

A Monte Carlo estimator of flow fields for sampling and noise problems

Michael S. Albergo^a and Gurtej Kanwar^{b,*}

^a*Society of Fellows, Harvard University*

78 Mount Auburn St, Cambridge, MA United States

^b*Higgs Centre for Theoretical Physics, School of Physics and Astronomy,
University of Edinburgh, EH9 3FD Edinburgh, United Kingdom*

E-mail: gkanwar@ed.ac.uk

Learned field transformations may help address ubiquitous critical slowing down and signal-to-noise problems in lattice field theory. In the context of an annealed sequence of distributions, field transformations are defined by integrating flow fields that exactly solve a local transport problem. These proceedings discuss a new Monte Carlo approach to evaluating these flow fields, which can then be used directly in such contexts or as a means of generating unbiased training data for machine learning approaches. By defining the Monte Carlo estimator using coupled Langevin noise, the statistical noise in the required integrals is significantly mitigated. Demonstrations of the method include a $U(1)$ transport problem and an $SU(N)$ glueball correlator.

*The 42nd International Symposium on Lattice Field Theory (LATTICE2025)
2-8 November 2025
Tata Institute of Fundamental Research, Mumbai, India*

*Speaker

1. Introduction

We are interested in the problem of dynamically annealing from a simple prior density $\rho_0(\phi) := e^{-S_0(\phi)+F_0}$ through a sequence of densities parameterized by a fictitious time $t \geq 0$ as

$$\rho_t(\phi) := e^{-S_t(\phi)+F_t}, \quad F_t = -\log \left(\int d[\phi] e^{-S_t(\phi)} \right). \quad (1)$$

This can be achieved by deterministic evolution of the fields, starting from samples $\phi_{t=0}$ distributed according to $\rho_0(\phi)$, which are evolved in time via a *flow*,

$$\dot{\phi}_t = b_t(\phi_t). \quad (2)$$

The density of samples ϕ_t follow the distribution ρ_t , for all t , as long as the time-dependent vector field $b_t(\phi)$ satisfies the associated *continuity equation*,

$$\nabla \cdot (b_t \rho_t) = -\partial_t \rho_t \quad \implies \quad \nabla \cdot b_t - b_t \cdot \nabla S_t = \partial_t S_t - \partial_t F_t. \quad (3)$$

Depending on the choice of ρ_t , this paradigm has various use cases in the lattice field theory computational workflow. Most generally, the computation of expected values of operators under ρ_t can be rewritten as an expectation value with respect to ρ_0 as

$$\langle O \rangle_t := \int d[\phi] O(\phi) \rho_t(\phi) = \int d[\phi_0] O(\phi_t) \rho_0(\phi_0), \quad (4)$$

where on the right side ϕ_t is defined by solving (2) with ϕ_0 as the initial condition. This change-of-variables relation has been explored as an option to reduce the cost of Monte Carlo integration [4, 7, 11], which may have significant computational benefits when one is interested in target distributions ρ_t that are more challenging to simulate than the prior density ρ_0 .

More recently, other works have explored applying a derivative with respect to t in (4) to insert operators into correlation functions [1, 3, 6]. Specifically, for correlation functions involving an operator Q , we define the time-dependent probability density in (1) using the action $S_t(\phi) = S(\phi) + tQ(\phi)$. Differentiating any n -point correlation function with respect to this time dependence leads to a connected $(n + 1)$ -point correlation function,

$$\partial_t \langle O_1 \dots O_n \rangle_t \Big|_{t=0} = \langle O_1 \dots O_n Q \rangle_0 - \langle O_1 \dots O_n \rangle_0 \langle Q \rangle_0. \quad (5)$$

Assuming the flow field b_t satisfies (3), the derivative at $t = 0$ can be directly calculated as

$$\partial_t \langle O_1 \dots O_n \rangle_t \Big|_{t=0} = \langle b_{t=0} \cdot \nabla (O_1 \dots O_n) \rangle_0. \quad (6)$$

This provides a separate approach to estimating the $(n + 1)$ -point correlation function, with the possibility of better signal-to-noise scaling.

