






Learning minimal representations of many-body physics from snapshots of a quantum simulator

Frederik Møller ^{1,2} Gabriel Fernández-Fernández ³ Thomas Schweigler,² Paulin de Schoulepnikoff ⁴
Jörg Schmiedmayer ^{2,*} and Gorka Muñoz-Gil ^{4,†}

¹*Institute of Science and Technology Austria, Am Campus 1, 3400 Klosterneuburg, Austria*

²*Vienna Center for Quantum Science and Technology (VCQ), Atominstitut, TU Wien, Vienna, Austria*

³*ICFO—Institut de Ciències Fotòniques, The Barcelona Institute of Science and Technology,
Av. Carl Friedrich Gauss 3, 08860 Castelldefels (Barcelona), Spain*

⁴*Department for Theoretical Physics, University of Innsbruck, Technikerstrasse. 21a, A-6020 Innsbruck, Austria*



(Received 26 September 2025; accepted 25 February 2026; published 29 April 2026)

Analog quantum simulators provide access to many-body dynamics beyond the reach of classical computation. However, extracting physical insights from experimental data is often hindered by measurement noise, limited observables, and incomplete knowledge of the underlying microscopic model. Here, we develop a machine learning approach based on a variational autoencoder (VAE) to analyze interference measurements of tunnel-coupled one-dimensional Bose gases, which realize the sine-Gordon quantum field theory. Trained in an unsupervised manner, the VAE learns a minimal latent representation that strongly correlates with the equilibrium control parameter of the system. Applied to nonequilibrium protocols, the latent space uncovers signatures of frozen-in solitons following rapid cooling, and reveals anomalous postquench dynamics not captured by conventional correlation-based methods. These results demonstrate that generative models can extract physically interpretable variables directly from noisy and sparse experimental data, providing complementary probes of equilibrium and nonequilibrium physics in quantum simulators. More broadly, our work highlights how machine learning can supplement established field-theoretical techniques, paving the way for scalable, data-driven discovery in quantum many-body systems.

DOI: [10.1103/r7pj-gl7r](https://doi.org/10.1103/r7pj-gl7r)

I. INTRODUCTION

Quantum analog simulators provide a powerful platform to investigate complex many-body systems by engineering physical setups whose intrinsic dynamics implement the desired model Hamiltonian [1–3]. In these experiments, engineered quantum systems evolve naturally under the governing physical laws, and relevant observables are extracted through measurement. This approach has been fruitfully applied across diverse platforms, from ultracold and neutral atoms [4–6] to trapped ions [7,8] and superconducting circuits [9], enabling insight into strongly correlated phases, nonequilibrium dynamics, and topological phenomena beyond the reach of classical computation [10]. In terms of continuous systems, a particularly well-studied example is the realization of one-dimensional (1D) superfluids, where coherent condensates or coupled arrays serve as testbeds for low-energy quantum field theories describing collective fluctuations [11–16].

The readout from quantum simulators, however, poses unique challenges compared to digital simulations. For instance, measurement constraints limit access to only a subset of observables [17], often further degraded by finite imaging resolution and noise. Consequently, extracting comprehensive information demands both tailored measurement protocol and advanced data analysis methods [18–21]. These include full counting statistics [22,23], higher-order correlation functions [13,24,25], and increasingly, machine learning (ML) techniques that leverage modern statistical inference to infer hidden degrees of freedom and dynamical rules from experimental trajectories [26–32].

Another central difficulty arises when the simulator’s effective Hamiltonian or noise processes are not fully known a priori or when experimental datasets are limited, making direct parameter estimation unreliable or biased. Therefore, methods that are agnostic to detailed microscopic assumptions yet sufficiently expressive and interpretable to reveal underlying physics are highly desirable [33–35].

In this work, we address this challenge by employing a custom variational autoencoder (VAE) trained on the relative phase field measurements of a quantum simulator composed of two tunnel-coupled 1D ultracold Bose gases. The proposed VAE, an unsupervised generative model [36] equipped with an autoregressive architecture [37], identifies a minimal parametric representation of the stochastic process governing phase fluctuations of the system without relying on prior knowledge of the Hamiltonian or noise sources. We

*Contact author: schmiedmayer@atomchip.org

†Contact author: munoz.gil.gorka@gmail.com

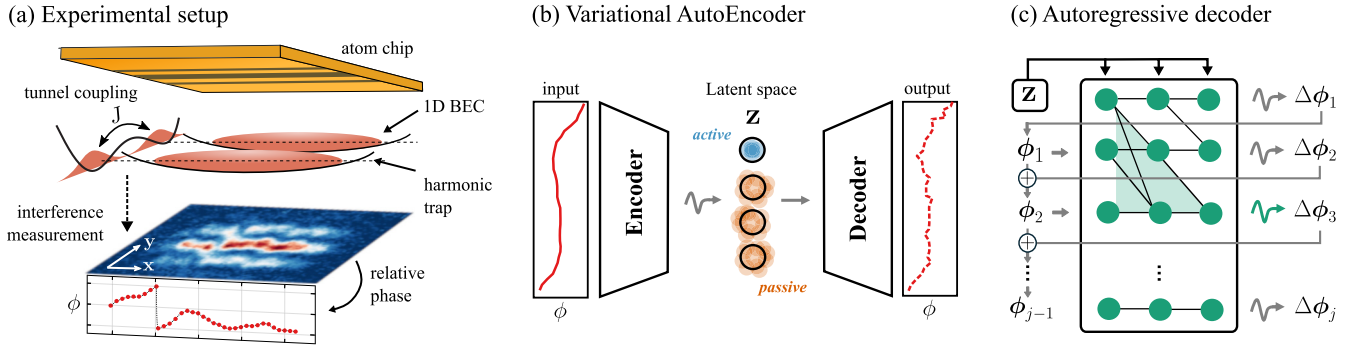


FIG. 1. Experimental and computational pipeline. (a) Two quasi-one-dimensional superfluids in a magnetic double-well created by an atom chip with tunable tunnel coupling J . The relative phase ϕ is extracted from absorption imaging of the matter-wave interference pattern. (b) The encoder maps each phase trajectory to a lower-dimensional latent vector \mathbf{z} . The model compresses the input into a minimal set of neurons (active), effectively suppressing noncontributing ones (passive). The decoder reconstructs then the phase trajectory from \mathbf{z} . (c) The decoder predicts the conditional probability of the next phase increment $\Delta\phi_j$ based on \mathbf{z} and the previous phases $\phi_{i < j}$. Its autoregressive nature ensures causal prediction (green shaded area)

demonstrate that the representation learned by the model strongly correlates with physical parameters controlling the effective sine-Gordon field theory realized in the experiment. Furthermore, the model's representation provides a sensitive probe of nonequilibrium phenomena, including the detection of topological soliton defects and the characterization of postquench dynamics where conventional observables fail to discriminate equilibrium from nonequilibrium states.

II. METHODS

A. Experimental realization of the sine-Gordon quantum simulator

We study two quasi-one-dimensional, tunnel-coupled ^{87}Rb superfluids prepared in a magnetic double-well potential created by an atom chip (see Fig. 1 and also Refs. [13,38] for further details). At low energies, their relative degrees of freedom are captured by the sine-Gordon Hamiltonian [39]:

$$H_{\text{SG}} = \int_0^L dx \left[g \delta\rho^2 + \frac{\hbar^2 n}{4m} (\partial_x \phi)^2 - 2\hbar J n \cos(\phi) \right], \quad (1)$$

where $\delta\rho(x)$ and $\phi(x)$ are canonically conjugate relative density and phase fields, m is the atomic mass, g is the 1D contact-interaction parameter, $n(x)$ is the linear density, which varies slowly due to the weak axial confinement, and J is the tunnel coupling. The relative phase field $\phi(x)$ is accessed through matter-wave interference [40], recorded with a CCD camera of pixel size $\Delta x \approx 2 \mu\text{m}$, and henceforth treated as a discrete quantity ϕ_j , which we refer to as a *phase trajectory*. This destructive measurement requires repeated preparation and imaging, causing that each realization yields a single sample trajectory subject to shot-to-shot fluctuations of both temperature and atom number arising from technical imperfections.

The sine-Gordon model provides a unifying description of one-dimensional systems: It can be viewed as a Luttinger liquid [41,42] perturbed by a cosine interaction potential, admits solitonic excitations familiar from classical nonlinear dynamics [43], and has long served as a minimal quantum field theory relevant in condensed matter [44], gauge theories

[45], and cosmology [46]. In our setting, the strength of the cosine interaction is tunable by controlling the double-well separation, and thus J , through radio-wave dressing, thereby enabling simulation of various interaction regimes.

In thermal equilibrium, the coherence properties are governed by the dimensionless coupling $Q = \lambda_T/l_J$, with $l_J = \sqrt{\hbar/(4mJ)}$ the Josephson length and $\lambda_T = 2\hbar^2 n/(mk_B T)$ the thermal coherence length [47]. Systems with equal Q exhibit identical behavior when expressed in units of λ_T , so Q fully controls the phase statistics. While not directly measurable, Q is related to the coherence factor $\langle \cos \phi \rangle$ (see Appendix C). In the strong-coupling regime ($Q \gg 1$), tunneling locks the relative phase near $\phi \approx 0$, yielding $\langle \cos \phi \rangle \approx 1$ and long-range correlations. When thermal fluctuations dominate ($Q \lesssim 1$), the relative phase is randomized, $\langle \cos \phi \rangle \rightarrow 0$, and correlations decay exponentially over λ_T [48].

