Submission code: NG7X96AYV9

See this and other abstracts on our website: https://home.sbgf.org.br/Pages/resumos.php

Quantum Layer Integration for Parameter-Efficient PINNs in Wave Equations

Bruno Fernandez (Senai-Cimatec), Daniel Revelo (SENAI CIMATEC), Caio Leão (SENAI CIMATEC), Gleydson Fernandes de Jesus (Senai Cimatec), Yan Alef Chagas Silva (QuIIN - SENAI CIMATEC)

SBGf Conference
18-20 NOV | Rio'25

Quantum Layer Integration for Parameter-Efficient PINNs in Wave Equations

Summary

Physics-Informed Neural Networks (PINNs) embed the governing differential equation in their loss function, but often require thousands of trainable parameters. We introduce a Quantum-enhanced PINN (QPINN) in which a single classical hidden layer is replaced by a four-qubit Variational Quantum Circuit (VQC). Two one-dimensional wave-propagation problems validate the approach. In a Hardy-constrained wave-equation test, the QPINN achieved lower error rates while using 71% fewer parameters than the classical PINN. A second experiment with Gaussian-pulse excitation showed similar results: the hybrid model obtained improved accuracy with approximately 70% parameter reduction. These results demonstrate the viability of quantum layers as feature extractors in PINN architectures.

Introduction

The modeling of partial differential equations (PDEs) with traditional methods like the Finite Difference Method (FDM) faces challenges in complex or high-dimensional settings. Physics-Informed Neural Networks (PINNs) have emerged as a mesh-free, learning-based alternative that approximates PDE solutions by embedding physical laws directly into the network's loss function (Raissi et al., 2019). Boundary conditions in PINNs approach can be enforced through approaches such as "Hard constraints," where they are structurally imposed on the network output (Alkhadhr and Almekkawy, 2023). These well-defined initial profiles and homogeneous boundary conditions enable precise evaluation of the models' ability to capture spatiotemporal wave evolution. Such simulations are crucial for applications like seismic wave propagation modeling and subsurface property inversion, where accuracy and computational efficiency are paramount. By employing physics-based formulations and hybrid architectures, we explore how quantum layers can reduce model complexity while preserving or enhancing solution fidelity.

Despite their advantages, PINNs often involve large numbers of trainable parameters, leading to high computational costs and long training times. This motivates the search for more efficient architectures that balance expressiveness with reduced complexity. In this work, we investigate PINNs and their quantum extensions (QPINNs) applied to the one-dimensional wave equation in geophysically relevant scenarios, expanding the experimental analysis performed in Fernandez et al. (2025). Two experiments are considered: a sinusoidal initial condition and a transient Gaussian pulse. Both applicable fundamental cases in acoustic wave propagation through homogeneous media, commonly used to model seismic wave dynamics.

Theory

Wave propagation in a one-dimensional medium without energy loss follows the classical wave equation governing small-amplitude acoustic disturbances:

$$\partial_t^2 u(x,t) = c^2\, \partial_x^2 u(x,t), \tag{1}$$

where $u(x,t)$ denotes the wavefield and c is the constant wave speed. A well-posed problem requires an initial profile $u(x,0) = u_0(x)$, optionally the initial velocity $\partial_t u(x,0)$, and boundary conditions on the spatial domain $\Omega = [0,1]$.

To evaluate the ability of PINNs and QPINNs to solve this equation, we consider two excitation scenarios. In the first, the model is initialized with a known spatial profile and homogeneous boundary conditions. Following the approach of embedding the initial and boundary conditions into the network output, we define the solution as

$$\hat{u}(x,t) = t\, x(1-x)\, \mathcal{N}_\theta(x,t) + u_0(x),$$

where $\mathcal{N}_\theta(x,t)$ is the raw output of the neural network. This formulation ensures that the boundary conditions $\hat{u}(0,t) = \hat{u}(1,t) = 0$ and the initial condition $\hat{u}(x,0) = u_0(x)$ are automatically satisfied, reducing the loss function to the residual of the PDE. As a reference profile, we use a smooth sinusoidal function $u_0(x) = \sin(\pi x)$, which generates a standing wave and allows for analytical comparison.

