Quantum Federated Learning for Image Classification

Leo Sünkel, Philipp Altmann, Michael Kölle and Thomas Gabor
Institute for Informatics, LMU Munich, Germany

Keywords: Federated Learning, Quantum Machine Learning, Distributed Learning, Quantum Networks.

Abstract: Federated learning is a technique in classical machine learning in which a global model is collectively trained by a number of independent clients, each with their own datasets. Using this learning method, clients are not required to reveal their dataset as it remains local; clients may only exchange parameters with each other. As the interest in quantum computing and especially quantum machine learning is steadily increasing, more concepts and approaches based on classical machine learning principles are being applied to the respective counterparts in the quantum domain. Thus, the idea behind federated learning has been transferred to the quantum realm in recent years. In this paper, we evaluate a straightforward approach to quantum federated learning using the widely used MNIST dataset. In this approach, we replace a classical neural network with a variational quantum circuit, i.e., the global model as well as the clients are trainable quantum circuits. We run three different experiments which differ in number of clients and data-subsets used. Our results demonstrate that basic principles of federated learning can be applied to the quantum domain while still achieving acceptable results. However, they also illustrate that further research is required for scenarios with increasing number of clients.

1 INTRODUCTION

Privacy in the context of machine learning models has become a greater concern in recent years and one approach to enhance the privacy of user data is *federated learning* (McMahan et al., 2017). Using this technique, a global model (e.g., a neural network) is collectively trained by a number of client models. Clients do not reveal their data; instead they synchronize and train the global model by other means (for example by aggregating parameters or weights) while their individual dataset is kept local, i.e., private. The concept of federated learning is not new in classical machine learning, the basic concept has been thoroughly discussed (Konečnỳ et al., 2016; McMahan et al., 2017; Zhao et al., 2018; Yang et al., 2019) as well as a range of problems, challenges and potential solutions identified and proposed (Geyer et al., 2017; Li et al., 2020; Mammen, 2021; Lyu et al., 2020). While the field is still advancing in the classical domain, interest in extending these ideas into the quantum realm has been steadily growing in recent years (Chen and Yoo, 2021; Li et al., 2021; Chehimi and Saad, 2022; Kumar et al., 2023). This allows for the rise of several possible architectures and approaches, each with their own set of challenges, problems and potential advantages. For instance, clients could still communicate via classical networks, however, it would also be feasible to incorporate the methodology into a quantum communication network (Chehimi et al., 2023), which could provide certain benefits such as secure quantum communication channels. However, the specifics require further investigation as both *quantum federated learning* and quantum communication networks are still in their infancy.

Whether quantum federated learning is able to provide advantages besides privacy or security is another open question demanding more research. For instance, how does the approach effect the trainability of a quantum circuit, or how do both approaches compare in terms of their ability to generalize to unseen data? These questions are already relevant and subject to research for traditional training and learning approaches, so investigating these questions in a federated learning setting will most likely also be important.

In this paper, we evaluate a simple quantum federated learning approach on an image classification problem, namely the task of recognizing images of digits from the MNIST dataset. We divide the dataset such that each subset contains images for two classes. A global model is trained in a federated manner where each client is a variational quantum circuit used for binary classification. We compare this approach to a

Sünkel, L., Altmann, P., Kölle, M. and Gabor, T.
Quantum Federated Learning for Image Classification.
DOI: 10.5220/0012421200003636

In *Proceedings of the 16th International Conference on Agents and Artificial Intelligence (ICAART 2024) - Volume 3*, pages 936-942
ISBN: 978-989-758-680-4; ISSN: 2184-433X

regular variational quantum circuit trained in a traditional, i.e., non federated or distributed manner.

This paper is structured as follows. In Section 2 we discuss the background of quantum machine learning and (quantum) federated learning. Related work is discussed in Section 3 while we present our experimental setup in Section 4. Results are presented in Section 5. We conclude and give an outlook for future work in Section 6.

2 BACKGROUND

In this section, we briefly recap the fundamentals of quantum machine learning (QML) and discuss federated learning (FL) as well as quantum federated learning (QFL).

