QEUI

Q//HORIZON 2026 · TEAM QEUI · EGYPT UNIVERSITY OF INFORMATICS

Can a quantum model learn grid stability, and when is it worth it?

Ziad Ahmed
Data and classical baselines

Seif Haytham
Encodings and circuit resources

AbdUlRahman Hany
Kernels and variational classifier

Abdelrahman Gabr
Integration, analysis, report

THE PROBLEM

Will this grid operating point stay stable?

10,000 simulated operating points of a four-node star grid: one producer and three consumers. Each point is labelled stable or unstable.

10,000
operating points, UCI dataset 471

63.8%
labelled unstable (6,380 rows)

We keep 8 inputs: each node’s reaction time (tau1 to tau4) and price coefficient (g1 to g4). The brief caps circuits at 10 qubits, so eight numbers fit at one qubit each, or fewer.

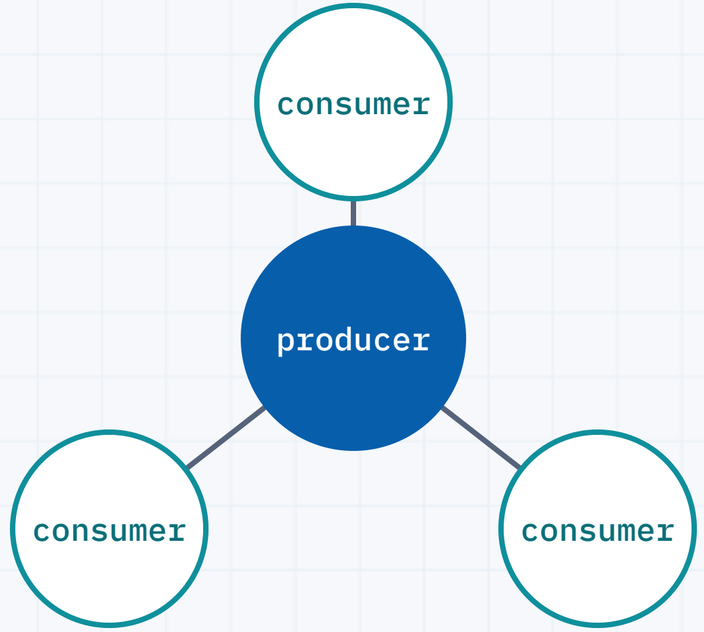

Each node has a reaction time (tau) and a price coefficient (g).

A FAIR FIGHT

Same rows, same folds, same features for every model

10,000 rows
UCI simulated grid

400 + 100
train and holdout rows

3 folds
scaling refit inside each

Pick by CV
macro-F1, never the holdout

Score once
on the 100 holdout rows

Same inputs

The classical control receives exactly the eight angle-scaled features that the circuits receive.

Same data budget

400 training rows for every model. The 8,000-row classical run is context, never the opponent.

Same scoring

Macro-F1 first, because 64% of rows are unstable. Row-level predictions are saved for paired tests.

ENCODINGS

Three ways to load eight numbers into qubits



8

qubits for 8 features

Angle

One feature per qubit, as a rotation angle.

depth 4 · 32 gates · 0 two-qubit



4

qubits for 8 features

Dense-angle

Two features per qubit: one RY, then one RZ.

depth 4 · 16 gates · 0 two-qubit



3

qubits for 8 features

Amplitude

The normalized vector becomes the state’s amplitudes.

depth 17 · 23 gates · 4 two-qubit

Counts after transpiling to rz, sx, x, cx. Fewer qubits is not cheaper: amplitude is the deepest circuit.

MODELS

Two quantum models built on those encodings

Fidelity-kernel SVM

$$K(x, z) = |\langle \varphi(x) | \varphi(z) \rangle|^2$$

Encode two rows, measure how much their quantum states overlap, and hand the overlap matrix to a standard SVM.

