

Adaptive Fidelity Estimation for Quantum Programs with Graph-Guided Noise Awareness

Tingting Li^{1,2}, Ziming Zhao^{2,*}, Jianwei Yin^{2,*}

¹Shanghai Qi Zhi Institute, Shanghai, China

²Zhejiang University, Hangzhou, China

{litt2020, zhaoziming}@zju.edu.cn, zjuyjw@cs.zju.edu.cn

Abstract

Fidelity estimation is a critical yet resource-intensive step in testing quantum programs on noisy intermediate-scale quantum (NISQ) devices, where the required number of measurements is difficult to predefine due to hardware noise, device heterogeneity, and transpilation-induced circuit transformations. We present QuFid, an adaptive and noise-aware framework that determines measurement budgets online by leveraging circuit structure and runtime statistical feedback. QuFid models a quantum program as a directed acyclic graph (DAG) and employs a control-flow-aware random walk to characterize noise propagation along gate dependencies. Backend-specific effects are captured via transpilation-induced structural deformation metrics, which are integrated into the random-walk formulation to induce a noise-propagation operator. Circuit complexity is then quantified through the spectral characteristics of this operator, providing a principled and lightweight basis for adaptive measurement planning. Experiments on 18 quantum benchmarks executed on IBM Quantum backends show that QuFid significantly reduces measurement cost compared to fixed-shot and learning-based baselines, while consistently maintaining acceptable fidelity bias.

Introduction

Quantum computing exploits quantum-mechanical phenomena such as superposition and entanglement to enable advances in optimization (Urbanek, Nachman, and de Jong 2020; Ayanzadeh et al. 2023), simulation (Zhu et al. 2024; Xiang, Chen et al. 2024), and material science (Nielsen and Chuang 2001). Promising applications have already emerged in quantum simulation (Preskill 2018), quantum machine learning (Li, Zhao, and Yin 2024; Biamonte et al. 2017), and cryptography (Shor 1999; Xu et al. 2024a).

However, in the current Noisy Intermediate-Scale Quantum (NISQ) era, quantum devices remain severely constrained by decoherence, gate errors, limited qubit counts, and the absence of full-scale error correction (Preskill 2018; Hu et al. 2023; Liang et al. 2023; Campbell 2024). These noise sources degrade computational accuracy and make reliable state evaluation inherently challenging, hindering practical deployment and necessitating careful device characterization prior to execution. Benchmarking techniques

such as randomized benchmarking (RB) (Magesan, Gambetta, and Emerson 2012) and cross-entropy benchmarking (XEB) (Boixo et al. 2018) assess device performance through repeated circuit executions and statistical analysis (Sun et al. 2014). In practice, however, the number of repetitions (shots) is typically chosen heuristically: insufficient shots lead to inaccurate fidelity estimates, while excessive shots waste scarce quantum resources.

Recent efforts have explored fidelity prediction and reliability estimation using pre-characterized noise models or learning-based approaches (Tan et al. 2023; Wang et al. 2022b). While effective in controlled settings, these methods rely on prior knowledge of device-specific noise characteristics, which are often time-varying and difficult to obtain in real-world environments (Li et al. 2025c). As a result, there remains a pressing need for adaptive fidelity estimation strategies that can dynamically determine measurement budgets while accounting for both circuit structure and hardware-induced effects.

Through our investigation, we identify three fundamental challenges in determining the number of repetitions required for quantum benchmarking. (i) *Noise heterogeneity*. Quantum noise, including both Markovian and non-Markovian components (Breuer and Petruccione 2002), varies across devices and over time (Kandala et al. 2019), making static or universal repetition rules unreliable. (ii) *Circuit-hardware discrepancy*. During transpilation, qubit mapping, circuit rewriting, and optimization (McKay et al. 2018) reshape the circuit’s dependency structure and noise exposure (Li et al. 2025a). Existing approaches lack a unified way to quantify the resulting structural deformation and its impact on measurement reliability. (iii) *Precision-latency trade-off*. Fidelity evaluation must balance statistical accuracy (*e.g.*, for validation and verification) against rapid turnaround (*e.g.*, for calibration and testing), calling for adaptive and uncertainty-aware measurement strategies.

To address these challenges, we propose QuFid, an adaptive and noise-aware fidelity measurement framework. QuFid models a quantum program as a directed acyclic graph (DAG) and constructs a control-flow-aware random walk to characterize noise propagation along gate dependencies. Backend-specific effects introduced by transpilation are captured through *structural deformation metrics*, which quantify changes in effective connectivity and de-

*Corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

pendency structure after mapping. These metrics are integrated into the random-walk formulation, inducing a noise-propagation operator whose spectral characteristics provide a principled and lightweight measure of circuit complexity. Based on this structural characterization and runtime statistical feedback, QuFid adaptively allocates measurement budgets and applies confidence-driven early stopping to avoid redundant sampling while preserving estimation accuracy.

In summary, this paper makes the following contributions.

- We identify fundamental limitations of existing fidelity estimation strategies and highlight the need for adaptive, structure-aware measurement planning under realistic noise conditions.
- We propose QuFid, a novel framework that integrates control-flow-aware noise propagation, transpilation-induced structural deformation modeling, and spectral complexity analysis to guide adaptive fidelity measurement for quantum circuits.
- We evaluate QuFid on 18 diverse quantum benchmarks executed on real IBM Quantum backends, demonstrating a significant reduction in measurement cost while maintaining high accuracy. Additional studies and extended analysis further validate its scalability and adaptability.

Background and Related Work

In this section, we review relevant background on quantum computing and fidelity measurement, survey existing approaches to measurement optimization, and motivate the need for adaptive fidelity estimation strategies.

Quantum Computing in the NISQ Era. Classical computers process information using deterministic bits and logic gates, whereas quantum computers operate on qubits that can exist in superposition, $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ with $|\alpha|^2 + |\beta|^2 = 1$, and exploit superposition, entanglement, and interference to explore multiple computational paths simultaneously (Huang et al. 2021). In the Noisy Intermediate-Scale Quantum (NISQ) era (Li et al. 2025b), quantum devices are constrained by limited qubit counts, high gate error rates, decoherence, and the absence of fault-tolerant error correction (Preskill 2018). As a result, quantum computations are inherently noisy and probabilistic, requiring repeated measurements and statistical analysis to obtain reliable results (Preskill 2018), making accurate state and fidelity evaluation both essential and resource-intensive.

