

Bandwidth, Not Circuit Scaling, Drives Quantum-Kernel Recovery in Cross-Quality Face Verification

Rohan Ramakrishna , Ying Zhe George Loh , Bing Hui Terence Goh* 

KLASS Engineering and Solutions Pte. Ltd., Singapore

*Corresponding author: innovation@klases.com.sg

Abstract

Quantum fidelity kernels are often judged by circuit architecture, but their first test is simpler: does the encoding create the right similarity geometry for the task? We answer this question in cross-quality face verification using fixed AdaFace embeddings and the official XQLFW fold structure. A default four-qubit ZZ fidelity kernel at the unscaled input bandwidth is near chance. Calibrating only the input bandwidth moves the same circuit family into a competitive regime, and the recovery persists under nested, training-only bandwidth selection. The effect is therefore a genuine model-selection result, not a post-hoc choice from held-out test scores. We then test two common architectural instincts: add entangling connectivity, or increase qubit/PCA width. In this benchmark, denser ZZ connectivity does not improve performance, and scaling from four to eight qubits produces a plateau rather than a gain. Matched linear and RBF SVMs remain at least as strong as the best quantum-kernel configurations, giving a hard classical reference point. The message is direct: for quantum kernels on frozen classical embeddings, bandwidth is a first-order design choice, while neither more entanglement nor more qubits improves the inductive bias here.

1 Motivation

Quantum kernel methods encode a classical input x into a quantum state $|\phi(x)\rangle$ and pass state overlaps to a classical learner as a kernel [1, 3]. The promise is not merely a larger feature space, but a task-relevant inductive bias that is hard for classical kernels to reproduce. Recent theory and benchmarking make the counterpoint equally clear: a large Hilbert space is not automatically useful. Quantum kernels can concentrate, align poorly with labels, or be matched closely by classical kernels on classical data [4–7, 11].

Input bandwidth is the most direct handle on this failure mode. Scaling the data before circuit binding changes the state-overlap geometry, moving a kernel between underfitting, useful generalization, and concentration regimes [8–10]. We use cross-quality face verification as a controlled diagnostic setting because it forces the kernel to preserve identity-relevant similarity under image-quality mismatch. XQLFW is built for this stress test [12], and AdaFace provides strong quality-aware embeddings [13]. The study asks a narrow but important question: when a fidelity kernel is applied to frozen biometric representations, which design choice actually makes it informative?

2 Protocol

XQLFW contains 6,000 verification pairs arranged into ten official folds. Each image is represented by a fixed 512-dimensional AdaFace embedding. For kernel and SVM experiments, each pair is converted into the symmetric descriptor $[|e_1 - e_2|, e_1 \odot e_2]$. Standardization and PCA are fitted using training data only. The principal experiments use 200 balanced training pairs for each outer fold and subset seed, and evaluate on all 600 held-out pairs in the corresponding official fold. Three independently sampled training subsets are used per fold.

For bandwidth λ , the quantum kernel is

$$K_Q(x, z) = |\langle \phi(\lambda x) | \phi(\lambda z) \rangle|^2.$$

We evaluate a local product-state feature map and a ZZ feature map with exact statevector simulation. The main ZZ configuration is a linear chain (E2); the topology ablation also tests no entanglement (E0), matching connectivity (E1), and full connectivity (E3). Precomputed Gram matrices are supplied to an SVM [14]. Classical references include AdaFace cosine similarity on the original embeddings, and linear and RBF SVMs on matched PCA inputs.

Model selection is fully nested inside each outer fold. For each model, fold, and training seed, the 200 training pairs are split into deterministic stratified five-fold inner validation splits. Standardization and PCA are refitted inside every inner split, and $\lambda \in \{1, 1/2, 1/4, 1/8, 1/16, 1/32, 1/64\}$ is selected using inner-validation ROC-AUC only. The selected model is then retrained on all 200 outer-training pairs and evaluated once on the untouched 600-pair test fold. AUCs are computed separately for each trained model; seed results are averaged within fold, and the ten folds are used for paired bootstrap intervals and sign-permutation tests.

3 Results

Bandwidth calibration rescues an uninformative ZZ kernel. At the unscaled input bandwidth, the four-qubit ZZ linear-chain kernel is essentially uninformative. Reducing the input scale produces a monotone recovery across the main grid, reaching the competitive regime around $\lambda = 1/16$ and peaking diagnostically near $\lambda = 1/32$ (Fig. 1a). On the complete held-out folds, the fixed-grid improvement from $\lambda = 1$ to $\lambda = 1/16$ is +0.419 AUC, with a 95% paired bootstrap interval of $[0.409, 0.429]$. This is the central result: the same circuit family changes from near chance to competitive once its similarity geometry is calibrated.




