The Princeton Journal of Interdisciplinary Research
ISSN 3069-8200

Empirical Evaluation of Quantum Support Vector Machines with Shallow Feature Maps on Classical Benchmark Datasets

Nathan Liang (author) [1]

[1]Cambridge Centre for International Research, United Kingdom

Abstract

Quantum machine learning seeks to integrate quantum-mechanical principles with statistical inference to construct models capable of representing complex data beyond the limits of classical feature mappings. This study investigates whether quantum support vector machines (QSVMs) provide measurable performance or representational advantages over established classical classifiers on standard biomedical benchmark datasets. Quantum and classical models were trained and tested under equivalent preprocessing and validation protocols to ensure methodological parity in the controlled evaluation. QSVMs are implemented using parameterized quantum feature maps and evaluated through kernel-based classification. A support vector classifier (SVC) was trained with a precomputed kernel, with kernel entries computed using a statevector simulator to isolate representational capacity from sampling and hardware noise. Across datasets, optimized classical models consistently achieve higher predictive accuracy than their quantum counterparts, although QSVMs perform competitively in lower-dimensional regimes. Qualitative analysis of the resulting kernel matrices reveals structured off-diagonal correlations indicative of higher-order feature interactions introduced by quantum entanglement, suggesting that QSVMs capture non-separable relationships that are not directly accessible to standard Gaussian kernels. Despite this richer representational structure, current simulated QSVM implementations do not surpass tuned classical baselines in overall classification accuracy. These findings indicate that the primary near-term value of QSVMs lies in their distinctive feature-space geometry rather than in immediate performance gains, and that QSVMs implicitly learn a richer set of feature interactions from circuit topology rather than through explicit hyperparameter optimization. This highlights the importance of improved encoding strategies and hardware-aware optimization for realizing practical quantum advantage in supervised learning.

Keywords: quantum, kernels, support vector machines

Introduction

Quantum computing's new computational paradigms provide opportunities to speed up machine learning tasks (Liu, Arunachalam, and Temme, 2021). In particular, quantum kernel methods have gained attention as a way to harness quantum computers for classification problems by embedding data into high-dimensional quantum Hilbert spaces and using a classical algorithm (e.g., SVMs) to find decision boundaries (Havlíček et al., 2019). An SVM's power comes from the kernel, which implicitly maps data into a nonlinear feature space. Quantum computing offers a means of representing feature spaces of extremely high dimension. Havlíček et al. (2019) demonstrated the first QSVM on a superconducting quantum processor, showing that quantum-enhanced feature spaces may outperform classical kernels. To date, no practically relevant quantum advantage has been demonstrated for QSVMs on real-world datasets, although theoretical and contrived-problem separations do exist.

Kernel methods are a cornerstone of classical machine learning, with SVMs being a prime example (Scholkopf and Smola, n.d.). A rich body of theory exists on different kernel functions (e.g., Gaussian RBF, polynomial) and their effectiveness across various data distributions. Schuld (2019) popularized the idea of using quantum computing for kernel evaluation, noting that a quantum computer can implicitly implement feature maps by preparing quantum states. In effect, quantum models can be viewed as kernel machines operating in Hilbert space. Havlíček et al. showed that a quantum circuit could estimate inner products of quantum states, yielding a kernel value that can be fed into classical algorithms.

Havlíček et al. (2019) provided the first proof-of-concept of a QSVM, using a shallow quantum feature map (entangling Z rotations) on a superconducting qubit device for a binary classification task. Following that, several works have explored quantum kernels. For instance, Peters et al. (2021) and Hubregtsen et al. (2021) investigated kernel alignment and the expressive power of various quantum feature maps on small datasets. Their findings indicate that while quantum kernels can fit training data well, their advantage is highly dependent on choosing a feature map aligned with the target function.

On the theoretical front, Liu et al. (2021) proved that quantum kernels can, in principle, lead to exponential speed-ups in learning for specific, contrived classification problems (e.g., based on discrete log problems) and thereby separate complexity classes. Prior work by Rebentrost et al. (2014) proposed a QSVM using amplitude encoding and HHL-style matrix inversion, though assumptions were restrictive and not yet hardware practical. Jäger and Krems (2023) further showed that a sufficiently expressive quantum kernel (or a variational quantum classifier) can represent very complex decision boundaries (even BQP-complete classification problems), highlighting that quantum models have a broad expressive hypothesis class. These results, however, are asymptotic and do not immediately translate to practical superiority on typical datasets. Furthermore, their results relate to expressivity and do not imply trainability on noisy devices.

Empirically, most demonstrations of QSVM have been conducted on small scales so far. For example, Blank et al. (2020) compared a variational quantum classifier and QSVM on a few small UCI datasets, finding no clear quantum advantage but noting QSVM's training is deterministic once the kernel is computed, unlike variational models, which require iterative training. On the more applied side, Blance and Spannowsky (2021) applied a QSVM to LHC data for particle classification, and Suzuki et al. (2024) ran QSVM on a trapped-ion hardware for fraud detection and image data, finding that the QSVM matched classical SVM performance on small images but with noise impacts on hardware. These studies highlight that QSVMs often perform as well as, but not necessarily better than, classical SVMs.

This paper is a benchmarking study comparing classical and quantum SVM. Classical SVMs (linear and RBF kernels) and other classical models will be evaluated against QSVMs using two types of quantum feature maps (ZZFeatureMap and PauliFeatureMap) with varying circuit depths. This goes beyond prior comparisons by including multiple classical baselines and quantum feature map variants. The algorithms will be compared based on performance on two datasets: Iris (small, balanced, multi-class) and Breast Cancer (larger, binary, higher feature count). For the latter, PCA was used to reduce the 30-dimensional feature space to accommodate the quantum kernel. Focusing on small tabular datasets provides an accessible yet rigorous baseline. This work is intended to serve as a stepping stone towards more complex experiments designed to gauge the state of quantum kernels in comparison to classical methods.

