

Quantum delegated and federated learning via quantum homomorphic encryption

Weikang Li¹  and Dong-Ling Deng^{1,2,3}

¹Center for Quantum Information, Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, China; ²Shanghai Qi Zhi Institute, Shanghai, China and ³Hefei National Laboratory, Hefei, China

Results

Cite this article: Li W and Deng D-L (2025). Quantum delegated and federated learning via quantum homomorphic encryption. *Research Directions: Quantum Technologies*. **3**, e3, 1–6. <https://doi.org/10.1017/qut.2025.2>

Received: 26 October 2024

Accepted: 20 January 2025

Keywords:

Quantum machine learning; quantum homomorphic encryption; security and privacy; federated learning

Corresponding authors:

Weikang Li;
Email: lwk20@mails.tsinghua.edu.cn;
Dong-Ling Deng;
Email: dldeng@tsinghua.edu.cn

Abstract

Quantum learning models hold the potential to bring computational advantages over the classical realm. As powerful quantum servers become available on the cloud, ensuring the protection of clients' private data becomes crucial. By incorporating quantum homomorphic encryption schemes, we present a general framework that enables quantum delegated and federated learning with a computation-theoretical data privacy guarantee. We show that learning and inference under this framework feature substantially lower communication complexity compared with schemes based on blind quantum computing. In addition, in the proposed quantum federated learning scenario, there is less computational burden on local quantum devices from the client side, since the server can operate on encrypted quantum data without extracting any information. We further prove that certain quantum speedups in supervised learning carry over to private delegated learning scenarios employing quantum kernel methods. Our results provide a valuable guide toward privacy-guaranteed quantum learning on the cloud, which may benefit future studies and security-related applications.

Introduction

Quantum machine learning exhibits a novel paradigm of learning and inference based on data (Biamonte et al., 2017; Dunjko and Briegel, 2018; Das Sarma et al., 2019; Cerezo et al., 2022), which is promising to bring advantages over classical methods for certain learning tasks (Rebentrost et al., 2014; Liu et al., 2021; Huang et al., 2022; Molteni et al., 2024). These advantages mainly focus on the computational complexity (Anshu and Arunachalam, 2024; Banchi et al., 2024), for example the running time and number of samples required to build the learning model. To obtain such quantum-versus-classical learning advantages, it is usually required to make hardness assumptions on certain computational problems (Liu et al., 2021; Gyurik and Dunjko, 2023), or utilize quantum correlations for unconditional proofs (Gao et al., 2022; Zhang et al., 2024). With a fully-fledged quantum computer featuring a large number of individually addressable and high-fidelity qubits, such learning advantages could be experimentally demonstrated. Along this line, a long-term goal is to achieve advantageous applications for practical tasks and benefit other fields (Daley et al., 2022).

However, aside from the computational aspect, the security issue is also crucial for near-term and future quantum applications (Sheng and Zhou, 2017; Liu and Jiang, 2024; Hai et al., 2024; Caro et al., 2024; Flöther, 2023). As this field progresses, the early generations of publicly available quantum computers are most likely expensive and only presented in the form of cloud quantum servers. For learning tasks, a client could upload the training data as well as the learning algorithm to the quantum server. After the server completes the learning procedure, the results will be sent back to the client for further use. In this delegated learning scenario, a natural question arises: How can one ensure that the client's data or computation is kept private from the server? Indeed, a malicious server may try to infer from the learning procedure or even disobey the instructions from the client to extract sensitive information. It is therefore of both theoretical and practical importance to develop privacy-preserving delegated learning protocols.

The exploration of private delegated quantum computations dates back to an interactive protocol (Childs, 2005), where a client, Alice, with limited quantum capabilities, delegates a task to a more powerful quantum server, Steve. Within the delegation procedure, a quantum one-time pad scheme is applied to hide the information of a quantum state (Ambainis et al., 2000). The following works along this direction are mainly divided into two categories—blind quantum computing and quantum homomorphic encryption. In blind quantum computing, the client utilizes interactive protocols to hide all information from the server, including the input, output, and computation (Broadbent et al., 2009). This framework has been developed and applied to enhance the security of quantum learning tasks (Li et al., 2021; Li, Li et al., 2024). In comparison, quantum homomorphic encryption focuses on quantum operations to be performed on encrypted data in such a way that the underlying plaintext data remains hidden from the server. During the whole delegated computation process, there is only one round of

