

RAPID QUANTUM PROTOTYPING (RQP)

The RQP Suite Overview

One honest methodology, seven instruments - portfolio optimization, index replication, feature selection, credit allocation, collateral allocation, quantum machine learning, and quantum Monte Carlo. What each tool does, which quantum methods it runs, what the runs actually delivered, how to take the trained circuits into your own stack, and how to try it on your own data.

September 2026

qubit-lab.ch

qubit-lab.ch
contact@qubit-lab.ch
© 2026 qubit-lab.ch · All rights reserved · qubit-lab.ch/legal

Contents

1. What this suite is for . . . 2
2. The seven instruments . . . 3
3. Portfolio Optimization RQP Pro . . . 3
4. Index Replication RQP . . . 3
5. Feature Selection RQP . . . 4
6. Credit Allocation RQP Pro . . . 4
7. Robust Stochastic Collateral Allocation RQP . . . 4
8. QML RQP Pro - quantum machine learning . . . 4
9. QMC RQP Lab - quantum Monte Carlo . . . 4
10. The quantum method, in four switches . . . 5
11. Take the circuit with you - code exports . . . 6
12. Worked examples, in numbers . . . 7
13. The chain - a complete, honest arc on real data . . . 7
14. Six method studies, published against ourselves . . . 8
15. What you get with every run . . . 8
For the technically curious: how the quantum engine works . . . 8
16. Trying it and working with qubit-lab.ch . . . 9
Appendix - technical parameters, tool by tool . . . 9

© qubit-lab.ch

1. What this suite is for

Quantum computing in finance is at the stage where the right move is neither to believe nor to dismiss, but to measure. The Rapid Quantum Prototyping (RQP) suite is that measuring instrument. You bring your own data in an Excel workbook. The suite runs the best current quantum method and the best classical method on the identical problem, and hands back three things: a report, the code, and a verdict you can defend.

Four kinds of people use it, for four different reasons:

- **The senior manager** gets a vendor-neutral answer on the firm's own use case - advantage, parity, or classical wins - with a certified referee behind the number. Enough to decide what to fund and what to park, in hours, without a vendor slide and without a coding project.
- **The quant team lead** gets a benchmarking bench: every quantum result next to the strongest classical baseline, same objective, same data, simulator or real hardware. The team's own use cases, measured the same way, comparable across quarters.
- **The quant** gets tuning that takes minutes instead of days: flip a switch, rerun, compare the two reports side by side. Through the MCP interface, activated on request, an AI assistant can drive whole parameter sweeps and summarize what moved the result.
- **The learner** gets a working reference: a trained circuit exported as Qiskit, PennyLane, Cirq, CUDA-Q or Classiq code. Compare it with your own implementation, rerun it, then change it with your own ideas and see what happens.

Four commitments hold across all seven products, whoever is using them:

- **Like-for-like scoring.** Quantum and classical methods are judged on the identical objective, on the identical data. No cherry-picked baselines.
- **A referee wherever possible.** Certified exact solvers (SCIP) prove the optimum for the optimization tools; tuned classical baselines and pre-registered protocols referee the ML and Monte Carlo tools. A "0.00% gap" in an RQP report is a mathematical statement.
- **The gate that says no.** Assessments run before any paid quantum run and are allowed to conclude "Do Not Prioritize Quantum". A tool that can never say no is marketing, not measurement.
- **Verifiable artifacts.** Every claim lands in a report, a notebook, or a machine-readable file that reruns offline. No claims you cannot check.

2. The seven instruments

Everything runs on your own data, uploaded as a simple Excel workbook: portfolios, index compositions, credit books, collateral inventories, labelled datasets, payoff scenarios. Current focus is finance; energy and other domains are covered by the shipped demonstrators and can be prepared jointly.

