
Denoising Autoencoders for High-Qubit Quantum Dynamics Simulations on Quantum Computers

Connor Powers

Collaboratory for Advanced Computing
and Simulations
University of Southern California
Los Angeles, CA 90089
cdpowers@usc.edu

Lindsay Bassman

Lawrence Berkeley National Laboratory
Berkeley, CA 94720
lbassman@lbl.gov

Yifan Geng

Collaboratory for Advanced Computing
and Simulations
University of Southern California
Los Angeles, CA 90089
yifangen@usc.edu

Rajiv Kalia

Collaboratory for Advanced Computing
and Simulations
University of Southern California
Los Angeles, CA 90089
rkalia@usc.edu

T. K. Satish Kumar

University of Southern California
Los Angeles, CA 90089
tkskwork@gmail.com

Thomas Linker

Collaboratory for Advanced Computing
and Simulations
University of Southern California
Los Angeles, CA 90089
tlinker@usc.edu

Kuang Liu

Collaboratory for Advanced Computing
and Simulations
University of Southern California
Los Angeles, CA 90089
liukuang@usc.edu

Aiichiro Nakano

Collaboratory for Advanced Computing
and Simulations
University of Southern California
Los Angeles, CA 90089
anakano@usc.edu

Pankaj Rajak

Argonne National Laboratory
Argonne, IL 60439
prajak@anl.gov

Priya Vashishta

Collaboratory for Advanced Computing
and Simulations
University of Southern California
Los Angeles, CA 90089
priyav@usc.edu

Abstract

Quantum computers hold great promise for discovering new materials behavior by performing dynamic simulations of quantum materials that are intractable on classical computers. However, high-fidelity results from these simulations are currently hindered by high levels of device noise. Here, we present an autoencoder, trained with quantum simulations of small systems, that is capable of filtering noise from dynamic simulations of larger systems run on quantum computers. We thus

show a lightweight and feasible route for higher-fidelity quantum simulations of large systems beyond points of classical intractability.

1 Introduction

Quantum computers of the near-future are expected to be able to perform dynamic simulations of quantum materials at system sizes and levels of accuracy not achievable on the most advanced classical computers [1, 2, 3]. Using quantum bits, or qubits, for information processing, quantum computers can take advantage of purely quantum effects like entanglement and superposition to overcome the exponential scaling of time and memory resources that simulations of quantum materials face on classical computers. In particular, dynamic simulations of quantum materials, which map very naturally onto quantum computers, promise to unlock new discoveries in myriad areas of science including condensed matter physics, materials science, and quantum chemistry. Unfortunately, currently available quantum computers, widely known as Noisy Intermediate-Scale Quantum (NISQ) computers, suffer from low fidelity due to qubit decoherence, logic gate error rates, and read-out noise [4]. As a result, simulations on NISQ computers require either very short quantum circuits (i.e. the schedule of quantum logic gates performed on the qubits), or error mitigation techniques to reduce noisy output. While reduction of circuit depths, using various circuit synthesis and compilation protocols [5, 6, 7, 8, 9], is an active field of research aiming to minimize error prior to running simulations, less well-studied are approaches for filtering out noise post-simulation.

Already proven to be effective on image [10] and audio [11] de-noising tasks, machine learning (ML) is thus a promising tool for filtering noise from simulation results from NISQ computers. While ML methods based on neural networks have been used to facilitate quantum state tomography on NISQ devices [12, 13, 14, 15], applying neural networks to noise-reduction of dynamic materials simulations on NISQ computers remains a less explored area. A major difficulty for this task lies in the collection of training data, specifically noise-free simulation results. Such data can be computed on classical computers simulating a noiseless quantum computer, however, the complexity grows exponentially with simulation system size [16], capping systems sizes at tens of particles, even on the largest supercomputers [17]. As the materials systems we wish to simulate on quantum computers are precisely those which cannot be simulated on classical computer, a ML model that can filter noise from dynamic simulations of a large system trained by those of smaller systems is highly desirable. Here, we present an autoencoder that can successfully filter NISQ-device noise from dynamic simulations of systems with larger numbers of particles than the simulations with which the model was trained. The ability to filter noise from dynamic simulations performed on NISQ computers with particle counts too large for processing on classical computers paves the way towards new discoveries in the complex behaviour of quantum materials.

2 Methods

2.1 Quantum simulations of transverse field Ising models

The time-dependent Transverse Field Ising Model (TFIM) is a paradigmatic model for understanding quantum many-body magnetism in materials [18]. It is defined by quantum spins on a lattice with exchange-interaction coupling between nearest-neighbors, in the presence of a transverse magnetic field. Figure 1a depicts a schematic of spin degrees of freedom in a material interacting with an external magnetic field. These spin degrees of freedom are mapped to the qubits of a quantum computer, which can simulate their evolution through time, depicted in Figure 1b.

