

Sample-based training of quantum Boltzmann machines

Maria Demidik^{1,2}, Cenk Tüysüz³, Michele Grossi³, and Karl Jansen^{2,1}

¹ Deutsches Elektronen-Synchrotron DESY, Zeuthen, Germany

² Computation-Based Science and Technology Research Center, The Cyprus Institute, Nicosia, Cyprus

³ European Organisation for Nuclear Research (CERN), Geneva, Switzerland

Introduction. Certain families of quantum circuits can sample from probability distributions that are believed to be classically intractable under standard complexity-theoretic assumptions [1–3]. Such hardness results have motivated quantum generative models, which represent and sample from probability distributions beyond the reach of classical generative methods [4]. However, trainability rather than expressivity has emerged as a central bottleneck. Parameter-shift and finite-difference methods estimate gradients by executing separate circuits for each trainable parameter [5–9]. Finite-shot noise increases the required number of circuit executions further [10], and barren plateaus can raise the measurement cost of gradient estimation exponentially [11]. Unlike classical neural networks, general quantum circuit models do not admit a backpropagation mechanism that computes all gradients in a single backward pass [12, 13], and scalable training has so far been limited to specific generative architectures [14, 15].

In this work, we propose a generalized contrastive divergence (CD) algorithm for training quantum models. In analogy with classical backpropagation, each update requires $\mathcal{O}(1)$ forward passes independent of the parameter count. The update rule is estimated from data and model samples alone, without explicit access to the model distribution. We apply this procedure to quantum Boltzmann machines (QBMs), yielding a training loop that requires only conditional sampling. The full version of the manuscript [16] contains the proofs and further details.

Background. CD emerged in classical machine learning as a training scheme for energy-based models, updating all parameters from data and model samples at a constant cost per step, independent of the parameter count [17–19]. The common learning rule [17, 20] for an energy-based model defined through Hamiltonian H can be written as $\partial_{\theta_i} \mathcal{L} = \langle H_i \rangle_{\text{data}} - \langle H_i \rangle_{\text{model}}$, where $\langle \cdot \rangle_{\text{data}}$ and $\langle \cdot \rangle_{\text{model}}$ denote expectation values under the empirical data and the model. This learning rule is commonly applied to Boltzmann machines (BMs), which consist of n binary visible units representing the data and m hidden units capturing correlations among them [21]. At thermal equilibrium, a BM is described by the Gibbs state $\rho = e^{-\beta H} / \mathcal{Z}$, $\mathcal{Z} = \text{Tr}[e^{-\beta H}]$, where H is a parametrized Hamiltonian, β is the inverse temperature, and $\mathcal{Z}$ is the partition function. The model distribution over visible configurations follows from a projective measurement on the visible subsystem, $p(v) = \text{Tr}[(\Pi_v \otimes \mathbb{1}^{\otimes m}) \rho]$ with $\Pi_v = |v\rangle\langle v|$ and $v \in \{0, 1\}^n$.

Estimating the expectations in the learning rule amounts to sampling from the Gibbs state, which is classically intractable for general architectures [22, 23]. Restricting the connectivity to a bipartite structure between the visible and hidden layers yields the restricted Boltzmann machine (RBM) [24], whose conditional distributions factorize as $p(h | v) = \prod_{j=1}^m p(h_j | v)$ and $p(v | h) = \prod_{i=1}^n p(v_i | h)$, where v_i and h_j are the states of visible unit i and hidden unit j . These conditionals can be sampled exactly and in parallel, which makes Gibbs sampling [25] tractable. However, running Markov chains to equilibrium remains computationally expensive and often yields high-variance gradient estimates [26]. CD [17] avoids equilibrium sampling by comparing the data with reconstructions obtained after k Gibbs updates started from data samples, where small k is sufficient to approximate gradient directions [20, 27]. The resulting updates take the form $\Delta\theta_i \propto \langle H_i \rangle_{\text{data}} - \langle H_i \rangle_{p^k}$, where p^k denotes the model distribution after k sampling steps. These updates are biased [28] and do not follow the gradient of any objective function [27]. Despite these limitations, CD remains the primary heuristic for training RBMs [29–31].

