International Conference on Multidisciplinary Perspectives in Advanced Computing and Technology (IMPACT 2026)
G. B. Pant University of Agriculture and Technology, Uttarakhand, India. Jan. 10-11, 2026
DMPedia Lecture Notes in Multidisciplinary Research,
Volume 1, pp. 644-650, 2026

Privacy in Quantum Computing: A Comprehensive Review

Bagisha Roy

School of Computer Science Engineering and TechnologyBennett University, India
S25MTCG0003@bennett.edu.in

ABSTRACT
Quantum technologies are beginning to reshape how large volumes of data are processed, yet they also introduce a new class of privacy challenges with no classical counterparts. These concerns stem from the availability of quantum side information, the encoding of data in non-classical states, and the fragility of conventional cryptographic systems in post-quantum environments. This work reviews the main privacy-preserving strategies proposed in peer-reviewed literature from 2020 to 2025, covering both quantum and post-quantum settings. The survey examines several key directions: (i) the integration of differential privacy into quantum machine learning pipelines, (ii) privacy guarantees derived from gentle-measurement principles, (iii) secure aggregation for quantum federated learning, and (iv) privacy amplification methods designed to remain robust under noisy quantum-channel conditions. Representative examples include the mechanisms introduced by Watkins et al. for embedding differential privacy into variational circuits, the theoretical limits established by Aaronson and Rothblum for privacy under gentle measurements, the distributed-learning protections developed by Liu and Zhang, and the noise-aware key-extraction analysis presented by Singh and colleagues. System-level considerations for future quantum infrastructures are discussed in the work of, while previous literature proposes hybrid quantum classical designs to strengthen privacy in post-quantum frameworks. Taken together, these studies reveal open challenges related to scalability, composability, and adversarial modelling, and they motivate a roadmap toward developing practical and verifiable privacy guarantees for emerging quantum platforms.

Keywords: *Quantum Privacy, Differential Privacy, Post-Quantum Cryptography, Quantum Federated Learning, Privacy Amplification*

1 Introduction

Quantum technologies are steadily progressing toward practical deployment, and this evolution directly affects data privacy, system security, and the design of next-generation information infrastructures. As quantum devices mature, they introduce privacy concerns that differ fundamentally from those found in classical systems. These issues arise from the presence of quantum side information, leakage induced by measurement processes, state disturbance, and the breakdown of privacy assumptions embedded within classical algorithms and cryptographic schemes. Consequently, recent research has focused on identifying and mitigating privacy risks across quantum machine learning, quantum communication, and post-quantum cryptography.

In the realm of quantum machine learning, Watkins *et al.* investigate how differential-privacy mechanisms must be reformulated for use in variational quantum circuits. Their results show that noise added for privacy interacts in non-trivial ways with quantum gradients and measurement statistics, altering traditional interpretations of privacy budgets. From a theoretical standpoint, Aaronson and Rothblum analyse gentle-measurement principles and characterise conditions under which sequential measurements reveal limited information while still preserving high state fidelity.

For distributed-learning scenarios, Liu and Zhang examine quantum federated learning and demonstrate that quantum-encoded gradients may still leak client information if aggregation is not performed securely. Their findings highlight the need for quantum-aware aggregation protocols capable of protecting against both semi-honest and malicious adversaries. In quantum communication, Singh and co-authors study privacy-amplification procedures under realistic noise models and quantify the impact of

channel noise on extractable key rates, showing that practical systems must account for reduced entropy and adversarial quantum side information.

At a broader architectural level, Malina *et al.* argue that future intelligent infrastructures must embed post-quantum privacy protections throughout their communication, authentication, and data-handling layers. Their system-level analysis suggests that quantum-capable adversaries necessitate structural redesigns rather than incremental improvements. Complementing this view, Kumar and co-authors propose hybrid quantum classical privacy frameworks that combine quantum randomness sources with classical cryptographic primitives to enhance entropy and improve privacy resilience during the transition to fully post-quantum environments.

Taken together, these studies indicate that privacy in quantum settings requires rethinking underlying models, assumptions, and system mechanisms. The reviewed literature reveals persistent gaps in scalability, composability, and robustness to noise, motivating the development of more comprehensive and verifiable privacy frameworks for emerging quantum technologies.

2 Scope and Contributions

The scope of this survey is limited to peer-reviewed work published between 2020 and 2025 that investigates privacy risks and privacy-preserving approaches in quantum and post-quantum environments. The review focuses on six foundational studies covering quantum differential privacy, gentle-measurement privacy bounds, quantum federated learning, privacy amplification under noise, post-quantum infrastructure security, and hybrid quantum–classical privacy frameworks. Based on these works, the paper articulates four principal contributions.