Using the transfer matrix formalism [49,50], for homogeneous condensates in thermal equilibrium, the fluctuating thermal phase field of the sine-Gordon model can be expressed using the stochastic Itô equation [51,52], where position takes the role usually taken by time:

$$d\phi = A[\phi]dx + \sqrt{2D}dx \mathcal{N}(0, 1), \quad (2)$$

where A is the deterministic Q -dependent drift term obtained by diagonalizing the transfer operator (see Appendix B) and $D = 2\lambda_T^{-1}$ is the diffusion constant. Heuristically, the drift term encodes the restoring force of the cosine potential through the tunnel coupling, while D encodes the thermal randomization of the phase. In practice, the measured phase trajectories deviate from the ideal Itô process of Eq. (2), as the finite resolution and coarse-graining of the imaging system smooth the underlying fluctuations. Further, the harmonic confinement of the magnetic trap introduces a weak density variation; we focus on the central $\approx 70 \mu\text{m}$ (35 pixels) of relative-phase images, where the linear density n varies between 70 and $140 \mu\text{m}^{-1}$.

B. Variational autoencoder architecture

To extract low-dimensional, physically interpretable representations from our experimental data, we employed a VAE

architecture [36], which has proven effective in capturing the essential features of stochastic physical dynamics [53,54]. A VAE comprises three main components: an encoder, a latent space, and a decoder (see Fig. 1). The encoder, built here via a stack of convolutional and feed-forward layers, compresses each input phase trajectory into a small set of latent variables. Such latent space is parametrized by a multivariate Gaussian distribution, where the mean μ_i and variance σ_i of each latent variable z_i are predicted by the encoder. Samples \mathbf{z} drawn from this distribution are then passed to the decoder.

The decoder reconstructs the input from the latent variables. Importantly, as shown in Refs. [53,54], the decoder must faithfully reproduce the statistical properties of the training data to yield a meaningful latent representation. Conceptually, the decoder learns the physical mapping from the system's underlying degrees of freedom to the experimentally observed trajectories. To achieve this, we design it to model the conditional probability of the next phase increment given the current phase history and the latent variables,

$$p_\theta(\Delta\phi_{j+1} | \phi_j, \phi_{j-1}, \dots, \mathbf{z}), \quad (3)$$

where θ denotes the set of trainable parameters of the decoder. In this way, the model is explicitly encoding the stochastic nature of the dynamics in its output distribution. In our system, we expect the next increment $\Delta\phi_{j+1}$ to depend directly on the present phase ϕ_j , given that the deterministic drift term in the effective Itô process is a function of the current phase [see Eq. (2)]. Consequently, we adopt an autoregressive structure conditioned on past phase values, rather than past increments. In practice, the conditional distribution is parametrized using a WaveNet architecture [37], which employs dilated convolutions throughout the previous phase points [shaded region in Fig. 1(c)] to correctly capture correlations of the data.

The VAE is trained by maximizing the Evidence Lower Bound (ELBO) [36]. For a dataset of N phase trajectories $\phi^{(n)}$, the loss function reads as

$$\mathcal{L} = \frac{1}{N} \sum_{n=1}^N \mathbb{E}_q[\ln p_\theta(\phi^{(n)} | \mathbf{z})] - \beta \text{KL}[q_{\tilde{\theta}}(\mathbf{z} | \phi^{(n)}) || p(\mathbf{z})]. \quad (4)$$

This expression contains two contributions: (1) a reconstruction loss, which maximizes the likelihood of the input data under the decoder p_θ ; and (2) a regularization term, which aligns the latent distribution $q_{\tilde{\theta}}(z_i)$ with the standard Gaussian prior $p(z_i) = \mathcal{N}(0, 1)$ through the Kullback-Leibler (KL) divergence. Here, $q_{\tilde{\theta}}$ denotes the encoder with trainable parameters $\tilde{\theta}$. The weighting factor β balances the relative strength of the two terms. Following Ref. [55], we adopt here the Total Correlation VAE formulation, which further encourages conditional independence among latent variables (see Appendix A), typically leading to more interpretable latent spaces.

In practice, the competition between the reconstruction and regularization terms in Eq. (4) drives the model to encode only the minimal set of latent neurons required to accurately reconstruct the data. Those neurons that do not meaningfully contribute to the reconstruction collapse to the prior, i.e., $\mu_i \rightarrow 0$ and $\sigma_i \rightarrow 1$, becoming *passive* neurons. This process effectively filters out noise and redundancies in the representation. When the dataset is governed by a small

number of hidden degrees of freedom—as in many physical systems—the *active* latent variables naturally align with these underlying parameters [56,57], providing a minimal and physically interpretable coordinate system for the observed dynamics. As we will see below, this mechanism enables the VAE to distill complex, noisy phase trajectories into a reduced description that correlates strongly with the control parameters of the sine-Gordon model.

III. RESULTS

A. Training VAEs on experimental interference measurements

We train the autoregressive VAE via Ref. [58] on experimental equilibrium phase trajectories obtained across several tunnel couplings J [13]. The datasets contain around 10^3 experimental trajectories per J . Starting from a six-dimensional latent space, the VAE robustly collapses all but one neurons to the uninformed prior. This is shown in Fig. 2(a), where, for 20 different training initializations, the variances σ_i of the collapsed neurons converge to one (orange lines) while the remaining active neuron remains at values $\sigma_a \ll 1$ (blue line). Hence, for the rest of this work, we will only analyze the active neuron and also consider that $z_a \leftarrow \mu_a$ due to its small variance. This result identifies a single dominant degree of freedom in the input data, consistent with the experiment effectively varying one control parameter (the coupling J) and with the sine-Gordon equilibrium process being governed by the single dimensionless parameter Q . That only one neuron is active despite experimental imperfections and imaging noise highlights the ability of the VAE to robustly learn the physical model underlying the input data even in such challenging conditions. Indeed, when analyzing the values of the active latent neuron for the training experimental trajectories, we observe a broad distribution for all J s, reflecting shot-to-shot parameter fluctuations of temperature and linear density.

B. Interpreting the physical model learned in the latent space

To interpret the physics encoded in the active latent neuron, we use the generative capabilities of the model and sample from the decoder ensembles of trajectories conditioned on a wide range of latent values z_a , keeping in all cases the value of the collapsed neurons to zero.

As shown in Fig. 2(b) (gray line), the coherence factor $\langle \cos \phi \rangle$ computed from the generated samples depends continuously on z_a , even when the original training data only covered, in principle, a discrete set of values (blue horizontal lines). Notably, z_a dependence resembles the theoretical relation between $\langle \cos \phi \rangle$ and Q , up to a sign and an overall scale (see Fig. 5), a feature that was not enforced but autonomously learned by the VAE.

To further assess the relation between the active neuron and the generated phases' properties, we analyze their distribution at representative values of z_a [marked by colored points in Fig. 2(b)]. Figure 2(c) shows the expected behavior: At negative z_a (related to high coherence) phases concentrate near $\phi = 0$, while at larger z_a the distribution broadens and weight near $\pm\pi$ appears, indicating trajectories that cross into neighboring vacua of the interaction potential.

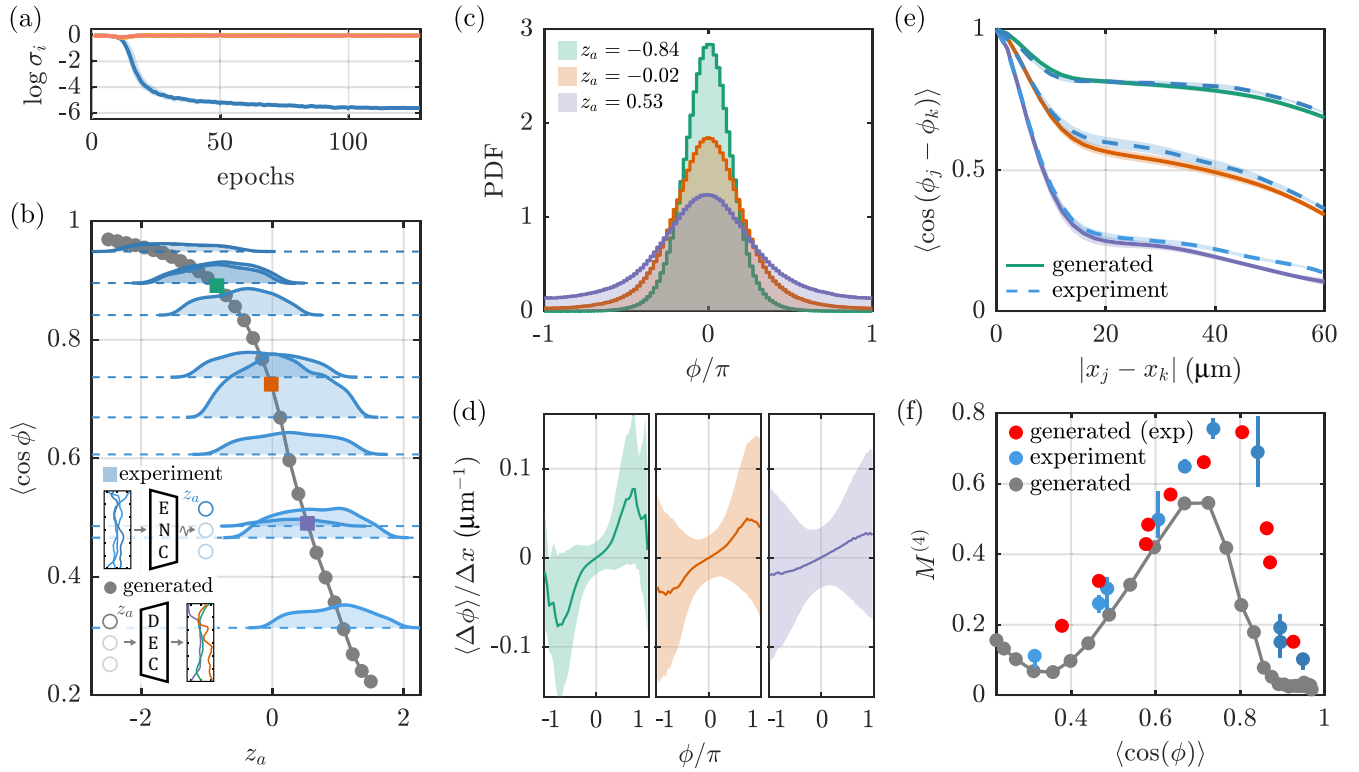


FIG. 2. Interpreting the latent space of the trained VAE. (a) Median variance of active (blue) and passive (orange) latent neurons over 20 trained models. (b) Coherence factor $\langle \cos \phi \rangle$ computed from trajectories generated at fixed values of the active neuron z_a (gray points) and from subsets of the experimental training data (blue horizontal markers). For each subset, the distribution of activation obtained by passing the data through the encoder is shown vertically; the distribution area reflects the varying number of experimental samples. (c) Histograms of generated phases ϕ for three representative z_a values. (d) Mean (solid) and standard deviation (shaded) of generated phase increments $\Delta \phi_{j+1}$ as a function of ϕ_j . (e) Circular two-point correlations as a function of distance for the same three representative z_a values and for experimental subsets with matching coherence factors. Shaded areas indicate standard error. (f) Connected fourth moment $M^{(4)}$ vs coherence factor for experimental data (blue) and samples generated from the decoder either from activation z_a (gray) or values sampled from the distribution of experimental activations (red). Experimental error bars indicate 80% bootstrap confidence intervals.

