

Coherent-State Propagation: A Computational Framework for Simulating Bosonic Quantum Systems

Full version: <https://arxiv.org/abs/2604.19625>. QML applications are given in Section 2

Nikita Guseynov¹

Zoë Holmes²

Armando Angrisani²

¹*SJTU, China*

²*EPFL, Switzerland*

Abstract. We develop *coherent-state propagation*, a computational framework for simulating bosonic systems. We focus on bosonic circuits composed of displaced linear optics augmented by Kerr nonlinearities, a universal model of bosonic quantum computation that is also physically motivated by driven Bose–Hubbard dynamics. The method works in the Schrödinger picture, representing the evolving state as a sparse superposition of coherent states. We develop approximation strategies that keep the simulation cost tractable in physically relevant regimes, notably when the number of Kerr gates is small or the Kerr nonlinearities are weak, and prove rigorous guarantees for both observable estimation and sampling. In particular, bosonic circuits with logarithmically many Kerr gates admit quasi-polynomial-time classical simulation with exponentially small error in trace distance. We further identify a weak-nonlinearity regime in which the runtime is polynomial for arbitrarily small constant precision. We complement these results with numerical benchmarks on the Bose–Hubbard model with all-to-all connectivity. Finally, we discuss applications to quantum machine learning, where coherent-state propagation extends the capabilities of classically trainable Bosonic Born Machines by incorporating controlled non-Gaussian nonlinearities.

Keywords: Coherent-State Propagation, Bosonic Simulation, Kerr Nonlinearity, Bosonic Born Machines

Bosonic degrees of freedom are native to many physical systems, including phonons, photonic and cavity lattices, superconducting resonators, and ultracold gases. In these platforms, bosonic models describe nonlinear response, coherent transport, Josephson dynamics, and driven Bose–Hubbard physics. Practical methods for simulating interacting bosons are therefore of broad relevance: they enable quantitative modelling across a wide range of experimental platforms, while also providing a natural arena in which to delineate the boundary between classical and quantum computational power.

Resource-sensitive classical simulation asks what dynamical features control the cost of representing, propagating, and measuring a quantum state. For qubits, stabilizer dynamics give an efficiently simulable baseline, while simulation overhead is often tied to non-Clifford structure [1–3]. Propagation-based methods give a complementary viewpoint: Pauli propagation tracks observables through a circuit and truncates the growing Pauli expansion when it remains sparse [4, 5], with related Majorana-propagation ideas for fermions [6, 7]. These approaches suggest a useful principle: choose a representation adapted to a tractable baseline, propagate in that representation, and quantify when the generated nonclassical structure remains controlled.

For continuous-variable bosonic systems, Gaussian states, Gaussian operations, and Gaussian measurements are efficiently simulable [8], while non-Gaussianity is required for universal continuous-variable computation. Existing classical methods approaching non-Gaussianity include quasiprobability techniques, and displacement propagation in noisy circuits [9–11]. By contrast, methods for noiseless bosonic circuits with dynamically generated non-Gaussianity remain less developed. This raises the key question: **how much dynamically generated bosonic non-Gaussianity can be tolerated while retaining controlled classical simulation?**

We study circuits built from displaced passive linear optics and local Kerr nonlinearities $K_\kappa = e^{i\kappa\hat{n}^2}$, as shown in Fig. 1. This architecture is universal for continuous-

variable computation in the energy-cutoff sense [12, 13]. Displaced passive linear optics map coherent states to coherent states, while Kerr gates are the only operations that branch the coherent-state representation, making the model well suited for isolating the simulation cost of non-Gaussianity.

We develop **coherent-state propagation**, a Schrödinger-picture framework that represents the evolving state as a sparse superposition of multimode coherent-product states. The method updates coherent amplitudes exactly under displaced passive linear optics and approximates Kerr branching through controlled coherent-state reconstructions.

Main results:

- **Framework:** we develop an end-to-end coherent-state propagation for DPLO+Kerr circuits.
- **Theoretical guarantees:** in regimes with few Kerr gates or weak non-Gaussianity we provide polynomial (quasi-poly) complexity guarantees.
- **Readout:** we derive readout costs for observables estimation and quadrature sampling.
- **Numerics:** the theoretical guarantees are supported by numerical benchmarks on driven Bose–Hubbard dynamics against MPS and Fock simulation.
- **Applications:** our methods broaden the scope of classically trainable Bosonic Born Machines [14], by extending Gaussian models to weakly nonlinear non-Gaussian regimes.