Fidelity Measurement and Optimization. Fidelity measurement is a fundamental technique for validating and characterizing quantum devices and programs (Gilyén and Poremba 2022; Wang et al. 2022c; Huang, Kueng, and Preskill 2020). Due to the high cost of repeated measurements, prior work has explored fidelity prediction and measurement optimization using machine learning (Zhang et al. 2021; Yu, Shang, and Gühne 2022; Liu and Zhou 2020), variational algorithms (Cerezo et al. 2020; Chen et al. 2021; Tan and Volkoff 2021), and classical-shadow-based techniques (Huang, Kueng, and Preskill 2020). For example, QuEst (Wang et al. 2022a) applies graph transformers to predict circuit fidelity, while QuCT (Tan et al. 2023) employs a unitary-to-vector representation. While effective in

certain scenarios, these approaches typically rely on pre-characterized noise models or historical training data and focus on predicting fidelity values rather than determining the minimal number of measurements required to achieve a desired accuracy, limiting their applicability under dynamic and heterogeneous hardware conditions.

Random Walk Models and Structural Analysis. Random walks model stochastic traversal over graph structures based on transition probabilities (Tong, Faloutsos, and Pan 2006) and are widely used to capture both local and global structural properties. In quantum computing (Li et al. 2024), random-walk-based techniques have been applied to analyze state evolution, circuit structure, and dependency patterns (Qiang et al. 2016), as well as to support noise-aware synthesis and optimization (Bergholm et al. 2018). More broadly, random walks are extensively studied in classical graph domains for information diffusion, graph property estimation, and Monte Carlo optimization (Tong, Faloutsos, and Pan 2006; Li et al. 2015; Craswell and Szummer 2007). Most existing applications use random walks as heuristic tools for structural characterization or sampling, whereas our work leverages them to model noise propagation along control-flow dependencies, enabling an operator-level abstraction that connects circuit structure, hardware-induced deformation, and measurement uncertainty.

Motivation and Positioning. Determining an appropriate number of measurements remains a central challenge in quantum benchmarking and testing. This challenge arises from heterogeneous and time-varying noise (Breuer and Petruccione 2002; Kandala et al. 2019), discrepancies between logical circuits and their transpiled implementations (McKay et al. 2018), and the trade-off between statistical accuracy and measurement latency. In practice, shot counts are often chosen heuristically, leading to unreliable estimates or unnecessary resource consumption (Wang et al. 2022a). Existing approaches (Tan et al. 2023; Wang et al. 2022b) largely assume static noise characteristics or offline training, limiting their effectiveness in realistic NISQ environments. Motivated by these limitations, we propose an adaptive fidelity measurement framework that integrates structural circuit analysis with runtime statistical feedback. By modeling transpilation-induced structural deformation and noise propagation within a unified graph-based abstraction, our approach enables principled, backend-aware measurement budgeting that adapts to both circuit structure and observed convergence behavior.

Methodology

We model fidelity estimation as a structure-aware analysis problem that connects circuit dependencies, noise propagation, and statistical convergence. Rather than using random walks as a direct complexity estimator (Tong, Faloutsos, and Pan 2006), we employ a control-flow-aware random walk as an abstraction for how noise diffuses through gate dependencies in a quantum circuit (Masuda, Porter, and Lambiotte 2017; Fontana et al. 2025). This formulation induces a noise-propagation operator whose structural properties reflect the circuit’s intrinsic complexity after transpila-

Algorithm 1: Adaptive Fidelity Measurement Pipeline

Input: Quantum circuit Q , target error bound δ , confidence level α

Output: Estimated fidelity \hat{F}

- 1: $\mathcal{G} \leftarrow$ Convert Q into a DAG
 - 2: Construct graph-structural deformation metrics on \mathcal{G}
 - 3: Build control-flow-aware noise-propagation operator P
 - 4: Compute spectral complexity $C(\mathcal{G})$ from P
 - 5: $\mathcal{D} \leftarrow$ Depth of transpiled circuit
 - 6: $\mathcal{P} \leftarrow C(\mathcal{G}) \cdot \log(\mathcal{D})$ ▷ Adaptive batch size
 - 7: Initialize measurement set $T \leftarrow \emptyset$
 - 8: **repeat**
 - 9: Execute \mathcal{P} measurements $\rightarrow \{T_i\}$
 - 10: $T \leftarrow T \cup \{T_i\}$
 - 11: $\hat{F} \leftarrow \text{mean}(T)$
 - 12: $\sigma \leftarrow \text{std}(T)$
 - 13: $CI \leftarrow z_\alpha \cdot \frac{\sigma}{\sqrt{|T|}}$
 - 14: **if** $CI \leq \delta$ **then**
 - 15: **return** \hat{F} ▷ Confidence-driven early stopping
 - 16: **end if**
 - 17: **until** $|T| \geq \mathcal{P}_{\max}$
 - 18: **return** \hat{F}
-

tion. Circuit complexity is then quantified through the spectral characteristics of this operator, providing a principled and lightweight basis for determining adaptive measurement iterations (Coifman and Lafon 2006). The overall adaptive measurement pipeline is summarized in Algorithm 1.

Quantum Circuit Graph Construction. We represent a quantum program as a directed acyclic graph (DAG) $\mathcal{G} = (V, E)$, where each node $v_i \in V$ corresponds to a quantum gate and each directed edge $(v_i, v_j) \in E$ encodes a non-commutative control-flow dependency induced by shared qubits. A directed edge indicates that gate v_j cannot be executed before v_i without altering circuit semantics, thereby explicitly capturing execution-order constraints and inter-gate interactions. This DAG representation provides a unified abstraction for both logical circuit structure and dependency-induced error propagation. As shown in Figure 1, we first convert the Bernstein-Vazirani circuit into a DAG (Aleksandrowicz et al. 2019) and then into a MultiDi-Graph to capture topological connections. Each gate corresponds to a node, and qubits define edges. Nodes store gate type, target qubits, and parameters, while edges encode non-commutative relationships (Smith et al. 2023).

