Quantum Machine Learning

- QML fuses quantum computing with machine learning
- Course scope: models, benchmarks, honest evaluation
- 2026 status: quantum-data advantage proven; business-data edge unproven
- Hybrid quantum-classical pipelines are the practical norm
- Roadmap: foundations → algorithms → applications → adoption

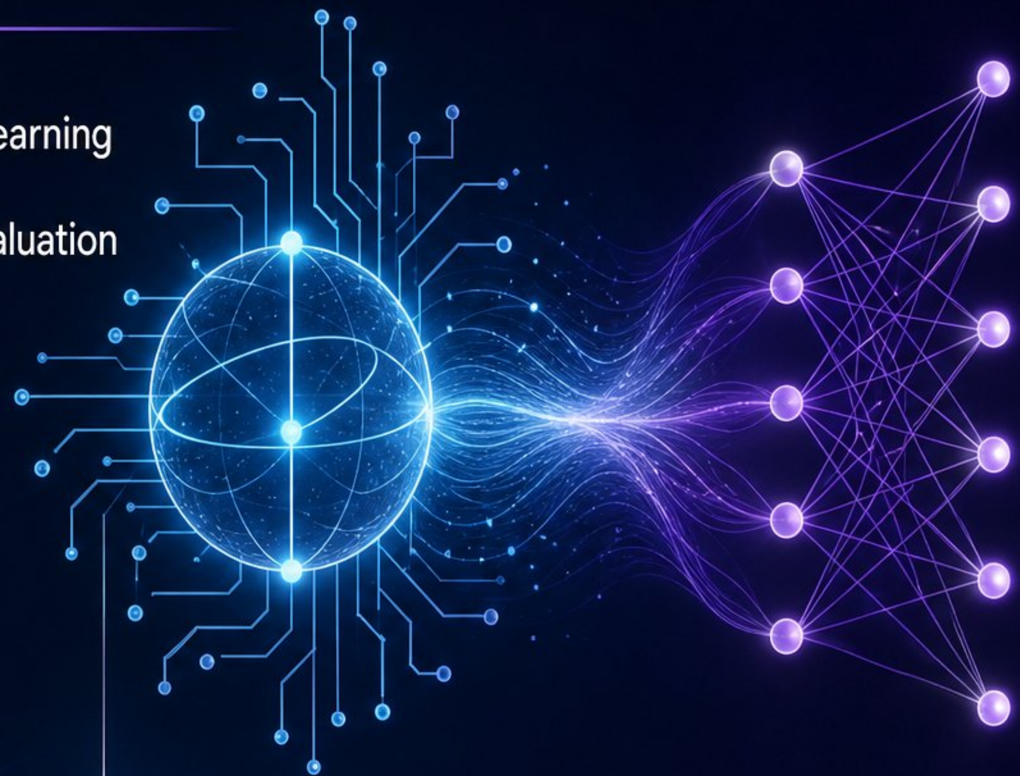

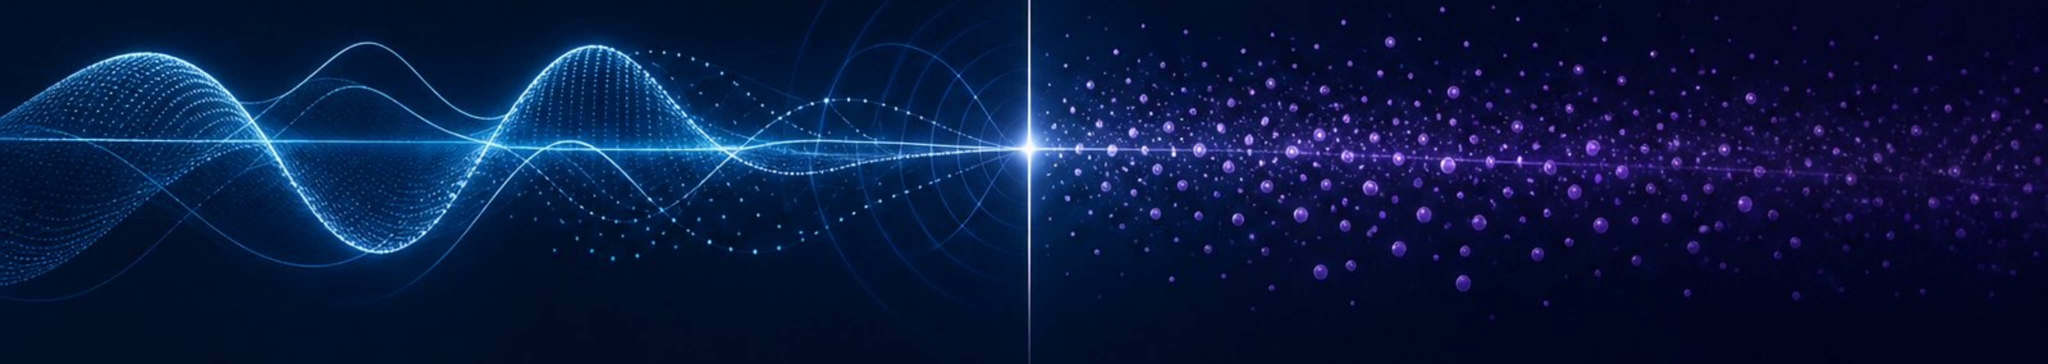

Why Quantum Machine Learning Matters Now

- The open question: when do quantum resources improve learning?
- Quantum-data advantage is experimental; classical-data advantage is unproven
- Hybrid classical-quantum pipelines dominate practice
- Balancing real evidence against vendor hype

