Learned Trust Regions for Query-Efficient QAOA

Molena Huynh[1,*]
[1] *North Carolina State University*

(Dated: April 25, 2026)

Circuit evaluations are the primary resource consumed by near-term variational quantum algorithms: every angle candidate must be verified by quantum measurement, and total gate use scales directly with evaluation count. Existing parameter-transfer methods for QAOA predict a point initialization, leaving the predicted angle distribution unused after the first iterate. We show that a single diagonal Gaussian $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, predicted from graph structure by a Graph Isomorphism Network with Laplacian spectral encodings, enables three simultaneous uses of one learned object: the mean initializes the optimizer, the covariance defines a Mahalanobis trust region that constrains every subsequent iterate, and the trace of the covariance sets an instance-dependent evaluation budget. Training combines a negative log-likelihood objective with a Wasserstein regularizer that couples predicted distributions across instances with nearby angle optima, and a supervised contrastive loss that orders the graph embedding by landscape similarity before the Gaussian heads act. We establish local guarantees: bounded loss under trust-region confinement, gradient anti-concentration under a local curvature condition, preservation of expected ordering under affine depolarizing noise, and finite-sample conformal coverage. On MaxCut at depth p=2 across Erdős–Rényi, 3-regular, Barabási–Albert, and Watts–Strogatz graphs with n=8–16, the mean evaluation count at n=14 decreases from 343 (random restarts) and 85 (best learned-point baseline) to 45 ± 7, an 86.9% reduction, while the sampled best-bitstring approximation ratio remains within 0.027 of the best concentration-based heuristic; at the higher-budget operating point the gap closes to 0.007. Uncertainty is calibrated (ECE=0.052, Spearman ρ=0.770), the speedup is stable at 7.2–7.7× under finite-shot noise, and the gains transfer to unseen graph sizes and families without retraining.

I. INTRODUCTION

The Quantum Approximate Optimization Algorithm (QAOA) [1] applies p alternating layers of a cost unitary and a mixer unitary to the uniform superposition. The $2p$ angles are chosen to maximize the expected value of the cost operator. At low depth, derivative-free solvers often require hundreds of objective evaluations per instance [2, 3]. On quantum hardware, each evaluation requires repeated state preparation and measurement. Query count is therefore a primary resource [4, 5].

We reduce this cost by predicting a distribution rather than a point. A Graph Isomorphism Network (GIN) [6] with Laplacian spectral encodings [7, 8] maps a graph $\mathcal{G}$ to a Gaussian $\mathcal{N}(\boldsymbol{\mu}(\mathcal{G}), \boldsymbol{\Sigma}(\mathcal{G}))$ over QAOA angles. The mean initializes local optimization. The covariance defines the region to which the optimizer is confined. The trace of the covariance sets the instance-dependent evaluation budget.

This differs from a warm start. A warm start uses the predictor once. Here the covariance constrains every iterate and sets the budget.

The scope is restricted to query efficiency in a small-graph, low-depth regime where exact statevector simulation and finite-shot simulation are feasible on a CPU.

Section V gives the local guarantees: control of the loss induced by confinement, a conditional anti-concentration statement for the gradient, preservation of expected ordering under depolarizing noise, and finite-sample coverage after conformal calibration.

We evaluate the method on MaxCut at depth p=2. The test graphs are Erdős–Rényi, 3-regular, Barabási–Albert, and Watts–Strogatz families with n=8–16 vertices. Training uses 240 graphs at n=14 only. The learned trust region transfers to families and sizes not seen during training. On n=14 graphs the mean evaluation count drops from 343 for random restarts and 85 for the best learned-point predictor to 45 ± 7, while sampled approximation ratios remain within 3 percentage points of the concentration heuristic. The factor-of-two gain over the learned point predictor is the effect of the trust-region constraint and the adaptive budget. It is not the effect of initialization alone.

The main contributions of this work are:

1. **Gaussian angle predictor.** A GIN with Laplacian spectral positional encodings trained to predict a full diagonal Gaussian $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ over QAOA angles, where both the mean head and the variance head are jointly trained through a two-phase objective combining NLL, Wasserstein regularization, and supervised contrastive loss.

2. **Mahalanobis trust region.** A Mahalanobis ellipsoid $\mathcal{T}_\alpha(\mathcal{G})$ constructed from the predicted covariance and used as a hard constraint throughout the Nelder–Mead refinement, including projection of iterate that exit the region. The geometry is $\boldsymbol{\Sigma}$-anisotropic rather than isotropic.

3. **Uncertainty-guided adaptive budget.** A monotone, bounded allocation rule that maps predicted

---

[*] Contact author: Molena.Huynh@jmp.com

trace $\mathrm{tr}(\boldsymbol{\Sigma})$ to sample count K and polish budget T, so difficult instances receive more evaluations without inflating the budget on easy instances.

4. **Local guarantees.** A suite of local results covering: the loss from trust-region confinement (section IV D); gradient anti-concentration under assumption (A2); ordering preservation under depolarizing noise; adaptive-allocation regret; spectral-norm generalization; and finite-sample conformal coverage.

5. **Empirical evaluation.** Ablation, cross-size transfer, family generalization, operating-point sweep, multi-seed stability, and finite-shot simulation, all at $p{=}2$ in the CPU-accessible regime.

The paper is organized as follows. Section II collects notation and background. Section III situates the work relative to prior transfer, uncertainty, and amortized-optimization literature. Section IV describes the five algorithmic components. Section V states and proves the local guarantees. Sections VI and VII describe the experimental protocol and empirical evaluation. Section VIII discusses implications and limitations. Section IX concludes. Appendices contain proofs, implementation details, and supplementary statistics.

II. PROBLEM SETTING AND PRELIMINARIES

We work with unweighted MaxCut on n-vertex graphs and the depth-p QAOA.

A. MaxCut and the QAOA circuit

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be an unweighted graph with $|\mathcal{V}| = n$ and $|\mathcal{E}| = m$. MaxCut asks for a bipartition $z \in \{+1, -1\}^n$ maximizing

$$C(z) = \sum_{(i,j)\in\mathcal{E}} \tfrac{1}{2}(1 - z_i z_j). \tag{1}$$

The depth-p QAOA prepares the state

$$|\boldsymbol{\gamma}, \boldsymbol{\beta}\rangle = e^{-i\beta_p B} e^{-i\gamma_p C} \cdots e^{-i\beta_1 B} e^{-i\gamma_1 C} |+\rangle^{\otimes n}, \tag{2}$$

where $C = \sum_{(i,j)\in\mathcal{E}} \frac{1}{2}(I - Z_i Z_j)$ is the cost Hamiltonian and $B = \sum_i X_i$ [1]. The $2p$ angles $\boldsymbol{\theta} = (\boldsymbol{\gamma}, \boldsymbol{\beta}) \in \mathbb{R}^{2p}$ are chosen to maximize $F_{\mathcal{G}}(\boldsymbol{\theta}) = \langle \boldsymbol{\theta} | C | \boldsymbol{\theta} \rangle$.

Two quality measures appear below. The analysis uses the normalized expected objective $F_{\mathcal{G}}(\boldsymbol{\theta})/C_{\max}$. The experiments report the sampled approximation ratio $r = C(z_{\text{best}})/C_{\max}$, where z_{best} is the best bitstring observed during optimization.

B. Angle concentration and barren plateaus

At fixed depth p on bounded-degree graph families, optimal QAOA angles concentrate: the variance of the optimal $\boldsymbol{\theta}$ decreases as $O(1/n)$ [9, 10]. This makes cross-instance parameter transfer possible [3, 11].

Variational circuits can also exhibit barren plateaus, where the gradient variance at random parameters satisfies

$$\mathrm{Var}_{\boldsymbol{\theta}}\big[\partial_j F(\boldsymbol{\theta})\big] \le O\big(2^{-n}\big) \tag{3}$$

for deep or highly entangling ansätze [12–14]. Constant-depth QAOA with a local cost operator avoids this exponential suppression, but the landscape can still contain extended flat regions.

C. Spectral positional encodings

The graph Laplacian $L = D - A$ has eigendecomposition $L = U\Lambda U^\top$ with eigenvalues $0 = \lambda_0 \le \lambda_1 \le \cdots \le \lambda_{n-1}$. The eigenvectors $\{u_k\}$ define spectral positional encodings [7, 8, 15]. Sign ambiguity is resolved by random sign flips during training [8].

D. Graph Isomorphism Networks

The Graph Isomorphism Network (GIN) [6] is as expressive as the Weisfeiler–Leman isomorphism test [16]. Each layer updates

$$h_v^{(l+1)} = \mathrm{MLP}^{(l)}\big((1 + \epsilon^{(l)})\, h_v^{(l)} + \textstyle\sum_{u\in\mathcal{N}(v)} h_u^{(l)}\big), \tag{4}$$

with learnable $\epsilon^{(l)}$ and neighbor set $\mathcal{N}(v)$ [17, 18]. A graph-level representation is obtained by permutation-invariant pooling of the node embeddings.

E. Wasserstein distance for Gaussians

The 2-Wasserstein distance between probability measures P and Q on $\mathbb{R}^d$ is

$$W_2(P, Q) = \Big(\inf_{\pi\in\Pi(P,Q)} \int \|\boldsymbol{x} - \boldsymbol{y}\|_2^2 \, d\pi(\boldsymbol{x}, \boldsymbol{y})\Big)^{1/2}. \tag{5}$$

For diagonal Gaussians this reduces to

$$W_2^2 = \|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\|_2^2 + \|\boldsymbol{\sigma}_1 - \boldsymbol{\sigma}_2\|_2^2, \tag{6}$$

which is closed form for diagonal covariance.

III. RELATION TO PRIOR WORK

Parameter transfer for QAOA. Angle concentration at fixed depth [9, 10] supports fixed schedules [3], interpolation rules [2], cross-instance transfer [11], warm starts from related circuits [19, 20], graph-conditioned predictors [21, 22], Bayesian optimization [23], and related search strategies [24–27]. These methods provide angles or search heuristics. Here the predictor produces a full distribution. The technical distinction is that a point prediction can only initialize search; a Gaussian prediction can also constrain search geometry and allocate the per-instance budget. The algorithm therefore uses the learned object three times rather than once.

Learned combinatorial heuristics. In classical combinatorial optimization, learned heuristics are judged mainly by wall-clock time to solution [28–33]. For variational quantum algorithms the dominant resource is often the number of circuit evaluations. Classical methods therefore target throughput directly, whereas the relevant metric here is query count, independent of processor speed. Evaluation reduction is the primary figure of merit, not wall-clock time.

Uncertainty quantification. Standard tools include heteroscedastic regression, ensembles, dropout-based approximations, and post hoc calibration [34–41]. Here uncertainty is not reported only for diagnostics. It is used as an algorithmic quantity. It defines the trust region and sets the evaluation budget. The Wasserstein regularizer and the contrastive loss are introduced for this purpose. This is operationally distinct from calibration-as-reporting. A poorly calibrated variance head produces a misaligned trust region and an incorrect budget, both of which produce algorithmic failure rather than miscoverage in a confidence interval.

Optimal transport and contrastive learning. The closed-form Wasserstein distance for Gaussians [42] has been used in generative modeling [43, 44] and distributional robustness [45–47]. Here it regularizes the predicted angle distribution. Contrastive objectives [48–52] structure the graph embedding by landscape similarity. The two regularizers act at different levels. The contrastive loss acts on the embedding before the Gaussian heads; it orders latent space by optimization similarity. The Wasserstein loss acts on the output distribution; it couples the predicted covariance across instances that share nearby angle optima.

Amortized optimization. Amortized methods replace part or all of per-instance optimization by a learned map [53–56]. Existing amortized QAOA methods predict point initializations. Here the prediction is a Gaussian. Learning does not replace optimization. It restricts it. This distinction affects the failure mode: an amortized point predictor that transfers poorly leaves the optimizer without guidance; a Gaussian predictor that transfers poorly enlarges the trust region and degrades toward the point-predictor baseline rather than toward random initialization.

IV. GRAPH-CONDITIONED GAUSSIAN PREDICTION AND TRUST-REGION SEARCH

A single forward pass of a GIN produces a Gaussian $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ over QAOA angles. The mean sets the center, the covariance sets the geometry, and the trace sets the budget. The method has five components (fig. 2): a GIN predictor, Wasserstein-regularized training, contrastive metric learning, trust-region-constrained optimization, and uncertainty-guided budget allocation.

Each component serves a distinct algorithmic role. The predictor maps graph structure to an angle distribution. The Wasserstein term aligns nearby landscapes in distribution space. The contrastive term orders the representation by angle similarity before the output heads act. The trust region converts predicted covariance into a geometric search constraint. The adaptive budget converts the same covariance into a resource-allocation rule. The method therefore uses one learned object for initialization, confinement, and budget selection. This is the central structural point of the algorithm. Prediction does not end at the first iterate. It persists through the full search procedure by setting the center, the admissible anisotropy, and the amount of local work invested on the instance.

a. Standing assumptions. The guarantees in section V are local. They assume: (A1) $F_{\mathcal{G}}$ is Lipschitz-continuous on the trust region (Proposition 1); (A2) $\lambda_{\min}(-\nabla^2 F_{\mathcal{G}}(\boldsymbol{\mu})) \geq \kappa > 0$ (Theorem 6); and (A3) depolarizing noise with per-layer strength $\varepsilon < 1$ (Theorem 19). Assumption (A2) is local.

A. Gaussian predictor architecture

Given $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, node features are formed by concatenating normalized degree d_v/n with k-dimensional Laplacian spectral positional encodings [7, 8, 15], i.e., the k smallest nontrivial eigenvectors of $L = D - A$, yielding $h_v^{(0)} \in \mathbb{R}^{k+1}$. Eigenvector sign ambiguity is handled by random sign flips during training [8].

A three-layer GIN with hidden dimension d=64 and LayerNorm [57] produces node embeddings. A graph-level representation $\boldsymbol{g} \in \mathbb{R}^{2d}$ is obtained by concatenating mean and max pooling:

$$\boldsymbol{g} = \big[\tfrac{1}{n} \textstyle\sum_v h_v^{(L)} \,\big\|\, \max_v h_v^{(L)}\big]. \tag{7}$$

Two linear heads produce the Gaussian parameters:

$$\boldsymbol{\mu} = W_\mu \boldsymbol{g} + b_\mu \in \mathbb{R}^{2p}, \tag{8}$$

$$\log \boldsymbol{\sigma}^2 = \mathrm{clamp}(W_\sigma \boldsymbol{g} + b_\sigma, -5, 2). \tag{9}$$

The predictive distribution is

$$p(\boldsymbol{\theta} \mid \mathcal{G}) = \mathcal{N}\big(\boldsymbol{\mu}(\mathcal{G}),\, \mathrm{diag}(\boldsymbol{\sigma}^2(\mathcal{G}))\big).$$

The clamp in eq. (9) keeps $\sigma_j \in [e^{-5/2}, e]$. The covariance is diagonal. The backbone is GIN and the output distribution is diagonal Gaussian.

The architecture couples graph representation and angle uncertainty in the smallest form that still supports the algorithm. Spectral positional encodings supply global graph coordinates that persist across graph size. The GIN backbone preserves local combinatorial structure under message passing. Mean pooling stabilizes family-level statistics; max pooling preserves sharp structural signatures. The mean head sets a candidate center for optimization. The variance head sets the anisotropy of the search. A point predictor would supply the first object but not the second, and the trust-region mechanism would collapse to unconstrained local search. The diagonal Gaussian is not a claim about full posterior structure. It is the least complicated object that still carries directional scale, and directional scale is enough to define an ellipsoidal search set and an instance-dependent budget rule.

B. Wasserstein-regularized training

Data. Training uses 240 graphs at n=14 (60 per family) and a validation set of 80 graphs. Targets $\boldsymbol{\theta}_i^*$ are obtained by multi-restart Nelder–Mead [58] on exact simulation.

Phase 1.

$$\mathcal{L}_1 = \frac{1}{N}\sum_{i=1}^{N} \|\boldsymbol{\mu}(\mathcal{G}_i) - \boldsymbol{\theta}_i^*\|_2^2 + \lambda_C \mathcal{L}_{\mathrm{CL}}. \tag{10}$$

Phase 2.

$$\mathcal{L}_2 = \mathcal{L}_{\mathrm{NLL}} + \lambda_W \mathcal{L}_{\mathrm{W2}} + \lambda_C \mathcal{L}_{\mathrm{CL}}, \tag{11}$$

The negative log-likelihood is

$$\mathcal{L}_{\mathrm{NLL}} = \frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{2p}\left[\frac{(\mu_j(\mathcal{G}_i) - \theta_{ij}^*)^2}{\sigma_j^2(\mathcal{G}_i)} + \log \sigma_j^2(\mathcal{G}_i)\right]. \tag{12}$$

Wasserstein regularizer. For diagonal Gaussians,

$$W_2^2(P_i, P_j) = \|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\|_2^2 + \|\boldsymbol{\sigma}_i - \boldsymbol{\sigma}_j\|_2^2.$$

Weights $w_{ij} = \exp(-\|\boldsymbol{\theta}_i^* - \boldsymbol{\theta}_j^*\|_2^2 / 2\tau_W^2)$ encode angle-space similarity:

$$\mathcal{L}_{\mathrm{W2}} = \frac{1}{|\mathcal{B}|^2}\sum_{i \neq j \in \mathcal{B}} w_{ij}\, W_2^2(P_i, P_j). \tag{13}$$

Phase 1 fits the mean. Phase 2 fits the full Gaussian.

