What do AI practitioners need to know about Quantum AI?

HIBA EL OIRGHI IAN WHITEHOUSE SARAH H. MILLER
PAUL LOPATA WILLIAM REGLI

University of Maryland, College Park
eloirghi@umd.edu

Bluesky

Motivation & Contributions

To make Quantum AI accessible and actionable for AI professionals and help empower them to understand the value of quantum and hybrid approaches as the field continues to progress

1. **The Quantum AI Technology Readiness Levels (QAITRLs)** framework, adapting TRLs in aerospace and defense to the challenges of Quantum AI. QAITRLs is a comprehensive maturity index enabling practitioners to situate any given quantum advantage claim along a transparent, evidence-based spectrum.
2. **A parallel glossary between quantum and classical machine learning concepts**, emphasizing that QAI is best understood as a continuum with, rather than a replacement for, modern machine learning.

The Quantum–Classical Parallel Glossary

Quantum Concept	**Classical / Deep Learning Analog**	**Implication for Practitioners**
Hamiltonian, H (energy operator)	Loss function, energy-based model	Designing H is analogous to designing a loss; poor Hamiltonians lead to bad minima.
Ground state of H	Global minimum of loss	Evolving/annealing to a (family of) ground state(s) parallels training to find a global optimum.
Unitary operator, U	Normalization mechanism (e.g, orthogonal matrices in RNNs)	When U acts on a quantum state it preserves the L_2 norm ($U^\dagger U = UU^\dagger = I$); note these are reversible transformations and can be visualized as length-preserving rotations in complex vector space.
Parameterized unitary, $U(\theta)$	Linear layer with weights	VQCs and QNNs without mid-circuit measurements are structured like purely linear versions of deep nets with θ as learnable weights; note that (parameterized) unitaries are fundamentally linear, requiring measurement to realize any non-linearity.
Measurement / observable	Readout layer / logits + non-linear activation	Measurement defines what the model predicts; choice of observable is analogous to output heads; note this is a primary source of non-linearity in otherwise often linear quantum transformations.
Barren plateau	Vanishing gradient	Training pathologies are shared: both quantum and classical networks can suffer from flat landscapes. For a more detailed comparison see McClean et al. [83]'s work.

