



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Purifying Deep Boltzmann Machines for Thermal Quantum States

Yusuke Nomura, Nobuyuki Yoshioka, and Franco Nori

Phys. Rev. Lett. **127**, 060601 — Published 4 August 2021

DOI: [10.1103/PhysRevLett.127.060601](https://doi.org/10.1103/PhysRevLett.127.060601)

Purifying Deep Boltzmann Machines for Thermal Quantum States

Yusuke Nomura^{1*}

¹*RIKEN Center for Emergent Matter Science, 2-1 Hirosawa, Wako, Saitama 351-0198, Japan*

Nobuyuki Yoshioka^{2,3†} and Franco Nori^{3,4,5}

²*Department of Applied Physics, University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-8656, Japan*

³*Theoretical Quantum Physics Laboratory, RIKEN Cluster for Pioneering Research (CPR), Wako-shi, Saitama 351-0198, Japan*

⁴*RIKEN Center for Quantum Computing (RQC), Wako-shi, Saitama 351-0198, Japan and*

⁵*Physics Department, University of Michigan, Ann Arbor, Michigan 48109-1040, USA*

(Dated: June 9, 2021)

We develop two cutting-edge approaches to construct deep neural networks representing the purified finite-temperature states of quantum many-body systems. Both methods commonly aim to represent the Gibbs state by a highly expressive neural-network wave function, exemplifying the idea of purification. The first method is an entirely deterministic approach to generate deep Boltzmann machines representing the purified Gibbs state exactly. This strongly assures the remarkable flexibility of the ansatz which can fully exploit the quantum-to-classical mapping. The second method employs stochastic sampling to optimize the network parameters such that the imaginary time evolution is well approximated within the expressibility of neural networks. Numerical demonstrations for transverse-field Ising models and Heisenberg models show that our methods are powerful enough to investigate the finite-temperature properties of strongly correlated quantum many-body systems, even when the problematic effect of frustration is present.

Introduction.— The thermal behavior of quantum many-body systems is one of the most fundamental problems in physics. Statistical mechanics states that the density matrix describing a system in thermal equilibrium, governed by a Hamiltonian \mathcal{H} at an inverse temperature β , is given by the Gibbs state $\rho = e^{-\beta\mathcal{H}}/\text{Tr}[e^{-\beta\mathcal{H}}]$. Computing and extracting physical properties from the Gibbs state is a significant challenge to understand natural phenomena, which in reality all occur at finite β .

One of the most celebrated numerical techniques in lattice systems is the quantum Monte Carlo (QMC) method [1–4], typically based on the path integral formalism of the partition function. The QMC method yields numerically exact solutions when the positive definiteness is assured; otherwise, the infamous negative sign problem arises. Many physically intriguing system falls into the latter category, and therefore various efforts have been devoted to overcome this difficulty: tensor-network-based algorithms [5–8] mostly applied to one-dimensional (1D) systems, dynamical mean-field theory [9] which becomes exact in the infinite coordination-number limit, diagrammatic Monte Carlo methods [10, 11], to name a few [12–14]. In another notable approach [15–18] using the thermal pure quantum (TPQ) states, one can extract the ensemble property from a single pure state that represents the thermal equilibrium. We point out, however, that it remains extremely challenging to establish a methodology that is both reliable and scalable for finite-temperature calculations in two-dimensional (2D) systems—the most exotic and intriguing realm in quantum many-body problems.

Neural networks, initially developed for classical data processing in the context of machine learning, offer a very strong methodology for quantum physics [19–25]. As was firstly demonstrated by Carleo and Troyer [19], neural networks applied as variational wave functions, commonly dubbed as the neural quantum states, are capable of simulating ground

states [19, 26–28], excited states [29–31], and even out-of-equilibrium property [19, 32–38] of strongly correlated systems up to unprecedentedly large size. Among the tremendous variety of network structures, Boltzmann machines with restricted connectivity, known to be universal approximators for arbitrary real/complex amplitudes [39, 40], are useful for statistical mechanics and quantum information. The Boltzmann machines with the shallowest structure are already powerful enough to compactly express complex quantum states with extensively-growing quantum entanglement [41, 42]. Furthermore, deep Boltzmann machines (DBMs), i.e., the ones with multiple hidden layers, are guaranteed to provide an efficient description for an even wider range of quantum states [43]. Strongly motivated by their extremely high representability, the ground states in quantum many-body spin systems have been successfully simulated by DBMs [44].

In the present study, we provide two state-of-the-art methods to construct DBMs that capture the finite-temperature behavior of quantum many-body systems. Both methods share the strategy of employing DBMs to express the purified Gibbs state. Namely, a mixed state under imaginary time evolution is compactly encoded as a pure DBM wave function in the enlarged Hilbert space. In the first method, we find a completely deterministic way to construct DBMs, realizing the exact purified expression of finite temperature states. This proves the remarkable flexibility and power of DBMs for investigating finite-temperature many-body phenomena. In the second method, we provide a stochastic way to simulate the imaginary time evolution which exploits the versatile expressive power of DBMs as approximators. For demonstration, we apply these methods to the 1D transverse-field Ising (TFI) model and the 2D J_1 – J_2 Heisenberg model on the square lattice to find surprisingly high accuracy compared to numerically exact methods. We emphasize that only a polynomial number of auxiliary spins suffices in both approaches, yield-

