↑
view our
paper here

Quantum Privacy Aggregation of Teacher Ensembles (Q-PATE) for Privacy Preserving Quantum Machine Learning

William Watkins[1], Heehwan Wang[2], Sangyoon Bae[2], Huan-Hsin Tseng[4], Jiook Cha[2], Samuel Yen-Chi Chen[3], Shinjae Yoo[4]

Johns Hopkins University[1], Seoul National University[2], Wells Fargo[3], Brookhaven National Lab[4]

↑
view our
paper here

Introduction

Differential Privacy (DP)

- A mathematical way to protect individuals when their data is used in data sets.
- Seeks to address privacy concerns through the privacy-loss framework.
- DP has two hyperparameters
  - ε : the privacy budget (=privacy loss *L*, explicates the differences in the distributions characterized by two similar queries) to the system
  - δ : probability of leaking more information than allowed by the privacy budget (=privacy cutoff)

Quantum Machine Learning (QML)

- Quantum machine learning implement machine learning algorithms by utilizing quantum computing via quantum circuit.
- Quantum computing utilizes qubits, which can exist in multiple states simultaneously due to the principles of superposition and entanglement.
- As quantum circuits are *differentiable*, and a quantum computer itself can compute the change in control learnable parameters.

Method

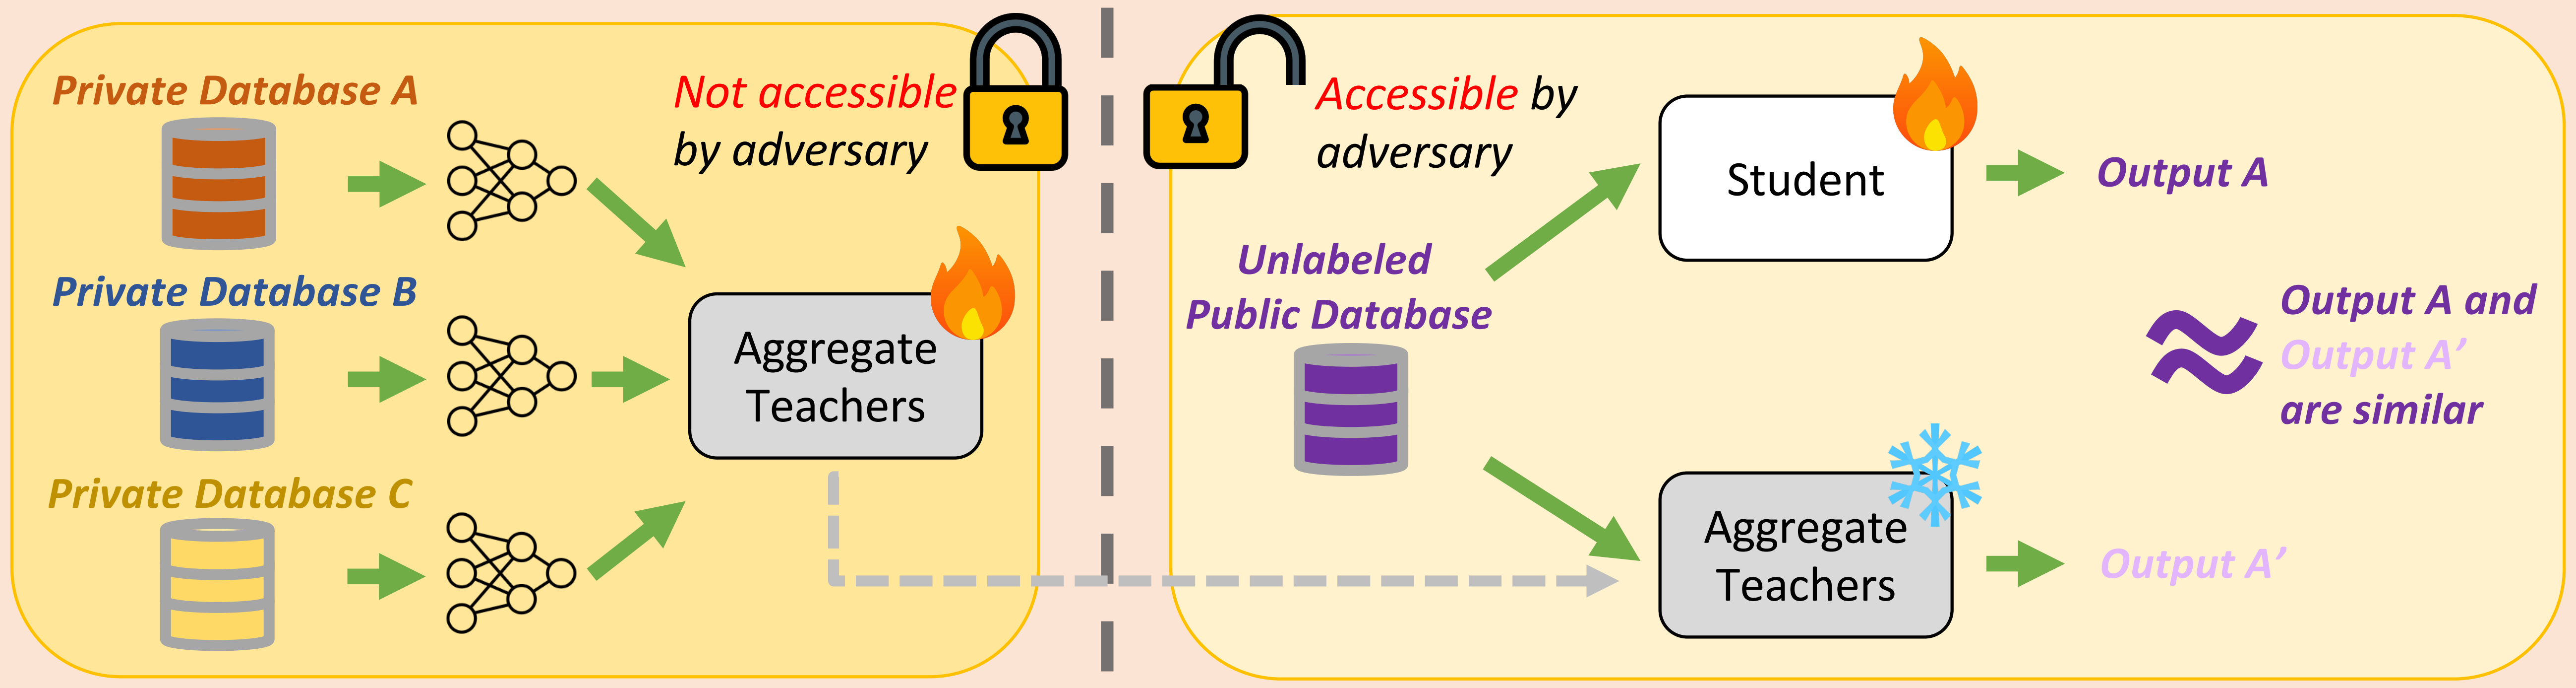

Figure 1. Schematic overview of PATE

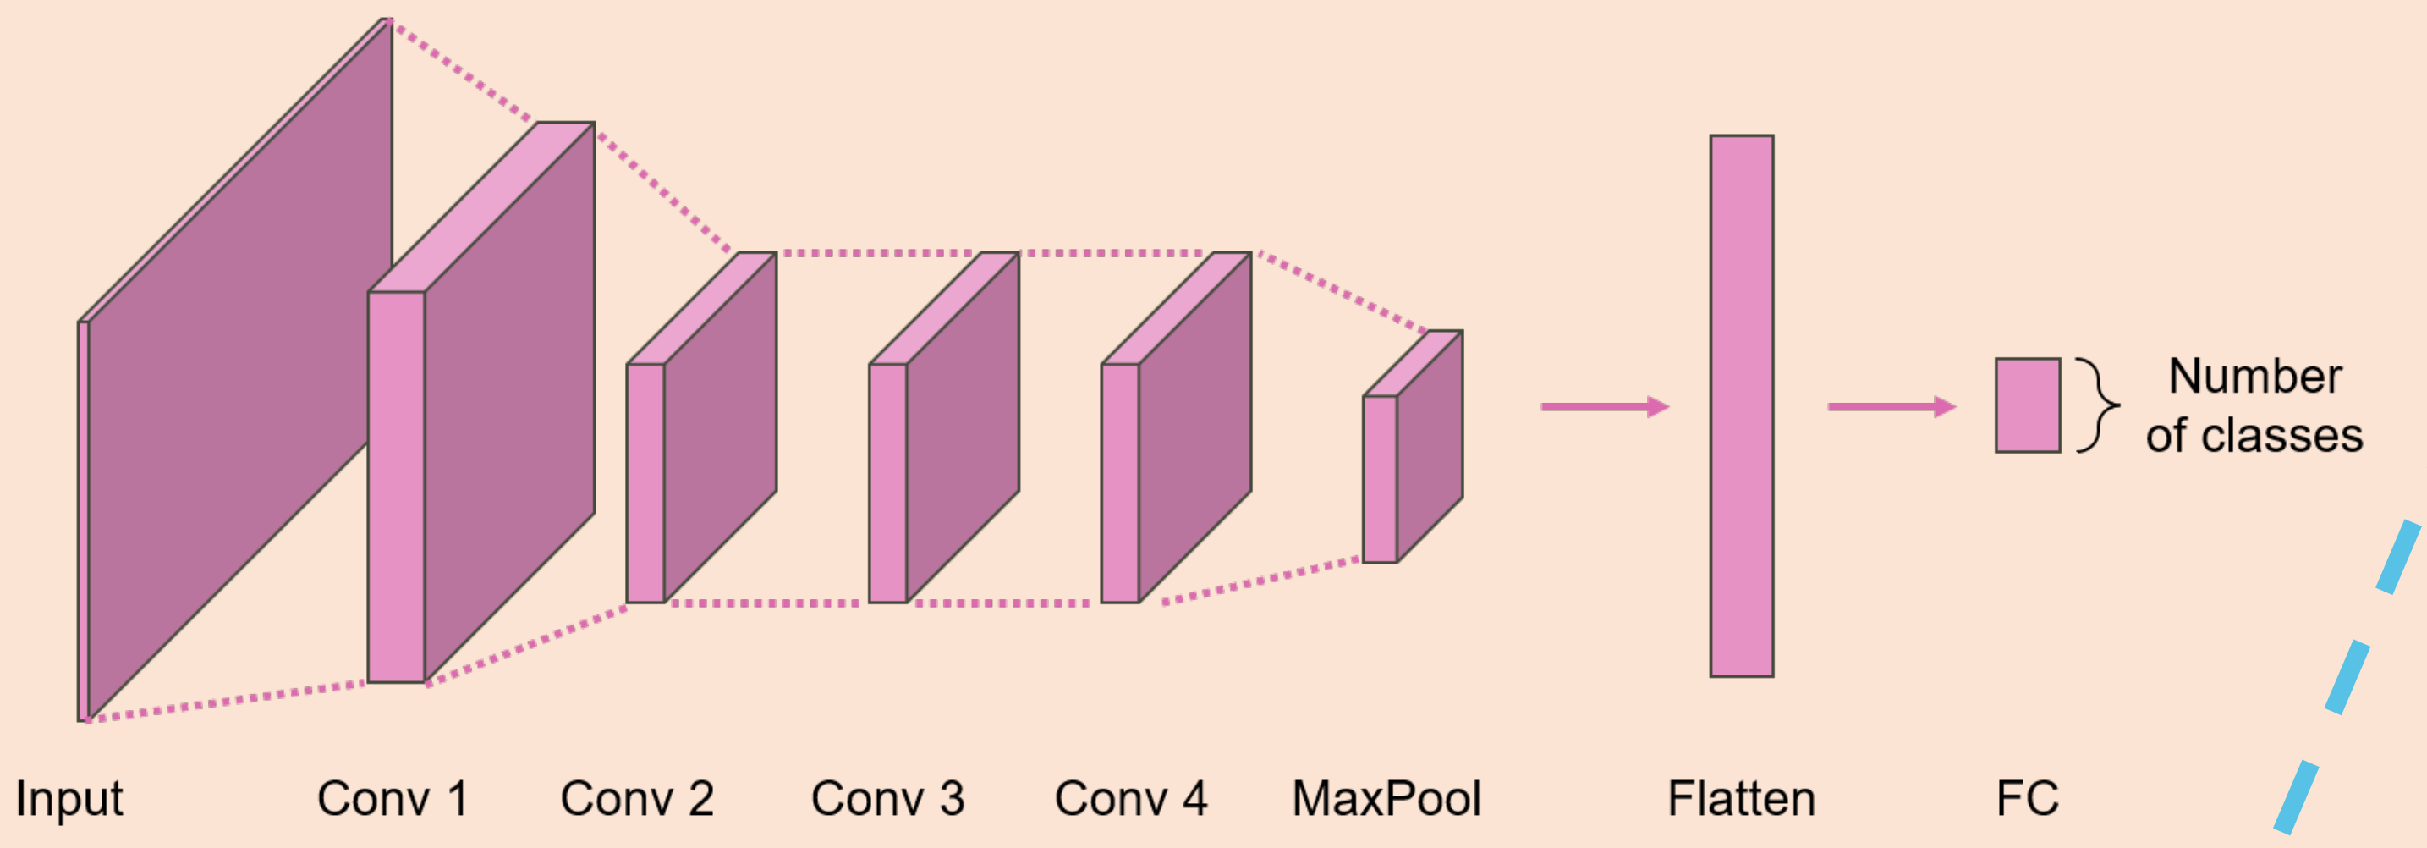

Figure 2. Classical PATE network architecture. Classical PATE uses four convolution blocks