Quantum Computing Foundations for ML Practitioners

- \- Qubits exploit superposition and entanglement for computation
- \- Quantum gates and circuits act as computational graphs
- \- Measurement collapses quantum states into classical outputs
- \- BQP captures problems efficiently solvable on quantum computers
- \- Key differences from classical bits shape model design

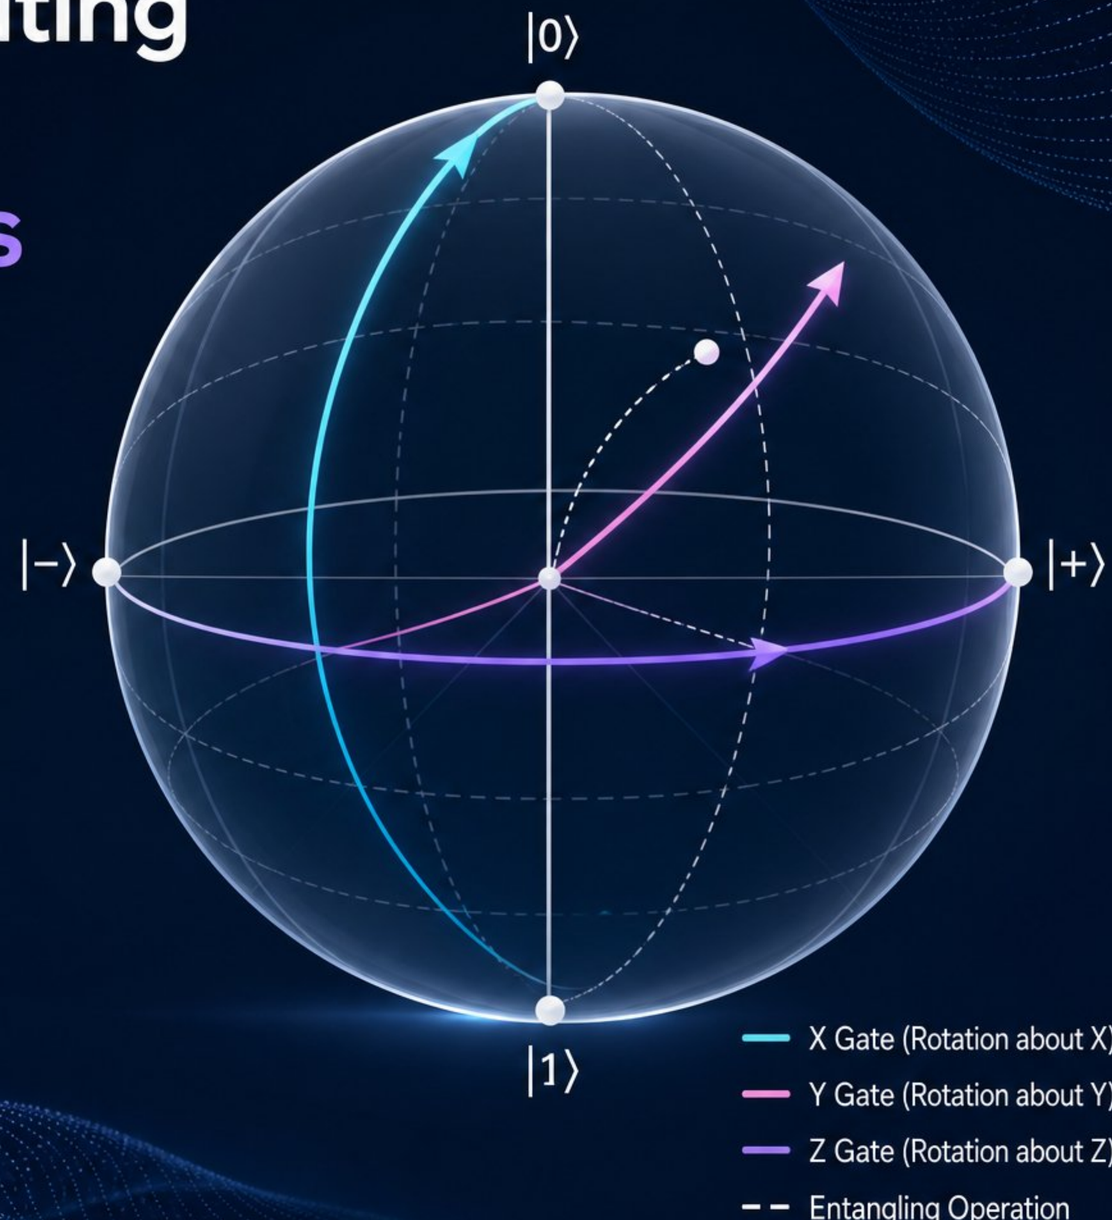

From Classical ML Concepts to Quantum Models

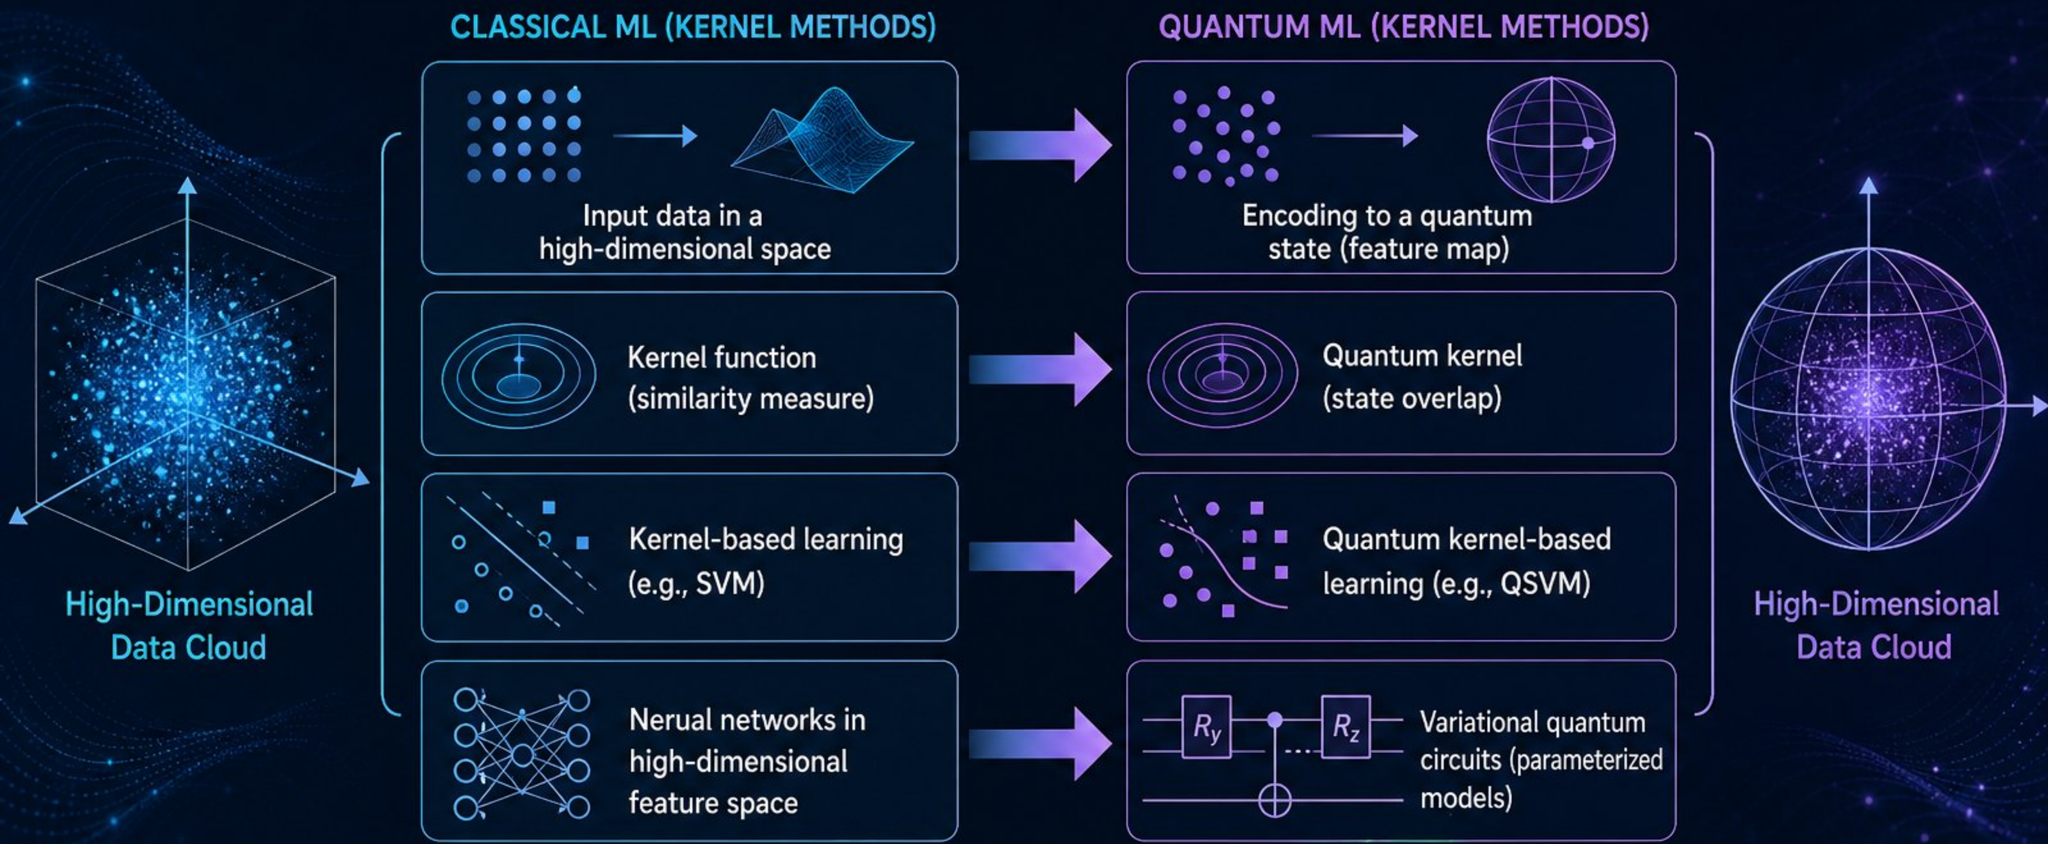

- Encoding maps classical data into quantum states, a critical first design choice
- Quantum kernels and feature maps: direct analogues to classical kernels
- Variational quantum circuits: parameterized models for learning
- Extend SVM, neural network, and kernel intuition to quantum circuits

Data Encoding and Feature Map Design

ANGLE ENCODING

Data embedded as rotation angles on qubits.

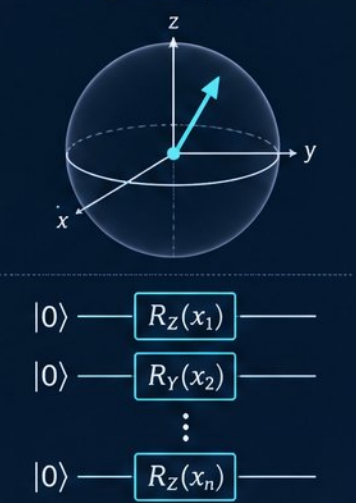

AMPLITUDE ENCODING

Data embedded in the amplitudes of a quantum state.

