

Leveraging neural control variates for enhanced precision in lattice field theory

Paulo F. Bedaque^{✉*} and Hyunwoo Oh^{✉†}

Department of Physics and Maryland Center for Fundamental Physics, University of Maryland, College Park, Maryland 20742, USA

 (Received 20 December 2023; accepted 8 May 2024; published 30 May 2024)

Results obtained with stochastic methods have an inherent uncertainty due to the finite number of samples that can be achieved in practice. In lattice QCD this problem is particularly salient in some observables like, for instance, observables involving one or more baryons and it is the main problem preventing the calculation of nuclear forces from first principles. The method of control variables has been used extensively in statistics and it amounts to computing the expectation value of the difference between the observable of interest and another observable whose average is known to be zero but is correlated with the observable of interest. Recently, control variates methods emerged as a promising solution in the context of lattice field theories. In our current study, instead of relying on an educated guess to determine the control variate, we utilize a neural network to parametrize this function. Using $1 + 1$ dimensional scalar field theory as a testbed, we demonstrate that this neural network approach yields substantial improvements. Notably, our findings indicate that the neural network ansatz is particularly effective in the strong coupling regime.

DOI: [10.1103/PhysRevD.109.094519](https://doi.org/10.1103/PhysRevD.109.094519)

I. INTRODUCTION

Monte Carlo methods have achieved enormous success in studying nonperturbative field theory phenomena numerically. Still, it faces problems in some models and/or observables where the statistical noise overwhelms the signal. That is the case of theories with a sign problem (see for instance [1–3] for reviews), including all real-time (as opposed to imaginary time) calculations and models/observables with infinite variance [4–6]. It is also the case of certain correlators in lattice QCD whose signal-to-noise ratio decreases exponentially with (imaginary) time, making the extraction of energy levels extremely challenging [7,8]. The signal-to-noise ratio of n -baryon states correlators, for instance, decays as $\sim e^{-(M-3m_\pi/2)nt}$ (M , m_π are the baryon and pion masses). Given that excited states contaminate the correlator at small t , this decaying signal-to-noise ratio at larger t is a serious problem and is, in fact, the main obstacle one faces in computing the nuclear forces from lattice QCD.

In this paper we study the use of control variates to minimize the variance of lattice observables. The basic idea is simple and well known [9,10]. The expectation value $\langle \mathcal{O} \rangle$ of an observable \mathcal{O} can be computed as $\langle \mathcal{O} - f \rangle = \langle \mathcal{O} \rangle$, where f is known to have vanishing expectation value $\langle f \rangle = 0$. While the expectation value of the observables \mathcal{O} and $\mathcal{O} - f$ are the same, their variance is not:

$$\langle (\mathcal{O} - f)^2 \rangle = \langle \mathcal{O}^2 \rangle + \langle f^2 \rangle - 2\langle f\mathcal{O} \rangle. \quad (1)$$

If a control variate f can be found that is strongly correlated with \mathcal{O} while maintaining $\langle f \rangle = 0$, the variance of $\mathcal{O} - f$ is smaller than the variance of \mathcal{O} and so it is a better estimator of $\langle \mathcal{O} \rangle$. This basic strategy has many forms depending on how one goes about finding a suitable f . For instance, recently in [11] control variates for a scalar field theory calculation was found by expressing it as a linear combination of optimal control variates for a free field theory. It has also been argued that control variates can be a possible solution to remove the sign problem exactly [12,13].

In this work, we find control variates for lattice field theory observables using machine learning techniques, namely, a very general function f is parametrized by a feed-forward neural network in such a way that the condition $\langle f \rangle = 0$ is automatically satisfied. Then, standard minimization techniques are used to minimize the variance. In the language of machine learning, the variance becomes the “cost function” whose minimization can be thought of as a form of unsupervised learning. The fact that general

*bedaque@umd.edu

†hyunwooh@umd.edu

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

neural networks are universal function approximators (in the sense of being able to approximate any function with a sufficiently large network) suggests we are looking for a control variate within a very large class of functions. A similar approach has been used in Monte Carlo integration in small dimensional space [14–17].

