

QUANTUM GRAPH NEURAL NETWORKS

Anonymous authors

Paper under double-blind review

ABSTRACT

We introduce Quantum Graph Neural Networks (QGNN), a new class of quantum neural network ansatzes which are tailored to represent quantum processes which have a graph structure, and are particularly suitable to be executed on distributed quantum systems over a quantum network. Along with this general class of ansatzes, we introduce further specialized architectures, namely, Quantum Graph Recurrent Neural Networks (QGRNN) and Quantum Graph Convolutional Neural Networks (QGCNN). We provide four example applications of QGNN's: learning Hamiltonian dynamics of quantum systems, learning how to create multipartite entanglement in a quantum network, unsupervised learning for spectral clustering, and supervised learning for graph isomorphism classification.

1 INTRODUCTION

Variational Quantum Algorithms are a promising class of algorithms rapidly emerging as a central subfield of Quantum Computing (McClean et al., 2016; Farhi et al., 2014; Farhi & Neven, 2018). Similar to parameterized transformations encountered in deep learning, these parameterized quantum circuits are often referred to as Quantum Neural Networks (QNNs). Recently, it was shown that QNNs that have no prior on their structure suffer from a quantum version of the no-free lunch theorem (McClean et al., 2018) and are exponentially difficult to train via gradient descent. Thus, there is a need for better QNN ansatzes. One popular class of QNNs has been Trotter-based ansatzes (Farhi et al., 2014; Hadfield et al., 2019). The optimization of these ansatzes has been extensively studied in recent works, and efficient optimization methods have been found (Verdon et al., 2019b; Li et al., 2019). On the classical side, graph-based neural networks leveraging data geometry have seen some recent successes in deep learning, finding applications in biophysics and chemistry (Kearnes et al., 2016). Inspired from this success, we propose a new class of Quantum Neural Network ansatz which allows for both quantum inference and classical probabilistic inference for data with a graph-geometric structure. In the sections below, we introduce the general framework of the QGNN ansatz as well as several more specialized variants and showcase four potential applications via numerical implementation.

2 BACKGROUND

2.1 CLASSICAL GRAPH NEURAL NETWORKS

Graph Neural Networks (GNNs) date back to Sperduti & Starita (1997) who applied neural networks to acyclic graphs. Gori et al. (2005) and Scarselli et al. (2008) developed methods that learned node representations by propagating the information of neighbouring nodes. Recently, GNNs have seen great breakthroughs by adapting the convolution operator from CNNs to graphs (Bruna et al., 2013; Henaff et al., 2015; Defferrard et al., 2016; Kipf & Welling, 2016; Niepert et al., 2016; Hamilton et al., 2017; Monti et al., 2017). Many of these methods can be expressed under the message-passing framework (Gilmer et al., 2017).

Let graph $\mathcal{G} = (\mathbf{A}, \mathbf{X})$ where $\mathbf{A} \in \mathbb{R}^{n \times n}$ is the adjacency matrix, and $\mathbf{X} \in \mathbb{R}^{n \times d}$ is the node feature matrix where each node has d features.

$$\mathbf{H}^{(k)} = P(\mathbf{A}, \mathbf{H}^{(k-1)}, \mathbf{W}^{(k)}) \quad (1)$$

where $\mathbf{H}^{(k)} \in \mathbb{R}^{n \times d}$ are the node representations computed at layer k , P is the message propagation function and is dependent on the adjacency matrix, the previous node encodings and some learnable parameters $\mathbf{W}^{(k)}$. The initial embedding, $\mathbf{H}^{(0)}$ is naturally \mathbf{X} . One popular implementation of this framework is the GCN (Kipf & Welling, 2016) which implements it as follows:

$$\mathbf{H}^{(k)} = P(\mathbf{A}, \mathbf{H}^{(k-1)}, \mathbf{W}^{(k)}) = \text{ReLU}(\tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}} \mathbf{H}^{(k-1)} \mathbf{W}^{(k-1)}) \quad (2)$$

where $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}$ is the adjacency matrix with inserted self-loops, $\tilde{\mathbf{D}} = \sum_j \tilde{\mathbf{A}}_{ij}$ is the re-normalization factor (degree matrix).

2.2 NETWORKED QUANTUM SYSTEMS

Consider a graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, where \mathcal{V} is the set of vertices (or nodes) and \mathcal{E} the set of edges. We can assign a quantum subsystem with Hilbert space \mathcal{H}_v for each vertex in the graph, forming a global Hilbert space $\mathcal{H}_{\mathcal{V}} \equiv \bigotimes_{v \in \mathcal{V}} \mathcal{H}_v$. Each of the vertex subsystems could be one or several qubits, a qudit, a qumode (Weedbrook et al., 2012), or even an entire quantum computer. One may also define a Hilbert space for each edge and form $\mathcal{H}_{\mathcal{E}} \equiv \bigotimes_{e \in \mathcal{E}} \mathcal{H}_e$. The total Hilbert space for the graph would then be $\mathcal{H}_{\mathcal{E}} \otimes \mathcal{H}_{\mathcal{V}}$. For the sake of simplicity and feasibility of numerical implementation, we consider this to be beyond the scope of the present work. The edges of the graph dictate the communication between the vertex subspaces: couplings between degrees of freedom on two different vertices are allowed if there is an edge connecting them. This setup is called a quantum network (Kimble, 2008; Qian et al., 2019) with topology given by the graph \mathcal{G} .

