Advancing EEG Signal Analysis with Quantum Machine Learning

Stephanie Anne Murray[*]
University of Washington Bothell
Computer Science & Software Engineering
Bothell, WA, USA
stephanie@stephaniemurray.org

Erika Parsons
University of Washington Bothell
Computer Science & Software Engineering
Bothell, WA, USA
efuente@uw.edu

Abstract

Electroencephalography (EEG) offers a noninvasive way to monitor brain activity, however, signals are weak, variable, and often masked by noise. Classical machine-learning pipelines such as Random Forests with Principal Component Analysis (PCA) + Common Spatial Patterns (CSP) features can classify movement-related EEG, yet they sometimes fail to capture subtle cross-channel structure. This poster investigates whether a ten-qubit Variational Quantum Classifier (VQC) can serve as a viable alternative. Using a curated subset of the PhysioNet Motor Movement/Imagery dataset, with features reduced to 10 dimensions, we trained and evaluated both a tuned Random Forest (RF) and a VQC implemented in Qiskit. Across 40 simulated runs, the VQC achieved a best macro-F1 of 0.95 (compared to 0.70 for RF), and significantly higher precision and recall for the movement class ($p < 0.001$). While the VQC showed higher run-to-run variance, it consistently reached stronger top-end performance. Feasibility of hardware execution was checked in limited IBM runs. All main results are from Qiskit simulations.

Keywords

EEG, quantum machine learning, variational quantum circuits, brain–computer interface, rehabilitation, Qiskit

ACM Reference Format:
Stephanie Anne Murray and Erika Parsons. 2025. Advancing EEG Signal Analysis with Quantum Machine Learning. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC '25 Posters).* ACM, New York, NY, USA, 2 pages.

Introduction

Brain–computer interfaces (BCIs) enable direct translation of neural signals into device control, with applications in rehabilitation and assistive technologies. Electroencephalography (EEG) is particularly attractive because of its millisecond resolution and patient's non-invasiveness, but movement-related signals are low in amplitude and easily contaminated by artifacts. Classical pipelines rely on preprocessing, feature extraction, and classifiers such as RFs or Support Vector Machines (SVMs), but they may miss cross-channel dependencies. Quantum machine learning (QML) offers a fundamentally different approach by embedding features in high-dimensional Hilbert spaces. We evaluate a compact ten-qubit Variational Quantum Classifier (VQC) for binary classification of movement versus rest, asking whether such a model can compete with, or exceed, a tuned RF baseline.

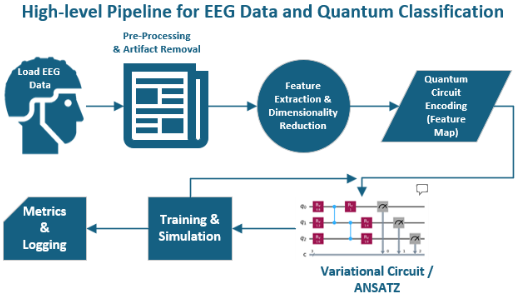

Figure 1: High-level overview of EEG data processing pipeline with VQC

Data and Preprocessing

We use the PhysioNet EEG Motor Movement/Imagery dataset [2][4]. A curated subset was filtered and segmented into 1.88 s windows. Eight motor/central channels (FC5, FC6, CP3, Fz, C1, Cz, C3, C4) were used. Each window was summarized by seven time-domain features per channel (56 total). Features were z-scored, then reduced by PCA (6 components) followed by CSP (4 components), resulting in a compact 10-dimensional vector aligned with the ten-qubit model.

Models

Baseline: A tuned RF with 300 trees and balanced class weights, using feature subsampling at each split. **Quantum:** A VQC implemented in Qiskit [3], using Angle (RY) encoding with a RealAmplitudes ansatz (2 repetitions, linear entanglement, ~40 parameters). Training uses COBYLA with warm starts. Both models use an 80/20 stratified split; SMOTE balances the training set. Each configuration is repeated 40 times to assess variance.

Evaluation Protocol

Metrics: Accuracy, macro-F1 (primary), AUROC, and precision/recall on the movement class. Paired tests use Shapiro–Wilk for normality, then paired t-test or Wilcoxon with Bonferroni correction. All main results are from Qiskit Aer (state-vector, NVIDIA A100 GPU). A

---

[*] Current affiliation: University of Hawaiʻi at Mānoa. Email: stephanie.murray@hawaii.edu.

SC '25 Posters, St. Louis, MO, USA

subset of runs was executed on IBM superconducting hardware to confirm feasibility.

Results

The best VQC configuration (Angle + RealAmplitudes + COBYLA) achieved macro-F1 = 0.95, AUROC = 0.83, and accuracy ≈0.76. The RF baseline reached macro-F1 ≈0.70 and AUROC ≈0.75. Across 40 runs, the VQC averaged macro-F1 ≈0.75 and AUROC ≈0.80, compared to RF at 0.70 and 0.76. Gains in macro-F1 and AUROC were modest overall, but precision and recall for the movement class improved significantly ($p < 0.001$, very large effect sizes). The VQC exhibited higher run-to-run variance but consistently reached stronger best-case performance.

