Neural QAOA^2: Differentiable Joint Graph Partitioning and Parameter Initialization for Quantum Combinatorial Optimization

Zubin Zheng*[1], Jiahao Wu*[1], Shengcai Liu[1]

Accepted to ICML 2026. *Equal contribution [1]Guangdong Provincial Key Laboratory of Brain-Inspired Intelligent Computation, Department of CSE, SUSTech. Correspondence to: Shengcai Liu

Outline

- Motivation
- Neural QAOA^2
- Experiments

Motivation

Background

Quantum approximate optimization algorithm (QAOA) holds promise for combinatorial optimization but is constrained by limited qubits

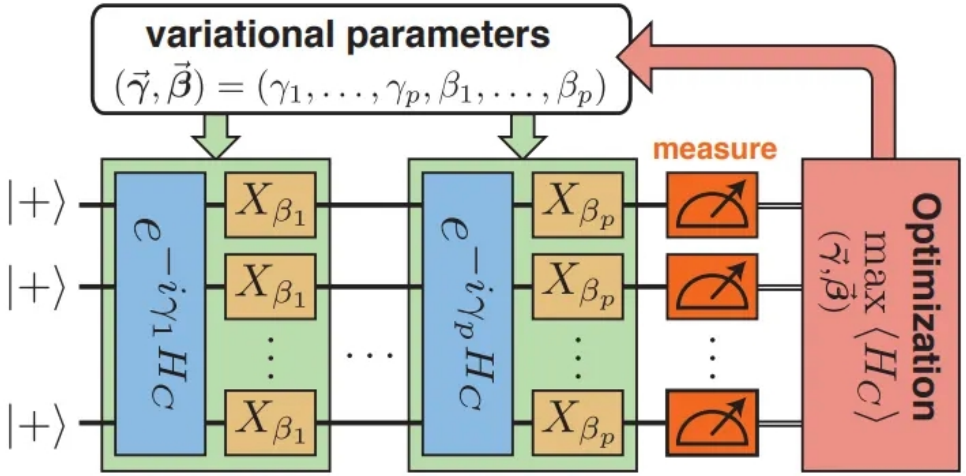

Divide-and-conquer frameworks like QAOA-in-QAOA (QAOA²) address scalability by partitioning graphs into subgraphs

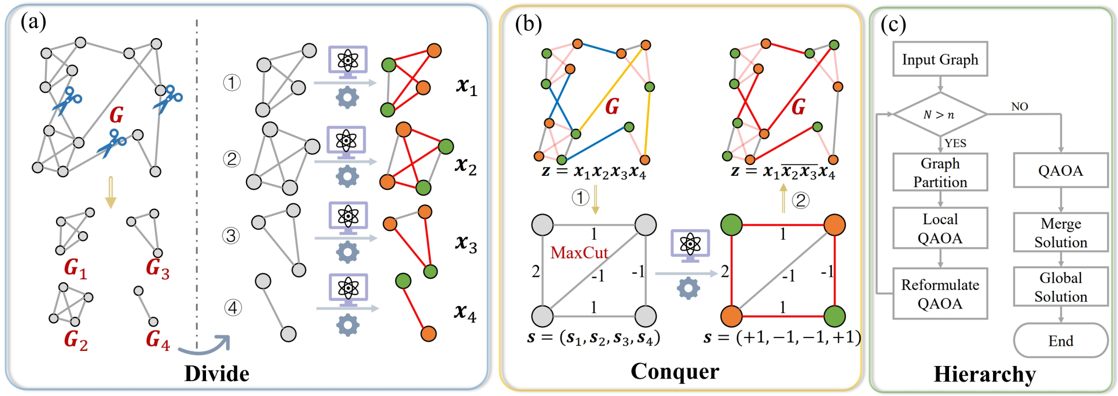

Farhi et al. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028*, 2014.
image sources:
Zhou et al. Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices. *Physical Review X*, 2020.
Zhou et al. QAOA-in-QAOA: solving large-scale MaxCut problems on small quantum machines. *Physical Review Applied*, 2023.