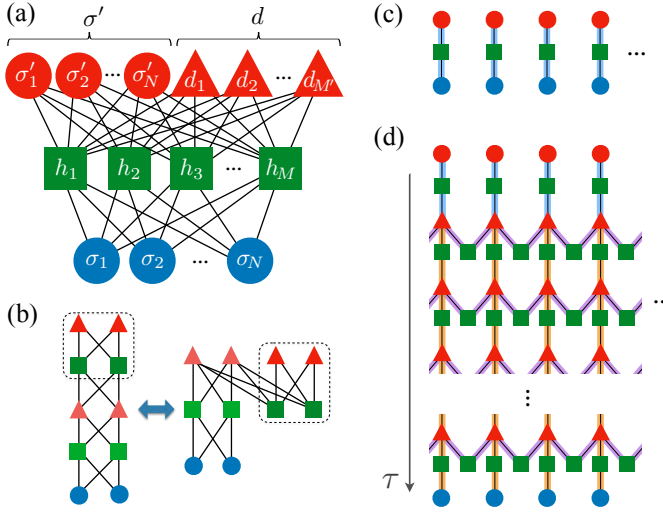


FIG. 1. (a) Structure of the three-layer DBM used in the present study. The visible layer (blue) corresponds to physical spins σ_i . Two hidden layers are distinguished as hidden (green) and deep (red) layers; while the hidden layer simply consists of hidden spins h_j , the deep layer is composed of the deep spins d_k and ancilla spins σ'_i , which are introduced to purify the Gibbs-state density matrix. Here, the number of physical/ancilla, hidden, and deep spins are denoted as N , M , M' , respectively. (b) Two different ways of depicting the identical DBM structure. A network with seemingly many hidden layers can always be recast into that with only two hidden layers. (c) The DBM representing the infinite-temperature state $|\Psi(T = \infty)\rangle$ for quantum spin-1/2 systems. (d) The DBM construction for the finite-temperature states of the 1D TFI model $|\Psi(T)\rangle$ in Eq. (3). The arrow denotes the growth of the DBM structure along the imaginary-time τ propagation. Light blue, orange, and purple bonds in (c) and (d) have couplings $i\frac{\pi}{4}$, $\frac{1}{2}\text{arcosh}\left(\frac{1}{\tanh(\Gamma\delta\tau)}\right)$, and $\frac{1}{2}\text{arcosh}(e^{2J\delta\tau})$, respectively.

ing a huge computational advantage even under the presence of the problematic effect of frustration.

DBM for purification.— Throughout this work, we utilize the idea of purification to represent the Gibbs state. Namely, the finite-temperature density matrix ρ of a target system \mathcal{S} is encoded as a pure state in an extended system $\mathcal{S} + \mathcal{A}$, such that ρ is obtained by tracing out the ancillary system \mathcal{A} [45]. For instance, the purification of an infinite-temperature state can be generated by the superposition $\sum_x |x\rangle_{\mathcal{S}} |a_x\rangle_{\mathcal{A}}$, where $\{|x\rangle_{\mathcal{S}}\}$ is the complete orthonormal basis set of the target system, while $\{|a_x\rangle_{\mathcal{A}}\}$ is an orthonormal but not necessarily complete basis set of ancillary system.

For the sake of concreteness, let us consider a quantum many-body spin-1/2 system. We introduce N_{site} binary degrees of freedom $\{\sigma_i\}$ so that $|\sigma\rangle = |\sigma_1, \dots, \sigma_{N_{\text{site}}}\rangle$ spans the Hilbert space of the target system \mathcal{S} . Hereafter we call them physical spins. As the ancillary system \mathcal{A} , we introduce an identical number of ancilla spins $\{\sigma'_i\}$ so that an arbitrary mixed state can be purified in principle. While the purification of a mixed state is not unique, here we exclu-

sively take the purified infinite-temperature state as $|\Psi(T = \infty)\rangle = \bigotimes_{i=1}^{N_{\text{site}}} (|\uparrow\uparrow'\rangle + |\downarrow\downarrow'\rangle)_i$, and perform the imaginary time evolution as $|\Psi(T)\rangle = e^{-\beta\mathcal{H}/2} \otimes \mathbb{1}' |\Psi(T = \infty)\rangle$, with $\beta = 1/T$ to simulate the finite-temperature Gibbs state. Note that the infinite-temperature state $\rho_{\infty} = \mathbb{1}/2^{N_{\text{site}}}$ and the finite-temperature state $\rho_T = e^{-\beta\mathcal{H}}/\text{Tr}[e^{-\beta\mathcal{H}}]$ are reproduced by tracing out the ancilla spins.

