

Flow-based Nonperturbative Simulation of First-order Phase Transitions

Yang Bai and Ting-Kuo Chen

Department of Physics, University of Wisconsin-Madison, Madison, WI 53706, USA

Abstract

We present a flow-based method for simulating and calculating nucleation rates of first-order phase transitions in scalar field theory on a lattice. Motivated by recent advancements in machine learning tools, particularly normalizing flows for lattice field theory, we propose the “partitioning flow-based Markov chain Monte Carlo (PFMCMC) sampling” method to address two challenges encountered in normalizing flow applications for lattice field theory: the “mode-collapse” and “rare-event sampling” problems. Using a (2+1)-dimensional real scalar model as an example, we demonstrate the effectiveness of our PFMCMC method in modeling highly hierarchical order parameter probability distributions and simulating critical bubble configurations. These simulations are then used to facilitate the calculation of nucleation rates. We anticipate the application of this method to (3+1)-dimensional theories for studying realistic cosmological phase transitions.

Contents

1	Introduction	2
2	Nucleation theory for first-order phase transitions	4
2.1	A review of classical nucleation theory	4
2.2	Nonperturbative calculations based on lattice field theory	6
3	A flow-based simulation algorithm	8
3.1	Flow-based sampling in lattice field theory	8
3.2	Mode-collapse and rare-event sampling	10
3.3	Partitioning flow-based MCMC	12
3.3.1	A toy 2D ϕ^4 theory	13
3.3.2	A toy 2D ϕ^6 theory	16
4	(2+1)D $\mathbb{Z}_2$-symmetric real scalar theory	17
4.1	Finite-temperature (2+1)D $\mathbb{Z}_2$ -symmetric theory	18
4.2	Benchmark test: bubble simulation	18
4.3	Dynamical factor and nucleation rate calculations	22
5	Discussion and conclusions	24
A	The nucleation theory for relativistic quantum fields	25
B	Finite-temperature (2+1)D $\mathbb{Z}_2$-symmetric scalar theory	26

1 Introduction

Cosmological first-order phase transitions (FOPTs) in the early universe have long been an intriguing field of study. With a wide range of extensions beyond the Standard Model (SM) of particle physics suggesting the possibility of cosmological FOPTs, there is a constant demand for rigorous predictions and measurements of the associated signatures. For instance, electroweak FOPTs are essential for explaining the matter-antimatter asymmetry through the mechanism of electroweak baryogenesis [1, 2]. In addition, QCD could also undergo a FOPT with the assistance of new physics beyond the SM. For instance, late-collapsed domain walls with corresponding discrete symmetries anomalous under QCD can create a different thermal environment to induce QCD FOPTs [3]. Such transitions could potentially be probed by several pulsar timing array experiments aimed at detecting stochastic gravitational wave backgrounds (SGWBs) [4–7]. Beyond the symmetries in the SM, FOPTs may occur in the dark matter sector, Grand Unified Models, quark and neutrino flavor models, with the produced gravitational waves serving as important messengers to probe early-universe physics (see Ref. [8] for a recent review).

The traditional method to study cosmological FOPT is based on the nucleation theory pioneered by Langer [9] for non-relativistic classical systems, which was later generalized by Coleman [10] and Linde [11, 12] for zero- and finite-temperature relativistic quantum fields. The general picture is to identify the probability flux flowing from the metastable vacuum phase, through the saddle-point or critical-bubble configuration, to the stable vacuum phase as the nucleation rate. This picture allows one to factor the nucleation problem into a statistical (equilibrium) part and a dynamical (nonequilibrium) part and solve it using the method of saddle-point approximation. However, this semi-classical approach falls short of the following matters: 1) It cannot account for the fluctuations (and thus non-sphericity) of the bubbles in a realistic way. 2) The dynamical part requires the real-time evolution input of the system according to, *e.g.*, the Langevin equation. 3) In the case of strong-coupling models, such as QCD-type models, numerical nonperturbative tools are necessary for performing the calculations. Nevertheless, as the nucleation picture given by Langer is still generally true for a system with hierarchical FOPT and particle scattering timescales, one can simply generalize the numerical method from the saddle-point approximation to the nonperturbative lattice field theory. For such studies of electroweak FOPTs, one can refer to Refs. [13, 14]. More recently, Ref. [15] has shown up to compare the nucleation rates calculated from the lattice formalism and the perturbative formalism for a real scalar theory in (3+1) dimensions.

In the literature, lattice-based Markov chain Monte Carlo (MCMC) simulations of FOPTs continually encounter the following challenges:

$\mathcal{A}$: critical slowing down [16–20];

$\mathcal{B}$: mode collapse for multimodal distributions [21, 22];

$\mathcal{C}$: rare-event sampling [23].

To tackle challenge $\mathcal{A}$, a recent method that utilizes normalizing flows (NFs), a class of generative machine learning models, has been demonstrated to be highly effective [24]. However, direct application of this flow-based method to FOPT simulations turns out to be ineffective

because of challenges $\mathcal{B}$ and $\mathcal{C}$. On the other hand, the method of multicanonical MCMC (MMCMC) [25] was proposed to address challenges $\mathcal{B}$ and $\mathcal{C}$. Even though MMCMC has been widely adopted in the nonperturbative FOPT studies, it was noted in a recent paper, Ref. [14], that the development of more efficient algorithms is warranted.

In this study, we propose a flow-based algorithm capable of simultaneously addressing all three challenges, which we term the “partitioning flow-based (PF) MCMC” algorithm. The concept behind this algorithm is to train the NFs based on a reweighted target distribution (similar to the MMCMC algorithm), but instead of attempting to construct the complete distribution in one step, we “slice” or “partition” the phase space into multiple partitions, address them individually, and then “glue” them together. This method effectively tackles the three challenges as follows:

- $\mathcal{A}$: Flow-based MCMC algorithms inherently resolve the critical slowing down issue by decoupling the autocorrelation time of a Markov chain from the problem (or lattice) size [24] through the adoption of the independence Metropolis sampling algorithm [26].
- $\mathcal{B}$: While conventional flow-based modeling struggles to capture the multimodal structure of the target distribution, the PFMCMC method treats it as several unimodal or monotonic modeling problems (except for regions around local minima, which are assumed to have a small enough hierarchy to be properly managed), thus enhancing stability and effectiveness.
- $\mathcal{C}$: By partitioning the model into several sub-models, one can efficiently sample even from the rarest phase space regions using the corresponding sub-model.

To demonstrate the capability of the PFMCMC method, we first apply it to several toy 2D models. We show that it yields results consistent with those from the literature and outperforms certain revised flow-based algorithms. Additionally, we employ it in a realistic finite-temperature (2+1)D $\mathbb{Z}_2$ -symmetric real scalar theory. Using the resulting probability distribution and bubble configurations, we illustrate a complete nucleation rate calculation for a benchmark scenario. Although our ultimate goal is to apply this method to a (3+1)D model for studying realistic cosmological FOPTs, the (2+1)D model we consider is highly motivated. For instance, in condensed matter physics, certain studies on the superfluidity of a 2D Bose gas rely on thermal mean field theory in (2+1) dimensions (see, for example, Refs. [27–30]).

Our paper is organized as follows. In Section 2, we review the classical nucleation theory outlined by Langer and its nonperturbative generalization based on lattice field theory. The semi-classical application of this theory to relativistic quantum fields is briefly reviewed in Appendix A. The flow-based algorithm for lattice field simulations is overviewed in Section 3, after which we discuss the challenges of bubble simulations and how the PFMCMC method can address these issues, demonstrated within a few toy 2D models. We introduce the finite-temperature (2+1)D $\mathbb{Z}_2$ -symmetric real scalar theory, apply the PFMCMC method on one benchmark, and demonstrate a complete nucleation rate calculation based on the simulation results in Section 4, with the details of the scalar field theory given in Appendix B. Finally, we discuss and conclude our study in Section 5.

2 Nucleation theory for first-order phase transitions

In this section, we first briefly review the classical nucleation theory outlined by Langer in Ref. [9] to calculate the nucleation rates for the decays of a metastable or false vacuum. We then describe a convenient way to project the multi-dimensional functional space onto a one-dimensional order-parameter functional direction, following Ref. [13]. After that, we introduce a method based on lattice field theory to calculate the nucleation rates nonperturbatively.

2.1 A review of classical nucleation theory

For early-universe FOPTs, we are interested in the homogeneous nucleation or calculating the probability of a true-vacuum-phase bubble or droplet appearing in a system initially in the metastable vacuum phase. For a non-relativistic system, the theory of nucleation was developed by Langer through introducing a set of variables, η_i , to describe the collective degrees of freedom of the system [9]. The distribution function or the probability density over the configurations and time, $\rho(\{\eta\}, t)$, is related to the free-energy of the system for equilibrium configurations (this is justified as long as the associated timescale for bubble nucleation is much longer than the particle thermalization timescale) by $\rho_{\text{eq}}(\{\eta\}) \propto \exp[-F(\{\eta\})/T]$. The phase transition begins from the system sitting at some initial point in the metastable region, later passing through the region of the saddle point, and then reaches the stable region. As pointed out by Langer, one could identify the probability flux through the saddle point with the nucleation rate without having to solve the complete time-dependent problem.

Based on the saddle-point approximation, the nucleation rate can be expressed as the product of two factors: the dynamical factor F_D and the statistical factor F_S . The statistical factor F_S can be calculated using the saddle-point approximation in Ref. [9], while the dynamical factor F_D is the exponential growth rate of a bubble that just exceeds the critical size, and is proportional to the negative eigenvalue of the free energy second-order differential matrix in η_i , evaluated at the saddle point. The dynamical factor was calculated in Ref. [9] for a non-relativistic system and in Ref. [31] for a relativistic system.

For a relativistic quantum field theory, which is more relevant when describing early-universe FOPTs, Coleman has developed the formula for zero-temperature theories [10] which was later generalized for finite-temperature theories by Linde [11, 12] (see Appendix A for details). For our numerical lattice simulation of a simple scalar field model in $(2 + 1)$ dimensions, the general picture depicted by Langer is still applicable, although the lattice simulation of the phase transition does not need to rely on the saddle-point approximation.

