Estimating the Euclidean quantum propagator with deep generative modeling of Feynman paths

Yanming Che^{1,2,*} Clemens Gneiting,^{2,3,†} and Franco Nori^{1,2,3,‡}

¹Department of Physics, University of Michigan, Ann Arbor, Michigan 48109-1040, USA

²Theoretical Quantum Physics Laboratory, RIKEN Cluster for Pioneering Research, Wako-shi, Saitama 351-0198, Japan

³RIKEN Center for Quantum Computing (RQC), 2-1 Hirosawa, Wako-shi, Saitama 351-0198, Japan

(Received 8 February 2022; revised 1 June 2022; accepted 1 June 2022; published 15 June 2022)

Feynman path integrals provide an elegant, classically inspired representation for the quantum propagator and the quantum dynamics, through summing over a huge manifold of all possible paths. From computational and simulational perspectives, the ergodic tracking of the whole path manifold is a hard problem. Machine learning can help, in an efficient manner, to identify the relevant subspace and the intrinsic structure residing at a small fraction of the vast path manifold. In this work, we propose the Feynman path generator for quantum mechanical systems, which efficiently generates Feynman paths with fixed endpoints, from a (low-dimensional) latent space and by targeting a desired density of paths in the Euclidean space-time. With such path generators, the Euclidean propagator as well as the ground-state wave function can be estimated efficiently for a generic potential energy. Our work provides an alternative approach for calculating the quantum propagator and the ground-state wave function, paves the way toward generative modeling of quantum mechanical Feynman paths, and offers a different perspective to understand the quantum-classical correspondence through deep learning.

DOI: 10.1103/PhysRevB.105.214205

I. INTRODUCTION

Feynman path (FP) integrals provide an elegant, classically inspired representation for the quantum propagator and the quantum dynamics, through summing over a huge manifold of all possible paths connecting two fixed endpoints [1–3]. From the perspective of the path integral, quantum dynamics arises from coherent contributions including both the classical path and quantum fluctuations. It thus also provides an intuitive framework for understanding the quantum-classical correspondence.

Generally, the ergodic tracking of each path contribution to the quantum propagator is a computationally hard problem. In this respect, a similar situation occurs in the classical simulation of many-body quantum systems, whose Hilbert space is exponentially large, prohibiting exact and efficient methods in the generic case. However, in many practical scenarios, physically relevant features and structures may only reside at a small fraction of the vast Hilbert space; this enables efficient machine learning techniques for quantum-physics problems, like the neural-network approach for quantum many-body systems [4–18], neural network quantum state tomography [19–24], manifold learning and clustering of quantum phases [25–30], and physical concepts rediscovered with neural networks [31,32].

Similarly, in the path integral formalism, contributions from large deviations with respect to the classical path generically cancel each other due to rapid oscillations (in the

2469-9950/2022/105(21)/214205(8)

Lorentzian space-time) or are exponentially suppressed (in the Euclidean space-time) [1-3]. Therefore, dominant contributions to the propagator may only reside at a corner of the huge path manifold. Then one may naturally ask:

(1) Can machine learning capture these structures in FPs?(2) Can the propagator be efficiently estimated with neural networks?

So far, the literature relating machine learning to FP integrals either has not fully investigated these questions [33–37], or has mainly focused on quantum field-theoretical systems [38–45]. In this paper, we address these questions by proposing the FP generator for quantum mechanical systems, which provides a fresh perspective on efficiently calculating the propagator as well as the ground-state wave function, and on studying FPs with deep generative modeling.

II. LORENTZIAN AND EUCLIDEAN FEYNMAN PATHS

We start from a brief overview of the real- and imaginarytime FP integral. The quantum unitary time evolution of the wave function is given by

$$\psi(x_f, t_f) = \int dx_i \ K(x_f, t_f; x_i, t_i) \psi(x_i, t_i), \qquad (1)$$

where (x, t) denotes the (d + 1)-dimensional space-time, and $\psi(x_i, t_i)$ is the initial wave function. In the FP integral formulation, the propagator (or kernel) is given by summing over all possible paths starting from x_i at time $t_i = 0$ and ending at x_f at time t_f [1–3,46]:

$$K(x_f, t_f; x_i, 0) = A(t_f) \sum_{\{x(t)\}} \exp(iS[x(t)]/\hbar), \qquad (2)$$

^{*}yanmingche01@gmail.com

[†]clemens.gneiting@riken.jp

^{*}fnori@riken.jp



FIG. 1. Schematics of Feynman paths and path integrals in (a) Lorentzian and (b) Euclidean space-times. The times τ and tare connected via the Wick rotation $it \rightarrow \tau$. The integrands of the propagators are plotted with the respective representative action curves S[x(t)] and $S_E[x(\tau)]$. In (a), Re and Im denote the real and imaginary paths, respectively, and $x_{cl}(t)$ stands for the classical path (stationary path), which satisfies $\delta S = 0$ [3]. In (b), $x_{cl}(\tau)$ is the path with the least Euclidean action. It can be seen that such paths and their neighborhoods contribute dominantly to the propagators, while large deviations away from them cancel each other through rapid oscillations in (a), and are exponentially suppressed in (b). In (b), the exponential decay factor defines a probability distribution function for the Feynman paths in Euclidean space-time. In the case when $x_{cl}(\tau)$ is identical to the classical Euclidean path, the red cloud around it represents quantum fluctuations.

where $A(t_f)$ is the normalization factor independent of the specific path; $\{x(t)\}$ is the manifold of all possible FPs connecting $(x_i, 0)$ and (x_f, t_f) in space-time; S[x(t)] is the classical action of the path x(t); and \hbar is the Planck constant.