Structural Deformation Metrics. The transpilation transforms a logical circuit into a hardware-executable form by reordering gates, inserting auxiliary operations, and remapping qubits. These transformations expose and reshape control-flow dependencies, often introducing additional interactions that affect noise accumulation and propagation. Let $\mathcal{G}_0 = (V_0, E_0)$ and $\mathcal{G}_t = (V_t, E_t)$ denote the circuit graphs before and after transpilation, respectively. Transpilation induces *structural deformation* by introducing auxiliary gates, modifying dependency paths, and altering connectivity patterns. We quantify such deformation using graph-level

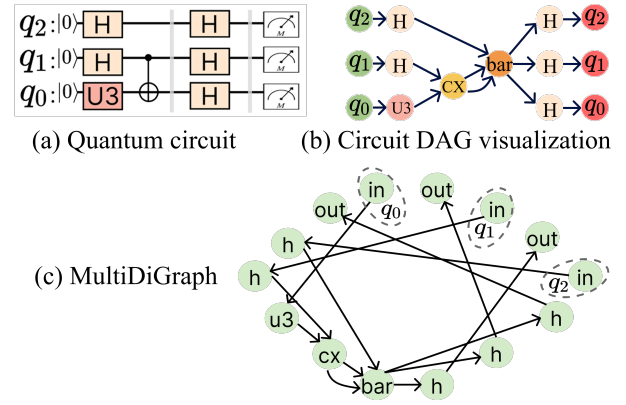


Figure 1: The DAG conversion and multi-directed graph of the BV algorithm quantum circuit.




Figure 2: Illustrative explanation of circuit analysis.

metrics defined as relative structural changes:

$$\Delta_{\text{struct}}(\mathcal{G}_0, \mathcal{G}_t) = \{\Delta_{\text{deg}}, \Delta_{\text{path}}, \Delta_{\text{conn}}\}, \quad (1)$$

where Δ_{deg} measures shifts in node degree distribution, Δ_{path} captures the expansion of critical paths and long dependency chains, and Δ_{conn} reflects effective connectivity inflation. These metrics summarize how transpilation reshapes the circuit’s structural footprint in a backend-aware yet noise-model-agnostic manner.

Control-Flow-Aware Noise Propagation Model. Based on the deformed circuit graph \mathcal{G}_t , we model noise propagation using a control-flow-aware random walk (Figure 2) that captures how local gate errors accumulate and spread along dependency chains during circuit execution. Let $A \in \mathbb{R}^{n \times n}$ denote a weighted adjacency matrix incorporating structural deformation, where $n = |V_t|$. Each weight A_{ij} reflects the strength of the control-flow dependency from gate v_i to v_j , modulated by deformation-induced factors such as dependency expansion and effective connectivity inflation. We define the corresponding noise-propagation operator

$$P = D^{-1}A, \quad D = \text{diag}(d_1, \dots, d_n), \quad d_i = \sum_j A_{ij}, \quad (2)$$

which normalizes outgoing dependencies to form a row-stochastic transition matrix. Each entry P_{ij} represents the conditional probability that noise originating at gate v_i influences gate v_j through valid control-flow paths. Repeated application of P thus models a Markovian diffusion process over the circuit’s dependency structure, providing an

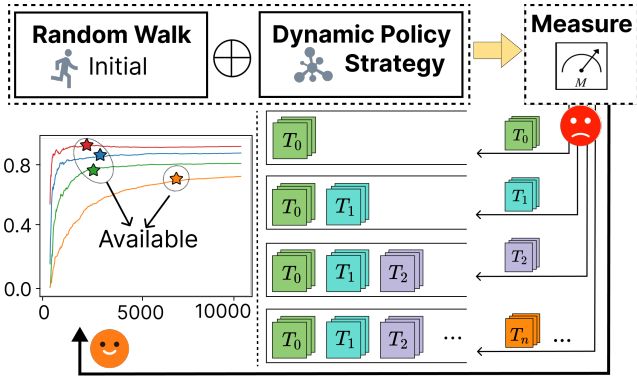


Figure 3: Adaptive measurement planning procedure.

abstraction of cumulative noise propagation across execution steps. Unlike uniform random walks, the deformation-aware weighting in A biases propagation toward structurally critical paths, such as long dependency chains or high fan-in regions. As a result, the induced operator P reflects not only logical control flow but also backend-induced structural effects that amplify or attenuate noise transmission, making it a suitable foundation for structure-aware fidelity analysis.

Spectral Complexity Estimation. We quantify circuit complexity through the spectral characteristics of the noise-propagation operator P , which provide an intrinsic and time-independent description of how noise diffuses through circuit dependencies. Let $\{\lambda_i\}_{i=1}^n$ denote the eigenvalues of P , ordered by decreasing magnitude. Unlike procedural convergence metrics that depend on iterative dynamics, spectral components characterize global structural properties of the underlying graph and capture the strength, persistence, and reach of dependency-driven noise diffusion. In particular, dominant eigenvalues correspond to slowly decaying propagation modes, indicating long-range coupling and strong interdependence among gates. Such modes imply that local noise can influence a larger portion of the circuit and persist across execution steps, increasing the uncertainty of fidelity estimation. We therefore define the circuit complexity as

$$C(\mathcal{G}) = \sum_{i=1}^k |\lambda_i|, \quad (3)$$

where $k \ll n$ retains only the dominant spectral modes that contribute most significantly to noise propagation. This truncated spectral aggregation yields a compact yet expressive complexity measure that is robust to graph size and minor structural variations.

Intuitively, a larger spectral mass indicates stronger long-range dependencies and slower attenuation of noise influence, suggesting that more measurements are required for empirical fidelity estimates to converge. Conversely, circuits with rapidly decaying spectra exhibit limited noise diffusion and require fewer samples to achieve stable estimates. By grounding complexity estimation in spectral properties, our approach provides a principled and backend-aware basis for adaptive measurement planning.