© The Author(s), 2025. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

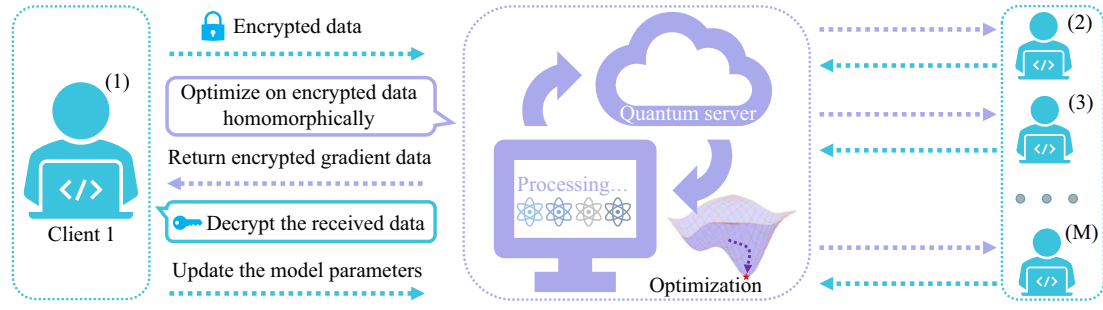


Figure 1. A schematic illustration of quantum delegated and federated learning adapting quantum homomorphic encryption techniques. On the left side, we exhibit single-client quantum delegated learning. For the training data in the form of quantum states or classical bits, the client applies a quantum or classical one-time pad to encrypt it, respectively. Upon receiving the data, the server homomorphically operates on the encrypted data and returns the encrypted results, which contain the information for model optimization, to the client. After decrypting the results, the client could then update the model parameters. On the right side, the protocol is extended to the multi-party federated learning scenario, where different clients, each holding their private data, can collaboratively train a shared model.

interaction between the two parties (Broadbent and Jeffery, 2015; Mahadev, 2020; Dulek et al., 2016; Brakerski, 2018; Ouyang et al., 2018; Tham et al., 2020; Ma and Li, 2022) and it is shown feasible to support variational quantum algorithms (Li, Quan et al., 2024). We emphasize a key difference between the two approaches. In blind quantum computing, the server is treated as a “dumb” entity that is completely unaware of the computations performed. In contrast, quantum homomorphic encryption allows the server to execute computations on encrypted quantum data, where the server knows the algorithm being applied without gaining any information about the processed data.

In this work, we leverage ideas of quantum homomorphic encryption and present a general framework for both quantum delegated learning and federated learning as illustrated in Figure 1, which further supports delegated inference after the learning process. We first construct a quantum classification model and adapt it to the quantum homomorphic encryption scenario. For the training samples in the form of general quantum states, it is required to encrypt these states before sending them to the quantum server. The non-trivial role of quantum homomorphic encryption is reflected by the fact that it allows the quantum server to manipulate the encrypted states, and further provide the desired outputs in an encrypted form. After receiving the outputs from the server, the client can efficiently recover the correct results by decryption. This feature enables us to design delegated optimization strategies in a privacy-preserving fashion. We further discuss several intriguing aspects of learning under this framework, including lower communication complexity, less demand for local computational power, higher compatibility with error correction schemes, and provable quantum speedups with kernel methods.

Framework and theoretical background

We start with the theoretical framework for quantum delegated learning with quantum homomorphic encryption. Given a general n -qubit quantum state $|\psi_x\rangle$, an information-theoretical secure encryption way is to apply the quantum one-time pad to this state:

$$|\psi_x\rangle \rightarrow (Z^{a_1} \otimes \dots \otimes Z^{a_n})(X^{b_1} \otimes \dots \otimes X^{b_n})|\psi_x\rangle, \quad (1)$$

where Z and X denote Pauli- Z and Pauli- X gates respectively, a_i and b_i are encryption keys chosen from $\{0, 1\}$ randomly and independently (Ambainis et al., 2000). For anyone without the keys a_i and b_i , this encrypted state is equivalent to a maximally mixed state and thus no information can be extracted. Since

homomorphic encryption works with encrypted data, it would be desirable if compatible schemes could be designed for the above one-time-padded data to achieve high-level security. Yet, this is challenging unless either (1) the client sends over certain quantum states which contain information about the keys and can be exploited to implement quantum operations on the server, or (2) the information-theoretical security, which is strong, is converted to a computation-theoretical one, assuming the hardness of certain problems, such as learning with errors (Broadbent and Jeffery, 2015; Mahadev, 2020; Dulek et al., 2016; Brakerski, 2018; Fisher et al., 2014; Ouyang et al., 2018; Tham et al., 2020; Ma and Li, 2022).

