


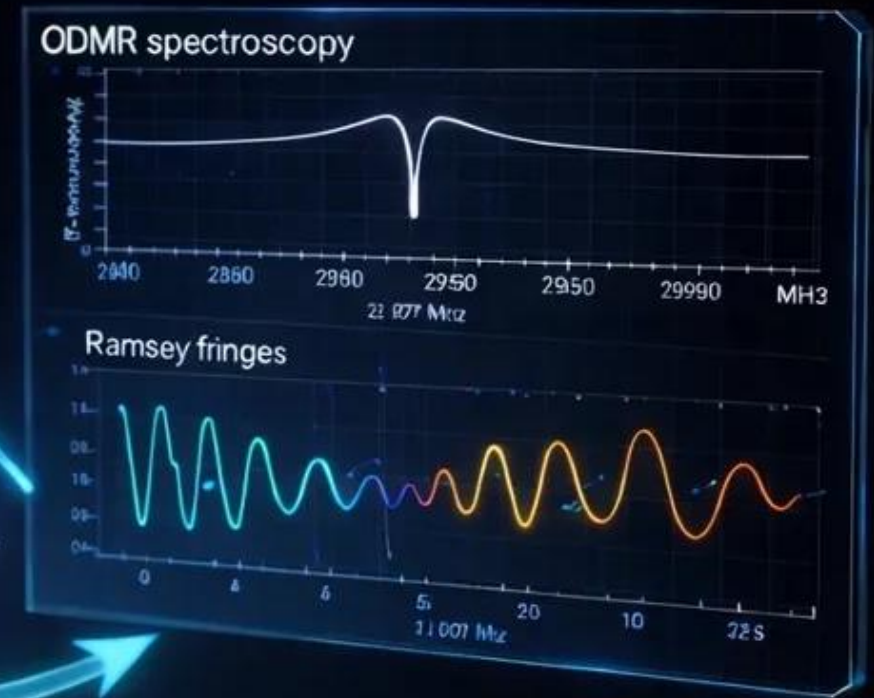
When does adaptation win? Scaling laws for meta-learning in quantum control

Nima Leclerc¹, Chris Miller, and Nicholas Brawand

¹Principal Investigator and Research Scientist | Adaptive Quantum Sensing Program


The MITRE Corporation | Boston, MA USA

International Conference on Machine Learning 2026 | Seoul, South Korea

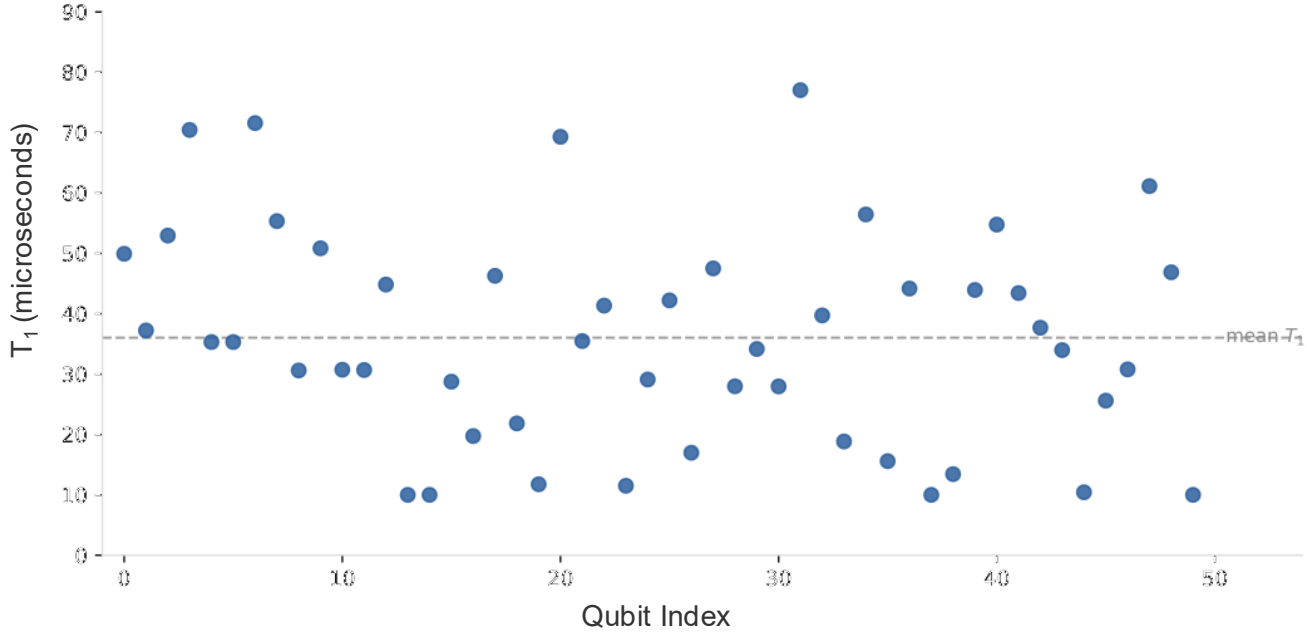


The calibration bottleneck: fixed controllers become stale

Qubit parameters vary across the chip and drift over time, so a controller tuned to yesterday's or the average device is quickly off-target.