Limitations in Existing Methods

However, QAOA² suffers from two limitations:

- Misalignment between heuristic partitioning metrics and quantum optimization goal
- Topology-blind parameter initialization that leads to optimization cold starts

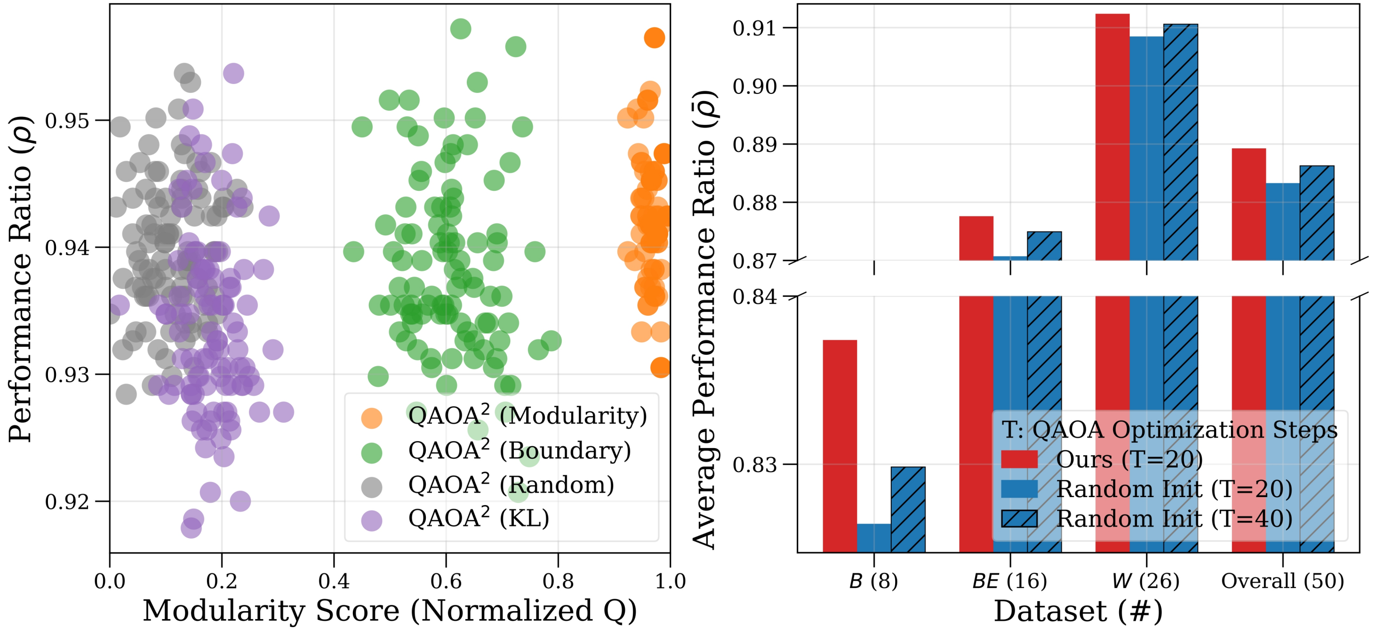

We propose **Neural QAOA²** to address them

Neural QAOA^2

Neural QAOA^2

A unified divide-and-conquer framework that reimagines the paradigm by integrating a generative evaluative network (GEN)

GEN replaces heuristic partitioning and random parameter initialization with a single neural backbone that jointly generates graph partitions and QAOA initial parameters

This end-to-end design aligns partitioning with quantum optimization objectives and enables topology-aware warm starts