An equivalent but more convenient formalism for numerical evaluation is to work in the Euclidean space-time, via the Wick rotation $it \rightarrow \tau$. Then the system propagates along the imaginary time τ , and the propagator at $\tau = T$ reads

$$K_E(x_f, T; x_i, 0) = A_E(T) \sum_{\{x(\tau)\}} \exp(-S_E[x(\tau)]/\hbar), \quad (3)$$

where $x(\tau)$ is the Euclidean FP and

$$S_E[x(\tau)] = \int_0^T H_E[x(\tau)] d\tau$$
(4)

is the Euclidean action, with

$$H_E[x(\tau)] = \frac{1}{2}m\left(\frac{dx}{d\tau}\right)^2 + V[x(\tau)]$$
(5)

the Hamiltonian and V(x) the potential energy.

From the Lorentzian and Euclidean path integrals, one can find sparse features and patterns in the path manifold (Fig. 1). The dominant contributions to the propagator come from paths that are located at the basin around the stationary path (which satisfies $\delta S = 0$) in the action landscape [47]. The stationary path together with quantum fluctuations in the vicinity of it dictate the quantum kernel, while large deviations cancel each other due to rapid oscillations in the real-time path integral, and are exponentially suppressed in the Euclidean version (3).

Equation (3) can be interpreted as a probability distribution of FPs in Euclidean space-time,

$$P_E[x(\tau)] = Z^{-1} \exp(-S_E[x(\tau)]/\hbar), \qquad (6)$$

where

$$Z = \sum_{\{x(\tau)\}} \exp(-S_E[x(\tau)]/\hbar)$$
(7)

is the partition function. Note that here Z depends on the two endpoints. When the normalization factor $A_E(T) = 1$, Z is the imaginary-time propagator in Eq. (3). Here $P_E[x(\tau)]$ is the target distribution of our FP generator.

III. VARIATIONAL FEYNMAN PATH GENERATOR

We now show that the Euclidean quantum propagator can be efficiently estimated through the generative modeling of FPs. We first introduce the concept of the FP generator, which produces FPs with fixed endpoints from a latent space, in which tracking and sampling are made simple. The statistics of the latent variables can be modeled by simple distributions such as the standard Gaussian. Such FP generators can be realized through variational recurrent neural networks (RNN) [48] or through normalizing flows [49,50].

In this work, we will restrict ourselves to the Euclidean space-time, where the density of paths is described by the target distribution (6), and we will use a variational FP generator (VFPG) of deep neural-net architecture, to approximate the target distribution and generate FPs via an efficient parallel



FIG. 2. (a) A variational realization of the Feynman path generator (VFPG) with a recurrent neural network (RNN). The input of the VFPG is a two-dimensional (2D) latent vector z sampled from the standard Gaussian distribution $\mathcal{N}(z; \mathbf{0}, \mathbf{I})$, followed by a repeat operation (Repeat) that extends the dimension to match the input shape of the RNN. The output sequence is fed into densely connected layers (Dense) to produce the variational parameters ϕ for the path distribution $q_{\phi}[x(\tau)]$, which is modeled, at each time stamp, by a Gaussian mixture (see Appendix A for details). The network is trained to target the path distribution in Eq. (6). (b) The structure of the RNN in (a), where $X(\tau)$ and $Y(\tau)$ are the input and output sequences, respectively; the blue squares on the right denote the RNN units, and the blue arrows denote the hidden state transfer [55]. In this work we use the long short-term memory (LSTM) [56] for the RNN.

sampling. As shown in Fig. 2, the architecture of the VFPG is similar to that of the decoder in the variational autoencoder (VAE), a well-known generative model for learning image and time-series signals [51–54].

Here a FP $x(\tau)$ is represented by a discrete N_{τ} -dimensional sequence (trajectory) $\{x(\tau_k)\}_{k=0}^{N_{\tau}-1}$, with $\tau_k = k \times \Delta \tau$, where $\Delta \tau = T/(N_{\tau} - 1)$ is the time step (see Appendix A for details). The VFPG samples 2D latent vectors z from the standard Gaussian distribution $\mathcal{N}(z; \mathbf{0}, \mathbf{I})$ as its input and generates the path time series $x(\tau)$ with fixed endpoints $(x_i, 0)$ and (x_f, T) , respectively, according to the output variational density distribution of paths $q_{\phi}[x(\tau)]$, with ϕ the output of the VFPG. The most probable output path should be the one that minimizes S_E .

As pictured in Fig. 2, the output sequence of the RNN is fed into a densely connected layer to produce the variational parameters ϕ , and at each time stamp, $q_{\phi}[x(\tau)]$ is modeled as a Gaussian mixture model, which is a universal approximator of densities [55] (see Appendix A). We use the long short-term memory (LSTM) [56] for the RNN to build the generative network.