Device-wide heterogeneity + drift across a quantum processor



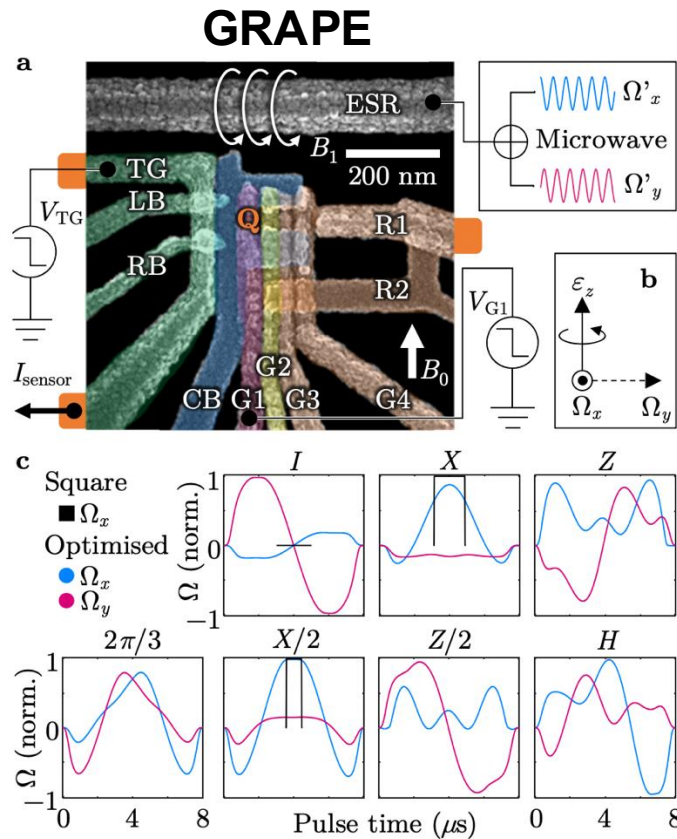
Mean calibration misses many qubits

Need: adaptive control, not one-size-fits-all calibration

Use a few calibration measurements to infer the current drift state, then update the controller for that device — avoiding both fixed-controller failure and full re-optimization from scratch.

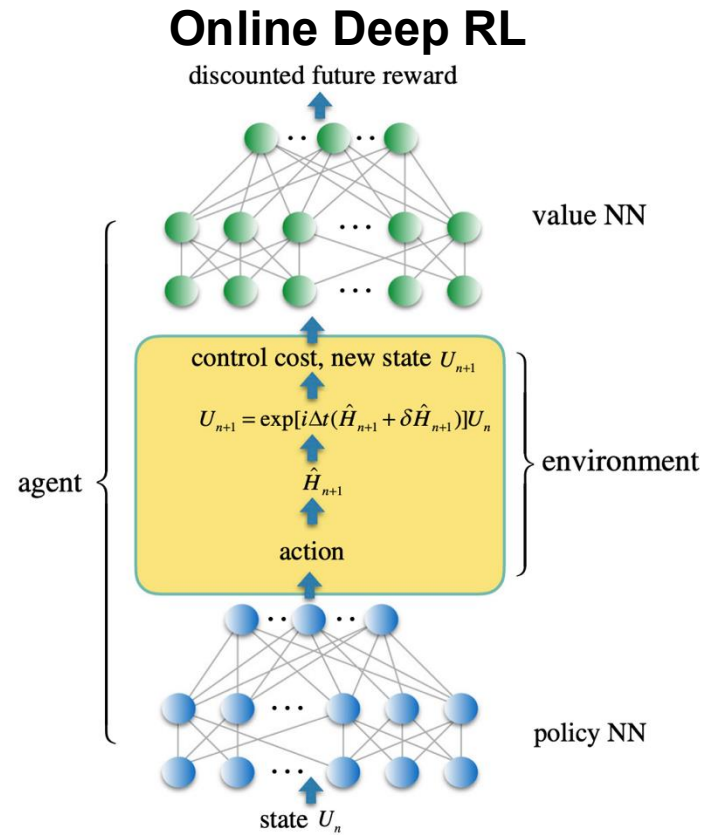
This work asks: when does adaptation beat a robust fixed controller, and how many adaptation steps are worth paying for?