QBMs are natural candidates for such a training scheme. They define the model distribution through the Gibbs state of a non-commuting Hamiltonian [32] and can exceed the expressivity of their classical counterparts [33]. Likelihood-based training, however, estimates $\langle H_i \rangle_{\text{data}} - \langle H_i \rangle_{\text{model}}$ for each Hamiltonian term, so the measurement cost grows with the number of parameters. These expectation values further require explicit access to target data probabilities [32, 34–36], which precludes training from samples. The common workaround of a maximum mean discrepancy (MMD) loss performs poorly for complex distributions [37, 38]. No method has so far trained a QBM from data samples alone. The semi-quantum restricted Boltzmann




Fig. 1. Summary of generalized contrastive divergence. Short Gibbs chains start from data samples. For each hidden Pauli basis, conditional samples of hidden and visible units give the positive and negative phase bitstrings that update all parameters via Eq. (4). For fixed k and $\mathcal{W}_h$, the number of forward passes per update is independent of the parameter count.

machine (sqRBM) unifies the classical and quantum settings for learning from classical data. We restate its definition from [39].

Definition 1 (Semi-quantum RBM [39]). A semi-quantum restricted Boltzmann machine with n visible and m hidden units, denoted $\text{sqRBM}_{n,m}$, is described by a parametrized Hamiltonian of the form

$$H = \sum_{i=1}^n a_i^Z \sigma_i^Z + \sum_{P \in \mathcal{W}_h} \left(\sum_{j=1}^m b_j^P \sigma_{n+j}^P + \sum_{i=1}^n \sum_{j=1}^m w_{ij}^P \sigma_i^Z \sigma_{n+j}^P \right), \quad (1)$$

where $\mathbf{a} \in \mathbb{R}^n$, $\mathbf{b} \in \mathbb{R}^{|\mathcal{W}_h| \cdot m}$ and $\mathbf{w} \in \mathbb{R}^{|\mathcal{W}_h| \cdot n \cdot m}$ are the parameter vectors of the model. $\mathcal{W}_h$ is a non-commuting set of one-qubit Pauli operators that non-trivially act only on the hidden units. σ_i^P is a length $n + m$ Pauli string acting non-trivially with Pauli operator P on the i -th qubit, $\sigma_i^P = \mathbb{1}^{\otimes(i-1)} \otimes P \otimes \mathbb{1}^{\otimes(n+m-i)}$.

The Hamiltonian in Eq. (1) can be written compactly as $H = \sum_i \theta_i H_i$, where $\boldsymbol{\theta} = \{\mathbf{a}, \mathbf{b}, \mathbf{w}\}$ and H_i are the associated Pauli strings. Choosing $\mathcal{W}_h = \{\mathbb{Z}\}$ recovers the classical RBM, in which all Hamiltonian terms commute. The full set $\mathcal{W}_h = \{\mathbb{X}, \mathbb{Y}, \mathbb{Z}\}$, used throughout this work, introduces non-commuting interactions and matches the learning performance of RBMs with three times fewer hidden units [39].

Theoretical results. For sqRBMs, we establish that the conditional probabilities central to CD training take a simple and exact form, which translates directly into a circuit-level implementation.

