HubbardNet: Efficient Predictions of the Bose-Hubbard Model Spectrum with Deep Neural Networks

Ziyan Zhu
Stanford Institute of Materials and Energy Science,
SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA
ziyanzhu@stanford.edu

Marios Mattheakis
John A. Paulson School of Engineering and Applied Sciences,
Harvard University, Cambridge, Massachusetts 02138, USA
mariosmat@gmail.com

Weiwei Pan
John A. Paulson School of Engineering and Applied Sciences,
Harvard University, Cambridge, Massachusetts 02138, USA
weiweipan@g.harvard.edu

Efthimios Kaxiras
Department of Physics, Harvard University, Cambridge, Massachusetts 02138, USA
John A. Paulson School of Engineering and Applied Sciences,
Harvard University, Cambridge, Massachusetts 02138, USA
kaxiras@g.harvard.edu

Abstract

We present a deep neural network (DNN)-based model, the `HubbardNet`, to variationally solve for the ground state and excited state wavefunctions of the one-dimensional and two-dimensional Bose-Hubbard model on a square lattice. Using this model, we obtain the Bose-Hubbard energy spectrum as an analytic function of the Coulomb parameter, U, and the total number of particles, N, from a single training, bypassing the need to solve a new hamiltonian for each different input. We show that the DNN-parametrized solutions have excellent agreement with exact diagonalization while outperforming exact diagonalization in terms of computational scaling, suggesting that our model is promising for efficient, accurate computation of exact phase diagrams of many-body lattice hamiltonians.

1 Introduction

Many-body lattice models such as Hubbard-like models have produced a wealth of results for understanding the phase diagrams and excitation spectra of several exotic condensed-matter systems, including high-temperature superconductors [1, 2, 3], moiré twisted multilayered van der Waals heterostructures [4, 5, 6, 7, 8], and the superfluidity to insulator transition in bosonic gases [9, 10, 11]. Solving these lattice models remains a challenge due to size of the basis needed to fully represent the many-body wavefunction: as the system size increases, the number of possible states in which the

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

system can exist rises exponentially. Typically, a full solution relies on the exact diagonalization of the hamiltonian matrix which yields all the eigenvalues and eigenstates of the many-body system.

Main Contributions: We propose a deep learning approach (referred to as the `HubbardNet`) that minimizes the expectation value of the Bose-Hubbard hamiltonian with respect to the deep neural network (DNN)-parametrized solutions. We demonstrate the efficacy of the model by obtaining the ground state and excited states of the one-dimensional (1D) and two-dimensional (2D) square lattices. Our DNN-parametrized solutions are analytic functions of the Hubbard U parameter and the total number of particles N, which circumvents the need to perform an optimization for every new hamiltonian. We also show that our model exhibits better performance than exact diagonalization for large system size without sacrificing accuracy.

Related Work: Methods based on machine learning (ML) with variational freedom to minimize the energy have proven useful for solving the ground state of interacting spin systems [12, 13], as well as bosonic systems [14, 15, 16] and fermionic systems [17, 18, 19]. In principle, the amount of encoded information in these ML-based models scales linearly with the number of particles and sites. However, the training of a DNN is known to be computationally costly and in many cases it can even far exceed the cost of exact diagonalization. Moreover, only ground states have been solved using this approach to date.

2 `HubbardNet`: DNN-parametrized Solution to the Bose-Hubbard model

Bose-Hubbard Model: The Bose-Hubbard hamiltonian is given by

$$\hat{\mathcal{H}} = -t \sum_{\langle ij \rangle} \hat{a}_i \hat{a}_j^\dagger + \frac{U}{2} \sum_i \hat{n}_i (\hat{n}_i - 1), \tag{1}$$

where t is the hopping parameter for particles moving between neighboring sites labeled as i and j, U describes the on-site Coulomb interaction, $\hat{a}_i, \hat{a}_j^\dagger$ are the creation and destruction operators for particles on the lattice sites, and $\hat{n} = \hat{a}_i^\dagger \hat{a}_i$ is the number operator, counting the number of particles at site i. The sum on $\langle ij \rangle$ runs over all nearest-neighbor pairs, while the sum on i runs over all lattice sites. Two other important features of the model are the number of particles N and the total number of states $\mathcal{N}_B$ (the basis), which are given, respectively, by the expressions