II. METHOD

The purpose of this section is to review the basic knowledge of the control variates and to explain the way to parametrize the control variates through neural networks.

A. Control variates

Let us denote by $\langle \cdot \rangle$ the average with respect to the Boltzmann factor:

$$\langle \mathcal{O}(\phi) \rangle = \frac{1}{Z} \int D\phi \mathcal{O}(\phi) e^{-S[\phi]}, \quad (2)$$

with $Z = \int D\phi e^{-S[\phi]}$. If $\langle f \rangle = 0$, \mathcal{O} and $\tilde{\mathcal{O}} = \mathcal{O} - f$ will have the same expectation value:

$$\langle \tilde{\mathcal{O}} \rangle = \langle \mathcal{O} - f \rangle = \langle \mathcal{O} \rangle. \quad (3)$$

However, their variances differ:

$$\begin{aligned} \text{Var}(\tilde{\mathcal{O}}) &= \langle (\mathcal{O} - f)^2 \rangle - \langle \mathcal{O} - f \rangle^2 \\ &= \text{Var}(\mathcal{O}) + \langle f^2 \rangle - 2\langle \mathcal{O}f \rangle. \end{aligned} \quad (4)$$

The function f is called a control variate. The goal is then to find a function f that is highly correlated with the observable \mathcal{O} in order to minimize the variance of $\tilde{\mathcal{O}}$.

We write our control variate candidate f as

$$f(\phi) = \sum_x \left(\frac{\partial g[\phi]_x}{\partial \phi_x} - g[\phi]_x \frac{\partial S}{\partial \phi_x} \right), \quad (5)$$

where x indexes the sites on the spacetime lattice and $g[\phi]: \mathbb{R}^V \rightarrow \mathbb{R}^V$ (V is the spacetime volume) has yet to be defined. For any $g_x[\phi]$, where $g = (g_1, \dots, g_V)$, integration by parts shows that

$$\left\langle \frac{\partial g_x}{\partial \phi_x} \right\rangle = \left\langle g_x \frac{\partial S}{\partial \phi_x} \right\rangle, \quad (6)$$

and therefore, $\langle f \rangle = 0$ by construction.

B. Machine learning

While Eq. (5) does not define the most general function $f[\phi]$, we aim at having a universal representation of $g[\phi]$. The only constraint we will impose is space-time translational symmetry

$$f(T_y[\phi]) = f(\phi), \quad (7)$$

where the translation operator T_y displaces a field configuration by y : $T_y[\phi_x] = \phi_{x+y}$. This can be achieved if g is covariant:

$$g[T_y[\phi]]_x = g[\phi]_{x+y}. \quad (8)$$

We can impose translational invariance by defining a function $g_0: \mathbb{R}^V \rightarrow \mathbb{R}$ from which we define a $g[\phi]_x$ by

$$g[\phi]_x = g_0(T_x[\phi]). \quad (9)$$

It can be easily shown that the control variate f is translational invariant:

$$\begin{aligned} f(T_y[\phi]) &= \sum_x \left(\frac{\partial g[T_y[\phi]]_x}{\partial \phi_{x+y}} - g[T_y[\phi]]_x \frac{\partial S(T_y[\phi])}{\partial \phi_{x+y}} \right) \\ &= \sum_x \left(\frac{\partial g_0(T_{x+y}[\phi])}{\partial \phi_{x+y}} - g_0(T_{x+y}[\phi]) \frac{\partial S(T_{x+y}[\phi])}{\partial \phi_{x+y}} \right) \\ &= \sum_{x'} \left(\frac{\partial g_0(T_{x'}[\phi])}{\partial \phi_{x'}} - g_0(T_{x'}[\phi]) \frac{\partial S(T_{x'}[\phi])}{\partial \phi_{x'}} \right) \\ &= f(\phi). \end{aligned} \quad (10)$$

We will define $g_0[\phi]$ by a fully connected feed-forward neural network with V inputs and only one output. For theories with parity symmetry $\phi \rightarrow -\phi$, as the model we will consider in the next section, we choose the activation function $\sigma(x) = \text{arcsinh}(x)$ which is odd. In addition, we remove the bias term in the linear transformation so that the network is an odd function. In this way, we can reduce the number of parameters, which makes training faster.

