
Federated Learning with Quantum Computing and Fully Homomorphic Encryption: A Novel Computing Paradigm Shift in Privacy-Preserving ML

Siddhant Dutta

DJ Sanghvi College of Engineering
siddhant.dutta180@svkmmumbai.onmicrosoft.com

Pavana P Karanth

GSSS Institute of Engineering
pavana.karanth17@gmail.com

Pedro Maciel Xavier

Davidson School of Chemical Engineering
Purdue University
pmaciex@purdue.edu

Iago Leal de Freitas

Davidson School of Chemical Engineering
Purdue University
ilealdef@purdue.edu

Nouhaila Innan

eBRAIN Lab, NYU Abu Dhabi (NYUAD)
Center for Quantum and Topological Systems
NYUAD Research Institute
nouhaila.innan@nyu.edu

Sadok Ben Yahia

The Maersk Mc-Kinney Moller Institute
University of Southern Denmark
say@mimi.sdu.dk

Muhammad Shafique

eBRAIN Lab, NYU Abu Dhabi (NYUAD)
Center for Quantum and Topological Systems
NYUAD Research Institute
muhammad.shafique@nyu.edu

David E. Bernal Neira

Davidson School of Chemical Engineering
Purdue University
dbernaln@purdue.edu

Abstract

The widespread deployment of products powered by machine learning models is raising concerns around data privacy and information security worldwide. To address this issue, Federated Learning was first proposed as a privacy-preserving alternative to conventional methods that allow multiple learning clients to share model knowledge without disclosing private data. A complementary approach known as Fully Homomorphic Encryption (FHE) is a quantum-safe cryptographic system that enables operations to be performed on encrypted weights. However, implementing such mechanisms in practice often has significant computational overhead and can expose potential security threats. Novel computing paradigms, such as analog, quantum, and specialized digital hardware, present opportunities for implementing privacy-preserving machine learning systems while enhancing security and mitigating performance loss. This work instantiates these ideas by applying the FHE scheme to a Federated Learning Neural Network architecture that integrates both classical and quantum layers.

1 Introduction

With the widespread deployment of Machine Learning (ML) applications, the level of direct human-machine interaction has increased rapidly. This surge has raised concerns and increased user awareness about the capabilities and limitations of these technologies. As the impact of Machine Learning (ML)-based gadgets becomes more relevant in the public discourse, governmental agencies begin

to develop and implement regulatory policies regarding the fair use and overall protection of user data. For example, the 2016 EU Data Regulation Act [1] and the Brazilian LGPD from 2018 [2] establish guidelines for the processing of personal data, setting the security requirements for safe information storage and delimiting the scope of use of these data. To address this issue, the Federated Learning (FL) framework was proposed [3, 4] as a mechanism for coordinating multiple independent clients that cooperate in a shared learning task by transmitting only the “knowledge” of the trained model to their peers while keeping private data stored locally. In contrast to conventional ML, where data is often centralized in a single server for model training, this distributed approach is suitable for addressing data privacy concerns. Each independent client shares their model updates with a server responsible for combining and broadcasting the aggregated model back to its clients. This allows one to benefit from the insights produced through other clients’ data without ever having direct access to it. However, these benefits are not free. The siloing of the data compromises the speed with which “knowledge” diffuses through the network, affecting the efficiency of training in the form of computational overheads and communication inefficiencies [5]. Moreover, exposing model data to potentially vulnerable communication channels between clients and the server could defeat FL’s original privacy goal. Even with the distribution of data for each client in an FL framework, the privacy of this data can be threatened by the interception of messages between the clients and the server. Once these messages are intercepted, the original private data can be inferred. To this end, an extra layer of quantum-safe privacy protection is implemented by encrypting the model updates before reaching the central server so that its aggregation operations are performed on this data without decrypting it. This technique is known as Fully Homomorphic Encryption (FHE) [6]. Implementing this technique prevents the server from accessing direct model updates, which prevents any potential intercept of the messages sent by the clients. Clients can then trust that the aggregation technique is performed without exposing their local data or the resulting learning model. Since the encryption occurs before the model updates are communicated, each client’s local ML model is not restricted, yet it is protected by FHE. This layer of protection comes at the expense of higher resource consumption to manipulate encrypted model updates [6].