In either the sampling or signal-to-noise context, if the flow field b_t only approximately satisfies (3), exact corrective factors can be included to provide unbiased results. However, the computational benefits obtained in either case depend on the quality of the solution, and it is important to solve this high-dimensional partial differential equation to good accuracy. Several approaches have been devised to approximately solve for b_t in the context of lattice field theory, including analytical small- t

expansion [11], optimization of a physics-inspired ansatz [4], and machine learning parameterization and optimization [2, 3].

In this work, we propose a new Monte Carlo estimator for $b_t(\phi)$ which directly satisfies (3) in expectation. While this introduces noise, an unbiased estimator of b_t can be directly used in evaluating higher-point functions via the relation (6). It also serves as a “ground truth” estimate which can either be compared to machine learning outputs or used as training data for such methods. A complementary approach, based on stochastic automatic differentiation, was developed by Catumba and Ramos in Ref. [6]; Section 4 makes a concrete connection between the present work and the Catumba-Ramos method.

2. Monte Carlo estimates of flow fields

To begin, we write the flow field in terms of a time-dependent scalar potential $\varphi_t(\phi)$,

$$b_t(\phi) = \nabla\varphi_t(\phi), \quad (7)$$

in terms of which the continuity equation (3) reduces to the nonlinear Poisson equation

$$\nabla^2\varphi_t - \nabla\varphi_t \cdot \nabla S_t = \partial_t S_t - \partial_t F_t. \quad (8)$$

While still a challenging partial differential equation, a solution to (8) is formally given for each value of t by the Feynman-Kac formula [9],

$$\varphi_t(\phi) = - \int_0^\infty d\tau \mathbb{E}[\partial_t S_t(\Phi^\tau) - \partial_t F_t \mid \Phi^0 = \phi]. \quad (9)$$

Here, Φ^τ is a random variable sampled by Langevin evolution over simulation time τ from the starting configuration $\Phi^0 = \phi$, according to¹

$$\partial_\tau \Phi^\tau = -\nabla S_t(\Phi^\tau) + \eta^\tau, \quad \langle \eta^\tau \eta^{\tau'} \rangle = 2\delta(\tau - \tau'). \quad (10)$$

The expectation value in (9) is taken over realizations of the stochastic noise η^τ . A proof that this solves the relevant Poisson equation can for example be found in Ref. [2].

Equation (10) is the standard Langevin process for the action S_t [5], and the integrand therefore converges asymptotically to its expectation value under S_t . Under the assumption of geometric ergodicity, the integrand by definition converges exponentially to zero:

$$\lim_{\tau \rightarrow \infty} \mathbb{E}[\partial_t S_t(\Phi^\tau) - \partial_t F_t] = \langle \partial_t S_t \rangle_t - \partial_t F_t = 0. \quad (11)$$

Equation (9) is then a well-defined integral representation for the potential and correspondingly the flow field. Practically speaking, however, integrating over noisy estimates of the expectation value leads to growing noise in this estimate of φ_t as $\tau \rightarrow \infty$. This is illustrated in Figure 1 for a simple time-dependent action $S_t = \frac{1}{2}\phi^2 + t\phi$ over a single variable $\phi \in \mathbb{R}$; in this example, the numerical results are compared against the exact solution $\varphi_t(\phi) = -\phi$ and $b_t(\phi) = -1$.

¹More precisely, this stochastic differential equation can be understood in the Itô sense as $d\Phi^\tau = -\nabla S_t d\tau + \sqrt{2}dW_\tau$ in terms of the Wiener process W_τ .




Figure 1: Evaluation of $\varphi_t(\phi = 0.5)$ for the one-variable Gaussian distribution at $t = 0$ described in the main text. Results are shown for 1024 Langevin walkers Φ^τ initialized from $\Phi^0 = 0.5$. The right plot compares the cumulative integral as a function of the maximum integration limit T against the true value, demonstrating both convergence and the growing statistical noise expected from integrating the noisy integrand.




Figure 2: Evaluation of $b_t(\phi = 0.5) = \nabla\varphi_t(0.5)$ for the same theory as Figure 1. For this Gaussian theory, there is no statistical variation in b_t using coupled Brownian noise. Evaluation is shown for one walker initialized from $\Phi^0 = 0.5$ and is compared against the true value.