Method

Datasets and Preprocessing

The Iris dataset contains 150 samples, evenly split across three classes: Setosa, Versicolor, and Virginica. It includes four features: Sepal Length, Sepal Width, Petal Length, and Petal Width. Classes are roughly linearly separable in certain 2D projections, so Iris is often used as a small starting dataset for classification algorithms. It is small enough that full kernel matrices (classical or quantum) can be computed quickly, enabling clear timing comparisons. Breast cancer is a binary classification dataset (malignant vs. benign tumors) with 569 samples and 30 continuous features (cell nucleus measurements). Angle/feature encoding typically uses one qubit per encoded feature, so this dataset highlights the dimensionality challenge for quantum encodings and motivates careful choices of the maximum number of qubits and circuit depth. PCA was used to reduce 30 features to 6 principal components; $n = 6$ qubits was chosen as a compromise between simulation time and preserving variance. Classical models were trained and tested on the PCA-reduced features dataset to maintain parity with quantum experiments, as well as the full dataset for additional comparison. For each dataset, a single stratified train/test split is created with 70% of the data used for training and 30% held out for final evaluation. Random seed is fixed at 42, where stratification ensures class balance is preserved in both splits. The 30% test split is never used during hyperparameter tuning, preprocessing, or model selection.

Classical baselines

As classical baselines, the following models are trained and tuned: linear SVM, RBF SVM, logistic regression, decision tree, and k-nearest neighbors (K-NN). Each model was implemented using scikit-learn. For each model, hyperparameter tuning was performed using GridSearchCV with stratified K-fold cross-validation on the training data only. The best configuration from this tuning is then refit on the full training split and evaluated once on the held-out test split. Each classical model is wrapped in a scikit-learn pipeline that includes feature scaling (via StandardScaler) and, where specified, principal component analysis (PCA). This entire pipeline (scaler, optional PCA, and classifier) is passed directly to GridSearchCV, ensuring that scaling and PCA are always fit only on the training folds within cross-validation and never on the held-out test data. This design prevents any train/test leakage; therefore, no information from the 30% test split influences preprocessing, PCA components, or hyperparameter choices.

Quantum Kernel and QSVM

A quantum kernel matrix is created by computing the pairwise state fidelities between embedded states. For two samples x_i, x_j, a quantum kernel is defined as:

$k(x_i, x_j) = |\langle\phi(x_i)|\phi(x_j)\rangle|^2$ where ϕ is a data-encoding unitary feature-map circuit which translates classical feature vectors into rotations of qubits, where the choice of feature map defines the shape of the quantum kernel. In this experiment, two kernel matrices are built: $K_{train} \in \mathbb{R}^{n_{tr} \times n_{tr}}$, and cross-kernel $K_{test} \in \mathbb{R}^{n_{te} \times n_{tr}}$. Each kernel is normalized so every self-similarity term has unit magnitude and then centered in feature space using an HKH transformation. After centering, a ridge term is added to improve numerical stability.

The QSVM model was trained as a standard scikit-learn SVC with the precomputed K_{train} kernel. As with the classical models, hyperparameter tuning was performed using GridSearchCV on the training portion only. Two feature maps were evaluated:

1. ZZFeatureMap with pairwise entangling $Z_i Z_j$ interactions with depths of $p \in \{1, 2, 3\}$
2. PauliFeatureMap uses rotations generated by X, Z, ZX with Pauli terms matched in depth to ZZ

In quantum feature-map constructions, classical scalars are encoded as parameters of single-qubit rotation gates. Features were scaled to a range of $[0, \pi]$ using only the training data to avoid leakage and reduce periodic aliasing from angle encoding. Min-max scaling computed on the training split only; test values are clipped to $[0, \pi]$. Calibrated class probabilities can be obtained (via Platt scaling) to compute probabilistic quality metrics using the SVM. The statevector simulator was used for quantum evaluation, providing exact noise-free overlaps without finite-sampling effects. This yields analytically exact kernel values for our circuit sizes and avoids stochastic variance associated with sampling-based execution. The configuration, therefore, reflects the idealized quantum feature maps under perfect unitary evolution, suitable for benchmarking against classical baselines and isolating the representational properties of the feature maps without conflating shot noise or hardware noise.

Model evaluation

The models' performance was evaluated with accuracy and F1 score. For Iris (multi-class), overall accuracy was used. This is also the primary metric for model selection in CV. F1 was computed for each class and average, giving equal weight to each class. For Iris (3 classes), macro-F1 penalizes if a model does poorly on one of the species, even if overall accuracy is high. For breast cancer (binary), AUROC and PR-AUC were reported as primary discriminatory metrics, alongside F1 (positive class), and macro-F1. Where relevant, reporting at fixed precision and precision at fixed recall is performed to characterise operational trade-offs. In addition, the probabilistic quality of each model is assessed via log loss (cross-entropy) and the Brier score. All metrics are computed on the held-out test split with stratified sampling. The following efficiency metrics were also measured: training time (cross-validation), inference time, kernel build time/kernel entries computed per second (QSVM only).

Results

Iris

When testing on the Iris dataset, decision tree achieved the highest accuracy of 97.78% and macro-F_1 of 97.78%, while QSVM with Pauli and ZZ feature maps both achieved accuracies of 88.89% and macro-F_1 scores of 88.88% at a circuit depth of 1. Decision tree had a log loss of 0.0696 and a Brier score of 0.0150, while QSVM had a log loss of 0.287 and a Brier Score of 0.0513 (Figures 1 & 2).