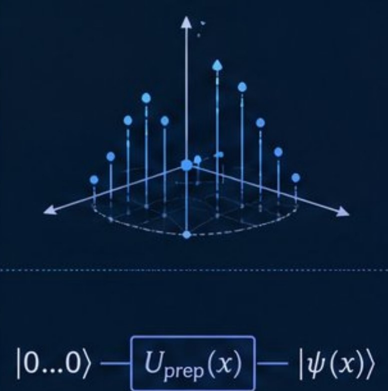

DATA RE-UPLOADING

Data encoded iteratively across multiple layers.

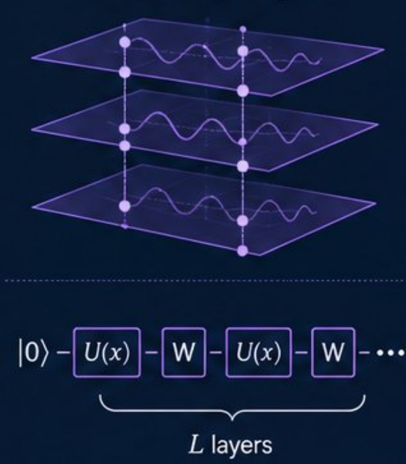

- Angle, amplitude, and basis encoding offer fundamental capacity–depth tradeoffs
- Encoding choice controls expressivity, trainability, and noise vulnerability
- Feature map structure dictates the model's learnable frequency spectrum
- Correlation-aware maps build data dependencies into quantum circuits
- 2026 advances include iterative and shot-based alternatives to gate-heavy methods