Here, we adapt the schemes from Refs. (Mahadev, 2020; Dulek et al., 2016; Brakerski, 2018; Ma and Li, 2022) as technical subroutines for our framework. Instead of only encrypting the quantum data, the homomorphic encryption scheme includes two parallel computations. Let $\hat{f}(\mathbf{a}, \mathbf{b})$ denote the quantum encryption operation in Equations (1) and $\text{Enc}(\mathbf{a}, \mathbf{b})$ be a classical encryption function, where \mathbf{a}, \mathbf{b} denote the classical keys a_1, \dots, a_n and b_1, \dots, b_n , respectively. For a given quantum sample $|\psi_x\rangle$, the client generates classical keys \mathbf{a}, \mathbf{b} , and sends both the encrypted state $\hat{f}(\mathbf{a}, \mathbf{b})|\psi_x\rangle$ and encrypted classical keys $\text{Enc}(\mathbf{a}, \mathbf{b})$ to the server. By applying the quantum homomorphic encryption protocol, the server homomorphically applies a target operation U to the original state after receiving these two pieces of information. The data processing can be illustrated as

$$\begin{aligned} \text{Enc}(\mathbf{a}, \mathbf{b}) &\rightarrow \text{Enc}(\mathbf{a}', \mathbf{b}'), \\ \hat{f}(\mathbf{a}, \mathbf{b})|\psi_x\rangle &\rightarrow \hat{f}(\mathbf{a}', \mathbf{b}')U|\psi_x\rangle, \end{aligned} \quad (2)$$

where the encryption key is updated to \mathbf{a}', \mathbf{b}' and the server can homomorphically compute $\text{Enc}(\mathbf{a}', \mathbf{b}')$ without access to the true values of \mathbf{a}' and \mathbf{b}' (Mahadev, 2020). Besides, the setting for handling general quantum states can be relaxed to states on the computational basis in certain scenarios, which allows a purely classical client and will be discussed in the section on quantum learning advantages with kernel methods.

Delegated learning with an untrusted server

We first consider variational quantum classifiers, which is applicable to near-term quantum devices, in our delegated learning setting (Cerezo et al., 2021). The learning model is built on a parameterized quantum circuit denoted by U_θ , where θ represents

the collection of trainable parameters. For a training sample $|\psi_x(i)\rangle$ indexed by i , the output of this model is an expectation value $\langle\psi_x(i)|U_{\theta}^{\dagger}OU_{\theta}|\psi_x(i)\rangle$ for a certain observable O . By designing appropriate cost functions to measure the distance between the current output value and the target one, an optimization procedure can be applied to capture the data pattern and update circuit parameters. Assuming the label of this sample is $y(i)$, the mean square error can be used as a cost function:

$$C_{\text{MSE}} = \frac{1}{N} \sum_i (\langle\psi_x(i)|U_{\theta}^{\dagger}OU_{\theta}|\psi_x(i)\rangle - y(i))^2, \quad (3)$$

where N denotes the size of the training set.

The delegated learning procedure begins with encrypting the training data according to Equation (2) on the client's side. To compile the variational quantum learning model in the homomorphic encryption scenario, it is worth considering the encryption protocol allowing efficient implementation of arbitrary single-qubit gates (Ma and Li, 2022). After the server homomorphically processes the data, the output quantum state becomes $\hat{f}(\mathbf{a}', \mathbf{b}')U_{\theta}|\psi_x\rangle$ with updated classical keys. In general, the prediction of the learned model is made according to the expectation value of certain local observables. Without loss of generality, we choose a two-label classification task and a Pauli-Z measurement on qubit indexed by k , that is Z_k as the prediction. The classification can be made by defining the expectation value of Z_k over zero as class 1, and otherwise as class 2. However, since the output state is encrypted by \mathbf{a}' and \mathbf{b}' , we need to decode the measured values from the server accordingly. We note that $\langle\psi_x(i)|U_{\theta}^{\dagger}\hat{f}^{\dagger}(\mathbf{a}', \mathbf{b}')Z_k\hat{f}(\mathbf{a}', \mathbf{b}')U_{\theta}|\psi_x(i)\rangle$ can be simplified to $(-1)^{b'_k}\langle\psi_x(i)|U_{\theta}^{\dagger}Z_kU_{\theta}|\psi_x(i)\rangle$, where b'_k is the k -th element in \mathbf{b}' . This indicates that the sign of the target output is also encrypted. On the server's side, only the encrypted classical keys and encrypted states are available and it is impossible to obtain the correct output. On the contrary, the client holds the decryption key to the function $\text{Enc}()$ and thus can efficiently decrypt the value of b'_k after receiving $\text{Enc}(\mathbf{a}', \mathbf{b}')$, after which the correct output value can be deciphered from the encrypted one.