The decoder output constitutes an effective Itô-type process. To assess the nature of such process and its relation to the Itô process in Eq. (2), we analyze the generated increments $\Delta \phi$ as a function of the current phase $\phi \pmod{2\pi}$. Importantly, this analysis is made possible by the model's ability to generate a large number of trajectories, providing sufficient statistics that the limited experimental dataset alone could not offer. As shown in Fig. 2(d), the mean increment produces a restoring drift comparable to Eq. (2), with its strength increasing for negative z_a . Interestingly, the diffusion term, emerging here from the increment variances (shaded areas), shows a notable ϕ dependence, which is absent in the ideal Itô process. We attribute this dependence to the experimental imaging process, effectively mixing the terms of the underlying process. Indeed, simulations of the ideal process accounting for finite imaging resolution and coarse-graining [38,59] reproduce this effect qualitatively (see Appendix B).

An essential feature of this process is the emergence of long-range phase order associated with finite tunnel coupling, which is directly reflected in the two-point correlation function [48]. As illustrated in Fig. 2(e), these correlations decay exponentially at short distances, then approach a nonzero plateau set by Q . At larger separations, a slow residual

decay appears, reflecting the weak density inhomogeneity of the trapped condensates. Phase trajectories generated by the decoder reproduce these behaviors, confirming that the latent representation captures both short-range thermal fluctuations and long-range phase locking.

Beyond such two-point correlations, a defining feature of this system is the presence of high-order non-Gaussian phase correlations between the coupled condensates [13,60]. To quantify these correlations, we compute the normalized connected fourth moment

$$M^{(4)} = \frac{\sum_{\mathbf{j}} | \langle (\phi_{j_1} - \phi_0) \cdots (\phi_{j_4} - \phi_0) \rangle_c |}{\sum_{\mathbf{j}} | \langle (\phi_{j_1} - \phi_0) \cdots (\phi_{j_4} - \phi_0) \rangle |}, \quad (5)$$

where $\mathbf{j} = \{j_1, j_2, j_3, j_4\}$ and $j_1 \leq j_2 \leq j_3 \leq j_4$, ϕ_0 indicates a fixed reference point, and $\langle \cdot \rangle_c$ is the connected correlator. In Fig. 2(f), we plot $M^{(4)}$ versus $\langle \cos \phi \rangle$ for both the experimental phase trajectories (blue) and trajectories generated by the VAE for the same range of z_a shown in Fig. 2(b). Remarkably, both ensembles showcase the same qualitative behavior: In the low-coherence limit, $\langle \cos \phi \rangle \rightarrow 0$, the cosine interaction potential in the sine-Gordon Hamiltonian becomes negligible and the system reduces to a quadratic Tomonaga-Luttinger liquid

[41,42]. Hence, higher-order connected correlations, and thus $M^{(4)}$, vanish. At the opposite extreme, $\langle \cos \phi \rangle \rightarrow 1$, the tunnel coupling strongly locks the relative phase near $\phi = 0$. In this regime, the cosine potential can be expanded to quadratic order, yielding the massive Klein-Gordon model [61], which is again Gaussian and produces vanishing $M^{(4)}$. Only at intermediate coherence factors, where the cosine potential is neither negligible nor purely harmonic, do genuine non-Gaussian states contribute, leading to finite values of $M^{(4)}$.

Interestingly, the generated trajectories produce systematically lower $M^{(4)}$ than experiment. We attribute such behavior to the sampling of the generated trajectories being performed at fixed z , whereas each experimental dataset spans a broad range of latent activations [see Fig. 2(b)], reflecting variation of physical parameters within the set. Indeed, when sampling latent variables from the experimentally observed distribution of encoded activation values [blue distributions in Fig. 2(b)], and generating trajectories from the decoder accordingly, the resulting connected fourth moment $M^{(4)}$ quantitatively reproduces the experimental trend.

C. Detecting topological defects via latent representations

Having established that the VAE faithfully reproduces the equilibrium, measurement-limited sine-Gordon statistics, we next use the previously trained model as a discriminative probe for nonequilibrium states. The goal is to assess to what extent a model trained in a well-characterized regime can be repurposed to probe more complex and unknown scenarios. For this purpose, we consider a second experimental dataset in which the final cooling ramp is ten times faster than in the equilibrium runs [13]. Such rapid cooling can freeze in domain regions that settle into different minima of the cosine interaction potential. Hence, the phase must wind by 2π across the domain boundaries, creating topological solitons of the sine-Gordon model. In this regime, the solitons are static on experimental timescales and the field fluctuates around them, representing excitations above topologically distinct “false” vacua [62]. Although rare at strong coupling, even a few of such defects markedly alter higher-order moments [13].

We input the fast-cooled trajectories to the equilibrium-trained VAE’s encoder and compare, in Figs. 3(a)–3(d), the resulting latent activation histograms (red) to those from equilibrium data with similar coherence factors (blue). For weak coupling [Figs. 3(a)–3(c)], the two datasets are nearly indistinguishable, as thermally excited solitons are already common in equilibrium. However, at stronger coupling [Fig. 3(d)], the fast-cooled data develop a clear second peak at positive activations in addition to the equilibriumlike negative peak. As analyzed in Fig. 2(c), negative activations correspond to trajectories localized near the main vacuum $\phi \approx 0$, whereas positive activations signal large phase increments, which, at strong coupling, is associated with solitonic defects. To confirm such behavior, we examine samples assigned to both peaks, both from its original absorption images [Figs. 3(e) and 3(f)] and their corresponding phase trajectories [Figs. 3(g) and 3(h)]. As shown, while samples with negative z_a stay in the main vacuum, positive values lead to trajectories exhibiting localized 2π rotations that connect neighboring vacua, with their absorption images revealing corresponding phase slips

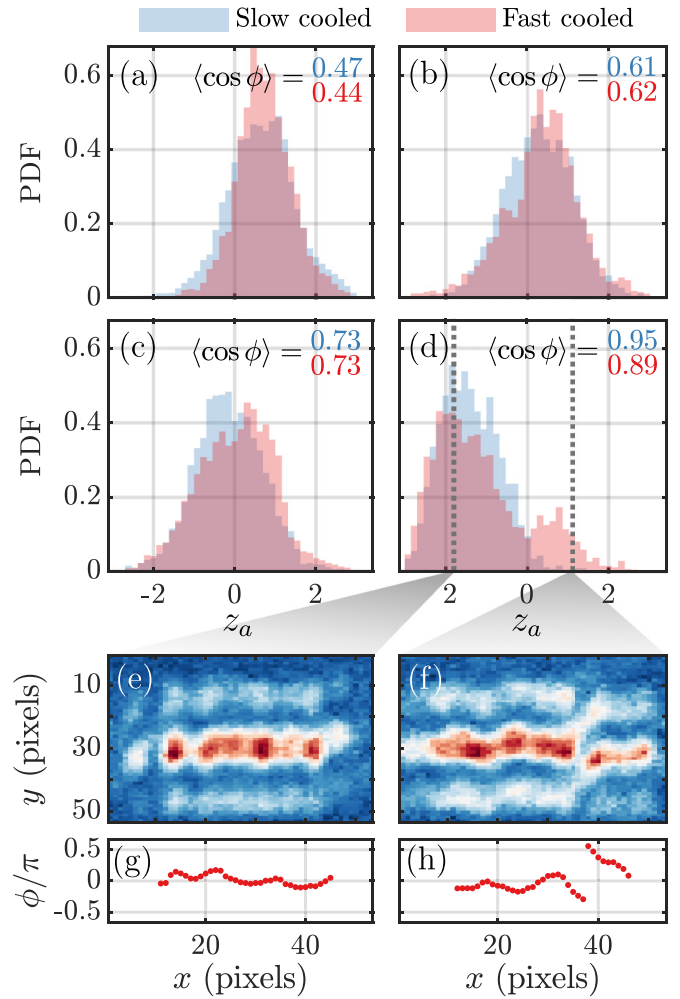


FIG. 3. Latent-space discrimination of solitonic defects in fast-cooled data. (a)–(d) Activation histograms of the active latent neuron for equilibrium (slow-cooled) and fast-cooled trajectories with matched coherence factors. (e)–(h) Representative phase trajectories and corresponding absorption images from the negative-activation (equilibriumlike) and positive-activation (solitonic excitation) peaks.

in the interference fringes. Thus, the encoder trained solely on equilibrium data can identify topologically nontrivial excitations in the quantum simulator.

D. Probing quench dynamics beyond equilibrium field theory

Finally, we turn to a third experimental protocol involving a sudden quench of the tunnel coupling [38,63]. In this sequence, the system is first prepared in a fully uncoupled state $J = 0$. The magnetic trapping potential is then rapidly quenched to a configuration with finite coupling J , and measurements are taken at a series of evolution times t after the quench.

