

PERSPECTIVE | OCTOBER 01 2026

Recoverable quantum computation: An information-centric paradigm for quantum computing with errors

Shengwang Du 



AVS Quantum Sci. 8, 040501 (2026)

<https://doi.org/10.1116/5.0353214>



Articles You May Be Interested In

The future of quantum computing with superconducting qubits

J. Appl. Phys. (October 2022)

Perspective: Coarse-grained models for biomolecular systems

J. Chem. Phys. (September 2013)

Large language models as AI agents for digital atoms and molecules: Catalyzing a new era in computational biophysics

APL Computational Physics (September 2025)



Advance your science and your career as a member of AVS

[LEARN MORE](#)

Recoverable quantum computation: An information-centric paradigm for quantum computing with errors

Cite as: AVS Quantum Sci. **8**, 040501 (2026); doi: [10.1116/5.0353214](https://doi.org/10.1116/5.0353214)

Submitted: 20 July 2026 · Accepted: 15 September 2026 ·

Published Online: 1 October 2026



View Online



Export Citation



CrossMark

Shengwang Du^{1,2,3,a)} 

AFFILIATIONS

¹Elmore Family School of Electrical and Computer Engineering, Purdue University, West Lafayette, Indiana 47907, USA

²Department of Physics and Astronomy, Purdue University, West Lafayette, Indiana 47907, USA

³Purdue Quantum Science and Engineering Institute, Purdue University, West Lafayette, Indiana 47907, USA

^{a)}Electronic mail: dusw@purdue.edu

ABSTRACT

Quantum computing promises transformative advances in computation, communication, sensing, and machine learning. Yet, the realization of large-scale fault-tolerant quantum computers remains hindered by the enormous overhead required for quantum error correction. This challenge raises a fundamental question: *Must useful quantum computing wait until fully fault-tolerant quantum hardware becomes available?* In this Perspective, we propose *recoverable quantum computation* (RQC), an information-centric paradigm for quantum computing with errors. Rather than requiring faithful preservation of the complete quantum state, RQC focuses on preserving the computational information required to accomplish a given task. A noisy quantum computation is task-level recoverable when its specified output can be inferred with prescribed accuracy and confidence using a finite, quantifiable recovery procedure. RQC may preserve quantum advantages if its overall computational cost is still lower than that of the best known classical method. We characterize recoverability through recovery overhead, task error, success probability, and end-to-end resource cost. We illustrate the framework using quantum Fourier transform period estimation on IBM quantum hardware and a conceptual example from quantum machine learning, demonstrating that useful computational information may remain recoverable despite significant physical errors. Building on these examples, we propose a preliminary classification of quantum applications according to their expected recoverability and outline a research roadmap toward a predictive theory of recoverability. RQC complements fault-tolerant quantum computing, error mitigation, quantum utility, and application-oriented benchmarking by separating the recovery of task-relevant information from the stronger claim of computational advantage.

© 2026 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1116/5.0353214>

I. INTRODUCTION

Quantum computing promises transformative advances in computation, communication, sensing, and machine learning.^{1–3} Over the past decades, remarkable progress has been achieved across a broad range of hardware platforms, including superconducting circuits, trapped ions, neutral atoms, photonic systems, and semiconductor quantum devices.^{4–11} Quantum processors containing tens to thousands of physical qubits have now been demonstrated, while landmark advances in quantum advantage, quantum simulation, variational algorithms, and quantum machine learning (QML) have fueled growing expectations for practical quantum technologies.^{8,12–14} Despite this remarkable progress, a fundamental challenge remains. Quantum systems are inherently susceptible to errors arising from

decoherence, imperfect control, environmental coupling, crosstalk, state-preparation imperfections, and measurement noise.^{3,15} As quantum processors scale to larger systems and deeper circuits, these errors accumulate and eventually degrade computational performance.

The prevailing solution is fault-tolerant quantum computing (FTQC), in which logical quantum information is protected by encoding it into many physical qubits through quantum error correction.^{16–18} Supported by the quantum threshold theorem,^{19–21} this framework provides a rigorous path toward arbitrarily accurate quantum computation. However, this capability comes at a substantial cost. Depending on the hardware platform and physical error rates, a single logical qubit may require hundreds or even thousands of physical qubits, while a practical fault-tolerant quantum computer capable

of solving large-scale problems may ultimately require millions of physical qubits,²² which significantly increases the challenges in connectivity.²³ Although impressive progress toward logical-qubit demonstrations continues,^{10,24} the hardware overhead required for fault tolerance remains one of the greatest obstacles to practical quantum computing.

This situation naturally raises a fundamental question:

Must useful quantum computing wait for fully fault-tolerant quantum hardware?

FTQC provides the general route to scalable and arbitrarily accurate quantum computation. At the same time, extensive research on noisy intermediate-scale quantum (NISQ) algorithms, quantum error mitigation, quantum utility, and application-oriented benchmarking asks whether useful task-level results may be obtained before full fault tolerance is available.^{12,25–29} The question is therefore not whether fault tolerance remains essential for general large-scale computation, but whether a scientifically and technologically meaningful intermediate regime can be characterized systematically.

The history of information science suggests a broader perspective.³⁰ Reliable communication routinely occurs over noisy channels because the transmitted information can be recovered despite imperfections.³¹ Likewise, experimental measurements in physics are almost never noise free, yet meaningful information can often be extracted through repeated measurements, statistical averaging, parameter estimation, and inference. In both cases, the essential question is not whether the underlying process is perfectly error free, but whether the desired information remains recoverable with acceptable overhead.

Quantum computing may also admit such an intermediate regime. Most quantum algorithms do not require complete reconstruction of

the final quantum state. Instead, they seek specific computational information encoded in measurement outcomes, such as a frequency, phase, classification label, optimized parameter, or physical observable. Although an n -qubit pure state is described by exponentially many amplitudes, the reported task-level output may be much smaller. Preserving a general quantum state can therefore constitute a stronger requirement than recovering a specified output, even though the latter may still be computationally demanding.

Motivated by this observation, we propose *recoverable quantum computation* (RQC), an information-centric framework for quantum computing with errors. Rather than using faithful preservation of the complete quantum state as the sole criterion, RQC focuses on whether the computational information required for a specified task can be extracted from noisy quantum outputs with prescribed accuracy and confidence using a finite, quantifiable recovery procedure. We distinguish the task-level recoverability from the advantage-preserving recoverability, which additionally requires the total quantum-classical workflow to outperform the best available classical method under matched problem and performance requirements.

RQC is closely related to NISQ computing, quantum error mitigation, quantum utility, and application-oriented benchmarking.^{12,25–29,32,33} These approaches already emphasize task performance, classical comparison, or the resources required to obtain useful results from imperfect quantum hardware. RQC is intended not to supersede them, but to isolate computational-information recovery as an explicit object of study: what task-relevant information survives noisy quantum execution, what recovery procedure extracts it, and how does the required overhead scale with circuit and problem size?

Figure 1 illustrates the conceptual relationship among classical computing, RQC, and FTQC. Classical computation is limited by




Fig. 1. Conceptual relationship among classical computing, recoverable quantum computation (RQC), and fault-tolerant quantum computing (FTQC). Classical computing is limited by classical computational complexity, whereas FTQC achieves arbitrarily accurate quantum computation through active error detection and correction. RQC describes an intermediate regime in which task-relevant computational information may remain recoverable from noisy quantum outputs through repeated executions and classical post-processing.

classical computational complexity, whereas FTQC seeks arbitrarily accurate quantum information processing through active quantum error correction. RQC focuses on the intermediate regime in which task-relevant computational information remains accessible despite noisy quantum execution. Repeated executions, statistical inference, error mitigation, and classical post-processing may improve recoverability, but their costs must be included in any subsequent claim of quantum advantage.