Despite growing noise in this estimator of (9), we here propose to evaluate $b_t(\phi)$ by analytically differentiating (9) with respect to the initial condition $\Phi^0 = \phi$. This amounts to using linearity to move the gradient into the expectation value,

$$b_t(\phi) = - \int_0^\infty d\tau \nabla_\phi \mathbb{E}[\partial_t S_t(\Phi^\tau) - \partial_t F_t | \Phi^0 = \phi] = - \lim_{T \rightarrow \infty} \mathbb{E}[\int_0^T \nabla_\phi \partial_t S_t(\Phi^\tau(\phi))]. \quad (12)$$

Above, $\Phi^\tau(\phi)$ is written as a function of the initial condition ϕ , with respect to which the derivative is taken. For each Langevin noise sample η^τ to be used in the average, the gradient in (12) can for example be evaluated analytically by applying backpropagation to the numerical integration of $\hat{\varphi}_t^{[0,T]}$, where

$$\hat{\varphi}_t^{[a,b]} := - \int_a^b d\tau \partial_t S_t(\Phi^\tau). \quad (13)$$

This gradient is straightforward to evaluate using standard automatic differentiation packages; Pytorch [12] is used for all numerical results in this work.

The key insight is that differentiating with respect to the initial condition for fixed Langevin noise can yield low-noise estimates of $b_t(\phi)$ even when (9) leads to noisy estimates for $\varphi_t(\phi)$. This is demonstrated in Figure 2 for the one-variable example, in which case this gives a *zero-noise*




Figure 3: Left: Illustration of Langevin evolution with coupled noise in a quartic potential, showing eventual convergence of various initial conditions. Right: The dependence on the initial conditions, $\partial\Phi^\tau/\partial\phi$, which decreases exponentially when measured for the same trajectories.

estimator for $b_t(\phi)$. We note also the additional benefit of not having to estimate $\partial_t F_t = \langle \partial_t S_t \rangle_t$ because it drops out under the gradient.

This estimate of the gradient conditioned on fixed samples of the noise process η^τ can be considered a specific case of the more general approach of *coupled-noise* Langevin processes, in which correlated Langevin noise is used across different initial conditions. Remarkably, coupled-noise Langevin evolution for many physical potentials results in convergence to a path independent of the initial condition ϕ for *each noise sample* η^τ . An example of this behavior is shown for a simple quartic potential in Figure 3, using a simple coupling where identical noise η^τ is applied for all initial conditions ϕ . Using any coupled noise with this convergence property leads to the field dependence $\nabla_\phi \Phi^\tau$ vanishing asymptotically, as shown in the right panel of the figure. Evaluating (12) in this way translates to the integral converging asymptotically for each noise sample, suggesting an asymptotically constant growth of noise in estimates of $b_t(\phi)$ with this approach, if such a coupled Langevin process can be defined.

Finally, we describe an adjoint-sensitivity method of evaluating (12) without backpropagation [10]. The *adjoint state* at the reversed-time coordinate $\xi := (T - \tau)$ is defined to be

$$A^\xi := \frac{\partial}{\partial \Phi^\xi} \hat{\varphi}_t^{[T-\xi, T]}. \quad (14)$$

It evolves according to a stochastic differential equation (SDE) for ξ running from $0 \rightarrow T$,

$$\partial_\xi A^\xi = -\nabla \partial_t S_t(\Phi^\xi) - (A^\xi \cdot \nabla) \nabla \partial_t S_t(\Phi^\xi), \quad (15)$$

where sample paths of Φ^ξ are obtained as the reverse of the sample paths Φ^τ . The result in (12) is recovered by taking the expectation value and infinite-time limit,

$$b_t(\phi) = \nabla \varphi_t = \lim_{T \rightarrow \infty} \mathbb{E}[A^{\xi=T}]. \quad (16)$$

While numerical results in the remainder of this work are obtained from the more direct backpropagation method, this adjoint sensitivity method is used for theoretical results below and may provide a more efficient route to evaluating b_t in future work.