2.1 Quantum Machine Learning

Over the years several QML algorithms have been proposed by the research community, however, the so-called variational quantum algorithm (VQA) approach based on variational quantum circuits (VQCs) (Mitarai et al., 2018; Schuld and Killoran, 2019) appears to be the most popular and relevant one in the current NISQ-era of quantum computing (QC), and this approach is the one we employ as part of this work. Thus, when we use the term QML we refer to this approach.

A VQC is a parameterized quantum circuit consisting of the following parts: *(i) feature map*, *(ii) entanglement, and rotation* and *(iii) measurement*. In (i) the classical data, i.e., features, are encoded into a quantum state through the use of rotation gates where a feature corresponds to the angle of rotation. This is followed by a series of repeating layers consisting of parameterized rotations and entangling gates (e.g., CNOT gates), where the parameters of the rotation gates are the weights to be optimized by a classical optimization algorithm. In the last step a number of qubits are measured, resulting in classical values which can in turn be used to derived the prediction for a classification task. An example VQC with 1 layer is depicted in Figure 3. The design of the architecture of the circuit is its own research question and we consider circuits of the basic architecture described above in this paper.

The optimization algorithm runs on a classical computer, making this an hybrid iterative approach. The circuit is initialized with features and weights, executes on a quantum computer and returns some measurement results that are interpreted in order to establish a prediction, which can then be fed into the optimizer resulting in a set of new updated weights. The process continues until the preset number of iterations have been reached or some other termination criteria is met. For a more in depth discussion of this topic we refer to (Mitarai et al., 2018) and (Schuld et al., 2020).

2.2 (Quantum) Federated Learning

We will summarize the main idea behind FL in this section and refer for a more in depth discussion to (Yang et al., 2019). After establishing FL in the classical setting, we discuss how to transfer these concepts to the QC domain.

In a FL setting, a global model is collectively trained by a number of client-models that have some means of communication, for example over a communication network. The clients itself may be trained on different data-subsets or even entirely different datasets. The data, however, is kept private, i.e., local; the clients only exchange parameters or gradients with the global model. Note that the details of the exchange depends on the implementation, there exist various techniques and approaches in the literature. One of the main advantages and motivation is the possibility of collaboratively training a shared model by multiple, potentially unknown clients while still keeping sensitive data private.

The main approach is as follows. The global model distributes its parameters (i.e., weights) to each client. Then a client trains its model on its local dataset for a number of epochs. After a defined number of epochs, clients send their weights to the global model which then aggregates all weights and updates its own weights. The global model then distributes the updated weights among the clients and the process repeats until the maximal number of epochs has been reached. The overall architecture of FL is illustrated in Figure 1.

A straightforward approach to transfer these ideas to the quantum realm would be to use a VQC as the global model and a number of VQCs as client models. This is similar to the approach in (Chen and Yoo, 2021), however, they use a hybrid model. These VQCs may or may not have the same architecture. The models weights can be exchanged over classical communication channels. However, quantum communication networks allow the transfer of quantum states and could also be incorporated into the approach, yielding further possible advantages such as secure quantum communication channels. Further extension to include *blind quantum computing* is a another possible pathway to ensure privacy, as discussed by (Li et al., 2021).

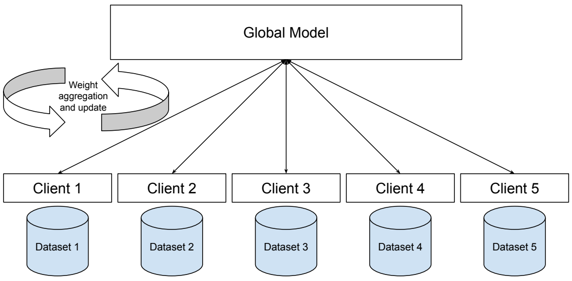

Figure 1: Overview of an example FL architecture. In this example, 5 clients train their model on their local dataset and send their updated weights to the global model. The global model aggregates all weights collected and updates the global model which is then distributed among the clients for the next iteration of training.