We anticipate that RQC could provide a useful conceptual framework for understanding the broad landscape between today's noisy quantum processors and tomorrow's fault-tolerant quantum computers. By shifting attention from state fidelity alone to the survival and recovery cost of task-relevant information, RQC offers a complementary way to characterize which computations remain operationally useful and which lose their intended output under realistic noise.

It is important to emphasize that RQC does not propose new quantum algorithms or new computational objectives. Quantum algorithms have always been designed to estimate parameters, identify solutions, generate samples, or produce decisions, and existing quantum-utility and benchmarking frameworks evaluate such outputs. The contribution proposed here is a focused organization of noisy quantum computation around computational-information recoverability: separating recovery from advantage, quantifying the additional resources required for recovery, identifying transitions from recoverable to unrecovered behavior, and comparing applications according to the structure of the information that must survive. The central contribution of this Perspective is therefore not a new algorithm or a standalone proof of quantum utility, but an information-centric framework and research program for analyzing when task-relevant outputs survive noisy quantum execution.

The remainder of this Perspective is organized as follows. Section II discusses why fault tolerance is not the only possible path toward useful quantum computation. Section III introduces RQC and develops a hierarchy of recoverability. Section IV presents operational metrics for quantifying recoverability. Sections V and VI illustrate the concept through quantum Fourier transform (QFT) period estimation and quantum machine learning, respectively. Section VII proposes a preliminary classification of quantum applications according to their expected recoverability. Section VIII outlines a research roadmap toward a predictive theory of recoverability, and Sec. IX concludes with a broader perspective on quantum computation with errors.

II. WHY FAULT TOLERANCE IS NOT THE ONLY PATH

FTQC represents one of the greatest achievements in quantum information science. The development of quantum error correction and the quantum threshold theorem established that, in principle, arbitrarily long quantum computations can be performed reliably despite the presence of physical errors.^{16–21} Consequently, fault tolerance has become the dominant paradigm for scalable quantum computation and is widely regarded as the ultimate destination of quantum computing technology.

For general scalable quantum computation, accumulated errors eventually prevent computational correctness unless they are actively controlled, and much of today's hardware development is therefore devoted to reducing physical error rates and achieving fault-tolerant operation. This does not exclude more restricted pre-fault-tolerant

tasks whose relevant outputs may tolerate or be recovered from imperfect execution.

Quantum error correction deliberately imposes a strong and general requirement: protect encoded logical information against a sufficiently broad class of physical errors throughout the computation. For many specific applications, however, successful completion is judged by a smaller task-level output encoded in the final measurement statistics rather than by complete knowledge of the underlying quantum state. The distinction between preserving a general logical state and recovering a specified computational output motivates the RQC viewpoint.

We believe this distinction is important. An n -qubit pure state is described by exponentially many amplitudes, whereas the reported output of a particular computational task may consist of only a frequency, phase, decision label, optimized parameter, observable, or sample. This difference does not imply that such outputs are automatically easy to recover, but it suggests that their robustness under noise should be studied directly rather than inferred solely from full-state fidelity.

A useful analogy can be drawn from classical communication theory. Communication channels are inherently noisy, yet reliable communication remains possible because the transmitted information can be recovered despite imperfections.³¹ The objective is not to eliminate all noise from the communication channel, but rather to ensure that the desired information remains recoverable with acceptable overhead. Experimental physics provides another familiar example. Photon shot noise, thermal fluctuations, detector imperfections, and environmental disturbances routinely affect measurements. Nevertheless, meaningful physical information can often be extracted through repeated measurements, statistical averaging, parameter estimation, and statistical inference. Experimental physicists rarely ask whether a measurement is perfectly noise free; instead, they ask whether the quantity of interest can be determined with sufficient confidence.

Quantum computing may admit a similar intermediate regime. Between today's noisy quantum processors and tomorrow's fully fault-tolerant quantum computers lies a broad region in which task-relevant computational information may remain recoverable despite imperfect quantum execution. In this regime, hardware errors degrade the quality of the quantum state without necessarily destroying the information relevant to the computational task. Repeated quantum executions, statistical inference, and classical post-processing may therefore compensate for imperfect hardware. It may still preserve computational quantum advantage if the total cost is lower than that of the best known classical methods.

Developments in NISQ computing,¹² quantum error mitigation,^{25–27,32} quantum utility,²⁸ application-oriented benchmarking,²⁹ variational quantum algorithms,^{13,34} quantum machine learning,^{14,35,36} and quantum sensing already investigate useful task-level information from imperfect hardware. RQC builds on this literature by centering the analysis on the survival, extraction, and scaling of task-specific computational information.

This observation motivates the central question addressed in this Perspective:

Under what conditions can useful computational information remain recoverable even when the underlying quantum computation is imperfect?

RQC addresses this question by shifting the focus from state fidelity alone to the computational information required by the task. The first criterion is whether that information can be recovered with specified accuracy and confidence and with quantifiable overhead. Whether the resulting workflow also preserves quantum advantage is a second criterion requiring end-to-end classical comparison.

This viewpoint naturally leads to three complementary computational regimes:

1. **Classical computation**, in which performance is fundamentally limited by classical computational complexity.
2. **Recoverable quantum computation (RQC)**, in which useful computational information remains recoverable despite noisy quantum execution.
3. **Fault-tolerant quantum computation (FTQC)**, in which active quantum error correction enables arbitrarily accurate quantum information processing.

FTQC remains the general solution for scalable, arbitrarily accurate quantum information processing. RQC neither replaces fault tolerance nor diminishes the importance of quantum error correction. Rather, it seeks to characterize a potentially broad intermediate regime in which specified task-level outputs may remain recoverable before fully fault-tolerant hardware becomes available.

The existence, boundaries, and practical significance of this intermediate regime remain open questions. Addressing these questions requires a framework for defining recoverability, identifying the computational information relevant to different quantum algorithms, and quantifying the overhead required to recover that information. Developing such a framework is the objective of Secs. III and IV.

III. RECOVERABLE QUANTUM COMPUTATION (RQC)

The discussion in Sec. II suggests that useful quantum computation does not necessarily require faithful preservation of the complete quantum state. Instead, many quantum algorithms ultimately require only the computational information extracted from the final measurement outcomes. This observation motivates a shift in perspective. Rather than asking whether a quantum computation faithfully preserves the entire quantum state, RQC asks whether the computational information required to accomplish a given task remains recoverable despite imperfect quantum execution.

Throughout this Perspective, we use the term *computational information* to denote the minimum task-specific information required to successfully complete a computational objective. Depending on the application, this information may consist of a physical parameter estimated by a quantum Fourier transform, a classification label produced by a quantum-machine-learning algorithm, an optimized variable obtained through variational optimization, or a physical observable measured in quantum sensing. Importantly, computational information should be distinguished from the complete quantum state. Although an n -qubit pure state is described by exponentially many amplitudes, the reported output of many quantum algorithms is much smaller. Identifying rigorous mathematical definitions of computational information for general quantum algorithms remains an open problem and represents an important direction for future research. From this perspective, computational information is

an application-dependent concept rather than an intrinsic property of a quantum state.

Traditional FTQC seeks to preserve the entire logical quantum state throughout a computation. RQC adopts a different evaluation criterion. Rather than evaluating noisy quantum computation primarily through preservation of the complete quantum state, it evaluates whether the computational information required by the task remains recoverable. From this perspective, errors become significant not simply because they reduce quantum-state fidelity, but because they may prevent recovery of the information needed to complete the computational task.

A. A hierarchy of recoverability

This distinction motivates three complementary notions of recoverability.

1. State recoverability

The strongest notion of recoverability concerns the quantum state itself. FTQC seeks to preserve the encoded logical quantum state throughout the computation by means of active quantum error correction. This objective forms the foundation of scalable fault-tolerant quantum computation.

