Letter

Lattice real-time simulations with learned optimal kernels

Daniel Alvestad,[1] Alexander Rothkopf,[1] and Dénes Sexty[2]
[1]*Department of Mathematics and Physics, University of Stavanger, 4021 Stavanger, Norway*
[2]*Institute of Physics, NAWI Graz, University of Graz, Universitätsplatz 5, Graz, Austria*

(Received 20 October 2023; accepted 4 January 2024; published 5 February 2024)

We present a simulation strategy for the real-time dynamics of quantum fields, inspired by reinforcement learning. It builds on the complex Langevin approach, which it amends with system-specific prior information, a necessary prerequisite to overcome this exceptionally severe sign problem. The optimization process underlying our machine-learning approach is made possible by deploying inherently stable solvers of the complex Langevin stochastic process and a novel optimality criterion derived from insight into so-called boundary terms. This conceptual and technical progress allows us to both significantly extend the range of real-time simulations in $1+1$d scalar field theory beyond the state of the art and to avoid discretization artifacts that plagued previous real-time field theory simulations. Limitations of and promising future directions are discussed.

DOI: 10.1103/PhysRevD.109.L031502

Introduction. What unites many of the pressing open questions in modern physics, irrespective of whether they relate to eV- (condensed matter), MeV- (nuclear), or GeV- (particle physics) energy scales, is the need to access the dynamics of strongly correlated quantum many-body systems in Minkowski time. Concretely, as, e.g., outlined in the recent Snowmass community review [1,2], an *ab initio* understanding of transport properties of nuclear matter at high temperature and density, as well as the scattering of showers of high-energy partons still remain out of reach of state-of-the-art analytic and numerical Monte Carlo methods. First principles insight into real-time transport of nonrelativistic fermions [3] and their interaction with gauge fields is a key puzzle piece in understanding high-temperature superconductivity (e.g., in the Hubbard model [4]). Fission and fusion dynamics [5], too, remain currently out of reach of fully *ab initio* field-theoretic approaches, requiring model input.

Vital *ab initio* insight into the static (thermodynamic) properties of strongly correlated many-body systems has been achieved in the past through Monte Carlo simulations of Feynman's path integral [6]. These numerical techniques rely on analytic continuation to an unphysical Euclidean time. In turn, the extraction of relevant real-time dynamics becomes an ill-posed inverse problem [7], which severely affects the accurate determination of central quantities of interest: from transport coefficients [8,9] to in-medium decay rates [10,11] to vacuum parton distribution functions [12,13]. Developing a direct simulation approach in Minkowski time is thus called for.

Direct simulations of real-time dynamics suffer from the so-called sign problem [14,15]. Feynman's path integral is formulated as a sum over field configurations weighted by a complex phase. A minute signal emerges from the sum of a vast number of almost canceling phases, overwhelming otherwise efficient Markov-chain sampling based approaches. Some sign problems have been proven [16] to belong to the class of nondeterministic-polynomial-hard computational problems, which entails that no generic solution method in polynomial time exists on a classical computer. Various approaches have been proposed to tackle the sign problem, such as reweighting (RW), extrapolation [17–20], density of states (DSs) [21–23], tensor networks (TNs) [24,25], Lefschetz thimbles (LTs) [26–29], and complex Langevin (CL) [30,31]. They all propose a "system-agnostic" recipe to the estimation of observables in the presence of a sign problem. Without a system-specific component, each of these methods are destined to eventually fail, be it that their computational cost scales unfavorably when applied to systems in realistic volumes in $3+1$ dimensions (RW, TNs, DSs, LTs) or they suffer from convergence to an incorrect solution (CL).

Quantum computing offers a different angle of attack to the sign problem [32], as in principle it can compute the unitary time evolution of a spin system. The mapping of a realistic field theory to spin systems remains an open challenge, especially if gauge degrees of freedom are involved [2]. The necessity to derive a Hamiltonian for implementation on a quantum computer, to date, requires

truncation of the continuous state space and in the case of photons and gluons fixing to a particular gauge. It is acknowledged in the quantum computing community (see, e.g., [33,34]) that, with near-future noisy intermediate scale devices, many physical systems of interest remain too complex to be modeled with quantum circuits.

Therefore, progress in the short term requires innovation among "system-specific" real-time simulation techniques on classical computers.

Real-time complex Langevin. Here we build upon the complex Langevin approach, which is one of the complexification strategies to the sign problem, with similarities, but important differences to contour deformations (LTs) (see discussion in [35,36]). In conventional stochastic quantization [37,38] one proves that the expectation values of a Euclidean quantum field theory $\langle\mathcal{O}\rangle(\tau) = \int \mathcal{D}\phi_E \mathcal{O}(\phi_E)\exp[-S_{\mathrm{E}}(\phi_E)]$ can be reproduced by simulating a stochastic process in an additional Langevin time τ_L direction, using the Langevin equation $\partial_{\tau_L}\phi(\tau_L,\tau) = -\delta S_{\mathrm{E}}/\delta\phi_E + \eta(\tau_L,\tau)$ with Gaussian noise $\langle\eta(\tau_L,\tau)\eta(\tau_L',\tau')\rangle = 2\delta(\tau_L-\tau_L')\delta(\tau-\tau')$.

