Fundamentals of quantum Boltzmann machine learning with visible and hidden units

Mark M. Wilde
School of Electrical and Computer Engineering, Cornell University, Ithaca, New York 14850, USA

Background and motivation: Boltzmann machines constitute one of the earliest neural-network based approaches to generative modeling [AHS85, HS86], in which the goal is to simulate a target probability distribution $q(v)$ by means of a parameterized model distribution $p_\theta(v)$, where θ is a parameter vector. The traditional approach to Boltzmann machine learning is to employ a classical Hamiltonian as an energy function, with some of the variables called visible and the others called hidden (often referred to as visible and hidden units). The probability $p_\theta(v,h)$ for a particular configuration of visible and hidden variables to occur is then proportional to a Boltzmann weight, establishing a strong link between Boltzmann machine learning and statistical physics. The marginal distribution $p_\theta(v)$ on the visible variables should be tuned to closely approximate the target probability distribution, while the role of the hidden variables is to mediate complex correlations between the visible variables by means of simple interactions between the visible and hidden variables. This is similar to how the Hamiltonian of mean force represents complex interactions between system variables in classical thermodynamics [Jar04, TH20], which are mediated by the interactions between a thermodynamic system and an external reservoir.

The conventional approach to generative modeling with classical Boltzmann machines is to tune θ so that the relative entropy $D(q\|p_\theta)$ between the target distribution $q(v)$ and the model distribution $p_\theta(v)$ is minimized. In order to do so, it is helpful to have a formula for the elements of the gradient $\nabla_\theta D(q\|p_\theta)$, which is well known to be given by [AHS85]

$$\frac{\partial}{\partial \theta_j} D(q\|p_\theta) = \sum_{v,h} q(v) p_\theta(h|v) G_j(v,h) - \sum_{v,h} p_\theta(v,h) G_j(v,h), \tag{1}$$

when the classical Hamiltonian has the form $G_\theta(v,h) \coloneqq \sum_j \theta_j G_j(v,h)$. As such, one can estimate the elements of the gradient by means of Monte Carlo sampling, i.e., sampling from $q(v)p_\theta(h|v)$ and $p_\theta(v,h)$. In conjunction, one can employ a standard stochastic gradient descent algorithm in order to minimize $D(q\|p_\theta)$. While this optimization problem is known to be non-convex in general, in practice, a local minimum of this objective function often works well. After tuning θ in this way, one can then sample from the model distribution p_θ at will.

Roughly a decade ago now, quantum Boltzmann machines were proposed for generative modeling [AAR$^+$18, KW17, BRGBPO17], with the idea being that distributions arising from measuring quantum states might be able to capture more complex correlations in target probability distributions than would be possible with classical Boltzmann machines alone. Beyond this application, quantum Boltzmann machines can alternatively be used for the non-classical task of quantum state learning [KW17], in which the goal is to model a target quantum state ρ by means of a parameterized state $\sigma(\theta)$. It was realized in [AAR+18, KW17, BRGBPO17] that there are difficulties in arriving at simple expressions like that in (1) and procedures for estimating it when trying to develop quantum extensions of generative modeling, while using both visible and hidden variables. A notable exception occurs for quantum state learning when using quantum Boltzmann machines with visible units only: indeed, the authors of [KW17] established a simple expression for the gradient (see [KW17, Eq. (4)]), while the authors of [CB24] proved that the corresponding minimization problem is convex in the parameter vector θ. Thus, for this particular problem, a hybrid quantum–classical algorithm, which uses a quantum computer for estimating the gradient and stochastic gradient descent for optimization, is guaranteed to converge to a global minimum with efficient sample complexity [CB24, Theorem 1].

What has remained open since these prior contributions is to determine an analytical expression for the gradient that is amenable to estimation on a quantum computer when using quantum Boltzmann machines, with both visible and hidden units, for generative modeling and quantum state learning. Under the assumption that the visible units of the quantum Boltzmann machine are classical while the hidden units are quantum, there has been recent progress on developing algorithms for optimizing them for generative modeling [PKPW25, DTP+25, KKH25, VSKB25, DTGJ25, Wil25b]. However, these works do not address the aforementioned open question for the task of quantum state learning with both visible and hidden units, for which a quantum Boltzmann machine with quantum visible units is required.

Summary of contributions: In this submission, I address this open question by determining an analytical expression for the gradient in quantum state learning, which is amenable to estimation on a quantum computer. In more detail, suppose that ρ is a target quantum state and $\sigma(\theta)$ is a model quantum state, based on a quantum Boltzmann machine with visible and hidden units, and the goal is to minimize the quantum relative entropy $D(\rho\|\sigma(\theta))$. Here I establish an analytical expression for the gradient $\nabla_\theta D(\rho\|\sigma(\theta))$, which can be viewed as a quantum generalization of (1), incorporating aspects to account for the non-commutativity inherent in quantum mechanics (see [Wil25a, Theorem 8]). Furthermore, I develop a quantum algorithm for estimating the gradient when given sample access to ρ and $\sigma(\theta)$, implying that one can find a local minimum of $D(\rho\|\sigma(\theta))$ by means of a hybrid quantum–classical algorithm that uses a quantum computer for estimating the gradient and a classical computer for performing the stochastic gradient descent optimization algorithm. A subroutine of this quantum algorithm involves simulating what is called modular flow, and by invoking [QKR25, Lemma 3] and [GSLW19, Corollary 60 and Lemma 61], I establish an improvement of the recent quantum algorithm from [LK25] for simulating modular flow.