The above idea has a direct application in delegated inference on a powerful quantum server which, for example, is equipped with a large-scale and fine-tuned quantum learning model. Such models may be built on sophisticated platforms and thus expensive, only deployed on the cloud and allowing clients to send computation queries remotely. For a general quantum state sample, a quantum one-time pad encrypts it first and the quantum learning model can make predictions on the encrypted data following the quantum homomorphic encryption protocol. Assume the output obtained by the server is w . With decryption keys of $\text{Enc}()$, only the client can decipher the classical keys \mathbf{a}' , \mathbf{b}' and make correct classification of this sample according to $(-1)^{b'_k}\text{sign}(w)$.

In addition to delegated inference, learning in a delegated fashion is also important in data-sensitive scenarios, for example an institution holds private health data while a server holds high computational power. This task reduces to implementing optimization under the homomorphic encryption scheme. Back to Equation (3), our goal is to optimize the variational parameters θ to minimize the cost function over the training set. During this procedure, a crucial step is to calculate the gradients of the cost function with respect to these parameters. To this end, we can apply finite difference methods or the parameter-shift rule (Mitarai et al., 2018; Schuld et al., 2019), both of which boil down

Algorithm 1. Quantum federated learning

Input The model h with parameters θ_0 , the number of clients M , the number of iterations T
Output The trained model $h(\theta^*)$
 Initialization: Generate a length- random string R ; each element in R corresponds to an index of a client
for $i \in [T]$ **do**
 1) Choose the R_i -th client and randomly choose a subset of training samples held by the client
 2) Calculate the average gradients over the selected training samples according to the protocol based on quantum homomorphic encryption
 3) Update the model parameters according to the obtained gradients $\theta_{i+1} \leftarrow \theta_i$, where strategies such as differential privacy can be adapted to reduce privacy leakage in model updates
end for
 Output the trained model $h(\theta_T)$

to measuring the expectation values of the same observable by shifting certain parameters in the quantum circuit. Under the proposed delegated learning framework, the server can flexibly shift the model parameters and obtain the desired gradients in an encrypted form. The client decrypts the gradients and uploads the updated parameters to the server, which completes an optimization iteration.

Extension to federated learning

Federated machine learning, also known as collaborative training, allows the training of a shared model across multiple parties while keeping their private training data safe (Kairouz et al., 2021). Developing techniques along this line is crucial for scenarios where the data is sensitive and each party only holds a small share insufficient for training, such as limited patients data held by different hospitals. Federated learning is traditionally achieved by aggregating model updates computed by local parties to a central server (Kairouz et al., 2021), which has been extended to the quantum domain recently (Li et al., 2021; Chen and Yoo, 2021; Du et al., 2022; Zhao, 2023; Chu et al., 2023; Ren et al., 2024; Chehimi et al., 2024). Yet, such a conventional approach would suffer from a trade-off between privacy preserving and local computational power requirement: if each party computes on their own devices and uploads the gradients to collaboratively train a model, there are high requirements for local computational power; Instead, if they upload their data directly to a central server to release local computational burden, a malicious server may violate their data privacy.

With the federated learning framework built upon quantum homomorphic encryption and delegated learning, we overcome the above trade-off and achieve both data privacy and the utilization of remote computational power. Suppose there are M independent parties, each holding their dataset with a limited size. A learning epoch proceeds by enumerating all the parties. For each party, a subset of samples are randomly selected. The samples are encrypted using the quantum one-time pad before being sent to the quantum server with the encrypted classical keys $\text{Enc}(\mathbf{a}, \mathbf{b})$. With a similar idea in delegated learning, by decrypting $\text{Enc}(\mathbf{a}', \mathbf{b}')$ to obtain the updated classical keys, the party can recover the correct gradient values of a given cost function. After taking a weighted

sum of the calculated gradients over the selected samples in this party's dataset, an update of model parameters is then uploaded to the quantum server. Such optimization can be iteratively applied until the training converges (Algorithm 2).

It is worth mentioning that, since the model parameters are publicly trained, there is inevitable private information flowing from training samples to the server. Previous works from both quantum (Li et al., 2021; Li, Kumar et al., 2024) and classical (Zhu et al., 2019) community have shown that the uploaded gradients can be utilized to infer the input data through reverse engineering. This is an intrinsic feature for training a public model. To defend against potential attacks, various strategies can be applied as additional subroutines in the original protocol. An effective way is to adapt differential privacy and add appropriate noise to the calculated gradients, which provides an information-theoretical bound on the leaked privacy (Li et al., 2021; Abadi et al., 2016). In the quantum learning setting, secure inner product estimation and secure weighted gradient summation techniques have also been explored and can be utilized to enhance privacy protection (Li, Kumar et al., 2024).