Following the quench, the coherence factor immediately increases sharply, then continues to grow gradually [see Fig. 4(a) inset], qualitatively similar to theoretical predictions [64]. At first glance, conventional analysis based on the connected fourth moment $M^{(4)}$ [Eq. (5)] yields a

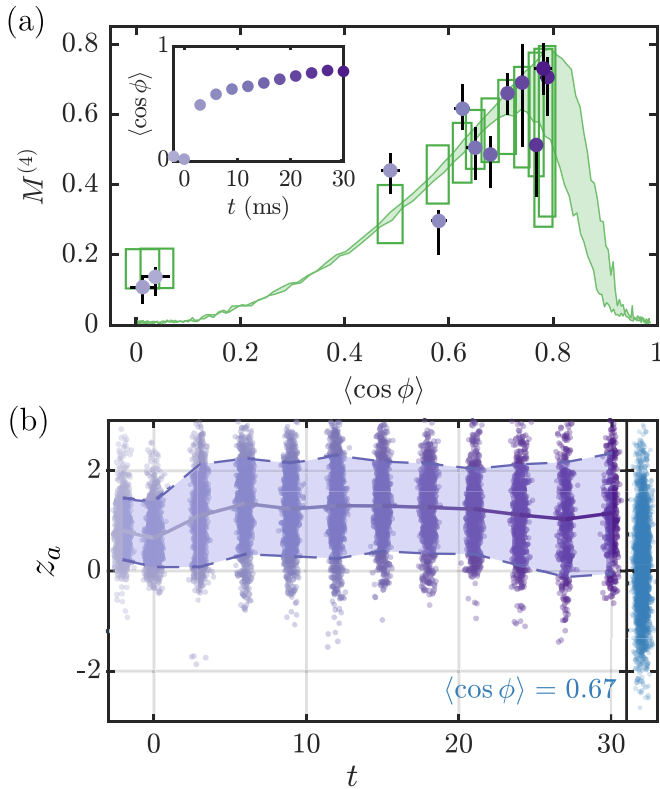


FIG. 4. Latent-space analysis of quench dynamics. (a) Fourth connected moment $M^{(4)}$ for postquench experimental trajectories, compared to equilibrium sine-Gordon predictions from Itô-process sampling with imaging effects. The shaded area spans $\lambda_T = 20 \mu\text{m}$ (top) to $\lambda_T = 30 \mu\text{m}$ (bottom). Rectangles show theory for $\lambda_T = 20 \mu\text{m}$ with the same number of samples as in the experiment; widths and heights indicate 80% confidence intervals. Inset shows measured coherence factor as a function of time. (b) Latent neuron activation vs evolution time t for the quench data, with solid line and shaded area giving the median and 80% confidence interval. Blue distribution: equilibrium subset with coherence factor matched to the final quench value.

somewhat surprising result: The postquench trajectories appear to follow the same functional dependence as equilibrium sine-Gordon ensembles (see Appendix C), with an effective coupling strength that evolves in time. Taken at face value, this would suggest that the system rapidly equilibrates to a thermal state despite the substantial energy injected by the quench. Yet such an interpretation is counterintuitive. Starting from $J = 0$ means that the relative phase is initially random across $[-\pi, \pi)$, whereby quenching to finite interaction excites high-energy modes, likely beyond the low-energy sector usually captured by the sine-Gordon description. In this regime, mixing between the relative and common degrees of freedom is expected to play an important role, potentially invalidating the assumption of an equilibrium sine-Gordon model [63,65,66].

The latent-space analysis of our VAE provides a complementary perspective [Fig. 4(b)]. Before the quench, the activations are centered around positive values of z_a , consistent with weak interactions and large thermal fluctuations. After the quench, rather than drifting toward negative values

as one would expect from the increased coherence factor, the latent activations instead remain pinned around large positive values, even shifting slightly upward. This behavior lies outside the range spanned by equilibrium datasets (rightmost, blue points), indicating that the postquench trajectories differ markedly from equilibrium statistics. Inspecting the measured trajectories confirm this: Large-amplitude fluctuations seeded by the initial uncoupled state persist well beyond the quench, more pronounced than in equilibrium or fast-cooled conditions (see Fig. 7). At larger scales, the phase distribution evolves from a nearly uniform profile to one peaked around $\phi \sim 0$ with small side peaks near $\pm 2\pi$, indicating the presence of solitonlike windings similar to those observed after fast cooling.

One possible interpretation is that the system enters a *prethermal* state [22,24,67]. Prethermalization refers to the formation of a long-lived, quasistationary regime where a restricted set of observables rapidly relax to stationary values, while full thermalization is delayed by approximate conservation laws or weak couplings to additional degrees of freedom. Such behavior is particularly common in nearly integrable systems, where weak integrability breaking can confine dynamics near a constrained manifold for extended times [68,69]. In our setting, the persistence of stable latent activations in Fig. 4(b), together with gradually evolving phase distributions, is consistent with this scenario: The relative phase field could quickly settle into a prethermal equilibrium characterized by large fluctuations and possible topological excitations, while full thermalization proceeds only slowly through coupling to the common degrees of freedom of the two superfluids [64].

At the same time, alternative explanations cannot be excluded. The discrepancy between correlation-function analysis and latent-space signatures may also arise from experimental limitations, such as ambiguities in phase unwrapping for increments near $\Delta\phi \sim \pm\pi$, which can propagate errors through entire trajectories (see Ref. [38] and Appendix C). Although we assessed the robustness of our results by discarding problematic trajectories, such effects remain a potential source of bias. Furthermore, given the substantial energy injected by the quench, it is plausible that the postquench dynamics simply lie outside the sine-Gordon regime altogether, meaning that neither $M^{(4)}$ nor the latent representation can be mapped to equilibrium physics.

Regardless of the precise interpretation, our analysis highlights the complementary nature of the two approaches. Whereas $M^{(4)}$ suggests an equilibriumlike description, the VAE exposes latent features inconsistent with equilibrium ensembles, pointing to either prethermalization or a breakdown of the sine-Gordon model. This divergence underscores the usefulness of ML-based methods as additional probes of quantum simulators: They provide access to trajectory-level features that conventional observables may average out, thereby offering handles on out-of-equilibrium physics. In regards of the system studied here, and to establish unambiguous conclusions, further experiments are needed. For instance by introducing smaller quenches that inject less energy into the system, or with protocols that perform opposite quenches, from strong to weak tunnelings. Moreover, recent techniques now permit simultaneous extraction of both relative and

common phases [20], which would allow a direct test of whether coupling between the two sectors underlies the observed dynamics. Extending our VAE to a multichannel architecture could then provide a natural framework for analyzing such joint data.

IV. DISCUSSION

In this work, we demonstrated that a variational autoencoder can extract a low-dimensional and physically interpretable representation of equilibrium phase trajectories from interference snapshots of tunnel-coupled 1D superfluids, a quantum simulator of the sine-Gordon field theory. By balancing its reconstruction error with the latent space regularization, the model autonomously identified the minimal parameter set characterizing the system, namely, the dimensionless coupling Q , without any prior knowledge on the underlying physical model. For this, it is crucial that the generative capabilities of the VAE's decoder faithfully reproduce the processes' statistics, such that the latent space correctly encodes its physical parameters and not other spurious information. We achieved this by employing an autoregressive decoder, ensuring that the resulting phase trajectories can converge to the expected class of stochastic processes. We showed that the decoder correctly recovers the main properties of the process, namely, its long-range correlations and non-Gaussian higher-order moments.

We then applied the equilibrium-trained model to nonequilibrium regimes. In the case of rapid cooling, the latent representation revealed soliton defects frozen during the cooling ramp, enabling their identification at the level of individual trajectories. In a different scenario, where the condensates undergo a sudden quench of the tunnel coupling, the VAE exposed a qualitatively distinct latent signature: Despite the connected fourth moments appearing equilibriumlike, the latent activation departed strongly from the equilibrium sine-Gordon snapshots, indicating that the postquench dynamics may lie outside the effective low-energy theory. Together, these results show that the VAE provides a sensitive probe of both equilibrium and nonequilibrium many-body physics from limited and noisy measurements.

More broadly, this work establishes VAEs as powerful tools for the analysis of analog quantum simulation. They address practical experimental challenges such as finite resolution, uncontrolled fluctuations, and scarce sampling, while remaining agnostic to microscopic modeling assumptions. This capability paves the way toward data-driven discovery of physical regimes, particularly in emerging generations of quantum simulators with unprecedented control and tunability. While latent representations may not provide a complete description on their own, their integration with expert physical knowledge and complementary ML methods such as symbolic regression offers a scalable route toward interpretable characterization of unknown quantum dynamics.

ACKNOWLEDGMENTS

We thank Sebastian Erne and Igor Mazets for helpful discussions and sharing codes for the transfer matrix sampling. This research was funded in part by the European Research

Council: ERC Advanced Grant ‘‘Emergence in Quantum Physics’’ (EmQ) under Grant Agreement No. 101097858 and ERC Advanced Grant ‘‘Artificial agency and learning in quantum environments’’ (QuantAI) under Grant Agreement No. 101055129. This work was also supported by the Austrian Science Fund (FWF) (SFB BeyondC F7102, 10.55776/F71). G.F.-F. acknowledges the European Research Council AdG NOQIA; MCIN/AEI [PGC2018-0910.13039/501100011033, CEX2019-000910-S/10.13039/501100011033, Plan National FIDEUA PID2019-106901GB-I00, Plan National STAMEENA PID2022-139099NB, I00, project funded by MCIN/AEI/10.13039/501100011033 and by the ‘‘European Union NextGenerationEU/PRTR’’ (PRTR-C17.I1), FPI]; QUANTERA DYNAMITE PCI2022-132919 under Grant Agreement No. 101017733; Ministry for Digital Transformation and of Civil Service of the Spanish Government through the QUANTUM ENIA project call—Quantum Spain project, and by the European Union through the Recovery, Transformation and Resilience Plan—NextGenerationEU within the framework of the Digital Spain 2026 Agenda; Fundació Cellex; Fundació Mir-Puig; Generalitat de Catalunya (European Social Fund FEDER and CERCA program); Barcelona Supercomputing Center MareNostrum (FI-2023-3-0024); (HORIZON-CL4-2022-QUANTUM-02-SGA PASQuanS2.1, 101113690, EU Horizon 2020 FET-OPEN OPTologic, Grant No. 899794, QU-ATTO, 101168628), EU Horizon Europe Program (This project has received funding from the European Union's Horizon Europe research and innovation program under Grant Agreement No. 101080086 NeQST); ICFO Internal ‘‘QuantumGaudi’’ project. This research was funded in whole or in part by the Austrian Science Fund (FWF) [10.55776/COE1] through the Cluster of Excellence quantA (Quantum Science Austria).