3 QUANTUM GRAPH NEURAL NETWORKS

3.1 GENERAL QUANTUM GRAPH NEURAL NETWORK ANSATZ

The most general Quantum Graph Neural Network ansatz is a parameterized quantum circuit on a network which consists of a sequence of Q different Hamiltonian evolutions, with the whole sequence repeated P times:

$$\hat{U}_{\text{QGNN}}(\boldsymbol{\eta}, \boldsymbol{\theta}) = \prod_{p=1}^P \left[\prod_{q=1}^Q e^{-i\eta_{pq} \hat{H}_q(\boldsymbol{\theta})} \right], \quad (3)$$

where the product is time-ordered (Poulin et al., 2011), the $\boldsymbol{\eta}$ and $\boldsymbol{\theta}$ are variational (trainable) parameters, and the Hamiltonians $\hat{H}_q(\boldsymbol{\theta})$ can generally be any parameterized Hamiltonians whose topology of interactions is that of the problem graph:

$$\hat{H}_q(\boldsymbol{\theta}) \equiv \sum_{\{j,k\} \in \mathcal{E}} \sum_{r \in \mathcal{I}_{jk}} W_{qrjk} \hat{O}_j^{(qr)} \otimes \hat{P}_k^{(qr)} + \sum_{v \in \mathcal{V}} \sum_{r \in \mathcal{J}_v} B_{qrv} \hat{R}_j^{(qv)}. \quad (4)$$

Here the W_{qrjk} and B_{qrv} are real-valued coefficients which can generally be independent trainable parameters, forming a collection $\boldsymbol{\theta} \equiv \cup_{q,j,k,r} \{W_{qrjk}\} \cup_{q,v,r} \{B_{qrv}\}$. The operators $\hat{R}_j^{(qv)}$, $\hat{O}_j^{(qr)}$, $\hat{P}_k^{(qr)}$ are Hermitian operators which act on the Hilbert space of the j^{th} node of the graph. The sets \mathcal{I}_{jk} and \mathcal{J}_v are index sets for the terms corresponding to the edges and nodes, respectively. To make compilation easier, we enforce that the terms of a given Hamiltonian \hat{H}_q commute with one another, but different \hat{H}_q 's need not commute.

In order to make the ansatz more amenable to training and avoid the barren plateaus (quantum parametric circuit no free lunch) problem (McClean et al., 2018), we need to add some constraints and specificity. To that end, we now propose more specialized architectures where parameters are tied spatially (convolutional) or tied over the sequential iterations of the exponential mapping (recurrent).

3.2 QUANTUM GRAPH RECURRENT NEURAL NETWORKS (QGRNN)

We define quantum graph recurrent neural networks as ansatz of the form of equation 3 where the temporal parameters are tied between iterations, $\eta_{pq} \mapsto \eta_q$. In other words, we have tied the parameters between iterations of the outer sequence index (over $p = 1, \dots, P$). This is akin to

classical recurrent neural networks where parameters are shared over sequential applications of the recurrent neural network map. As η_q acts as a time parameter for Hamiltonian evolution under \hat{H}_q , we can view the QGRNN ansatz as a Trotter-based (Lloyd, 1996; Poulin et al., 2011) quantum simulation of an evolution $e^{-i\Delta\hat{H}_{\text{eff}}}$ under the Hamiltonian $\hat{H}_{\text{eff}} = \Delta^{-1} \sum_q \eta_q \hat{H}_q$ for a time step of size $\Delta = \|\boldsymbol{\eta}\|_1 = \sum_q |\eta_q|$. This ansatz is thus specialized to learn effective quantum Hamiltonian dynamics for systems living on a graph. In Section 3 we demonstrate this by learning the effective real-time dynamics of an Ising model on a graph using a QGRNN ansatz.

3.3 QUANTUM GRAPH CONVOLUTIONAL NEURAL NETWORKS (QGCNN)

Classical Graph Convolutional neural networks rely on a key feature: that of permutation invariance. In other words, the ansatz should be invariant under permutation of the nodes. This is analogous to translational invariance for ordinary convolutional transformations. In our case, permutation invariance manifests itself as a constraint on the Hamiltonian, which now should be devoid of *local* trainable parameters, and should only have global trainable parameters. The θ parameters thus become tied over indices of the graph: $W_{q_r j k} \mapsto W_{q_r}$ and $B_{q_r v} \mapsto B_{q_r}$. A broad class of graph convolutional neural networks we will focus on is the set of so-called Quantum Alternating Operator Ansatz (Hadfield et al., 2019), the generalized form of the Quantum Approximate Optimization Algorithm ansatz (Farhi et al., 2014).

