_{j-1}(2x+n), \tag{C.24}$$

$$\bar{\phi}_j(x) = \sum_{n=0}^{L_{j-1}-1} \bar{g}_j(n)\, \phi_{j-1}(2x+n), \tag{C.25}$$

where $x = 0, \ldots, L_j - 1$, $L_j = L_{j-1}/2$, and all indices are understood modulo L_{j-1}. These filter maps provide the discrete realization of the abstract operators W_j and $\overline{W}_j$ introduced in Eq. (7.2). Their orthogonality and the corresponding inverse reconstruction are given in Eqs. (7.3)–(7.5). The change of variables has unit Jacobian,

$$\mathrm{d}\phi_{j-1} = \mathrm{d}\phi_j\, \mathrm{d}\bar{\phi}_j. \tag{C.26}$$

The filter sums satisfy

$$\sum_n g_j(n) = \sqrt{2}, \qquad \sum_n \bar{g}_j(n) = 0. \tag{C.27}$$

The second equality expresses the zero spatial average of the wavelet sector. At fixed coarse field, conditional updates of the wavelet variables therefore do not modify the uniform component of the reconstructed field. This is why the magnetization cannot be used to monitor the decorrelation of an individual conditional chain in Sec. 8.3.2.

The two-dimensional transform is obtained by applying the one-dimensional filters along the two lattice directions. The low-low component defines the coarse field,

$$\phi_j(x,y) = \sum_{m,n} g_j(m) g_j(n) \phi_{j-1}(2x+m, 2y+n), \tag{C.28}$$

whereas the remaining combinations define three wavelet channels,

$$\bar{\phi}_{j,1}(x,y) = \sum_{m,n} g_j(m) \bar{g}_j(n) \phi_{j-1}(2x+m, 2y+n), \tag{C.29}$$

$$\bar{\phi}_{j,2}(x,y) = \sum_{m,n} \bar{g}_j(m) g_j(n) \phi_{j-1}(2x+m, 2y+n), \tag{C.30}$$

$$\bar{\phi}_{j,3}(x,y) = \sum_{m,n} \bar{g}_j(m) \bar{g}_j(n) \phi_{j-1}(2x+m, 2y+n). \tag{C.31}$$

A field on the lattice Λ_{j-1} has $4L_j^2$ degrees of freedom. After the transform, these are divided into L_j^2 coarse variables and $3L_j^2$ wavelet variables, so the number of degrees of freedom is preserved.

We use Haar filters for the FM, PM, and critical configurations and DB4 filters for the MPM, IC, and AF configurations. Explicit standard filter coefficients can be found in Refs. [106, 107].

C.4 Learning the probability distributions

At every scale $j = 1, \ldots, J$, the training set consists of pairs defined in Eq. (7.53) obtained by applying the wavelet transform to equilibrium microscopic configurations.

C.4.1 Coarsest energy

The WCRG reconstruction starts from the coarsest distribution $p_{\theta_J}(\phi_J)$, whose general exponential form is given in Eq. (7.35). In the present implementation the decomposition is iterated until $L_J = 1$, so that ϕ_J consists of a single degree of freedom. We specialize the linear ansatz in Eq. (7.32) to a one-dimensional local potential expanded on N_s fixed functions s_ℓ. The optimization determines their coefficients $\theta_{J,\ell}$. Following Ref. [167], we choose this finite family to be made of translated sigmoids,

$$s_\ell(u) = \sigma\left(\frac{u - u_\ell}{\Delta_\ell}\right) \quad \text{with} \quad \sigma(z) = \frac{1}{1+e^{-z}}. \tag{C.32}$$

The centers u_ℓ are chosen as N_s equally spaced points covering the empirical support of the field variable. The widths are fixed proportionally to the spacing between consecutive centers, i.e., $\Delta_\ell = (3/2)\,(u_{\ell+1} - u_\ell)$, with the endpoint widths chosen by extending the same spacing convention. More generally, the wavelet decomposition discussed above, which is stopped at $L_J = 1$, could be stopped at a finite coarse size $L_J > 1$. In that case, E_{θ_J} is no longer a one-dimensional potential, but an effective energy defined on the remaining coarse lattice. It could

then be parametrized according to the same "quadratic-plus-local" structure used below for the free-energy terms, which encompasses quadratic terms involving different lattice sites and a local potential involving also higher powers of the local field. In the numerical results presented here, however, we always use system sizes for which the hierarchy reaches $L_J = 1$.

The parameters θ_J are estimated from the empirical distribution of the coarsest field by applying the non-conditional version of the score-matching procedure described in Sec. 7.2.3.

C.4.2 Conditional-energy ansatz

At fixed coarse field ϕ_j, the parametric conditional distribution of the wavelet variables is given by Eq. (7.24). In the present implementation, we first reconstruct the finer field through Eq. (7.5) and then evaluate the linear conditional energy ansatz of Eq. (7.23) on that field. Thus $\bar{E}_{\bar{\theta}_j}$ is a function of the wavelet variables $\bar{\phi}_j$ through the inverse wavelet reconstruction, while the coarse field ϕ_j is kept fixed.

The conditional operator vector $\bar{\Psi}_j$ contains a translationally invariant quadratic part and a local scalar-potential part. The quadratic part is built from products of the reconstructed field at displaced sites. The local part is expanded on the sigmoid functions introduced in Eq. (C.32), with $N_s = 20$. The retained displacements are collected in $\mathcal{R}_j$, which contains first- and second-neighbor displacements on the reconstructed lattice Λ_{j-1}, with periodic boundary conditions. The resulting conditional operator vector is

$$\bar{\Psi}_j(\bar{\phi}_j;\phi_j) = \left(\left\{ \frac{1}{2} \sum_{i\in\Lambda_{j-1}} \phi_{j-1}(i)\,\phi_{j-1}(i+r) \right\}_{r\in\mathcal{R}_j}, \left\{ \sum_{i\in\Lambda_{j-1}} s_\ell\big(\phi_{j-1}(i)\big) \right\}_{\ell=1}^{N_s} \right). \tag{C.33}$$

Here ϕ_{j-1} denotes the reconstructed field. The parameter vector $\bar{\theta}_j$ contains the quadratic coefficients $\{\bar{K}_j(r)\}_{r\in\mathcal{R}_j}$ followed by the local coefficients $\{\bar{\gamma}_{j,\ell}\}_{\ell=1}^{N_s}$, in the same order as the components of $\bar{\Psi}_j$. Although the model is linear in these parameters, the components of $\bar{\Psi}_j$ are nonlinear functions of the reconstructed field.