Simulated exactly: no shot noise, no hardware noise.

tested with angle, dense-angle and amplitude

Variational classifier

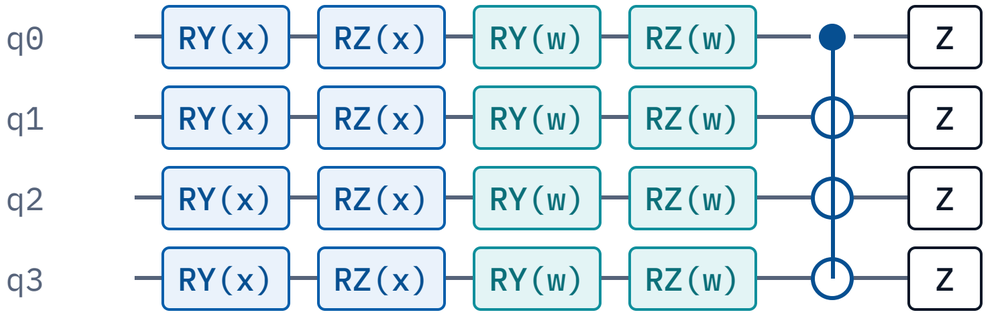

blue: data · teal: trained · one of 3 layers

One qubit per grid node, q0 the producer. Each layer re-uploads the node's two features with a trained scale and shift, rotates, then entangles from the producer. 77 parameters, Adam.

RESULTS

Results on the same 100 held-out rows

HOLDOUT MACRO-F1

0.902

classical RBF SVM, accuracy 0.91

0.830

variational circuit, accuracy 0.84

0.813

angle-kernel SVM, accuracy 0.84

Chart: cross-validation macro-F1, mean of 3 folds, all 18 models.