$$N = \sum_i n_i, \quad \mathcal{N}_B = \binom{M+N-1}{N-1}. \tag{2}$$

The hamiltonian matrix can be constructed by listing all possible vectors $\boldsymbol{n}$ and taking the inner product $\langle \boldsymbol{n}|\mathcal{H}|\boldsymbol{n}'\rangle$ for every pair of $\boldsymbol{n}$ and $\boldsymbol{n}'$. Each $\boldsymbol{n}$ is associated with an integer label obtained using the Ponomarev ordering [20, 21, 22]. The ground-truth solution can be found by exact diagonalization of the matrix $\langle \Psi_i|\hat{\mathcal{H}}|\Psi_j\rangle$ with eigenvalues E_j and eigenvectors $|\Psi_j\rangle$ which are linear superpositions of the vectors $\boldsymbol{n}$ with coefficients $\psi_j(\boldsymbol{n})$.

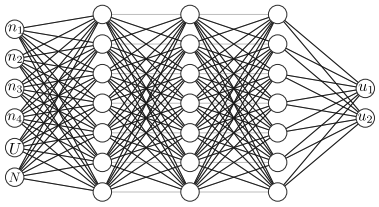

Figure 1: Schematic diagram of the data-free neural network to solve the Bose-Hubbard model.

Architecture: We parametrize the wavefunction Ψ_j by a fully connected DNN $\Psi_j = g(\boldsymbol{W})$ where $\boldsymbol{W}$ denotes the network parameters (Fig. 1). The network takes $(M+2)$ inputs, which are the number of particles n_j at all M sites, U, and N. Note that U and N are input parameters such that the trained network can make efficient predictions at different values of U and N's at inference, unlike all existing techniques which require a new solution for each new parameter. We use a hyperbolic tangent activation function for all hidden layers. The wavefunction is generally complex, and thus we need two outputs, u_1 and u_2, for the real and imaginary parts. For the ground state energy, we choose

the activation function of the output layer to be an exponential, $\psi_j(\boldsymbol{n}) = \exp(u_1 + \mathrm{i}u_2)$ because we observed better convergence. For excited states, we take a linear activation function for the output layer, $\psi_j(\boldsymbol{n}) = u_1 + \mathrm{i}u_2$. Note that although the solution to the 1D and 2D square lattices is purely real, we keep u_2 for generality.

Optimization: From the variational principle, the ground state energy E_{gs} can be obtained by minimizing the expectation value of the hamiltonian,

$$\langle\hat{\mathcal{H}}\rangle = \frac{\sum_{\boldsymbol{n},\boldsymbol{n}'} \psi_j^*(\boldsymbol{n})\langle\boldsymbol{n}|\hat{\mathcal{H}}|\boldsymbol{n}'\rangle\psi_j(\boldsymbol{n}')}{\sum_{\boldsymbol{n}} |\psi_j(\boldsymbol{n})|^2}, \tag{3}$$

with respect to the NN weights $\boldsymbol{W}$. Equation 3 represents the loss function, that is, we minimize the expectation value of the hamiltonian directly by optimizing the weights of the DNN. By using multiple values of U and N as inputs, the loss function can be modified to be the energy averaged over U and N.