Tool	Tier	The business question
Portfolio Optimization RQP Pro	Pro	Which investment blocks should the portfolio hold, under budget, risk, and type constraints?
Index Replication RQP	Pro	How well can a large index be tracked with a small basket - and what does each mandate rule cost?
Feature Selection RQP	Pro	Which K features may the model use, under fairness, bundle, and budget constraints?
Credit Allocation RQP Pro	Pro	Which facilities enter the book when exposures move as one transaction, under capital and concentration risk?
Robust Stochastic CLA RQP	Pro	Where should each free collateral line be posted - at what cost, and how stable is the coverage under stress?
QML RQP Pro	Pro	Can a quantum kernel or variational classifier beat tuned classical models on your labelled data?
QMC RQP Lab	Lab	Where does quantum amplitude estimation stand against classical Monte Carlo on real payoffs?

Pro versus Lab is an honesty distinction, not a quality one: Pro tools create suspense on your own data - the answer is genuinely unknown before the run. The Lab tool delivers transparency on estimator convergence - the scientific state of play, measured cleanly.

3. Portfolio Optimization RQP Pro

Cardinality-style portfolio construction as a QUBO: costs, expected returns, covariance, budget, and up to five subtype budget constraints (sector, region, asset class, sustainability bucket). Four independent routes attack the same objective: a classical heuristic, the QAOA sampler, the certified SCIP referee, and - on request - a real IBM quantum computer as a **hardware second opinion**, whose noisy readout is reported as honestly as the simulator's clean one. The four method switches of section 10 and the code exports of section 11 apply in full. Four measured switches shape the quantum method - training objective (expected energy or CVaR), ansatz (fixed depth or the adaptive, seed-independent second opinion), start state (uniform or the classical seed start after Egger, with blur epsilon), and the penalty weights - and every report prints how much probability the trained circuit put on the certified optimum next to the certified gap.

4. Index Replication RQP

Replicate a large benchmark with a small basket - fixed core holdings, basket-traded building blocks, single lines - under exclusions and type budgets. The referee prices every mandate rule: the report states the certified tracking-error cost of each ESG exclusion, fixed holding, and band. The cardinality-conserving XY mixer keeps quantum readout feasible by construction; enhanced-indexing mode adds a return tilt on the same certified ladder.

5. Feature Selection RQP

Upstream of any model: which K of N candidate features may the model use, under must-include and must-exclude policy, all-or-nothing bundles (one-hot families, sin/cos phase pairs), exclusive variant groups, and type budgets - exact or band, conserved by the quantum mixer with zero slack qubits. The solver ladder ends in a certified optimum plus a ranked **solution pool**, because the best generalizer is not always the objective winner - choosing rank 5 is a governance decision, and the manifest records it. One click exports a colored QML handoff pack for the next tool in the chain.

6. Credit Allocation RQP Pro

Batch credit selection when exposures move as one transaction - portfolio acquisitions, significant-risk-transfer pools, participation pipelines: which facilities enter the book, under a capital budget, sector caps, and the fact that defaults cluster. The quadratic heart is correlated unexpected loss - the concentration cross-terms additive capital formulas assume away; connected-client groups decide as one (bundles), the committed book enters as constants and cross-terms, and the certified SCIP referee proves the optimum. The report adds the committee view - expected loss, book unexpected loss with diversification benefit, RAROC - and a classical Monte Carlo loss tail, labeled as such. The same four quantum switches as the Portfolio instrument (objective, adaptive ansatz, classical warm start, penalties) are available, and hardware replay and notebook exports reproduce the trained circuit exactly.

7. Robust Stochastic Collateral Allocation RQP

Assign free collateral lines to margin obligations at minimum posting cost while a scenario-covariance term keeps the coverage stable under stress. The inventory is taken as it actually is: locked pre-posted lines (constants with covariance cross-terms), pledge pools and baskets that move whole, per-venue eligibility with haircut and FX pricing, and concentration limits priced or hard. A single knob - lambda_risk - trades funding cost against coverage stability, and the classical frontier sweep shows the whole certified curve. The referee enumerates every valid allocation exhaustively; the QAOA method is warm-started from the classical incumbent and disclosed as such. The adaptive ansatz is available as the second opinion that does not inherit that incumbent, alongside the CVaR objective and the penalty weights.