Breast Cancer

When testing on the Breast Cancer dataset, logistic regression with 6-component PCA achieved the highest accuracy of 95.91% and macro-F_1 of 95.56%, while QSVM with Pauli and ZZ feature maps both achieved accuracies of 90.06% and macro-F_1 scores of 89.35% at a circuit depth of 1. Logistic regression had a log loss of 0.108 and a Brier score of 0.0292, while QSVM had a log loss of 0.222 and a Brier Score of 0.0716. Additionally, logistic regression achieved an accuracy of 98.25% when trained on the full dataset (Figures 1 & 2).

Figure 1

Comparative performance of classical models and quantum kernel SVMs

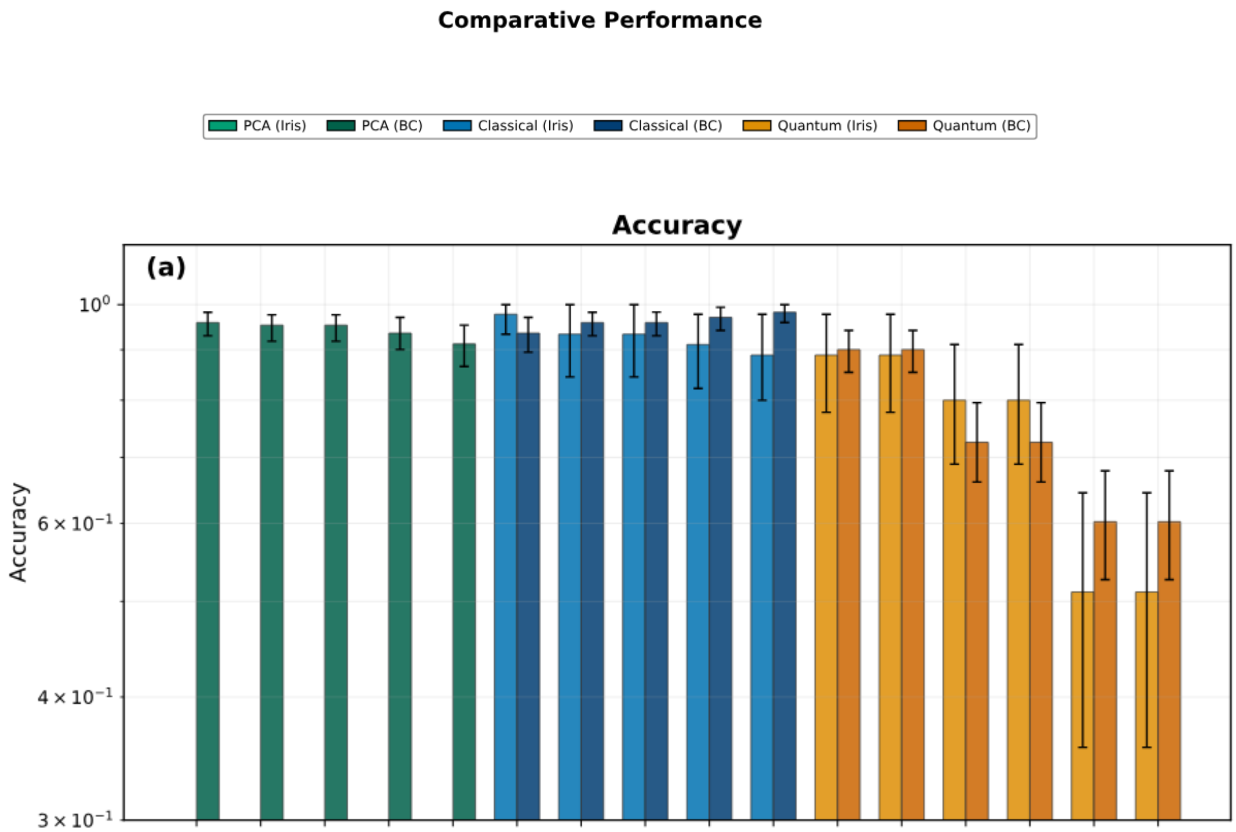

The Princeton Journal of Interdisciplinary Research
ISSN 3069-8200

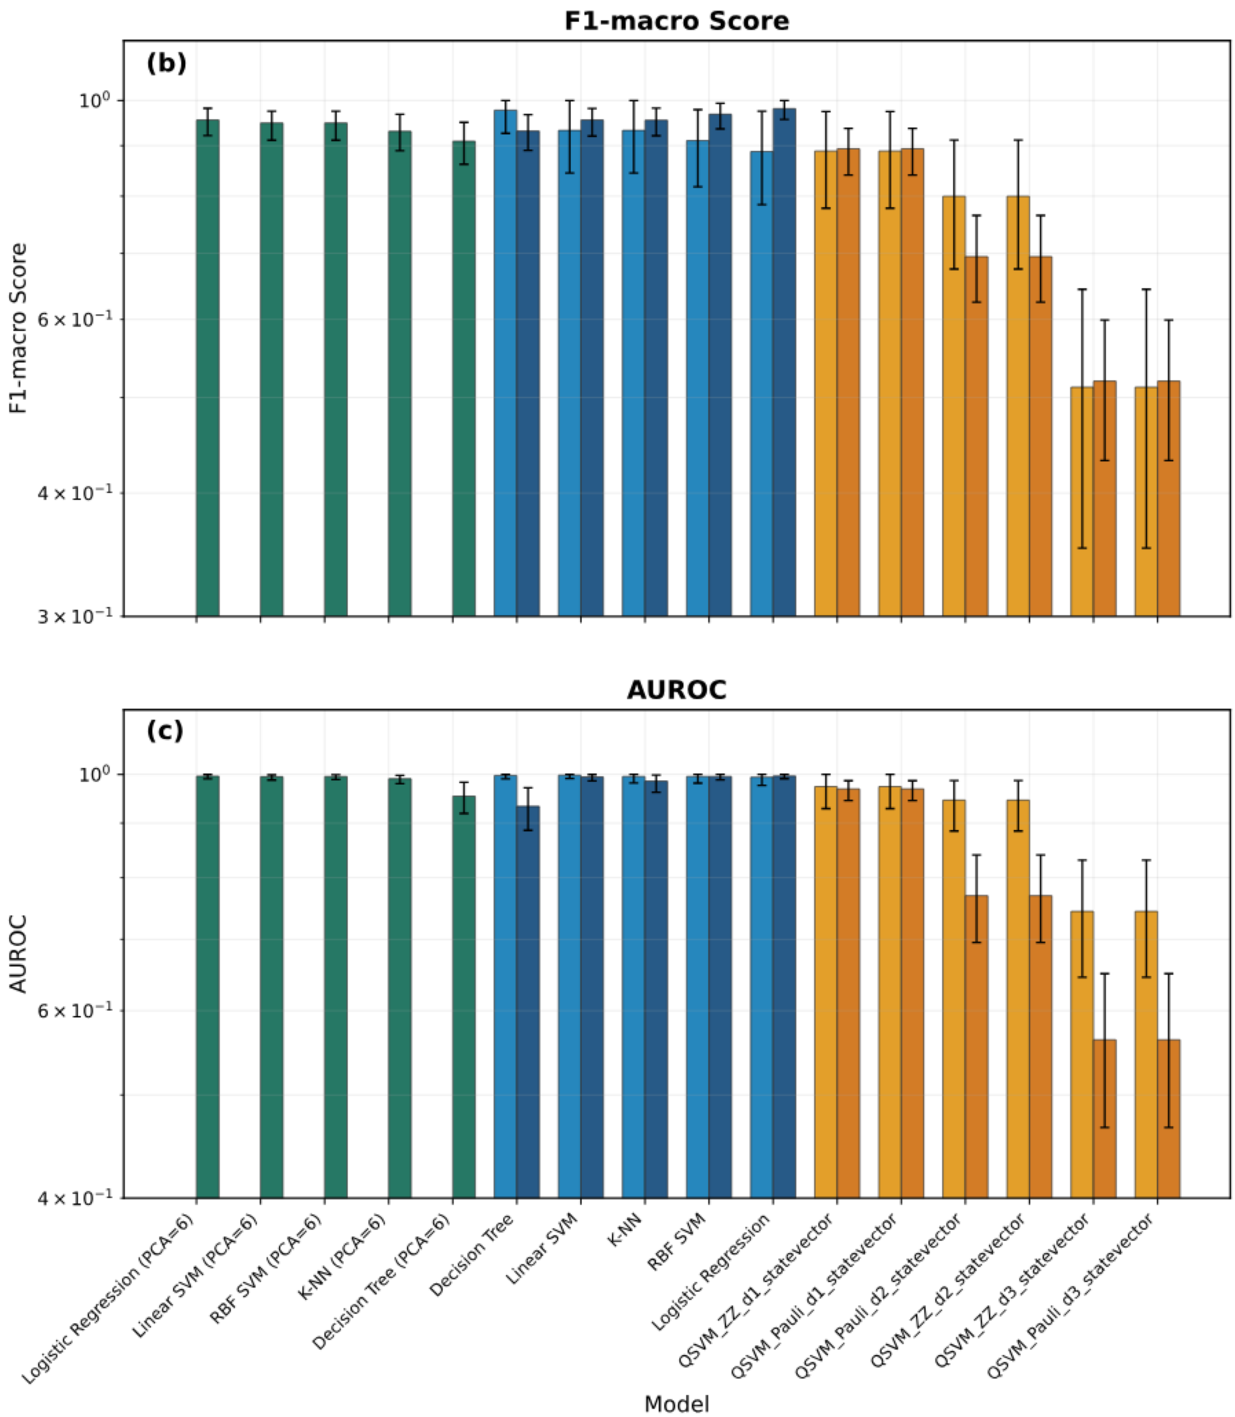

Note. (a) Test accuracy for each model with error bars showing 95% bootstrap confidence intervals. Results are grouped by dataset and by model family (PCA baselines, classical models,

The Princeton Journal of Interdisciplinary Research
ISSN 3069-8200

and QSVMs at depths 1-3 for both ZZ and Pauli feature maps). (b) Macro-F1 scores for the same set of models and configurations, illustrating class-balance sensitivity. (c) AUROC values showing discriminatory quality of each classifier.