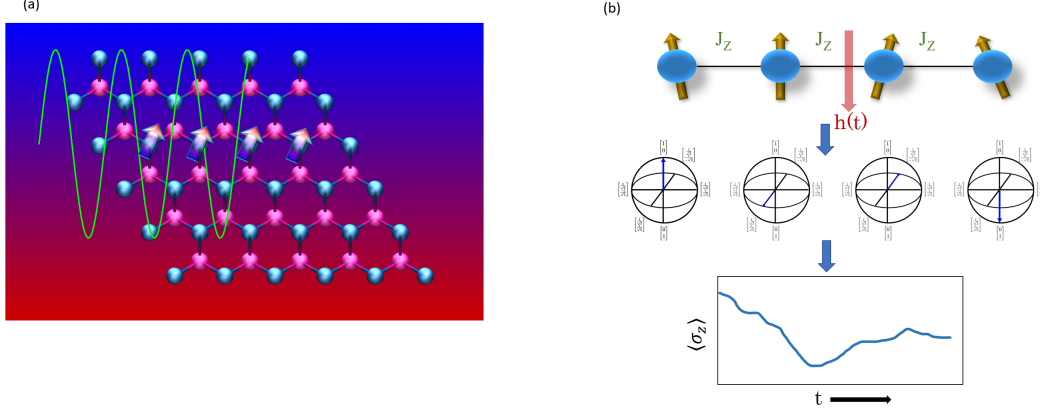


Figure 1: (a) Schematic of the spin degrees of freedom in a material (represented by arrows) interacting with an external magnetic field (represented by the sinusoidal curve). (b) A spin chain governed by Equation 1 is mapped to the qubits of a quantum computer, which can then simulate evolution of the spin chain through time.

The dynamics of the TFIM are governed by a Hamiltonian that can be written as:

$$H(t) = -J_z \sum_{i=1}^{N-1} \sigma_i^z \sigma_{i+1}^z - h_x(t) \sum_{i=1}^N \sigma_i^x \quad (1)$$

Here, the strength of nearest-neighbor interactions are determined by J_z , and the amplitude of the externally applied magnetic field is determined by the time-dependent factor $h_x(t)$. For the models utilized in this paper, $h_x(t) = J_z \cos(\omega_{ph}t)$, where ω_{ph} represents the frequency of an excited phonon in the material. σ_i^z and σ_i^x represent the Z-Pauli and X-Pauli matrices acting on qubit i , respectively.

2.2 Networks

We apply fully-connected autoencoders to the task of reconstructing noiseless results from quantum simulation data from noisy quantum devices. Autoencoders minimize the cost function: $L(\mathbf{x}, g(f(\mathbf{x}|\Theta_1)|\Theta_2))$ where \mathbf{x} is the input example, Θ is the learnable weights of the model, $f(\cdot)$ is the encoder that learns the hidden vectors of the input, and $g(\cdot)$ is the decoder that learns to map the vector back to the values of the training examples. A denoising autoencoder, shown in Figure 2 with fully connected layers, minimizes a modified cost function: $L(\mathbf{x}, g(f(\tilde{\mathbf{x}}|\theta_1)|\theta_2)) = \frac{1}{N} \sum_{i=0}^N |x_i - g(f(\tilde{x}_i|\theta_1)|\theta_2)|$ where $\tilde{\mathbf{x}}$ is a copy of \mathbf{x} that has been corrupted by noise. Training a denoising autoencoder forces the encoder and decoder to learn the structure of provided input examples, subsequently learning the noise. The input feature length, shown as K in Figure 2, is fixed to the number of timesteps in the quantum simulation, and the hidden vector length, shown as J in Figure 2, is a hyperparameter of the model.




Figure 2: General architecture of neural networks used to filter quantum simulation spin data.

2.3 Data generation and training

For each quantum simulation size, or number of spins being simulated, Qiskit’s quantum simulator was used to generate the noiseless ground truth data required for supervised learning. The same quantum simulator was then configured to simulate the noise expected in IBM’s 16-qubit "Melbourne" quantum computer, and then used to generate 100 noisy simulations for each size of simulated spin chain. In practice, we expect this noisy input data to be computed on real quantum hardware. Data from the three largest systems was set aside for testing, and the model was trained on the remaining data for 500 epochs with a 20% validation split via the ADAM optimizer with a learning rate of 0.001. Each constructed network is trained and tested on data for a single spin index within all of the simulated spin chains (e.g. the fifth spin), subsequently "specializing" in a specific spin location. Here, L was chosen to be 9, based on constraints in generating testing data for high qubit simulations.