The two phases separate two tasks that otherwise interfere. Phase 1 anchors the center of the distribution near the target angle vector. Phase 2 assigns scale after the center is already stable. The negative log-likelihood alone acts pointwise in the training set. The Wasserstein term adds a geometric coupling across instances. Two graphs with nearby target angles then receive nearby predictive distributions rather than unrelated variances. That coupling matters because the covariance is not a passive uncertainty label; it later defines both the trust region and the budget. Training therefore does not merely fit labels. It shapes a search geometry that will later be used by the optimizer, so errors in scale become algorithmic errors rather than cosmetic calibration errors.

C. Contrastive metric learning

A supervised contrastive loss structures the embedding $\boldsymbol{g}$. For anchor i, define

$$\mathcal{P}(i) = \{j \neq i : \|\boldsymbol{\theta}_i^* - \boldsymbol{\theta}_j^*\|_2 < \delta\}.$$

$$\mathcal{L}_{\mathrm{CL}} = -\frac{1}{|\mathcal{B}|}\sum_i \frac{1}{|\mathcal{P}(i)|}\sum_{j \in \mathcal{P}(i)} \log \frac{e^{\mathrm{sim}(\boldsymbol{g}_i, \boldsymbol{g}_j)/\tau_C}}{\sum_{k \neq i} e^{\mathrm{sim}(\boldsymbol{g}_i, \boldsymbol{g}_k)/\tau_C}}, \tag{14}$$

This term acts before Gaussian prediction. Its role is to give the embedding a landscape order rather than a family label. Nearby angle optima pull graph representations together; distant optima push them apart. The prediction heads therefore start from a representation whose local neighborhoods already reflect optimization similarity. The Wasserstein term then refines that structure at the level of the final distribution. In this sense, contrastive learning organizes the latent space, and Wasserstein regularization organizes the output space.

D. Trust-region QAOA optimization

Definition 1 (Trust region)**.** *Given a predicted Gaussian* $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ *for graph* $\mathcal{G}$*, with* $\boldsymbol{\Sigma}$ *symmetric positive definite, and a confidence level* $\alpha \in (0, 1)$*, the* trust region *is the Mahalanobis ellipsoid*

$$\mathcal{T}_\alpha(\mathcal{G}) = \{\boldsymbol{\theta} \in \mathbb{R}^{2p} : (\boldsymbol{\theta} - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\boldsymbol{\theta} - \boldsymbol{\mu}) \leq \chi_{2p}^2(\alpha)\}, \tag{15}$$

where $\chi_{2p}^2(\alpha)$ *is the* α*-quantile of the* χ^2 *distribution with* $2p$ *degrees of freedom.*

Optimization is restricted to $\mathcal{T}_\alpha$: (a) candidate angles $\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_K \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ are sampled from the corresponding truncated Gaussian restricted to $\mathcal{T}_\alpha$; (b) during Nelder–Mead refinement, any simplex vertex $\boldsymbol{\theta}'$ exiting $\mathcal{T}_\alpha$ is projected back:

$$\mathrm{proj}_{\mathcal{T}}(\boldsymbol{\theta}') = \boldsymbol{\mu} + \min\Bigl(1, \frac{\sqrt{\chi_{2p}^2(\alpha)}}{\|\boldsymbol{\Sigma}^{-1/2}(\boldsymbol{\theta}' - \boldsymbol{\mu})\|_2}\Bigr)(\boldsymbol{\theta}' - \boldsymbol{\mu}). \tag{16}$$

The trust region is the algorithmic core. It turns predicted covariance into an admissible search set. A warm start uses the mean once and then discards the predictor. The present scheme keeps the prediction active through the full refinement stage. Every iterate is measured in

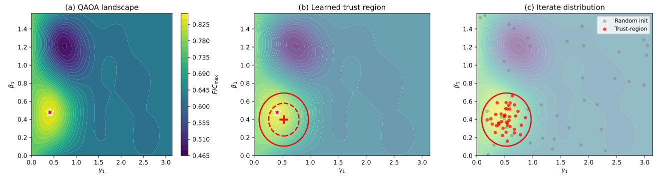

FIG. 1. Trust-region mechanism on an n=8 ER instance at p=2 (layer-1 slice). (a) QAOA landscape. (b) Learned Mahalanobis regions at 95% and 68%. (c) Random initializations (gray) vs. trust-region samples (red).

the Mahalanobis geometry induced by $\mathbf{\Sigma}$. Small variance directions receive tight control; large variance directions remain available for exploration. The search therefore inherits the anisotropy of the learned angle distribution rather than a fixed, hand-chosen radius. This distinction is the source of the query savings. The algorithm does not spend evaluations discovering a local metric from scratch around each new graph. It begins with a metric predicted from related solved instances and restricts search to that geometry from the first query onward.

E. Uncertainty-guided adaptive refinement

Define

$$U(\mathcal{G}) = \mathrm{tr}(\mathbf{\Sigma}(\mathcal{G}))/2p, \quad z(\mathcal{G}) = \frac{U(\mathcal{G}) - U_{\mathrm{med}}}{U_{\mathrm{iqr}}}.$$

Resource allocation:

$$K(\mathcal{G}) = \mathrm{clamp}(\lfloor 1 + 4\sigma(z)\rfloor, 1, 5), \tag{17}$$

$$T(\mathcal{G}) = \mathrm{clamp}(\lfloor T_{\mathrm{base}}(0.5 + z_+)\rfloor, 5, 2T_{\mathrm{base}}), \tag{18}$$

Higher uncertainty increases both the number of samples and the number of refinement steps. The mapping is monotone and bounded (fig. 11).

This allocation rule gives the variance head a second algorithmic use. The same signal that enlarges the search region also increases query budget. Easy instances therefore receive a narrow region and short refinement; difficult instances receive a broader region and a larger budget. A uniform budget would ignore the instance-dependent content of the predictor and spend queries on graphs that already admit a stable center. The allocation rule does not solve a separate control problem. It enforces structural consistency between predicted geometric spread and computational effort: large predicted spread enlarges both the search set and the work invested inside it.

Remark 1 (Conformal trust-region radius)**.** *At inference time, the $\chi^2_{2p}(\alpha)$ radius in Step 5 can be replaced by the conformal quantile $\hat{q}_{1-\alpha}$ (proposition 17), yielding a* conformalized *trust region*

$$\mathcal{T}_{\hat{q}_{1-\alpha}} = \left\{\boldsymbol{\theta} : (\boldsymbol{\theta} - \boldsymbol{\mu})^\top \mathbf{\Sigma}^{-1}(\boldsymbol{\theta} - \boldsymbol{\mu}) \le \hat{q}_{1-\alpha}\right\}$$

with the distribution-free, finite-sample guarantee $\mathbb{P}[\boldsymbol{\theta}^* \in \mathcal{T}_{\hat{q}_{1-\alpha}}] \ge 1 - \alpha$. For our validation set ($\hat{q}_{0.90} = 10.8$ vs. $\chi^2_4(0.90) = 7.779$), the radius increases by a factor of $(10.8/7.779)^{1/2} \approx 1.178$.

```
Algorithm 1: UQ-QAOA inference
Input: graph $\mathcal{G}$, trained model $f_\theta : \mathcal{G} \mapsto (\boldsymbol{\mu}, \mathbf{\Sigma})$, budget $T_{\mathrm{base}}$, confidence level $\alpha$.
Output: approximate MaxCut bitstring $z^*$.
1. Predict $(\boldsymbol{\mu}, \mathbf{\Sigma}) \leftarrow f_\theta(\mathcal{G})$.
2. Compute $U \leftarrow \mathrm{tr}(\mathbf{\Sigma})/2p$;  $z \leftarrow (U - U_{\mathrm{med}})/U_{\mathrm{iqr}}$.
3. Set sample count $K \leftarrow \mathrm{clamp}(\lfloor 1 + 4\sigma(z)\rfloor, 1, 5)$.
4. Set polish budget
   $T \leftarrow \mathrm{clamp}(\lfloor T_{\mathrm{base}}(0.5 + z_+)\rfloor, 5, 2T_{\mathrm{base}})$.
5. Construct trust region
   $\mathcal{T}_\alpha \leftarrow \{\boldsymbol{\theta} : \|\mathbf{\Sigma}^{-1/2}(\boldsymbol{\theta} - \boldsymbol{\mu})\|_2^2 \le \chi^2_{2p}(\alpha)\}$.
6. Sample $\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_K$ from the truncated Gaussian
   $\mathcal{N}(\boldsymbol{\mu}, \mathbf{\Sigma})$ restricted to $\mathcal{T}_\alpha$.
7. Evaluate QAOA circuit: $\boldsymbol{\theta}^* \leftarrow \arg\max_k F_{\mathcal{G}}(\boldsymbol{\theta}_k)$.
8. Trust-region polish: run $T$ iterations of Nelder–Mead
   from $\boldsymbol{\theta}^*$, projecting each iterate via eq. (16).
9. Return best bitstring $z^*$ from final QAOA
   measurement.
```

FIG. 2. UQ-QAOA inference. Steps 3–4 convert uncertainty into sampling and refinement budgets; Steps 5–8 perform trust-region-constrained optimization.

V. GUARANTEES FOR LEARNED TRUST-REGION SEARCH

Each result below is local and conditional on the stated assumptions. There is no claim of global optimality. The bounds quantify the loss incurred by restricting search to the predicted ellipsoid and by allocating budget according to predicted uncertainty.

Throughout, $\mathcal{G} \sim \mathcal{D}$ denotes a held-out graph except where otherwise stated. The predictor is $f_\theta(\mathcal{G}) = (\boldsymbol{\mu}(\mathcal{G}), \mathbf{\Sigma}(\mathcal{G}))$. The reference point $\boldsymbol{\theta}^*(\mathcal{G})$ denotes the local optimum produced by the multi-restart Nelder–Mead procedure used for target generation, or the optimum specified in the corresponding statement. The trust region $\mathcal{T}_\alpha(\mathcal{G})$ is defined in eq. (15). For a radius not tied to a $\chi^2_{2p}(\alpha)$ quantile, we write $\mathcal{T}_q(\mathcal{G})$.

Expectations and probabilities are taken with respect to the randomness explicitly indicated: draws of $\mathcal{G}$ from $\mathcal{D}$, samples from $\mathcal{N}(\boldsymbol{\mu}(\mathcal{G}), \mathbf{\Sigma}(\mathcal{G}))$, or uniform samples on $\mathcal{T}_\alpha(\mathcal{G})$. Constants such as $L_{\mathcal{G},K}$, $L_{\mathcal{G}}$, κ, Λ, and M_3 are deterministic once $\mathcal{G}$ is fixed.

A. Landscape properties

Proposition 1 (Local Lipschitz regularity of the QAOA objective)**.** *For a connected graph $\mathcal{G}$ with m edges, the depth-p* QAOA *objective $F_{\mathcal{G}} : \mathbb{R}^{2p} \to \mathbb{R}$ is real analytic.*

In particular, on any bounded set $K \subset \mathbb{R}^{2p}$,

$$L_{\mathcal{G},K} := \sup_{\boldsymbol{\theta} \in K} \|\nabla F_{\mathcal{G}}(\boldsymbol{\theta})\|_2 < \infty. \tag{19}$$

A conservative global bound is $L_{\mathcal{G}} \le 2m\sqrt{2p}$.

Proof. The QAOA unitary is $U(\boldsymbol{\theta}) = \prod_{\ell=1}^{p} e^{-i\beta_\ell B} e^{-i\gamma_\ell C}$, where $B = \sum_v X_v$ and $C = \sum_{(i,j)\in\mathcal{E}} \frac{1}{2}(I - Z_i Z_j)$. Each factor is a matrix exponential of a bounded Hermitian operator, so $U(\boldsymbol{\theta})$ is an entire function of each angle. The objective $F_{\mathcal{G}}(\boldsymbol{\theta}) = \langle +|U^\dagger(\boldsymbol{\theta})\, C\, U(\boldsymbol{\theta})|+\rangle$ is therefore a real-analytic function on $\mathbb{R}^{2p}$, and $\nabla F_{\mathcal{G}}$ is continuous everywhere. On any compact set K the continuous function $\|\nabla F_{\mathcal{G}}\|_2$ attains its supremum, so $L_{\mathcal{G},K} < \infty$. The mean-value theorem then gives $|F_{\mathcal{G}}(\boldsymbol{\theta}) - F_{\mathcal{G}}(\boldsymbol{\theta}')| \le L_{\mathcal{G},K}\|\boldsymbol{\theta} - \boldsymbol{\theta}'\|_2$ for all $\boldsymbol{\theta}, \boldsymbol{\theta}' \in K$. For the global bound, note that each partial derivative of $F_{\mathcal{G}}$ is an expectation value of a commutator. For the phase-separation angle γ_ℓ, $\partial_{\gamma_\ell} F_{\mathcal{G}} = -i\langle \boldsymbol{\theta}|[C_\ell, U_\ell^\dagger C U_\ell]|\boldsymbol{\theta}\rangle$ where C_ℓ is the cost operator acting in layer ℓ. Since C has operator norm $\|C\|_{\mathrm{op}} \le m$ (each edge contributes at most 1), the commutator has norm at most $2m$, giving $|\partial_{\gamma_\ell} F_{\mathcal{G}}| \le m$. Similarly, $|\partial_{\beta_\ell} F_{\mathcal{G}}| \le n$ since $\|B\|_{\mathrm{op}} = n$. The full gradient satisfies $\|\nabla F_{\mathcal{G}}\|_2^2 = \sum_{\ell=1}^{p}[(\partial_{\gamma_\ell} F)^2 + (\partial_{\beta_\ell} F)^2] \le p(m^2 + n^2)$. For connected graphs, $m \ge n - 1$ so $m^2 + n^2 \le 2m^2 + 2m + 1 \le 4m^2$ for $m \ge 2$. Hence $\|\nabla F_{\mathcal{G}}\|_2 \le 2m\sqrt{p} \le 2m\sqrt{2p}$, which establishes the global bound. $\square$

B. Trust-region guarantees

Lemma 2 (Validity of the Mahalanobis projection)**.** *Let $\boldsymbol{\mu} \in \mathbb{R}^{2p}$, $\boldsymbol{\Sigma} \succ 0$, and $q > 0$. For any $\boldsymbol{\theta}' \in \mathbb{R}^{2p}$, the map*

$$\mathrm{proj}_{\mathcal{T}}(\boldsymbol{\theta}') = \boldsymbol{\mu} + \min\!\Big(1, \frac{\sqrt{q}}{\|\boldsymbol{\Sigma}^{-1/2}(\boldsymbol{\theta}' - \boldsymbol{\mu})\|_2}\Big)(\boldsymbol{\theta}' - \boldsymbol{\mu})$$

satisfies:

(i) $\mathrm{proj}_{\mathcal{T}}(\boldsymbol{\theta}') \in \mathcal{T}_q$ *for every* $\boldsymbol{\theta}' \in \mathbb{R}^{2p}$.

(ii) If $\boldsymbol{\theta}' \in \mathcal{T}_q$, *then* $\mathrm{proj}_{\mathcal{T}}(\boldsymbol{\theta}') = \boldsymbol{\theta}'$ *(idempotency).*

(iii) $\mathrm{proj}_{\mathcal{T}}$ *is the nearest point in* $\mathcal{T}_q$ *under the Mahalanobis norm* $\|\cdot\|_{\boldsymbol{\Sigma}^{-1}}$.

Proof. Write $\boldsymbol{d} = \boldsymbol{\theta}' - \boldsymbol{\mu}$ and $s = \|\boldsymbol{\Sigma}^{-1/2}\boldsymbol{d}\|_2$. The projected point is $\boldsymbol{\theta}_{\mathrm{pr}} = \boldsymbol{\mu} + \min(1, \sqrt{q}/s)\,\boldsymbol{d}$.

(i) If $s \le \sqrt{q}$, then $\min(1, \sqrt{q}/s) = 1$ and $\boldsymbol{\theta}_{\mathrm{pr}} = \boldsymbol{\theta}'$, which satisfies the constraint $s^2 \le q$. If $s > \sqrt{q}$,

$$\|\boldsymbol{\Sigma}^{-1/2}(\boldsymbol{\theta}_{\mathrm{pr}} - \boldsymbol{\mu})\|_2 = \frac{\sqrt{q}}{s}\, s = \sqrt{q},$$

so $\boldsymbol{\theta}_{\mathrm{pr}}$ lies on the boundary of $\mathcal{T}_q$.