Quantum field theory with a mixed initial density matrix in Minkowski time constitutes an initial value problem and is formulated on the Schwinger-Keldysh contour $\mathcal{C}_{\mathrm{SK}}$ with a forward and backward branch, housing the fields ϕ_1 and ϕ_2, respectively,

$$\langle\mathcal{O}\rangle(t,\mathbf{x}) = \int d[\phi_1^{(i)},\phi_2^{(i)}]\rho(\phi_1^{(i)},\phi_2^{(i)})\int_{\phi_1^{(i)}}^{\phi_2^{(i)}}\mathcal{D}\phi\mathcal{O}e^{iS_{\mathcal{C}_{\mathrm{SK}}}} \tag{1}$$

$$\overset{\rho_\beta}{=}\int\mathcal{D}\phi_E e^{-S_E}\int_{\phi_E^{(f)}}^{\phi_E^{(i)}}\mathcal{D}\phi\mathcal{O}e^{iS_{\mathcal{C}_{\mathrm{SK}}}} = \int_{\mathcal{S}_{\mathrm{SKt}}}\mathcal{D}\phi\mathcal{O}e^{iS_{\mathcal{C}_{\mathrm{SKt}}}}. \tag{2}$$

In a thermal system at $T = 1/\beta$, we have $\rho_\beta \sim \exp[-\beta H]$ and sampling over initial conditions can be written as a path integral on a compact imaginary time domain of length β, connecting the real-time branches as closed contour $\mathcal{C}_{\mathrm{SKt}}$. We parametrize $\mathcal{C}_{\mathrm{SKt}}$ in the complex time plane with the real contour parameter γ: $t(\gamma)$. Cauchy's theorem allows us to deform the integration contour and we choose the convention sketched in Fig. 1, where the downward portion of $\mathcal{C}_{\mathrm{SKt}}$ is divided into two pieces at $t = t_{\max}$ and $t = t_0$. On $\mathcal{C}_{\mathrm{SKt}}$, Eq. (2) is highly oscillatory and requires regularization, which is implemented conventionally as a regulator term in $S_{\mathcal{C}_{\mathrm{SK}}}$ or as a tilt of the contour into the complex plane (see, e.g., [39]).

The naive CL method proposes to complexify the field degree of freedom (d.o.f.) $\varphi = \phi_R + i\phi_I$ and to carry out the following coupled stochastic process [38] for ϕ_R and ϕ_I in Langevin time τ_L:

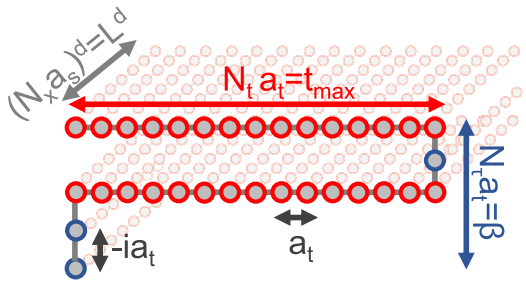

FIG. 1. Geometry of discretized $(d+1)$-dimensional scalar field theory on the Schwinger-Keldysh contour. Here we use $N_t = 32$, $N_\tau = 4$, and $N_x = 8$ and $a_t m = 1/10, a_s m = 2/10$.

$$\frac{d\varphi(\tau_L,x)}{d\tau_L} = i\frac{\delta S[\varphi]}{\delta\varphi(\tau_L,x)} + \eta(\tau_L,x), \tag{3}$$

estimating observables $\langle\mathcal{O}\rangle$ from the ensemble average over analytically continued observables $\langle\mathcal{O}\rangle[\varphi]$. While significant progress has been made in application of CL in various model systems [40,41] and even to the theory of strong interactions at finite Baryon-chemical potential [42–46], the simulation of real-time dynamics so far has been hampered by various hurdles: divergencies (runaways) and convergence to incorrect solutions as the real-time extent of the contour is increased [47–49].

The runaway problem leads to a breakdown of the numerical solver for the Langevin equation when the process explores regions of the complexified manifold far from the origin. It has been shown in Ref. [39], that it can also be understood as a consequence of the stiffness of the nonlinear CL dynamics. This practical problem is solved by either using an adaptive step size control [50] or through the use of inherently stable implicit discretization schemes, such as the Euler-Maruyama method, for Eq. (3) [39]. Inherent regularization provided by the implicit scheme also makes it possible to directly simulate with CL on the real-time axis of the $\mathcal{C}_{\mathrm{SKt}}$ contour, without tilt.