8. QML RQP Pro - quantum machine learning

Quantum kernel methods (QSVM) and variational classifiers (VQC) against seriously tuned classical baselines, on a temporal-honest split. The staged pipeline - prepare, assess, plan, execute - puts two assessments before any paid run, including the **Fourier Wall analysis** (order-matched periodic twin with joint terms to 3rd order, off-grid indicator, entanglement lift, joint-spectrum density, and the five-condition quantum-territory checklist; independent implementation after arXiv:2607.15815) that predicts whether quantum feature maps have anything to grip on your data. The verdict metric is the Quantum Advantage Factor with a full confidence interval; the whole-CI rule is pre-registered, and "Do Not Prioritize Quantum" is a common, correct answer. Since September 2026 an optional **quantum feature augmentation** adds referee rows: the classical baselines are also trained on the quantum model's own encoded features, so a quantum win must beat the strongest classical use of the same information.

9. QMC RQP Lab - quantum Monte Carlo

Amplitude-estimation pricing of option payoffs (calls, puts, digitals, barriers, baskets) benchmarked against classical Monte Carlo on shared budget schedules: absolute error versus oracle calls, both estimators, same axes. The quadratic-speedup claim, measured instead of cited. Estimators include low-depth maximum-likelihood amplitude estimation and iterative amplitude estimation; every readout is checked against analytic and grid-based references before it is compared.

10. The quantum method, in four switches

Five of the seven tools run the same quantum engine, QAOA (explained in plain language at the end of this document). Four switches shape how it is run, and every report prints the positions that were used:

- **Training objective.** Expected energy, or **CVaR**: train on the best tail of the readout, so the circuit concentrates probability on good solutions instead of on a good average.
- **Ansatz.** Fixed depth, or **adaptive** (ADAPT-QAOA): the circuit grows its own layers from a pool of candidate operators. A seed-independent second opinion on the fixed-depth result.
- **Start state.** Uniform, or the **classical seed start** after Egger: the circuit begins near the classical incumbent, blurred by a disclosed epsilon so that it can still leave it.
- **Second opinions.** A noisy simulator calibrated on a **noise ladder** (or on a real IBM backend's calibration data), a run on **IBM hardware**, and the **certified exact run**.

Not every switch exists on every tool: the constraint-conserving tools keep readout feasible with a mixer instead of a penalty, and the collateral tool has no hardware path yet.

Tool	Objective	Ansatz	Start state	Feasibility	Second opinions
Portfolio Optimization	energy / CVaR	fixed / adaptive	uniform, angle warm start, classical seed	penalty; matched mixer with the seed	noisy simulator, IBM hardware, certified exact
Index Replication	energy / CVaR	fixed	uniform	X mixer with penalty, or XY mixer that conserves the basket size	certified exact, greedy and random baselines
Feature Selection	energy	fixed	uniform	XY mixer conserving type budgets, or X mixer with repair	certified exact plus ranked solution pool
Credit Allocation	energy / CVaR	fixed / adaptive	uniform, angle warm start, classical seed	penalty; matched mixer with the seed	noisy simulator, IBM hardware, certified exact
Robust Stochastic CLA	energy / CVaR	fixed / adaptive	uniform, classical seed	penalty; lambda_risk frontier	exhaustive certified referee, frontier sweep

The two other tools have switches of their own. **QML** chooses between quantum kernel and variational classifier, then feature map, ansatz, entanglement and measurement - with the Fourier Wall assessment and the optional augmentation referee rows in front of any run. **QMC** chooses the amplitude-estimation estimator, the register sizes and the budget schedules. The technical parameters of all seven tools are listed in the appendix.