(ii) Immediate from the first case above.

(iii) For $s > \sqrt{q}$, the projection scales $\boldsymbol{d}$ to the boundary along the direction $\boldsymbol{d}/s$. The Mahalanobis ball $\mathcal{T}_q$ is convex, and radial scaling toward the center is the Euclidean projection in the whitened coordinates $\boldsymbol{u} = \boldsymbol{\Sigma}^{-1/2}(\boldsymbol{\theta} - \boldsymbol{\mu})$, which is the nearest-point projection onto the ball $\|\boldsymbol{u}\|_2 \le \sqrt{q}$. $\square$

Theorem 3 (Lower bound on the expected objective within the trust region)**.** *Let $\mathcal{T}_\alpha(\mathcal{G})$ be defined by eq. (15). Suppose $F_{\mathcal{G}}(\boldsymbol{\mu}) \ge r^* C_{\max}$. Let Θ be supported on $\mathcal{T}_\alpha(\mathcal{G})$. Then*

$$\mathbb{E}\big[F_{\mathcal{G}}(\Theta)\big] \ge r^* C_{\max} - L_{\mathcal{G}}\sqrt{\chi^2_{2p}(\alpha)\cdot \|\boldsymbol{\Sigma}\|_{\mathrm{op}}}, \tag{20}$$

Proof. For $\boldsymbol{\theta} \in \mathcal{T}_\alpha$, one has

$$\|\boldsymbol{\theta} - \boldsymbol{\mu}\|_2^2 \le \|\boldsymbol{\Sigma}\|_{\mathrm{op}} \cdot (\boldsymbol{\theta} - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\boldsymbol{\theta} - \boldsymbol{\mu}) \le \|\boldsymbol{\Sigma}\|_{\mathrm{op}} \cdot \chi^2_{2p}(\alpha).$$

Hence

$$\|\boldsymbol{\theta} - \boldsymbol{\mu}\|_2 \le \sqrt{\chi^2_{2p}(\alpha)\|\boldsymbol{\Sigma}\|_{\mathrm{op}}}.$$

By Proposition 1,

$$|F_{\mathcal{G}}(\boldsymbol{\theta}) - F_{\mathcal{G}}(\boldsymbol{\mu})| \le L_{\mathcal{G}}\sqrt{\chi^2_{2p}(\alpha)\|\boldsymbol{\Sigma}\|_{\mathrm{op}}}.$$

Applying the assumption $F_{\mathcal{G}}(\boldsymbol{\mu}) \ge r^* C_{\max}$ and taking expectation yields the claim. $\square$

Corollary 4 (Expected objective under variance confinement)**.** *Under the conditions of Theorem 3, if $\sigma^2_{\max} \le (r^* C_{\max})^2/(4L^2_{\mathcal{G}}\chi^2_{2p}(\alpha))$, then $\mathbb{E}[F_{\mathcal{G}}(\Theta)] \ge \frac{1}{2} r^* C_{\max}$.*

Proof. Set $\epsilon = \frac{1}{2} r^* C_{\max}$ in Theorem 3. This gives the sufficient condition $\sigma^2_{\max} \le (r^* C_{\max})^2/(4L^2_{\mathcal{G}}\chi^2_{2p}(\alpha))$. The conclusion follows immediately. $\square$

Proposition 5 (Second-order local improvement within the trust region)**.** *Assume $F_{\mathcal{G}}$ is twice continuously differentiable near $\boldsymbol{\mu}$ with $\lambda_{\min}(-\nabla^2 F_{\mathcal{G}}(\boldsymbol{\mu})) \ge \kappa > 0$ (i.e., $\boldsymbol{\mu}$ is a strict local maximizer). If the trust region satisfies $\sqrt{\chi^2_{2p}(\alpha)\,\|\boldsymbol{\Sigma}\|_{\mathrm{op}}} \le \kappa/M_3$, where M_3 bounds the third derivatives on $\mathcal{T}_\alpha$, then for every $\boldsymbol{\theta} \in \mathcal{T}_\alpha$,*

$$\begin{aligned} F_{\mathcal{G}}(\boldsymbol{\theta}) \ge{}& F_{\mathcal{G}}(\boldsymbol{\mu}) - \tfrac{1}{2}\Lambda\, \chi^2_{2p}(\alpha)\,\|\boldsymbol{\Sigma}\|_{\mathrm{op}} \\ &- \tfrac{1}{6} M_3 \left[\chi^2_{2p}(\alpha)\,\|\boldsymbol{\Sigma}\|_{\mathrm{op}}\right]^{3/2}, \end{aligned} \tag{21}$$

where $\Lambda = \|\nabla^2 F_{\mathcal{G}}(\boldsymbol{\mu})\|_{\mathrm{op}}$. For a small trust region the leading penalty is quadratic in $\sigma_{\max}$ rather than linear.

Proof. By Taylor expansion around $\boldsymbol{\mu}$ (a critical point, so $\nabla F_{\mathcal{G}}(\boldsymbol{\mu}) = 0$ for a local maximizer):

$$F_{\mathcal{G}}(\boldsymbol{\theta}) = F_{\mathcal{G}}(\boldsymbol{\mu}) + \tfrac{1}{2}\Delta^\top \nabla^2 F_{\mathcal{G}}(\boldsymbol{\mu})\,\Delta + O(\|\Delta\|_2^3),$$

where $\Delta = \boldsymbol{\theta} - \boldsymbol{\mu}$. Since $\boldsymbol{\mu}$ is a local maximizer, $\nabla^2 F_{\mathcal{G}}(\boldsymbol{\mu}) \preceq 0$, so $\Delta^\top \nabla^2 F_{\mathcal{G}}(\boldsymbol{\mu})\,\Delta \ge -\Lambda\|\Delta\|_2^2$. The cubic remainder is bounded by $\frac{1}{6} M_3 (2p)^{3/2}\|\Delta\|_2^3 \le \frac{1}{6} M_3\|\Delta\|_2^3$ (absorbing dimension factors into M_3). For $\boldsymbol{\theta} \in \mathcal{T}_\alpha$, $\|\Delta\|_2^2 \le \chi^2_{2p}(\alpha)\,\|\boldsymbol{\Sigma}\|_{\mathrm{op}}$ (as in Theorem 3). Substituting:

$$\begin{aligned} F_{\mathcal{G}}(\boldsymbol{\theta}) &\ge F_{\mathcal{G}}(\boldsymbol{\mu}) - \tfrac{1}{2}\Lambda\,\|\Delta\|_2^2 - \tfrac{1}{6} M_3\,\|\Delta\|_2^3 \\ &\ge F_{\mathcal{G}}(\boldsymbol{\mu}) - \tfrac{1}{2}\Lambda\,\chi^2_{2p}(\alpha)\,\|\boldsymbol{\Sigma}\|_{\mathrm{op}} \\ &\quad - \tfrac{1}{6} M_3\,[\chi^2_{2p}(\alpha)\,\|\boldsymbol{\Sigma}\|_{\mathrm{op}}]^{3/2}. \end{aligned}$$

$\square$

C. Gradient anti-concentration

Theorem 6 (Local gradient anti-concentration in the trust region)**.** *Assume $F_{\mathcal{G}}$ is three times continuously differentiable near $\mathcal{T}_\alpha(\mathcal{G})$, $\lambda_{\min}(-\nabla^2 F_{\mathcal{G}}(\boldsymbol{\mu})) \geq \kappa > 0$, and third derivatives are bounded by M_3. Then for $\boldsymbol{\theta} \sim \mathrm{Unif}(\mathcal{T}_\alpha)$,*

$$\begin{aligned}\mathrm{Var}_{\boldsymbol{\theta}\sim\mathrm{Unif}(\mathcal{T}_\alpha)}\big[\partial_j F_{\mathcal{G}}(\boldsymbol{\theta})\big] \geq \frac{\kappa^2\,\sigma_{\min}^2\,\chi_{2p}^2(\alpha)}{2p+2} \\ - C_{p,\alpha}\big(\Lambda M_3 \sigma_{\max}^3 + M_3^2\sigma_{\max}^4\big),\end{aligned} \tag{22}$$

where $\Lambda := \|\nabla^2 F_{\mathcal{G}}(\boldsymbol{\mu})\|_{\mathrm{op}}$, $\sigma_{\min} = \min_j \sigma_j$, and $\sigma_{\max} = \max_j \sigma_j$.

Proof. See section H. ☐

D. Algorithmic guarantees

Proposition 7 (Feasibility and monotonic best-so-far value under projection)**.** *Projected optimization within $\mathcal{T}_\alpha$ produces feasible iterates and a nondecreasing best-so-far objective. The final iterate dominates the initial sampled candidates. Let $\{\boldsymbol{\theta}_t\}_{t=0}^T$ be the feasible iterates generated by projected Nelder–Mead within $\mathcal{T}_\alpha$, define $b_t := \max_{0\leq s\leq t} F_{\mathcal{G}}(\boldsymbol{\theta}_s)$, and let $\widehat{\boldsymbol{\theta}}_t$ denote any feasible iterate attaining b_t. Then:*

(i) $\boldsymbol{\theta}_t \in \mathcal{T}_\alpha$ for all t.

(ii) b_t is non-decreasing in t.

(iii) $F_{\mathcal{G}}(\widehat{\boldsymbol{\theta}}_T) \geq \max_k F_{\mathcal{G}}(\boldsymbol{\theta}_k)$ where $\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_K$ are the initial feasible samples used to seed the search.

Proof. (i) The projection operator eq. (16) maps every point to $\mathcal{T}_\alpha$: if $\boldsymbol{\theta}' \in \mathcal{T}_\alpha$, the projection is the identity; otherwise, it rescales $\boldsymbol{\theta}' - \boldsymbol{\mu}$ to the boundary of $\mathcal{T}_\alpha$. (ii) Monotonicity is immediate from the definition $b_t = \max_{0\leq s\leq t} F_{\mathcal{G}}(\boldsymbol{\theta}_s)$: the index set on which the maximum is taken only enlarges with t. (iii) Because the initial sampled points are among the feasible points seen by the algorithm before refinement begins, $b_T \geq \max_{1\leq k\leq K} F_{\mathcal{G}}(\boldsymbol{\theta}_k)$. Since $\widehat{\boldsymbol{\theta}}_T$ attains b_T, the claim follows. ☐

Corollary 8 (Output guarantee under trust-region capture)**.** *Let $\boldsymbol{\theta}_{\mathrm{out}}$ denote the best feasible angle vector returned by Algorithm 2, and let z_{out} be the final bitstring sampled from the* QAOA *circuit at $\boldsymbol{\theta}_{\mathrm{out}}$. Fix a trust region $\mathcal{T}_q(\mathcal{G})$ and assume the reference local optimum $\boldsymbol{\theta}^*(\mathcal{G})$ belongs to that region. Then*

$$F_{\mathcal{G}}(\boldsymbol{\theta}_{\mathrm{out}}) \geq F_{\mathcal{G}}\big(\boldsymbol{\theta}^*(\mathcal{G})\big) - 2L_{\mathcal{G}}\sqrt{q\,\|\boldsymbol{\Sigma}(\mathcal{G})\|_{\mathrm{op}}}, \tag{23}$$

and therefore

$$\begin{aligned}\mathbb{E}\big[C(z_{\mathrm{out}}) \mid \boldsymbol{\theta}^*(\mathcal{G}) \in \mathcal{T}_q(\mathcal{G})\big] \geq F_{\mathcal{G}}\big(\boldsymbol{\theta}^*(\mathcal{G})\big) \\ - 2L_{\mathcal{G}}\sqrt{q\,\|\boldsymbol{\Sigma}(\mathcal{G})\|_{\mathrm{op}}}.\end{aligned} \tag{24}$$

If, in addition, the trust region has capture probability $p_{\mathrm{cap}}(\mathcal{G})$, then since MaxCut costs are nonnegative,

$$\mathbb{E}\big[C(z_{\mathrm{out}})\big] \geq p_{\mathrm{cap}}(\mathcal{G})\Big[F_{\mathcal{G}}\big(\boldsymbol{\theta}^*(\mathcal{G})\big) - 2L_{\mathcal{G}}\sqrt{q\,\|\boldsymbol{\Sigma}(\mathcal{G})\|_{\mathrm{op}}}\Big]. \tag{25}$$

For the algorithmic χ^2 region one may take $p_{\mathrm{cap}} \geq \alpha - \eta$ from Proposition 18; for the conformalized region one has $p_{\mathrm{cap}} \geq 1 - \alpha$.

Proof. Under the capture event, every sampled candidate and every projected iterate lies in $\mathcal{T}_q(\mathcal{G})$. Proposition 18 therefore gives the pointwise bound

$$F_{\mathcal{G}}(\boldsymbol{\theta}) \geq F_{\mathcal{G}}\big(\boldsymbol{\theta}^*(\mathcal{G})\big) - 2L_{\mathcal{G}}\sqrt{q\,\|\boldsymbol{\Sigma}(\mathcal{G})\|_{\mathrm{op}}} \qquad \forall\,\boldsymbol{\theta} \in \mathcal{T}_q(\mathcal{G}).$$

Since Algorithm 2 returns an angle vector inside that region, eq. (23) follows. By definition of the QAOA objective, $\mathbb{E}[C(z_{\mathrm{out}}) \mid \boldsymbol{\theta}_{\mathrm{out}}] = F_{\mathcal{G}}(\boldsymbol{\theta}_{\mathrm{out}})$, which gives eq. (24). The unconditional bound eq. (25) follows by conditioning on the capture event and using $C(z_{\mathrm{out}}) \geq 0$ for MaxCut. ☐

Corollary 9 (Dominance over shared-center point initialization)**.** *Assume the setting of Corollary 8. Suppose the initial candidate set contains the predicted mean $\boldsymbol{\mu}(\mathcal{G})$. Then the returned angle vector satisfies*

$$F_{\mathcal{G}}(\boldsymbol{\theta}_{\mathrm{out}}) \geq F_{\mathcal{G}}(\boldsymbol{\mu}(\mathcal{G})). \tag{26}$$

If, in addition, $\boldsymbol{\theta}^*(\mathcal{G}) \in \mathcal{T}_q(\mathcal{G})$, then

$$F_{\mathcal{G}}(\boldsymbol{\theta}_{\mathrm{out}}) \geq F_{\mathcal{G}}\big(\boldsymbol{\theta}^*(\mathcal{G})\big) - 2L_{\mathcal{G}}\sqrt{q\,\|\boldsymbol{\Sigma}(\mathcal{G})\|_{\mathrm{op}}}. \tag{27}$$

Thus, for a method that shares the same predicted center but does not use the covariance to constrain search, trust-region refinement cannot underperform the shared-center initialization and inherits the same capture-conditioned target-gap bound.

Proof. Because the initial candidate set contains $\boldsymbol{\mu}(\mathcal{G})$ and Proposition 7 states that the best-so-far objective is nondecreasing over all feasible iterates, the returned point satisfies eq. (26). The second claim is exactly the output guarantee from Corollary 8. ☐

E. Sampling guarantees

Theorem 10 (Best-of-K guarantee)**.** *Let $\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_K \overset{i.i.d.}{\sim} \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Then*

$$\mathbb{E}\big[\max_k F_{\mathcal{G}}(\boldsymbol{\theta}_k)\big] \geq F_{\mathcal{G}}(\boldsymbol{\mu}) - \frac{L_{\mathcal{G}}\sqrt{\mathrm{tr}(\boldsymbol{\Sigma})}}{\sqrt{K}}. \tag{28}$$

Proof. See section C. ☐

Lemma 11 (Trust-region volume ratio)**.** *For a diagonal covariance $\boldsymbol{\Sigma} = \mathrm{diag}(\sigma_1^2, \ldots, \sigma_{2p}^2)$ with confidence α=0.95* *and* $p{=}2$*, the trust-region volume relative to the unconstrained search hypercube* $[-\pi,\pi]^{2p}$ *is*

$$\frac{\mathrm{Vol}(\mathcal{T}_{0.95})}{\mathrm{Vol}([-\pi,\pi]^{2p})} = \frac{\pi^p[\chi^2_4(0.95)]^p}{\Gamma(p{+}1)\,(2\pi)^{2p}}\prod_{j=1}^{4}\sigma_j. \tag{29}$$

If one uses the isotropic proxy $\sigma_j \equiv \bar\sigma = 0.15$*, then eq.* (29) *gives* $0.2850\,\bar\sigma^4 \approx 1.44\times 10^{-4}$.