New techniques must be implemented to address all the considerations of data privacy on the scale on which these FL models will be deployed. Tackling this challenge requires improvements in how efficiently each client can learn individually, reducing the number of communication rounds, and in how to encrypt the messages exchanged with a server. We consider that novel computational paradigms can be the answer to these challenges. In particular, we claim that each client can address their learning tasks with enhanced architectures that leverage this hardware. The compositional nature of deep learning models allows some of its layers to be implemented using new computing paradigms such as photonic [7], neuromorphic [8, 9] or quantum computers [10]. These machines consist of specialized hardware with promising computational speedup capabilities relevant to ML, such as matrix multiplication or calculating gradients.

One paradigm that has received particular attention in the context of privacy is *Quantum Computing* [11]. Quantum computing (QC) is the processing of information using phenomena explained through quantum mechanics. The basic information unit for QC is the quantum bit or *qubit*, which can be a superposition of the states 0 or 1. Processing over qubits subject to quantum mechanics allows one to accelerate specific computational tasks, even exponentially [10]. Quantum computers, that is, devices capable of implementing quantum computations, are still limited in size and capabilities primarily concerned with handling unintended interactions with their environment [10], which has the same effect of projecting the quantum states onto a classical one, known as *measurement*. However, they are steadily gaining traction, with the potential expectation that in the future they can surpass classical machines [12, 13] in tasks related to combinatorial optimization [14] and cryptography [15]. This opens the doors for hybrid ML algorithms that, acknowledging and accounting for the limitations of each computational paradigm, take advantage of the classical and quantum layers [16, 17].

2 Related Work

Since its inception, FL has focused on communication-efficient learning with applications to data privacy [3]. Subsequent research has expanded on this foundation, with different applications [4, 18, 19] and addressing challenges arising from practice such as training over non-IID data [20]. FL combined with FHE has gained prominence as a privacy-preserving approach to ML [21]. In particular, there has been an increase in applications in healthcare [22, 23], where preserving privacy

is crucial when working with data from medical diagnosis and imaging [24–26]. Recent work has focused on tackling the computational inefficiencies inherent in FHE [27, 28].

Another prominent area consists of adding quantum computing layers to the distributed clients’ architecture, giving rise to an extension of Quantum Machine Learning (QML) [29, 30] known as Quantum Federated Learning (QFL) [16, 31]. Special attention has been paid to QFL on quantum [32] or decentralized data [33]. Extensions of QFL considering other classical ML architectures, such as convolutional networks [34, 35], have been proposed. Applications of QFL in healthcare have been explored using existing quantum hardware [36] or classical simulations of quantum computers [23]. Other applications in finance [37] and Internet-of-Things (IoT) security [17] have been explored. There are also previous developments for integrating it with encrypted weights [38]. Finally, as a novel technology, it faces implementation and resource allocation challenges [39, 40].

3 Problem Description and Methodology

Despite its applications in distributed ML, classical FL still faces challenges. Among these are communication bottlenecks across large networks [41], privacy concerns during model updates [42], especially sensitive data such as those found in healthcare applications [23], and computational inefficiencies due to training large datasets on devices with limited resources [43]. QFL is a distributed learning paradigm capable of tackling some of these challenges. It consists of clients capable of accessing quantum computers that collaboratively train a global model while communicating with a centralized classical server. Since the local computations found in FL tend to be smaller than centralized ML datasets, it becomes a great use case for the still resource-constrained quantum devices available today [10, 44]. Furthermore, QFL is capable of using quantum-enhanced communication protocols that offer inherent privacy advantages over the classical ones [15].

In a standard neural network, the weights are iteratively updated by gradient descent. The weight update at time step $t+1$ is $W_i^{t+1} = W_i^t - \eta \nabla \mathcal{L}(W_i^t)$, where W_i^t denotes the weight at time t for client i , η is the learning rate, and $\nabla \mathcal{L}(W_i^t)$ is the loss function gradient with respect to the weight. This update minimizes the loss by adjusting the weights accordingly. However, in a scenario using FHE, the process changes since weights are encrypted on each client. FHE allows performing the same calculations on the encrypted weights $\mathcal{E}(W_i^t)$, generating new encrypted weights as $\mathcal{E}(W_i^{t+1}) = \mathcal{E}(W_i^t - \eta \nabla \mathcal{L}(W_i^t))$, leveraging FHE’s ability to perform operations directly on encrypted parameters. FL with FHE aggregates encrypted weights across multiple clients without revealing