Figure 1: Main empirical pattern. (a) Bandwidth calibration changes the four-qubit ZZ kernel from near-chance performance at the unscaled encoding to a competitive regime at smaller input scales. (b) Increasing qubit/PCA width from four to eight does not produce a meaningful gain for either quantum model; all observed changes remain within the predefined 0.005 AUC practical-effect threshold.

Training-only selection recovers the same small-bandwidth regime. The nested selector finds the same region without seeing the outer test folds. For the four-qubit ZZ kernel, the extended selector places most choices at $1/16$, $1/32$, and $1/64$, and achieves an outer-test AUC within 0.001 of the best diagnostic fixed-grid value. The exact selected bandwidth varies across training subsets, but the selected region is stable. Table 1 shows the practical consequence: the default ZZ kernel at $\lambda = 1$ is near chance, nested bandwidth selection restores competitive quantum-kernel performance, and matched classical baselines remain at least as strong.

Table 1: Compact summary of complete-fold and nested-selection results. Intervals are fold-bootstrap 95% CIs. The qubit-scaling rows use nested, training-only bandwidth selection.

Method	Input	Mean AUC	95% CI
AdaFace cosine	512D embeddings	0.9225	[0.9149, 0.9300]
Linear SVM	4 PCA dims	0.9207	[0.9130, 0.9282]
RBF-SVM	4 PCA dims	0.9207	[0.9130, 0.9282]
Local, nested	4 qubits / 4 PCA dims	0.9202	[0.9126, 0.9277]
ZZ, nested	4 qubits / 4 PCA dims	0.9179	[0.9113, 0.9246]
ZZ, fixed $\lambda = 1$	4 qubits / 4 PCA dims	0.4999	[0.4924, 0.5079]
ZZ E2, fixed $\lambda = 1/16$	4 qubits / 4 PCA dims	0.9190	[0.9119, 0.9260]
Local, nested	6 qubits / 6 PCA dims	0.9204	[0.9126, 0.9281]
ZZ, nested	6 qubits / 6 PCA dims	0.9182	[0.9098, 0.9266]
Local, nested	8 qubits / 8 PCA dims	0.9198	[0.9120, 0.9276]
ZZ, nested	8 qubits / 8 PCA dims	0.9173	[0.9094, 0.9251]

More connectivity does not explain the recovery. Topology ablation separates the bandwidth effect from entanglement structure. At $\lambda = 1/8$, the non-entangling ZZ backbone has the strongest mean AUC among the tested ZZ topologies, while denser connectivity has lower mean AUC. Relative to E0, E2 is lower by about 0.0055 AUC and E3 by about 0.0133 AUC on the complete folds. At $\lambda = 1/16$, these differences shrink substantially. Added entangling connectivity is therefore not the source of the recovery, consistent with broader evidence that entangling capability alone need not determine classifier quality [15]. The geometry diagnostics support the same conclusion: reducing bandwidth increases centered kernel-target alignment [2] and decreases normalized effective rank of the training Gram matrix, so the recovered kernel is better aligned with the verification labels and less concentrated near the identity.

Increasing to eight qubits produces a plateau. Increasing qubit/PCA width also fails to create a meaningful gain (Fig. 1b and Table 1). Repeating nested bandwidth selection at four, six, and eight qubits gives Local nested AUCs of 0.9202, 0.9204, and 0.9198, and ZZ nested AUCs of 0.9179, 0.9182, and 0.9173. Matched linear and RBF SVMs show the same flat pattern over 4, 6, and 8 PCA dimensions. Under the predefined practical-effect threshold of 0.005 AUC, every qubit-scaling comparison is a plateau. We therefore treat the 4–8 qubit range as a controlled diagnostic regime rather than a scaling claim; importantly, the matched classical SVMs also plateau over the same 4–8 PCA dimensions, suggesting that wider low-dimensional PCA inputs do not reveal a hidden classical performance trend in this protocol.

4 Conclusion

This study gives a direct answer to a common quantum-kernel design question: when a fidelity kernel fails on structured classical data, the first remedy should be geometry calibration, not simply more circuit structure. In cross-quality face verification, bandwidth calibration changes a four-qubit ZZ kernel from near chance to competitive performance, and the effect survives nested training-only selection. Denser ZZ connectivity and increasing exact-simulation width from four to eight qubits do not produce meaningful gains. Matched classical baselines remain at least as strong, which sharpens the conclusion rather than weakening it: feature-scale geometry is the decisive design variable in this benchmark.