2.4 Model Evaluation and Selection

Many different autoencoders were found to successfully perform the extrapolation task explored in this study. Subsequently, minimal successful networks, defined as the networks with the least trainable parameters that successfully perform the extrapolation task to within a set evaluation error bound, were found via a separate evaluation metric defined as Mean Maximum Error (MME) over a test dataset consisting of M noisy simulations of length K : $L(y, \hat{y}) = \frac{1}{M} \sum_{i=0}^M \max_{j \in K} |y_j^i - \hat{y}_j^i|$

This error metric was chosen to identify the most useful models because it penalizes large deviations from the truth at any point in the reconstruction, which in practice would entail a model suddenly diverging from the underlying truth.

3 Results

To study the performance of autoencoders in the extrapolation task, the models were trained on the dynamic evolution of the fifth spin (σ_z^5) in the spin chain from noisy simulations of 5- to 9-spin systems. Setting $L = 9$, the model was tasked with filtering noise from the dynamic evolution of σ_z^5 in quantum simulations of $L + 1$, $L + 2$, and $L + 3$ systems, with results shown in Figure 3a-c, respectively. Absolute model errors as a function of simulation time-step for the $L + 1$, $L + 2$, and $L + 3$ extrapolation tasks are presented in Figure 3d. As shown, extrapolating to simulations of up to three more qubits does not lead to a significant decrease in autoencoder performance. This indicates that extending this noise filtration technique to NISQ simulations with particles counts far larger than what is tractable on classical computers may lead to new discoveries in materials dynamics by allowing noise to be removed from simulations that were heretofore impossible to perform.

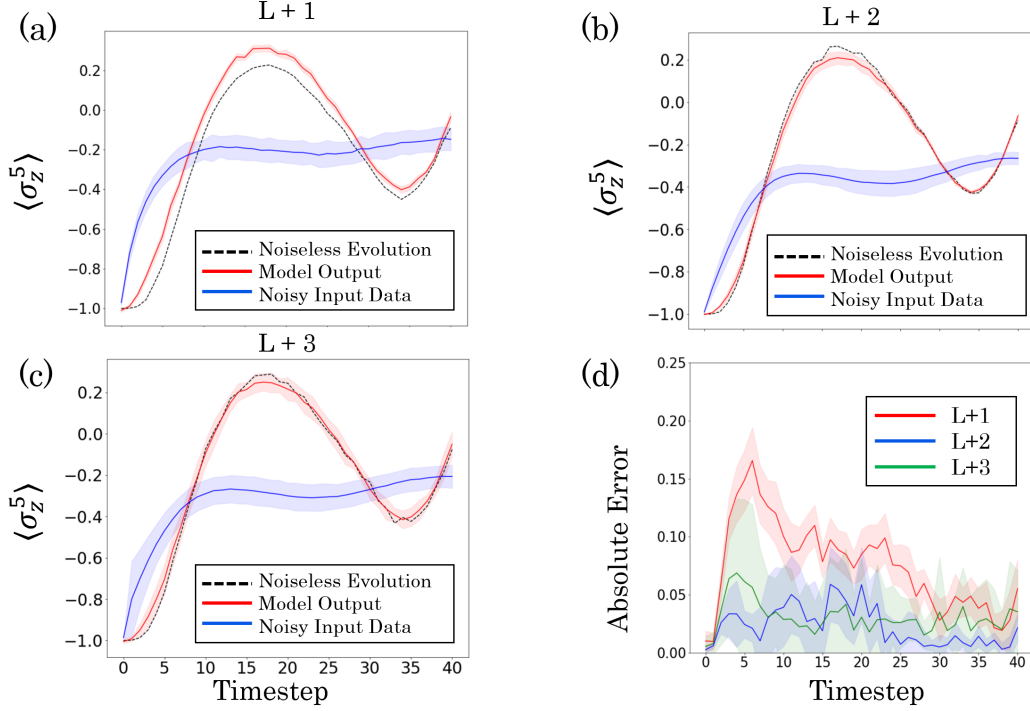


Figure 3: Fifth qubit noise filtration performance of an autoencoder with a single hidden layer, $J = 10$, and 20% dropout in the extrapolation task. (a) Model performance on data from a noisy $L + 1$ qubit simulation (b) Model performance on data from a noisy $L + 2$ qubit simulation performance (c) Model performance on data from a noisy $L + 3$ qubit simulation (d) Absolute model error as a function of simulation timestep

It is noted that the models presented in this paper require no pre-selection or manual filtering of noisy data from lower-qubit simulations, and no such pre-processing was done for the results presented. After performing a minimal successful network search as described in Section 2.4, it was found that a network with two hidden layers consisting of a single node could achieve a MME of 0.1 with 1000 epochs of training, and cutting the target MME to 0.05 only resulted in adding a single node and 20% dropout to each hidden layer.

