

Ravines in quantum cost landscapes: opportunities for enhanced VQA predictions

Felix J. Beckmann^{1,2} and João F. Bravo¹

¹ Fraunhofer IAO, Stuttgart, Germany

² Karlsruhe Institute of Technology, Karlsruhe, Germany

Abstract. The geometric and topological structure of quantum cost landscapes (QCLs) governs the optimization and thus the predictive power of variational quantum algorithms (VQAs). We systematically analyze ravines — low-cost paths connecting local minima — using an adapted version of the nudged elastic band (NEB) algorithm, a method originating from theoretical chemistry. By training quantum neural networks (QNNs) to classify the concentratable entanglement of quantum states, we apply the NEB algorithm and numerically identify ravine structures in QCLs of hardware-efficient ansatzes. Beyond visualizing these ravines, we construct an ensemble prediction framework by averaging predictions from QNNs parameterized along the low-cost NEB path.

We introduce a resource-light pre-training metric which quantifies local-prediction variability and serves as a strong performance indicator for VQAs, even beyond the scope of this study. When base classifiers are drawn from circuit and weight initializations exhibiting high local-prediction variability, the quantum-based NEB ensembles outperform both classical and naive quantum alternatives. Moreover, a complexity analysis shows that leveraging the ravine-like structure of QCLs with the QNN NEB approach substantially reduces computational costs compared to naive QNN ensembling.

A depth and qubit scaling analysis indicates that ravines persist across both scalings, and that, despite the expected growth in resource requirements with the qubit scaling, the NEB approach also accelerates convergence over the naive alternative.

Based on the preprint [1], accessible at <https://arxiv.org/abs/2607.01329>.

Keywords: quantum machine learning · nudged elastic band · quantum cost landscape · ravines · mode connectivity · quantum data

1 Extended abstract

Quantum computing has been theoretically shown to offer significant computational advantages and is already being explored across a range of applications. This work focuses on the potential gains that quantum computing could bring to machine learning. Specifically, we analyze the structure of the quantum cost landscape (QCL) of quantum neural networks (QNNs) and show how this structure can be leveraged to enhance their prediction capabilities.

The structure of classical cost landscapes has been analyzed extensively, yielding valuable insights into their geometry and how it can be exploited [4,2]. For the QCL, research has focused predominantly on landscape flatness, namely barren plateaus (BPs). Comparatively little attention has been paid to other geometric and topological properties, despite their relevance in the NISQ era and beyond. Hamilton et al. [3] analyzed ravines in the QCL for a regression task on classical data, finding numerical evidence for their existence in certain configurations.

The present work extends this line of research by employing a native quantum dataset and scaling QNN circuit sizes up to 9 qubits and 288 parameters. After confirming the existence of ravines across the vast majority of tested configurations (Figs. 1 and 2), we leverage this geometrical structure to build a resource-efficient ensemble predictor following the framework presented in Figure 3.

Following the landscape fluctuation metric introduced by Zhang et al. [5], this study proposes an light-weight pre-training metric ζ that evaluates local prediction correlation. This filtering metric ζ captures the most important landscape properties as it serves as a strong indicator of classifier performance with its applicability extending well beyond the scope of this study. As ζ is specifically designed for QNNs, multi-layer perceptrons (MLPs) are instead filtered by χ , the minimum validation cost averaged across the classifiers of one ensemble. This ensures that any performance differences between QNNs are rooted in their prediction structure rather than the choice of filtering metric.

This study considers two axes of comparison: quantum versus classical architectures (QNN vs. MLP) and naive versus NEB ensembling. This setup gives direct access to the sources of performance differences.

Our results show that individual QNN classifiers are less correlated than their classical counterparts, and their individual accuracy lags behind, particularly for high values of ζ and χ . Nevertheless, this lower correlation enables




Fig. 1: Visualization of a QNN NEB path between two initialization points. The path is straightened along the x -axis, with the y -axis representing the direction orthogonal to it. The color encodes the training cost $\mathcal{C}_{\text{tr}}$, and the gray and white crosses mark the initialization and pivot points, respectively.