11. Take the circuit with you - code exports

Every completed run can leave the platform as code. Each product page carries a compact export table, one row per environment, with a **Rerun** notebook and, where it applies, a **Retrain** notebook. Rerun replays the trained circuit with its trained angles in your own stack and compares the readout with the RQP result. Retrain re-optimizes the angles there from the exported problem coefficients and compares the gap; the QMC tool re-estimates the run's budgets instead.

Environment	Rerun	Retrain
Qiskit	notebook, plus a plain Python script on the Pro optimization tools	notebook
PennyLane	notebook	notebook
Google Cirq	notebook	notebook
NVIDIA CUDA-Q	notebook	notebook
Classiq (Qmod)	notebook	notebook
OpenQASM 2.0	circuit file	-

All seven tools carry the table. Every notebook embeds its code-export package - circuit specification, trained parameters and the coefficients of your problem - so it runs without the platform. The Qiskit, PennyLane and Cirq notebooks are execution-verified with pinned package versions; the Classiq rerun path was executed in Classiq Studio for all seven tools, including both QML models (on QML the Classiq notebook is a rerun notebook, because the platform's optimizer targets energies rather than classifier losses); the CUDA-Q notebooks are gate-verified against a recorded run, and execution on a CUDA-Q runtime was not part of the verification. Feature Selection additionally exports a QML handoff pack for the next tool in the chain. The exports embed workbook-derived coefficients: treat them like the workbook. On the Portfolio, Credit and Collateral tools a Tester key unlocks the exports.

© qubit-lab.ch

12. Worked examples, in numbers

All numbers below come from actual RQP runs; the full reports, workbooks, and notebooks behind them are available on the product pages and on request.

Allocating new capital for a private bank (Portfolio Optimization, demo case). A global equity model portfolio with a Europe/Switzerland anchor: 78 positions evaluated, 23 investment decisions mapped to 23 qubits, a USD 1.013m budget, a technology band of USD 225k-305k, and six pre-approved thematic bundles. QAOA at depth p = 4 with layerwise warm start reached, in CVaR objective mode, **100% of the certified optimum**. The problem size is deliberate: large enough to carry the real mandate structure, small enough that the exact referee can still certify the optimum - which is what makes the benchmark meaningful.

The S&P 500 with 20 stocks (Index Replication, demo case). The full 500-name benchmark; an instrument pool built the way mandates actually trade: five fixed mega-caps, CIO sector blocks, AMC-style theme baskets (climate transition, dividend quality, AI infrastructure), a completion basket of the smallest index weights, and ten single lines - 17 qubits in total. Certified optimum: **4.42% tracking error with exactly 20 ISINs**. The two mandate rules - a technology cap and a climate-transition floor - cost a certified 0.40% per year, and the report shows what the ESG rule displaced (the AI basket). QAOA reproduces the certified optimum exactly - and the referee proves it.

Governed feature selection on one million rows (Feature Selection, real public data). The NeurIPS Bank Account Fraud dataset: 32 raw columns reduced to 15 governed features under a fairness exclusion, a mandated bureau score, one-hot bundles, and a behavioral budget. QAOA matched the certified optimum at **0.00% gap** with every constraint conserved by construction - feasible readout mass 1.0000, zero repairs.

Buying a competitor's SME book (Credit Allocation, demo case). A regional bank evaluates a competitor's loan book offered as one package: 45 facilities, 8 connected-client groups that transfer whole or not at all, 23 decision qubits, a board envelope covering 62% of demand, and a 25% commercial-real-estate cap that bites. Over 8.4 million possible baskets the exact referee proves the optimum in minutes - and the classical heuristic misses it by **0.19%**, the first scale where the certificate visibly catches money left on the table. On the shipped sample book the tuned CVaR configuration reached the certified optimum at **0.00% gap** - after default settings failed honestly at 40.8%, and the report published both.