An excited state is a local minimum of $\langle\hat{\mathcal{H}}\rangle$. We can obtain excited states iteratively by minimizing Eq. 3 for the j^{th} state Ψ_j with the constraint that Ψ_j is orthogonal to $\Psi_0, ..., \Psi_{j-1}$ because of the orthonormality of states $\langle\Psi_i|\Psi_j\rangle = \delta_{ij}$ for $i, j \leq \mathcal{N}_B$. The constraint is achieved iteratively through the Gram-Schmidt orthogonalization process. Assuming that $\boldsymbol{v}_0, \boldsymbol{v}_2, ..., \boldsymbol{v}_j$ are linearly independent outputs from the neural network, with the subscript being the excited state index, the orthogonal set of wavefunctions can be obtained at each epoch as follows:

$$\Psi_0 = \boldsymbol{v}_0, \quad \Psi_{j\neq 0} = \boldsymbol{v}_j - \sum_{k=1}^{j-1} \mathrm{proj}_{\boldsymbol{u}_k}(\boldsymbol{v}_k), \quad \text{where } \mathrm{proj}_{\boldsymbol{u}} = \frac{\langle\boldsymbol{u}|\boldsymbol{v}\rangle}{\langle\boldsymbol{u}|\boldsymbol{u}\rangle}\boldsymbol{u}. \tag{4}$$

Moreover, it is effective to add a penalty term to the loss function, $\mathcal{L}_P = -\exp\big[-|\bar{E}_s - \bar{E}_0|\big]$, to encourage the average excited state energy, $\bar{E}_s$, to be close to the average ground state energy, $\bar{E}_0$.

3 Results

Unless otherwise noted, we choose the network width to consist of 4 hidden layers with $D = 400$ nodes per layer. We use stochastic gradient descent with momentum $= 0.9$, and a cosine annealing scheme for varying the learning rate, lr, starting from $lr = 0.01$, and reset every 1000 epochs. We train the network until the variance of the loss function during the last 200 steps is less than 1×10^{-7} (Fig. 2(a)). Additionally, for the excited states, we apply an L2 regularization with a coefficient of 10^{-4} to the network weights in order to discourage energies from being drastically different for similar values of the parameter U. To overcome the limitation of a small training set, we randomly perturb the value of the input values of U during every epoch by an amplitude less than 0.01. The code is written in PyTorch [23] and is publically available [1]. we performed our calculations using a single Nvidia K80 GPU.

Ground State: We first solve for the ground state of the Bose-Hubbard model. For example, we train the network with three different values of U for a 4×3 2D system ($M = 12$) and $N = 3$ with an open boundary condition, resulting in $\mathcal{N}_B = 1365$. Figure 2(b) shows that `HubbardNet` is capable of predicting the energy as a function U exactly. Figure 2(c) shows the average occupation number, $\langle\hat{n}_i\rangle$, for a given site i, for $U = 5.5$, which is defined as $\langle\hat{n}_i\rangle = \sum_{\boldsymbol{n}} n_i|\psi_i(\boldsymbol{n})|^2$. Figure 2(d)-(e) compares the wavefunction as obtained from `HubbardNet` and from the exact diagonalization, both in-distribution ($U = 2.0$) and out-of-distribution ($U = 4.5$) agree well with the ground truth. Using as few as 3 values of U for training, `HubbardNet` is capable of obtaining the ground state wavefunctions and energies of the Bose-Hubbard model for a wide range of U values accurately, as Figure 2(b) demonstrates.

Excited States: Figure 3(a) shows the energy spectra for 4 different values of U from one training at inference for $M = 5, N = 3$ with an open boundary condition, and `HubbardNet` finds the correct energy for all excited states. The single training allows us to obtain the full spectrum of the Bose-Hubbard model for an arbitrary U. Figure 3(b) shows the spectrum produced at inference. The spectrum agrees well with the exact diagonalization, except for a small deviation at $U \lesssim 2$, most likely because they are out-of-distribution.

[1] https://drive.google.com/drive/folders/19PKPxQvaqHyVl0YRME7invmczSMDDcIK?usp=sharing

