Computer Science Capstone: Learning Quantum Spin Liquids with Neural Quantum States

Adam Remels

Abstract

In collaboration with Lucas Brito and Sam Bear, this capstone was completed through the CSCI1470/2470 (Deep Learning) final project. Quantum spin liquids are a theorized state of matter proposed as a platform for quantum computing. We investigated whether neural quantum states—a form of Boltzmann machine designed to model quantum states with reinforcement learning (arXiv:1606.02318) —are capable of learning quantum spin liquid phases. We first investigated this in one dimension (the Haldane-Shastry model), then in two dimensions (Z_2 gauge theory), then incorporated matter (Z_2 gauge theory with matter).

The models are variations on the restricted Boltzmann machine, an energy-based model consisting of a visible layer and a hidden layer. Symmetries of the energy function (the Hamiltonian) such as translation invariance simply the architecture of the Boltzmann machine. Its parameters are trained by minimizing a given Hamiltonian with respect to the hidden layer. The descent is performed with stochastic reconfiguration, a form of Monte Carlo sampling.

We first created a simple, 1D version of the Haldane-Shastry model, to get our feet under us as it was the easiest to directly evaluate for accuracy. From that, we moved to implementing a gauge equivariant neural network (GENN), which allowed us to replicate our 1D results on a chain to a 2D lattice. Then, we incorporated matter and the corresponding real degrees of freedom by combining these two architectures into a new model of our own design, which we named GERBIL (Gauge Equivariant Restricted Boltzmann Machine - It Learns!). The architecture is below in Figure 1.

We ended up getting largely successful results from all three models. The below graphs display our results on the Haldane-Shastry (Figure 2), Z_2 (Figure 3), and Z_2 with matter (Figure 4) models, respectively. We found that our models were largely consistent with our calculated values. Our final GERBIL model did result in a final stable energy prediction that was consistently about 1J higher than the true value, but this was only a 5% relative difference and could likely be finetuned out.

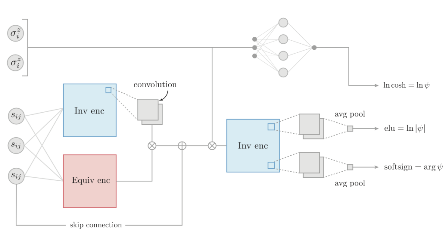

Figure 1

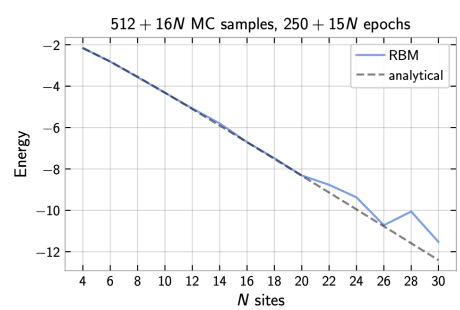

Figure 2

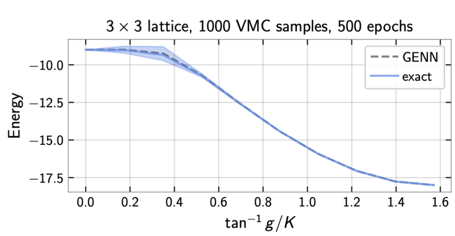

Figure 3

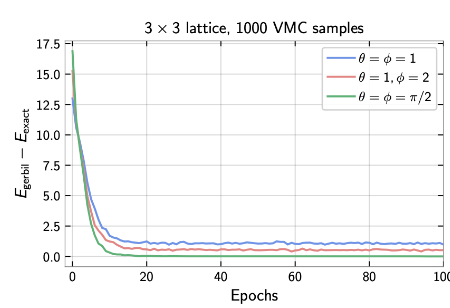

Figure 4