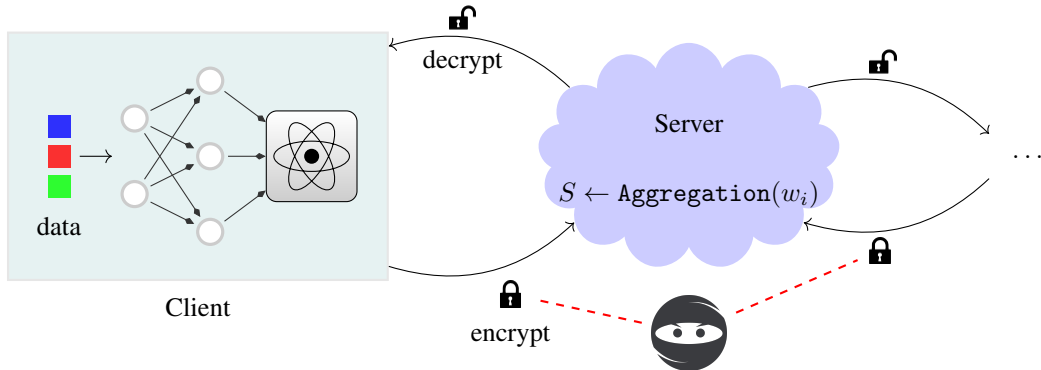


Figure 1: The client nodes utilize local data to train a model that incorporates both classical and quantum layers. After training, these models are encrypted and transmitted to a central server for aggregation into the average of all encrypted models. This new global model is then distributed to all clients. This global model is decrypted on each client, initiating another round of training. Intruders to the client-server communication would only intercept of quantum-safe encrypted models.

individual weights. The server operates on encrypted weights as $\mathcal{E}(S^{t+1}) = \sum_{i=1}^N c_i \mathcal{E}(W_i^{t+1})$, where c_i is a proportion parameter for the i -th client. Upon receiving the updated weights, the clients decrypt them and set their weights for the next model as $W_i^{t+1} = \mathcal{E}^{-1}(\mathcal{E}(S^{t+1}))$.

QFL with FHE works by training and encrypting the models for all clients in parallel. At the same time, a centralized server receives models, aggregates them, and redistributes the new model to all clients according to the procedure just described. This process repeats until convergence or meeting

any other stopping criterion. Algorithm 1 gives a pseudocode description of the procedure, and Fig. 1 shows a high-level visualization.

Algorithm 1 Quantum Federated Learning with Fully Homomorphic Encryption

```

1: Require:
2:    $ctx$ : Fully homomorphic encryption context
3:    $N$ : Number of federated clients
4:    $params$ : Encryption parameters
5:    $G$ : Quantum gate set
6:    $D$ : Parameterized Quantum Circuit (PQC) depth
Ensure: Aggregated global model  $w_g$ 
7: Initialization:
8: Generate CKKS context  $ctx \leftarrow \text{CKKSContext}(params)$ 
9: Generate Galois keys for rotations keys  $\leftarrow ctx.generate\_galois\_keys()$ 
10: Initialize global QNN model  $w_g \leftarrow \text{InitializeQNN}(D, G)$ 
11: Client-Side QNN Training and Encryption:
12: for each client  $k \in \{1, \dots, N\}$  in parallel do
13:   Prepare quantum dataset  $\mathcal{D}_k \leftarrow \text{PrepareQuantumDataset}(k)$ 
14:   Train local QNN  $w_k \leftarrow \text{TrainQNN}(\mathcal{D}_k, w_g, D, G)$ 
15:   Quantize and encrypt the local model  $w_k^{enc} \leftarrow \text{Encrypt}(\text{Quantize}(w_k), ctx)$ 
16:   Send encrypted model  $w_k^{enc}$  to the server
17: end for
18: Server-Side Aggregation:
19: Initialize  $S \leftarrow 0$  ▷ Accumulator for weighted sum
20:  $n_{total} \leftarrow \sum_{k=1}^N n_k$  ▷ Total number of samples across all clients
21: for each client  $k \in \{1, \dots, N\}$  do
22:   Receive  $w_k^{enc}$  from client  $k$ 
23:   Aggregate encrypted weights  $S \leftarrow S + w_k^{enc} \cdot \frac{n_k}{n_{total}}$ 
24: end for
25: Client-Side Decryption and Global Model Update:
26: for each client  $k \in \{1, \dots, N\}$  in parallel do
27:   Decrypt aggregated model  $w_g \leftarrow \text{Decrypt}(S, secret\_key)$ 
28:   Update global QNN model  $w_g$  on the client
29: end for
30: PQC Update:
31: Adjust PQC parameters and architecture  $w_g \leftarrow \text{OptimizePQC}(w_g, D, G)$ 
32: Model Distribution:
33: for each client  $k \in \{1, \dots, N\}$  in parallel do
34:   Send global model  $w_g$  to client  $k$ 
35: end for
36: Repeat from step 11 until maximum communication rounds
37: return  $w_g$ 