The ideal cost function for the training of g_0 is the variance of $\mathcal{O} - f$. We will estimate the variance of $\tilde{\mathcal{O}}$ by the variance of a sample $\{\phi^a\}$, $a = 1, \dots, \mathcal{N}$, of \mathcal{N} field configurations:

$$L(w) = \frac{1}{\mathcal{N}} \sum_{a=1}^{\mathcal{N}} \left(\mathcal{O}(\phi^a) - f_w(\phi^a) \right)^2, \quad (11)$$

where f_w , given by Eqs. (5) and (9), depends on the parameters w of the network defining g_0 . Note that the term from $\langle \mathcal{O} - f \rangle^2$ can be omitted since Eq. (3) does not affect the new variance.

Since, in practice, \mathcal{N} will be small (typically of the order of hundreds in lattice QCD) while the neural network can easily contain a much larger number of parameters, the risk of overfitting exists. In that case, even though $\langle f \rangle = 0$ by construction, $\tilde{\mathcal{O}}$ is minimized by having $f(\phi^a) \approx \mathcal{O}(\phi^a)$ for every ϕ^a in the sample, leading to the sample average $\bar{f} \approx \bar{\mathcal{O}}$. This is not a good approximation to the ideal control variate f . This problem was recognized in [14] where the

authors suggest to minimize instead the quantity:

$$L(w) = \frac{1}{\mathcal{N}} \sum_{a=1}^{\mathcal{N}} \left(\mathcal{O}(\phi^a) - f_w(\phi^a) - \mu_0 \right)^2. \quad (12)$$

where μ_0 is the sample average of the observable $\mu_0 = \bar{\mathcal{O}}$. In this case, overfitting leads to $f_w(\phi^a) \approx \mathcal{O}(\phi^a) - \mu_0$ for every ϕ^a in the sample and $\bar{f}_w \approx 0$, a much better approximation to the ideal control variate. In fact, [14,18] suggested to use, instead of a fixed value μ_0 , a variable μ that is updated during the training process just like the parameters w of f_w . We verified that in our examples this last procedure was significantly better and all results presented below are obtained with this last procedure.¹ Finally, to further avoid overfitting we applied the L_2 regularization

$$L(w, \mu) = \frac{1}{\mathcal{N}} \sum_{a=1}^{\mathcal{N}} (\mathcal{O}(\phi^a) - f_w(\phi^a) - \mu)^2 + \delta \sum w^2, \quad (13)$$

where $\sum w^2$ is the sum of the squares of all neural network parameters and δ is the regularization parameter.

III. RESULTS

In order to test our approach we use a scalar field theory in $1 + 1$ dimensions. Due to its simplicity, the same model has been used as a testbed for other approaches to the signal-to-noise problem and sign problem [11,19–21], so a direct comparison is possible.

We will consider a $L_0 \times L_1$ lattice (so $V = L_0 L_1$). We work in units where the lattice spacing is 1. Then the action on the lattice can be written as

$$S = \sum_{x,\mu} \left(\frac{(\phi_{x+\mu} - \phi_x)^2}{2} + \frac{m^2}{2} \phi_x^2 + \frac{\lambda}{4!} \phi_x^4 \right), \quad (14)$$

where $\mu = 0, 1$ indexes the directions on the lattice and $x = (x_0, x_1)$. The observable we choose to consider is the two-point correlator at momentum $p = 0$:

$$\mathcal{O}(t) = \frac{1}{L_0} \sum_{y_0} \left[\left(\frac{1}{L_1} \sum_{x_1} \phi_{y_0+t, x_1} \right) \left(\frac{1}{L_1} \sum_{y_1} \phi_{y_0, y_1} \right) \right]. \quad (15)$$