IV. LOSS FUNCTION AND TRAINING

The loss function of the network is given by the Kullback-Leibler (KL) divergence between $q_{\phi}[x(\tau)]$ and $P_E[x(\tau)]$ in Eq. (6):

$$\mathcal{L}_{\text{VFPG}} = D_{\text{KL}} \{ q_{\phi}[x(\tau)] \mid \mid P_E[x(\tau)] \}, \tag{8}$$

subject to the constraints

$$x(0) = x_i, \ x(T) = x_f.$$
 (9)

The definition of the KL divergence is

$$D_{\text{KL}}[q(x) || p(x)] = \mathop{\mathbb{E}}_{x \sim q(x)} \ln \frac{q(x)}{p(x)},$$
 (10)

where \mathbb{E} denotes the expectation value.

.

After some straightforward derivations with the definition of the KL divergence and the expression of the target probability $P_E[x(\tau)]$ in Eq. (6), the VFPG loss can be rewritten as

$$\mathcal{L}_{\text{VFPG}} = \hbar^{-1} \left\{ \mathop{\mathbb{E}}_{x \sim q_{\phi}} S_E[x(\tau)] + \hbar \mathop{\mathbb{E}}_{x \sim q_{\phi}} \ln q_{\phi}[x(\tau)] \right\} + \ln Z.$$
(11)

The last term in Eq. (11) is independent of the network parameters and does not contribute to the gradient of the loss. The second expectation term is the (negative) differential entropy of the variational distribution q_{ϕ} .

The generator is trained to effectively find an equilibrium for the competition between minimizing the averaged action and maximizing the differential entropy. Viewing the action as an energy functional of a path configuration, and \hbar as the temperature (setting the Boltzmann constant $k_B = 1$), the quantity in the large brackets in Eq. (11) describes the variational free energy of the path manifold (ensemble):

$$\mathcal{F}_{\phi} = \mathop{\mathbb{E}}_{x \sim q_{\phi}} F_{\phi}[x(\tau)], \tag{12}$$

where

$$F_{\phi}[x(\tau)] = S_E[x(\tau)] + \hbar \ln q_{\phi}[x(\tau)]. \tag{13}$$

The true free energy is given by $\mathcal{F} = -\hbar \ln Z$. By taking the normalization in Eq. (3) to be $A_E(T) = 1$, we can estimate the Euclidean quantum propagator with $K_E(x_f, T; x_i, 0) = Z_{\phi}$, that is,

$$K_E(x_f, T; x_i, 0) = \exp(-\mathcal{F}_{\phi}/\hbar). \tag{14}$$

This means that we used an estimator Z_{ϕ} for Z which sets $\mathcal{L}_{\text{VFPG}} = \ln(Z/Z_{\phi})$.

Note that the generator loss (11), in terms of the gap between the variational and the true free energies, was also used for addressing the statistical mechanics of lattice spin models [38,57,58], for sampling molecular structures [59] and lattice field theories [38–45], and for the variational neural annealing [60]. For the VFPG here, the latent space is low dimensional and the path manifold is continuous, bringing new challenges for modeling the output density distribution of paths. In addition, the constraints on the two endpoints of each path will be considered as penalties during the training process [61] (see Appendix B for details).

V. GENERATING FEYNMAN PATHS AND ESTIMATING THE QUANTUM PROPAGATOR

Once the VFPG is trained, it can generate FPs by sampling the latent variables $z \sim \mathcal{N}(z; \mathbf{0}, \mathbf{I})$, giving an estimation of the Euclidean propagator $K_E(x_f, T; x_i, 0)$, which is the kernel of the imaginary-time propagation. In the demonstrated examples, we note that the choice of the number of time stamps, N_{τ} , only affects the normalization prefactor of the propagator that is independent of FPs. Thus the normalized propagator will not depend on the value of N_{τ} .

The spectral representation of the kernel is given by [62]

$$K_E(x_f, T; x_i, 0) = \sum_{n=0}^{\infty} e^{-TE_n/\hbar} \psi_n(x_i) \psi_n^*(x_f), \quad (15)$$

where E_n is the *n*th eigenenergy and $\psi_n(x)$ is the corresponding eigenstate at *x*. In the case $T \gg \hbar/\Delta E$, where ΔE is the energy gap between the ground state and the first excited state, we have

$$|\psi_0(x)|^2 \propto K_E(x, T; x, 0).$$
 (16)

Therefore, the normalized Euclidean propagator leads to the ground-state probability density $|\psi_0(x)|^2$.

A. Harmonic oscillator

First in Fig. 3 we present the results for the prototypical harmonic oscillator. Figure 3(a) shows the estimated Euclidean propagator (or the kernel) after the trace normalization (blue dots), with fixed $x_i = 0$. The propagator trace is defined as

$$\operatorname{tr} K_E = \int dx \ K_E(x, T; x, 0), \tag{17}$$

which is estimated through integrating over the smooth fit of VFPG results for the diagonal propagator. Figure 3(b) plots the ground-state probability density obtained from the diagonal propagator (blue dots). The blue dashed line is a Gaussian fit of the VFPG observations. In Figs. 3(a) and 3(b), the VFPG observations exhibit good agreement with the (red dashed)