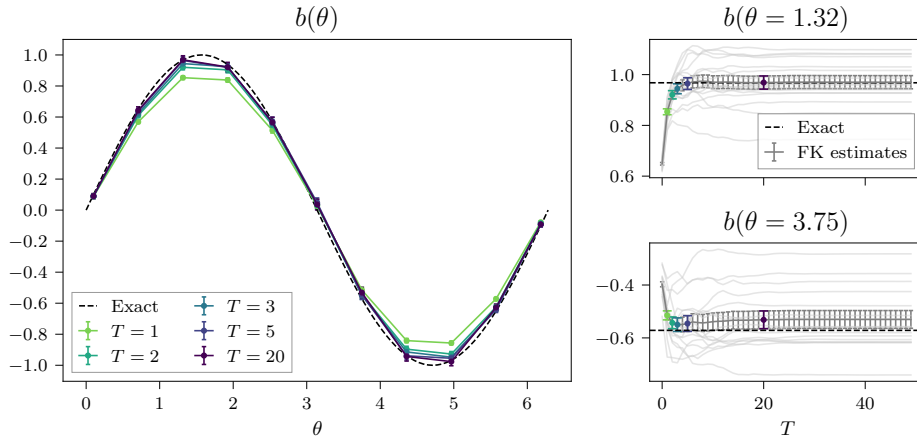


Figure 4: Estimate of $b(\theta)$ using the Feynman-Kac method for the von Mises distribution defined in the main text. Each point in the left panel is evaluated using 4096 Langevin walkers and is plotted for a range of upper integration limits T . The right panels show detailed convergence as a function of T for two values of θ . Evolution of block-averaged sample estimates are shown by light gray traces, and the bootstrap average over all walkers is indicated by the dark gray error bars.

3. Non-Euclidean domains

We introduce two improvements that are helpful in the subsequent numerical demonstrations on compact, non-Euclidean domains.

First, as usual for Langevin evolution, (10) can be modified with a noise kernel K to²

$$\partial_\tau \Phi^\tau = -K(\Phi^\tau) \nabla S_t(\Phi^\tau) + \nabla \cdot K(\Phi^\tau) + \sqrt{K(\Phi^\tau)} \eta^\tau, \quad (17)$$

without affecting the equilibrium distribution. The solution $\tilde{b}_t(\phi)$ obtained by using these modified dynamics in the evaluation of (12) satisfies a modified continuity equation

$$\nabla \cdot (K \tilde{b}_t) - (K \tilde{b}_t) \cdot \nabla S_t = \partial_t S_t - \partial_t F_t. \quad (18)$$

A valid flow field is then obtained in this case by assigning $b_t = K \tilde{b}_t$. Depending on the choice of kernel, the convergence and precision of estimates of b_t may be improved. This approach is demonstrated below for a single $U(1)$ variable.

Second, we describe a specific coupling scheme for the Langevin noise in evaluating Feynman-Kac path integrals for $SU(N)$ theories. This approach is demonstrated below for alleviating the noise problem in a glueball correlator for a 2+1D $SU(N)$ lattice gauge theory.

3.1 $U(1)$ von Mises distribution

As a first example, we consider a single $U(1)$ variable $\theta \in [0, 2\pi]$ with time-dependent action

$$S_t(\theta) = t \cos(\theta), \quad (19)$$

which is a von Mises distribution with concentration parameter $\kappa = -t$. At $t = 0$, the flow field can be analytically solved to give $b(\theta) := b_{t=0}(\theta) = \sin(\theta)$. This exact solution is compared against a

²This stochastic differential equation is to be understood in the Itô sense.

kernel-modified Feynman-Kac estimate in Figure 4, with the noise kernel $K(\theta) = 0.5 + \sin(\theta/2)$, which is chosen to be concave almost everywhere to improve the rate of convergence of the coupled-noise Langevin process. As shown in the figure, the integral for b_t converges *per Monte Carlo sample*, giving asymptotically constant noise as the integration cutoff is removed. In contrast, we have found that evaluation without a kernel suffers from growing noise, requiring careful tuning of the cutoff time to balance systematic and statistical error.