Table I summarizes the bare parameters and neural network parameters used in this paper. We use two values of the coupling λ , couplings 0.5 and 24.0 in order to explore both the weak and strong regimes. The number of hidden layers is chosen from 1 to 5, and the number of neurons for hidden layers varies from 1 to 32. The training of neural

TABLE I. Bare parameters of scalar field and neural network parameters in Sec. III.

	Lattice	m^2	λ	Layers	Neurons	δ
Fig. 1	20×20	0.1	0.5, 24.0	5, 5	4, 4	0
Fig. 2	40×40	0.1	0.5, 24.0	1, 5	1, 32	0.05, 0.0005
Fig. 3	40×10	0.01	0.1	1	16	0.1

networks is performed with the usual stochastic gradient method and the *ADAM* optimizer [22] implemented with the help of the *JAX* [23] and *Flax* [24] libraries.

The choice of learning rate η is critical to the success of the training. If the learning rate is too small, the training might be stuck at local minima. A too large learning rate leads to a process that misses minima of $L(w, \mu)$ altogether. After a lengthy trial-and-error process we find that it is efficient to use an exponentially decreasing learning rate in the beginning of the training and a constant one after that:

$$\eta(n) = \begin{cases} 10^{-3} \times 0.99^{\frac{n}{1000}}, & \text{if } n \leq n_0 \\ 10^{-6}, & \text{if } n \geq n_0 \end{cases}, \quad (16)$$

where n is the training step and n_0 is in the range $(6.8 \times 10^5, 6.9 \times 10^5)$.

Since the gradient of $L(w, \mu)$ is estimated stochastically, we need to decide the number n of samples used. The stochastic error in every step in the gradient descent scales as $\sim \eta / \sqrt{n}$ while the computational cost of following the gradient descent path by unit of “time” scales as $\sim n / \eta$. Therefore, the optimal choice is $n = 1$, the value we used in all our calculations. We have tried other choices, $n \neq 1$, and confirmed $n = 1$ works better empirically.

In Fig. 1 we show the training history on 20×20 lattice, both at small and large couplings and with networks with and without hidden layers. The observable we aim at improving was a mid-lattice correlator, that is, Eq. (15) with $L_0 = L_1 = 20, t = L_0/2 = 10$. Figure 1 shows the ratio of standard deviations (the improvement in the uncertainty) as a function of the training step for small and large couplings. We used 10^4 and 10^3 fully decorrelated field configurations to train the network and 10^3 samples to estimate the variance. An epoch is defined as $\mathcal{N} = 10^4$ or 10^3 gradient descent steps, that is, one step with each training configuration. The result shows that the variance can be greatly improved, but it also implies that the result of training depends on the size of the training set. These results can be compared to the ones in [11] where a different, more direct method was used to obtain control variates for the same theory. Their method is exact for a free theory and it performs better than at weak coupling. Neural control variates also work better as weak coupling but outperforms the method in [11] at strong coupling. As Fig. 1 shows, a linear transformation (network with no hidden layers) performs well at small coupling since a

¹Notice that $\langle \mathcal{O} \rangle$ is still estimated by $\overline{(\mathcal{O} - f_w)}$, not $\overline{(\mathcal{O} - f_w - \mu)}$.

