

Sampling algorithms in statistical physics: a guide for statistics and machine learning

Michael F. Faulkner and Samuel Livingstone

Abstract. We discuss several algorithms for sampling from unnormalized probability distributions in statistical physics, but using the language of statistics and machine learning. We provide a self-contained introduction to some key ideas and concepts of the field, before discussing three well-known problems: phase transitions in the Ising model, the melting transition on a two-dimensional plane and simulation of an all-atom model for liquid water. We review the classical Metropolis, Glauber and molecular dynamics sampling algorithms before discussing several more recent approaches, including cluster algorithms, novel variations of hybrid Monte Carlo and Langevin dynamics and piece-wise deterministic processes such as event chain Monte Carlo. We highlight cross-over with statistics and machine learning throughout and present some results on event chain Monte Carlo and sampling from the Ising model using tools from the statistics literature. We provide a simulation study on the Ising and XY models, with reproducible code freely available online, and following this we discuss several open areas for interaction between the disciplines that have not yet been explored and suggest avenues for doing so.

Key words and phrases: Statistical physics, sampling algorithms, Markov chain Monte Carlo, Ising model, Potts model, XY model, hard-disk model, molecular simulation, Metropolis, Glauber dynamics, molecular dynamics, hybrid Monte Carlo, Langevin dynamics, event chain Monte Carlo.

1. INTRODUCTION

Sampling algorithms are commonplace in statistics and machine learning – in particular, in Bayesian computation – and have been used for decades to enable inference, prediction and model comparison in many different settings. They are also widely used in statistical physics, where many popular sampling algorithms first originated ([Metropolis et al., 1953](#); [Alder and Wainwright, 1957, 1959, 1960](#)). At a high level, the goals within each discipline are the same – to sample from and approximate expectations with respect to some probability distribution – but the motivations, nomenclature and methods of explanation differ significantly.

Practitioners in Bayesian inference estimate parameter expectations based on fixed hyper-parameters and input data. To provide for this, researchers in Bayesian computation typically strive to establish general-purpose sampling algorithms (most notably Markov chain Monte Carlo) and therefore develop theory concerning how a given sampler behaves in a variety of different settings, characterised by features such as how the tails of a distribution decay (e.g. [Järner and Hansen \(2000\)](#)) or how much the sampler exploits some particular structure of the model (e.g. [Papaspiliopoulos, Roberts and Sköld \(2007\)](#)). The main concern for a given algorithm is often the extent to which it can be widely implemented with little problem-specific

HH Wills Physics Laboratory, University of Bristol, UK, (e-mail: michael.faulkner@bristol.ac.uk).
Department of Statistical Science, University College London, UK, (e-mail: samuel.livingstone@ucl.ac.uk).

tuning. Different samplers are compared by assessing how performance depends on the dimension of the parameter space (e.g. [Roberts and Rosenthal \(2001\)](#)) where ‘performance’ is typically defined as either the mixing time or the asymptotic variance of ergodic averages. Comparisons are usually based on theoretical results, which are complemented with numerical studies to corroborate the theory.

In statistical physics, expectations are studied as *functions* of the hyperparameters (e.g. the temperature) in order to predict the effect of their variation on the physical system of interest. The primary goal is to describe complex many-particle phenomena in terms of a reduced set of simplified particle–particle interactions, typically using a Boltzmann–Gibbs distribution. Unlike the Bayesian posterior, these distributions do not depend on input data, but the normalising constant is nonetheless typically intractable. Physicists compare estimated expectations with experimental data – both as functions of the relevant hyperparameters – which leads either to model modifications if there are discrepancies between the results, or to a successful description of the complex phenomena in terms of the simplified set of interactions; the latter may be followed by further predictions that experimentalists then attempt to confirm or refute. As regards sampling algorithms, there is of course concern for wide applicability, but another and perhaps stronger imperative is to assess sampler performance on a class of important benchmark models. Algorithm performance is often defined as the number of computational steps required to generate independent samples, though mixing times are also measured (e.g. [Lei and Krauth \(2018\)](#)). Comparisons are made based on the scaling of performance with the number of particles, which is proportional to the dimension of the parameter.

In addition, distributions of interest in statistical physics very often exhibit multi-modality or anisotropy at low temperature, and the ability of an algorithm to navigate this is also a key measure of sampling efficiency. By contrast, while anisotropy is a common feature of Bayesian posterior distributions – typically caused by parameter dependencies induced by the data or built into the prior – multi-modality is usually confined to specific classes of models that have known non-identifiability issues, such as mixture models ([Jasra, Holmes and Stephens, 2005](#)) or neural networks ([Neal, 2012](#)). Nevertheless, anisotropy in statistical physics informs Bayesian computation, and indeed Bayesians should also be wary of the critical slowing down that can accompany either feature near the critical temperature. Examples of multi-modality, anisotropy and critical slowing down are given in Section 6.

The objective of this work is therefore to review some model problems and sampling algorithms used in statistical physics, but from the perspective of the statistician or machine learner. The timing of our contribution is pertinent, as there have been recent parallel advances in nonreversible sampling algorithms in both Bayesian computation and statistical physics. Statisticians have established much theory assessing the merits of these algorithms (e.g. [Bierkens, Kamatani and Roberts \(2018\)](#); [Bierkens, Roberts and Zitt \(2019\)](#); [Andrieu et al. \(2021\)](#); [Andrieu and Livingstone \(2021\)](#); [Deligiannidis et al. \(2021\)](#)), while physicists have applied them to great effect in many practical scenarios of interest ([Bernard, Krauth and Wilson, 2009](#); [Bernard and Krauth, 2011](#); [Michel, Kapfer and Krauth, 2014](#); [Kapfer and Krauth, 2015](#); [Kampmann, Boltz and Kierfeld, 2015a](#); [Michel, Mayer and Krauth, 2015](#); [Faulkner et al., 2018](#); [Höllmer et al., 2020](#); [Faulkner, 2022](#)). Interdisciplinary understanding has been at times lacking, however, so that one goal of the present contribution is to support improved communication between these fields – to aid the cross-pollination of ideas and innovations.

We do not aim to provide an exhaustive review, as this would be impossible within the confines of a single article. Instead we give a brief overview of statistical physics in Section 2, before focusing attention on three well-known problems in Section 3: phase transitions in the Ising model, the melting transition on a two-dimensional plane and an all-atom model of water. In Section 4 we discuss three classical sampling algorithms used in statistical physics: the Metropolis algorithm ([Metropolis et al., 1953](#)), Glauber dynamics ([Glauber, 1963](#)) and molecular dynamics ([Alder and Wainwright, 1959, 1960](#)). In Section 5 we review some more modern alternatives, before presenting some simulation studies in Section 6 and a discussion in Section 7, in which we suggest open areas for potential collaboration between disciplines.

2. STATISTICAL PHYSICS

2.1 Microscopic statistical models

The fundamental aim of statistical physics is to predict macroscopic physical phenomena using statistical models of microscopic particle–particle interactions. Physical systems of interest tend to be viewed as collections of particles either restricted to locations on a shared d -dimensional lattice (as in Figure 1) or moving around on a shared compact d -dimensional manifold (as in Figure 3). Models of the former are used in *hard condensed matter* to study solid materials, as well as lattice-confined quantum gases (Roskilde et al., 2016) and other similar systems. Their constituent particles typically remain fixed to each lattice site and interact as a function of their positions and/or some other microscopic quantity, such as their *spin*. Models of the latter are predominantly used in *soft-matter physics* to study and compare the solid, liquid and gaseous states of a variety of materials, and their constituent particles typically interact as a function of their positions. Statistical physics is therefore the bridge between microscopic particle–particle interactions and the macroscopic world.

A microscopic statistical model consists of a collection of particles and a set of simplified rules governing their interactions, all of which is encoded in a joint probability distribution for the particle positions or spins. The *state* of an N -particle model encodes the microscopic information and is represented by the parameter $x := (x_1, \dots, x_N)^T \in \mathcal{M}^N$, where \mathcal{M}^N is the *configuration space* and \mathcal{M} is the *one-particle configuration space*. In hard condensed matter, $x_i \in \mathcal{M}$ typically describes the spin (or some other microscopic quantity) of particle i , \mathcal{M} is often a subset of \mathbb{Z} , \mathbb{R} or \mathbb{R}^2 , and $N = N_1 N_2 \dots N_d$ with N_i the number of lattice sites along the i^{th} dimension of the lattice. In soft-matter physics, the picture is somewhat simpler, with $x_i \in \mathcal{M}$ the position (and occasionally the spin) of particle i and \mathcal{M} a compact d -dimensional manifold. In both cases, the model is then defined by the Boltzmann–Gibbs probability distribution

$$(1) \quad \pi(x; \beta, \theta, N) \propto e^{-\beta U(x; \theta, N)},$$

where the *inverse temperature* $\beta > 0$ is the reciprocal of the system temperature with units such that βU is dimensionless, and the potential energy $U : \mathcal{M}^N \rightarrow \mathbb{R}$ encodes the particle–particle interactions, the number of particles N and a vector of additional hyperparameters θ . In soft matter, the particle density $\eta := \gamma N/V$ is always a component of θ , with V the volume of \mathcal{M} and $\gamma > 0$ a dimensionless constant.

In addition, physicists view some microscopic model recast in terms of different hyperparameters as the same model but in a different *statistical ensemble*. For example, the soft-matter model described at the end of the previous paragraph is in the *canonical ensemble* in which N , V and β are fixed, but it may be re-expressed in some other statistical ensemble, such as the *grand canonical ensemble* in which V and β are fixed, but the number of particles can fluctuate with some fixed potential cost. Thermodynamic theory then provides a bridge between different statistical ensembles, which can be useful when comparing numerical simulations with physical experiments. In the remainder, we assume N , V and β are fixed, unless otherwise stated.

2.2 Periodic boundary conditions and the thermodynamic limit

To remove boundary effects in hard condensed matter, researchers typically apply *periodic boundary conditions* by choosing the shared d -dimensional lattice (on which particle locations are restricted) to have d -dimensional toroidal topology: for a lattice of $N_1 \dots N_d$ particle sites, we identify lattice site (y_1, y_2, \dots, y_d) with $(y_1 + N_1, y_2, \dots, y_d) \sim \dots \sim (y_1, y_2, \dots, y_d + N_d)$. For example, the one-dimensional Ising configuration in Figure 1 is on a ring lattice. This better reflects the macroscopic systems under consideration, in which boundary effects are usually negligible compared to the large bulk of the system. Analysis is performed on the $N_1 \dots N_d$ -site lattice before the *thermodynamic limit* is taken by letting $N_1, \dots, N_d \rightarrow \infty$ with the ratio $N_1 : \dots : N_d$ fixed.

In soft-matter physics, the one-particle configuration space \mathcal{M} is typically chosen to be the d -dimensional torus \mathbb{T}^d of volume $V = L^d$, which can be defined as d -dimensional

Euclidean space under the identification $(x_1, x_2, \dots, x_d) \sim (x_1 + L, x_2, \dots, x_d) \sim \dots \sim (x_1, x_2, \dots, x_d + L)$, where $L > 0$ is the linear size of the torus (note that differing linear sizes can also be chosen along each dimension). For example, the right and left / top and bottom ‘boundaries’ in Figure 3 are identified with each other. More formally, this is the quotient space $\mathbb{R}^d / \mathbb{L}^d$, where $\mathbb{L} := LZ$. Again, this better reflects the negligible boundary effects in macroscopic systems. Analysis is typically performed on N particles on a torus of finite volume V , and the thermodynamic limit is then taken by setting $N \propto V$ and letting $N \rightarrow \infty$. In this paper, the *minimal separation vector* x_{ij} between particles i and j is defined as the shortest vector from x_j to x_i on \mathbb{T}^d , and

$$(2) \quad \mathfrak{d}(x_i, x_j) := \|x_{ij}\|_2$$

is their *minimal separation distance*. The metric \mathfrak{d} can also be induced from the quotient space representation of \mathbb{T}^d given above.

2.3 Observables, phase transitions and fundamental axiom

An *observable* of the model is any function of its state and hyperparameters, and well-chosen observables allow for the study of thermodynamic phase space and phase transitions. For an observable $\chi(x; \beta, \theta, N)$ the *expected observable*¹ will be denoted

$$(3) \quad \check{\chi}(\beta, \theta, N) := \mathbb{E}[\chi(x; \beta, \theta, N)].$$

The *thermodynamic phase space* of some observable is the space of all possible values of the *thermodynamic observable*

$$(4) \quad \lim_{N \rightarrow \infty} \check{\chi}(\beta, \theta, N).$$

A *thermodynamic phase* is any open and connected region of thermodynamic phase space in which the thermodynamic observable is analytic in both β and θ , and a *phase transition* is any boundary between two distinct thermodynamic phases. Different thermodynamic phases therefore correspond to strikingly different values of some thermodynamic observable, and distinct thermodynamic phases are separated by one or more non-analytic boundaries, each of which indicates a phase transition. For example, two thermodynamic observables associated with the two-dimensional Ising model (introduced in Section 3.1) are presented in Figure 2, both of which exhibit a phase transition.

Typically when studying phase transitions an observable is a sum of $\mathcal{O}(1)$ random variables per particle (e.g. the magnetic density of the Ising model in (11)). This leads us to a fundamental axiom of statistical physics. For any such observable, if the ratio of the standard deviation and expectation of its norm can be made arbitrarily small with increasing particle number N , then there exists some finite particle number at which the expectation is considered to have ‘reached the thermodynamic limit,’ as fluctuations from the thermodynamic value are immeasurably small. This tends to apply to macroscopic physical systems composed of large numbers of particles, though exceptions do occur near phase transitions and other regions of thermodynamic phase space that exhibit power-law correlations (e.g. Archambault, Bramwell and Holdsworth (1997); Faulkner (2022)). The consequence of this axiom is that simulations based on a large but finite number of particles can approximate behaviour in the thermodynamic limit.

2.4 Entropy, free energy and equation of state

The (dimensionless) *entropy* $S(\beta, \theta, N) := -\mathbb{E}[\log \pi(x; \beta, \theta, N)]$ is well-known to statisticians as a measure of the uncertainty associated with the probability distribution π . Boltzmann–Gibbs distributions can exhibit varying degrees of multi-modality or anisotropy depending on the values of hyperparameters such as the temperature, meaning phase transitions are often captured by changes in entropy. Indeed, if not carefully designed, a sampling

¹Rather than $\mathbb{E}[\cdot]$, physicists tend to represent expectations using the notation $\langle \cdot \rangle$.

algorithm might lose access to certain regions of non-negligible probability mass in some low-entropy thermodynamic phase (on a timescale that diverges with particle number N). This may also reflect a loss of *physical ergodicity* (Palmer, 1982) if the dynamics of the sampling algorithm are sufficiently similar to those found in nature.

The *free energy* $F(\beta, \theta, N) := \check{U}(\beta, \theta, N) - \beta^{-1}S(\beta, \theta, N)$ describes the competition between the expected potential and the entropy. It provides a toolkit for making predictions about thermodynamic phases because it can be expressed analogously to a marginal log-likelihood function: $F(\beta, \theta, N) = -\beta^{-1} \log [Z(\beta, \theta, N)]$, where the *partition function* $Z(\beta, \theta, N) := \int \exp[-\beta U(x; \theta, N)] dx$ is the normalising constant of the Boltzmann–Gibbs distribution. Free energies at different fixed values of θ can then be compared to predict the most likely state of matter at any given temperature β^{-1} . For example, simulations of a model fluid can be performed at some fixed temperature with two different values of the particle density η (with the rest of θ fixed) where one is a possible density of the gaseous phase (η_g) and the other is a possible density of the liquid phase ($\eta_l > \eta_g$). If the free energy at η_g is less than the free energy at η_l , it would follow that the fluid is more likely to be in a gaseous state than a liquid state at the chosen temperature. This is analogous to marginal log-likelihood model comparison in Bayesian inference, where the two fixed values of η are equivalent to the two models being compared, and analogies can be drawn between the single fixed temperature and the fixed data of the Bayesian statistical model.

The free energy is also a cumulant generating function, so that it can be used to derive useful expected observables such as the expected (dimensionless) *specific heat*

$$(5) \quad \check{C}(\beta, \theta, N) = -\beta^2 \partial_\beta^2 (\beta F(\beta, \theta, N)) = \beta^2 \text{Var} [U(x; \theta, N)].$$

This can be useful when classifying phase transitions. In addition, the *equation of state* in a soft-matter model is defined via the expected *pressure*

$$(6) \quad \check{p}(\beta, \theta, N) := -\partial_V F(\beta, \theta(N, V), N).$$

This may be familiar to the statistician, as the original work of Metropolis et al. (1953) entitled ‘Equation of State Calculations by Fast Computing Machines’ applied the Metropolis algorithm to the two-dimensional hard-disk model in an attempt to estimate its equation of state. Moreover, higher quality estimations of this equation of state were used to identify the fluid–hexatic phase transition described in Section 3.2.3 (Bernard and Krauth, 2011).

3. SOME EXAMPLE MODELS

In this section, we present some common models from statistical physics. In Section 3.1, we present the *Ising model* and an analysis of its thermodynamic phases. This is a hard-condensed-matter model of magnetism and provides an example of both analytically tractable free energies and *critical slowing down* at a phase transition. We then comment on the Potts and XY models before devoting the remainder of the section to soft-matter physics. In Section 3.2, we present the *two-dimensional hard-disk model*. This is possibly the simplest model of particle–particle interactions in soft matter and captures the physics of the *melting transition* in two spatial dimensions. With this basis, we then move on to the more complex physical interactions used in molecular modelling, where researchers model materials such as pure water or a collection of polymers in a liquid solvent. The full potential of an *all-atom model* of some molecular system is typically formed by combining several sub-potentials, each of which models a specific force. In Sections 3.3 and 3.4, we present (respectively) the *Lennard-Jones potential* and the *Coulomb potential*. We then introduce two common potentials used to represent molecular-bond bending and stretching in Section 3.5, before using all four sub-potentials to construct a full all-atom model of water in Section 3.6.

3.1 Ising model

The d -dimensional Ising model is possibly the most well-known model of statistical physics and describes a collection of N particles fixed at the sites y_1, \dots, y_N of a regular

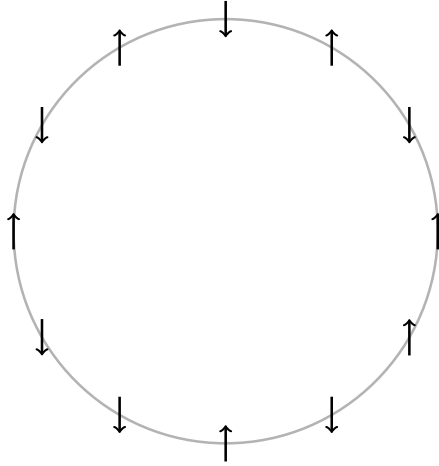


FIG 1. Example configuration of the one-dimensional Ising model on a circular ring lattice. Each up/down arrow represents the spin value $x_i = \pm 1$ of some particle i .

d -dimensional cubic lattice (with toroidal topology). It was originally constructed as a simple model of d -dimensional magnets (Ising, 1925), but has since been implemented to model many other physical and non-physical systems. Most importantly, it is commonly viewed as a paradigmatic model of phase transitions as it is possible to compute its free energy analytically in $d = 1$ (Ising, 1925) and $d = 2$ (Onsager, 1944) dimensions. Physicists refer to analytical free-energy computation as ‘exactly solving’ the model in question (Baxter, 2008).

The ferromagnetic Ising model is defined by the potential

$$(7) \quad U_{\text{Ising}}(x; J, h, N) := -\frac{J}{2} \sum_{i=1}^N \sum_{j \in S_i} x_i x_j - h \sum_{i=1}^N x_i,$$

where $J > 0$ is the *exchange constant*, $h \in \mathbb{R}$ controls the strength of an external magnetic field, S_i is the set of the $2d$ neighbours of particle i , and $x_i = \pm 1$ is the *spin* of particle i .² An example configuration in $d = 1$ dimensions is shown in Figure 1. The exchange constant J controls the level of correlation between spin values at neighbouring sites. Setting $J < 0$ defines the antiferromagnetic Ising model, in which neighbouring spin values are negatively correlated.