At fixed ϕ_j, this ansatz defines an effective conditional energy for the wavelet variables $\bar{\phi}_j$. Additive terms depending only on ϕ_j are absorbed into the induced free energy $\bar{F}^{\mathrm{ind}}_{\bar{\theta}_j}$ and therefore do not affect either the conditional probability or the score with respect to $\bar{\phi}_j$.

C.4.3 Score-matching linear system

The general derivation of conditional score matching, including the conditional Fisher divergence, its empirical form, and the resulting quadratic objective, is given in Eqs. (7.63)–(7.68). For the present application, all derivatives act only on the wavelet variables and are evaluated by automatic differentiation of the operator vector in Eq. (C.33) on the training pairs of Eq. (7.53). The empirical matrix $\bar{M}_j$ and vector $\bar{b}_j$ are therefore those defined in Eqs. (7.67) and (7.68). In the numerical implementation, the parameters are obtained from the regularized linear system

$$\big(\bar{M}_j + \rho_j I\big)\,\bar{\theta}_j = \bar{b}_j. \tag{C.34}$$

The regularization strength ρ_j stabilizes the inversion of the empirical matrix. It is chosen in a regime in which moderate variations do not modify the observables of the synthetic configurations

within statistical uncertainty.

C.4.4 Conditional free energy and reconstruction of microscopic energy

The induced free energy $\bar{F}^{\mathrm{ind}}_{\bar{\theta}_j}(\phi_j)$ in Eq. (7.25) is fixed by the normalization of the conditional probability. It is not required for conditional sampling, because ϕ_j is kept fixed during the sampling of $\bar{\phi}_j$. However, it is needed if one wants to reconstruct an explicit microscopic energy functional associated with the learned WCRG hierarchy.

Following the WCRG construction of Ref. [167], we approximate the induced free energy by a finite-dimensional ansatz,

$$\bar{F}^{\mathrm{ind}}_{\bar{\theta}_j}(\phi_j) \simeq \bar{F}_{\alpha_j}(\phi_j) = \alpha_j^{\mathrm{T}} \Phi_j(\phi_j), \tag{C.35}$$

where Φ_j depends only on the coarse field. We use the same quadratic-plus-local structure as for the conditional energy. The retained coarse-lattice displacements are collected in $\mathcal{R}_j^F$, periodic boundary conditions are understood, and the local terms use the sigmoid functions in Eq. (C.32). Accordingly, the operator vector $\Phi_j(\phi_j)$ has the same quadratic-plus-local structure as $\bar{\Psi}_j$ in Eq. (C.33), with the reconstructed field ϕ_{j-1} replaced by the coarse field ϕ_j and $\mathcal{R}_j$ replaced by $\mathcal{R}_j^F$. The coefficient vector α_j contains the quadratic parameters $\{\widetilde{K}_j(r)\}_{r\in\mathcal{R}_j^F}$ followed by the local parameters $\{\widetilde{\gamma}_{j,\ell}\}_{\ell=1}^{N_s}$, in the same order as the components of Φ_j.

The parameters α_j are obtained by matching the coarse-field gradient of the ansatz $\bar{F}_{\alpha_j}$ to the gradient of the induced free energy $\bar{F}^{\mathrm{ind}}_{\bar{\theta}_j}$. Differentiating Eq. (7.25) gives

$$\nabla_{\phi_j} \bar{F}^{\mathrm{ind}}_{\bar{\theta}_j}(\phi_j) = \left\langle \nabla_{\phi_j} \bar{E}_{\bar{\theta}_j}(\bar{\phi}_j;\phi_j) \right\rangle_{\bar{p}_{\bar{\theta}_j}(\bar{\phi}_j|\phi_j)}. \tag{C.36}$$

The average is formally over wavelet variables sampled from the learned conditional distribution $\bar{p}_{\bar{\theta}_j}(\bar{\phi}_j \mid \phi_j)$ at fixed coarse field ϕ_j. In the implementation used in this work, we approximate this gradient-matching problem using the empirical wavelet–coarse pairs of Eq. (7.53). Accordingly, no additional wavelet variables are generated from $\bar{p}_{\bar{\theta}_j}(\bar{\phi}_j \mid \phi_j^{(n)})$ during the free-energy fit.

The free-energy parameters are then estimated by minimizing the gradient-matching loss

$$\mathcal{L}_j^F(\alpha_j) = \left\langle \left| \nabla_{\phi_j} \bar{F}_{\alpha_j}(\phi_j) - \nabla_{\phi_j} \bar{E}_{\bar{\theta}_j}(\bar{\phi}_j;\phi_j) \right|^2 \right\rangle_{\mathrm{data}}. \tag{C.37}$$

Since $\bar{F}_{\alpha_j}$ is linear in α_j, minimizing Eq. (C.37) gives a linear regression problem. Writing $\Phi_{j,a}$ and $\bar{\Psi}_{j,b}$ for the components of the two operator vectors, we define

$$\left(\widetilde{M}_j\right)_{ab} = \left\langle \nabla_{\phi_j}\Phi_{j,a}(\phi_j) \cdot \nabla_{\phi_j}\Phi_{j,b}(\phi_j) \right\rangle_{\mathrm{data}}, \tag{C.38}$$

and

$$\left(\widetilde{B}_j\right)_{ab} = \left\langle \nabla_{\phi_j}\Phi_{j,a}(\phi_j) \cdot \nabla_{\phi_j}\bar{\Psi}_{j,b}(\bar{\phi}_j;\phi_j) \right\rangle_{\mathrm{data}}. \tag{C.39}$$

The fitted parameters are obtained from the regularized normal equations

$$\left(\widetilde{M}_j + \rho_j^F I\right) \alpha_j = \widetilde{B}_j \bar{\theta}_j. \tag{C.40}$$

The parameter ρ_j^F regularizes the empirical regression matrix $\widetilde{M}_j$, in the same way as ρ_j regularizes the score-matching system for the conditional energies.