Intriguingly, the purified Gibbs state at each temperature can be efficiently expressed by the DBM. In particular, we use the DBM with two hidden layers [see Fig. 1(a)] to represent a purified wave function whose amplitude $\Psi(\sigma, \sigma') = \langle \sigma, \sigma' | \Psi \rangle$ is parametrized as

$$\Psi(\sigma, \sigma') = \sum_{h,d} \phi(\sigma, \sigma'; h, d), \quad (1)$$

$$\phi(\sigma, \sigma'; h, d) = \exp \left[\sum_j b_j h_j + \sum_{ji} h_j (W_{ji} \sigma_i + W'_{ji} \sigma'_i) + \sum_{jk} W'_{jk} h_j d_k \right], \quad (2)$$

where we have introduced hidden spins $\{h_j\}$ and deep spins $\{d_k\}$, in addition to the physical spins $\{\sigma_i\}$ constituting the visible layer and ancilla spins $\{\sigma'_i\}$ allocated in the second hidden layer (we define all spins to be either +1 or -1). This structure is “universal” in terms of the DBM architecture; the arbitrary multi-hidden layer structure can be rearranged to have only two hidden layers as shown in Fig 1(b). The number of complex variational parameters $\mathcal{W} = \{b, W, W'\}$ [46] are directly related to the number of h and d spins, which controls the representability of the DBM wave function. Note that the purification technique for neural networks has been considered in the context of quantum tomography [47] and dissipative quantum physics [34–37].

Method (I): Analytic purification using DBM.— In Method (I), we analytically construct DBMs that exactly reproduce the behaviours of Gibbs states. First, to represent the infinite-temperature state, we introduce a DBM with N_{site} hidden and ancilla spins [Fig. 1(c)]. By setting the parameters \mathcal{W}_{∞} as $b_j = 0$ and $W_{ji} = W'_{ji} = i\frac{\pi}{4}\delta_{ji}$, we find that the analytical expression of the DBM wave function becomes $\Psi(\sigma, \sigma') = \prod_i 2 \cosh [i\frac{\pi}{4}(\sigma_i + \sigma'_i)]$. This exactly reproduces the $|\Psi(T = \infty)\rangle$ described above.

Next, to express finite-temperature states $|\Psi(T)\rangle$ analytically, we introduce the Suzuki-Trotter decomposition [1]. Namely, given the Hamiltonian $\mathcal{H} = \sum_{\nu} \mathcal{H}_{\nu}$, the purified Gibbs state up to the Trotter error is expressed as

$$|\Psi(T)\rangle := \left(\left[\prod_{\nu} e^{-\delta\tau \mathcal{H}_{\nu}} \right]^{N_{\tau}} \otimes \mathbb{1}' \right) |\Psi(T = \infty)\rangle, \quad (3)$$

where $\delta\tau = \beta/2N_{\tau}$ is the propagation time step. Remarkably, thanks to the flexible representability of the DBM, we can find analytical solutions for Eq. (3): Starting from the initial DBM state $|\Psi(T = \infty)\rangle$ given by the \mathcal{W}_{∞} parameters, each short-time propagation $e^{-\delta\tau \mathcal{H}_{\nu}}$ can be implemented exactly by modifying the parameters and structure of

the DBM [48], whose the explicit expression depends on the form of the Hamiltonian.

As a concrete example, let us consider the TFI model on the N_{site} -spin chain under periodic boundary condition. The Hamiltonian is given by $\mathcal{H} = \mathcal{H}_1 + \mathcal{H}_2$, with $\mathcal{H}_1 = -J \sum_i \sigma_i^z \sigma_{i+1}^z$ and $\mathcal{H}_2 = -\Gamma \sum_i \sigma_i^x$, where σ_i^a ($a = x, y, z$) denotes the Pauli matrix operating on the i -th site. We take the Ising-type interaction J as the energy unit ($J = 1$), and Γ as the strength of the transverse magnetic field. Using Method (I), we can analytically construct the finite-temperature state using the DBM. Solutions to realize the propagation by updating the DBM parameters from \mathcal{W} to $\tilde{\mathcal{W}}$, i.e., $e^{-\delta\tau \mathcal{H}_\nu} |\Psi_{\mathcal{W}}\rangle = C |\Psi_{\tilde{\mathcal{W}}}\rangle$ (C : a constant) can be sketched as follows [44]: For $\nu = 1$ (interaction propagator), we add a single hidden spin for each neighboring visible spins σ_i and σ_{i+1} . For $\nu = 2$ (transverse-field propagator), we add new hidden- and deep-spin layers between the visible and neighboring hidden layers. All the DBM parameters are determined analytically [see Supplementary Materials (SM) for detailed implementation [49]]. As a consequence of propagations, the DBM architecture grows as in Fig. 1(d), and the number of hidden and deep spins scale as $\mathcal{O}(N_\tau N_{\text{site}})$, which is also common among general local Hamiltonians.

Now that the Gibbs states are represented, let us discuss how the physical quantities are computed using the DBM framework in general. Since the expectation value of a physical observable $\langle \mathcal{O} \rangle = \frac{\langle \Psi(T) | \mathcal{O} \otimes \mathbb{1}' | \Psi(T) \rangle}{\langle \Psi(T) | \Psi(T) \rangle}$ becomes analytically intractable, we use the Monte Carlo (MC) method for its numerical estimation. The sampling weight is based on the expression of the normalization factor of the DBM state [Eq. (1)] given as

$$\begin{aligned} \langle \Psi | \Psi \rangle &= \sum_{\sigma, \sigma'} |\Psi(\sigma, \sigma')|^2 \\ &= \sum_{\sigma, \sigma'} \sum_{h_1, d_1, h_2, d_2} \phi^*(\sigma, \sigma'; h_1, d_1) \phi(\sigma, \sigma'; h_2, d_2). \end{aligned} \quad (4)$$