```

3.1 Quantum Neural Network Initialization and Client-Side Training

The Quantum Neural Network (QNN) is initialized by constructing a variational Parameterized Quantum Circuit (PQC) with a specified depth D and a set of quantum gates G . An example of this PQC that we used in our illustrative test cases is given in Fig. 2. This variational PQC, which will be trained on quantum data, uses parameters such as quantum gate angles and biases that can either be initialized randomly or based on pre-trained values. After initialization, each client prepares its quantum dataset \mathcal{D}_k , which could be based on local quantum measurements or preexisting datasets. Using this data, the client trains its local QNN through variational quantum algorithms. Once training is complete, the model weights are quantized to a fixed precision format to reduce the size of the encrypted parameters before encrypting the weights for secure transmission to the central server. The CKKS scheme is configured with a polynomial modulus degree of 8192, which defines the ring size $\mathbb{Z}[X]/(X^n + 1)$ with $n = 8192$, providing a security level of 128 bits. The coefficient modulus is split into four primes with bit sizes [60, 40, 40, 60], resulting in a total modulus size of 200 bits, balancing security and computational efficiency. The global scaling factor of 2^{40} ensures sufficient precision for fixed-point arithmetic.

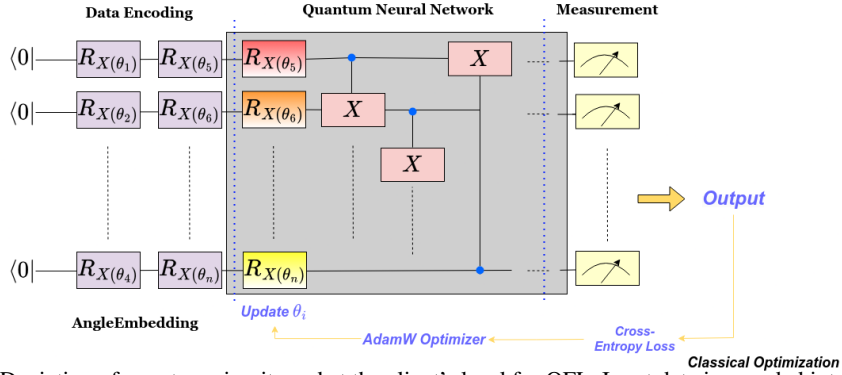


Figure 2: Depiction of quantum circuit used at the client’s level for QFL. Input data is encoded into a quantum state using angle embedding via parameterized rotation gates $R_X(\theta_i)$. The encoded quantum states are then processed by a PQC, where the weights & parameters of the PQC are encrypted using FHE. The encrypted quantum states undergo operations involving parameterized rotation gates & Controlled-NOT gates, facilitating entanglement & complex quantum state manipulation in the encrypted domain. After measurement, classical outputs are obtained, and the Cross-Entropy loss is computed. The encrypted parameters θ_i are updated during training using a classical optimizer.

3.2 Server-Side Aggregation and Global Model Update

On the server, encrypted QNN weights from all clients are aggregated using a weighted summation method, where each client’s contribution is proportional to the size of their data set. This ensures that the global model reflects each client’s training effort. Further optimizations to the PQC, such as modifying the depth D or adjusting the gate set G , may be performed to enhance model performance. After these updates, the global model is distributed back to the clients, and this iterative process continues until convergence, ultimately resulting in a privacy-preserving QNN model.

4 Computational Results

Training times for FHE-FedQNN models are notably extended due to the combined computational demands of quantum simulation and FHE. This increased duration is particularly evident with datasets like CIFAR-10, where the use of 6 qubits to represent 10 classes adds to the computational burden. Similarly, Brain MRI, which requires 4 qubits for 4 classes, and PCOS, with only 2 qubits for 2 classes, reflect varying computation times based on the number of qubits utilized. The choice of batch size is adapted to the dataset’s size. For CIFAR-10, with a substantial number of images (48k training and 12k testing), a batch size of 128 is used. In contrast, smaller datasets such as Brain MRI and PCOS, with 5.7k/1.3k and 2.56k/0.64k samples, respectively, used smaller batch sizes of 32. These adjustments help to optimize computation within the FL framework according to the size of the dataset.