Overview of GEN

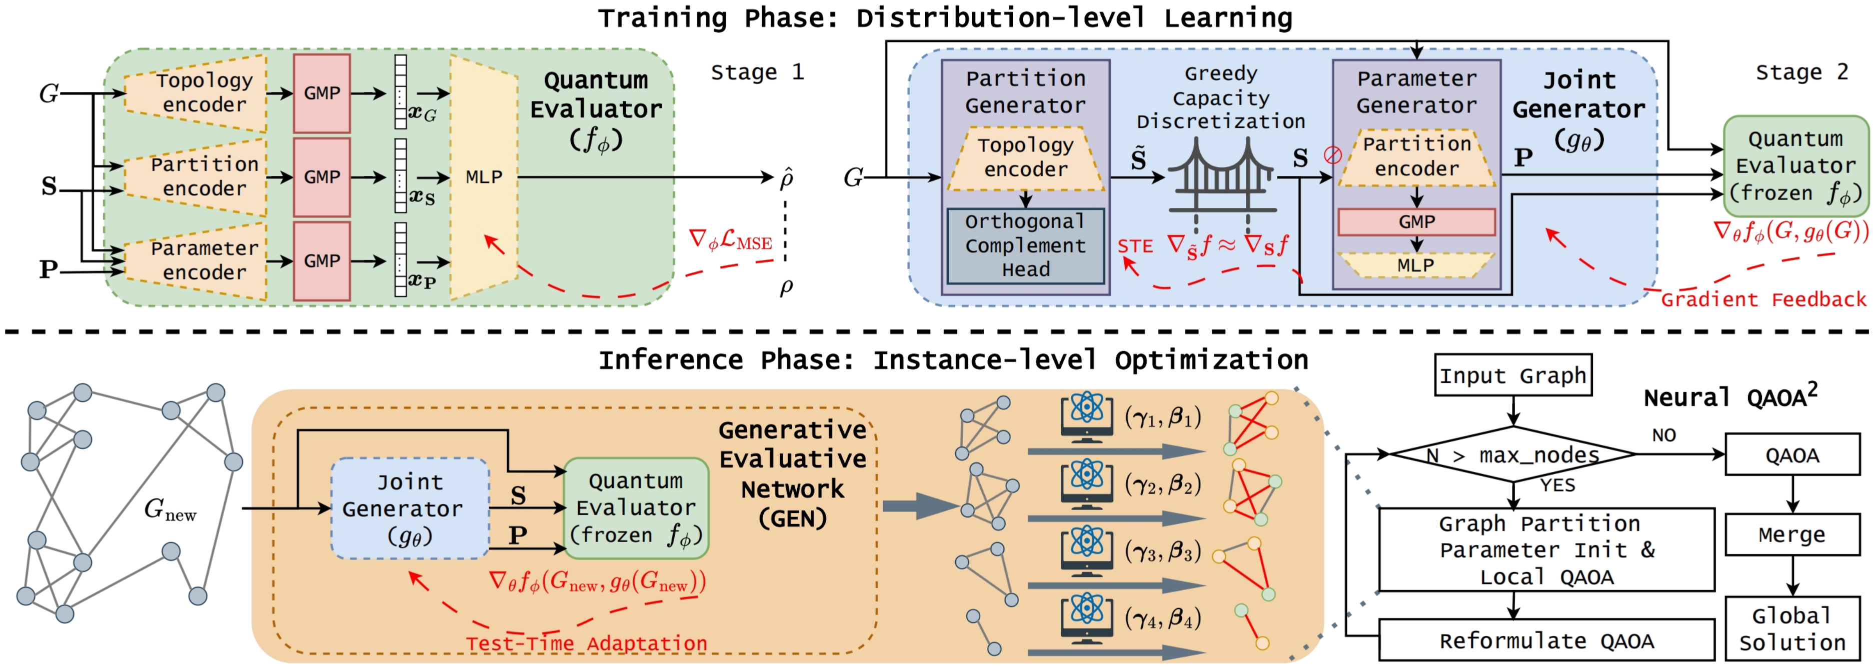

Training Phase: Distribution-level Learning

Stage 1: $\min_{\phi} \mathbb{E}_{(G,\mathbf{S},\mathbf{P},\rho)\sim\mathcal{D}_{\text{offline}}} \left[\mathcal{L}_{\text{MSE}}(f_\phi(G,\mathbf{S},\mathbf{P}),\rho)\right]$