Existing AI approaches to quantum control — and why meta-Learning is different



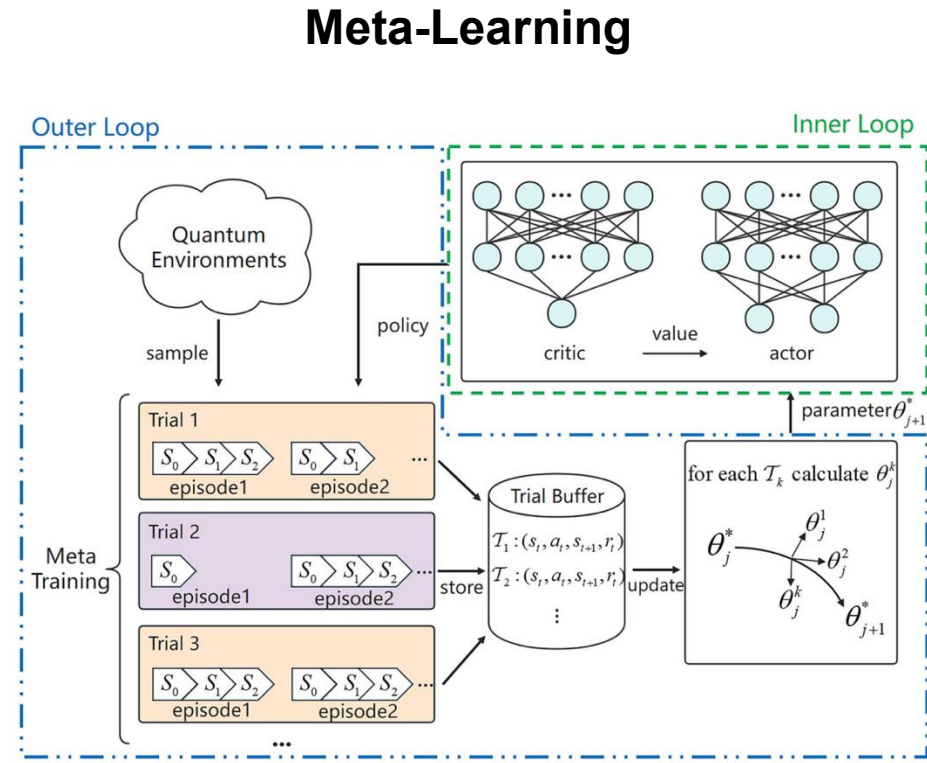
Yang et al., *Nat. Electron.* **2**, 151–158 (2019)

Optimal pulses for a fixed Hamiltonian. Fast at deployment but requires full re-optimization when device parameters drift.



Niu et al., *npj Quantum Inf.* **5**, 33 (2019)

Adapts in real time, but trains from scratch on each device. Sample-inefficient for hardware with limited coherence time.