Concerning Table 1, although the introduction of FHE results in computational overhead, the impact on test accuracy for FHE-FedQNN models is minimal. The difference compared to standard FedQNN models is around 1-2%, suggesting that the benefits of enhanced data security and quantum processing can outweigh this slight accuracy trade-off. Upon evaluating the FHE-FedQNN model, it was observed that there was improved performance in the PCOS dataset, resulting in a 4% gain in classification accuracy. This progress suggests that the FHE scheme could potentially assist the model in managing the noise introduced by encryption, thereby improving its generalization capabilities. Despite the increased test loss in the FHE-FedQNN model, which is likely due to noise amplification from FHE, the quantum model demonstrates superior generalizability. All models achieved near-optimal training accuracies, a typical outcome in FL settings since each client trains on a subset of data. However, test accuracy, which is measured on a separate set of test images, more accurately reflects the performance of the aggregated federated model.

5 Discussion

In the end, the more complicated architecture of FL with FHE induces a trade-off from speed to privacy. We have shown that new computing paradigms, and in particular a quantum computer, can be

Table 1: Comprehensive performance analysis of Fully Homomorphic Encryption-Enabled Federated Quantum Neural Networks (FHE-FedQNN) and Federated Neural Networks (FHE-FedNN) versus their Standard counterparts (non-FHE) on CIFAR-10, Brain MRI, and PCOS Datasets. Each model was trained over 20 Rounds with 20 clients and 10 epochs per round.

Dataset	FHE-FedQNN Models				Standard FedQNN Models			
	Train. Acc.	Test. Acc.	Test. Loss	Time (min)	Train. Acc.	Test. Acc.	Test. Loss	Time (min)
CIFAR-10 [45]	99.10%	70.12%	1.240	156.5	97.15%	72.16%	1.202	151.5
Brain MRI [46]	99.60%	88.75%	0.360	116.5	100.00%	89.71%	0.338	110.6
PCOS [47]	100%	70.15%	1.09	87.2	100%	66.19%	0.611	70.9

Dataset	FHE-FedNN Models				Standard FedNN Models			
	Train. Acc.	Test. Acc.	Test. Loss	Time (min)	Train. Acc.	Test. Acc.	Test. Loss	Time (min)
CIFAR-10 [45]	100%	68.53%	1.322	136.4	100%	71.09%	1.257	128.9
Brain MRI [46]	100%	88.4%	0.402	98.4	100.00%	90.36%	0.298	89.3
PCOS [47]	100%	64.11%	1.379	84.3	100%	65.37%	0.813	68.6

used in these ML models as a tool to accelerate local computations. The small-scale clients’ models in FL are more amenable to the limits encountered on current quantum devices. However, considering the current quantum hardware scale, this approach still has limitations. These can be mitigated in some cases by classical simulation of quantum systems via tensor networks [23, 48], although only practical until a certain scale. In the future, advances in quantum hardware, qubit error correction, and encryption techniques are expected to make QFL practical for real-world applications.

For future work, an in-detailed study of the loss flow & the gradients flow rate is necessary to provide conclusive evidence on the performance impact of QNNs integrated with FHE. This investigation will help quantify the trade-offs between encryption and model accuracy. Additionally, exploring more advanced quantum circuit designs is crucial to mitigate the issue of barren plateaus, which can hinder optimization and training in quantum neural networks. These efforts will enhance both the efficiency and scalability of FHE-enabled QNNs. All code for the results in the methodology is open source and available in the repository <https://github.com/elucidator8918/QFL-MLNCP-NeurIPS>.

6 Perspectives

FL has emerged as a viable technology for machine learning in domains where data privacy is important. Challenges related to training efficiency and vulnerability to eavesdropping have spurred a number of developments, including FHE. Leveraging the composability of current deep learning methods, some proposals have integrated classical and novel computational paradigms to satisfy the ever-growing requirements of FL applications. In particular, quantum computing has been successfully integrated with FL in this work and others [23, 31, 34, 40]. Our contribution was to show the potential of combining FHE with quantum FL and provide an implementation of these methods that is replicable on classical computers through efficient simulation of quantum circuits. The results obtained suggest that incorporating both quantum layers and FHE does not significantly increase the training time, and in some cases, it even improves the learning metrics. More importantly, it shows how new computing paradigms can already aid in relevant ML tasks.

