Risk-Controlled Regime-Adaptive Optimization for Variational Quantum Algorithms

Michael Paley

St Paul's School, London SW13 9JT, United Kingdom

23 November 2025

Abstract

Variational quantum algorithms can become difficult to train when barren plateaus suppress gradients as system size grows. Lie-algebraic theory helps explain when these regions arise, but an optimizer must still decide, from finite and noisy measurements, when to explore and when to exploit. We develop a measurement-efficient scheduler with explicit, anytime-valid risk control. It combines a χ^2-calibrated sequential test for distinguishing plateaus from informative regions, a probabilistic trust-region strategy with one-sided acceptance under shot noise, and a variance-aware exploitation phase. We establish geometric bounds on plateau exit times and linear convergence in informative basins, and we show that Lie-algebraic variance proxies can improve test power without disturbing statistical calibration.

1 Introduction

While theory has clarified when and why plateaus appear—including precise Lie-algebraic criteria for trainability [1]—turning these insights into practical, finite-sample optimisation strategies remains challenging. Existing approaches typically (i) attempt to avoid BPs via ansatz design, or (ii) identify them post hoc once training stalls. In contrast, the measurement-frugal line of work focuses on allocating shots adaptively during optimisation, e.g., iCANS/gCANS and related batch-sizing ideas [2, 3, 4], often coupled with parameter-shift gradient estimators [5, 6]. Yet most methods assume smoothness/curvature conditions without offering explicit sequential guarantees for deciding when to explore vs. exploit under tight budgets [7, 8, 9, 10].

Our contribution. We introduce a scheduler that combines regime detection with shot-efficient exploitation under finite-sample control. Plateau-versus-informative discrimination is formulated as a sequential hypothesis test on a whitened gradient-norm statistic, with anytime-valid error control from classical sequential analysis [11, 12]. Under H_0 (plateau), the statistic is χ^2; under H_1 (informative), it is non-central χ^2. This supports likelihood-ratio decisions with calibrated Type-I/II risks, while Welch and Clopper–Pearson components provide conservative one-sided bounds for trust-region moves [13, 14]. When H_0 is favoured, the scheduler uses probabilistic trust-region (PTR) exploration with one-sided acceptance and geometric radius updates [15]. When H_1 is favoured, it switches to gCANS-style exploitation with variance-proportional shot allocation, recovering linear, or geometric, convergence under standard local smoothness and PL conditions [3, 9, 10]. Gradients are obtained by parameter shift and measured with weighted operator sampling (WRS) for shot efficiency [5, 16].

Relation to prior work. The scheduler brings together three strands: (i) sequential testing for principled, anytime regime decisions [11, 12]; (ii) trust-region exploration with explicit statistical accept/reject rules [15, 13, 14]; and (iii) measurement-efficient exploitation

through iCANS/gCANS and batch-size/learning-rate coupling [2, 3, 4, 17]. Complementing Lie-algebraic theory that predicts where barren plateaus arise [1], the method supplies finite-sample, implementation-ready controls for deciding when to explore or exploit and how to allocate shots near the fundamental noise floor of a fixed budget.

The method alternates between regime detection and optimization. When a plateau is detected, it performs probabilistic trust-region (PTR) exploration with geometric radius expansion. When an informative region is identified, it executes gCANS-style gradient descent with variance-proportional shot allocation. Algorithm 1 summarizes one outer iteration.

Algorithm 1 Regime-adaptive scheduler (one outer iteration)

Require: current θ; risk (α, β); PTR params $(R_0, K, m, \tau, \alpha_{\text{acc}})$; step size η; exploitation rule $S_i \propto \sigma_i$

Pilot: estimate σ_i and optionally Lie proxy V_i; set exploration shots $B_i^{\text{expl}} \propto V_i/\sigma_i^2$ (default $V_i \equiv 1$)
Regime test: draw $s_t = \sum_i (B_i^{\text{expl}}/\sigma_i^2)\widehat{g}_{i,t}^2$; update log-likelihood Λ_k
if $\Lambda_k \le B$ **then** ▷ plateau
  for $k = 0, \dots, K-1$ **do**
    for $j = 1, \dots, m$ **do**
      $R \leftarrow R_0 2^k$; sample unit v; $\theta^+ \leftarrow \theta + Rv$
      Estimate $\Delta = f(\theta^+) - f(\theta)$; compute one-sided Welch UCB
      if $\text{UCB}(\Delta) \le -\tau R^2$ **then**
        $\theta \leftarrow \theta^+$; **break**
else ▷ informative
  Allocate $S_i \propto \sigma_i$ (fixed total S); compute $\widehat{g}(\theta)$
  $\theta \leftarrow \theta - \eta\, \widehat{g}(\theta)$

The pilot phase is primarily required during initial exploration when the optimization landscape structure is unknown. When available, Lie-algebraic information about the circuit generators can inform the pilot through variance proxies $V_i \propto \|[H_i, O]\|_F^2$, which identify parameters where gradient signal is more concentrated according to the DLA's purity structure. This allows the regime test to allocate exploration shots preferentially to informative directions while maintaining proper χ^2 calibration, as the whitening procedure uses the same shot allocation B_i^{expl} regardless of whether uniform ($V_i \equiv 1$) or Lie-informed weights are employed.

2 Problem Setup and Statistical Framework

2.1 Variational Quantum Optimization

Consider a parameterized quantum circuit $U(\theta) = \prod_{l=1}^{L} e^{iH_l\theta_l}$ with objective function $f : \mathbb{R}^d \to \mathbb{R}$ given by

$$f(\theta) = \text{Tr}[U(\theta)\rho U^\dagger(\theta) O], \tag{1}$$

where ρ is the input state and O is the observable. Our goal is to minimize $f(\theta)$ using gradient-based methods under shot-limited measurements.

2.2 Shot Noise Model

Gradient estimation via parameter-shift or similar schemes produces noisy measurements. We formalize this setting through the following assumptions.

Assumption 1 (Shot noise and coordinate structure). For shot allocation $\mathbf{B} = (B_1, \dots, B_d)$, the gradient estimator $\widehat{g}(\theta)$ satisfies three properties. First, it is unbiased: $\mathbb{E}[\widehat{g}_i(\theta)] = \nabla f_i(\theta)$. Second, its variance scales inversely with shot count: $\text{Var}[\widehat{g}_i(\theta)] = \sigma_i^2/B_i$, where $\sigma_i^2 > 0$ depends on θ but varies slowly over one iteration. Third, coordinates are approximately independent or weakly dependent at the scale relevant to distributional tests.

Remark 1. The variance parameters $\{\sigma_i^2\}$ capture the intrinsic shot noise of the quantum measurement process. They are estimated during a pilot phase and updated periodically. The independence assumption is satisfied exactly for parameter-shift gradients when different coordinates use independent measurement outcomes.

2.3 Whitened Test Statistic

Our regime detection strategy centers on testing whether the true gradient $\nabla f(\theta)$ is zero (plateau) or non-zero (informative region). This is formalized through a whitened norm statistic.