Important insight into the convergence properties of CL have been gained in [51,52] through analysis of the relation between the real probability distribution $P[\phi_R,\phi_I]$ sampled by Eq. (3) and the complex Feynman weights $\exp[iS]$ in Eq. (2). Connection is made via the real Fokker-Planck operator L and its complex generalization $\mathcal{L}$. For CL to correctly reproduce expectation values, the sampled distribution $P[\phi_R,\phi_I]$ must fall off sufficiently fast in ϕ_I, to enable integration by parts. At the same time, the spectrum of $\mathcal{L}$ must have negative real parts. Based on this insight, criteria for correct convergence have been developed: boundary terms [53,54] and two improvements to the complex Langevin method have been proposed: gauge cooling [55], where gauge freedom is exploited to keep the d.o.f. close to the original real-valued manifold, and dynamic stabilization [56], which introduces an additional drift term into Eq. (3). The drawback of the latter is that the new drift term is nonholomorphic and thus at odds with the

proof of convergence of the CL and it might introduce a bias in the results.

It is long known [38] that the real Langevin method can be modified by a kernel K, without changing its stationary distribution. This freedom has been exploited to improve autocorrelation in Euclidean theories [57]. In the complex Langevin method, one may introduce a complex kernel $K = K[\varphi; \tau_{\mathrm{L}}]$ [58–60]. K must be a holomorphic function and be factorizable as $K = H^T H$, but can otherwise be an arbitrary (matrix) function of the fields. It can encode transformations [61] such as in coordinates, contour deformations, or redefinition of variables. The general kerneled CL evolution equation for discretized spatial coordinates $\varphi(\tau_L, x_j) = \varphi^j_{\tau_L}$ reads

$$\frac{d\varphi^j}{d\tau_L} = \left[iK_{jk}(\varphi) \frac{\partial S(\varphi)}{\partial \varphi^k_{\tau_L}} + \frac{\partial K_{jk}(\varphi)}{\partial \varphi^k_{\tau_L}} \right] + H_{jk}(\varphi)\eta_k. \quad (4)$$

In the past, a few system-specific transformations have been found that softened (see, e.g., [62]) or even avoided (see, e.g., [59–61]) the sign problem in model systems (see also reformulation strategies, e.g., [63–65]). In real-time gauge theory, kernels have recently been explored in [66,67]. For realistic systems, success has been limited and no systematic recipe is known to extend results from simpler systems. This study instead uses machine-learning (ML) techniques to systematically learn optimal kernels, based on system-specific prior information.

Machine-learning assisted kerneled Langevin. Our machine-learning strategy for kerneled Langevin is inspired by reinforcement learning (RL) [68]. RL underlies recent advances in diverse fields: beating computer games or steering autonomous vehicles. It is based on an agent, endowed with a set of limited actions, placed in a predefined environment. Success of the agent is encoded in a cost/policy functional defined from environment variables and the internal state of the agent. A mathematical representation of the actions of the agent allows the use of differential programming techniques [69] to evaluate the gradients of the cost functional with respect to those actions. Challenges are the robust detection of failure modes of the agent and the trade-off between generality of the actions of the agent and learning efficiency.

Specifying to real-time simulations, we define our agent as the controller of the kernel K, which allows it to explore the abstract space of stationary distributions of the stochastic process Eq. (4). A crucial ingredient is our use of system-specific prior information to define the cost functional $\mathbb{p}[K]$, used to assess the success of CL convergence. As was shown, e.g., in [36] the failure of convergence of CL on the Schwinger-Keldysh (SK) contour occurs globally, i.e., it affects correlators on all branches. In a thermal setting, time translation invariance requires equal-time correlation functions to be constant on the whole complex time contour. Conventional simulations in the Euclidean domain in addition give access to their values, as well as to the Euclidean unequal-time correlators. Deviations from this prior knowledge are easily assessed within the CL simulations. To test for successful convergence, we deploy $\mathbb{p}[K] = \{\sum_{ij}(\mathrm{Re}\langle(\varphi^i)^2\rangle - \langle\phi^2\rangle_{\mathrm{HMC}})(C^{\mathrm{Re}})^{-1}_{ij}(\mathrm{Re}\langle(\varphi^j)^2\rangle - \langle\phi^2\rangle_{\mathrm{HMC}}) + (\mathrm{Im}\langle(\varphi^i)^2\rangle)(C^{\mathrm{Im}})^{-1}_{ij}(\mathrm{Im}\langle(\varphi^j)^2\rangle)\}$, which consists of two likelihood terms involving the covariance matrices $C^{\mathrm{Re/Im}}$ of the CL equal-time correlators $\mathrm{Re/Im}\langle\varphi^2\rangle$. The sum over i, j refers to all space-time coordinates and expectation values denote the average over τ_L. It thus assesses the constancy and agreement with *a priori* known values. $\mathbb{p}[K]$ makes reference to expectation values involving the fields φ and thus implicitly depends on K. To obtain robust gradients with respect to the entries of K, we must take derivatives over the whole stochastic dynamics. While adjoint [70] and shadowing methods [71] are promising to compute gradients, we find from CL Lyapunov exponents [72] that the dynamics actually becomes chaotic, degrading the performance of conventional differential programming techniques.