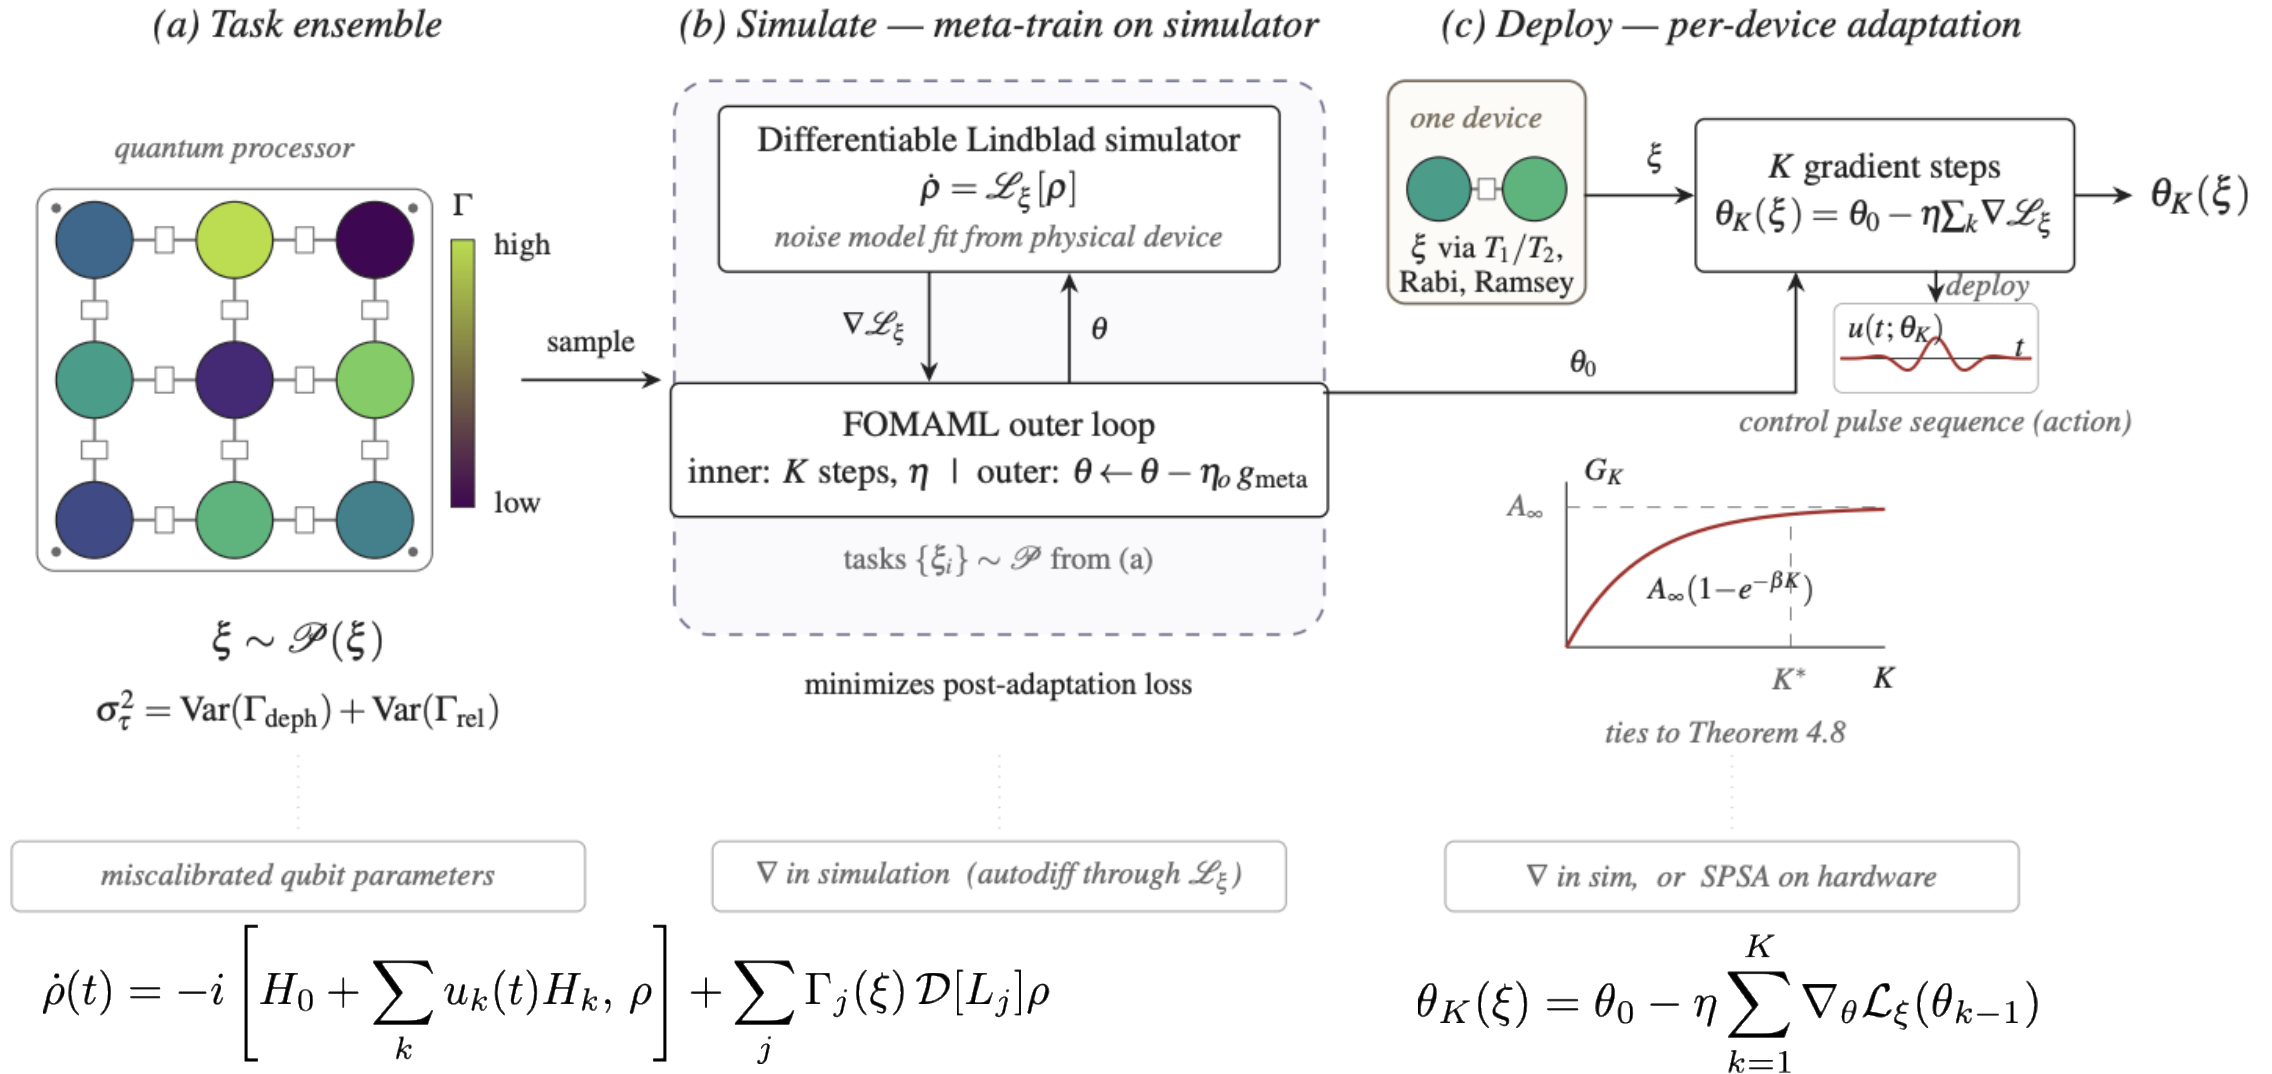
Zhang et al., *npj Quantum Inf.* **11**, 81 (2025)

Learns to learn: meta-trains across device environments so the agent adapts to new control tasks in a few steps without retraining.

Adaptation is inevitable but its cost is not. This work asks: how much is enough?

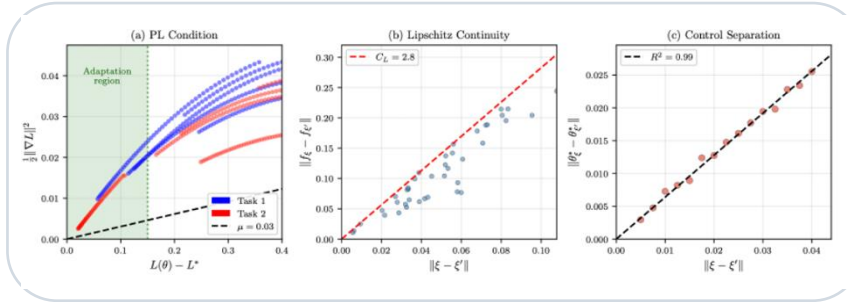
Simulate once, deploy everywhere: the digital twin for quantum control

We meta-train once on a differentiable Lindblad simulator (digital twin), then deploy with K device-specific steps. The central question: *how many steps K do you need, and when does adaptation justify its cost?*



From optimization geometry to adaptation gap scaling

1 Geometry checks



PL-like progress: gradients improve loss predictably near optima.