2. Computational information recoverability

A weaker, yet considerably more general, notion concerns the recoverability of the computational information required by a specific algorithm. Many quantum algorithms ultimately seek only a frequency, a phase, a classification label, an optimized parameter, or another task-specific quantity. In such cases, faithful reconstruction of the complete quantum state is unnecessary provided that the desired computational information remains recoverable.

3. Advantage-preserving recoverability

A separate and more demanding notion concerns the preservation of quantum advantage. Computational information may remain recoverable even when the resources required for its recovery eliminate any advantage over classical computation. We therefore call a computation advantage preserving only when the total end-to-end cost of preparing the input, executing and repeating the quantum circuit, reading out the results, and performing classical post-processing remains lower than that of the best available classical method under matched input, output, accuracy, and confidence requirements.

Accordingly, RQC distinguishes two questions that should not be conflated: whether task-relevant computational information remains recoverable, and whether recovering that information preserves a quantum advantage. The first concerns the survival and extraction of useful information under noise; the second requires a separate end-to-end comparison with the best available classical method.

B. Definition of recoverable quantum computation

Motivated by the preceding discussion, we propose the following working definition.

Recoverable Quantum Computation (RQC).

A noisy quantum computation is computationally recoverable for a specified task if the task-relevant information can be extracted from its outputs with prescribed accuracy and confidence using a finite, quantifiable recovery procedure. It is advantage preserving only if the total cost of the resulting quantum-classical workflow remains lower than that of the best available classical method under matched problem and performance requirements.

Several aspects of this definition merit emphasis. First, recoverability is inherently application dependent. Different quantum algorithms encode computational information in different ways and therefore exhibit different sensitivities to hardware noise. Second, recoverability is fundamentally distinct from quantum-state fidelity. A computation with relatively low state fidelity may nevertheless preserve the computational information required by the application. Conversely, a computation exhibiting high state fidelity may still fail to provide useful results if the desired computational information cannot be recovered efficiently. Third, recoverability does not require active quantum error correction. Repeated quantum executions, statistical estimation, classical inference, machine learning, or other post-processing techniques may all contribute to recovering useful computational information. Finally, recoverability alone is not a claim of quantum advantage. Such a claim requires a fully specified classical baseline and consistent accounting of all relevant quantum and classical resources.

C. Computational information as a resource

The RQC framework proposed here motivates a different perspective on computational resources. In conventional FTQC, the principal resource being protected is the logical quantum state. Considerable physical resources are devoted to maintaining high quantum-state fidelity throughout the computation. RQC instead treats computational information as the primary resource. Rather than correcting every physical error, one allows imperfect quantum execution and compensates through repeated quantum measurements combined with classical post-processing. The trade-off therefore shifts from physical-qubit overhead to recovery overhead. Conceptually, this distinction resembles the difference between error correction and statistical estimation. Error correction seeks to prevent information loss, whereas statistical estimation seeks to recover useful information despite imperfect measurements. RQC adopts the latter perspective for quantum computation.

D. Relation to existing paradigms

RQC is related to several existing paradigms and is intended as a complementary organizing framework rather than a replacement for them. FTQC seeks to preserve encoded logical quantum information through active quantum error correction. Quantum error mitigation seeks to suppress or compensate for the effects of hardware noise on measured observables. NISQ computing describes a hardware regime characterized by limited qubit counts and imperfect operations.¹² Quantum utility and application-oriented benchmarking, in turn, evaluate quantum processors through task-level

performance, comparison with classical methods, and practical resource requirements.^{28,29,33} RQC shares this application-level orientation but isolates a particular question within it: what task-relevant computational information survives noisy quantum execution, and what additional resources are required to recover that information?

The specific contribution of RQC is therefore not the general principle that quantum hardware should be assessed through useful tasks, which is already central to quantum utility and application-oriented benchmarking. Rather, RQC treats computational-information recovery as an explicit object of study. It separates information recoverability from end-to-end computational advantage, characterizes the recovery overhead required to reach a prescribed task-level accuracy and confidence, and organizes applications according to whether their outputs are parameters, decisions, observables, or more complete state information. This viewpoint emphasizes recoverable-to-unrecoverable transitions as noise, circuit size, and problem size increase.

The distinction between RQC and quantum error mitigation deserves particular emphasis. Quantum error mitigation comprises a family of techniques that estimate or approximate ideal quantum outcomes by compensating for hardware errors through methods such as probabilistic error cancellation, zero-noise extrapolation, virtual distillation, symmetry verification, or other classical post-processing strategies. RQC addresses a broader and different question: under what conditions does a noisy quantum computation remain practically useful? From the perspective of RQC, quantum error mitigation may serve as one possible mechanism for improving recoverability, but it is neither required nor sufficient to define recoverability itself. A quantum computation may remain recoverable without explicit error-mitigation techniques. Conversely, successful error mitigation does not by itself establish advantage-preserving recoverability, because mitigation, sampling, and post-processing costs must be included in the end-to-end comparison. In this sense, recoverability is an application-level concept, whereas error mitigation is a hardware- and algorithm-level methodology.

E. Toward quantifying recoverability

The concepts introduced above remain qualitative. Transforming RQC into a predictive framework requires quantitative measures of task error, success probability, and recovery cost. The central question is not simply how accurately a processor preserves a quantum state, but how reliably and efficiently the desired information can be recovered from noisy outputs. A modest recovery cost may permit, but does not establish, an end-to-end quantum advantage; a rapidly growing cost may eliminate an otherwise available advantage. These observations motivate the operational quantities introduced next. The QFT (Sec. V) and QML (Sec. VI) examples then illustrate parameter recovery and decision recovery, respectively.

IV. RECOVERABILITY METRICS

Section III introduced RQC as an information-centric framework for recovering the computational information required by a specified task. The goal here is not to develop a complete mathematical theory, but to introduce operational quantities for asking: *How much additional effort is required to recover the desired information from noisy quantum outputs at a specified accuracy and confidence?* We call this additional effort the *recovery overhead*.

A. Recovery overhead

Consider a quantum algorithm that requires N_{ideal} executions under ideal conditions to reach a specified accuracy and confidence. On a noisy processor, hardware imperfections can broaden output distributions, bias measured quantities, or reduce decision margins. Consequently, N_{rec} executions may be required by an adopted recovery procedure to meet the same performance target.

We refer to N_{rec} as the *recovery execution requirement* and define the dimensionless execution-overhead factor,

$$\Gamma_{\text{rec}} = \frac{N_{\text{rec}}}{N_{\text{ideal}}}. \quad (1)$$

This distinction matters because ideal quantum algorithms are often probabilistic and may themselves require repeated measurements. In the QFT benchmark of Sec. V, $N_{\text{ideal}} = 1$, so N_{rec} and Γ_{rec} coincide numerically. More generally, recoverability depends on how the additional requirement Γ_{rec} and other recovery resources scale with problem size.

B. Recoverability and scaling

The practical accessibility of task-relevant information may depend critically on the scaling of the execution-overhead factor. Suppose that

$$\Gamma_{\text{rec}}(n) = \text{poly}(n), \quad (2)$$

we refer to the behavior as polynomial recoverability. By contrast, if

$$\Gamma_{\text{rec}}(n) \sim e^{\alpha n}, \quad (3)$$

where $\alpha > 0$, we refer to exponential recoverability. These labels describe the additional execution requirement relative to the ideal quantum procedure. Polynomial overhead does not automatically preserve quantum advantage, particularly when the ideal speedup is polynomial or when preparation and post-processing dominate. Conversely, exponential overhead does not automatically eliminate advantage if the best classical cost grows exponentially with a larger rate.

This observation naturally suggests two qualitative recovery-overhead regimes:

- **Polynomially recoverable regime:** Γ_{rec} scales polynomially with problem size.
- **Exponentially recoverable regime:** Γ_{rec} scales exponentially with problem size.