Adaptive Measurement Planning. Fidelity estimation is formulated as an adaptive testing process guided by structural complexity. Given the estimated spectral complexity $C(\mathcal{G})$ and the transpiled circuit depth \mathcal{D} , we determine the batch size as

$$\mathcal{P} = C(\mathcal{G}) \cdot \log(\mathcal{D}), \quad (4)$$

which assigns larger measurement budgets to circuits exhibiting stronger structural deformation and deeper dependency chains. This logarithmic scaling with depth prevents over-allocation for deep but weakly coupled circuits, while ensuring sufficient sampling for circuits with pronounced long-range dependencies. Measurements are executed iteratively, and fidelity statistics are updated after each batch, allowing sampling effort to be dynamically adjusted based on structural characteristics and convergence behavior. Figure 3 illustrates this adaptive batching process. At each iteration, a batch of \mathcal{P} measurements is executed and aggregated with previous results, and the process continues until the uncertainty-aware stopping criterion is satisfied.

Uncertainty-Aware Termination. To ensure statistical reliability, we employ a confidence-interval-based stopping criterion that explicitly controls estimation uncertainty. Let \hat{F} be the empirical fidelity estimate and σ its standard deviation. Sampling terminates once

$$z_\alpha \cdot \frac{\sigma}{\sqrt{|T|}} \leq \delta, \quad (5)$$

ensuring that the estimation error does not exceed δ with confidence level $(1 - \alpha)$. This uncertainty-aware policy complements the structure-guided batching strategy by halting measurements as soon as statistical stability is achieved, thereby avoiding redundant sampling while preserving rigorous accuracy guarantees.

Overall, our methodology formulates fidelity estimation as a structure-aware and uncertainty-driven test planning problem. By integrating structural deformation analysis, control-flow-aware noise propagation, spectral complexity estimation, and adaptive stopping, QuFid provides an efficient fidelity oracle for noisy quantum programs.

Evaluation

This section describes the experimental design and infrastructure used to evaluate QuFid. We first outline the evaluation objectives, followed by the experimental environment, benchmark circuits, and implementation details. Our evaluation is designed to assess QuFid from multiple complementary perspectives. Specifically, we examine whether QuFid can reduce the number of fidelity measurements required to achieve a fixed precision target, and how its adaptive behavior trades off estimation accuracy against measurement cost under varying test budgets. We further analyze the quality of measurement allocation, including test effectiveness, scalability, and adaptability.

Experimental Platform and Settings. All experiments are conducted on IBM Quantum hardware platforms (IBM 2022) via the Qiskit SDK (Aleksandrowicz et al. 2019). Specifically, we use three IBM backends: *Sherbrooke*, *Kyiv*, and *Brisbane*, each offering up to 127 qubits.

Notation	Quantum Circuit	Notation	Quantum Circuit	Notation	Quantum Circuit
BV	Bernstein-Vazirani	Clifford	Random Clifford circuit	Ising	Linear Ising model
QAOA	Quantum Approx. Optimization	VQE	Variational Quantum Eigensolver	QFT	Quantum Fourier Transform
QKNN	Quantum k-NN	QNN	Quantum Neural Network	QPE	Quantum Phase Estimation
QSVM	Quantum Support Vector Machine	QuGAN	Quantum GAN	RB	Randomized Benchmarking
Amplitude	Amplitude Estimation	Shor	Shor's factoring algorithm	Simon	Simon's algorithm
SU2	SU(2) circuits	VQC	Variational Quantum Circuit	XEB	Cross-Entropy Benchmarking

Table 1: Quantum circuit benchmarks used in our experiments.

Iter. Times	BV				QKNN				QPE			
	4 qubits	6 qubits	8 qubits	10 qubits	4 qubits	6 qubits	8 qubits	10 qubits	4 qubits	6 qubits	8 qubits	10 qubits
50	0.0062	0.0308	0.1416	0.0247	0.3889	0.5318	0.9321	0.8609	0.0516	0.0728	0.0648	0.0700
550	0.0150	0.0077	0.0358	0.0035	0.0255	0.0429	0.1646	0.2574	0.0080	0.0022	0.0179	0.0162
1050	0.0075	0.0014	0.0083	0.0028	0.0135	0.0260	0.1084	0.1290	0.0092	0.0047	0.0156	0.0063
1550	0.0056	0.0025	0.0008	0.0004	0.0100	0.0151	0.0662	0.0891	0.0070	0.0044	0.0010	0.0141
2050	0.0012	0.0051	0.0079	0.0020	0.0069	0.0117	0.0392	0.0634	0.0048	0.0024	0.0010	0.0051
2550	0.0001	0.0078	0.0060	0.0027	0.0002	0.0110	0.0264	0.0506	0.0040	0.0042	0.0050	0.0033
3050	0.0026	0.0056	0.0018	0.0021	0.0006	0.0064	0.0211	0.0441	0.0030	0.0036	0.0041	0.0034
3550	0.0014	0.0049	0.0010	0.0048	0.0013	0.0045	0.0180	0.0326	0.0029	0.0012	0.0051	0.0018
4050	0.0005	0.0042	0.0006	0.0047	0.0009	0.0028	0.0153	0.0273	0.0001	0.0004	0.0039	0.0001
4550	0.0016	0.0049	0.0009	0.0032	0.0018	0.0034	0.0138	0.0226	0.0002	0.0018	0.0027	0.0058
5050	0.0008	0.0035	0.0024	0.0064	0.0013	0.0036	0.0109	0.0189	0.0018	0.0029	0.0007	0.0015
6050	0.0003	0.0014	0.0018	0.0027	0.0008	0.0020	0.0073	0.0161	0.0006	0.0005	0.0002	0.0029
7050	0.0009	0.0000	0.0026	0.0014	0.0006	0.0019	0.0052	0.0081	0.0004	0.0016	0.0019	0.0019
8050	0.0013	0.0016	0.0020	0.0011	0.0007	0.0002	0.0047	0.0045	0.0024	0.0026	0.0014	0.0022
9050	0.0005	0.0005	0.0017	0.0002	0.0004	0.0005	0.0022	0.0011	0.0008	0.0021	0.0002	0.0004
Our bias	0.0071	0.0017	0.0084	0.0033	0.0076	0.0086	0.0092	0.0082	0.0039	0.0022	0.0008	0.0064
Our iter.	592	792	860	981	1192	2827	5584	6992	173	954	1570	1906

Table 2: Comparison of the vanilla iterative processes and the proposed scheme for 4 qubit sizes in 3 circuits, *i.e.*, ‘BV’, ‘QKNN’, and ‘QPE’, with setting the fidelity bias to less than 0.01. ‘Our iter’ and ‘Our bias’ correspond to required iterations and achieved fidelity bias. ‘Iter. Times’ refers to the number of vanilla iterations.