Margin calls under stress (Robust Stochastic CLA, demo case). Tuesday morning at a Swiss bank treasury: three venues calling at once, 23 collateral lines - ten locked, three pledge pools moving whole - 50 stress scenarios, 24 QUBO variables. The warm-started CVaR QAOA matches the certified optimum at **0.00% gap, with the optimum as the single most likely readout of the whole register** (2.6% over 16.8 million states). The frontier prices coverage stability: on the full-scale variant, roughly CHF 19,000 a year of extra posting cost buys the coverage volatility down from 8.3M to 7.6M - the risk-return trade of a treasury book, certified point by point.

Quantum kernel classification (QML). On an engineered 8-qubit benchmark designed to favour quantum kernels, QSVM reached **0.87 test accuracy and 0.95 ROC-AUC versus 0.73 and 0.85** for the best classical baseline on identical data - so you can see what winning conditions look like. On plain tabular data, classical baselines often still win, and the reports say so openly (see the chain below).

Pricing readout (QMC Lab). Quantum amplitude-estimation readout reproduces the analytic and grid-based references across European, digital, capped, basket, and barrier payoffs, plus rare-event probability estimation, with confidence intervals and convergence curves against classical Monte Carlo on shared oracle budgets.

13. The chain - a complete, honest arc on real data

The suite's tools compose, and the flagship demonstrator runs the full chain twice. **Act one**: on the public Bank Account Fraud dataset, Feature Selection solved the constrained 15-of-32 selection to a certified 0.00% gap and handed off to QML - whose own gate said Do Not Prioritize Quantum. We ran the kernel anyway, for the record: classical won, exactly as predicted. **Act two**: on two years of real German power-market data, phase-encoded cycle features chosen by the assessment's own mechanism reasoning reached **certified statistical parity** - Quantum Advantage Factor 0.99, confidence interval 0.975 to 1.011 on purely future test hours - under a pre-registered six-configuration protocol with a hard stop.

The dataset where the gate said no. The dataset that reached certified parity. The engineered dataset where the quantum kernel wins. All three exhibits are downloadable with reports, workbooks, and rerunnable notebooks - the loss published next to the win, because the loss is what makes the win credible.

14. Six method studies, published against ourselves

Between August and September 2026 the same machinery was pointed at the suite itself. Six method studies are public on qubit-lab.ch/evidence, each with a design locked before the first run, certified classical referees on every run, committed scripts that reproduce every published number, and revision notes that disclose corrections: **quantum concentration** (how much probability a trained circuit puts on the certified optimum, and which switches move it), the **noise ladder** (a calibrated simulator ladder with a hardware rung on ibm_kingston and calibrated digital twins), the two **advantage-island maps** (where quantum feature maps have territory on real data, for kernels and for variational classifiers) and the two **configuration atlases** (which kernel and classifier configurations ever win against tuned classical baselines, and which never do). The companion paper The RQP Evidence Engine describes the machinery and shows how to run a question of your own through it.

15. What you get with every run

- **A branded PDF report** written for a committee, not only for quants - verdict first: every report opens with the verdict and the certified gap, then the plain-language lead-ins, solver comparisons, circuit resources, and a methodology primer.
- **Runnable notebooks** in six environments (section 11) with the full result embedded - they rerun offline and survive cloud record expiry.
- **Machine-readable results** (JSON/CSV) for your own tooling, plus the exact workbook and settings for reproduction.
- **An honest verdict** - advantage, parity, or classical win - stated in pre-registered terms you can defend internally.

For the technically curious: how the quantum engine works

Five of the seven instruments - portfolio optimization, index replication, feature selection, credit allocation, and collateral allocation - run on the same quantum engine, the Quantum Approximate Optimization Algorithm (QAOA). In brief:

Unlike a classical optimizer, QAOA does not examine one candidate portfolio after another. Instead, it creates a single quantum wave that spans the entire search space, giving every possible solution the same initial probability.