An operationally unrecoverable regime may additionally arise when the specified accuracy and confidence cannot be reached within the available resources or when noise removes the statistical distinguishability needed by the adopted recovery procedure. None of these recovery-overhead classifications, by itself, establishes or excludes quantum advantage; that question requires comparison of total quantum and classical costs.

C. End-to-end advantage comparison

Recoverability should first be quantified independently of quantum advantage through task-level quantities such as N_{rec} , estimation error, and success probability. To determine whether recoverability

also preserves a computational advantage, one must compare complete workflows. Let C_{setup} denote one-time compilation or setup cost, C_{prep} the input-state preparation cost per execution, C_{exec} the circuit-execution cost, C_{read} the readout cost, and C_{post} the cost of classical recovery and post-processing. A general end-to-end quantum-classical cost is

$$C_{\text{RQC}} = C_{\text{setup}} + N_{\text{rec}}(C_{\text{prep}} + C_{\text{exec}} + C_{\text{read}}) + C_{\text{post}}. \quad (4)$$

Here, “cost” may denote runtime, query complexity, energy, financial cost, or another explicitly specified resource, but the same cost model must be used for both quantum and classical workflows. If $C_{\text{C}}^{\text{best}}$ denotes the cost of the best available classical method for the same problem under matched input-access, output-accuracy, and confidence requirements, an advantage ratio may be defined as

$$\eta_{\text{adv}} = \frac{C_{\text{C}}^{\text{best}}}{C_{\text{RQC}}}. \quad (5)$$

A value $\eta_{\text{adv}} > 1$ indicates an advantage within the stated cost model, whereas $\eta_{\text{adv}} < 1$ does not. This ratio is meaningful only when the classical baseline and all material quantum and classical overheads are specified. It should therefore be regarded as an application-dependent resource comparison rather than a universal recoverability metric.

D. Recoverability versus fidelity

Traditional analyses of noisy quantum computation typically emphasize quantities such as state fidelity, process fidelity, logical error rates, and gate fidelities. These metrics are indispensable for characterizing quantum hardware, but they do not directly measure whether the computational information required by a particular application remains recoverable. Recoverability and fidelity are closely related but conceptually distinct. A computation may exhibit relatively low state fidelity while still preserving the computational information required by the application. Conversely, a computation with high state fidelity may nevertheless fail to produce useful results if the desired computational information cannot be recovered efficiently.

This distinction is particularly important for applications such as quantum machine learning, optimization, parameter estimation, and quantum sensing, where the final objective is often a decision, an optimized parameter, or a physical observable rather than complete reconstruction of the quantum state. From the perspective of RQC, fidelity measures the quality of the quantum evolution, whereas recoverability measures the usefulness of the computational information that survives.

E. Toward application-specific recoverability

The quantities introduced above are intentionally general. Different algorithms encode their outputs in different statistical features and can therefore exhibit different recovery behavior. In the present QFT benchmark, information is carried by a spectral peak; in a classifier, it may be carried by a probability margin; and in sensing or metrology, it may be carried by an estimator of a physical parameter.

Recoverability should therefore not be viewed as a single universal quantity. It is characterized by application-dependent measures defined for a specified task, accuracy, confidence, recovery rule, and cost model. Sections V and VI consider parameter recovery using the

QFT and decision recovery in QML. The former is an experimental benchmark; the latter is a conceptual construction.

These operational quantities are a starting point rather than a complete mathematical theory. A predictive formulation will require precise task-dependent definitions of computational information, error and confidence targets, recovery procedures, and resource costs under general noise models, together with connections to complexity theory, quantum information theory, and statistical decision theory.

V. EXAMPLE I: QUANTUM FOURIER TRANSFORM AND PERIOD ESTIMATION

One central hypothesis of the RQC framework is that useful computational information may remain recoverable even after the underlying quantum state has been substantially degraded by hardware noise. Quantum Fourier transform (QFT) period estimation provides a natural benchmark for investigating this hypothesis because its computational objective—the estimation of a single frequency—is precisely defined, quantitatively measurable, and theoretically well understood. Moreover, period estimation forms the computational core of several foundational quantum algorithms, including Shor's factoring algorithm and related hidden-subgroup algorithms.^{37,38} It therefore provides a natural platform for investigating how quantum noise influences the recoverability of computational information.

The purpose of the present experiment is not to demonstrate an end-to-end computational speedup, but to determine how the measurement resources required to recover a specified Fourier component change as hardware noise and circuit size increase. It should therefore be understood as a controlled QFT recoverability benchmark rather than a complete implementation of period finding in Shor's algorithm.

For the controlled period-estimation problem considered here, the desired computational information is a specified signal-frequency bin. Although the underlying state is described by exponentially many amplitudes, the reported output is one frequency. This separation makes the experiment a useful test of task-level recoverability, while not by itself establishing computational advantage.

We consider a complex signal,

$$s(t) = e^{i(2\pi f_0 t + \phi)}, \quad (6)$$

where f_0 is the signal frequency and ϕ is the initial phase. Throughout this example, we choose $f_0 = R/4$ and $\phi = \pi/10$, where R denotes the sampling rate. A conventional 2^n -point fast Fourier transform (FFT) uses 2^n uniformly sampled values to obtain the frequency-bin spacing $\Delta f = R/2^n$. This provides a convenient sampling reference for the present benchmark, but it should not be interpreted as the complexity of the best possible classical estimator for the restricted single-frequency problem considered here. Specialized classical estimators may exploit prior knowledge of the signal model, noise, or allowed frequencies. By contrast, an ideal error-free quantum computer encodes the sampled signal into an n -qubit quantum state,

$$|\psi_{f_0}\rangle = \frac{1}{\sqrt{2^n}} \sum_{x=0}^{2^n-1} e^{i(\pi x/2 + \phi)} |x\rangle, \quad (7)$$

and applies the inverse QFT. The correct frequency bin, $m_0 = 2^n f_0 / R = 2^{n-2}$, is then obtained with unit probability in a single execution, $N_{\text{QFT}} = 1$.

In this controlled benchmark, the target phase pattern is directly encoded into the input circuit. For the chosen value $f_0 = R/4$, this state has a particularly simple phase structure and can be prepared efficiently using single-qubit operations. The experiment does not, however, address the general cost of loading an unknown classically sampled signal into a quantum state. Any end-to-end comparison with classical frequency estimation would need to specify the input-access model and include state preparation together with circuit execution, readout, and classical post-processing costs.

The information-centric perspective developed in Secs. II–IV naturally leads to the central question of this experiment:

Can the desired frequency remain recoverable when the QFT is implemented on a realistic noisy quantum processor?

Hardware noise inevitably broadens the Fourier spectrum and reduces the probability associated with the correct frequency bin. From the perspective of RQC, however, this degradation is not the primary issue. The key question in this experiment is whether the desired computational information—the signal frequency—can still be recovered and how its measurement-shot overhead changes with circuit size.

To investigate this question, we implemented QFT circuits on IBM superconducting quantum processors for $n = 4, 8, 10, 12, 14$, and 16 qubits. For each circuit size, the target frequency corresponds to the Fourier peak located at $m_0 = 2^n f_0 / R = 2^{n-2}$. Figure 2 summarizes the experimental results. Panels (a1)–(a5) present the measured probability spectra obtained directly from IBM quantum hardware. The horizontal axis is expressed in the normalized frequency coordinate $2^n f / R$, which coincides with the discrete frequency-bin index. Because the input signal satisfies $f_0 = R/4$, the correct frequency always appears at the position 2^{n-2} .

As the circuit size increases, the measured probability spectra evolve in a systematic and physically intuitive manner. Figures 2(a1)–2(a5) show that the dominant Fourier peak remains correctly located at the target frequency for $n = 4$ through $n = 14$. At the same time, however, the peak probability decreases steadily while an increasingly broad background of competing peaks emerges. These observations are consistent with the cumulative effects of hardware noise during quantum computation.