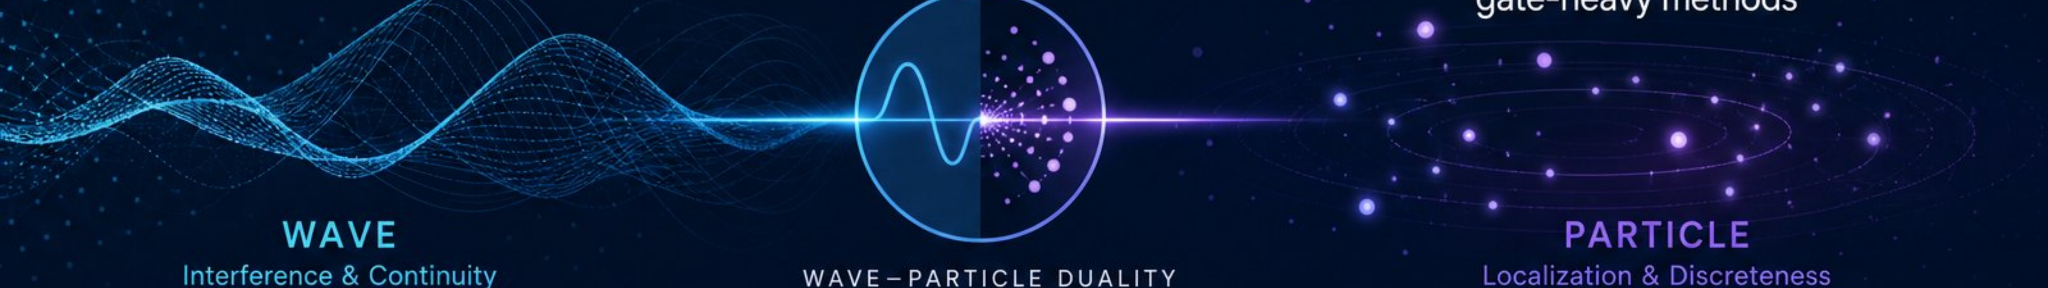

Core QML Algorithm Families

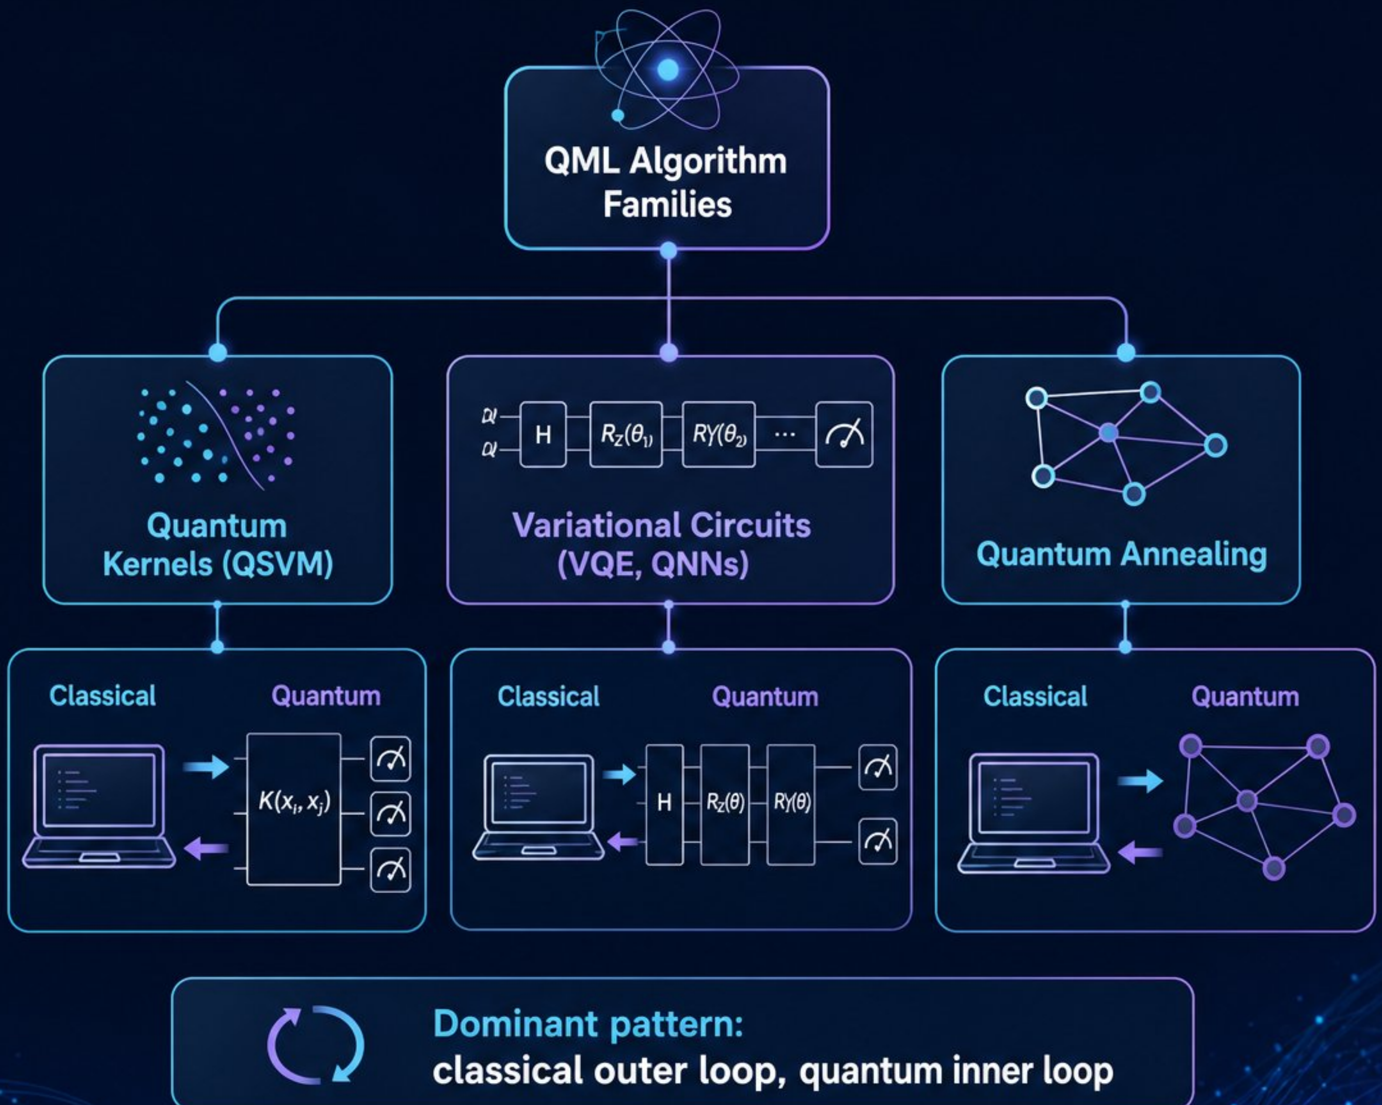

- Quantum kernels (QSVM) embed data in Hilbert space; train classically
- Variational circuits (VQE, QNNs) train via hybrid classical-quantum loops
- Quantum kernel training via kernel-target alignment: best results, ~2,000× overhead
- Quantum annealing maps learning tasks to Ising models
- Dominant pattern: classical outer loop, quantum inner loop

BENCHMARK SCOREBOARD

Benchmarks and Realistic Performance Expectations

- 970-experiment study: no significant QSVM advantage on tabular data
- Classical kernels dominate 6/7 datasets; one small dataset favored quantum
- QKT matched classical accuracy but at ~2,000× computational overhead
- Reported quantum wins often shrink with tuned baselines and matched search budgets