Instead we use a low-cost optimization functional $\mathbb{I}[K]$, which provides gradients $\nabla_K \mathbb{I}[K]$ that significantly reduce the values of the actual cost functional $\mathbb{p}[K]$. We find that

$$\mathbb{I}[K] = \int d^d x d\gamma \mathrm{Im}[\varphi(\tau_L, \gamma, \mathbf{x})]^2, \quad (5)$$

proposed in [73], offers the best performance in minimizing $\mathbb{p}[K]$, compared to earlier choices in [36]. The formulation of $\mathbb{I}[K]$, which leads to improved performance in locating optimal kernels, is in alignment with the insight gained by studying the improved correctness criterion in [74], as they require quick decay of the distribution of fields in imaginary directions. In general, using a low-cost functional will drive K toward the optimum of $\mathbb{p}[K]$ only at intermediate iteration steps. Hence our optimal K is taken at the smallest overall $\mathbb{p}[K]$ achieved. Interestingly, minimizing our specific cost functional $\mathbb{I}[K]$ we found that it always decreases the functional $\mathbb{p}[K]$.

We restrict ourselves to the simplest type of a field- and τ_L-independent kernel. Note that, even though the optimization functional may contain nonholomorphic terms, the kernel does not and thus Eq. (4) is compatible with the proof for correct convergence. The potentially costly Jacobian $\delta K/\delta\varphi$ also does not need to be computed.

Numerical results. Let us apply our machine-learning assisted complex Langevin approach to thermal scalar field theory in $1+1$ dimensions with $m = 1$ and a quartic self-coupling $(\lambda/4!)\phi^4$ with $\lambda = 1$ at $\beta m = 4/10$, a benchmark also used in [75]. The field is discretized on the SK contour sketched in Fig. 1 with $N_t = 32$ points along the real-time branches each and $N_\tau = 4$ steps along the imaginary time direction. The spatial dimension is resolved with $N_x = 8$

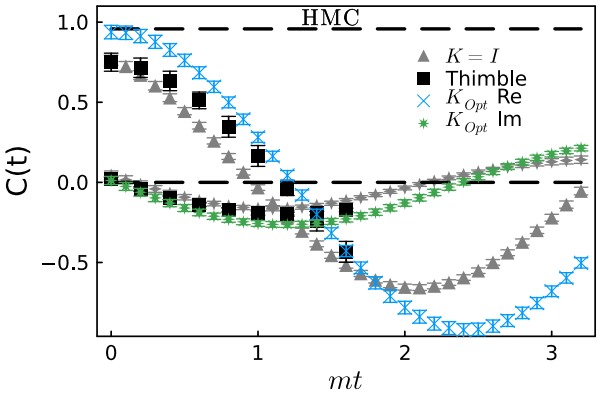

FIG. 2. Re[C] and Im[C] in 1 + 1d field theory from the naive (gray triangle, stars) and optimal kernel CL (blue crosses, green stars) with $N_t = 32$. Result from contour deformation [75] as black squares with $N_t = 8$. The value of the correlator at $t = 0$ from hybrid Monte Carlo simulations is given as gray solid line.

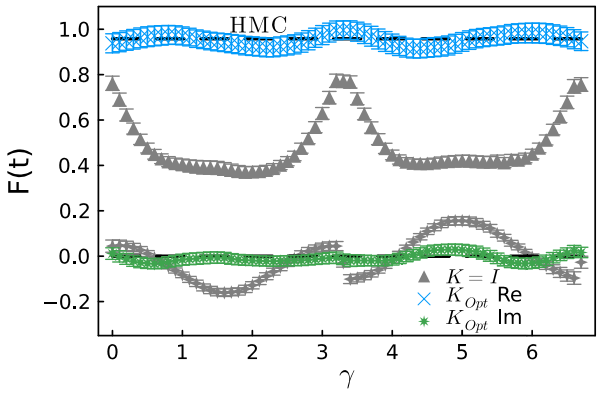

FIG. 3. Re$[F](\gamma)$ and Im$[F](\gamma)$ in 1 + 1d scalar field theory for the naive CL (gray triangle, rhombuses) and for optimal learned kernels (blue crosses, green stars).

points. As lattice spacing we use $a_s m = 2/10$ and a finer $a_t m = 1/10$, to avoid discretization artifacts from the corners of the SK contour. The discretized action is identical to the one of Ref. [75]. (The code is available at [76].)

