Where quantum kernels can win at all

— a **measured map** between the classical wall and the concentration wall. The literature names two reasons quantum machine learning should fail. We measured where they actually bite — and how much winning territory lies **between** them.

996
refereed arms
pre-specified; design locked before the grids

10
classical variants
each threshold-tuned — the strongest bar we could mount

≥94%
wins up to 12 qubits
at 512 shots as at exact simulation

THE SETUP

Two walls — and the question of what lies between.

Wall one, classical learnability: **if ordinary structure explains the labels, a tuned classical model wins** — no quantum method needed. Wall two, kernel concentration: as registers grow, quantum kernel entries shrink toward noise — **at some size, shots can't resolve them.**

Every arm faces ten tuned classical variants — including an **order-matched twin** built to imitate the quantum model classically. A win counts only when the **whole 95% confidence interval** clears parity.

Advantage isn't claimed. It's certified — or it isn't counted.

The territory is real — and larger than the back-of-envelope says.

Wins in ≥94% of arms, up to 12 qubits

On quantum-native labels the kernel classifier beats the full classical bar at every size up to 12 qubits — **at 512 measurement shots just as at exact simulation**. Erosion begins at 14–16 qubits, where the concentration wall bites.

The standard heuristic is ~2 sizes too pessimistic

The per-entry shot-noise rule predicts failure about two register sizes before it happens. We measured why by ablation: **learners aggregate thousands of entries, and the information is redundant** — single-entry pessimism doesn't survive contact with the full pipeline.

The classical wall, located

Zero wins on purely classical label structure — as it should be. Wins begin at **75% quantum-native mixing**. The map has a coastline, and now it has coordinates.

READ THE LABEL

Best case, by construction — and disclosed as such.

The winning labels are **quantum-native by construction** — generated by the same kernel family the classifier uses. This is a **benchmark envelope**: it measures where the method *can* win, not a claim that your dataset lives there.

That is exactly what makes it useful: the map tells you **which measurable properties** a dataset needs before the quantum side deserves your compute budget — and the next study in this series asks the same question for the **trainable** quantum model.

A best-case map, honestly labeled, beats a real-world claim you can't check.

PRINTED, ON PURPOSE

The study refuted our own instruments — twice.

Our checklist: 83 of 84 wrong vetoes

The screening checklist our own product shipped **wrongly vetoed 83 of 84 datasets** on which the kernel then won. We measured it, printed it, and **replaced it with calibrated measurements** that rank outcomes at AUC up to 0.97.

Our favorite explanation: refuted by its own ablation

Our working conjecture for *why* shot noise was survivable failed its dedicated test. **The refutation is in the study** — alongside what the ablation showed instead.

The expanded referee flipped zero verdicts

Adding a tuned RBF kernel SVM and a neural network to the bar — the reviewer's challenge — changed **no verdict**. The wins survive the strongest bar we could mount.

WHAT IT CHANGES

The map ships in the product.

Every QML assessment now prints an **island placement**: where your dataset sits on the measured map, its **kernel-concentration indicator**, and a **shots-to-resolve** estimate — the study's instruments, as standard diagnostics in every report.

No categorical yes/no from a checklist — **calibrated measurements with the study behind them**, so the "is our data quantum-suited" conversation starts from coordinates instead of vibes.

From "maybe, somewhere" to a placement on a measured map.

THE STUDY

The full study, unedited.

The **complete 14-page study** is public — the map, the referee table, the shot-noise ablation, both refutations, and every interval. No summary layer: **the evidence itself.**

The study — 14 pages

qubit-lab.ch/evidence

Study PDF, published benchmark records and the QML Rapid Quantum Prototype are on qubit-lab.ch/evidence. Evidence, not promises.