

Double-bracket quantum algorithms for thermal state preparation

Andrew Wright, Reyhaneh Aghaei Saem, Supanut Thanasilp,
Yudai Suzuki, Zoë Holmes

This work is based on Ref. [1]

Thermal state preparation is a critical primitive in quantum computation, with direct applications in quantum machine learning and optimization [2, 3], such as training Quantum Boltzmann Machines (QBM) [4, 5] for generative modeling. While the preparation of thermal states is QMA-hard in the worst case [6], quantum computing may still offer significant advantages over classical algorithms for typical problem instances, making the development of such quantum algorithms a major subject of research. While fault-tolerant algorithms based on quantum singular value transformation and linear combinations of unitaries offer rigorous guarantees, they require prohibitive resources like block-encodings [7–13], making them unsuitable for near-term devices. Conversely, variational quantum imaginary-time evolution (VarQITE) and related approaches [14–19] operate well within near-term constraints but are hindered by shot noise, numerical instability, and the fact that good ansatzes are largely unknown.

To bridge the gap between resource-heavy fault-tolerant methods and heuristic variational approaches, we introduce the *double-bracket thermofield double (DB-TFD)* framework, a novel, non-variational method that prepares thermal states deterministically. We expand on recent results of double-bracket quantum algorithms [20], which have previously been shown to successfully approximate imaginary-time evolution [21] and apply general polynomial transformations [22]. Building on these capabilities, our approach utilizes the double-bracket iterative framework to simulate imaginary-time evolution on a TFD state to achieve scalable thermal state preparation. We present two variants: *vanilla* DB-TFD and *poly* DB-TFD, and provide a rigorous complexity analysis showing its scalability. Furthermore, we demonstrate the practical utility of DB-TFD by applying it to train QBM for generative modeling, outperforming standard variational approaches under realistic noise models. Specifically, DB-TFD achieves lower generalization error and exhibits strong resilience to shot noise, a major bottleneck for existing variational techniques.

The DB-TFD Framework and Complexity

The core idea of DB-TFD is to prepare a highly entangled TFD state at infinite temperature, $|\text{TFD}(0)\rangle$, which is a maximally entangled state between a target subsystem A and an auxiliary subsystem B of the same size, and subsequently cool it down to the target inverse temperature β via imaginary-time evolution. This yields the target thermal state:

$$|\text{TFD}(\beta)\rangle = \frac{1}{\sqrt{\mathcal{N}}} e^{-\beta(H \otimes \mathbb{1})/2} |\text{TFD}(0)\rangle, \quad (1)$$

where $\mathcal{N}$ is the normalization factor. Tracing out the auxiliary register yields the thermal Gibbs state of H . We implement this non-unitary imaginary-time evolution deterministically by formulating it as a double-bracket quantum algorithm. Originally developed for continuous matrix flow problems [20], double-bracket iterations apply sequences of unitary transformations driven by the commutator between the projector onto the current state $\psi_k = |\psi_k\rangle\langle\psi_k|$ and the Hamiltonian H . The state is updated recursively as:

$$|\psi_{k+1}\rangle = e^{s_k[\psi_k, H]} |\psi_k\rangle, \quad (2)$$

which coincides with imaginary-time evolution in the continuous limit.

Because exactly exponentiating the commutators $e^{s_k[\psi_k, H]}$ is intractable natively, we approximate these double-bracket rotations (setting $s_k = t^2$) using m -th order product formulas [23–25] defined as:

$$e^{t^2[\psi_k, H]} = e^{i\alpha_1 t H} e^{i\alpha_2 t \psi_k} e^{i\alpha_3 t H} \dots e^{i\alpha_p t \psi_k} + \mathcal{O}(t^{m+1}). \quad (3)$$

This mathematical decomposition allows us to break down the complex evolution into alternating sequences of Hamiltonian evolutions $e^{i\alpha t H}$ and state reflections $e^{i\alpha t \psi_k}$. However, each state reflection $e^{i\alpha t \psi_k}$ requires executing the exponential of the density matrix, which inherently depends on preparing the state $|\psi_k\rangle$. Within the product formula, this recursive structure induces an exponential growth in the number of required queries with respect to the number of steps k . Consequently, this limits the number of iterations to only a few in practice. In our work, we develop two distinct variants to approximate the ITE flow:

1. **vanilla DB-TFD:** This approach implements a first-order approximation of imaginary-time evolution by double-bracket flow. The target state is approached via iterations:

$$|\psi_k^{(\text{vanilla})}\rangle = \prod_{i=0}^{k-1} e^{\frac{\beta}{2k} [\psi_i^{(\text{vanilla})}, H]} |\text{TFD}(0)\rangle, \quad (4)$$

where $\psi_i^{(\text{vanilla})} = |\psi_i^{(\text{vanilla})}\rangle\langle\psi_i^{(\text{vanilla})}|$. In the limit $k \rightarrow \infty$, this sequence of double-bracket rotations exactly reproduces the continuous imaginary-time evolution.

2. **poly DB-TFD:** To achieve better complexity, this variant employs double-bracket quantum signal processing to approximate the exponential function of imaginary-time evolution using a polynomial transformation of degree k , leveraging the fact that polynomial approximations of the exponential operator are already well known [26]. The iterations take the form:

$$|\psi_k^{(\text{poly})}\rangle = \prod_{i=0}^{k-1} e^{i\theta_i \psi_i^{(\text{poly})}} e^{s_i [\psi_i^{(\text{poly})}, H]} |\text{TFD}(0)\rangle, \quad (5)$$

where z_i are the complex roots of an efficient polynomial approximation of the exponential [26], and the parameters s_i and θ_i are determined by the energy $E_i = \langle\psi_i|H|\psi_i\rangle$ and variance $V_i = \langle\psi_i|(H - E_i)^2|\psi_i\rangle$ of the intermediate states according to:

$$s_i = -\frac{1}{\sqrt{V_i}} \arccos\left(\frac{|E_i - z_i|}{\sqrt{V_i + |E_i - z_i|^2}}\right), \quad \theta_i = \arg\left(\frac{E_i - z_i}{|E_i - z_i|}\right). \quad (6)$$

Crucially, because these parameters depend on the intermediate state properties, the energy E_i and variance V_i must be measured after each iteration. However, this method requires exponentially fewer iterations compared to the *vanilla* approach, significantly reducing the overall circuit depth at the cost of these intermediate measurements.

We can formally state the main error bound regarding the required number of oracle calls as a theorem:

Theorem 1 (Oracle calls to Hamiltonian evolutions and reflection operators) *To achieve an ε -precision approximation of the thermofield double state at inverse temperature β using the poly DB-TFD algorithm, at least N_{tot} oracle calls to Hamiltonian evolutions and reflection operators are required. Under a broad practical regime, the scalings of the poly approach is given by:*

$$N_{\text{tot}} \in \exp\left(\tilde{\mathcal{O}}(\beta, \log(1/\varepsilon))\right),$$

This query complexity for *poly* DB-TFD is exponential in β , which is in line with prior state-of-the-art methods [10, 12, 13]. This result is derived by accounting for the errors induced by the product formula and polynomial approximation error.

Application to Quantum Generative Modeling

We apply DB-TFD to generative modeling tasks using QBMs. The goal of generative modeling is to generate a target distribution η by training a parameterized quantum state. This is typically achieved by minimizing the quantum relative entropy $S(\eta||\rho_\theta)$ between the target and a model thermal state $\rho_\theta = e^{-H_\theta}/Z$, where $Z = \text{Tr}(e^{-H_\theta})$ is the partition function. Parameter updates in this gradient descent optimization require estimating the expectation values of all Hamiltonian generators over the thermal state ρ_θ .

We benchmark the performance of *poly* DB-TFD against VarQITE [19, 27] across three generative modeling tasks: an XXZ spin chain, a Bernoulli distribution, and a 2×2 Bars and Stripes dataset (see Fig. 1). The observed performance gap between *poly* DB-TFD and VarQITE under noisy settings can be attributed to their different resource requirements. We note that both methods use N_{shots} number of measurements shots per observable to estimate the set of expectation values $\{\langle G_i \rangle_{\rho_\theta}\}_{i=1}^M$, which are used to update the parameters θ . This suggests that the difference ultimately reduces to the cost of the thermal state preparation required at each step.