Stage 2: $\max_{\theta} \mathbb{E}_{G\sim\mathcal{D}_{\text{graph}}} \left[f_\phi(G, g_\theta(G))\right]$

Inference Phase: Instance-level Optimization

$$\theta^* = \arg\max_{\theta} f_\phi(G_{\text{new}}, g_\theta(G_{\text{new}}))$$

$$(\mathbf{S}^*, \mathbf{P}^*) = g_{\theta^*}(G_{\text{new}})$$

Quantum Evaluator $f_\phi(G, \mathbf{S}, \mathbf{P})$

a differentiable surrogate to circumvent expensive physical simulations of QAOA² and provide gradient guidance for the upstream generator

encode heterogeneous inputs (discrete topologies/binary partitions/continuous parameters)

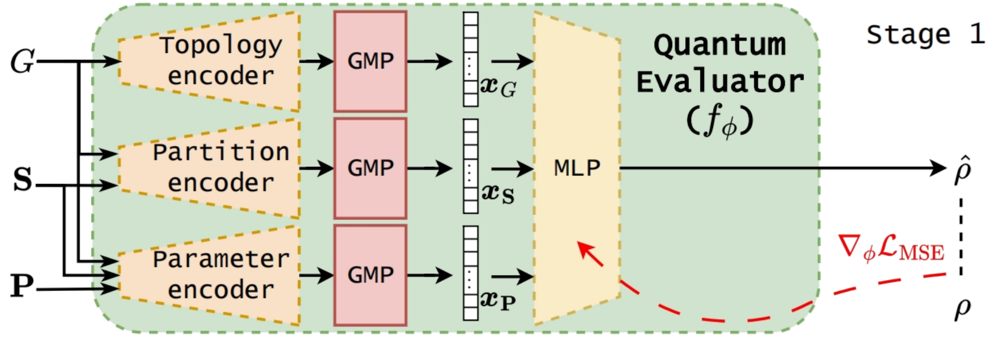

Global Topology View

$$\mathbf{H}_{\text{topology}} = \text{Encoder}_{\text{topology}}(\mathbf{X}, \mathbf{A}) \in \mathbb{R}^{N\times h}$$

Partition-Induced Subgraph View

$$\mathbf{A}_{\text{sub}} = \mathbf{A} \odot (\mathbf{S}\mathbf{S}^\top) \quad \mathbf{H}_{\text{partition}} = \text{Encoder}_{\text{partition}}(\mathbf{X}, \mathbf{A}_{\text{sub}}) \in \mathbb{R}^{N\times h}$$

Quantum Parameter View

$$\mathbf{X}_{\text{param}} = \mathbf{S}\mathbf{P}^\top \in [0,2\pi)^{N\times 2p} \quad \widetilde{\mathbf{X}}_{\text{param}} = \text{concat}[\sin(\mathbf{X}_{\text{param}}),\cos(\mathbf{X}_{\text{param}})] \in [-1,1]^{N\times 4p}$$

$$\mathbf{H}_{\text{param}} = \text{Encoder}_{\text{param}}(\widetilde{\mathbf{X}}_{\text{param}}, \mathbf{A}_{\text{sub}}) \in \mathbb{R}^{N\times h}$$

Joint Generator $g_\theta(G)$

generate discrete partition matrix $\mathbf{S}$ constrained by qubit capacity + variational parameters $\mathbf{P}$

