

Local Fidelities for Enhanced Quantum Kernels

Valeria Repetto^{1,2}, André Manuel Sequeira³, and Romina D’Aurizio¹

¹Institute of Informatics and Telematics (IIT), National Research Council of Italy (CNR), Pisa, Italy

²Computational Biomedicine Unit, Department of Medical Sciences, University of Turin, Torino, Italy

³High-Assurance Software Laboratory (HASLab), INESC TEC, Braga, Portugal

Quantum kernels (QKs) have emerged as a promising method for learning with quantum-enhanced feature spaces. However, they suffer from a fundamental limitation known as the *vanishing similarity* phenomenon, or kernel concentration. The effect originates from the global measurement underlying the kernel: as the number of qubits grows, the variance of the kernel matrix entries decays exponentially, requiring an exponential number of measurements for accurate estimation. At the same time, the resulting Gram matrix progressively converges to the identity matrix, limiting its generalization capability on real-world, high-dimensional datasets. In this work we introduce the Local Fidelity Quantum Kernels (LFQKs), which aim to mitigate kernel concentration by replacing the global fidelity measurement of the quantum circuit with a combination of fidelity measurements over local subsystems.

The fidelity-based QK associated to an N -qubit system can be expressed by decomposing the global all zero state projector in the Pauli-basis as $|0\rangle\langle 0|^{\otimes N} = \frac{1}{2^N} \sum_j P_j$ where $P_j \in \{I, Z\}^N$. This exponentially large decomposition explicitly leads to the concentration observed in QK. Building on this observation, we propose the LFQK, in which we consider a subset of $M \in O(\text{poly}(N))$ local projectors of locality at most $k = O(\log(N))$. The LFQK is then defined as $k(x, x') = \frac{1}{M} \sum_{i=0}^{M-1} \langle \psi(x, x') | |0\rangle\langle 0|_i^{\otimes k} | \psi(x, x') \rangle$ with $|0\rangle\langle 0|_i^{\otimes k} = \frac{1}{2^k} \sum_j P_j$. We show that, assuming the data is uniformly random, the variance of the proposed local kernel decays polynomially with the number of qubits.

We first validate the theoretical predictions by estimating the variance of kernel values over randomly sampled input pairs. We consider IQP feature maps with both linear and all-to-all entanglement, circuit depths ranging from one to four repetition layers, system sizes between 2 and 20 qubits, and two locality configurations, where M corresponds to the number of contiguous or of all possible k -qubit subsystems, for several values of the locality parameter k . Across all tested architectures, the variance of the global kernel decays exponentially as 2^{-N} , whereas LFQKs exhibit polynomial variance decay, confirming the theoretical analysis. Using the same feature maps and circuit configurations, we evaluate the practical implications of this improved scaling: we benchmark LFQKs on synthetic classification tasks whose dimensionality scales with the number of qubits. This setup enables a direct comparison between global and local kernels as concentration becomes increasingly severe. LFQKs outperform the global kernel, with feature-map-dependent accuracy gains that become more pronounced as the number of qubits increases, reaching a maximum of ~ 30 percentage points. Such effect is most marked for the full-entanglement feature map. Additionally, for this feature map, above 18 qubits, LFQK accuracy approaches that of the classical RBF kernel. Finally, replicating these experiments under finite sampling with a quadratic number of measurement shots, the overall trend remains consistent with the ideal-state case. Mitigating kernel concentration does not by itself guarantee a performance improvement, which depends also largely on data and embedding choice. Nonetheless, these findings show that LFQKs remain effective where global kernels fail, leaving room for quantum enhancements at larger qubit numbers or in quantum experiments at scale. A remaining consideration is the classical hardness of the embedding, as an excessively local construction may compromise it, provided that every projector in the subset has a backwards light-cone of $O(\log(N))$. Nevertheless, since the LFQK combines multiple projectors with different light-cones, depending on the embedding and locality structure, some may remain classically hard to estimate. Ongoing work is thus focused on characterizing this trade-off and extending LFQKs to mixtures of localities, investigating whether combining projectors with different localities could enhance performance whilst avoiding classical simulability. LFQKs are also being explored within the framework of trainable QKs, where locality can be jointly optimized with circuit parameters to improve data-dependent expressivity. Finally, we are investigating their effectiveness on high-dimensional multi-omics biomedical datasets, whose complex correlation structure provides a natural benchmark for scalable quantum similarity learning.