The main distinction between the quantum expression reported in [Wil25a, Theorem 8] and the classical one in (1) is that the conditional probability distribution $p_\theta(h|v)$ in (1) is replaced with a Hermiticity-preserving, trace-preserving map $\Sigma^\theta_{v\to vh}$ that is reminiscent of, yet different from, the rotated Petz recovery map of [Wil15], the twirled Petz recovery map of [JRS+18], and the state-over-time of [FP22] (see also [LN24]). In the case that both the visible and hidden units are classical, the expression reduces to the classical expression in (1).

I also consider the case of quantum visible units and classical hidden units, a case of interest and possibly easier to implement in practice than fully quantum Boltzmann machines. For this scenario, I report an analytical expression for the gradient (see [Wil25a, Theorem 14]), along with a quantum algorithm for estimating it. The main distinction between this quantum–classical case and the fully classical case is that the conditional probability distribution $p_\theta(h|v)$ in (1) is replaced with the Hermiticity-preserving, trace-preserving map $\sum_x p_x(\theta)\Sigma_v^{\theta,x} \otimes |x\rangle\langle x|_h$, as defined in [Wil25a, Theorem 14]. This map is reminiscent of, yet different from, the rotated pretty good instrument introduced in [Wil15, Remark 5.7].

Finally, I extend all of the results to apply to the case when the objective function is the Petz–Tsallis relative entropy of order $q \in (0,1) \cup (1,2]$, instead of the quantum relative entropy. Doing so involves developing an independent derivation of a formula for the matrix derivative of the power function, as given in [Wil25a, Lemma 4] below. This formula could find alternative applications in quantum information theory, and it is connected to the probability density functions that appeared in [JRS$^+$18, Lemma 3.2], the latter building on Hirschman's improvement of Hadamard's three-line theorem [Hir52].

Paper organization: The paper [Wil25a] is organized as follows:

Section 2 provides some preliminary material used throughout the paper. This includes a brief subsection on notation (Section 2.1) and a review of classical Boltzmann machines and the derivation of the gradient of the relative entropy (Section 2.2). Also, Section 2.3 presents various formulas for matrix derivatives, including an independent derivation of a formula for the matrix derivative of the power function (see Lemma 4).

Section 3 contains some of the main aforementioned results, including an analytical expression for the gradient of the quantum relative entropy between a target state ρ and the reduced state $\sigma_v(\theta)$ of the visible units of a quantum Boltzmann machine (Theorem 8), as well as a quantum algorithm for estimating it (Section 3.2). As noted above, a subroutine of this quantum algorithm simulates modular flow, and I establish an improvement of this subroutine in Appendix B.1, by invoking [QKR25, Lemma 3] and [GSLW19, Corollary 60 and Lemma 61]. I also evaluate the gradient expression for a model called restricted quantum Boltzmann machines (Section 3.3), and then I show how it reduces to the known expression for the gradient [KW17, Eq. (4)] when there are no hidden units (Section 3.4).

Section 4 specializes the results of Section 3 to the case when the visible units are quantum and the hidden units are classical. In particular, Section 4.1 provides a formula for the gradient of the quantum relative entropy in this case, Section 4.2 details a quantum algorithm for estimating the gradient, and Section 4.3 evaluates the formula for restricted quantum–classical Boltzmann machines.

Section 5 specializes the results of Section 3 to the case when the visible units are classical and the hidden units are quantum. The findings here are related to the recent findings of [DTP$^+$25], but it is instructive to see how this reduction follows from the fully quantum case presented in Section 3.

The final technical section of this paper is Section 6, in which I show how to generalize the findings in the whole paper to the case when the objective function is the Petz–Tsallis relative entropy with parameter $q \in (0,1) \cup (1,2]$. To derive the results here, I provide an independent derivation of a formula for the derivative of the matrix power function (see Lemma 4, with its proof in Appendix A).

In the conclusion (Section 7), I provide a brief summary of the main findings of this paper, followed by suggestions for future research.

References

[AAR+18] Mohammad H. Amin, Evgeny Andriyash, Jason Rolfe, Bohdan Kulchytskyy, and Roger Melko. Quantum Boltzmann machine. *Physical Review X*, 8:021050, May 2018. URL: https://link.aps.org/doi/10.1103/PhysRevX.8.021050, doi:10.1103/PhysRevX.8.021050.

[AHS85] David H. Ackley, Geoffrey E. Hinton, and Terrence J. Sejnowski. A learning algorithm for Boltzmann machines. *Cognitive Science*, 9(1):147–169, 1985. doi:10.1016/S0364-0213(85)80012-4.