The views and opinions expressed in this article are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council—neither the European Union nor the granting authority can be held responsible for them.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available. The data are available from the authors upon reasonable request.

APPENDIX A: MACHINE LEARNING MODEL

The architecture of our model follows closely the convolutional VAE with an autoregressive decoder introduced in Ref. [53]. For completeness, we summarize the specific design and hyperparameters used in this work.

1. Architecture

The model consists of three components: (1) an encoder that compresses each one-dimensional input trajectory of length $L = 35$ into a latent representation, (2) a latent space made of six probabilistic neurons $\mathbf{z} \in \mathbb{R}^6$, sampled via the reparametrization trick, and (3) a decoder consisting of a convolutional upsampling stage followed by an autoregressive

TABLE I. Summary of model architecture and hyperparameters. $B = 512$ denotes the batch size. The WaveNet outputs two parameters $(\tilde{\mu}_j, \tilde{\sigma}_j)$, parametrizing a Gaussian distribution.

Layer type	Output size
Input	$B \times 1 \times L$
Encoder	
4 × 1D conv. (kernel = 3 and stride = 1)	$B \times 16 \times (L - 8)$
Adaptive pooling (avg + max)	$B \times 16 \times (16 + 16)$
Flatten	$B \times 512$
MLP (200 and 100 neurons)	$B \times 100$
Latent distribution	$B \times 12$
Latent layer ($ z = 6$)	$B \times 6$
Decoder	
MLP (100, 200, and 512 neurons)	$B \times 512$
Reshape	$B \times 16 \times 32$
Interpolation	$B \times 16 \times (L - 8)$
3 × 1D transposed conv.	$B \times 16 \times L$
1D transposed conv. ($N_c = 6$)	$B \times 6 \times L$
WaveNet (kernel = 4 and layers = 3)	$B \times 2 \times L$

module. The architecture of the model is summarized in Table I.

The encoder comprises four 1D convolutional layers with 16 channels each, followed by an adaptive pooling and a two-layer multilayer perceptron (200 and 100 neurons). Each latent neuron is parametrized as a Gaussian distribution with mean and variance predicted by the encoder.

The decoder mirrors the encoder structure with multilayer perceptron (MLP) layers (100, 200, and 512), channel reshaping, interpolation, and transposed convolutions. The upsampled latent representation conditions all layers of the autoregressive module, which is implemented as a WaveNet [37]. This module uses a single stack of three layers of dilated convolutions with kernel size 4 that yield a receptive field of $RF = 23$. We employ zero padding so that the autoregressive outputs have the same length as the input trajectories.

2. Loss function

The model is trained by maximizing the ELBO within a Total Correlation VAE (TC-VAE) objective [55]. Given the increments of an input phase trajectory $\Delta\phi$, the overall loss reads

$$\mathcal{L} = \mathbb{E}_q[\ln p_\theta(\Delta\phi|\mathbf{z})] - \beta \text{KL}[q(\mathbf{z}|\phi)\|p(\mathbf{z})] - \gamma \text{TC}(\mathbf{z}) - \alpha \text{MI}(\mathbf{z}; \phi), \quad (\text{A1})$$

with fixed parameters $\beta = 3$, $\alpha = 10^{-4}$, and $\gamma = 0.1$, and where MI is the mutual information between the data and latent variables based on q .

Here, each input to the model is a full wrapped phase trajectory $\phi = (\phi_1, \dots, \phi_L)$ with $\phi_j \in [-\pi, \pi]$, while the decoder predicts the sequence of phase increments $\Delta\phi_j = \phi_j - \phi_{j-1}$. This choice avoids issues with phase wrapping at the $\pm\pi$ boundary, since the increment variables are unbounded. Thus, the reconstruction term is evaluated with respect to the probability distribution of increments rather than the wrapped phases themselves.

The decoder parametrizes the conditional probability of the increments in an autoregressive manner,

$$p_\theta(\Delta\phi) = \prod_{j=1}^{L-1} p_\theta(\Delta\phi_j | \phi_{<j}, \mathbf{z}), \quad (\text{A2})$$

where each conditional distribution is taken to be Gaussian,

$$p_\theta(\Delta\phi_j | \phi_{<j}, \mathbf{z}) = \mathcal{N}(\tilde{\mu}_j(\mathbf{z}, \phi_{<j}), \tilde{\sigma}_j^2(\mathbf{z}, \phi_{<j})). \quad (\text{A3})$$

The reconstruction loss is therefore the negative log-likelihood of the predicted increments,

$$-\ln p_\theta(\Delta\phi|\mathbf{z}) = -\sum_{i=1}^N \sum_{j=1}^{L-1} \ln \mathcal{N}(\Delta\phi_j^{(i)} | \tilde{\mu}_j^{(i)}, \tilde{\sigma}_j^{2(i)}), \quad (\text{A4})$$

which, for a single data point, expands to

$$-\ln \mathcal{N}(\Delta\phi_j | \tilde{\mu}_j, \tilde{\sigma}_j^2) = \ln(\tilde{\sigma}_j \sqrt{2\pi}) + \frac{(\Delta\phi_j - \tilde{\mu}_j)^2}{2\tilde{\sigma}_j^2}. \quad (\text{A5})$$

The minimum of this loss is achieved when the predicted mean $\tilde{\mu}_j$ matches the observed increment $\Delta\phi_j$, while a smaller predicted variance $\tilde{\sigma}_j^2$ reduces the loss further. In this way, the network learns to output both the most likely increment and an uncertainty estimate that reflects the stochastic nature of the experimental trajectories. The TC and MI penalty terms act to encourage disentangled, minimally correlated latent variables, while the KL term aligns the approximate posterior with the unit Gaussian prior.

While the present system is governed by a single effective parameter, the approach naturally extends to systems with multiple independent control parameters. In such cases, the VAE is expected to activate multiple latent dimensions, provided the reconstruction and regularization terms of Eq. (A1) are properly balanced.

3. Training

The model is trained on experimental equilibrium trajectories across several tunnel couplings. Each training consists of 128 epochs with a maximum learning rate of 5×10^{-4} , optimized using Adam [70] and the one-cycle learning rate [71,72] scheduler provided by `fastai` [73]. Batch size is $B = 512$. All model parameters are initialized using Kaiming initialization [74] (fan-out mode), with latent variances initialized to zero.

The specific values of the regularization parameters used throughout this work were determined empirically by training models with different fixed values of the regularization weights from the start of training. We selected values for which the reconstruction quality remained high while the latent space exhibited a clear collapse of inactive dimensions. Importantly, we verified that the qualitative structure of the latent space and all physical conclusions reported in this work are robust against variations of training hyperparameters. In particular, models trained with $\beta \in [1, 5]$, different random initializations, and different latent dimensionalities consistently yield a single active latent variable correlated with the sine-Gordon control parameter and




FIG. 5. Coherence factor $\langle \cos \phi \rangle$ as a function of effective interaction strength Q of the sine-Gordon model. Phase trajectories are sampled using the Itô process [Eq. (2)]. The experimental imaging process is modeled by convolving ϕ with a Gaussian, then coarse-graining. On the top x axis, the coherence factor as a function of the active latent variable z is plotted for phase trajectories generated with the trained model.

reproduce the same separation between equilibrium and quench data.

APPENDIX B: TRANSFER MATRIX FORMALISM AND ITÔ PROCESS FOR THE SINE-GORDON MODEL

The stochastic Itô process used to sample equilibrium phase trajectories of the sine-Gordon model [Eq. (2)] follows from the transfer matrix formalism (see, e.g., Refs. [13,51]). More generally, this approach applies to a broad class of one-dimensional bosonic systems in thermal equilibrium, for which the relevant low-energy degrees of freedom can be described by a classical field theory with local interactions. Under these conditions, the spatial correlations of the field can be mapped onto a stochastic process, where position plays the role of an effective time variable. Here, we briefly summarize the steps specific to the sine-Gordon case, referring to the general derivation to these references.

For two tunnel-coupled one-dimensional superfluids, the classical Hamiltonian reduces under the standard approximations (decoupling of relative and common degrees of freedom) to the sine-Gordon Hamiltonian given in the main text. Within the transfer matrix formalism, thermal correlation functions are generated by the auxiliary operator

$$\hat{K}_{\text{SG}} = \frac{1}{\lambda_T} \left(-2\partial_\phi^2 - 2\frac{Q^2}{8}(\cos \phi - 1) \right), \quad (\text{B1})$$

where $\lambda_T = 2\hbar^2 n / (mk_B T)$ is the thermal coherence length and Q is the effective interaction strength. The equilibrium distribution of relative phases is determined by the ground state $\Psi_0(\phi)$ of \hat{K}_{SG} , which obeys a Mathieu-type equation.