Smooth task dependence: nearby device parameters induce nearby dynamics.

Control separation: different devices require different optimal pulses.

2 Scaling-law consequence

$$G_K = \mathbb{E}_{\xi \sim P} [\mathcal{L}_\xi(\theta_0) - \mathcal{L}_\xi(\theta_K(\xi))]$$

local PL geometry gives finite-step improvement

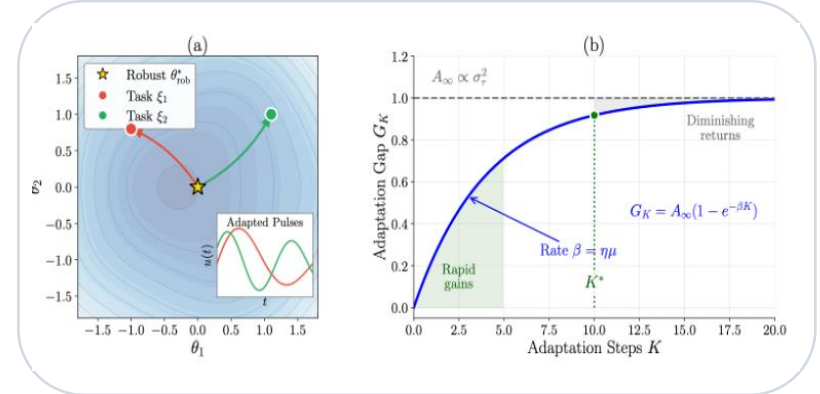
$$\text{PL geometry} \Rightarrow G_K \approx A_\infty (1 - e^{-\beta K})$$

ceiling and rate are interpretable

$$A_\infty \propto \sigma_\tau^2, \quad \beta = \eta \mu_{\min}$$

Task variance sets payoff; curvature sets speed.

3 Adaptation gap intuition



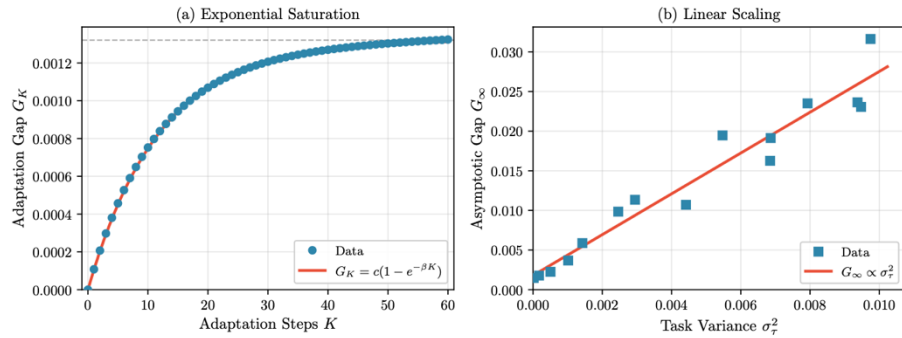

Robust control: average-optimal, not device-optimal.

Adaptation: rapid early gains; K^* marks diminishing returns.

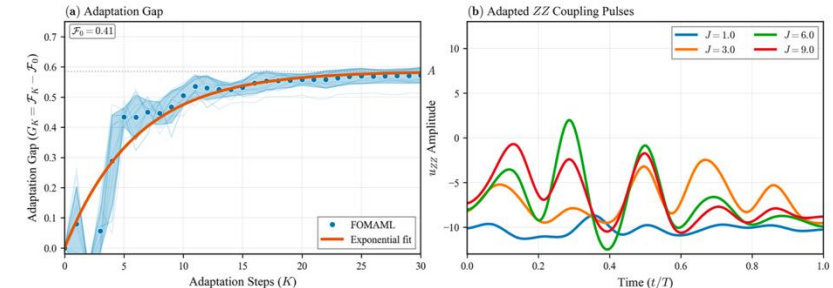



The scaling law follows from optimization geometry — not quantum-specific physics.

Validation: the scaling law generalizes across quantum gate primitives

Single-qubit X gate: exponential saturation and variance scaling holds.