Theorem 1 (Conditional probabilities of sqRBMs, informal). Let H be the Hamiltonian of an sqRBM acting on $\mathcal{H}_{\text{vis}} \otimes \mathcal{H}_{\text{hid}}$ with n visible and m hidden qubits, and let $\beta > 0$. For $v \in \{0, 1\}^n$, let $\Pi_v^Z := |v\rangle\langle v|$ be the projector onto the visible computational basis state. For the hidden subsystem, fix a Pauli basis P and let Π_h^P be the projector onto the basis state $|h\rangle_P$ of $\mathcal{H}_{\text{hid}}$. For any non-zero positive semidefinite operator X on $\mathcal{H}_{\text{vis}} \otimes \mathcal{H}_{\text{hid}}$, define the completely positive map $\mathcal{G}_\beta(X) := e^{-(\beta/2)H} X e^{-(\beta/2)H}$ and its trace normalization $\Gamma_\beta(X) := \mathcal{G}_\beta(X) / \text{Tr}[\mathcal{G}_\beta(X)]$. Then the conditional distributions of the sqRBM Gibbs state are given by

$$p(h^P | v) = \text{Tr}[(\mathbb{1}^{\otimes n} \otimes \Pi_h^P) \Gamma_\beta(\Pi_v^Z \otimes 2^{-m} \mathbb{1}^{\otimes m})], \quad (2)$$

$$p(v | h^P) = \text{Tr}[(\Pi_v^Z \otimes \mathbb{1}^{\otimes m}) \Gamma_\beta(2^{-n} \mathbb{1}^{\otimes n} \otimes \Pi_h^P)]. \quad (3)$$

To sample $p(h^P | v)$ on a quantum device, one prepares the visible register in $|v\rangle$ and the hidden register in the maximally mixed state, applies the imaginary-time evolution operator $\exp(-(\beta/2)H)$, and measures the hidden register in basis P . Sampling $p(v | h^P)$ proceeds analogously, with the roles of the registers exchanged and the visible register measured in the $\mathbb{Z}$ basis. The proof and the circuits are given in Ref. [16].

Generalized contrastive divergence algorithm. Theorem 1 gives the conditional probabilities of sqRBMs in closed form, which allows us to formulate a generalized CD- k procedure, summarized in Fig. 1. The hidden operator pool $\mathcal{W}_h = \{\mathbb{X}, \mathbb{Y}, \mathbb{Z}\}$ partitions into commuting subsets. We perform a CD update in each basis $P \in \mathcal{W}_h$ and accumulate the full gradient. Starting from a visible configuration v^{data} drawn from the dataset,




Fig. 2. Sampling cost of generalized contrastive divergence (CD) and negative log-likelihood (NLL) training. KL divergence to the target distribution (top row) and ratio of valid images among generated samples (bottom row) as a function of the total number of samples drawn during training, for a discretized Gaussian ($n = 9$), bars and stripes (BAS, $n = 3 \times 3$), and Tetris ($n = 4 \times 4$). Green curves show generalized CD, pink curves NLL training with 10^0 to 10^4 shots per expectation-value estimate, darker shades corresponding to larger shot budgets. Within each panel, all curves correspond to the same model.

we sample the positive phase $h^{(0),P} \sim p(h^P | v^{\text{data}})$ and run a Gibbs chain $v^{(s)} \sim p(v | h^{(s-1),P})$, $h^{(s),P} \sim p(h^P | v^{(s)})$ for $s = 1, \dots, k$. The resulting bitstrings yield the parameter updates

$$\Delta \mathbf{w}^{Z,P} = v^{\text{data}} \cdot (h^{(0),P})^\top - v^{(k)} \cdot (h^{(k),P})^\top, \quad \Delta \mathbf{a} = v^{\text{data}} - v^{(k)}, \quad \Delta \mathbf{b}^P = h^{(0),P} - h^{(k),P}. \quad (4)$$

Since the visible subspace of an sqRBM is commuting, the hidden conditionals $p(h^P | v)$ can be sampled classically and in parallel, as in standard RBMs. Only $p(v | h^P)$ requires quantum resources. One CD- k update therefore costs $|\mathcal{W}_h| \cdot k$ circuit executions, each sampled once. We refer to a single evaluation of $p(v | h^P)$ as a *forward pass*, since it is the cost of sampling the visible distribution $p(v)$. For fixed k and $|\mathcal{W}_h| = 3$, gradient computation costs $\mathcal{O}(1)$ forward passes per update.