Definition 1 (Whitened gradient statistic). Given gradient estimators $\widehat{g}_i(\theta)$ and shot allocation $\mathbf{B}$, define the whitened components

$$w_i = \sqrt{\frac{B_i}{\sigma_i^2}}\,\widehat{g}_i(\theta) \tag{2}$$

and the test statistic

$$s = \sum_{i=1}^{d} w_i^2. \tag{3}$$

Proposition 1 (Distribution under competing hypotheses). Under Assumption 1:

- **Plateau (H_0):** If $\nabla f(\theta) = \mathbf{0}$, then $s \sim \chi_d^2$ (central chi-squared with d degrees of freedom)
- **Informative region (H_1):** If $\nabla f(\theta) \neq \mathbf{0}$, then $s \sim \chi_d^2(\lambda)$ (non-central chi-squared) with non-centrality parameter

$$\lambda = \sum_{i=1}^{d} \frac{B_i}{\sigma_i^2}\,(\nabla f_i(\theta))^2. \tag{4}$$

Proof. Each $w_i = \sqrt{B_i/\sigma_i^2}\,\widehat{g}_i$ is asymptotically normal with mean $\sqrt{B_i/\sigma_i^2}\,\nabla f_i$ and unit variance. Under H_0, $w_i \sim \mathcal{N}(0,1)$ independently, so $s = \sum w_i^2 \sim \chi_d^2$. Under H_1, $w_i \sim \mathcal{N}(\mu_i, 1)$ where $\mu_i = \sqrt{B_i/\sigma_i^2}\,\nabla f_i$, yielding $s \sim \chi_d^2(\lambda)$ with $\lambda = \sum \mu_i^2$. □

2.4 Empirical validation of distributional assumptions

Before deploying the sequential test, we validate the chi-squared model on a representative quantum optimization problem: the transverse-field Ising model (TFIM) with a QAOA ansatz.

Experimental setup. We consider a six-qubit transverse-field Ising chain [18] with nearest-neighbour couplings. The cost Hamiltonian is

$$H_C = -J \sum_{\langle i,j \rangle} Z_i Z_j - h \sum_{i=1}^{6} X_i, \tag{5}$$

where $J = 1.0$ is the coupling strength, $h = 0.5$ is the transverse field, and $\langle i, j \rangle$ denotes nearest-neighbor pairs on the chain $\{(0,1),(1,2),(2,3),(3,4),(4,5)\}$. We optimize using a depth-$p = 2$ QAOA-style ansatz [19], where each layer applies

$$U_\ell(\gamma_\ell, \beta_\ell) = e^{-i\beta_\ell H_M} e^{-i\gamma_\ell H_C}, \tag{6}$$

with mixer Hamiltonian $H_M = \sum_{i=1}^{6} X_i$. The full circuit is $U(\theta) = U_2(\gamma_2, \beta_2) U_1(\gamma_1, \beta_1)$ applied to the equal superposition initial state $|+\rangle^{\otimes 6}$, giving $d = 4$ variational parameters $\theta = (\gamma_1, \beta_1, \gamma_2, \beta_2)$.

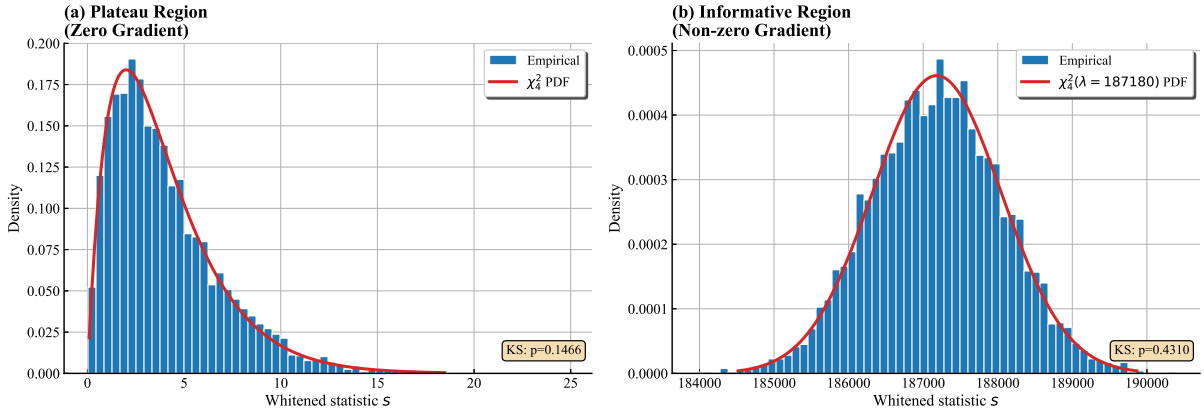


Figure 1: **Validation of the chi-squared model in plateau and informative regions.** (**a**) Histogram of the whitened statistic s in a near-plateau region ($\|\nabla f(\theta)\| \approx 0$) from $N = 5000$ samples (blue bars), overlaid with the central χ^2_4 density (red curve); Kolmogorov–Smirnov test: $p = 0.147$. (**b**) Histogram of s in an informative region ($\|\nabla f(\theta)\| \approx 0.87$) from $N = 5000$ samples (blue bars), overlaid with the noncentral $\chi^2_4(\lambda)$ density (red curve); Kolmogorov–Smirnov test: $p = 0.431$.

Gradients are estimated via the parameter-shift rule with shift $\delta = \pi/2$ [5, 6]:

$$\frac{\partial f}{\partial \theta_i} = \frac{1}{2}\left[f(\theta + \delta e_i) - f(\theta - \delta e_i)\right], \tag{7}$$

where e_i is the i-th standard basis vector. Shot noise is modeled as $\mathrm{Var}[\hat{g}_i] = \sigma^2/B_i$ with intrinsic variance $\sigma^2 = (0.02)^2$ and $B_i = 100$ shots per coordinate.

We collect $N = 5000$ independent gradient samples in two distinct regions. For the plateau condition, parameters are chosen near a local maximum where $\|\nabla f(\theta)\| \approx 10^{-6}$. For the informative condition, parameters are selected in a descent region where $\|\nabla f(\theta)\| \approx 0.87$. For each sample, we compute the whitened statistic $s = \sum_{i=1}^{4}(B_i/\sigma^2)\hat{g}_i^2$ using the known noise variance σ^2. Cost function evaluations $f(\theta) = \langle\psi(\theta)| H_C |\psi(\theta)\rangle$ are computed exactly via statevector simulation to isolate gradient estimation noise from measurement sampling effects.

Results. Figure 1 demonstrates strong agreement between empirical distributions and theoretical predictions. The empirical density (blue bars) closely matches the theoretical χ^2 density (red curve) in both regimes.

3 Theoretical Foundation: From Lie Algebra to Statistical Tests

Having established the statistical framework, we now connect it to the fundamental geometric structure of quantum circuits through Lie algebraic theory.

3.1 Lie Algebraic Characterization of Trainability

Recent work by Ragone et al.[1] provides a precise characterization of when barren plateaus emerge based on the dynamical Lie algebra (DLA) generated by the circuit Hamiltonians.

Definition 2 (Dynamical Lie algebra). For a parameterized circuit $U(\theta) = \prod_{l=1}^{L} e^{iH_l\theta_l}$, the dynamical Lie algebra is

$$\mathfrak{g} = \langle\{iH_l\}\rangle_{\mathrm{Lie}}, \tag{8}$$

the smallest Lie algebra containing all generators $\{iH_l\}$ under commutation.