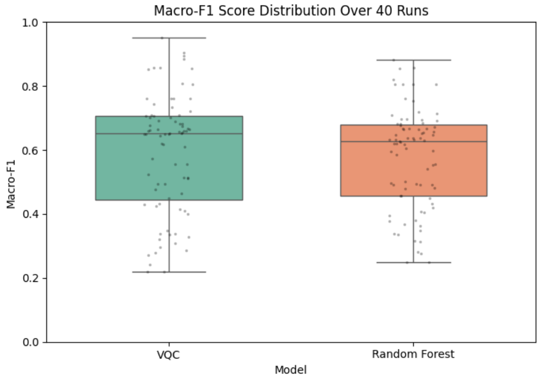

Figure 2: Macro-F1 distribution across 40 runs. VQC shows higher variance but stronger best-case performance than Random Forest.

Relative to Random Forest, VQC improved movement-class precision by +17 points (0.71 vs 0.54) and recall by +19 points (0.61 vs 0.42), $p < 0.001$; see Fig. 2.

Design Insights

Angle encoding outperformed Z, ZZ, and Pauli maps in both stability and accuracy [1]. COBYLA provided the most stable convergence; SPSA was highly variable; gradient descent plateaued early. The circuit remained shallow (∼30 CNOTs), making hardware execution feasible, as confirmed in limited IBM Quantum runs.

Conclusions

With PCA+CSP preprocessing, a ten-qubit VQC can match or exceed a tuned Random Forest on movement-related EEG. The quantum model provides stronger detection of true movement events, an important property for rehabilitation-oriented BCIs.

Limitations and Future Work

All main results are from simulations, with limited hardware runs. Current analysis is within-subject and binary. Future directions include cross-subject evaluation, frequency-domain or hybrid features, adaptive optimizers to reduce variance, larger-scale hardware testing, and exploration of real-time feasibility.

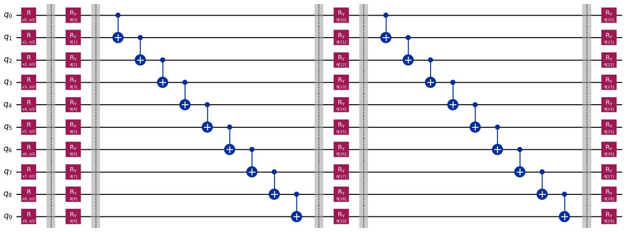

Figure 3: Ten-qubit Variational Quantum Classifier (VQC): Angle (RY) feature map + RealAmplitudes ansatz (2 reps, linear entanglement). Implemented in Qiskit; measured observable is Pauli-Z parity.

Reproducibility Notes

All code was written in Python 3.11. Classical baselines used scikit-learn, and quantum models were implemented in Qiskit. Experiments were run on an NVIDIA A100 GPU with fixed random seeds for splits and initialization, and each configuration was repeated 40 times. Scripts and configurations are available on request.

Table 1: Key hyperparameters for baselines and quantum model.

Random Forest	VQC (Qiskit)
n_estimators = 300	qubits = 10
max_features = sqrt(d)	feature map = Angle (RY)
class_weight = balanced	ansatz = RealAmplitudes, reps = 2
split = 80/20 stratified	entanglement = linear
SMOTE on train fold	optimizer = COBYLA, maxiter = 90
runs = 40	runs = 40 (warm start)

Acknowledgments

We thank Dr. Pierre Mourad, Dr. Michael Stiber, and Dr. Wooyoung Kim for their valuable input and feedback on this thesis work, which was conducted as part of the Master's program in Computer Science & Software Engineering at the University of Washington Bothell.

References

[1] Marcello Benedetti, Ethan Lloyd, Stefan Sack, and Matteo Fiorentini. 2019. Parameterized Quantum Circuits as Machine-Learning Models. *Quantum Science and Technology* 4, 4 (2019), 043001.
[2] Ary L. Goldberger, Luis A. N. Amaral, Leon Glass, and et al. 2000. PhysioBank, PhysioToolkit, and PhysioNet: Components of a New Research Resource for Complex Physiologic Signals. *Circulation* 101, 23 (2000), e215–e220.
[3] Qiskit Development Team. 2021. Qiskit: An Open-source Framework for Quantum Computing. https://qiskit.org/.
[4] Zaid Shuqfa, Abderrahmane Lakas, and Abdelkader N. Belkacem. 2024. Increasing Accessibility to a Large BCI Dataset: Curation of PhysioNet EEG Motor Movement/Imagery Dataset. *Data in Brief* 54 (2024), 110181.