The dependence of α_j on $\bar{\theta}_j$ simply reflects the fact that the target gradient in Eq. (C.37) is computed from the conditional energy model learned in the previous step. The free-energy fit therefore does not introduce an independent conditional model; it provides an explicit parametrization of the normalization induced by $\bar{E}_{\bar{\theta}_j}$.

The fitted functions $\bar{F}_{\alpha_j}$ can then be inserted into Eq. (7.28) to obtain an explicit approximation to the implicit energy of the learned hierarchy. They are not required for sampling, which instead generates the normalized synthetic measure in Eq. (7.26). Because $\bar{F}_{\alpha_j}$ only approximates $\bar{F}^{\mathrm{ind}}_{\bar{\theta}_j}$, the explicit energy in Eq. (7.28) need not coincide exactly with the implicit synthetic energy in Eq. (7.27).

C.5 Conditional sampling and reconstruction

Once $\bar{\theta}_j$ is known, the conditional distribution is sampled with local Metropolis updates. A single wavelet coefficient, whose index a contains both its spatial and channel labels, is proposed to change according to

$$\bar{\phi}'_{j,a} = \bar{\phi}_{j,a} + \eta, \qquad \eta \sim \mathcal{U}[-\delta_j, \delta_j]. \tag{C.41}$$

The proposal is accepted with probability

$$P_{\mathrm{acc}} = \min\left\{1, \exp[-\Delta \bar{E}_{\bar{\theta}_j}]\right\}. \tag{C.42}$$

During a short warm-up stage, δ_j is adjusted to obtain an acceptance rate close to 50%. The original microscopic configurations are generated with the analogous single-site Metropolis rule using changes of Eq. (8.1).

A synthetic configuration is generated by first sampling ϕ_J, then sampling $\bar{\phi}_j$ conditionally for $j = J, J-1, \ldots, 1$, and applying the inverse wavelet transform after each stage. Only the coarsest and conditional energies are required for this procedure; the conditional free energies do not enter the acceptance ratios.

The choice of local conditional updates should be contrasted with global Langevin-based proposals, such as MALA [168, 169], which have also been used in related learned multiscale sampling contexts [115]. In such methods, the update of the full set of variables is proposed at once, then accepted or rejected. Global proposals are naturally parallelizable in contrast to local Metropolis sweeps, but their acceptance and decorrelation properties may degrade as the number of degrees of freedom increases. For example, optimal scaling results for MALA give a characteristic scaling of order $N^{1/3}$ decorrelation steps for a target distribution of N variables. In two spatial dimensions, where $N = L^2$, this corresponds to $L^{2/3}$.

Here we use local conditional Metropolis updates for a different reason. They allow us to directly measure the decorrelation time of the conditional chain at each reconstruction scale in units of sweeps. The price to pay is that a Monte Carlo sweep consists of successive local proposals and is therefore not a single fully parallel global move.

Bibliography

[1] G. Bandini, D. Venturelli, S. A. M. Loos, A. Jelic, and A. Gambassi, "The XY model with vision cone: non-reciprocal vs. reciprocal interactions", J. Stat. Mech. **2025**, 053205 (2025).

[2] G. Bandini, A. Jelic, and A. Gambassi, *Proactivity and pinning in the non-reciprocal XY model with vision anisotropy*, 2026, arXiv:2606.27310 [cond-mat.stat-mech].

[3] G. Bandini, G. Biroli, P. Charbonneau, and A. Gambassi, "Overcoming critical slowing down in frustrated spin systems by learned multiscale sampling", To appear, 2026.

[4] K. Huang, *Statistical mechanics, 2nd ed* (Wiley India Pvt. Limited, 2008).

[5] N. Goldenfeld, *Lectures on phase transitions and the renormalization group* (CRC Press, 2018).

[6] J. Cardy, *Scaling and renormalization in statistical physics*, Cambridge Lecture Notes in Physics (Cambridge University Press, Cambridge, 1996).

[7] K. G. Wilson and J. Kogut, "The renormalization group and the ϵ expansion", Phys. Rep. **12**, 75 (1974).

[8] K. G. Wilson, "The renormalization group: critical phenomena and the kondo problem", Rev. Mod. Phys. **47**, 773 (1975).

[9] L. P. Kadanoff, "Scaling laws for ising models near T_c", Phys. Phys. Fiz. **2**, 263 (1966).

[10] U. Täuber, *Critical dynamics: a field theory approach to equilibrium and non-equilibrium scaling behavior* (Cambridge University Press, 2014).

[11] J. Marro and R. Dickman, *Nonequilibrium phase transitions in lattice models*, Collection Alea-Saclay: Monographs and Texts in Statistical Physics (Cambridge University Press, Cambridge, 1999).

[12] S. Ramaswamy, "The mechanics and statistics of active matter", Annu. Rev. Condens. Matter Phys. **1**, 323 (2010).

[13] M. C. Marchetti, J. F. Joanny, S. Ramaswamy, T. B. Liverpool, J. Prost, M. Rao, and R. A. Simha, "Hydrodynamics of soft active matter", Rev. Mod. Phys. **85**, 1143 (2013).

[14] C. Bechinger, R. Di Leonardo, H. Löwen, C. Reichhardt, G. Volpe, and G. Volpe, "Active particles in complex and crowded environments", Rev. Mod. Phys. **88**, 045006 (2016).

[15] T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, and O. Shochet, "Novel type of phase transition in a system of self-driven particles", Phys. Rev. Lett. **75**, 1226 (1995).

[16] J. Toner and Y. Tu, "Long-range order in a two-dimensional dynamical XY model: how birds fly together", Phys. Rev. Lett. **75**, 4326 (1995).

[17] H. Chaté, F. Ginelli, G. Grégoire, and F. Raynaud, "Collective motion of self-propelled particles interacting without cohesion", Phys. Rev. E **77**, 046113 (2008).