Numerical results. We compare generalized CD to standard NLL training on three targets, a discretized Gaussian on $n = 9$ units, 3×3 bars and stripes (BAS), and a Tetris dataset of all tetromino images under rotation and translation on a 4×4 grid. The NLL gradient requires expectation values, estimated with 10^0 to 10^4 measurement shots each. Generalized CD uses $k = 10$ Gibbs steps with ten persistent chains [40]. As shown in Fig. 2, CD reaches its final accuracy with orders of magnitude fewer samples than NLL, at a nearly constant sample budget while the cost of NLL grows with n . With a sufficient shot budget, NLL attains a lower KL divergence, as the two differ only by a constant. Although CD-trained models retain a higher KL divergence than the best NLL runs, they generate valid BAS and Tetris images at a comparable rate.

Conclusion. We introduced a generalized CD algorithm that trains sqRBMs from data samples alone, without explicit access to target data probabilities, at a cost of $\mathcal{O}(1)$ forward passes per update. Our numerical results show that it reduces the sampling cost of training by orders of magnitude at a moderate loss in final accuracy. Generalized CD remains a heuristic without convergence guarantees. A hardware demonstration requires implementing the imaginary-time evolution $e^{-(\beta/2)H}$. Quantum Gibbs samplers with a system-size-independent spectral gap exist for k -local Hamiltonians, including those of sqRBMs, at sufficiently high temperatures [41], and adapting them to our conditional sampling is left for future work.