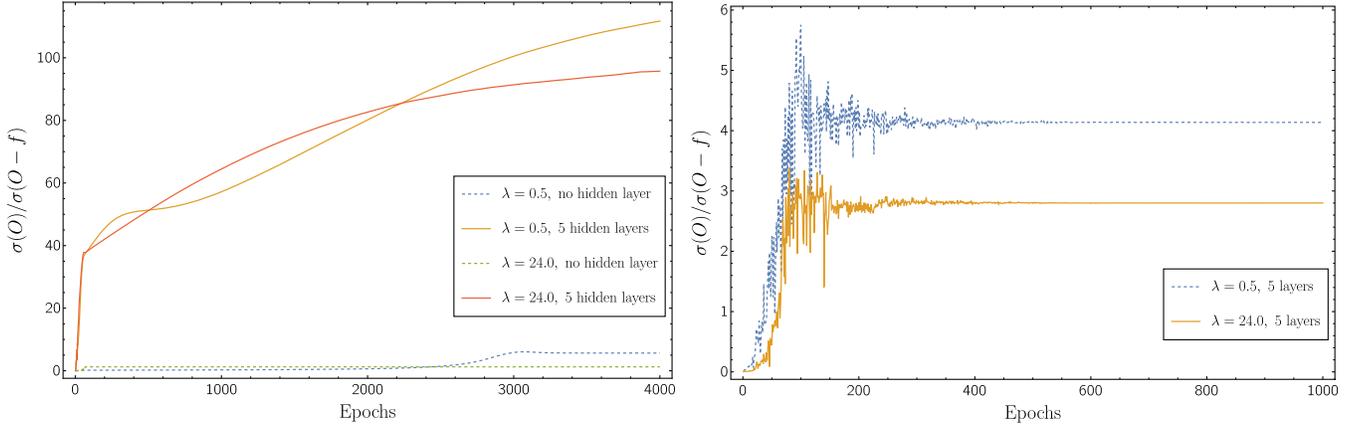


FIG. 1. Training histories of the small and large couplings on 20×20 lattice. The figure shows the improvement of standard deviation using control variates with respect to the raw one ($\sigma_{\text{Raw}}/\sigma_{\text{CV}}$). The dashed line represents the zero hidden layer result (linear transformation), and the solid line represents the result with hidden layers. Networks for small and large couplings have 5 hidden layers and each has 4 neurons. For left and right panels, 10^4 and 10^3 samples are reserved for training the network respectively, and 10^3 samples are used to estimate the variance.

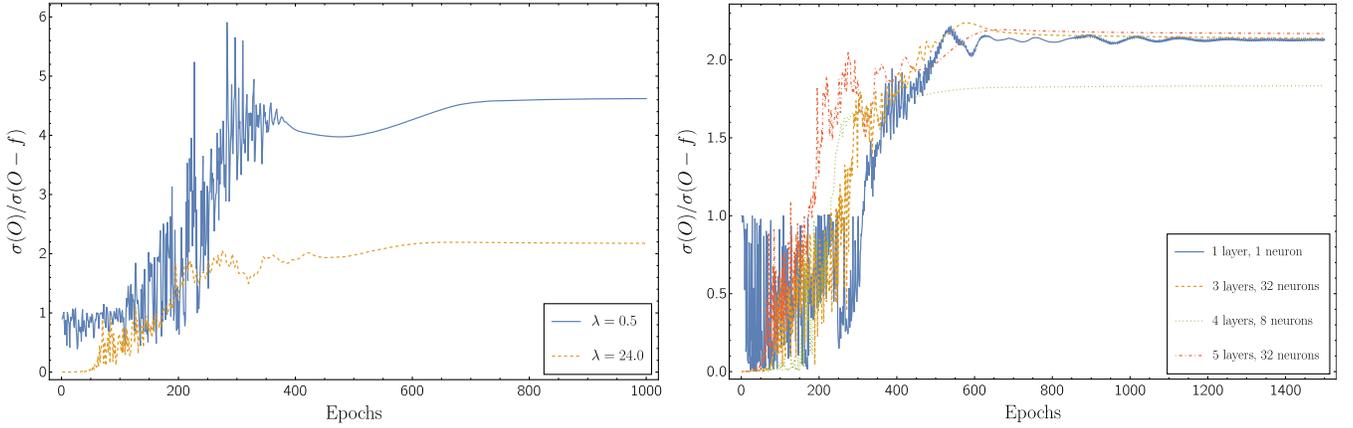


FIG. 2. Training histories of control variates at $t = L_0/2 = 20$ on 40×40 lattice. The left plot shows the improvement of standard deviation with hidden layers. The right panel displays the training histories of different depths of networks with the large coupling, $\lambda = 24.0$. 10^3 samples are used to train the neural control variates and 10^3 samples are employed to estimate the standard deviation.