As the number of qubits and quantum gates increases, accumulated gate errors redistribute probability amplitude from the target Fourier mode into many other computational basis states. Because the dimension of the Hilbert space grows exponentially with the number of qubits, the number of accessible error states also increases rapidly. The experimental consequence is a progressive broadening of the Fourier spectrum accompanied by a steady reduction of the probability associated with the correct frequency. The measured peak probability decreases from approximately 0.87 for $n = 4$ to 0.51 for $n = 8$, 0.25 for $n = 10$, 0.15 for $n = 12$, and only 0.025 for $n = 14$. By $n = 16$ (not shown in Fig. 2), the correct peak is no longer distinguishable from competing peaks generated by accumulated hardware errors, indicating that the QFT spectrum itself has become severely degraded.




FIG. 2. Experimental demonstration of RQC using QFT period estimation on IBM quantum hardware. (a1)–(a5) Measured probability spectra for $n = 4, 8, 10, 12$, and 14 qubits. The horizontal axis is the normalized frequency coordinate $2^n f / R$, which coincides with the discrete frequency-bin index. The correct signal frequency corresponds to $f_0 = R/4$ and is located at 2^{n-2} . (b1)–(b5) Recovery of the computational information (signal frequency). The vertical axis shows the normalized frequency error, $2^n (\bar{f} - f_0) / R$, as a function of the number of recorded shots, N_{rec} . Error bars represent one standard deviation obtained from independent subsets of the experimental data.

Viewed solely from the perspective of quantum-state fidelity, these observations might appear discouraging. The central question posed by RQC, however, is fundamentally different. Rather than asking whether the probability spectrum remains nearly ideal, we ask

whether the desired computational information—the signal frequency—can still be recovered despite the degradation of the underlying quantum state. To answer this question, we analyze the complete measurement record rather than individual probability spectra. The

experimental data are divided into independent groups containing N_{rec} repeated quantum executions. For each group, the recovered frequency is defined as the frequency bin with the largest number of recorded counts, $\hat{f} = (R/2^n) \arg\max_k N_k$, where N_k denotes the number of occurrences measured in the k th frequency bin. Repeating this procedure for all independent groups yields a distribution of recovered frequencies, from which the mean recovered frequency $\bar{f}$ and its standard deviation σ_f are determined. Figures 2(b1)–2(b5) plot the normalized frequency error, $2^n(\bar{f} - f_0)/R$, as a function of the number of recorded shots. The error bars represent one standard deviation obtained from independent subsets of the experimental data.

Several important observations immediately emerge. Although the probability spectra shown in Figs. 2(a1)–2(a5) become progressively broader with increasing circuit size, the recovered frequency converges steadily toward the correct value as additional quantum executions are accumulated. Even after the probability associated with the correct Fourier peak has been reduced substantially, repeated measurements together with straightforward classical post-processing remain sufficient to recover the desired frequency with high accuracy over a broad range of circuit sizes. This behavior highlights the central distinction between quantum-state fidelity and computational-information recoverability. The degradation of the quantum state does not immediately imply the loss of the computational information. Instead, useful information may remain recoverable long after substantial distortion has appeared in the measured probability spectrum.

For $n = 4$ –12, we define the recoverability threshold, N_{rec}^* , as the smallest value among the evaluated group sizes at which the root-mean-square frequency error becomes smaller than one frequency bin. In normalized bin units,

$$\text{RMSE}_m = \sqrt{\left[\frac{2^n(\bar{f} - f_0)}{R} \right]^2 + \left(\frac{2^n \sigma_f}{R} \right)^2} < 1. \quad (8)$$

This criterion accounts for both bias and statistical dispersion. Direct reanalysis of the raw shot records gives $N_{\text{rec}}^* = 3, 19, 80$, and 121 for $n = 4, 8, 10$, and 12, respectively, retaining the thresholds obtained in the original analysis.

For descriptive purposes, we also define an FFT-referenced sampling ratio,

$$\rho_{\text{samp}} = \frac{2^n}{N_{\text{rec}}^*}. \quad (9)$$

This dimensionless quantity compares the number of samples in the conventional 2^n -point FFT reference with the noisy-QFT measurement shots required by Eq. (8). It characterizes shot recovery within this controlled benchmark and is not an end-to-end quantum-advantage metric. The experimentally determined results are summarized in Table I.

Several important conclusions can be drawn from these results. First, for $n = 4$ –12, the target frequency satisfies the root-mean-square error (RMSE) recovery criterion with relatively modest measurement-shot overhead. Although the measured probability spectra are progressively distorted by hardware noise, the task-relevant Fourier component remains recoverable. The values $\rho_{\text{samp}} > 1$ describe this shot-count comparison with the conventional FFT reference; they do not establish computational quantum advantage.

TABLE I. Recoverability threshold N_{rec}^* and FFT-referenced sampling ratio ρ_{samp} for QFT period estimation on IBM quantum hardware. For $n = 4$ –12, N_{rec}^* satisfies the RMSE criterion in Eq. (8). For $n = 14$, approximately 10 200 accumulated shots return the target as the modal bin, but the available record does not contain enough independent groups to establish an RMSE-based threshold. The $n = 16$ target is not recovered under the tested conditions.

n	N_{rec}^* or scale	ρ_{samp}
4	3	5.33
8	19	13.47
10	80	12.80
12	121	33.85
14	$\sim 10\,200^a$	...
16	Not recovered	...

^aObserved cumulative-shot recovery scale; statistical confidence cannot be independently determined from the available 16 384-shot record.

Second, a pronounced transition occurs at $n = 14$. Although the target is the largest peak in the full probability record, several competing peaks have comparable amplitudes. The accumulated modal estimate returns the target bin at approximately 10 200 shots. However, because the complete record contains only 16 384 shots, this accumulation length leaves only one nonoverlapping group and therefore cannot provide an independent estimate of the RMSE or recovery confidence. We consequently identify $n = 14$ as a marginal-recovery regime and report 10 200 only as an observed cumulative-shot recovery scale, not as a statistically established N_{rec}^* .

Finally, at $n = 16$, the correct frequency is not recovered from the available 65 536-shot record using the adopted modal estimator. The target peak is no longer distinguishable from stronger competing peaks generated by accumulated coherent and stochastic errors. We therefore describe the result as not recovered under the tested hardware, dataset, and recovery procedure, rather than as fundamentally unrecoverable.

These observations naturally divide the experimental results into three computational regimes:

1. **Efficiently recoverable regime ($n = 4$ –12):** The target frequency satisfies the RMSE criterion with relatively modest measurement-shot overhead.
2. **Marginal-recovery regime ($n = 14$):** The accumulated record returns the target only after approximately 10^4 shots, while the available data are insufficient to establish recovery confidence from independent groups.
3. **Not recovered under the tested conditions ($n = 16$):** The target frequency is not identified reliably from the available measurement record using the adopted estimator.

Taken together, these observations provide experimental support for the central distinction underlying RQC in the specific context of QFT-based period estimation: substantial degradation of the measured output distribution does not immediately imply loss of the task-relevant information encoded in that distribution. The experiment reveals a progression from efficient recovery to marginal recovery and ultimately unsuccessful recovery under the tested conditions. It does not establish an end-to-end computational quantum advantage, which would require a matched comparison including input access,

state preparation, circuit execution, readout, repetition, classical post-processing, and the best available classical estimator.

The QFT experiment illustrates a clear distinction between quantum-state fidelity and computational-information recoverability. The former characterizes the quality of the quantum evolution, whereas the latter determines whether the computation remains practically useful. This distinction forms the conceptual foundation of RQC and motivates the broader question of how recoverability manifests across different classes of quantum algorithms.

Section VI considers how the same information-centric questions may apply to a different output type. Whereas the QFT benchmark recovers a frequency bin, quantum classification seeks a discrete decision. This comparison motivates a general study of task-specific outputs without assuming that the two applications share the same recovery mechanism or resource scaling.