[18] Y. Fily and M. C. Marchetti, "Athermal phase separation of self-propelled particles with no alignment", Phys. Rev. Lett. **108**, 235702 (2012).

[19] I. Theurkauff, C. Cottin-Bizonne, J. Palacci, C. Ybert, and L. Bocquet, "Dynamic clustering in active colloidal suspensions with chemical signaling", Phys. Rev. Lett. **108**, 268303 (2012).

[20] G. S. Redner, M. F. Hagan, and A. Baskaran, "Structure and dynamics of a phase-separating active colloidal fluid", Phys. Rev. Lett. **110**, 055701 (2013).

[21] M. E. Cates and J. Tailleur, "Motility-induced phase separation", Annu. Rev. Condens. Matter Phys. **6**, 219 (2015).

[22] V. Narayan, S. Ramaswamy, and N. Menon, "Long-lived giant number fluctuations in a swarming granular nematic", Science **317**, 105 (2007).

[23] S. Ramaswamy, R. A. Simha, and J. Toner, "Active nematics on a substrate: giant number fluctuations and long-time tails", Europhys. Lett. **62**, 196 (2003).

[24] R. Aditi Simha and S. Ramaswamy, "Hydrodynamic fluctuations and instabilities in ordered suspensions of self-propelled particles", Phys. Rev. Lett. **89**, 058101 (2002).

[25] L. Giomi, M. J. Bowick, X. Ma, and M. C. Marchetti, "Defect annihilation and proliferation in active nematics", Phys. Rev. Lett. **110**, 228101 (2013).

[26] L. Giomi, M. J. Bowick, P. Mishra, R. Sknepnek, and M. C. Marchetti, "Defect dynamics in active nematics", Philos. Trans. R. Soc. A **372**, 20130365 (2014).

[27] S. Shankar, S. Ramaswamy, M. C. Marchetti, and M. J. Bowick, "Defect unbinding in active nematics", Phys. Rev. Lett. **121**, 108002 (2018).

[28] P. Digregorio, D. Levis, L. F. Cugliandolo, G. Gonnella, and I. Pagonabarraga, "Unified analysis of topological defects in 2d systems of active and passive disks", Soft Matter **18**, 566 (2022).

[29] J. Toner and Y. Tu, "Flocks, herds, and schools: a quantitative theory of flocking", Phys. Rev. E **58**, 4828 (1998).

[30] R. Garcia-Millan and G. Pruessner, "Run-and-tumble motion in a harmonic potential: field theory and entropy production", J. Stat. Mech. **2021**, 063203 (2021).

[31] L. Cocconi, G. Salbreux, and G. Pruessner, "Scaling of entropy production under coarse graining in active disordered media", Phys. Rev. E **105**, L042601 (2022).

[32] G. Pruessner and R. Garcia-Millan, "Field theories of active particle systems and their entropy production", Rep. Prog. Phys. **88**, 097601 (2025).

[33] M. Fruchart and V. Vitelli, *Nonreciprocal many-body physics*, 2026, arXiv:2602 . 11111 [cond-mat].

[34] A. Melzer, V. A. Schweigert, and A. Piel, "Transition from Attractive to Repulsive Forces between Dust Molecules in a Plasma Sheath", Phys. Rev. Lett. **83**, 3194 (1999).

[35] A. V. Ivlev, J. Bartnick, M. Heinen, C.-R. Du, V. Nosenko, and H. Löwen, "Statistical mechanics where newton's third law is broken", Phys. Rev. X **5**, 011035 (2015).

[36] R. Niu, A. Fischer, T. Palberg, and T. Speck, "Dynamics of Binary Active Clusters Driven by Ion-Exchange Particles", ACS Nano **12**, 10932 (2018).

[37] S. Saha, S. Ramaswamy, and R. Golestanian, "Pairing, waltzing and scattering of chemotactic active colloids", New J. Phys. **21**, 063006 (2019).

[38] R. Soto and R. Golestanian, "Self-assembly of catalytically active colloidal molecules: tailoring activity through surface chemistry", Phys. Rev. Lett. **112**, 068301 (2014).

[39] C. Wu, J. Dai, X. Li, L. Gao, J. Wang, J. Liu, J. Zheng, X. Zhan, J. Chen, X. Cheng, M. Yang, and J. Tang, "Ion-exchange enabled synthetic swarm", Nat. Nanotechnol. **16**, 288 (2021).

[40] A. V. Zampetaki, B. Liebchen, A. V. Ivlev, and H. Löwen, "Collective self-optimization of communicating active particles", Proc. Natl. Acad. Sci. U.S.A. **118**, e2111142118 (2021).

[41] N. J. Cira, A. Benusiglio, and M. Prakash, "Vapour-mediated sensing and motility in two-component droplets", Nature **519**, 446 (2015).

[42] C. H. Meredith, P. G. Moerman, J. Groenewold, Y.-J. Chiu, W. K. Kegel, A. van Blaaderen, and L. D. Zarzar, "Predator–prey interactions between droplets driven by non-reciprocal oil exchange", Nat. Chem. **12**, 1136 (2020).

[43] E. Theveneau, B. Steventon, E. Scarpa, S. Garcia, X. Trepat, A. Streit, and R. Mayor, "Chase-and-run between adjacent cell populations promotes directional collective migration", Nat. Cell Biol. **15**, 763 (2013).

[44] J. Agudo-Canalejo and R. Golestanian, "Active phase separation in mixtures of chemically interacting particles", Phys. Rev. Lett. **123**, 018101 (2019).

[45] S. Saha, J. Agudo-Canalejo, and R. Golestanian, "Scalar active mixtures: the nonreciprocal Cahn-Hilliard model", Phys. Rev. X **10**, 041009 (2020).

[46] Z. You, A. Baskaran, and M. C. Marchetti, "Nonreciprocity as a generic route to traveling states", Proc. Natl. Acad. Sci. U.S.A. **117**, 19767 (2020).

[47] Y. Duan, J. Agudo-Canalejo, R. Golestanian, and B. Mahault, "Dynamical pattern formation without self-attraction in quorum-sensing active matter: the interplay between nonreciprocity and motility", Phys. Rev. Lett. **131**, 148301 (2023).