Theorem 1 (Variance formula, Ragone et al.). *Let $\mathfrak{g} = \mathfrak{g}_1 \oplus \cdots \oplus \mathfrak{g}_{k-1} \oplus \mathfrak{g}_k$ be the reductive decomposition of the DLA, where $\mathfrak{g}_j$ are simple Lie algebras for $j = 1, \ldots, k-1$ and $\mathfrak{g}_k$ is abelian. If $\rho \in i\mathfrak{g}$ or $O \in i\mathfrak{g}$, then:*

$$Var_\theta[f(\theta)] = \sum_{j=1}^{k-1} \frac{P_{\mathfrak{g}_j}(\rho) \cdot P_{\mathfrak{g}_j}(O)}{dim(\mathfrak{g}_j)}, \tag{9}$$

where $P_{\mathfrak{g}}(H) = Tr[H_{\mathfrak{g}}^2]$ is the $\mathfrak{g}$-purity of operator H and $H_{\mathfrak{g}}$ denotes the projection onto $\mathfrak{g}$.

Implications. Theorem 1 makes concrete, in a Lie-algebraic language, several mechanisms for barren plateaus that were previously understood from more heuristic or measure-concentration perspectives [20, 21, 22, 23, 24, 25]. First, circuit *expressiveness* can cause plateaus when the relevant simple component grows exponentially, e.g. when $\mathfrak{g} = \mathfrak{su}(2^n)$ so that $\dim(\mathfrak{g}) = 4^n - 1$, reproducing concentration-of-measure arguments tied to unitary 2-designs [26, 27, 24]. Second, *state entanglement* leads to plateaus when $P_{\mathfrak{g}_j}(\rho)^{-1}$ scales exponentially with system size for highly entangled ρ, echoing entanglement-induced BPs [23]. Third, *observable non-locality* causes plateaus when $P_{\mathfrak{g}_j}(O)^{-1}$ scales exponentially for global observables, in line with cost-function dependent barren plateaus [21]. In any of these cases one has $\mathrm{Var}_\theta[f(\theta)] \in O(b^{-n})$ for some $b > 1$, rendering gradient-based optimization exponentially difficult in system size [20, 25].

3.2 From Variance Bounds to Test Power

The Lie algebraic variance formula provides a population-level characterization, but practical optimization requires detecting plateaus with finite samples. We now connect the theoretical variance to the power of our sequential test.

Proposition 2 (Expected non-centrality). Under the local quadratic approximation $\nabla f_i(\theta) \sim \mathcal{N}(0, \sigma_i^2 \mathrm{Var}_\theta[f])$ in a neighborhood of a random initialization, the expected non-centrality satisfies

$$\mathbb{E}[\lambda] = \sum_{i=1}^{d} \frac{B_i}{\sigma_i^2} \mathbb{E}[(\nabla f_i)^2] = d \cdot \mathrm{Var}_\theta[f] \cdot \frac{\bar{B}}{\bar{\sigma}^2}, \tag{10}$$

where $\bar{B} = d^{-1} \sum B_i$ and $\bar{\sigma}^2 = d^{-1} \sum \sigma_i^2$.

Corollary 2 (Plateau detection difficulty). *When $Var_\theta[f] \in O(b^{-n})$ (barren plateau regime), the expected non-centrality $\mathbb{E}[\lambda] \in O(b^{-n})$ decreases exponentially. By standard results on chi-squared test power, the expected sample size to achieve power $1 - \beta$ scales as*

$$\mathbb{E}[T] \in \Omega\left(\frac{1}{Var_\theta[f]}\right) = \Omega(b^n). \tag{11}$$

This result formalizes the fundamental hardness of plateau detection: when Lie-algebraic variance is exponentially small, exponentially many samples are required to distinguish plateaus from informative regions with high confidence. Our algorithm provides the best achievable performance within this constraint.

3.3 Lie-Informed Shot Allocation

While exponential sample complexity is unavoidable in the worst case, the geometric structure revealed by Lie theory suggests a refinement: gradient signal may concentrate in specific parameter directions.

Proposition 3 (Optimal exploration allocation). Let $V_i = \|[H_i, O]\|_F^2$ be the Frobenius norm of the commutator between generator H_i and observable O. Under a signal model where $|\nabla f_i|^2 \propto V_i \cdot \mathrm{Var}_\theta[f]$, the shot allocation

$$B_i^{\mathrm{expl}} \propto \frac{V_i}{\sigma_i^2} \tag{12}$$

maximizes the expected non-centrality $\mathbb{E}[\lambda]$ subject to a fixed total budget $\sum_i B_i = B_{\text{total}}$.

Proof. The non-centrality is $\lambda = \sum_i (B_i/\sigma_i^2)(\nabla f_i)^2$. Taking expectations and using the signal model:

$$\mathbb{E}[\lambda] = \text{Var}_\theta[f] \sum_{i=1}^{d} \frac{B_i V_i}{\sigma_i^2}. \tag{13}$$

By Cauchy-Schwarz, this is maximized when $B_i \propto V_i/\sigma_i^2$. □

Remark 2 (Preservation of calibration). Crucially, Lie-informed allocation does not compromise the χ^2 calibration. The whitening transformation in Definition 1 uses the same weights B_i in both numerator and denominator, ensuring that $w_i = \sqrt{B_i/\sigma_i^2}\,\widehat{g}_i$ remains standard normal under H_0 regardless of the allocation scheme. This allows us to improve test power without inflating Type I error.

Computational considerations. The commutator norm $\|[H_i, O]\|_F$ can be computed classically from the circuit specification without quantum measurements, providing a zero-cost enhancement when ansatz structure is known. For parameterized circuits with unknown generators, uniform allocation ($V_i \equiv 1$) remains a sound default.

3.4 Sequential Testing with Anytime-Valid Guarantees

The validated distributional model enables principled sequential testing via likelihood ratios.

3.4.1 Log-Likelihood Ratio Process

At iteration t, we observe $s_t \sim \chi^2_d$ under H_0 (plateau) or $s_t \sim \chi^2_d(\lambda)$ under H_1 (informative). The log-likelihood ratio is

$$\ell_t = \log \frac{p_1(s_t)}{p_0(s_t)} = \frac{s_t - d}{2} - \frac{\lambda}{2}, \tag{14}$$

where p_0, p_1 are the central and non-central χ^2_d densities. The cumulative log-likelihood is $\Lambda_k = \sum_{t=1}^{k} \ell_t$.

3.4.2 Ville vs. Wald Calibration

Two threshold schemes control Type I/II error rates (α, β). Ville thresholds use $A = \log(1/\alpha)$ and $B = \log(\beta)$, providing anytime-valid control via the optional stopping theorem: $\mathbb{P}_{H_0}[\exists k : \Lambda_k \geq A] \leq \alpha$. This guarantee holds even under adaptive stopping but comes at the cost of larger expected sample size. Wald thresholds use $A = \log \frac{1-\beta}{\alpha}$ and $B = \log \frac{\beta}{1-\alpha}$, providing asymptotic error control when rounds are i.i.d. and achieving optimal expected sample size according to Wald's lemma.