3.4 QUANTUM SPECTRAL GRAPH CONVOLUTIONAL NEURAL NETWORKS (QSGCNN)

We can take inspiration from the continuous-variable quantum approximate optimization ansatz introduced in Verdon et al. (2019a) to create a variant of the QGCNN: the Quantum Spectral Graph Convolutional Neural Network (QSGCNN). We show here how it recovers the mapping of Laplacian-based graph convolutional networks (Kipf & Welling, 2016) in the Heisenberg picture, consisting of alternating layers of message passing, node update, and nonlinearities.

Consider an ansatz of the form from equation 3 with four different Hamiltonians ($Q = 4$) for a given graph. First, for a weighted graph \mathcal{G} with edge weights Λ_{jk} , we define the *coupling Hamiltonian* as

$$\hat{H}_C \equiv \frac{1}{2} \sum_{\{j,k\} \in \mathcal{E}} \Lambda_{jk} (\hat{x}_j - \hat{x}_k)^2.$$

The Λ_{jk} here are the *weights* of the graph \mathcal{G} , and are *not* trainable parameters. The operators denoted here by \hat{x}_j are quantum continuous-variable position operators, which can be implemented via continuous-variable (analog) quantum computers (Weedbrook et al., 2012) or emulated using multiple qubits on digital quantum computers (Somma, 2015; Verdon et al., 2018). After evolving by \hat{H}_C , which we consider to be the message passing step, one applies an exponential of the *kinetic* Hamiltonian, $\hat{H}_K \equiv \frac{1}{2} \sum_{j \in \mathcal{V}} \hat{p}_j^2$. Here \hat{p}_j denotes the continuous-variable momentum (Fourier conjugate) of the position, obeying the canonical commutation relation $[\hat{x}_j, \hat{p}_j] = i\delta_{jk}$. We consider this step as a node update step. In the Heisenberg picture, the evolution generated by these two steps maps the position operators of each node according to

$$e^{-i\gamma\hat{H}_K} e^{-i\alpha\hat{H}_C} : \hat{x}_j \mapsto \hat{x}_j + \gamma\hat{p}_j - \alpha\gamma \sum_{k \in \mathcal{V}} L_{jk} \hat{x}_k,$$

where

$$L_{jk} = \delta_{jk} \left(\sum_{v \in \mathcal{V}} \Lambda_{jv} \right) - \Lambda_{jk}$$

is the *Graph Laplacian* matrix for the weighted graph \mathcal{G} . We can recognize this step as analogous to classical spectral-based graph convolutions. One difference to note here is that *momentum* is free to accumulate between layers.

Next, we must add some non-linearity in order to give the ansatz more capacity.¹ The next evolution is thus generated by an *anharmonic* Hamiltonian $\hat{H}_A = \sum_{j \in \mathcal{V}} f(\hat{x}_j)$, where f is a nonlinear function of degree greater than 2, e.g., a quartic potential of the form $f(\hat{x}_j) = ((\hat{x}_j - \mu)^2 - \omega^2)^2$ for some

¹From a quantum complexity standpoint, adding a nonlinear operation (generated by a potential of degree superior to quadratic) creates states that are *non-Gaussian* and hence are non efficiently simulatable on classical computers (Bartlett et al., 2002), in general composing layers of Gaussian and non-Gaussian quantum transformations yields quantum computationally universal ansatz. (Lloyd & Braunstein, 1999).

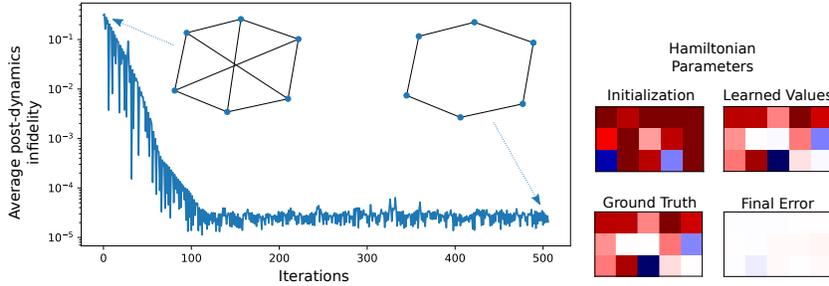


Figure 1: Left: Batch average infidelity with respect to ground truth state sampled at 15 randomly chosen times of quantum Hamiltonian evolution. We see the initial guess has a densely connected topology and the QGRNN learns the ring structure of the true Hamiltonian. Right: Ising Hamiltonian parameters (weights & biases) on a color scale.