Rather than serving as a quantum-advantage or hardware-validation claim, these results provide a controlled diagnostic benchmark for quantum kernels on frozen classical embeddings. All quantum kernels are evaluated by exact statevector simulation on a single face-verification benchmark. The contribution is a reproducible diagnostic result for credible quantum-kernel benchmarking: calibrate the induced similarity geometry before attributing performance to entanglement, circuit width, or quantum inductive bias. Operationally, this gives a practical protocol: before increasing circuit depth, entangling connectivity, width, or QPU shot budget, quantum-kernel studies should first check whether feature scale is distorting the induced similarity geometry.

References

- [1] Maria Schuld and Nathan Killoran. Quantum machine learning in feature Hilbert spaces. *Physical Review Letters*, 122(4):040504, 2019. doi:10.1103/PhysRevLett.122.040504.
- [2] Nello Cristianini, John Shawe-Taylor, André Elisseeff, and Jaz S. Kandola. On kernel-target alignment. In *Advances in Neural Information Processing Systems*, volume 14, pages 367–373, 2002.
- [3] Vojtěch Havlíček, Antonio D. Córcoles, Kristan Temme, Aram W. Harrow, Abhinav Kandala, Jerry M. Chow, and Jay M. Gambetta. Supervised learning with quantum-enhanced feature spaces. *Nature*, 567(7747):209–212, 2019. doi:10.1038/s41586-019-0980-2.
- [4] Hsin-Yuan Huang, Michael Broughton, Masoud Mohseni, Ryan Babbush, Sergio Boixo, Hartmut Neven, and Jarrod R. McClean. Power of data in quantum machine learning. *Nature Communications*, 12:2631, 2021. doi:10.1038/s41467-021-22539-9.
- [5] Jonas M. Kübler, Simon Buchholz, and Bernhard Schölkopf. The inductive bias of quantum kernels. In *Advances in Neural Information Processing Systems*, volume 34, pages 12661–12673, 2021.
- [6] Lucas Slattey, Ruslan Shaydulin, Shouvanik Chakrabarti, Marco Pistoia, Sami Khairy, and Stefan M. Wild. Numerical evidence against advantage with quantum fidelity kernels on classical data. *Physical Review A*, 107(6):062417, 2023. doi:10.1103/PhysRevA.107.062417.
- [7] Supanut Thanasilp, Samson Wang, M. Cerezo, and Zoë Holmes. Exponential concentration in quantum kernel methods. *Nature Communications*, 15:5200, 2024. doi:10.1038/s41467-024-49287-w.
- [8] Ruslan Shaydulin and Stefan M. Wild. Importance of kernel bandwidth in quantum machine learning. *Physical Review A*, 106(4):042407, 2022. doi:10.1103/PhysRevA.106.042407.
- [9] Abdulkadir Canatar, Evan Peters, Cengiz Pehlevan, Stefan M. Wild, and Ruslan Shaydulin. Bandwidth enables generalization in quantum kernel models. *Transactions on Machine Learning Research*, 2023. URL: <https://openreview.net/forum?id=A1N2qp4yAq>.
- [10] Sebastian Egginger, Alona Sakhnenko, and Jeanette Miriam Lorenz. A hyperparameter study for quantum kernel methods. *Quantum Machine Intelligence*, 6:44, 2024. doi:10.1007/s42484-024-00172-1.
- [11] Jan Schnabel and Marco Roth. Quantum kernel methods under scrutiny: A benchmarking study. *Quantum Machine Intelligence*, 7:58, 2025. doi:10.1007/s42484-025-00273-5.
- [12] Martin Knoche, Stefan Hörmann, and Gerhard Rigoll. Cross-quality LFW: A database for analyzing cross-resolution image face recognition in unconstrained environments. In *2021 16th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021)*, pages 1–5, 2021. doi:10.1109/FG52635.2021.9666960.
- [13] Minchul Kim, Anil K. Jain, and Xiaoming Liu. AdaFace: Quality adaptive margin for face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18750–18759, 2022. doi:10.1109/CVPR52688.2022.01819.
- [14] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 20:273–297, 1995. doi:10.1007/BF00994018.
- [15] Thomas Hubregtsen, Josef Pichlmeier, Patrick Stecher, and Koen Bertels. Evaluation of parameterized quantum circuits: On the relation between classification accuracy, expressibility, and entangling capability. *Quantum Machine Intelligence*, 3:9, 2021. doi:10.1007/s42484-021-00038-w.