Quantum Machine Learning Data Check: Is a quantum experiment on this table worth setting up?

17 November 2026. Table: qml_data_check_sample_table.csv. The table has 1,200 rows and 5 numeric feature columns and 3 classes; the smallest class holds 33 %. After the standard preparation a quantum model would need 5 qubits, which fits the simulator.

In short: this check cannot say yes or no for this table; a small experiment can.

This is a quick first look using classical models only. It is not a quantum run and no promise about one.

QML Suitability

Open (score band 41–60 out of 100)

Classical first	**Open**	Worth an experiment	Strong candidate
0–40	41–60	61–80	81–100

Your data carry a signal: the classical models score above a model that always predicts the largest class. The table is compact and fits a first experiment on 5 qubits. This check alone cannot decide for or against a quantum model here. A small diagnostic experiment, measured against the classical models, is the way to find out.

What decides it: the best classical model reaches a macro F1 of 0.63 out of 1.00, and neither the simple nor the nonlinear classical model stands out, so the question stays open.

What speaks for this data

Signal: the best classical model reaches a macro F1 of 0.63, against 0.17 for a model that always predicts the largest class.

Size: the prepared table has 5 features and fits 5 qubits.

Balance: the classes are the same size.

What speaks against it

No nonlinear structure found: a nonlinear classical model reaches 0.61, the best simple classical model 0.63.

Scores are macro F1 on 240 validation rows of the table, one split; 1.00 is perfect. The band rests on a score that weighs these findings and further checks.

From the published records

On two public credit datasets a plain quantum circuit did as well as deeper, costlier ones. Against logistic regression the best of 27 quantum configurations was on par on one dataset (macro F1: quantum 0.661, logistic regression 0.648) and below on the other (quantum 0.616, logistic regression 0.668); both gaps are smaller than the difference between two training runs. No advantage is claimed on real data (QL-ST-009).

All records and studies, with their numbers: qubit-lab.ch/evidence

With an access key to the RQP platform

- The exact score with its uncertainty, and every finding for and against, with the measure behind it.
- The Deep Assessment: quantum-kernel diagnostics on your own data.
- A run of the quantum model beside the classical models on identical splits, with a verdict.
- The full report, and the run as notebooks for five quantum frameworks.
- The first engagement is the Use Case Check: one use case, two weeks, fixed price.

If you want to talk this result through: thirty minutes on this table and your question.

Book a 30-minute call: calendly.com/daniel-hug-qubit-lab/new-meeting · or write to daniel.hug@qubit-lab.ch

Run the check on a table of your own: qubit-lab.ch/data-check · Daniel Hug, founder: ETH physicist, MBA, formerly UBS, Credit Suisse, SIX.

qubit-lab.ch held the table in memory only for this check, in the Google Cloud region Zurich, and deleted it when the answer was returned; nothing was written to disk or kept. If the table was pasted or attached in a chat, it also passed through that AI assistant, whose data terms apply. Use public, synthetic or masked data for this check, never client data.

This page is a first indication from classical models. It is not a certificate, not a quantum result and no forecast of one. The pages that follow describe the RQP platform.

qubit-lab.ch

Quantum Computing for Financial Services.
The RQP platform: vendor-neutral tests on your own problems and data.

qubit-lab.ch

Test before you invest

Every bank is asked what quantum computing will do for its business. Vendor claims do not answer that. Only a measured result on the bank's own problem does.

The RQP platform gives that answer. A team loads its own data, runs today's quantum algorithms on it, and gets a report that states plainly how the quantum result compares with the best classical one.

What the bank gets

A verdict per use case
Each run ends in a report with the verdict first: advantage, parity or no advantage, with the numbers behind it.

Reports for two audiences
Every run produces a two-page overview for management and a full report for specialists, both as PDF.

Runs on quantum hardware
Simulators, and for several RQPs IBM hardware on the team's own account: the bank sees what today's machines deliver.