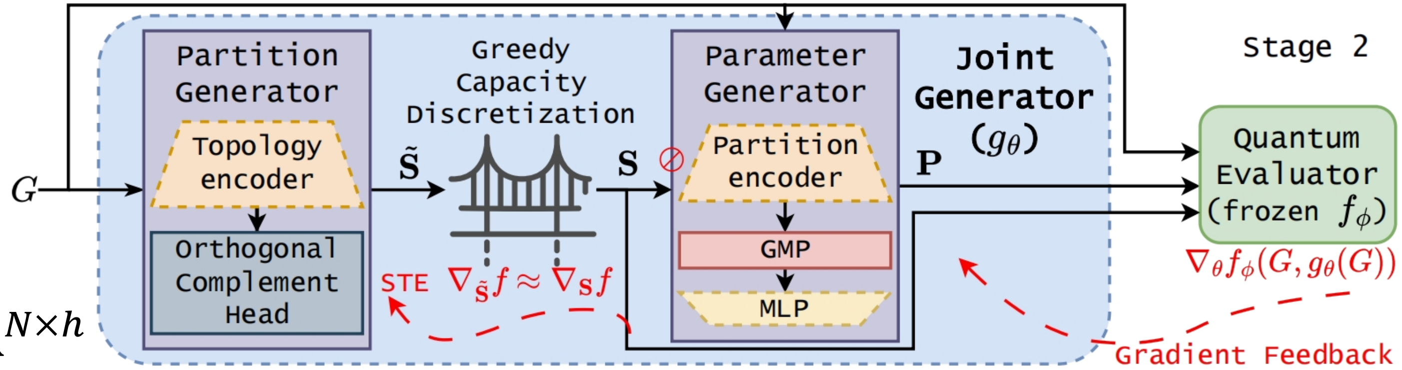

partition-first, parameter-second

Soft Partition via OCH

$$\mathbf{H}_{\text{topology}} = \text{Encoder}_{\text{topology}}(\mathbf{X}, \mathbf{A}) \in \mathbb{R}^{N \times h}$$

$$\tilde{\mathbf{S}} = \text{softmax}(\mathbf{H}_{\text{topology}} \mathbf{C}^\top) \in [0,1]^{N \times k} \text{ s.t. } \mathbf{C}\boldsymbol{g} = \mathbf{0} \text{ and } \mathbf{C}\mathbf{C}^\top = \mathbf{I} \text{ where } \boldsymbol{g} = \text{GMP}(\mathbf{H}_{\text{topology}}) \in \mathbb{R}^h$$

Differentiable Discretization

Forward: $\mathbf{S} = \text{GCD}(\tilde{\mathbf{S}}) \in \{0,1\}^{N \times k}$, $\mathbf{A}_{\text{sub}} = \mathbf{A} \odot (\mathbf{S}\mathbf{S}^\top)$ Backward: $\nabla_{\tilde{\mathbf{S}}} f \approx \nabla_{\mathbf{S}} f$ (Bengio, 2013)

Parameter Generation

$$\mathbf{H}_{\text{partition}} = \text{Encoder}_{\text{partition}}(\mathbf{X}, \text{sg}(\mathbf{A}_{\text{sub}})) \qquad \mathbf{H}_{\text{sub}} = \text{GMP}(\mathbf{H}_{\text{partition}}) \in \mathbb{R}^{h \times k}$$

$$(\mathbf{x}, \mathbf{y}) = \mathcal{F}_{proj}(\text{MLP}(\mathbf{H}_{\text{sub}})), \ \ s.t.\, \mathbf{x}, \mathbf{y} \in \mathbb{R}^{2p \times k} \quad \mathbf{P} = \text{arctan2}(\mathbf{y}, \mathbf{x}) \ (\text{mod } 2\pi) \in [0, 2\pi)^{2p \times k}$$

Experiments

Experimental Setup

Datasets a public benchmark library[1], includes QUBO (datasets *B*, *BE*, *GKA*), Ising spin glass (datasets *L*, *BMZ*), and MaxCut (datasets *W*, *HR*) instances; all exceed maximum qubit limit 10

Training Details 80% of instances from datasets *B*, *BE*, and *W* (problem sizes range from 51 to 501), which consist of QUBO and MaxCut instances only, to form the training graph set $\mathcal{D}_{\text{graph}}$. Ising spin glass instances are excluded from training

Baselines QAOA[2] framework equipped with four partitioning heuristics: Random, Modularity (Clauset et al., 2004), Boundary (Guerreschi, 2021), KL (Kernighan & Lin, 1970)

Metrics average performance ratio $\rho = \frac{\text{Cut}(G) - \text{Neg}(G)}{\text{OPT}(G) - \text{Neg}(G)}$, average rank in 10 independent quantum simulation runs