Figure 2

Discrimination and probability calibration for classical models and QSVMs

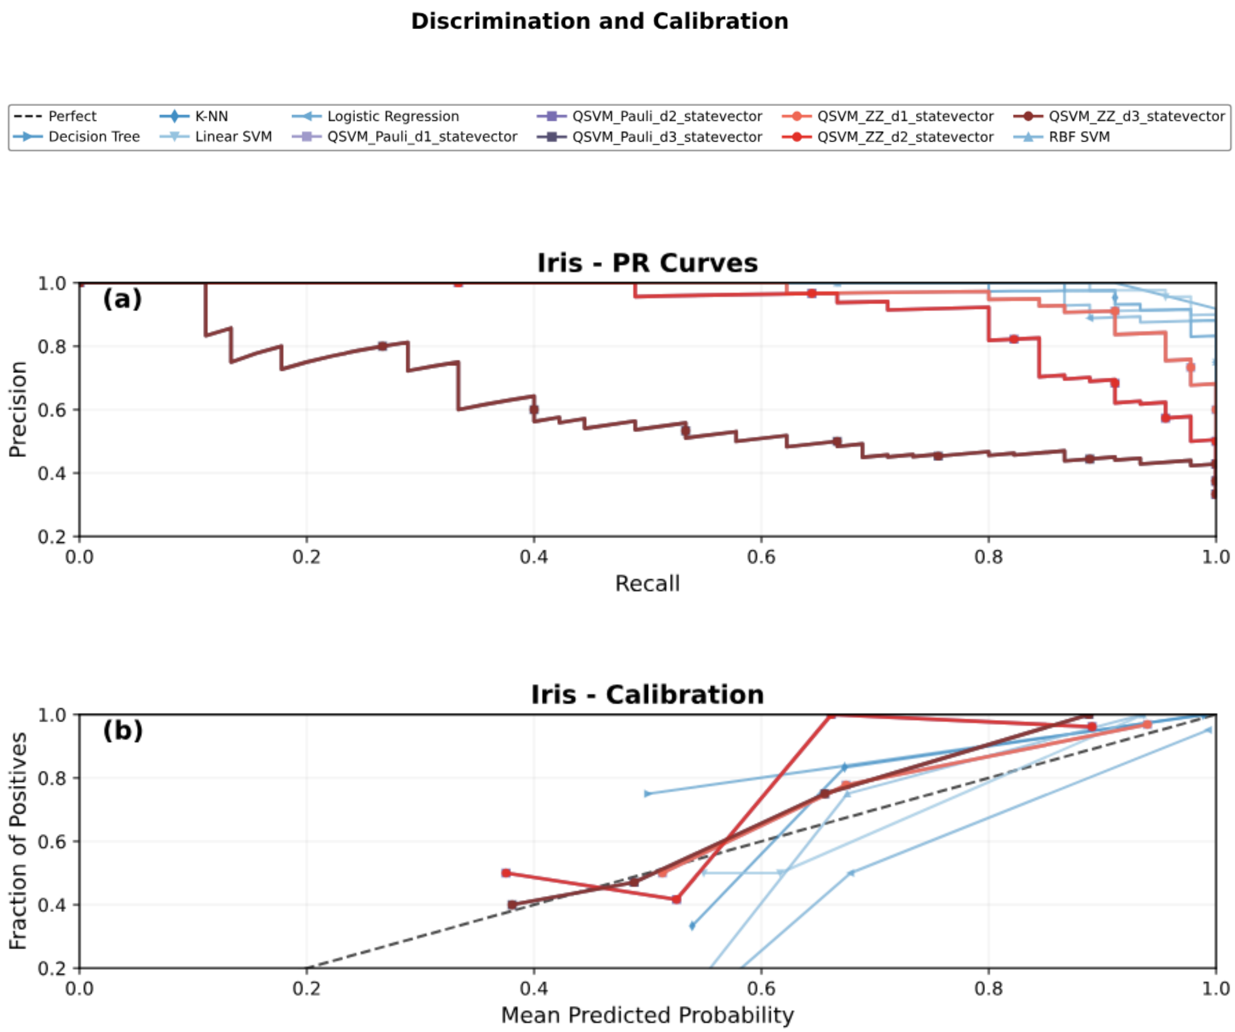

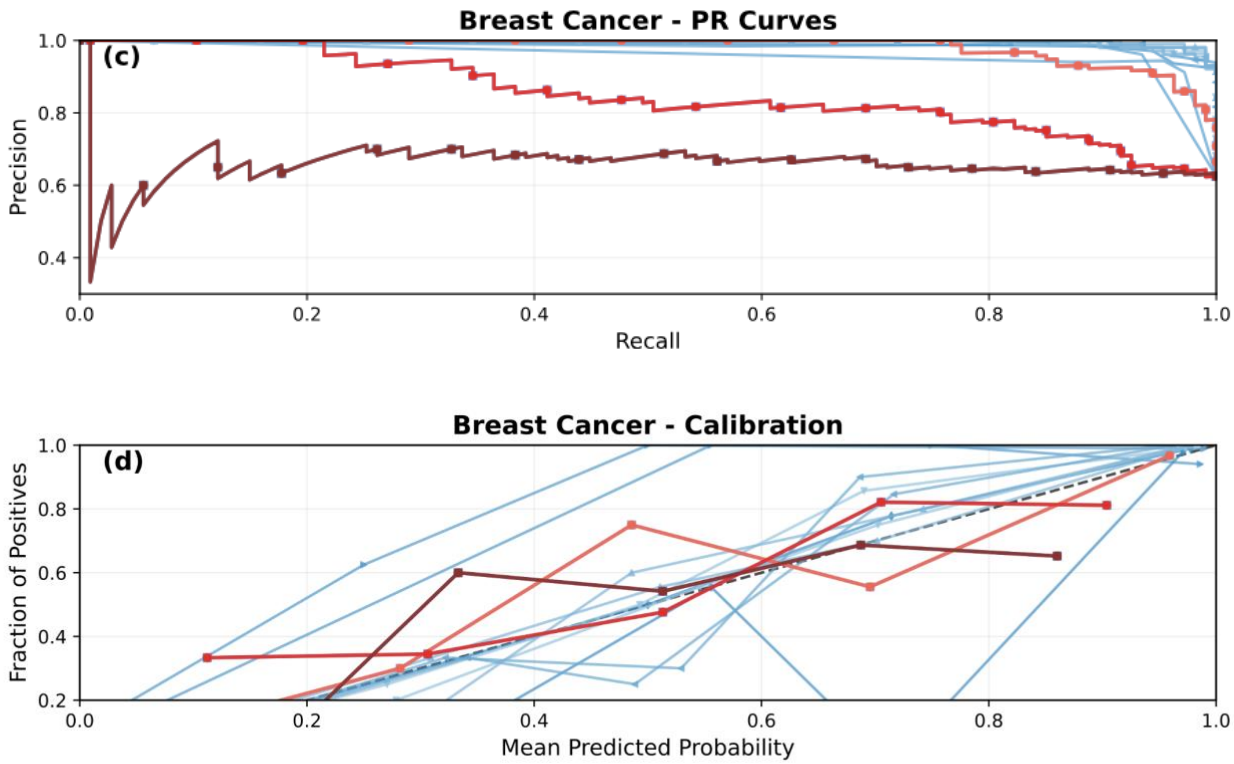

Note. (a) Precision-recall curves for the Iris dataset; comparison classical baselines with QSVMs at multiple circuit depths (b) Calibration curves (fraction of positives vs. mean predicted probability) for the Iris dataset, illustrating reliability of predicted probabilities relative to the perfect-calibration diagonal (c) Precision-recall curves for the Breast Cancer dataset (including PCA-reduced classical models) (d) Calibration curves for Breast Cancer.

Circuit depth

The effect of circuit depth on QSVM performance is consistent across both datasets and both feature-map families (ZZFeatureMap and PauliFeatureMap), revealing a clear inverse relationship between expressivity and generalization. On the Iris dataset, depth-1 configurations (both ZZ and Pauli) achieve 88.89% accuracy and 0.974 macro-AUROC. Performance declines

The Princeton Journal of Interdisciplinary Research
ISSN 3069-8200

sharply with additional layers: depth 2 yields 80.00% accuracy, and depth 3 collapses to 51.11 % - a value statistically indistinguishable from random guessing (i.e. 33.3 % expected for balanced three-class data) where this drop is highly significant (McNemar $p < 0.001$ vs. best classical baseline after Bonferroni correction). The Breast Cancer dataset exhibits the same systematic deterioration. Depth-1 QSVMs reach 90.06 % accuracy and 0.969 - 0.982 macro-AUROC. At depth 2, accuracy falls to 72.51%-77.8%, and at depth 3 it reduces further to only 60.23% ± 7.61% (AUROC 56.34 in the worst case); again, significantly worse than all classical models (McNemar $p < 0.001$).