4 Conclusion

We have demonstrated the ability of denoising autoencoders to successfully perform the extrapolation task of filtering noise from quantum simulations involving higher spin counts than those trained upon. Specifically, we found that autoencoders trained on 5- to 9-spin simulations were able to successfully filter noise from 10-, 11-, and 12-spin simulations without significant degradation in performance with increasing system size.

Broader Impact

The ability of denoising autoencoders to extrapolate to higher-qubit simulations provides an accessible path towards noise reduction in high-qubit quantum simulations on NISQ-era devices beyond the limits of classical intractability. With the recently announced quantum technology roadmap from IBM to realizing 1000-qubit quantum computers by 2023, denoising techniques for high-qubit simulations will be increasingly crucial to unlocking the full potential of these new devices. Filtering out noise in quantum simulations of systems too large for classical computers may herald new discoveries of complex behaviour in quantum materials, with broad applications in materials science, chemistry, and condensed matter physics.

Acknowledgements

This work was supported as part of the Computational Materials Sciences Program funded by the U.S. Department of Energy, Office of Science, Basic Energy Sciences, under Award Number DE-SC0014607. CP was supported by the USC Merit Research Award.

References

- [1] Feynman, R. P. Simulating physics with computers. *Int. J. Theor. Phys* **21**, 467–488 (1982). URL <https://doi.org/10.1007/BF02650179>.
- [2] Lloyd, S. Universal quantum simulators. *Science* **273**, 1073–1078 (1996). URL <https://dx.doi.org/10.1126/science.273.5278.1073>.
- [3] Abrams, D. S. & Lloyd, S. Simulations of many-body fermi systems on a universal quantum computer. *Physical Review Letters* **79**, 2586–2589 (1997). URL <https://dx.doi.org/10.1103/PhysRevLett.79.2586>.
- [4] Preskill, J. Quantum computing in the nisq era and beyond. *Quantum* **2**, 79 (2018). URL <https://dx.doi.org/10.22331/q-2018-08-06-79>.
- [5] Khatri, S. *et al.* Quantum-assisted quantum compiling. *Quantum* **3**, 140 (2019). URL <http://dx.doi.org/10.22331/q-2019-05-13-140>.
- [6] Murali, P., Baker, J. M., Javadi-Abhari, A., Chong, F. T. & Martonosi, M. Noise-adaptive compiler mappings for noisy intermediate-scale quantum computers. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, 1015–1029 (2019). URL <http://dx.doi.org/10.1145/3297858.3304075>.
- [7] Younis, E., Sen, K., Yelick, K. & Iancu, C. Qfast: Quantum synthesis using a hierarchical continuous circuit space. *arXiv preprint arXiv:2003.04462* (2020).
- [8] Cincio, L., Rudinger, K., Sarovar, M. & Coles, P. J. Machine learning of noise-resilient quantum circuits. *arXiv preprint arXiv:2007.01210* (2020).
- [9] Bassman, L. *et al.* Domain-specific compilers for dynamic simulations of quantum materials on quantum computers. *arXiv preprint arXiv:2004.07418* (2020).
- [10] Gondara, L. Medical image denoising using convolutional denoising autoencoders. *2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW)* (2016). URL <http://dx.doi.org/10.1109/ICDMW.2016.0041>.
- [11] Lu, X., Tsao, Y., Matsuda, S. & Hori, C. Speech enhancement based on deep denoising auto-encoder. *Proc. Interspeech* 436–440 (2013).
- [12] Beach, M. J. S. *et al.* Qcucumber: wavefunction reconstruction with neural networks. *SciPost Physics* **7** (2019). URL <http://dx.doi.org/10.21468/SciPostPhys.7.1.009>.
- [13] Cai, Z. & Liu, J. Approximating quantum many-body wave functions using artificial neural networks. *Physical Review B* **97** (2018). URL <http://dx.doi.org/10.1103/PhysRevB.97.035116>.
- [14] Xu, Q. & Xu, S. Neural network state estimation for full quantum state tomography. *arXiv preprint arXiv:1811.06654* (2018).
- [15] Torlai, G. *et al.* Neural-network quantum state tomography. *Nature Physics* **14**, 447–450 (2018). URL <http://dx.doi.org/10.1038/s41567-018-0048-5>.
- [16] Jones, T., Brown, A., Bush, I. & Benjamin, S. C. Quest and high performance simulation of quantum computers. *Scientific Reports* **9** (2019). URL <http://dx.doi.org/10.1038/s41598-019-47174-9>.
- [17] Wu, X.-C. *et al.* Full-state quantum circuit simulation by using data compression. *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (2019). URL <http://dx.doi.org/10.1145/3295500.3356155>.
- [18] Bassman, L. *et al.* Towards simulation of the dynamics of materials on quantum computers. *Physical Review B* **101** (2020). URL <http://dx.doi.org/10.1103/PhysRevB.101.184305>.