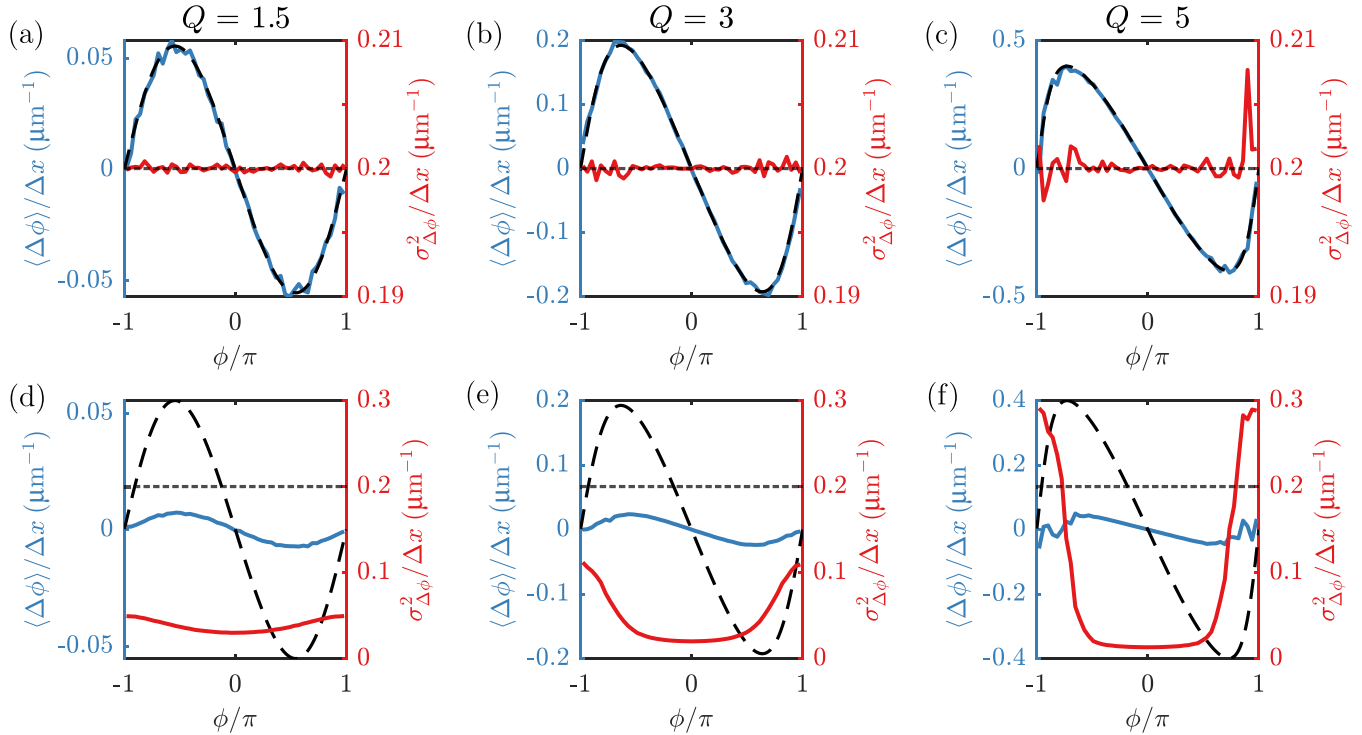


FIG. 6. Phase increments from the Itô process of Eq. (2). The left vertical axis shows the mean phase increment $\langle \Delta \phi \rangle$ as a function of the instantaneous phase ϕ (solid blue), compared to the theoretical drift term A [Eq. (B3)] (long-dashed black). The right vertical axis shows the standard deviation of the phase increments (solid red), together with the theoretical diffusion coefficient $2D = 4/\lambda_T$ (short-dashed black). Panels (a)–(c) (top row) present results obtained directly from the Itô process for three values of the effective interaction strength Q . Panels (d)–(f) (bottom row) show the corresponding results after accounting for imaging effects, implemented by convolving the phase trajectories with a Gaussian point-spread function to model finite resolution, followed by coarse-graining.

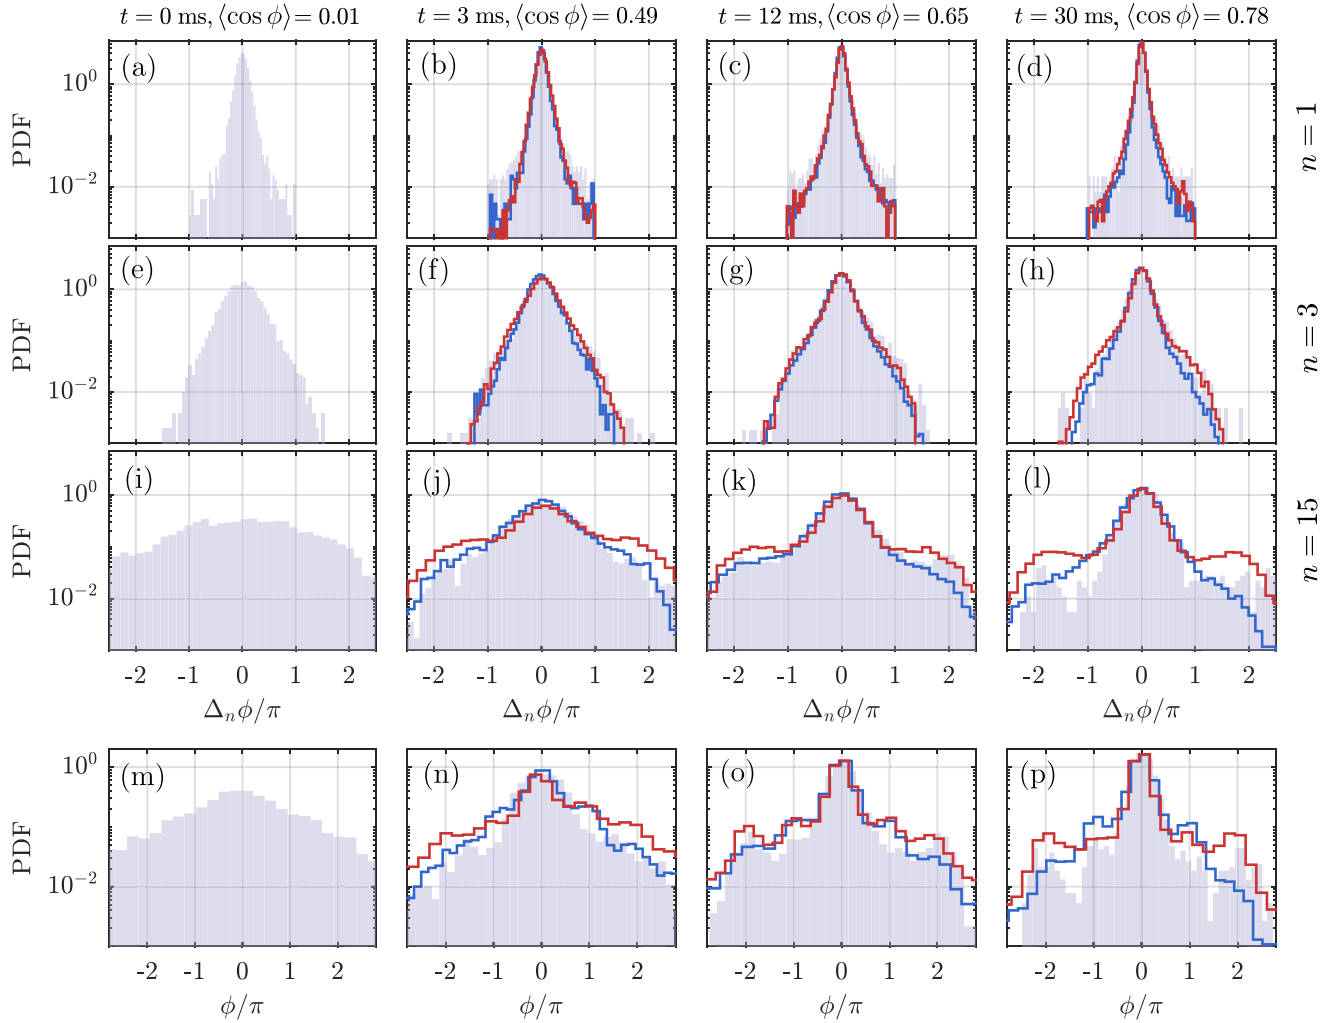


FIG. 7. Distributions of phase increments and phase values following quench of the tunnel coupling J (purple shaded) compared to slow-cooled (blue lines) and fast-cooled (red lines) data at similar coherence factor $\langle \cos \phi \rangle$. The phase increments are computed at various distances, such that $\Delta_n \phi_i = \phi_i - \phi_{i-n}$. (a)–(d) Phase increments between neighboring points $n = 1$. (e)–(h) Phase increments at 3 pixels separation $n = 3$. (i)–(l) Phase increments at 15 pixels separation $n = 15$. (m)–(p) Distributions of phase values ϕ .

By general arguments (see Ref. [52]), the corresponding Fokker-Planck equation is equivalent to an Itô process of the form

$$d\phi = A[\phi] dx + \sqrt{2D} dx \mathcal{N}(0, 1), \quad (\text{B2})$$

with diffusion constant $D = 2/\lambda_T$. The deterministic drift is given by

$$A[\phi] = -2D \frac{\partial}{\partial \phi} \ln |\Psi_0(\phi)|, \quad (\text{B3})$$

which encodes the effect of the tunnel coupling through the cosine potential.

Thus, equilibrium phase trajectories of the sine-Gordon model can be generated efficiently by integrating the Itô process using the drift determined from $\Psi_0(\phi)$ and the universal diffusion constant set by λ_T . In practice, we use numerical solutions of the Mathieu equation to evaluate $\Psi_0(\phi)$ and hence the drift $A[\phi]$ [47].

APPENDIX C: IMAGING EFFECTS AND SIMPLIFIED MODELING

The relative phase field $\phi(x)$ is experimentally accessed through matter-wave interference imaged on a CCD camera with finite pixel size $\Delta x \approx 2 \mu\text{m}$. This process inevitably distorts the underlying phase fluctuations of the condensates. Two main effects are relevant: (1) the imaging system has a finite point spread function (PSF), effectively a Gaussian blurring with width $\sigma \sim 3 \mu\text{m}$, which smooths the phase trajectory over neighboring points; and (2) coarse-graining occurs when the CCD integrates over the size of the individual pixels, further reducing resolution at length scales smaller than Δx . Together, these effects suppress short-wavelength fluctuations of the phase field, modify the measured correlation functions, and can mix drift and diffusion contributions of the Itô process for the sine-Gordon phase field [Eq. (2)], leading to an apparent phase-dependent diffusion term [see Fig. 2(d)].