3.1.1 One-dimensional case. Ising showed that the one-dimensional Ising model can be solved analytically (Ising, 1925). We show in the supplement that the free energy is

$$(8) \quad F_{\text{Ising}, d=1}(\beta, J, h, N) = -\beta^{-1} \log [\lambda_+^N(\beta, J, h) + \lambda_-^N(\beta, J, h)].$$

where

$$\lambda_{\pm}(\beta, J, h) = e^{\beta J} \left[\cosh(\beta h) \pm \sqrt{\sinh^2(\beta h) + e^{-4\beta J}} \right].$$

The free energy and all of its derivatives are therefore analytic. It follows that no thermodynamic observable constructed from derivatives of the free energy exhibits a phase transition.

3.1.2 Two-dimensional case. Building on the initial work of Kramers and Wannier (1941a,b), Onsager (1944) showed that the two-dimensional Ising model can also be solved analytically. The calculations are more involved, but the thermodynamic zero-field ($h = 0$) specific heat (see (5)) per particle is

$$(9) \quad \lim_{N \rightarrow \infty} \left(\frac{1}{N} \check{C}_V(\beta, J, h = 0, N) \right) = \beta^2 \partial_{\beta}^2 \gamma(\beta J),$$

²The notation $2 \sum_{\langle i, j \rangle}$ is often used in place of $\sum_{i=1}^N \sum_{j \in S_i}$.

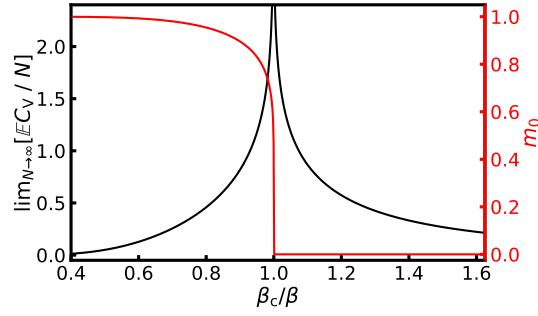


FIG 2. Thermodynamic zero-field ($h = 0$) specific heat per particle (black curve; left-hand axis; see (9)) and spontaneous magnetic density (red curve; right-hand axis; see (12)) of the two-dimensional Ising model, both as functions of β_c/β .

where

$$(10) \quad \gamma(\beta J) := \ln(2 \cosh(2\beta J)) + \frac{1}{\pi} \int_0^{\pi/2} \ln \left[\frac{1}{2} \left(1 + \sqrt{1 - \frac{4 \sinh^2(2\beta J) \sin^2 w}{\cosh^4(2\beta J)}} \right) \right] dw.$$

It then follows that the thermodynamic zero-field specific heat per particle diverges logarithmically at the *inverse critical temperature* $\beta_c := \ln(1 + \sqrt{2})/(2J)$. This predicts a phase transition at $\beta = \beta_c, h = 0$, as supported by the black curve in Figure 2. In addition, the *magnetic density*

$$(11) \quad m(x; \beta, J, h, N) := \frac{1}{N} \sum_i x_i$$

can also be used to demonstrate the phase transition, where [Onsager \(1949\)](#) and [Yang \(1952\)](#) proved a non-differentiability in the *spontaneous magnetic density*

$$(12) \quad m_0(\beta J) := \lim_{h \downarrow 0} \lim_{N \rightarrow \infty} \check{m}(\beta, J, h, N) = \begin{cases} (1 - (\sinh(2\beta J))^{-4})^{1/8} & \text{for } \beta > \beta_c, \\ 0 & \text{for } \beta < \beta_c. \end{cases}$$

This provides further evidence of the phase transition at $\beta = \beta_c, h = 0$ (as supported by the red curve in Figure 2) with the additional insight that it is one between a low-temperature ($\beta > \beta_c$) *ferromagnetic* (ordered) phase and a high-temperature ($\beta < \beta_c$) *paramagnetic* (disordered) one. We present a detailed simulation study of this model in Section 6.1.

3.1.3 Comments on Potts and XY models. The d -dimensional Potts model ([Potts, 1952](#)) is a generalisation of the d -dimensional Ising model, this time with $x_i \in \{1, 2, \dots, q\}$ ($q \geq 2$ is an integer) and potential

$$(13) \quad U_{\text{Potts}}(x; J, N) := -\frac{J}{2} \sum_{i=1}^N \sum_{j \in S_i} \mathbb{I}[x_i = x_j].$$

When $q = 2$, the Potts model is equivalent to the zero-field Ising model. As well as to phase transitions, the Potts model has been successfully applied to image processing ([Storath et al., 2015](#)).

The d -dimensional XY model can be thought of as another generalisation of the d -dimensional Ising model. Rather than on $\{-1, +1\}$, each XY spin is contained in $[0, 2\pi)$, and the XY potential is

$$(14) \quad U_{\text{XY}}(x; J, h, N) := -\frac{J}{2} \sum_{i=1}^N \sum_{j \in S_i} \cos(x_i - x_j) - h_{\text{XY}} \cdot \sum_i \begin{pmatrix} \cos x_i \\ \sin x_i \end{pmatrix},$$

where $h_{\text{XY}} \in \mathbb{R}^2$. The $d = 2$ case leads to incredibly rich physics which has been a significant focus of theoretical statistical-physics research since the 1960s ([Salzberg and Prager,](#)

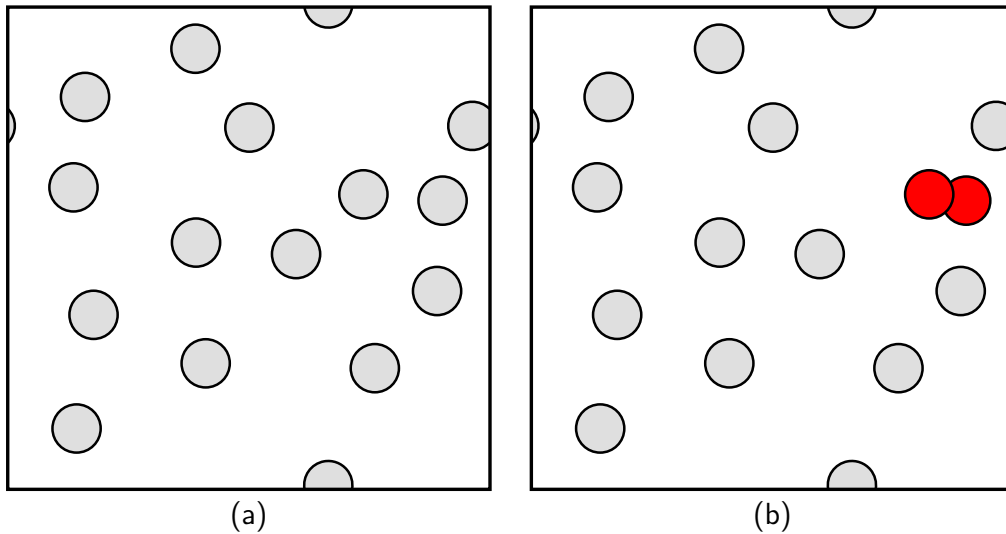


FIG 3. Examples of configurations that have non-zero (a) and zero (b) probability density in the two-dimensional hard-disk model. Two disks pass through the periodic ‘boundaries’ in each example. The red disks in (b) induce its zero-valued probability density. These configurations are accepted (a) and rejected (b) by the Metropolis algorithm described in Section 4.1.

1963; Mermin and Wagner, 1966; Hohenberg, 1967; Berezinskii, 1971; Kosterlitz and Thouless, 1973; Kosterlitz, 1974; José et al., 1977; Bramwell and Holdsworth, 1993; Vallat and Beck, 1994; Archambault, Bramwell and Holdsworth, 1997; Bramwell, Holdsworth and Pinton, 1998; José et al., 2013; Faulkner, Bramwell and Holdsworth, 2015; Faulkner, 2022). We present a simulation study of this model in Section 6.2.

3.2 Two-dimensional hard-disk model

The two-dimensional hard-disk model is perhaps the simplest approach to modelling short-range repulsive particle–particle interactions in soft-matter physics. It is defined by the probability density

$$(15) \quad \pi(x; \eta, N) \propto \prod_{1 \leq i < j \leq N} \mathbb{I}[\mathfrak{d}(x_i, x_j) > 2\sigma]$$

and describes a collection of N identical circular disks of radius $\sigma > 0$, where each exists on the compact manifold \mathbb{T}^2 and $\eta = N\pi\sigma^2/L^2$ is the disk density. All configurations in which no two disks overlap are equally likely, while all others have zero probability density (examples of two such configurations are shown in Figure 3). The model is therefore independent of the inverse temperature β and the sole hyperparameter of interest is the disk density η . It can be viewed as the $k \rightarrow \infty$ limit of the *two-dimensional soft-disk model*, which is defined on the same parameter space but with the potential $U_{\text{soft-disk}}(x; k, \eta, \varepsilon, N) := \sum_{i < j} U_{\text{sd}}(x_i, x_j; k, \eta, \varepsilon)$, where

$$(16) \quad U_{\text{sd}}(x_i, x_j; k, \eta, \varepsilon) := \varepsilon \left[\frac{2\sigma}{\mathfrak{d}(x_i, x_j)} \right]^k$$

is the two-particle soft-disk potential, with $\varepsilon > 0$ some constant with units of energy. The hard-disk model is usually studied on its own, but is also used as a sub-potential in more complex models of attractive particle–particle interactions in which the other sub-potentials contain divergences at $\mathfrak{d}(x_i, x_j) = 0$.

Despite its simplicity, the hard-disk model is able to recreate both fluid and solid structures on a two-dimensional plane, and can therefore be used to investigate the *melting transition* in two spatial dimensions. In the three-dimensional analogue of hard spheres (each) on \mathbb{T}^3 , the surface area through which each particle can move is sufficiently large that the effect of




FIG 4. A close-packed configuration of hard disks on a two-dimensional torus of volume $(16\sigma) \times (8\sqrt{3}\sigma)$. This configuration has six-point rotational symmetry about the centre of any disk in the analogous close-packed limit on \mathbb{R}^2 .

movement on other particles diminishes rapidly with $\mathfrak{d}(x_i, x_j)$, so that the three-dimensional model can easily form a solid at densities below the high-density limit. Conversely, in the one-dimensional analogue, each particle can move only along the single axis of the one-dimensional torus \mathbb{T} , so that the effect of movement is significant at all separation distances, and the system cannot form a solid at any density below the high-density limit. The two-dimensional hard-disk model is an intermediate case whose complete thermodynamic behaviour had eluded researchers for around sixty years (Metropolis et al., 1953; Alder and Wainwright, 1957; Kosterlitz and Thouless, 1973; Halperin and Nelson, 1978; Young, 1979) until event chain Monte Carlo simulations (Bernard and Krauth, 2011) led to a theory for its melting transition. The theory was then corroborated by both molecular dynamics and massively parallel Metropolis simulations (Engel et al., 2013) and subsequently confirmed in physical experiments on a collection of colloids on a two-dimensional plane (Thorneywork et al., 2017). This provides a basis for the melting transition in more complex soft-matter systems, such as films, suspensions and the crossover between two- and three-dimensional behaviour (Peng et al., 2010). Here we review the theory.

3.2.1 *Positional correlations and solid phase.* The positional correlation function is

$$(17) \quad g_p(x; r, \epsilon, \eta, N) := \sum_{i < j} \mathbb{I}[|r - \mathfrak{d}(x_i, x_j)| < \epsilon]$$

for all $r > 2\sigma$. For a fixed configuration x and a chosen r and $\epsilon > 0$, it measures the number of particle pairs whose separation distance is within ϵ of r . This provides information about the transition into the solid phase. Its expected value exhibits a drastic change in behaviour as a function of r at the particle density $\eta = \eta_s \simeq 0.720$. More precisely, Figures 3(b) and S8 of Bernard and Krauth (2011) show that for suitably small ϵ and $N = 1024^2$

$$(18) \quad \check{g}_p(r, \epsilon, \eta, N) \propto \begin{cases} \exp(-r/\xi_p(\eta)) & \text{for } \eta = 0.718, r > 10\sigma, \\ r^{-1/3} & \text{for } \eta = 0.720, r > 10\sigma, \end{cases}$$

where $\xi_p(\eta)$ is the *positional correlation length* of the non-solid phases and $x \propto y$ implies that $x = Cy$ for some $C \neq 0$. This demonstrates exponentially decaying positional correlations for all $r > 10\sigma$ at $\eta = 0.718$ and power-law decaying positional correlations for all $r > 10\sigma$ at $\eta = 0.720$. These results are consistent with a phase transition from positional disorder to a solid phase with quasi-long-range positional order as η increases through $\eta = \eta_s \simeq 0.720$, in agreement with earlier analytical thermodynamic predictions (Kosterlitz and Thouless, 1973; Halperin and Nelson, 1978; Young, 1979).

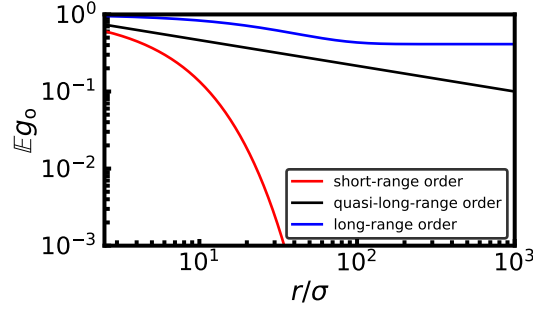


FIG 5. A schematic illustration of the three possible forms of long-distance behaviour of the expected orientational correlation function $\check{g}_o(r, \epsilon, \eta, N)$ defined in (20). Short-range / quasi-long-range / long-range order corresponds to exponentially decaying / power-law / constant correlations on distances comparable to the linear system size L ($L > 10^3\sigma$ here).

3.2.2 Orientational correlations and non-solid phases. The exponentially decaying positional correlations at $\eta = 0.718$ indicate a non-solid phase, but this is not enough to characterise the fluid phase, as the model also possesses *orientational correlations*. To quantify this, we define the *local orientation* Ψ_i of particle i via the complex number

$$(19) \quad \Psi_i(x) := \frac{1}{|\tilde{S}_i|} \sum_{j \in \tilde{S}_i} \exp(6i\phi_{ij}).$$

Here \tilde{S}_i is the set of *neighbours* of particle i , with particles i and j defined as neighbours if the centre of their minimal separation vector x_{ij} is closer to particles i and j than to any other particle. The angle $\phi_{ij} \in [0, 2\pi)$ is found by expressing x_{ij} in polar coordinates $x_{ij} := (r_{ij}, \phi_{ij})$. The factor of 6 ensures that the local orientation $\Psi_i(x)$ preserves the six-point rotational symmetry of the close-packed limit (see Figure 4). For a depiction of the local orientation $\Psi_i(x)$ within an example hard-disk configuration, see Figures 1(b-d) of [Bernard and Krauth \(2011\)](#).

The *orientational correlation function* is then defined as

$$(20) \quad g_o(x; r, \epsilon, \eta, N) := \frac{1}{\mathbb{E}|\Psi_i|^2} \sum_{i < j} \mathbb{I}[|r - \mathfrak{d}(x_i, x_j)| < \epsilon] \Psi_i^*(x) \Psi_j(x)$$

for all $r > 2\sigma$. Figure S9 of [Bernard and Krauth \(2011\)](#) and Figure 4.13 of [Bernard \(2011\)](#) show that for $N = 1024^2$ particles

$$(21) \quad \check{g}_o(r, \epsilon, \eta, N) \propto \begin{cases} \exp(-r/\xi_o(\eta)), & \text{for } \eta = 0.700, r > 200\sigma \\ r^{-\alpha_o(\eta)}, & \text{for } \eta = 0.718, r > 100\sigma \\ \tilde{C} & \text{for } \eta = 0.720, r > 100\sigma, \end{cases}$$

where $\xi_o(\eta) > 0$ is the *orientational correlation length* of the fluid phase, $\alpha_o(\eta) > 0$ is an orientational exponent and $\tilde{C} > 0$ is some constant. This indicates exponentially decaying orientational correlations for all $r > 200\sigma$ at $\eta = 0.700$, power-law decaying orientational correlations for all $r > 100\sigma$ at $\eta = 0.718$ and non-decaying orientational correlations for all $r > 100\sigma$ at $\eta = 0.720$. These results are consistent with i) an orientationally (and positionally) disordered fluid phase for all $\eta < \eta_f \simeq 0.700$, ii) quasi-long range orientational order (and positional disorder) in an *hexatic phase* at $\eta = 0.718$, and iii) long-range orientational order (and quasi-long range positional order) in the solid phase ($\eta > \eta_s$).

For clarity, Figure 5 provides a schematic illustration of the three possible forms of long-distance behaviour (short-range order, quasi-long-range order and long-range order) of the expected orientational correlation function. It is worth noting that the expected positional correlation function also exhibits long-range order (constant positive long-distance correlations) in the close-packed limit, where the particles form the precise hexagonal lattice with

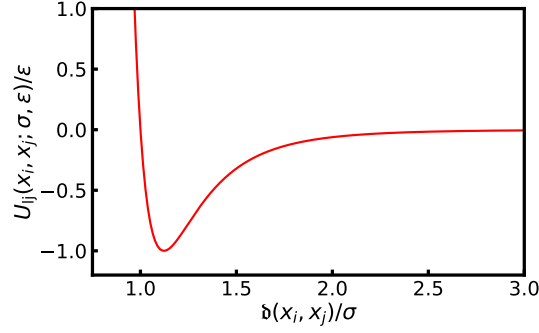


FIG 6. Two-particle Lennard–Jones potential.

six-point rotational symmetry shown in Figure 4. Moreover, in the two-dimensional zero-field Ising model presented in Section 3.1.2, an analogously defined spin–spin correlation function exhibits long/short-range order at low/high temperature, and quasi-long-range order at the phase transition.

3.2.3 Equation of state and fluid–hexatic phase transition. The above characterises the fluid, hexatic and solid phases, but does not identify a fluid–hexatic phase transition, as this requires analysis of the equation of state (6). For the two-dimensional hard-disk model, this is given by (Metropolis et al., 1953; Engel et al., 2013)

$$(22) \quad \beta \check{p}(\eta, N) = \frac{\eta}{\pi \sigma^2} \left(1 + 2\eta \lim_{r \downarrow 2\sigma} \check{g}_p(r, \eta, N) \right).$$

Equation-of-state calculations then show that, upon transforming to a model in which the pressure p is a hyperparameter, continuously increasing the pressure leads to a discontinuous jump in the expected particle density from η_f to $\eta_{\text{hex}} \simeq 0.716$ at some critical value of the pressure (Bernard and Krauth, 2011). This is the fluid–hexatic phase transition and the interval $(\eta_{\text{hex}}, \eta_s)$ is the hexatic phase.

3.3 Lennard–Jones model

The Lennard–Jones model describes soft-matter systems composed of N electrically charge-neutral atoms, each on the compact manifold \mathbb{T}^3 . It can be viewed as a more sophisticated version of the soft-disk model presented in Section 3.2 and is defined by the potential $U_{\text{LJ}}(x; \eta, \sigma, \varepsilon, N) = \sum_{i < j} U_{\text{LJ}}(x_i, x_j; \sigma, \varepsilon)$, where

$$U_{\text{LJ}}(x_i, x_j; \sigma, \varepsilon) := 4\varepsilon \left[\left(\frac{\sigma}{\mathfrak{d}(x_i, x_j)} \right)^{12} - \left(\frac{\sigma}{\mathfrak{d}(x_i, x_j)} \right)^6 \right]$$

is the two-particle Lennard–Jones potential between atoms (or particles) i and j (see Figure 6). Here, $\sigma > 0$ determines the most probable minimal separation distance between any two atoms and $\varepsilon > 0$ is the *potential-well depth*.³ The $\mathfrak{d}(x_i, x_j)^{-12}$ *Pauli-repulsion term* is the three-dimensional analogue of a two-particle soft-disk potential and represents the Pauli repulsion between the composite electrons of each atom. This quantum effect is significant for nearby particles, but diminishes rapidly at larger $\mathfrak{d}(x_i, x_j)$. The attractive $\mathfrak{d}(x_i, x_j)^{-6}$ *dispersion term* describes the electrical atom–atom attraction due to the instantaneous *electric-dipole moment* of each atom, where the electric-dipole moments are caused by electron-density fluctuations within each atom. The resultant regions of high electron density within one atom are attracted to resultant regions of low electron density in another.