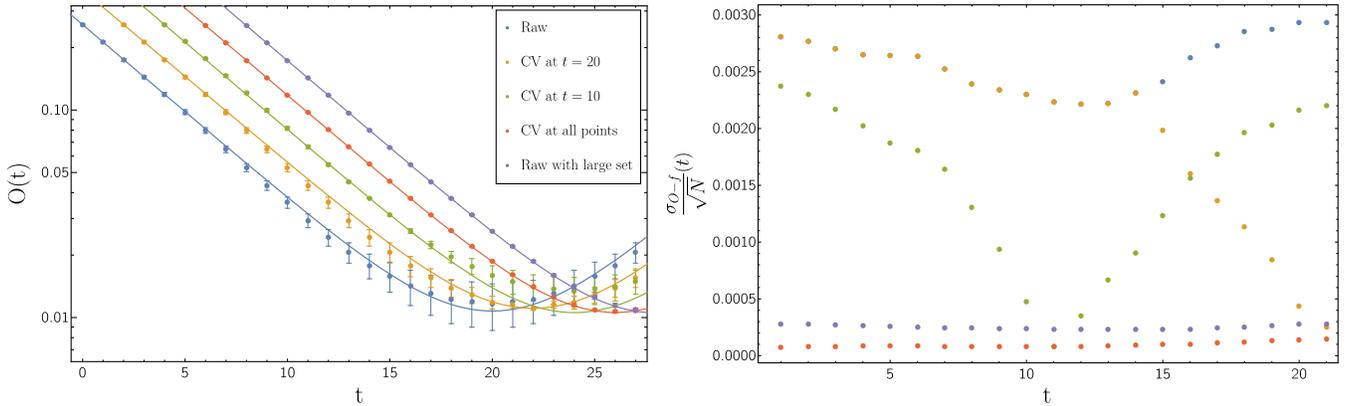


FIG. 3. Correlation functions with $m^2 = 0.01$ and $\lambda = 0.1$ on 40×10 lattice. The raw result and the result with control variates are shown. 2×10^3 samples are used in total and for the control variate result, 10^3 samples are used for training and the whole samples are used for estimating observables. For the raw result with large sets, the correlators are calculated with 2×10^5 samples. The left plot shows the correlation functions with their fitting. Results are shifted horizontally for better visualization. The right plot displays the errors of the correlators in the left plot.

control variate linear on the fields is sufficient for a free theory. At strong coupling, the variance is not reduced much with just a linear combination ansatz, which is shown as the dashed line in the figure. However, a better control variate is found by introducing more nonlinear terms through increasing hidden layers of neural networks.

The same method, applied to larger 40×40 lattices, shows a smaller improvement than the 20×20 lattice case, both at weak and strong coupling (see Fig. 2). It is notable that networks of different sizes lead to similar results and, in some cases, larger networks do not lead to a larger improvement in the variance. This is in contradiction to the theoretical expectation since, as one increases the number of hidden layers or the number of neurons at each hidden layer, the representability of the neural network increases and, therefore, the larger network should have a better performance than the smaller one. Our conclusion is that there is room for improvement in the training of large networks which we have not yet been able to accomplish. Future work will concentrate on that.