The default fidelity measurement for baselines is set to $\text{shots}=10000$, with sampling intervals of $\text{step}=20$ for iteration comparisons. We set the confidence level $\alpha = 0.05$ and the default fidelity bias threshold $\delta = 0.01$.

Benchmark Circuits. We evaluate QuFid on 18 representative quantum algorithms (Zhang et al. 2023; Li, Ding, and Xie 2019, 2020b,a) widely used in software verification, optimization, and simulation. As summarized in Table 1, these include classical algorithms (BV, Simon, QPE, QFT, Shor, Clifford), variational algorithms (QAOA, VQE, VQC, SU2), machine learning models (QKNN, QNN, QSVM, QuGAN), and benchmarking circuits (RB, XEB, and Amplitude Estimation). Each circuit is evaluated under qubit sizes $\{4, 6, 8, 10\}$ and Figure 4 illustrates example circuits.

Efficiency Evaluations. We evaluate 18 benchmarks to estimate the number of measurements required to achieve a fidelity deviation below 0.01, as shown in Figure 5. The results show that required iterations generally increase with qubit count. Most algorithms, except BV, QPE, and Simon, need around 8000 measurements for 10-qubit circuits due to higher structural complexity, whereas 4-qubit circuits typically converge within a few hundred iterations. This variation indicates that using a fixed shot count can either waste resources or reduce accuracy, highlighting the need for adaptive measurement. Across all settings, QuFid reduces measurement shots compared to the fixed-shot baseline.

Fidelity Measurement Trends. To further analyze fidelity

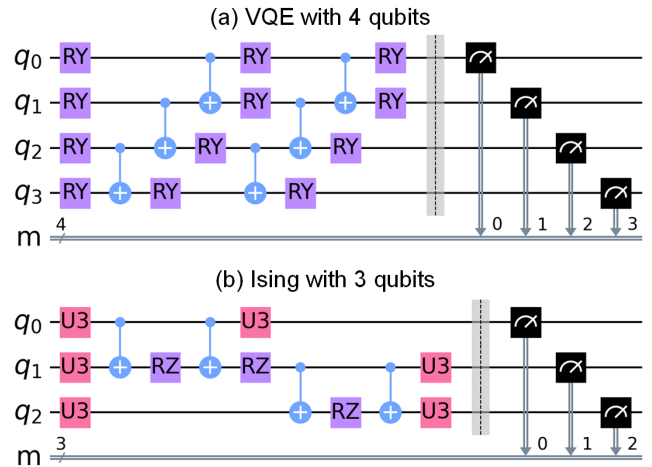


Figure 4: Some instances of quantum circuits.

evolution with increasing measurements, Table 2 reports results for three representative algorithms, Bernstein-Vazirani (BV), Quantum k-Nearest Neighbor (QKNN), and Quantum Phase Estimation (QPE), on circuits with 4-10 qubits. Experiments are run on three IBM devices, and results are averaged. Absolute fidelity deviation is measured using the Hellinger distance between the estimated distribution and




Figure 5: The number of iterations required for 18 circuits with 4 qubit sizes sets the fidelity bias to less than 0.01.

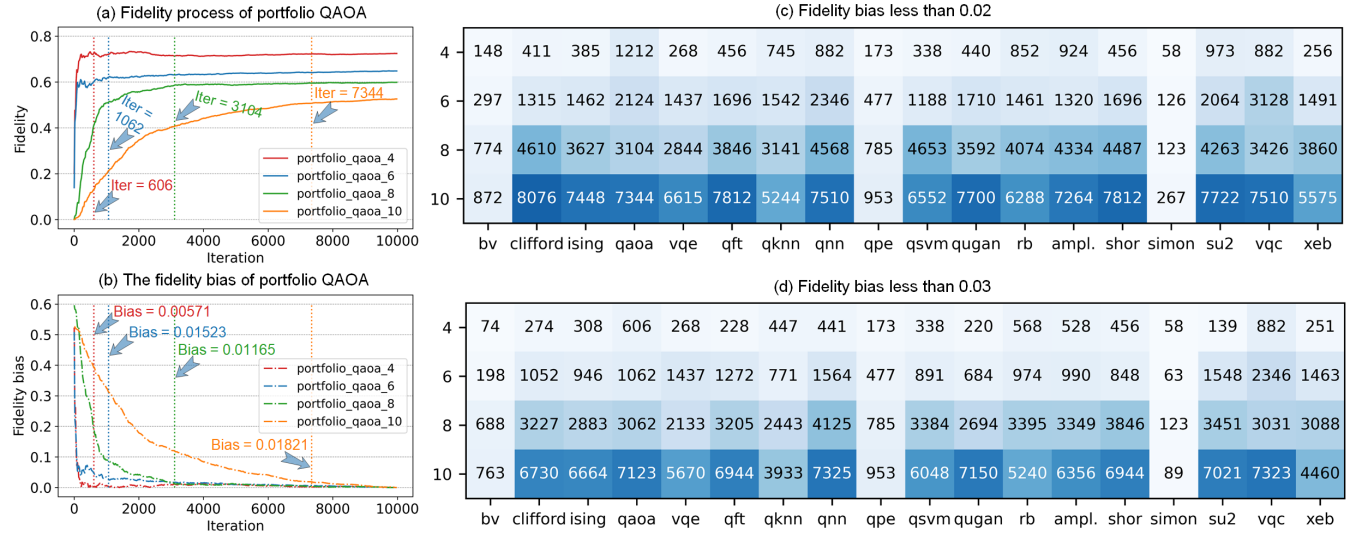


Figure 6: Case study and evaluation with different measurement biases.

the “true” distribution obtained from 10,000 shots. Vanilla iteration uses a fixed sampling interval of 20 up to 10,000 iterations, with our evaluation sampled every 500 iterations. Gray shading denotes the estimated iteration bounds. All deviations are within 0.01. As shown in Table 2, BV and QPE require relatively stable measurement counts as qubit size grows, while QKNN exhibits a steep increase: about 1,550, 2,550, 5,050, and 7,050 measurements for 4, 6, 8, and 10 qubits, respectively, reflecting higher structural complexity. In all cases, QuFid maintains an absolute fidelity deviation below 0.01 while using significantly fewer measurements than baseline approaches, demonstrating robustness and efficiency for resource-constrained quantum experiments.