FIG. 3. Estimating the Euclidean quantum propagator through the variational Feynman path generator (VFPG) for a harmonic oscillator, where $V(x) = \frac{1}{2}m\omega^2 x^2$. (a) The trace-normalized propagator vs x_f , with $x_i = 0$, $\omega = 1$, and T = 0.5. (b) The ground-state probability density obtained from the Euclidean propagator. The blue dashed line is fit from the VFPG observations (blue dots), and we have used $\omega = 5$ and T = 1. In (a) and (b), the error bars denote the two standard deviations of the plotted quantity, evaluated from the training noise (standard deviation) of \mathcal{F}_{ϕ} . See the main text. (c) An example of the distribution of 10⁴ generated Feynman paths in the plane of the variational log probability and the action S_E , where both axes are shifted, with S_{Emin} the minimum action and $b = (-\mathcal{F}_{\phi} + S_{Emin})/\hbar$. The red solid line represents the exact exponential distribution (6). It is also the location of Feynman paths satisfying $F_{\phi}[x(\tau)] = \mathcal{F}_{\phi}$, and the color (from blue to red) indicates the distance between $F_{\phi}[x(\tau)]$ and its expectation value \mathcal{F}_{ϕ} . (d) An example of the free energy \mathcal{F}_{ϕ} (blue line) during the training process for estimating the quantum propagator in (a), where the width of the red shaded area denotes its two standard deviations in the generated path ensemble. The parameters used in (c) and (d) are the same as that in (a), except that the final positions are fixed at $x_f = 1$. Other representative parameters used are $N_{\tau} = 32$, $\hbar = 1$, and m = 1, respectively.

analytic exact results, which are obtained via Wick rotating the real-time propagator [2].

Note that we use a stochastic gradient descent algorithm to train the network (with Adam optimizer and learning rate $\eta = 1 \times 10^{-4}$), with 2048 data points sampled from the 2D latent space and a batch size of 128. Such a training algorithm causes statistical noises in the final result of \mathcal{F}_{ϕ} . Following the error analysis formula in Ref. [42] and with Eq. (14), the error bars (two standard deviations) in Figs. 3(a) and 3(b) are estimated from the training noises (standard deviation $\delta \mathcal{F}_{\phi}$) of the free energy, with $\delta K_E = K_E \ \delta \mathcal{F}_{\phi} / \sqrt{N_r}$, in $N_r = 10$ independent runs of the training with each up to 3000 epochs.

For the purpose of an intuitive understanding, we plot in Fig. 3(c) an example of the distribution of the generated FPs in the plane of the variational log probability vs the action. The approximate paths are distributed along the red solid line, which is the training target of the VFPG (in the case that

 \mathcal{F}_{ϕ} is unbiased for $-\hbar \ln Z$), i.e., the exponential distribution (6). The red solid line is also the location of FPs that exactly satisfy $F_{\phi}[x(\tau)] = \mathcal{F}_{\phi}$, and the color of the images of the generated FPs (from blue to red) encodes the distance $d = |F_{\phi}[x(\tau)] - \mathcal{F}_{\phi}|$ from the exact distribution. Note that both the two axes in Fig. 3(c) are shifted for a universal target distribution, where $S_{E\min}$ is the smallest Euclidean action and $b = (-\mathcal{F}_{\phi} + S_{E\min})/\hbar$. Figure 3(d) shows an example of the behavior of the free energy (blue line) during the training process (only 1000 epochs are shown), where the width of the red shaded area is its two-standard derivations in the generated path ensemble.

B. Double-well potential

In Fig. 4 we show the results from the VFPG for a double-well potential $V(x) = \alpha x^4 + \beta x^2$, with $\alpha = 0.05$, and



FIG. 4. Double-well potential $V(x) = \alpha x^4 + \beta x^2$. (a) Ground state wave function obtained from the Euclidean quantum propagator estimated via the variational Feynman path generator (VFPG), where the red dashed line is the result of the exact diagonalization. The green dashed line denotes the shifted and rescaled potential energy $\tilde{V}(x) = [V(x) + 5]/100$. (b) An example of the distribution of 10⁴ generated Feynman paths (with $x_i = x_f = x = 3$) in the plane of the log-probability and the action S_E , where both axes are shifted, with $S_{E\min}$ the minimum action and $b = (-\mathcal{F}_{\phi} + S_{E\min})/\hbar$. The red solid line represents the training target (6). It is also the location of Feynman paths that exactly satisfy $F_{\phi}[x(\tau)] = \mathcal{F}_{\phi}$, and the color (from blue to red) indicates the distance $d = |F_{\phi}[x(\tau)] - \mathcal{F}_{\phi}|$ from the exact distribution. Representative parameters used are: $\alpha = 0.05$, $\beta = -1$, $N_{\tau} = 32$, $\hbar = 1$, m = 1 and T = 2, respectively.

 $\beta = -1$. Figure 4(a) plots the estimated ground-state probability densities (blue dots), which again exhibit good agreement with the result from the exact diagonalization (red dashed). It has two peaks centered at the two respective minima of the double-well potential. Details for training and the error estimation formula are the same as those stated above for the harmonic oscillator. Again, the error bars (two standard deviations) here are estimated from the training noises (standard deviation $\delta \mathcal{F}_{\phi}$) of the free energy, in $N_r = 10$ independent runs of the training with each up to 3000 epochs, as in Figs. 3(a) and 3(b).