2.1 Development of a Comprehensive Taxonomy of Quantum Privacy Threats

Privacy leakage in quantum machine learning is closely linked to gradient estimation and measurement statistics. Watkins *et al.* show that variational circuits expose sensitive gradient information during training [1]. Sequential or adaptive measurements introduce additional leakage channels; Aaronson and Rothblum demonstrate that even individually gentle measurements may accumulate non-trivial pri- vacy loss when composed [3]. In distributed environments, Liu and Zhang reveal that quantum states exchanged in federated learning can leak client-level information in both semi-honest and malicious set- tings [5]. In quantum communication, Singh *et al.* observe that realistic channel noise reduces secrecy levels and affects adversarial advantage in key extraction [6]. These findings collectively motivate a unified threat model spanning computation, communication, and distributed learning.

2.2 Detailed Review and Synthesis of Quantum-Ready Privacy-Preserving Mechanisms

Watkins *et al.* adapt differential-privacy mechanisms to variational quantum circuits by calibrating noise to quantum measurement variance [1]. Aaronson and Rothblum formalize privacy guarantees under gentle-measurement conditions and derive bounds on information leakage in quantum algorithms [3]. Liu and Zhang evaluate secure aggregation techniques for quantum federated learning and identify when classical aggregation remains privacy-preserving in quantum settings [5]. This synthesis organizes cur- rent tools and highlights gaps in the design of quantum-aware mechanisms.

2.3 Comparative Analysis of Privacy Amplification and Post-Quantum Privacy Resilience

Singh *et al.* quantify the impact of depolarizing and amplitude-damping noise on extractable key rates, showing that practical systems must employ stronger extractors to withstand quantum side information [6]. Malina *et al.* provide a broader system-level analysis of post-quantum infrastructure require-

ments, emphasizing the need to redesign communication, authentication, and storage layers to withstand quantum adversaries [2]. This contribution integrates algorithmic and architectural perspectives.

2.4 Formation of a Unified Research Agenda for the Quantum Era

Kumar *et al.* propose hybrid quantum–classical privacy architectures that combine quantum randomness with classical cryptographic primitives to enhance resilience during the transition to post-quantum systems [4]. Synthesizing insights from all six studies, the paper identifies key unresolved challenges in scalability, adversarial modeling, and composability that must be addressed for real-world deployment.

3 Background

3.1 Quantum Information Basics

Quantum information is encoded in qubits, which store data through complex probability amplitudes rather than fixed binary values. A generic qubit can be written as

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \qquad |\alpha|^2 + |\beta|^2 = 1, \tag{1}$$

with the amplitudes determining the distribution observed during measurement. Because qubits maintain superposition and phase relations, they exhibit behaviors that lack classical analogs. Measurement, however, collapses the state into a definite outcome, and this collapse can inadvertently reveal structural information about the computation, creating potential privacy concerns.

Entanglement introduces further complexity. When qubits are entangled, their measurement outcomes become correlated even when spatially separated. This enables distributed processing but also creates subtle channels for quantum side information. An adversary observing part of an entangled system may infer information about the remainder. Aaronson and Rothblum show that repeated gentle measurements can cumulatively leak information and introduce measurable state disturbance, posing privacy risks for iterative algorithms [3].

Quantum algorithms rely on unitary transformations that preserve state norm but may increase the distinguishability of intermediate states. Watkins *et al.* demonstrate that variational quantum circuits, which depend on repeated sampling to estimate gradients, reveal information through their measurement statistics, linking privacy behavior directly to circuit structure [1]. In distributed settings, Liu and Zhang show that quantum states exchanged for federated learning can leak client-specific details when provided to an aggregator [5]. Together, these observations emphasize that information flow in quantum systems is inseparable from the physics of state manipulation and measurement, making privacy a core design challenge.

3.2 Classical vs. Quantum Privacy

Classical privacy mechanisms such as differential privacy and cryptographic primitives assume deterministic data representations and predictable access patterns. In quantum settings, however, measurement collapses the state, and adversaries may exploit intermediate quantum states to extract partial information. We model privacy leakage as the information gained about a classical or quantum secret S given quantum side information ρ_E.

A Taxonomy of Quantum Privacy Threats

Quantum systems introduce privacy risks at multiple stages of data handling due to the physical behavior of qubits, measurement processes, and channel characteristics. The main categories of threats are as follows:

1. **Encoding leakage:** During classical-to-quantum encoding, amplitude and phase values may reveal structural information about input data. Variational circuits inherit this exposure because repeated sampling reveals statistics linked to encoded features, as demonstrated by leakage observed in quantum gradient estimation [1].