3 RELATED WORK

A federated QML approach based on a hybrid quantum models is discussed in (Chen and Yoo, 2021). The authors evaluate their approach on binary classification tasks and show that their approach yields similar results than regular training. Incorporating blind quantum computing is discussed in (Li et al., 2021). They show, among other things, the training of a VQC using blind quantum computing in the context of FL. Privacy in QFL is discussed in (Kumar et al., 2023) and (Rofougaran et al., 2023) while (Chehimi and Saad, 2022) propose a quantum federated learning framework. In (Wang et al., 2023) the authors discuss quantum federated learning over quantum networks, where the weights are communicated via teleportation. They evaluate their approach on a binary classification task. An overview of challenges and opportunities is given in (Chehimi et al., 2023).

4 EXPERIMENTAL SETUP

We discuss our approach to QFL and experimental setup in this section. In our experiments, the global model as well as all clients are VQCs, each with the same circuit design. However, the specific qubits measured vary, as discussed in detail in the following sections. Each VQC is its own model with its own optimizer and set of parameters. Each client is trained on its own data sub-set, we will discuss this point in more detail in Section 5.

4.1 Approach

Our approach revolves around the Clients training their model on their own dataset and update their local parameters independently of each other. Every n epochs, the parameters of every client are aggregated and used update the global model. The updated parameters are determined by calculating the mean over all clients. The parameters of the updated global model are then used to also update the clients parameters, i.e., all clients are synchronized with the global model. The overall process for the experiment with 5 clients and subsets is depicted in Figure 2 while the VQC architecture employed is shown in Figure 3 and is discussed below.

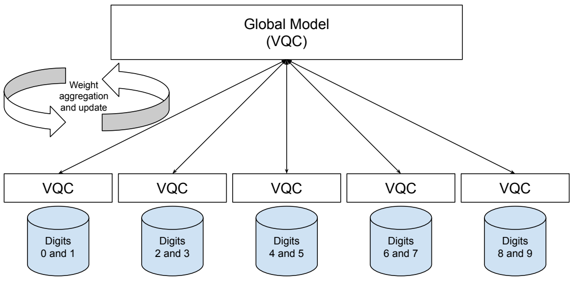

Figure 2: Example QFL architecture.

Note that we only exchange models parameters and a classical communication channel is sufficient for this. More specifically, a network is not necessary either; the approach can be executed entirely local, as is done in our experiments. Incorporating the approach into a quantum network is out of scope and is a potential avenue for future work.

We evaluate our approach for a classification task using the MNIST dataset, which contains images of hand-written digits. Each client is given a subset of this dataset containing the images for two digits and each client is given a distinct dataset. In our experiments, we train several clients where the first one uses digits 0 and 1, the second one 2 and 3 and so forth.

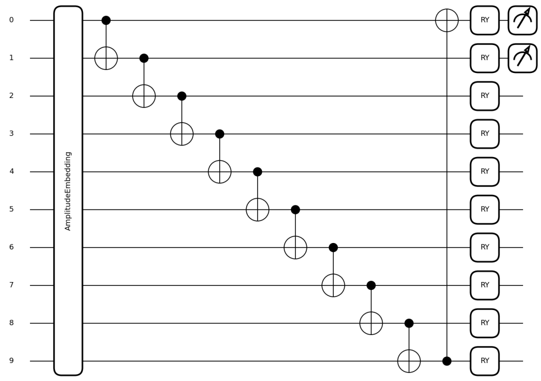

Figure 3: Example circuit architecture employed in our experiments with depth 1. Note that in our experiments we used circuits with a higher depth, however, the same overall architecture pattern.

Note that each VQC is used for binary classification while the global model is evaluated on all classes after training has completed.