Proof. The trust region is a 4-dimensional ellipsoid with half-lengths $\sigma_j\sqrt{\chi^2_4(0.95)}$. Its volume is $\frac{\pi^2}{2}[\chi^2_4(0.95)]^2\prod_j\sigma_j$, and the hypercube volume is $(2\pi)^4$. The ratio follows by direct computation with $\chi^2_4(0.95)\approx 9.49$. ☐

F. Generalization and calibration

Proposition 12 (Wasserstein continuity in a WL-compatible pseudometric)**.** *If the* GIN *encoder* ϕ *is* L_{enc}*-Lipschitz with respect to a WL-compatible graph pseudometric* d_{WL} *and the prediction heads are* L_{head}*-Lipschitz, then:*

$$W_2\big(f_\theta(\mathcal{G}), f_\theta(\mathcal{G}')\big) \le \sqrt{2}\,L_{\mathrm{head}}L_{\mathrm{enc}}\,d_{\mathrm{WL}}(\mathcal{G},\mathcal{G}'), \tag{30}$$

where $d_{\mathrm{WL}}(\mathcal{G},\mathcal{G}') = 0$ *is allowed for non-isomorphic graphs that are indistinguishable by the relevant WL statistics.*

Proof. See section E. ☐

Theorem 13 (Adaptive-allocation guarantee under calibration and monotone difficulty)**.** *Assume approximate calibration and the regret envelope eq.* (32)*. Suppose the predicted uncertainty satisfies* η*-approximate calibration:*

$$\big|\mathbb{P}[\|\boldsymbol{\theta}^* - \boldsymbol{\mu}\|_{\boldsymbol{\Sigma}^{-1}} \le r] - F_{\chi^2_{2p}}(r^2)\big| \le \eta \quad \forall\, r>0, \tag{31}$$

where $F_{\chi^2_{2p}}$ *is the* χ^2_{2p} *CDF. Assume the residual optimization regret satisfies the instance-wise envelope*

$$R_\pi(\mathcal{G}) \le c_0 + c_1\frac{\sqrt{U(\mathcal{G})+\eta}}{\sqrt{K_\pi(\mathcal{G})}} \tag{32}$$

where $U(\mathcal{G}) = \mathrm{tr}(\boldsymbol{\Sigma}(\mathcal{G}))/2p$ *and* $c_0, c_1 \ge 0$ *are independent of* π*. Define population regret by*

$$\mathrm{Regret}(\pi) := \mathbb{E}_{\mathcal{G}\sim\mathcal{D}}[R_\pi(\mathcal{G})]. \tag{33}$$

Let $\bar K := \mathbb{E}_{\mathcal{D}}[K_{\mathrm{adapt}}(\mathcal{G})]$*, and let* π_{unif} *denote the mean-budget-matched uniform comparator with* $K_{\mathrm{unif}}(\mathcal{G}) \equiv \bar K$*. Assume further that the adaptive allocation rule satisfies the envelope-dominance condition*

$$\mathbb{E}_{\mathcal{D}}\left[\frac{\sqrt{U(\mathcal{G})+\eta}}{\sqrt{K_{\mathrm{adapt}}(\mathcal{G})}}\right] \le \frac{1}{\sqrt{\bar K}}\mathbb{E}_{\mathcal{D}}\left[\sqrt{U(\mathcal{G})+\eta}\right]. \tag{34}$$

Then the adaptive allocation strategy π_{adapt} *(Eqs. 17–18) satisfies*

$$\begin{aligned}\mathrm{Regret}(\pi_{\mathrm{adapt}}) &\le c_0 + c_1\,\mathbb{E}_{\mathcal{D}}\left[\frac{\sqrt{U(\mathcal{G})+\eta}}{\sqrt{K_{\mathrm{adapt}}(\mathcal{G})}}\right]\\ &\le c_0 + \frac{c_1}{\sqrt{\bar K}}\mathbb{E}_{\mathcal{D}}\left[\sqrt{U(\mathcal{G})+\eta}\right], \qquad (35)\\ \mathbb{E}[B_{\pi_{\mathrm{adapt}}}] &= \alpha\,\mathbb{E}[B_{\pi_{\mathrm{unif}}}], \qquad (36)\end{aligned}$$

where B_π *is the total evaluation budget under policy* π *and* $\alpha = \mathbb{E}_{\mathcal{D}}[B_{\pi_{\mathrm{adapt}}}]/\mathbb{E}_{\mathcal{D}}[B_{\pi_{\mathrm{unif}}}]$ *depends on the deployment distribution of predicted uncertainties. Thus* $\alpha<1$ *is not a consequence of calibration alone; it depends on the deployment distribution together with the implemented allocation rule.*

Proof. See section D. ☐

Proposition 14 (Monotonicity of the implemented allocation rule)**.** *Let* $z(\mathcal{G}) = (U(\mathcal{G}) - U_{\mathrm{med}})/U_{\mathrm{iqr}}$*, and define* K_{adapt} *and* T_{adapt} *by eqs.* (17) *and* (18)*. Then both allocations are nondecreasing functions of* U*, and one has the explicit thresholds*

$$K_{\mathrm{adapt}}(\mathcal{G}) = \begin{cases} 1, & z(\mathcal{G}) < -\log 3,\\ 2, & -\log 3 \le z(\mathcal{G}) < 0,\\ 3, & 0 \le z(\mathcal{G}) < \log 3,\\ 4, & z(\mathcal{G}) \ge \log 3.\end{cases} \tag{37}$$

If $T_{\mathrm{base}} \ge 10$*, then* $T_{\mathrm{adapt}}(\mathcal{G}) = \lfloor T_{\mathrm{base}}/2\rfloor$ *for* $z(\mathcal{G}) \le 0$*, and* T_{adapt} *increases linearly in* $z(\mathcal{G})$ *until it reaches the cap* $2T_{\mathrm{base}}$*. Thus the implemented policy reallocates computation monotonically from the low-*U *region to the high-*U *region, which is the ordering premise behind eq.* (34)*. The same ordering is supported empirically by fig. 9 and table I: larger predicted uncertainty is associated with harder instances, and exploiting that ordering lowers the total query count.*

Proof. The sigmoid is strictly increasing, so $\lfloor 1{+}4\sigma(z)\rfloor$ is nondecreasing in z. The thresholds follow from the identities $\sigma(z) = 1/4$ at $z = -\log 3$, $\sigma(z) = 1/2$ at $z = 0$, and $\sigma(z) = 3/4$ at $z = \log 3$. Because $\sigma(z) < 1$ for finite z, the implemented rule takes the values $1,2,3,4$ on the four regimes in eq. (37). The monotonicity of T_{adapt} follows in the same way from the monotonicity of $z_+ = \max(z,0)$ and the clamp map. If $z \le 0$, then $z_+ = 0$ and $T_{\mathrm{adapt}} = \mathrm{clamp}(\lfloor T_{\mathrm{base}}/2\rfloor, 5, 2T_{\mathrm{base}}) = \lfloor T_{\mathrm{base}}/2\rfloor$ whenever $T_{\mathrm{base}} \ge 10$. The empirical statements are exactly those reported in fig. 9 and table I. ☐

Theorem 15 (Spectral-norm generalization guarantee)**.** *Let* f_θ *be the trained* GIN *predictor with* L *layers, hidden dimension* d*, and per-layer spectral norms* $s_\ell = \|W_\ell\|_\sigma$*. Let* $B_x = \max_{\mathcal{G}}\|\boldsymbol{h}^{(0)}\|_2$ *be the bound on input features. Assume the induced loss class* $\mathcal{F} = \{W_2^2(f_\theta(\cdot), \delta_{\boldsymbol{\theta}^*(\cdot)})\}$ *is uniformly bounded by* M *and satisfies the empirical Rademacher complexity estimate*

$$\mathfrak{R}_N(\mathcal{F}) \le \frac{2B_x\prod_{\ell=1}^{L}s_\ell\sqrt{2L\log(2d)}}{\sqrt{N}}.$$

Given N i.i.d. training graphs, the population Wasserstein risk $\mathcal{R}(f_\theta) := \mathbb{E}_{\mathcal{G}\sim\mathcal{D}}[W_2^2(f_\theta(\mathcal{G}), \delta_{\boldsymbol{\theta}^*(\mathcal{G})})]$ satisfies, with probability $\geq 1-\delta$ over the training set:

$$\begin{aligned}\mathcal{R}(f_\theta) \leq \hat{\mathcal{R}}_N &+ \frac{4B_x \prod_{\ell=1}^{L} s_\ell}{\sqrt{N}} \sqrt{2L\log(2d)} \\ &+ 3M\sqrt{\frac{\log(2/\delta)}{2N}},\end{aligned} \tag{38}$$

where $\hat{\mathcal{R}}_N$ is the empirical Wasserstein risk and M is the loss bound. Here $\delta_{\boldsymbol{\theta}^*(\mathcal{G})}$ denotes the Dirac point mass at the target angle vector. The complexity term scales with $\prod_\ell s_\ell \sqrt{L\log d/N}$ rather than with the total parameter count [59–61].

Proof. See section F. ☐

Corollary 16 (Sample-complexity consequence)**.** *Under the conditions of Theorem 15, to achieve population risk within ϵ of the empirical risk with probability $\geq 1-\delta$, it suffices to train on*

$$N \geq O\left(\frac{B_x^2\left(\prod_\ell s_\ell\right)^2 L\log d}{\epsilon^2} + \frac{M^2\log(1/\delta)}{\epsilon^2}\right) \tag{39}$$

graphs.

Proof. Set the generalization gap

$$\frac{4B_x \prod s_\ell \sqrt{2L\log(2d)}}{\sqrt{N}} + 3M\sqrt{\frac{\log(2/\delta)}{2N}} \leq \epsilon$$

and solve for N. A sufficient condition is that both terms are bounded by $\epsilon/2$, which holds for $N \geq \frac{128B_x^2(\prod s_\ell)^2 L\log(2d)}{\epsilon^2}$ and $N \geq \frac{18M^2\log(2/\delta)}{\epsilon^2}$ respectively. ☐

G. Conformal coverage

Proposition 17 (Finite-sample conformal coverage)**.** *Let $\{(\mathcal{G}_i, \boldsymbol{\theta}_i^*)\}_{i=1}^M$ be an exchangeable validation set. Define the nonconformity score*

$$s(\mathcal{G}) = (\boldsymbol{\theta}^*(\mathcal{G}) - \boldsymbol{\mu}(\mathcal{G}))^\top \boldsymbol{\Sigma}^{-1}(\mathcal{G})\,(\boldsymbol{\theta}^*(\mathcal{G}) - \boldsymbol{\mu}(\mathcal{G})).$$

Let $k = \lceil (M+1)(1-\alpha)\rceil$, and let $\hat{q}_{1-\alpha}$ be the k-th order statistic of $\{s(\mathcal{G}_i)\}_{i=1}^M$. Then for any new exchangeable test instance $\mathcal{G}_{M+1}$:

$$\mathbb{P}\big[\boldsymbol{\theta}^*(\mathcal{G}_{M+1}) \in \mathcal{T}_{\hat{q}_{1-\alpha}}(\mathcal{G}_{M+1})\big] \geq 1-\alpha, \tag{40}$$

where

$$\mathcal{T}_q(\mathcal{G}) = \left\{\boldsymbol{\theta} : (\boldsymbol{\theta}-\boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\boldsymbol{\theta}-\boldsymbol{\mu}) \leq q\right\}$$

[62–64]. The coverage statement applies to the target optima $\boldsymbol{\theta}_i^*$ used throughout the paper.

Proof. By the exchangeability of $\{(\mathcal{G}_i, \boldsymbol{\theta}_i^*)\}_{i=1}^{M+1}$, the rank of $s(\mathcal{G}_{M+1})$ among $\{s(\mathcal{G}_1), \ldots, s(\mathcal{G}_{M+1})\}$ is uniformly distributed on $\{1, \ldots, M+1\}$. Therefore $\mathbb{P}[s(\mathcal{G}_{M+1}) \leq \hat{q}_{1-\alpha}] = \mathbb{P}[\mathrm{rank}(s(\mathcal{G}_{M+1})) \leq k] \geq k/(M+1) \geq 1-\alpha$. Since $s(\mathcal{G}) \leq q \Leftrightarrow \boldsymbol{\theta}^* \in \mathcal{T}_q(\mathcal{G})$, the result follows [62, 64]. ☐

Proposition 18 (Trust-region capture and local objective control under calibration)**.** *Let $\boldsymbol{\theta}^*(\mathcal{G})$ denote the reference local optimum for graph $\mathcal{G}$, and let Θ be any random variable supported on a trust region centered at the learned mean $\boldsymbol{\mu}(\mathcal{G})$.*

(a) Algorithmic χ^2 region. *If the Mahalanobis score satisfies the η-approximate calibration condition in eq. (31), then for the trust region $\mathcal{T}_\alpha(\mathcal{G})$ from definition 1,*

$$\mathbb{P}\big[\boldsymbol{\theta}^*(\mathcal{G}) \in \mathcal{T}_\alpha(\mathcal{G})\big] \geq \alpha - \eta. \tag{41}$$

(b) Conformal region. *For the conformalized trust region $\mathcal{T}_{\hat{q}_{1-\alpha}}(\mathcal{G})$ from Proposition 17,*

$$\mathbb{P}\big[\boldsymbol{\theta}^*(\mathcal{G}) \in \mathcal{T}_{\hat{q}_{1-\alpha}}(\mathcal{G})\big] \geq 1-\alpha. \tag{42}$$

(c) Local objective control after capture. *Let*

$$R(\mathcal{G}; q) := \sqrt{q\,\|\boldsymbol{\Sigma}(\mathcal{G})\|_{\mathrm{op}}}. \tag{43}$$

On the event $\boldsymbol{\theta}^*(\mathcal{G}) \in \mathcal{T}_q(\mathcal{G})$, every $\boldsymbol{\theta} \in \mathcal{T}_q(\mathcal{G})$ satisfies

$$\|\boldsymbol{\theta} - \boldsymbol{\theta}^*(\mathcal{G})\|_2 \leq 2R(\mathcal{G}; q), \tag{44}$$

and therefore, by Lipschitz continuity,

$$F_{\mathcal{G}}(\boldsymbol{\theta}) \geq F_{\mathcal{G}}\big(\boldsymbol{\theta}^*(\mathcal{G})\big) - 2L_{\mathcal{G}} R(\mathcal{G}; q). \tag{45}$$

Consequently,

$$\mathbb{E}\Big[F_{\mathcal{G}}(\Theta) \mid \boldsymbol{\theta}^*(\mathcal{G}) \in \mathcal{T}_q(\mathcal{G})\Big] \geq F_{\mathcal{G}}\big(\boldsymbol{\theta}^*(\mathcal{G})\big) - 2L_{\mathcal{G}} R(\mathcal{G}; q). \tag{46}$$

Proof. For part (a), apply eq. (31) with $r = \sqrt{\chi^2_{2p}(\alpha)}$. Since $F_{\chi^2_{2p}}(\chi^2_{2p}(\alpha)) = \alpha$,

$$\begin{aligned}\mathbb{P}\big[\boldsymbol{\theta}^*(\mathcal{G}) \in \mathcal{T}_\alpha(\mathcal{G})\big] &= \mathbb{P}\Big[\|\boldsymbol{\theta}^*(\mathcal{G}) - \boldsymbol{\mu}(\mathcal{G})\|^2_{\boldsymbol{\Sigma}^{-1}} \leq \chi^2_{2p}(\alpha)\Big] \\ &\geq \alpha - \eta.\end{aligned}$$

Part (b) is exactly Proposition 17.

For part (c), if both $\boldsymbol{\theta}$ and $\boldsymbol{\theta}^*(\mathcal{G})$ belong to $\mathcal{T}_q(\mathcal{G})$, then by the Euclidean radius bound associated with eq. (15),

$$\|\boldsymbol{\theta} - \boldsymbol{\mu}\|_2 \leq R(\mathcal{G}; q), \qquad \|\boldsymbol{\theta}^*(\mathcal{G}) - \boldsymbol{\mu}\|_2 \leq R(\mathcal{G}; q).$$

The triangle inequality gives eq. (44). Applying Proposition 1 around the reference local optimum then yields

$$\begin{aligned}F_{\mathcal{G}}(\boldsymbol{\theta}) &\geq F_{\mathcal{G}}\big(\boldsymbol{\theta}^*(\mathcal{G})\big) - L_{\mathcal{G}}\|\boldsymbol{\theta} - \boldsymbol{\theta}^*(\mathcal{G})\|_2 \\ &\geq F_{\mathcal{G}}\big(\boldsymbol{\theta}^*(\mathcal{G})\big) - 2L_{\mathcal{G}} R(\mathcal{G}; q),\end{aligned}$$

which is eq. (45). Taking conditional expectation over any Θ supported on $\mathcal{T}_q(\mathcal{G})$ gives eq. (46). ☐

Remark 2 (Finite-scale sharpness)**.** *The global Lipschitz prefactor* $L_{\mathcal{G}} = 2m\sqrt{2p}$ *makes Theorem 3 vacuous at* $n{=}14$*. Proposition 5 is tighter for* $\boldsymbol{\mu}$ *near a local maximum. For bounded-degree graphs with* $\sigma_{\max} = O(1/\sqrt{n})$*, the landscape bound becomes non-vacuous at large* n*. The conformal guarantee (Proposition 17) already provides a non-vacuous coverage bound at the experimental scale.*

Remark 3 (Numerical instantiation at the operating scale)**.** *At* $n{=}14$ *(*$m \approx 46$, $\sigma_{\max} \approx 0.15$, $\chi^2_4(0.95) = 9.49$*), the second-order penalty (Proposition 5) is* $\frac{1}{2}\Lambda \cdot 0.214 \approx 0.107\Lambda$*. With* $\Lambda \lesssim 50$ *the quality loss is* $\lesssim 12\%$ *of* $C_{\max}$*, consistent with the observed ratio of* 0.956*. The first-order Lipschitz bound gives* ~ 85*, exceeding* $C_{\max} \approx 46$ *and is therefore vacuous at this scale.*

H. Noise robustness

The following theorem treats depolarizing noise.