Using adaptive step size with maximum Langevin step $d\tau_L = 10^{-3}$, we simulate at a real-time extent of $t_{\max} = 3.2$, which lies deep in the region where naive CL ($K = I$) fails to converge correctly, as shown by the gray triangles in Fig. 2, denoting the real and imaginary part of the unequal time, momentum zero correlator $C(t) = \langle \varphi(t, p = 0)\varphi(0, p = 0)\rangle$. The failure manifests in the deviation of $C(0)$ at $K = I$ from the value of the equal-time correlation function $F(\gamma) = \langle \varphi(\gamma, p = 0)\varphi(\gamma, p = 0)\rangle$ at $\gamma = 0$. Its values are known from conventional hybrid Monte Carlo (HMC) simulations and indicated by the black dashed line [in 1 + 1d, $F(0)$ only carries a minute lattice spacing dependence].

We parametrize a fully dense, constant, complex kernel via $H = A + iB$ with real matrices A and B, each of which have $[(2N_t + N_\tau)N_x]^2$ entries. Learning of the optimal kernel starts from the trivial choice $H = \mathbb{1}$, i.e., $K = H^T H = \mathbb{1}$. The resulting stiff dynamics is solved with an implicit Euler-Maruyama integrator for which we use the implementation in the `DifferentialEquations.jl` library [77] of the Julia language. After a Langevin time of $\tau_L^{\text{opt}} = 5$ we compute the gradient of the discrete Eq. (5) using the automatic differentiation capability of Julia. Based on the Adam algorithm [78] with learning rate $r_l = 10^{-3}$ the entries of A and B are iteratively updated reducing the initial $\mathbb{p}[K = \mathbb{1}] \approx 1200$ to $\mathbb{p}[K_{\text{opt}}] \approx 7.2$. Observables for K_{opt} are obtained from three streams with $\tau_L^{\text{obs}} = 5000$.

The central result of our study, the unequal-time correlation function $C(t)$ from optimal learned kernels in 1 + 1d, is shown as blue crosses Re$[C]$ and green stars Im$[C]$ in Fig. 2. We reach a real-time extent at $t_{\max} = 3.2$ which is twice that previously achieved in the literature using contour deformations [75]. Note that our simulation results for Re$[C](0)$ and Im$[C](0)$ both agree with $F(0)$ from HMC simulations, given by given by the dashed black lines.

We emphasize that the advantageous scaling properties, which our approach inherits from the CL, enable us to deploy a finer grid here. Implicit methods do not pose a problem, as highly optimized implementations of solvers are readily available. Since we restrict ourselves to field-independent kernels so far, we avoid the need for Jacobians, whose computational cost is a central limiting factor for contour deformation methods.

We showed in [36] that the equal-time correlator $F(\gamma)$ is more difficult to reproduce in the CL method than the unequal-time correlator $C(t)$ and thus plot the former against the contour parameter γ in Fig. 3. Note that both Re$[F](\gamma)$ and Im$[F](\gamma)$, while showing minute oscillations, agree with the *a priori* known values (gray dashed lines). The optimal kernel CL results are in stark contrast to the naive CL ($K = I$), for which $F(\gamma)$ (gray triangles, rhombuses) clearly deviates from the HMC simulation. This crosscheck provides convincing support for the correctness of convergence.

Further support for correct convergence is provided from the absence of boundary terms, such as B_1 (see [53,54]) for

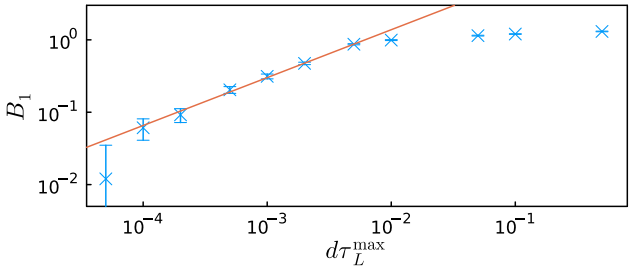

FIG. 4. Magnitude of the dominant boundary term B_1, based on the $\langle \phi^2 \rangle$ observable. Note the best fit power-law dependence on $(d\tau_L^{\max})^{0.66}$ consistent with correct convergence.