4.2 Variational Quantum Circuit

All models in our experiments are VQCs with the same architecture. All features are encoded using amplitude embedding, this is followed by a simple ansatz consisting of repeating layers of CNOT and parameterized rotation gates. We measure two qubits in the binary classification case. In each VQC, the digits it classifies also correspond to the index of the qubit that is measured. That is, the VQC trained on the subset consisting of images of digits 0 and 1 measures qubits 0 and 1 while the VQC trained on digits 2 and 3 measures qubits 2 and 3. The measurement results are interpreted as class probabilities and used for label prediction. The number of layers is determined by the depth parameter, the parameters used in our experiments are discussed in the next section.

4.3 Data and Experiments

We briefly review the configuration used in our experiments while an overview of the parameters is given in Table 1. We used the MNIST dataset in our training, which contains images of size 28x28, resulting in 784 features for each image. With amplitude embedding, a circuit with 10 qubits is sufficient to embed all features. Note that we do not use all training data in our experiments, instead roughly 1000 samples were used per epoch for each model. Pytorch and PennyLane (Bergholm et al., 2018) were used to implement our experiments. Using a circuit with 10 qubits and a depth of 15 results in 150 trainable parameters in our approach, as we use a single *Ry* rotation per qubit in each layer. We evaluate our approach in three different experiments. In the first, a global model is trained by two clients with their own respective data subset for binary classification. For instance, one client is trained on a data subset consisting of images depicting the digits 0 and 1 while the other clients subset contains digits 2 and 3, thus the global model is trained on four classes in total. In the second experiment we use three clients, the additional client is then trained on digits 4 and 5 while in the third experiment we use all 10 digits and train 5 clients.

5 RESULTS

Table 1: Experiment configuration.

Parameter	
Qubits	10
Depth	15
Parameters	150
Epochs	30
Batch size	40
Models synchronized every n generations	3
Seeds	4
Clients	2, 3 or 5

In this section, we discuss the results of experiments conducted as part of this work. We first present the results from the baseline model and then continue with the QFL approach and conclude with a comparison of both.

5.1 Baseline

As baseline we use VQCs for binary classification, each trained on a subset of the data. More specifically, one baseline model is trained to classify the digits 0 and 1, another to the digits 2 and 3 and so forth. Note that each model is trained individually, i.e., not in a federated or distributed manner, and uses the same hyper-parameters as in the QFL experiments. The mean training accuracy for each of this models is depicted in Figure 4 and the loss is shown in Figure 5, both aggregated over all seeds. Test results of the VQC trained on all classes are also discussed below.

5.2 Quantum Federated Learning

We discuss the results from our QFL experiments next. However, first a note on how the results are aggregated. Recall that client models are trained on their subset of the training data, weights are aggregated to form the global model which is subsequently distributed to synchronize the clients. The training results of the QFL approach depicted below are the

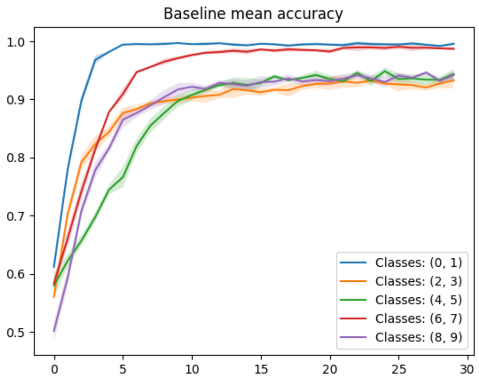

Figure 4: Comparison of training accuracy of binary classification from the baseline-model.

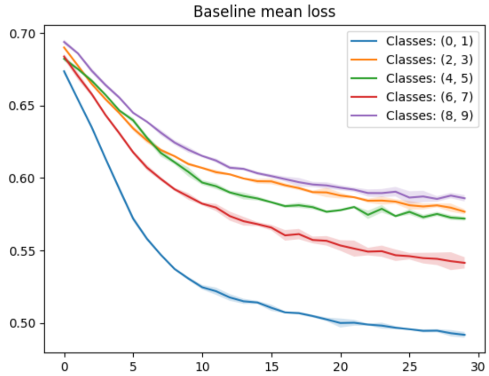

Figure 5: Comparison of training loss of binary classification from the baseline-model.

mean of all client models, aggregated over all seeds. However, we also depict and discuss the test results from the global model.