Analysis

The framework introduced for quantum delegated and federated learning through quantum homomorphic encryption bears several intriguing merits. First, this approach reduces local computational burdens, ensuring that training on edge quantum devices is not necessary for data owners. For this point, delegated learning based on blind quantum computing also provides a solution. The difference lies in the fact that, unlike blind quantum computing where the server is treated as an unknowing entity, the server in our framework is fully capable of performing computations and knowing the algorithm without access to the original data. From an information-theoretical point of view, if the client hides all the information about the data and computation from the server - as achieved in the protocols based on blind computing - then for each round of delegated learning, the client has to transfer all the information about the quantum learning model to the server which might be a huge amount even in a classical description. In contrast, here in the proposed scheme we only focus on protecting the private data information, while the server holds the model to manipulate the encrypted data. In this way, the client only needs to send the training samples to the server and participate in the optimization procedure, while there is no need to hold and transfer the model information. A direct consequence of this feature is that there is substantially less communication required between the server and clients during the learning and inference processes, minimizing overhead in a distributed network.

Error correction

Aside from the communication complexity aspect, since the server directly works on the encrypted data and is in full control of executing the quantum circuit, quantum error correction schemes can be naturally applied without involving the client (Terhal, 2015). For protocols based on blind computing subroutines, hiding all the computational details from the server becomes a double-edged sword: On the one hand, perfect privacy can be achieved in single-party delegated learning. On the other hand, however, the client does not have detailed device information on the server, and the server has no knowledge about the delegated computation. Thus, error correction becomes comparably more demanding, with additional overhead and efforts on both the client and server

for fault-tolerant delegated quantum learning (Morimae and Fujii, 2012; Gheorghiu et al., 2015; Chien et al., 2015).

Computational learning advantage

In the delegated learning scenario, designing quantum learning tasks with computational advantages is more challenging. On the one hand, transferring training data to the server or preparing initial states remotely will bring additional computational costs, and quantifying quantum speedups becomes more subtle: For example, in the well-known quantum support vector machine protocol (Rebentrost et al., 2014), by querying a data oracle, the learning model finally achieves the running time scaling as $\mathcal{O}(\log(Nd))$, where N and d denote the number of samples and dimension of the feature space, respectively. This features an exponential speedup over the classical methods. However, in the delegated learning setting, transferring such a dataset immediately brings $\mathcal{O}(Nd)$ complexity, which nullifies all the speedups. On the other hand, cryptographic methods are applied to operate on encrypted data, which brings additional overheads on the server's side. To exhibit speedups in quantum learning under our framework, we consider the case of learning with classical data (Liu et al., 2021; Sweke et al., 2021; Jerbi et al., 2024). More specifically, we consider a concept class defined based on the discrete logarithm problem (Liu et al., 2021; Gyurik and Dunjko, 2023):

Definition 1. For an n -bit prime number p with a generator a of \mathbb{Z}_p^* , we define the concept class based on the discrete logarithm problem as $\mathcal{C}_n^{\text{DLP}} = \{c_i\}_{i \in \mathbb{Z}_p^*}$, where

$$c_i(x) = \begin{cases} +1, & \text{if } \log_a x \in [i, i + \frac{p-3}{2}], \\ -1, & \text{otherwise.} \end{cases} \quad (4)$$

Assuming the hardness of the discrete logarithm problem, it is classically intractable to learn the above concept class (Liu et al., 2021). For a quantum learner, for any sample x , it can efficiently prepare the feature state $|\phi(x)\rangle = \frac{1}{\sqrt{2^k}} \sum_{i=0}^{2^k-1} |xa^i\rangle$ where $k = n - t \log n$ for a constant t (Liu et al., 2021). By estimating the inner products of these feature states, we obtain a kernel matrix that further allows a classical computer to build a linear classifier.

In our proposed framework, since the data is in the form of bit strings, the encrypted state $\hat{f}(\mathbf{a}, \mathbf{b})|\psi_x\rangle$ is also in the computational basis. Thus, we can instead send the corresponding encrypted classical bit string to the server and let the server prepare the initial state, releasing the requirement of limited quantum power for the clients. In other words, the clients can be purely classical without any quantum power in our scenario. By sending both the data and ancillary bits in encrypted forms and homomorphically applying the feature mappings, the server obtains the encrypted feature state $\frac{1}{\sqrt{2^k}} \sum_{i=0}^{2^k-1} \hat{f}(\mathbf{a}, \mathbf{b})|xa^i\rangle$. To estimate the kernel element between two samples x_1 and x_2 , we apply the same quantum one-time pad $\hat{f}(\mathbf{a}, \mathbf{b})$ so that the inner product of the two encrypted output states is exactly the desired kernel element. During this process, sending classical bit strings takes $\mathcal{O}(N^2 d / \epsilon^2)$ additional complexity with N , d , and ϵ being the number of samples, the data dimensionality, and the target additive error, respectively. Meanwhile, quantum homomorphic encryption brings a constant overhead for each gate. Thus in all, the delegated learning framework takes at most polynomial computational overheads, and the exponential quantum-classical learning separation still

holds. We mention that the above computational learning advantage is not limited to the discrete logarithm concept class. By designing proper learning problems, similar protocols may be implemented, for example for the discrete cube root problems (Gyurik and Dunjko, 2023; Kearns and Vazirani, 1994), with a final goal towards practical and real-life applications.