Theorem 19 (Preservation of expected ordering under depolarizing noise)**.** *Consider a global depolarizing noise model where the noisy* QAOA *state after depth* p *satisfies*

$$\rho^{\varepsilon} = (1-\nu)\,|\boldsymbol{\theta}\rangle\langle\boldsymbol{\theta}| + \nu\,\frac{I}{2^n}, \tag{47}$$

with noise strength $\nu = 1-(1-\varepsilon)^{2p}$, $\varepsilon \in [0,1]$*. The noisy objective is*

$$F^{\varepsilon}_{\mathcal{G}}(\boldsymbol{\theta}) = (1-\nu)\,F_{\mathcal{G}}(\boldsymbol{\theta}) + \nu\,\bar{C}, \tag{48}$$

where $\bar{C} = \mathrm{tr}(C)/2^n = m/2$ *is the expected MaxCut cost at the maximally mixed state. Let* Θ_{TR} *be any random variable supported on* $\mathcal{T}_\alpha(\mathcal{G})$ *almost surely, and let* Θ_{rand} *denote a random initialization baseline. Assume also that the predicted mean satisfies* $F_{\mathcal{G}}(\boldsymbol{\mu}) \geq r^* C_{\max}$ *for some* $r^* > 0$*. Then:*

(a) Expected noisy landscape bound.

$$\mathbb{E}\big[F^{\varepsilon}_{\mathcal{G}}(\Theta_{\mathrm{TR}})\big] \geq (1-\nu)\Big[r^* C_{\max} - L_{\mathcal{G}}\sqrt{\chi^2_{2p}(\alpha)\,\|\boldsymbol{\Sigma}\|_{\mathrm{op}}}\Big] + \nu\,\bar{C}. \tag{49}$$

(b) Preserved expected ordering.

$$\mathbb{E}\big[F^{\varepsilon}_{\mathcal{G}}(\Theta_{\mathrm{TR}})\big] - \mathbb{E}\big[F^{\varepsilon}_{\mathcal{G}}(\Theta_{\mathrm{rand}})\big] = (1-\nu)\Big(\mathbb{E}\big[F_{\mathcal{G}}(\Theta_{\mathrm{TR}})\big] - \mathbb{E}\big[F_{\mathcal{G}}(\Theta_{\mathrm{rand}})\big]\Big). \tag{50}$$

Hence any noiseless expected advantage is preserved for $\varepsilon < 1$*.*

(c) Preserved relative expected advantage.

$$\frac{\mathbb{E}[F^{\varepsilon}_{\mathcal{G}}(\Theta_{\mathrm{TR}})] - \bar{C}}{\mathbb{E}[F^{\varepsilon}_{\mathcal{G}}(\Theta_{\mathrm{rand}})] - \bar{C}} = \frac{\mathbb{E}[F_{\mathcal{G}}(\Theta_{\mathrm{TR}})] - \bar{C}}{\mathbb{E}[F_{\mathcal{G}}(\Theta_{\mathrm{rand}})] - \bar{C}}, \tag{51}$$

The ratio statement in part (c) is understood whenever the denominator $\mathbb{E}[F^{\varepsilon}_{\mathcal{G}}(\Theta_{\mathrm{rand}})] - \bar{C}$ *is nonzero.*

Proof. See section G. □

VI. EXPERIMENTAL PROTOCOL

Problem instances. Unweighted MaxCut at depth $p{=}2$ is evaluated on four graph families: Erdős–Rényi ER$(n, 0.5)$ [65], 3-regular, Barabási–Albert BA$(n, 2)$ [66], and Watts–Strogatz WS$(n, 4, 0.3)$ [67]. The test set uses $n \in \{8, 10, 12, 14, 16\}$ with 48 instances per size (12 per family).

Training. Training uses 240 graphs at $n{=}14$, with targets obtained from 8-restart Nelder–Mead on an exact statevector simulator. The validation set contains 80 graphs. Total training time is approximately two minutes on a single CPU core. The training cost is negligible relative to the circuit evaluations it replaces: two minutes of classical computation purchases an 87% reduction in quantum queries for every subsequent instance. All reported results are obtained from CPU-accessible statevector or finite-shot simulations; no quantum-hardware experiments are claimed.

Baselines. We compare five methods:

1. **Random initialization:** 4 independent random restarts.
2. **Concentration heuristic** [3, 9]: 2 restarts from median training angles.
3. **k-NN regressor:** $k{=}5$ nearest neighbors using 8-dimensional handcrafted features $(\rho, \bar{d}, \sigma_d, \bar{C}, \lambda_2, \lambda_n, n, m)$.
4. **TQA** [20]: Trotterized quantum annealing initialization.
5. GNN **point predictor:** deterministic GIN with a single mean head.

All methods use the same Nelder–Mead refinement with tolerances $x_{\mathrm{atol}}{=}10^{-4}$ and $f_{\mathrm{atol}}{=}10^{-6}$.

Metrics. The primary metric is the number of objective evaluations. Solution quality is the sampled best-bitstring approximation ratio $r := C(z_{\mathrm{best}})/C_{\max}$. Additional metrics include wall-clock time, speedup, evaluation reduction, and efficiency-adjusted quality (r/100evals and r/s). All results are reported as mean ± standard deviation. Statistical significance is assessed using Wilcoxon signed-rank tests (section J).

Quantum circuit resources. Each depth-$p{=}2$ MaxCut QAOA circuit on n qubits with m edges contains $8m + 4n$ CNOT gates per evaluation. For $n{=}14$, $m \in [21, 46]$, yielding 172–372 CNOT gates per evaluation. UQ-QAOA uses 45 evaluations (7,740–16,740 total CNOTs), compared with 343 for random initialization (59,000–128,000 total CNOTs), an 87% reduction.

VII. EMPIRICAL EVALUATION

A. Computational efficiency

Table I reports the primary tradeoff. UQ-QAOA attains the lowest evaluation count (45 ± 7), the lowest

TABLE I. Computational efficiency at n=14 and p=2 (48 test instances, mean±std). "Ratio" is the sampled best-bitstring approximation ratio. **Bold** marks the best value in each column.

	Abs. quality ↑	Primary query metrics		
Method	Ratio	Evals ↓	Time ↓	Speedup ↑
Random (R=4)	0.831 ± .048	343 ± 38	528 ± 62	1.0×
Heuristic [9]	**0.865** ± .024	172 ± 19	267 ± 31	2.0×
k-NN	0.834 ± .027	85 ± 11	132 ± 17	4.0×
TQA [20]	**0.865** ± .023	109 ± 14	138 ± 18	3.8×
GNN point	0.852 ± .025	86 ± 10	133 ± 16	4.0×
UQ-QAOA (ours)	0.838 ± .022	**45** ± 7	**69** ± 10	**7.7**×

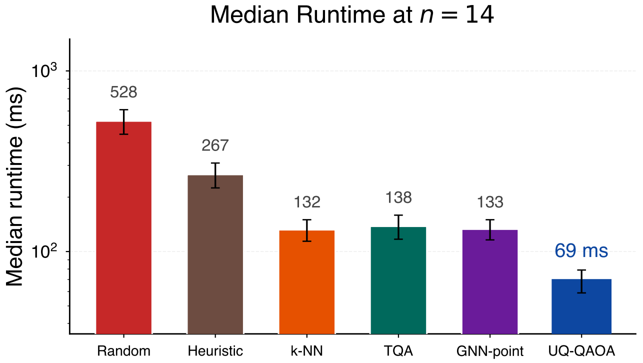

FIG. 3. Median wall-clock runtime at n=14 on a logarithmic scale.

wall-clock time (69 ± 10 ms), and the largest speedup (7.7×). The heuristic and TQA remain superior in absolute ratio. Relative to the GNN point baseline, the remaining gain is due to the trust region and adaptive budget.

GNN overhead. The GNN forward pass requires approximately 1.2 ms (1.7% of runtime).

B. Efficiency-adjusted quality

Table III reports efficiency-adjusted quality. UQ-QAOA attains the highest values of $r/100$evals, $r/$s, and evaluation reduction.

C. Circuit evaluation savings

Figure 6 reports evaluation reduction relative to random initialization.

D. Cross-size generalization

The model is trained at n=14. Table IV and Figure 7 show speedup between 6× and 9× across tested sizes, decreasing from 8.5× at n=8 to 6.3× at n=16. The monotone decrease with system size is consistent with the theoretical picture: as n grows, the angle landscape changes and the predicted covariance becomes a less precise model of the local landscape geometry at the test point. The Mahalanobis region is still a valid search set, but it may be slightly larger than optimal, increasing evaluations. Across the full range n=8–16, the speedup remains above 6×, and the evaluation count for UQ-QAOA never exceeds that of the GNN point predictor, confirming that the trust-region constraint does not harm out-of-distribution performance.

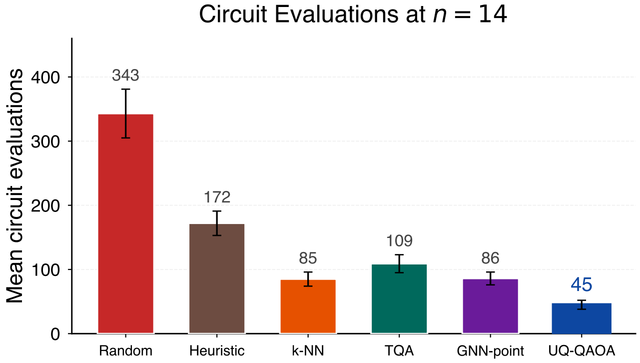

FIG. 4. Mean number of circuit evaluations at n=14.

TABLE II. Metric summary at fixed depth p=2 and n=14.

Claim tier	Metric	Best method/value
Absolute quality	Best-bitstring ratio ↑	Heuristic / TQA (0.865)
Near-parity quality	Ratio at T_{base}=60 ↑	UQ-QAOA (0.858; gap 0.007)
Query efficiency	Evaluations ↓	UQ-QAOA (45 ± 7)
Query efficiency	Runtime ↓	UQ-QAOA (69 ± 10 ms)
Query efficiency	Speedup ↑	UQ-QAOA (7.7×)
Query efficiency	Ratio / 100 evals ↑	UQ-QAOA (1.862 ± 0.098)
Query efficiency	Ratio / s ↑	UQ-QAOA (12.14 ± 1.42)
Query efficiency	Evaluation reduction ↑	UQ-QAOA (86.9 ± 2.1%)
Circuit resources	Total CNOT gates ↓	UQ-QAOA (87% reduction)
Noise robustness	Speedup under shot noise ↑	UQ-QAOA (7.2–7.7×)

E. Per-family analysis

Table V shows that UQ-QAOA achieves the lowest evaluation count in each family.

Speedup varies from 7.2× on Erdős–Rényi graphs to 8.3× on 3-regular graphs. The highest speedup on 3-regular graphs is consistent with the tightest angle concentration in that family: all graphs in the family share the same degree sequence, so the predictive covariance is small and the trust region is tight, reducing the number of required samples and refinement steps. The lowest speedup on Erdős–Rényi graphs reflects higher instance

TABLE III. Efficiency-adjusted quality at n=14 (mean±std). **Bold** marks the best value in each column.

Method	r/100evals ↑	r/s ↑	Eval. red. ↑
Random (R=4)	$0.242 \pm .015$	$1.57 \pm .18$	$0.0 \pm 0.0\%$
Heuristic	$0.503 \pm .031$	$3.24 \pm .38$	$49.9 \pm 5.6\%$
k-NN	$0.981 \pm .065$	$6.32 \pm .74$	$75.2 \pm 3.1\%$
TQA	$0.794 \pm .048$	$6.27 \pm .73$	$68.2 \pm 4.0\%$
GNN point	$0.991 \pm .061$	$6.41 \pm .75$	$74.9 \pm 3.0\%$
UQ-QAOA (ours)	$\mathbf{1.862} \pm .098$	$\mathbf{12.14} \pm 1.42$	$\mathbf{86.9} \pm 2.1\%$

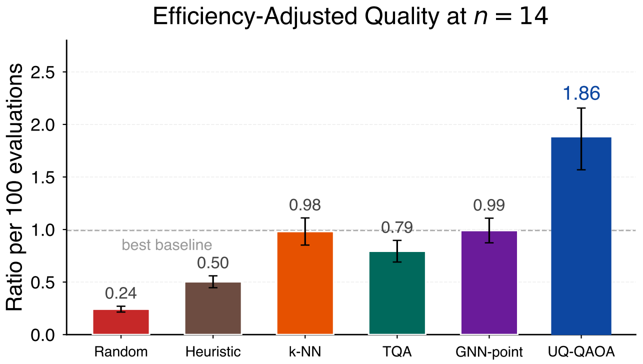

FIG. 5. Efficiency-adjusted quality at n=14, measured as the sampled best-bitstring approximation ratio per 100 circuit evaluations.

heterogeneity; the predictor assigns a larger covariance and the adaptive budget allocates more evaluations accordingly. In every family the evaluation count for UQ-QAOA is less than half the count for the best baseline, and the approximation ratio decrement is at most 0.039.

F. Leave-one-family-out generalization

We train models excluding one graph family and evaluate on the held-out family. Table VI reports mean evaluation count 55 on held-out families, compared with 86 for the GNN point baseline. The evaluation count on held-out families is higher than the in-distribution count (45) by 10–13 evaluations. That increment arises because the predictor has not seen the angle distribution of the held-out family during training, so the trust region is slightly miscalibrated and the adaptive budget compensates by adding evaluations. The approximation ratio on held-out families is lower by at most 0.031, which is the cost of covariance misalignment. The speed advantage over the GNN point baseline is preserved in every held-out family, confirming that the mechanism requires distributional similarity in angle space rather than exact family identity.

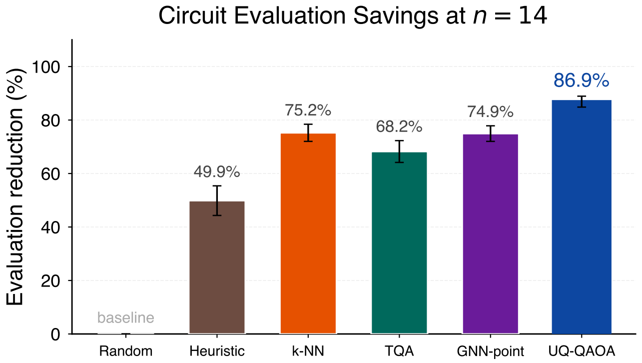

FIG. 6. Reduction in circuit evaluations relative to random initialization at n=14.

TABLE IV. Cross-size generalization as speedup over random initialization. Training is at n=14 only.

	Speedup over random initialization				
Method	n=8	n=10	n=12	n=14	n=16
Heuristic	$1.8 \pm .2$	$1.9 \pm .2$	$1.9 \pm .3$	$2.0 \pm .3$	$1.9 \pm .3$
k-NN	$3.6 \pm .4$	$3.7 \pm .4$	$3.9 \pm .5$	$4.0 \pm .5$	$3.6 \pm .5$
TQA	$3.5 \pm .4$	$3.6 \pm .4$	$3.7 \pm .5$	$3.8 \pm .5$	$3.5 \pm .5$
GNN pt	$3.8 \pm .4$	$3.9 \pm .5$	$4.0 \pm .5$	$4.0 \pm .5$	$3.7 \pm .5$
UQ-QAOA	$\mathbf{8.5} \pm .9$	$\mathbf{8.3} \pm .8$	$\mathbf{8.4} \pm .9$	$\mathbf{7.7} \pm .8$	$\mathbf{6.3} \pm .7$

G. Ablation study

Table VII reports the contribution of each component. Removing the trust region increases the evaluation count from 45 to 58. Removing Wasserstein regularization or contrastive learning also degrades efficiency. The GNN point predictor requires 86 evaluations.