In VarQITE, a parametrized quantum circuit with L parameters is iteratively applied for $k^{(\text{V})}$ steps to prepare an approximation of the thermal state. This requires estimating all entries of the $L \times L$ matrix A (the




Figure 1: Comparison of final training losses for QBM generative modeling tasks using the exact thermal state, *poly* DB-TFD, and VarQITE under noiseless, shot noise, and depolarizing noise conditions. For DB-TFD, k denotes the number of double-bracket iterations and m indicates the order of the product formula used. For VarQITE, l represents the number of ansatz layers. All tasks prepare a thermal state for 4 target qubits (an 8-qubit total system size). We observe that shot noise heavily affects VarQITE due to the much larger number of parameters and circuit repetitions it requires, while depolarizing noise primarily affects DB-TFD due to its greater circuit depth. Moreover, a very large number of shots is required to achieve a good loss for VarQITE. While we have here chosen to display the effects of shot noise and hardware noise independently to illustrate their respective effects, we expect that DB-TFD outperforms VarQITE for a wide parameter range where noise levels are modest and shot noise is substantial.

Fubini-Study metric tensor) and the vector C of size L (the energy gradient) via McLachlan’s variational principle. Therefore, obtaining the information required to update the parameters entails $\mathcal{O}((L^2 + LM)k^{(V)})$ circuit execution at each optimization step. In contrast, *poly* DB-TFD avoids this highly parameterized ansatz optimization altogether. It only requires estimating the system’s overall energy and its variance at each step, demanding measurements of the Pauli operators associated with the Hamiltonian H and its square H^2 . The corresponding measurement cost therefore scales as $\mathcal{O}((M^2 + M)k^{(\text{poly})})$, where $k^{(\text{poly})}$ denotes the required polynomial degree. In practice, the number of parameters L in VarQITE may exceed the number of Hamiltonian terms M to faithfully reproduce the imaginary-time evolution dynamics. Under this condition, and for a fixed iteration count, *poly* DB-TFD requires fewer measurements.

To assess performance under realistic hardware imperfections, we further subject the algorithms to depolarizing noise. In our noise model, a depolarizing channel acts independently on each elementary gate layer of the VarQITE ansatz, and on each bulk Hamiltonian evolution or reflection gate for DB-TFD, with error rates $q \in \{10^{-6}, 10^{-5}\}$. While VarQITE’s fixed circuit depth affords it some resilience against depolarizing noise, its optimization trajectory degrades severely under limited measurement budgets (N_{shots}). Conversely, our numerical experiments demonstrate that *poly* DB-TFD achieves lower generalization error and is much more resilient to shot noise. Note that due to computational constraints, we do not simulate both sources of noise simultaneously. However, given how severely VarQITE is affected by shot noise, we expect DB-TFD to outperform it overall. Although DB-TFD is more affected by gate noise, this can potentially be solved by quantum error correction, whereas the shot-noise bottleneck faced by variational methods constitutes a more fundamental limit.

Conclusion

We have introduced the DB-TFD framework as a deterministic, non-variational approach to thermal state preparation. For Quantum Machine Learning applications, such as generative modeling with QBMs, DB-TFD offers compelling advantages over traditional variational methods. By eliminating the need for heuristic ansatz selection and avoiding the measurement overhead required for iterative parameter updates, our method demonstrates significant resilience to statistical shot noise. Although DB-TFD requires larger circuit depths, its robustness to shot-noise positions it as a powerful and scalable tool for advancing generative Quantum Machine Learning on late-NISQ and early fault-tolerant quantum processors.

References

- [1] Andrew Wright, Reyhaneh Aghaei Saem, Supanut Thanasilp, Yudai Suzuki, and Zoë Holmes. Double-bracket quantum algorithms for thermal state preparation, 2026. URL: <https://arxiv.org/abs/>

2606.05947, [arXiv:2606.05947](https://arxiv.org/abs/2606.05947).