[BRGBPO17] Marcello Benedetti, John Realpe-Gómez, Rupak Biswas, and Alejandro Perdomo-Ortiz. Quantum-assisted learning of hardware-embedded probabilistic graphical models. *Physical Review X*, 7(4):041052, November 2017. URL: http://dx.doi.org/10.1103/PhysRevX.7.041052, doi:10.1103/physrevx.7.041052.

[CB24] Luuk Coopmans and Marcello Benedetti. On the sample complexity of quantum Boltzmann machine learning. *Communications Physics*, 7(1):274, 2024. doi:10.1038/s42005-024-01763-x.

[DTGJ25] Maria Demidik, Cenk Tüysüz, Michele Grossi, and Karl Jansen. Sample-based training of quantum generative models, 2025. URL: https://arxiv.org/abs/2511.11802, arXiv:2511.11802.

[DTP+25] Maria Demidik, Cenk Tüysüz, Nico Piatkowski, Michele Grossi, and Karl Jansen. Expressive equivalence of classical and quantum restricted Boltzmann machines. *Communications Physics*, 8(1):413, 2025. doi:10.1038/s42005-025-02353-1.

[FP22] James Fullwood and Arthur J. Parzygnat. On quantum states over time. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 478(2264):20220104, 08 2022. doi:10.1098/rspa.2022.0104.

[GSLW19] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, page 193–204, New York, NY, USA, 2019. Association for Computing Machinery. arXiv:1806.01838, doi:10.1145/3313276.3316366.

[Hir52] Isidore Isaac Hirschman. A convexity theorem for certain groups of transformations. *Journal d'Analyse Mathématique*, 2(2):209–218, 1952. doi:10.1007/BF02825637.

[HS86] Geoffrey E. Hinton and Terrence J. Sejnowski. Learning and relearning in Boltzmann machines. In David E. Rumelhart and James L. McClelland, editors, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, volume 1, pages 282–317. MIT Press, 1986. doi:10.7551/mitpress/5236.003.0011.

[Jar04] Christopher Jarzynski. Nonequilibrium work theorem for a system strongly coupled to a thermal environment. *Physical Review E*, 69(2):021101, 2004. doi:10.1103/PhysRevE.69.021101.

[JRS+18] Marius Junge, Renato Renner, David Sutter, Mark M. Wilde, and Andreas Winter. Universal recovery maps and approximate sufficiency of quantum relative entropy. *Annales Henri Poincaré*, 19(10):2955–2978, 2018. doi:10.1007/s00023-018-0716-0.

[KKH25] Takeshi Kimura, Kohtaro Kato, and Masahito Hayashi. Structured quantum learning via em algorithm for Boltzmann machines, 2025. URL: https://arxiv.org/abs/2507.21569, arXiv:2507.21569.

[KW17] Maria Kieferova and Nathan Wiebe. Tomography and generative data modeling via quantum Boltzmann training. *Physical Review A*, 96(6):062327, December 2017. arXiv:1612.05204, doi:10.1103/PhysRevA.96.062327.

[LK25] Ian T. Lim and Isaac H. Kim. A quantum algorithm for modular flow, 2025. URL: https://arxiv.org/abs/2508.16826v1, arXiv:2508.16826v1.

[LN24] Seok Hyung Lie and Nelly H. Y. Ng. Quantum state over time is unique. *Physical Review Research*, 6:033144, August 2024. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.6.033144, doi:10.1103/PhysRevResearch.6.033144.

[PKPW25] Dhrumil Patel, Daniel Koch, Saahil Patel, and Mark M. Wilde. Quantum Boltzmann machine learning of ground-state energies, 2025. URL: https://arxiv.org/abs/2410.12935v3, arXiv:2410.12935v3.

[QKR25] Yixian Qiu, Kelvin Koor, and Patrick Rebentrost. The quantum Esscher transform, 2025. URL: https://arxiv.org/abs/2401.07561v2, arXiv:2401.07561v2.

[TH20] Peter Talkner and Peter Hänggi. Colloquium: Statistical mechanics and thermodynamics at strong coupling: Quantum and classical. *Reviews of Modern Physics*, 92(4):041002, 2020. doi:10.1103/RevModPhys.92.041002.

[VSKB25] Adit Vishnu, Abhay Shastry, Dhruva Kashyap, and Chiranjib Bhattacharyya. Density operator expectation maximization, 2025. URL: https://arxiv.org/abs/2507.22786, arXiv:2507.22786.

[Wil15] Mark M. Wilde. Recoverability in quantum information theory. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 471(2182):20150338, 10 2015. doi:10.1098/rspa.2015.0338.

[Wil25a] Mark M. Wilde. Fundamentals of quantum Boltzmann machine learning with visible and hidden units, 2025. URL: https://arxiv.org/abs/2512.19819v1, arXiv:2512.19819v1.

[Wil25b] Mark M. Wilde. Generative modeling using evolved quantum Boltzmann machines, 2025. URL: https://arxiv.org/abs/2512.02721, arXiv:2512.02721.