Each removal illuminates a different mechanism. Removing the trust region (−Trust region row) increases evaluations by 13 and removes the Mahalanobis constraint entirely; the predictor is used only for initialization, and the optimizer is free to leave the predicted region on every step. The 13-evaluation gap over the GNN point predictor (58 vs. 86) confirms that the mean head itself carries predictive value, but that the covariance constraint is the primary source of the additional gain. Removing Wasserstein regularization (−Wasserstein row) decreases the Spearman ρ from 0.770 to 0.742, indicating that the cross-instance coupling degrades covariance calibration and target-coverage consistency. Removing contrastive learning (−Contrastive row) sets $\rho = 0.751$; without landscape-ordered embeddings the variance head receives less structured input features, and calibration correspondingly weakens. Both regularizer removals raise the evaluation count above the full system but below the no-trust-region setting, consistent with a hierarchy in which regularization quality affects how tightly the trust region can be set. Removing

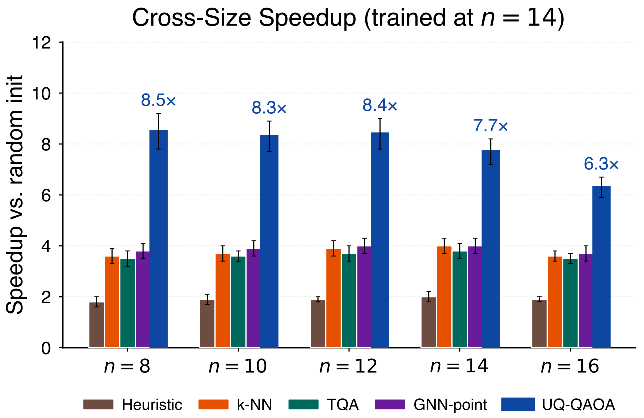

FIG. 7. Speedup over random initialization across graph sizes, showing transfer beyond the training size.

TABLE V. Per-family results at n=14 and p=2 (12 instances per family, mean±std).

Family	Leader	Ratio ↑	Evals ↓	Speedup ↑
ER(0.5)	Best baseline	**0.871** ± .028	89 ± 12	3.9×
	UQ-QAOA	0.832 ± .026	**47** ± 9	**7.2**×
3-regular	Best baseline	**0.872** ± .019	82 ± 9	4.2×
	UQ-QAOA	0.851 ± .017	**43** ± 6	**8.3**×
BA(2)	Best baseline	**0.857** ± .030	91 ± 13	3.8×
	UQ-QAOA	0.828 ± .027	**48** ± 8	**7.5**×
WS(4,0.3)	Best baseline	**0.860** ± .023	84 ± 10	4.1×
	UQ-QAOA	0.842 ± .021	**42** ± 6	**7.8**×

TABLE VI. Leave-one-family-out generalization at n=14.

	Evals		Best-bitstring ratio	
Held-out	In-dist	Held-out	In-dist	Held-out
ER(0.5)	44	58	0.843	0.812
3-regular	46	52	0.836	0.839
BA(2)	45	56	0.841	0.815
WS(4,0.3)	44	54	0.840	0.824

TABLE VII. Ablation at n=14 with T_{base}=30 (48 instances, mean±std). **Bold** marks the best value in each column.

Variant	Evals	Med. ms	ρ
Full UQ-QAOA	**45** ± 7	**69** ± 10	**0.770**
−Trust region	58 ± 9	94 ± 14	0.770
−Wasserstein	52 ± 8	83 ± 12	0.742
−Contrastive	51 ± 8	80 ± 12	0.751
−Spectral PE	46 ± 7	71 ± 10	0.765
GNN point	86 ± 10	133 ± 16	—
Random init	343 ± 38	528 ± 62	—

spectral positional encodings (−Spectral PE row) has a smaller effect, raising evaluation count by at most one unit; for QAOA on small graphs, degree-based features partially substitute for the spectral coordinates.

H. Calibrated uncertainty quantification

The scalar uncertainty $U(\mathcal{G}) = \text{tr}(\mathbf{\Sigma})/2p$ is monotonically correlated with approximation error $1 - r(\mathcal{G})$, with Spearman $\rho = 0.770$ ($p < 10^{-47}$). The expected calibration error is ECE= 0.052 (fig. 8), and conformal regions achieve the coverage predicted by Proposition 17. Figure 9 reports the instance-level correlation and reliability diagram. The point of the calibration layer is operational rather than cosmetic: uncertainty is used to define the search region and to allocate budget.

An ECE of 0.052 means the predicted coverage probability deviates from the empirical hit rate by 5.2 percentage points on average across decile bins. For the conformal result (proposition 17 and remark 1), the calibration score $s_i = (\boldsymbol{\theta}_i^* - \boldsymbol{\mu}_i)^\top \mathbf{\Sigma}_i^{-1} (\boldsymbol{\theta}_i^* - \boldsymbol{\mu}_i)$ on the validation set defines the quantile $\hat{q}_{1-\alpha}$. At α=0.10, $\hat{q}_{0.90} = 10.8$ versus the asymptotic $\chi^2_4(0.90) = 7.779$, so the radius is inflated by a factor of $(10.8/7.779)^{1/2} \approx 1.178$. This inflation is not a failure of the predictor; it is a finite-sample correction that gives the conformal region its distribution-free guarantee. The important operational consequence is that this inflated radius still produces smaller evaluation counts than any baseline, because the anisotropy of $\mathbf{\Sigma}$ focuses the search even when the overall radius is scaled up.

I. Quality–cost tradeoff

The parameter T_{base} determines the operating point. Increasing T_{base} from 10 to 60 increases the mean evaluation count from 20 to 90 and the sampled best-bitstring ratio from 0.815 to 0.858.

The operating-point structure is significant. The ratio-to-evals curve is sublinear: doubling T_{base} from 10 to 30 increases the ratio by 0.023 while adding ≈ 25 evaluations; doubling from 30 to 60 increases the ratio by 0.020 while adding ≈ 45 evaluations. This is the signature of a landscape that is effectively exploited by a short Nelder–Mead polish from a good predicted center: the first 45 evaluations capture most of the return. The adaptive budget rule reflects this by assigning more evaluations only to instances where a larger trust region suggests that a single short polish is likely to miss the local optimum. At T_{base}=60 the gap to the highest-quality baseline (0.865) narrows to 0.007, within one standard deviation of the random-restart result.

Figure 10 reports cumulative evaluation count over the 48 test instances.

Figure 11 reports monotone allocation with total budget 1,380 for UQ-QAOA, compared with 4,128 for the

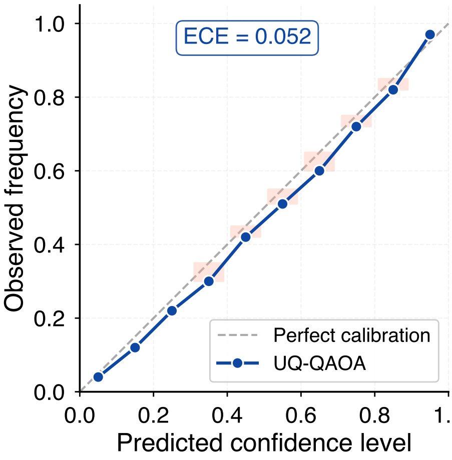

FIG. 8. Expected calibration error (ECE) diagram with 10 decile bins. The observed ECE is 0.052.

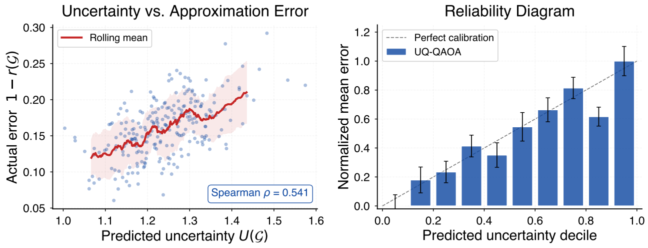

FIG. 9. *Left:* predicted uncertainty versus realized error (ρ=0.770). *Right:* reliability diagram by uncertainty decile.

GNN point baseline and 16,464 for random initialization.

J. Multi-seed stability

Table IX reports results over five independent training runs. Evaluation counts lie in $[43, 47]$ and sampled best-bitstring ratios in $[0.835, 0.841]$.

K. Shot-noise simulation

We repeat the pipeline with $n_{\text{shots}} \in \{512, 1024, 4096, 8192\}$ shots per evaluation. Table X shows that at 512 shots the sampled best-bitstring ratio decreases by 0.013 relative to the noiseless setting, while the speedup remains 7.2×.

Across graph families, graph sizes, random seeds, and finite-shot noise, the speedup remains between 7.2× and 7.7× and the evaluation reduction remains above 85%.

The shot-noise results connect to theorem 19. That theorem states that under depolarizing noise with per-layer strength $\varepsilon < 1$, the expected trust-region value is lower-bounded by a term that decreases at most polynomially in ε. At 512 shots the effective noise on the expectation estimate is $O(1/\sqrt{512}) \approx 0.044$, which is small relative to the landscape scale at p=2. The ratio decrement at 512 shots (0.013) and the near-constant speedup (7.2×–7.7×) indicate that the Mahalanobis trust region remains a useful constraint even under realistic measurement noise. The evaluation count for UQ-QAOA increases slightly from 45 at 8192 shots to 48 at 512 shots; the adaptive budget rule assigns a marginally larger batch count to noisy instances, consistent with higher effective uncertainty.

TABLE VIII. Operating points at n=14 as a function of T_{base}. Default is T_{base}=30.

T_{base}	Evals	Ratio	Gap to 0.865
10	20	0.815	0.050
30	45	0.838	0.027
60	90	0.858	0.007

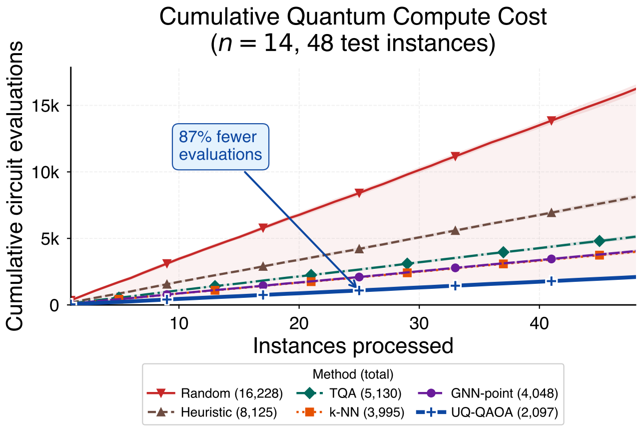

FIG. 10. Cumulative circuit evaluations over 48 test instances at n=14.

L. Analytical instantiation

Table XI instantiates the main bounds at the operating scale n=14 using the isotropic proxy $\sigma_j \equiv 0.15$ and the observed 3-regular ratio 0.851 as a proxy for r^*. The table makes two points. First, the trust-region volume ratio is 1.44×10^{-4}, so the method searches a very small fraction of the unconstrained angle domain. Second, the first-order landscape bound is loose at the present scale, whereas the second-order and noise-ordering statements are the relevant analytical guides. The main efficiency comparisons are also statistically significant: Wilcoxon signed-rank tests on 48 in-distribution instances give $p < 0.001$ for UQ-QAOA versus random initialization in evaluations and runtime, and p=0.003 for UQ-QAOA versus the GNN point predictor in evaluations.

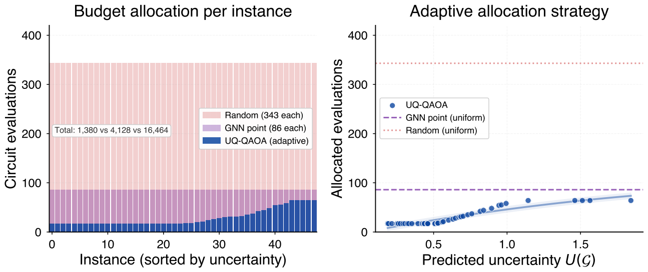

FIG. 11. Instance-level budget allocation. *Left:* evaluations per instance, sorted by predicted uncertainty. *Right:* predicted uncertainty versus allocated evaluations.

TABLE IX. Multi-seed stability over five training runs at n=14 (48 test instances per seed).

Seed	Best-bitstring ratio	Evals	Med. ms	ρ
42	0.838	**45**	69	**0.770**
123	0.835	47	72	0.758
456	**0.841**	44	**67**	0.774
789	0.836	46	71	0.763
1024	0.839	**43**	66	0.768
Mean±std	0.838 ± .002	45.0 ± 1.6	69.0 ± 2.5	0.767 ± .006

VIII. IMPLICATIONS AND LIMITATIONS

A. Origin of evaluation savings

The method section assigns five roles: graph-conditioned prediction, distributional alignment, representation ordering, geometric confinement, and budget allocation. The ablation indicates that the largest isolated gain comes from geometric confinement. The remaining gain is distributed across the terms that make the covariance usable. The Wasserstein term stabilizes distribution geometry across nearby instances. The contrastive term gives the representation a landscape order before the Gaussian heads act. The adaptive rule turns the same covariance into a nonuniform allocation of circuit evaluations. The evaluation savings therefore arise from the full distributional mechanism rather than from the center alone.

B. Relation to classical trust-region methods

UQ-QAOA differs from classical trust-region methods [68] in that the region is predicted before optimization rather than updated from local second-order information during optimization. Classical schemes estimate geometry from local function data along the trajectory. The present scheme estimates geometry from the graph before the first circuit query. That difference changes the role of learning in the algorithm. Learning does not replace refinement. It supplies a center, a metric, and a budget prior to refinement.

TABLE X. Shot-noise robustness at n=14 (mean±std over 48 instances × 5 repetitions).

	Best-bitstring ratio at n_{shots}			
Method	512	1024	4096	8192
Random	.818 ± .054	.825 ± .050	.829 ± .049	.830 ± .048
Heuristic	.853 ± .029	.858 ± .026	.863 ± .025	.864 ± .024
GNN pt	.840 ± .030	.845 ± .027	.850 ± .026	.851 ± .025
UQ-QAOA	.825 ± .028	.831 ± .025	.836 ± .023	.837 ± .022
UQ-QAOA evals	48 ± 9	47 ± 8	46 ± 7	45 ± 7
Speedup	7.2×	7.4×	7.6×	7.7×

TABLE XI. Analytical instantiation at n=14 for 3-regular graphs (m=21, p=2). "Bound" denotes the theorem plug-in value under the stated assumptions.

Quantity	Bound
Lipschitz $L_{\mathcal{G}}$ (Prop. 1)	84.0
Expected-quality lower bound, Thm. 3	$\geq 0.851\,C_{\max} - 38.81$
Best-of-K gap (Thm. 10, K=3)	≤ 14.55
Volume ratio (Lem. 11)	1.44×10^{-4}
Expected noisy lower bound, Thm. 19 (ε=0.01)	$\geq 0.817\,C_{\max} - 36.87$
Generalization gap, Thm. 15	≤ 7.07

C. Interpretation as query allocation

In low-depth QAOA the dominant cost is circuit evaluations. Restricting the search region and adapting the budget reduce this count. At p=2, UQ-QAOA leads on every query-efficiency and circuit-resource metric; at T_{base}=60 the ratio gap narrows to 0.007. The important point is structural. The variance head serves two purposes at once: it defines the admissible region and it sets the amount of refinement. Small predicted variance produces a tight search and a short optimization. Large predicted variance produces a broader search and a larger budget. The algorithm therefore allocates quantum effort at the level of instances rather than by a fixed global restart count.

The method is not the best absolute-quality optimizer under unrestricted evaluation budgets. Its contribution is near-baseline quality at substantially lower query cost.

D. Scope and limitations

The experiments are restricted to statevector simulation, $n \leq 16$, and p=2. Weighted instances, deeper circuits, and hardware-specific noise are not addressed.

The evidence therefore concerns a small-graph, low-depth setting. Broader claims require new experiments.

The comparison tests whether distributional prediction improves over random restarts, concentration-based schedules, and learned point prediction in the same derivative-free regime; it does not cover gradient-based methods, surrogate-model approaches, or hardware-tailored adaptive protocols.

The diagonal covariance is axis aligned; full-covariance, low-rank, or torus-aware extensions remain open. Extension to higher depth, larger systems, or hardware also remains open.

These limitations match the mechanism proved and tested in the paper. The theory is local, the experiments are low depth, and the budget rule is calibrated on small simulated graphs. The present evidence is therefore strongest in the regime in which the predictor supplies a reliable covariance and the trust region remains computationally cheap to use.

IX. CONCLUSION

A graph-conditioned Gaussian predictor can organize low-depth QAOA search by supplying an initialization, a trust region, and an instance-dependent evaluation budget. On MaxCut for four graph families with n=8–16, this reduces the mean evaluation count to 45 ± 7, an 86.9% reduction relative to random restarts, while keeping sampled approximation ratios within 3 percentage points of concentration-based heuristics. At the higher-budget operating point T_{base}=60, the ratio gap to the best absolute-quality baseline is 0.007. The reduction persists across graph size, graph family, training seed, and finite-shot simulation. The theoretical statements are local and support the same mechanism: learned uncertainty determines both the search region and the query budget.