3.2 $SU(N)$ glueball correlator

As a second example, we apply the Feynman-Kac path integral to $SU(N)$ variables using coupled noise. In this context, the intrinsic positive curvature of the $SU(N)$ manifold provides an opportunity to design a useful coupling that is convergent even without a convex potential. Here, we choose the noise coupling to minimize the average distance between the evolving $SU(N)$ variable V^τ and a reference variable U^τ , where the Hilbert-Schmidt distance is used,

$$d_{HS}^2(U^\tau, V^\tau) := 1 - \frac{1}{N} \text{ReTr}[U^\tau (V^\tau)^\dagger]. \quad (20)$$

We first consider left-acting Brownian noise, which can be written in terms of Lie algebra components as $\eta^\tau = \eta_a^\tau T^a$. Any rotation $\xi_a^\tau = O_{ab} \eta_b^\tau$ preserves the correct statistics, so that a coupling can be designed by choosing the orthogonal matrix O as a function of U^τ and V^τ , with noise samples η and ξ respectively applied to U^τ and V^τ . Considering the noise average of $d_{HS}^2(e^{i\eta^\tau \sqrt{d\tau}} U^\tau, e^{i\xi^\tau \sqrt{d\tau}} V^\tau)$ for $d\tau \rightarrow 0$, the rotation O that minimizes this distance is [8]

$$O = \underset{O}{\text{argmax}} \{O^T M\} = M \sqrt{M^T M}, \quad \text{where } M^{ab} := \text{ReTr}[T^b T^a U^\tau (V^\tau)^\dagger]. \quad (21)$$

A right-acting, coupled Brownian noise can be defined by an analogous derivation.

The remaining choice is which trajectory or configuration U^τ to use as a reference. We here apply a *self-coupling* strategy, since we are only interested in the derivative with respect to the initial condition: we take $U^\tau = \mathring{V}^\tau$ itself as the reference, where the circle indicates that we consider \mathring{V}^τ to be independent of the initial condition for the purpose of differentiating.³

To test this approach, we evaluate a 0^{++} glueball correlator in an $SU(2)$ Yang-Mills theory with Wilson gauge action at inverse gauge coupling $\beta = 4.0$ on a $4^2 \times 8$ lattice. To improve convergence of the Feynman-Kac integral for b_t in this context, we apply the $SU(N)$ coupled noise defined above acting on both the left and right side of each gauge link. We further apply a constant kernel, $K = \alpha P + (1 - P)$, to rescale noise along the pure-gauge directions by a factor $\alpha = 10$, where P is a projector onto the pure-gauge direction acting in the space of left and right acting noise at each site x . The observable is determined by applying the derivative trick to a one-point function,

$$\langle O(x)O(0) \rangle - \langle O(x) \rangle \langle O(0) \rangle = \langle b_{t=0} \cdot (\nabla O(x)O(0)) \rangle, \quad (22)$$

where the glueball interpolator is chosen to be a local plaquette operator $O(x) = \text{ReTr}[P_{01}(x)]$. The results are compared to a standard vacuum-subtracted estimate in Figure 5; significantly more precise results are obtained, while using approximately $8\times$ fewer configurations.

³This can be implemented in Pytorch using `detach()`.




Figure 5: Evaluation of a 0^{++} glueball correlator as a function of separation x/a for an $SU(2)$ lattice gauge theory as described in the main text. Estimates using a standard vacuum-subtracted estimator are compared to the flow-field approach with a Feynman-Kac estimator integrated up to a variety of integration limits T .

4. Relation to Catumba-Ramos

A similar method to estimate flow fields was recently introduced by Catumba and Ramos in Ref. [6], in which the final state of a continuous Markovian process is differentiated with respect to the time dependence of the action. This can be related to the present work by taking the Markovian process in their method to be an overdamped Langevin simulation for the action S_t with a fixed choice of parameter t ,

$$\partial_\xi \bar{\Phi}^\xi = -\nabla S_t(\bar{\Phi}^\xi) + \eta^\xi. \quad (23)$$