References

1. M. J. Bremner, R. Jozsa, and D. J. Shepherd, “Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy,” *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 467, no. 2126, p. 459–472, Aug. 2010. [Online]. Available: <http://dx.doi.org/10.1098/rspa.2010.0301>
2. F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. V. Klimov, S. Knysh, A. Korotkov, F. Kostitsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis, “Quantum supremacy using a programmable superconducting processor,” *Nature*, vol. 574, no. 7779, pp. 505–510, Oct. 2019.
3. S. Aaronson and A. Arkhipov, “The computational complexity of linear optics,” in *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, ser. STOC ’11. New York, NY, USA: Association for Computing Machinery, 2011, p. 333–342.
4. H.-Y. Huang, M. Broughton, N. Eassa, H. Neven, R. Babbush, and J. R. McClean, “Generative quantum advantage for classical and quantum problems,” 2025.
5. J. Li, X. Yang, X. Peng, and C.-P. Sun, “Hybrid quantum-classical approach to quantum optimal control,” *Phys. Rev. Lett.*, vol. 118, p. 150503, Apr 2017. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.118.150503>
6. K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, “Quantum circuit learning,” *Phys. Rev. A*, vol. 98, p. 032309, Sep 2018. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.98.032309>
7. M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, “Evaluating analytic gradients on quantum hardware,” *Phys. Rev. A*, vol. 99, p. 032331, Mar 2019. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.99.032331>
8. L. Banchi and G. E. Crooks, “Measuring Analytic Gradients of General Quantum Evolution with the Stochastic Parameter Shift Rule,” *Quantum*, vol. 5, p. 386, Jan. 2021. [Online]. Available: <https://doi.org/10.22331/q-2021-01-25-386>
9. D. Wierichs, J. Izaac, C. Wang, and C. Y.-Y. Lin, “General parameter-shift rules for quantum gradients,” *Quantum*, vol. 6, p. 677, Mar. 2022. [Online]. Available: <https://doi.org/10.22331/q-2022-03-30-677>
10. D. A. Kreplin and M. Roth, “Reduction of finite sampling noise in quantum neural networks,” *Quantum*, vol. 8, p. 1385, Jun. 2024. [Online]. Available: <https://doi.org/10.22331/q-2024-06-25-1385>
11. J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, “Barren plateaus in quantum neural network training landscapes,” *Nature Communications*, vol. 9, no. 1, p. 4812, Nov. 2018.
12. A. Abbas, R. King, H.-Y. Huang, W. J. Huggins, R. Movassagh, D. Gilboa, and J. McClean, “On quantum backpropagation, information reuse, and cheating measurement collapse,” in *Advances in Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, Eds., vol. 36. Curran Associates, Inc., 2023, pp. 44 792–44 819. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2023/file/8c3cae2f725c8e2a55ecd600563d172-Paper-Conference.pdf
13. J. Bowles, D. Wierichs, and C.-Y. Park, “Backpropagation scaling in parameterised quantum circuits,” *Quantum*, vol. 9, p. 1873, Oct. 2025. [Online]. Available: <https://doi.org/10.22331/q-2025-10-02-1873>
14. E. Recio-Armengol, S. Ahmed, and J. Bowles, “Train on classical, deploy on quantum: scaling generative quantum machine learning to a thousand qubits,” 2026. [Online]. Available: <https://arxiv.org/abs/2503.02934>
15. C. Tüysüz, O. Kyriienko, and M. Grossi, “Quantum fourier generative models trainable at large scale,” 2026. [Online]. Available: <https://arxiv.org/abs/2606.28483>
16. M. Demidik, C. Tüysüz, M. Grossi, and K. Jansen, “Sample-based training of quantum generative models,” 2025. [Online]. Available: <https://arxiv.org/abs/2511.11802>
17. G. E. Hinton, “Training products of experts by minimizing contrastive divergence,” *Neural Comput.*, vol. 14, no. 8, p. 1771–1800, Aug. 2002. [Online]. Available: <https://doi.org/10.1162/089976602760128018>
18. R. Salakhutdinov and G. Hinton, “An Efficient Learning Procedure for Deep Boltzmann Machines,” *Neural Computation*, vol. 24, no. 8, pp. 1967–2006, Aug. 2012. [Online]. Available: <https://direct.mit.edu/neco/article/24/8/1967-2006/7793>
19. Y. Du, S. Li, J. Tenenbaum, and I. Mordatch, “Improved Contrastive Divergence Training of Energy Based Models,” Jun. 2021.
20. M. A. Carreira-Perpiñán and G. Hinton, “On contrastive divergence learning,” in *Proceedings of the Tenth International Workshop on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, R. G. Cowell and Z. Ghahramani, Eds., vol. R5. PMLR, 06–08 Jan 2005, pp. 33–40, reissued by PMLR on 30 March 2021. [Online]. Available: <https://proceedings.mlr.press/r5/carreira-perpinan05a.html>
21. N. Le Roux and Y. Bengio, “Representational Power of Restricted Boltzmann Machines and Deep Belief Networks,” *Neural Computation*, vol. 20, no. 6, pp. 1631–1649, Jun. 2008.
22. D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, *A learning algorithm for Boltzmann machines*. Cambridge, MA, USA: MIT Press, 1988, p. 635–649.
23. P. M. Long and R. A. Servedio, “Restricted boltzmann machines are hard to approximately evaluate or simulate,” in *Proceedings of the 27th International Conference on International Conference on Machine Learning*, 2010. [Online]. Available: http://www.phillong.info/publications/LS10_rbm.pdf
24. P. Smolensky, *Information processing in dynamical systems: foundations of harmony theory*. Cambridge, MA, USA: MIT Press, 1986, p. 194–281.
25. G. E. Hinton and T. J. Sejnowski, “Optimal perceptual inference,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1983. [Online]. Available: <https://api.semanticscholar.org/CorpusID:10379672>