VI. EXAMPLE II: RECOVERABILITY IN QUANTUM MACHINE LEARNING

The QFT example presented in Sec. V experimentally illustrates recovery of a specified Fourier component from a strongly degraded output distribution. RQC, however, is not restricted to parameter-estimation problems. Many quantum algorithms seek a discrete computational decision rather than a parameter. Quantum machine learning (QML) provides a conceptual example of this broader class of applications; unlike the QFT study, no hardware demonstration or quantitative advantage claim is made here.

In supervised quantum classification, the reported objective is typically an output label rather than reconstruction of the complete quantum state. From the perspective of RQC, the desired computational information is therefore the classification decision itself. A compact output does not guarantee robustness or computational advantage, because feature encoding, circuit depth, decision margin, noise, training, sampling, and classical alternatives all matter. It nevertheless motivates studying decision recovery directly rather than inferring it solely from state fidelity.

This observation motivates asking whether classification decisions can remain stable when noise perturbs the output distribution.^{14,35,36} The relevant mechanism and overhead must be determined for each classifier, dataset, and noise process.

A. Decision recovery versus state recovery

The distinction between state recovery and decision recovery becomes particularly clear in QML. Consider an n -qubit quantum processor described by

$$|\Psi\rangle = \sum_{k=0}^{2^n-1} c_k |k\rangle, \quad (10)$$

where the quantum state is described by 2^n complex probability amplitudes. FTQC protects the encoded logical state against physical errors, thereby preserving all subsequent logical operations and measurements. In contrast, a classification task reports only a label derived from the final measurement statistics. For a classification problem involving C classes, the desired output is simply

$$\hat{y} \in \{1, 2, \dots, C\}, \quad (11)$$

rather than the complete quantum state.

RQC therefore asks whether the decision derived from the noisy output statistics remains recoverable at a specified confidence. This task does not require reconstructing the complete state, but it may still demand substantial resources if the class-probability margin is small or noise changes the leading class.

B. A minimal quantum classifier

To illustrate this idea, we consider a minimal quantum classifier consisting of three stages: quantum feature encoding, quantum evolution, and measurement. An input feature vector $\mathbf{x}$ is first encoded into an n -qubit quantum state,

$$|\psi(\mathbf{x})\rangle = U_{\text{enc}}(\mathbf{x})|0\rangle^{\otimes n}, \quad (12)$$

where U_{enc} denotes the feature-encoding circuit. The encoded state subsequently evolves according to

$$|\phi(\mathbf{x})\rangle = U|\psi(\mathbf{x})\rangle, \quad (13)$$

where U may represent either a fixed quantum circuit or a trained variational quantum model.

The probability associated with class i is obtained from the final measurement,

$$P_i(\mathbf{x}) = \sum_{z \in C_i} |\langle z | \phi(\mathbf{x}) \rangle|^2, \quad (14)$$

where C_i denotes the subset of computational basis states assigned to class i . The predicted class is therefore

$$\hat{y} = \arg \max_i P_i. \quad (15)$$

Unlike FTQC, successful quantum classification does not require preservation of every probability amplitude. It requires only that the correct decision remain identifiable from the measured probability distribution.

C. Interference-dominated decision making

For classifiers in which class i is associated with a designated readout amplitude, one possible physical interpretation becomes transparent in terms of quantum interference. That probability amplitude may be expressed as

$$A_i(\mathbf{x}) = \sum_P a_i(P; \mathbf{x}), \quad (16)$$

where P labels the different computational pathways contributing coherently to the final outcome. The corresponding class probability is

$$P_i(\mathbf{x}) = |A_i(\mathbf{x})|^2. \quad (17)$$

For the correct class, constructive interference enhances the total probability amplitude, whereas destructive interference suppresses competing classes. We refer to this process as *decision interference*, emphasizing that the interference determines the computational decision rather than the complete quantum state.

On a noisy quantum processor, each computational pathway is perturbed,

$$a_i(P; \mathbf{x}) \rightarrow a_i(P; \mathbf{x}) + \delta_i(P; \mathbf{x}), \quad (18)$$

leading to a redistribution of probability among different measurement outcomes. Nevertheless, the correct decision may remain recoverable if the constructive interference responsible for the correct classification continues to dominate the accumulated perturbations. This observation motivates the following working principle.

Interference-Dominance Hypothesis.

For algorithms whose desired outcome is produced by a dominant set of coherent pathways, decision recoverability may persist while the interference margin favoring that outcome remains larger than the perturbations induced by noise.

This statement is a heuristic hypothesis, not a general recoverability criterion. Realistic stochastic and nonunitary noise is more generally described by quantum channels and cannot always be represented as additive pathway amplitudes. A rigorous treatment would need to relate the decision margin or other task-specific statistic to the noisy output distribution and to the number of samples required for a specified confidence.

D. Recoverability of classification decisions

Because individual quantum measurements are probabilistic, repeated executions may be used to recover the correct classification decision. Suppose that N_{rec} repeated measurements produce N_i occurrences of class i . The recovered decision is then

$$\hat{y} = \arg \max_i N_i. \quad (19)$$

This recovery procedure is directly analogous to the frequency recovery introduced in Sec. V. In the QFT example, repeated measurements identify the dominant Fourier component. Here, repeated measurements identify the dominant classification outcome.

Let $N_{\text{rec}}^*(\epsilon, \delta)$ denote the minimum number of repeated measurements required to recover a decision whose error is at most ϵ with confidence at least $1 - \delta$. This quantity characterizes task-level sampling overhead. Whether the corresponding classifier is advantage preserving must be evaluated separately using the end-to-end cost in Eq. (4), including feature encoding, training or parameter optimization, repeated circuit execution, readout, and classical post-processing, and comparing against the best available classical classifier under matched data-access and performance requirements.

Although a detailed quantitative analysis lies beyond the scope of this Perspective, one possible behavior parallels the QFT example. If noise reduces but does not reverse the probability margin between the leading and competing classes, repeated measurements may recover the same decision with increasing sampling overhead. If noise changes the underlying ordering of class probabilities, additional shots alone will converge to the wrong decision and cannot restore the ideal classification. Quantifying this transition requires a specified dataset, circuit, noise model, recovery rule, and confidence criterion.

Together, the QFT experiment and QML construction illustrate two complementary forms of task-specific information: parameter recovery and decision recovery. The former provides hardware data; the latter identifies a hypothesis and a quantitative program for future study. Both motivate evaluating the task-relevant output directly, but

neither implies that output recoverability alone establishes quantum advantage.

VII. A PRELIMINARY CLASSIFICATION OF QUANTUM APPLICATIONS BY RECOVERABILITY

The preceding examples suggest organizing applications partly by the structure of the information they ultimately report. Recoverability is not a property of the output type alone: it belongs to a specified task, algorithm, circuit, noise process, recovery procedure, and performance target. Nevertheless, distinguishing parameters, decisions, collections of observables, and richer state information may help formulate testable comparisons.

This observation motivates a preliminary classification of quantum applications according to their expected recoverability. The classification proposed here is intended not as a definitive taxonomy, but as a conceptual framework for organizing future theoretical and experimental investigations of RQC, as shown in Fig. 3.

A. Parameter-recovery applications

Many quantum algorithms ultimately seek to recover one or a few continuous physical parameters. Representative examples include QFT-based period estimation, quantum phase estimation, quantum sensing, quantum metrology, and Hamiltonian parameter estimation.^{37,38}

The QFT benchmark in Sec. V illustrates this class: despite substantial spectral broadening, the target frequency satisfies the RMSE criterion for $n = 4$ – 12 , becomes marginal at $n = 14$, and is not recovered at $n = 16$ under the tested conditions. Quantum sensing and metrology likewise estimate physical parameters from repeated measurements, although their noise models, estimators, and resource comparisons differ from those of the present benchmark.