1 Algorithm

We now describe the simulation data and update rules. Coherent-state propagation follows the wavefunction in a moving phase-space expansion that branches at nonlinear Kerr updates. For a single bosonic mode, a coherent state is

$$|\alpha\rangle = D(\alpha)|0\rangle, \quad D(\alpha) = \exp(\alpha\hat{a}^\dagger - \alpha^*\hat{a}). \quad (1)$$




Fig. 1: Circuit architecture. The input is a finite superposition of multimode coherent-product states. Each layer applies a displaced passive linear-optical unitary followed by a local Kerr gate $K_\kappa = e^{i\kappa\hat{n}^2}$, which branches the coherent-state representation. A final Gaussian unitary maps the propagated branches to Gaussian components.




Fig. 2: Benchmark. Six-mode all-to-all driven Bose-Hubbard model, with $\hat{X}_j$ the position quadrature. (a) $\langle \hat{X}_0 \hat{X}_1 \rangle$ trajectory and absolute error for branch cutoff $S = 2^{12}$ and finite-Fourier order $M = 3$. (b) Relative error of $\langle \hat{X}_0^2 \rangle$ at $T = 1$ versus S , for different M . (c) Same for $\langle \hat{X}_0 \hat{X}_1 \rangle$. Here M is the finite-Fourier Kerr parameter: one local Kerr update uses $2M$ phase-orbit coherent states.

For an m -mode amplitude vector $\alpha_k = (\alpha_k^{(1)}, \dots, \alpha_k^{(m)})$, we use the product coherent state

$$|\alpha_k\rangle := \bigotimes_{j=1}^m |\alpha_k^{(j)}\rangle. \quad (2)$$

Coherent-state propagation stores the state as

$$|\psi\rangle = \sum_{k=1}^N C_k |\alpha_k\rangle, \quad (3)$$

where the simulation data are the coefficients C_k and amplitudes α_k . The main bookkeeping variable is the number N of branches: if N remains moderate, overlaps, observables, and sampling distributions can be evaluated from coherent- or Gaussian-state matrix elements.

A circuit layer consists of a displaced passive linear-optical unitary followed by a local Kerr gate, as in Fig. 1. The displaced linear step maps each amplitude vector by an affine update

$$\alpha_k \mapsto e^{-iKt} \alpha_k + \gamma(t), \quad (4)$$

so the coherent-product form is preserved. The Kerr gate

$$K_\kappa = \exp(i\kappa\hat{n}^2) \quad (5)$$

is the only step that changes a single incoming coherent component into a superposition. For arbitrary Kerr strength, its action on a coherent state is reconstructed on a phase orbit,

$$K_\kappa|\alpha\rangle \approx \sum_{r=0}^{N_F} c_r |\alpha\omega_r\rangle, \quad \omega_r = \exp\left(2\pi i \frac{r}{N_F + 1}\right), \quad (6)$$

with coefficients obtained from a discrete Fourier trans-

form of the Kerr phases. The required local reconstruction size scales as

$$N_F = |\alpha|^2 + \mathcal{O}(\log(1/\varepsilon_{\text{kerr}})). \quad (7)$$

The full algorithm propagates a list of coherent branches through the circuit

$$U = U_G \prod_{\ell=1}^L U^{(\ell)} := U_G \prod_{\ell=1}^L U_{\text{Kerr}}^{(\ell)} U_{\text{DPLO}}^{(\ell)}, \quad (8)$$

after which the final Gaussian unitary maps each surviving coherent branch to a Gaussian component.

2 QML application

We clarify here how coherent-state propagation applies to quantum machine learning by extending the classically trainable Bosonic Born Machines [14].