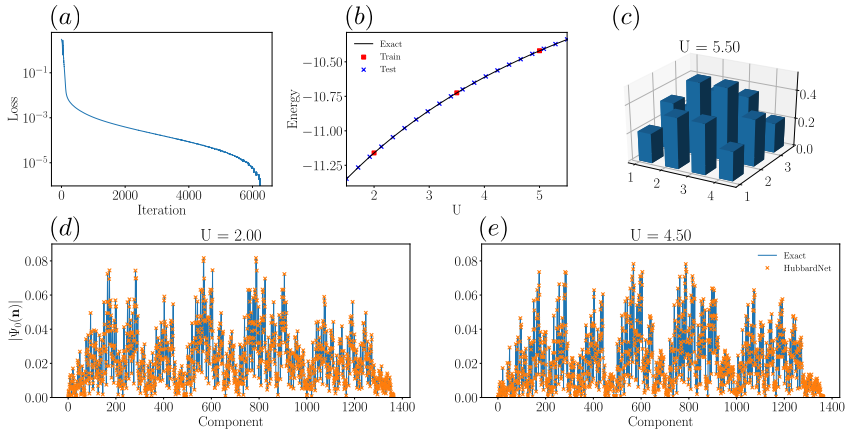

Figure 2: Solution to the Bose-Hubbard model with an open boundary condition and $M = 12, N = 4$. (a) Loss function trace offset by the minimum total energy. (b) Ground state energies from exact diagonalization (blue dashed line) and `HubbardNet` for the training set (red crosses) and prediction (blue scattered points). (c) Occupation number $\langle \boldsymbol{n} \rangle$ for $U = 5.5$ from `HubbardNet`. (d)-(f) Wavefunction magnitude $|\psi_0(\boldsymbol{n})|$ labeled using Ponomarev ordering from the exact diagonalization (blue lines) and `HubbardNet` (orange crosses) for (d) $U = 2.0$, (e) $U = 4.5$.

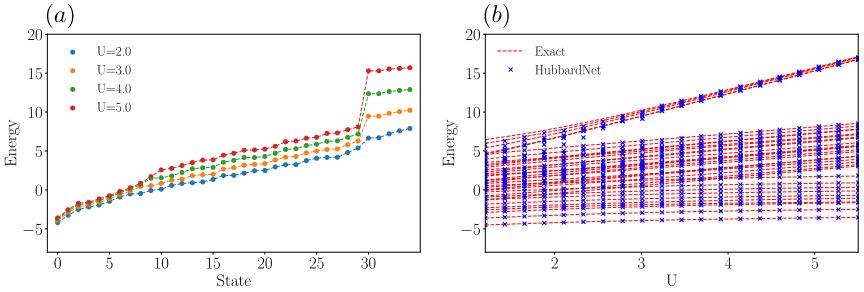

Figure 3: Energy spectrum for $M = 5, N = 3$ and an open boundary condition. (a) Energy from the `HubbardNet` (scattered points) and exact diagonalization (dashed lines) for the training set which includes 4 values of U (color coded); the horizontal axis is the state index with 0 corresponding to the ground state. (b) Energy spectrum for different values of U in the testing set (blue crosses). Eigenvalues from exact diagonalization are shown in red dashed lines.

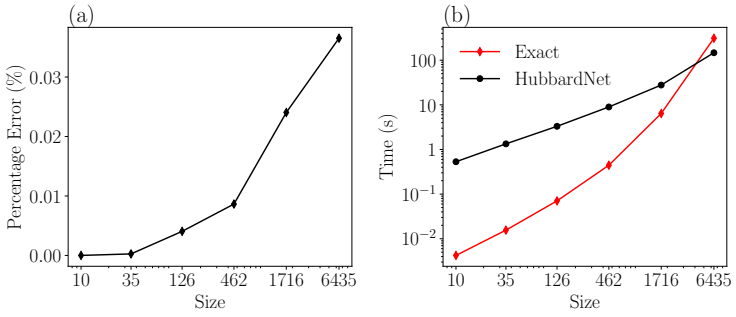

Figure 4: Performance comparison between exact diagonalization and `HubbardNet` using a 1D chain (for $M = N = 3, \ldots, 8$) and periodic boundary conditions, as a function of the system basis-size $\mathcal{N}_B$, see Eq. (2). (a) Percentage error of the ground state energies obtained with `HubbardNet` versus exact diagonalization (ground truth). (b) Computation time for `HubbardNet` (black) and exact diagonalization (red).