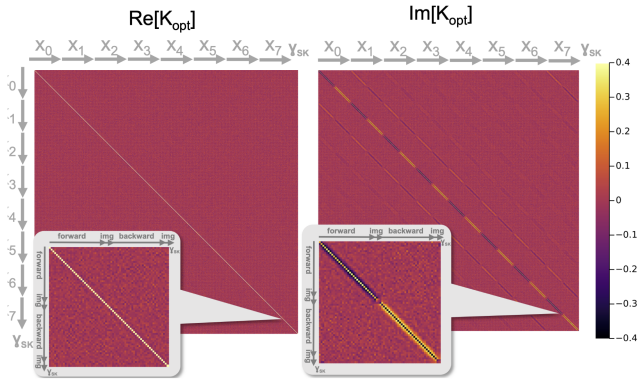

FIG. 5. Top: entries of the real (left) and imaginary part (right) of the optimal kernel underlying Figs. 2 and 3. Bottom insets: enlargement of (left) $\mathrm{Re}[K_{\mathrm{opt}}]$ and (right) $\mathrm{Im}[K_{\mathrm{opt}}]$ on a single spatial slice with the subsections of $\mathcal{C}_{\mathrm{SKt}}$ labeled on gray arrows. (see text for detailed discussion).

the $\langle\phi^2\rangle$ observable. As shown in Fig. 4, B_1 exhibits a power-law dependence on $d\tau_L$, consistent with vanishing boundary terms in the continuum limit.

In the top row of Fig. 5 we plot the values of the optimal learned kernel underlying Figs. 2 and 3. Color coding resolves values $|K_{\mathrm{opt}}| \leq 0.4$, sufficient for all but the diagonal entries of $\mathrm{Re}[K_{\mathrm{opt}}^{\mathrm{diag}}] \approx 1.025$. Each pixel corresponds to a component of K_{opt} that connects two space-time points, ordered such that, in each spatial slice x_i denoted by a gray arrow, the parameter γ traverses the full contour $\mathcal{C}_{\mathrm{SKt}}$. $\mathrm{Re}[K_{\mathrm{opt}}]$ is dominated by the diagonal alone, while $\mathrm{Im}[K_{\mathrm{opt}}]$ shows a distinct banded structure in space with diminishing amplitude farther away from the diagonal, which represents the spatially nonlocal nature of the transformation implemented by K_{opt}. In the two insets, we highlight the behavior of the kernel in a single spatial slice, where the corresponding sections of the SK contour connected by the kernel entries are indicated by the gray arrows. Similar to results in $0+1$d, we find a characteristic finite-difference-like behavior. Entries of opposite sign on the sub- and supradiagonal accompany those on the diagonal, indicating a Fourier filter.

Our ML approach is able to identify a much simpler structure than what a naive extension of the free theory kernel suggests (cf. [36]), a testament to the efficacy of the learning strategy, which holds potential for analytic insight into convergence restoring transformations.

While success of the ML approach is encouraging, it too will fail at larger real-time extents. Establishing an exact range of validity is work in progress. One reason is that we only optimize based on the low-cost functional $\mathbb{I}[K]$ and not directly on $\mathbb{p}[K]$. The development of robust gradient estimators for chaotic stochastic systems (see, e.g., NILSAS for the Lorenz system [72]) is called for. Another reason is the thimble structure of the theory (see also [79]). One d.o.f. models tell us that parameters exist in which a field-dependent kernel is needed to capture the physics of multiple contributing thimbles. In that case the kernel can be systematically expanded via a rational approximation involving φ, limiting computational cost. Transfer learning between kernels of different expressivity will be key to keep cost in check.

In summary, we have presented a machine-learning approach to direct real-time simulations on the lattice, in which system-specific prior information is incorporated into the CL via iterative ML of an optimal field-independent kernel. Using a novel low-cost functional for the computation of gradients for learning, we achieve efficient convergence in $1+1$d field theory to at least twice the real-time extent previously accessible. Because of the efficiency of the approach, we can access fine grids to avoid discretization artifacts affecting previous studies. Work is ongoing to extend the results to realistic $(3+1)$ dimensions and we explore the inclusion of field-dependent kernels.

The code for this paper is publicly available from the Zenodo repository [76].

Acknowledgments. D. A. and A. R. thank Rasmus Larsen for helpful discussions and acknowledge support by the Research Council of Norway under the FRIPRO YRT Grant No. 286883. D. S. acknowledges support of the Austrian Science Fund (FWF) through the stand-alone project P36875. The study benefited from computing resources from Sigma2 under project NN9578K-QCDrtX, as well as from the computing cluster of the University of Graz (GSC).

[1] Z. Davoudi *et al.*, arXiv:2209.10758.
[2] B. Nachman, D. Provasoli, W. A. De Jong, and C. W. Bauer, Phys. Rev. Lett. **126**, 062001 (2021).
[3] D. Hangleiter, I. Roth, D. Nagaj, and J. Eisert, Sci. Adv. **6**, eabb8341 (2020).
[4] M. Qin, T. Schäfer, S. Andergassen, P. Corboz, and E. Gull, Annu. Rev. Condens. Matter Phys. **13**, 275 (2022).
[5] M. Bender *et al.*, J. Phys. G **47**, 113002 (2020).
[6] J. Smit, *Introduction to Quantum Fields on a Lattice: A Robust Mate* (Cambridge University Press, Cambridge, England, 2011), Vol. 15.
[7] A. Rothkopf, Front. Phys. **10** (2022).
[8] N. Brambilla, V. Leino, J. Mayer-Steudte, and P. Petreczky (TUMQCD Collaboration), Phys. Rev. D **107**, 054508 (2023).

