

Control variates for lattice field theory

Tanmoy Bhattacharya^{1,*}, Scott Lawrence^{2,†} and Jun-Sik Yoo^{1,‡}¹Group T-2, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA²Department of Physics, University of Colorado, Boulder, Colorado 80309, USA

(Received 8 August 2023; accepted 30 January 2024; published 23 February 2024)

In most lattice field theories, correlators are plagued by a signal-to-noise problem of exponential difficulty in the time separation. We propose a method for improving the signal-to-noise ratio, in which control variates are systematically constructed from lattice Schwinger-Dyson relations. The method is demonstrated on various two-dimensional lattices in scalar field theory, and a strategy for scaling to larger systems is explored.

DOI: [10.1103/PhysRevD.109.L031505](https://doi.org/10.1103/PhysRevD.109.L031505)

With the advent of supercomputing, numerical lattice field theory has become the method of choice for first-principles calculations beyond perturbation theory. In particular, lattice QCD calculations provide the most reliable estimates for quantities of interest in hadronic physics [1]. It is a systematically improvable scheme in which the only uncertainties stem from the extrapolations to the chiral-continuum-infinite-volume limit and from statistical noise. For situations where the measure of the path integral—determined by the Euclidean action of the theory—is real and positive, it has become routine to carry out high-statistics, large-volume simulations of lattice QCD with physical quark masses at a number of lattice spacings, leading to reliable extrapolations. This is especially true for pseudoscalar meson correlators, whose signal-to-noise ratio stays constant at long Euclidean time separation.

Extending these successes to precise calculations with other mesons and baryons is, however, meeting barriers. In particular, since the interpolating fields used in an interacting field theory couple to multiple asymptotic states, calculations from short separations in Euclidean time have irreducible systematic uncertainties due to contributions involving excited states. For most of these states, at the Euclidean distances needed to isolate the lowest state of interest, the signal-to-noise degrades and the statistical noise becomes insurmountable with today's resources [2]. A canonical example of the signal-to-noise problem is in the baryon-baryon correlator of QCD. An argument of

Parisi and Lepage [3,4] shows that the signal and noise both fall exponentially, but the exponential fall of the signal is faster than that of the noise, and therefore the signal-to-noise ratio falls exponentially as well. The situation for gluonic correlators is similar: the accuracy with which we can evaluate the expectation values of large Wilson loops or the static-quark potential is severely limited by noise. Since the most accurate determination of the lattice scale is from such gluonic observables [5], this is a serious problem for any precise lattice calculations [6][as an example]. In this work we will examine the simpler situation of a scalar field theory in two spacetime dimensions, where the signal falls exponentially while the noise does not; though the methods developed will be generic and expected to be generalizable to other theories.

Previous work has addressed this problem and developed techniques for variance reduction. Of particular interest is a class of techniques in which the observable of interest can be approximated by an observable that is less expensive to evaluate but whose fluctuations are highly correlated [7,8]. The statistical fluctuations in both the approximant and the difference can then be evaluated with smaller effort, leading to an improvement in the signal-to-noise performance. A different class of methods starts from the observation that the signal-to-noise problem is closely connected to the sign problem, and the same methods have often been applied to both. In particular, the method of contour deformations used in [9,10] has its history in approaches to the sign problem; see [11] for a review. It relies on the fact that the path integrals defining the expectation values of the observables are analytic functions, as such a continuation of the domain of the path integral to complex field space can keep the value of the integral unchanged. The statistical variance, however, is not analytic, and it can be reduced.

In this work, we take an alternative course to variance reduction that merges the strengths of the two methods. It is similar to the correlated approximant method in that we

*tanmoy@lanl.gov

†scott.lawrence-1@colorado.edu

‡junsik@lanl.gov

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). 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.

find a correction to the integrand whose fluctuations are correlated with the observable. But, like the contour deformation method, the subtraction is of a term that can be guaranteed to integrate to zero. This general method has previously been explored for the fermion sign problem, via perturbative expansion [12] and machine learning [13], as well as for spin systems [14,15].