³System sizes are typically chosen such that the probability density is negligible where $\mathfrak{d}(x_i, x_j) \sim L$. This ensures that the periodic boundaries do not qualitatively affect the physics. For computational efficiency, the two-particle potential is then often set to zero for all $\mathfrak{d}(x_i, x_j) > 2\sigma$.

Using techniques similar to those presented in Section 3.2, the model can be studied on its own in order to analyse the liquid–gas phase transition in simple three-dimensional fluids and other physical phenomena. The two-particle potential is also used as a sub-potential in all-atom models of more complex fluids such as water, as described in Section 3.6.

3.4 Electrostatic Coulomb potential

The toroidal Coulomb potential models electrostatic interactions between N electrically charged particles, each on the compact manifold \mathbb{T}^3 . It is derived from the Coulomb law of electrostatics, which states that each Cartesian component of the electrostatic force between particles i and j is proportional to $c_i c_j / \mathfrak{d}(x_i, x_j)^2$, where $c_i \in \mathbb{R}$ is the electric charge of particle i (Coulomb, 1785). On \mathbb{R}^3 , the two-particle Coulomb potential is given by

$$U_c(x_i, x_j; c_i, c_j) = \frac{1}{4\pi\epsilon_0} \frac{c_i c_j}{\mathfrak{d}(x_i, x_j)},$$

with $\epsilon_0 > 0$ the *vacuum permittivity*. This is the solution of the Poisson equation on \mathbb{R}^3 . In simulation, one must use the more involved solution $U_c(x_i, x_j; c_i, c_j) \propto c_i c_j G(x_i, x_j)$ to the toroidal Poisson equation $\int_{\mathbb{T}^3} \nabla_x^2 G(x, x') f(x') dx' = -f(x)$ for all test functions $f : \mathbb{R}^3 \rightarrow \mathbb{R}$, as presented in de Leeuw, Perram and Smith (1980). Charge neutrality is also required on the torus, either via $\sum_i c_i = 0$ or some charge-neutralisation technique.

If particles i and j have charge values of opposite sign, then $c_i c_j < 0$ and $U_c(x_i, x_j; c_i, c_j) \rightarrow -\infty$ as $\mathfrak{d}(x_i, x_j) \rightarrow 0$, which strongly encourages particles of opposite charge to exist arbitrarily closely together. When combined with the two-particle Lennard–Jones potential (or a suitable alternative), however, the total potential $U_c(x_i, x_j; c_i, c_j) + U_{\text{lj}}(x_i, x_j; \sigma, \epsilon) \rightarrow \infty$ as $\mathfrak{d}(x_i, x_j) \rightarrow 0$, due to the Pauli-repulsion term. This combination regularises the Coulomb potential and allows for the simulation of collections of particles with charge values of opposite sign, such as the all-atom model of water presented in Section 3.6. We add that, when applied to collections of particles with charge values of the same sign, the electrostatic Coulomb potential can also be studied on its own.

3.5 Bonded potentials

A molecule is an electrically charge-neutral group of atoms held together by chemical bonds. Molecular fluids are composed of collections of molecules that interact via intermolecular potentials, such as the Lennard–Jones and Coulomb potentials. In addition, intramolecular or bonded potentials describe the chemical-bond interactions between the composite atoms of each molecule. Two of the most common types of bonded potential are bond-stretching and bond-angle potentials, where the former dictate atom–atom minimal separation distances and the latter dictate the angle formed by the positions of three atoms (see Figure 7). For three bonded atoms i, j and k on \mathbb{T}^3 , the *harmonic bond-stretching potential* is

$$U_s(x_i, x_j; r_0, k_b) := \frac{1}{2} k_b (\mathfrak{d}(x_i, x_j) - r_0)^2,$$

and the *harmonic bond-angle potential* is

$$U_a(x_i, x_j, x_k; \phi_0, k_a) := \frac{1}{2} k_a (\phi(x_i, x_j, x_k) - \phi_0)^2.$$

Here, $r_0 > 0$, $k_b > 0$, $\phi_0 > 0$ and $k_a > 0$ are constants that depend on the molecular fluid, and

$$\phi(x_i, x_j, x_k) := \arccos \left(\frac{x_{ij}^T x_{jk}}{\mathfrak{d}(x_i, x_j) \mathfrak{d}(x_j, x_k)} \right)$$

is the *bond angle* between three bonded atoms i, j and k .

The quadratic bond-stretching potential is derived from *Hooke's law*, but non-quadratic bond-stretching potentials are also used. For example, graphene typically uses quartic bond-stretching potentials to reflect the enhanced strength of its atomic bonds (Wei, Song and Wang, 2011). Similarly, non-quadratic bond-angle potentials are also used.

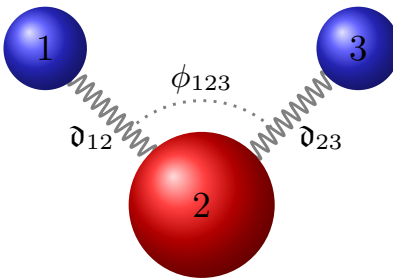


FIG 7. Water molecule with indexing as described in Section 3.6. Oxygen/hydrogen atoms are red/blue. The bond angle $\phi_{123} := \phi(x_1, x_2, x_3)$ and the minimal separation distance $d_{ij} := d(x_i, x_j)$.

3.6 An all-atom model of water

The electrostatic Coulomb and Lennard–Jones potentials can be combined with the bonded potentials described above to produce an all-atom model of water. In molecular modelling, an all-atom model is any microscopic model that accounts for the interactions between the individual atoms that form the molecules. This projects the fundamental quantum-mechanical many-body system onto a simplified classical model of the atomic positions. Given this significant simplification, many different all-atom models of water exist, where the most suitable model is situation-dependent.

The simple point-charge water model with flexible molecules (Wu, Tepper and Voth, 2006) is composed of sub-potentials that describe i) two-body oxygen–hydrogen bond stretching, ii) three-body hydrogen–oxygen–hydrogen angle bending, iii) oxygen–oxygen Lennard–Jones interactions, and iv) electrostatic Coulomb interactions between all (intermolecular) atoms, so that its two-molecule potential is given by

$$U_{\text{mol-mol}}(x_1, \dots, x_6) = \sum_{j \in \{1,3\}} U_s(x_2, x_j) + \sum_{j \in \{4,6\}} U_s(x_5, x_j) + U_a(x_1, x_2, x_3) \\ + U_a(x_4, x_5, x_6) + U_{\text{lj}}(x_2, x_5) + \sum_{i=1}^3 \sum_{j=4}^6 U_c(x_i, x_j; c_i, c_j),$$

where particles 2 and 5 are oxygen atoms, particles 1, 3, 4 and 6 are hydrogen atoms, and this two-molecule potential generalises to an arbitrary number of molecules. Similar techniques to those presented above are used to analyse the various thermodynamic phases of the model.

4. CLASSICAL SAMPLING ALGORITHMS

4.1 Metropolis

In the first instance of Markov chain Monte Carlo, Metropolis et al. (1953) developed an algorithm to sample from the Boltzmann–Gibbs distribution and applied it to the two-dimensional hard-disk model. The algorithm is typically referred to as either ‘the Metropolis algorithm’ or ‘Monte Carlo’ within the statistical physics community. Modern-day statisticians, however, may feel more comfortable with the term Metropolis–within–Gibbs, as only a subset of the state is updated at each iteration. Evolving a single particle at each iteration of the algorithm is generally preferred in high-density particle systems, as evolving all particles at once very often leads to slow mixing.

When applied to some d -dimensional soft-matter model, each iteration of the Metropolis algorithm consists of proposing the movement of some particle i to a new candidate position

$x'_i = x_i \oplus u$, where $x_i \in \mathbb{T}^d$ is the original position of particle i , each $u_j \sim \mathcal{U}[-\epsilon, \epsilon]$ for $j \in \{1, \dots, d\}$ for some appropriately chosen $\epsilon > 0$, and \oplus is addition on \mathbb{T}^d . This candidate move is then accepted with probability

$$\alpha(x_i, x'_i) := \min \left(1, e^{-\beta \Delta U(x_i, x'_i)} \right),$$

where $\Delta U(x_i, x'_i)$ denotes the change in the potential when replacing x_i with x'_i . This is simply a ratio of Boltzmann–Gibbs distributions. The original algorithm is a systematic scan sampler, meaning the particles are cycled through in a deterministic fashion, rather than randomly selected at each iteration. When applied to the hard-disk model this probability becomes

$$\alpha(x_i, x'_i) = \min \left(1, \frac{\prod_{j \neq i} \mathbb{I}[\mathfrak{d}(x'_i, x_j) > 2\sigma]}{\prod_{j \neq i} \mathbb{I}[\mathfrak{d}(x_i, x_j) > 2\sigma]} \right).$$

Since this target distribution is uniform, $\alpha(x_i, x'_i)$ simplifies to being one / zero if particles do not / do overlap in the proposed configuration (see Figure 3 for examples of accepted and rejected configurations).

Metropolis *et al.* investigated the melting transition in two spatial dimensions by simulating the hard-disk model at various choices of disk density. In the simulations the linear torus size $L = 1$ and the number of particles $N = 224$, and the particle density η was varied by adjusting the disk diameter σ . At each chosen disk density the initial state was set to be a 14×16 hexagonal grid, and the simulations consisted of 16 burn-in sweeps followed by another 48 – 64 sampling sweeps, where a single sweep is N iterations of the algorithm. Each sweep took around 3 minutes, meaning a total running time of 4-5 hours using the MANIAC computer at Los Alamos National Laboratory. In fact no evidence of a thermodynamic phase transition was found in the simulations. This is due to the Metropolis algorithm exhibiting extremely slow mixing in the vicinity of both the liquid-hexatic and hexatic-solid phase transitions, because particle moves will very often result in disk overlaps at high particle density, leading to rejections. To alleviate this, one must choose a very small step size ϵ , which typically leads to very high auto-correlation within the chain and slow convergence to equilibrium. Similar results were found when applying the Metropolis algorithm to the two-dimensional Lennard–Jones potential and the three-dimensional hard-spheres model in Rosenbluth and Rosenbluth (1954). Wood and Parker (1957) found, however, some evidence of a phase transition when applying the Metropolis algorithm to the three-dimensional Lennard–Jones potential.

4.2 Glauber dynamics

The Metropolis algorithm can also be applied to the Ising model. At each iteration a candidate move is generated by randomly selecting a particle (meaning a site on the lattice) and flipping the sign of the spin of that particle. This proposal distribution is symmetric and hence the Metropolis rule can be used to accept or reject the move. Another very similar algorithm introduced in Glauber (1963) and now known as *Glauber dynamics* is also commonly used for this application.

Glauber dynamics is most easily understood by the statistician as a random scan Gibbs sampler for the Ising model. At each iteration of the algorithm a particle is selected uniformly and a new value for the spin at that site is drawn from the conditional distribution given the spin values of neighbouring particles.

It is natural to compare the two approaches, and in fact this can be done straightforwardly using some well-known tools of the statistician. To do this consider the Glauber dynamics transition as proceeding in three stages. In the first a particle i is randomly selected. In the second a candidate move is considered in which the spin of that particle is changed. In the third the candidate move is accepted with probability

$$\alpha_{GD}(x_i, x'_i) = \frac{e^{-\beta \Delta U(x_i, x'_i)}}{1 + e^{-\beta \Delta U(x_i, x'_i)}}.$$

From this representation, it can be seen that Glauber dynamics can also be viewed as a version of Metropolis–Hastings, whose acceptance rate has been replaced with that advocated by

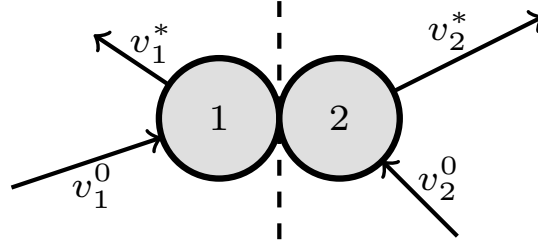


FIG 8. A collision of two equal-mass hard disks in a molecular-dynamics simulation. The dashed line is the perpendicular bisector of the minimal separation vector $x_{21}(t^*) = (2\sigma, 0)$ at the collision time t^* ($x_{21}(t^*)$ is not shown). The black arrows represent the particle velocities before and after the collision, with arrow length proportional to the norm of the vector. All changes in velocity are proportional to the component of $v_{21}(t^*) = v_2(t^*) - v_1(t^*)$ perpendicular to the dashed line: $v_1^0 := v_1(t_0) = (1, 1/3)\iota$, $v_1^* := v_1(t^*) = (-1/2, 1/3)\iota$, $v_2^0 := v_2(t_0) = (-1/2, 1/2)\iota$ and $v_2^* := v_2(t^*) = (1, 1/2)\iota$, where $\iota > 0$ has units of velocity.

[Barker \(1965\)](#). Using this observation the superiority of Metropolis in terms of asymptotic variance can be established. The proposition below is an immediate consequence of Theorem 4 in [Latuszyński and Roberts \(2013\)](#).

PROPOSITION 1. *Let P_M denote the Markov chain produced by the Metropolis algorithm and P_G denote that produced by Glauber dynamics for the Ising model. For any f such that $\sum_{x \in \mathcal{M}^N} f(x)^2 e^{-\beta U_{\text{ising}}(x)} < \infty$ it holds that*

$$\nu(P_M, f) \leq \nu(P_G, f) \leq 2\nu(P_M, f) + \text{Var}_\pi(f)$$

where $\nu(P, f) := \lim_{n \rightarrow \infty} n \text{Var}(\hat{f}_n)$ is the asymptotic variance of the ergodic average $\hat{f}_n := n^{-1} \sum_{i=1}^n f(X_i)$, with $X_i | X_{i-1} \sim P(X_{i-1}, \cdot)$ and X_1 a sample from the stationary distribution of P .

4.3 Molecular dynamics

In a molecular dynamics simulation Newton's equations of motion are (approximately) solved to directly compute all particle trajectories, after setting random initial velocities. This approach differs in many ways from the Metropolis algorithm, in which only a single particle is perturbed at each iteration and only the equilibrium behaviour of the system is modelled. The approach constitutes a direct numerical solution to the N -body problem, in order to understand dynamical properties of the system. The samples produced from certain molecular dynamics simulations can still be used to estimate expected observables at equilibrium, provided the simulation is run for a sufficiently long time.

The molecular dynamics algorithm (abbreviated MD) was first applied to the two-dimensional hard-disk model ([Alder and Wainwright, 1957](#)) before a general method was developed by [Alder and Wainwright \(1959, 1960\)](#). Unlike the Metropolis algorithm, MD did eventually find convincing evidence of a phase transition in the two-dimensional hard-disk model ([Alder and Wainwright, 1962](#)). This empirical finding motivated [Kosterlitz and Thouless \(1973\)](#); [Halperin and Nelson \(1978\)](#); [Young \(1979\)](#) to develop a two-step theory for the melting transition in two spatial dimensions, which predicted two phase transitions through an intermediate hexatic phase. Strong particle–particle positional correlations in the vicinity of the transition, however, meant that contemporary molecular dynamics simulations could neither disprove nor corroborate this more nuanced theory.

4.3.1 Molecular dynamics for hard disks. The hard-disk model is composed of N particles of equal mass, which we set to unity for simplicity. Newtonian dynamics take a simple form: particles move at constant velocity until two collide, at which point the velocities are updated based on the speed and angle of the collision. In a typical molecular dynamics simulation the particles are initialised at some chosen positions (e.g. a hexagonal lattice) and given random initial velocities $v_i \in \mathbb{R}^2$ for $i \in \{1, \dots, N\}$ at time t_0 (often the magnitudes are set

to be equal and the directions are sampled uniformly). After this point the system evolves deterministically, with the position of particle i at time $t \geq t_0$ calculated as

$$(23) \quad x_i(t) = x_i(t_0) + (t - t_0)v_i(t), \quad v_i(t) = v_i(t_0).$$

The above equation is correct assuming that no boundary has been crossed. To incorporate periodic boundary conditions the positions must be adjusted modulo L upon hitting a boundary. Each particle evolves simultaneously in this manner until two collide. The time of the first pair-wise collision between any two particles is again completely pre-determined, and is the first time t^* at which

$$(24) \quad \mathfrak{d}(x_i(t^*), x_j(t^*)) = 2\sigma$$

for some (i, j) pair. This can be calculated exactly. Assuming no boundary has been crossed and no other collisions have occurred, the minimal separation vector between particles i and j (i.e. the shortest vector from x_j to x_i) is $x_{ij}(t) = x_{ij}(t_0) + (t - t_0)(v_i(t_0) - v_j(t_0))$ for any $t \leq t^*$. In this case (24) becomes a simple quadratic in t^* , and denoting $v_{ij}(t) := v_i(t) - v_j(t)$ it has minimum (positive) solution

$$t^* = t_0 - \frac{x_{ij}(t_0)^T v_{ij}(t_0) + \sqrt{(x_{ij}(t_0)^T v_{ij}(t_0))^2 - v_{ij}(t_0)^2 (x_{ij}(t_0) - 4\sigma^2)}}{\|v_{ij}(t_0)\|_2^2},$$

provided that the expression inside the square root is positive and $x_{ij}(0)^T v_{ij}(0) < 0$. To account for periodic boundary conditions, the above procedure can be straightforwardly modified by re-evaluating t^* each time a particle passes through a boundary. When a collision occurs the velocities of the two involved particles are updated such that the total energy (or the quadratic kinetic energy in this case) is conserved, using the formulae

$$(25) \quad v_i(t^*) = v_i(t_0) - \left(\frac{v_{ij}(t_0)^T x_{ij}(t^*)}{4\sigma^2} \right) x_{ij}(t^*),$$

$$(26) \quad v_j(t^*) = v_j(t_0) + \left(\frac{v_{ij}(t_0)^T x_{ij}(t^*)}{4\sigma^2} \right) x_{ij}(t^*).$$

Note that $\|x_{ij}(t^*)\| = 2\sigma$ at all collision times t^* . The second term in each equation therefore contains two $x_{ij}(t^*)/\|x_{ij}(t^*)\|$ terms and all changes in velocity are proportional to the component of v_{ij} parallel to $x_{ij}(t^*)$, or perpendicular to the dashed line in the example collision depicted in Figure 8. After any such collision, t_0 is updated to t^* and the process is repeated. An example implementation of molecular dynamics for the hard-disk model is given in Algorithm 1 below.

Algorithm 1: Molecular simulation for the hard-disk model

```

1 Require  $\{(x_i(0), v_i(0)) : 1 \leq i \leq N\}$ , desired collisions  $C^* < \infty$ ;
2 Set  $t_0 \leftarrow 0$ ,  $C \leftarrow 0$ ;
3 for  $(i, j) \in \{1, \dots, N\}^2$  do
4   | Compute next collision time  $t_{ij}^*$ ;
5 end
6 Set  $t^* \leftarrow \min_{(i,j)} t_{ij}^*$ ;
7 if  $t^* = t_{ij}^*$  then
8   | Update velocities for particles  $i$  and  $j$  using (25)-(26);
9 end
10 Set  $t_0 \leftarrow t_0 + t^*$ ,  $C \leftarrow C + 1$ ;
11 for  $i \in \{1, \dots, N\}$  do
12   | Compute  $x_i(t^*)$  using (23) with boundary corrections;
13 end
14 if  $C < C^*$  then
15   | Return to line 3;
16 end

```

Algorithm 1 is an exact description of the dynamics of the hard-disk system. The only numerical errors introduced into molecular dynamics simulation of the model are from floating point arithmetic calculations at collision times. It may be surprising, therefore, to learn that such errors can sometimes accumulate rapidly. The reason, put simply, is due to the nature of pair-wise collisions between particles, in which small differences in calculation of the angle of refraction can result in amplified differences in the positions of particles at the next collision. Further discussion of this phenomenon is provided in Section 2.1.2 of Krauth (2006). The hard-disk dynamics are often called *event driven*, as ballistic movement of particles is interrupted by collision events. The idea of using event driven dynamics has more recently been applied in statistics to construct sampling algorithms based on Hamiltonian dynamics in the presence of general discontinuous distributions in Nishimura, Dunson and Lu (2020).