The algorithm then repeatedly applies two global transformations to this wave. The cost step causes the wave to evolve according to the optimization objective. The mixer step allows the wave to evolve further, enabling different parts of the wave to interfere. Through constructive and destructive interference, the probability distribution is gradually reshaped. Regions of the search space containing good solutions become increasingly likely, while regions containing poorer solutions become less likely.

After a small number of these alternating steps, the quantum state is measured. Each measurement returns a single portfolio; in practice the circuit is read out many times, yielding a ranked list of candidates. There is no guarantee that the best candidate is the optimum, but it is significantly more likely to be a high-quality solution than one selected at random.

The only quantities that must be adjusted are the strengths of the cost and mixer steps. These are tuned by a conventional classical optimizer, which repeatedly executes the quantum circuit, evaluates the measured results, and updates the parameters until the best setting is found. In the adaptive variant the circuit also chooses its own next mixer step, each time the candidate with the steepest gradient, and stops when no candidate improves the result - so the depth is measured, not guessed.

A useful intuition is to imagine a hall filled with tuning forks. Nobody walks through the hall striking or tuning individual forks. Instead, the entire hall is exposed to a carefully designed sequence of sound fields. The forks respond collectively. As the sound field changes, resonance and interference gradually reshape the overall pattern of vibration

until certain resonances dominate while others fade away. Listening to the hall is like measuring the quantum state: the strongest resonance is now much more likely to correspond to a good solution than to a poor one.

The analogy is not exact, but it captures the essential idea. QAOA does not search through candidate portfolios one by one. It repeatedly reshapes a single collective wave, using interference to make good solutions more likely to be observed.

The full treatment - the picture, the mathematics, and each tool's constraint-conserving variant - is in every product manual's "The algorithm in two passes" section.

16. Trying it and working with qubit-lab.ch

Each product page on qubit-lab.ch carries sample workbooks, sample reports from the runs above, and demo access; an LLM prompt pack lets any capable AI assistant build a valid workbook from your own data against the live validator, and the MCP interface, activated on request, lets the assistant validate, run and sweep directly. For teams, the **Lean Quantum Exploration** process on the services page takes you from a first briefing through a hands-on kickstart to a six-week guided exploration of your own use cases on the suite - fixed fees, agreed up front. Swiss-based, hands-on, one counterpart.

Intro call, demo keys, walkthroughs: contact@qubit-lab.ch

Appendix - technical parameters, tool by tool

Small print for the technically minded: workbook keys, allowed values and defaults as implemented in September 2026. The product manuals carry the full semantics. Common to the five QAOA tools: depth p (layers), optimizer iterations (maxiter), multistart restarts, exact readout or shots, random seed, and the referee's time limit and number of certified solutions.

Portfolio Optimization RQP Pro. Mode classical_only / qaoa_lightning_sim (exact probabilities up to 24 qubits) / qaoa_tensor_sim (sampling). Ansatz standard / adaptive with adapt max layers 1-24 (default 12; adaptive is exact-simulation only, ignores p and restarts, and maxiter becomes the per-layer budget). CVaR mode off / polish / objective, alpha in (0, 1] with default 0.10, polish iterations 30. Layer-wise angle warm start on/off, restart perturbation 0.05. Warm start mode none / classical (Egger seed start, matched mixer) with epsilon 0.01-0.49, default 0.25. Second opinion internal_only / qiskit_export / qiskit_noisy_sim / ibm_external_run / certified_exact_run. Noise source parametric / ibm_calibration; two-qubit fidelity F2 in [0.9, 1.0], default 0.995, presets 0.995 / 0.999 / 0.9999; single-qubit, two-qubit and readout error rates derived from F2 unless set; noisy shots 256-200,000, default 4,096, up to 22 qubits. IBM hardware: token, instance, backend override, hardware shots, fractional RZZ gates, parallelization. Certified exact on/off, solutions 5, time limit 60 s. Penalty weights lambda_budget 50 and lambda_variance 6. Classical optimizer COBYLA, not selectable.