μ, ω hyperparameters. Finally, we apply another evolution according to the kinetic Hamiltonian. These last two steps yield an update

$$e^{-i\beta\hat{H}_K} e^{-i\delta\hat{H}_A} : \hat{x}_j \mapsto \hat{x}_j + \beta\hat{p}_j - \delta\beta f'(\hat{x}_j),$$

which acts as a nonlinear mapping. By repeating the four evolution steps described above in a sequence of P layers, i.e.,

$$\hat{U}_{\text{QSGCNN}}(\alpha, \beta, \gamma, \delta) = \prod_{j=1}^P e^{-i\beta_j\hat{H}_K} e^{-i\delta_j\hat{H}_A} e^{-i\gamma_j\hat{H}_K} e^{-i\alpha_j\hat{H}_C}$$

with variational parameters $\theta = \{\alpha, \beta, \gamma, \delta\}$, we then recover a quantum-coherent analogue of the node update prescription of Kipf & Welling (2016) in the original graph convolutional networks paper.²

4 APPLICATIONS & EXPERIMENTS

4.1 LEARNING QUANTUM HAMILTONIAN DYNAMICS WITH QUANTUM GRAPH RECURRENT NEURAL NETWORKS

Learning the dynamics of a closed quantum system is a task of interest for many applications (Wiebe et al., 2014), including device characterization and validation. In this example, we demonstrate that a Quantum Graph Recurrent Neural Network can learn effective dynamics of an Ising spin system when given access to the output of quantum dynamics at various times.

Our target is an Ising Hamiltonian with transverse field on a particular graph,

$$\hat{H}_{\text{target}} = \sum_{\{j,k\} \in \mathcal{E}} J_{jk} \hat{Z}_j \hat{Z}_k + \sum_{v \in \mathcal{V}} Q_v \hat{Z}_v + \sum_{v \in \mathcal{V}} \hat{X}_v.$$

We are given copies of a fixed low-energy state $|\psi_0\rangle$ as well as copies of the state $|\psi_T\rangle \equiv \hat{U}(T)|\psi_0\rangle = e^{-iT\hat{H}_{\text{target}}}$ for some known but randomly chosen times $T \in [0, T_{\text{max}}]$. Our goal is to learn the target Hamiltonian parameters $\{J_{jk}, Q_v\}_{j,k,v \in \mathcal{V}}$ by comparing the state $|\psi_T\rangle$ with the state obtained by evolving $|\psi_0\rangle$ according to the QGRNN ansatz for a number of iterations $P \approx T/\Delta$ (where Δ is a hyperparameter determining the Trotter step size). We achieve this by training the parameters via Adam (Kingma & Ba, 2014) gradient descent on the average infidelity $\mathcal{L}(\theta) = 1 - \frac{1}{B} \sum_{j=1}^B |\langle \psi_{T_j} | U_{\text{QGRNN}}^j(\Delta, \theta) | \psi_0 \rangle|^2$ averaged over batch sizes of 15 different times T . Gradients were estimated via finite difference differentiation with step size $\epsilon = 10^{-4}$. The fidelities (quantum state overlap) between the output of our ansatz and the time-evolved data state were estimated via the quantum swap test (Cincio et al., 2018). The ansatz uses a Trotterization of a random densely-connected Ising Hamiltonian with transverse field as its initial guess, and successfully learns the Hamiltonian parameters within a high degree of accuracy as shown in Fig. 1a.

²For further physical intuition about the behaviour of this ansatz, note that the sum of the coupling and kinetic Hamiltonians $\hat{H}_K + \hat{H}_C$ is equivalent to the Hamiltonian of a network of quantum harmonic oscillators coupled according to the graph weights and network topology. By adding a quartic \hat{H}_A , we are thus emulating parameterized dynamics on a harmonically coupled network of anharmonic oscillators.

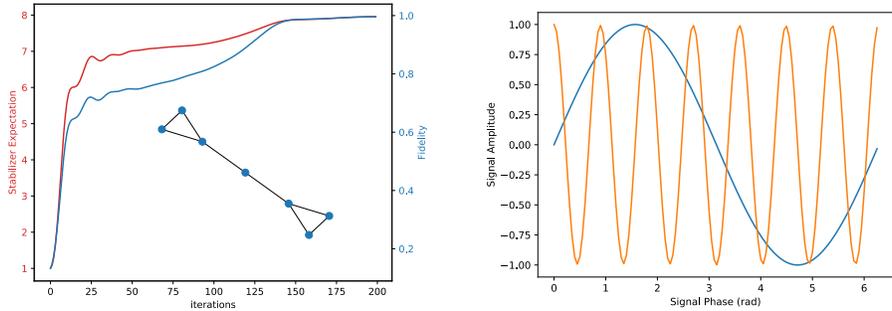


Figure 2: Left: Stabilizer Hamiltonian expectation and fidelity over training iterations. A picture of the quantum network topology is inset. Right: Quantum phase kickback test on the learned GHZ state. We observe a 7x boost in Rabi oscillation frequency for a 7-node network, thus demonstrating we have reached the Heisenberg limit of sensitivity for the quantum sensor network.