In the left panel of Fig. 1, we show a cartoon picture that illustrates various paths along which the configurations $\{\phi(\vec{x})\}$ flow from the metastable region to the stable region. Between these two regions, field configurations are supposed to pass through the critical-bubble region (for simplicity, only one critical-bubble configuration is assumed to exist). The phase space of the critical-bubble configuration or the saddle-point approximation region is smaller than those of the metastable and stable regions because the critical bubble sits at a local maximum of the free energy spectrum at least in one direction of the configuration space. A field configuration sitting at a point slightly above the critical bubble is expected to flow along the direction that reduces the free energy, or the “flux” direction, in Fig. 1. Perpendicular to the flux direction,




Figure 1: *Left:* A cartoon illustration of the flow of field configurations from the metastable region to the stable region through the transition surface. Here, ϕ_F and ϕ_T represent the field configurations in the metastable (false) phase and the stable (true) phase, respectively. ϕ_C denotes the critical-bubble configuration, with “flux” indicating the flow direction of the field configuration at ϕ_C , and “transition surface” referring to the hyper-dimensional surface perpendicular to the flux direction. *Right:* A similar illustration to the left panel, but with the hyper-dimensional functional space projected onto the one-dimensional measurable $\mathcal{O}$ space. The relative heights of the metastable, stable, and transition-point regions indicate their different statistical probabilities. Depending on the choice of the functional $\mathcal{O}[\phi(\vec{x})]$, the locations of $\mathcal{O}_{F,C,T}$ may be close to but do not exactly match to the configurations $\phi_{F,C,T}$ in the left panel.

one could define the transition surface that is co-dimension one in the configuration space. The nucleation rate could be obtained by calculating both the statistical factor (the rareness of the critical-bubble configuration) and the dynamical factor (the quickness of the probability passing the critical-bubble point) around the saddle-point approximation region.

In the right panel of Fig. 1, we show the projected one-dimensional subspace of $\{\phi(\vec{x})\}$ defined through a “measurable” operator $\mathcal{O}[\phi(\vec{x})]$ that is a functional of the field configurations. The motivation for introducing such a measurable is for the convenience of later numerical simulations of FOPT. The independence of nucleation on a particular choice of $\mathcal{O}$ is demonstrated in Ref. [13] (see also Ref. [15] for the explorations of numerically friendly measurables). For a suitable choice of $\mathcal{O}$, there are still three regions: the metastable region, the transition point, and the stable region. The flow of field configurations in the left panel becomes a simple straight flow from the metastable region (left) to the stable region (right) that passes through the transition point in between. The relative heights of different regions indicate their different statistical probabilities. For a linear projection from $\{\phi(\vec{x})\}$ to $\mathcal{O}$, the locations of $\phi_{F,C,T}$ in the left panel should match to those of $\mathcal{O}_{F,C,T}$ in the right panel. This is not generally true for a non-linear projection. However, for a suitable choice of $\mathcal{O}$, the three regions on the left panel should have large overlaps with the corresponding regions on the right panel. Therefore, one can focus on the transition point region around $\mathcal{O}_C$ to calculate the nucleation rate.

Based on the probability distribution $P(\mathcal{O})$ of the measurable, one can define a “working region”, $[\mathcal{O}_c - \epsilon/2, \mathcal{O}_c + \epsilon/2]$ with a small ϵ , around the saddle-point approximation region (the numerical calculation can extend beyond the saddle-point approximation) to calculate the

nucleation rate. First, the statistical factor as defined in the Langer’s method becomes

$$F_S = \frac{P([\mathcal{O}_c - \epsilon/2, \mathcal{O}_c + \epsilon/2])}{\epsilon P(\mathcal{O} < \mathcal{O}_c - \epsilon/2)} , \quad (1)$$

which is the ratio of the probability density containing the critical bubble over the probability of the metastable phase configurations.

Second, the dynamical factor is proportional to the probability flux, $|d\mathcal{O}/dt|_{\mathcal{O}_c}$, at the transition point or the “critical separatrix”. In a practical numerical simulation, however, one needs to multiply the flux by a “success factor”, S , to account for two facts: 1) Not every field evolution trajectory that passes through the critical point will eventually reach the stable region. 2) Some trajectories may have multiple turnarounds around the critical point before reaching the stable region, but they only overall contribute to one tunneling. This factor depends on either even or odd number of times that the trajectory crosses the critical point, $N_{\text{crossings}}$, and is given by $S = (N_{\text{crossings}} \bmod 2)/N_{\text{crossings}}$. As one does not distinguish between the crossing in one direction and the opposite during the simulation, a factor of $\frac{1}{2}$ should be further multiplied in order to only account for metastable $\rightarrow$ stable trajectories. Consequently, the dynamical factor is [13]

$$F_D = \left\langle \frac{1}{2} S \left| \frac{d\mathcal{O}}{dt} \right|_{\mathcal{O}_c} \right\rangle = \left\langle \frac{1}{2} \frac{(N_{\text{crossings}} \bmod 2)}{N_{\text{crossings}}} \left| \frac{d\mathcal{O}}{dt} \right|_{\mathcal{O}_c} \right\rangle , \quad (2)$$

where $\langle \rangle$ means averaging over numerical samples.

Combining the two factors in Eqs. (1) and (2), one has the homogeneous nucleation rate per volume as

$$\Gamma = \frac{F_D F_S}{V} = \left\langle \frac{1}{2} \frac{(N_{\text{crossings}} \bmod 2)}{N_{\text{crossings}}} \left| \frac{d\mathcal{O}}{dt} \right|_{\mathcal{O}_c} \right\rangle \times \frac{P([\mathcal{O}_c - \epsilon/2, \mathcal{O}_c + \epsilon/2])}{\epsilon P(\mathcal{O} < \mathcal{O}_c - \epsilon/2) V} , \quad (3)$$

where V stands for a large enough volume that contains the critical-bubble configuration. We remark that the exact value of ϵ does not affect the overall result much as long as it is reasonably small. There are counteracting dependencies on ϵ for the statistical and dynamical factors which eventually cancel out. The details of the argument can be found in Ref. [13].

2.2 Nonperturbative calculations based on lattice field theory

Using lattice field theory as a nonperturbative method to study FOPTs has a few advantages: 1) One is not restricted to perturbative theories, but could also study theories with nonperturbative couplings such as QCD-type models. 2) The calculations for nucleation rates or other quantities can go beyond the saddle-point approximation as in Langer’s method. 3) The field configurations sampled numerically could be used to study other physics related to the FOPT, including the properties of non-spherical critical bubbles and gravitational wave productions. In this work, we focus on calculating the bubble nucleation rates of FOPTs, which is one of the most important quantities pertaining to the history of cosmological phase transitions. Our following lattice field theory setup follows closely to the general strategy in Ref. [13], where the detailed numerical implementations for various terms in Eq. (3) have been addressed. Some recent studies following

Ref. [13] can be found in Refs. [14, 15, 32]. Our main goal is to incorporate modern machine-learning tools into the studies of FOPTs.

The calculation for the dynamical factor in Eq. (2) is fairly straightforward at least for a simple model with only a few fields. One could make a large sample of field configurations around the transition point (see the $\mathcal{O}_c$ region in the right panel of Fig. 1) and use the real-time Langevin equation to evolve the configurations following [33–35]

$$\frac{d\phi}{dt} = -\frac{\delta S[\phi]}{\delta\phi} + \nu(t, \vec{x}) . \quad (4)$$

Here, $S[\phi] = T \times S_E[\phi]$ with $S_E[\phi]$ being the temperature-reduced Euclidean action for studies of thermal FOPT [see Eq. (30) for the example of a scalar field theory]; $\nu(t, x)$ is a Gaussian white noise satisfying

$$\langle \nu(t, \vec{x}) \rangle_\nu = 0, \quad \langle \nu(t, \vec{x}) \nu(t', \vec{x}') \rangle_\nu = 2T \delta^{(d-1)}(\vec{x} - \vec{x}') \delta(t - t') , \quad (5)$$

with $\langle \rangle_\nu$ representing the ensemble average over ν samples and $d - 1$ as the number of spacial dimensions, which follows a Gaussian distribution with a width of $\sqrt{2T}$, *i.e.*,

$$d\rho[\nu(t, \vec{x})] = d\nu(t, \vec{x}) e^{-\int d^{d-1}x dt \nu^2(t, \vec{x})/(4T)} . \quad (6)$$

For the probability flux, one starts with a field configuration in the working region $[\mathcal{O}_c - \epsilon/2, \mathcal{O}_c + \epsilon/2]$ and evolves it both backward and forward in time according to Eq. (6). One can define the probability flux $|d\mathcal{O}/dt|_{\mathcal{O}_c} \approx |\Delta\mathcal{O}/\Delta t|_{\mathcal{O}_c}$ for a small time interval Δt , which does not distinguish the metastable $\rightarrow$ stable direction from the opposite direction. The success factor S in Eq. (1) depends on whether the evolution trajectory eventually ends up in the opposite region (metastable vs. stable). If not, $N_{\text{crossings}}$ is an even number and $(N_{\text{crossings}} \bmod 2) = 0$, indicating that this trajectory does not contribute to the dynamical factor. On the contrary, $(N_{\text{crossings}} \bmod 2) = 1$ for an odd $N_{\text{crossings}}$, and this trajectory will contribute to the dynamical factor (see Fig. 14 for two examples of the trajectories from our numerical simulations). Finally, a factor of $1/N_{\text{crossings}}$ is multiplied to account for the fact that there is only one effective crossing, as mentioned previously.

The calculation of the statistical factor in Eq. (1) requires a large sample of field configurations to obtain a reliable probability distribution, $\rho_{\text{eq}}(\{\phi\}) \propto \exp[-F(\{\phi\})/T]$ (see the right panel of Fig. 1 in the one-dimensional measurable space). In the context of cosmological phase transition, the critical-bubble configuration with $\mathcal{O}_c$ could be highly suppressed by $\mathcal{O}(e^{-100})$ as a result of the hierarchy between the phase transition temperature and the Hubble scale. On top of this huge suppression factor, the multimodal structure of the probability distributions (see the right panel of Fig. 1) also makes it difficult for traditional MCMC methods to model them properly due to their “mode-seeking” behavior [36]. Furthermore, traditional MCMC methods usually suffer from the critical slowing down issue that imposes a barrier to scalability. In lattice field theory, this issue is usually manifest in the lattice-size-dependence of the autocorrelation time of a constructed Markov chain [24]. For convenience, we have named the three challenges that traditional MCMC methods are usually faced with when calculating the probability distribution of measurable values as: $\mathcal{A}$: critical slowing down; $\mathcal{B}$: mode collapse for multimodal distributions; $\mathcal{C}$: rare-event sampling.

Before we discuss our algorithm that adopts modern machine learning to overcome the three challenges, we also briefly comment on the existing method of MMCMC (see also Sec. 3.2) introduced in Ref. [25] where a weight function in terms of the measurable is introduced to flatten the target probability distribution for more efficient samplings. After the samplings are finished, the results are then reweighted back to retrieve the correct statistics. This method requires some initial (“well educated”) guess for the right weight function and was already used and commented in Ref. [13] (see Ref. [14] for a recent application and request for the development of more efficient algorithms).

3 A flow-based simulation algorithm

In this section, we develop a flow-based algorithm to address all three challenges mentioned previously and outline the procedure to calculate the nucleation rates of FOPTs. For challenge $\mathcal{A}$, it has been shown in the literature that an independence Metropolis algorithm based on NFs naturally solves the critical slowing down issue [24, 37]. As a family of generative models, NFs can model the intricate probability distributions by transforming a simple base distribution into a complex distribution via a series of invertible and differentiable transformations. Usually, these transformations are implemented with trainable deep neural networks. When applied to an independence Metropolis sampler (see Sec. 3.1), the autocorrelation time of the constructed Markov chain does not scale with the lattice size but only depends on how closely the trained NF model matches to the target distribution [24]. However, the simple application of NFs to lattice field theory does not solve challenge $\mathcal{B}$ (the mode-collapse problem), which has been discussed in Refs. [21, 22], or challenge $\mathcal{C}$ (the rare-event sampling problem), which has been discussed in Ref. [23]. We will discuss these issues in detail and propose a partitioning flow-based method to overcome these challenges and apply it to bubble nucleation rate calculations.

