

Exploiting biased noise in supervised variational quantum models

Connor van Rossum

Riddhi Gupta

Sally Shrapnel

July 7, 2026

Abstract

Variational quantum algorithms (VQAs) have dominated literature as tools for demonstrating quantum utility on near-term quantum hardware, with applications in optimisation, quantum simulation, and machine learning. While researchers have studied how easy VQAs are to train, the effect of quantum noise on the classical optimisation process is still not well understood. Contrary to expectations, we find that twirling, which is commonly used in standard error-mitigation strategies to symmetrise noise, actually degrades performance in the variational setting, whereas preserving biased or non-unital noise can help classical optimisers find better solutions. Analytically, we study a universal quantum regression model and demonstrate that relatively uniform Pauli channels suppress gradient magnitudes and reduce expressivity, making optimisation more difficult. Conversely, asymmetric noise such as amplitude damping or biased Pauli channels introduces directional bias that can be exploited during optimisation. Numerical experiments on a variational eigensolver for the transverse-field Ising model confirm that non-unital noise yields lower-energy states compared to twirled noise. Finally, we show that coherent errors are fully mitigated by re-parameterisation. These findings challenge conventional noise-mitigation strategies and suggest that preserving noise biases may enhance VQA performance.

1 Introduction & Setup

Variational quantum algorithms (VQAs) feature predominantly in quantum applications development for noisy intermediate-scale quantum (NISQ) devices [1, 2, 3]. Typically, the first step in combating noise in quantum circuits is to symmetrise it via twirling, which reduces complex noise to more analytically convenient channels [4, 5]. Twirling over the Pauli group reduces noise to stochastic Pauli errors [6], while twirling over the Clifford group further simplifies noise to a depolarising channel [7, 8]. However, despite its central role in state-of-the-art error mitigation protocols, it remains unclear how noise symmetrisation interacts with classical optimisation.

To systematically characterise how biased or non-unital noise affects variational learning, we first establish a variational learning problem whose noiseless behaviour is analytically exact. While such analytical guarantees are rare for general VQAs, recent results in quantum machine learning provide them for a class of universal quantum models used in classification and regression tasks. These so-called data re-uploading circuits (DRC) can learn any square-integrable function [9], offering a setting where the ideal, zero-error solution is known. This makes them an ideal testbed for isolating the influence of biased and non-unital noise from the inherent difficulty of the optimisation process.

The output of DRCs admit a truncated Fourier series,

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

where the frequency spectrum Ω is determined by the eigenvalues of H , our data encoding Hamiltonian, and the coefficients c_{ω} depend on the trainable parameters $\boldsymbol{\theta}$ and the observable M .

We employ two key properties to assess the practical utility of a variational quantum machine learning model: *expressivity* and *trainability*. Expressivity refers to the model’s capacity to represent a broad class of functions or quantum states, which we quantify by evaluating whether the quantum model can generate a sufficiently broad range of predictions to inform the loss calculation. Randomly generated target functions g' are scaled by a factor a to maximise their output range, such that there still exists a parameter $\boldsymbol{\theta} \in \Theta$ for which the loss $\mathcal{L}(\boldsymbol{\theta})$ falls below the threshold ϵ . Under ideal conditions, this range equals 2, reflecting the bounded spectrum of quantum observables. For trainability, we characterise the distribution of gradient magnitudes of the model output, $|\partial f / \partial \theta_i|$, with respect to each trainable parameter θ_i .

2 Results

While the Fourier structure of DRCs in the noiseless regime is well known, we reformulate this description using a Pauli Transfer Matrix (PTM) formalism, reproducing Eq. (1) via channel representations. We extend this channel representation to then include noise that previously could not be analysed using unitary representations alone. A key insight emerges: noise alters the magnitude and distribution of Fourier coefficients but not the frequency spectrum, which remains determined by the data-encoding Hamiltonians. Leveraging this property, we systematically analyse how representative noise channels reshape the output range and gradient landscape, thereby impacting expressivity and trainability.

Impact of noise on expressivity. For Pauli and coherent noise, the impact on the output function manifests as altered Fourier coefficients in Eq. (1), while amplitude damping introduces additive terms that preclude a simple closed-form expression. When noise channels are composed with the channel representation of the data re-uploading circuit, the resulting noisy output function can be expressed as

$$\tilde{f}(x, \boldsymbol{\theta}) = \sum_{\omega \in \Omega} \tilde{c}_{\omega}(\boldsymbol{\theta}) e^{i\omega x}, \quad (2)$$

which retains the structure of Eq. (1), with noise effects isolated in $\tilde{c}_{\omega}(\boldsymbol{\theta})$, and Ω is unchanged. For noise diagonal in the PTM representation, such as Pauli noise, the circuit output inherits multiplicative factors $\in [0, 1]$ that scale the Fourier coefficients, $\tilde{c}_{\omega}(\boldsymbol{\theta}) = n_{\omega}(\boldsymbol{\theta}) c_{\omega}(\boldsymbol{\theta})$, where $0 \leq n_{\omega}(\boldsymbol{\theta}) \leq 1$ quantifies the attenuation.