A comprehensive treatment of the imaging process has been developed in Ref. [59] and in the detailed discussion of Ref. [38], where the full absorption imaging dynamics,

integration along the imaging axis, and technical noise sources are modeled. Such approaches are necessary for precise quantitative comparison of microscopic observables, in particular when extracting parameters like the effective diffusion constant from first principles.

For the purposes of the present work, however, a simplified but well-controlled approximation is sufficient. We account for imaging resolution by convolving the generated phase trajectories with a Gaussian of width σ corresponding to the effective PSF, followed by coarse-graining to the experimental pixel size Δx . This procedure reproduces the main observable consequences of imaging: The change of the coherence factor is modest (Fig. 5), but higher-order quantities, such as the variance of increments, are more strongly affected. In particular, the convolution mixes neighboring phases such that the measured Itô process acquires an effective ϕ -dependent diffusion (see Fig. 6), consistent with the trends seen in experiment.

Because our VAE is trained directly on experimental trajectories, which already include these imaging effects, the simplified model is only used when making equilibrium checks against theory (e.g., Itô sampling). In these cases, the Gaussian-convolution approach captures the essential impact

of imaging without requiring the full machinery of absorption-imaging simulations.

Additional effects/defects may arise when the relative phase is extracted from the absorption imaging (see Ref. [38] for full details). To extract the phase, the interference pattern is fitted with a cosine function with a Gaussian envelope, from which the phase modulo 2π is extracted. In order to ensure a continuous phase trajectory, it is assumed that the phase difference between neighboring pixels does not exceed π . However, in the presence of large fluctuations, found, for instance, following the quench of the tunnel coupling, errors in the extracted phase trajectory may occur. In Fig. 7, distributions of the phase increments before and following the quench of the tunnel coupling are shown; indeed, no increments with amplitude greater than π is found. Comparing with phase increment distributions of the slow and fast cooled data, we do observe a slightly larger number of increments near the values $\phi = \pm 1$, which could indicate an increased number of errors in the extracted phase trajectories. Nevertheless, compared to the total number of increments, the number of potential errors is rather low. Indeed, filtering out profiles containing increments $|\Delta\phi| \geq 0.9\pi$ does not significantly change our results.

-
- [1] D. Jaksch, C. Bruder, J. I. Cirac, C. W. Gardiner, and P. Zoller, Cold bosonic atoms in optical lattices, *Phys. Rev. Lett.* **81**, 3108 (1998).
- [2] I. M. Georgescu, S. Ashhab, and F. Nori, Quantum simulation, *Rev. Mod. Phys.* **86**, 153 (2014).
- [3] D. J. Gross, A. Sevriin, and P. Zoller, *Physics of Quantum Information, Proceedings of the 28th Solvay Conference on Physics* (World Scientific Publishing Company, Singapore, 2023).
- [4] I. Bloch, J. Dalibard, and S. Nascimbène, Quantum simulations with ultracold quantum gases, *Nat. Phys.* **8**, 267 (2012).
- [5] S. Ebadi, T. T. Wang, H. Levine, A. Keesling, G. Semeghini, A. Omran, D. Bluvstein, R. Samajdar, H. Pichler, W. W. Ho, S. Choi, S. Sachdev, M. Greiner, V. Vuletić, and M. D. Lukin, Quantum phases of matter on a 256-atom programmable quantum simulator, *Nature (London)* **595**, 227 (2021).
- [6] D. Bluvstein, H. Levine, G. Semeghini, T. T. Wang, S. Ebadi, M. Kalinowski, A. Keesling, N. Maskara, H. Pichler, M. Greiner, V. Vuletić, and M. D. Lukin, A quantum processor based on coherent transport of entangled atom arrays, *Nature (London)* **604**, 451 (2022).
- [7] R. Blatt and C. F. Roos, Quantum simulations with trapped ions, *Nat. Phys.* **8**, 277 (2012).
- [8] C. Monroe, W. C. Campbell, L.-M. Duan, Z.-X. Gong, A. V. Gorshkov, P. W. Hess, R. Islam, K. Kim, N. M. Linke, G. Pagano, P. Richerme, C. Senko, and N. Y. Yao, Programmable quantum simulations of spin systems with trapped ions, *Rev. Mod. Phys.* **93**, 025001 (2021).
- [9] A. A. Houck, H. E. Türeci, and J. Koch, On-chip quantum simulation with superconducting circuits, *Nat. Phys.* **8**, 292 (2012).
- [10] A. J. Daley, I. Bloch, C. Kokail, S. Flannigan, N. Pearson, M. Troyer, and P. Zoller, Practical quantum advantage in quantum simulation, *Nature (London)* **607**, 667 (2022).
- [11] M. A. Cazalilla, R. Citro, T. Giamarchi, E. Orignac, and M. Rigol, One dimensional bosons: From condensed matter systems to ultracold gases, *Rev. Mod. Phys.* **83**, 1405 (2011).
- [12] C. Gross and I. Bloch, Quantum simulations with ultracold atoms in optical lattices, *Science* **357**, 995 (2017).
- [13] T. Schweigler, V. Kasper, S. Erne, I. Mazets, B. Rauer, F. Cataldini, T. Langen, T. Gasenzer, J. Berges, and J. Schmiedmayer, Experimental characterization of a quantum many-body system via higher-order correlations, *Nature (London)* **545**, 323 (2017).
- [14] M. Tajik, I. Kukuljan, S. Sotiriadis, B. Rauer, T. Schweigler, F. Cataldini, J. Sabino, F. Møller, P. Schüttelkopf, S.-C. Ji, D. Sels, E. Demler, and J. Schmiedmayer, Verification of the area law of mutual information in a quantum field simulator, *Nat. Phys.* **19**, 1022 (2023).
- [15] M. Tajik, M. Gluza, N. Sebe, P. Schüttelkopf, F. Cataldini, J. Sabino, F. Møller, S.-C. Ji, S. Erne, G. Guarnieri, S. Sotiriadis, J. Eisert, and J. Schmiedmayer, Experimental observation of curved light-cones in a quantum field simulator, *Proc. Natl. Acad. Sci. USA* **120**, e2301287120 (2023).
- [16] B. Doyon, S. Gopalakrishnan, F. Møller, J. Schmiedmayer, and R. Vasseur, Generalized hydrodynamics: A perspective, *Phys. Rev. X* **15**, 010501 (2025).
- [17] T. Barthel and J. Lu, Fundamental limitations for measurements in quantum many-body systems, *Phys. Rev. Lett.* **121**, 080406 (2018).
- [18] P. Kunkel, M. Prüfer, S. Lannig, R. Rosa-Medina, A. Bonnin, M. Gärttner, H. Strobel, and M. K. Oberthaler, Simultaneous readout of noncommuting collective spin observables beyond the standard quantum limit, *Phys. Rev. Lett.* **123**, 063603 (2019).