3.1 Flow-based sampling in lattice field theory

In lattice field theory, one uses MCMC approaches to generate field configurations according to the target distribution

$$p(\phi) = e^{-S_E[\phi]}/Z, \quad \text{with} \quad Z = \int \prod_{j=1}^D d\phi_j e^{-S_E[\phi]}, \quad (7)$$

where j indexes the D components of $\phi \in \mathbb{R}^D$ and ϕ is defined to be a vector with D real components matching the spacetime degrees of freedom for a real scalar $\phi(\vec{x})$ evaluated on a finite, discrete spacetime lattice. Within the flow-based framework, one constructs a Markov chain according to $p(\phi)$ via the independence Metropolis sampler [26]. For each step i of the chain, an updated proposal ϕ' is generated by sampling from $\tilde{p}(\phi)$ [a function to be optimized to approximately match the true $p(\phi)$], independent of the previous configuration $\phi^{(i-1)}$. The step is accepted with probability

$$A(\phi^{(i-1)}, \phi') = \min \left[1, \frac{\tilde{p}(\phi^{(i-1)}) p(\phi')}{p(\phi^{(i-1)}) \tilde{p}(\phi')} \right]. \quad (8)$$




Figure 2: A schematic plot of the flow model used in this study. In each iteration, a set of random variables on the lattice z (a vector of $\mathbb{R}^D$) is drawn from the prior distribution $r(z)$. When the random variables pass through each (bi-)layer of the flow model, the white and black lattice sites are updated according to $g_{i,W}$ and $g_{i,B}$ sequentially. In each $g_{i,W(B)}$, the white(black) sites are frozen, and black(white) sites are updated according to the non-linear squared (NLSq) function. The parameters, a, b, c, d, g , of the NLSq function are to be optimized and are parametrized with a convolutional neural network (CNN) using the inputs from the frozen sites. The output is then forwarded to the next (bi-)layer until the last layer is reached. The final outputs are the transformed random variables $\phi = f^{-1}(z) \sim \tilde{p}(\phi)$. By minimizing the loss function $D_{\text{KL}}(\tilde{p}||p)$, the model is trained to sample ϕ from $\tilde{p}(\phi) \approx p(\phi)$. Our particular model structure has roughly 500K trainable hyperparameters.

If the proposal is accepted, $\phi^{(i)} = \phi'$, otherwise $\phi^{(i)} = \phi^{(i-1)}$.

To construct $\tilde{p}(\phi)$, we construct a NF f that maps from $z \sim r(z)$ to $\phi = f^{-1}(z) \sim \tilde{p}(\phi)$, where from using Jacobian one has

$$\tilde{p}(\phi) = r(f(\phi)) \left| \det \frac{\partial f(\phi)}{\partial \phi} \right|. \quad (9)$$

Here, the prior distribution function $r(z)$ with $z \in \mathbb{R}^D$ could be some simple known function (Gaussian distribution for instance). The functional mapping from $z \xrightarrow{f^{-1}} \phi$ or the inverted one $\phi \xrightarrow{f} z$ could be very complicated. One could construct f by the composition of *coupling layers* and choose a triangular matrix for the transformation of each layer such that the determinant of the Jacobian can be efficiently computed [38]. A schematic illustration of the flow model used in our study is shown in Fig. 2.

The model is optimized through minimizing the (approximate) shifted reverse Kullback-

Leibler (KL) divergence [39],

$$\begin{aligned} L(\tilde{p}, p) &\equiv D_{\text{KL}}(\tilde{p}||p) - \log Z \\ &= \int \mathcal{D}\phi \tilde{p}(\phi) [\log \tilde{p}(\phi) + S_{\text{E}}[\phi]] \approx \frac{1}{N} \sum_{\phi \in \text{batch}} [\log \tilde{p}(\phi) + S_{\text{E}}[\phi]] . \end{aligned} \quad (10)$$

Physically, shifting the loss function by the constant Z , which is in general intractable, does not affect the simulation outcome, thus justifying the choice of the loss function. In the last step, the loss is calculated by using a batch of N samples at each epoch from the model $\{\phi \sim \tilde{p}\}$. As proven and experimented in Ref. [24], this flow-based MCMC sampler is immune to the critical slowing down problem, and the integrated autocorrelation time is only determined by the acceptance rate of the Markov chain, which is in turn determined by how well $\tilde{p}(\phi)$ assimilates $p(\phi)$. Moreover, training with the reverse KL divergence does not require a pre-generated dataset, which is another appealing feature of this method.

3.2 Mode-collapse and rare-event sampling

One major challenge for this flow-based method is the “mode-seeking” behavior of the reverse KL self-training [36, 40] that will possibly lead to the “mode-collapse” problem (one mode with a higher probability is identified while the additional lower-probability modes are ignored or underestimated). The dynamics of reverse KL training cause initial overconcentration on peaks, which has been observed to manifest as a tendency for the sampler to initially collapse onto a subset of modes. If the model never samples from a given mode, it is not penalized for badly mismodeling that excluded mode, since we cannot detect whether a mode has been missed as the normalizing factor Z is unknown based on the loss function defined in Eq. (10). Ref. [21] has proposed several methods to mitigate the problem, and Ref. [22] has further compared the scalabilities of reverse and forward KL training schemes¹ while also proposed several mitigation strategies. We have tested two methods from Ref. [21] that seem to provide better results:

- **Adiabatic retraining:** The idea of adiabatic retraining is derived from transfer learning, wherein a model trained targeting some easier-to-approximate set of parameters is used to initialize the training for harder ones. Adiabatic retraining generalizes this idea to vary the parameters epoch-by-epoch over some trajectory through parameter space over the course of training. As long as the variation is sufficiently slow, the model being trained will remain a good approximation of the target distribution throughout training. The way in Ref. [21] applies this method is to start the training from modeling a unimodal distribution and then adiabatically transform it into a multimodal one, which requires certain knowledge of the target distribution before training.
- **Flow-distance regularization:** The flow-distance regularizer is a term added to the loss function to penalize flow functions f that transform prior samples z to significantly different

¹These refer to the training using the reverse and forward (or normal) KL divergence, respectively. For the forward KL, the samples are prepared before the training using some other MCMC algorithms, while for the reverse KL (which is our case), the samples are drawn from the model itself.

output samples ϕ , as given by

$$L(\tilde{p}, p)_{\text{fd}} = L(\tilde{p}, p) + k [1 - g(t)] \frac{1}{D} \sum_{i=1}^D \|\phi_i - z_i\|_2, \quad (11)$$

where the L^2 or $\|\cdot\|_2$ norm implies a sum over all degrees of freedom on the lattice, k is the coupling strength of this constraint, and $g(t) \in [0, 1]$ is a scheduler function of training time t that gradually turns off the regularization. To some extent, this method is applicable to any theory without a priori knowledge of the mode structure and relies only on the broadness of the prior.

In practice, one can combine these two methods as they are both t -dependent. Ref. [21] has shown that this method can avoid the problem of undersampling a minimum for a (slightly-broken) $\mathbb{Z}_2$ -symmetric ϕ^4 theory compared to the other flow-based methods that intrinsically relies on the $\mathbb{Z}_2$ symmetry.

In the context of cosmological phase transition, however, the relative weights among the stable and metastable vacua as well as the critical-bubble configuration are usually more than several orders of magnitude, making the problem almost unsolvable using even the revised flow-based multimodal modeling method mentioned above. There are at least two reasons for this: 1) The large hierarchy among the probability weights makes challenge $\mathcal{B}$ or the mode-collapse problem much more serious and cannot be dealt with easily. 2) More fundamentally, even if we have a perfectly trained flow model, it still cannot sample from the minimum of the probability distribution efficiently, as described by challenge $\mathcal{C}$ or the “rare-event sampling” problem.

To address these challenges, Ref. [13] uses the method of MMCMC (as briefly mentioned at the end of Sec. 2.2), which is in essence a reweighted sampling method based on a modification to the action

$$S_E[\phi] \rightarrow S_E[\phi] + W(\mathcal{O}(\phi)), \quad (12)$$

that aims to uniformly sample across the measurable $\mathcal{O}(\phi)$ space by choosing $W(\mathcal{O}(\phi)) \approx -\log P(\mathcal{O}(\phi))$. Here, $P(\mathcal{O}(\phi))$ is the true probability distribution of $\mathcal{O}(\phi)$. After that, one can retrieve the correct weights by performing a post-reweighting. However, one cannot in principle know $W(\mathcal{O}(\phi))$ a priori. Ref. [13] resolves this issue by performing two rounds of sampling: in the first round, a piecewise construction of $W(\mathcal{O}(\phi))$ is performed on the fly during an evolving MMCMC sampling process; after the first round, one then uses the constructed $W(\mathcal{O}(\phi))$ to perform a full MMCMC sampling. More recently, Ref. [23] proposes the “normalizing flow assisted importance sampling (NOFIS)” algorithm to estimate the rare event probability. The basic idea is similar to that of MMCMC, which aims to train a NF model based on a reweighted target distribution. Given that the region of rare events to be modeled is known a priori (but it is either hard or time-consuming to compute the associated probability directly), one can divide the region into a series of nested sub-regions with strictly decreasing probabilities. First, one should focus on modeling the smallest sub-region by “suppressing” the probabilities outside. Based on this initial model, one can then proceed to model the next-level sub-region. This process is repeated until the entire rare-event region is properly modeled.

3.3 Partitioning flow-based MCMC

Inspired by the ideas of MMCMC and NOFIS, we propose the method of “partitioning flow-based (PF) MCMC” (PFMCMC) to address the critical bubble sampling problem. Similar to MMCMC and NOFIS, we train our flow models based on reweighted actions; however, instead of trying to model the target distribution globally, we “slice” the space of $\mathcal{O}(\phi)$ into several partitions and then “glue” them together to obtain the “overall” $P(\mathcal{O}(\phi))$. Furthermore, since we cannot identify the rare-event region a priori (and thus we cannot directly adopt the NOFIS method), our slicing is done in a sequential rather than in a nested fashion for the purpose of searching for the minimum.