4.2 QUANTUM GRAPH CONVOLUTIONAL NEURAL NETWORKS FOR QUANTUM SENSOR NETWORKS

Quantum Sensor Networks are a promising area of application for the technologies of Quantum Sensing and Quantum Networking/Communication (Kimble, 2008; Qian et al., 2019). A common task considered where a quantum advantage can be demonstrated is the estimation of a parameter hidden in weak qubit phase rotation signals, such as those encountered when artificial atoms interact with a constant electric field of small amplitude (Qian et al., 2019). A well-known method to achieve this advantage is via the use of a quantum state exhibiting multipartite entanglement of the Greenberger–Horne–Zeilinger kind, also known as a GHZ state (Greenberger et al., 1989). Here we demonstrate that, without global knowledge of the quantum network structure, a QGCNN ansatz can learn to prepare a GHZ state. We use a QGCNN ansatz with $\hat{H}_1 = \sum_{\{j,k\} \in \mathcal{E}} \hat{Z}_j \hat{Z}_k$ and $\hat{H}_2 = \sum_{j \in \mathcal{V}} \hat{X}_j$. The loss function is the negative expectation of the sum of stabilizer group generators which stabilize the GHZ state (Tóth & Gühne, 2005), i.e.,

$$\mathcal{L}(\eta) = - \langle \bigotimes_{j=0}^n \hat{X} + \sum_{j=1}^{n-1} \hat{Z}_j \hat{Z}_{j+1} \rangle_{\eta}$$

for a network of n qubits. Results are presented in Fig. 1b. Note that the advantage of using a QGNN ansatz on the network is that the number of quantum communication rounds is simply proportional to P , and that the local dynamics of each node are independent of the global network structure.

In order to further validate that we have obtained an accurate GHZ state on the network after training, we perform the quantum phase kickback test on the network’s prepared approximate GHZ state (Wei et al., 2019).³ We observe the desired frequency boost effect for our trained network preparing an approximate GHZ state at test time, as displayed in Figure 2.

4.3 UNSUPERVISED GRAPH CLUSTERING WITH QUANTUM GRAPH CONVOLUTIONAL NETWORKS

As a third set of applications, we consider applying the QSGCNN from Section 2 to the task of spectral clustering (Ng et al., 2002). Spectral clustering involves finding low-frequency eigenvalues of the graph Laplacian and clustering the node values in order to identify graph clusters. In Fig. 3 we present the results for a QSGCNN for varying multi-qubit precision for the representation of the continuous values, where the loss function that was minimized was the expected value of the anharmonic potential $\mathcal{L}(\eta) = \langle \hat{H}_C + \hat{H}_A \rangle_{\eta}$. Of particular interest to near-term quantum computing with low numbers of qubits is the single-qubit precision case, where we modify the QSGCNN construction as $\hat{p}_j^2 \mapsto \hat{X}_j$,

³For this test, one applies a phase rotation $\bigotimes_{j \in \mathcal{V}} e^{-i\varphi \hat{Z}_j}$ on all the qubits in parallel, then one applies a sequence of CNOT’s (quantum adder gates) such as to concentrate the phase shifts onto a single collector node, $m \in \mathcal{V}$. Given that one had a GHZ state initially, one should then observe a phase shift $e^{-in\varphi \hat{Z}_m}$ where $n = |\mathcal{V}|$. This boost in frequency of oscillation of the signal is what gives quantum multipartite entanglement its power to increase sensitivity to signals to super-classical levels (Degen et al., 2017).