[48] K. L. Kreienkamp and S. H. L. Klapp, "Nonreciprocal alignment induces asymmetric clustering in active mixtures", Phys. Rev. Lett. **133**, 258303 (2024).

[49] K. L. Kreienkamp and S. H. L. Klapp, "Clustering and flocking of repulsive chiral active particles with non-reciprocal couplings", New J. Phys. **24**, 123009 (2022).

[50] A. K. Dutta, S. Chatterjee, M. Mangeat, and R. Paul, *Stability and breakdown of chiral motion in non-reciprocal flocking*, 2026, arXiv:2604.18125 [cond-mat].

[51] S. M. Chakraborty, B. Sahoo, P. Sollich, and R. Mandal, *Effects of Non-reciprocity on Coupled Kuramoto Oscillators*, 2025, arXiv:2511.15845 [cond-mat].

[52] G. Pisegna, S. Saha, and R. Golestanian, "Emergent polar order in nonpolar mixtures with nonreciprocal interactions", Proc. Natl. Acad. Sci. U.S.A. **121**, e2407705121 (2024).

[53] S. Osat and R. Golestanian, "Non-reciprocal multifarious self-organization", Nat. Nanotechnol. **18**, 79 (2023).

[54] S. Osat, J. Metson, M. Kardar, and R. Golestanian, "Escaping kinetic traps using nonreciprocal interactions", Phys. Rev. Lett. **133**, 028301 (2024).

[55] S. A. M. Loos and S. H. L. Klapp, "Irreversibility, heat and information flows induced by non-reciprocal interactions", New J. Phys. **22**, 123051 (2020).

[56] Z. Zhang and R. Garcia-Millan, "Entropy production of nonreciprocal interactions", Phys. Rev. Res. **5**, L022033 (2023).

[57] T. Suchanek, K. Kroy, and S. A. M. Loos, "Entropy production in the nonreciprocal Cahn-Hilliard model", Phys. Rev. E **108**, 064610 (2023).

[58] A. Dinelli, J. O'Byrne, A. Curatolo, Y. Zhao, P. Sollich, and J. Tailleur, "Non-reciprocity across scales in active mixtures", Nat. Commun. **14**, 7035 (2023).

[59] G. G. Lorenzana, D. Martin, Y. Avni, D. S. Seara, M. Fruchart, G. Biroli, and V. Vitelli, *When is nonreciprocity relevant?*, 2025, arXiv:2509.17972 [cond-mat.stat-mech].

[60] A. Garcés and D. Levis, "Phase transitions in single species Ising models with non-reciprocal couplings", J. Stat. Mech. **2025**, 043205 (2025).

[61] M. Fruchart, R. Hanai, P. B. Littlewood, and V. Vitelli, "Non-reciprocal phase transitions", Nature **592**, 363 (2021).

[62] T. Suchanek, K. Kroy, and S. A. M. Loos, "Time-reversal and parity-time symmetry breaking in non-hermitian field theories", Phys. Rev. E **108**, 064123 (2023).

[63] K. L. Kreienkamp and S. H. L. Klapp, "Synchronization and exceptional points in nonreciprocal active polar mixtures", Commun. Phys. **8**, 307 (2025).

[64] C. Weis, M. Fruchart, R. Hanai, K. Kawagoe, P. B. Littlewood, and V. Vitelli, "Generalized exceptional points in nonlinear and stochastic dynamics", Phys. Rev. Res. **7**, 043157 (2025).

[65] Y. Rouzaire, D. J. G. Pearce, I. Pagonabarraga, and D. Levis, "Nonreciprocal interactions reshape topological defect annihilation", Phys. Rev. Lett. **134**, 167101 (2025).

[66] Y. Rouzaire, D. J. Pearce, I. Pagonabarraga, and D. Levis, *Dynamics of O(2) excitations in a non-reciprocal medium*, 2026, arXiv:2603.23225 [cond-mat].

[67] D. H. Kelley and N. T. Ouellette, "Emergent dynamics of laboratory insect swarms", Sci. Rep. **3**, 1073 (2013).

[68] M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, V. Lecomte, A. Orlandi, G. Parisi, A. Procaccini, M. Viale, and V. Zdravkovic, "Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study", Proc. Natl. Acad. Sci. U.S.A. **105**, 1232 (2008).

[69] M. Nagy, Z. Ákos, D. Biro, and T. Vicsek, "Hierarchical group dynamics in pigeon flocks", Nature **464**, 890 (2010).

[70] A. Cavagna, A. Cimarelli, I. Giardina, G. Parisi, R. Santagati, F. Stefanini, and M. Viale, "Scale-free correlations in starling flocks", Proc. Natl. Acad. Sci. U.S.A. **107**, 11865 (2010).

[71] A. Cavagna, L. Del Castello, S. Dey, I. Giardina, S. Melillo, L. Parisi, and M. Viale, "Short-range interactions versus long-range correlations in bird flocks", Phys. Rev. E **92**, 012705 (2015).

[72] Y. Katz, K. Tunstrøm, C. C. Ioannou, C. Huepe, and I. D. Couzin, "Inferring the structure and dynamics of interactions in schooling fish", Proc. Natl. Acad. Sci. U.S.A. **108**, 18720 (2011).

[73] J. E. Herbert-Read, A. Perna, R. P. Mann, T. M. Schaerf, D. J. T. Sumpter, and A. J. W. Ward, "Inferring the rules of interaction of shoaling fish", Proc. Natl. Acad. Sci. U.S.A. **108**, 18726 (2011).

[74] F. Ginelli, F. Peruani, M.-H. Pillot, H. Chaté, G. Theraulaz, and R. Bon, "Intermittent collective dynamics emerge from conflicting imperatives in sheep herds", Proc. Natl. Acad. Sci. U.S.A. **112**, 12729 (2015).

[75] D. Helbing, "Traffic and related self-driven many-particle systems", Rev. Mod. Phys. **73**, 1067 (2001).

[76] I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and N. R. Franks, "Collective memory and spatial sorting in animal groups", J. Theor. Biol. **218**, 1 (2002).