There are at least three advantages of this approach: 1) If chosen wisely, all of the partitions will be either unimodal (around the two vacua) or monotonic except for the one containing the critical-bubble configuration, which, however, should have a small enough hierarchy that can be properly managed. Therefore, within each partition of the order-parameter space, flow-based MCMC can easily outperform traditional MCMC as has been demonstrated in many previous works. 2) Compared to modeling a multimodal distribution using one single flow model directly, modeling several unimodal and/or monotonic distributions (plus a small-hierarchy U-shaped distribution) can in general be more efficient and stable. 3) When it comes to bubble nucleation rate calculation, we only need the information between $\mathcal{O}_F$ and $\mathcal{O}_C$, and thus we can ignore much of the order-parameter space from around $\mathcal{O}_C$ to $\mathcal{O}_T$ using this method, even though there is no obstacle to construct the remaining space. The flip side of our method is that we will need to train multiple models for one single benchmark. On the other hand, since most of the partitions will either be monotonic or unimodal, the training of each model will be easier to be managed.²

We now explicitly explain the methodology. For each model benchmark, we choose a wide range of the simulation region in terms of the measurable $[\mathcal{O}_L, \mathcal{O}_R]$ to include both $\mathcal{O}_F$ and $\mathcal{O}_C$, or $\mathcal{O}_F, \mathcal{O}_C \in [\mathcal{O}_L, \mathcal{O}_R]$. We partition the simulation region $[\mathcal{O}_L, \mathcal{O}_R]$ into several mildly hierarchical overlapping partitions with each individual partition defined by $[\mathcal{O}_i, \mathcal{O}_j]$. Targeting one individual partition, we train a flow model according to the reweighted action

$$S_{\text{E}}^{ij}[\phi] = \begin{cases} S_{\text{E}}[\phi] + W_{\text{L}}^i \times \left(e^{B_{\text{L}}^i(\mathcal{O}_i - \mathcal{O})} - 1 \right), & \mathcal{O} \leq \mathcal{O}_i \\ S_{\text{E}}[\phi], & \mathcal{O}_i < \mathcal{O}(\phi) < \mathcal{O}_j \\ S_{\text{E}}[\phi] + W_{\text{R}}^j \times \left(e^{B_{\text{R}}^j(\mathcal{O} - \mathcal{O}_j)} - 1 \right), & \mathcal{O}_j \leq \mathcal{O} \end{cases}, \quad (13)$$

where $W_{\text{L}}^i, W_{\text{R}}^j, B_{\text{L}}^i, B_{\text{R}}^j > 0$ are manually chosen parameters to “suppress” the irrelevant order-parameter space, thus effectively making a “partition”. After sampling from each partition, we glue the partitions together following this procedure:

1. Starting from the partition containing the local minimum, we simply take the histogram values of the partition in focus within the non-overlapping regions.
2. When two partitions overlap, the endpoint of one of the partitions in that region must be a (local) maximum of the original partition. We denote the probability distribution of this

²We also remark that one can try to construct a mixture model based on the sub-models of these partitions as mentioned in Ref. [21], but that is not necessary for the calculation of bubble nucleation rates.

first partition (partition 1) as $P_1(\langle\phi\rangle)$, while the other (partition 2) as $P_2(\langle\phi\rangle)$. We scale the entire partition 2 so that the endpoint of partition 1 matches to the value at the same point of partition 2.

3. For the remaining points in the overlapping region, we take a simple weighted average of the histogram values from the two partitions $[w_1 P_1(\langle\phi\rangle) + w_2 P_2(\langle\phi\rangle)] / (w_1 + w_2)$, $w_1 \geq w_2$. In our study, this is done by human intervention, while a more sophisticated gluing algorithm is certainly possible.

The idea behind steps 2. and 3. is that it is more likely for a flow model to undersample or mismodel a minimum than a maximum, which is a compelling reason for making overlapping partitions than mutually exclusive ones.

3.3.1 A toy 2D ϕ^4 theory

As an example, we use a toy 2D (only two spacial dimensions) ϕ^4 theory on a lattice of size $L \times L$ to illustrate our PFMCMC algorithm. The action of this model is described by

$$S_{\phi^4}[\phi] = \int d^2x \left[\frac{1}{2} (\partial^\mu \phi)^2 + V_{\phi^4}(\phi) \right], \quad V_{\phi^4}(\phi) = \alpha \phi + m^2 \phi^2 + \kappa \phi^3 + \lambda \phi^4, \quad (14)$$

where the discretized Laplacian operator is given to the leading order by

$$(\nabla_L^2)_x^{(1)} = \sum_{i=a\hat{x}, a\hat{y}} \frac{1}{a^2} (\phi_{x+i} + \phi_{x-i} - 2\phi_x) . \quad (15)$$

The measurable $\mathcal{O}$ is chosen to be the spacial mean field value or

$$\mathcal{O}(\phi) = \langle\phi\rangle \equiv \frac{1}{L^2} \int d^2x \phi(x) . \quad (16)$$

We choose three model benchmarks with the common choice of $(m^2, \lambda) = (-2, 1)$, and with (I) $(\alpha, \kappa) = (0, 0)$, (II) $(\alpha, \kappa) = (0.008, 0)$, and (III) $(\alpha, \kappa) = (0, -0.05)$, respectively. Their potential profiles are shown in Fig. 3.

For the flow model, we use real-valued non-volume preserving (real NVP) flows [38] constructed from non-linear squared couplings (NLSq) [41]

$$g(x) = a + b x + \frac{c}{1 + (d x + g)^2}, \quad (17)$$

where a, b, c, d, g are model parameters to be optimized. Following Ref. [24], each coupling layer checkerboards the sites and updates one parity (white or black) of sites conditioned on the frozen values of the other parity, with the parities switched between layers (see Fig. 2). We use a convolutional neural network (CNN) of 4 hidden layers with 12 nodes and a kernel size of 5 to parametrize the transformations in each coupling layer. All hidden layers have a leaky ReLU activation of negative slope 0.01, *i.e.*

$$\text{Leaky ReLU}_{0.01}(x) = \begin{cases} x, & x > 0 \\ 0.01 x, & \text{otherwise} \end{cases}, \quad (18)$$




Figure 3: The potential profiles of the three model benchmarks of the 2D ϕ^4 theory [see Eq. (14)]. Here, all benchmarks have the common choice of $(m^2, \lambda) = (-2, 1)$, with $(\alpha, \kappa) =$ (I) $(0, 0)$, (II) $(0.008, 0)$, and (III) $(0, -0.05)$, respectively.

except for the last one which has a tanh activation. The entire model consists of 40 such coupling layers, resulting in roughly 500K trainable hyperparameters. We choose the lattice size $L = 10$, lattice spacing $a = 1$, and batch size $N = 1000$. We split the $\langle\phi\rangle$ space into three partitions: $[-1.5, 0]$, $[-0.5, 0.5]$, and $[0, 1.5]$, with a common choice of $W_L^i = W_R^i = 1000$ and $B_L^i = B_R^j = 1$. We train the models based on the loss function given in Eq. (10).




Figure 4: *Left:* The simulated probability distribution $P(\langle\phi\rangle)$ for benchmark (I) of the V_{ϕ^4} model is shown in the blue and solid line and is normalized by the value at the (global) maximum. The red, green and yellow dashed lines are the distributions of the individual partitions. *Right:* The same as the left one, but for benchmark (II).

Using the PFMCMC method described above, we first show the probability distributions of $\langle\phi\rangle$ [normalized by the value at the (global) maximum] for benchmarks (I) and (II) of the 2D ϕ^4 model in Fig. 4. For this toy model, we compare the distributions of benchmarks (I) and (II) obtained through our method to those through the existing two methods reported in Ref. [21] in Fig. 5: the adiabatic retraining + flow-distance regularization (AR+FR) method

and the augmented hybrid MC (AHMC) method, the latter of which can be thought of (not guaranteed) as the “truth distributions”.³ We can see that for both benchmarks, the PFMCMC distributions are very well aligned with the AHMC distributions except for the local minimum of benchmark (II), which is slightly underestimated. This latter issue can in principle be resolved by making finer partitions. On the contrary, the AR+FR method consistently overestimates the local minima. These comparisons demonstrate the success of our PFMCMC algorithm to obtain the right probability distributions and the reliability of the PFMCMC method when constructing multimodal distributions.




Figure 5: *Left*: The probability distributions of benchmark (I) obtained through the PFMCMC, AR+FR, and AHMC methods. *Right*: The same as the left one, but for benchmark (II). The AR+FR and AHMC distributions are directly quoted from Ref. [21].




Figure 6: The same as Fig. 4 but for benchmark (III).

We further show the probability distribution of benchmark (III) obtained through the PFMCMC method in Fig. 6. For this benchmark point, the probability of the metastable vacuum is suppressed by $\sim \mathcal{O}(10^{-5})$ compared to the stable vacuum. Our PFMCMC method can still

³One can view the AHMC method as directly sampling from the target distribution, whereas the other two methods (PFMCMC and AR+FR) might exhibit bias toward the proposal distributions learned by the flow models. Therefore, if AHMC is executed perfectly, it should generate the “truth distributions,” yet in principle, there is no analytical means to confirm this.

capture the bimodal structure, which is difficult to obtain using the straightforward flow-based sampling methods.

3.3.2 A toy 2D ϕ^6 theory

As a demonstration of how the PFMC MC method can be used to model a largely hierarchical distribution and as a prelude to our study of the finite-temperature (2+1)D scalar field theory in Sec. 4, we study one benchmark of the 2D $\mathbb{Z}_2$ -symmetric ϕ^6 theory described by

$$S_{\phi^6}[\phi] = \int d^2x \left[\frac{1}{2}(\partial^\mu \phi)^2 + V_{\phi^6}(\phi) \right], \quad V_{\phi^6}(\phi) = \frac{m^2}{2}\phi^2 + \frac{\lambda}{4!}\phi^4 + \frac{\eta}{6!}\phi^6, \quad (19)$$

with $(m^2, \lambda, \eta) = (6.50, -7.53, 4.87)$. The potential profile is shown in Fig. 7. We still choose the same measurable $\mathcal{O}(\phi) = \langle \phi \rangle$, but normalize the probability distribution by the value at $\langle \phi \rangle = 0$, which is the metastable vacuum. Moreover, while this benchmark has a two-fold stable vacuum degeneracy, we only focus on transition to the positive stable vacuum. The same order parameter, normalization scheme, and choice of the stable vacuum are also used in the later study of the finite-temperature theory (see Sec. 4). For this particular benchmark, we choose a




Figure 7: The potential profile of the 2D $\mathbb{Z}_2$ -symmetric ϕ^6 theory [see Eq. (19)] benchmark with $(m^2, \lambda, \eta) = (6.50, -7.53, 4.87)$.

slightly smaller lattice size $L = 8$ with the same lattice spacing $a = 1$. Because this benchmark has a probability hierarchy of $\sim \mathcal{O}(10^{-50})$ between the metastable vacuum and the critical-bubble configuration, we deploy roughly 20 partitions to apply the PFMC MC method. For each partition, we adopt the same model and train upon the same loss function as before and choose $W_L^i = W_R^i = 1000$, $B_L^i = 6$, and $B_R^j = 4$. The different choices for $B_{L,R}$ is because we only focus on scanning the region between the metastable vacuum and the critical configuration and thus do not expect as much effect from the stable vacuum.