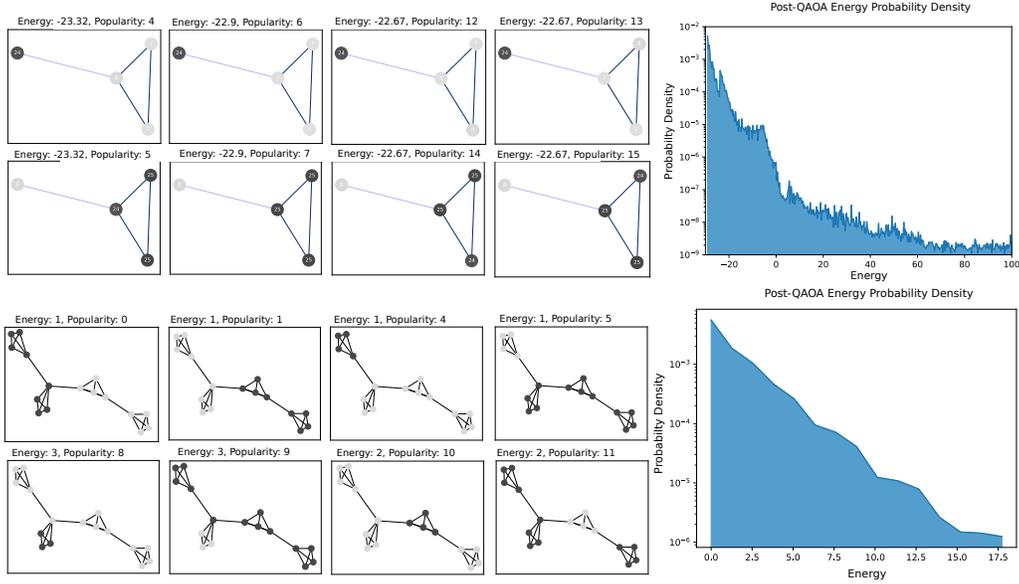


Figure 3: QSGCNN spectral clustering results for 5-qubit precision (top) with quartic double-well potential and 1-qubit precision (bottom) for different graphs. Weight values are represented as opacity of edges, output sampled node values as grayscale. Lower precision allows for more nodes in the simulation of the quantum neural network. The graphs displayed are the most probable (populated) configurations, and to their right is the output probability distribution over potential energies. We see lower energies are most probable and that these configurations have node values clustered.

$\hat{H}_A \mapsto I$ and $\hat{x}_j \mapsto |1\rangle\langle 1|_j$ which transforms the coupling Hamiltonian as

$$\hat{H}_C \mapsto \frac{1}{2} \sum_{\{j,k\} \in \mathcal{E}} \Lambda_{jk} (|1\rangle\langle 1|_j - |1\rangle\langle 1|_k)^2 = \sum_{jk} L_{jk} |1\rangle\langle 1|_j \otimes |1\rangle\langle 1|_k, \quad (5)$$

where $|1\rangle\langle 1|_k = (\hat{I} - \hat{Z}_k)/2$. We see that using a low-qubit precision yields sensible results, thus implying that spectral clustering could be a promising new application for near-term quantum devices.

4.4 GRAPH ISOMORPHISM CLASSIFICATION VIA QUANTUM GRAPH CONVOLUTIONAL NETWORKS

Recently, a benchmark of the representation power of classical graph neural networks has been proposed (Xu et al., 2018) where one uses classical GCN’s to identify whether two graphs are isomorphic. In this spirit, using the QSGCNN ansatz from the previous subsection, we benchmarked the performance of this Quantum Graph Convolutional Network for identifying isomorphic graphs. We used the single-qubit precision encoding in order to simulate the execution of the quantum algorithms on larger graphs.

Our approach was the following, given two graphs \mathcal{G}_1 and \mathcal{G}_2 , one applies the single-qubit precision QSGCNN ansatz $\prod_{j=1}^P e^{i\eta_j \hat{H}_K} e^{i\gamma_j \hat{H}_C}$ with $\hat{H}_K = \sum_{j \in \mathcal{V}} \hat{X}_j$ and \hat{H}_C from equation 5 in parallel according to each graph’s structure. One then samples eigenvalues of the coupling Hamiltonian \hat{H}_C on both graphs via standard basis measurement of the qubits and computation of the eigenvalue at each sample of the wavefunction. One then obtains a set of samples of “energies” of this Hamiltonian. By comparing the energetic measurement statistics output by the QSGCNN ansatz applied with identical parameters $\theta = \{\eta, \gamma\}$ for two different graphs, one can then infer whether the graphs are isomorphic.

We used the Kolmogorov-Smirnoff test (Lilliefors, 1967) on the distribution of energies sampled at the output of the QSGCNN to determine whether two given graphs were isomorphic. In order to determine the binary classification label deterministically, we considered all KS statistic values above 0.4 to indicate that the graphs were non-isomorphic. For training and testing purposes, we set the