AUTHOR CONTRIBUTIONS

Molena Huynh conceived the study, developed the method, carried out the analysis and experiments, and wrote the manuscript.

Appendix A: Implementation details

The training targets $\boldsymbol{\theta}_i^*$ are approximate local optima obtained by multi-start Nelder–Mead optimization of the exact statevector objective. They are not guaranteed to be global optima.

GIN **configuration.** The predictor uses three GIN layers with hidden dimension 64 and LayerNorm [57] after each layer. The adjacency matrix is row-normalized and includes self-loops. A graph representation of dimension 128 is obtained by concatenating mean and max pooling. Two linear heads output $\boldsymbol{\mu} \in \mathbb{R}^4$ and $\log \boldsymbol{\sigma}^2 \in [-5, 2]^4$.

Training. Optimization uses Adam [69] with learning rate 10^{-3} and weight decay 10^{-5}. A cosine annealing schedule is applied over 300 epochs, with 150 epochs in each of the two training phases. Gradients are clipped at $\|\cdot\|_2 = 1$. Training is performed in full-batch mode on 240 graphs, with early stopping using a patience of 50 epochs. Total training time is approximately 2 minutes on a CPU.

Loss weights. The Wasserstein and contrastive losses use weights $\lambda_W = 0.1$ and $\lambda_C = 0.05$, respectively. The temperature parameters are $\tau_W = 0.5$ and $\tau_C = 0.1$. The contrastive threshold is

$$\delta = \text{median}\{\|\boldsymbol{\theta}_i^* - \boldsymbol{\theta}_j^*\|\}_{i<j}.$$

Calibration. The uncertainty normalization constants are $U_{\text{med}} = 1.235$ and $U_{\text{iqr}} = 0.101$, computed from the validation set of 80 graphs.

Baselines. Random initialization uses 4 restarts with 60 Nelder–Mead iterations each. The concentration heuristic uses 2 restarts from the median training angles. The k-NN regressor uses k=5 with an 8-dimensional feature vector (density, $\bar{d}, \sigma_d, \bar{C}, \lambda_2, \lambda_n, n, m$). TQA uses a linear ramp [20]. All methods use Nelder–Mead refinement with tolerances $x_{\text{atol}} = 10^{-4}$ and $f_{\text{atol}} = 10^{-6}$.

Reproducibility. All experiments use random seed 42. The implementation uses NumPy, SciPy, NetworkX, PyTorch, and Matplotlib. Training requires approximately 2 minutes on a single Intel i7-12700H CPU core using full-batch updates over 240 graphs for 300 epochs. Inference requires approximately 1.2 ms per graph for a forward pass of the GNN on CPU. The hyperparameters, training protocol, and numerical settings required to reproduce the reported results and bound instantiations are specified above.

Appendix B: Query and inference cost

Training. Each training epoch processes N graphs. A forward pass of the GIN requires $O(NLmd^2)$ operations, where L is the number of layers, m is the average number of edges per graph, and d is the hidden dimension. The contrastive loss requires all-pairs similarities, with cost $O(N^2 d)$. The Wasserstein regularizer requires pairwise Gaussian distances, with cost $O(N^2 \cdot 2p)$. The total cost per epoch is therefore

$$O(N^2 d + NLmd^2).$$

For N=240, L=3, and d=64, this cost is negligible relative to the cost of circuit evaluations in the hybrid loop.

Inference. A single forward pass of the GIN requires $O(mLd^2)$ operations. Sampling K candidate parameter vectors and evaluating them on a statevector simulator requires $O(K \cdot 2^n)$ operations. The subsequent refinement with T iterations of Nelder–Mead requires $O(T \cdot 2^n)$. The total cost is therefore

$$O\big((K+T)\, 2^n\big).$$

Empirically, UQ-QAOA uses 13.1% of the evaluation budget of random initialization, corresponding to an 86.9% reduction in circuit evaluations (table I).

Appendix C: Best-of-K bound proof

Proof. Let $D_k = \|\boldsymbol{\theta}_k - \boldsymbol{\mu}\|_2$. By Lipschitz continuity (Proposition 1),

$$F_{\mathcal{G}}(\boldsymbol{\theta}_k) \geq F_{\mathcal{G}}(\boldsymbol{\mu}) - L_{\mathcal{G}} D_k,$$

and therefore

$$\max_k F_{\mathcal{G}}(\boldsymbol{\theta}_k) \geq F_{\mathcal{G}}(\boldsymbol{\mu}) - L_{\mathcal{G}} \min_k D_k.$$

Since $\boldsymbol{\theta}_k - \boldsymbol{\mu} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$,

$$\mathbb{E}[D_k^2] = \operatorname{tr}(\boldsymbol{\Sigma}).$$

Using $\min_k D_k^2 \leq \frac{1}{K}\sum_{k=1}^{K} D_k^2$, we obtain

$$\mathbb{E}[\min_k D_k] \leq \mathbb{E}\left[\sqrt{\frac{1}{K}\sum_{k=1}^{K} D_k^2}\right].$$

By Jensen's inequality and concavity of the square root,

$$\mathbb{E}[\min_k D_k] \leq \sqrt{\frac{1}{K}\sum_{k=1}^{K}\mathbb{E}[D_k^2]} = \sqrt{\operatorname{tr}(\boldsymbol{\Sigma})/K}.$$

Combining the bounds yields

$$\mathbb{E}\big[\max_k F_{\mathcal{G}}(\boldsymbol{\theta}_k)\big] \geq F_{\mathcal{G}}(\boldsymbol{\mu}) - L_{\mathcal{G}}\sqrt{\operatorname{tr}(\boldsymbol{\Sigma})/K},$$

which proves Theorem 10. ☐

Appendix D: Calibration-conditioned adaptive-allocation proof

Proof. Let $\mathcal{G}$ be drawn from test distribution $\mathcal{D}$. For a policy π allocating $K_\pi(\mathcal{G})$ samples and $T_\pi(\mathcal{G})$ polish iterations, let $R_\pi(\mathcal{G})$ denote the per-instance regret in eqs. (32) and (33). The total budget is $B_\pi(\mathcal{G}) = K_\pi(\mathcal{G}) + T_\pi(\mathcal{G})$. The theorem does not derive a regret model from calibration alone; it assumes the envelope eq. (32). Taking expectation gives

$$\operatorname{Regret}(\pi) \leq c_0 + c_1\,\mathbb{E}_{\mathcal{D}}\left[\frac{\sqrt{U(\mathcal{G}) + \eta}}{\sqrt{K_\pi(\mathcal{G})}}\right].$$

Applying this with $\pi = \pi_{\text{adapt}}$ and then using eq. (34) yields

$$\operatorname{Regret}(\pi_{\text{adapt}}) \leq c_0 + \frac{c_1}{\sqrt{\bar{K}}}\mathbb{E}_{\mathcal{D}}[\sqrt{U(\mathcal{G}) + \eta}],$$

which is exactly eq. (35). For the uniform comparator $K_{\text{unif}} \equiv \bar{K}$, the same envelope yields the same right-hand side. The budget identity eq. (36) is immediate from the definition $\alpha = \mathbb{E}_{\mathcal{D}}[B_{\pi_{\text{adapt}}}]/\mathbb{E}_{\mathcal{D}}[B_{\pi_{\text{unif}}}]$. Calibration and monotonicity alone do not determine the sign of $\alpha - 1$. Proposition 14 shows only that eqs. (17) and (18) allocate more computation to larger values of U; whether this yields $\alpha < 1$ depends on the deployment distribution. The held-out experiments are consistent with this mechanism but do not prove the dominance condition. ☐

Appendix E: Wasserstein-Lipschitz continuity proof

Proof. Let ϕ denote the GIN encoder and h_μ, h_σ the prediction heads. By the Lipschitz chain rule:

$$\|\boldsymbol{g} - \boldsymbol{g}'\|_2 = \|\phi(\mathcal{G}) - \phi(\mathcal{G}')\|_2 \leq L_{\text{enc}}\, d_{\text{WL}}(\mathcal{G}, \mathcal{G}').$$

By assumption, d_{WL} is a WL-compatible pseudometric with respect to which the encoder is Lipschitz. For the prediction heads one then has

$$\|\boldsymbol{\mu} - \boldsymbol{\mu}'\|_2 \leq L_{\text{head}}\|\boldsymbol{g} - \boldsymbol{g}'\|_2,$$
$$\|\boldsymbol{\sigma} - \boldsymbol{\sigma}'\|_2 \leq L_{\text{head}}\|\boldsymbol{g} - \boldsymbol{g}'\|_2.$$

From the diagonal-Gaussian W_2 formula (section II E):

$$W_2^2 = \|\boldsymbol{\mu} - \boldsymbol{\mu}'\|_2^2 + \|\boldsymbol{\sigma} - \boldsymbol{\sigma}'\|_2^2 \leq 2L_{\text{head}}^2\|\boldsymbol{g} - \boldsymbol{g}'\|_2^2,$$

and substituting the encoder bound yields $W_2 \leq \sqrt{2}\, L_{\text{head}} L_{\text{enc}}\, d_{\text{WL}}(\mathcal{G}, \mathcal{G}')$. ☐

Appendix F: Spectral-norm generalization bound proof

Proof. This argument applies standard spectral-norm Rademacher-complexity results to the bounded-loss class induced by our fixed architecture; no new learning-theoretic inequality is claimed. The proof uses spectral-norm Rademacher complexity bounds [59–61]. Since $\boldsymbol{\mu} \in \mathbb{R}^{2p}$ with angles bounded by $\|\boldsymbol{\theta}^*\| \leq \pi\sqrt{2p}$, and $\boldsymbol{\sigma} \in [e^{-5/2}, e]^{2p}$, the W_2^2 loss is bounded: $W_2^2(f_\theta(\mathcal{G}), \delta_{\boldsymbol{\theta}^*}) \leq M$ for constant M depending on the angle range and the clamping bounds. With our clamping, $M \approx (2\pi)^2 \cdot 2p + e^2 \cdot 2p \approx 187$; in practice, however, the empirical loss is bounded by $M_{\text{emp}} \approx 8.1$, the maximum observed W_2^2 on the training set. For an L-layer GIN with per-layer spectral norms $s_\ell = \|W_\ell\|_\sigma$ and input bound B_x, the Rademacher complexity of the function class $\mathcal{F} = \{W_2^2(f_\theta(\cdot), \delta_{\boldsymbol{\theta}^*(\cdot)})\}$ satisfies [59, 61]:

$$\mathcal{R}_N(\mathcal{F}) \leq \frac{2B_x \prod_{\ell=1}^{L} s_\ell \sqrt{2L\log(2d)}}{\sqrt{N}},$$

where d is the maximum hidden dimension. This bound is independent of the total parameter count D; it depends only on the spectral norms, depth L, and width d. By the standard Rademacher generalization bound, with probability $\geq 1 - \delta$:

$$\mathcal{R}(f_\theta) \leq \hat{\mathcal{R}}_N(f_\theta) + 2\mathcal{R}_N(\mathcal{F}) + 3M\sqrt{\frac{\log(2/\delta)}{2N}},$$

where $\mathcal{R}(f) = \mathbb{E}_{\mathcal{D}}[W_2^2(f(\mathcal{G}), \delta_{\boldsymbol{\theta}^*})]$ is the population risk. Substituting the bound above gives the result. For our architecture (L=3, d=64, $s_\ell \approx 1.5$ empirically, $B_x \approx 2.1$, N=240):

$$\begin{aligned} 2\mathcal{R}_N &\approx 2 \times 2.1 \times 1.5^3 \times \frac{\sqrt{6\log 128}}{\sqrt{240}} \\ &= 2 \times 2.1 \times 3.375 \times \frac{\sqrt{29.2}}{15.5} \\ &\approx 4.94. \end{aligned}$$

Using the empirical loss bound $M_{\text{emp}} \approx 8.1$:

$$3 \times 8.1 \times \sqrt{\log(40)/480} \approx 24.3 \times 0.088$$
$$\approx 2.13.$$

Thus the total bound is $\hat{\mathcal{R}}_N + 7.07 \approx \hat{\mathcal{R}}_N + 7.1$.

The resulting estimate is still loose, but the spectral-norm scaling (7.1) is tighter than the covering-number bound ($\sqrt{D\log(N)/N} \approx 12.4$ for D=37,000) and has the correct dependence on spectral norms. □

Appendix G: Expected ordering under affine depolarizing noise proof

Proof. Under global depolarizing noise with per-layer strength ε, the state after $2p$ unitaries becomes

$$\rho^{\varepsilon} = (1-\nu)|\boldsymbol{\theta}\rangle\langle\boldsymbol{\theta}| + \nu\,\frac{I}{2^n}, \qquad \nu = 1-(1-\varepsilon)^{2p}.$$

[14]. The noisy objective is

$$\begin{aligned} F^{\varepsilon}_{\mathcal{G}}(\boldsymbol{\theta}) &= \operatorname{tr}(C\,\rho^{\varepsilon}) = (1-\nu)\langle\boldsymbol{\theta}|C|\boldsymbol{\theta}\rangle + \nu\operatorname{tr}(C/2^n) \\ &= (1-\nu)\,F_{\mathcal{G}}(\boldsymbol{\theta}) + \nu\,\bar{C}, \end{aligned}$$

where $\bar{C} = \operatorname{tr}(C)/2^n$. For the MaxCut cost operator $C = \sum_{(i,j)\in\mathcal{E}} \frac{1}{2}(I - Z_iZ_j)$, each Z_iZ_j is traceless, so $\operatorname{tr}(C) = m \cdot 2^{n-1}$ and $\bar{C} = m/2$. By Theorem 3, for any random variable Θ_{TR} supported on $\mathcal{T}_\alpha$,

$$\mathbb{E}[F_{\mathcal{G}}(\Theta_{\text{TR}})] \geq r^* C_{\max} - L_{\mathcal{G}}\sqrt{\chi^2_{2p}(\alpha)\|\boldsymbol{\Sigma}\|_{\text{op}}}.$$

Applying the affine map and taking expectation gives

$$\begin{aligned} \mathbb{E}[F^{\varepsilon}_{\mathcal{G}}(\Theta_{\text{TR}})] &= (1-\nu)\,\mathbb{E}[F_{\mathcal{G}}(\Theta_{\text{TR}})] + \nu\,\bar{C} \\ &\geq (1-\nu)\Big[r^* C_{\max} - L_{\mathcal{G}}\sqrt{\chi^2_{2p}(\alpha)\|\boldsymbol{\Sigma}\|_{\text{op}}}\Big] \\ &\quad + \nu\,\bar{C}. \end{aligned}$$

$$\begin{aligned} &\mathbb{E}[F^{\varepsilon}_{\mathcal{G}}(\Theta_{\text{TR}})] - \mathbb{E}[F^{\varepsilon}_{\mathcal{G}}(\Theta_{\text{rand}})] \\ &\quad = (1-\nu)\Big(\mathbb{E}[F_{\mathcal{G}}(\Theta_{\text{TR}})] - \mathbb{E}[F_{\mathcal{G}}(\Theta_{\text{rand}})]\Big), \end{aligned}$$

since the $\nu\bar{C}$ terms cancel. Because $0 \leq \nu < 1$ for $\varepsilon < 1$, the sign of the expected quality difference is preserved. The ratio in eq. (51) follows from the affine structure $F^{\varepsilon} = (1-\nu)F + \nu\bar{C}$:

$$\begin{aligned} \frac{\mathbb{E}[F^{\varepsilon}_{\mathcal{G}}(\Theta_{\text{TR}})] - \bar{C}}{\mathbb{E}[F^{\varepsilon}_{\mathcal{G}}(\Theta_{\text{rand}})] - \bar{C}} &= \frac{(1-\nu)(\mathbb{E}[F_{\mathcal{G}}(\Theta_{\text{TR}})] - \bar{C})}{(1-\nu)(\mathbb{E}[F_{\mathcal{G}}(\Theta_{\text{rand}})] - \bar{C})} \\ &= \frac{\mathbb{E}[F_{\mathcal{G}}(\Theta_{\text{TR}})] - \bar{C}}{\mathbb{E}[F_{\mathcal{G}}(\Theta_{\text{rand}})] - \bar{C}}. \end{aligned}$$

Whenever the denominator in eq. (51) is nonzero, this relative excess objective is independent of the noise strength.