It is natural to consider the ergodic properties of this approach, which in its simplest form is completely deterministic apart from the random choice of initial particle velocities. Ergodic properties of various forms of the hard-disk model have now been established under mild conditions by Simányi (2003) after pioneering earlier work by Sinai (1970) establishing ergodicity for the case of two particles. The result is significant in forging a concrete connection between Newtonian dynamics (on which the equations of motion are based) and the Boltzmann distribution (from which the equilibrium properties of the system are deduced). Among mathematicians, models such as the hard-disk system are often referred to as the study of dynamical billiards (e.g. Tabachnikov (2005)).

4.3.2 Smooth potentials: the microcanonical ensemble. Newton’s equations of motion for smooth potentials are given by the dynamical system

$$(27) \quad m_i \frac{d^2 x_i}{dt^2} = F_i(x),$$

where $m_i \in [0, \infty)$ denotes the mass of particle i and $F_i(x) := -\nabla_i U(x)$ is the total force acting on particle i . Here ∇_i is the gradient operator of particle i . It is common to introduce an auxiliary velocity variable $v_i := \dot{x}_i$, reducing the second order system (27) into a first order system

$$(28) \quad \begin{aligned} \dot{x}_i &= v_i, \\ \dot{v}_i &= F_i(x)/m_i. \end{aligned}$$

The above dynamics can also be described in terms of position and momenta $p_i := m_i v_i$. This *Hamiltonian* formulation of classical mechanics gives rise to a dynamical system with several appealing features, such as volume preservation and invariance of the Hamiltonian function $H(x, p) = U(x) + \sum_i p_i^T p_i / (2m_i)$, which describes the total energy of the system by combining the potential $U(x)$ with a quadratic kinetic energy term.

The system (28) cannot usually be solved analytically, but in many cases numerical integrators that preserve many geometrical features of the original system exist. A general survey is beyond the scope of this article, but see Hairer, Lubich and Wanner (2003); Bou-Rabee and Sanz-Serna (2018) for comprehensive reviews or Leimkuhler and Reich (2004); Hairer, Lubich and Wanner (2006) for book-length treatments. The most popular algorithm in use today is the *velocity Verlet* algorithm (Verlet, 1967), in which the dynamics are approximated by first taking a half-step in the momentum component $p_i(\varepsilon/2) = p_i(0) + (\varepsilon/2)F_i(x(0))$ for each particle (where $\varepsilon > 0$ is the step-size) and then iterating the *leapfrog* dynamics for each $n \in \{1, \dots, \ell\}$

$$(29) \quad \begin{aligned} x_i(n\varepsilon) &= x_i((n-1)\varepsilon) + \varepsilon p_i((n-1/2)\varepsilon)/m_i \\ p_i((n+1/2)\varepsilon) &= p_i((n-1/2)\varepsilon) + \varepsilon F_i(x(n\varepsilon)), \end{aligned}$$

before a half-step in the momentum component is taken at the final iteration, in order to generate a skeleton trajectory up to time $\ell\varepsilon$ (the alternative *position Verlet* algorithm begins and ends with a half-step in the position). The momentum update is referred to as the ‘kick’ and the position update the ‘drift’. The algorithm is also known as the *leapfrog* scheme, owing

to the intermediate leapfrogging action of the position and momentum coordinates, and is popular because only one force evaluation $F(x)$ is needed per time step (ignoring the initial and final kicks) while achieving $\mathcal{O}(\varepsilon^2)$ global error over fixed time scales (Bou-Rabee and Sanz-Serna, 2018). The algorithm was originally used within the contemporary molecular dynamics literature by Loup Verlet to simulate a system of 864 particles interacting under a Lennard–Jones potential, but had been used much earlier than this by Störmer (1907), and is sometimes called Störmer–Verlet integration for this reason. The positions of particles are then updated again after each leapfrog step to incorporate periodic boundary conditions by applying the modular transformation described in Section 2.2.

The molecular dynamics algorithm described above is restricted to exploring the *micro-canonical ensemble*, meaning the space of possible states

$$\{(x, p) \in \mathcal{M}^N \times \mathbb{R}^{Nd} : H(x, p) = H_0\},$$

where $H_0 := H(x(0), p(0))$, combined with the uniform probability measure over these states. The total energy of the system hence remains constant. On configuration space \mathcal{M}^N , this usually will not correspond to the support of the Boltzmann–Gibbs distribution $e^{-\beta U(x)}$, as the potential is restricted to the set $\{U(x) \leq H_0\}$ owing to the non-negativity of the kinetic energy. This was of no concern for the hard-disk model as the potential is almost everywhere constant, meaning that each level set of H allows exploration of the entire space. For general potentials, this requires the ability to move between contours of the Hamiltonian. The collection of states

$$\{(x, p) \in \mathcal{M}^N \times \mathbb{R}^{Nd} : e^{-\beta H(x, p)} > 0\}$$

combined with the probability measure $\propto e^{-\beta H(x, p)} dx dp$ over these states, is known as the *canonical ensemble*. In some physical settings the microcanonical ensemble is of direct interest, but if the canonical ensemble is desired then the above approach to molecular simulation is no longer sufficient.

4.3.3 Smooth potentials: the canonical ensemble. There are various approaches to simulating the canonical ensemble, and hence exploring the entirety of the Boltzmann–Gibbs distribution $\pi(x) \propto e^{-\beta U(x)}$. From the physical perspective sampling from the canonical ensemble can be understood as allowing the system of particles to exchange energy with the outside world. The system is often assumed to be contained within a *heat bath* or *thermal reservoir* meaning that the system temperature β^{-1} can be controlled whilst still allowing for heat exchange with the external environment. Here we will primarily discuss *Langevin dynamics* to sample from the canonical ensemble, although there are many other approaches for this task (Leimkuhler and Matthews, 2016, Chapter 6).

Langevin dynamics consist of adding some stochasticity to Newton’s equations of motion (27), which allows the total system energy to fluctuate over time. The deterministic system (28) is combined with an Ornstein–Uhlenbeck process on the momentum coordinate, resulting in the system of stochastic differential equations

$$(30) \quad \begin{aligned} dx_i(t) &= m_i^{-1} p_i(t) dt \\ dp_i(t) &= F_i(x(t)) dt - \gamma m_i^{-1} p_i(t) dt + \sqrt{2\gamma\beta^{-1}} dW_i(t), \end{aligned}$$

where each $(W_i(t))_{t \geq 0}$ is a standard Wiener process on \mathbb{R}^d , and $\gamma > 0$ controls the strength of frictional forces. The linear drift term $-\gamma m_i^{-1} p_i(t) dt$ taken in isolation results in an exponential decay in the momentum, with γ dictating the rate at which energy dissipates from the system due to friction. The final term $dW_i(t)$ represents an injection of stochastic force, and its coefficient can be determined using the *fluctuation-dissipation theorem* (Pavliotis, 2014, Chapter 9). We direct the interested reader to Section 6.3.2 of Leimkuhler and Matthews (2016) for more precise physical intuition.

The system (30) is commonly known as *underdamped* Langevin dynamics. It can be solved numerically in various ways. One popular approach is to split the dynamics into two separate

systems, the first being simply the Hamiltonian system (28), and the second the Ornstein–Uhlenbeck process $dp_i(t) = -\gamma m_i^{-1} p_i(t) dt + \sqrt{2\gamma\beta^{-1}} dW_i(t)$, which has explicit solution

$$(31) \quad p_i(t) | p_i(0) \sim N \left(p_i(0) e^{-\gamma m_i^{-1} t}, m_i \beta^{-1} (1 - e^{-2\gamma m_i^{-1} t}) \right).$$

The full system (30) can then be approximated by iterating between numerically solving (29) for a short time increment and exactly solving the Ornstein–Uhlenbeck dynamics for the same length of time. Justification for solving in this manner is given by the Baker–Campbell–Hausdorff formula (e.g. Leimkuhler and Reich (2004)). Ergodic properties and mixing times of underdamped Langevin dynamics are an area of constant study in the applied mathematics literature (e.g. Wu (2001); Talay (2002); Mattingly, Stuart and Higham (2002)) as well as the behaviour of discretisation schemes (e.g. Mattingly, Stuart and Higham (2002); Durmus et al. (2021)). Interest in the approach has also grown recently within the machine learning community in the context of establishing non-asymptotic mixing-time bounds for approximate sampling algorithms (e.g. Dalalyan and Riou-Durand (2020)).

A simplified form of the above dynamics that will be very familiar to statisticians and machine learners can be found by considering a particular limiting regime of (30). Taking $m_i = 1$ for simplicity, in the *overdamped limit* $\gamma \rightarrow \infty$ the right-hand side of (31) becomes $N(0, \beta^{-1})$, thereby flushing all memory of previous momenta from the system. Alternating (31) and (29) in this limit results in the transition $x_i(\varepsilon) = x_i(0) + \varepsilon^2 F_i(x(0))/2 + \varepsilon \zeta_i$, where $\zeta_i \sim N(0, \beta^{-1})$, which is the Euler–Maruyama numerical scheme applied to the stochastic differential equation

$$(32) \quad dy_i(t) = F_i(y(t)) dt + \sqrt{2\beta^{-1}} dB_i(t),$$

where $y(t) := x(\gamma t)$. A rigorous derivation of the above is provided in Section 6.5 of Pavliotis (2014), see also Section 1.2.2 of Stoltz (2021). Among the statistical physics community (32) is often called *Brownian dynamics*, as introduced by Rosicky, Doll and Friedman (1978). Among statisticians (32) is the starting point of the Metropolis-adjusted Langevin algorithm popularised by Roberts and Tweedie (1996) and used extensively since this point.

4.4 Hybrid algorithms

We end this section by briefly mentioning hybrid Monte Carlo, which is now more commonly known as Hamiltonian Monte Carlo (Neal, 2011). Following the above discussion the original name should seem natural, as among physicists the algorithm is easily understood as a hybrid of the molecular dynamics and Metropolis approaches. The algorithm was introduced in lattice field theory by Duane et al. (1987), and was popularised in the statistics literature by Neal (1993). Today it is considered to be among the state-of-the-art approaches to sampling in many statistical applications. Two notable examples where the method has found success are posterior distributions for Bayesian neural networks (Neal, 2012) and hierarchical models (Betancourt and Girolami, 2015).

Today it is widely used in both disciplines and has been extensively studied. Indeed, many recent algorithmic innovations have been made by statisticians and machine learners, such as the automated tuning of the time for which Hamilton’s equations should be simulated before a Metropolis step is applied and the momentum is re-sampled (Hoffman and Gelman, 2014; Sherlock, Urbas and Ludkin, 2021), and the incorporation of geometric ideas to allow for position-dependent masses that use local information about π for sampling highly anisotropic distributions (Girolami and Calderhead, 2011; Betancourt et al., 2017).

5. ADVANCED ALGORITHMS

5.1 Cluster algorithms for lattice models

The potential associated with the Ising model induces strong correlations between spin values at neighbouring lattice sites when the inverse temperature β is large. This can make the sampling task very challenging using a site-by-site updating strategy as employed by the

Metropolis algorithm and Glauber dynamics, leading to poorly mixing Markov chains. An alternative approach is to consider changing the spins of several particles in a single step of the algorithm. A popular strategy for doing this was proposed by [Swendsen and Wang \(1987\)](#), and later modified by [Wolff \(1989a\)](#). In the following we set $h = 0$ (see (8)) for brevity.

The Swendsen–Wang algorithm is based on the idea of changing the spin values of entire clusters of particles together. The key ingredient is to introduce an auxiliary variable for each edge joining adjacent lattice sites (excluding edges passing through the periodic boundaries). Consider a lattice of N sites (corresponding to N particles) and label each with an index $1 \leq i \leq N$. Each edge in the lattice can then be assigned an auxiliary *bond* variable indexed by the two particles that it connects. In the one-dimensional Ising model this equates to introducing $N - 1$ auxiliary variables, and for a d -dimensional lattice $dN^{(d-1)/d}(N^{1/d} - 1)$ such variables. For two neighbouring particles i and j we denote the associated bond variable $b_{ij} \in \{0, 1\}$. If $b_{ij} = 1$ then particles i and j are grouped in the same cluster, and if not they belong to different clusters. The edge variables therefore partition the set of particles. If $x_i \neq x_j$ then $b_{ij} = 0$, meaning only particles with the same spin can be in the same cluster. If $x_i = x_j$ then $b_{ij} = 0$ with probability $e^{-2\beta J}$, meaning

$$\mathbb{P}[b_{ij} = 1 | x_i, x_j] = q_{ij}(x) := 1 - \exp\{-2\beta J \mathbb{I}(x_i = x_j)\}.$$

Once all edge variables have been sampled, all spins within each cluster are flipped with probability $1/2$, and for each cluster the decision of whether or not to flip the spins is taken independently of all other clusters. In this way, large numbers of particles can be flipped simultaneously. A proof of the following well-known result is provided in the supplement.

PROPOSITION 2. *The Markov chain induced by the Swendsen–Wang algorithm targeting the Boltzmann–Gibbs distribution $\pi(x) \propto e^{-\beta U_{\text{ising}}(x; J, 0, N)}$ is ergodic for any choice $\beta \in (0, \infty)$, $J > 0$ and $N \in \mathbb{N}$.*

Conditions for rapid mixing of the algorithm are discussed in [Gore and Jerrum \(1999\)](#), and convergence has also been considered by [Huber \(2003\)](#). Generalizations and further discussion are provided in [Edwards and Sokal \(1988\)](#) and elsewhere.

The Wolff algorithm ([Wolff, 1989a](#)) differs from the approach of Swendsen & Wang in that only a single cluster is flipped at each iteration. The approach can be uncovered by sampling each bond variable as in the Swendsen–Wang algorithm, but then simply choosing a particle uniformly at random and flipping the spins within the cluster to which that particle belongs. There is, however, another mathematically equivalent way to construct the cluster to be flipped that does not require every bond variable to be sampled, and is therefore computationally more efficient. We do not provide details here but refer the interested reader to Section 5.2.3 of [Krauth \(2006\)](#). [Wolff \(1989b\)](#) reports superiority of the single cluster approach through simulations on a 64^3 lattice near the critical temperature.

[Nott and Green \(2004\)](#) applied the Swendsen–Wang approach to Bayesian variable selection. In Bayesian variable selection auxiliary variables $\gamma_j \in \{0, 1\}^p$ are often introduced, allowing a spike-and-slab prior to be placed on each β_j by specifying that if $\gamma_j = 0$ then $\beta_j = 0$ with probability 1, and otherwise $\beta_j | (\gamma_j = 1)$ has a continuous prior distribution. When appropriate priors are chosen the marginal posterior distribution for $\gamma \in \{0, 1\}^p$ can often be obtained directly, meaning that Markov chain Monte Carlo methods can be employed directly on this space. As in the Ising model, the result is a distribution to be sampled from that is defined over a large space of correlated random variables that can each take two possible values. [Nott and Green \(2004\)](#) found that the Swendsen–Wang approach can yield substantial improvements compared to component-wise Metropolis in the presence of high multi-collinearity. Improvements were not always observed, however. The models are not of course identical. In particular the correlations between γ_j variables in Bayesian variable selection are dictated by the data, and therefore some can be high and others low in an unstructured manner. This contrasts with the rigid correlations imposed by the lattice structure of the Ising model discussed here.

Following the success of the cluster approach, similar ideas were used to design algorithms for the hard-disk and other soft-matter models ([Dress and Krauth, 1995](#)). A different strategy has, however, proved more successful in these systems, which we turn to next.

5.2 Jaster's algorithm for hard disks and spheres

A key weakness of the Metropolis algorithm applied to the hard-disk model at high particle density is that randomly perturbing the position of a particle is very likely to cause overlap with another, leading to a rejected move. Jaster (1999a,b) proposed a simple approach to combat this. In Jaster's algorithm an initial uniform innovation u is drawn and a particle i is selected uniformly at random. The first proposal is then to move particle i to $x_i \oplus u$, as in the Metropolis algorithm. If this results in an overlap with particle j , then a new position $x_j \oplus u$ is proposed for this particle. The process continues until either a configuration is found in which no particles overlap, a particle overlaps with more than one other, or a pre-specified maximum number of attempted moves have been made without finding a new non-overlapping configuration. If the first of these three scenarios occurs then the new configuration is accepted, and in either of the others it is rejected.

Jaster's algorithm can be described as a particular case of the delayed-rejection algorithm (Mira, 2001a) from statistics. It can also be cast in the more recent sequential proposal Markov chain Monte Carlo framework of Park and Atchadé (2020). Full details of this are given in Section 7.5 of Andrieu, Lee and Livingstone (2020).

Jaster's algorithm is an improvement on Metropolis in the sense that the probability of rejection is strictly lower. This in turn improves the asymptotic variance of the resulting Markov chain per iteration (e.g. Mira (2001b)), although each iteration is now also more expensive. The improvements, however, are often small in the case of high-density particle systems, in which it is very likely that the first stage proposal will result in the active particle i overlapping with more than one other. Jaster acknowledges this and suggests a modification in which the same particle is moved by very small amounts in one direction at each iteration, in order to reduce the chances of multiple particle overlaps. The idea has since been fully developed and will be introduced in the next section.

5.3 Event chain Monte Carlo

Event chain Monte Carlo can be viewed as a natural innovation of Jaster's algorithm to alleviate the issue of collisions involving more than two disks. The central idea is embedded in Jaster's remark that if disks are perturbed by smaller increments then configurations involving multiple disk overlaps are less likely to occur. The same logic suggests that the limiting continuous-time algorithm in which a single particle makes infinitesimally small moves (until collision) would completely remove the danger of collisions involving more than two particles. The resulting sampling algorithm applied to hard-disks is the event chain Monte Carlo algorithm of Bernard, Krauth and Wilson (2009), which was later extended to general potentials by Michel, Kapfer and Krauth (2014).

The event chain Monte Carlo algorithm simulates a continuous-time stochastic process known as a piecewise deterministic Markov process (PDMP), which involves deterministic dynamics and jumps, but no diffusive behaviour. This form often lends itself well to exact simulation. PDMPs have also been proposed as sampling algorithms in statistics, notably by Bouchard-Côté, Vollmer and Doucet (2018); Bierkens, Fearnhead and Roberts (2019). A recent review is given in Fearnhead et al. (2018).

5.3.1 Event chain Monte Carlo for the hard-disk model. To simulate event chain Monte Carlo for the hard-disk model an initial configuration and active particle $i \in \{1, \dots, N\}$ must be chosen. A single-particle velocity u is then simulated from some initial two-dimensional distribution on the unit circle. The configuration of the process $x(t)$ at time $t \geq 0$ is then determined in a manner that has some parallels with molecular dynamics simulations for the hard-disk model. The active particle i moves at unit speed in the direction u while all other particles remain still. The time to the first collision can then be calculated using (24) as in a molecular dynamics simulation. When a collision occurs the active particle is updated.