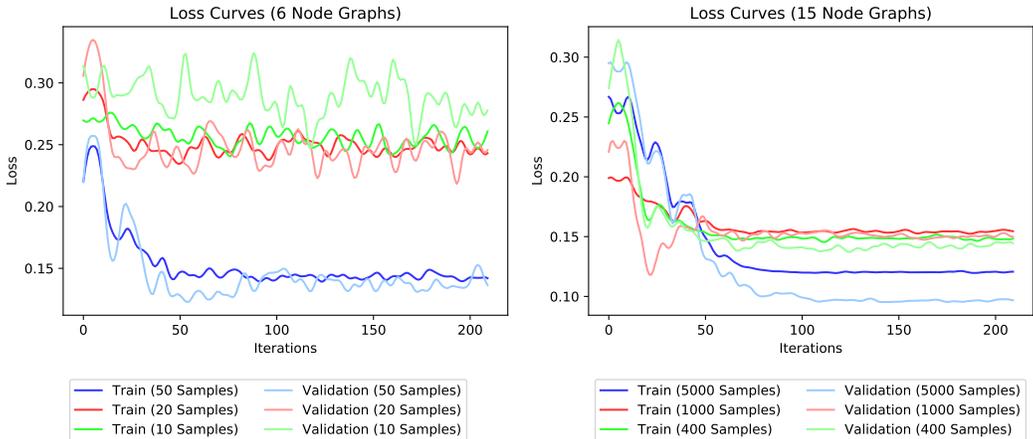


Figure 4: Graph isomorphism loss curves for training and validation for various numbers of samples. Left is for 6 node graphs and right is for 15 node graphs. The loss is based on the Kolmogorov-Smirnov statistic comparing the sampled distribution of energies of the QGCNN output on two graphs.

Table 1: Classification Accuracy for 15 Node Graphs

Samples	Test	Validation
5000	100.0	100.0
1000	100.0	100.0
400	100.0	100.0

Table 2: Classification Accuracy for 6 Node Graphs

Samples	Test	Validation
50	100.0	100.0
20	90.0	100.0
10	100.0	80.0

loss function to be $\mathcal{L}(y, \text{KS}) = (1 - y)(1 - \text{KS}) + y \text{KS}$, where $y = 1$ if graphs are isomorphic, and $y = 0$ otherwise.

For the dataset, graphs were sampled uniformly at random; to prepare a balanced dataset, we selected isomorphic and non-isomorphic pairs. In all of our experiments, we had 100 pairs of graphs for training, 50 for validation, 50 for testing, and in all cases there are balanced isomorphic and non-isomorphic pairs. The networks were trained via Adam gradient-based optimizer with batches of size 50.

Presented in Figure 4 is the training and testing losses for various graph sizes and numbers of energetic samples. In Tables 1 and 2, we present the graph isomorphism classification accuracy for the training and testing sets using the trained QGCNN with the previously described thresholded KS statistic as the label. We see we get highly accurate performance even at low sample sizes. This seems to imply that the QGCNN is fully capable of identifying graph isomorphism, as desired for graph convolutional network benchmarks. We leave a comparison to similar scale classical graph convolutional networks to future work.

5 CONCLUSION & OUTLOOK

Results featured in this paper should be viewed as a promising set of first explorations of the potential applications of QGNNs. Through our numerical experiments, we have shown the use of these QGNN ansatz in the context of quantum dynamics learning, quantum sensor network optimization, unsupervised graph clustering, and supervised graph isomorphism classification. Given that there is a vast set of literature on the use of Graph Neural Networks and their variants to quantum chemistry, future works should explore hybrid methods where one can learn a graph-based hidden quantum representation (via a QGNN) of a quantum chemical process. As the true underlying process is quantum in nature and has a natural molecular graph geometry, the QGNN could serve as a more accurate model for the hidden processes which lead to perceived emergent chemical properties. We seek to explore this in future work. Other future work could include generalizing the QGNN to include

quantum degrees of freedom on the edges, include quantum-optimization-based training of the graph parameters via quantum phase backpropagation (Verdon et al., 2018), and extending the QSGCNN to multiple features per node.

ACKNOWLEDGMENTS

Numerics in this paper were executed using a custom interface between Google’s Cirq and TensorFlow (Abadi et al., 2016). The authors would like to thank Edward Farhi, Hartmut Neven, Jae Yoo, and Li Li for useful discussions. GV, EL, and VS would like to thank the team at X for the hospitality and support during their respective Quantum@X and AI@X residencies where this work was completed. GV acknowledges funding from NSERC.

REFERENCES

- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.
- Stephen D Bartlett, Barry C Sanders, Samuel L Braunstein, and Kae Nemoto. Efficient classical simulation of continuous variable quantum information processes. *Physical Review Letters*, 88(9): 097904, 2002.
- Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. Spectral networks and locally connected networks on graphs. *arXiv preprint arXiv:1312.6203*, 2013.
- Lukasz Cincio, Yiğit Subaşı, Andrew T Sornborger, and Patrick J Coles. Learning the quantum algorithm for state overlap. *New Journal of Physics*, 20(11):113022, 2018.
- Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in neural information processing systems*, pp. 3844–3852, 2016.
- Christian L Degen, F Reinhard, and P Cappellaro. Quantum sensing. *Reviews of modern physics*, 89(3):035002, 2017.
- Edward Farhi and Hartmut Neven. Classification with quantum neural networks on near term processors. *arXiv preprint arXiv:1802.06002*, 2018.
- Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028*, 2014.
- Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1263–1272. JMLR. org, 2017.
- Marco Gori, Gabriele Monfardini, and Franco Scarselli. A new model for learning in graph domains. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, volume 2, pp. 729–734. IEEE, 2005.
- Daniel M Greenberger, Michael A Horne, and Anton Zeilinger. Going beyond bell’s theorem. In *Bell’s theorem, quantum theory and conceptions of the universe*, pp. 69–72. Springer, 1989.
- Stuart Hadfield, Zihui Wang, Bryan O’Gorman, Eleanor G Rieffel, Davide Venturelli, and Rupak Biswas. From the quantum approximate optimization algorithm to a quantum alternating operator ansatz. *Algorithms*, 12(2):34, 2019.
- Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, pp. 1024–1034, 2017.
- Mikael Henaff, Joan Bruna, and Yann LeCun. Deep convolutional networks on graph-structured data. *arXiv preprint arXiv:1506.05163*, 2015.