Table 1

Performance Comparison of Classical and Quantum Support Vector Machines

Dataset	Model	Accuracy (%)	Macro-F1 (%)	AUROC	PR-AUC	Log Loss	Brier	Time (s)	Sig.
Iris	**Decision Tree**	**97.78 ± 3.33×10^{-2}**	**97.78 ± 3.69×10^{-2}**	**0.9978 ± 4.6×10^{-3}**	0.9921	0.070	0.015	0.14	—
Iris	Linear SVM	93.33 ± 7.78×10^{-2}	93.27 ± 7.78×10^{-2}	**0.9985 ± 4.4×10^{-3}**	**0.9972**	0.146	0.023	3.24	—
Iris	QSVM (d=1)	88.89 ± 1.00×10^{-1}	88.88 ± 9.85×10^{-2}	0.9741 ± 3.56×10^{-2}	0.9525	0.287	0.051	0.39	ns

Iris	QSVM (d=2)	80.00 ± 1.11×10^{-1}	79.98 ± 1.19×10^{-1}	0.9459 ± 5.10×10^{-2}	0.9182	0.485	0.087	0.51	ns
Iris	QSVM (d=3)	51.11 ± 1.44×10^{-1}	51.23 ± 1.46×10^{-1}	0.7437 ± 9.30×10^{-2}	0.6385	0.886	0.181	0.70	***
Breast	**Logistic Regression**	**98.25 ± 2.05×10^{-2}**	**98.12 ± 2.15×10^{-2}**	**0.9966 ± 4.6×10^{-3}**	**0.9980**	0.104	0.027	0.03	—
Breast	RBF SVM	97.08 ± 2.63×10^{-2}	96.85 ± 2.88×10^{-2}	0.9956 ± 5.9×10^{-3}	0.9974	0.088	0.025	0.12	—
Breast	QSVM (d=1)	90.06 ± 4.39×10^{-2}	89.35 ± 4.85×10^{-2}	0.9692 ± 2.06×10^{-2}	0.9824	0.222	0.072	3.03	**
Breast	QSVM (d=2)	72.51 ± 6.73×10^{-2}	69.47 ± 7.00×10^{-2}	0.7694 ± 7.24×10^{-2}	0.8548	0.583	0.194	4.43	***
Breast	QSVM (d=3)	60.23 ± 7.61×10^{-2}	51.95 ± 8.38×10^{-2}	0.5634 ± 9.19×10^{-2}	0.6635	0.690	0.243	5.78	***

Note. Comparison of classical machine learning baselines and QSVM with ZZ feature maps across Iris and Breast Cancer datasets. Results show mean ± 95% bootstrap confidence interval

(1000 iterations). All experimental data, including raw predictions, timings, and kernel matrices, are available in the associated repository (Liang, 2025).