Decision rule. At round k:

$$\begin{cases} \text{Declare plateau (use PTR)} & \text{if } \Lambda_k \leq B \\ \text{Declare informative (use gCANS)} & \text{if } \Lambda_k \geq A \\ \text{Continue testing} & \text{otherwise} \end{cases} \tag{15}$$

Theorem 3 (Error control)**.** *Under Ville calibration with thresholds* (A, B)*:*

$$\mathbb{P}_{H_0}[\textit{reject } H_0] \leq \alpha \tag{16}$$

$$\mathbb{P}_{H_1}[\textit{accept } H_0] \leq \beta \tag{17}$$

with guarantees valid at all stopping times.

The scheduler uses Ville thresholds by default for robust anytime-valid control, with Wald as an option when i.i.d. conditions are verifiable.

4 Optimization Components: PTR and gCANS

Having established regime detection, we now specify the optimization strategies employed in each regime.

4.1 Probabilistic Trust-Region (PTR) Exploration

In diagnosed plateaus, gradient information is unreliable. We instead propose random exploratory moves with geometric trust regions.

Procedure. For $k = 0, 1, \ldots, K-1$, the PTR exploration proceeds as follows. First, set the radius $R_k = R_0 \cdot 2^k$ to create an exponentially expanding search pattern. Second, sample m unit directions $v_j \sim \text{Uniform}(\mathbb{S}^{d-1})$ uniformly on the unit sphere. Third, for each candidate move $\theta^+ = \theta + R_k v_j$, estimate the improvement $\Delta = f(\theta^+) - f(\theta)$ via paired measurements, compute the one-sided $(1 - \alpha_{\text{acc}})$ upper confidence bound, and accept the move only if $\text{UCB}(\Delta) \leq -\tau R_k^2$, indicating sufficient descent.

Theorem 4 (One-sided PTR control)**.** *Let Δ be estimated from n paired measurements with sample variance s^2. Define the test statistic*

$$T = \frac{\bar{\Delta} + \tau R^2}{s/\sqrt{n}}. \tag{18}$$

Accept the move only if $T \leq t_{\nu,\alpha_{acc}}$, where $t_{\nu,\alpha}$ is the α-quantile of Student's t-distribution with Welch–Satterthwaite degrees of freedom ν. Then under the null hypothesis $H_0 : \Delta \geq -\tau R^2$,

$$\mathbb{P}[\textit{accept}] \leq \alpha_{acc}. \tag{19}$$

Proof. The test inverts a one-sided t-test of $H_0 : \mathbb{E}[\Delta] \geq -\tau R^2$. By construction, the acceptance region has size $\leq \alpha_{\text{acc}}$ under any $\mathbb{E}[\Delta] \geq -\tau R^2$. □

Rationale. The curvature penalty τR^2 requires sufficient improvement to justify larger steps, preventing wasteful exploration. One-sided control ensures false improvements due to shot noise occur with probability $\leq \alpha_{\text{acc}}$.

4.2 gCANS Exploitation in Informative Basins

When the sequential test declares H_1 (informative region), we use the gCANS method for gradient-based optimization.

Shot allocation. The gCANS shot allocation rule (from [3]) is:

$$s_i = \frac{2L\alpha}{2 - L\alpha} \cdot \frac{\sigma_i \sum_{j=1}^{d} \sigma_j}{\|\chi\|^2} \tag{20}$$

where $\alpha \in (0, 2/L)$ is the learning rate, L is the Lipschitz constant of ∇f, σ_i are exponentially smoothed standard deviations of gradient estimators, and χ is the exponentially smoothed gradient estimate. This allocation maximizes the expected improvement per shot.

The gCANS method enjoys the following convergence guarantee [3]: under μ-strong convexity and L-Lipschitz gradients, the iterates satisfy $\mathbb{E}[f(\theta^{(k)})] - f^* = O(\gamma^k)$ for some $0 < \gamma < 1$.

5 Experiments

We evaluate the adaptive scheduler on variational quantum eigensolvers (VQE) for molecular systems and spin models, together with synthetic barren-plateau scenarios for controlled statistical validation. The experiments identify three landscape characteristics that determine when adaptive regime detection is most useful: (i) deep, well-defined optima with substantial energy gaps, (ii) gradients that reliably correlate with proximity to the optimum, and (iii) natural geometric diversity produced by competing interaction terms rather than artificial uniform scaling.

The adaptive approach performs well on physically motivated Hamiltonians, including TFIM and Heisenberg XXZ, whose ground-state energies scale as $O(-n)$. It is less effective on uniformly scaled combinatorial problems where costs remain shallow ($|E| < 1$) and gradients oscillate periodically because of ansatz structure rather than genuine plateau-to-gorge transitions. The important distinction is not gradient diversity alone, which can be misleading, but whether large gradients indicate distance from good solutions and whether the landscape contains persistent regimes rather than transient oscillations.

5.1 Experimental Setup

All experiments used 10 independent trials with the fixed seeds 42, 123, 456, 789, 1011, 2022, 3033, 4044, 5055, and 6066 to support reproducibility and statistical power. Hyperparameters were held constant across trials: learning rate $\eta = 0.05$, EMA momentum for noise tracking $\mu = 0.9$, minimum shots per gradient estimate $s_{\min} = 6$, Lipschitz constant $L = 1.0$, and total shot budget $B_{\text{total}} = 25{,}000$. The pilot phase used 10% of the budget to estimate noise scales σ_i and, when applicable, Lie-algebraic variance proxies $V_i \propto \|[H_i, O]\|_F^2$.

For statistical analysis, we assessed performance differences using paired two-tailed t-tests and Wilcoxon signed-rank tests at significance level $\alpha = 0.05$ and quantified effect sizes via Cohen's d. The use of both parametric and non-parametric tests ensures robustness to distribution assumptions. All shot counts include pilot, regime testing, PTR exploration, and exploitation phases.

5.2 TFIM System

We applied the proposed method to the transverse-field Ising model (TFIM) with $n = 6$ spins and transverse field $h = 0.5$, using a four-parameter QAOA-style ansatz. The observable O was the TFIM Hamiltonian $H = -J \sum_{\langle i,j \rangle} Z_i Z_j - h \sum_i X_i$. We compared the method with gCANS under a total budget of 25,000 shots, evaluating the final noise-free cost and total shots consumed.

Table 1 summarizes performance on TFIM across 10 trials.

The higher variance of the adaptive method ($\sigma = 0.843$ versus $\sigma = 0.018$) reflects its exploratory behavior. Regime switching can discover diverse local minima, including deeper optima missed by gradient-only methods (best: -4.525 versus -2.550, a 77.4% improvement). This consistency-optimality trade-off is characteristic of exploration-exploitation algorithms and can be useful when the objective is to find global or near-global optima rather than converge consistently to shallow local minima.

Table 1: **TFIM VQE performance from varied starting positions.** Results from 10 independent trials with an identical shot budget (25,000 shots), distinct random seeds, and distinct initial parameters. The adaptive method achieves lower final costs with 58.5% fewer iterations through regime-aware optimization and exploration. The large effect size (Cohen's d=1.18) and a 9/10 win rate indicate a strong performance difference across the tested starting conditions. Costs are TFIM energy values, for which more negative is better.