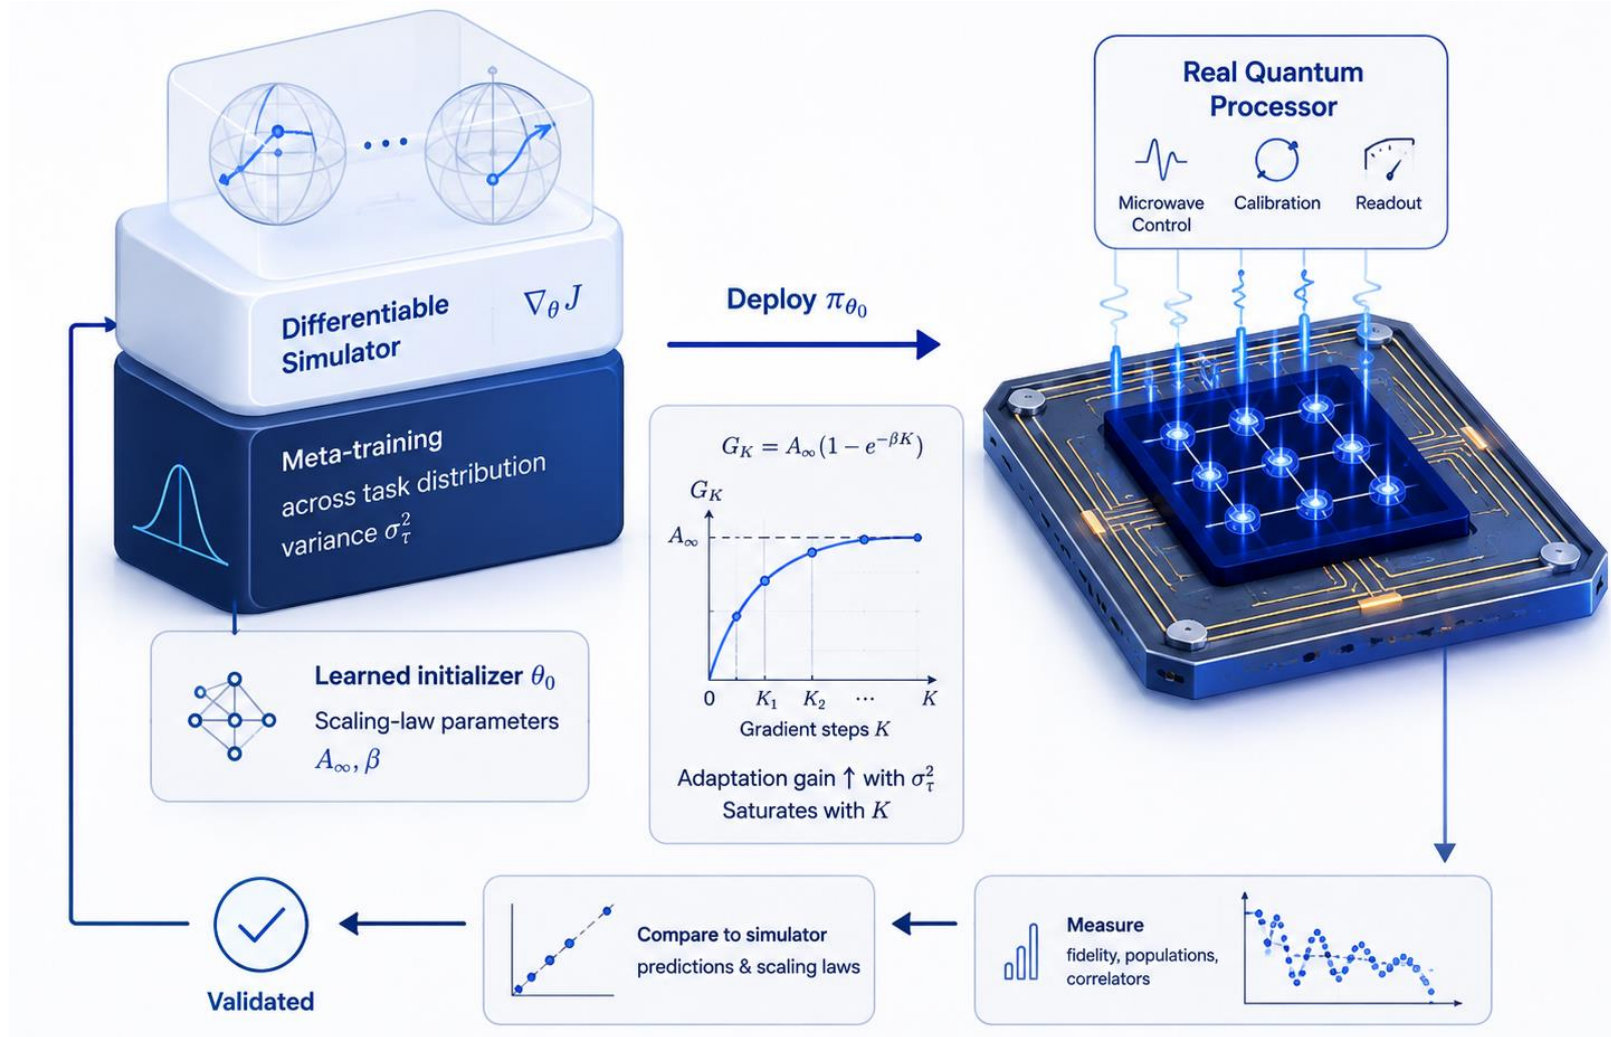
Two-qubit CZ gate: same scaling law, larger payoff.



FOMAML already starts near the task optimum, so $K = 10$ steps only marginally improve fidelity; fixed-average control benefits more when initialization is poor.

Under strong OOD noise, adaptation improves fidelity from 54.2% to 95.7%, showing that harder entangling gates are where adaptive calibration becomes valuable.

Simulate, then deploy in practice with meta-learning



Hardware validation turns the theory from a simulation result into a calibration policy: adapt only when device variance is high enough to justify the cost.