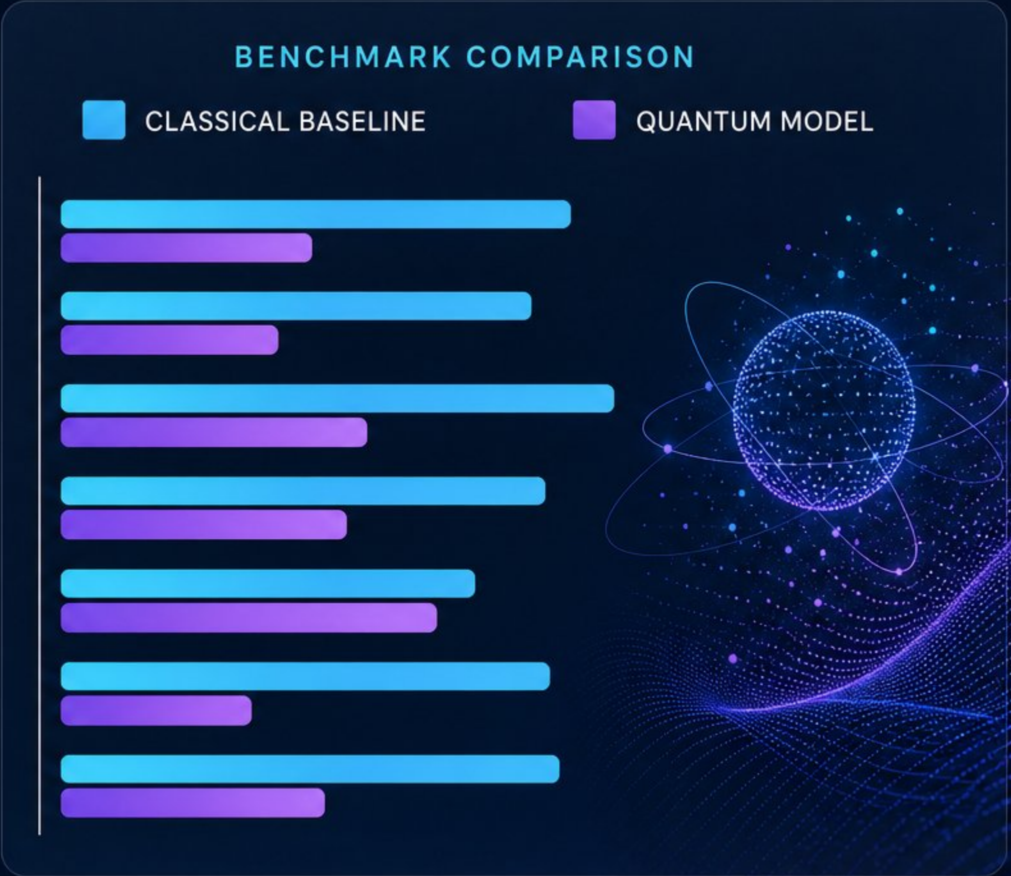

Training and Optimization Challenges

- Barren plateaus: exponentially vanishing gradients with qubit count
- Risk: random initialization, deep circuits, global cost functions
- Mitigated by local costs, structured ansätze, and informed initialization
- Noise-induced plateaus persist even with local costs in NISQ hardware
- Finite-shot noise demands high measurement budgets for stable gradients