- [2] Fernando G. S. L. Brandao and Krysta Svore. Quantum speed-ups for semidefinite programming, 2017. URL: <https://arxiv.org/abs/1609.05537>, [arXiv:1609.05537](https://arxiv.org/abs/1609.05537).
- [3] Lester Ingber. Adaptive simulated annealing (asa): Lessons learned. *arXiv preprint cs/0001018*, 2000.
- [4] Mohammad H. Amin, Evgeny Andriyash, Jason Rolfe, Bohdan Kulchytskyy, and Roger Melko. Quantum boltzmann machine. *Physical Review X*, 8(2), May 2018. URL: <http://dx.doi.org/10.1103/PhysRevX.8.021050>, [doi:10.1103/physrevx.8.021050](https://doi.org/10.1103/physrevx.8.021050).
- [5] Luuk Coopmans and Marcello Benedetti. On the sample complexity of quantum boltzmann machine learning. *Communications Physics*, 7(1), August 2024. URL: <http://dx.doi.org/10.1038/s42005-024-01763-x>, [doi:10.1038/s42005-024-01763-x](https://doi.org/10.1038/s42005-024-01763-x).
- [6] Dorit Aharonov, Daniel Gottesman, Sandy Irani, and Julia Kempe. The power of quantum systems on a line. *Communications in Mathematical Physics*, 287(1):41–65, January 2009. URL: <http://dx.doi.org/10.1007/s00220-008-0710-3>, [doi:10.1007/s00220-008-0710-3](https://doi.org/10.1007/s00220-008-0710-3).
- [7] Chi-Fang Chen, Michael J. Kastoryano, Fernando G. S. L. Brandão, and András Gilyén. Quantum thermal state preparation, 2023. URL: <https://arxiv.org/abs/2303.18224>, [arXiv:2303.18224](https://arxiv.org/abs/2303.18224).
- [8] Patrick Rall, Chunhao Wang, and Pawel Wocjan. Thermal state preparation via rounding promises. *Quantum*, 7:1132, October 2023. URL: <http://dx.doi.org/10.22331/q-2023-10-10-1132>, [doi:10.22331/q-2023-10-10-1132](https://doi.org/10.22331/q-2023-10-10-1132).
- [9] Zhiyan Ding, Yongtao Zhan, John Preskill, and Lin Lin. End-to-end efficient quantum thermal and ground state preparation made simple, 2025. URL: <https://arxiv.org/abs/2508.05703>, [arXiv:2508.05703](https://arxiv.org/abs/2508.05703).
- [10] 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 '19, page 193–204. ACM, June 2019. URL: <http://dx.doi.org/10.1145/3313276.3316366>, [doi:10.1145/3313276.3316366](https://doi.org/10.1145/3313276.3316366).
- [11] Zoe Holmes, Gopikrishnan Muraleedharan, Rolando D. Somma, Yigit Subasi, and Burak Şahinoğlu. Quantum algorithms from fluctuation theorems: Thermal-state preparation. *Quantum*, 6:825, October 2022. URL: <http://dx.doi.org/10.22331/q-2022-10-06-825>, [doi:10.22331/q-2022-10-06-825](https://doi.org/10.22331/q-2022-10-06-825).
- [12] Anirban Narayan Chowdhury and Rolando D. Somma. Quantum algorithms for gibbs sampling and hitting-time estimation, 2016. URL: <https://arxiv.org/abs/1603.02940>, [arXiv:1603.02940](https://arxiv.org/abs/1603.02940).
- [13] David Poulin and Pawel Wocjan. Sampling from the thermal quantum gibbs state and evaluating partition functions with a quantum computer. *Physical Review Letters*, 103(22), November 2009. URL: <http://dx.doi.org/10.1103/PhysRevLett.103.220502>, [doi:10.1103/physrevlett.103.220502](https://doi.org/10.1103/physrevlett.103.220502).
- [14] Mario Motta, Chong Sun, Adrian T. K. Tan, Matthew J. O’Rourke, Erika Ye, Austin J. Minnich, Fernando G. S. L. Brandão, and Garnet Kin-Lic Chan. Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution. *Nature Physics*, 16(2):205–210, November 2019. URL: <http://dx.doi.org/10.1038/s41567-019-0704-4>, [doi:10.1038/s41567-019-0704-4](https://doi.org/10.1038/s41567-019-0704-4).
- [15] Jingxiang Wu and Timothy H Hsieh. Variational thermal quantum simulation via thermofield double states. *Physical review letters*, 123(22):220502, 2019.
- [16] D. Zhu, S. Johri, N. M. Linke, K. A. Landsman, C. Huerta Alderete, N. H. Nguyen, A. Y. Matsuura, T. H. Hsieh, and C. Monroe. Generation of thermofield double states and critical ground states with a quantum computer. *Proceedings of the National Academy of Sciences*, 117(41):25402–25406, September 2020. URL: <http://dx.doi.org/10.1073/pnas.2006337117>, [doi:10.1073/pnas.2006337117](https://doi.org/10.1073/pnas.2006337117).
- [17] Troy J. Sewell, Christopher David White, and Brian Swingle. Thermal multi-scale entanglement renormalization ansatz for variational gibbs state preparation, 2022. URL: <https://arxiv.org/abs/2210.16419>, [arXiv:2210.16419](https://arxiv.org/abs/2210.16419).
- [18] R. Sagastizabal, S. P. Premaratne, B. A. Klaver, M. A. Rol, V. Negîrneac, M. S. Moreira, X. Zou, S. Johri, N. Muthusubramanian, M. Beekman, C. Zachariadis, V. P. Ostroukh, N. Haider, A. Bruno, A. Y. Matsuura, and L. DiCarlo. Variational preparation of finite-temperature states on a quantum computer. *npj Quantum Information*, 7(1), August 2021. URL: <http://dx.doi.org/10.1038/s41534-021-00468-1>, [doi:10.1038/s41534-021-00468-1](https://doi.org/10.1038/s41534-021-00468-1).
- [19] Christa Zoufal, Aurélien Lucchi, and Stefan Woerner. Variational quantum boltzmann machines. *Quantum Machine Intelligence*, 3(1), February 2021. URL: <http://dx.doi.org/10.1007/s42484-020-00033-7>, [doi:10.1007/s42484-020-00033-7](https://doi.org/10.1007/s42484-020-00033-7).
- [20] Marek Gluza. Double-bracket quantum algorithms for diagonalization. *Quantum*, 8:1316, April 2024. URL: <http://dx.doi.org/10.22331/q-2024-04-09-1316>, [doi:10.22331/q-2024-04-09-1316](https://doi.org/10.22331/q-2024-04-09-1316).