Fig. 2: Both Subfigures show the QNN training cost normalized to its value at $\alpha = 0$, evaluated along both a linear interpolation and an optimized NEB path between the initialization points. Subfigures (a) and (b) correspond to filtering by the 25th and 100th percentile of the prediction correlation metric ζ , respectively, computed prior to optimization. The solid lines represent the average across configurations, while the shaded areas indicate the corresponding $\pm 1\sigma$ band.




Fig. 3: Illustration of the introduced NEB ensemble prediction framework. Pivot points along the NEB path each define a classifier configuration. A subset of these classifiers is selected and their classification of the input states $|\Psi\rangle$ with respect to their concentratable entanglement (CE) are aggregated across the ensemble, here via averaging.

statistically superior ensemble prediction for the QNN NEB at $\zeta_{\text{pct}}, \chi_{\text{pct}} \in \{0.25, 0.5\}$ (Fig. 4), while reducing computational cost by about 27% as our complexity analysis demonstrates.




Fig. 4: Ensemble test accuracy distributions for both network types (QNN and MLP) and ensembling approaches (NEB and naive), under three filtering thresholds: Subfigures (a), (b), and (c) correspond to the 25th, 50th, and 100th percentile of the prediction correlation metric ζ , computed prior to optimization, or the MLP validation cost filter χ , computed after the optimization, depending on the network type. Dashed vertical lines indicate the median of each distribution.

To assess scalability to deeper and wider circuits, the study concludes with a scaling analysis in both depth and qubits. Depth scaling results are consistent with the existence of ravines. While individual classifier correlation increases due to improved individual performance, ensemble accuracy rises to 91%, underscoring the strength of the approach. Qubit scaling requires adapting the dataset and circuit architecture, yet features typical of ravines persist. Increasing the number of qubits raises resource demands regardless of ensembling strategy as shown in Figure 5; however, our analysis presented in Figure 5 and 6 suggests that NEB ensembling scales more favorably in terms of convergence speed than naive ensembling, making it the superior choice.




Fig. 5: Number of iterations required to achieve a 95% reduction in QNN training cost relative to the convergence limit estimated by an exponential fit.




Fig. 6: Decay exponent b of the exponential fit $f(x) = ae^{bx} + c$ to the average QNN training cost decay.

Overall, our results confirm the existence of ravines across the vast majority of tested configurations and demonstrate a geometry-informed optimization framework for QNNs that reduces resource requirements while preserving high ensemble performance. The introduced pre-training metric ζ , capturing local prediction correlation, further proved to be a strong performance indicator *prior* to optimization, with potential applicability well beyond the scope of this study.

References

1. Beckmann, F.J., Bravo, J.F.: Ravines in quantum cost landscapes: opportunities for improved VQA predictions (2026), <https://arxiv.org/abs/2607.01329>
2. Draxler, F., Veschini, K., Salmhofer, M., Hamprecht, F.: Essentially no barriers in neural network energy landscape. In: Dy, J., Krause, A. (eds.) Proceedings of the 35th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 80, pp. 1309–1318. PMLR, Stockholm, Sweden (2018), <https://proceedings.mlr.press/v80/draxler18a.html>
3. Kathleen E. Hamilton, E. Lynn, R. Pooser: Mode connectivity in the loss landscape of parameterized quantum circuits (2021). <https://doi.org/10.1007/s42484-021-00059-5>
4. Keskar, N.S., Mudigere, D., Nocedal, J., Smelyanskiy, M., Tang, P.T.P.: On Large-Batch Training for Deep Learning: Generalization Gap and Sharp Minima (2017). <https://doi.org/10.48550/arXiv.1609.04836>
5. Zhang, H.K., Zhu, C., Wang, X.: Predicting quantum learnability from landscape fluctuation (2024), <https://doi.org/10.48550/arXiv.2406.11805>