These novel computational paradigms still have significant untapped potential. We highlight that FL can be the meeting point of two branches of quantum information sciences: quantum computing and quantum communication. To achieve exponential speedups using QML, it has been shown that one can operate directly over quantum data, without the need for encoding [33, 49]. At the same time, the advantages of quantum communication arise only when transmitting qubits. Exploring the simultaneous usage of both technologies presents a fascinating application of this technology, with federated and machine learning being the use case that requires them together.

Acknowledgment

This work was supported in part by the NYUAD Center for Quantum and Topological Systems (CQTS), funded by Tamkeen under the NYUAD Research Institute grant CG008, and the Center for Cyber Security (CCS), funded by Tamkeen under the NYUAD Research Institute Award G1104. P.M.X., I.L.d.F, and D.E.B.N graciously acknowledge the support of the Purdue Davidson School of Chemical Engineering startup grant.

References

- [1] European Union. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation). Official Journal of the European Union, 2016. <https://eur-lex.europa.eu/eli/reg/2016/679/oj>.
- [2] Brasil. Lei No. 13.709, de 2018 (Lei Geral de Proteção de Dados Pessoais). Diário Oficial da União, 2018. http://www.planalto.gov.br/ccivil_03/_Ato2015-2018/2018/Lei/L13709.htm.
- [3] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-Efficient Learning of Deep Networks from Decentralized Data. In Aarti Singh and Jerry Zhu, editors, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 1273–1282. PMLR, 20–22 Apr 2017.
- [4] Jakub Konečný, H. Brendan McMahan, Felix X. Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated learning: Strategies for improving communication efficiency, 2016.
- [5] Werner Zellinger, Volkmar Wieser, Mohit Kumar, David Brunner, Natalia Shepeleva, Rafa Gálvez, Josef Langer, Lukas Fischer, and Bernhard Moser. Beyond federated learning: On confidentiality-critical machine learning applications in industry. *Procedia Computer Science*, 180:734–743, January 2021.
- [6] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, STOC '09, page 169–178, New York, NY, USA, 2009. Association for Computing Machinery.
- [7] Kirill P. Kalinin, George Mourgias-Alexandris, Hitesh Ballani, Natalia G. Berloff, James H. Clegg, Daniel Cletheroe, Christos Gkantsidis, Istvan Haller, Vassily Lyutsarev, Francesca Parmigiani, Lucinda Pickup, and Antony Rowstron. Analog iterative machine (aim): using light to solve quadratic optimization problems with mixed variables, 2023.
- [8] C. Mead. Neuromorphic electronic systems. *Proceedings of the IEEE*, 78(10):1629–1636, 1990.
- [9] Catherine D. Schuman, Thomas E. Potok, Robert M. Patton, J. Douglas Birdwell, Mark E. Dean, Garrett S. Rose, and James S. Plank. A survey of neuromorphic computing and neural networks in hardware, 2017.
- [10] John Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79, August 2018.
- [11] Nicolas Gisin and Rob Thew. Quantum communication. *Nature Photonics*, 1(3):165–171, 2007.
- [12] Richard P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6–7):467–488, June 1982.
- [13] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O’Brien. Quantum computers. *Nature*, 464(7285):45–53, March 2010.
- [14] Amira Abbas, Andris Ambainis, Brandon Augustino, Andreas Bäertschi, Harry Buhrman, Carleton Coffrin, Giorgio Cortiana, Vedran Dunjko, Daniel J. Egger, Bruce G. Elmegreen, Nicola Franco, Filippo Fratini, Bryce Fuller, Julien Gacon, Constantin Gonciulea, Sander Gribling, Swati Gupta, Stuart Hadfield, Raoul Heese, Gerhard Kircher, Thomas Kleinert, Thorsten Koch, Georgios Korpas, Steve Lenk, Jakub Marecek, Vanio Markov, Guglielmo Mazzola, Stefano Mensa, Naeimeh Mohseni, Giacomo Nannicini, Corey O’Meara, Elena Peña Tapia, Sebastian Pokutta, Manuel Proissl, Patrick Rebentrost, Emre Sahin, Benjamin C. B. Symons, Sabine Törnøw, Victor Valls, Stefan Woerner, Mira L. Wolf-Bauwens, Jon Yard, Sheir