Shown in Fig. 4(b) is an example of the cluster of generated FPs in the plane of the variational log probability vs the action,

where the representative position value is taken as x = 3. As in Fig. 3(c), the generated paths are distributed around the training target of the VFPG (red solid line). Details and other physical parameters can be found in the caption of Fig. 4.

VI. DISCUSSION AND OUTLOOK

The propagator plays the role of the kernel for the quantum evolution. By virtue of the FP integral representation, here the kernel as well as the ground-state density can be efficiently estimated by generating FPs via parallel sampling. Instead of the generative modeling of quantum fields [38-45], here we focus on generating Feynman paths for quantum mechanical systems (i.e., trajectories or sequences as in Refs. [63,64]), with a fixed starting and terminating position. There are two advantages of the proposed VFPG: First, the sample complexity in the 2D latent space is lower compared with normalizing-flow models [39,41-45,49,50], which are bijections between the latent space and the path space and require a much higher latent dimension (same as the output dimension). Second, the sampling of the output paths is performed in a parallel manner and therefore is more efficient than Markov chain Monte Carlo (MCMC) methods [62], which iteratively generate path samples from the exact target path distribution in Eq. (6). Moreover, the Monte Carlo method cannot give a direct estimation of the partition function (i.e., the propagator) as in our work.

In the limit $\hbar \rightarrow 0$, the loss function in Eq. (11) will be dominated by the action term, and we find that Feynman paths generated by our RNN decoder (the VFPG) collapse to the minimal-action path, which can give a different perspective to understand the quantum-classical correspondence. In contrast, normalizing flows are bijections, which cannot be used for this purpose. As a matter of fact, we also produced ourselves some results on estimating the propagator of a harmonic oscillator with flow-based models (e.g., the real-valued non-volume preserving (real NVP) transformations in Ref. [50]), but so far the performance is not as good as the results presented in this paper (therefore those are not shown here). More elaborations in this respect and detailed comparisons between the current VFPG and normalizing flows are left for future work.

In addition, the variance reduction for discrete variable systems as in Ref. [57], inspired from reinforcement learning, does not apply well to the continuous variables considered here. An alternative systematic variance reduction is required, which is particularly important when the numerical value of the final free energy is small. Other possible generalizations may include extending the current results to higher spatial dimensions and to more complex systems, as well as investigating the generative modeling of FPs in the Lorentzian space-time.

VII. SUMMARY

We delivered the concept of FP generators for modeling Euclidean quantum mechanical FPs with *fixed endpoints* from the latent space. A variational realization with the recurrent neural network is performed, and as a proof-of-principle demonstration, the quantum propagators (or kernel functions) are efficiently estimated for both a harmonic oscillator and an anharmonic potential. Our work paves the way toward deep generative modeling of FPs with fixed starting and terminating points, respectively, and can provide a future fresh perspective to understand the quantum-classical correspondence through deep learning.

All the data and the code (in PYTHON and TENSORFLOW) for generating the results in this work are available upon request to the authors.

Note added. Recently, we noted the generative flow network (GFlowNet) [63,64], which generates a set of paths from a starting state to a terminating one on a graph, but with a different learning objective (compared to the path generator in our work) and for different tasks (in particular, molecule generation).

ACKNOWLEDGMENTS

The computation was performed in the RIKEN supercomputer (HOKUSAI) system. We acknowledge Tao Liu, Zheng-Yang Zhou, and Yu-Ran Zhang for helpful discussions. We thank Enrico Rinaldi for critical reading of the manuscript. F.N. is supported in part by Nippon Telegraph and Telephone Corporation (NTT) Research, the Japan Science and Technology Agency (JST) [via the Quantum Leap Flagship Program (Q-LEAP)], the Moonshot R&D Grant No. JPMJMS2061, the Japan Society for the Promotion of Science (JSPS) [via the Grants-in-Aid for Scientific Research (KAK-ENHI) Grant No. JP20H00134], the Army Research Office (ARO) (Grant No. W911NF-18-1-0358), the Asian Office of Aerospace Research and Development (AOARD) (via Grant No. FA2386-20-1-4069), and the Foundational Questions Institute Fund (FQXi) via Grant No. FQXi-IAF19-06.

APPENDIX A: MODELING THE DENSITY OF PATHS WITH VARIATIONAL RECURRENT NEURAL NETWORKS

Here the Feynman path $x(\tau)$ is approximately represented by a discrete N_{τ} -dimensional vector of a time series $\{x(\tau_k)\}_{k=0}^{N_{\tau}-1}$, with $\tau_k = k \times \delta \tau$, where $\delta \tau$ is a fixed time step. So the time interval $\tau \in [0, T]$ is sliced into N_{τ} discrete points with $T = (N_{\tau} - 1)\delta \tau$, while the spatial dimensions are continuous, i.e., not discretized. This is a time-domain lattice approximation of the Feynman path, which is suitable for machine learning models. It can be a good approximation when a small time step $\delta \tau$ is used, and the dimension N_{τ} of the time series is proportional to the total time T.