Figure 6 shows the mean training accuracy achieved with the QFL approach while the loss is depicted in Figure 7. In both figures (i.e., accuracy and loss plots) it can be seen that the accuracy and loss ”jump” every few epochs. This might be due to fact that every n generations (parameter value is shown in Table 1), the global model is updated and distributed. Adjusting this parameter may improve the models performance.

The mean training accuracy of the baseline and global model for 4 digits is shown in Figure 8 and for 6 digits is shown in Figure 9. Also not that here the results of binary classification for each subset are aggregated for both models.

In Figure 11 the test accuracy aggregated from the performance on each subset from the global model is shown. That is, the figure depicts the mean test accuracy for each binary classification task (i.e., 0 vs 1, 2 vs 3 and so forth) from the global model. In Figure 12 the test results of the global model on all classes is depicted.

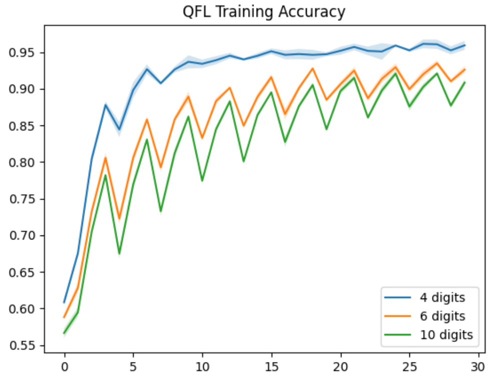

Figure 6: Mean train accuracy (QFL).

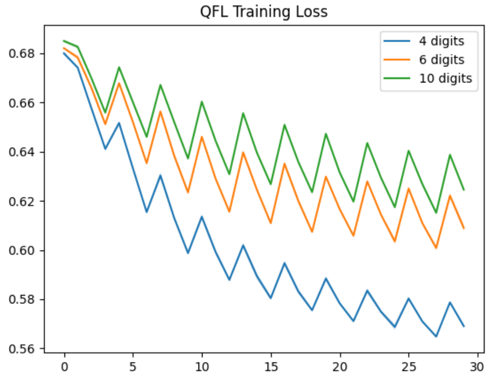

Figure 7: Mean train loss (QFL).

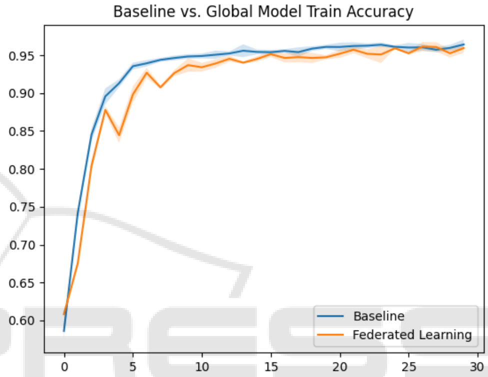

Figure 8: Comparison of aggregated mean training accuracy between baseline and QFL for 4 digits.

5.3 Discussion

From these results one can see that while the baseline achieves the best performance, the QFL approach also seems to achieve acceptable results, especially in the

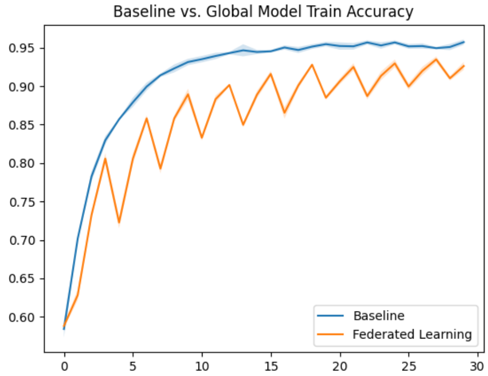

Figure 9: Comparison of aggregated mean training accuracy between baseline and QFL for 6 digits.