Problem statement. Given samples from a target distribution $p(\mathbf{x})$, with $\mathbf{x} \in \{0, 1\}^m$, train a Bosonic CV state ρ_w whose measured distribution reproduces the target:

$$q_w(\mathbf{x}) = \text{Tr}(\rho_w M_{\mathbf{x}}), \quad \mathbf{w}_* = \arg \min_{\mathbf{w}} L(p, q_w), \quad (9)$$

where $M_{\mathbf{x}} = \bigotimes_{j=1}^m M_{x_j}^{(j)}$ is a product of local measurements. The model focuses on two binary coarse grainings of Fock-basis outcomes. Parity readout is defined by

$$M_0^{\text{par}} = \sum_{n=0}^{\infty} |2n\rangle\langle 2n|; \quad M_1^{\text{par}} = I - M_0^{\text{par}}, \quad (10)$$

whereas threshold readout is defined by

$$M_0^{\text{th}} = |0\rangle\langle 0|; \quad M_1^{\text{th}} = I - M_0^{\text{th}}. \quad (11)$$

Classical learning. In the original paper [14] the ansatz prepares a parameterized Gaussian state using interferometers, squeezing, and displacements. Its mean vector and covariance matrix can therefore be propagated classically, allowing L and its gradients to be evaluated without quantum-hardware calls.

Quantum deployment. After optimization, $\mathbf{w}_*$ specifies a photonic circuit that is suggested to be run on hardware to sample from $q_{\mathbf{w}_*}$. The potential quantum advantage lies in this sampling step: in the relevant Gaussian-boson-sampling regimes, *producing samples from the underlying photonic distribution is believed to be classically hard* for POVMs (10), (11). Thus training is classical, but generation is delegated to the photonic device.

Coherent-state generalization. Coherent-state propagation suggests a route from Gaussian to controlled small-non-Gaussian bosonic Born machines. In the Gaussian model, classical trainability comes from propagating only the mean vector and covariance matrix, but this also restricts the ansatz to Gaussian states. We propose to replace this backend by coherent-state propagation, so that Kerr nonlinearities from the Gaussian+Kerr architecture of [15] can be included while retaining a classical description in the regimes of Theorems 1 and 2. This is promising because Kerr gates add a non-Gaussian resource and increase the expressivity of the learned distribution. The measurement part remains compatible with the same learning objective: parity and threshold readouts only require coherent-state matrix elements $\langle \alpha | M_x^s | \beta \rangle$, which have closed analytical forms. Hence the operator-string expectations entering L can still be evaluated classically, and the trained model can be deployed with the same even/odd or click/no-click photonic measurements.

3 Main Guarantees

The cost is governed by the number of retained coherent branches and by the local coherent amplitudes before Kerr updates. The next two theorems formalize two regimes: arbitrary Kerr strength with few nonlinear layers, and weak Kerr strength with branch truncation.

Let

$$\lambda := \max_{\ell} \max_{k \in [N_{\ell}]} |\alpha_{k,\ell}^{(1)}|^2. \quad (12)$$

This is the largest local mean photon number on the Kerr mode among the coherent branches encountered during propagation.

Theorem 1 (Circuits with few Kerr layers) Let $|\psi_0\rangle$ be a superposition of $N_0 = \text{poly}(m)$ coherent states and let $|\psi\rangle = U|\psi_0\rangle$, where U is the circuit defined in Eq. (8). For any $\varepsilon > 0$, coherent-state propagation outputs a state $|\widetilde{\psi}\rangle$ such that

$$\| |\psi\rangle - |\widetilde{\psi}\rangle \|_2 \leq \varepsilon \quad (13)$$

in time

$$\mathcal{O}\left(Lm^3 + m^2(\lambda + \log(L/\varepsilon))^L\right). \quad (14)$$

Consequently, for circuits with $L = \mathcal{O}(\log m)$ layers:

1. if $\lambda = \text{poly}(m)$ and $\varepsilon = \exp(-\mathcal{O}(m))$, then the simulation runs in quasi-polynomial time,

$$m^{\mathcal{O}(\log m)}; \quad (15)$$

2. if $\lambda = \mathcal{O}(\log m)$ and $\varepsilon = 1/\text{poly}(m)$, then the simulation runs in almost-polynomial time,

$$m^{\mathcal{O}(\log \log m)}. \quad (16)$$

This result gives the general resource scaling: at fixed coherent energy scale, the overhead is controlled by the number of nonlinear branch-generating layers and by the local accuracy required at each Kerr reconstruction.