Performance Evaluation: We compare the performance of `HubbardNet` with exact diagonalizationin in Figure 4, on a 2.7 GHz Quad-Core Intel Core i7. We take the network width $D = 200$. We use $lr = 0.01$ and converge the loss function to 1×10^{-6}. Figure 4(a) shows that the percentage error of the ground state energies between exact diagonalization and `HubbardNet` is less than 0.05 % for all system sizes, showing that `HubbardNet` is capable of predicting the energy accurately. Figure 4(b) shows that `HubbardNet` has a better computation time scaling than exact diagonalization. While exact diagnolization is faster for a smaller system, `HubbardNet` outperforms exact diagonalization, in terms of computational speed, when $\mathcal{N}_B \geq 6435$ or $M \geq 8$. We note, however, that while the error remains small, the error grows as the system size increases. We also expect the errors to grow faster for excited states due to the error accumulation of the projection (Fig. 3).

4 Conclusions

In this work, we present a DNN-parametrized solution to the ground state and excited states of the Bose-Hubbard model. We not only successfully find the exact solution but also show a better scaling than exact diagonalization for a single state. The DNN solution is an analytic function of the Coulomb parameter, U, allowing for the efficient calculation of the full spectrum from a single training. While we only present the solution as a function of U in this manuscript, we have successfully used `HubbardNet` to obtain energy and wavefunctions as a function of both N and U Therefore, `HubbardNet` is an excellent candidate for mapping the exact phase boundary as a function of the Coulomb interaction strength and the filling factor, especially for a large system size.

Broader Impact

Many-body lattice models like the Hubbard model are of intense theoretical and experimental interest in condensed matter physics. They are commonly employed to study exotic correlated states and phase transitions. The neural network-based approach to solve the Hubbard model that we propose here has several advantages: not only it outperforms exact diagonalization for a large system but also provides an exact analytic dependence of the solution on the Coulomb interaction parameter U and the system size N, potentially allowing for the accurate prediction of the exact phase diagram with a single training. To further improve the computational efficiency, a Monte Carlo sampler can be used to stochastically evaluate the expectation value in Eq. 1 [24]. In the long term, `HubbardNet` can be used to obtain solutions to more realistic hamiltonians including fermionic systems, hamiltonians with long-range Coulomb interactions, and systems in the presence of an external potential and a nonzero chemical potential. The improved efficiency and the exact nature of the DNN solution could be of great benefit to the condensed matter physics community and related fields: by reducing the computational resources needed to obtain the phase diagram of many-body systems, it enhances our ability to understand exotic phases of matter.

Acknowledgments and Disclosure of Funding

Z.Z. and E.K. acknowledge the STC Center for Integrated Quantum Materials, NSF Grant No. DMR-1922172, ARO MURI Grant No. W911NF-2120147, and NSF DMREF Grant No. 1922165, and Simons Foundation, Award No. 896626.

References

[1] PW Anderson. The theory of superconductivity in the high t. *Tc*, 1998.

[2] Edwin W. Huang, Christian B. Mendl, Shenxiu Liu, Steve Johnston, Hong-Chen Jiang, Brian Moritz, and Thomas P. Devereaux. Numerical evidence of fluctuating stripes in the normal state of high-T_c cuprate superconductors. *Science*, 358(6367):1161–1164, December 2017.

[3] Kevin S Bedell, D Coffey, DE Meltzer, D Pines, and JR Schrieffer. High temperature superconductivity: Proceedings. Technical report, Los Alamos National Lab., NM (USA), 1990.

[4] Fengcheng Wu, Timothy Lovorn, Emanuel Tutuc, and A. H. MacDonald. Hubbard model physics in transition metal dichalcogenide moiré bands. *Phys. Rev. Lett.*, 121:026402, Jul 2018.

[5] Haining Pan, Fengcheng Wu, and Sankar Das Sarma. Band topology, hubbard model, heisenberg model, and dzyaloshinskii-moriya interaction in twisted bilayer wse_2. *Phys. Rev. Research*, 2:033087, Jul 2020.

