Growing Circuit Analysis

Callum Duffy, Smit Chaudhary, Gergana V. Velikova
PASQAL, 2 av. Augustin Fresnel Palaiseau, 91120, France

Quantum scientific machine learning (QSciML) addresses complex science and engineering problems, including solving partial differential equations (PDEs). These problems often challenge conventional computational methods, where, in some cases, solutions become intractable. A notable QSciML method using quantum devices involves differentiable quantum circuits (DQCs) [1]. Inspired by physics-informed neural networks (PINNs), which utilize classical neural networks to represent PDE solutions, DQCs extends this concept to quantum neural networks (QNNs). DQCs can exploit a spectral basis-set size that scales exponentially with the number of qubits for accurate high-dimensional solutions on near-term devices. This study pragmatically analyzes the DQC algorithm, wherein the spectral basis can vary in size during training due to the circuit being dynamic with respect to the number of gates present, potentially yielding compact high performant QNNs.

The chosen architecture of a QNN influences its corresponding inductive biases, convergence time and generalization capabilities. Traditional approaches, such as neural architecture search [2] and evolutionary methods [3], have been widely explored. In this study, we focus on a novel strategy for growing a given QNN from an initial set of small gates by strategically adding in more gates during training. The idea of a dynamic neural network has been previously studied and shown to be effective for finding compact models [4]. Specifically, we concentrate on QNNs that take the form of reuploader circuits, with our growing method stemming from considering these circuits as truncated Fourier series [5]. As a result, we analyze how the addition of feature map gates, which increase the number of available frequency modes available to the QNN, affects the resulting performance.

We consider various schemes for growing reuploader-based QNNs, taking into account parameter initialization, gradient clipping as well as where and when to add additional reuploader blocks. As an initial proof of principle we demonstrate the efficacy of the design choices by performing regression based tasks on 1D and 2D truncated Fourier series, since the exact solutions to these problems are contained in the QNN's solution space. We then apply these techniques to solving partial differential equations (PDEs) such as the heat equation for finding efficient compact models which are known to have Fourier-like solutions. Moreover, we assess the potential regularisation properties of these grown QNNs through regression tasks on noisy truncated Fourier series. In scenarios where deep overparameterized models overfit and fail to generalize, misinterpreting noise as high-frequency modes, grown models better approximate the underlying Fourier series by dynamically selecting the appropriate number of frequencies.

The reuploader models in this study are formed of feature embedding gates $S(x)$ and trainable gates $W(\theta)$, where x is a feature and θ is a trainable parameter. A single reuploader block can then be represented as $S(x)W(\theta)$, which is then repeated a user-specified number of times L before measurement. Parameterized quantum circuits of this form can be expressed as truncated Fourier series

$$f(x) = \sum_{\omega \in \Omega} c_\omega(\theta) e^{i\omega x}, \tag{1}$$

where Ω is the set of available frequencies to the model determined by the feature encoding gates $S(x)$.

Figure 1 shows how a growing QNN can aid regression problems. The regression problems we present here have target functions that take the form of a truncated Fourier series

$$f(x) = \sum_{n=0}^{N} c_n e^{inx}, \tag{2}$$

the target functions we consider have coefficients $\{0.1, 0.15 + 0.15i\}$ and $\{0.1, 0.15 + 0.15i, 0.10 + 0.10i, 0.10 + 0.10i, 0.05 + 0.05i\}$ respectively. By choosing target functions of this form, we can ensure the exact solution to the problem lies within the solution space of the reuploader models, allowing us to analyze their ability to converge. For each of the two regression problems, we have trained a model labelled as the 'exact model', the number of feature mapping gates $S(x)$ is equal to the degree of the truncated Fourier series, thus has the same number of coefficients $(c_\omega('theta))$ as the truncated Fourier series it is trying to fit. As a result, these models easily fit their respective target functions, achieving losses on the order of 10^{-13} relatively quickly. Deep models are defined to have 50 reuploader layers and can be seen in each case to struggle to converge to loss values that are competitive with the other models; this can be attributed to how the model is far too expressive for the problem at hand. Lastly, we present two versions of the growing QNN, the first dubbed the 'growing model', which begins with a single feature map gate $S(x)$ and 50 trainable gates. As training progresses, more feature map gates are inserted between trainable gates until a model which matches the architecture of the 'deep model' is created. This model, despite having an identical form to the 'deep model' by the end of training, achieves a lower loss, potentially due to being able to focus on training lower frequency modes at the beginning of training. The second of the growing models, 'growing model v2', begins with a single reuploader block and adds more blocks during training, but only once the loss has plateaued and ceases to add more blocks if the loss does not improve with further training. This form of growing model can be seen to be highly effective in producing a model which performs as well as the 'exact model'. The final architecture of the model is chosen depending on the number of reuploader blocks that have been added and, thus, the number of frequency modes it has access to.