[77] H. Hildenbrandt, C. Carere, and C. Hemelrijk, "Self-organized aerial displays of thousands of starlings: a model", Behav. Ecol. **21**, 1349 (2010).

[78] P. T. Nguyen, S.-H. Lee, and V. T. Ngo, "Effect of vision angle on the phase transition in flocking behavior of animal groups", Phys. Rev. E **92**, 032716 (2015).

[79] M. Durve, A. Saha, and A. Sayeed, "Active particle condensation by non-reciprocal and time-delayed interactions", Eur. Phys. J. E **41**, 49 (2018).

[80] T. Bagarti and S. N. Menon, "Milling and meandering: flocking dynamics of stochastically interacting agents with a field of view", Phys. Rev. E **100**, 012609 (2019).

[81] M. Gaur, A. Saha, and S. Paul, *From Global Flocking to Local Clustering: Interplay between Velocity Alignment and Visual Perception of Active Particles*, 2026, arXiv:2602.21338 [cond-mat].

[82] W.-C. Guo, Z. Wang, P.-F. Wu, L.-J. Lang, B.-Q. Ai, and L. He, *Emergent Macroscopic Nonreciprocity from Identical Active Particles via Spontaneous Symmetry Breaking*, 2026, arXiv:2604.00552 [cond-mat].

[83] Q.-s. Chen, A. Patelli, H. Chaté, Y.-q. Ma, and X.-q. Shi, "Fore-aft asymmetric flocking", Phys. Rev. E **96**, 020601 (2017).

[84] L. P. Dadhichi, J. Kethapelli, R. Chajwa, S. Ramaswamy, and A. Maitra, "Nonmutual torques and the unimportance of motility for long-range order in two-dimensional flocks", Phys. Rev. E **101**, 052601 (2020).

[85] L. Barberis and F. Peruani, "Large-scale patterns in a minimal cognitive flocking model: incidental leaders, nematic patterns, and aggregates", Phys. Rev. Lett. **117**, 248001 (2016).

[86] R. S. Negi, R. G. Winkler, and G. Gompper, "Collective behavior of self-steering active particles with velocity alignment and visual perception", Phys. Rev. Res. **6**, 013118 (2024).

[87] Z. Liu and M. Dijkstra, "Collective dynamics of intelligent active Brownian particles with visual perception and velocity alignment in 3D: spheres, rods, and worms", Soft Matter **21**, 1529 (2025).

[88] P. Iyer, D. A. Fedosov, and G. Gompper, "Dynamics of active swarms at the edge of disorder", Commun. Phys. **9**, 37 (2026).

[89] S. A. M. Loos, S. H. L. Klapp, and T. Martynec, "Long-range order and directional defect propagation in the nonreciprocal *XY* model with vision cone interactions", Phys. Rev. Lett. **130**, 198301 (2023).

[90] Y. Avni, M. Fruchart, D. Martin, D. Seara, and V. Vitelli, "Nonreciprocal ising model", Phys. Rev. Lett. **134**, 117103 (2025).

[91] N. D. Mermin and H. Wagner, "Absence of ferromagnetism or antiferromagnetism in one- or two-dimensional isotropic Heisenberg models", Phys. Rev. Lett. **17**, 1133 (1966).

[92] D. Dopierala, H. Chaté, X.-q. Shi, and A. Solon, "Inescapable anisotropy of nonreciprocal XY models", Phys. Rev. Lett. **135**, 088302 (2025).

[93] P. Popli, A. Maitra, and S. Ramaswamy, "Ordering and defect cloaking in nonreciprocal lattice XY models", Phys. Rev. Lett. **135**, 088303 (2025).

[94] T. Welker and R. Alert, "Lattice-dependent orientational order in active crystals", Soft Matter **21**, 7228 (2025).

[95] P. C. Hohenberg and B. I. Halperin, "Theory of dynamic critical phenomena", Rev. Mod. Phys. **49**, 435 (1977).

[96] A. Sokal, "Monte Carlo methods in statistical mechanics: foundations and new algorithms", in *Functional integration: basics and applications*, edited by C. DeWitt-Morette, P. Cartier, and A. Folacci (Springer US, 1997), pp. 131–192.

[97] D. Landau and K. Binder, *A guide to Monte Carlo simulations in statistical physics*, 5th ed. (Cambridge University Press, 2021).

[98] K. Binder and D. Heermann, *Monte Carlo simulation in statistical physics: an introduction*, Graduate Texts in Physics (Springer Berlin Heidelberg, 2010).

[99] R. Agrawal, L. F. Cugliandolo, L. Faoro, L. B. Ioffe, and M. Picco, "Nonequilibrium critical dynamics of the two-dimensional $\pm J$ Ising model", Phys. Rev. E **108**, 064131 (2023).

[100] R. Agrawal, L. F. Cugliandolo, L. Faoro, L. B. Ioffe, and M. Picco, "Dynamical critical behavior on the Nishimori point of frustrated Ising models", Phys. Rev. E **110**, 034120 (2024).

[101] J.-S. Wang and R. H. Swendsen, "Cluster Monte Carlo algorithms", Physica A **167**, 565 (1990).

[102] U. Wolff, "Collective Monte Carlo updating for spin systems", Phys. Rev. Lett. **62**, 361 (1989).

[103] G. A. Miranda, M. Zheng, P. Charbonneau, A. Coniglio, L. F. Cugliandolo, and M. Tarzia, *Percolation and criticality of systems with competing interactions on bethe lattices: limitations and potential strengths of cluster schemes*, 2025, arXiv:2510.02961 [cond-mat.stat-mech].

[104] G. Alfaro Miranda, in preparation, 2026.

[105] M. Zheng, M. Tarzia, and P. Charbonneau, "Communication: weakening the critical dynamical slowing down of models with SALR interactions", J. Chem. Phys. **157**, 181103 (2022).

[106] S. Mallat, "A theory for multiresolution signal decomposition: the wavelet representation", IEEE Trans. Pattern Anal. Mach. Intell. **11**, 674 (1989).

[107] I. Daubechies, "Orthonormal bases of compactly supported wavelets", Commun. Pure Appl. Math. **41**, 909 (1988).