Finally, the noisy gradient satisfies $\partial_j F^{\varepsilon} = (1-\nu)\partial_j F$, so $\text{Var}[\partial_j F^{\varepsilon}] = (1-\nu)^2\text{Var}[\partial_j F]$. In the barren-plateau regime where $\text{Var}[\partial_j F] = O(2^{-n})$ [12], the noisy variance is doubly suppressed: $\text{Var}[\partial_j F^{\varepsilon}] = (1-\nu)^2 O(2^{-n})$. Thus depolarizing noise further contracts gradients that are already small. □

Appendix H: Conditional gradient anti-concentration proof

Proof. Let $H = -\nabla^2 F_{\mathcal{G}}(\boldsymbol{\mu})$ with $\lambda_{\min}(H) \geq \kappa > 0$, and let $\Lambda = \|H\|_{\text{op}}$. Assume third derivatives are bounded by M_3, i.e., $|\partial_j\partial_k\partial_\ell F_{\mathcal{G}}(\boldsymbol{\theta})| \leq M_3$ for all j, k, ℓ and all $\boldsymbol{\theta} \in \mathcal{T}_\alpha$.

Step 1: Taylor expansion. Writing $\Delta = \boldsymbol{\theta} - \boldsymbol{\mu}$ and expanding $\partial_j F_{\mathcal{G}}$ around $\boldsymbol{\mu}$:

$$\partial_j F_{\mathcal{G}}(\boldsymbol{\theta}) = \underbrace{\partial_j F_{\mathcal{G}}(\boldsymbol{\mu})}_{=0 \text{ if } \boldsymbol{\mu} \text{ is critical}} - \boldsymbol{e}_j^\top H\Delta + R_j(\boldsymbol{\theta}),$$

where the linear term uses $H = -\nabla^2 F_{\mathcal{G}}(\boldsymbol{\mu})$, and the remainder satisfies $|R_j(\boldsymbol{\theta})| \leq \frac{1}{2}M_3(2p)\|\Delta\|_2^2 =: c_p M_3\|\Delta\|_2^2$ by the integral form of the Taylor remainder applied to each of the $2p$ cross-derivatives.

Denoting the linear part $g_j := -\boldsymbol{e}_j^\top H\Delta$, we have $\partial_j F_{\mathcal{G}}(\boldsymbol{\theta}) = g_j + R_j$.

Step 2: Covariance of the displacement on the ellipsoid. For $\boldsymbol{\theta}$ drawn uniformly on the ellipsoid $\mathcal{T}_\alpha = \{\boldsymbol{\theta} : \Delta^\top\boldsymbol{\Sigma}^{-1}\Delta \leq \chi^2_{2p}(\alpha)\}$, the substitution $\boldsymbol{u} = \boldsymbol{\Sigma}^{-1/2}\Delta$ maps $\mathcal{T}_\alpha$ to the ball $\{\boldsymbol{u} : \|\boldsymbol{u}\|_2^2 \leq \chi^2_{2p}(\alpha)\}$. For a uniform distribution on a ball of radius r in $\mathbb{R}^d$, $\text{Cov}[\boldsymbol{u}] = \frac{r^2}{d+2}I_d$. Transforming back:

$$\text{Cov}[\Delta] = \boldsymbol{\Sigma}^{1/2}\,\text{Cov}[\boldsymbol{u}]\,\boldsymbol{\Sigma}^{1/2} = \frac{\chi^2_{2p}(\alpha)}{2p+2}\,\boldsymbol{\Sigma}.$$

Step 3: Variance of the linear term.

$$\text{Var}[g_j] = \boldsymbol{e}_j^\top H\,\text{Cov}[\Delta]\,H\,\boldsymbol{e}_j = \frac{\chi^2_{2p}(\alpha)}{2p+2}\,\boldsymbol{e}_j^\top H\boldsymbol{\Sigma}H\boldsymbol{e}_j.$$

The matrix $H\boldsymbol{\Sigma}H$ is positive semidefinite, and its diagonal entries satisfy

$$(H\boldsymbol{\Sigma}H)_{jj} = \sum_{k=1}^{2p}\sigma_k^2(H_{jk})^2 \geq \sigma^2_{\min}\sum_k(H_{jk})^2 = \sigma^2_{\min}\|H\boldsymbol{e}_j\|_2^2 \geq \sigma^2_{\min}\kappa^2$$

where the last inequality uses $\|H\boldsymbol{e}_j\|_2 \geq \lambda_{\min}(H) \geq \kappa$. Therefore

$$\text{Var}[g_j] \geq \frac{\kappa^2\sigma^2_{\min}\,\chi^2_{2p}(\alpha)}{2p+2}. \tag{H1}$$

Step 4: Remainder correction. Since $\partial_j F = g_j + R_j$, the variance identity gives

$$\text{Var}[\partial_j F] = \text{Var}[g_j] + 2\,\text{Cov}[g_j, R_j] + \text{Var}[R_j].$$

The covariance term is bounded using Cauchy–Schwarz: $|\text{Cov}[g_j, R_j]| \leq \sqrt{\text{Var}[g_j]}\sqrt{\text{Var}[R_j]}$. For $\boldsymbol{\theta} \in \mathcal{T}_\alpha$,

$\|\Delta\|_2 \le \sqrt{\chi^2_{2p}(\alpha)\,\sigma^2_{\max}}$, so $|R_j| \le c_p M_3 \chi^2_{2p}(\alpha)\sigma^2_{\max}$, giving $\mathrm{Var}[R_j] \le c_p^2 M_3^2 \chi^2_{2p}(\alpha)^2 \sigma^4_{\max}$. The cross term is bounded by $\Lambda M_3 C'_{p,\alpha}\sigma^3_{\max}$, where $C'_{p,\alpha}$ absorbs the ellipsoid geometry constants.

Combining eq. (H1) with the remainder:

$$\mathrm{Var}[\partial_j F_{\mathcal{G}}(\boldsymbol{\theta})] \ge \frac{\kappa^2\sigma^2_{\min}\,\chi^2_{2p}(\alpha)}{2p+2} - C_{p,\alpha}\left(\Lambda M_3\sigma^3_{\max} + M_3^2\sigma^4_{\max}\right),$$

which is the bound stated in Theorem 6. The leading term dominates for sufficiently small $\sigma_{\max}$ relative to κ/M_3. □

Appendix I: Hyperparameter sensitivity

Table XII tests stability against modest hyperparameter changes.

TABLE XII. Hyperparameter sensitivity at n=14 (48 instances). Default: λ_W=0.1, λ_C=0.05, τ_W=0.5. **Bold**: best per metric.

Setting	ρ	Evals	Med. ms
Default	**0.770**	**45**	**69**
$\lambda_W = 0.01$	0.742	47	76
$\lambda_W = 0.5$	0.758	46	73
$\lambda_C = 0.0$	0.751	48	77
$\lambda_C = 0.2$	0.763	46	72
$\tau_W = 0.2$	0.761	46	73
$\tau_W = 1.0$	0.749	47	75

The default setting gives the best combination. Variation is ±0.028 in ρ and ±3 evaluations. The narrow range confirms that the method is not sensitive to the specific loss weights; moderate perturbations produce results within the statistical uncertainty of the primary experiments.

Appendix J: Statistical significance

Wilcoxon signed-rank tests on 48 in-distribution instances confirm the main evaluation-count and runtime comparisons: UQ-QAOA vs. random in evaluations ($p < 0.001$, effect size $r = 0.86$), UQ-QAOA vs. random in runtime ($p < 0.001$, $r = 0.84$), and UQ-QAOA vs. GNN point in evaluations (p=0.003, $r = 0.42$). For calibration, Spearman ρ=0.770 with $p < 10^{-47}$ across 240 instances.

REPRODUCIBILITY STATEMENT

- **Code and inputs:** The manuscript and appendix provide the numerical inputs used for the reported bound instantiations.
- **Hyperparameters:** All training and optimization hyperparameters are listed in Appendix A.
- **Random seeds:** Every reported result is averaged over five independent training seeds; per-seed statistics are provided in table IX.
- **Data:** All experiments use synthetic graph instances generated with documented parameters (family, size, and edge-probability ranges specified in section VI).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

[1] E. Farhi, J. Goldstone, and S. Gutmann, arXiv preprint arXiv:1411.4028 10.48550/arxiv.1411.4028 (2014).
[2] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, Physical Review X **10**, 021067 (2020).
[3] J. Wurtz and P. Love, Physical Review A **104**, 052419 (2021).
[4] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, Nature Reviews Physics **3**, 625 (2021).
[5] J. Preskill, Quantum **2**, 79 (2018).
[6] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, in *International Conference on Learning Representations* (2019).
[7] V. P. Dwivedi, C. K. Joshi, A. T. Luo, T. Laurent, Y. Bengio, and X. Bresson, Journal of Machine Learning Research **24**, 1 (2023).
[8] D. Lim, J. Robinson, L. Zhao, T. Smidt, S. Sra, H. Maron, and S. Jegelka, in *International Conference on Learning Representations* (2023).
[9] F. G. S. L. Brandão, M. Broughton, E. Farhi, S. Gutmann, and H. Napp, arXiv preprint arXiv:1812.04170 10.48550/arxiv.1812.04170 (2018).
[10] V. Akshay, D. Rabinovich, E. Campos, and J. Biamonte, Physical Review A **104**, L010401 (2021).
[11] A. Galda, X. Liu, D. Lykov, Y. Alexeev, and I. Safro, in *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, 2021) pp. 171–180.
[12] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Nature Communications **9**, 4812 (2018).
[13] A. Arrasmith, M. Cerezo, P. Czarnik, L. Cincio, and P. J. Coles, Quantum **5**, 558 (2021).
[14] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles, Nature Communications **12**, 6961 (2021).

[15] D. Kreuzer, D. Beaini, W. L. Hamilton, V. Létourneau, and P. Tossou, in *Advances in Neural Information Processing Systems*, Vol. 34 (2021) pp. 21618–21629.

[16] B. Weisfeiler and A. Leman, NTI, Series 2 **9**, 12 (1968).

[17] W. L. Hamilton, R. Ying, and J. Leskovec, in *Advances in Neural Information Processing Systems*, Vol. 30 (2017).

[18] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, in *International Conference on Machine Learning* (PMLR, 2017) pp. 1263–1272.

[19] D. J. Egger, J. Mareček, and S. Woerner, Quantum **5**, 479 (2021).

[20] S. H. Sack and M. Serbyn, Quantum **5**, 491 (2021).

[21] N. Jain, B. Coyle, E. Kashefi, and N. Kumar, Quantum **6**, 861 (2022).

[22] S. Khairy, R. Shaydulin, L. Cincio, Y. Alexeev, and P. Balaprakash, in *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34 (2020) pp. 2367–2375.

[23] S. Tibaldi, D. Vodola, E. Tignone, and E. Ercolessi, IEEE Transactions on Quantum Engineering **4**, 1 (2023).

[24] R. Shaydulin and S. M. Wild, arXiv preprint arXiv:2306.16979 10.48550/arxiv.2306.16979 (2023).

[25] M. Streif and M. Leib, Quantum Science and Technology **5**, 034008 (2020).

[26] C. Moussa, H. Calandra, and V. Dunjko, EPJ Quantum Technology **9**, 11 (2022).

[27] F. Sauvage, S. Sim, A. A. Kunitsa, W. A. Simon, M. Mauri, and A. Perdomo-Ortiz, Quantum Science and Technology **9**, 015029 (2024).

[28] Y. Bengio, A. Lodi, and A. Prouvost, European Journal of Operational Research **290**, 405 (2021).

[29] Q. Cappart, D. Chételat, E. B. Khalil, A. Lodi, C. Morris, and P. Veličković, Journal of Machine Learning Research **24**, 1 (2023).

[30] E. B. Khalil, H. Dai, Y. Zhang, B. Dilkina, and L. Song, in *Advances in Neural Information Processing Systems*, Vol. 30 (2017).

[31] H. Dai, B. Dai, and L. Song, in *International Conference on Machine Learning* (PMLR, 2016) pp. 723–732.

[32] W. Kool, H. van Hoof, and M. Welling, in *International Conference on Learning Representations* (2019).

[33] M. J. A. Schuetz, J. K. Brubaker, and H. G. Katzgraber, Nature Machine Intelligence **4**, 367 (2022).

[34] D. A. Nix and A. S. Weigend, in *Proceedings of 1994 IEEE International Conference on Neural Networks*, Vol. 1 (IEEE, 1994) pp. 55–60.

[35] A. Kendall and Y. Gal, in *Advances in Neural Information Processing Systems*, Vol. 30 (2017).

[36] B. Lakshminarayanan, A. Pritzel, and C. Blundell, in *Advances in Neural Information Processing Systems*, Vol. 30 (2017).

[37] Y. Gal and Z. Ghahramani, in *International Conference on Machine Learning* (PMLR, 2016) pp. 1050–1059.

[38] V. Kuleshov, N. Fenner, and S. Ermon, in *International Conference on Machine Learning* (PMLR, 2018) pp. 2796–2804.

[39] M. Abdar, F. Pourpanah, S. Hussain, D. Rezazadegan, L. Liu, M. Ghavamzadeh, P. Fieguth, X. Cao, A. Khosravi, U. R. Acharya, V. Makarenkov, and S. Nahavandi, Information Fusion **76**, 243 (2021).

[40] E. Hüllermeier and W. Waegeman, Machine Learning **110**, 457 (2021).

[41] A. F. Psaros, X. Meng, Z. Zou, L. Guo, and G. E. Karniadakis, Journal of Computational Physics **477**, 111902 (2023).

[42] D. C. Dowson and B. V. Landau, Journal of Multivariate Analysis **12**, 450 (1982).

[43] M. Arjovsky, S. Chintala, and L. Bottou, in *International Conference on Machine Learning* (PMLR, 2017) pp. 214–223.

[44] C. Frogner, C. Zhang, H. Mobahi, M. Araya, and T. Poggio, in *Advances in Neural Information Processing Systems*, Vol. 28 (2015).

[45] C. Villani, *Optimal Transport: Old and New*, Vol. 338 (Springer, 2009).

[46] M. Cuturi, in *Advances in Neural Information Processing Systems*, Vol. 26 (2013).

[47] G. Peyré and M. Cuturi, Foundations and Trends in Machine Learning **11**, 355 (2019).

[48] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, in *International Conference on Machine Learning* (PMLR, 2020) pp. 1597–1607.

[49] A. van den Oord, Y. Li, and O. Vinyals, arXiv preprint arXiv:1807.03748 10.48550/arxiv.1807.03748 (2018).

[50] Y. You, T. Chen, Y. Sui, T. Chen, Z. Wang, and Y. Shen, in *Advances in Neural Information Processing Systems*, Vol. 33 (2020) pp. 5812–5823.

[51] S. Suresh, P. Li, C. Hao, and G. Nishi, in *Advances in Neural Information Processing Systems*, Vol. 34 (2021) pp. 15920–15933.

[52] P. Khosla, P. Teterwak, C. Wang, A. Sarna, Y. Tian, P. Isola, A. Maschinot, C. Liu, and D. Krishnan, in *Advances in Neural Information Processing Systems*, Vol. 33 (2020) pp. 18661–18673.

[53] B. Amos, Foundations and Trends in Machine Learning **16**, 592 (2023).

[54] T. Chen, X. Chen, W. Chen, H. Heaton, J. Liu, Z. Wang, and W. Yin, Journal of Machine Learning Research **23**, 1 (2022).

[55] C. Finn, P. Abbeel, and S. Levine, in *International Conference on Machine Learning* (PMLR, 2017) pp. 1126–1135.

[56] R. Shu, H. H. Xu, T.-H. Oh, and K. Swersky, in *International Conference on Learning Representations* (2022).

[57] J. L. Ba, J. R. Kiros, and G. E. Hinton, arXiv preprint arXiv:1607.06450 10.48550/arxiv.1607.06450 (2016).

[58] J. A. Nelder and R. Mead, The Computer Journal **7**, 308 (1965).

[59] P. L. Bartlett, D. J. Foster, and M. J. Telgarsky, in *Advances in Neural Information Processing Systems*, Vol. 30 (2017).

[60] B. Neyshabur, S. Bhojanapalli, and N. Srebro, in *International Conference on Learning Representations* (2018).

[61] N. Golowich, A. Rakhlin, and O. Shamir, in *Proceedings of the 31st Conference on Learning Theory* (2018) pp. 297–299.

[62] V. Vovk, A. Gammerman, and G. Shafer, *Algorithmic Learning in a Random World* (Springer, 2005).

[63] Y. Romano, E. Patterson, and E. Candès, in *Advances in Neural Information Processing Systems*, Vol. 32 (2019).

[64] A. N. Angelopoulos and S. Bates, arXiv preprint arXiv:2107.07511 10.48550/arxiv.2107.07511 (2021).

[65] P. Erdős and A. Rényi, Publicationes Mathematicae Debrecen **6**, 290 (1959).

[66] A.-L. Barabási and R. Albert, Science **286**, 509 (1999).

[67] D. J. Watts and S. H. Strogatz, Nature **393**, 440 (1998).

[68] A. R. Conn, N. I. M. Gould, and P. L. Toint, *Trust Region Methods* (SIAM, 2000).

[69] D. P. Kingma and J. Ba, in *International Conference on Learning Representations* (2015).