Theorem 2 (Circuits with small nonlinearity)

Let $|\psi_0\rangle$ be a superposition of $N_0 = \text{poly}(m)$ coherent states and let $|\psi\rangle = U|\psi_0\rangle$, where U is the circuit defined in Eq. (8). Assume the strength of the Kerr gates is upper bounded by

$$\kappa \leq \frac{\delta}{L(1 + \lambda^2)}, \quad (17)$$

for a sufficiently small constant δ .

Then, for every $\varepsilon \in (0, 1)$, coherent-state propagation with a small-angle truncation scheme outputs a state $|\widetilde{\psi}\rangle$ such that

$$\| |\psi\rangle - |\widetilde{\psi}\rangle \|_2 \leq \mathcal{O}(\delta) + \varepsilon, \quad (18)$$

with runtime

$$t = \mathcal{O}\left(Lm^3 + m^2L^{1+\mathcal{O}(\log(1/\varepsilon))}\right). \quad (19)$$

In particular, for constant ε , the runtime is polynomial in m and L .

The weak-nonlinearity guarantee is closely aligned with short-time Trotterized dynamics. The small branch in the two-term Kerr approximation can be pruned systematically, replacing a binary tree by $\mathcal{O}(N_0L^s)$ retained terms for fixed s , analogous in spirit to small-angle qubit simulation [16].

After the final Gaussian layer, the **output** has the form $|\psi_{\text{out}}\rangle = \sum_j C_j |G_j\rangle$. The readout guarantees cover polynomial quadrature observables

$$\hat{O} = \sum_{r=1}^{\eta} \left(o_r \hat{\mathbf{X}}^{q^{(r)}} \hat{\mathbf{P}}^{s^{(r)}} + o_r^* \hat{\mathbf{P}}^{s^{(r)}} \hat{\mathbf{X}}^{q^{(r)}} \right), \quad (20)$$

with cost $\mathcal{O}(N^2(m^3 + \eta m^2))$. We can also sample rotated quadratures

$$\hat{X}_{\theta} = \hat{X} \cos \theta + \hat{P} \sin \theta. \quad (21)$$

Sampling N_{shot} outcomes costs $\mathcal{O}(N^2 m^2 \log(1/\varepsilon) + N_{\text{shot}} N m^2)$.

4 Numerical Experiments

We validate coherent-state propagation on the driven Bose-Hubbard model in the rotating frame,

$$\begin{aligned} \hat{H}_{\text{rot}} = & -J \sum_{\langle i,j \rangle} \left(\hat{a}_i^{\dagger} \hat{a}_j + \hat{a}_j^{\dagger} \hat{a}_i \right) - \Delta \sum_i \hat{n}_i \\ & + \frac{U}{2} \sum_i \hat{n}_i^2 + \sum_i \left(\Omega_i \hat{a}_i^{\dagger} + \Omega_i^* \hat{a}_i \right). \end{aligned} \quad (22)$$

The hopping, detuning, and drive terms form the DPLO part, while the onsite interaction gives local Kerr evolution [17–20]. For the six-mode all-to-all system in Fig. 2, we compare against matrix-product-state reference data [21, 22]. The agreement and cutoff trends support the central picture: coherent-state propagation is effective when weak Kerr steps generate an organized, tunable coherent-state expansion.