[6] Masayuki Ochi, Mikito Koshino, and Kazuhiko Kuroki. Possible correlated insulating states in magic-angle twisted bilayer graphene under strongly competing interactions. *Phys. Rev. B*, 98:081102, Aug 2018.

[7] J. F. Dodaro, S. A. Kivelson, Y. Schattner, X. Q. Sun, and C. Wang. Phases of a phenomenological model of twisted bilayer graphene. *Phys. Rev. B*, 98(7):075154, August 2018.

[8] Pawel Potasz, Ming Xie, and A. H. MacDonald. Exact Diagonalization for Magic-Angle Twisted Bilayer Graphene. *Phys. Rev. Lett.*, 127(14):147203, October 2021.

[9] D. J. Scalapino, S. R. White, and S. C. Zhang. Superfluid density and the drude weight of the hubbard model. *Phys. Rev. Lett.*, 68:2830–2833, May 1992.

[10] J. K Freericks and H Monien. Phase diagram of the bose-hubbard model. *Europhysics Letters (EPL)*, 26(7):545–550, jun 1994.

[11] T. D. Kühner and H. Monien. Phases of the one-dimensional bose-hubbard model. *Phys. Rev. B*, 58:R14741–R14744, Dec 1998.

[12] Giuseppe Carleo and Matthias Troyer. Solving the quantum many-body problem with artificial neural networks. *Science*, 355(6325):602–606, February 2017.

[13] Sergey Bravyi, David Gosset, Robert König, and Kristan Temme. Approximation algorithms for quantum many-body problems. *Journal of Mathematical Physics*, 60(3):032203, March 2019.

[14] Hiroki Saito. Solving the Bose-Hubbard Model with Machine Learning. *Journal of the Physical Society of Japan*, 86(9):093001, September 2017.

[15] Hiroki Saito and Masaya Kato. Machine Learning Technique to Find Quantum Many-Body Ground States of Bosons on a Lattice. *Journal of the Physical Society of Japan*, 87(1):014001, January 2018.

[16] Kristopher McBrian, Giuseppe Carleo, and Ehsan Khatami. Ground state phase diagram of the one-dimensional Bose-Hubbard model from restricted Boltzmann machines. In *Journal of Physics Conference Series*, volume 1290 of *Journal of Physics Conference Series*, page 012005, October 2019.

[17] Javier Robledo Moreno, Giuseppe Carleo, Antoine Georges, and James Stokes. Fermionic Wave Functions from Neural-Network Constrained Hidden States. *arXiv e-prints*, page arXiv:2111.10420, November 2021.

[18] Nobuyuki Yoshioka, Wataru Mizukami, and Franco Nori. Solving quasiparticle band spectra of real solids using neural-network quantum states. *Communications Physics*, 4(1):106, December 2021.

[19] Elizabeth R. Bennewitz, Florian Hopfmueller, Bohdan Kulchytskyy, Juan Carrasquilla, and Pooya Ronagh. Neural Error Mitigation of Near-Term Quantum Simulations. *arXiv e-prints*, page arXiv:2105.08086, May 2021.

[20] A. V. Ponomarev, S. Denisov, and P. Hänggi. ac-Driven Atomic Quantum Motor. *Phys. Rev. Lett.*, 102(23):230601, June 2009.

[21] A. V. Ponomarev, S. Denisov, and P. Hänggi. Lévy distribution in many-particle quantum systems. *Phys. Rev. A*, 81(4):043615, April 2010.

[22] David Raventós, Tobias Graß, Maciej Lewenstein, and Bruno Juliá-Díaz. Cold bosons in optical lattices: a tutorial for exact diagonalization. *Journal of Physics B Atomic Molecular Physics*, 50(11):113001, June 2017.

[23] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary Devito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NeurIPS*, 2017.

[24] 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 Théveniaut, Giacomo Torlai, Filippo Vicentini, and Alexander Wietek. Netket: A machine learning toolkit for many-body quantum systems. *SoftwareX*, page 100311, 2019.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See the excited states in Sec. 3
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [N/A] Our results do not require any proof but we provided references to justify our design choices.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]