- Yarkoni, Dirk Zechiel, Sergiy Zhuk, and Christa Zoufal. Quantum optimization: Potential, challenges, and the path forward, 2023.
- [15] S. Pirandola, U. L. Andersen, L. Banchi, M. Berta, D. Bunandar, R. Colbeck, D. Englund, T. Gehring, C. Lupo, C. Ottaviani, J. L. Pereira, M. Razavi, J. Shamsul Shaari, M. Tomamichel, V. C. Usenko, G. Vallone, P. Villoresi, and P. Wallden. Advances in quantum cryptography. *Adv. Opt. Photon.*, 12(4):1012–1236, Dec 2020.
- [16] Samuel Yen-Chi Chen and Shinjae Yoo. Federated Quantum Machine Learning. *Entropy*, 23(4):460, April 2021. Number: 4 Publisher: Multidisciplinary Digital Publishing Institute.
- [17] Danish Javeed, Muhammad Shahid Saeed, Ijaz Ahmad, Muhammad Adil, Prabhat Kumar, and A.K.M. Najmul Islam. Quantum-empowered federated learning and 6g wireless networks for iot security: Concept, challenges and future directions. *Future Generation Computer Systems*, 160:577–597, 2024.
- [18] Li Li, Yuxi Fan, Mike Tse, and Kuo-Yi Lin. A review of applications in federated learning. *Computers & Industrial Engineering*, 149:106854, November 2020.
- [19] M. Victoria Luzón, Nuria Rodríguez-Barroso, Alberto Argente-Garrido, Daniel Jiménez-López, Jose M. Moyano, Javier Del Ser, Weiping Ding, and Francisco Herrera. A Tutorial on Federated Learning from Theory to Practice: Foundations, Software Frameworks, Exemplary Use Cases, and Selected Trends. *IEEE/CAA Journal of Automatica Sinica*, 11(4):824–850, April 2024.
- [20] Yang Liu, Tian Chen, Qing Li, Yang Qin, Haibo Yu, and Qiang Yang. Federated learning on non-iid data silos: An experimental study. *arXiv preprint arXiv:2006.14090*, 2020.
- [21] Haokun Fang and Quan Qian. Privacy preserving machine learning with homomorphic encryption and federated learning. *Future Internet*, 13(4):94, 2021.
- [22] Dinh C Nguyen, Quoc-Viet Pham, Pubudu N Pathirana, Ming Ding, Aruna Seneviratne, Zihuai Lin, Octavia Dobre, and Won-Joo Hwang. Federated learning for smart healthcare: A survey. *ACM Computing Surveys (Csur)*, 55(3):1–37, 2022.
- [23] Amandeep Singh Bhatia and David E. Bernal Neira. Federated hierarchical tensor networks: a collaborative learning quantum ai-driven framework for healthcare, 2024.
- [24] Li Zhang, Jianbo Xu, Pandi Vijayakumar, Pradip Kumar Sharma, and Uttam Ghosh. Homomorphic encryption-based privacy-preserving federated learning in iot-enabled healthcare system. *IEEE Transactions on Network Science and Engineering*, 10(5):2864–2880, 2022.
- [25] Hanchao Ku, Willy Susilo, Yudi Zhang, Wenfen Liu, and Mingwu Zhang. Privacy-preserving federated learning in medical diagnosis with homomorphic re-encryption. *Computer Standards & Interfaces*, 80:103583, 2022.
- [26] Xavier Lessage, Leandro Collier, Charles-Henry Bertrand Van Ouysel, Axel Legay, Saïd Mahmoudi, and Philippe Massonet. Secure federated learning applied to medical imaging with fully homomorphic encryption. In *2024 IEEE 3rd International Conference on AI in Cybersecurity (ICAIC)*, pages 1–12. IEEE, 2024.
- [27] Qipeng Xie, Siyang Jiang, Linshan Jiang, Yongzhi Huang, Zhihe Zhao, Salabat Khan, Wangchen Dai, Zhe Liu, and Kaishun Wu. Efficiency optimization techniques in privacy-preserving federated learning with homomorphic encryption: A brief survey. *IEEE Internet of Things Journal*, 11(14):24569–24580, 2024.
- [28] Sakib Anwar Rieyan, Md Raisul Kabir News, ABM Muntasir Rahman, Sadia Afrin Khan, Sultan Tasneem Jawad Zaarif, Md Golam Rabiul Alam, Mohammad Mehedi Hassan, Michele Ianni, and Giancarlo Fortino. An advanced data fabric architecture leveraging homomorphic encryption and federated learning. *Information Fusion*, 102:102004, 2024.
- [29] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. Quantum machine learning. *Nature*, 549:195–202, 2017.
- [30] Patrick Rebentrost, Masoud Mohseni, and Seth Lloyd. Quantum support vector machine for big data classification. *Physical Review Letters*, 113:130503, 2014.
- [31] Nouhaila Innan, Muhammad Al-Zafar Khan, Alberto Marchisio, Muhammad Shafique, and Mohamed Bennai. FedQNN: Federated learning using quantum neural networks. In *2024 International Joint Conference on Neural Networks (IJCNN)*, pages 1–9, 2024.