We show the simulated final distribution $P_{V_{\phi^6}}(\langle \phi \rangle)$ (left panel) and the zoom-in view around the critical separatrix (right panel) in Fig. 8. In the right panel, we show both the individual partition and combined distributions. We note that the sharp drops at the end points of the green and red dashed curves are the results of partitioning, whose values are not used in the analysis. For this benchmark, we choose the critical separatrix to be $\langle \phi \rangle \in [0.870, 0.880]$ in

accordance with the histogram bin width, with the local minimum sitting at $\langle\phi\rangle = 0.875$ and $\epsilon = 0.01$. The statistical factor of this particular benchmark is given by

$$F_S = \frac{1}{2} \times \frac{P_{V_{\phi^6}}(\langle\phi\rangle \in [0.870, 0.880])}{0.01 \times P_{V_{\phi^6}}(\langle\phi\rangle \in [0, 0.870])} = 8.60 \times 10^{-53},$$

where the factor of $1/2$ accounts for the negative side of the metastable vacuum for this $\mathbb{Z}_2$ -symmetric model. We remark that the large hierarchy in this benchmark is unlikely to be captured using the straightforward flow-based sampling methods introduced before.




Figure 8: *Left*: The overall view of the $P_{V_{\phi^6}}(\langle\phi\rangle)$ distribution [normalized by $P_{V_{\phi^6}}(0)$] for the V_{ϕ^6} model benchmark. *Right*: The zoom-in view of the left panel around the critical separatrix. The green band given by $\langle\phi\rangle \in [0.870, 0.880]$ stands for the critical separatrix. The dashed lines in the right panel stand for the individual partition distributions while the blue and solid line for the combined distribution. The sharp drops at the endpoints of the green and red dashed curves are the results of partitioning, whose values are not used in the analysis.

4 (2+1)D $\mathbb{Z}_2$ -symmetric real scalar theory

We apply the PFMC MC method to study the finite-temperature (2+1)D $\mathbb{Z}_2$ -symmetric real scalar theory. We first summarize the properties of the theory and choose one benchmark with a manageable probability hierarchy. Then, we calculate the statistical factor of the benchmark using the PFMC MC method with a few different lattice setups. To demonstrate the nucleation rate calculation procedure, we choose the setups with unit lattice spacing $a = 1$ and calculate the corresponding dynamical factors. After that, we “compare” the calculated nucleation rates of the benchmark to the one obtained using *CosmoTransitions* [42] and *BubbleDet* [43], packages based on the bounce solution action using the saddle-point approximation, in a way that takes into consideration the finite-spacing effect.

4.1 Finite-temperature (2+1)D $\mathbb{Z}_2$ -symmetric theory

The finite-temperature Euclidean action of the system is defined in a 2D space and is given by

$$S_E[\phi] = \frac{1}{T} \int d^2x \left[\frac{1}{2} \partial_\mu \phi \partial^\mu \phi + V_{\text{eff}}(\phi) \right] \equiv \frac{S[\phi]}{T} . \quad (20)$$

Here, the finite-temperature effective potential $V_{\text{eff}}(\phi)$ is calculated at one-loop level (see Appendix B for detailed calculations) and is

$$V_{\text{eff}}(\phi) = V_0(\phi) + V_1(\phi, T) + V_\beta(\phi, T) , \quad (21)$$

with

$$V_0(\phi) = \frac{m^2}{2} \phi^2 + \frac{\lambda}{4!} \phi^4 + \frac{\eta}{6!} \phi^6 , \quad V_1(\phi, T) = -\frac{1}{12\pi} m_{\text{eff}}^3(\phi, T) , \quad (22)$$

$$V_\beta(\phi, T) = -\frac{T^3}{2\pi} \left[\frac{m_{\text{eff}}}{T} \text{Li}_2 \left(e^{-\frac{m_{\text{eff}}}{T}} \right) + \text{Li}_3 \left(e^{-\frac{m_{\text{eff}}}{T}} \right) \right] , \quad (23)$$

with Li_n denoting the polylogarithmic functions. Here m_{eff} stands for the thermal effective mass and is given by

$$m_{\text{eff}}^2(\phi, T) = m_T^2 + \frac{\lambda}{2} \phi^2 + \frac{\eta}{4!} \phi^4 + \frac{\eta}{4} \phi^2 \tau_1(T, m_T) , \quad (24)$$

where m_T is the IR-cutoff parameter and

$$\tau_1(T, m_T) = -\frac{T}{2\pi} \ln \left(2 \sinh \frac{m_T}{2T} \right) . \quad (25)$$

Though the non-analytic structure of the potential can induce imaginary components, the benchmark that we choose for the remaining study is strictly real, as we will demonstrate in the next section. The details of the theory can be found in Appendix B. While this theory has a two-fold degeneracy in the stable vacua, we only focus on the transition to the positive vacuum in the remaining study.

4.2 Benchmark test: bubble simulation

For the benchmark test, we choose $(m^2, \lambda, \eta, T) = (1, -5, 10, 7.2)$,⁴ which corresponds to $m_T^2 \approx 1.296$. The potential profile and the corresponding bounce solution are shown in Fig. 9. From the bounce solution, one can see that the spherical bubble radius $r_{b,\text{eff}} \approx 2.5 - 3$. The bounce action of this benchmark is $S_{b,\text{eff}}/T \approx 16.60$. We remark that for this benchmark, the potential is strictly real, as can be seen from the imaginary component plotted in the figure. Also, this benchmark is chosen to have a smaller probability hierarchy in the measurable space for the sake of simulation efficiency (there should be no obstacle to simulate a benchmark with a more hierarchical probability distribution, as we demonstrated in Sec. 3.3). As a proof of concept of the PFMC MC method, our discretization scheme simply takes the first-order discretized Laplacian given in Eq. (15) and uses the continuum effective potential for the lattice action. For a rigorous study of the lattice perturbation theory of scalar fields, one can refer to Ref. [44].

⁴We keep the unit of energy to be 1.




Figure 9: *Left*: the potential profile for the benchmark point of the finite-temperature effective potential of the (2+1)D $\mathbb{Z}_2$ real scalar theory with $(m^2, \lambda, \eta, T) = (1, -5, 10, 7.2)$, which corresponds to $m_T^2 \approx 1.296$. Note that the imaginary component of the potential is zero. *Right*: the corresponding bounce solution for this benchmark potential. The bounce action is $S_{b,\text{eff}}/T \approx 16.60$.

For the PFMCMC simulation, we use the same model and loss function as the ones described in Sec. 3.3. In principle, a proper choice of a should be much smaller than the correlation length ξ of the system, *i.e.*, $a \ll \xi \sim 1/m_\phi$, m_ϕ being the mass of the scalar field [44]. Furthermore, one should extract the continuum limit result by sampling a few different a 's and taking $a \rightarrow 0$ based on a function fit (such as $b_0 + b_1 a^3$ used in Ref. [15]). In a system with a finite lattice spacing, the high-momentum modes are cut off, leading to a hard upper bound on the bubble wall tension scale. In the thin-wall limit, the critical-bubble radius scales roughly as $\sigma/\Delta V$, σ being the one-dimensional wall tension and ΔV the two-dimensional vacuum pressure difference, and thus one direct outcome of a finer lattice spacing is the increase of the bubble radius, which will become much milder when the previous condition $a \ll \xi$ is satisfied. This requirement will in general lead to a high training cost of the flow models, which we will discuss in Sec. 5. Since this work is a proof of concept of the PFMCMC method, we will not restrict ourselves to the mentioned condition, but will only demonstrate the finite spacing effect on the simulation results without performing any function fits or extrapolations.

For this purpose, we fix the lattice size $L = 10$ and choose a few different lattice spacings $a = 1, \frac{5}{6}, \frac{2}{3}$. We deploy 1, 1, 2 partitions for the three lattice spacings within the ranges of $\langle \phi \rangle \in [0, 1], [0, 1.2], [0, 1.0] \cup [0.7, 1.7]$, respectively, with a common choice of $W_L^i = W_R^i = 1000$ and $B_L^i = B_R^i = 2$ for all partitions. The calculated probability distributions with respect to $\langle \phi \rangle$ [$P_{V_{\text{eff}}}(\langle \phi \rangle)$] are shown in Fig. 10. The critical separatrices are given by

$$\langle \phi \rangle_{a=1} \in [0.518, 0.528], \quad \langle \phi \rangle_{a=\frac{5}{6}} \in [0.781, 0.791], \quad \langle \phi \rangle_{a=\frac{2}{3}} \in [1.491, 1.501]. \quad (26)$$

It can be seen that the position of the critical separatrix shifts to the right as a decreases, which implies the increase of the critical-bubble size. Furthermore, the probability hierarchy also increases along with the decrease in a , which is another result of including higher-momentum modes as they can contribute to larger free energy differences. We also show one critical-bubble configuration for each of the three setups in Fig. 11, from which one can see the increase in the bubble size as a decreases.




Figure 10: The $P_{V_{\text{eff}}}(\langle\phi\rangle)$ distributions for $L = 10$ and $a = 1, \frac{5}{6}, \frac{2}{3}$, respectively. The critical separatrices are shown in terms of the colored bands with the numerical values given in Eq. (26).




Figure 11: Example critical-bubble configurations properly contained in the lattice of $L = 10$ and (left) $a = 1$, (middle) $a = 5/6$, and (right) $a = 2/3$, respectively. Their $\langle\phi\rangle$ values fall within the critical separatrices defined in Eq. (26).

On the other hand, to determine a proper lattice size, we follow the argument given in Ref. [13] to derive the condition. First, the lattice size L should be larger than the bubble size or $L > 2r_b$ to contain the whole bubble. Here r_b stands for the radius of a generic simulated bubble and is measured through real simulations, which could differ from the one obtained from the bounce solution profile, $r_{b,\text{eff}}$ (see the right panel of Fig. 9). Furthermore, to avoid the dominance of unphysical periodic “bubbles” (bubbles which will minimize their free energies by connecting through the unphysical periodic boundaries), one needs a more stringent condition: $L \gtrsim \pi r_b$. Nevertheless, for a too large L , the probability of obtaining the interesting bubble configurations will be diluted by noises outside the bubble region.⁵ It has been argued and shown in Refs. [13, 14] that as long as the lattice size satisfies the previous condition, the nucleation rate should only have exponentially small volume-dependence, which can be shown through a fit based on the function $\Gamma_\infty + c_0 \exp(-c_1 L)$ with $c_{0,1} > 0$, and taking $L \rightarrow \infty$ will give the infinite-volume result. Moreover, because of the fixed-size nature of the critical bubbles, another signature of a properly chosen lattice size is the mildly volume-dependent distribution of

⁵See Ref. [15] for a better choice of the measurable to address this issue.

the volume-integrated measurables, such as $L^2\langle\phi\rangle$. For demonstration purpose, we choose three different lattice sizes $L = 9, 10, 11$ that are not significantly larger than πr_b with unit lattice spacing $a = 1$, for which we find that $r_b \lesssim 3$ for the majority of bubbles. We will explicitly show the mild volume-dependence of the $L^2\langle\phi\rangle$ distribution, but will not perform any function fits or extrapolations.