Thus, the core of the approach is the construction of a *control variate*: a function of the fields which is correlated with the observable of interest, but can be proven to have vanishing expectation value. Naming the control variate f , we define an improved observable $\tilde{\mathcal{O}} \equiv \mathcal{O} - f$. First note that since $\langle f \rangle = 0$, this new observable obeys $\langle \tilde{\mathcal{O}} \rangle = \langle \mathcal{O} \rangle$. The variance of $\tilde{\mathcal{O}}$, however, has changed:

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

It is apparent that, when the correlation $\langle \mathcal{O}f \rangle$ is large and positive, the variance of the improved observable is substantially reduced.

It remains to minimize the variance $\langle \tilde{\mathcal{O}}^2 \rangle$ over a set of functions f that can be guaranteed to integrate to 0. The main approach we put forward is to obtain a large basis $\{F_i\}$ of functions such that $\langle F_i \rangle = 0$, and then optimize the coefficients c_i that define the control variate $f \equiv \sum_i c_i F_i$. This optimization can be performed efficiently as follows. Measure the covariance matrix of the F_i , and the cross-correlations between F_i and the target observable $\mathcal{O}$, defining

$$M_{ij} = \langle F_i F_j \rangle \quad \text{and} \quad v_i = \langle F_i \mathcal{O} \rangle. \quad (2)$$

The optimal control variate is defined to be the linear combination of F_i that minimizes the variance of Eq. (1). Assuming that M and v have been measured accurately, this is readily shown to be given by coefficients

$$c = M^{-1}v. \quad (3)$$

To obtain a basis $\{F_i\}$, we note that for any function g of fields the integral of ∂g (where “ ∂ ” indicates any derivative with respect to the fields) reduces to boundary terms, and therefore 0 under mild assumptions about the asymptotic behavior.¹

Having obtained these coefficients and therefore the definition of the control variate f , we may compute the expectation value of $\tilde{\mathcal{O}}$ (and therefore $\langle \mathcal{O} \rangle$) by standard Monte Carlo methods. It is important, to avoid introducing a bias, that this final computation of the expectation value is

¹For the example of scalar field theory used in this paper, g can be constructed from polynomials in the fields as explained below; in the context of compact gauge theories there is no boundary and therefore no restriction on g .

performed over a set of samples uncorrelated with those used to determine the coefficients c in the first place.

Let us see how this method works in scalar field theory in two spacetime dimensions. This model is defined by the action

$$S = \sum_{\langle r, r' \rangle} \frac{(\phi(r) - \phi(r'))^2}{2} + \sum_r \left[\frac{m^2}{2} \phi(r)^2 + \frac{\lambda}{4!} \phi(r)^4 \right], \quad (4)$$

where m^2 (not necessarily non-negative) defines the bare mass and λ the bare coupling constant. The first summation is taken over pairs of neighboring lattice sites. Lattice units of $a = 1$ are assumed throughout, and we note that the action has a $\mathbb{Z}_2$ symmetry $\phi \rightarrow -\phi$.

The object of primary interest for us is the correlation function in the $\mathbb{Z}_2$ -vector channel, defined by

$$C(\tau) = \sum_x \langle \phi(\tau, x) \phi(0, 0) \rangle. \quad (5)$$

As examined in [9], this correlator has a signal-to-noise problem that grows exponentially at large values of τ . We will reduce this signal-to-noise problem by identifying a set of observables that have vanishing expectation values, and constructing a linear combination that is maximally covariant with $\phi(\tau, x) \phi(0, 0)$.