Here we model the output density $q_{\phi}[x(\tau)]$ of Feynman paths with a recurrent neural network (see Fig. 2 in the main text), where ϕ represents a set of neural-network parameters. In such a model, the probability density of a discrete path $\{x(\tau_k)\}_{k=0}^{N_r-1}$ is given by an autoregressive form [55]

$$q_{\phi}[x(\tau)] = \prod_{k=0}^{N_{\tau}-1} q_k[x(\tau_k)|x(\tau_{< k})], \qquad (A1)$$

- R. P. Feynman, Space-time approach to non-relativistic quantum mechanics, Rev. Mod. Phys. 20, 367 (1948).
- [2] R. MacKenzie, Path integral methods and applications, arXiv:quant-ph/0004090.

where the factor distribution $q_k[x(\tau_k)|x(\tau_{< k})]$ denotes the probability density of positions at τ_k conditioned on its previous times stamps, with < k denoting i = 0, 1, ..., k - 1, and the first and last positions are constrained to match the specified endpoints, x_i and x_f , respectively. This factor distribution at each time stamp is parametrized by a Gaussian mixture model (GMM), which is a universal approximator of densities [55]. In the GMM, we have

$$q_k[x(\tau_k)|x(\tau_{< k})] = \sum_{j=1}^{N_c} \gamma_j^{(k)} \,\mathcal{N}\big(x(\tau_k); \mu_j^{(k)}, \sigma_j^{(k)}\big), \quad (A2)$$

where the mixing weight $\gamma_j^{(k)}$, the mean $\mu_j^{(k)}$, and the standard deviation $\sigma_j^{(k)}$ of the component Gaussian are the output of the recurrent neural network at the *k*th unit [see Fig. 2(b) in the main text], and therefore they intrinsically have the conditional dependence on the values in the previous time stamps. The number of Gaussian components is N_c (we set N_c to be the batch size in our code). The mixing weight satisfies $0 \leq \gamma_j^{(k)} \leq 1$ and $\sum_{j=1}^{N_c} \gamma_j^{(k)} = 1$, which is realized by a softmax function. Compared to the MCMC sampling, here the generation of Feynman paths can be obtained by combinations of parallel samplings at each time stamps.

APPENDIX B: TOTAL LOSS FUNCTION

The training of the Feynman path generator involves the minimization of the KL divergence between the variational density $q_{\phi}[x(\tau)]$ of paths and the target distribution $P_E[x(\tau)]$, which equals to the (shifted) variational free energy in the main text. The constraints on the two endpoints are added to the KL loss as penalties. In the GMM formulation of $q_{\phi}[x(\tau)]$, the total loss can be written as

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{KL}} + \mathcal{L}_1 + \mathcal{L}_2, \tag{B1}$$

where

$$\mathcal{L}_{\mathrm{KL}} = \mathop{\mathbb{E}}_{x \sim q_{\phi}} F_{\phi}[x(\tau)]/\hbar, \tag{B2}$$

with $F_{\phi}[x(\tau)] = S_E[x(\tau)] + \hbar \log q_{\phi}[x(\tau)]$, and

$$\mathcal{L}_{1} = N_{\tau} \sum_{j=1}^{N_{c}} \left[\mu_{j}^{(k=0)} - x_{i} \right]^{2} + \left[\sigma_{j}^{(k=0)} \right]^{2},$$
$$\mathcal{L}_{2} = N_{\tau} \sum_{j=1}^{N_{c}} \left[\mu_{j}^{(k=N_{\tau}-1)} - x_{f} \right]^{2} + \left[\sigma_{j}^{(k=N_{\tau}-1)} \right]^{2}, \quad (B3)$$

where x_i and x_f are the initial and final positions of all possible paths, respectively.

- [3] R. P. Feynman, A. R. Hibbs, and D. F. Styer, *Quantum Mechan*ics and Path Integrals (Dover, New York, 2010).
- [4] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, Science 355, 602 (2017).