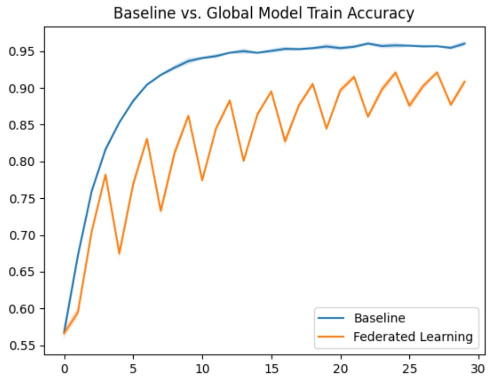

Figure 10: Comparison of aggregated mean training accuracy between baseline and QFL for 10 digits.

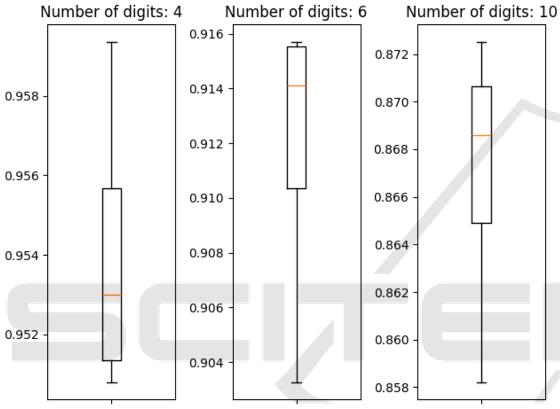

Figure 11: Mean test accuracy over all subsets (Global model.)

experiment with 2 clients and data-subsets. However, the more clients and data-subsets are added, the discrepancy increases, compare Figures 8 and 10. This is also illustrated in Figures 12 and 13. Increasing the number of epochs or adjusting other hyper-parameters may improve the performance though. Also recall that in all experiments we limited the number of training samples per epoch to roughly 1000, allowing more samples per epoch may also yield better results. The main aim of this study was to evaluate the potential of QFL on a simple example with a straightforward approach by adopting methods from classical FL. Incorporating more advanced methods from FL or investigating new techniques suitable for the quantum domain may be required to achieve superior performance and results.

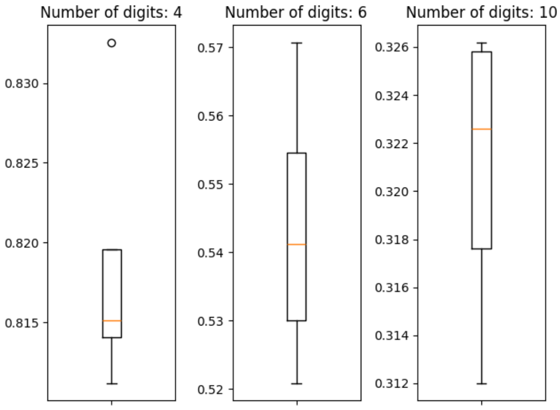

Figure 12: Test accuracy of global model on all classes.

6 CONCLUSION

In this paper, we discussed and evaluated a basic approach that transfers concepts from classical FL into the domain of quantum computing. We applied an approach in which the global model as well as the clients are VQCs that transfer their parameters (i.e., weights) in a classical manner. We ran three experiments in the domain of image classification on the MNIST dataset. The nature of the experiments varied in terms of number of clients and data-subsets, that is, we used 2, 3, and 5 clients in different experiments to train a global model. Overall these experiments yield acceptable results, however, as the number of clients increase (and thus the number of data-subsets), the performance drops steadily. There could be numerous reasons and countermeasures for this. For instance, increasing the number of training epochs as well as the training samples may increase the performance. Using different circuits with more parameters may also help. Though more research is required in

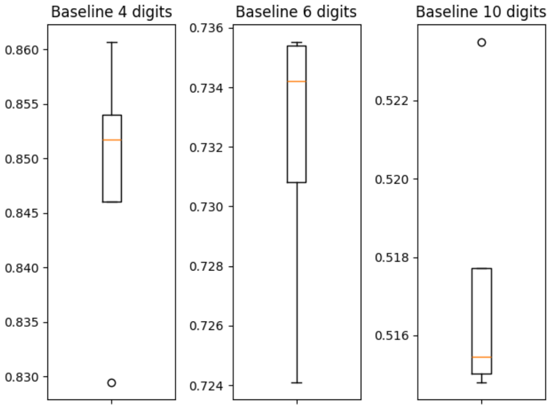

Figure 13: Test accuracy of the baseline on all classes.

this direction.