- [19] M. Gluza, T. Schweigler, B. Rauer, C. Krumnow, J. Schmiedmayer, and J. Eisert, Quantum readout for cold atomic quantum simulators, *Commun. Phys.* **3**, 12 (2020).
- [20] T. Murtadho, F. Cataldini, S. Erne, M. Gluza, M. Tajik, J. Schmiedmayer, and N. H. Y. Ng, Measurement of total phase fluctuation in cold-atomic quantum simulators, *Phys. Rev. Res.* **7**, L022031 (2025).
- [21] M. Prüfer, Y. Minoguchi, T. Zhang, Y. Kuriatnikov, M. V. Ramana, and J. Schmiedmayer, Quantum-limited generalized measurement for tunnel-coupled condensates, *Phys. Rev. Lett.* **133**, 250403 (2024).
- [22] M. Gring, M. Kuhnert, T. Langen, T. Kitagawa, B. Rauer, M. Schreitl, I. Mazets, D. A. Smith, E. Demler, and J. Schmiedmayer, Relaxation and prethermalization in an isolated quantum system, *Science* **337**, 1318 (2012).
- [23] J. F. Wienand, S. Karch, A. Impertrio, C. Schweizer, E. McCulloch, R. Vasseur, S. Gopalakrishnan, M. Aidelsburger, and I. Bloch, Emergence of fluctuating hydrodynamics in chaotic quantum systems, *Nat. Phys.* **20**, 1732 (2024).
- [24] T. Langen, S. Erne, R. Geiger, B. Rauer, T. Schweigler, M. Kuhnert, W. Rohringer, I. E. Mazets, T. Gasenzer, and J. Schmiedmayer, Experimental observation of a generalized Gibbs ensemble, *Science* **348**, 207 (2015).
- [25] T. V. Zache, T. Schweigler, S. Erne, J. Schmiedmayer, and J. Berges, Extracting the field theory description of a quantum many-body system from experimental data, *Phys. Rev. X* **10**, 011020 (2020).
- [26] J. Wang, S. Paesani, R. Santagati, S. Knauer, A. A. Gentile, N. Wiebe, M. Petruzzella, J. L. O'Brien, J. G. Rarity, A. Laing, and M. G. Thompson, Experimental quantum Hamiltonian learning, *Nat. Phys.* **13**, 551 (2017).
- [27] A. Bohrdt, C. S. Chiu, G. Ji, M. Xu, D. Greif, M. Greiner, E. Demler, F. Grusdt, and M. Knap, Classifying snapshots of the doped Hubbard model with machine learning, *Nat. Phys.* **15**, 921 (2019).
- [28] R. Ott, T. V. Zache, M. Prüfer, S. Erne, M. Tajik, H. Pichler, J. Schmiedmayer, and P. Zoller, Hamiltonian learning in quantum field theories, *Phys. Rev. Res.* **6**, 043284 (2024).
- [29] H. Lange, M. Kebrič, M. Buser, U. Schollwöck, F. Grusdt, and A. Bohrdt, Adaptive quantum state tomography with active learning, *Quantum* **7**, 1129 (2023).
- [30] P. Schüttelkopf, M. Tajik, N. Bazhan, F. Cataldini, S.-C. Ji, J. Schmiedmayer, and F. Møller, Characterising transport in a quantum gas by measuring Drude weights, *Science* **391**, 290 (2026).
- [31] H. Lange, G. Bornet, G. Emperauger, C. Chen, T. Lahaye, S. Kienle, A. Browaeys, and A. Bohrdt, Transformer neural networks and quantum simulators: A hybrid approach for simulating strongly correlated systems, *Quantum* **9**, 1675 (2025).
- [32] O. Simard, A. Dawid, J. Tindall, M. Ferrero, A. M. Sengupta, and A. Georges, Learning interactions between Rydberg atoms, *PRX Quantum* **6**, 030324 (2025).
- [33] S. J. Wetzel, S. Ha, R. Iten, M. Klopotek, and Z. Liu, Interpretable machine learning in physics: A review, *arXiv:2503.23616*.
- [34] C. Miles, R. Samajdar, S. Ebadi, T. T. Wang, H. Pichler, S. Sachdev, M. D. Lukin, M. Greiner, K. Q. Weinberger, and E.-A. Kim, Machine learning discovery of new phases in programmable quantum simulator snapshots, *Phys. Rev. Res.* **5**, 013026 (2023).
- [35] R. Verdel, V. Vitale, R. K. Panda, E. D. Donkor, A. Rodriguez, S. Lannig, Y. Deller, H. Strobel, M. K. Oberthaler, and M. Dalmonte, Data-driven discovery of statistically relevant information in quantum simulators, *Phys. Rev. B* **109**, 075152 (2024).
- [36] D. P. Kingma and M. Welling, Auto-encoding variational Bayes, *arXiv:1312.6114*.
- [37] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, WaveNet: A generative model for raw audio, *arXiv:1609.03499*.
- [38] T. Schweigler, Correlations and dynamics of tunnel-coupled one-dimensional Bose gases, Dissertation, Technische Universität Wien, 2019.
- [39] V. Gritsev, A. Polkovnikov, and E. Demler, Linear response theory for a pair of coupled one-dimensional condensates of interacting atoms, *Phys. Rev. B* **75**, 174511 (2007).
- [40] T. Schumm, S. Hofferberth, L. M. Andersson, S. Wildermuth, S. Groth, I. Bar-Joseph, J. Schmiedmayer, and P. Krüger, Matter-wave interferometry in a double well on an atom chip, *Nat. Phys.* **1**, 57 (2005).
- [41] S.-i. Tomonaga, Remarks on Bloch's method of sound waves applied to many-fermion problems, *Prog. Theor. Phys.* **5**, 544 (1950).
- [42] J. Luttinger, An exactly soluble model of a many-fermion system, *J. Math. Phys.* **4**, 1154 (1963).
- [43] L. Faddeev and V. Korepin, Quantum theory of solitons, *Phys. Rep.* **42**, 1 (1978).
- [44] J. B. Kogut, An introduction to lattice gauge theory and spin systems, *Rev. Mod. Phys.* **51**, 659 (1979).
- [45] S. Coleman, More about the massive Schwinger model, *Ann. Phys.* **101**, 239 (1976).
- [46] *The Sine-Gordon Model and its applications: From Pendula and Josephson Junctions to Gravity and High-Energy Physics*, Non-linear Systems and Complexity, edited by J. Cuevas-Maraver, P. G. Kevrekidis, and F. Williams (Springer, Cham, 2014), Vol. 10.
- [47] P. Grišins and I. E. Mazets, Coherence and Josephson oscillations between two tunnel-coupled one-dimensional atomic quasicondensates at finite temperature, *Phys. Rev. A* **87**, 013629 (2013).
- [48] T. Betz, S. Manz, R. Bücker, T. Berrada, C. Koller, G. Kazakov, I. E. Mazets, H.-P. Stimming, A. Perrin, T. Schumm, and J. Schmiedmayer, Two-point phase correlations of a one-dimensional bosonic Josephson junction, *Phys. Rev. Lett.* **106**, 020407 (2011).
- [49] J. A. Krumhansl and J. R. Schrieffer, Dynamics and statistical mechanics of a one-dimensional model Hamiltonian for structural phase transitions, *Phys. Rev. B* **11**, 3535 (1975).
- [50] J. F. Currie, J. A. Krumhansl, A. R. Bishop, and S. E. Trullinger, Statistical mechanics of one-dimensional solitary-wave-bearing scalar fields: Exact results and ideal-gas phenomenology, *Phys. Rev. B* **22**, 477 (1980).
- [51] S. Beck, I. E. Mazets, and T. Schweigler, Nonperturbative method to compute thermal correlations in one-dimensional systems, *Phys. Rev. A* **98**, 023613 (2018).

- [52] H. Risken and H. Haken, *The Fokker-Planck Equation: Methods of Solution and Applications*, 2nd ed. (Springer, Heidelberg, Germany, 1989).
- [53] G. Fernández-Fernández, C. Manzo, M. Lewenstein, A. Dauphin, and G. Muñoz-Gil, Learning minimal representations of stochastic processes with variational autoencoders, *Phys. Rev. E* **110**, L012102 (2024).
- [54] P. de Schoulepnikoff, G. Muñoz-Gil, H. P. Nautrup, and H. J. Briegel, Interpretable representation learning of quantum data enabled by probabilistic variational autoencoders, *Phys. Rev. A* **112**, 062423 (2025).
- [55] R. T. Q. Chen, X. Li, R. B. Grosse, and D. K. Duvenaud, Isolating sources of disentanglement in variational autoencoders, in *Advances in Neural Information Processing Systems*, edited by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Curran Associates, Inc., La Jolla, California, 2018), Vol. 31.
- [56] R. Iten, T. Metger, H. Wilming, L. Del Rio, and R. Renner, Discovering physical concepts with neural networks, *Phys. Rev. Lett.* **124**, 010508 (2020).
- [57] H. P. Nautrup, T. Metger, R. Iten, S. Jerbi, L. M. Trenkwalder, H. Wilming, H. J. Briegel, and R. Renner, Operationally meaningful representations of physical systems in neural networks, *Mach. Learn.: Sci. Technol.* **3**, 045025 (2022).
- [58] G. Fernández-Fernández, SPIVAE, GitHub (2024), <https://github.com/GabrielFernandezFernandez/SPIVAE>
- [59] F. Møller, T. Schweigler, M. Tajik, J. Sabino, F. Cataldini, S.-C. Ji, and J. Schmiedmayer, Thermometry of one-dimensional Bose gases with neural networks, *Phys. Rev. A* **104**, 043305 (2021).
- [60] T. Schweigler, M. Gluza, M. Tajik, S. Sotiriadis, F. Cataldini, S.-C. Ji, F. S. Møller, J. Sabino, B. Rauer, J. Eisert, and J. Schmiedmayer, Decay and recurrence of non-Gaussian correlations in a quantum many-body system, *Nat. Phys.* **17**, 559 (2021).
- [61] N. K. Whitlock and I. Bouchoule, Relative phase fluctuations of two coupled one-dimensional condensates, *Phys. Rev. A* **68**, 053609 (2003).
- [62] S. Coleman, Fate of the false vacuum: Semiclassical theory, *Phys. Rev. D* **16**, 1248(E) (1977).
- [63] M. Pigneur, T. Berrada, M. Bonneau, T. Schumm, E. Demler, and J. Schmiedmayer, Relaxation to a phase-locked equilibrium state in a one-dimensional bosonic Josephson junction, *Phys. Rev. Lett.* **120**, 173601 (2018).
- [64] E. G. Dalla Torre, E. Demler, and A. Polkovnikov, Universal rephasing dynamics after a quantum quench via sudden coupling of two initially independent condensates, *Phys. Rev. Lett.* **110**, 090404 (2013).
- [65] J.-F. Mennemann, I. E. Mazets, M. Pigneur, H. P. Stimming, N. J. Mauser, J. Schmiedmayer, and S. Erne, Relaxation in an extended bosonic Josephson junction, *Phys. Rev. Res.* **3**, 023197 (2021).
- [66] A. Bastianello, Sine-Gordon model from coupled condensates: A generalized hydrodynamics viewpoint, *Phys. Rev. B* **109**, 035118 (2024).
- [67] J. Berges, S. Borsányi, and C. Wetterich, Prethermalization, *Phys. Rev. Lett.* **93**, 142002 (2004).
- [68] F. Møller, C. Li, I. Mazets, H.-P. Stimming, T. Zhou, Z. Zhu, X. Chen, and J. Schmiedmayer, Extension of the generalized hydrodynamics to the dimensional crossover regime, *Phys. Rev. Lett.* **126**, 090602 (2021).
- [69] F. Cataldini, F. Møller, M. Tajik, J. Sabino, S.-C. Ji, I. Mazets, T. Schweigler, B. Rauer, and J. Schmiedmayer, Emergent Pauli blocking in a weakly interacting Bose gas, *Phys. Rev. X* **12**, 041032 (2022).
- [70] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [71] L. N. Smith, A disciplined approach to neural network hyperparameters: Part 1—Learning rate, batch size, momentum, and weight decay, [arXiv:1803.09820](https://arxiv.org/abs/1803.09820).
- [72] L. N. Smith and N. Topin, Super-convergence: Very fast training of neural networks using large learning rates, [arXiv:1708.07120](https://arxiv.org/abs/1708.07120).
- [73] J. Howard and S. Gugger, *fastai*: A layered API for deep learning, *Information* **11**, 108 (2020).
- [74] K. He, X. Zhang, S. Ren, and J. Sun, Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification, in *International Conference on Computer Vision (ICCV)* (IEEE, Washington, DC, 2015), pp. 1026–1034.