QAITRLs

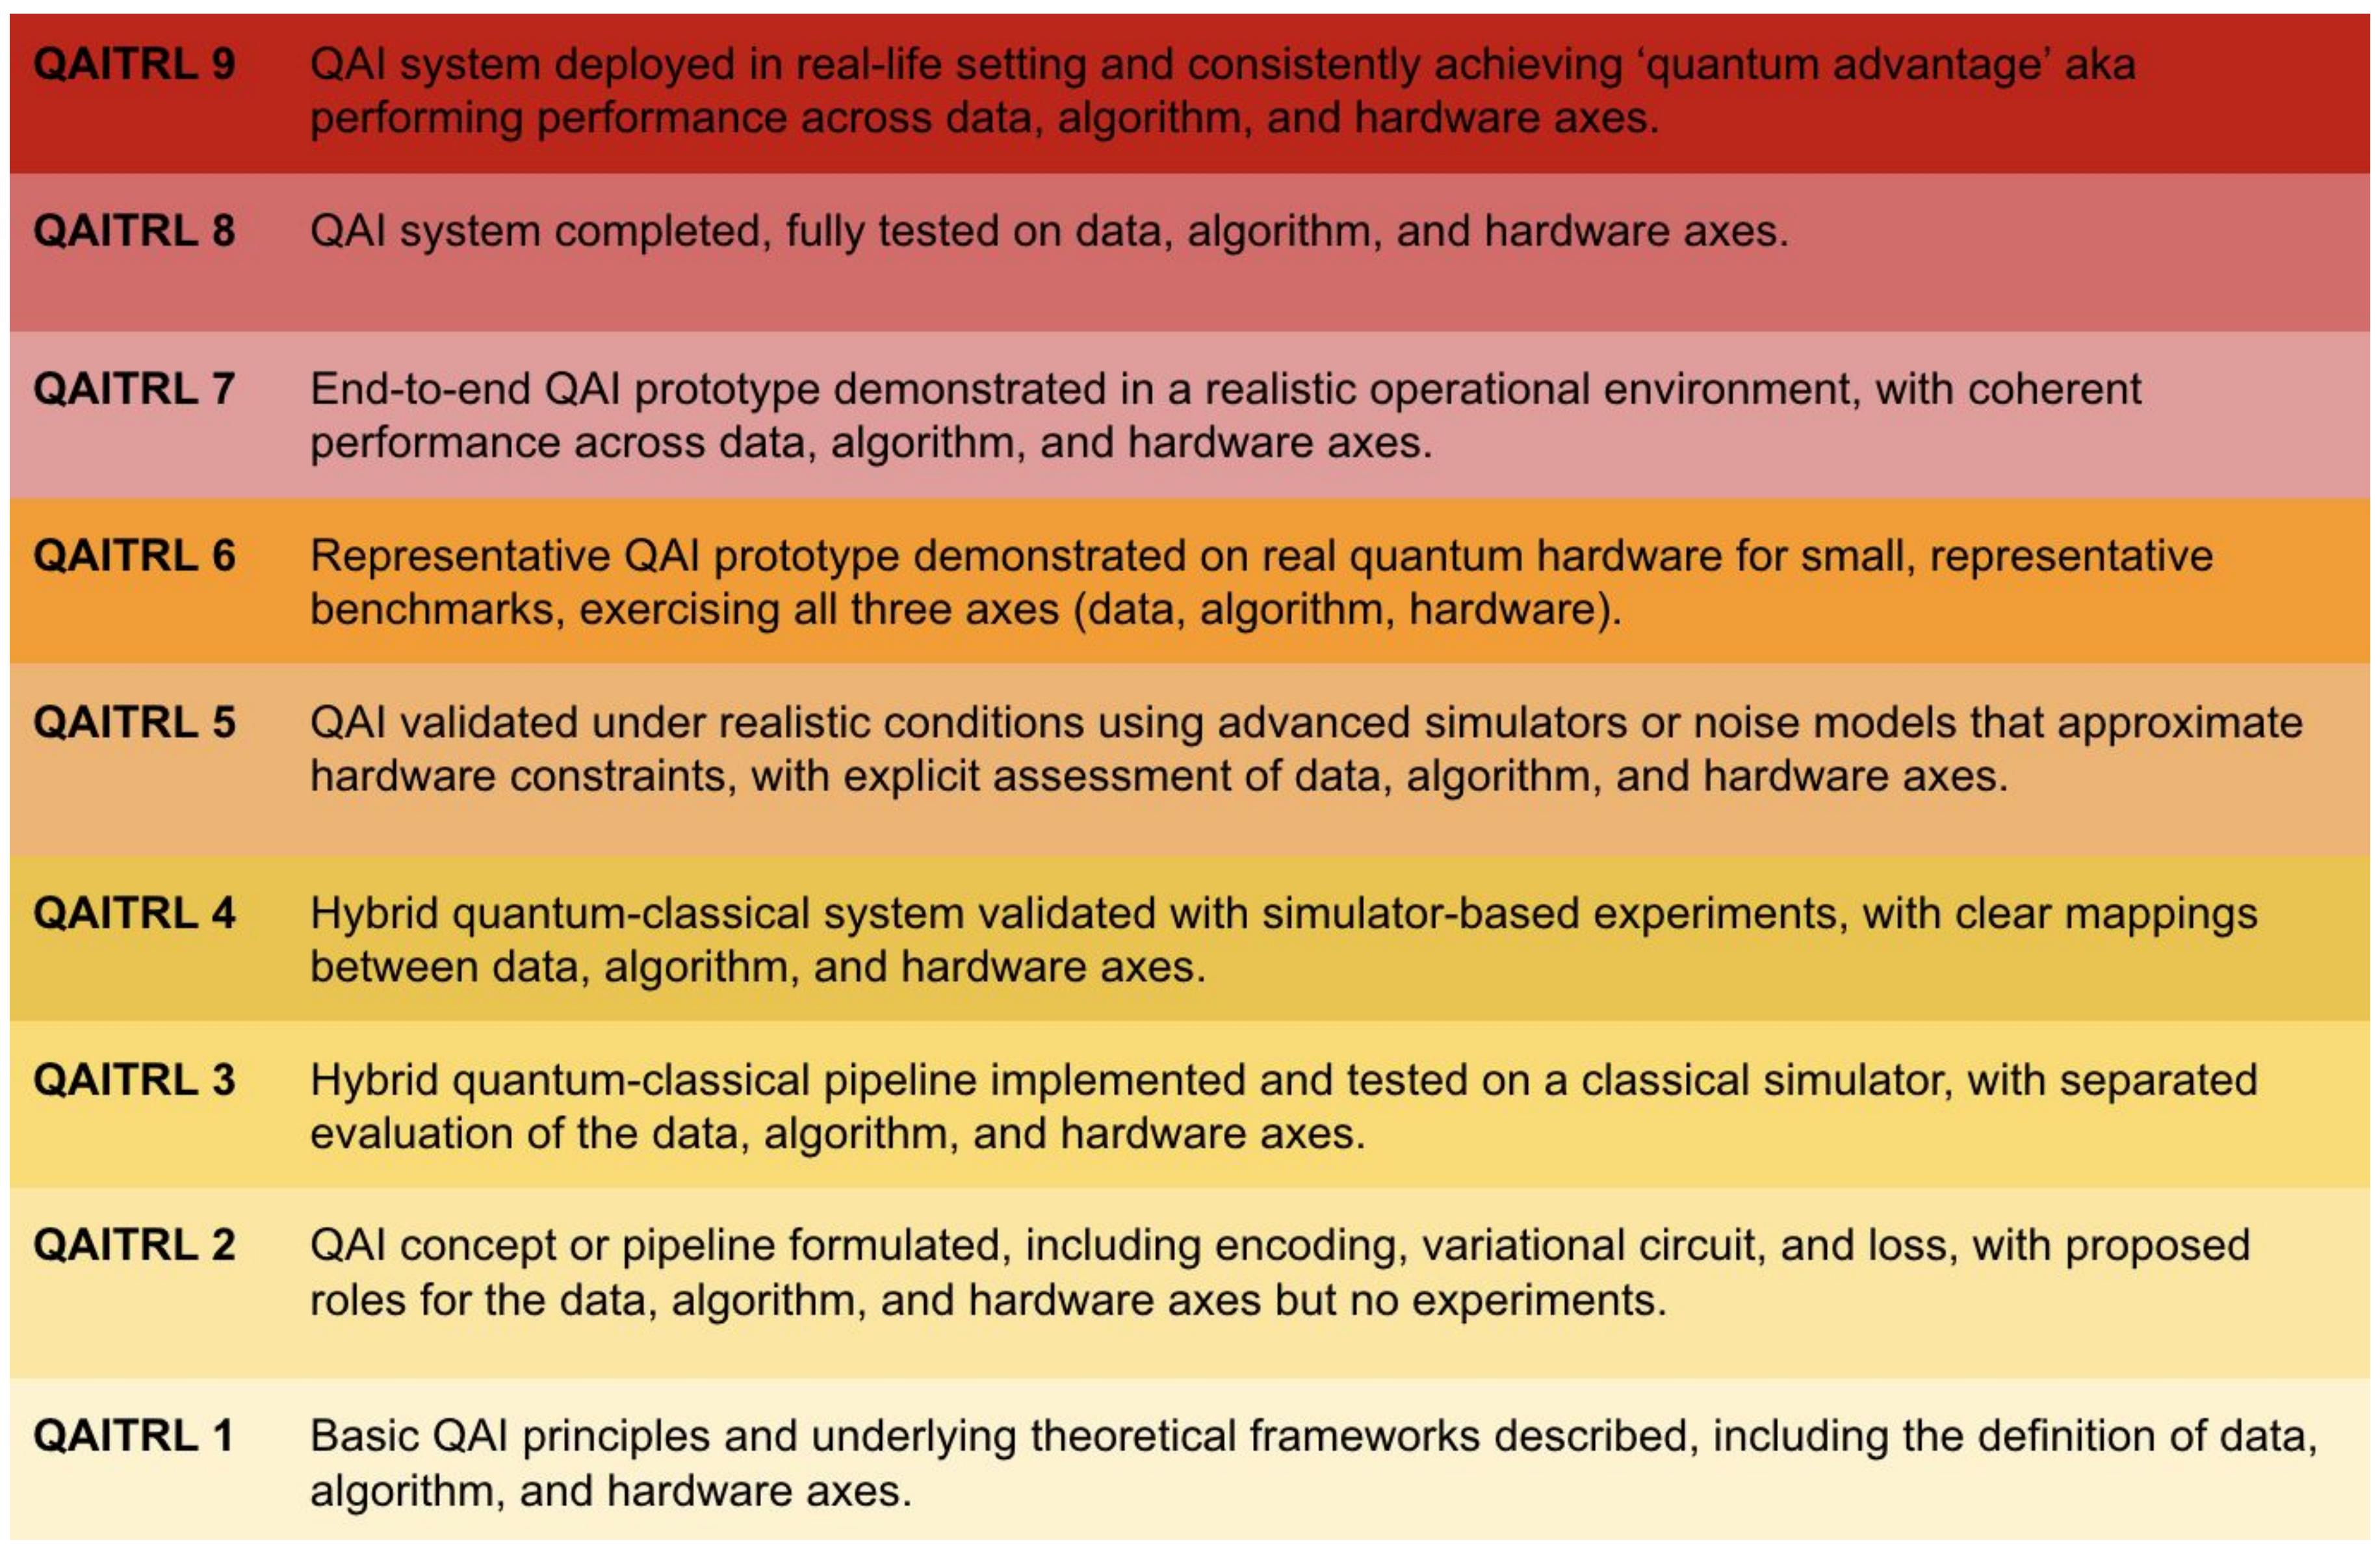

QAITRL 9	QAI system deployed in real-life setting and consistently achieving 'quantum advantage' aka performing performance across data, algorithm, and hardware axes.
QAITRL 8	QAI system completed, fully tested on data, algorithm, and hardware axes.
QAITRL 7	End-to-end QAI prototype demonstrated in a realistic operational environment, with coherent performance across data, algorithm, and hardware axes.
QAITRL 6	Representative QAI prototype demonstrated on real quantum hardware for small, representative benchmarks, exercising all three axes (data, algorithm, hardware).
QAITRL 5	QAI validated under realistic conditions using advanced simulators or noise models that approximate hardware constraints, with explicit assessment of data, algorithm, and hardware axes.
QAITRL 4	Hybrid quantum-classical system validated with simulator-based experiments, with clear mappings between data, algorithm, and hardware axes.
QAITRL 3	Hybrid quantum-classical pipeline implemented and tested on a classical simulator, with separated evaluation of the data, algorithm, and hardware axes.
QAITRL 2	QAI concept or pipeline formulated, including encoding, variational circuit, and loss, with proposed roles for the data, algorithm, and hardware axes but no experiments.
QAITRL 1	Basic QAI principles and underlying theoretical frameworks described, including the definition of data, algorithm, and hardware axes.

QAITRLs are defined along three axes that reflect the core concerns of an AI practitioner:

Complexity-Theoretic Foundation (A1): Is there a well-defined quantum algorithm or architecture that targets a specific AI workload (e.g., kernel evaluation, combinatorial optimization, feature extraction), grounded in formal asymptotic analysis or complexity-class separations (BPP vs. BQP)? This axis evaluates whether potential quantum advantage has rigorous theoretical justification.

Physical Realization & Hardware Fidelity (A2): Has the method been instantiated on realistic hardware, explicitly addressing constraints of the NISQ era? This includes finite qubit counts, connectivity constraints, gate error rates, decoherence times, and measurement noise. This axis ensures physical feasibility.

Statistical Validity & Benchmarking (A3): Has the method been rigorously evaluated using objective metrics (accuracy, geometric difference, TEI, BPI) and compared to state-of-the-art classical baselines (SVMs, deep networks, specialized solvers) on realistic datasets with transparent reporting? This axis ensures empirical validation with reproducible metrics.

Practitioner Decision Framework

Step 1: Is your workload classically hard? If your problem is solved well by existing classical methods (neural nets, gradient boosting, specialized solvers) with reasonable computational budget, stop here. Quantum approaches are unlikely to help in the near term.

Step 2: What is the data source? If data is natively quantum (molecular properties, quantum state tomography, materials simulation), proceed to Step 3a. If data is classical, proceed to Step 3b.