To obtain the basis vectors $\{F_i\}$, we note that for any polynomial $P(\phi)$ of the fields and any choice of site x , we have

$$\left\langle \frac{\partial P(\phi)}{\partial \phi_x} \right\rangle = \left\langle P(\phi) \frac{\partial S(\phi)}{\partial \phi_x} \right\rangle. \quad (6)$$

As mentioned above, this identity follows straightforwardly from the fact that the integral of a total derivative reduces to a sum of boundary terms, all of which vanish as long as $P(\phi)$ is subexponential. This identity is the lattice version of the Schwinger-Dyson equations described in [16].²

We will restrict ourselves to low-order polynomials. Due to the $\mathbb{Z}_2$ symmetry of the action, Schwinger-Dyson relations derived from even-order polynomials are uninteresting, relating only $\mathbb{Z}_2$ -odd expectation values which are known to vanish by symmetry. First- and third-order monomials at site y yield the following relations for all pairs of sites x, y :

$$\delta_{x,y} = \left\langle \phi_y \frac{\partial S(\phi)}{\partial \phi_x} \right\rangle, \quad (7a)$$

$$3\langle \phi_y^2 \rangle \delta_{x,y} = \left\langle \phi_y^3 \frac{\partial S(\phi)}{\partial \phi_x} \right\rangle. \quad (7b)$$

²More elaborate identities can be derived; we have deliberately taken the simplest construction.

Note that, in the case of $\lambda = 0$, the lattice Schwinger-Dyson relations Eq. (7a) are closed. In particular, they entirely determine the correlation function $C(\tau)$ being studied. A minimization of the variance in this case would be tantamount to solving the theory exactly, reducing the statistical noise to zero. At any nonzero coupling, this is no longer true. Even the full infinite set of Schwinger-Dyson relations is no longer sufficient to uniquely fix expectation values. This can be seen by noticing that Eq. (6) is satisfied on any complex integration contour, as long as the integral of e^{-S} along that contour converges and is nonzero. Multiple nonhomologous integration contours are available in the presence of a ϕ^4 interaction, and therefore the solution to the Schwinger-Dyson relations is not unique.

Concretely, our procedure for constructing a control variate from the Schwinger-Dyson relations Eqs. (7a) and (7b) is as follows. Supposed we have K samples on an $L \times L$ lattice, and the observable of interest is $\mathcal{O}$. We divide the samples into a training and an evaluation set; here we choose them to be of equal size $\frac{K}{2}$. The training set is used in determining the coefficients defining the control variate, and the evaluation set for computing expectation values.³ From Eqs. (7a) and (7b), we define the following zero-mean functions, a linear combination of which will become control variates:

$$F_{xy}^{(1)} = \delta_{xy} - \phi_x \frac{\partial S}{\partial \phi_y}, \quad (8a)$$

$$F_{xy}^{(3)} = 3\delta_{xy}\phi_x^2 - \phi_x^3 \frac{\partial S}{\partial \phi_y}. \quad (8b)$$

The first set of observables will be referred to throughout as first-order and the second set as the third-order candidates. For convenience, we will denote all the candidate basis functions being used in the optimization as F_i for $i = 1, \dots, B$.

We now compute, over the training set, the covariance matrix M of the candidates and the correlation v of the candidates with the target observable, as defined by Eq. (2). With this done, Eq. (3) defines the coefficients of the optimal coefficients c_i defining the control variate, and the variance-reduced unbiased observable

$$\tilde{\mathcal{O}} \equiv \mathcal{O} - \sum_i c_i F_i. \quad (9)$$

Note that this full procedure must be repeated, yielding a separate control variate, for each observable of interest. In particular, when computing a correlator $\langle \phi(z)\phi(0) \rangle$, a

³The purpose of this division into two sets is to avoid biasing the final measurement; if the full set of samples is used for both training and evaluation, final measurements are typically biased towards 0.




FIG. 1. Demonstration of variance reduction via Schwinger-Dyson relations on an 8×8 lattice at $m^2 = 0.1$ and $\lambda = 0.5$. The variance-reduced data is compared to raw correlators with 100 times as many samples as a check of the correctness of the procedure. Note that the same set of samples is reused between different values of τ , and therefore the displayed one-sigma uncertainties are correlated.

separate control variate must be optimized for each possible value of z .