2. **Computation leakage:** Quantum operations generate intermediate states that can leak information when measured or reused. Aaronson and Rothblum show that even gentle measurements accumulate privacy loss when performed sequentially, enabling adversaries to infer properties of the underlying state [3]. In quantum machine learning, ancilla reuse and iterative sampling amplify this effect.

3. **Communication leakage:** Quantum communication channels are vulnerable to noise-induced privacy degradation. Singh *et al.* demonstrate that depolarizing and amplitude damping noise reduce secret key rates and increase adversarial advantage in privacy amplification [6]. In federated settings, quantum gradients transmitted between clients and an aggregator may expose local model information [5].

4. **Post-processing leakage:** After quantum computation, classical reconstruction, gradient aggregation, or model updating can expose sensitive attributes. Liu and Zhang identify post-processing leakage in quantum federated learning where classical interpretation of quantum outputs allows inference of client-specific data [5]. System-wide post-quantum infrastructures must therefore redesign storage, authentication, and logging mechanisms to limit classical leakage [2].

4 Quantum Differential Privacy (QDP)

Differential privacy (DP) provides probabilistic guarantees that a single record's inclusion has a limited effect on output. QDP generalizes DP for quantum algorithms; several definitions exist based on the adversary's access model (classical outputs vs quantum side information).

4.1 Definitions

A quantum mechanism $\mathcal{M}$ acting on states is (ε, δ)-QDP if for any two neighboring datasets D, D' and any measurement outcome set E:

$$\Pr[\mathcal{M}(D) \in E] \leq e^{\varepsilon} \Pr[\mathcal{M}(D') \in E] + \delta.$$

When the adversary holds quantum side information, definitions use trace distance between output density matrices:

$$\| \rho_D - \rho_{D'} \|_1 \leq f(\varepsilon, \delta).$$

5 Gentle Measurement and Measurement-Aware Privacy

The behavior of measurements in quantum systems plays a central role in understanding privacy leakage during computation, learning, and communication. Aaronson and Rothblum formalize how measure- ment operations that minimally disturb a quantum state can offer privacy guarantees when used within larger algorithms [3]. Their framework analyzes how much information an adversary can extract from repeated or adaptive measurements and shows that even in cases where the disturbance is small, cumu- lative leakage must be carefully controlled. This perspective is essential for quantum machine learning and protocols that rely on iterative circuit evaluations, because measurements are the primary source of observable information and therefore the primary source of privacy exposure.

5.1 Gentle Measurement Lemma

The gentle measurement lemma provides a quantitative tool for characterizing privacy preservation in quantum algorithms. For a measurement defined by POVM elements $\{M_m\}$, if an outcome occurs with probability at least $1-\varepsilon$, then the post-measurement state remains close to the original state in trace distance on the order of $\sqrt{\varepsilon}$. This result implies that high-probability measurement events cause minimal disturbance and reveal limited information about the underlying state. Aaronson and Rothblum demonstrate that when such measurements are composed sequentially, the overall disturbance grows predictably, enabling bounded privacy leakage even in long algorithmic sequences [3]. This principle has practical implications for quantum learning, where circuit evaluations often rely on repeated mea- surements of related states. It also informs the design of privacy mechanisms in distributed quantum protocols, where measurement patterns must maintain accuracy while limiting adversarial inference. Overall, gentle measurement theory establishes fundamental limits on how information flows through quantum systems and provides the basis for measurement-aware privacy techniques.

6 Privacy in Quantum Federated Learning

Federated learning enables multiple clients to collaboratively train a global model without sharing raw data. Quantum federated learning extends this paradigm to settings where clients generate, process, or encode information using quantum circuits. Liu and Zhang analyze how privacy risks emerge when quantum states or measurement results are exchanged during distributed optimization [5]. Because quantum gradients are typically represented as quantum states or sampled measurement statistics, an aggregator may infer sensitive properties about a client's local data from the structure of these shared states. Their work identifies conditions under which privacy breaches occur and provides guidelines for designing secure quantum aware aggregation protocols.

6.1 Threat Model and Assumptions

The quantum federated learning setting includes both semi honest and malicious aggregators. In the semi honest case, the aggregator follows the protocol but attempts to infer information from the received quantum states. In the malicious case, the aggregator may actively manipulate measurement settings or introduce additional queries to extract hidden structure. Clients typically compute local updates by eval- uating parameter shifted quantum circuits and encode these updates as quantum states or measurement distributions. The threat arises because these states may contain correlations that reveal client specific data features [6].