[9] L. Altenkort, O. Kaczmarek, R. Larsen, S. Mukherjee, P. Petreczky, H.-T. Shu, and S. Stendebach (HotQCD Collaboration), Phys. Rev. Lett. **130**, 231902 (2023).
[10] S. Kim, P. Petreczky, and A. Rothkopf, J. High Energy Phys. 11 (2018) 088.
[11] R. Larsen, S. Meinel, S. Mukherjee, and P. Petreczky, Phys. Lett. B **800**, 135119 (2020).
[12] J. Karpie, K. Orginos, A. Rothkopf, and S. Zafeiropoulos, J. High Energy Phys. 04 (2019) 057.
[13] J. Liang, T. Draper, K.-F. Liu, A. Rothkopf, and Y.-B. Yang (XQCD Collaboration), Phys. Rev. D **101**, 114503 (2020).
[14] C. Gattringer and K. Langfeld, Int. J. Mod. Phys. A **31**, 1643007 (2016).
[15] G. Pan and Z. Y. Meng, Encycl. Condens. Matter Phys. **1**, 879 (2024).
[16] M. Troyer and U.-J. Wiese, Phys. Rev. Lett. **94**, 170201 (2005).
[17] P. de Forcrand and O. Philipsen, Phys. Rev. Lett. **105**, 152001 (2010).
[18] J. Braun, J.-W. Chen, J. Deng, J. E. Drut, B. Friman, C.-T. Ma, and Y.-D. Tsai, Phys. Rev. Lett. **110**, 130404 (2013).
[19] J. Braun, J. E. Drut, and D. Roscher, Phys. Rev. Lett. **114**, 050404 (2015).
[20] J. Guenther, R. Bellwied, S. Borsányi, Z. Fodor, S. Katz, A. Pásztor, C. Ratti, and K. Szabó, Nucl. Phys. **A967**, 720 (2017).
[21] F. Wang and D. P. Landau, Phys. Rev. Lett. **86**, 2050 (2001).
[22] K. Langfeld, B. Lucini, and A. Rago, Phys. Rev. Lett. **109**, 111601 (2012).
[23] C. Gattringer and P. Törek, Phys. Lett. B **747**, 545 (2015).
[24] R. Orus, Nat. Rev. Phys. **1**, 538 (2019).
[25] Y. Meurice, J. C. Osborn, R. Sakai, J. Unmuth-Yockey, S. Catterall, and R. D. Somma, arXiv:2203.04902.
[26] N. Rom, D. Charutz, and D. Neuhauser, Chem. Phys. Lett. **270**, 382 (1997).
[27] M. Cristoforetti, F. Di Renzo, and L. Scorzato (AuroraScience Collaboration), Phys. Rev. D **86**, 074506 (2012).
[28] A. Alexandru, G. Basar, P. F. Bedaque, and N. C. Warrington, Rev. Mod. Phys. **94**, 015006 (2022).
[29] J. Nishimura, K. Sakai, and A. Yosprakob, J. High Energy Phys. 09 (2023) 110.
[30] J. R. Klauder, Acta Phys. Austriaca Suppl. **25**, 251 (1983).
[31] G. Parisi, Phys. Lett. **131B**, 393 (1983).
[32] J. Preskill, Quantum **2**, 79 (2018).
[33] C. W. Bauer *et al.*, PRX Quantum **4**, 027001 (2023).
[34] A. M. Dalzell *et al.*, arXiv:2310.03011.
[35] G. Aarts, L. Bongiovanni, E. Seiler, and D. Sexty, J. High Energy Phys. 10 (2014) 159.
[36] D. Alvestad, R. Larsen, and A. Rothkopf, J. High Energy Phys. 04 (2023) 057.
[37] P. H. Damgaard and H. Hüffel, Phys. Rep. **152**, 227 (1987).
[38] M. Namiki, I. Ohba, K. Okano, Y. Yamanaka, A. K. Kapoor, H. Nakazato, and S. Tanaka, Lect. Notes Phys. Monogr. **9**, 1 (1992).
[39] D. Alvestad, R. Larsen, and A. Rothkopf, J. High Energy Phys. 08 (2021) 138.
[40] E. Seiler, EPJ Web Conf. **175**, 01019 (2018).
[41] C. E. Berger, L. Rammelmüller, A. C. Loheac, F. Ehmann, J. Braun, and J. E. Drut, Phys. Rep. **892**, 1 (2021).
[42] D. Sexty, Phys. Lett. B **729**, 108 (2014).
[43] D. Sexty, Phys. Rev. D **100**, 074503 (2019).
[44] M. Scherzer, D. Sexty, and I. O. Stamatescu, Phys. Rev. D **102**, 014515 (2020).
[45] Y. Ito, H. Matsufuru, Y. Namekawa, J. Nishimura, S. Shimasaki, A. Tsuchiya, and S. Tsutsui, J. High Energy Phys. 10 (2020) 144.