In practice the above procedure has high memory consumption. To make it more feasible on large datasets, we make two simplifications. First, only translationally invariant sums of F_{xy} need to be used,⁴ reducing the size of the basis from $\sim V^2$ to $\sim V$. Second, before computing the correlations M and v , we group the samples, still maintaining a sufficiently large number of blocks—specifically, eight times the size of the basis being used.

Figure 1 shows the efficacy of this approach to variance reduction. On an 8×8 lattice with bare parameters $m^2 = 0.1$ and $\lambda = 0.5$, we obtain 10^5 decorrelated samples by Markov chain Monte Carlo (MCMC) techniques. The correlator determined from the first 10^3 samples is plotted with and without the variance reduction procedure, using only the first-order candidates $F^{(1)}$. To demonstrate that the correlator is accurate, the results without variance reduction from the full set of 10^5 samples is also plotted. The variance-reduced correlator is in agreement with the high-statistics correlator, indicating the correctness of the calculation.

To illustrate the need for higher Schwinger-Dyson equations as we go to more correlated systems, we show in Fig. 2 the same procedure being performed on a larger, 24×24 lattice, with bare parameters $m^2 = 0$ and $\lambda = 2$. Variance reduction is performed with $F^{(1)}$ alone and with the combination of $F^{(1)}$ and $F^{(3)}$. At this strong coupling, the use of the additional third-order control variates make a substantial difference at large time separations.

⁴Recall that the correlator is translationally invariant, being averaged over all possible sources for any fixed displacement. Were the observable not translationally invariant, this reduction in the space of candidate control variates would not be possible.




FIG. 2. Scalar correlator on a 24×24 lattice, with bare parameters $m^2 = 0$ and $\lambda = 2$. The raw result—for which the correlator cannot be distinguished from zero at large time separations—is compared to correlators obtained with first-order control variates alone, and those using both first- and third-order control variates. A total of 10^4 samples were used, with half being reserved for optimization when a control variate is used.

This is not surprising. As mentioned above, the first-order Schwinger-Dyson relations of Eq. (7a) give perfect variance elimination at zero coupling, and continue to perform well at weak couplings. At larger couplings, the remaining statistical fluctuations obtained with these relations become larger, and inclusion of higher-order relations [in this work, only Eq. (7b) was investigated] becomes more important. Figure 3 shows this in more detail: the estimates obtained as a function of lattice coupling λ (on a 20×20 lattice with bare mass parameters $m^2 = 0.1$) clearly show the need for the third-order term as the coupling gets stronger. At the largest couplings, using the third-order term reduces the number of samples required to get the same size of statistical error by a factor of $\sim 10^2$.

When optimizing the control variate, it is important to have more samples available than degrees of freedom in the control variate. Otherwise, the optimization results in an overfit to the training data, and the control variate does not generalize to the samples used to compute the expectation value. The control variates discussed so far have V or $2V$ degrees of freedom, and therefore at least that many samples must be used in the optimization procedure. On the other hand, there is no requirement that the training and evaluation sets need to be of equal size—with limited sampling, it may be advantageous to use a larger training set. Nevertheless, this presents a problem for the scaling of this method to larger systems, where it is common to have far fewer samples available than the volume of the lattice. Thus, many lattice QCD calculations use a $O(10^3)$ configurations when the number of field degrees of freedom is $\gtrsim 10^8$ [17]. A full treatment of this problem is left for future work, but here we show one approach in the present context.




FIG. 3. For a fixed time separation of $\tau = 10$, the scalar correlator on a 20×20 lattice with $m^2 = 0.1$ as a function of the bare coupling. A total of 10^4 samples were used, with half being reserved for optimization when a control variate is used. Inset is plotted the size of the error bars alone, for the same data.