Namely, we sample over the configurations of $(\sigma, \sigma', h_1, h_2, d_1, d_2)$ weighted by the product of amplitudes as $w(\sigma, \sigma'; h, d) \equiv \phi^*(\sigma, \sigma'; h_1, d_1) \phi(\sigma, \sigma'; h_2, d_2)$. Alternatively, it is possible to trace out the hidden spins h analytically, and use $\sum_{h_1, h_2} \phi^*(\sigma, \sigma'; h_1, d_1) \phi(\sigma, \sigma'; h_2, d_2)$ as the MC sampling weight over configurations of $(\sigma, \sigma', d_1, d_2)$ (or trace out d spins and sample over h spins). See SM for more details of the method [49].

To verify the construction of the DBM and the proposed MC sampling strategy, we apply the method to the 16-site TFI model. Figure 2 shows the DBM results (symbols) for the temperature dependence of the (a) energy, (b) specific heat, and (c) susceptibility. As expected, the DBM results follow the exact temperature evolution (solid curves). This confirms the remarkable representability of the DBM not only at zero temperature [44] but also at finite temperatures, offering an intriguing quantum-to-classical mapping.

Method (II): Numerical purification using DBM.— When the weight of individual spin configuration can be taken to




FIG. 2. Finite-temperature calculations for the 1D TFI model on a 16-site chain ($N_{\text{site}} = 16$) with periodic boundary condition: (a) Energy E , (b) specific heat C , and (c) susceptibility $\chi = \frac{1}{T} \sum_i \langle \sigma_i \sigma_i \rangle$. The symbols denote the DBM results [Method (I) with $\delta\tau = 0.05$], which agree well with the exact-diagonalization results (solid curves).

be always positive, Method (I) is quite useful and provides numerically exact finite-temperature results. However, when the frustration exists in the spin Hamiltonian, for instance, we cannot avoid the existence of negative weights as in other finite-temperature calculations based on the QMC method. While it is possible to construct finite-temperature states analytically, the estimation of physical quantities becomes extremely difficult because of the negative sign problem.

To make the application to frustrated models possible, we propose an alternative method which employs DBMs with only ancilla spins σ' in the second hidden layer. By tracing out hidden spins h , such a purified DBM wave function has a simple form:

$$\Psi(\sigma, \sigma') = \prod_j 2 \cosh \left[b_j + \sum_i (W_{ji} \sigma_i + W'_{ji} \sigma'_i) \right]. \quad (5)$$

Then, we can avoid negative signs by simply employing $|\Psi(\sigma, \sigma')|^2$ as the weight for sampling over σ and σ' spins. However, in this case, differently from Method (I), the imaginary-time evolution cannot be followed analytically. Instead, we need to update parameters numerically at each time so that the DBM obeys the imaginary-time evolution starting from the infinite-temperature state $\Psi_\infty(\sigma, \sigma') = \prod_i 2 \cosh \left[\frac{i\pi}{4} (\sigma_i + \sigma'_i) \right]$ (recall that this infinite-temperature DBM does not require d spins). For this purpose, we employ the stochastic reconfiguration (SR) method [49–51] [52]. The SR optimization is designed to minimize the distance between quantum states following the exact and variational




FIG. 3. Finite-temperature calculations for the 2D J_1 - J_2 Heisenberg model ($J_1 = 1$) on the 6×6 square lattice with periodic boundary condition: (a) Energy E , (b) specific heat C , and (c) z component of the spin structure factor $S^z(\mathbf{q}) = \frac{1}{N_{\text{site}}} \sum_{ij} e^{i\mathbf{q} \cdot (\mathbf{R}_i - \mathbf{R}_j)} \langle S_i^z S_j^z \rangle$ at $\mathbf{q} = (\pi, \pi)$. The symbols denote Method (II) results with $8N_{\text{site}}$ hidden spins, which show a good agreement with the numerically exact references (solid curves) obtained by Method (I) with $\delta_\tau = 0.005$ ($J_2 = 0$) and the TPQ method ($J_2 = 0.5$). The TPQ calculations are performed with $\mathcal{H}\Phi$ [53]. The shaded regions show the size of the error bars of the TPQ results [in Method (I), the size of the error bars is small].

imaginary-time evolution, as much as possible within the expressive power of the DBM wave function in Eq. (5). The expressive power is systematically controlled by the number of h spins; It is ensured that any quantum states can be represented exactly (universal approximation) by an infinitely wide network structure [39, 40].

As a demonstration of Method (II), we turn to a highly challenging problem: frustrated spin systems. As a representative, here, we consider the 2D antiferromagnetic J_1 - J_2 Heisenberg model on the $L \times L$ square lattices with periodic boundary condition. The Hamiltonian reads $\mathcal{H} = J_1 \sum_{\langle i,j \rangle} \mathbf{S}_i \cdot \mathbf{S}_j + J_2 \sum_{\langle\langle i,j \rangle\rangle} \mathbf{S}_i \cdot \mathbf{S}_j$. Here, \mathbf{S}_i is the spin-1/2 operator at site i , and $J_1 (= 1)$ and J_2 are the nearest-neighbor and next-nearest-neighbor couplings, respectively. When J_2 is finite, the spin configuration cannot satisfy the energy gain by the J_1 and J_2 interactions simultaneously (frustration). Around $J_2 = 0.5$, where the frustration is strong, an exotic state of matter, quantum spin liquid without any symmetry breaking, might be stabilized as the ground state [21, 54–59]. The model also attracts attention because of its possible relevance to the physics of high- T_c cuprates [60–62]. However, because numerically exact QMC results are not available due to the sign problem, the ground-state phase diagram is still under active debate.