This can be described mathematically by defining the flow operator

$$(33) \quad \phi_t(x, v, i) := (x + tv, v, i)$$

for configuration x with N -particle velocity vector $v \in \mathbb{R}^{2N}$, which will have only two non-zero entries corresponding to the velocity u of the active particle i . We define the colliding particle as

$$c(x, i) := \begin{cases} \arg \min_{j \neq i} \mathfrak{d}(x_i, x_j) & \text{if } \min_{j \neq i} \mathfrak{d}(x_i, x_j) \leq 2\sigma \\ i & \text{otherwise.} \end{cases}$$

It is of course still possible that two particles could be equidistant from i , meaning $c(x, i)$ takes multiple values. At equilibrium this is a measure-zero event, but care must be taken to initialise the algorithm so that this does not occur. Upon collision the velocity can be updated in numerous ways. Of course it must be transferred such that the active particle is now $c(x, i)$. But the non-zero part of the velocity u can also be modified. In *reflected event chain Monte Carlo* it is adjusted according to the angle of the collision, as in a molecular dynamics simulation. This strategy did not perform as well, however, as the version known as *straight event chain Monte Carlo* in which u is not modified after collisions. In this case the new N -particle velocity $s(v, i, c(x, i))$ is calculated by simply swapping the i th and $c(x, i)$ th two-dimensional components of v , which results in the $c(x, i)$ th component becoming u and the i th component becoming zero. The *swap-upon-collision* operator can therefore be defined as

$$S(x, v, i) := (x, s(v, i, c(x, i)), c(x, i)).$$

Iterating the maps ϕ_ε for some small $\varepsilon > 0$ and S then leads to an algorithm in which a particle moves in a straight line until collision, at which point the active particle and velocity vector are updated. Defining the augmented state $z := (x, v, i)$, event chain dynamics can therefore be formally defined as $z(t) := \xi_t(z(0))$, where

$$\xi_t(z(0)) := \lim_{\varepsilon \rightarrow 0} [S \circ \phi_\varepsilon]^{\lfloor t/\varepsilon \rfloor}(z(0))$$

for any $t \geq 0$, which is a completely deterministic trajectory through time.

The event chain algorithm also involves a final ‘refreshment’ step, which plays a similar role to momentum refreshment in hybrid Monte Carlo. At certain times the single-particle velocity u is changed in some way. This is to help the process reach equilibrium. This refreshment can be done in several ways, one example being complete uniform re-sampling from the unit circle. In the best-performing implementation, however, the initial value of u is chosen to be either $(1, 0)$ or $(0, 1)$, each with probability $1/2$, and each refreshment simply entails swapping the elements of u (Bernard, Krauth and Wilson, 2009). This implementation is called the *xy*-version of event chain Monte Carlo, as the active particle will always be travelling parallel to either the x - or y -axis.

If the single-particle velocity refreshments are implemented according to the *xy*-version of the algorithm at fixed times, then the event chain algorithm for hard-disk systems is in fact completely deterministic. If, however, the refreshment times occur according to a Poisson process, then some stochasticity is introduced.

5.3.2 (Generalized) Event chain Monte Carlo for smooth potentials. In the hard-disk model the form of the potential leads to a natural definition of both collisions and updates of the active particle. It is not immediately obvious, however, how to extend these ideas to sampling from continuous potentials. Fortunately these hurdles were overcome by Michel, Kapfer and Krauth (2014) following earlier work from Peters and de With (2012).

We will consider the case of a potential composed of generic pairwise components $U_g(x_i, x_j)$. As in the case of hard-disks let $z := (x, v, i)$ be the augmented state of the process. The algorithm proceeds by defining a Poisson process for each pair of particles. The process associated with particles j and k has event rate defined by the function

$$\lambda_{jk}(z) := \beta \mathbb{I}(i = j) \max(0, \langle \nabla_x U_g(x_j, x_k), v \rangle).$$

This will only be non-zero if j is the current active particle. If an event associated with this process occurs then the active particle is swapped from j to k . As in the hard-disk model this

is achieved by swapping the j th and k th d -dimensional components of the velocity vector v . In the absence of an event the process advances deterministically according to the flow map (33), which results in the active particle moving with constant velocity v with all others remaining in place. The algorithm also includes a velocity refreshment step.

5.3.3 The generator and its properties. Recall that the infinitesimal generator of a continuous-time Markov process is defined point-wise as $Lf(z) := \lim_{\delta \rightarrow 0} \delta^{-1} (\mathbb{E}[f(z(t+\delta)) | z(t) = z] - f(z))$, whenever such a limit is well-defined (we will skip the technicalities here in favour of a more intuitive discussion). The generalized event chain algorithm can be described through the PDMP with infinitesimal generator (34)

$$Lf(z) := \langle v, \nabla_x f(z) \rangle + \sum_{j < k} \lambda_{jk}(z) [f(S_{jk}(z)) - f(z)] + \lambda_{\text{ref}} \int [f(z') - f(z)] R(z, dz')$$

where

$$S_{jk}(z) := (x, s(v, j, k), k)$$

and the function $s(v, j, k)$ swaps the j th and k th (d -dimensional) elements of v . In the final term on the right-hand side $\lambda_{\text{ref}} > 0$ denotes the refreshment rate for the active velocity component, and $R(z, dz')$ is a Markov kernel only changing the active component of the velocity. In different implementations R can either perform uniform refreshment for $v_i = u$ on the unit sphere or apply the xy transformation, which in the d -dimensional case we will treat as shifting each element of u one space to the right modulo d .

Properties of the Markov process associated with the event chain algorithm can be extracted by studying (34) as an operator on a suitably defined Hilbert space. Here we will consider the space $L^2(\mu)$, where $\mu(dz)$ is the product measure formed by combining $\pi(dx)$ with the discrete uniform distribution on $\{1, \dots, N\}$ for i and a uniform distribution on the $(Nd - 1)$ unit sphere for v (an alternative choice is the conditional probability measure for $v|i$ for which the i th d -dimensional component of v is uniform on the unit $(d - 1)$ -sphere and all other elements are 0). Recall that on such a space a μ -reversible process will have the property that the associated generator is self-adjoint, meaning $\langle f, Lg \rangle_\mu = \langle Lf, g \rangle_\mu$, where $\langle f, g \rangle_\mu := \int f(z)g(z)\mu(dz)$ is the $L^2(\mu)$ inner product. The operator (34) is not self-adjoint, meaning the process is not μ -reversible. We can, however, consider a more general property introduced in [Andrieu and Livingstone \(2021\)](#) known as (μ, Q) -self-adjointness. This means that there is another operator Q on $L^2(\mu)$ satisfying $Q^2 = I$ and $\langle f, g \rangle_\mu = \langle Qf, Qg \rangle_\mu$ (called an *isometric involution*) and for which

$$(35) \quad \langle f, Lg \rangle_\mu = \langle QLQf, g \rangle_\mu.$$

Clearly making the choice $Q = I$ equates (35) with μ -reversibility, but other choices are possible. It is shown in Appendix E of [Andrieu, Lee and Livingstone \(2020\)](#) that the generator associated with event chain Monte Carlo satisfies (35) with $Qf(x, v, i) := f(x, -v, i)$ when the refresh kernel R is taken to be uniform. We extend this to the xy implementation below when $d = 2$ (a proof is provided in the supplement). For this particular choice (μ, Q) -self-adjointness can be related to the notion of skew-detailed balance for a discrete time Markov chain with a velocity component (e.g. [Vucelja \(2016\)](#); [Turitsyn, Chertkov and Vucelja \(2011\)](#)), since the associated transition kernel $P_t(z, dz)$ satisfies skew-detailed balance for any choice of $t \geq 0$ (see Theorem 9 in [Andrieu and Livingstone \(2021\)](#)).

PROPOSITION 3. *The event chain Monte Carlo infinitesimal generator (34) with xy -refreshments is (μ, Q) -self-adjoint with the choice $Qf(x, v, i) := f(x, -v, i)$ when $d = 2$.*

REMARK 1. Note that Proposition 3 is not wholly satisfactory because in the xy -version of event chain Monte Carlo the active velocity component u only ever takes one of d values on the $(d - 1)$ -dimensional unit sphere (in contrast with uniform refreshment on the unit sphere). The algorithm therefore satisfies (μ, Q) -self-adjointness with the indicated μ when

$d = 2$, but is not μ -irreducible, meaning that μ is not the limiting distribution for the chain. More generally no rigorous proof of ergodicity for the xy -version of the straight event chain algorithm is known to the authors at the time of writing.

REMARK 2. The potential associated with the hard-disk model is not smooth, meaning a generator-level definition akin to (34) is not straightforward. One route to such an object is to consider a sequence of processes associated with the soft-disk potential (16) indexed by k , and then letting $k \rightarrow \infty$. Similar ideas underpin the extension of Hamiltonian Monte Carlo to non-smooth models discussed in Nishimura, Dunson and Lu (2020). Recent work by Monemvassitis, Guillin and Michel (2023) has provided a generator-level description via a different approach and established conditions under which ergodicity can be shown for the uniform refreshment process.

5.3.4 *Implementation details.* The generator-level description given in (34) assumes that the potential can be broken into factors at the pair-wise level, but in practice other factorisation schemes are possible. In the case of pair-wise factorisation the continuous-time algorithm can be derived by taking an appropriate limit of a discrete-time Metropolis algorithm in which the usual acceptance rate is replaced with the *factorized Metropolis filter* $\prod_{i \neq j} \min(1, e^{-\beta(U_g(x'_i, x'_j) - U_g(x_i, x_j))})$ (Michel, Kapfer and Krauth, 2014). This is of course less efficient in the sense of Peskun as fewer proposed moves will be accepted, but has a computational advantage as each component of the acceptance rate/event rate only requires evaluation of one pair-wise interaction.

To simulate a PDMP in practice it must be possible to either directly simulate from a Poisson process with intensity $\sum_{j \neq k} \lambda_{ij}(z)$, or to establish a tractable upper bound and then perform thinning. One approach to the latter is called the *cell veto* method (Kapfer and Krauth, 2016), in which the manifold \mathcal{M} is partitioned into cells and a local upper bound is found within each cell. This approach has proved to be particularly effective for systems in which particles interact over long distances and has been applied by Faulkner et al. (2018) to the all-atom model of water presented in Section 3.6. In this case the Metropolis algorithm requires prohibitively high per iteration costs, and molecular dynamics approaches are numerically unstable unless a very small step-size is chosen. Event chain Monte Carlo circumvents both issues, although work is still ongoing to improve molecular rotational mixing within the algorithm. The cell-veto method has connections to a recently proposed approach for simulating PDMPs in the statistics literature by Corbella, Spencer and Roberts (2022).

Alternative events are possible other than simply swapping the active particle using either the straight or reflected event chain strategies described above. Michel, Durmus and Sénécal (2020) introduce *forward* event chain Monte Carlo, in which the velocity v is stochastically perturbed in a prescribed way when a collision event occurs. The motivation is that if enough randomness is introduced during this step then the algorithm can perform well even without introducing additional velocity refreshment events. Klementa and Engel (2019) and Höllmer, Maggs and Krauth (2022) employ similar ideas in *Newtonian* event chain Monte Carlo. We add that parallel implementations of event chain Monte Carlo involving multiple active particles have also been considered in Kampmann, Boltz and Kierfeld (2015b).

5.4 The Xtra chance algorithm

The philosophy of continuing on the same path upon a rejection in the hope of reaching acceptance has also been proposed by Campos and Sanz-Serna (2015) in the context of hybrid Monte Carlo. In this algorithm Hamiltonian dynamics are numerically simulated for a period of time $T := \ell\varepsilon$ and then an accept-reject decision is taken by sampling $u \sim \mathcal{U}[0, 1]$ and assessing whether or not $u \leq \exp[-\beta(H(x', p') - H(x, p))]$, where $(x', p') := \psi_\ell^\varepsilon(x, p)$ is the Hamiltonian proposal. Upon rejection, however, a second-stage proposal is computed as $(x'', p'') := \psi_\ell^\varepsilon(x', p')$. In other words, the dynamics are simulated for an additional time $T = \ell\varepsilon$. This second stage proposal is accepted if $u < \exp[-\beta(H(x'', p'') - H(x, p))]$, where u is the same uniform random variable used in the first stage accept-reject decision. The

scheme has the appealing interpretation that first a u is simulated, and then proposals are repeatedly tested until one is encountered for which the Metropolis ratio is larger than u . Typically a maximum number of attempts are tried until this goal is attained, otherwise all proposals are rejected.

The Xtra chance algorithm was designed so that the expensive computations associated with simulating Hamiltonian dynamics are re-used in the case of a rejection, as the end point of this simulation is used as the starting point for the next stage proposal. As in the case of Jaster’s algorithm the probability of rejection is strictly decreased compared to ordinary hybrid Monte Carlo, but similarly more computation is associated with each iteration. The Xtra Chance algorithm can also be regarded as a particular case of the sequential proposals algorithm of [Park and Atchadé \(2020\)](#), as is remarked in that work. In addition it can again be fitted into the delayed rejection framework of [Mira \(2001a\)](#). One surprising feature of the Xtra chance algorithm when viewed through the lens of delayed rejection is its attractively simple acceptance rate at the second stage (and beyond). In contrast, the usual acceptance rate for a delayed rejection algorithm is

$$\alpha_2(x, x', x'') = \min \left(1, \frac{\pi(x'')q_1(x'', x')q_2(x'', x', x)(1 - \alpha_1(x'', x'))}{\pi(x)q_1(x, x')q_2(x, x', x'')(1 - \alpha_1(x, x'))} \right),$$

where $q_1(x, \cdot)$ and $q_2(x, x', \cdot)$ are the first and second stage proposal kernels and $\alpha_1(x, x')$ is the usual Metropolis–Hastings acceptance probability for the first stage proposal x' . In addition, a fresh uniform random variable must be drawn to decide whether or not to accept the second stage proposal as compared to that used in the first stage. The reason that a much simpler algorithm can be used in both the Xtra chance and Jaster algorithm is in part owing to the symmetries of the dynamics of the transition, but also to the augmented *slice sampler* target density $\mu(x, u) := \mathbb{I}(u < \pi(x))$ ([Neal, 2003](#)). When viewing delayed rejection with μ as the target distribution, the acceptance rates for these algorithms reduce to being either 1 or 0, and the persistent uniform sample that determines when a proposal is accepted is nothing more than a transformed sample from the conditional distribution of $u|x$. [Andrieu and Livingstone \(2021\)](#) prove that taking extra chances in this manner reduces the asymptotic variance of ergodic averages (the result does not, however, account for computational cost).

5.5 Shadow hybrid Monte Carlo

[Izaguirre and Hampton \(2004\)](#) introduced a modification to the hybrid Monte Carlo method that was later developed and introduced to the statistics community by [Radivojević and Akhmatkaya \(2020\)](#). The shadow hybrid Monte Carlo method is motivated by the field of backward error analysis for ordinary differential equations (ODEs). Given an ODE $\dot{x}(t) = f(x(t))$ and a numerical scheme for simulating this ODE, *forward* error analysis is concerned with understanding how far apart the numerical and exact solutions of the ODE are at some time step t , usually as a function of the numerical step-size $\varepsilon > 0$. A scheme is called p th order accurate if this global error can be bounded by some function $C(t)\varepsilon^p$. The idea of *backward* error analysis is to instead seek a *modified* ODE system $\dot{x}(t) = \tilde{f}_\varepsilon(x(t))$ for which the numerical scheme that is used for the original ODE will be accurate to a higher order. The differences between \tilde{f}_ε and f can then be studied to understand qualitative differences in behaviour between the numerical scheme and the true solution to the original ODE. As a simple example, consider numerically simulating the system $\dot{x}(t) = f(x(t))$ using a first order scheme. Under sufficient smoothness assumptions on f , finding a modified system $\dot{x}(t) = \tilde{f}_\varepsilon(x(t))$ for which the scheme is second order accurate involves setting $\tilde{f}_\varepsilon(x) = f(x) + \varepsilon f_1(x)$ and then considering a Taylor series expansion of the one step error $\tilde{x}(t + \varepsilon) - \tilde{x}(t)$ to choose an f_1 that results in a cancellation of the relevant lower order terms in ε (here $\tilde{x}(t)$ denotes the true solution of the modified system). See e.g. Chapter 5 of [Leimkuhler and Reich \(2004\)](#) for a more detailed explanation and examples.

Backward error analysis has found success in the study of symplectic numerical schemes for Hamiltonian systems, such as the velocity Verlet approach introduced in Section 4.3. Using this approach it is possible to show that in many cases the modified system is also Hamiltonian, and that the numerical scheme preserves the value of the modified or *shadow* Hamiltonian over surprisingly long time scales ([Leimkuhler and Reich, 2004](#), Chapter 5). This gives

some insight into the success of the hybrid Monte Carlo method. It also presents opportunities to develop alternative approaches. The shadow hybrid Monte Carlo method works by simulating a hybrid Monte Carlo algorithm, but in place of the true Hamiltonian H in the acceptance rate a chosen shadow Hamiltonian \tilde{H}_k of order k is used (the choice of order usually depends on trading off accuracy with computational cost). The shadow Hamiltonian \tilde{H}_k will typically also depend on the momentum in a less straightforward way than H , meaning that in place of re-sampling the momentum directly from its marginal distribution a proposed change in momentum is drawn (from its marginal distribution in the true Hamiltonian system) and then accepted or rejected. [Izaguirre and Hampton \(2004\)](#) propose to use rejection sampling, whereas [Radiwojević and Akhmatskaya \(2020\)](#) introduce a Metropolis step. Rather than simple Monte Carlo averages of the algorithm output, which would give expectations with respect to the modified distribution with density $\propto e^{-\beta\tilde{H}_k}$, importance sampling estimators using weights $e^{-\beta(H-\tilde{H}_k)}$ can be employed, allowing expectations with respect to the true distribution of interest to be computed. [Radiwojević and Akhmatskaya \(2020\)](#) consider several further modifications to the original algorithm, including different numerical integrators as introduced in [Radiwojević et al. \(2018\)](#) and partial momentum refreshment as proposed in ordinary hybrid Monte Carlo by [Horowitz \(1991\)](#).

REMARK 3. Combining importance sampling and Markov chain Monte Carlo in this way has also been proposed and studied in the statistics literature (e.g. [Franks and Vihola \(2020\)](#); [Vihola, Helske and Franks \(2020\)](#)). The intriguing property of the above scheme is the level of stability provided to the importance weights by the modified Hamiltonian. Traditional importance sampling typically performs poorly in high-dimensional settings, but there is much numerical evidence that this is not true of shadow hybrid Monte Carlo.

5.6 Sampling with modified kinetic energies

The underdamped Langevin diffusion (30) can be generalized in such a way that the invariant density for momentum is changed while that for the position variable x remains $\pi(x) \propto e^{-\beta U(x)}$. The resulting system of stochastic differential equations can be written analogously to (30) as

$$(36) \quad \begin{aligned} dx_i(t) &= \nabla_i K(p(t)) dt \\ dp_i(t) &= F_i(x(t)) dt - \gamma \nabla_i K(p(t)) dt + \sqrt{2\gamma\beta^{-1}} dW_i(t). \end{aligned}$$

where K is some *kinetic energy* function. The choice $K(p) = p^T M^{-1} p / 2$ leads to (30), where M is a diagonal matrix with entries $M_{ii} = m_i$. [Artemova and Redon \(2012\)](#) propose to modify the standard quadratic form choice of K in such a way that a particle does not move if it has momentum below a chosen threshold $p_{\min} > 0$, with the aim of reducing computational cost by freezing particles in place unless they will move by an appreciable amount. This is achieved by setting $K(p) := \sum_i k(p_i)$ where

$$k(p_i) := \begin{cases} 0 & \|p_i\|_2 < p_{\min} \\ \frac{1}{2m_i} p_i^T p_i & \|p_i\|_2 > p_{\max}. \end{cases}$$

In between the constants p_{\min} and p_{\max} the function smoothly interpolates between 0 and $p_{\max}^2 / (2m_i)$ ([Stoltz and Trstanova, 2018](#), Section 4.1.1). Ergodicity properties of (36) are studied in [Redon, Stoltz and Trstanova \(2016\)](#) and guidelines for the choice of step-size as compared to the standard choice of kinetic energy are given in [Stoltz and Trstanova \(2018\)](#). A parallel implementation is introduced in [Singh, Marin and Redon \(2017\)](#).