[108] T. Marchand, M. Ozawa, G. Biroli, and S. Mallat, "Multiscale data-driven energy estimation and generation", Phys. Rev. X **13**, 041038 (2023).

[109] M. Gabrié, G. M. Rotskoff, and E. Vanden-Eijnden, "Adaptive Monte Carlo augmented with normalizing flows", Proc. Natl. Acad. Sci. U.S.A. **119**, e2109420119 (2022).

[110] D. Wu, L. Wang, and P. Zhang, "Solving statistical mechanics using variational autoregressive networks", Phys. Rev. Lett. **122**, 080602 (2019).

[111] D. Ghio, Y. Dandi, F. Krzakala, and L. Zdeborová, "Sampling with flows, diffusion, and autoregressive neural networks from a spin-glass perspective", Proc. Natl. Acad. Sci. U.S.A. **121**, e2311810121 (2024).

[112] L. M. Del Bono, G. Biroli, P. Charbonneau, and M. Gabrié, *The critical slowing down in diffusion models*, 2026, arXiv:2605.12597 [cond-mat.stat-mech].

[113] L. M. Del Bono, F. Ricci-Tersenghi, and F. Zamponi, "Performance of machine-learning-assisted Monte Carlo in sampling from simple statistical physics models", Phys. Rev. E **112**, 045307 (2025).

[114] R. M.-P. Lallena, J. Martín-Roca, and C. Valeriani, *Partial vision leads to an unexpected emergent collective behavior in active aligning particles*, 2026, arXiv:2607.27819 [cond-mat.soft].

[115] F. Guth, E. Lempereur, J. Bruna, and S. Mallat, "Conditionally strongly log-concave generative models", in Proceedings of the 40th international conference on machine learning, Vol. 202, edited by A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, Proceedings of Machine Learning Research (2023), pp. 12224–12251.

[116] F. Koehler, A. Heckett, and A. Risteski, *Statistical efficiency of score matching: the view from isoperimetry*, 2022, arXiv:2210.00726 [cs.LG].

[117] V. L. Berezinskii, "Destruction of long-range order in one-dimensional and two-dimensional systems having a continuous symmetry group i. classical systems", Sov. Phys. JETP **32**, 493 (1971).

[118] J. M. Kosterlitz and D. J. Thouless, "Ordering, metastability and phase transitions in two-dimensional systems", J. Phys. C: Solid State Phys. **6**, 1181 (1973).

[119] J. M. Kosterlitz, "The critical properties of the two-dimensional XY model", J. Phys. C: Solid State Phys. **7**, 1046 (1974).

[120] D. S. Seara, A. Piya, and A. P. Tabatabai, "Non-reciprocal interactions spatially propagate fluctuations in a 2D Ising model", J. Stat. Mech. **2023**, 043209 (2023).

[121] J. Tobochnik and G. V. Chester, "Monte Carlo study of the planar spin model", Phys. Rev. B **20**, 3761 (1979).

[122] P. Archambault, S. T. Bramwell, and P. C. W. Holdsworth, "Magnetic fluctuations in a finite two-dimensional model", J. Phys. A: Math. Gen. **30**, 8363 (1997).

[123] F. Cooper, B. Freedman, and D. Preston, "Solving $\phi^4_{1,2}$ field theory with Monte Carlo", Nucl. Phys. B **210**, 210 (1982).

[124] D. J. Amit and V. Martin-Mayor, *Field theory, the renormalization group, and critical phenomena*, 3rd (World Scientific, 2005).

[125] T. Martynec, S. H. L. Klapp, and S. A. M. Loos, "Entropy production at criticality in a nonequilibrium Potts model", New J. Phys. **22**, 093069 (2020).

[126] C. E. F. Noa, P. E. Harunari, M. J. de Oliveira, and C. E. Fiore, "Entropy production as a tool for characterizing nonequilibrium phase transitions", Phys. Rev. E **100**, 012104 (2019).

[127] T. Tomé and M. J. de Oliveira, "Entropy production in nonequilibrium systems at stationary states", Phys. Rev. Lett. **108**, 020601 (2012).

[128] J. Schnakenberg, "Network theory of microscopic and macroscopic behavior of master equation systems", Rev. Mod. Phys. **48**, 571 (1976).

[129] Z.-Q. Li, L.-P. Yang, Z. Y. Xie, H.-H. Tu, H.-J. Liao, and T. Xiang, "Critical properties of the two-dimensional q-state clock model", Phys. Rev. E **101**, 060105 (2020).

[130] R. Baxter, *Exactly solved models in statistical mechanics*, Dover books on physics (Dover Publications, 2007).

[131] T. Surungan, S. Masuda, Y. Komura, and Y. Okabe, "Berezinskii–Kosterlitz–Thouless transition on regular and Villain types of q-state clock models", J. Phys. A: Math. Theor. **52**, 275002 (2019).

[132] L. M. Tuan, T. T. Long, D. X. Nui, P. T. Minh, N. D. Trung Kien, and D. X. Viet, "Binder ratio in the two-dimensional q-state clock model", Phys. Rev. E **106**, 034138 (2022).

[133] V. Drouin-Touchette, *The Kosterlitz-Thouless phase transition: an introduction for the intrepid student*, 2022, arXiv:2207.13748 [cond-mat.stat-mech].

[134] M. Hasenbusch, "The two-dimensional XY model at the transition temperature: a high-precision Monte Carlo study", J. Phys. A: Math. Gen. **38**, 5869 (2005).

[135] R. G. Melko, A. Paramekanti, A. A. Burkov, A. Vishwanath, D. N. Sheng, and L. Balents, "Supersolid order from disorder: hard-core bosons on the triangular lattice", Phys. Rev. Lett. **95**, 127207 (2005).

[136] A. Honecker, D. C. Cabra, H.-U. Everts, P. Pujol, and F. Stauffer, "Order by disorder and phase transitions in a highly frustrated spin model on the triangular lattice", Phys. Rev. B **84**, 224410 (2011).

[137] I. Rousochatzakis, J. Reuther, R. Thomale, S. Rachel, and N. B. Perkins, "Phase diagram and quantum order by disorder in the Kitaev $K_1 - K_2$ honeycomb magnet", Phys. Rev. X **5**, 041035 (2015).