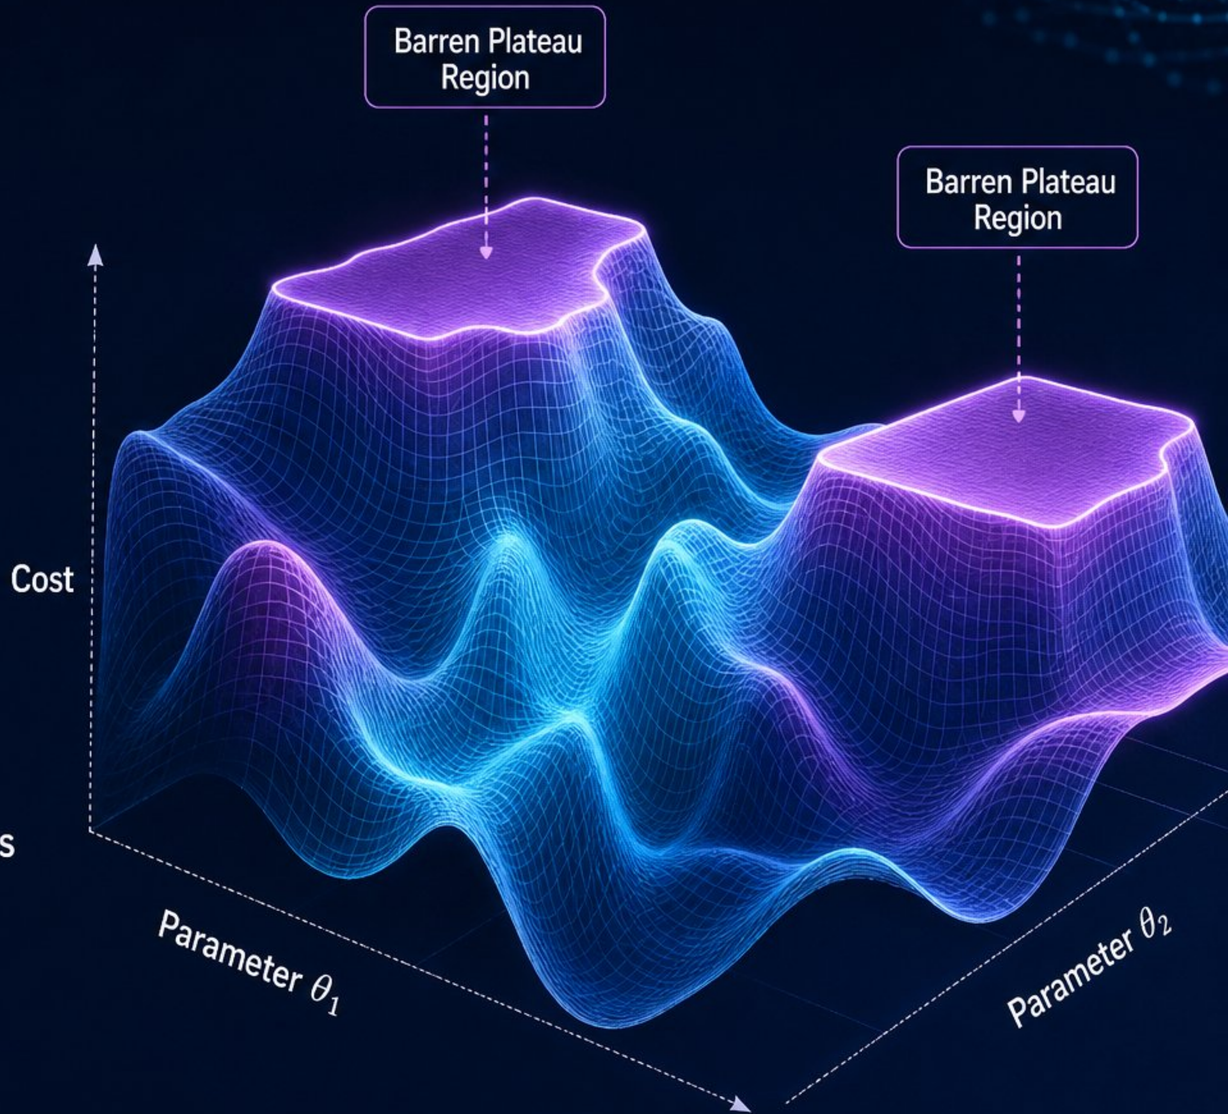

Tools, Frameworks, and Simulation Environments

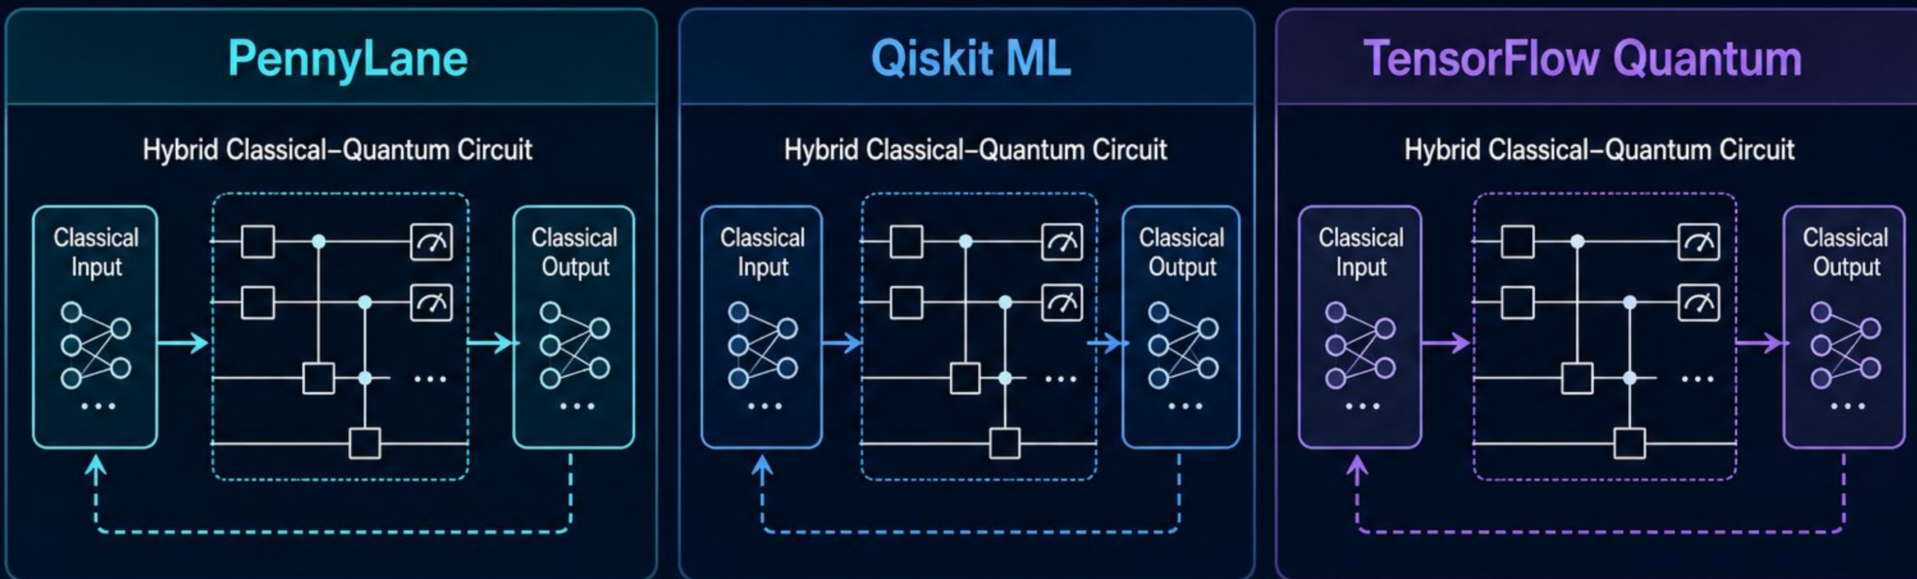

- PennyLane: hybrid experimentation, differentiable programming, multi-backend flexibility
- Qiskit ML: natural fit for Qiskit-centered, IBM-adjacent workflows
- TensorFlow Quantum: specialized choice for TensorFlow-native teams
- Simulator-first workflows dominate; hardware access remains limited
- Pick the framework that matches your existing stack and next project

Case Studies and Application Patterns

- Medical imaging: hybrid CNN-QML matches deep CNNs (~93–98%) and improves data-scarce performance
- HEP event classification: quantum kernels reach 99.99% recall in ideal simulations; noise attenuates but preserves gains
- Sports analytics: QSVM competitive with classical SVM (AUC 0.915 vs 0.942)
- Fraud detection: apparent quantum advantage in clustering disappears under budget-matched benchmarks
- Criteria: quantum-native data, scarce data, or structured problems with meaningful geometry

Reproducibility and Experimental Discipline

- Quantum kernels: 264 configurations vs. 105 classical; all effects below 0.013 ARI
- 29 paired tests in 970-experiment benchmark: no significant quantum advantage on tabular data
- Document backends, seeds, datasets, and circuit configs to tame stochasticity
- Compare against tuned classical baselines with matched search effort
- Report shot counts, noise models, and kernel-fidelity limits