Experiments in [Artemova and Redon \(2012\)](#) indicate that the resulting *adaptively restrained* Langevin dynamics induce stronger autocorrelations over time among the particle positions, which is intuitive given that their movement is restricted. This is more than offset, however, by the seven-fold computational speed up exhibited by the restrained dynamics when compared to the standard underdamped Langevin approach, resulting in around four

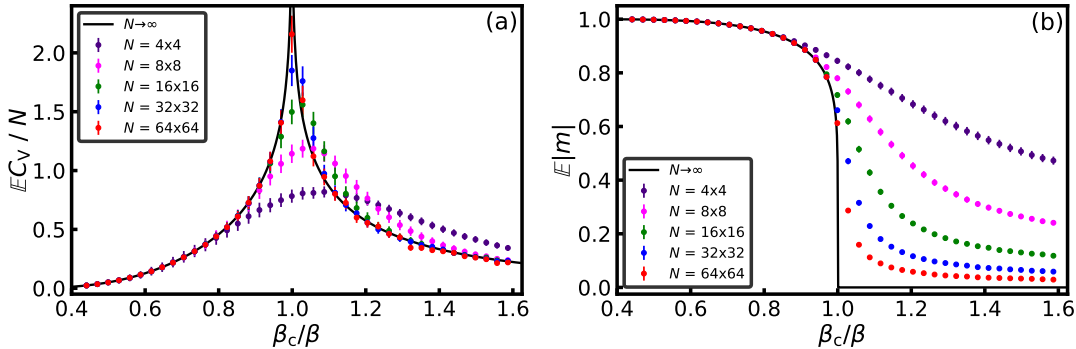


FIG 9. Estimations of the expected specific heat per particle (a) and expected absolute magnetic density (b) of the two-dimensional Ising model as functions of β_c/β and number of particles N at $h = 0$. Results were generated from 10^4 Wolff samples (10^4 burn-in iterations were discarded) and then averaged over 28 simulations. The black curves in (a) and (b) correspond to the analytical thermodynamic predictions in (9) and (12).

times better accuracy overall when comparing the error in estimating some chosen test functions in an example simulation of 343 Argon particles in which intra-molecular interactions were modelled using a Lennard–Jones potential. [Stoltz and Trstanova \(2018\)](#) also consider the problem of metastability when sampling from multi-modal distributions, and provide numerical evidence that this can be reduced through an appropriate choice of kinetic energy.

REMARK 4. The closely related idea of modifying the choice of kinetic energy within hybrid Monte Carlo has been considered in the statistics and machine learning literature in [Lu et al. \(2017\)](#); [Zhang et al. \(2016\)](#) and studied in terms of ergodicity properties in [Livingstone, Faulkner and Roberts \(2019\)](#). The main focus of these works, however, is on the increased robustness and numerical stability that can be achieved by using a slower growing kinetic energy than the standard choice. [Nishimura, Dunson and Lu \(2020\)](#) also consider different choices of kinetic energy in order to sample from discrete distributions and those with discontinuous potentials.

6. SIMULATION STUDIES

In this section we present simulation studies of the two-dimensional Ising and XY models, comparing the Metropolis algorithm with two of the advanced algorithms presented in Section 5. For storage reasons, only an N -skeleton of the Metropolis chain is retained, corresponding to storing the system state after each N -particle sweep. The elapsed simulation time during each N -particle sweep defines one unit of Metropolis simulation time Δt_{Metrop} . The elapsed simulation time during each Wolff iteration defines one unit of Wolff simulation time Δt_{Wolff} .

6.1 Two-dimensional Ising model

Figure 9 shows estimates of the expected zero-field ($h = 0$) specific heat per particle (see (5)) and expected zero-field absolute magnetic density (see (11)) of the two-dimensional Ising model, as functions of β_c/β and N (recall that $\beta_c := \ln(1 + \sqrt{2})/(2J)$ is the inverse critical temperature). The output was generated using the Wolff algorithm and provides evidence for the phase transition predicted by [Onsager \(1944\)](#) at $\beta = \beta_c, h = 0$. The specific-heat output in Figure 9(a) appears to approach the analytical thermodynamic prediction in (9) with increasing N . The magnetic-density output in Figure 9(b) suggests that the expected zero-field absolute magnetic density tends to the spontaneous magnetic density $m_0(\beta J)$ defined in (12) and corresponding to the solid black line in the figure. The output tends to one for all N as $\beta \rightarrow \infty$ because the Boltzmann–Gibbs distribution puts all probability mass on two equally likely states ($x_i = 1$ for all i and $x_i = -1$ for all i) in this limit, while it appears to tend to zero in the thermodynamic limit for all $\beta < \beta_c$ because the magnetic density m satisfies

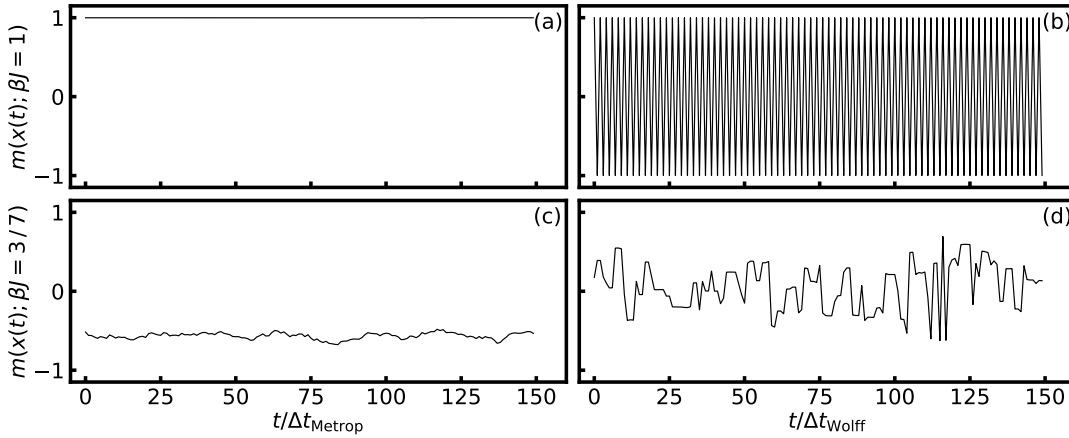


FIG 10. Zero-field magnetic density $m(x(t); \beta, J, h = 0, N)$ versus normalised simulation time $t/\Delta t_{\text{Metrop/Wolff}}$ for $N = 64 \times 64$ particles at low temperature ($\beta J = 1 > \beta_c J$) and near β_c ($\beta J = 3/7$) using both the Metropolis (a, c) and Wolff (b, d) algorithms. In each case 10^4 burn-in iterations were discarded.

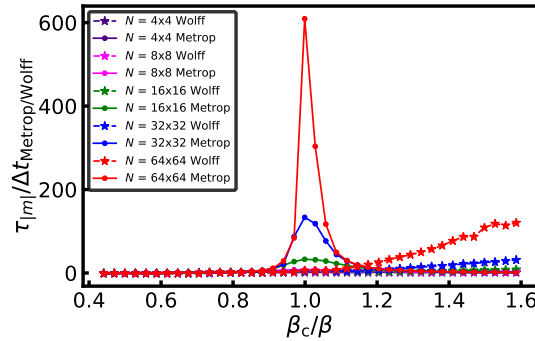


FIG 11. Estimations of the normalised magnetic-norm integrated autocorrelation time $\tau_{|m|}/\Delta t_{\text{Metrop/Wolff}}$ of the two-dimensional Ising model as a function of β_c/β and number of particles N , with respect to the Metropolis and Wolff algorithms. Autocorrelation functions were generated from 10^4 samples (10^4 burn-in iterations were discarded) and then averaged over 28 simulations.

$\mathbb{P}[m(x; \beta, J, h = 0, N) \neq 0] \rightarrow 0$ as $N \rightarrow \infty$ for all $\beta < \beta_c$. Both outputs become noisier as β becomes smaller and N larger, which we discuss below.

We now compare the Metropolis and Wolff algorithms in the context of *spontaneous symmetry breaking*. At $h = 0$ the potential is symmetric in x for all β, J, N , but numerical simulations (of the system) constrained to single spin flips spontaneously break this Z_2 symmetry at finite $\beta > \beta_c$, leaving the system stuck close to one of the two $\beta \rightarrow \infty$ states on a timescale that diverges with N . This is an example of spontaneous symmetry breaking and is reflected in the low-temperature ($\beta > \beta_c$) zero-field magnetic-density trace plots in Figures 10(a) and (b). On the presented simulation timescale, the Metropolis simulation starts at $m = 1$ and stays in this state, while the Wolff simulation mixes between $m = 1$ and $m = -1$. This case of symmetry breaking is caused by the multi-modality of the Boltzmann–Gibbs distribution, but anisotropic distributions can also result in the phenomenon (e.g. the asymmetric Metropolis simulation output in Section 6.2).

The Wolff algorithm also combats the related phenomenon of *critical slowing down*. As β_c is approached from small values of β , strong particle–particle correlations begin to set in on increasingly long lattice-site-separation distances. Near β_c , this results in increasingly large clusters of particles with the same spin value, which slows mixing significantly for simulations constrained to single spin flips, resulting in very noisy \tilde{m} statistics. This is remedied by flipping large clusters of spins, as reflected in the output in Figures 10(c-d) and 11. Figures 10(c) and (d) show (respectively) trace plots of the magnetic density at small $|\beta - \beta_c|$

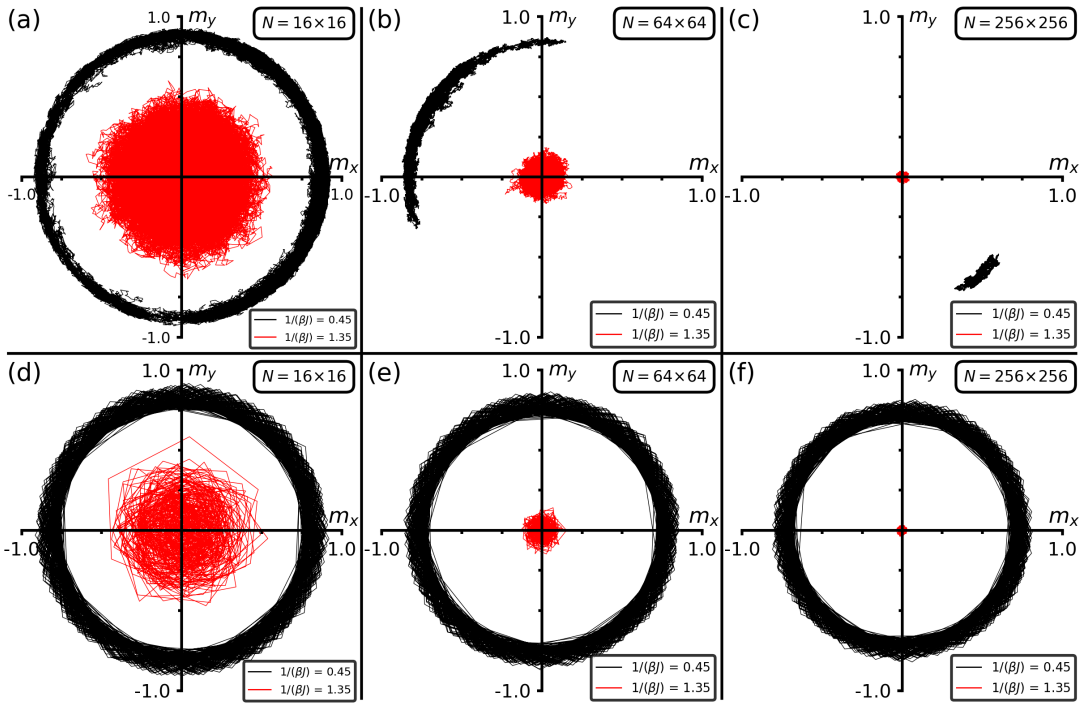


FIG 12. Evolution of the magnetic density m of the two-dimensional XY model using the Metropolis (a-c) and event-chain (d-f) algorithms. The size of the Metropolis skeleton chain is 10^5 (with acceptance probability $\simeq 0.6$) while that of the event-chain algorithm is 10^3 .

using the Metropolis and Wolff algorithms. The Metropolis output is strongly time correlated, while the Wolff output mixes on the presented simulation timescale. Figure 11 shows estimates of the magnetic-norm integrated autocorrelation times $\tau_{|m|}$ of the zero-field two-dimensional Ising model as functions of β_c/β and N , using both the Metropolis and Wolff algorithms. The Metropolis estimates appear to diverge with system size as $\beta \rightarrow \beta_c$, while those of the Wolff algorithm depend only weakly on N for all $\beta_c/\beta < 1.2$, before developing a stronger N -dependence at smaller β , due to smaller typical cluster sizes. This results in the development of noise in the output in Figure 9 at small β . Physicists tend to remove any N -dependence from the Wolff timescale by multiplying it by some metric for the typical cluster density at any given temperature (e.g. Tamayo, Brower and Klein (1990)).

6.2 Two-dimensional XY model

In addition to its use in the development of the theory of the two-dimensional melting transition presented in Section 3.2, event chain Monte Carlo applied to the XY model has also enjoyed success analogous to that of the Wolff algorithm for the Ising model. Figure 12 shows the evolution of the two-dimensional magnetic density vector (as defined in (11)) of the zero-field two-dimensional XY model (presented in Section 3.1.3) using both the Metropolis and event chain Monte Carlo algorithms. The expected value is $(0, 0)$ at all nonzero temperatures, but the low-temperature Metropolis mean converges to this expected value on a timescale that diverges with N , as described in detail in Faulkner (2022) and reflected in Figures 12(a-c) ('low temperature' corresponds to all finite $\beta > \beta_{\text{BKT}}$, where $\beta_{\text{BKT}} \simeq 1.13/J$ is the inverse critical temperature of the model). The event-chain output (Figures 12(d-f)) by contrast suggests N -independence for all β . The event chain algorithm was simulated for a time period proportional to N .

7. DISCUSSION

It is important to acknowledge that there has been much successful cross-pollination of ideas between statistical physics and statistics/machine learning for many decades now. The

perhaps unfortunately named *Gibbs sampler* is so-called from its initial use in sampling from Ising models applied to pixel distributions in black and white images (Geman and Geman, 1984). The now well-established use of energy-based models in image processing (LeCun et al., 2006) is an example of how such cross-pollination can lead to new approaches and insights. Nonetheless, we believe that there is still much that each field can learn from the other, and speculate on some possible avenues below.

The ergodicity properties of the most successful version of (generalized) event chain Monte Carlo are, to the best knowledge of the authors, still largely unknown (see Remark 1 of Section 5.3). Statisticians have recently found success in studying ergodicity properties of piecewise deterministic Markov processes (e.g. Bierkens, Roberts and Zitt (2019)) and may have something to offer here. In addition, there is a need to establish theoretical results comparing the straight and reflected versions of event chain Monte Carlo, as well as uniform versus xy refreshment of particle velocities, to support empirical findings. Statisticians have some track record in establishing orderings among sampling algorithms, and the accumulated knowledge within the field could be suitable for this task.

Often when hybrid (or *Hamiltonian*) Monte Carlo is described in statistics and machine learning, an analogy is given to the movement of a single Nd -dimensional fictitious ‘particle’ (e.g. Neal (2011)). Physicists, by contrast, think separately about the movement of each individual d -dimensional *physical* particle within a system. With this mindset interactions between particles become a natural consideration, and it may be that such a mindset is beneficial in the context of probabilistic models. One example is sampling under ordering constraints, such as learning the positions of histogram breaks (Worrall et al., 2022) or knot positions in a nonparametric regression model (Smith and Kohn, 1996). One-dimensional implementations of molecular or event chain dynamics would trivially preserve the necessary monotonicity requirements imposed by such models. Repulsive potentials such as Lennard–Jones could also be incorporated into prior distributions for parameters in a Bayesian setting to ensure that sufficiently different posterior distributions are attained, such as in the estimation of parameters in mixture models.

In the field of dimension reduction ideas from statistical physics are implicitly used. In multi-dimensional scaling (Torgerson, 1952) d -dimensional representations x_i for $1 \leq i \leq N$ of some higher dimensional data points y_1, \dots, y_N are sought in order to reduce the dimension of the problem. In metric multi-dimensional scaling, estimates for x_1, \dots, x_N are found by minimising the *stress* function $S(x_1, \dots, x_N) := \sum_{i \neq j} (d_{ij} - \mathfrak{d}(x_i, x_j))^2$, where d_{ij} is some appropriate notion of dissimilarity between y_i and y_j computed in the original high-dimensional space.⁴ The stress function, when viewed through the lens of statistical physics, is nothing more than a bonded potential. This observation has in fact already been noted (Andrecut, 2009), but there is still much scope to build on the connection. In particular there is recent interest in quantifying uncertainty in the parameters x_1, \dots, x_N through a Bayesian treatment (e.g. Ren et al. (2017)), and viewing the problem through the lens of Boltzmann–Gibbs distributions, in particular when combined with recent advances in model-free Bayesian inference (Bissiri, Holmes and Walker, 2016; Jewson, Smith and Holmes, 2018; Knoblauch, Jewson and Damoulas, 2022), may be a fruitful avenue for this problem.

This article has emphasised that the study of phase transitions is central to statistical physics. In a typical problem, samples are drawn from several different Boltzmann–Gibbs distributions, each corresponding to a different fixed temperature (or some other fixed hyperparameter of the system), and the qualitative behaviour of some observable is studied as a function of temperature. One example of a ‘phase transition’ in statistics/machine learning is the estimator for a coefficient value in L1-penalised regression, as a function of the weight λ . The inverse temperature parameter is also present (albeit with the addition of a prior) as the *learning rate* in Gibbs posterior distributions (e.g. Haddouche et al. (2021); Bissiri, Holmes and Walker (2016); Syring and Martin (2020); Grünwald and Van Ommen (2017)). Typically the goal has been to choose a unique optimal value for this parameter (Wu and Martin,

⁴Often the square root of S is defined as the stress, but this operation has no impact on the minima.

2022), but summary statistics of interest could also be treated as ‘observables’ and studied for evidence of phase transitions as the learning rate varies.

There are numerous approaches to sampling from Boltzmann–Gibbs distributions that were omitted for brevity. Nothing is mentioned of thermostats, such as Nosé–Hoover dynamics and extensions (e.g. [Evans and Holian \(1985\)](#); [Martyna, Klein and Tuckerman \(1992\)](#); [Leimkuhler, Noorzadeh and Theil \(2009\)](#)). Details of this approach and a simulation study comparing to hybrid Monte Carlo are provided in [Cances, Legoll and Stoltz \(2007\)](#). Biasing methods such as the Wang–Landau algorithm ([Wang and Landau, 2001](#)) and the adaptive biasing force method ([Darve, Rodríguez-Gómez and Pohorille, 2008](#)) were also not discussed. We refer the interested reader to book-level treatments of the subject such as [Lelièvre, Rousset and Stoltz \(2010\)](#); [Leimkuhler and Matthews \(2016\)](#); [Krauth \(2006\)](#).

We have also omitted the estimation of dynamical properties. Barrier-crossing techniques such as metadynamics ([Barducci, Bonomi and Parrinello, 2011](#)) or multicanonical sampling ([Berg and Neuhaus, 1991](#)) can be integrated into sampling algorithms to estimate the depths and shapes of potential wells, providing a route to order-of-magnitude estimates of chemical-reaction rates, rare-event timescales, and other similar phenomena. More direct dynamical quantities may also be estimated when simulating steady-state systems (such as transport coefficients describing particle flow in response to an applied field). These are typically the solution of a partial differential equation, though Metropolis simulations of such systems often produce remarkably accurate results (e.g. [Kaiser et al. \(2013, 2015\)](#)).

Finally, we have said very little here of algorithms that have been developed in statistics/machine learning but may have uses in statistical physics. Recent advances in sampling on discrete state spaces such as the locally-informed approach of [Zanella \(2020\)](#) and the non-reversible strategies outlined in [Power and Goldman \(2019\)](#) are a good example of this. Several methods for estimating normalising constants have also been developed, such as Annealed Importance Sampling ([Neal, 2001](#)). These methods could well be useful in the context of free energy calculations. We have deliberately restricted ourselves to algorithms that originated from statistical physics, but a future article could certainly be written introducing these methods to an audience of physicists.

Code and data availability. All code used in this article is freely available on GitHub at <https://github.com/michaelfaulkner/super-aLby>, commit hash [1c014ca](#) (Ising simulations) and <https://github.com/michaelfaulkner/xy-type-models>, commit hash [95093aa](#) (XY simulations). All published data can be reproduced using these applications (as outlined in each README) and are available at the University of Bristol data repository, [data.bris](#), at <https://doi.org/10.5523/bris.sju7uasr7e2b2n518hk72p3ur>.