Discussion

QSVMs with depth-1 ZZ or Pauli feature maps achieve respectable test performance - 88.89 % accuracy on Iris and approximately 90.06 % on Breast Cancer - demonstrating that shallow quantum embeddings may produce structured kernels with visible class-dependent interference patterns. However, even at this optimal depth, QSVMs are outperformed by multiple classical baselines across all reported metrics (accuracy, macro-F1, AUROC, log loss, and Brier score). The gap is particularly pronounced on the Beast Cancer dataset, where classical models in the full 30-dimensional space exceed 97.5% accuracy, whilst the best depth-1 QSVM remains below 91%.

Increasing depth from 1 to 3 produces a catastrophic and monotonic degradation in generalization performance on both datasets. Accuracy collapses to random-guessing levels (51.11% on Iris, 60.23 % on Breast Cancer at depth 3), accompanied by near-complete loss of block structure in the centered kernel matrices. This behavior is the empirical signature of the over-expressivity and concentration-of-measure phenomenon repeatedly predicted in literature - as the number of entangling layers grows, transition amplitudes between distinct data points decay exponentially, thereby driving kernel entries towards zero (or δ_{ij} after normalization) and rendering the embedded points effectively orthogonal.

In the case where classical models are restricted to the same reduced dimensionality imposed by qubit count (e.g. 6 principal components for Breast Cancer), they continue to dominate QSVMs by a few percentage points in accuracy and substantially larger margins in probabilistic metrics. This confirms that the quantum feature space - despite its exponential theoretical dimension - does not discover more separable representations than well-tuned classical kernels (e.g., linear or RBF) for these datasets.

Even under ideal statevector simulation, quantum kernel evaluation is up to 100× slower than training and inference with classical models (i.e., kernel build times of several seconds vs. sub-second end-to-end classical baselines). On real NISQ hardware, the disparity would be orders of magnitude larger due to shot noise and latency, further widening the gap. Classical models consistently exhibit lower log-loss and Brier scores, indicating superior probability estimates. QSVM calibration curves display systematic overconfidence at higher depths, consistent with noise domination over kernels rather than signal.

Under the rigorous, leakage-free protocol employed, standard quantum kernel SVMs offer no performance benefit over classical baselines on the evaluated datasets and exhibit markedly worse scaling with circuit expressivity and computational cost.

Conclusion

Although this study finds no advantage for standard QSVM implementation(s) on the evaluated tabular datasets, several well-established strategies - many already commonplace in classical kernel methods - offer clear routes to potentially stronger quantum kernels. These are summarized below and represent immediate, practical directions for future work.

The catastrophic degradation observed beyond depth 1 underscores that unrestricted repetitions of entangling layers rapidly induce over-expressivity and overfitting in the kernel feature space (Thanasilp et al., 2024). Fixing depth to 1, or explicitly optimizing depth as a hyperparameter with early-stopping on a validation kernel-alignment score, is essential for stable generalization. A more principled approach would be to employ model selection via cross-validated structural risk minimization, where circuit depth is treated as a discrete hyperparameter jointly optimized with the SVM regularization parameter C. This contrasts with the current fixed-depth evaluation, enabling adaptive complexity calibration matched to the dataset's intrinsic dimensionality. Alternatively, quantum neural architecture search (QNAS), as proposed by Ostaszewski et al. (2021), may determine more optimal gate topologies and repetition counts via differentiable circuit design.

The current grid search employs identical search spaces parameterized only by the regularization parameter C for classical and quantum, which could potentially disadvantage quantum models since kernel spectra may differ fundamentally from RBF/polynomial ones. Extended grids (e.g., non-linear shrinkage of the kernel matrix, class-weighted SVM, or v-SVM) and structured regularization of the quantum kernel itself routinely yield several gains in classical kernel settings and should be adopted as standard practice (Schnabel and Roth, 2025).

Fixed reduction to six dimensions for the breast cancer dataset imposed by six-qubit circuit limits is ultimately arbitrary and not guided by quantum considerations. Classical PCA retains variance but ignores how data geometry interacts with the quantum kernel. A quantum-aware approach that selects the number of components to maximize validation kernel-target alignment, rather than relying on classically defined variance criteria, is likely to produce kernels better suited to the subsequent linear separation. Wang (2023) proposed a QSVM-based feature-selection method that jointly optimizes accuracy, subset size, and circuit complexity, thereby showing that quantum kernel results depend strongly on which features are encoded and how many. Such findings suggest that feature selection driven by a quantum kernel importance may outperform conventional variance-based PCA in identifying features most suitable for quantum encoding.

Min-max scaling to $[0, \pi]$ may introduce periodic wrapping artifacts, particularly for features with outliers. This effectively compresses feature space information near the distribution boundaries. Scaling to $[-\pi/2, \pi/2]$ or using arcsin-based encodings that respect the geometry of the Bloch sphere has been shown to mitigate overwrapping and improve kernel structure (Schuld, 2021). The current ridge is small which serves as a purely numerical purpose and lacks statistical regularization effect. Cross-validated Tikhonov regularization, spectral truncation, or Frobenius-norm penalties may provide genuine regularization, combat concentration of measure, and improve conditioning of quantum kernels by several percent on noisy benchmarks (Thanasilp et al. 2024; Schnabel et al. 2024).

Kernel target alignment is a computationally cheap, label-aware proxy for expected accuracy. Maximizing KTA via gradient descent over feature-map parameters enables fast architecture selection and reveals depth-1 circuits as near-optimal before concentration destroys alignment (Coelho et al. 2025; Schnabel et al. 2024).