Case Study. We use the portfolio Quantum Approximate Optimization Algorithm (QAOA) as a case study to visualize how QuFid adapts measurement counts and fidelity bias as iterations increase. Experiments are conducted on 4, 6, 8, and 10 qubits, with vertical dashed lines in Figures 6(a)-(b) indicating the iteration points selected by QuFid. *Measurement iterations.* As shown in Figure 6(a), larger qubit sizes require more iterations to stabilize fidelity. For all qubit settings, fidelity increases initially, fluctuates briefly, and then converges. Stabilization occurs at approximately 606, 1062, 3104, and 7344 iterations for 4-, 6-, 8-, and 10-qubit circuits, respectively, consistent with the estimates made by QuFid.

Fidelity bias. Figure 6(b) shows that fidelity bias decreases sharply at first, fluctuates, and then converges as iterations grow. Smaller qubit circuits converge faster: the 4-qubit circuit reaches a bias of 0.00571 in 606 iterations, while the 10-qubit circuit reaches 0.01821 after 7344 iterations. Overall, QuFid achieves stable fidelity with fewer tests than fixed-shot methods, balancing accuracy and resource usage.

Different Measurement Biases. To study how bias tolerance affects measurement cost, we analyze the number of required measurements under different fidelity bias thresholds. Figures 6(c)-(d) report the measurement counts when the fidelity bias is constrained to be below 0.02 and 0.03, respectively. Across all benchmarks and circuit sizes, relaxing the bias tolerance consistently reduces the required number of measurements, often by several thousand shots. The reduction becomes more pronounced for deeper and structurally complex circuits. For example, for the Ising benchmark with 10 qubits, the required measurement count decreases from 7448 under a bias threshold of 0.02 to 6664 when the threshold is relaxed to 0.03. Similar trends are observed for Clifford, QFT, and QKNN circuits, indicating a clear trade-off between estimation accuracy and measurement cost. These results confirm that QuFid can effectively exploit looser accuracy requirements to substantially reduce testing overhead, while maintaining controlled fidelity bias.




Figure 7: Comparison of measurement shots and bias.

Comparison with Existing Methods. We compare QuFid with representative learning-based fidelity predictors, QuCT (Tan et al. 2023) and QuEst (Wang et al. 2022a). Unlike QuCT and QuEst, which predict fidelity values from circuit features, QuFid adaptively determines the minimal number of measurements needed to achieve reliable fidelity estimation. Figure 7 reports both the measurement shots consumed (bar plots) and the resulting fidelity bias (line plots) across representative benchmark circuits. Across all cases, QuFid consistently requires substantially fewer measurement shots while achieving lower or comparable fidelity bias. In contrast, QuCT incurs higher measurement cost due to its reliance on fine-grained feature modeling and validation, while QuEst exhibits the largest shot consumption and higher bias, reflecting the overhead and generalization limitations of graph-based prediction models. These results highlight the advantage of QuFid in jointly optimizing measurement efficiency and estimation accuracy under realistic NISQ settings.

Discussion and Future Work

This work reframes fidelity estimation for quantum programs as a structure-aware and uncertainty-driven test planning problem. Rather than treating the number of measurements as a fixed hyperparameter or relying on pre-trained predictors, QuFid dynamically allocates measurement budgets by jointly considering circuit structure, transpilation-induced deformation, and runtime statistical feedback. The experimental results demonstrate that this perspective is both effective and practical under realistic NISQ conditions.

About Circuit Representation. A key insight of QuFid is that circuit structure, especially after transpilation, plays a central role in determining how noise accumulates and propagates. Our results show that circuits with similar depth or qubit counts can exhibit markedly different convergence behaviors due to differences in dependency patterns and effective connectivity. By modeling these effects through a control-flow-aware noise-propagation operator and spectral complexity, QuFid captures global structural characteristics that are missed by depth-based or gate-count-based heuristics. This explains why QuFid consistently adapts measurement budgets more accurately than fixed-shot strategies

across diverse circuit families.

Interpretability and Practicality. Compared to learning-based approaches such as QuCT and QuEst, QuFid does not require offline training, historical execution traces, or explicit noise models. This makes QuFid particularly well-suited for cold-start scenarios and time-varying hardware environments, where prior knowledge may be unavailable or outdated. While learning-based methods aim to predict fidelity values directly, QuFid focuses on determining the minimal number of measurements needed to achieve a desired accuracy, resulting in lower effective measurement cost and stronger robustness guarantees. Each component of the framework, structural deformation metrics, spectral complexity, and confidence-driven stopping, has a clear physical or statistical meaning. This transparency makes the system easier to reason about and debug, and allows practitioners to adjust accuracy targets or confidence levels in practice.

Limitations and Future Directions. Despite its effectiveness, QuFid has several limitations that suggest directions for future work. First, the current formulation assumes a Markovian abstraction of noise propagation, which may not fully capture long-range temporal correlations in certain devices. Second, while spectral complexity provides a compact global measure, incorporating localized or qubit-specific deformation metrics may further improve adaptivity. Finally, extending QuFid to support other frameworks, *e.g.*, MindSpore Quantum (Xu et al. 2024b), and multi-objective optimization, such as jointly optimizing fidelity, latency, and energy consumption, remains promising directions.