As a straightforward approach to avoid overfitting, we might try to design a family of control variates with relatively few parameters and optimize over that family. However, that assumes *a priori* knowledge of which basis functions are heavily correlated with the desired observable. In this work we take the less opinionated approach of requiring the vector of coefficients c to be sparse. Each of the $\sim V$ coefficients begins unknown, but we demand that only a small number of them are sizable in the final result. To learn such a sparse control variate, we introduce an additional term into the objective function:

$$E_\mu^{(p)}(c) = c^T M c - 2c^T v + \mu \sum_i |c_i|^p. \quad (10)$$

This is the L_p regularization, and small positive values of p enforce high sparsity. The optimization is nonconvex for $p < 1$, so we choose to work with the L_1 norm. For well-chosen values of μ , the minimum of this objective will be near a minimum of the original objective function, while having only a small number of contributing entries. This sparsity acts to prevent overfitting. Heuristically, rather than requiring a number of samples greater than the total number of degrees of freedom, we need only to have as many samples as the number of *non-negligible* degrees of freedom. Of course, this objective function no longer has a closed-form solution for the minimum. Instead, we train the coefficients using the iterative shrinkage/thresholding algorithm [18], which performs substantially better than naïve gradient descent for such nonsmooth objectives.

Sparsity is not a property of a vector alone, but rather of the components of the vector as written in a certain basis. Equivalently, the L_1 norm added to the objective function of Eq. (10) is not basis independent. We chose to define the L_1 norm with respect to momentum basis; that is, the




FIG. 4. Correlator on a 50×50 lattice with $m^2 = \lambda = 0.1$, trained under L^1 regularization with $\mu = 10^{-4}$ as described in the text. A total of 10^3 samples are used. For clarity, points within 2σ of 0 are not plotted.

Fourier transform of the position basis in which $F_{xy}^{(1)}$ was defined above. This approach is motivated by the limit of weak coupling, as at $\lambda = 0$ only a single component need be nonzero to control the fluctuations.

Putting this together, we define a control variate

$$f = \frac{1}{L^2} \sum_{p,q} \tilde{c}_{p,q} \sum_{x,y} e^{ipx+iqy} F_{xy}^{(1)}, \quad (11)$$

where p, q are summed over all momentum modes on the lattice while x, y are summed over all L^2 lattice sites—where, as before, we have used translational invariance to reduce the implemented basis size. We optimize the coefficients $\tilde{c}_{p,q}$ to minimize the objective function $E_\mu^{(1)}(c)$ defined in Eq. (10). Different values of μ must be tried to find a control variate that makes an acceptable tradeoff between sparsity (avoiding overfitting) and quality; there is no guarantee that such a tradeoff is available.

Figure 4 shows the result of this optimization on a 50×50 lattice at bare lattice parameters $m^2 = \lambda = 0.1$. Without L_1 regularization, at least 2.5×10^4 samples would be required to avoid overfitting; here only 5×10^2 samples were used in the optimization of the control variates with an L_1 regularization strength of $\mu = 10^{-4}$, with another 5×10^2 used to compute the expectation value of the improved observable. Each data point corresponds to a different observable and therefore a different set of coefficients, but the typical vector of coefficients had ~ 100 elements of appreciable size, thus allowing optimization with no noticeable overfitting with only 500 samples.

Thus, we have developed a method in which straightforward optimization techniques can be used to choose combinations of Schwinger-Dyson equations to reduce the variance of observables. The strength of the method lies in the fact that existing ensembles of field configurations can be used, the improved observable is exactly equivalent to the original, and the path-integral measure is not distorted alleviating any concerns about overlap of distributions. Additionally, the use of Schwinger-Dyson equations allows physical intuition to guide the selection of the sparsity penalty that makes the method workable. The examination of the efficacy of the method for multipoint functions and for nonscalar fields is currently under investigation.

We are indebted to Rajan Gupta, Frederic Koehler, and Boram Yoon for helpful discussions over the course of this work. S. L. is grateful for the hospitality of the Los Alamos National Laboratory, where this project began. S. L. is supported by the U.S. Department of Energy under Contract No. DE-SC0017905, and T. B. and J.-S. Y. under Contract No. DE-AC52-06NA25396. T. B. and J.-S. Y. were also supported by the LDRD program of the Los Alamos National Laboratory. Code used in the demonstrations in this paper is available at [19].