For this challenging problem, the wave functions using neural networks have started to be applied to the zero temperature calculation [59, 63–68]. However, to detect a hallmark of the possible quantum spin liquid phase experimentally, the finite-temperature behavior needs to be elucidated. Here, we apply Method (II) to perform the finite-temperature calculations for $J_2 = 0$ and 0.5. To check its accuracy, we compare the results with numerically exact ones obtained by Method (I) for the non-frustrated case ($J_2 = 0$). For the frustrated case ($J_2 = 0.5$), Method (I) suffers from the sign problem, but, up to a 6×6 lattice, the TPQ results are available, which are also numerically exact. Therefore, we perform calculations using the 6×6 lattice with the total magnetization restricted to be zero ($\sum_i S_i^z = 0$). To further improve the accuracy of the calculation, we utilize the translational and point-group symmetry of the extended system [68]. See SM for the practical details [49].

Figure 3 shows Method (II) results (symbols) for the temperature dependence of the (a) energy, (b) specific heat, and (c) z component of spin structure factor $S^z(\pi, \pi)$, which quantifies the Néel-type antiferromagnetic correlation. We can see that, by the frustration, the antiferromagnetic correlation is largely suppressed, and the entropy release slows down. Method (II) results accurately reproduce the exact imaginary-time evolution, showing its reliability even in the frustrated regime.

In Method (II), by optimizing parameters numerically, we obtain a more compact and dense network to represent finite-temperature states compared to the analytically derived network in Method (I): the number of hidden units is $N_h = 8N_{\text{site}}$ in this case, which is in contrast to $\mathcal{O}(N_\tau N_{\text{site}})$ in Method (I). The compactness of the network without d spins results in the absence of the negative weights in the MC sampling. Considering the observation that the N_h scales polynomially with respect to the system size, our results are strongly encouraging to expect a computational advantage in even more dedicated simulations for larger systems.

From the application to the Heisenberg model, we see that both the implementation of symmetrization and increasing N_h contributes to enhancing the accuracy of calculations [49]. It is an important future task to check the reliability of Method (II) for other Hamiltonians, since the convergence of the error in the imaginary-time evolution with respect to N_h should be model dependent.

Summary and outlook.— In summary, we have proposed two cutting-edge approaches that utilize DBMs to simulate the finite-temperature properties of quantum many-body systems. In the first approach, we provide a deterministic construction of DBMs that exactly represents Gibbs states, which proves the suitability and flexibility of neural networks for encoding thermal properties. In the second approach, the DBM network parameters are optimized stochastically so that the imaginary time evolution can be approximated efficiently, even for one of the most challenging 2D problems, such as frustrated systems.

Several future directions can be envisioned. It is an interesting question how the neural-network quantum states perform

under other schemes of finite-temperature calculation such as TPQ methods. All variational ansätze are, by construction, not powerful enough to express Haar random states, which are taken as the initial states in TPQ calculations. Nonetheless, results by tensor-network-based algorithms [8] imply that using the truncated Hilbert space is sufficient in practical simulations. The major obstacle is considered to be the entanglement growth along the time propagation, which we expect to be simulated well by neural networks, based on previous works on real-time evolution [19, 32, 33]. Also, it is natural to explore the scalability of our methods in larger and/or more complex systems, or ask whether other network structures (e.g. deep feed-forward networks) are suited for the finite-temperature calculations; the trade-off relationship between the representability and trainability of shallow/deep neural networks remains to be open.

ACKNOWLEDGEMENTS

Y.N. is supported by Grant-in-Aids for Scientific Research (JSPS KAKENHI) (Grants No. 16H06345, No. 17K14336, No. 18H01158, and No. 20K14423) and MEXT as “Program for Promoting Researches on the Supercomputer Fugaku” (Basic Science for Emergence and Functionality in Quantum Matter —Innovative Strongly-Correlated Electron Science by Integration of “Fugaku” and Frontier Experiments—) (Project ID: hp200132). N.Y. is supported by the Japan Science and Technology Agency (JST) (via the Q-LEAP program). F.N. is supported in part by: NTT Research, Army Research Office (ARO) (Grant No. W911NF-18-1-0358), Japan Science and Technology Agency (JST) (via the CREST Grant No. JPMJCR1676), JSPS (via the KAKENHI Grant No. JP20H00134 and the JSPS-RFBR Grant No. JPJSBP120194828), the Asian Office of Aerospace Research and Development (AOARD) (via Grant No. FA2386-20-1-4069), and the Foundational Questions Institute Fund (FQXi) via Grant No. FQXi-IAF19-06. Part of numerical calculations were performed using NetKet [69]. Some calculations were performed using the computational resources of supercomputers in RIKEN (Fugaku and HOKUSAI) and the Institute of Solid State Physics at the University of Tokyo.