Acknowledgements This work was conceived at the [Scalable inference; statistical, algorithmic, computational aspects](#) workshop (part of *i-like*, an EPSRC programme grant) at the Isaac Newton Institute for Mathematical Sciences, University of Cambridge. All simulations were performed on BlueCrystal 4 at the Advanced Computing Research Centre, University of Bristol. MFF acknowledges support from EPSRC fellowship EP/P033830/1. SL acknowledges support from EPSRC grant EP/V055380/1.

REFERENCES

- ALDER, B. J. and WAINWRIGHT, T. E. (1957). Phase Transition for a Hard Sphere System. *J. Chem. Phys.* **27** 1208.
- ALDER, B. J. and WAINWRIGHT, T. E. (1959). Studies in Molecular Dynamics. I. General Method. *J. Chem. Phys.* **31** 459.
- ALDER, B. J. and WAINWRIGHT, T. E. (1960). Studies in Molecular Dynamics. II. Behavior of a Small Number of Elastic Spheres. *J. Chem. Phys.* **33** 1439.
- ALDER, B. J. and WAINWRIGHT, T. E. (1962). Phase Transition in Elastic Disks. *Phys. Rev.* **127** 359.
- ANDRECUT, M. (2009). Molecular dynamics multidimensional scaling. *Physics Letters A* **373** 2001–2006.
- ANDRIEU, C., LEE, A. and LIVINGSTONE, S. (2020). A general perspective on the Metropolis-Hastings kernel. *arXiv preprint arXiv:2012.14881*.

- ANDRIEU, C. and LIVINGSTONE, S. (2021). Peskun–Tierney ordering for Markovian Monte Carlo: Beyond the reversible scenario. *Ann. Statist.* **49** 1958.
- ANDRIEU, C., DURMUS, A., NÜSKEN, N. and ROUSSEL, J. (2021). Hypocoercivity of piecewise deterministic Markov process-Monte Carlo. *The Annals of Applied Probability* **31** 2478–2517.
- ARCHAMBAULT, P., BRAMWELL, S. T. and HOLDSWORTH, P. C. W. (1997). Magnetic fluctuations in a finite two-dimensional XY model. *J. Phys. A* **30** 8363.
- ARTEMOVA, S. and REDON, S. (2012). Adaptively restrained particle simulations. *Physical Review Letters* **109** 190201.
- BARDUCCI, A., BONOMI, M. and PARRINELLO, M. (2011). Metadynamics. *WIREs: Comput. Mol. Sci.* **1** 826.
- BARKER, A. A. (1965). Monte Carlo calculations of the radial distribution functions for a proton-electron plasma. *Australian Journal of Physics* **18** 119–134.
- BAXTER, R. J. (2008). *Exactly Solved Models in Statistical Mechanics*. Dover Publications Inc.
- BEREZINSKII, V. L. (1971). 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.
- BERG, B. A. and NEUHAUS, T. (1991). Multicanonical algorithms for first order phase transitions. *Phys. Lett. B* **267** 249.
- BERNARD, E. P. (2011). Algorithms and applications of the Monte Carlo method : Two-dimensional melting and perfect sampling, PhD thesis, Université Pierre-et-Marie-Curie - Paris VI.
- BERNARD, E. P., KRAUTH, W. and WILSON, D. B. (2009). Event-chain Monte Carlo algorithms for hard-sphere systems. *Phys. Rev. E* **80** 056704.
- BERNARD, E. P. and KRAUTH, W. (2011). Two-Step Melting in Two Dimensions: First-Order Liquid-Hexatic Transition. *Phys. Rev. Lett.* **107** 155704.
- BETANCOURT, M. and GIROLAMI, M. (2015). Hamiltonian Monte Carlo for hierarchical models. *Current trends in Bayesian methodology with applications* **79** 2–4.
- BETANCOURT, M., BYRNE, S., LIVINGSTONE, S. and GIROLAMI, M. (2017). The geometric foundations of Hamiltonian Monte Carlo.
- BIERKENS, J., FEARNHEAD, P. and ROBERTS, G. (2019). The zig-zag process and super-efficient sampling for Bayesian analysis of big data. *The Annals of Statistics* **47** 1288–1320.
- BIERKENS, J., KAMATANI, K. and ROBERTS, G. O. (2018). High-dimensional scaling limits of piecewise deterministic sampling algorithms. *arXiv preprint arXiv:1807.11358*.
- BIERKENS, J., ROBERTS, G. O. and ZITT, P.-A. (2019). Ergodicity of the zigzag process. *The Annals of Applied Probability* **29** 2266–2301.
- BISSIRI, P. G., HOLMES, C. C. and WALKER, S. G. (2016). A general framework for updating belief distributions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **78** 1103–1130.
- BOU-RABEE, N. and SANZ-SERNA, J. M. (2018). Geometric integrators and the Hamiltonian Monte Carlo method. *Acta Numerica* **27** 113–206.
- BOUCHARD-CÔTÉ, A., VOLLMER, S. J. and DOUCET, A. (2018). The Bouncy Particle Sampler: A Nonreversible Rejection-Free Markov Chain Monte Carlo Method. *J. Am. Stat. Assoc.* **113** 855.
- BRAMWELL, S. T. and HOLDSWORTH, P. C. W. (1993). Magnetization and universal sub-critical behaviour in two-dimensional XY magnets. *J. Phys.: Condens. Matter* **5** L53.
- BRAMWELL, S. T., HOLDSWORTH, P. C. W. and PINTON, J. F. (1998). Universality of rare fluctuations in turbulence and critical phenomena. *Nature* **396** 552.
- CAMPOS, C. M. and SANZ-SERNA, J. M. (2015). Extra chance generalized hybrid Monte Carlo. *Journal of Computational Physics* **281** 365–374.
- CANCES, E., LEGOLL, F. and STOLTZ, G. (2007). Theoretical and numerical comparison of some sampling methods for molecular dynamics. *ESAIM: Mathematical Modelling and Numerical Analysis* **41** 351–389.
- CORBELLA, A., SPENCER, S. E. and ROBERTS, G. O. (2022). Automatic Zig-Zag sampling in practice. *arXiv preprint arXiv:2206.11410*.
- COULOMB, C. A. (1785). Second mémoire sur l'électricité et le magnétisme. *Histoire de l'Académie Royale des Sciences* 578.
- DALALYAN, A. S. and RIOU-DURAND, L. (2020). On sampling from a log-concave density using kinetic Langevin diffusions. *Bernoulli* **26** 1956–1988.
- DARVE, E., RODRÍGUEZ-GÓMEZ, D. and POHORILLE, A. (2008). Adaptive biasing force method for scalar and vector free energy calculations. *The Journal of chemical physics* **128** 144120.
- DE LEEUW, S. W., PERRAM, J. W. and SMITH, E. R. (1980). Simulation of electrostatic systems in periodic boundary conditions. I. Lattice sums and dielectric constants. *Proc. R. Soc. Lond. A* **373** 27.
- DELIGIANNIDIS, G., PAULIN, D., BOUCHARD-CÔTÉ, A. and DOUCET, A. (2021). Randomized Hamiltonian Monte Carlo as scaling limit of the bouncy particle sampler and dimension-free convergence rates. *The Annals of Applied Probability* **31** 2612–2662.
- DRESS, C. and KRAUTH, W. (1995). Cluster algorithm for hard spheres and related systems. *Journal of Physics A: Mathematical and General* **28** L597.
- DUANE, S., KENNEDY, A. D., PENDLETON, B. J. and ROWETH, D. (1987). Hybrid Monte Carlo. *Physics letters B* **195** 216–222.

- DURMUS, A., ENFROY, A., MOULINES, É. and STOLTZ, G. (2021). Uniform minorization condition and convergence bounds for discretizations of kinetic Langevin dynamics. *arXiv preprint arXiv:2107.14542*.
- EDWARDS, R. G. and SOKAL, A. D. (1988). Generalization of the Fortuin–Kasteleyn–Swendsen–Wang representation and Monte Carlo algorithm. *Physical Review D* **38** 2009.
- ENGEL, M., ANDERSON, J. A., GLOTZER, S. C., ISOBE, M., BERNARD, E. P. and KRAUTH, W. (2013). Hard-disk equation of state: First-order liquid-hexatic transition in two dimensions with three simulation methods. *Phys. Rev. E* **87** 042134.
- EVANS, D. J. and HOLIAN, B. L. (1985). The Nosé–Hoover thermostat. *The Journal of Chemical Physics* **83** 4069–4074.
- FAULKNER, M. F. (2022). General symmetry breaking at a topological phase transition. *arXiv (preprint arXiv:2209.03699)*.
- FAULKNER, M. F., BRAMWELL, S. T. and HOLDSWORTH, P. C. W. (2015). Topological-sector fluctuations and ergodicity breaking at the Berezinskii–Kosterlitz–Thouless transition. *Phys. Rev. B* **91** 155412.
- FAULKNER, M. F., QIN, L., MAGGS, A. C. and KRAUTH, W. (2018). All-atom computations with irreversible Markov chains. *J. Chem. Phys.* **149** 064113.
- FEARNHEAD, P., BIERKENS, J., POLLOCK, M. and ROBERTS, G. O. (2018). Piecewise deterministic Markov processes for continuous-time Monte Carlo. *Statist. Sci.* **33** 386–412.
- FRANKS, J. and VIHOLA, M. (2020). Importance sampling correction versus standard averages of reversible MCMCs in terms of the asymptotic variance. *Stochastic Processes and their Applications* **130** 6157–6183.
- GEMAN, S. and GEMAN, D. (1984). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on pattern analysis and machine intelligence* **6** 721–741.
- GIROLAMI, M. and CALDERHEAD, B. (2011). Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73** 123–214.
- GLAUBER, R. J. (1963). Time-dependent statistics of the Ising model. *Journal of mathematical physics* **4** 294–307.
- GORE, V. K. and JERRUM, M. R. (1999). The Swendsen–Wang process does not always mix rapidly. *Journal of Statistical Physics* **97** 67–86.
- GRÜNWARD, P. and VAN OMMEN, T. (2017). Inconsistency of Bayesian inference for misspecified linear models, and a proposal for repairing it. *Bayesian Analysis* **12** 1069–1103.
- HADDOUCHE, M., GUEDJ, B., RIVASPLATA, O. and SHAWE-TAYLOR, J. (2021). PAC-Bayes unleashed: generalisation bounds with unbounded losses. *Entropy* **23** 1330.
- HAIRER, E., LUBICH, C. and WANNER, G. (2003). Geometric numerical integration illustrated by the Störmer–Verlet method. *Acta numerica* **12** 399–450.
- HAIRER, E., LUBICH, C. and WANNER, G. (2006). *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations*. Springer–Verlag.
- HALPERIN, B. I. and NELSON, D. R. (1978). Theory of Two-Dimensional Melting. *Phys. Rev. Lett.* **41** 121–124.
- HOFFMAN, M. D. and GELMAN, A. (2014). The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.* **15** 1593–1623.
- HOHENBERG, P. C. (1967). Existence of Long-Range Order in One and Two Dimensions. *Phys. Rev.* **158** 383.
- HÖLLMER, P., MAGGS, A. C. and KRAUTH, W. (2022). Hard-disk dipoles and non-reversible Markov chains. *J. Chem. Phys.* **156** 084108.
- HÖLLMER, P., QIN, L., FAULKNER, M. F., MAGGS, A. C. and KRAUTH, W. (2020). JeLLyFysh-Version1.0 – a Python application for all-atom event-chain Monte Carlo. *Comput. Phys. Commun.* **253** 07168.
- HOROWITZ, A. M. (1991). A generalized guided Monte Carlo algorithm. *Physics Letters B* **268** 247–252.
- HUBER, M. (2003). A bounding chain for Swendsen–Wang. *Random Structures & Algorithms* **22** 43–59.
- ISING, E. (1925). Beitrag zur Theorie des Ferromagnetismus. *Z. Physik* **31** 253.
- IZAGUIRRE, J. A. and HAMPTON, S. S. (2004). Shadow hybrid Monte Carlo: an efficient propagator in phase space of macromolecules. *Journal of Computational Physics* **200** 581–604.
- JARNER, S. F. and HANSEN, E. (2000). Geometric ergodicity of Metropolis algorithms. *Stoch. Process. Their Appl.* **85** 341.
- JASRA, A., HOLMES, C. C. and STEPHENS, D. A. (2005). Markov chain Monte Carlo methods and the label switching problem in Bayesian mixture modeling.
- JASTER, A. (1999a). Computer simulations of the two-dimensional melting transition using hard disks. *Physical Review E* **59** 2594.
- JASTER, A. (1999b). An improved Metropolis algorithm for hard core systems. *Physica A: Statistical Mechanics and its Applications* **264** 134–141.
- JEWSON, J., SMITH, J. Q. and HOLMES, C. (2018). Principles of Bayesian inference using general divergence criteria. *Entropy* **20** 442.
- JOSÉ, J. V., KADANOFF, L. P., KIRKPATRICK, S. and NELSON, D. R. (1977). Renormalization, vortices, and symmetry-breaking perturbations in the two-dimensional planar model. *Phys. Rev. B* **16** 1217.
- JOSÉ, J. V., KOSTERLITZ, J. M., THOULESS, D. J., ORTIZ, G., COBANERA, E., NUSSINOV, Z., GOLDMAN, A. M., BENFATTO, L., CASTELLANI, C., GIAMARCHI, T., TEITEL, S., FAZIO, R., SCHÖN, G., VIGNOKUR, V. M., BATURINA, T. I., HADZIBABIC, Z., DALIBARD, J., FERTIG, H. A. and MURTHY, G. (2013). *40 Years of Berezinskii–Kosterlitz–Thouless Theory*. World Scientific.

- KAISER, V., BRAMWELL, S. T., HOLDSWORTH, P. C. W. and MOESSNER, R. (2013). Onsager's Wien effect on a lattice. *Nat. Mater.* **12** 1033.
- KAISER, V., BRAMWELL, S. T., HOLDSWORTH, P. C. W. and MOESSNER, R. (2015). ac Wien Effect in Spin Ice, Manifest in Nonlinear, Nonequilibrium Susceptibility. *Phys. Rev. Lett.* **115** 037201.
- KAMPMANN, T. A., BOLTZ, H. H. and KIERFELD, J. (2015a). Monte Carlo simulation of dense polymer melts using event chain algorithms. *J. Chem. Phys.* **143** 044105.
- KAMPMANN, T. A., BOLTZ, H. H. and KIERFELD, J. (2015b). Parallelized event chain algorithm for dense hard sphere and polymer systems. *J. Comput. Phys.* **281** 864.
- KAPFER, S. C. and KRAUTH, W. (2015). Two-Dimensional Melting: From Liquid-Hexatic Coexistence to Continuous Transitions. *Phys. Rev. Lett.* **114** 035702.
- KAPFER, S. C. and KRAUTH, W. (2016). Cell-veto Monte Carlo algorithm for long-range systems. *Physical Review E* **94** 031302.
- KLEMENTA, M. and ENGEL, M. (2019). Efficient equilibration of hard spheres with Newtonian event chains. *J. Chem. Phys.* **150** 174108.
- KNOBLAUCH, J., JEWSON, J. and DAMOULAS, T. (2022). An Optimization-centric View on Bayes' Rule: Reviewing and Generalizing Variational Inference. *Journal of Machine Learning Research* **23** 1–109.
- KOSTERLITZ, J. M. (1974). The critical properties of the two-dimensional XY model. *J. Phys. C: Solid State Phys.* **7** 1046.
- KOSTERLITZ, J. M. and THOULESS, D. J. (1973). Ordering, metastability and phase transitions in two-dimensional systems. *J. Phys. C: Solid State Phys.* **6** 1181.
- KRAMERS, H. A. and WANNIER, G. H. (1941a). Statistics of the Two-Dimensional Ferromagnet. Part I. *Phys. Rev.* **60** 252.
- KRAMERS, H. A. and WANNIER, G. H. (1941b). Statistics of the Two-Dimensional Ferromagnet. Part II. *Phys. Rev.* **60** 263.
- KRAUTH, W. (2006). *Statistical Mechanics: Algorithms and Computations*. Oxford University Press.
- ŁATUSZYŃSKI, K. and ROBERTS, G. O. (2013). CLTs and asymptotic variance of time-sampled Markov chains. *Methodology and Computing in Applied Probability* **15** 237–247.
- LECUN, Y., CHOPRA, S., HADSELL, R., RANZATO, M. and HUANG, F. (2006). A tutorial on energy-based learning. *Predicting structured data* **1**.
- LEI, Z. and KRAUTH, W. (2018). Irreversible Markov chains in spin models: Topological excitations. *EPL (Europhys. Lett.)* **121** 10008.
- LEIMKUHLER, B. and MATTHEWS, C. (2016). *Molecular Dynamics*. Springer.
- LEIMKUHLER, B., NOORIZADEH, E. and THEIL, F. (2009). A gentle stochastic thermostat for molecular dynamics. *Journal of Statistical Physics* **135** 261–277.
- LEIMKUHLER, B. and REICH, S. (2004). *Simulating Hamiltonian Dynamics* **14**. Cambridge university press.
- LELIEVRE, T., ROUSSET, M. and STOLTZ, G. (2010). *Free Energy Computations: A Mathematical Perspective*. World Scientific.
- LIVINGSTONE, S., FAULKNER, M. F. and ROBERTS, G. O. (2019). Kinetic energy choice in Hamiltonian/hybrid Monte Carlo. *Biometrika* **106** 303–319.
- LU, X., PERRONE, V., HASENCLEVER, L., TEH, Y. W. and VOLLMER, S. (2017). Relativistic Monte Carlo. In *Artificial Intelligence and Statistics* 1236–1245. PMLR.
- MARTYNA, G. J., KLEIN, M. L. and TUCKERMAN, M. (1992). Nosé–Hoover chains: The canonical ensemble via continuous dynamics. *The Journal of chemical physics* **97** 2635–2643.
- MATTINGLY, J. C., STUART, A. M. and HIGHAM, D. J. (2002). Ergodicity for SDEs and approximations: locally Lipschitz vector fields and degenerate noise. *Stochastic processes and their applications* **101** 185–232.
- MERMIN, N. D. and WAGNER, H. (1966). Absence of Ferromagnetism or Antiferromagnetism in One- or Two-Dimensional Isotropic Heisenberg Models. *Phys. Rev. Lett.* **17** 1133.
- METROPOLIS, N., ROSENBLUTH, A. W., ROSENBLUTH, M. N., TELLER, A. H. and TELLER, E. (1953). Equation of State Calculations by Fast Computing Machines. *J. Chem. Phys.* **21** 1087.
- MICHEL, M., DURMUS, A. and SÉNÉCAL, S. (2020). Forward event-chain Monte Carlo: Fast sampling by randomness control in irreversible Markov chains. *Journal of Computational and Graphical Statistics* **29** 689–702.
- MICHEL, M., KAPFER, S. C. and KRAUTH, W. (2014). Generalized event-chain Monte Carlo: Constructing rejection-free global-balance algorithms from infinitesimal steps. *The Journal of chemical physics* **140** 054116.
- MICHEL, M., MAYER, J. and KRAUTH, W. (2015). Event-chain Monte Carlo for classical continuous spin models. *EPL (Europhys. Lett.)* **112** 20003.
- MIRA, A. (2001a). On Metropolis-Hastings algorithms with delayed rejection. *Metron* **59** 231–241.
- MIRA, A. (2001b). Ordering and improving the performance of Monte Carlo Markov chains. *Statist. Sci.* 340–350.
- MONEMVASSITIS, A., GUILLIN, A. and MICHEL, M. (2023). PDMP characterisation of event-chain Monte Carlo algorithms for particle systems. *J. Stat. Phys.* **190** 66.
- NEAL, R. M. (1993). Bayesian learning via stochastic dynamics. In *Advances in neural information processing systems* 475–482.
- NEAL, R. M. (2001). Annealed importance sampling. *Statistics and computing* **11** 125–139.
- NEAL, R. M. (2003). Slice sampling. *Annals of statistics* 705–741.