To investigate the network's ability to model transient and localized phenomena, we perform a second experiment using a compact, broadband Gaussian pulse as excitation. This is implemented either as a spatially localized initial condition,

$$u_0(x) = A\, e^{-\frac{(x-x_0)^2}{2\sigma_x^2}},$$

or as a boundary-driven source,

$$u(0,t) = A\, e^{-\frac{(t-t_0)^2}{2\sigma_t^2}}, \qquad u(1,t) = 0.$$

Gaussian pulses are standard benchmarks in wave propagation problems, as they allow for the evaluation of dispersion, reflection, and wavefront dynamics in a controlled setting.

Quantum-Enhanced Physics-Informed Neural Networks (QPINN)

We propose the use of a hybrid QPINN architecture where the first classical hidden layer is replaced by a quantum layer, a variational quantum circuit (VQC). Initial classical layers perform the pre-processing role where they receive the raw input coordinates, (x,t), and transform them into an abstract feature set. This prepares the information for the quantum encoding step, where the classically processed features parameterize the quantum layer. This is typically achieved by using the classical data to control the rotation angles of quantum gates, which effectively "imprints" the information onto a quantum state.

The quantum layer part acts as the trainable quantum analog of a classical hidden layer. The circuits consist of a sequence of quantum gates whose parameters are systematically optimized during the network's training phase. Its architecture is fundamental to its performance, often comprising alternating layers of single-qubit rotations and multi-qubit entangling gates, such as the CNOT gate. Entanglement is a key resource, as it creates complex correlations between qubits, enabling the circuit to explore the vast computational Hilbert space and learn relationships that might be difficult for a classical network to capture efficiently .

Following the quantum processing within the quantum layer, the final state is measured to collapse it into a set of classical values, such as operator expectation values. These outcomes form a quantum-processed feature set. These features are then passed to the final classical layers, which perform the post-processing task. Their function is to interpret the abstract, quantum-processed information and map it to the physically meaningful values of the final solution, $\hat{u}(x,t)$. The entire hybrid system is trained end-to-end; while the classical weights are updated with standard backpropagation, the quantum layer parameters are updated via specialized methods like the parameter-shift rule.

Results

All models were coded in `PyTorch` (version 2.4.1), with the quantum layer built in `PennyLane` (version 0.38.0). Training was performed using the Adam optimizer, with a cosine-annealing learning-rate schedule additionally employed in the Gaussian-pulse test. Each QPINN incorporated a four-qubit circuit, and we benchmarked four common VQCs (*alternate*, *cascade*, *cross-mesh*, and *layered*) retaining the best performer in every experiment.

In the first experiment with Hard constraints, a fully classical PINN containing $1\,630$ parameters was trained for $3\,000$ epochs and reached an $\mathrm{RMSE} = 5.91 \times 10^{-4}$. Replacing one hidden layer by a VQC (*layered* ansatz) reduced the parameter count to 363 (71%) and improved the error to 1.64×10^{-5} (Table 1).

Table 1: Performance comparison of various quantum layer approaches within the QPINN model for the hard constraints experiment.

Circuits	Parameters	RMSE
alternate	371	8.3662×10^{-5}
cascade	363	2.9872×10^{-4}
cross-mesh	363	7.5867×10^{-5}
layered	363	1.6395×10^{-5}

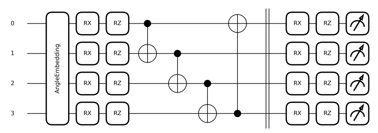

Figure 1: Representation of the circuit layered VQC approach.

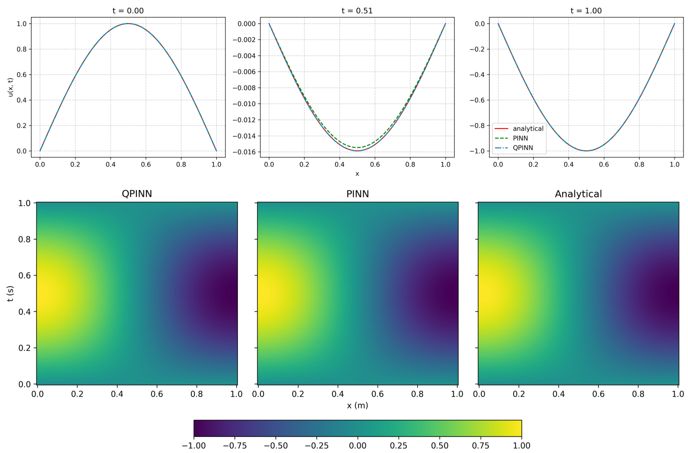

Figure 2: Comparison of wavefield (layered VQC) solutions for different configurations. Top: wave snapshots at different times. Bottom: full spatial wavefield distributions.