Incorporating QFML approaches that use quantum communication or quantum networks is a further important research direction. This is for instance discussed in (Chehimi et al., 2023) and (Wang et al., 2023). As the aim in this paper was to evaluate a straightforward approach, in future more elaborate schemes as well as different domains and use-cases in future communication networks should be explored.

ACKNOWLEDGEMENTS

Sponsored in part by the Bavarian Ministry of Economic Affairs, Regional Development and Energy as part of the 6GQT project.

REFERENCES

Bergholm, V., Izaac, J., Schuld, M., Gogolin, C., Ahmed, S., Ajith, V., Alam, M. S., Alonso-Linaje, G., Akash-Narayanan, B., Asadi, A., et al. (2018). Pennylane: Automatic differentiation of hybrid quantum-classical computations. *arXiv preprint arXiv:1811.04968*.

Chehimi, M., Chen, S. Y.-C., Saad, W., Towsley, D., and Debbah, M. (2023). Foundations of quantum federated learning over classical and quantum networks. *IEEE Network*.

Chehimi, M. and Saad, W. (2022). Quantum federated learning with quantum data. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8617–8621. IEEE.

Chen, S. Y.-C. and Yoo, S. (2021). Federated quantum machine learning. *Entropy*, 23(4):460.

Geyer, R. C., Klein, T., and Nabi, M. (2017). Differentially private federated learning: A client level perspective. *arXiv preprint arXiv:1712.07557*.

Konečnỳ, J., McMahan, H. B., Yu, F. X., Richtárik, P., Suresh, A. T., and Bacon, D. (2016). Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*.

Kumar, N., Heredge, J., Li, C., Eloul, S., Sureshbabu, S. H., and Pistoia, M. (2023). Expressive variational quantum circuits provide inherent privacy in federated learning. *arXiv preprint arXiv:2309.13002*.

Li, T., Sahu, A. K., Talwalkar, A., and Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE signal processing magazine*, 37(3):50–60.

Li, W., Lu, S., and Deng, D.-L. (2021). Quantum federated learning through blind quantum computing. *Science China Physics, Mechanics & Astronomy*, 64(10):100312.

Lyu, L., Yu, H., and Yang, Q. (2020). Threats to federated learning: A survey. *arXiv preprint arXiv:2003.02133*.

Mammen, P. M. (2021). Federated learning: Opportunities and challenges. *arXiv preprint arXiv:2101.05428*.

McMahan, B., Moore, E., Ramage, D., Hampson, S., and y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR.

Mitarai, K., Negoro, M., Kitagawa, M., and Fujii, K. (2018). Quantum circuit learning. *Physical Review A*, 98(3):032309.

Rofougaran, R., Yoo, S., Tseng, H.-H., and Chen, S. Y.-C. (2023). Federated quantum machine learning with differential privacy. *arXiv preprint arXiv:2310.06973*.

Schuld, M., Bocharov, A., Svore, K. M., and Wiebe, N. (2020). Circuit-centric quantum classifiers. *Physical Review A*, 101(3):032308.

Schuld, M. and Killoran, N. (2019). Quantum machine learning in feature hilbert spaces. *Physical review letters*, 122(4):040504.

Wang, T., Tseng, H.-H., and Yoo, S. (2023). Quantum federated learning with quantum networks. *arXiv preprint arXiv:2310.15084*.

Yang, Q., Liu, Y., Chen, T., and Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2):1–19.

Zhao, Y., Li, M., Lai, L., Suda, N., Civin, D., and Chandra, V. (2018). Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*.