* These two authors contribute equally to the paper. E-mail at: yusuke.nomura@riken.jp

† These two authors contribute equally to the paper. E-mail at: nyoshioka@ap.t.u-tokyo.ac.jp

- [1] Masuo Suzuki, “Relationship between d-Dimensional Quantal Spin Systems and (d+1)-Dimensional Ising Systems: Equivalence, Critical Exponents and Systematic Approximants of the Partition Function and Spin Correlations,” *Progress of Theoretical Physics* **56**, 1454–1469 (1976).
- [2] Masuo Suzuki, Seiji Miyashita, and Akira Kuroda, “Monte Carlo Simulation of Quantum Spin Systems. I,” *Progress of Theoretical Physics* **58**, 1377–1387 (1977).
- [3] Anders W. Sandvik, “Stochastic series expansion method with operator-loop update,” *Phys. Rev. B* **59**, R14157–R14160 (1999).
- [4] James Gubernatis, Naoki Kawashima, and Philipp Werner, *Quantum Monte Carlo Methods* (Cambridge University Press, 2016).
- [5] F. Verstraete, J. J. García-Ripoll, and J. I. Cirac, “Matrix product density operators: Simulation of finite-temperature and dissipative systems,” *Phys. Rev. Lett.* **93**, 207204 (2004).
- [6] Michael Zwolak and Guifré Vidal, “Mixed-state dynamics in one-dimensional quantum lattice systems: A time-dependent superoperator renormalization algorithm,” *Phys. Rev. Lett.* **93**, 207205 (2004).
- [7] Adrian E. Feiguin and Steven R. White, “Finite-temperature density matrix renormalization using an enlarged Hilbert space,” *Physical Review B* **72**, 220401 (2005).
- [8] Steven R. White, “Minimally entangled typical quantum states at finite temperature,” *Phys. Rev. Lett.* **102**, 190601 (2009).
- [9] Antoine Georges, Gabriel Kotliar, Werner Krauth, and Marcelo J. Rozenberg, “Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions,” *Rev. Mod. Phys.* **68**, 13–125 (1996).
- [10] Nikolai V. Prokof’ev and Boris V. Svistunov, “Polaron Problem by Diagrammatic Quantum Monte Carlo,” *Phys. Rev. Lett.* **81**, 25142517 (1998).
- [11] Kris Van Houcke, Evgeny Kozik, N. Prokof’ev, and B. Svistunov, “Diagrammatic Monte Carlo,” *Physics Procedia* **6**, 95–105 (2010), computer Simulations Studies in Condensed Matter Physics XXI.
- [12] 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**, 205–210 (2020).
- [13] Naoki Irikura and Hiroki Saito, “Neural-network quantum states at finite temperature,” *Phys. Rev. Research* **2**, 013284 (2020).
- [14] Toshiaki Iitaka, “Random phase product state for canonical ensemble,” (2020), [arXiv:2006.14459 \[cond-mat.str-el\]](https://arxiv.org/abs/2006.14459).
- [15] Sho Sugiura and Akira Shimizu, “Thermal pure quantum states at finite temperature,” *Phys. Rev. Lett.* **108**, 240401 (2012).
- [16] Sho Sugiura and Akira Shimizu, “Canonical thermal pure quantum state,” *Phys. Rev. Lett.* **111**, 010401 (2013).
- [17] Kensaku Takai, Kota Ido, Takahiro Misawa, Youhei Yamaji, and Masatoshi Imada, “Finite-Temperature Variational Monte Carlo Method for Strongly Correlated Electron Systems,” *Journal of the Physical Society of Japan* **85**, 034601 (2016).
- [18] Atsushi Iwaki, Akira Shimizu, and Chisa Hotta, “Thermal pure quantum matrix product states recovering a volume law entanglement,” (2020), [arXiv:2005.06829 \[cond-mat.str-el\]](https://arxiv.org/abs/2005.06829).
- [19] Giuseppe Carleo and Matthias Troyer, “Solving the quantum many-body problem with artificial neural networks,” *Science* **355**, 602–606 (2017).
- [20] Giacomo Torlai, Guglielmo Mazzola, Juan Carrasquilla, Matthias Troyer, Roger Melko, and Giuseppe Carleo, “Neural-network quantum state tomography,” *Nature Physics* **14**, 447–450 (2018).
- [21] Ling Wang and Anders W. Sandvik, “Critical Level Crossings and Gapless Spin Liquid in the Square-Lattice Spin-1/2 J_1 – J_2 Heisenberg Antiferromagnet,” *Phys. Rev. Lett.* **121**, 107202 (2018).
- [22] Giuseppe Carleo, Ignacio Cirac, Kyle Cranmer, Laurent Daudet, Maria Schuld, Naftali Tishby, Leslie Vogt-Maranto, and Lenka Zdeborová, “Machine learning and the physical sci-