[1] A. S. Kronfeld, T. Bhattacharya, T. Blum, N. H. Christ, C. DeT ar *et al.* (USQCD Collaboration), Lattice QCD and particle physics, [arXiv:2207.07641](#).
[2] Y.-C. Jang, R. Gupta, B. Yoon, and T. Bhattacharya, Axial vector form factors from lattice QCD that satisfy the PCAC relation, *Phys. Rev. Lett.* **124**, 072002 (2020).
[3] G. Parisi, The strategy for computing the hadronic mass spectrum, *Phys. Rep.* **103**, 203 (1984).
[4] G. P. Lepage, The analysis of algorithms for lattice field theory, in *Theoretical Advanced Study Institute in Elementary Particle Physics* (World Scientific, Singapore, 1990).

[5] Y. Aoki, T. Blum, G. Colangelo, S. Collins, M. Della Morte *et al.* (Flavour Lattice Averaging Group (FLAG) Collaboration), FLAG review 2021, *Eur. Phys. J. C* **82**, 869 (2022).
[6] C. Aubin, T. Blum, C. Tu, M. Golterman, C. Jung, and S. Peris, Light quark vacuum polarization at the physical point and contribution to the muon $g - 2$, *Phys. Rev. D* **101**, 014503 (2020).
[7] T. Blum, T. Izubuchi, and E. Shintani, New class of variance-reduction techniques using lattice symmetries, *Phys. Rev. D* **88**, 094503 (2013).

- [8] B. Yoon, T. Bhattacharya, and R. Gupta, Machine learning estimators for lattice QCD observables, *Phys. Rev. D* **100**, 014504 (2019).
- [9] W. Detmold, G. Kanwar, M.L. Wagman, and N.C. Warrington, Path integral contour deformations for noisy observables, *Phys. Rev. D* **102**, 014514 (2020).
- [10] G. Kanwar, A. Lovato, N. Rocco, and M. Wagman, Mitigating Green's function Monte Carlo signal-to-noise problems using contour deformations, [arXiv:2304.03229](https://arxiv.org/abs/2304.03229).
- [11] A. Alexandru, G. Basar, P.F. Bedaque, and N.C. Warrington, Complex paths around the sign problem, *Rev. Mod. Phys.* **94**, 015006 (2022).
- [12] S. Lawrence, Perturbative removal of a sign problem, *Phys. Rev. D* **102**, 094504 (2020).
- [13] S. Lawrence and Y. Yamauchi, Deep learning of fermion sign fluctuations, *Phys. Rev. D* **107**, 114505 (2023).
- [14] L. Fernandez and V. Martin-Mayor, Mean-value identities as an opportunity for Monte Carlo error reduction, *Phys. Rev. E* **79**, 051109 (2009).
- [15] M. Weigel and W. Janke, Error estimation and reduction with cross correlations, *Phys. Rev. E* **81**, 066701 (2010).
- [16] M. E. Peskin and D. V. Schroeder, *An Introduction to Quantum Field Theory* (Addison-Wesley, Reading, USA, 1995).
- [17] A. Bazavov, C. Bernard, C. DeTar, W. Freeman, S. Gottlieb *et al.* (MILC Collaboration), Lattice QCD ensembles with four flavors of highly improved staggered quarks, *Phys. Rev. D* **87**, 054505 (2013).
- [18] A. Beck and M. Teboulle, A fast iterative shrinkage-thresholding algorithm for linear inverse problems, *SIAM J. Imaging Sci.* **2**, 183 (2009).
- [19] T. Bhattacharya, S. Lawrence, and J.-S. Yoo (2023), <https://gitlab.com/s.lawrence/latticecv> (source on Gitlab) (2023).