Code Availability https://github.com/0SliverBullet/Neural-QAOA-Squared

[1]http://bqp.cs.uni-bonn.de/library/html/instances.html

Main Results

- 1.74 Best Overall Average Rank
- 28/50 Win Rate
- Stability

Table 1. **Main Results on 50 Test Instances.** Cell layout: mean $\bar{\rho}$ (top), $\pm$ std (middle), and average rank with w/l counts (bottom) across different datasets (# denotes the instance count). **Bold** indicates the best performance (mean $\bar{\rho}$ $\uparrow$, rank $\downarrow$). Neural QAOA^2 consistently outperforms heuristics, achieving the best average rank of **1.74** overall.

Dataset (#)	QAOA^2 with Heuristic Partition				Neural QAOA^2
	Random	Modularity	Boundary	KL	GEN
B (8)	0.8047 $\pm$ 0.0226 (4.75) 0/8	0.8351 $\pm$ 0.0304 (2.38) 2/6	0.8246 $\pm$ 0.0326 (2.63) 1/7	0.8092 $\pm$ 0.0177 (3.75) 0/8	**0.8417** $\pm$ 0.0256 (**1.50**) 5/3
BE (16)	0.8626 $\pm$ 0.0344 (4.81) 0/16	0.8692 $\pm$ 0.0342 (3.13) 0/16	0.8722 $\pm$ 0.0359 (2.31) 1/15	0.8672 $\pm$ 0.0342 (3.69) 0/16	**0.8824** $\pm$ 0.0304 (**1.06**) 15/1
W (26)	0.8962 $\pm$ 0.0543 (3.23) 2/24	0.9137 $\pm$ 0.0304 (**2.23**) 9/17	0.9114 $\pm$ 0.0335 (2.96) 6/20	0.8934 $\pm$ 0.0537 (4.27) 1/25	**0.9153** $\pm$ 0.0280 (**2.23**) 8/18
Overall (50)	0.8708 $\pm$ 0.0552 (3.98) 2/48	0.8869 $\pm$ 0.0437 (2.54) 11/39	0.8850 $\pm$ 0.0465 (2.70) 8/42	0.8716 $\pm$ 0.0529 (4.00) 1/49	**0.8930** $\pm$ 0.0390 (**1.74**) 28/22

Generalization Analysis

Distribution Generalization

- 1.46 Overall Out-of-Distribution Rank
- 60/93 Win Rate

Problem Size Generalization

- Size Interpolation
- Size Extrapolation

Table 2. **Distribution Generalization on 93 Test Instances.** We evaluate on two out-of-distribution graph topologies: *GKA* (QUBO) and *L* (Ising spin glass). Cell layout: mean $\bar{\rho}$ (top), $\pm$ std (middle), and average rank with w/l counts (bottom). **Bold** indicates the best metric. Neural QAOA^2 achieves the best average rank of **1.46** overall.

Dataset (#)	QAOA^2 with Heuristic Partition				Neural QAOA^2
	Random	Modularity	Boundary	KL	GEN
GKA (45)	0.8478 ± 0.0441 (4.16) 2/43	0.8659 ± 0.0364 (2.40) 7/38	0.8601 ± 0.0414 (2.89) 4/41	0.8503 ± 0.0393 (4.04) 0/45	**0.8762** ± 0.0359 (**1.51**) 32/13
L (48)	0.6984 ± 0.0305 (4.65) 0/48	0.7391 ± 0.0446 (3.06) 0/48	**0.8205** ± 0.0440 (1.60) 20/28	0.7022 ± 0.0235 (4.27) 0/48	0.8160 ± 0.0285 (**1.42**) 28/20
Overall (93)	0.7707 ± 0.0837 (4.41) 2/91	0.8005 ± 0.0754 (2.74) 7/86	0.8397 ± 0.0471 (2.23) 24/69	0.7739 ± 0.0807 (4.16) 0/93	**0.8451** ± 0.0441 (**1.46**) 60/33