Figure 12: *Left:* The $P_{V_{\text{eff}}}(\langle\phi\rangle)$ distributions for $a = 1$ and $L = 9, 10, 11$, respectively. *Right:* The $P_{V_{\text{eff}}}(L^2\langle\phi\rangle)$ distributions for the same lattice setups. The critical separatrices are shown in terms of the colored bands in the left plot with the numerical values given in Eq. (27). The curves in the right panel are mildly volume-dependent because $L^2\langle\phi\rangle$ should have a fixed value for a fixed-size bubble.

For each lattice size, we deploy one partition of $\langle\phi\rangle \in [0, 1]$ with the common choice of $W_L^i = W_R^i = 1000$ and $B_L^i = B_R^i = 2$. First note that $\langle\phi\rangle$ scales as $1/L^2$, so the partition $[0, 1]$ can sufficiently cover the interesting region from $\mathcal{O}_F$ to $\mathcal{O}_C$. Secondly, this benchmark is not that hierarchic, so we do not need to choose many partitions. We show $P_{V_{\text{eff}}}(\langle\phi\rangle)$ and the calculated probability distribution with respect to $L^2\langle\phi\rangle$ [$P_{V_{\text{eff}}}(L^2\langle\phi\rangle)$] in Fig. 12 for $L = 9, 10, 11$, along with the chosen critical separatrices shown in the $P_{V_{\text{eff}}}(\langle\phi\rangle)$ plot. The critical separatrices are given by

$$\langle\phi\rangle_{L=9} \in [0.644, 0.654], \quad \langle\phi\rangle_{L=10} \in [0.518, 0.528], \quad \langle\phi\rangle_{L=11} \in [0.423, 0.433]. \quad (27)$$

As pointed out in Refs. [13, 14], $P_{V_{\text{eff}}}(L^2\langle\phi\rangle)$ is nearly identical between the false vacuum and the critical separatrix for different lattice sizes that satisfy the previously-mentioned minimal condition. The reason is that for a fixed-size bubble, $L^2\langle\phi\rangle$ is a fixed value and should exhibit mild volume dependence. For illustration purpose, we show two example critical-bubble configurations of the smallest lattice size $L = 9$ in Fig. 13, both of which are properly contained in the lattice.

The statistical factors of different lattice sizes are calculated to be

- $L = 9$

$$F_S = \frac{1}{2} \times \frac{P_{V_{\text{eff}}, L=9}(\langle\phi\rangle \in [0.644, 0.654])}{0.01 \times P_{V_{\text{eff}}, L=9}(\langle\phi\rangle \in [0, 0.644])} = 0.502.$$

- $L = 10$

$$F_S = \frac{1}{2} \times \frac{P_{V_{\text{eff}}, L=10}(\langle\phi\rangle \in [0.518, 0.528])}{0.01 \times P_{V_{\text{eff}}, L=10}(\langle\phi\rangle \in [0, 0.518])} = 0.622.$$




Figure 13: Example critical-bubble configurations properly contained in the lattice of $L = 9$ and $a = 1$. Their $\langle\phi\rangle$ values fall within the critical separatrix defined in Eq. (27).

- $L = 11$

$$F_S = \frac{1}{2} \times \frac{P_{V_{\text{eff}}, L=11}(\langle\phi\rangle \in [0.423, 0.433])}{0.01 \times P_{V_{\text{eff}}, L=11}(\langle\phi\rangle \in [0, 0.423])} = 0.762 .$$

A simple power-law function $0.501 \times (L/9)^{2.1}$ can provide a good fit to the three data points, whose L^2 -dependence is mostly cancelled after being divided by $V = L^2$ [see Eq. (3)]. The residual L^2 -dependence is expected, and should (along with the L^2 -dependence of the dynamical factors) be dealt with through the function fit mentioned previously.

4.3 Dynamical factor and nucleation rate calculations

To demonstrate the procedure of nucleation rate calculation, we finish the calculation of the dynamical factor using the setups with $a = 1$ and $L = 9, 10, 11$. Because the uncertainties of the statistical factors are much smaller than those of the dynamical factors, we will simply use their face values when calculating the nucleation rates. Since this paper mainly serves as a proof of principle to efficiently sample rare events and calculate the statistical factors, we will not perform the extrapolation to obtain the infinite-volume and continuum limits, which are necessary to calculate the nucleation rates of a realistic model.

The procedure to calculate the dynamical factor was described around Eq. (6). A proper timestep Δt for a Langevin evolution should be shorter than the timescale to go from $\langle\phi\rangle_C$ to another configuration where $\log P(\langle\phi\rangle)$ differs by order 1. As long as this constraint is satisfied, the Brownian random walk nature of the evolution will guarantee that the quantity $S|d\langle\phi\rangle/dt|_{\phi_C}$ is Δt -independent [13]. To demonstrate this, we perform the Langevin evolution for the finite-temperature benchmark using the fourth-order Runge-Kutta method in real time with $\Delta t = 0.010, 0.015, 0.020$. On the other hand, we terminate the simulations when $\langle\phi\rangle < 0$ (enters the metastable phase) or $\langle\phi\rangle > \langle\phi\rangle_{\text{th}}$ (enters the stable phase), with $\langle\phi\rangle_{\text{th}}$ chosen such that the majority of the trajectories that have “tunneled” into the stable phase will not turn back and cross the critical separatrix again. To determine the proper $\langle\phi\rangle_{\text{th}}$ for different lattice sizes, we try out a few different values and choose the smallest one for each size starting from which the dynamical factors obtained begin to converge. The explicit choices are $\langle\phi\rangle_{\text{th}, (L=9)} = 1.4$, $\langle\phi\rangle_{\text{th}, (L=10)} = 1.2$, $\langle\phi\rangle_{\text{th}, (L=11)} = 1.0$.




Figure 14: Two examples of Langevin evolution trajectories of the critical configurations with both forward and backward evolutions in time and ($L = 9$, $a = 1$, $\Delta t = 0.01$). The blue trajectory has successfully tunnelled from the metastable to the stable vacuum, while the orange trajectory has failed. The green band stands for the critical separatrix.

We show two example trajectories of $L = 9$, $a = 1$ and $\Delta t = 0.01$ in Fig. 14, one of which successfully tunnels (even $N_{\text{crossings}}$) while the other fails (odd $N_{\text{crossings}}$). To calculate the dynamical factors, we generate roughly 2K critical configurations and collect 20 trajectories for each configuration for all L 's and Δt 's. The dynamical factors obtained are summarized in Table 1. It can be seen that the dynamical factors of different Δt do agree with each other within uncertainties.

	$L = 9$	$L = 10$	$L = 11$
$\Delta t = 0.010$	$(14.8 \pm 1.22) \times 10^{-3}$	$(11.1 \pm 0.862) \times 10^{-3}$	$(8.45 \pm 0.751) \times 10^{-3}$
$\Delta t = 0.015$	$(14.5 \pm 1.22) \times 10^{-3}$	$(11.2 \pm 0.943) \times 10^{-3}$	$(8.05 \pm 0.715) \times 10^{-3}$
$\Delta t = 0.020$	$(14.1 \pm 1.20) \times 10^{-3}$	$(11.5 \pm 0.976) \times 10^{-3}$	$(8.36 \pm 0.775) \times 10^{-3}$

Table 1: The dynamical factors (F_D 's) for different values of L and Δt .

Finally, combining both the statistical and dynamical factors, we can obtain the nucleation rates via Eq. (3). Using the results of $\Delta t = 0.01$ for the dynamical factor, the nucleation rates for all three lattice sizes are

$$\Gamma_{L,\text{eff}} = \begin{cases} (L = 9) : & (9.2 \pm 0.76) \times 10^{-5} \\ (L = 10) : & (6.9 \pm 0.54) \times 10^{-5} \\ (L = 11) : & (5.3 \pm 0.47) \times 10^{-5} \end{cases}, \quad (28)$$

which has the unit of m^{-4} ($m = 1$ in the simulations). To make a “fair” comparison to the saddle-point approximation, we apply a hard cutoff at $\Lambda = 1/a = 1$ [45] on the loop-corrected potentials [see Eqs. (40) and (41)] and obtain $\Gamma_{b,\text{eff}}^{\text{cut}} = 2.22 \times 10^{-4}$ based on CosmoTransitions and BubbleDet. While this matches the lattice-simulated results at the order-of-magnitude

level, we remark that a rigorous comparison should be made by taking the infinite-volume and continuum limits on the lattice results and leave this type of exercises to our future studies for $(3 + 1)\text{D}$ models.

5 Discussion and conclusions

As briefly mentioned, a significant issue when increasing the lattice volume for bubble simulation is the dilution of the bubble configurations of interest due to bulk fluctuations. Specifically, the fluctuations of $\langle\phi\rangle$ in the purely metastable or stable phase scale roughly as $1/\sqrt{V}$, while the contribution of critical bubbles to $\langle\phi\rangle$ scales as $1/V$ due to their fixed size nature [13]. Therefore, attempting to measure and extract bubble properties in the large-volume regime using the measurable $\langle\phi\rangle$ is likely to result in numerous “false bubbles”. To extrapolate lattice results to the infinite-volume limit, a suitable measurable that is less affected by significant bulk noise is required. One such measurable is $\langle\phi^2\rangle - 2A\langle\phi\rangle$, where A is a manually chosen parameter. This measurable has been recently utilized in Ref. [15].

Another fundamental challenge not just for the PFMCMC method but also for the general flow-based simulation framework is the high training cost of NFs as the problem scale increases. In the context of the scalar FOPT problem we study, this challenge arises from several aspects: 1) the necessity of a small enough lattice spacing; 2) the need for a sufficiently large lattice size; and 3) the dimensions of the problem. The scalability issue of NF training has been discussed in the context of lattice QCD in Ref. [46], which presents inconclusive findings regarding the scaling nature of the problem but suggests several potential directions for future improvement. For instance, one direct approach is parallelizing the algorithm. Alternatively, if obtaining a high-quality flow model proves challenging due to resource constraints, a hybrid approach combining flow-based and traditional Monte Carlo methods may alleviate the issue by reducing the reliance on the model’s quality while still leveraging its modeling capabilities to some extent.

While this paper serves as a proof of concept for the PFMCMC method, the immediate next step is to apply this method to a $(3+1)\text{D}$ theory, such as the ϕ^4 theory examined in Ref. [15], or other phenomenologically interesting models. An intriguing avenue of investigation lies in studying the non-sphericity of cosmological FOPT bubbles. This analysis could shed light on single-bubble gravitational wave generation, as recently explored in Ref. [47]. Additionally, numerically simulating bubbles and observing their evolution could improve our understanding of super-critical bubble expansion, including their interactions with the plasma [48, 49]. These insights may contribute to more precise calculations of stochastic gravitational wave production from cosmological FOPTs and their detectability in gravitational wave experiments [8].