For the second experiment involving transient Gaussian-pulse propagation, we trained for $20\,000$ epochs. The classical PINN used $5\,755$ parameters and obtained $\mathrm{RMSE} = 0.192$. The hybrid QPINN, with $1\,638$ parameters, lowered the error to 0.041, a four-fold improvement while employing only 29% of the parameters (Table 2).

Figures 2 and 4 shows the spatiotemporal evolution of the wave for the Hardy and Gaussian tests, respectively. In the hard constraints experiment, the snapshot at $t = 0.51$ s reveals that the QPINN curve is virtually indistinguishable from the analytical reference, whereas the classical PINN exhibits a slight phase lag. The second experiment highlights an even clearer contrast: as the Gaussian pulse propagates, the classical PINN gradually deteriorates, especially near the domain boundaries, while the QPINN continues to track the ground truth with better fidelity across all time frames.

Table 2: Performance comparison of various quantum layer approaches within the QPINN model for the Gaussian pulse experiment

Circuits	Parameters	RMSE
alternate	1646	0.0557
cascade	1638	0.5071
cross-mesh	1638	0.0410
layered	1638	0.2637

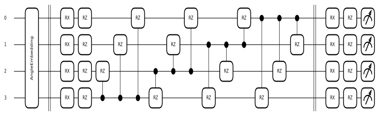

Figure 3: Representation of the cross-mesh VQC approach.

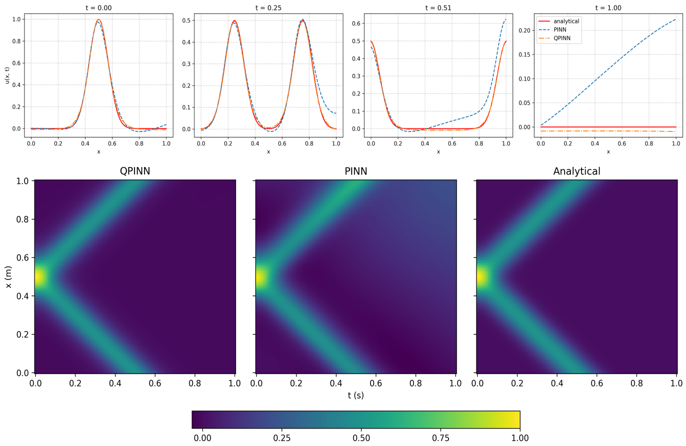

Figure 4: Gaussian pulse (cross-mesh VQC): wave snapshots (top) and 2D solutions at $t = 1.00\,\text{s}$ (bottom).

Conclusions

This work introduced a hybrid Quantum-enhanced Physics-Informed Neural Network in which a single classical hidden layer is replaced by a four-qubit quantum layer. Across two benchmark problems, a Hardy-constrained wave equation and a Gaussian-pulse propagation, the QPINN achieved superior accuracy while using 71% fewer parameters required by a fully classical PINN. The reduction was most pronounced with the layered ansatz, which, in the Hardy test, lowered the RMSE from 5.9×10^{-4} to 1.6×10^{-5} using just 363 trainable weights. These results highlight the potential of quantum layers to act as powerful feature extractors inside physics-guided machine-learning models. Future work will examine robustness on noisy quantum hardware and extension to higher-dimensional PDEs problems, paving the way for resource-efficient, quantum-assisted solvers in scientific computing.

Acknowledgments

This work has been partially supported by QuIIN - EMBRAPII CIMATEC Competence Center in Quantum Technologies, with financial resources from the PPI IoT/Manufatura 4.0 of the MCTI grant number 053/2023, signed with EMBRAPII.

References

Alkhadhr, S., and M. Almekkawy, 2023, Wave Equation Modeling via Physics-Informed Neural Networks: Models of Soft and Hard Constraints for Initial and Boundary Conditions: Sensors, **23**.

Fernandez, B., C. Leao, D. Revelo, and G. Fernandes, 2025, Exploring Hybrid Physics-Informed Neural Networks with Quantum Layers for Solving the Wave Equation: Fifth EAGE Digitalization Conference & Exhibition, European Association of Geoscientists & Engineers, 1–5.

Raissi, M., P. Perdikaris, and G. Karniadakis, 2019, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations: Journal of Computational Physics, **378**, 686–707.