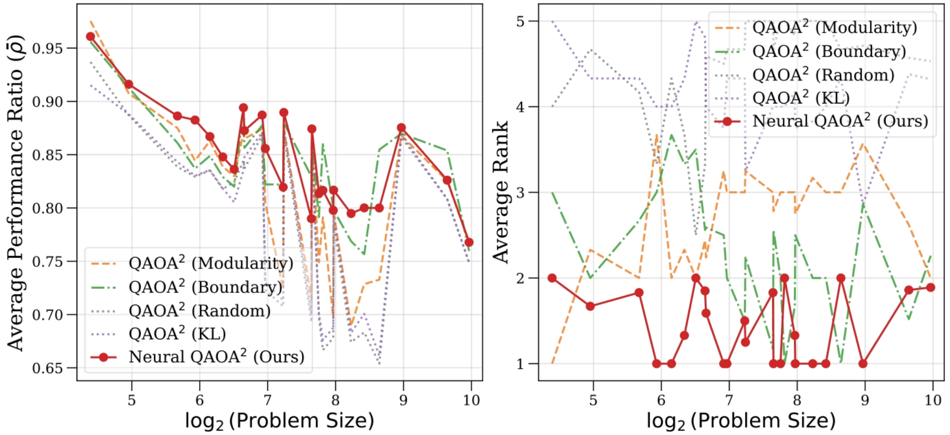

Figure 3. **Problem Size Generalization.** We evaluate Neural QAOA^2 on test instances ranging from $N = 21$ to $N = 1000$ (plotted on a $\log_2$ scale). The left panel shows the mean $\bar{\rho}$, and the right panel shows the average rank. Our method (red line) maintains superior ranking stability across the spectrum, effectively generalizing to instances with scales unseen during training.

Ablation Studies

Cost-Performance Trade-off

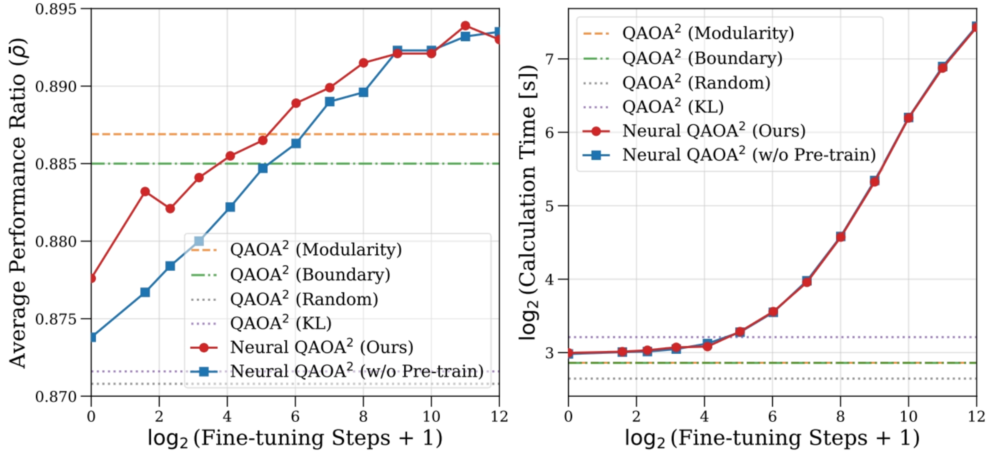

Figure 4. **Dynamics of Test-Time Adaptation.** Performance (left) and wall-clock time (right) vs. fine-tuning steps (plotted on a $\log_2$ scale). The red curve shows monotonic improvement, outperforming the strongest heuristic (modularity) after 64 steps. Comparison with the non-pretrained variant (blue) confirms that pre-training provides a critical initialization for rapid convergence.

Component Contribution Analysis

Table 3. **Ablation Study of Neural QAOA^2 Components.** $\Delta\%$ denotes the relative performance drop compared to Neural QAOA^2. Statistical significance is measured via a one-sided paired t-test. Differences are considered statistically significant at $p < 0.01$.