Vendor-neutral, no lock-in
Every run exports as notebooks for five quantum frameworks and as OpenQASM, so the bank's quants can rerun and extend it in their own environment.

Results in hours
A run takes hours, a full use case a few weeks. Seven integrated, ready-to-use Rapid Quantum Prototypes (RQPs) cover machine learning, optimisation and pricing.

A repeatable benchmark
The same workbook can be rerun as hardware improves, so the bank sees when the answer changes.

Evidence, not promises. Every quantum result is benchmarked against a classical referee, a certified optimum wherever one can be computed. We claim no quantum advantage today. On real financial data the published results show parity at best, and the platform shows exactly where each method stands. The studies and run records are public, with all numbers and data: **qubit-lab.ch/evidence**

How to start

1 · First call and demo	Thirty minutes to pick the right start.
2 · Personal access	Guest access to the platform on a sample workbook, at your own pace.
3 · Check or Sprint	One use case tested by us in two weeks, or your team on three to four in eight or ten weeks. Fixed price.

For a team new to quantum. The Quantum Capability Kickstart, five hands-on half-days, opens a Quantum Sprint or can be booked on its own.

Interested in a live demo? Contact me.
Daniel Hug, founder: ETH physicist, MBA, formerly UBS, Credit Suisse, SIX. · daniel.hug@qubit-lab.ch · qubit-lab.ch

qubit-lab.ch

Quantum Computing for Financial Services.
The RQP platform: vendor-neutral tests on your own problems and data.

qubit-lab.ch

The seven RQPs

RQP	The bank's question	Quantum method	Measured against
QML Classification	Fraud detection, credit scoring and downgrade prediction on the bank's own tables	Variational classifier, quantum kernel machine	Best standard models on the same data split
Portfolio Optimization	Which assets to hold, given risk, return, budget and type constraints?	QAOA	Heuristic benchmark and certified optimum (exact classical solver)
Index Replication	Which few names track the index best?	QAOA	Heuristic benchmark and certified optimum
Credit Allocation	Which facilities to take in a book purchase or risk transfer, given correlated losses?	QAOA	Heuristic benchmark and certified optimum
Collateral Allocation	Which collateral to post against which obligation, trading cost against coverage under stress?	QAOA	Certified optimum
Feature Selection	Which columns should feed a model?	QAOA	Heuristic benchmark and certified optimum
Quantum Monte Carlo	Derivatives pricing, including correlated baskets	Amplitude estimation	Classical Monte Carlo, with the crossover point in samples

How a run works

1 Workbook in
The team describes the problem in an Excel workbook, or a CSV table for classification.

2 Two runs, one referee
The quantum algorithm and the classical references solve the same problem under the same constraints.

3 Full report out
A full report and a two-page overview, with verdict, gaps, confidence intervals and seeds.

What a quant can check

- **The reference is exact.** For the optimisation RQPs an exact solver proves the optimum, so a gap is a fact, not an estimate.
- **Reports show the working.** The full report lists the circuit, settings, seeds, confidence intervals and the gap to each reference; the overview carries the verdict and the key numbers.
- **Notebook downloads.** Each run exports a rerun notebook, which reproduces the result, and a retrain notebook, which trains the circuit again from the data. Both come for Qiskit, PennyLane, Cirq, CUDA-Q and Classiq, with the circuit also as OpenQASM.
- **Hardware runs.** Several RQPs also run on IBM hardware through the team's own IBM account, and the report shows the hardware result next to the simulator result.
- **Failures are published.** Runs where the quantum method loses appear in the evidence records next to the ones where it matches. Numbers and data: qubit-lab.ch/evidence

Your data. First tests run on synthetic or masked data, and the exported notebooks rerun in-house, so nothing confidential has to leave the bank. Workbooks need no client identifiers. Platform runs with an access key are processed in the Google Cloud region Zurich, and their run artifacts are deleted after 24 hours. NDA on request.

qubit-lab.ch | daniel.hug@qubit-lab.ch