From scaling laws to calibration budgets

Few-shot decision rule

Probe $N = 3-5$ steps, fit the saturation curve, then choose the adaptation budget needed to reach the target gain.

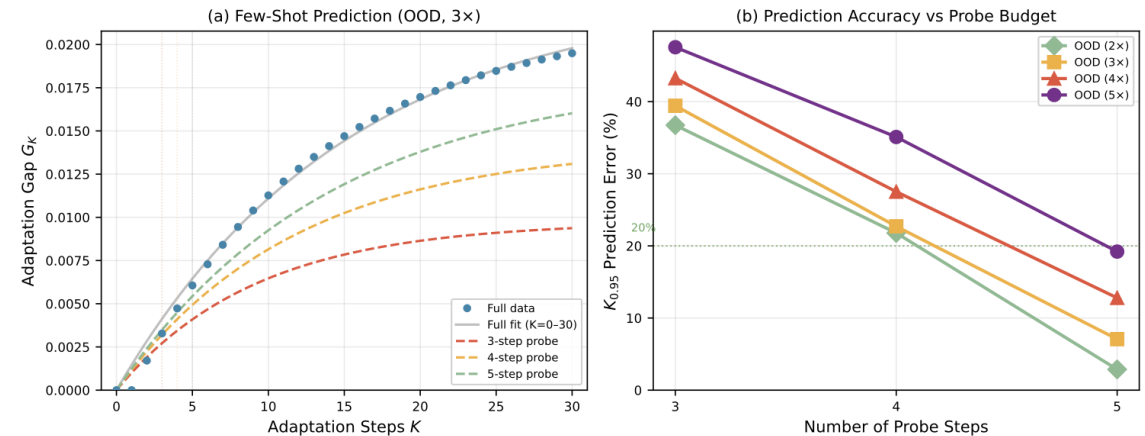
Algorithm 1 Few-Shot Pre-Adaptation Decision Protocol

Require: θ_0 , probe count $N \in \{3, 4, 5\}$, target gain fraction α (e.g., 0.95), benefit threshold ϵ , representative task ξ

- 1: Initialize $\theta^{(0)} \leftarrow \theta_0$
- 2: **for** $k = 1, \dots, N$ **do**
- 3: $\theta^{(k)} \leftarrow \theta^{(k-1)} - \eta \nabla_{\theta} L_{\xi}(\theta^{(k-1)})$ (inner-loop step)
- 4: $G_k \leftarrow L_{\xi}(\theta_0) - L_{\xi}(\theta^{(k)})$ (adaptation gap)
- 5: **end for**
- 6: Fit $G_k \approx \hat{A}_{\infty}(1 - e^{-\hat{\beta}k})$ to $\{(k, G_k)\}_{k=1}^N$, record R^2
- 7: **PL diagnostic:** if $R^2 < 0.9$, flag departure from local PL
- 8: Compute budget $\hat{K}_{\alpha} = \frac{1}{\hat{\beta}} \log \frac{1}{1-\alpha}$ (Corollary 4.9)
- 9: **return** don't adapt if $\hat{A}_{\infty} < \epsilon$; else *adapt with budget \hat{K}_{α}*

Accurate from few probes

With $N = 5$, the protocol estimates $K_{0.95}$ within 3-19% across $2 \times -5 \times$ OOD regimes.