Because the reported output consists of only a small number of parameters, these applications may provide favorable opportunities for recovery when the relevant statistical features remain identifiable under noise. A compact output alone is neither sufficient for recoverability nor evidence of quantum advantage; sensitivity, required precision, state preparation, circuit depth, estimator performance, and classical baselines must also be considered.

B. Decision-recovery applications

A second class of quantum algorithms seeks discrete computational decisions rather than continuous physical parameters. Representative examples include QML, quantum pattern recognition, quantum hypothesis testing, and related quantum decision-making algorithms.^{14,35,36}

As discussed in Sec. VI, quantum classifiers report a predicted label rather than the complete quantum state, so faithful reconstruction of every amplitude is unnecessary for successful classification. The interference-dominance hypothesis offers one possible mechanism for robustness in algorithms with a clear coherent-pathway interpretation. More generally, decision recoverability depends on whether the noisy output distribution retains a statistically resolvable margin favoring the desired class.

Decision-recovery applications may therefore offer favorable cases, but this remains a hypothesis requiring algorithm-specific theoretical, numerical, and experimental tests.

C. State-information applications

A third class of applications requires substantially richer information from the quantum state itself. Representative examples include quantum chemistry, general quantum simulation, many-body quantum dynamics, and quantum state tomography.³⁹ These applications often require recovery of a large collection of observables, correlation functions, or, in the limiting case of quantum state tomography, reconstruction of the complete quantum state.

Because these applications can require many observables, correlations, or state parameters, their sampling and mitigation costs may be larger. Repeated measurements reduce statistical uncertainty but cannot by themselves reverse a noise-induced bias or restore physically lost coherence; additional mitigation or error correction may be required.

Recoverability in these applications may therefore be more limited, and FTQC may ultimately remain indispensable for large-scale implementations. Nevertheless, identifying which observables remain recoverable despite imperfect quantum evolution remains an important open question deserving future investigation.

D. Recoverability as an organizing principle

We hypothesize that the nature and structure of the task-relevant information are important determinants of recoverability, alongside the algorithm, circuit, noise, and recovery procedure. Figure 3 summarizes this perspective conceptually. Tasks reporting a few parameters or decisions may offer higher recoverability when those features retain a sufficient statistical margin, while tasks requiring many observables or more complete state information may require

greater recovery overhead. The locations in Fig. 3 are qualitative and do not assert a universal ordering.

These observations motivate the following working hypothesis.

Recoverability Hypothesis.

We hypothesize that recoverability depends jointly on the amount and structure of the task-relevant information, the statistical margin by which it is encoded in the output, the algorithm and circuit, the noise process, and the adopted recovery procedure.

The hypothesis therefore treats the computational objective as one determinant among several, not as a sufficient predictor. Its value must be assessed by quantitative comparisons across algorithms and hardware.

The classification proposed here is intentionally preliminary. Many important quantum algorithms—including Grover search, variational quantum algorithms, large-scale quantum chemistry simulations, quantum communication, distributed quantum computing, and quantum networking—may occupy intermediate regions of the recoverability landscape whose properties remain largely unexplored. Developing quantitative criteria for locating quantum algorithms within this landscape represents one of the central challenges for future research.

VIII. RESEARCH ROADMAP

RQC provides an information-centric framework for studying task-level outputs in the presence of hardware imperfections. Sections III–VII suggest a research program for characterizing where computational information remains recoverable, where recovery becomes marginal, and where it is not recovered under specified conditions.




FIG. 3. Conceptual recoverability landscape of quantum applications. Recoverability is expected to decrease as the amount and complexity of the computational information to be recovered from the quantum state increase. Parameter-recovery and decision-recovery applications are anticipated to exhibit higher recoverability than applications requiring preservation of large amounts of quantum-state information. The placement of individual applications is illustrative and conceptual rather than quantitative.

Transforming this conceptual framework into a predictive theory will require substantial advances in both theory and experiment. We highlight several directions for testing, refining, or delimiting the proposed framework.

A. Toward a predictive theory of recoverability

At present, recoverability is characterized through operational quantities such as recovery overhead, task-level error, and success probability. An important theoretical challenge is to relate these quantities to circuit structure, hardware noise, the statistical distinguishability of relevant outputs, and the computational information required by different algorithms. Advantage preservation should then be evaluated through the separate end-to-end resource comparison defined in Sec. IV.

One particularly important question is whether recovery overhead and task error can be estimated before a quantum computation is performed. Such a capability would guide algorithm design, hardware development, and resource estimation and could identify boundaries among efficient recovery, marginal recovery, and failure to recover under specified noise, resource, and confidence conditions.

B. Recoverability-aware quantum algorithms

RQC suggests an application-level design objective: optimize the robustness and statistical distinguishability of task-relevant outputs, not only generic circuit fidelity. This could motivate algorithms, compilers, and error-mitigation strategies that reduce recovery overhead for a specified task under realistic hardware imperfections. In algorithms with a coherent-pathway interpretation, preserving a favorable interference margin may be useful; in more general settings, the appropriate target may be an observable error, decision margin, estimator variance, or success probability. This direction appears particularly relevant to variational and other hybrid quantum-classical approaches, whose performance depends on ansatz design, hardware noise, sampling overhead, and classical optimization.³⁴

C. Recoverability within application-oriented benchmarking

Hardware-oriented metrics such as gate fidelity, logical error rates, and quantum volume remain indispensable.^{10,40} At the same time, existing application-oriented benchmarks already evaluate task-level performance and resource requirements.²⁹ RQC should therefore be developed within, rather than presented as preceding, this broader benchmarking program.

A recoverability-oriented benchmark would add measurements of how task error, success probability, and recovery overhead evolve with noise, circuit size, and problem size. To enable meaningful hardware comparisons, it should specify input access, accuracy and confidence targets, execution and readout costs, classical post-processing, and strong classical baselines.

D. Exploring the recoverability landscape

The examples presented in this Perspective represent only the first steps toward understanding RQC. Many important quantum algorithms and applications—including Grover search, variational

quantum optimization, quantum chemistry, quantum simulation, quantum communication, distributed quantum computing, quantum networking, and quantum sensing^{39,41,42}—remain to be investigated from the perspective of recoverability.

An especially important direction is to develop quantitative recoverability maps for specified algorithms, noise models, and performance targets. Such maps could identify which tasks benefit from noisy hardware under a stated cost model and which require mitigation, partial error correction, or fault-tolerant operation.

More broadly, RQC should not be viewed as an alternative to FTQC, but as a complementary framework for characterizing task-level information in the pre-fault-tolerant and partially error-corrected regimes. Understanding recoverability can identify where noisy hardware retains useful outputs, where recovery overhead becomes prohibitive, and where fault-tolerant operation is required.

IX. CONCLUSIONS

FTQC provides the general route toward scalable, arbitrarily accurate quantum information processing through quantum error correction. This vision remains indispensable for general large-scale quantum computation. At the same time, the substantial resources required for fault tolerance motivate a complementary question: which specified task-level outputs remain recoverable on imperfect hardware, and at what cost?

In this Perspective, we have proposed RQC as an information-centric framework for quantum computing with errors. RQC asks whether the information required for a specified task can be recovered from noisy outputs with prescribed accuracy and confidence using a quantifiable recovery procedure. It complements quantum error correction, error mitigation, quantum utility, and application-oriented benchmarking by isolating computational-information recovery as an explicit object of study. Crucially, task-level recoverability is distinct from advantage-preserving recoverability, which requires a separate end-to-end comparison with the best available classical method.