Conclusions and outlooks

In this work, we have developed a quantum delegated and federated learning framework based on quantum homomorphic encryption, which guarantees computation-theoretic privacy while leveraging the computational power of quantum servers. Our framework allows for training and inference on encrypted data with reduced communication complexity, making it a promising approach for future quantum cloud services. The delegation of computation to quantum servers not only removes the computational burden for clients but also preserves provable quantum speedups in certain learning tasks employing kernel methods.

This work makes a primary step forward in enabling delegated, secure, scalable, and efficient quantum machine learning systems. Future studies could focus on extending the framework to more complex learning models and improving the privacy guarantees through advanced cryptographic techniques. In addition, integrating cutting-edge classical learning models into the framework could further enhance its real-world applicability. While promising, several challenges remain to be addressed. One issue involves managing the inevitable information leakage that occurs when publicly trained models, such as federated quantum learning models, are reverse-engineered to infer private training data. Although techniques like differential privacy can mitigate such risks as discussed above, additional work is needed to ensure robust protection in other scenarios. Finally, while the current framework reduces local computational requirements to a certain degree, the quantum homomorphic encryption adds a constant yet large prefactor to the overall complexity on the server's side. This renders its experimental demonstration with current noisy intermediate-scale quantum devices unattainable. Developing more efficient protocols that could further enhance scalability and minimize computational complexity would be crucial for both near-term and future applications.

Data availability statement. Data availability is not applicable to this article as no new data were created or analyzed in this study.

Acknowledgements. We thank Qi Ye, Si Jiang, Junyu Liu, and Dong Yuan for their helpful discussions.

Author contribution. WL and DLD conceived and designed the study and wrote the article.

Financial support. This work is supported by the National Natural Science Foundation of China (Grants No. 12075128 and No. T2225008), the Innovation Program for Quantum Science and Technology (No. 2021ZD0302203), Tsinghua University Dushi Program, and the Shanghai Qi Zhi Institute Innovation Program SQZ202318.

Competing interests. The authors declare no conflicts of interest.

Ethics statement. Ethical approval and consent are not relevant to this article type.

Connections references

D'Auria V and Teller M (2023) What are the priorities and the points to be addressed by a legal framework for quantum technologies? *Research Directions: Quantum Technologies*. 1, e9. <https://doi.org/10.1017/qut.2023.3>.