26. Y. Bengio and O. Delalleau, “Justifying and Generalizing Contrastive Divergence,” *Neural Computation*, vol. 21, no. 6, pp. 1601–1621, Jun. 2009. [Online]. Available: <https://doi.org/10.1162/neco.2008.11-07-647>
27. I. Sutskever and T. Tieleman, “On the convergence properties of contrastive divergence,” in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, Y. W. Teh and M. Titterton, Eds., vol. 9. Chia Laguna Resort, Sardinia, Italy: PMLR, 13–15 May 2010, pp. 789–795. [Online]. Available: <https://proceedings.mlr.press/v9/sutskever10a.html>
28. A. Yuille, “The convergence of contrastive divergences,” in *Proceedings of the 18th International Conference on Neural Information Processing Systems*, ser. NIPS’04. Cambridge, MA, USA: MIT Press, 2004, p. 1593–1600.
29. H. Chen and A. Murray, “A Continuous Restricted Boltzmann Machine with a Hardware- Amenable Learning Algorithm,” in *Artificial Neural Networks — ICANN 2002*, J. R. Dorronsoro, Ed. Berlin, Heidelberg: Springer, 2002, pp. 358–363.
30. Y. W. Teh, M. Welling, S. Osindero, and G. E. Hinton, “Energy-based models for sparse overcomplete representations,” *J. Mach. Learn. Res.*, vol. 4, no. null, p. 1235–1260, Dec. 2003.
31. Xuming He, R. Zemel, and M. Carreira-Perpinan, “Multiscale conditional random fields for image labeling,” in *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, vol. 2. Washington, DC, USA: IEEE, 2004, pp. 695–702. [Online]. Available: <http://ieeexplore.ieee.org/document/1315232/>
32. M. H. Amin, E. Andriyash, J. Rolfe, B. Kulchitsky, and R. Melko, “Quantum Boltzmann Machine,” *Physical Review X*, vol. 8, no. 2, p. 021050, May 2018.
33. C. Tüysüz, M. Demidik, L. Coopmans, E. Rinaldi, V. Croft, Y. Haddad, M. Rosenkranz, and K. Jansen, “Learning to generate high-dimensional distributions with low-dimensional quantum boltzmann machines,” 2024.
34. M. Kieferova and N. Wiebe, “Tomography and Generative Data Modeling via Quantum Boltzmann Training,” *Physical Review A*, vol. 96, no. 6, p. 062327, Dec. 2017.
35. H. J. Kappen, “Learning quantum models from quantum or classical data,” *Journal of Physics A: Mathematical and Theoretical*, vol. 53, no. 21, p. 214001, May 2020. [Online]. Available: <https://dx.doi.org/10.1088/1751-8121/ab7df6>
36. L. Coopmans and M. Benedetti, “On the sample complexity of quantum Boltzmann machine learning,” *Communications Physics*, vol. 7, no. 1, p. 274, Aug. 2024. [Online]. Available: <https://doi.org/10.1038/s42005-024-01763-x>
37. M. S. Rudolph, S. Lerch, S. Thanasilp, O. Kiss, O. Shaya, S. Vallecorsa, M. Grossi, and Z. Holmes, “Trainability barriers and opportunities in quantum generative modeling,” *npj Quantum Information*, vol. 10, no. 1, p. 116, Nov. 2024.
38. M. Herrero-Gonzalez, B. Coyle, K. McDowall, R. Grassie, S. Beentjes, A. Khamseh, and E. Kashefi, “The born ultimatum: Conditions for classical surrogation of quantum generative models with correlators,” 2025.
39. M. Demidik, C. Tüysüz, N. Piatkowski, M. Grossi, and K. Jansen, “Expressive equivalence of classical and quantum restricted Boltzmann machines,” *Communications Physics*, vol. 8, no. 1, p. 413, Oct. 2025.
40. T. Tieleman, “Training restricted boltzmann machines using approximations to the likelihood gradient,” in *Proceedings of the 25th International Conference on Machine Learning*, ser. ICML ’08. New York, NY, USA: Association for Computing Machinery, 2008, p. 1064–1071. [Online]. Available: <https://doi.org/10.1145/1390156.1390290>
41. T. Bergamaschi, “Fast mixing of all-to-all quantum systems at high temperatures,” 2026. [Online]. Available: <https://arxiv.org/abs/2606.26090>