- [5] J. Carrasquilla and R. G. Melko, Machine learning phases of matter, Nat. Phys. 13, 431 (2017).
- [6] X. Gao and L.-M. Duan, Efficient representation of quantum many-body states with deep neural networks, Nat. Commun. 8, 662 (2017).
- [7] J. Chen, S. Cheng, H. Xie, L. Wang, and T. Xiang, Equivalence of restricted Boltzmann machines and tensor network states, Phys. Rev. B 97, 085104 (2018).
- [8] E. P. L. Van Nieuwenburg, Y.-H. Liu, and S. D. Huber, Learning phase transitions by confusion, Nat. Phys. 13, 435 (2017).
- [9] M. Schuld, I. Sinayskiy, and F. Petruccione, An introduction to quantum machine learning, Contemp. Phys. 56, 172 (2015).
- [10] V. Dunjko and H. J. Briegel, Machine learning & artificial intelligence in the quantum domain: A review of recent progress, Rep. Prog. Phys. 81, 074001 (2018).
- [11] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, Machine learning and the physical sciences, Rev. Mod. Phys. **91**, 045002 (2019).
- [12] I. A. Luchnikov, A. Ryzhov, P.-J. Stas, S. N. Filippov, and H. Ouerdane, Variational autoencoder reconstruction of complex many-body physics, Entropy 21, 1091 (2019).
- [13] N. Yoshioka and R. Hamazaki, Constructing neural stationary states for open quantum many-body systems, Phys. Rev. B 99, 214306 (2019).
- [14] F. Vicentini, A. Biella, N. Regnault, and C. Ciuti, Variational Neural-Network Ansatz for Steady States in Open Quantum Systems, Phys. Rev. Lett. **122**, 250503 (2019).
- [15] A. Nagy and V. Savona, Variational Quantum Monte Carlo Method with a Neural-Network Ansatz for Open Quantum Systems, Phys. Rev. Lett. 122, 250501 (2019).
- [16] M. J. Hartmann and G. Carleo, Neural-Network Approach to Dissipative Quantum Many-Body Dynamics, Phys. Rev. Lett. 122, 250502 (2019).
- [17] A. Berezutskii, M. Beketov, D. Yudin, Z. Zimborás, and J. D. Biamonte, Probing criticality in quantum spin chains with neural networks, J. Phys.: Complexity 1, 03LT01 (2020).
- [18] Q.-S. Tan, M. Zhang, Y. Chen, J.-Q. Liao, and J. Liu, Generation and storage of spin squeezing via learning-assisted optimal control, Phys. Rev. A 103, 032601 (2021).
- [19] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Neural-network quantum state tomography, Nat. Phys. 14, 447 (2018).
- [20] A. Melkani, C. Gneiting, and F. Nori, Eigenstate extraction with neural-network tomography, Phys. Rev. A 102, 022412 (2020).
- [21] A. M. Palmieri, E. Kovlakov, F. Bianchi, D. Yudin, S. Straupe, J. D. Biamonte, and S. Kulik, Experimental neural network enhanced quantum tomography, npj Quantum Inf. 6, 20 (2020).
- [22] S. Lohani, B. T. Kirby, M. Brodsky, O. Danaci, and R. T. Glasser, Machine learning assisted quantum state estimation, Mach. Learn.: Sci. Technol. 1, 035007 (2020).
- [23] M. Neugebauer, L. Fischer, A. Jäger, S. Czischek, S. Jochim, M. Weidemüller, and M. Gärttner, Neural-network quantum state tomography in a two-qubit experiment, Phys. Rev. A 102, 042604 (2020).
- [24] S. Ahmed, C. Sánchez Muñoz, F. Nori, and A. F. Kockum, Quantum State Tomography with Conditional Generative Adversarial Networks, Phys. Rev. Lett. **127**, 140502 (2021); Classification and reconstruction of optical quantum states with deep neural networks, Phys. Rev. Research **3**, 033278 (2021).

- [25] Y. Che, C. Gneiting, T. Liu, and F. Nori, Topological quantum phase transitions retrieved through unsupervised machine learning, Phys. Rev. B 102, 134213 (2020).
- [26] M. S. Scheurer and R.-J. Slager, Unsupervised Machine Learning and Band Topology, Phys. Rev. Lett. 124, 226401 (2020).
- [27] Y. Long, J. Ren, and H. Chen, Unsupervised Manifold Clustering of Topological Phononics, Phys. Rev. Lett. 124, 185501 (2020).
- [28] T. Mendes-Santos, X. Turkeshi, M. Dalmonte, and A. Rodriguez, Unsupervised Learning Universal Critical Behavior via the Intrinsic Dimension, Phys. Rev. X 11, 011040 (2021).
- [29] P. Ronhovde, S. Chakrabarty, D. Hu, M. Sahu, K. K. Sahu, K. F. Kelton, N. A. Mauro, and Z. Nussinov, Detection of hidden structures for arbitrary scales in complex physical systems, Sci. Rep. 2, 329 (2012).
- [30] H.-Y. Huang, R. Kueng, G. Torlai, V. V. Albert, and J. Preskill, Provably efficient machine learning for quantum many-body problems, arXiv:2106.12627.
- [31] R. Iten, T. Metger, H. Wilming, L. del Rio, and R. Renner, Discovering Physical Concepts with Neural Networks, Phys. Rev. Lett. **124**, 010508 (2020).
- [32] C. Wang, H. Zhai, and Y.-Z. You, Emergent Schrödinger equation in an introspective machine learning architecture, Sci. Bull. 64, 1228 (2019).
- [33] X. Liu, G. Zhang, J. Li, G. Shi, M. Zhou, B. Huang, Y. Tang, X. Song, and W. Yang, Deep Learning for Feynman's Path Integral in Strong-Field Time-Dependent Dynamics, Phys. Rev. Lett. 124, 113202 (2020).
- [34] M. Yin, X. Li, Y. Zhang, and S. Wang, On the mathematical understanding of ResNet with Feynman path integral, arXiv:1904.07568.
- [35] A. Barr, W. Gispen, and A. Lamacraft, Quantum ground states from reinforcement learning, in *Proceedings of the First Mathematical and Scientific Machine Learning Conference* (PMLR, Princeton University, Princeton, NJ, 2020), Vol. 107, pp. 635–653.
- [36] Z. Jin, J. Y.-Y. Lin, and S.-F. Li, Learning Principle of Least Action with Reinforcement Learning, arXiv:2011.11891.
- [37] T. Mendes-Santos, A. Angelone, A. Rodriguez, R. Fazio, and M. Dalmonte, Intrinsic dimension of path integrals: Datamining quantum criticality and emergent simplicity, PRX Quantum 2, 030332 (2021).
- [38] S.-H. Li and L. Wang, Neural Network Renormalization Group, Phys. Rev. Lett. 121, 260601 (2018).
- [39] M. S. Albergo, G. Kanwar, and P. E. Shanahan, Flow-based generative models for Markov chain Monte Carlo in lattice field theory, Phys. Rev. D 100, 034515 (2019).
- [40] H.-Y. Hu, S.-H. Li, L. Wang, and Y.-Z. You, Machine learning holographic mapping by neural network renormalization group, Phys. Rev. Research 2, 023369 (2020).
- [41] G. Kanwar, M. S. Albergo, D. Boyda, K. Cranmer, D. C. Hackett, S. Racanière, D. J. Rezende, and P. E. Shanahan, Equivariant Flow-Based Sampling for Lattice Gauge Theory, Phys. Rev. Lett. **125**, 121601 (2020).
- [42] K. A. Nicoli, C. J. Anders, L. Funcke, T. Hartung, K. Jansen, P. Kessel, S. Nakajima, and P. Stornati, Estimation of Thermodynamic Observables in Lattice Field Theories with Deep Generative Models, Phys. Rev. Lett. **126**, 032001 (2021).