Here we use the barred variable $\bar{\Phi}^\xi$ to distinguish from the Langevin evolution Φ^τ used above and suggestively use ξ for Langevin time in this context. The sensitivity of $\bar{\Phi}^\xi$ to the t -dependence of the action follows an SDE,

$$\bar{A}^\xi := \frac{d}{dt} \bar{\Phi}^\xi \implies \partial_\xi \bar{A}^\xi = -\nabla \partial_t S_t(\bar{\Phi}^\xi) - (\bar{A}^\xi \cdot \nabla) \nabla S_t(\bar{\Phi}^\xi). \quad (24)$$

For a fixed starting distribution for $\bar{\Phi}^0$, ergodicity implies that the Markov chain evolution of $\bar{\Phi}^\xi$ will converge to the distribution ρ_t for $\xi \rightarrow \infty$. Perturbing the action appearing in the Langevin evolution to S_{t+dt} will similarly result in convergence to the distribution ρ_{t+dt} , such that a solution for the flow field that corresponds to this change in density is given by the conditional expectation value

$$b_t(\phi) = \lim_{T \rightarrow \infty} \mathbb{E}\left[\frac{d}{dt} \bar{\Phi}^T \mid \bar{\Phi}^T = \phi\right] = \lim_{T \rightarrow \infty} \mathbb{E}[\bar{A}^{\xi=T}]. \quad (25)$$

It is clear from (15) and (24) that there is an intimate connection between the Catumba-Ramos approach and the Feynman-Kac determination of the flow field. Here we sketch a proof that these formulations are in fact equivalent under reversal of Langevin time $\tau \leftrightarrow \xi$, when the starting condition for the former is the equilibrium distribution, ρ_t . The equivalence of these formulations is immediate if the following two distributions are equal to each other

1. The distribution of paths Φ^ξ obtained by reversing the paths Φ^τ generated by the standard Langevin simulation in (10), starting from $\Phi^0 = \phi$ and evolving to time $\tau = T$.
2. The conditional distribution of paths $\bar{\Phi}^\xi$ generated by the (same) Langevin simulation in (23), starting from $\bar{\Phi}^0$ sampled from the equilibrium distribution ρ_t and conditioning on the final state $\bar{\Phi}^T = \phi$.

We show equivalence of these distributions using detailed balance, $p(A \rightarrow B) = \frac{\rho_t(B)}{\rho_t(A)} p(B \rightarrow A)$, where $p(\cdot)$ is the transition kernel for Langevin evolution with timestep $dt = T/N$, and we use the fact that ρ_t is the equilibrium distribution.⁴ The pathwise probabilities forward and backwards can then be equated in the continuous-time limit,

$$\begin{aligned}
p(\bar{\Phi}^\xi) &= \lim_{N \rightarrow \infty} p(\bar{\Phi}^0 \rightarrow \bar{\Phi}^{dt}) p(\bar{\Phi}^{dt} \rightarrow \bar{\Phi}^{2dt}) \dots \frac{\rho_t(\bar{\Phi}^0) \delta(\bar{\Phi}^T - \phi)}{\rho_t(\phi)} \\
p(\Phi^\xi) &= \lim_{N \rightarrow \infty} p(\Phi^{dt} \rightarrow \Phi^0) p(\Phi^{2dt} \rightarrow \Phi^{dt}) \dots \delta(\Phi^T - \phi) \\
&= \lim_{N \rightarrow \infty} \frac{\rho_t(\Phi^0)}{\rho_t(\Phi^{dt})} p(\Phi^0 \rightarrow \Phi^{dt}) \frac{\rho_t(\Phi^{2dt})}{\rho_t(\Phi^{dt})} p(\Phi^{dt} \rightarrow \Phi^{2dt}) \dots \delta(\Phi^T - \phi) = p(\bar{\Phi}^\xi).
\end{aligned} \tag{26}$$

Under time reversal of the sampled processes, our approach is thus statistically equivalent to the method of Ref. [6], when overdamped Langevin is used. We however highlight that working with our formulation allows directly estimate the flow field for fixed input field configurations, as shown for example in Sec. 3.1. Moreover, framing the convergence problem in terms of convergence of coupled Langevin trajectories provides a new perspective that may yield progress on the challenges of evaluating flow fields in gauge theories.