References

- Abadi M, Chu A, Goodfellow I, McMahan HB, Mironov I, Talwar K and Zhang L (2016) Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS '16* (Association for Computing Machinery, New York, NY, USA, pp. 308–318.
- Ambaini A, Mosca M, Tapp A and De Wolf R (2000) Private quantum channels. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*, pp. 547–553.
- Anshu A and Arunachalam S (2024) A survey on the complexity of learning quantum states. *Nature Reviews Physics* 6(1), 59–69. <https://doi.org/10.1038/s42254-023-00662-4>.
- Banchi L, Pereira JL, Jose ST and Simeone O (2024) Statistical complexity of quantum learning. *Advanced Quantum Technologies* 2024, 2300311.
- Biamonte J, Wittek P, Pancotti N, Rebentrost P, Wiebe N and Lloyd S (2017) Quantum machine learning. *Nature* 549(7671), 195–202. <https://doi.org/10.1038/nature23474>.
- Brakerski Z (2018) Quantum FHE (almost) as secure as classical. In Shacham H and Boldyreva A (eds.), *Advances in Cryptology – CRYPTO 2018*. Cham: Springer International Publishing, pp. 67–95.
- Broadbent A, Fitzsimons J and Kashefi E (2009) Universal blind quantum computation. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pp. 517–526.
- Broadbent A and Jeffery S (2015) Quantum homomorphic encryption for circuits of low T-gate complexity. In Gennaro R and Robshaw M (eds.), *Advances in Cryptology – CRYPTO 2015*. Berlin, Heidelberg: Springer, pp. 609–629.
- Caro MC, Hinsche M, Ioannou M, Nietner A and Sweke R (2024) Classical verification of quantum learning. In *15th Innovations in Theoretical Computer Science Conference*, Schloss Dagstuhl Leibniz-Zentrum für Informatik.
- Cerezo M, Arrasmith A, Babbush R, Benjamin SC, Endo S, Fujii K, McClean JR, Mitarai K, Yuan X, Cincio L and Coles PJ (2021) Variational quantum algorithms. *Nature Reviews Physics* 3(9), 625–644. <https://doi.org/10.1038/s42254-021-00348-9>.
- Cerezo M, Verdon G, Huang H-Y, Cincio L and Coles PJ (2022) Challenges and opportunities in quantum machine learning. *Nature Computational Science* 2(9), 567–576. <https://doi.org/10.1038/s43588-022-00311-3>.
- Chehimi M, Chen SY-C, Saad W, Towsley D and Debbah M (2024) Foundations of quantum federated learning over classical and quantum networks. *IEEE Network* 38(1), 124–130. <https://doi.org/10.1109/MNET.2023.3327365>.
- Chen SY-C and Yoo S (2021) Federated quantum machine learning. *Entropy* 23(4), 460. <https://doi.org/10.3390/e23040460>.
- Chien C-H, Meter RV and Kuo S-Y (2015) Fault-tolerant operations for universal blind quantum computation. *Journal on Emerging Technologies in Computing Systems* 12(1), 1–26. <https://doi.org/10.1145/2700248>.
- Childs AM (2005) Secure assisted quantum computation. *Quantum Information & Computation* 5(6), 456–466.
- Chu C, Jiang L and Chen F (2023) CryptoQFL: quantum federated learning on encrypted data. In *IEEE International Conference on Quantum Computing and Engineering (QCE)*, vol. 01, pp. 1231–1237.
- Daley AJ, Bloch I, Kokail C, Flannigan S, Pearson N, Troyer M and Zoller P (2022) Practical quantum advantage in quantum simulation. *Nature* 607(7920), 667–676. <https://doi.org/10.1038/s41586-022-04940-6>.
- Du Y, Qian Y, Wu X and Tao D (2022) A distributed learning scheme for variational quantum algorithms. *IEEE Transactions on Quantum Engineering* 3, 1–16. <https://doi.org/10.1109/TQE.2022.3175267>.
- Dulek Y, Schaffner C and Speelman F (2016) Quantum homomorphic encryption for polynomial-sized circuits. In Robshaw M and Katz J (eds.),