References

- [1] Daniel Gottesman. The heisenberg representation of quantum computers, talk at. In *International Conference on Group Theoretic Methods in Physics*. Citeseer, 1998.
- [2] Scott Aaronson and Daniel Gottesman. Improved simulation of stabilizer circuits. *Physical Review A*, 70(5):052328, 2004.
- [3] Sergey Bravyi and David Gosset. Improved classical simulation of quantum circuits dominated by clifford gates. *Phys. Rev. Lett.*, 116:250501, Jun 2016.
- [4] Patrick Rall, Daniel Liang, Jeremy Cook, and William Kretschmer. Simulation of qubit quantum circuits via pauli propagation. *Physical Review A*, 99(6):062337, 2019.
- [5] Manuel S Rudolph, Tyson Jones, Yanting Teng, Armando Angrisani, and Zoë Holmes. Pauli propagation: A computational framework for simulating quantum systems. *arXiv preprint arXiv:2505.21606*, 2025.
- [6] Aaron Miller, Joachim Favre, Zoë Holmes, Özlem Salehi, Rahul Chakraborty, Anton Nykänen, Zoltan Zimboras, Adam Glos, and Guillermo García-Pérez. Simulation of fermionic circuits using majorana propagation. *arXiv preprint arXiv:2503.18939*, 2025.
- [7] Matteo D’Anna, Jannes Nys, and Juan Carrasquilla. Majorana string simulation of nonequilibrium dynamics in two-dimensional lattice fermion systems. *arXiv:2511.02809*, 2025.
- [8] Stephen D Bartlett and Barry C Sanders. Universal continuous-variable quantum computation: Requirement of optical nonlinearity for photon counting. *Phys. Rev. A*, 65:042304, Mar 2002.
- [9] Ulysse Chabaud, Giulia Ferrini, Frédéric Grosshans, and Damian Markham. Classical simulation of gaussian quantum circuits with non-gaussian input states. *Physical Review Research*, 3(3):033018, 2021.
- [10] Beatriz Dias and Robert König. Classical simulation of non-gaussian bosonic circuits. *Physical Review A*, 110(4):042402, 2024.
- [11] Varun Upreti, Ulysse Chabaud, Zoë Holmes, and Armando Angrisani. When quantum resources backfire: Non-gaussianity and symplectic coherence in noisy bosonic circuits. *arXiv preprint arXiv:2510.07264*, 2025.
- [12] Francesco Arzani, Robert I Booth, and Ulysse Chabaud. Can effective descriptions of bosonic systems be considered complete? *arXiv preprint arXiv:2501.13857*, 2025.
- [13] Seth Lloyd and Samuel L. Braunstein. Quantum computation over continuous variables. *Phys. Rev. Lett.*, 82:1784–1787, Feb 1999.
- [14] Zoltán Kolarovszki, Bence Bakó, Michał Oszmaniec, Changhun Oh, and Zoltán Zimborás. Generative modeling with Gaussian Boson Sampling: classically trainable Bosonic Born Machines, March 2026. arXiv:2603.11195.
- [15] Nathan Killoran, Thomas R. Bromley, Juan Miguel Arrazola, Maria Schuld, Nicolás Quesada, and Seth Lloyd. Continuous-variable quantum neural networks. *Physical Review Research*, 1(3):033063, October 2019.
- [16] Sacha Lerch, Ricard Puig, Manuel Rudolph, Armando Angrisani, Tyson Jones, M. Cerezo, Supanut Thanasilp, and Zoë Holmes. Efficient quantum-enhanced classical simulation for patches of quantum landscapes. *arXiv preprint arXiv:2411.19896*, 2024.
- [17] Matthew P. A. Fisher, Peter B. Weichman, G. Grinstein, and Daniel S. Fisher. Boson localization and the superfluid-insulator transition. *Phys. Rev. B*, 40:546–570, Jul 1989.
- [18] Markus Greiner, Olaf Mandel, Tilman Esslinger, Theodor W. Hänsch, and Immanuel Bloch. Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms. *Nature*, 415(6867):39–44, 2002.
- [19] Alexandre Le Boité, Giuliano Orso, and Cristiano Ciuti. Steady-state phases and tunneling-induced instabilities in the driven dissipative bose-hubbard model. *Phys. Rev. Lett.*, 110:233601, Jun 2013.
- [20] Iacopo Carusotto and Cristiano Ciuti. Quantum fluids of light. *Rev. Mod. Phys.*, 85:299–366, Feb 2013.
- [21] A J Daley, C Kollath, U Schollwöck, and G Vidal. Time-dependent density-matrix renormalization-group using adaptive effective hilbert spaces. *Journal of Statistical Mechanics: Theory and Experiment*, 2004(04):P04005, apr 2004.
- [22] Ulrich Schollwöck. The density-matrix renormalization group in the age of matrix product states. *Annals of Physics*, 326(1):96–192, 2011. January 2011 Special Issue.