Index Replication RQP. Methods qaoa_sim, greedy_forward, random_baseline, certified_exact. Layers 2, maxiter 60, restarts 3. Mixer x (transverse, with automatic cardinality penalty) or xy (ring-XY from a Dicke state, one ring per constraint group). Readout exact or shots (4,096). CVaR mode off / objective, alpha 0.10. Candidate pool 20, random subsets 200, refit iterations 600, random seed. Optimizer Adam for expected energy and COBYLA for CVaR, not selectable.

Feature Selection RQP. Mixer xy (default; conserves exact or band type budgets with zero slack qubits) or x (penalty plus repair; band constraints require xy). Layers 2, restarts 3, maxiter 150, shots 0 = statevector. Strict readout off/on (discard infeasible readout instead of repairing it). Type budget penalty auto. Referee time limit and number of certified solutions (top-N via no-good cuts, default 1). Random seed. Optimizer Adam, not selectable.

Credit Allocation RQP Pro. The Portfolio Optimization engine keys apply unchanged: ansatz standard / adaptive (max layers 1-24), CVaR off / polish / objective with alpha, p, maxiter, restarts, angle warm start, warm start mode none / classical with epsilon 0.25, lambda_budget, random seed, and the same second-opinion, noise-ladder and IBM hardware fields. In addition: exact probability export as shots on/off (draw the candidate list from a finite sample instead of the exact top-K list).

Robust Stochastic CLA RQP. Ansatz standard / adaptive (max layers 1-24). p default 1 (the measured configuration uses 3), maxiter 100, restarts 1, shots 1,024, random seed. CVaR mode off / objective, alpha 0.1. Warm start mode none / classical (seed = the classical incumbent) with epsilon 0.25, clamped to [0.01, 0.5]; the layer-wise angle warm start is not applicable. Run mode classical_only / qaoa_lightning_sim. lambda_risk sets the point on the cost-versus-stability frontier. Referee: exhaustive enumeration. No noise ladder and no hardware run on this tool.

QML RQP Pro. Model family vqc_classifier / quantum_kernel_svm. Feature map angle (rotation Y or X) / iqp / zz_like, repeats 1; input scaling robust_pi / robust_half_pi at percentile 95, bandwidth 0.05-4.0 (default 1.0). Ansatz hardware_efficient / basic_entangler / strongly_entangling / real_amplitudes_like / custom_rx_ry_cz, reps 2, entanglement linear / circular / full, measurement expval_z / expval_z_zz / expval_parity_all (below 12 qubits), readout head linear / quadratic, data re-uploading yes/no. Backend auto / pennylane_lightning / pennylane_default_qubit, analytic or finite shots. Optimizer adam / spsa, learning rate 0.01, batch size 64,

iterations 50, repeats, init strategy random / identity_block, early stopping with patience 20, threshold optimization on F1, validation metric PR-AUC, class weights auto, training cap 20,000 rows. Assessment fast / deep (with kernel sample cap); quantum feature augmentation referee rows on/off.

QMC RQP Lab. Run mode inspect_workbook / run_scenario / run_convergence. Estimator low_depth_ae (maximum-likelihood amplitude estimation) / iqae (experimental) / objective_shots / grover_objective_shots / objective_statevector / statevector / shots. Qubits per register 1-7, one or two uncertainty registers, eight qubits in total; encoding bounded_grid. Shots 2,048, confidence level 0.90 / 0.95 / 0.99, target error 0.01. Convergence: classical sample schedule, repetitions, quantum shot schedule, Grover power schedule, caps on total shots, oracle calls and depth. Simulator only, no noise model.

Legal notice. © 2026 qubit-lab.ch. All rights reserved. This document and the software it describes are proprietary to qubit-lab.ch and provided for authorized use only; reproduction or redistribution without written permission is prohibited. Outputs of the RQP tools are intended for technical review, experimentation, demonstration, and structured discussion - they are not financial advice. Full legal terms: qubit-lab.ch/legal.