- [43] M. S. Albergo, G. Kanwar, S. Racanière, D. J. Rezende, J. M. Urban, D. Boyda, K. Cranmer, D. C. Hackett, and P. E. Shanahan, Flow-based sampling for fermionic lattice field the-[54]
- ories, Phys. Rev. D 104, 114507 (2021).
 [44] M. Medvidovic, J. Carrasquilla, L. E. Hayward, and B. Kulchytskyy, Generative models for sampling of lattice field theories, arXiv:2012.01442.
- [45] M. Caselle, E. Cellini, A. Nada, and M. Panero, Stochastic normalizing flows as non-equilibrium transformations, arXiv:2201.08862.
- [46] Y.-L. Lin and F. Nori, Quantum interference in superconducting wire networks and Josephson junction arrays: An analytical approach based on multiple-loop Aharonov-Bohm Feynman path integrals, Phys. Rev. B 65, 214504 (2002); Strongly Localized Electrons in a Magnetic Field: Exact Results on Quantum Interference and Magnetoconductance, Phys. Rev. Lett. 76, 4580 (1996); Quantum interference on the kagome lattice, Phys. Rev. B 50, 15953 (1994).
- [47] For the Euclidean case, the most probable path should also satisfy $\delta^2 S > 0$.
- [48] A. van den Oord, N. Kalchbrenner, and K. Kavukcuoglu, Pixel recurrent neural networks, in *Proceedings of the 33rd International Conference on Machine Learning* (PMLR, New York, 2016), Vol. 48, pp. 1747–1756.
- [49] L. Dinh, D. Krueger, and Y. Bengio, NICE: Non-linear independent components estimation, arXiv:1410.8516.
- [50] L. Dinh, J. Sohl-Dickstein, and Y. Bengio, Density estimation using real NVP, arXiv:1605.08803.
- [51] D. P. Kingma and M. Welling, Auto-encoding variational Bayes, arXiv:1312.6114.
- [52] C. Doersch, Tutorial on variational autoencoders, arXiv:1606.05908.
- [53] N. C. Chung, B. Mirza, H. Choi, J. Wang, D. Wang, P. Ping, and W. Wang, Unsupervised classification of multi-omics data

during cardiac remodeling using deep learning, Methods **166**, 66 (2019).

- [54] J. Chung, K. Kastner, L. Dinh, K. Goel, A. C. Courville, and Y. Bengio, A recurrent latent variable model for sequential data, in *Advances in Neural Information Processing Systems* (Curran Associates, New York, 2015), Vol. 28, pp. 2980–2988.
- [55] I. J. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, MA, 2016), http://www.deeplearningbook.org.
- [56] S. Hochreiter and J. Schmidhuber, Long short-term memory, Neural Comput. 9, 1735 (1997).
- [57] D. Wu, L. Wang, and P. Zhang, Solving Statistical Mechanics Using Variational Autoregressive Networks, Phys. Rev. Lett. 122, 080602 (2019).
- [58] L. Wang, Y. Jiang, L. He, and K. Zhou, Continuous-mixture autoregressive networks for efficient variational calculation of many-body systems, arXiv:2005.04857.
- [59] F. Noé, S. Olsson, J. Köhler, and H. Wu, Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning, Science 365, eaaw1147 (2019).
- [60] M. Hibat-Allah, E. M. Inack, R. Wiersema, R. G. Melko, and J. Carrasquilla, Variational neural annealing, Nat. Mach. Intell. 3, 952 (2021).
- [61] S. K. Vadlamani, T. P. Xiao, and E. Yablonovitch, Physics successfully implements Lagrange multiplier optimization, Proc. Natl. Acad. Sci. USA 117, 26639 (2020).
- [62] K. Cranmer, S. Golkar, and D. Pappadopulo, Inferring the quantum density matrix with machine learning, arXiv:1904.05903.
- [63] Y. Bengio, T. Deleu, E. J. Hu, S. Lahlou, M. Tiwari, and E. Bengio, GFlowNet foundations, arXiv:2111.09266.
- [64] E. Bengio, M. Jain, M. Korablyov, D. Precup, and Y. Bengio, Flow network based generative models for non-iterative diverse candidate generation, arXiv:2106.04399.