5. Conclusions

This work introduced a Monte Carlo estimator for the flow field b_t that transports samples under a time-dependent action S_t . Such flow fields may be useful for both ensemble generation and noise problems in observables. By incorporating specific kernel transformations and an $SU(N)$ coupled-noise process, our method of estimating b_t was shown here to be effective for a $U(1)$ von Mises distribution and for measurements of a glueball correlation function in 2+1D $SU(N)$ Yang-Mills.

We have also demonstrated a close relation between our approach and the stochastic automatic differentiation method of Ref. [6]. While the overdamped Langevin dynamics used here generally have slower mixing time than the Hamiltonian evolution which has been recently applied in that context, the deep connection between the two approaches also suggests a Hamiltonian version of the Feynman-Kac path integral, which will be investigated in future work.

In the future, these Monte Carlo estimates of flow fields may provide useful training data for complementary machine learning or analytical approaches seeking to parameterize and learn the same flow fields. High-quality deterministic parameterizations based on this approach would further improve on the sampling and noise-reduction benefits conferred by using flow transport.

⁴Though typical discretizations of Langevin evolution do not exactly satisfy detailed balance, in the limit $N \rightarrow \infty$ corrections from violating detailed balance go to zero.

Acknowledgments

We are grateful to Ryan Abbott and Alberto Ramos for insightful discussions. MSA is supported by a Junior Fellowship at the Harvard Society of Fellows as well as the National Science Foundation under Cooperative Agreement PHY-2019786 (The NSF AI Institute for Artificial Intelligence and Fundamental Interactions⁵). This work has been made possible in part by a gift from the Chan Zuckerberg Initiative Foundation to establish the Kempner Institute for the Study of Natural and Artificial Intelligence, and was supported in part by the Schmidt Sciences AI2050 Early Career Fellowship.

References

- [1] R. Abbott, A. Botev, D. Boyda, D. C. Hackett, G. Kanwar, S. Racanière, D. J. Rezende, F. Romero-López, P. E. Shanahan, and J. M. Urban. Applications of flow models to the generation of correlated lattice QCD ensembles. *Phys. Rev. D*, 109(9):094514, 2024.
- [2] M. S. Albergo and E. Vanden-Eijnden. NETS: A Non-Equilibrium Transport Sampler. *arXiv:2410.02711*, 2024.
- [3] S. Bacchio. Novel approach for computing gradients of physical observables. *Phys. Rev. D*, 108(9):L091508, 2023.
- [4] S. Bacchio, P. Kessel, S. Schaefer, and L. Vaitl. Learning trivializing gradient flows for lattice gauge theories. *Phys. Rev. D*, 107(5):L051504, 2023.
- [5] G. G. Batrouni, G. R. Katz, A. S. Kronfeld, G. P. Lepage, B. Svetitsky, and K. G. Wilson. Langevin Simulations of Lattice Field Theories. *Phys. Rev. D*, 32:2736, 1985.
- [6] G. Catumba and A. Ramos. Stochastic automatic differentiation and the signal to noise problem. *Eur. Phys. J. C*, 85(9):1037, 2025.
- [7] G. P. Engel and S. Schaefer. Testing trivializing maps in the Hybrid Monte Carlo algorithm. *Comput. Phys. Commun.*, 182:2107–2114, 2011.
- [8] J. C. Gower and G. B. Dijkstra. *Procrustes Problems*. Oxford University Press, 2004.
- [9] M. Kac. On distributions of certain wiener functionals. *Transactions of the American Mathematical Society*, 65(1):1–13, 1949.
- [10] X. Li, T.-K. L. Wong, R. T. Chen, and D. Duvenaud. Scalable gradients for stochastic differential equations. In *International Conference on Artificial Intelligence and Statistics*, pages 3870–3882. PMLR, 2020.
- [11] M. Luscher. Trivializing maps, the Wilson flow and the HMC algorithm. *Commun. Math. Phys.*, 293:899–919, 2010.
- [12] A. Paszke et al. PyTorch: An Imperative Style, High-Performance Deep Learning Library. *arXiv:1912.01703*, 2019.

⁵<http://iaifi.org/>