Conclusion

In this paper, we present QuFid, an adaptive and noise-aware framework that determines measurement budgets online by integrating structural circuit analysis with runtime statistical guarantees. By modeling quantum programs as directed acyclic graphs and capturing transpilation-induced deformation through a control-flow-aware noise-propagation model, QuFid quantifies circuit complexity via spectral characteristics and uses this information to guide adaptive measurement planning. Extensive experiments on 18 representative quantum benchmarks executed on real IBM Quantum backends demonstrate that QuFid substantially reduces measurement cost while maintaining strict fidelity accuracy constraints. Compared with fixed-shot strategies and state-of-the-art learning-based predictors, QuFid achieves a superior balance between accuracy and resource efficiency. We believe that QuFid represents a step toward principled, interpretable, and resource-efficient testing methodologies for quantum programs.

Acknowledgments

This research was supported by the Shanghai Qi Zhi Institute Innovation Program (No. SQZ202318) and the CPS-Yangtze Delta Region Industrial Innovation Center of Quantum and Information Technology-MindSpore Quantum Open Fund.

References

- Aleksandrowicz, G.; et al. 2019. Qiskit: An open-source framework for quantum computing. <https://zenodo.org/records/2562111/>. Accessed: 2020-11-20.
- Ayazadeh, R.; Mousavi, A.; Alavisamani, N.; and Qureshi, M. 2023. Enigma: Privacy-Preserving Execution of QAOA on Untrusted Quantum Computers. *arXiv preprint arXiv:2311.13546*.
- Bergholm, V.; Izaac, J.; Schuld, M.; Gogolin, C.; Ahmed, S.; Ajith, V.; Alam, M. S.; Alonso-Linaje, G.; AkashNarayanan, B.; Asadi, A.; et al. 2018. Pennylane: Automatic differentiation of hybrid quantum-classical computations. *arXiv preprint arXiv:1811.04968*.
- Biamonte, J.; Wittek, P.; Pancotti, N.; Rebentrost, P.; Wiebe, N.; and Lloyd, S. 2017. Quantum machine learning. *Nature*, 549(7671): 195–202.
- Boixo, S.; Isakov, S. V.; Smelyanskiy, V. N.; Babbush, R.; Ding, N.; Jiang, Z.; Bremner, M. J.; Martinis, J. M.; and Neven, H. 2018. Characterizing quantum supremacy in near-term devices. *Nature Physics*, 14(6): 595–600.
- Breuer, H.-P.; and Petruccione, F. 2002. *The theory of open quantum systems*. Oxford University Press, USA.
- Campbell, E. 2024. A series of fast-paced advances in Quantum Error Correction. *Nature Reviews Physics*, 6(3): 160–161.
- Cerezo, M.; Poremba, A.; Cincio, L.; and Coles, P. J. 2020. Variational quantum fidelity estimation. *Quantum*, 4: 248.
- Chen, R.; Song, Z.; Zhao, X.; and Wang, X. 2021. Variational quantum algorithms for trace distance and fidelity estimation. *Quantum Science and Technology*, 7(1): 015019.
- Coifman, R. R.; and Lafon, S. 2006. Diffusion maps. *Applied and computational harmonic analysis*, 21(1): 5–30.
- Craswell, N.; and Szummer, M. 2007. Random walks on the click graph. In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, 239–246.
- Fontana, E.; Rudolph, M. S.; Duncan, R.; Rungger, I.; and Cirstoiu, C. 2025. Classical simulations of noisy variational quantum circuits. *npj Quantum Information*, 11(1): 84.
- Gilyén, A.; and Poremba, A. 2022. Improved quantum algorithms for fidelity estimation. *arXiv preprint arXiv:2203.15993*.
- Hu, Z.; et al. 2023. Battle Against Fluctuating Quantum Noise: Compression-Aided Framework to Enable Robust Quantum Neural Network. In *DAC*.
- Huang, H.-Y.; Kueng, R.; and Preskill, J. 2020. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10): 1050–1057.
- Huang, K.; Wang, Z.-A.; Song, C.; Xu, K.; Li, H.; Wang, Z.; Guo, Q.; Song, Z.; Liu, Z.-B.; Zheng, D.; et al. 2021. Quantum generative adversarial networks with multiple superconducting qubits. *npj Quantum Information*, 7(1): 165.
- IBM. 2022. IBMQ Quantum. <https://quantum.cloud.ibm.com/>. Accessed: 2022-03-15.
- Kandala, A.; Temme, K.; Córcoles, A. D.; Mezzacapo, A.; Chow, J. M.; and Gambetta, J. M. 2019. Error mitigation extends the computational reach of a noisy quantum processor. *Nature*, 567(7749): 491–495.
- Li, G.; Ding, Y.; and Xie, Y. 2019. Tackling the qubit mapping problem for NISQ-era quantum devices. In *ASPLOS*.
- Li, G.; Ding, Y.; and Xie, Y. 2020a. Eliminating redundant computation in noisy quantum computing simulation. In *2020 57th ACM/IEEE Design Automation Conference (DAC)*, 1–6. IEEE.
- Li, G.; Ding, Y.; and Xie, Y. 2020b. Towards efficient superconducting quantum processor architecture design. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, 1031–1045.
- Li, R.-H.; Yu, J. X.; Qin, L.; Mao, R.; and Jin, T. 2015. On random walk based graph sampling. In *2015 IEEE 31st international conference on data engineering*, 927–938. IEEE.
- Li, T.; Lu, L.; Zhao, Z.; Tan, Z.; Tan, S.; and Yin, J. 2024. QuST: Optimizing quantum neural network against spatial and temporal noise biases. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*.
- Li, T.; Zhao, Z.; and Yin, J. 2024. Minerva: Enhancing Quantum Network Performance for High-Fidelity Multimedia Transmission. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 3704–3712.
- Li, T.; et al. 2025a. Empowering quantum serverless circuit deployment optimization via graph contrastive learning and learning-to-rank co-designed approaches. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence*, 9250–9258.
- Li, T.; et al. 2025b. Fortuna: Towards Efficient Selection of High-Fidelity Link for Quantum Network in the Wild. In *IEEE INFOCOM 2025-IEEE Conference on Computer Communications*, 1–10. IEEE.
- Li, T.; et al. 2025c. Task-Driven Device Fingerprinting for Quantum Cloud Platforms via Modeling QNN Outcomes under Noise. *IEEE Transactions on Information Forensics and Security*.
- Liang, Z.; Song, Z.; Cheng, J.; He, Z.; Liu, J.; Wang, H.; Qin, R.; Wang, Y.; Han, S.; Qian, X.; et al. 2023. Hybrid gate-pulse model for variational quantum algorithms. In *2023 60th ACM/IEEE Design Automation Conference (DAC)*, 1–6. IEEE.
- Liu, J.; and Zhou, H. 2020. Reliability modeling of nisq-era quantum computers. In *2020 IEEE international symposium on workload characterization (IISWC)*, 94–105. IEEE.
- Magesan, E.; Gambetta, J. M.; and Emerson, J. 2012. Characterizing quantum gates via randomized benchmarking. *Physical Review A—Atomic, Molecular, and Optical Physics*, 85(4): 042311.
- Masuda, N.; Porter, M. A.; and Lambiotte, R. 2017. Random walks and diffusion on networks. *Physics reports*, 716: 1–58.
- McKay, D. C.; Alexander, T.; Bello, L.; Biercuk, M. J.; Bishop, L.; Chen, J.; Chow, J. M.; Córcoles, A. D.; Egger,