Variant (Fine-tuning Steps = 64)	Mean $\bar{\rho}$	$\Delta\%$	p-value
Neural QAOA^2	0.8892	–	–
Ablation on Quantum Evaluator Inputs			
w/o Global Topology View	0.8819	-0.83%	6.93×10^{-4}
w/o Partition-Induced Subgraph View	0.8797	-1.07%	7.44×10^{-7}
w/o Quantum Parameter View	0.8821	-0.80%	2.40×10^{-4}
Ablation on Joint Generator Mechanisms			
w/o Orthogonal Complement Head	0.8817	-0.85%	7.56×10^{-5}
w/o Stop-Gradient	0.8820	-0.81%	1.22×10^{-4}

A Deeper Look Into Why It Works

Resolving Metric Misalignment

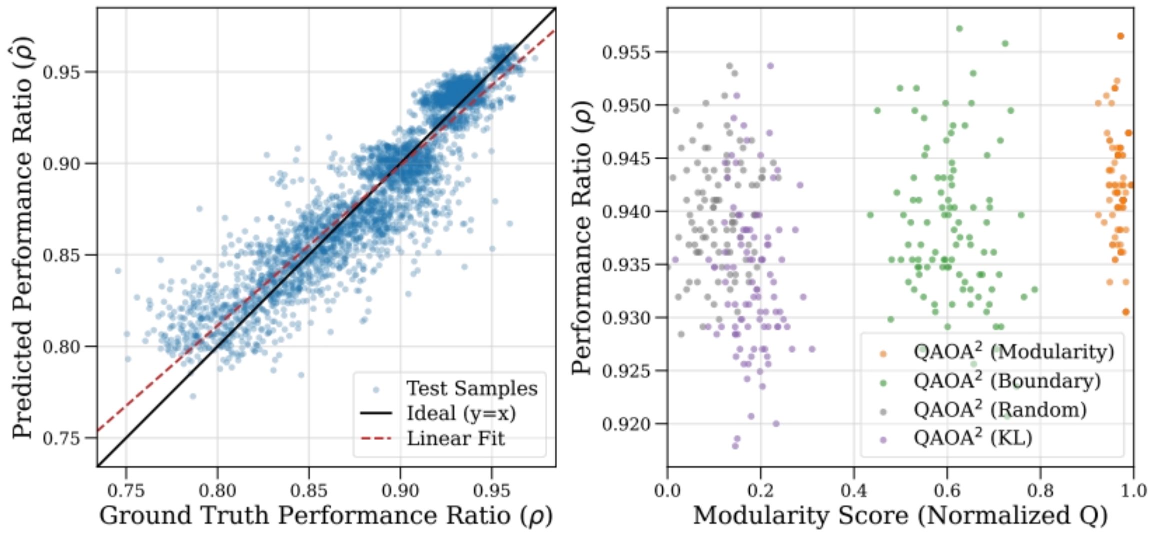

Figure 5. **Visual Analysis of Metric Misalignment.** Left: quantum evaluator predictions vs. ground truth across 3313 unseen test data points ($r = 0.9258$), showing high fidelity. Right: modularity vs. ground truth on the test instance *g05_100.1* ($r = 0.2859$), exposing a severe metric misalignment that lacks correlation with quantum performance.

Mitigating Cold Start

Table 5. **Impact of Joint Parameter Initialization.** Comparison against a variant using random parameter initialization. Neural QAOA^2 (20 steps) outperforms random initialization even when the latter is given double the optimization budget (40 steps), confirming the effectiveness of our topology-aware warm start. Statistical significance is measured via a one-sided paired t-test. Differences are considered statistically significant at $p < 0.01$.

Variant (Fine-tuning Steps = 64)	Mean $\bar{\rho}$	$\Delta\%$	Time [s]	p-value
QAOA Optimization Steps = 20				
Neural QAOA^2	0.8892	–	11.78	–
Neural QAOA^2 (Random **P**)	0.8833	-0.67%	11.75	1.89×10^{-5}
QAOA Optimization Steps = 40				
Neural QAOA^2 (Random **P**)	0.8863	-0.34%	17.49	8.22×10^{-3}