Reducing the uncertainty of one time slice of a propagator, even by a large factor, does not necessarily imply a reduction of the uncertainties of parameters extracted from it, like the value of masses. To investigate this question we applied our method to anisotropic 40×10 lattices with couplings fairly close to the continuum limit. Ideally, one would use a different control variate at each time slice but this procedure is too expensive. Instead, we found neural control variates optimizing the uncertainty at the $t = L_0/2 = 20$ or $t = L_0/4 = 10$ time slices. The correlators are shown in Fig. 3 and the results of the fits of the correlator to the form

$$C(t) = A(e^{-mt} + e^{-m(L_0-t)}) \quad (17)$$

are summarized in Table II. Note that the correlation between different points of the correlator is considered when fittings are implemented. The regularization strength δ is chosen as 0.1. The reduction of the uncertainty due to the use of a control variate is about a factor ≈ 20 at the time slice $t = L_0/2 = 20$ (or $t = L_0/4 = 10$). Since the data at different time slices are correlated, we find that reducing the variance of one time slice also reduces the variance of nearby time slices, albeit by a smaller amount. In the example we discussed the uncertainty in the mass estimate

TABLE II. The fitted values of correlators in Fig. 3.

	A	m	χ^2/dof
Raw, $\mathcal{N} = 2 \times 10^3$	0.000112(18)	0.1935(40)	0.55
CV at $t = 20$	0.0001191(51)	0.1920(13)	0.29
CV at $t = 10$	0.0001083(44)	0.1944(13)	0.55
CV at all points	0.0001096(7)	0.1938(2)	0.91
Raw, $\mathcal{N} = 2 \times 10^5$	0.0001088(18)	0.1940(4)	0.34

is reduced by a factor ≈ 3 corresponding to a configuration set ~ 9 times larger.

While we improved the error around 20 times at one point, the error of the fit giving the mass estimate is improved by a smaller factor (~ 3). A better mass estimate is obtained if a different control variate is used at each time slice. In order to reduce the computational cost of training we use as a starting point of the training at one time slice, the final result of a previous time slice (transfer training²). In doing so, we obtain around 20 times improvement of the standard deviation at every point and the uncertainty of the mass estimate is also improved with the same amount. We include in Table II the results of fitting the correlator obtained with much higher statistics in order to verify that the reduced variance result obtained with neural control variates and low statistics is consistent with it.

IV. DISCUSSION

We showed how control variates parametrized by neural networks reduce the variance in lattice field theories. By using the simple example of a $1 + 1$ dimensional ϕ^4 theory we established the feasibility of the method and learned some qualitative features that can be summarized as follows:

- (i) Reductions of variance by tens or hundreds are feasible.
- (ii) It is essential to design the neural network to automatically incorporate the symmetries of the model as those symmetries are difficult to “learn” from a sample of configurations of a realistic size (hundreds to thousands of configurations). Although we have not explored them yet, other techniques can plausibly work better.
- (iii) Techniques to avoid overfitting, like the ones we used, are essential for good results.
- (iv) More direct methods, like in [11], are more efficient at small coupling but neural networks tend to win out at larger couplings.

A natural question is whether how control variates results compare to more standard methods. After all, the variance of an observable can be reduced by just collecting a larger number \mathcal{N} of configurations. This process is slow as the uncertainties scale as $\sim 1/\sqrt{\mathcal{N}}$. On the other hand, the training of the neural control variate is computationally expensive, even if it needs to be performed only once per observable. The “break-even” point where our method wins out over the brute force increase in statistics depends on the cost of collecting a new independent configuration. For the scalar theory we studied, the computational cost of the configuration is very small, and for the precision we achieved, it is clearly more efficient simply to increase the statistics. In other theories, specially those with dynamical

²We thank the anonymous referee for insisting that we tried this method.

fermions, the configurations are very expensive and it makes sense to use substantial computing power extracting as much information as possible from them. Thus, the answer to whether (neural) control variates are more efficient or not has to be studied in every model separately.

Among the direction for future work, the most pressing is the extension of neural control variates to gauge theories. There are conceptual issues to be solved. Given our observation that it is important to impose the symmetries of the model on the neural network, a method to impose gauge invariance seems essential. Also, it is unclear what

an analog to Eq. (5) would be. Other issues are of a more practical nature and may be more important. For instance, finding more efficient ways of training neural networks will be crucial to extend the method to realistic, four-dimensional theories.