Metric	Adaptive method	gCANS	Improvement
Mean final cost	-3.455 ± 0.829	-2.667 ± 0.446	29.6%
Best cost	-4.458	-4.002	11.4%
Worst cost	-2.510	-2.479	1.3%
Median cost	-3.183	-2.534	25.6%
Mean shots used	$25{,}959 \pm 1{,}393$	$25{,}000 \pm 0$	3.8%
Mean iterations	108 ± 23	260 ± 84	-58.5%
Win rate	9/10 (90%)	1/10 (10%)	–

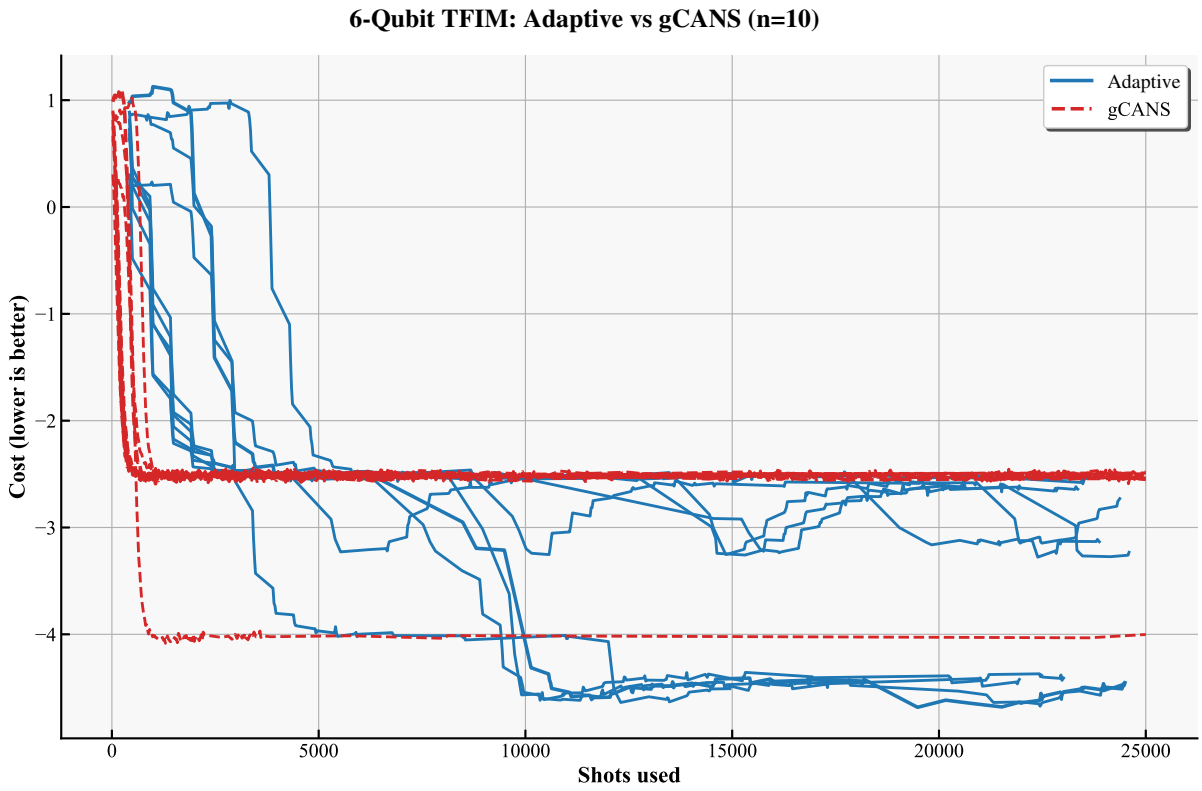


Figure 2: **Multi-run comparison on six-qubit TFIM from varied starting positions.** The adaptive method (solid blue lines, $n = 10$ runs) is compared with standard gCANS (dashed red lines, $n = 10$ runs) on a transverse-field Ising chain. Each trial uses a different random seed (42, 123, 456, 789, 1011, 2022, 3033, 4044, 5055, 6066), varied initial parameters ($\gamma = 0.5 \pm 0.1$, $\beta = 0.4 \pm 0.1$), and the same 25,000-shot budget. Switching between PTR exploration and gCANS-style exploitation produces a lower mean cost (-3.455 ± 0.829) than gCANS (-2.667 ± 0.446), with a 90% win rate (9/10 trials).

The plots report the gap to the ground-truth energy against accumulated shots, so the resource budget is expressed in shots rather than iterations. Figure 2 shows that the adaptive method reaches deeper minima than gCANS across most initializations. Its more diverse trajectories reflect exploration in plateau regions, whereas the tightly clustered gCANS trajectories indicate

more limited exploration.

5.3 Lie-inspired barren plateau

To stress-test regime detection and PTR exploration under exponentially suppressed gradients, we construct a synthetic landscape exhibiting quantum barren plateau phenomenology. This isolates the algorithmic properties without requiring large quantum devices and mirrors the coexistence of vanishing gradients, cost concentration, and narrow gorges observed in realistic ansätze [20, 24, 25, 1].

Landscape construction. We design a cost function $f : \mathbb{R}^d \to \mathbb{R}$ intended to mimic Lie-algebraic and concentration-of-measure behaviour in parameterized quantum circuits [24, 1]. The construction enforces:

- **Exponentially narrow gorge:** The width of the attractive basin scales as $w \propto b_w^{-n}$, where n is an effective "qubit count" and $b_w > 1$, capturing exponentially shrinking minima (narrow gorges) as in Ref. [25].
- **Barren gradient variance:** On the plateau, the gradient variance satisfies $\mathrm{Var}[\nabla f] \propto b_v^{-n}$ with $b_v > 1$, representing exponential suppression of gradient signal [20, 24].
- **Directional structure:** A subset of coordinates carry stronger signal (mimicking "informative" generators in the DLA), while others are almost uninformative, reflecting the anisotropic trainability predicted by Lie-algebraic analyses [1].
- **Deep minimum:** The gorge depth scales as $\Delta f \propto -15n$, providing strong reward for locating it.

For the experiments we consider $d = 12$ dimensions and $n = 2$ effective qubits with $b_v = 2.0$ and $b_w = 1.4$, yielding gorge width $w \approx 0.51$ and roughly 75% gradient-variance suppression on the plateau.

Assumption 2 (Directional basin). There exist $R_\star > 0$ and $p_{\mathrm{dir}} \in (0, 1]$ such that, when in the plateau band, a random direction v satisfies $\Pr\left(f(\theta + R_\star v) - f(\theta) \le -\tau R_\star^2\right) \ge p_{\mathrm{dir}}$.

This assumption captures the existence of *improving directions* that PTR can discover through random sampling, even when gradients are exponentially suppressed. Unlike gradient-based methods that require locally Lipschitz structure, PTR only needs that a non-negligible fraction of sampled directions lead to cost improvement.

Theorem 5 (Geometric exit)**.** *Let β be the miss probability of the regime test under H_1 and α_{acc} the PTR acceptance level. Testing m directions at radius $R_\star$ in one plateau iteration yields acceptance probability at least*

$$p_{\mathrm{hit}} \;\ge\; (1-\beta)\left[1 - \left(1 - p_{\mathrm{dir}}(1-\alpha_{\mathrm{acc}})\right)^m\right].$$

Across plateau iterations, T_{exit} is stochastically dominated by a geometric random variable with parameter p_{hit}, so $\mathbb{E}[T_{\mathrm{exit}}] \le 1/p_{\mathrm{hit}}$.