In contrast, amplitude damping introduces not only multiplicative attenuation but also a partial reduction in the effective circuit depth. For qubits affected by damping, the channel effectively resets them to the ground state, erasing earlier operations and enabling subsequent layers to act on a simplified state. This depth reduction enables damped qubits to still contribute meaningfully to lower-frequency components in the Fourier spectrum. Although the multiplicative attenuation of Fourier coefficients under amplitude damping mirrors that observed with Pauli noise, there are accompanying terms that additionally contribute to the lower frequency Fourier coefficients. Finally, we remark that coherent noise corresponds to unitary errors. Since the set of unitaries is closed under composition, coherent noise can be absorbed into the parameterisation of the model. For each $\boldsymbol{\theta}$, there exists a reparameterised $\boldsymbol{\theta}'$ such that $\tilde{c}_{\omega}(\boldsymbol{\theta}) \equiv c_{\omega}(\boldsymbol{\theta}')$.

We numerically quantify expressivity under noise in Fig. 1. Here, we measure the output range of the trained model under different noise conditions. For each seeded target function, we use binary search to find the largest output range for which the model converges. A trial succeeds if the final loss is below $\epsilon = 10^{-5}$ within the training budget of 150 steps. Fig. 1(a) shows the maximum range for which the model can successfully learn a target function under coherent noise, amplitude damping, and channels equivalent to Pauli-/Clifford-twirled amplitude damping, where we use the same set of seeded functions for each noise type. The twirled channels are simulated as equivalent noise channels rather than physically twirled circuits, due to the non-Clifford nature of the ansatz. These results confirm our theoretical analysis.

3 Conclusion

Our work counterintuitively indicates that bias-preserving noisy variational circuits may improve overall performance, whereas naively twirling noise may instead hinder classical optimisation of VQAs. Our findings suggest VQAs may benefit from retaining biases in intrinsic noise to preserve expressivity and trainability of quantum models. One implication appears to be that commonly used noise reshaping techniques, such as Pauli twirling, may inadvertently degrade the performance of variational algorithms. This perspective opens new avenues for designing noise-resilient quantum models that are better suited to the operational realities of NISQ hardware.




Figure 1: Impact of noise on model expressivity. (a) Output range of a single-qubit data re-uploading circuit under different noise channels, determined by increasing the amplitude of target functions until the model fails to fit the data. Error bars represent the mean and standard deviation across multiple samples, and lines of best fit were computed using linear regression. The comparison highlights the difference between biased noise (amplitude damping) and more uniform noise (Pauli noise, here simulated as the equivalent of Pauli- and Clifford-twirled amplitude damping), with the former preserving a broader output range and enabling more expressive behaviour. Coherent unitary errors, by contrast, show no degradation in expressivity, as they can be absorbed into the trainable parameters. (b) Upper and lower bounds of the model's output under different noise conditions. The same three example circuits are evaluated across multiple noise types to illustrate how each type deforms the achievable output range. (c) Output range comparison across circuits of increasing depth, where the degree of the target function matches the number of layers. As in (a), error bars denote the mean and standard deviation, and lines of best fit were obtained via linear regression. Results show that deeper circuits amplify the impact of noise, with more uniform channels causing more severe degradation.

References

- [1] John Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79, August 2018.
- [2] M. Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C. Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R. McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, and Patrick J. Coles. Variational Quantum Algorithms. *Nature Reviews Physics*, 3(9):625–644, August 2021.
- [3] Jules Tilly, Hongxiang Chen, Shuxiang Cao, Dario Picozzi, Kanav Setia, Ying Li, Edward Grant, Leonard Wossnig, Ivan Rungger, George H. Booth, and Jonathan Tennyson. The Variational Quantum Eigensolver: A review of methods and best practices. *Physics Reports*, 986:1–128, November 2022.
- [4] Akel Hashim, Ravi K. Naik, Alexis Morvan, Jean-Loup Ville, Bradley Mitchell, John Mark Kreikebaum, Marc Davis, Ethan Smith, Costin Iancu, Kevin P. O’Brien, Ian Hincks, Joel J. Wallman, Joseph Emerson, and Irfan Siddiqi. Randomized compiling for scalable quantum computing on a noisy superconducting quantum processor. *Physical Review X*, 11(4):041039, November 2021.
- [5] David Layden, Bradley Mitchell, and Karthik Siva. Theory of quantum error mitigation for non-Clifford gates, May 2025.
- [6] W. Dür, M. Hein, J. I. Cirac, and H.-J. Briegel. Standard forms of noisy quantum operations via depolarization. *Physical Review A*, 72(5):052326, November 2005.
- [7] Joseph Emerson, Robert Alicki, and Karol Życzkowski. Scalable Noise Estimation with Random Unitary Operators. *Journal of Optics B: Quantum and Semiclassical Optics*, 7(10):S347–S352, October 2005.
- [8] Christoph Dankert, Richard Cleve, Joseph Emerson, and Etera Livine. Exact and Approximate Unitary 2-Designs: Constructions and Applications. *Physical Review A*, 80(1):012304, July 2009.
- [9] Maria Schuld, Ryan Sweke, and Johannes Jakob Meyer. The effect of data encoding on the expressive power of variational quantum machine learning models. *Physical Review A*, 103(3), March 2021.