- In, *Advances in Cryptology – CRYPTO 2016*. Berlin, Heidelberg: Springer, pp. 3–32.
- Dunjko V and Briegel HJ** (2018) Machine learning & artificial intelligence in the quantum domain: A review of recent progress. *Reports on Progress in Physics* **81**(7), 074001. <https://doi.org/10.1088/1361-6633/aab406>.
- Fisher KAG, Broadbent A, Shalm LK, Yan Z, Lavoie J, Prevedel R, Jennewein T and Resch KJ** (2014) Quantum computing on encrypted data. *Nature Communications* **5**, 3074.
- Flöther FF** (2023) The state of quantum computing applications in health and medicine. *Research Directions: Quantum Technologies* **1**, e10.
- Gao X, Anschuetz ER, Wang S-T, Cirac JI and Lukin MD** (2022) Enhancing generative models via quantum correlations. *Physical Review X* **12**, 021037.
- Gheorghiu A, Kashefi E and Wallden P** (2015) Robustness and device independence of verifiable blind quantum computing. *New Journal of Physics* **17**(8), 083040. <https://doi.org/10.1088/1367-2630/17/8/083040>.
- Gyurik C and Dunjko V** (2023) Exponential separations between classical and quantum learners, arXiv: 2306.16028.
- Hai R, Hung S-H, Coopmans T and Geerts F** (2024) Data management in the noisy intermediate-scale quantum era, arXiv: 2409.14111.
- Huang H-Y, Broughton M, Cotler J, Chen S, Li J, Mohseni M, Neven H, Babbush R, Kueng R, Preskill J and McClean JR** (2022) Quantum advantage in learning from experiments. *Science* **376**(6598), 1182–1186. <https://doi.org/10.1126/science.abn7293>.
- Jerbi S, Gyurik C, Marshall SC, Molteni R and Dunjko V** (2024) Shadows of quantum machine learning. *Nature Communications* **15**(1), 5676. <https://doi.org/10.1038/s41467-024-49877-8>.
- Kairouz P, McMahan HB, Avent B, et al.** (2021) Advances and open problems in federated learning *Foundations and Trends® in Machine Learning* **14**(1–2), 1–210.
- Kearns MJ and Vazirani U** (1994) *An Introduction to Computational Learning Theory*. Cambridge, MA, USA: MIT Press.
- Li C, Kumar N, Song Z, Chakrabarti S and Pistoia M** (2024) Privacy-preserving quantum federated learning via gradient hiding. *Quantum Science and Technology* **9**, 035028.
- Li C, Li B, Amer O, Shaydulin R, Chakrabarti S, Wang G, Xu H, Tang H, Schoch I, Kumar N, Lim C, Li J, Cappellaro P and Pistoia M** (2024) Blind quantum machine learning with quantum bipartite correlator. *Physical Review Letters* **133**(12), 120602. <https://doi.org/10.1103/PhysRevLett.133.120602>.
- Li Q, Quan J, Shi J, Zhang S and Li X** (2024) Secure delegated variational quantum algorithms. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **43**(10), 3129–3142. <https://doi.org/10.1109/TCAD.2024.3391690>.
- Li W, Lu S and Deng D-L** (2021) Quantum federated learning through blind quantum computing. *Science China Physics, Mechanics & Astronomy* **64**(10), 100312. <https://doi.org/10.1007/s11433-021-1753-3>.
- Liu J and Jiang L** (2024) Quantum data center: Perspectives. *IEEE Network* **38**(5), 160–166. <https://doi.org/10.1109/MNET.2024.3397836>.
- Liu Y, Arunachalam S and Temme K** (2021) A rigorous and robust quantum speed-up in supervised machine learning. *Nature Physics* **17**(9), 1013–1017. <https://doi.org/10.1038/s41567-021-01287-z>.
- Ma G and Li H** (2022) Quantum fully homomorphic encryption by integrating pauli one-time pad with quaternions. *Quantum* **6**, 866. <https://doi.org/10.22331/q.22331/q>.
- Mahadev U** (2020) Classical homomorphic encryption for quantum circuits. *Siam Journal on Computing* **52**, FOCS18.
- Mitarai K, Negoro M, Kitagawa M and Fujii K** (2018) Quantum circuit learning. *Physical Review A* **98**(3), 032309. <https://doi.org/10.1103/PhysRevA.98.032309>.
- Molteni R, Gyurik C and Dunjko V** (2024) Exponential quantum advantages in learning quantum observables from classical data, arXiv: 2405.02027.
- Morimae T and Fujii K** (2012) Blind topological measurement-based quantum computation. *Nature Communications* **3**(1), 1036. <https://doi.org/10.1038/ncomms2043>.
- Ouyang Y, Tan S-H and Fitzsimons JF** (2018) Quantum homomorphic encryption from quantum codes. *Physical Review A* **98**(4), 042334. <https://doi.org/10.1103/PhysRevA.98.042334>.
- Rebentrost P, Mohseni M and Lloyd S** (2014) Quantum support vector machine for big data classification. *Physical Review Letters* **113**(13), 130503. <https://doi.org/10.1103/PhysRevLett.113.130503>.
- Ren C, Yu H, Yan R, Xu M, Shen Y, Zhu H, Niyato D, Dong ZY and Kwek LC** (2024) *Towards Quantum Federated Learning*, arXiv: 2306.09912.
- Sarma S Das, Deng D-L and Duan L-M** (2019) Machine learning meets quantum physics. *Physics Today* **72**(3), 48–54. <https://doi.org/10.1063/PT.3.4164>.
- Schuld M, Bergholm V, Gogolin C, Izaac J and Killoran N** (2019) Evaluating analytic gradients on quantum hardware. *Physical Review A* **99**(3), 032331. <https://doi.org/10.1103/PhysRevA.99.032331>.
- Sheng Y-B and Zhou L** (2017) Distributed secure quantum machine learning. *Science Bulletin* **62**(14), 1025–1029. <https://doi.org/10.1016/j.scib.2017.06.007>.
- Sweke R, Seifert J-P, Hangleiter D and Eisert J** (2021) On the quantum versus classical learnability of discrete distributions. *Quantum* **5**, 417. <https://doi.org/10.22331/q.22331/q>.
- Terhal BM** (2015) Quantum error correction for quantum memories. *Reviews of Modern Physics* **87**(2), 307–346. <https://doi.org/10.1103/RevModPhys.87.307>.
- Tham WK, Ferretti H, Bonsma-Fisher K, Brodutch A, Sanders BC, Steinberg AM and Jeffery S** (2020) Experimental demonstration of quantum fully homomorphic encryption with application in a two-party secure protocol. *Physical Review X* **10**, 011038.
- Zhang Z, Gong W, Li W and Deng D-L** (2024) Quantum-classical separations in shallow-circuit-based learning with and without noises. *Communications Physics* **7**, 290.
- Zhao H** (2023) Non-IID quantum federated learning with one-shot communication complexity. *Quantum Machine Intelligence* **5**, 3.
- Zhu L, Liu Z and Han S** (2019) Deep leakage from gradients. In *Advances in Neural Information Processing Systems*, vol. **32**. Vancouver, Canada: Curran Associates, Inc.