The trainable quantum model, mapped

— whose territory, at which size, and what a **fair attempt** costs. The companion map to the kernel study: same referee, same rows — plus the two quantum methods **head-to-head**. And this time, the design was **publicly tagged before the first run**.

65

fair-recipe trainings

of a registered ceiling of 150 — disclosed

0.98 → 1.18

the crossover

trainable vs kernel, 4 → 16 qubits, paired rows

1

registered bet — lost

against the folklore; printed with its interval

THE SETUP

Two quantum models. Do we need both?

The kernel method (QSVM) measures distances between quantum states; the **variational classifier (VQC)** trains a circuit end-to-end. Theory says they are cousins — the correspondence result even suggests the kernel might **replace** the trainable model outright.

So: same ten-variant classical referee, same data rows, and the two quantum methods scored **head-to-head in every cell** — across three label families and 4–16 qubits, under a fair, fixed training recipe. Hypotheses registered; chronology **verifiable, not asserted**.

Not a taste question — a measurement.

Territory is size-gated — and the crossover is monotone.

Small registers: the kernel wins everything

At 4 qubits the kernel method dominates all we measured — including a **perfect score on the trainable model's own native labels**.

Then the ratio climbs: 0.98 → 1.18

On its native structure the trainable model pulls ahead **monotonically** across 4–16 qubits — 18% ahead at sixteen — while kernel-native territory stays with the kernel **at every size**.

So: both — but not at every size

Which method a dataset deserves is **measurable**: structure decides, size gates it, and the map now says where. The correspondence holds; **predictive dominance reverses**.

The feared trainability wall did not appear — here.

The literature's vanishing-gradient warning comes with fine print: it targets deep circuits and global cost functions. In the regime the fine print permits — **shallow, local-cost, hardware-efficient** — we measured the gradients: **flat to rising through 16 qubits**, and the model's best score came at the largest size tested.

Our registered prediction of degradation at 14–16 qubits is **refuted on native territory** — and the bet we registered *against* the folklore lost too: initial-gradient variance **does** predict training outcomes (AUC 0.79). Both outcomes are printed.

The walls that exist are structural and configurational — and both are detectable before you spend.

FINDING 3 · THE COST

What a fair attempt actually costs.

Depth is the load-bearing knob

Full circuit depth matters most; a **near-identity warm start** recovers about half of what missing depth costs; one popular knob (periodic head refitting) **mildly hurts**.

Fairness is cheaper than feared

With a sane, proportionate stopping rule, **a quarter of the naive training budget already holds the verdicts** on matched structure. Fair ≠ expensive — fair = configured right.

No knob buys territory

On data whose structure belongs to another method, **no configuration at any price** produced a win. The data grants territory; the knobs only collect it.

WHAT IT CHANGES

The fair recipe ships as the default.

The study’s recipe — full depth, warm start where registers are large, proportionate early stopping — is now the QML product’s **training default**, and a **25-iteration probe** that predicts the full-run verdict (AUC 0.76, registered threshold met) runs before you commit a budget.

Together with the kernel study’s island map, the suite now answers the pair of questions that matter: **is this data quantum-suited at all — and if so, which quantum method deserves it?**

Screen first, train second, and let the referee keep the score.

THE STUDY

The full study, unedited.

The **complete 10-page study** is public — the crossover figures, the head-to-head tables, the knob probes, the lost bet with its interval, and the registration tag that **provably precedes every run**.

The study — 10 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.