- [32] Mahdi Chehimi and Walid Saad. Quantum federated learning with quantum data. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8627–8631. IEEE, 2022.
- [33] Hsin-Yuan Huang, Michael Broughton, Jordan Cotler, Sitan Chen, Jerry Li, Masoud Mohseni, Hartmut Neven, Ryan Babbush, Richard Kueng, John Preskill, et al. Quantum advantage in learning from experiments. *Science*, 376(6598):1182–1186, 2022.
- [34] Amandeep Singh Bhatia, Sabre Kais, and Muhammad Ashraful Alam. Federated quantum neural network: a new paradigm for collaborative quantum learning. *Quantum Science and Technology*, 8(4):045032, 2023.
- [35] Sabre Kais, Amandeep Bhatia, and Muhammad Alam. Quantum federated learning in healthcare: The shift from development to deployment and from models to data, 03 2023.
- [36] Luca Lusnig, Asel Saginalieva, Mikhail Surmach, Tatjana Protasevich, Ovidiu Michiu, Joseph McLoughlin, Christopher Mansell, Graziano de’Petris, Deborah Bonazza, Fabrizio Zanconati, et al. Hybrid quantum image classification and federated learning for hepatic steatosis diagnosis. *Diagnostics*, 14(5):558, 2024.
- [37] Nouhaila Innan, Alberto Marchisio, Muhammad Shafique, and Mohamed Bennai. QFNN-FFD: Quantum federated neural network for financial fraud detection. *arXiv preprint arXiv:2404.02595*, 2024.
- [38] Cheng Chu, Lei Jiang, and Fan Chen. CryptoQFL: quantum federated learning on encrypted data. In *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, volume 1, pages 1231–1237. IEEE, 2023.
- [39] Harashta Tatimma Larasati, Muhammad Firdaus, and Howon Kim. Quantum federated learning: Remarks and challenges. In *2022 IEEE 9th International Conference on Cyber Security and Cloud Computing (CSCloud)/2022 IEEE 8th International Conference on Edge Computing and Scalable Cloud (EdgeCom)*, pages 1–5, 2022.
- [40] Chao Ren, Han Yu, Rudai Yan, Minrui Xu, Yuan Shen, Huihui Zhu, Dusit Niyato, Zhao Yang Dong, and Leong Chuan Kwek. Towards quantum federated learning. *arXiv preprint arXiv:2306.09912*, 2023.
- [41] Osama Shahid, Seyedamin Pouriyeh, Reza M. Parizi, Quan Z. Sheng, Gautam Srivastava, and Liang Zhao. Communication efficiency in federated learning: Achievements and challenges. *CoRR*, abs/2107.10996, 2021.
- [42] Viraaji Mothukuri, Reza M. Parizi, Seyedamin Pouriyeh, Yan Huang, Ali Dehghantanha, and Gautam Srivastava. A survey on security and privacy of federated learning. *Future Generation Computer Systems*, 115:619–640, 2021.
- [43] Muhammad Umair, Wooi-Haw Tan, and Yee-Loo Foo. Challenges in federated learning for resource-constrained iot environments: Energy efficiency, privacy, and statistical heterogeneity. In *2023 IEEE 8th International Conference on Recent Advances and Innovations in Engineering (ICRAIE)*, pages 1–6, 2023.
- [44] Nature Machine Intelligence. Seeking a quantum advantage for machine learning. *Nature Machine Intelligence*, 5:813, 2023.
- [45] Alex Krizhevsky. Learning multiple layers of features from tiny images, 2009.
- [46] Msoud Nickparvar. Brain tumor MRI dataset, 2021.
- [47] Misa Hub, Palak Handa, Anushka Saini, Siddhant Dutta, Harsh Pathak, Nishi Choudhary, Nidhi Goel, and Jasdeep Kaur Dhanoo. Auto-pcos classification challenge, January 2024.
- [48] Román Orús. Tensor networks for complex quantum systems. *Nature Reviews Physics*, 1(9):538–550, August 2019.
- [49] Hsin-Yuan Huang, Michael Broughton, Masoud Mohseni, Ryan Babbush, Sergio Boixo, Hartmut Neven, and Jarrod R McClean. Power of data in quantum machine learning. *Nature communications*, 12(1):2631, 2021.