- [21] Marek Gluza, Jeongrak Son, Bi Hong Tiang, Yudai Suzuki, Zoë Holmes, and Nelly HY Ng. Double-bracket quantum algorithms for quantum imaginary-time evolution. *arXiv preprint arXiv:2412.04554*, 2024.
- [22] Yudai Suzuki, Bi Hong Tiang, Jeongrak Son, Nelly H. Y. Ng, Zoë Holmes, and Marek Gluza. Double-bracket algorithm for quantum signal processing without post-selection, 2025. URL: <https://arxiv.org/abs/2504.01077>, [arXiv:2504.01077](https://arxiv.org/abs/2504.01077).
- [23] Andrew M. Childs and Nathan Wiebe. Product formulas for exponentials of commutators. *Journal of Mathematical Physics*, 54(6), June 2013. URL: <http://dx.doi.org/10.1063/1.4811386>, doi: [10.1063/1.4811386](https://doi.org/10.1063/1.4811386).
- [24] F. Casas, A. Escorihuela-Tomàs, and P. A. Moreno Casares. Approximating exponentials of commutators by optimized product formulas. *Quantum Information Processing*, 24(2), January 2025. URL: <http://dx.doi.org/10.1007/s11128-025-04659-z>, doi: [10.1007/s11128-025-04659-z](https://doi.org/10.1007/s11128-025-04659-z).
- [25] Yu-An Chen, Andrew M. Childs, Mohammad Hafezi, Zhang Jiang, Hwanmun Kim, and Yijia Xu. Efficient product formulas for commutators and applications to quantum simulation. *Physical Review Research*, 4(1), March 2022. URL: <http://dx.doi.org/10.1103/PhysRevResearch.4.013191>, doi: [10.1103/physrevresearch.4.013191](https://doi.org/10.1103/physrevresearch.4.013191).
- [26] Sushant Sachdeva and Nisheeth Vishnoi. Approximation theory and the design of fast algorithms, 2013. URL: <https://arxiv.org/abs/1309.4882>, [arXiv:1309.4882](https://arxiv.org/abs/1309.4882).
- [27] Sam McArdle, Tyson Jones, Suguru Endo, Ying Li, Simon C. Benjamin, and Xiao Yuan. Variational ansatz-based quantum simulation of imaginary time evolution. *npj Quantum Information*, 5(1), September 2019. URL: <http://dx.doi.org/10.1038/s41534-019-0187-2>, doi: [10.1038/s41534-019-0187-2](https://doi.org/10.1038/s41534-019-0187-2).