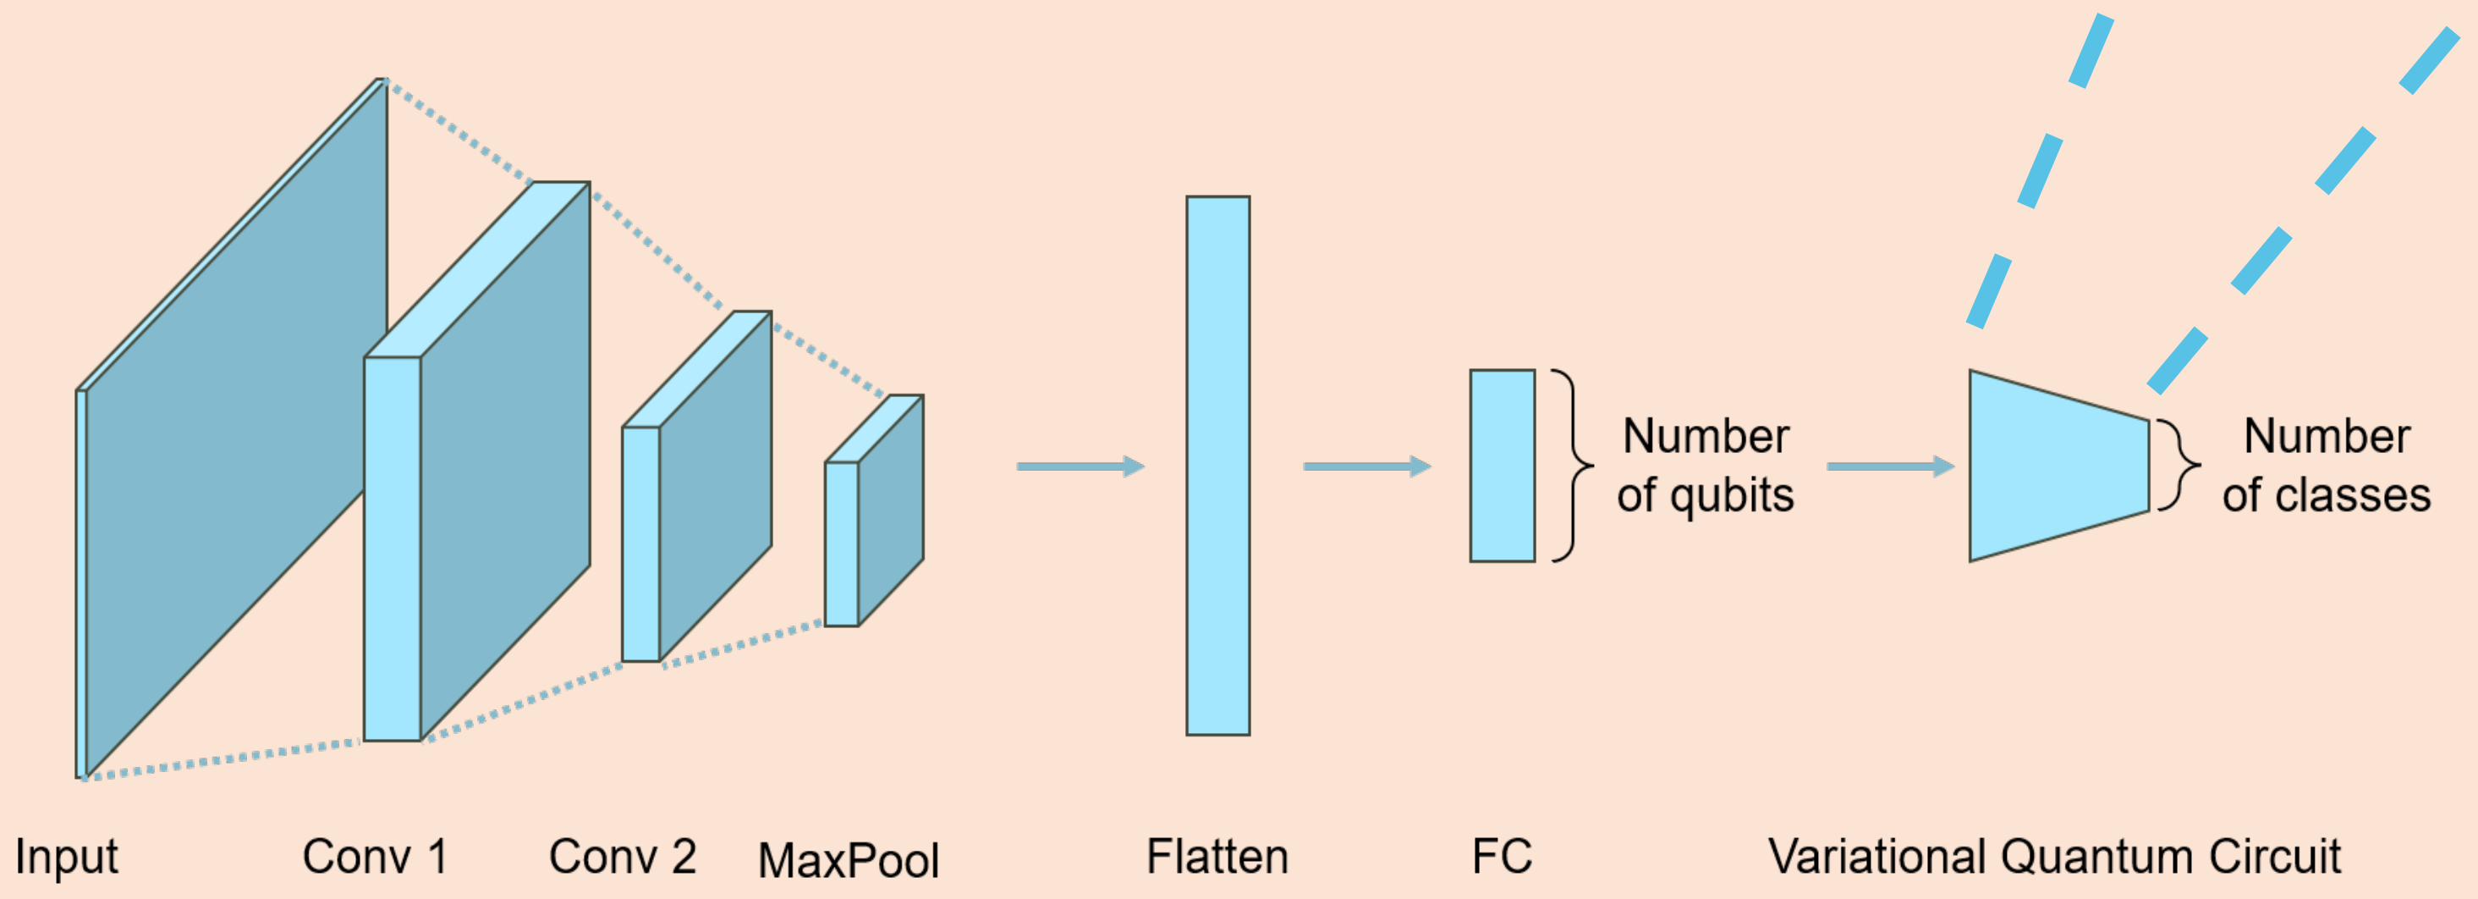

Figure 3. Quantum PATE network architecture. Quantum PATE uses two with the additional VQC blocks