In conclusion, we have demonstrated a novel flow-based algorithm, the PFMCMC method, capable of simulating cosmological FOPTs nonperturbatively, drawing on the nucleation framework outlined by Langer. Tested on various toy 2D models and the finite-temperature $(2+1)\text{D}$ $\mathbb{Z}_2$ -symmetric real scalar theory, our method yields results comparable to those found in the literature, effectively addressing highly hierarchical probability distribution challenges. Furthermore, we have illustrated how these results enable the calculation of the dynamical factor and nucleation rate of a thermal scalar field system. By integrating the power of NFs into the study of cosmological phase transitions, this method offers a promising approach for investigating higher-dimensional or more complex and realistic problems.

Acknowledgments

The authors thank Kyle Cranmer for useful discussion. The work is supported by the U.S. Department of Energy under the contract DE-SC-0017647. We thank the Center for High Throughput Computing at the University of Wisconsin-Madison (<https://chtc.cs.wisc.edu>) [50] for providing the computing resources.

A The nucleation theory for relativistic quantum fields

In this appendix, we review the nucleation theory for relativistic quantum fields, first derived by Coleman for zero-temperature theories [10] and later generalized for finite-temperature theories by Linde [11, 12]. In Sec. 2, we reviewed how the nucleation rate for a classical field can be decomposed into a dynamical factor F_D and a statistical factor F_S [see Eq. (3)]. By matching the free energy functional of the classical field $F[\phi]/T$ to the Euclidean quantum field action functional $S_E[\phi]$, F_S can be calculated from the fluctuations around the saddle point and is given up to the second order by [10, 51–55]

$$F_S = \left(\frac{S_E[\phi_b]}{2\pi} \right)^{(d-1)/2} \left| \frac{\det(-\nabla^2 + V''(\phi_F))}{\det'(-\nabla^2 + V''(\phi_b))} \right|^{1/2} e^{-(S_E[\phi_b] - S_E[\phi_F])} , \quad (29)$$

where S_E is the finite-temperature Euclidean action of the system and given by

$$S_E[\phi] = \frac{1}{T} \int d^{d-1}x \left[-\frac{1}{2} \phi \nabla^2 \phi + V(\phi) \right] \equiv \frac{S[\phi]}{T} . \quad (30)$$

Here, d is the spacetime dimension; ϕ_F is the metastable vacuum configuration; $\det'$ signifies that the zero modes are omitted from the functional determinant. Assuming spherical symmetry for the system, ϕ_b is solved using the equation of motion [10]

$$\left. \frac{\delta S[\phi]}{\delta \phi} \right|_{\phi_b} = -\partial_r^2 \phi_b(r) - \frac{d-2}{r} \partial_r \phi_b(r) + V'(\phi_b) = 0 , \quad (31)$$

subject to the boundary conditions

$$\lim_{r \rightarrow \infty} \phi_b(r) = \phi_F, \quad \partial_r \phi_b|_{r=0} = 0 . \quad (32)$$

On the other hand, for a real scalar theory, F_D can be calculated using Langevin dynamics and is given by [9, 56, 57]

$$F_D = \frac{1}{2\pi} \left(\sqrt{|\lambda_-| + \frac{\tilde{\eta}^2}{4}} - \frac{\tilde{\eta}}{2} \right) , \quad (33)$$

where λ_- is the negative eigenvalue of the functional second derivative $-\nabla^2 + V''(\phi_b)$ and $\tilde{\eta}$ is the Langevin damping coefficient, which has to be determined with real-time input.

B Finite-temperature (2+1)D $\mathbb{Z}_2$ -symmetric scalar theory

In this appendix, we study the finite-temperature (2+1)D $\mathbb{Z}_2$ -symmetric scalar theory. The most general renormalizable tree-level potential is given by

$$V_0(\phi) = \frac{m^2}{2}\phi^2 + \frac{\lambda}{4!}\phi^4 + \frac{\eta}{6!}\phi^6. \quad (34)$$

In order to calculate the thermal effective mass, we consider the self-energy corrections to ϕ up to two-loop since the six-point interaction in Eq. (34) will contribute at the leading perturbative order. Diagrammatically, this is given by

$$\begin{aligned} m_{\text{eff},0}^2(\phi, T) = & \text{---} \bigotimes \text{---} + \text{---} \text{V} \text{---} + \text{---} \text{X} \text{---} \\ & + \text{---} \text{O} \text{---} + \text{---} \text{X} \text{---} + \text{---} \text{O} \text{---} \\ & = m^2 + \frac{\lambda}{2}\phi^2 + \frac{\eta}{4!}\phi^4 + \frac{\lambda}{2}\tau_1(T, m) + \frac{\eta}{4}\phi^2\tau_1(T, m) + \frac{\eta}{8}\tau_1^2(T, m), \end{aligned} \quad (35)$$

where $\tau_1(T, m)$ is the thermal-loop bubble, as given by

$$\tau_1(T, m) = \bigcirc = -\frac{T}{2\pi} \ln \left(2 \sinh \frac{m}{2T} \right). \quad (36)$$

As pointed out in Ref. [58, 59], τ_1 is manifestly IR-divergent in the $m/T \rightarrow 0$ limit. Therefore, we introduce the thermal mass (Debye mass) m_T as the IR-cutoff, which can be obtained through solving the gap equation

$$\begin{aligned} m_T^2 = & \text{---} \bigotimes \text{---} + \text{---} \text{O} \text{---} + \text{---} \text{O} \text{---} \\ & = m^2 + \frac{\lambda}{2}\tau_1(T, m_T) + \frac{\eta}{8}\tau_1^2(T, m_T). \end{aligned} \quad (37)$$

Consequently, m_{eff} is given by

$$m_{\text{eff}}^2(\phi, T) = m_T^2 + \frac{\lambda}{2}\phi^2 + \frac{\eta}{4!}\phi^4 + \frac{\eta}{4}\phi^2\tau_1(T, m_T), \quad (38)$$

in contrast to the effective mass-squared $m_{\text{eff},0}^2$ without IR-regularization given in Eq. (35).

For a generic theory, one can in principle separate the loop corrections to the effective potential into two parts: the zero-temperature Coleman-Weinberg potential V_1 and the finite-temperature effective potential V_β . However, because m_T is required for IR regularization in this

model, the usual Coleman-Weinberg potential becomes temperature-dependent, $V_1 = V_1(\phi, T)$. Diagrammatically, these two corrections are given by

$$V_1(\phi, T) + V_\beta(\phi, T) = \text{diagram 1} + \text{diagram 2} + \text{diagram 3} + \text{diagram 4} + \text{diagram 5} + \dots, \quad (39)$$

where

$$V_1(\phi, T) = \frac{1}{2} \int \frac{d^3 p}{(2\pi)^3} \log [p^2 + m_{\text{eff}}^2(\phi, T)] = -\frac{1}{12\pi} m_{\text{eff}}^3(\phi, T), \quad (40)$$

and

$$V_\beta(\phi, T) = T \int \frac{d^2 p}{(2\pi)^2} \log \left[1 - e^{-\frac{1}{T} \sqrt{p^2 + m_{\text{eff}}^2(\phi, T)}} \right] = -\frac{T^3}{2\pi} \left[\frac{m_{\text{eff}}}{T} \text{Li}_2 \left(e^{-\frac{m_{\text{eff}}}{T}} \right) + \text{Li}_3 \left(e^{-\frac{m_{\text{eff}}}{T}} \right) \right], \quad (41)$$

where Li_n is the polylogarithmic function. Overall, the finite-temperature effective potential is given by

$$V_{\text{eff}}(\phi, T) = V_0(\phi) + V_1(\phi, T) + V_\beta(\phi, T). \quad (42)$$

We briefly discuss the perturbativity bounds on λ and η . For λ , we consider the diagram

$$I_\lambda = \text{diagram} \sim \frac{\lambda^2}{2} \int \frac{d^3 q}{(2\pi)^3} \frac{1}{(q^2 + m^2)^2} = \frac{\lambda^2}{2 \times 8\pi m}, \quad (43)$$

where we ignore all the $\mathcal{O}(1)$ factors. By requiring $|I_\lambda| \lesssim |\lambda|$, the perturbativity bound on λ is roughly given by $|\lambda/m| \lesssim 8\pi$. For η , we consider the diagram

$$I_\eta^{(2)} = \text{diagram} \sim \frac{\eta^2}{3} \int \frac{d^3 q_1 d^3 q_2}{(2\pi)^6} \frac{1}{q_1^2 q_2^2 (q_1 + q_2)^2} \sim \frac{\eta^2}{3 \times (2\pi^2)^2} \int \frac{dq}{q}, \quad (44)$$

where we ignore all the $\mathcal{O}(1)$ factors. By requiring $|I_\eta^{(2)}| \lesssim |\eta|$ (or the tree-level coupling), the perturbativity bound on η is roughly given by $|\eta| \lesssim (2\pi^2)^2$.

Before concluding, we remark that the non-analytic structure of $V_1(\phi, T)$ and $V_\beta(\phi, T)$ can possibly induce imaginary components, which suggests the instability of the homogeneous field state as pointed out in Refs. [60, 61]. Nevertheless, if the imaginary part of V_{eff} is sufficiently small compared to its real part, the decay timescale is then negligible compared to the transition timescale. In this case, we can safely extract the real part and use it to perform phase transition and nucleation calculations.