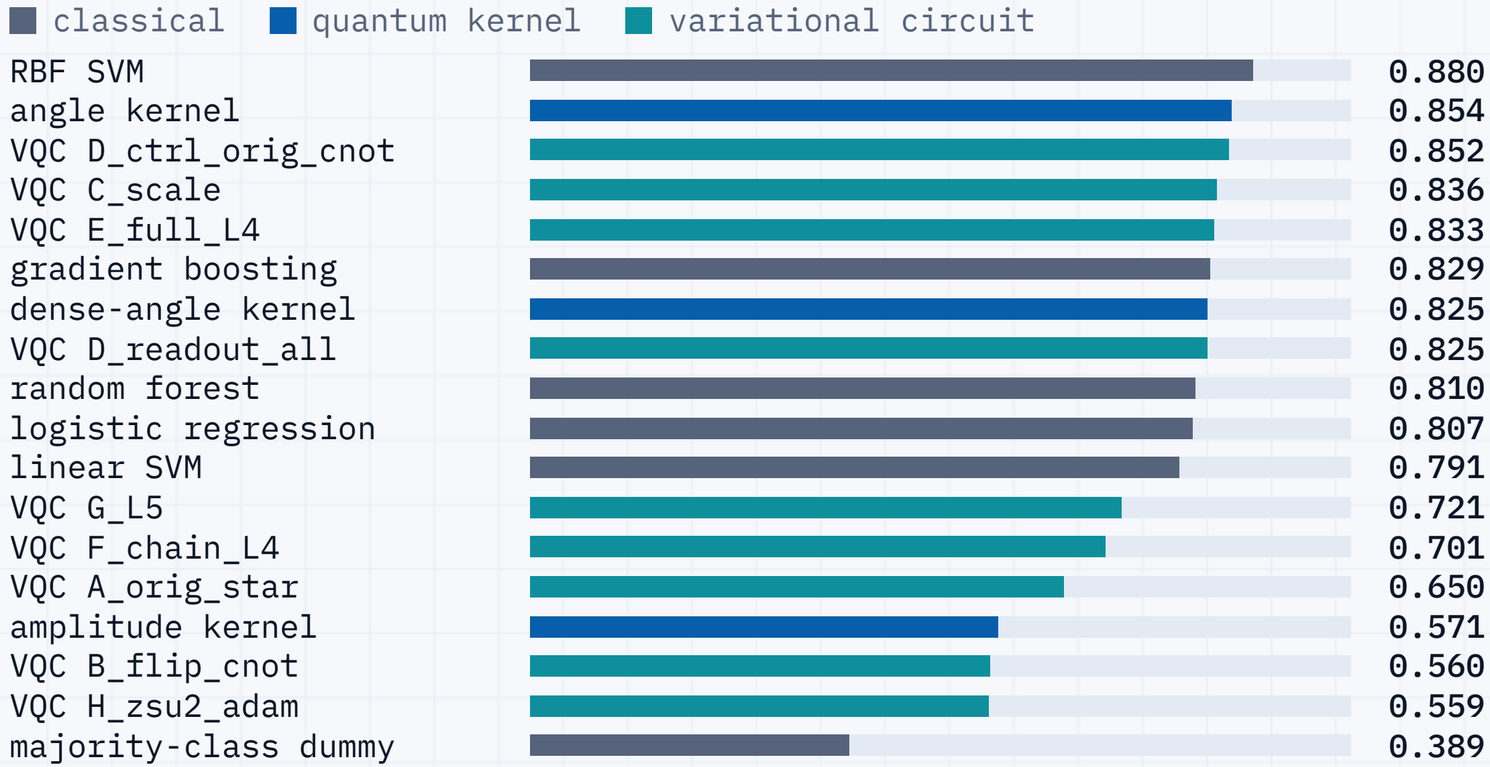

WHY, PART 1

The simple encodings are classical kernels

ANGLE KERNEL, 8 QUBITS

$$K = \prod \cos^2((x_i - z_i) / 2)$$

A product of one-qubit overlaps. No gate ever links two data qubits.

AMPLITUDE KERNEL, 3 QUBITS

$$K = (x \cdot z)^2$$

The squared dot product of the unit-norm inputs.

Simulator and formula agree to 2.4e-13 on 9,600 kernel entries each.

WHY AMPLITUDE FAILS · CV MACRO-F1 ON THE SAME INPUTS

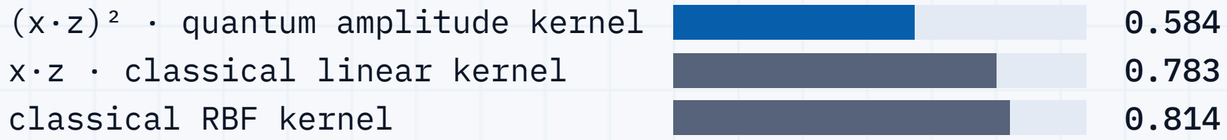

Cross-validation only, each kernel at its better C.

Squaring the overlap makes x and –x identical, so half the data folds onto the other half.

WHY, PART 2

The rotation scale is the knob that matters

Cross-validation macro-F1 on the training folds; kernels at the better of C = 1 and 10. Holdout not used.

Let the circuit learn it



Same variational circuit. A trained scale and shift per input lifts 0.560 to 0.836.

Rotate twice as far

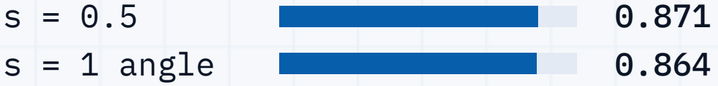

One separable kernel at three rotation scales. s = 2 is the Z feature map of our first kernel: that factor of two alone takes 0.864 to 0.698.

Add entangling gates

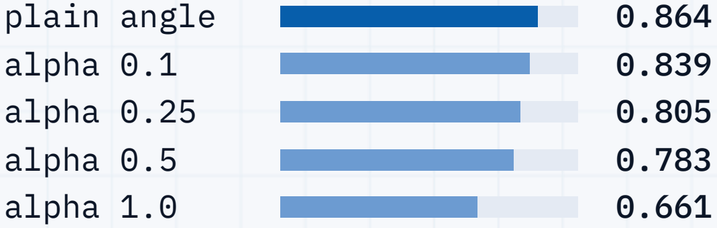

Grid-coupled map with its scale alpha tuned. Best 0.839 at the smallest scale, still below plain angle.

How far each input turns its qubit decided the score. Entanglement did not.

CONCLUSION

Classical wins here. The value is knowing why.

8 vs 1

Classical vs angle kernel: rows only one gets right. Exact McNemar p = 0.039.

13 vs 6

Classical vs variational circuit. p = 0.167: not separable on 100 rows.

10 vs 10

Variational circuit vs angle kernel. p = 1.0: a tie.

Limits: one simulated dataset, noiseless statevector simulation, a 100-row holdout, three folds. No hardware runs. Nothing here shows quantum advantage.

AI use: OpenAI Codex and Anthropic Claude assisted with code, analysis and documents. Every number comes from the project code.