- Steven Kearnes, Kevin McCloskey, Marc Berndl, Vijay Pande, and Patrick Riley. Molecular graph convolutions: moving beyond fingerprints. *Journal of computer-aided molecular design*, 30(8): 595–608, 2016.
- H Jeff Kimble. The quantum internet. *Nature*, 453(7198):1023, 2008.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- Li Li, Minjie Fan, Marc Coram, Patrick Riley, and Stefan Leichenauer. Quantum optimization with a novel gibbs objective function and ansatz architecture search. *arXiv preprint arXiv:1909.07621*, 2019.
- Hubert W Lilliefors. On the kolmogorov-smirnov test for normality with mean and variance unknown. *Journal of the American statistical Association*, 62(318):399–402, 1967.
- Seth Lloyd. Universal quantum simulators. *Science*, pp. 1073–1078, 1996.
- Seth Lloyd and Samuel L Braunstein. Quantum computation over continuous variables. In *Quantum Information with Continuous Variables*, pp. 9–17. Springer, 1999.
- Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics*, 18(2):023023, 2016.
- Jarrod R McClean, Sergio Boixo, Vadim N Smelyanskiy, Ryan Babbush, and Hartmut Neven. Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1):4812, 2018.
- Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodola, Jan Svoboda, and Michael M Bronstein. Geometric deep learning on graphs and manifolds using mixture model cnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5115–5124, 2017.
- Andrew Y Ng, Michael I Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm. In *Advances in neural information processing systems*, pp. 849–856, 2002.
- Mathias Niepert, Mohamed Ahmed, and Konstantin Kutikov. Learning convolutional neural networks for graphs. In *International conference on machine learning*, pp. 2014–2023, 2016.
- David Poulin, Angie Qarry, Rolando Somma, and Frank Verstraete. Quantum simulation of time-dependent hamiltonians and the convenient illusion of hilbert space. *Physical review letters*, 106(17):170501, 2011.
- Kevin Qian, Zachary Eldredge, Wenchao Ge, Guido Pagano, Christopher Monroe, James V Porto, and Alexey V Gorshkov. Heisenberg-scaling measurement protocol for analytic functions with quantum sensor networks. *arXiv preprint arXiv:1901.09042*, 2019.
- Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2008.
- Rolando D Somma. Quantum simulations of one dimensional quantum systems. *arXiv preprint arXiv:1503.06319*, 2015.
- Alessandro Sperduti and Antonina Starita. Supervised neural networks for the classification of structures. *IEEE Transactions on Neural Networks*, 8(3):714–735, 1997.
- Géza Tóth and Otfried Gühne. Entanglement detection in the stabilizer formalism. *Physical Review A*, 72(2):022340, 2005.
- Guillaume Verdon, Jason Pye, and Michael Broughton. A universal training algorithm for quantum deep learning. *arXiv preprint arXiv:1806.09729*, 2018.

- Guillaume Verdon, Juan Miguel Arrazola, Kamil Brádler, and Nathan Killoran. A quantum approximate optimization algorithm for continuous problems. *arXiv preprint arXiv:1902.00409*, 2019a.
- Guillaume Verdon, Michael Broughton, Jarrod R McClean, Kevin J Sung, Ryan Babbush, Zhang Jiang, Hartmut Neven, and Masoud Mohseni. Learning to learn with quantum neural networks via classical neural networks. *arXiv preprint arXiv:1907.05415*, 2019b.
- Christian Weedbrook, Stefano Pirandola, Raúl García-Patrón, Nicolas J Cerf, Timothy C Ralph, Jeffrey H Shapiro, and Seth Lloyd. Gaussian quantum information. *Reviews of Modern Physics*, 84(2):621, 2012.
- Ken X Wei, Isaac Lauer, Srikanth Srinivasan, Neereja Sundaresan, Douglas T McClure, David Toyli, David C McKay, Jay M Gambetta, and Sarah Sheldon. Verifying multipartite entangled ghz states via multiple quantum coherences. *arXiv preprint arXiv:1905.05720*, 2019.
- Nathan Wiebe, Christopher Granade, Christopher Ferrie, and David G Cory. Hamiltonian learning and certification using quantum resources. *Physical review letters*, 112(19):190501, 2014.
- Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*, 2018.