ACKNOWLEDGMENTS

This work was supported in part by the U.S. Department of Energy, Office of Nuclear Physics under Awards No. DE-SC0021143, and No. DE-FG02-93ER40762.

-
- [1] Z.-X. Li and H. Yao, *Annu. Rev. Condens. Matter Phys.* **10**, 337 (2019).
 - [2] P. de Forcrand, *Proc. Sci. LAT2009* (2009) 010 [arXiv:1005.0539].
 - [3] A. Alexandru, G. Basar, P.F. Bedaque, and N.C. Warrington, *Rev. Mod. Phys.* **94**, 015006 (2022).
 - [4] C. Yunus and W. Detmold, *Phys. Rev. D* **106**, 094506 (2022).
 - [5] H. Shi and S. Zhang, *Phys. Rev. E* **93**, 033303 (2016).
 - [6] A. Alexandru, P.F. Bedaque, A. Carosso, and H. Oh, *Phys. Rev. D* **107**, 094502 (2023).
 - [7] G. Parisi, *Phys. Rep.* **103**, 203 (1984).
 - [8] G.P. Lepage, in *Theoretical Advanced Study Institute in Elementary Particle Physics* (World Scientific, Singapore, 1990).
 - [9] R. Assaraf and M. Caffarel, *Phys. Rev. Lett.* **83**, 4682 (1999).
 - [10] A. Mira, R. Solgi, and D. Imperato, *Stat. Comput.* **23**, 653 (2012).
 - [11] T. Bhattacharya, S. Lawrence, and J.-S. Yoo, *Phys. Rev. D* **109**, L031505 (2024).
 - [12] S. Lawrence, *Phys. Rev. D* **102**, 094504 (2020).
 - [13] S. Lawrence and Y. Yamauchi, *Phys. Rev. D* **107**, 114505 (2023).
 - [14] R. Wan, M. Zhong, H. Xiong, and Z. Zhu, in *Machine Learning and Knowledge Discovery in Databases*, edited by U. Brefeld, E. Fromont, A. Hotho, A. Knobbe, M. Maathuis, and C. Robardet (Springer International Publishing, Cham, 2020), pp. 533–547.
 - [15] T. Müller, F. Rousselle, A. Keller, and J. Novák, *ACM Trans. Graph.* **39**, 1 (2020).
 - [16] Z. Sun, C.J. Oates, and F.-X. Briol, arXiv:2303.04756.
 - [17] K. Ott, M. Tiemann, P. Hennig, and F.-X. Briol, in *Proceedings of the Thirty-Ninth Conference on Uncertainty in Artificial Intelligence*, UAI '23 (JMLR.org, 2023).
 - [18] L.F. South, C.J. Oates, A. Mira, and C. Drovandi, *Bayesian Anal.* **18**, 865 (2023).
 - [19] W. Detmold, G. Kanwar, and M.L. Wagman, *Phys. Rev. D* **98**, 074511 (2018).
 - [20] W. Detmold, G. Kanwar, M.L. Wagman, and N.C. Warrington, *Phys. Rev. D* **102**, 014514 (2020).
 - [21] S. Lawrence, H. Oh, and Y. Yamauchi, *Phys. Rev. D* **106**, 114503 (2022).
 - [22] D. Kingma and J. Ba, in *International Conference on Learning Representations (ICLR)* (ICLR, San Diego, CA, USA, 2015).
 - [23] J. Bradbury, R. Frostig, P. Hawkins, M.J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, JAX: Composable transformations of Python + NumPy programs (2018), <http://github.com/google/jax>.
 - [24] J. Heek, A. Levskaia, A. Oliver, M. Ritter, B. Rondepierre, A. Steiner, and M. van Zee, Flax: A neural network library and ecosystem for JAX (2023), <http://github.com/google/flax>.