- ences,” *Rev. Mod. Phys.* **91**, 045002 (2019).
- [23] Roger G Melko, Giuseppe Carleo, Juan Carrasquilla, and J Ignacio Cirac, “Restricted Boltzmann machines in quantum physics,” *Nature Physics* **15**, 887–892 (2019).
- [24] Abhijeet Melkani, Clemens Gneiting, and Franco Nori, “Eigenstate extraction with neural-network tomography,” *Phys. Rev. A* **102**, 022412 (2020).
- [25] Shahnawaz Ahmed, Carlos Snchez Muoz, Franco Nori, and Anton Frisk Kockum, “Quantum state tomography with conditional generative adversarial networks,” (2020), [arXiv:2008.03240 \[quant-ph\]](https://arxiv.org/abs/2008.03240).
- [26] Yusuke Nomura, Andrew S. Darmawan, Youhei Yamaji, and Masatoshi Imada, “Restricted Boltzmann machine learning for solving strongly correlated quantum systems,” *Phys. Rev. B* **96**, 205152 (2017).
- [27] Hiroki Saito, “Solving the Bose-Hubbard Model with Machine Learning,” *J. Phys. Soc. Jpn.* **86**, 093001 (2017).
- [28] Zi Cai and Jinguo Liu, “Approximating quantum many-body wave functions using artificial neural networks,” *Phys. Rev. B* **97**, 035116 (2018).
- [29] Kenny Choo, Giuseppe Carleo, Nicolas Regnault, and Titus Neupert, “Symmetries and many-body excitations with neural-network quantum states,” *Phys. Rev. Lett.* **121**, 167204 (2018).
- [30] Yusuke Nomura, “Machine learning quantum states — extensions to fermion–boson coupled systems and excited-state calculations,” *J. Phys. Soc. Jpn.* **89**, 054706 (2020).
- [31] Nobuyuki Yoshioka, Wataru Mizukami, and Franco Nori, “Solving quasiparticle band spectra of real solids using neural-network quantum states,” *Communications Physics* **4**, 106 (2021).
- [32] Markus Schmitt and Markus Heyl, “Quantum many-body dynamics in two dimensions with artificial neural networks,” *Phys. Rev. Lett.* **125**, 100503 (2020).
- [33] Pinchen Xie and Weinan E, “Coarse-grained spectral projection: A deep learning assisted approach to quantum unitary dynamics,” *Physical Review B* **103**, 024304 (2021).
- [34] Nobuyuki Yoshioka and Ryusuke Hamazaki, “Constructing neural stationary states for open quantum many-body systems,” *Phys. Rev. B* **99**, 214306 (2019).
- [35] Alexandra Nagy and Vincenzo Savona, “Variational Quantum Monte Carlo Method with a Neural-Network Ansatz for Open Quantum Systems,” *Phys. Rev. Lett.* **122**, 250501 (2019).
- [36] Michael J. Hartmann and Giuseppe Carleo, “Neural-Network Approach to Dissipative Quantum Many-Body Dynamics,” *Phys. Rev. Lett.* **122**, 250502 (2019).
- [37] Filippo Vicentini, Alberto Biella, Nicolas Regnault, and Cristiano Ciuti, “Variational Neural-Network Ansatz for Steady States in Open Quantum Systems,” *Phys. Rev. Lett.* **122**, 250503 (2019).
- [38] Dong Yuan, Heran Wang, Zhong Wang, and Dong-Ling Deng, “Solving the Liouvillian Gap with Artificial Neural Networks,” (2020), [arXiv:2009.00019 \[quant-ph\]](https://arxiv.org/abs/2009.00019).
- [39] Nicolas Le Roux and Yoshua Bengio, “Representational Power of Restricted Boltzmann Machines and Deep Belief Networks,” *Neural Computation* **20**, 1631–1649 (2008).
- [40] Guido Montufar and Nihat Ay, “Refinements of universal approximation results for deep belief networks and restricted boltzmann machines,” *Neural Comput.* **23**, 1306–1319 (2011).
- [41] Dong-Ling Deng, Xiaopeng Li, and S. Das Sarma, “Quantum entanglement in neural network states,” *Phys. Rev. X* **7**, 021021 (2017).
- [42] Yoav Levine, Or Sharir, Nadav Cohen, and Amnon Shashua, “Quantum entanglement in deep learning architectures,” *Phys. Rev. Lett.* **122**, 065301 (2019).
- [43] Xun Gao and Lu-Ming Duan, “Efficient representation of quantum many-body states with deep neural networks,” *Nat. Commun.* **8**, 662 (2017).
- [44] Giuseppe Carleo, Yusuke Nomura, and Masatoshi Imada, “Constructing exact representations of quantum many-body systems with deep neural networks,” *Nature Communications* **9**, 5322 (2018).
- [45] Michael A Nielsen and Isaac L Chuang, *Quantum computation and quantum information* (Cambridge University Press, (2000)).
- [46] The renormalization factor is omitted in the expression. We ignore bias terms for visible and deep units because they are irrelevant in the present case.
- [47] Giacomo Torlai and Roger G. Melko, “Latent space purification via neural density operators,” *Phys. Rev. Lett.* **120**, 240503 (2018).
- [48] For spin-1/2 systems, imaginary-time propagators can be analytically represented by the DBM. This has been discussed in the context of classical many-body interacting systems as well [?].
- [49] See Supplementary Materials for more details of the method (URL to be added).
- [50] Sandro Sorella, “Green function monte carlo with stochastic reconfiguration,” *Phys. Rev. Lett.* **80**, 4558–4561 (1998).
- [51] Sandro Sorella, “Generalized Lanczos algorithm for variational quantum Monte Carlo,” *Phys. Rev. B* **64**, 024512 (2001).
- [52] The same optimization scheme has been developed independently in artificial-intelligence community, called “natural gradient” [? ?].
- [53] Mitsuaki Kawamura, Kazuyoshi Yoshimi, Takahiro Misawa, Youhei Yamaji, Syngge Todo, and Naoki Kawashima, “Quantum lattice model solver $\mathcal{H}\Phi$,” *Computer Physics Communications* **217**, 180–192 (2017).
- [54] Hong-Chen Jiang, Hong Yao, and Leon Balents, “Spin liquid ground state of the spin- $\frac{1}{2}$ square J_1 - J_2 Heisenberg model,” *Phys. Rev. B* **86**, 024424 (2012).
- [55] Wen-Jun Hu, Federico Becca, Alberto Parola, and Sandro Sorella, “Direct evidence for a gapless Z_2 spin liquid by frustrating Néel antiferromagnetism,” *Phys. Rev. B* **88**, 060402 (2013).
- [56] Satoshi Morita, Ryui Kaneko, and Masatoshi Imada, “Quantum Spin Liquid in Spin 1/2 J_1 - J_2 Heisenberg Model on Square Lattice: Many-Variable Variational Monte Carlo Study Combined with Quantum-Number Projections,” *J. Phys. Soc. Jpn.* **84**, 024720 (2015).
- [57] Francesco Ferrari and Federico Becca, “Gapless spin liquid and valence-bond solid in the J_1 - J_2 Heisenberg model on the square lattice: Insights from singlet and triplet excitations,” *Phys. Rev. B* **102**, 014417 (2020).
- [58] Wen-Yuan Liu, Shou-Shu Gong, Yu-Bin Li, Didier Poilblanc, Wei-Qiang Chen, and Zheng-Cheng Gu, “Gapless quantum spin liquid and global phase diagram of the spin-1/2 J_1 - J_2 square antiferromagnetic Heisenberg model,” (2020), [arXiv:2009.01821 \[cond-mat.str-el\]](https://arxiv.org/abs/2009.01821).
- [59] Yusuke Nomura and Masatoshi Imada, “Dirac-type nodal spin liquid revealed by machine learning,” (2020), [arXiv:2005.14142 \[cond-mat.str-el\]](https://arxiv.org/abs/2005.14142).
- [60] F Nori and G. T Zimanyi, “Comparison of Frustration and Doping Effects in Generalized Flux Phases,” *Europhysics Letters (EPL)* **16**, 397–403 (1991).
- [61] Franco Nori, Eduardo Gagliano, and Silvia Bacci, “Does frustration describe doping in models for high-temperature superconductors?” *Phys. Rev. Lett.* **68**, 240–243 (1992).