We introduced operational quantities based on recovery overhead, task error, success probability, and total workflow cost. In the QFT benchmark, a specified Fourier component remains recoverable by the RMSE criterion for $n = 4$ – 12 despite strongly degraded output distributions; $n = 14$ exhibits only marginal cumulative-record recovery, and the target is not recovered at $n = 16$ under the tested conditions. This experiment demonstrates a recoverability transition, not end-to-end quantum advantage. The QML discussion provides a complementary conceptual construction for decision recovery and identifies the decision margin and sampling confidence as subjects for future quantitative study.

We also proposed a deliberately preliminary application landscape. Parameter- and decision-recovery tasks may offer favorable cases when their relevant output features retain sufficient statistical margins, whereas tasks requiring many observables or richer state information may incur larger recovery costs. This is a working hypothesis, not a universal ordering: recoverability depends jointly on the output structure, algorithm, circuit, noise, accuracy and confidence requirements, input access, and recovery procedure.

RQC does not replace FTQC or diminish the importance of quantum error correction. FTQC protects encoded logical information to support general reliable computation; RQC evaluates whether a specified output remains recoverable under imperfect execution and quantifies the resources required. The two perspectives address different

levels of the same problem and can be applied across noisy, partially error-corrected, and fault-tolerant hardware.

Many important questions remain open. The ideas introduced in this Perspective are intended not as a complete theory, but as the starting point for a broader research program. Although only one representative application has been experimentally examined here, we anticipate that the proposed framework will motivate further theoretical and experimental investigations across diverse classes of quantum algorithms and help define the emerging science of recoverability in quantum computation. Ultimately, the future of quantum computing may depend not only on how well we preserve quantum states, but also on how effectively we preserve the computational information required to solve meaningful problems.

ACKNOWLEDGMENTS

We acknowledge the use of IBM Quantum services for this work. The views expressed are those of the author and do not reflect the official policy or position of IBM or the IBM Quantum team. The quantum circuits were executed using Qiskit v2.3.0 on the ibm_kingston processor. The work was supported by the NSF (Grant Nos. 2503230 and 2228725).

AUTHOR DECLARATIONS

Conflict of Interest

The author has no conflicts to disclose.

Author Contributions

Shengwang Du: Conceptualization (lead); Data curation (lead); Formal analysis (lead); Funding acquisition (lead); Investigation (lead); Methodology (lead); Project administration (lead); Resources (lead); Software (lead); Supervision (lead); Validation (lead); Visualization (lead); Writing – original draft (lead); Writing – review & editing (lead).

DATA AVAILABILITY

The data that support the findings of this study are openly available in Zenodo at <http://doi.org/10.5281/zenodo.21209526>, Ref. 43.

REFERENCES

- ¹R. P. Feynman, *Int. J. Theor. Phys.* **21**, 467 (1982).
- ²D. Deutsch, *Proc. R. Soc. A* **400**, 97 (1985).
- ³M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th Anniversary ed. (Cambridge University Press, 2010).
- ⁴J. M. Gambetta, J. M. Chow, and M. Steffen, *npj Quantum Inf.* **3**, 2 (2017).
- ⁵C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, *Appl. Phys. Rev.* **6**, 021314 (2019).
- ⁶D. S. Weiss and M. Saffman, *Phys. Today* **70**(7), 44 (2017).
- ⁷L. Isenhour, E. Urban, X. L. Zhang, A. T. Gill, T. Henage, T. A. Johnson, T. G. Walker, and M. Saffman, *Phys. Rev. Lett.* **104**, 010503 (2010).
- ⁸F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao *et al.*, *Nature* **574**, 505 (2019).
- ⁹H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding *et al.*, *Science* **370**, 1460 (2020).
- ¹⁰Google Quantum AI, *Nature* **614**, 676 (2023).
- ¹¹D. Bluvstein, S. J. Evered, A. A. Geim, S. H. Li, H. Zhou, T. Manovitz, S. Ebadi, M. Cain, M. Kalinowski *et al.*, *Nature* **626**, 58 (2024).
- ¹²J. Preskill, *Quantum* **2**, 79 (2018).
- ¹³A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, *Nat. Commun.* **5**, 4213 (2014).
- ¹⁴J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, *Nature* **549**, 195 (2017).
- ¹⁵*Quantum Error Correction*, edited by D. A. Lidar and T. A. Brun (Cambridge University Press, 2013).
- ¹⁶P. W. Shor, *Phys. Rev. A* **52**, R2493 (1995).
- ¹⁷A. M. Steane, *Phys. Rev. Lett.* **77**, 793 (1996).
- ¹⁸D. Gottesman, *Phys. Rev. A* **57**, 127 (1998).
- ¹⁹D. Aharonov and M. Ben-Or, *SIAM J. Comput.* **38**, 1207 (2008).
- ²⁰E. Knill, R. Laflamme, and W. H. Zurek, *Science* **279**, 342 (1998).
- ²¹P. Aliferis, D. Gottesman, and J. Preskill, *Quantum Inf. Comput.* **6**, 97 (2006).
- ²²A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, *Phys. Rev. A* **86**, 032324 (2012).
- ²³S. Du, Y. Ding, and C. Qiao, *AVS Quantum Sci.* **7**, 013803 (2025).
- ²⁴A. Paetznick, B. W. Reichardt, M. P. da Silva, C. Ryan-Anderson, D. Aasen, J. M. Bello-Rivas, J. P. Campora III, R. Chao, A. Chernoguzov *et al.*, *Nature* **654**, 349 (2026).
- ²⁵K. Temme, S. Bravyi, and J. M. Gambetta, *Phys. Rev. Lett.* **119**, 180509 (2017).
- ²⁶S. Endo, S. C. Benjamin, and Y. Li, *Phys. Rev. X* **8**, 031027 (2018).
- ²⁷Z. Cai, *Quantum* **5**, 548 (2021).
- ²⁸Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. van den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel *et al.*, *Nature* **618**, 500 (2023).
- ²⁹T. Lubinski, S. Johri, P. Varosy, J. Coleman, L. Zhao, J. Necaie, C. H. Baldwin, K. Mayer, and T. Proctor, *IEEE Trans. Quantum Eng.* **4**, 3100332 (2023).
- ³⁰R. Landauer, *Phys. Today* **44**(5), 23 (1991).
- ³¹C. E. Shannon, *Bell Syst. Tech. J.* **27**, 379 (1948).
- ³²W. J. Huggins, S. McArdle, T. E. O'Brien, J. Lee, N. C. Rubin, S. Boixo, K. B. Whaley, R. Babbush, and J. R. McClean, *Phys. Rev. X* **11**, 041036 (2021).
- ³³E. T. Campbell, A. Khurana, and A. Montanaro, *Quantum* **3**, 167 (2019).
- ³⁴M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan *et al.*, *Nat. Rev. Phys.* **3**, 625 (2021).
- ³⁵V. Havlicek, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, *Nature* **567**, 209 (2019).
- ³⁶M. Schuld and N. Killoran, *Phys. Rev. Lett.* **122**, 040504 (2019).
- ³⁷P. W. Shor, "Algorithms for quantum computation: Discrete logarithms and factoring," in *Proceedings of the 35th Annual Symposium on Foundations of Computer Science* (IEEE, 1994), pp. 124–134.
- ³⁸D. Coppersmith, "An approximate Fourier transform useful in quantum factoring," IBM Research Report RC19642, 1994.
- ³⁹S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, *Rev. Mod. Phys.* **92**, 015003 (2020).
- ⁴⁰A. W. Cross, L. S. Sheldon, P. D. Nation, and J. M. Gambetta, *Phys. Rev. A* **100**, 032328 (2019).
- ⁴¹L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing* (ACM, 1996), pp. 212–219.
- ⁴²C. L. Degen, F. Reinhard, and P. Cappellaro, *Rev. Mod. Phys.* **89**, 035002 (2017).
- ⁴³S. Du (2026). "Recoverable quantum computation: An information-centric paradigm for quantum computing with errors," Zenodo. <https://doi.org/10.5281/zenodo.21209526>