- NEAL, R. M. (2011). MCMC using Hamiltonian dynamics. *Handbook of markov chain monte carlo* **2** 2.
- NEAL, R. M. (2012). *Bayesian learning for neural networks* **118**. Springer Science & Business Media.
- NISHIMURA, A., DUNSON, D. B. and LU, J. (2020). Discontinuous Hamiltonian Monte Carlo for discrete parameters and discontinuous likelihoods. *Biometrika* **107** 365–380.
- NOTT, D. J. and GREEN, P. J. (2004). Bayesian variable selection and the Swendsen–Wang algorithm. *Journal of computational and Graphical Statistics* **13** 141–157.
- ONSAGER, L. (1944). Crystal Statistics. I. A Two-Dimensional Model with an Order-Disorder Transition. *Phys. Rev.* **65** 117.
- ONSAGER, L. (1949). Discussion (comment on the spontaneous magnetisation of the two-dimensional Ising model). *Nuovo Cimento* **6** 251.
- PALMER, R. G. (1982). Broken ergodicity. *Adv. Phys.* **31** 669.
- PAPASPILIOPOULOS, O., ROBERTS, G. O. and SKÖLD, M. (2007). A General Framework for the Parametrization of Hierarchical Models. *Statist. Sci.* 59.
- PARK, J. and ATCHADÉ, Y. (2020). Markov chain Monte Carlo algorithms with sequential proposals. *Statistics and Computing* **30** 1325–1345.
- PAVLIOTIS, G. A. (2014). *Stochastic processes and applications: diffusion processes, the Fokker-Planck and Langevin equations. Texts in Applied Mathematics* **60**. Springer.
- PENG, Y., WANG, Z., ALSAYED, A. M., YODH, A. G. and HAN, Y. (2010). Melting of Colloidal Crystal Films. *Phys. Rev. Lett.* **104** 205703.
- PETERS, E. A. J. F. and DE WITH, G. (2012). Rejection-free Monte Carlo sampling for general potentials. *Phys. Rev. E* **85** 026703.
- POTTS, R. B. (1952). Some generalized order-disorder transformations. *Math. Proc. Camb. Philos. Soc.* **48** 106.
- POWER, S. and GOLDMAN, J. V. (2019). Accelerated sampling on discrete spaces with non-reversible Markov processes. *arXiv preprint arXiv:1912.04681*.
- RADIVOJEVIĆ, T. and AKHMATSKAYA, E. (2020). Modified Hamiltonian Monte Carlo for Bayesian inference. *Statistics and Computing* **30** 377–404.
- RADIVOJEVIĆ, T., FERNÁNDEZ-PENDÁS, M., SANZ-SERNA, J. M. and AKHMATSKAYA, E. (2018). Multi-stage splitting integrators for sampling with modified Hamiltonian Monte Carlo methods. *Journal of Computational Physics* **373** 900–916.
- REDON, S., STOLTZ, G. and TRSTANOVA, Z. (2016). Error analysis of modified Langevin dynamics. *Journal of Statistical Physics* **164** 735–771.
- REN, B., BACALLADO, S., FAVARO, S., HOLMES, S. and TRIPPA, L. (2017). Bayesian nonparametric ordination for the analysis of microbial communities. *Journal of the American Statistical Association* **112** 1430–1442.
- ROBERTS, G. O. and ROSENTHAL, J. S. (2001). Optimal scaling for various Metropolis-Hastings algorithms. *Statist. Sci.* **16** 351.
- ROBERTS, G. O. and TWEEDIE, R. L. (1996). Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli* **2** 341–363.
- ROSCILDE, T., FAULKNER, M. F., BRAMWELL, S. T. and HOLDSWORTH, P. C. W. (2016). From quantum to thermal topological-sector fluctuations of strongly interacting bosons in a ring lattice. *New J. Phys.* **18** 075003.
- ROSENBLUTH, M. N. and ROSENBLUTH, A. W. (1954). Further Results on Monte Carlo Equations of State. *J. Chem. Phys.* **22** 881.
- ROSSKY, P. J., DOLL, J. and FRIEDMAN, H. (1978). Brownian dynamics as smart Monte Carlo simulation. *The Journal of Chemical Physics* **69** 4628–4633.
- SALZBERG, A. M. and PRAGER, S. (1963). Equation of State for a Two-Dimensional Electrolyte. *J. Chem. Phys.* **38** 2587.
- SHERLOCK, C., URBAS, S. and LUDKIN, M. (2021). Apogee to Apogee Path Sampler. *arXiv preprint arXiv:2112.08187*.
- SIMÁNYI, N. (2003). Proof of the Boltzmann–Sinai ergodic hypothesis for typical hard disk systems. *Invent. Math.* **154** 123.
- SINAI, Y. G. (1970). Dynamical systems with elastic reflections. *Russian Mathematical Surveys* **25** 137.
- SINGH, K. K., MARIN, D. F. and REDON, S. (2017). Parallel adaptively restrained molecular dynamics. In *2017 International Conference on High Performance Computing & Simulation (HPCS)* 308–314. IEEE.
- SMITH, M. and KOHN, R. (1996). Nonparametric regression using Bayesian variable selection. *Journal of Econometrics* **75** 317–343.
- STOLTZ, G. (2021). Computational statistical physics and hypocoercivity. *arXiv preprint arXiv:2112.08221*.
- STOLTZ, G. and TRSTANOVA, Z. (2018). Langevin dynamics with general kinetic energies. *Multiscale Modeling & Simulation* **16** 777–806.
- STORATH, M., WEINMANN, A., FRIKEL, J. and UNSER, M. (2015). Joint image reconstruction and segmentation using the Potts model. *Inverse Problems* **31** 025003.
- STÖRMER, C. (1907). Sur les trajectoires des corpuscules électrisés dans l’espace. Applications à l’aurore boréale et aux perturbations magnétiques. *Le Radium* **4** 2–5.
- SWENDSEN, R. H. and WANG, J.-S. (1987). Nonuniversal critical dynamics in Monte Carlo simulations. *Physical Review Letters* **58** 86.

- SYRING, N. and MARTIN, R. (2020). Gibbs posterior concentration rates under sub-exponential type losses. *arXiv preprint arXiv:2012.04505*.
- TABACHNIKOV, S. (2005). *Geometry and billiards* **30**. American Mathematical Soc.
- TALAY, D. (2002). Stochastic Hamiltonian systems: exponential convergence to the invariant measure, and discretization by the implicit Euler scheme. *Markov Process. Related Fields* **8** 163–198.
- TAMAYO, P., BROWER, R. C. and KLEIN, W. (1990). Single-cluster Monte Carlo dynamics for the Ising model. *J. Stat. Phys.* **58** 1083.
- THORNEYWORK, A. L., ABBOTT, J. L., AARTS, D. G. A. L. and DULLENS, R. P. A. (2017). Two-Dimensional Melting of Colloidal Hard Spheres. *Phys. Rev. Lett.* **118** 158001.
- TORGERSON, W. S. (1952). Multidimensional scaling: I. Theory and method. *Psychometrika* **17** 401–419.
- TURITSYN, K. S., CHERTKOV, M. and VUCELJA, M. (2011). Irreversible Monte Carlo algorithms for efficient sampling. *Physica D: Nonlinear Phenomena* **240** 410–414.
- VALLAT, A. and BECK, H. (1994). Coulomb-gas representation of the two-dimensional XY model on a torus. *Phys. Rev. B* **50** 4015.
- VERLET, L. (1967). Computer “Experiments” on Classical Fluids. I. Thermodynamical Properties of Lennard–Jones Molecules. *Physical Review* **159** 98.
- VIHOLA, M., HELSKE, J. and FRANKS, J. (2020). Importance sampling type estimators based on approximate marginal Markov chain Monte Carlo. *Scandinavian Journal of Statistics* **47** 1339–1376.
- VUCELJA, M. (2016). Lifting—A nonreversible Markov chain Monte Carlo algorithm. *Am. J. Phys.* **84** 958.
- WANG, F. and LANDAU, D. P. (2001). Efficient, Multiple-Range Random Walk Algorithm to Calculate the Density of States. *Phys. Rev. Lett.* **86** 2050.
- WEI, D., SONG, Y. and WANG, F. (2011). A simple molecular mechanics potential for μm scale graphene simulations from the adaptive force matching method. *The Journal of Chemical Physics* **134** 184704.
- WOLFF, U. (1989a). Collective Monte Carlo updating for spin systems. *Physical Review Letters* **62** 361.
- WOLFF, U. (1989b). Comparison between cluster Monte Carlo algorithms in the Ising model. *Physics Letters B* **228** 379–382.
- WOOD, W. W. and PARKER, F. R. (1957). Monte Carlo Equation of State of Molecules Interacting with the Lennard–Jones Potential. I. A Supercritical Isotherm at about Twice the Critical Temperature. *J. Chem. Phys.* **27** 720.
- WORRALL, J., BROWNING, R., WU, P. and Mengersen, K. (2022). Fifty years later: new directions in Hawkes processes. *SORT-Statistics and Operations Research Transactions* 3–38.
- WU, L. (2001). Large and moderate deviations and exponential convergence for stochastic damping Hamiltonian systems. *Stochastic processes and their applications* **91** 205–238.
- WU, P.-S. and MARTIN, R. (2022). A Comparison of Learning Rate Selection Methods in Generalized Bayesian Inference. *Bayesian Analysis* **1** 1–28.
- WU, Y., TEPPER, H. L. and VOTH, G. A. (2006). Flexible simple point-charge water model with improved liquid-state properties. *J. Chem. Phys.* **124** 024503.
- YANG, C. N. (1952). The Spontaneous Magnetization of a Two-Dimensional Ising Model. *Phys. Rev.* **85** 808.
- YOUNG, A. P. (1979). Melting and the vector Coulomb gas in two dimensions. *Phys. Rev. B* **19** 1855.
- ZANELLA, G. (2020). Informed proposals for local MCMC in discrete spaces. *Journal of the American Statistical Association* **115** 852–865.
- ZHANG, Y., WANG, X., CHEN, C., HENAO, R., FAN, K. and CARIN, L. (2016). Towards unifying Hamiltonian Monte Carlo and slice sampling. *Advances in Neural Information Processing Systems* **29**.

SUPPLEMENTARY MATERIAL

APPENDIX A: FREE ENERGY DERIVATION FOR ONE-DIMENSIONAL ISING MODEL

For a two-particle system, the partition function of the one-dimensional Ising model is

$$\begin{aligned} Z_{\text{Ising},d=1}(\beta, J, h, N = 2) &= \sum_{x_1=\pm 1} \sum_{x_2=\pm 1} e^{\beta J x_1 x_2 + \frac{\beta h}{2}(x_1+x_2)} e^{\beta J x_2 x_1 + \frac{\beta h}{2}(x_2+x_1)} \\ &= e^{2\beta J} \left(e^{2\beta h} + e^{-2\beta h} \right) + 2e^{-2\beta J}. \end{aligned}$$

Defining the matrix

$$P(\beta, J, h) := \begin{bmatrix} e^{\beta(J+h)} & e^{-\beta J} \\ e^{-\beta J} & e^{\beta(J-h)} \end{bmatrix},$$

this can be rewritten as

$$Z_{\text{Ising},d=1}(\beta, J, h, N = 2) = \text{tr} [P^2(\beta, J, h)].$$

For $N > 1$ particles, this then generalises to

$$Z_{\text{Ising},d=1}(\beta, J, h, N) = \text{tr} [P^N(\beta, J, h)],$$

which follows from the factorisation of the partition function into the product of N two-particle terms:

$$Z_{\text{Ising},d=1}(\beta, J, h, N) = \sum_{x_1=\pm 1} \dots \sum_{x_N=\pm 1} e^{\beta J x_1 x_2 + \frac{\beta h}{2}(x_1+x_2)} \dots e^{\beta J x_N x_1 + \frac{\beta h}{2}(x_N+x_1)}.$$

Since $P(\beta, J, h)$ is a symmetric, real-valued matrix, $\text{tr} [P^N(\beta, J, h)] = \lambda_+^N + \lambda_-^N$, where

$$\lambda_{\pm}(\beta, J, h) = e^{\beta J} \left[\cosh(\beta h) \pm \sqrt{\sinh^2(\beta h) + e^{-4\beta J}} \right]$$

are the two eigenvalues of $P(\beta, J, h)$. It then follows that the free energy is

$$(37) \quad F_{\text{Ising},d=1}(\beta, J, h, N) = -\beta^{-1} \log [\lambda_+^N(\beta, J, h) + \lambda_-^N(\beta, J, h)].$$

APPENDIX B: MISCELLANEOUS PROOFS

PROOF OF PROPOSITION 2. Consider the augmented state space (x, b) and the joint distribution

$$(38) \quad \mu(x, b) \propto e^{-\beta U_{\text{Ising}}(x; J, 0, N)} \left(\prod_{i=1}^N \prod_{j \in S_i} q_{ij}(x)^{b_{ij}/2} (1 - q_{ij}(x))^{(1-b_{ij})/2} \right).$$

We first consider π -invariance, followed by irreducibility and aperiodicity. On this augmented state space the Swendsen-Wang transition can be viewed as the combination of two updates applied sequentially. In the first we simply re-sample $b|x$ from its conditional distribution, which is clearly a μ -preserving transition. In the second we update $x|b$ by flipping the signs of spins within each cluster with probability 1/2. The transition probability associated with this second step can be written

$$P_{\text{SW}}((x, b), (x', b')) = 2^{-C(b)} \mathbb{I}(b = b') \left(\prod_{i=1}^N \prod_{j \in S_i} b_{ij} \mathbb{I}(x'_i = x'_j) \mathbb{I}(x_i = x_j) \right)^{\frac{1}{2}}.$$

The first term on the right-hand side is a normalising constant, in which $C(b)$ denotes the number of clusters in the partition induced by b . The second term stipulates that b does not change. The third ensures that if a bond exists between particles i and j then they must take the same value. The final indicator function $\mathbb{I}(x_i = x_j)$ is not strictly necessary as provided b is drawn from its conditional distribution given x then $b_{ij} \mathbb{I}(x_i = x_j) = b_{ij}$, as a bond can only exist between particles i and j if they have the same spin. It does, however, make it clear that $P_{\text{SW}}((x, b), (x', b')) = P_{\text{SW}}((x', b'), (x, b))$, meaning that μ -reversibility follows from showing that

$$(39) \quad e^{-\beta(U_{\text{Ising}}(x'; J, 0, N) - U_{\text{Ising}}(x; J, 0, N))} = \prod_{i=1}^N \prod_{j \in S_i} \frac{q_{ij}(x)^{b_{ij}/2} (1 - q_{ij}(x))^{(1-b_{ij})/2}}{q_{ij}(x')^{b_{ij}/2} (1 - q_{ij}(x'))^{(1-b_{ij})/2}}.$$

This can be seen by first considering the left-hand side of (39) and noting that $x_i x_j = 2\mathbb{I}(x_i = x_j) - 1$ when x_i and x_j can only take the values $\{-1, +1\}$, meaning

$$U_{\text{Ising}}(x'; J, 0, N) - U_{\text{Ising}}(x; J, 0, N) = -J \sum_{i=1}^N \sum_{j \in S_i} [\mathbb{I}(x'_i = x'_j) - \mathbb{I}(x_i = x_j)].$$

This can be further modified by noting that under the Swendsen-Wang update the function $\mathbb{I}(x_i = x_j) - \mathbb{I}(x'_i = x'_j)$ can only be non-zero for neighbouring particles i and j if $b_{ij} = 0$, meaning that if x' is generated from such an update then

$$U_{\text{Ising}}(x'; J, 0, N) - U_{\text{Ising}}(x; J, 0, N) = -J \sum_{i=1}^N \sum_{j \in S_i} [\mathbb{I}(x'_i = x'_j) - \mathbb{I}(x_i = x_j)] (1 - b_{ij}).$$

Turning to the right-hand side of (39) notice first that since $q_{ij}(x) = 1 - e^{-2\beta J\mathbb{I}(x_i=x_j)}$ then $q_{ij}(x) = q_{ij}(x')$ when $b_{ij} = 1$, meaning that upon substituting in the definition of $q_{ij}(x)$ the fraction can be re-written

$$\prod_{i=1}^N \prod_{j \in S_i} \frac{e^{-\beta J\mathbb{I}(x_i=x_j)(1-b_{ij})}}{e^{-\beta J\mathbb{I}(x'_i=x'_j)(1-b_{ij})}},$$

which from the calculations above is clearly equal to $e^{-\beta(U_{\text{Ising}}(x'; J, 0, N) - U_{\text{Ising}}(x; J, 0, N))}$ as required.

Establishing irreducibility and aperiodicity is straightforward. Aperiodicity can be seen by noting that the algorithm has a positive probability of not moving. Irreducibility can be seen by simply noting that for any fixed N there is a positive probability that $\prod_{i=1}^N \prod_{j \in S_i} b_{ij} = 0$ regardless of the value of the current state x . In this instance each particle belongs to its own cluster and therefore x' can take any value in $\{-1, +1\}^N$ with the same (non-zero) probability. Any x' can be arrived at from any x in a single iteration of the algorithm. Similarly any b' can be arrived at from any b by first transitioning through an appropriate x' , since for every configuration of bonds b' there is a configuration of spins x' for which $b'|x'$ has non-zero probability. The chain therefore has limiting distribution $\mu(x, b)$, and the marginal process on x has limiting distribution $\pi(x) \propto e^{-\beta U_{\text{Ising}}(x; J, 0, n)}$, from which ergodicity follows. \square

PROOF OF PROPOSITION 3. The result has already been established for the first two parts of the right-hand side of (34) in Appendix E of [Andrieu, Lee and Livingstone \(2020\)](#). To extend to the xy -version, it suffices therefore to show (μ, Q) -self-adjointness of the last part when R implements the xy transformation.

For the case $d = 2$ note that $R(z, dz')$ simply swaps $(u_1, u_2) \rightarrow (u_2, u_1)$. We write $v \rightarrow \mathcal{S}(v)$ for this transformation and note that $\mathcal{S}(-v) = -\mathcal{S}(v)$ and $\mathcal{S} \circ \mathcal{S}(v) = v$, meaning \mathcal{S} is an involution. Since $\langle f, g \rangle_\mu = \langle Qf, Qg \rangle_\mu$ and $Q^2 = I$ then setting

$$(40) \quad L'f(x, v, i) := \lambda_{\text{ref}}[f(x, \mathcal{S}(v), i) - f(x, v, i)]$$

one can equivalently show that $\langle QL'f, g \rangle_\mu = \langle f, QL'g \rangle_\mu$, meaning that the Q -symmetrization QL' is μ -self-adjoint (see [Andrieu and Livingstone \(2021\)](#) for more detail). Direct calculation gives

$$(41) \quad \begin{aligned} \lambda_{\text{ref}}^{-1} \langle QL'f, g \rangle_\mu &= \int [f(x, -\mathcal{S}(v), i) - f(x, -v, i)]g(x, v, i)d\mu \\ &= \int f(x, -\mathcal{S}(v), i)g(x, v, i)d\mu - \int f(x, -v, i)g(x, v, i)d\mu. \end{aligned}$$

We can apply the change of variables $v \rightarrow -\mathcal{S}(v)$ to the first integral and $v \rightarrow -v$ to the second. Since μ is invariant to both, the expression becomes

$$(42) \quad \int [g(x, -\mathcal{S}(v), i) - g(x, -v, i)]f(x, v, i)d\mu = \lambda_{\text{ref}}^{-1} \langle f, QL'g \rangle_\mu,$$

which completes the proof. \square

REMARK 5. The $d > 2$ setting can also be considered. In this case \mathcal{S} is no longer an involution, but it is invertible, with $\mathcal{S}^{-1}(v)$ simply shifting each element of u one space to the left modulo d . Following Remark 5 of [Andrieu, Lee and Livingstone \(2020\)](#) we can therefore introduce the auxiliary $w \in \{-1, +1\}$ and define the involution $\mathcal{S}(v, w) := (\mathcal{S}^w(v), -w)$ on an extended space. We can then incorporate w into the augmented state (x, v, i, w) and augment the measure μ to include a symmetric component for w , then perform analogous calculations to those above to establish the result. We omit the details for brevity.

Supplementary materials

Additional derivations and results.