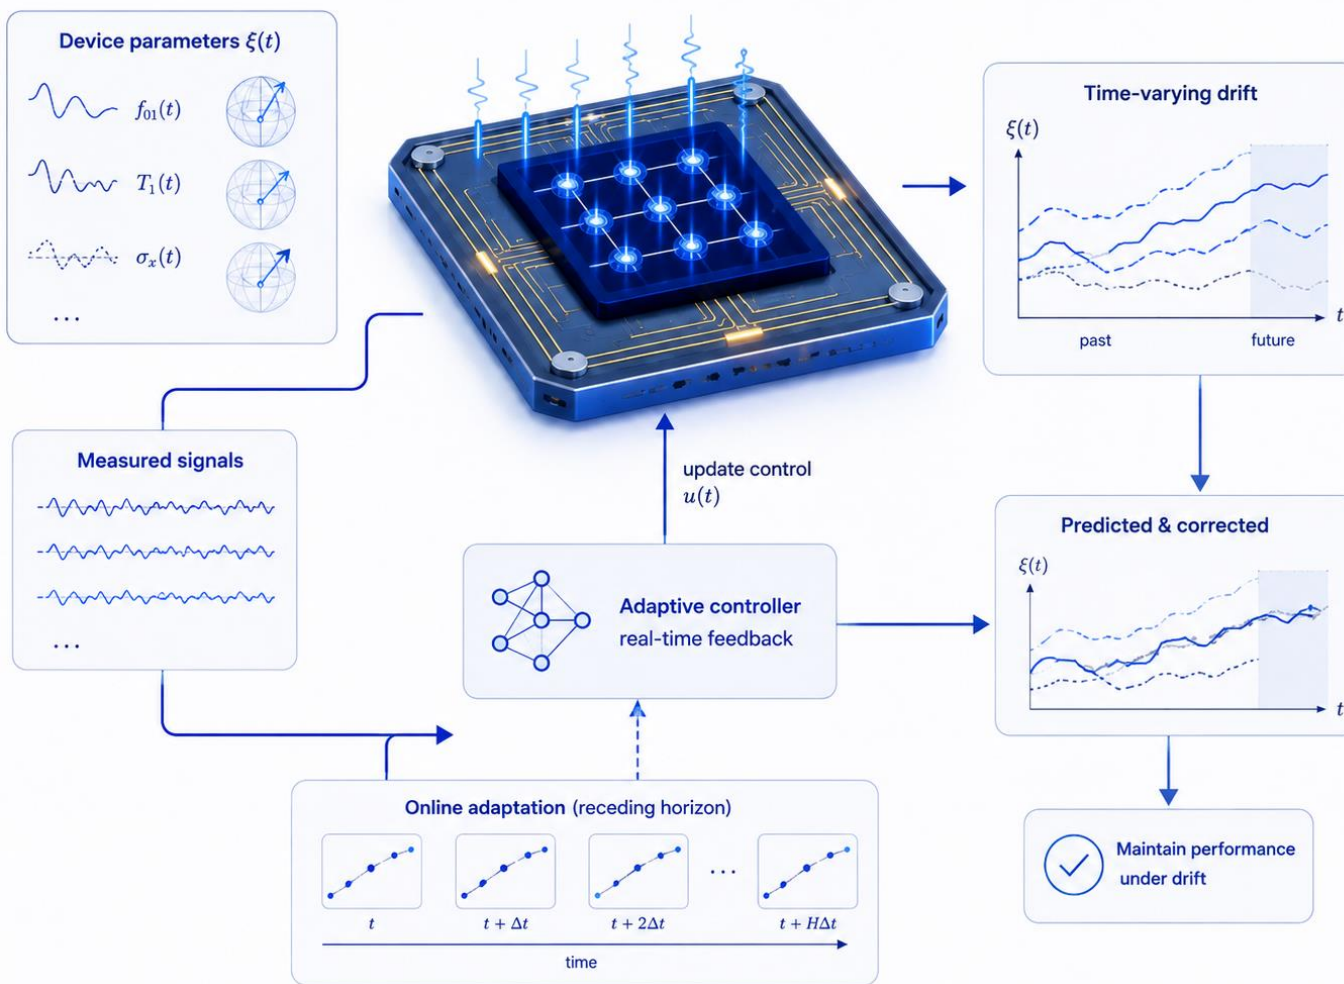


Regime	$K_{0.95}$	$N=3$	$N=4$	$N=5$
OOD (2x)	43.3	36.7%	21.8%	2.9%
OOD (3x)	45.4	39.4%	22.7%	7.1%
OOD (4x)	48.4	43.2%	27.5%	12.8%
OOD (5x)	52.4	47.5%	35.1%	19.2%


Adapt only when the predicted gain is large; stop once returns saturate.

Outlook: non-stationary control and many-qubit scaling

Non-stationary control



Many-qubit scaling



As processors scale, per-qubit heterogeneity decreases but total adaptation cost grows:

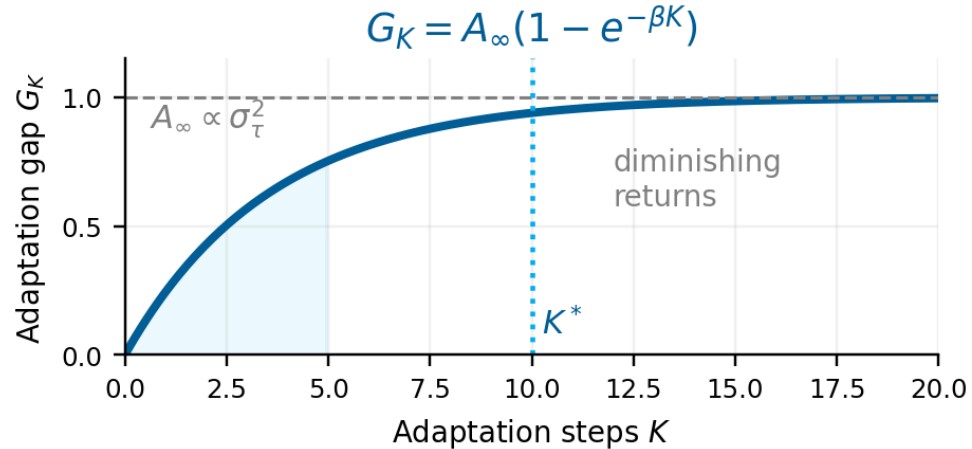
$$\sigma_{\tau}^2 \sim N_{\text{qubits}}^{-\alpha}, \quad K^* \sim K^*(N_{\text{qubits}})$$

Larger processors require more adaptation steps but the per-qubit payoff shrinks. The scaling law predicts the calibration budget before running a single experiment.

When Does Adaptation Win?

Scaling Laws for Meta-Learning in Quantum Control

Task variance sets the payoff. Curvature sets the speed.



Core idea

- Adaptation gain saturates exponentially with gradient steps.
- Maximum gain scales with task variance across devices.
- A few probe steps estimate the useful calibration budget.

Nima Leclerc, Chris Miller, Nicholas Brawand

www.nleclerc.com | nleclerc@mitre.org

ICML 2026



Paper

[arxiv.org/abs/
2601.18973](https://arxiv.org/abs/2601.18973)



Code

[github.com/nimalec/
metalearning-
quantum-control](https://github.com/nimalec/metalearning-quantum-control)