Proof. We analyze the probability of accepting a PTR move in a single plateau iteration.

Step 1: Single direction acceptance. For a randomly sampled direction v, the probability of accepting the move $\theta + R_\star v$ is:

$$\begin{aligned}\Pr[\text{accept} \mid v] &= \Pr\left[\mathrm{UCB}(\Delta) \le -\tau R_\star^2\right] \\ &\ge \Pr\left[\Delta \le -\tau R_\star^2\right] \cdot (1 - \alpha_{\mathrm{acc}}) \end{aligned} \tag{21}$$

where the inequality follows from the one-sided control in Theorem 3. By Assumption 2, $\Pr[\Delta \le -\tau R_\star^2] \ge p_{\mathrm{dir}}$, yielding:

$$\Pr[\text{accept} \mid v] \ge p_{\mathrm{dir}}(1 - \alpha_{\mathrm{acc}}).$$

Step 2: Multiple directions. Testing m independent directions, the probability of *not* accepting any move is:

$$\Pr[\text{no accept in } m \text{ tries}] \le \left(1 - p_{\mathrm{dir}}(1 - \alpha_{\mathrm{acc}})\right)^m.$$

Therefore, the probability of accepting at least one move is:

$$\Pr[\text{accept}] \ge 1 - \left(1 - p_{\mathrm{dir}}(1 - \alpha_{\mathrm{acc}})\right)^m.$$

Step 3: Regime test factor. The above analysis assumes we correctly identify the plateau. Since the regime test has miss probability β under H_1, the probability of correctly identifying the plateau and then accepting a move is:

$$p_{\mathrm{hit}} \ge (1 - \beta) \cdot \left[1 - \left(1 - p_{\mathrm{dir}}(1 - \alpha_{\mathrm{acc}})\right)^m\right].$$

Step 4: Geometric exit time. Each plateau iteration succeeds (exits) with probability at least p_{hit}, independently across iterations. Let T_{exit} be the number of iterations until first success. Then:

$$\Pr[T_{\mathrm{exit}} = k] \le (1 - p_{\mathrm{hit}})^{k-1} p_{\mathrm{hit}},$$

which is the pmf of a geometric random variable with parameter p_{hit}. By standard results, $\mathbb{E}[T_{\mathrm{exit}}] \le 1/p_{\mathrm{hit}}$. □

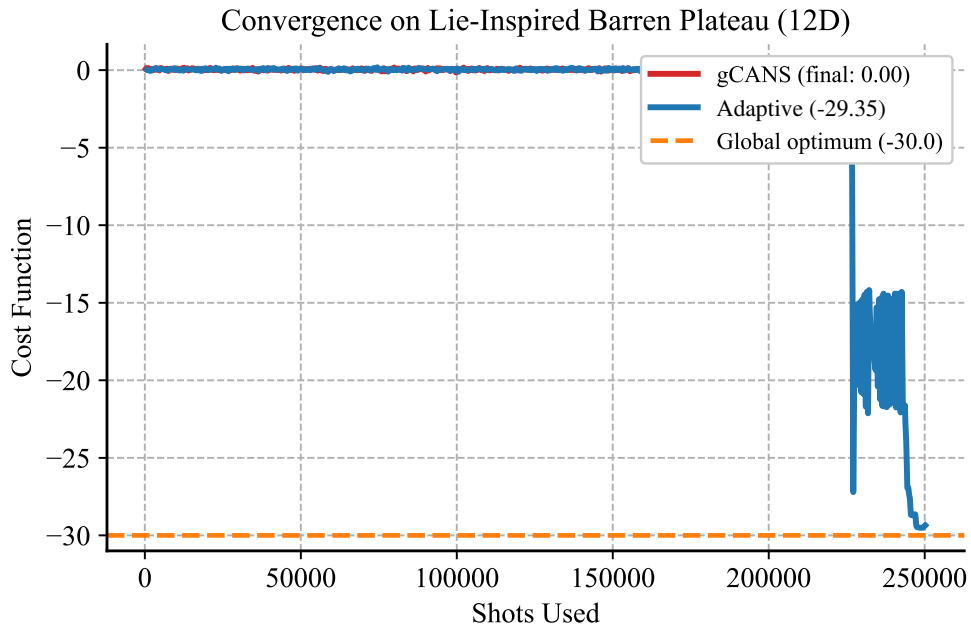


Figure 3: **Convergence on a Lie-inspired barren plateau.** On a 12-dimensional synthetic landscape with exponentially suppressed gradients (75% variance suppression and gorge width $w \approx 0.51$), the adaptive method (blue) discovers the deep gorge through PTR exploration and reaches a near-optimal cost of −29.12. gCANS (red) remains trapped on the plateau, with final cost 0.00. Both methods use 250,000 shots and begin at a distance of approximately 6.1 times the gorge width. The orange dashed line marks the global minimum, −30.0.