References

[1] O. Kyriienko, A. E. Paine, and V. E. Elfving, *Physical Review A*, vol. 103, no. 5, 2021.

[2] T. Duong, S. T. Truong, M. Tam, B. Bach, J.-Y. Ryu, and J.-K. K. Rhee, *Quantum neural architecture search with quantum circuits metric and bayesian optimization*, 2022. arXiv: `2206.14115 [quant-ph]`.

[3] M. Lourens, I. Sinayskiy, D. K. Park, C. Blank, and F. Petruccione, "Hierarchical quantum circuit representations for neural architecture search," *npj Quantum Information*, vol. 9, no. 1, pp. 1–15, 1 Aug. 5, 2023, ISSN: 2056-6387. DOI: `10.1038/s41534-023-00747-z`.

[4] U. Evci, M. Vladymyrov, T. Unterthiner, B. van Merrienboer, and F. Pedregosa, "Gradmax: Growing neural networks using gradient information," *ArXiv*, vol. abs/2201.05125, 2022.

[5] M. Schuld, R. Sweke, and J. J. Meyer, "Effect of data encoding on the expressive power of variational quantum-machine-learning models," *Physical Review A*, vol. 103, no. 3, Mar. 2021, ISSN: 2469-9934. DOI: `10.1103/physreva.103.032430`.

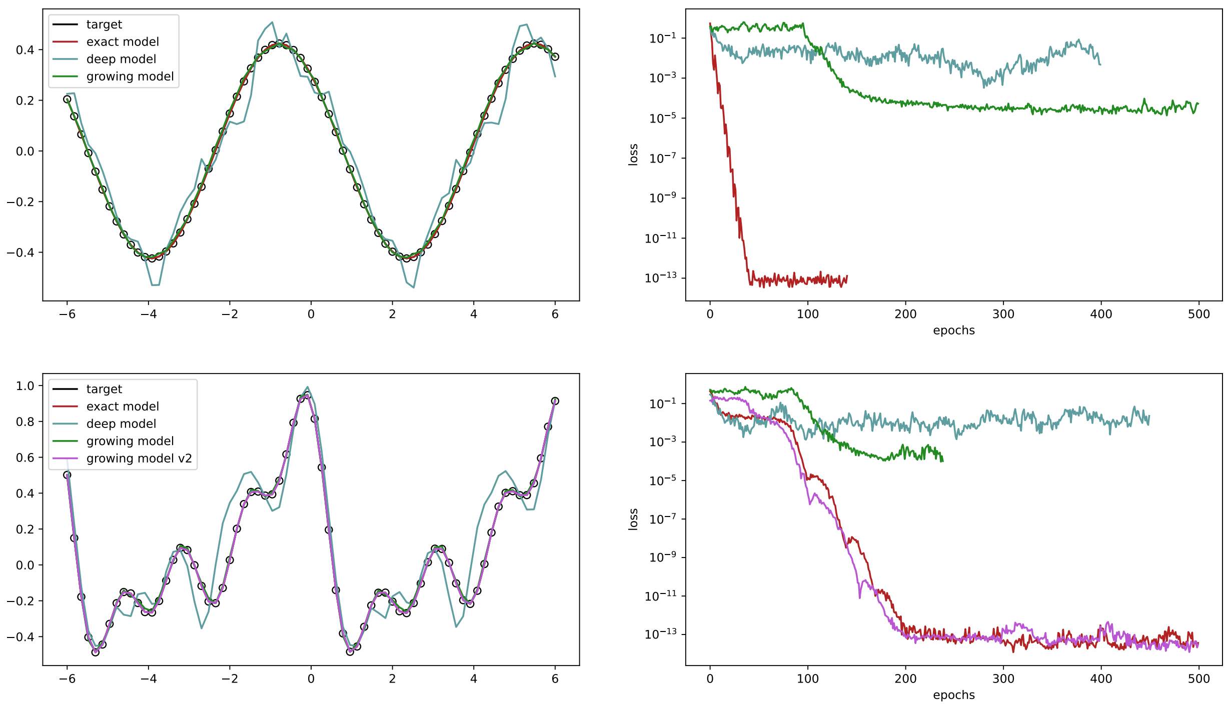

Figure 1: Various QNN designs perform a regression task on truncated Fourier series. (top row) The function to be fitted is a truncated Fourier series with coefficients $\{0.1, 0.15 + 0.15i\}$; on the left, we see the resulting fit of each of the models, the deep model fits most poorly while the exact and growing model fit far better. This can be corroborated by the corresponding losses seen throughout training on the right. (bottom row) The target function here is a degree four truncated Fourier series with coefficients $\{0.1, 0.15 + 0.15i, 0.10 + 0.10i, 0.10 + 0.10i, 0.05 + 0.05i\}$. On the left, we see the final fits with the deep model performing the most poorly again. On the right, we see how 'growing model v2' can converge to a solution as good as the 'exact model' in the same number of epochs.