References

- [1] V. A. Kuzmin, V. A. Rubakov, and M. E. Shaposhnikov, *On the Anomalous Electroweak Baryon Number Nonconservation in the Early Universe*, *Phys. Lett. B* **155** (1985) 36.
- [2] A. G. Cohen, D. B. Kaplan, and A. E. Nelson, *Progress in electroweak baryogenesis*, *Ann. Rev. Nucl. Part. Sci.* **43** (1993) 27–70, [[hep-ph/9302210](#)].
- [3] Y. Bai, T.-K. Chen, and M. Korwar, *QCD-collapsed domain walls: QCD phase transition and gravitational wave spectroscopy*, *JHEP* **12** (2023) 194, [[arXiv:2306.17160](#)].
- [4] **NANOGrav** Collaboration, Z. Arzoumanian et al., *The NANOGrav 12.5 yr Data Set: Search for an Isotropic Stochastic Gravitational-wave Background*, *Astrophys. J. Lett.* **905** (2020), no. 2 L34, [[arXiv:2009.04496](#)].
- [5] **EPTA** Collaboration, S. Chen et al., *Common-red-signal analysis with 24-yr high-precision timing of the European Pulsar Timing Array: inferences in the stochastic gravitational-wave background search*, *Mon. Not. Roy. Astron. Soc.* **508** (2021), no. 4 4970–4993, [[arXiv:2110.13184](#)].
- [6] B. Goncharov et al., *On the Evidence for a Common-spectrum Process in the Search for the Nanohertz Gravitational-wave Background with the Parkes Pulsar Timing Array*, *Astrophys. J. Lett.* **917** (2021), no. 2 L19, [[arXiv:2107.12112](#)].
- [7] B. B. P. Perera et al., *The International Pulsar Timing Array: Second data release*, *Mon. Not. Roy. Astron. Soc.* **490** (2019), no. 4 4666–4687, [[arXiv:1909.04534](#)].
- [8] P. Athron, C. Balázs, A. Fowlie, L. Morris, and L. Wu, *Cosmological phase transitions: From perturbative particle physics to gravitational waves*, *Prog. Part. Nucl. Phys.* **135** (2024) 104094, [[arXiv:2305.02357](#)].
- [9] J. S. Langer, *Statistical theory of the decay of metastable states*, *Annals Phys.* **54** (1969) 258–275.
- [10] S. R. Coleman, *The Fate of the False Vacuum. 1. Semiclassical Theory*, *Phys. Rev. D* **15** (1977) 2929–2936. [Erratum: *Phys.Rev.D* 16, 1248 (1977)].
- [11] A. D. Linde, *Fate of the False Vacuum at Finite Temperature: Theory and Applications*, *Phys. Lett. B* **100** (1981) 37–40.
- [12] A. D. Linde, *Decay of the False Vacuum at Finite Temperature*, *Nucl. Phys. B* **216** (1983) 421. [Erratum: *Nucl.Phys.B* 223, 544 (1983)].
- [13] G. D. Moore and K. Rummukainen, *Electroweak bubble nucleation, nonperturbatively*, *Phys. Rev. D* **63** (2001) 045002, [[hep-ph/0009132](#)].
- [14] O. Gould, S. Güyer, and K. Rummukainen, *First-order electroweak phase transitions: A nonperturbative update*, *Phys. Rev. D* **106** (2022), no. 11 114507, [[arXiv:2205.07238](#)].
- [15] O. Gould, A. Kormu, and D. J. Weir, *A nonperturbative test of nucleation calculations for strong phase transitions*, [arXiv:2404.01876](#).
- [16] U. Wolff, *Critical Slowing Down*, *Nucl. Phys. B Proc. Suppl.* **17** (1990) 93–102.
- [17] L. Del Debbio, G. M. Manca, and E. Vicari, *Critical slowing down of topological modes*, *Phys. Lett. B* **594** (2004) 315–323, [[hep-lat/0403001](#)].
- [18] H. B. Meyer, H. Simma, R. Sommer, M. Della Morte, O. Witzel, and U. Wolff, *Exploring*

- the HMC trajectory-length dependence of autocorrelation times in lattice QCD, *Comput. Phys. Commun.* **176** (2007) 91–97, [[hep-lat/0606004](#)].
- [19] S. Schaefer, R. Sommer, and F. Viotto, *Investigating the critical slowing down of QCD simulations*, *PoS LAT2009* (2009) 032, [[arXiv:0910.1465](#)].
 - [20] **ALPHA** Collaboration, S. Schaefer, R. Sommer, and F. Viotto, *Critical slowing down and error analysis in lattice QCD simulations*, *Nucl. Phys. B* **845** (2011) 93–119, [[arXiv:1009.5228](#)].
 - [21] D. C. Hackett, C.-C. Hsieh, M. S. Albergo, D. Boyda, J.-W. Chen, K.-F. Chen, K. Cranmer, G. Kanwar, and P. E. Shanahan, *Flow-based sampling for multimodal distributions in lattice field theory*, [arXiv:2107.00734](#).
 - [22] K. A. Nicoli, C. J. Anders, T. Hartung, K. Jansen, P. Kessel, and S. Nakajima, *Detecting and mitigating mode-collapse for flow-based sampling of lattice field theories*, *Phys. Rev. D* **108** (2023), no. 11 114501, [[arXiv:2302.14082](#)].
 - [23] Z. Gao, D. Zhang, L. Daniel, and D. S. Boning, *Rare event probability learning by normalizing flows*, [arXiv:2310.19167](#).
 - [24] M. S. Albergo, G. Kanwar, and P. E. Shanahan, *Flow-based generative models for Markov chain Monte Carlo in lattice field theory*, *Phys. Rev. D* **100** (2019), no. 3 034515, [[arXiv:1904.12072](#)].
 - [25] B. A. Berg and T. Neuhaus, *Multicanonical ensemble: A New approach to simulate first order phase transitions*, *Phys. Rev. Lett.* **68** (1992) 9–12, [[hep-lat/9202004](#)].
 - [26] L. Tierney, *Markov Chains for Exploring Posterior Distributions*, *The Annals of Statistics* **22** (1994), no. 4 1701 – 1728.
 - [27] C.-C. Chien, J.-H. She, and F. Cooper, *Mean-field description of pairing effects, bkt physics, and superfluidity in 2d bose gases*, *Annals of Physics* **347** (Aug., 2014) 192–206, [[arXiv:1203.3254](#)].
 - [28] P. Konietin and V. Pastukhov, *2d dilute bose mixture at low temperatures*, *Journal of Low Temperature Physics* **190** (Dec., 2017) 256–266, [[arXiv:1708.00432](#)].
 - [29] V. Pastukhov, *Polaron in dilute 2d bose gas at low temperatures*, *Journal of Physics B: Atomic, Molecular and Optical Physics* **51** (July, 2018) 155203, [[arXiv:1712.06978](#)].
 - [30] V. Pastukhov, *Ground-state properties of a dilute two-dimensional bose gas*, *Journal of Low Temperature Physics* **194** (Oct., 2018) 197–208, [[arXiv:1803.05242](#)].
 - [31] L. P. Csernai and J. I. Kapusta, *Nucleation of relativistic first order phase transitions*, *Phys. Rev. D* **46** (1992) 1379–1390.
 - [32] G. D. Moore, K. Rummukainen, and A. Tranberg, *Nonperturbative computation of the bubble nucleation rate in the cubic anisotropy model*, *JHEP* **04** (2001) 017, [[hep-lat/0103036](#)].
 - [33] N. Goldenfeld, *Lectures On Phase Transitions And The Renormalization Group*. Frontiers in physics. Westview Press, 1992.
 - [34] G. D. Moore, *Sphaleron rate in the symmetric electroweak phase*, *Phys. Rev. D* **62** (2000) 085011, [[hep-ph/0001216](#)].
 - [35] J. Zinn-Justin, *Quantum Field Theory and Critical Phenomena: Fifth Edition*. Oxford

University Press, 04, 2021.

- [36] T. Minka et al., *Divergence measures and message passing*, tech. rep., Microsoft Research, 2005.
- [37] I. Kobyzev, S. J. Prince, and M. A. Brubaker, *Normalizing flows: An introduction and review of current methods*, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **43** (Nov., 2021) 3964–3979.
- [38] L. Dinh, J. Sohl-Dickstein, and S. Bengio, *Density estimation using real nvp*, [arXiv:1605.08803](#).
- [39] S. Kullback and R. A. Leibler, *On Information and Sufficiency*, *The Annals of Mathematical Statistics* **22** (1951), no. 1 79 – 86.
- [40] L. Del Debbio, J. M. Rossney, and M. Wilson, *Efficient modeling of trivializing maps for lattice ϕ^4 theory using normalizing flows: A first look at scalability*, *Phys. Rev. D* **104** (2021), no. 9 094507, [[arXiv:2105.12481](#)].
- [41] Z. M. Ziegler and A. M. Rush, *Latent normalizing flows for discrete sequences*, [arXiv:1901.10548](#).
- [42] C. L. Wainwright, *CosmoTransitions: Computing Cosmological Phase Transition Temperatures and Bubble Profiles with Multiple Fields*, *Comput. Phys. Commun.* **183** (2012) 2006–2013, [[arXiv:1109.4189](#)].
- [43] A. Ekstedt, O. Gould, and J. Hirvonen, *BubbleDet: a Python package to compute functional determinants for bubble nucleation*, *JHEP* **12** (2023) 056, [[arXiv:2308.15652](#)].
- [44] O. Gould, *Real scalar phase transitions: a nonperturbative analysis*, *JHEP* **04** (2021) 057, [[arXiv:2101.05528](#)].
- [45] M. G. Alford and M. Gleiser, *Metastability in two-dimensions and the effective potential*, *Phys. Rev. D* **48** (1993) 2838–2844, [[hep-ph/9304245](#)].
- [46] R. Abbott et al., *Aspects of scaling and scalability for flow-based sampling of lattice QCD*, *Eur. Phys. J. A* **59** (2023), no. 11 257, [[arXiv:2211.07541](#)].
- [47] K. Blum and M. Mirbabayi, *A single-bubble source for gravitational waves in a cosmological phase transition*, [arXiv:2403.20164](#).
- [48] D. Bodeker and G. D. Moore, *Electroweak Bubble Wall Speed Limit*, *JCAP* **05** (2017) 025, [[arXiv:1703.08215](#)].
- [49] T. Krajewski, M. Lewicki, and M. Zych, *Bubble-wall velocity in local thermal equilibrium: hydrodynamical simulations vs analytical treatment*, [arXiv:2402.15408](#).
- [50] Center for High Throughput Computing, *Center for high throughput computing*, 2006.
- [51] C. G. Callan, Jr. and S. R. Coleman, *The Fate of the False Vacuum. 2. First Quantum Corrections*, *Phys. Rev. D* **16** (1977) 1762–1768.
- [52] A. I. Vainshtein, V. I. Zakharov, V. A. Novikov, and M. A. Shifman, *ABC’s of Instantons*, *Sov. Phys. Usp.* **25** (1982) 195.
- [53] A. Andreassen, D. Farhi, W. Frost, and M. D. Schwartz, *Precision decay rate calculations in quantum field theory*, *Phys. Rev. D* **95** (2017), no. 8 085011, [[arXiv:1604.06090](#)].
- [54] A. Andreassen, W. Frost, and M. D. Schwartz, *Scale Invariant Instantons and the Complete Lifetime of the Standard Model*, *Phys. Rev. D* **97** (2018), no. 5 056006,

- [[arXiv:1707.08124](#)].
- [55] W.-Y. Ai, B. Garbrecht, and C. Tamarit, *Functional methods for false vacuum decay in real time*, *JHEP* **12** (2019) 095, [[arXiv:1905.04236](#)].
 - [56] W.-Y. Ai, B. Garbrecht, and P. Millington, *Radiative effects on false vacuum decay in Higgs-Yukawa theory*, *Phys. Rev. D* **98** (2018), no. 7 076014, [[arXiv:1807.03338](#)].
 - [57] S. Chigusa, T. Moroi, and Y. Shoji, *Precise Calculation of the Decay Rate of False Vacuum with Multi-Field Bounce*, *JHEP* **11** (2020) 006, [[arXiv:2007.14124](#)].
 - [58] T. Appelquist and U. W. Heinz, *Three-dimensional $O(N)$ Theories at Large Distances*, *Phys. Rev. D* **24** (1981) 2169.
 - [59] F. A. Dilkes and D. G. C. McKeon, *Phase transitions in a scalar theory in $(2+1)$ -dimensions*, *Annals Phys.* **276** (1999) 1–11.
 - [60] E. J. Weinberg and A.-q. Wu, *Understanding Complex Perturbative Effective Potentials*, *Phys. Rev. D* **36** (1987) 2474.
 - [61] M. Sher, *Electroweak Higgs Potentials and Vacuum Stability*, *Phys. Rept.* **179** (1989) 273–418.