Real-world data are often heavily imbalanced. Quantum-compatible remedies include cost-sensitive SVM weighting, kernel-space synthetic minority over-sampling technique (SMOTE), quantum generative oversampling, and post-hoc threshold moving; improving minority-class recall by a considerable margin whilst preserving entanglement benefits (Kwon et al. 2025; Mohanty et al. 2024).

References

Blance, A. and Spannowsky, M., 2021. Quantum Machine Learning for Particle Physics using a Variational Quantum Classifier. *Journal of High Energy Physics* [Online], 2021(2), p.212. Available from: https://doi.org/10.1007/JHEP02(2021)212.

Blank, C., Park, D.K., Rhee, J.-K.K., and Petruccione, F., 2020. Quantum classifier with tailored quantum kernel. *npj Quantum Information* [Online], 6(1), p.41. Available from: https://doi.org/10.1038/s41534-020-0272-6.

Coelho, R., Kruse, G., Rosskopf, A., 2025. Quantum-Efficient Kernel Target Alignment. https://doi.org/10.48550/arXiv.2502.08225

Havlíček, V., Córcoles, A.D., Temme, K., Harrow, A.W., Kandala, A., Chow, J.M. and Gambetta, J.M., 2019. Supervised learning with quantum-enhanced feature spaces. *Nature* [Online], 567(7747), pp.209–212. Available from: https://doi.org/10.1038/s41586-019-0980-2.

Hubregtsen, T., Pichlmeier, J., Stecher, P. and Bertels, K., 2021. Evaluation of parameterized quantum circuits: on the relation between classification accuracy, expressibility, and entangling capability. *Quantum Machine Intelligence* [Online], 3(1), p.9. Available from: https://doi.org/10.1007/s42484-021-00038-w.

Jäger, J. and Krems, R.V., 2023. Universal expressiveness of variational quantum classifiers and quantum kernels for support vector machines. *Nature Communications* [Online], 14(1), p.576. Available from: https://doi.org/10.1038/s41467-023-36144-5.

Kwon, S., Huh, J., Kwon, S.J., Choi, S., Kwon, O., Kwon, S., Huh, J., Kwon, S.J., Choi, S., Kwon, O., 2025. Leveraging Quantum Machine Learning to Address Class Imbalance: A Novel Approach for Enhanced Predictive Accuracy. Symmetry 17. https://doi.org/10.3390/sym17020186

Liang, N. (2025). Comprehensive experimental results and raw data for "Classical vs Quantum-Kernel SVMs on Tabular Datasets": all_results.csv, breast_cancer_results.csv, iris_results.csv, kernel matrices, predictions, and full source code. GitHub repository, Cambridge Center for International Research (CCIR). https://github.com/swtchblade/CCIR (accessed 23 November 2025).

Liu, Y., Arunachalam, S. and Temme, K., 2021. A rigorous and robust quantum speed-up in supervised machine learning. *Nature Physics* [Online], 17(9), pp.1013–1017. Available from: https://doi.org/10.1038/s41567-021-01287-z.

Mohanty, N., Behera, B.K., Ferrie, C., Dash, P., 2025. A Quantum Approach to Synthetic Minority Oversampling Technique (SMOTE). Quantum Mach. Intell. 7, 38. https://doi.org/10.1007/s42484-025-00248-6

Ostaszewski, M., Trenkwalder, L.M., Masarczyk, W., Scerri, E. and Dunjko, V., 2021. *Reinforcement learning for optimization of variational quantum circuit architectures* [Online]. arXiv. Available from: https://doi.org/10.48550/arXiv.2103.16089 [Accessed 8 November 2025].

Peters, E., Caldeira, J., Ho, A., Leichenauer, S., Mohseni, M., Neven, H., Spentzouris, P., Strain, D. and Perdue, G.N., 2021. Machine learning of high dimensional data on a noisy quantum processor. *npj Quantum Information* [Online], 7(1), p.161. Available from: https://doi.org/10.1038/s41534-021-00498-9.

Rebentrost, P., Mohseni, M. and Lloyd, S., 2014. Quantum support vector machine for big data classification. *Physical Review Letters* [Online], 113(13), p.130503. Available from: https://doi.org/10.1103/PhysRevLett.113.130503.

Schnabel, J., Roth, M., 2025. Quantum Kernel Methods under Scrutiny: A Benchmarking Study. Quantum Mach. Intell. 7, 58. https://doi.org/10.1007/s42484-025-00273-5

Scholkopf, B. and Smola, A., n.d. *Support Vector Machines and Kernel Algorithms*.

Schuld, M., 2019. Quantum Machine Learning in Feature Hilbert Spaces. *Physical Review Letters* [Online], 122(4). Available from: https://doi.org/10.1103/PhysRevLett.122.040504.

Schuld, M., 2021. Supervised quantum machine learning models are kernel methods. https://doi.org/10.48550/arXiv.2101.11020

Suzuki, T., Hasebe, T. and Miyazaki, T., 2024. *Quantum support vector machines for classification and regression on a trapped-ion quantum computer* [Online]. arXiv. Available from: https://doi.org/10.48550/arXiv.2307.02091 [Accessed 14 October 2025].

Thanasilp, S., Wang, S., Cerezo, M., Holmes, Z., 2024. Exponential concentration in quantum kernel methods. https://doi.org/10.48550/arXiv.2208.11060

Wang, H., 2023. *A novel feature selection method based on quantum support vector machine* [Online]. arXiv. Available from: https://doi.org/10.48550/arXiv.2311.17646 [Accessed 8 November 2025].

Acknowledgments

I would like to thank Dr. Sergii Strelchuck and Dr. Sean Harvey for their mentorship, feedback, and guidance throughout the research process. I would also like to thank the CCIR team for supporting me during the research and publication process.