[138] Y.-B. Shi, R. Moessner, R. Alert, and M. Bukov, "Hamiltonian description of non-reciprocal interactions", Nature Physics, 10.1038/s41567-026-03317-0 (2026).

[139] M. Besse, H. Chaté, and A. Solon, "Metastability of constant-density flocks", Phys. Rev. Lett. **129**, 268003 (2022).

[140] A. Solon, "Flocking without moving", J. Club Condens. Matter Phys., 10.36471/jccm_december_2022_01 (2022).

[141] Y. Du, Y. Cao, W. Liu, and Y. Li, "Collective behavior of the non-reciprocal vision cone XY lattice-gas model", Newton **2**, 100291 (2026).

[142] A. Cavagna, L. Del Castello, I. Giardina, T. Grigera, A. Jelic, S. Melillo, T. Mora, L. Parisi, E. Silvestri, M. Viale, and A. M. Walczak, "Flocking and Turning: a New Model for Self-organized Collective Motion", J. Stat. Phys. **158**, 601 (2015).

[143] A. Haar, "Zur theorie der orthogonalen funktionensysteme", Math. Ann. **69**, 331 (1910).

[144] L. P. Kadanoff, "Scaling laws for ising models near T_c", Physics Physique Fizika **2**, 263 (1966).

[145] T. Cover and J. Thomas, *Elements of information theory* (Wiley, 2012).

[146] B. Polyak, "Gradient methods for the minimisation of functionals", USSR Comput. Math. Math. Phys. **3**, 864 (1963).

[147] Y. Nesterov, *Introductory lectures on convex optimization: a basic course*, Vol. 87, Applied Optimization (Springer, 2004).

[148] A. Hyvärinen, "Estimation of non-normalized statistical models by score matching", J. Mach. Learn. Res. **6**, 695 (2005).

[149] J. Kaupužs, R. V. N. Melnik, and J. Rimšāns, "Corrections to finite-size scaling in the $\phi 4$ model on square lattices", Int. J. Mod. Phys. C **27**, 1650108 (2016).

[150] Y. Hu and P. Charbonneau, "Numerical transfer matrix study of frustrated next-nearest-neighbor Ising models on square lattices", Phys. Rev. B **104**, 144429 (2021).

[151] M. P. Nightingale and H. W. J. Blöte, "Dynamic exponent of the two-dimensional Ising model and Monte Carlo computation of the subdominant eigenvalue of the stochastic matrix", Phys. Rev. Lett. **76**, 4548 (1996).

[152] M. P. Nightingale and H. W. J. Blöte, "Monte Carlo computation of correlation times of independent relaxation modes at criticality", Phys. Rev. B **62**, 1089 (2000).

[153] M. Akritidis, A. Garcés, A. Vasilopoulos, M. Carosi, D. Levis, and N. G. Fytas, *Fate of the Ising universality class under nonreciprocal interactions*, 2026, arXiv:2606.06981 [cond-mat.stat-mech].

[154] J. T. Young, A. V. Gorshkov, and M. Maghrebi, "Nonequilibrium universality of the nonreciprocally coupled $O(n_1) \times O(n_2)$ model", Phys. Rev. E **113**, 034132 (2026).

[155] E. Sezik, *Non-reciprocal yet equilibrium critical dynamics*, 2026, arXiv:2607.24252 [cond-mat.stat-mech].

[156] M. K. Johnsrud, G. Pisegna, and R. Golestanian, *Scaling behavior in non-reciprocal and odd conserved dynamics near criticality*, 2026, arXiv:2608.05027 [cond-mat.soft].

[157] A. Strandburg-Peshkin, C. R. Twomey, N. W. Bode, A. B. Kao, Y. Katz, C. C. Ioannou, S. B. Rosenthal, C. J. Torney, H. S. Wu, S. A. Levin, and I. D. Couzin, "Visual sensory networks and effective information transfer in animal groups", Curr. Biol. **23**, R709 (2013).

[158] M. Sandoval, *Information transfer enhanced by non-reciprocity in a model of turning flocks*, 2026, arXiv:2604.23808 [physics.bio-ph].

[159] Z. Kadkhodaie, F. Guth, S. Mallat, and E. P. Simoncelli, *Learning multi-scale local conditional probability models of images*, 2023, arXiv:2303.02984.

[160] M. E. J. Newman and G. T. Barkema, *Monte Carlo methods in statistical physics* (Oxford University Press, 1999).

[161] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, "Equation of state calculations by fast computing machines", J. Chem. Phys. **21**, 1087 (1953).

[162] B. Berche, A. I. Fariñas-Sánchez, Y. Holovatch, and R. Paredes, "Influence of quenched dilution on the quasi-long-range ordered phase of the 2d XY model", Eur. Phys. J. B **36**, 91 (2003).

[163] Y. Okabe and T. Surungan, "Phase transition of two-dimensional diluted XY and clock models", Prog. Theor. Phys. Suppl. **157**, 132 (2005).

[164] R. Pathria and P. D. Beale, "13 - Phase transitions: exact (or almost exact) results for various models", in *Statistical mechanics*, edited by R. Pathria and P. D. Beale, 4th ed. (Academic Press, 2022), pp. 487–554.

[165] V. Privman, *Finite size scaling and numerical simulation of statistical systems* (World Scientific, 1990).

[166] M. E. Fisher and M. N. Barber, "Scaling theory for finite-size effects in the critical region", Phys. Rev. Lett. **28**, 1516 (1972).

[167] A. Brossollet, E. Lempereur, S. Mallat, and G. Biroli, "Effective energy, interactions and out of equilibrium nature of scalar active matter", Commun. Phys. **8**, 499 (2025).

[168] N. S. Pillai, A. M. Stuart, and A. H. Thiéry, "Optimal scaling and diffusion limits for the Langevin algorithm in high dimensions", Ann. Appl. Probab. **22**, 2320 (2012).

[169] G. O. Roberts and J. S. Rosenthal, "Optimal scaling of discrete approximations to Langevin diffusions", J. R. Stat. Soc. Ser. B Stat. Methodol. **60**, 255 (1998).