6.2 Secure Aggregation Protocols

To mitigate these risks, Liu and Zhang examine secure aggregation methods adapted from classical multi party computation. They show that classical secure aggregation can be integrated with quantum encoded gradients when clients apply masking or sharing strategies before transmission [6,4]. In such protocols, each client encodes its update into a quantum or classical representation and applies additive masks that cancel when aggregated. The aggregator learns only the sum of all updates and cannot isolate any individual contribution. The protocol remains secure under standard assumptions such as limited adversarial access and bounded quantum storage. This demonstrates that combining classical cryptographic primitives with quantum state encoding enables practical privacy guarantees in distributed quantum learning.

7. Hybrid Quantum Classical Privacy Frameworks

Kumar and his team examine combining quantum randomness with conventional cryptography. The idea is simple. Use quantum bits to generate strong, random values, then feed them into classical tools we already trust. This can make privacy stronger, especially when the hardware is still new and not very stable. And since most current machines are small and noisy, you cannot rely on them alone.
So the point of these hybrid setups is to get something useful right now. You do not need a perfect quantum computer. You just need a source of randomness that is hard to guess. But there are limits. These systems still depend on classical steps, and those steps can leak information if they are not handled well. And the quantum parts need careful testing because near term hardware can behave in odd ways. Still, this approach gives a practical middle ground while we wait for better quantum devices.

8. Comparative Evaluation

This section provides a brief comparison of the different approaches we discuss. The idea is to look at how each method handles privacy, what kind of attacker it assumes, and what it needs to work. We also assess how computationally intensive the method is and whether it is useful in practice. Some tools give strong privacy but are slow. Some are fast but need strict assumptions. And some work well only in clean and controlled settings. This comparison helps show what each approach can and cannot do.

9. Conclusion

This survey pulls together what recent work says about privacy in quantum and post quantum systems. The short version is that we know a fair amount, but we still have a long way to go. The basic ideas are clear. Quantum measurements can leak information. Noisy channels can make privacy harder. And mixed quantum classical setups can fix some problems but also bring new ones. The papers we looked at give good starting points, but most of the methods are still early and not ready for large real systems.

So the honest truth is that we need better models, better testing, and clearer rules about what threats we care about. And we need approaches that work on today's hardware instead of only in perfect lab settings. Until then, privacy in quantum systems will stay a moving target. This survey helps map out what we know and where the gaps still are.

Acknowledgments

I want to thank my professor for guiding me through this work. Their advice helped me stay on track and understand the harder parts of the topic. I also want to thank my classmates for the small discussions that helped me clear doubts along the way. And I am grateful to the reviewers for their comments, which helped me fix mistakes and improve the paper.

References

[1] W. M. Watkins, S. Y. C. Chen, and S. Yoo, "Quantum machine learning with differential privacy," *Scientific Reports*, vol. 12, no. 1, 2023. Available: https://www.nature.com/articles/s41598-022-24082-z

[2] L. Malina, P. Dzurenda, S. Ricci, and J. Hajny, "Post-quantum era privacy protection for intelligent infrastructures," *IEEE Trans. Industrial Informatics*, vol. 17, no. 11, pp. 7632–7643, 2021. Available: https://ieeexplore.ieee.org/abstract/document/9363165/

[3] S. Aaronson and G. N. Rothblum, "Gentle measurement of quantum states and differential

privacy," in *Proc. 51st Annu. ACM Symp. Theory Comput. (STOC)*, 2020, pp. 322–333. Available: https://dl.acm.org/doi/abs/10.1145/3313276.3316378

[4] Suyal, H., Singh, A., & Shrivastava, G. (2025). Privacy Preserving Efficient Worker Selection in the Cloud-Based Crowdsourcing Platform. *Internet Technology Letters*, *8*(5), e70092.

[5] P. Kumar, R. Singh, and A. Sharma, "Quantum-enhanced privacy in post-quantum cryptographic frameworks," *Quantum Information Processing*, vol. 21, 2022. Available: https://link. springer.com/article/10.1007/s41403-022-00351-8

[6] Y. Liu and X. Zhang, "Quantum privacy in federated learning environments," *Applied Sciences*, vol. 12, no. 14, pp. 6893–6905, 2022. Available: https://www.mdpi.com/2076-3417/ 12/14/6893

[7] R. Singh, S. Sharma, and T. Gupta, "Towards secure quantum communication: privacy amplification in noisy channels," *SN Applied Sciences*, vol. 3, 2021. Available: https://link.