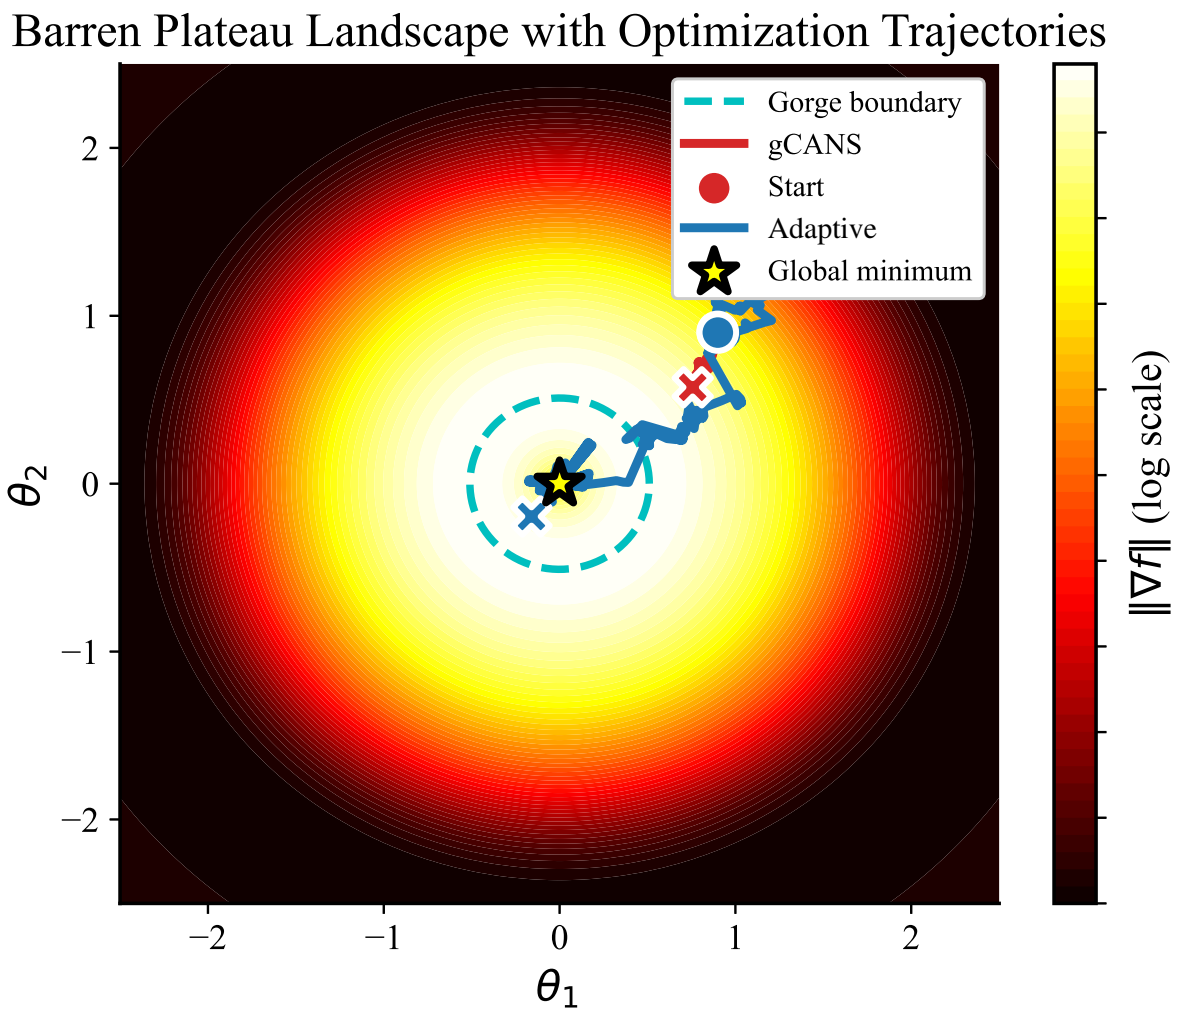


Figure 4: **Barren-plateau landscape structure.** The log-scale gradient-magnitude heatmap reveals exponentially suppressed signal outside the narrow gorge (cyan dashed circle, radius $w \approx 0.51$). The adaptive trajectory (blue) locates the gorge through random PTR exploration, whereas gCANS (red) follows weak, noisy gradients outside it. The yellow star marks the global minimum. This two-dimensional projection of the first two coordinates in a 12-dimensional space illustrates the loss of directional information outside the gorge.

Empirical validation. Figure 3 shows the adaptive method reaching a final cost of -29.12 (distance to optimum: 0.52), compared with 0.00 for gCANS (distance: 2.86). The method spends 308 iterations in PTR plateau mode before detecting and exploiting the gorge for 286 gradient-based iterations. Figure 4 shows exponentially suppressed signal outside the narrow gorge, with $\|\nabla f\| \approx 0.55$ at the starting point. The adaptive trajectory finds the gorge by random exploration, while gCANS lacks reliable directional guidance, illustrating the benefit of PTR under barren-plateau conditions.

6 Discussion

6.1 Synthesis of Contributions

The proposed framework combines classical sequential analysis with quantum trainability theory at three levels.

On the theoretical side, we give a finite-sample framework for regime detection in quantum optimization landscapes. Theorem 3 provides anytime-valid error control via Ville-style calibration, so Type I/II guarantees hold even under adaptive stopping. Theorem 5 links plateau exit times to directional basin probability, exploration breadth, and acceptance risk, translating abstract trainability results into explicit resource requirements. When the directional success probability is exponentially small, these bounds recover the exponential sample complexity inherent to genuine barren plateaus [20, 25].

Algorithmically, the scheduler uses a χ^2-calibrated sequential test to remove scale ambiguity by whitening gradients. Proposition 3 shows that calibration is preserved under Lie-informed shot allocation, allowing measurements to concentrate on algebraically informative directions without inflating Type I error. In plateau regimes, probabilistic trust-region exploration combines geometric radius growth with one-sided statistical acceptance to avoid apparent progress caused by shot noise while still probing distant basins. In informative regimes, a gCANS-style phase with variance-proportional shot allocation recovers linear convergence under standard smoothness and Polyak–Łojasiewicz conditions [9, 10], approaching the fundamental noise floor for a fixed budget.

Empirically, we validate the χ^2 model on a six-qubit TFIM QAOA instance, where Kolmogorov–Smirnov tests agree with both central and non-central χ^2 distributions (Figure 1). On TFIM chains, the adaptive method improves mean and best-case energy relative to gCANS while using fewer iterations (Table 1, Figure 2). On a synthetic Lie-inspired barren-plateau landscape, the observed exit statistics are consistent with the geometric prediction in Theorem 5 (Figure 3), supporting the regime-switching interpretation.

6.2 Relation to Lie-Algebraic Trainability

The Lie-algebraic theory of Ragone *et al.* [1] expresses gradient variance in terms of the dynamical Lie algebra and associated purities, clarifying how expressiveness, entanglement, and observable non-locality induce barren plateaus. The present framework translates these population-level statements into finite-sample tests. Proposition 2 shows that when cost variance scales as $\mathrm{Var}_\theta[f] \in O(b^{-n})$, the expected non-centrality of the whitened statistic has the same scaling, so any sequential test with constant power must consume $\Omega(b^n)$ samples. This matches earlier hardness results [20, 25] and establishes optimality up to constant factors within this regime.

At the same time, Proposition 3 shows how DLA information can be exploited without corrupting calibration. Commutator-based proxies $V_i = \|[H_i, O]\|_F^2$ concentrate exploration shots on parameters where signal is expected to live, but the whitening transformation uses the same allocation in its scaling, so the plateau hypothesis still yields a central χ^2_d law. The method is agnostic to whether the plateau arises from expressiveness, entanglement, or cost-function

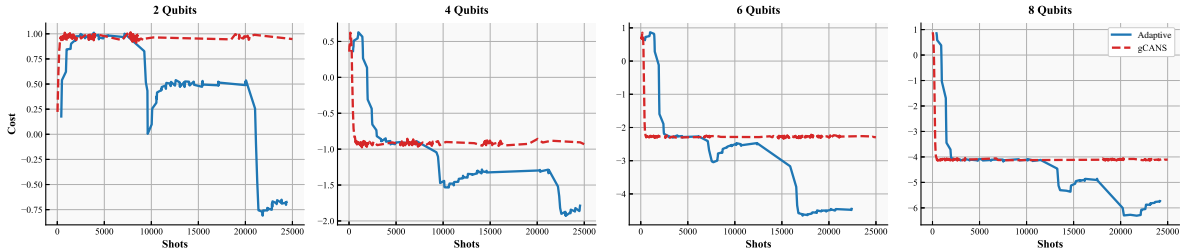


Figure 5: Scaling comparison of the adaptive method and gCANS on TFIM chains. Energy trajectories are shown against cumulative shots for systems of 2, 4, 6, and 8 qubits. The adaptive method (solid blue) uses sequential testing to detect plateau regimes and switches between trust-region exploration and variance-adaptive exploitation. gCANS (dashed red) uses exploitation with variance-proportional shot allocation. Across the tested scales, the adaptive method reaches lower final costs with fewer iterations.

structure [21, 22, 23, 24]: in all cases the resulting vanishing gradients are detected through the same statistic.