- D.; Filipp, S.; et al. 2018. Qiskit backend specifications for openqasm and openpulse experiments. *arXiv preprint arXiv:1809.03452*.
- Nielsen, M. A.; and Chuang, I. L. 2001. *Quantum computation and quantum information*, volume 2. Cambridge university press Cambridge.
- Preskill, J. 2018. Quantum computing in the NISQ era and beyond. *Quantum*.
- Qiang, X.; Loke, T.; Montanaro, A.; Aungskunsiri, K.; Zhou, X.; O'Brien, J. L.; Wang, J. B.; and Matthews, J. C. 2016. Efficient quantum walk on a quantum processor. *Nature communications*, 7(1): 11511.
- Shor, P. W. 1999. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM review*, 41(2): 303–332.
- Smith, K. N.; Perlin, M. A.; Gokhale, P.; Frederick, P.; Owusu-Antwi, D.; Rines, R.; Omole, V.; and Chong, F. 2023. Clifford-based circuit cutting for quantum simulation. In *Proceedings of the 50th Annual International Symposium on Computer Architecture*, 1–13.
- Sun, L.; Petrenko, A.; Leghtas, Z.; Vlastakis, B.; Kirchmair, G.; Sliwa, K.; Narla, A.; Hatridge, M.; Shankar, S.; Blumoff, J.; et al. 2014. Tracking photon jumps with repeated quantum non-demolition parity measurements. *Nature*, 511(7510): 444–448.
- Tan, K. C.; and Volkoff, T. 2021. Variational quantum algorithms to estimate rank, quantum entropies, fidelity, and Fisher information via purity minimization. *Physical Review Research*, 3(3): 033251.
- Tan, S.; Lang, C.; Xiang, L.; Wang, S.; Jia, X.; Tan, Z.; Li, T.; Yin, J.; Shang, Y.; Python, A.; et al. 2023. QuCT: A Framework for Analyzing Quantum Circuit by Extracting Contextual and Topological Features. In *Proceedings of the 56th Annual IEEE/ACM International Symposium on Microarchitecture*, 494–508.
- Tong, H.; Faloutsos, C.; and Pan, J.-Y. 2006. Fast random walk with restart and its applications. In *Sixth international conference on data mining (ICDM'06)*, 613–622. IEEE.
- Urbanek, M.; Nachman, B.; and de Jong, W. A. 2020. Error detection on quantum computers improving the accuracy of chemical calculations. *Physical Review A*, 102(2): 022427.
- Wang, H.; Liu, P.; Cheng, J.; Liang, Z.; Gu, J.; Li, Z.; Ding, Y.; Jiang, W.; Shi, Y.; Qian, X.; et al. 2022a. Quest: Graph transformer for quantum circuit reliability estimation. *arXiv preprint arXiv:2210.16724*.
- Wang, H.; et al. 2022b. QuantumNAS: Noise-adaptive search for robust quantum circuits. In *HPCA*.
- Wang, Q.; Zhang, Z.; Chen, K.; Guan, J.; Fang, W.; Liu, J.; and Ying, M. 2022c. Quantum algorithm for fidelity estimation. *IEEE Transactions on Information Theory*, 69(1): 273–282.
- Xiang, L.; Chen, J.; et al. 2024. Enhanced quantum state transfer by circumventing quantum chaotic behavior. *Nature Communications*, 15(1): 4918.
- Xu, S.; Zhao, Y.; Huang, L.; and Qiao, C. 2024a. Routing and Photon Source Provisioning in Quantum Key Distribution Networks. In *IEEE INFOCOM 2024-IEEE Conference on Computer Communications*, 1411–1420. IEEE.
- Xu, X.; et al. 2024b. MindSpore Quantum: a user-friendly, high-performance, and AI-compatible quantum computing framework. *arXiv:2406.17248*.
- Yu, X.-D.; Shang, J.; and Gühne, O. 2022. Statistical methods for quantum state verification and fidelity estimation. *Advanced Quantum Technologies*, 5(5): 2100126.
- Zhang, H.; Wu, A.; Wang, Y.; Li, G.; Shapourian, H.; Shabani, A.; and Ding, Y. 2023. Oneq: A compilation framework for photonic one-way quantum computation. In *Proceedings of the 50th Annual International Symposium on Computer Architecture*, 1–14.
- Zhang, X.; Luo, M.; Wen, Z.; Feng, Q.; Pang, S.; Luo, W.; and Zhou, X. 2021. Direct fidelity estimation of quantum states using machine learning. *Physical Review Letters*, 127(13): 130503.
- Zhu, Z.; Gao, L.; Bao, Z.; Xiang, L.; Song, Z.; Xu, S.; Wang, K.; Chen, J.; Jin, F.; Zhu, X.; et al. 2024. Quantum highway: Observation of minimal and maximal speed limits for few and many-body states. *arXiv preprint arXiv:2408.11900*.