[46] F. Attanasio, B. Jäger, and F. P. G. Ziegler, arXiv:2203.13144.
[47] J. Berges and I. O. Stamatescu, Phys. Rev. Lett. **95**, 202003 (2005).
[48] J. Berges, S. Borsanyi, D. Sexty, and I. O. Stamatescu, Phys. Rev. D **75**, 045007 (2007).
[49] J. Berges and D. Sexty, Nucl. Phys. **B799**, 306 (2008).
[50] G. Aarts, F. A. James, E. Seiler, and I.-O. Stamatescu, Phys. Lett. B **687**, 154 (2010).
[51] G. Aarts, F. A. James, E. Seiler, and I.-O. Stamatescu, Eur. Phys. J. C **71**, 1756 (2011).
[52] K. Nagata, J. Nishimura, and S. Shimasaki, Phys. Rev. D **94**, 114515 (2016).
[53] M. Scherzer, E. Seiler, D. Sexty, and I.-O. Stamatescu, Phys. Rev. D **99**, 014512 (2019).
[54] M. Scherzer, E. Seiler, D. Sexty, and I. O. Stamatescu, Phys. Rev. D **101**, 014501 (2020).
[55] E. Seiler, D. Sexty, and I.-O. Stamatescu, Phys. Lett. B **723**, 213 (2013).
[56] G. Aarts, F. Attanasio, B. Jäger, and D. Sexty, Acta Phys. Pol. B Proc. Suppl. **9**, 621 (2016).
[57] G. G. Batrouni, G. R. Katz, A. S. Kronfeld, G. P. Lepage, B. Svetitsky, and K. G. Wilson, Phys. Rev. D **32**, 2736 (1985).
[58] B. Soderberg, Nucl. Phys. **B295**, 396 (1988).
[59] H. Okamoto, K. Okano, L. Schulke, and S. Tanaka, Nucl. Phys. **B324**, 684 (1989).
[60] K. Okano, L. Schulke, and B. Zheng, Phys. Lett. B **258**, 421 (1991).
[61] G. Aarts, F. A. James, J. M. Pawlowski, E. Seiler, D. Sexty, and I.-O. Stamatescu, J. High Energy Phys. 03 (2013) 073.
[62] R. Levy and B. K. Clark, Phys. Rev. Lett. **126**, 216401 (2021).
[63] S. Chandrasekharan and U.-J. Wiese, Phys. Rev. Lett. **83**, 3116 (1999).
[64] Y. Delgado Mercado, H. G. Evertz, and C. Gattringer, Phys. Rev. Lett. **106**, 222001 (2011).
[65] T. Kloiber and C. Gattringer, Proc. Sci. LATTICE2013 (**2014**) 206 [arXiv:1310.8535].
[66] K. Boguslavski, P. Hotzy, and D. I. Müller, J. High Energy Phys. 06 (2023) 011.
[67] K. Boguslavski, P. Hotzy, and D. I. Müller, arXiv:2312.03063.
[68] R. Sutton and A. Barto, *Reinforcement Learning, Second Edition: An Introduction*, Adaptive Computation and Machine Learning Series (MIT Press, Cambridge, MA, 2018).
[69] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, J. Mach. Learn. Res. **18**, 1 (2018), https://jmlr.org/papers/v18/17-468.html.
[70] Y. Cao, S. Li, L. Petzold, and R. Serban, SIAM J. Sci. Comput. **24**, 1076 (2003).
[71] Q. Wang, R. Hu, and P. Blonigan, J. Comput. Phys. **267**, 210 (2014).
[72] D. Alvestad, Real-time complex Langevin—A differential programming perspective, Ph.D. thesis, University of Stavanger, 2023.
[73] N. M. Lampl and D. Sexty, arXiv:2309.06103.

[74] E. Seiler, D. Sexty, and I.-O. Stamatescu, arXiv:2304.00563.
[75] A. Alexandru, G. Basar, P. F. Bedaque, and G. W. Ridgway, Phys. Rev. D **95**, 114501 (2017).
[76] D. Alvestad, alvestad10/KernelCL: Towards learning optimized kernels for complex Langevin, Zenodo, 10.5281/zenodo.7373498 (2022).
[77] C. Rackauckas and Q. Nie, J. Open Res. Software **5** (2017).
[78] D. P. Kingma and J. Ba, arXiv:1412.6980.
[79] Z.-G. Mou, P. M. Saffin, A. Tranberg, and S. Woodward, J. High Energy Phys. 06 (2019) 094.