- [62] Franco Nori, R. Merlin, Stephan Haas, Anders W. Sandvik, and Elbio Dagotto, “Magnetic Raman Scattering in Two-Dimensional Spin-1/2 Heisenberg Antiferromagnets: Spectral Shape Anomaly and Magnetostrictive Effects,” *Phys. Rev. Lett.* **75**, 553–556 (1995).
- [63] Xiao Liang, Wen-Yuan Liu, Pei-Ze Lin, Guang-Can Guo, Yong-Sheng Zhang, and Lixin He, “Solving frustrated quantum many-particle models with convolutional neural networks,” *Phys. Rev. B* **98**, 104426 (2018).
- [64] Kenny Choo, Titus Neupert, and Giuseppe Carleo, “Two-dimensional frustrated J_1-J_2 model studied with neural network quantum states,” *Phys. Rev. B* **100**, 125124 (2019).
- [65] Francesco Ferrari, Federico Becca, and Juan Carrasquilla, “Neural Gutzwiller-projected variational wave functions,” *Phys. Rev. B* **100**, 125131 (2019).
- [66] Tom Westerhout, Nikita Astrakhantsev, Konstantin S. Tikhonov, Mikhail I. Katsnelson, and Andrey A. Bagrov, “Generalization properties of neural network approximations to frustrated magnet ground states,” *Nature Communications* **11**, 1593 (2020).
- [67] Attila Szabó and Claudio Castelnovo, “Neural network wave functions and the sign problem,” *Phys. Rev. Research* **2**, 033075 (2020).
- [68] Yusuke Nomura, “Helping restricted boltzmann machines with quantum-state representation by restoring symmetry,” *J. Phys.: Condens. Matter* **33**, 174003 (2021).
- [69] Giuseppe Carleo, Kenny Choo, Damian Hofmann, James E.T. Smith, Tom Westerhout, Fabien Alet, Emily J. Davis, Stavros Efthymiou, Ivan Glasser, Sheng-Hsuan Lin, Marta Mauri, Guglielmo Mazzola, Christian B. Mendl, Evert van Nieuwenburg, Ossian O’Reilly, Hugo Theveniaut, Giacomo Torlai, Filippo Vicentini, and Alexander Wietek, “Netket: A machine learning toolkit for many-body quantum systems,” *SoftwareX* **10**, 100311 (2019).