Practical Adoption Strategies for ML Teams

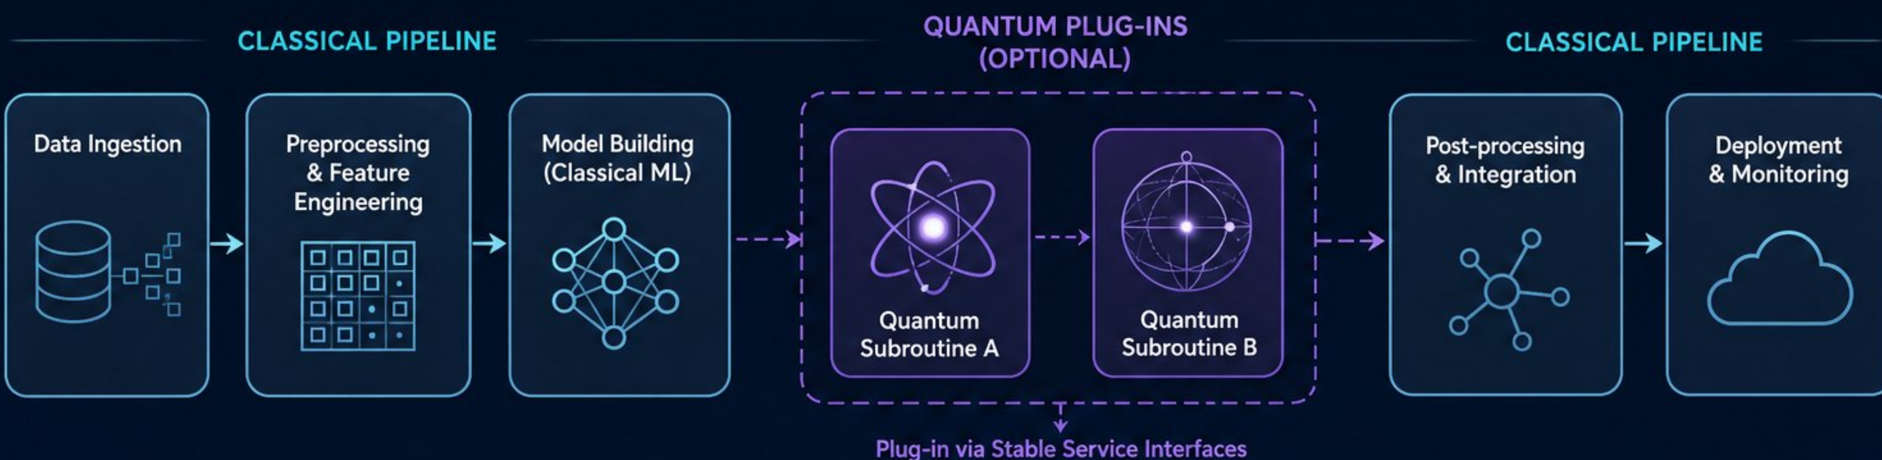

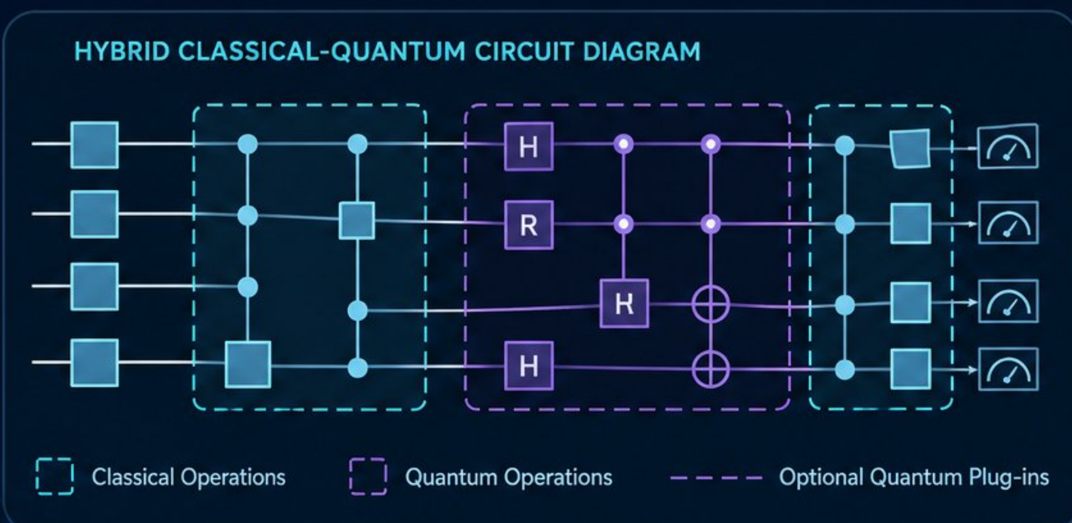

- Hybrid pipelines augment, not replace, classical ML
- Begin with narrow quantum subroutines and clear baselines
- Keep quantum components behind stable service interfaces
- Gate pilots on go/no-go metrics, not novelty

Open Problems and Research Directions

- Trainability, expressivity, and classical simulability remain coupled tensions
- Error correction and hardware roadmaps will shape ML workloads
- Quantum data and structured problems are the nearest advantage frontier
- QML parallels early deep learning's hardware-limited trajectory
- NISQ-to-fault-tolerant transition decides practical utility

Key Takeaways and Decision Framework

- Adopt QML for quantum-native data and structured problems
- Integrate as targeted hybrid components, not broad replacements
- Benchmark against strong, tuned classical models
- Build capability via simulator-first, small pilots
- Define clear go/no-go metrics before scaling

PersonWise.

PersonWise • AI digital-human course platform

Scan the code or click anywhere to watch this course live, with voice Q&A

personwise.ai/t/b3dd2b80/public