Step 3a (Quantum-native data): You are in the most favorable regime. Consider quantum simulation approaches (QAITRL 5–6 maturity). Investigate whether hybrid classical-quantum pipelines like Robledo-Moreno et al. [102] apply to your domain.

Step 3b (Classical data): Is your dataset small enough (roughly ≤1,000 samples, ≤20 features) to fit on current NISQ hardware without major dimensionality reduction? If yes, proceed to Step 4. If no, quantum approaches are likely premature for your workload, given current hardware.

Step 4: Problem structure. Does your problem reduce to a sparse QUBO/Ising formulation? Consider quantum annealing (QAITRL 5 maturity, but classical solvers are competitive) or emerging methods incorporating measurement induced phase transitions into QNN-like architectures (QAITRL 3-4, but may efficiently capture complexity impossible for classical methods). Does your problem involve kernel evaluation where classical kernels fail? Consider QSVT and quantum kernels (QAITRL 4–5). Otherwise, consider promising VQC/QNN approaches (QAITRL 3–4) with the understanding that these are exploratory research investments, not production-ready tools, and be particularly leery of dequantization.

Step 5: Timeline and risk tolerance. If you need production-ready results within 1–2 years, classical methods are likely most appropriate. If you are investing in 3–5 year research capabilities and can absorb the risk of a quantum approach not panning out, targeted prototyping is reasonable and may return tangible benefits over time.