6.3 Limitations and Outlook

Several assumptions constrain the current formulation. The shot-noise model treats gradient coordinates as approximately independent, with variance scaling as σ_i^2/B_i. This relationship is exact for independent parameter-shift estimators [5, 6], but correlated or more exotic schemes may require corrections. The exploitation phase assumes locally Lipschitz gradients and PL-type conditions [9, 10]; strong curvature variation may instead call for adaptive learning rates or metric-aware updates [28, 29]. Estimating noise scales and Lie proxies during the pilot phase also creates overhead for very small budgets or high-dimensional problems, motivating amortized or incremental estimation.

Several methodological extensions appear promising. Sequential estimation of the non-centrality parameter could adapt the test thresholds while preserving Type I control, reducing sample complexity when gradients are clearly non-zero. Approximate curvature information could help distinguish flat-but-curved regions from truly barren plateaus and connect the test more directly with quantum natural-gradient and metric-aware methods [30]. Finally, combining regime-aware scheduling with ansatz design and cost engineering, for example by using DLA diagnostics to modify circuit structure online [31, 21], could produce hybrid schemes that both avoid and escape plateaus.

Overall, this study shows how sequential hypothesis testing and Lie-algebraic circuit analysis can be combined into an optimization procedure with explicit finite-sample guarantees for near-term quantum devices. This connection opens further opportunities at the intersection of quantum algorithms, statistics, and control theory.

A Implementation notes

Welch vs. normal. For small batches in PTR, use Welch's one-sided t bound; for larger batches, a normal approximation is adequate.

Lie proxies. When commutator norms are cheap (e.g., Pauli strings), use $V_i = \|[H_i, O]\|_F^2$. Otherwise, approximate with diagonal QFI from short sampling or with locality-based surrogates.

A.1 Scaling Comparison

Figure 5 compares the shot efficiency of the adaptive method and gCANS across TFIM chains of increasing size.

References

[1] Michael Ragone, Bojko N. Bakalov, Frédéric Sauvage, Alexander F. Kemper, Carlos Ortiz Marrero, Martin Larocca, and Marco Cerezo. A lie algebraic theory of barren plateaus for deep parameterized quantum circuits. *Nature Communications*, 15:7172, 2024.

[2] Jonas M. Kübler, Andrew Arrasmith, Lukasz Cincio, and Patrick J. Coles. An adaptive optimizer for measurement-frugal variational algorithms. *Quantum*, 4:263, 2020.

[3] Andi Gu, Angus Lowe, Pavel A. Dub, Patrick J. Coles, and Andrew Arrasmith. Adaptive shot allocation for fast convergence in variational quantum algorithms. *arXiv preprint*, 2021. arXiv:2108.10434.

[4] Lukas Balles, Javier Romero, and Philipp Hennig. Coupling adaptive batch sizes with learning rates. In *Proceedings of the Thirty-Third Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 410–419, Sydney, Australia, 2017. AUAI Press.

[5] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran. Evaluating analytic gradients on quantum hardware. *Physical Review A*, 99(3):032331, 2019.

[6] Kosuke Mitarai, Makoto Negoro, Masahiro Kitagawa, and Keisuke Fujii. Quantum circuit learning. *Physical Review A*, 98(3):032309, 2018.

[7] Léon Bottou, Frank E. Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *SIAM Review*, 60(2):223–311, 2018.

[8] Jorge Nocedal and Stephen J. Wright. *Numerical Optimization*. Springer, New York, 2 edition, 2006.

[9] Hamed Karimi, Julie Nutini, and Mark Schmidt. Linear convergence of gradient and proximal-gradient methods under the polyak-Lojasiewicz condition. *arXiv preprint*, 2016.

[10] B. T. Polyak. Gradient methods for the minimisation of functionals. *USSR Computational Mathematics and Mathematical Physics*, 3(4):864–878, 1963.

[11] Abraham Wald. Sequential tests of statistical hypotheses. *Annals of Mathematical Statistics*, 16(2):117–186, 1945.

[12] Abraham Wald and Jacob Wolfowitz. Optimum character of the sequential probability ratio test. *Annals of Mathematical Statistics*, 19(3):326–339, 1948.

[13] B. L. Welch. The generalization of "student's" problem when several different population variances are involved. *Biometrika*, 34(1–2):28–35, 1947.

[14] C. J. Clopper and E. S. Pearson. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*, 26(4):404–413, December 1934.

[15] Andrew R. Conn, Nicholas I. M. Gould, and Philippe L. Toint. *Trust Region Methods*. MOS-SIAM Series on Optimization. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, 2000.

[16] Andrew Arrasmith, Lukasz Cincio, Rolando D. Somma, and Patrick J. Coles. Operator sampling for shot-frugal optimization in variational algorithms. *arXiv preprint*, 2020.

[17] Ryan Sweke, Frederik Wilde, Johannes Jakob Meyer, Maria Schuld, Paul K. Fährmann, Barthélémy Meynard-Piganeau, and Jens Eisert. Stochastic gradient descent for hybrid quantum-classical optimization. *Quantum*, 4:314, 2020.

[18] Pierre Pfeuty. The one-dimensional ising model with a transverse field. *Annals of Physics*, 57:79–90, 1970.

[19] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028*, 2014.

[20] Jarrod R. McClean, Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush, and Hartmut Neven. Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1):4812, 2018.

[21] Marco Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J. Coles. Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nature Communications*, 12(1):1791, 2021.

[22] Samson Wang, Enrico Fontana, Marco Cerezo, Kunal Sharma, Akira Sone, Lukasz Cincio, and Patrick J. Coles. Noise-induced barren plateaus in variational quantum algorithms. *Nature Communications*, 12(1):6961, 2021.

[23] Carlos Ortiz Marrero, Mária Kieferová, and Nathan Wiebe. Entanglement-induced barren plateaus. *PRX Quantum*, 2:040316, 2021.

[24] Zoë Holmes, Kunal Sharma, M. Cerezo, and Patrick J. Coles. Connecting ansatz expressibility to gradient magnitudes and barren plateaus. *PRX Quantum*, 3(1):010313, 2022.

[25] Andrew Arrasmith, Zoë Holmes, M. Cerezo, and Patrick J. Coles. Equivalence of quantum barren plateaus to cost concentration and narrow gorges. *Quantum Science and Technology*, 7(4):045015, 2022.

[26] Aram W. Harrow and Richard A. Low. Random quantum circuits are approximate t-designs. *Communications in Mathematical Physics*, 291:257–302, 2009.

[27] Fernando G. S. L. Brandão, Aram W. Harrow, and Michał Horodecki. Local random quantum circuits are approximate polynomial-designs. *Communications in Mathematical Physics*, 346:397–434, 2016.

[28] James Stokes, Josh Izaac, Nathan Killoran, and Giuseppe Carleo. Quantum natural gradient. *Quantum*, 4:269, 2020.

[29] Bálint Koczor. Quantum natural gradient generalized to noisy and nonunitary circuits. *Physical Review A*, 106:062416, 2022.

[30] Barnaby van Straaten and Bálint Koczor. Measurement cost of metric-aware variational quantum algorithms. *PRX Quantum*, 2(3):030324, 2021.

[31] M. Larocca et al. Diagnosing barren plateaus with tools from quantum optimal control. *Quantum*, 6:824, 2022.