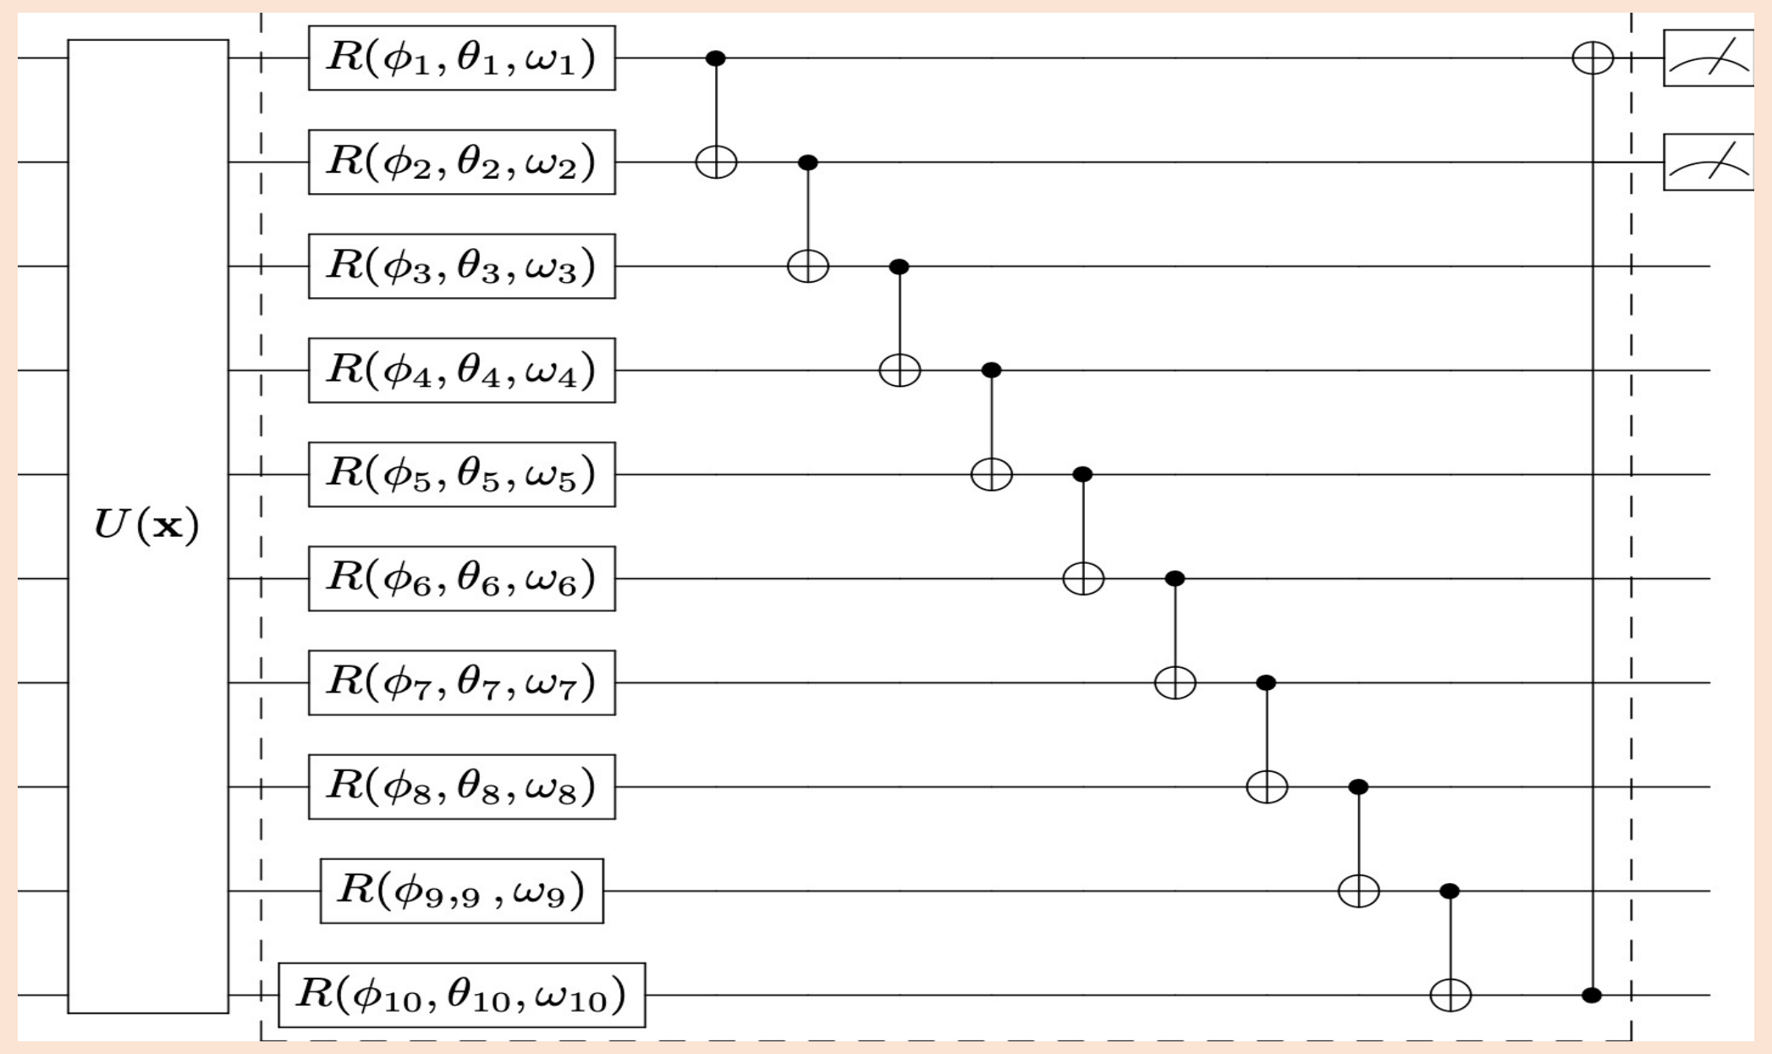
$U(\mathbf{x})$
$R(\phi_1, \theta_1, \omega_1)$
$R(\phi_2, \theta_2, \omega_2)$
$R(\phi_3, \theta_3, \omega_3)$
$R(\phi_4, \theta_4, \omega_4)$
$R(\phi_5, \theta_5, \omega_5)$
$R(\phi_6, \theta_6, \omega_6)$
$R(\phi_7, \theta_7, \omega_7)$
$R(\phi_8, \theta_8, \omega_8)$
$R(\phi_{9,9}, \omega_9)$
$R(\phi_{10}, \theta_{10}, \omega_{10})$

Figure 4. VQC block for MNIST classification. The VQC block encodes latent embeddings from convolution blocks within quantum PATE into quantum states represented by 10-qubits. $U(x)$ denotes the quantum algorithm for angle encoding. $\emptyset_i$, θ_i, and ω_i are the parameters to optimize. The dashed box denotes one subcircuit of the VQC block that is repeated two times. The dial to the far right represents that the circuit has two outputs. The expectation of σ_z is measured on two qubits.

Result

ϵ	δ	classical PATE	quantum PATE
0.01	10^{-5}	0.534 ± 0.0992	$\mathbf{0.688} \pm 0.0163$
0.1	10^{-5}	0.985 ± 0.0215	$\mathbf{0.992} \pm 0.0098$
1.0	10^{-5}	$\mathbf{0.997} \pm 0.0046$	0.99 ± 0.0134
10.0	10^{-5}	$\mathbf{0.997} \pm 0.0046$	0.991 ± 0.0137

Table 1. Results from binary MNIST classification. Accuracies of classical PATE and quantum PATE after 20 epochs. The private quantum classifier is more accurate and successful for ε between 0.01 and 0.1. The number of teachers is set as 4.

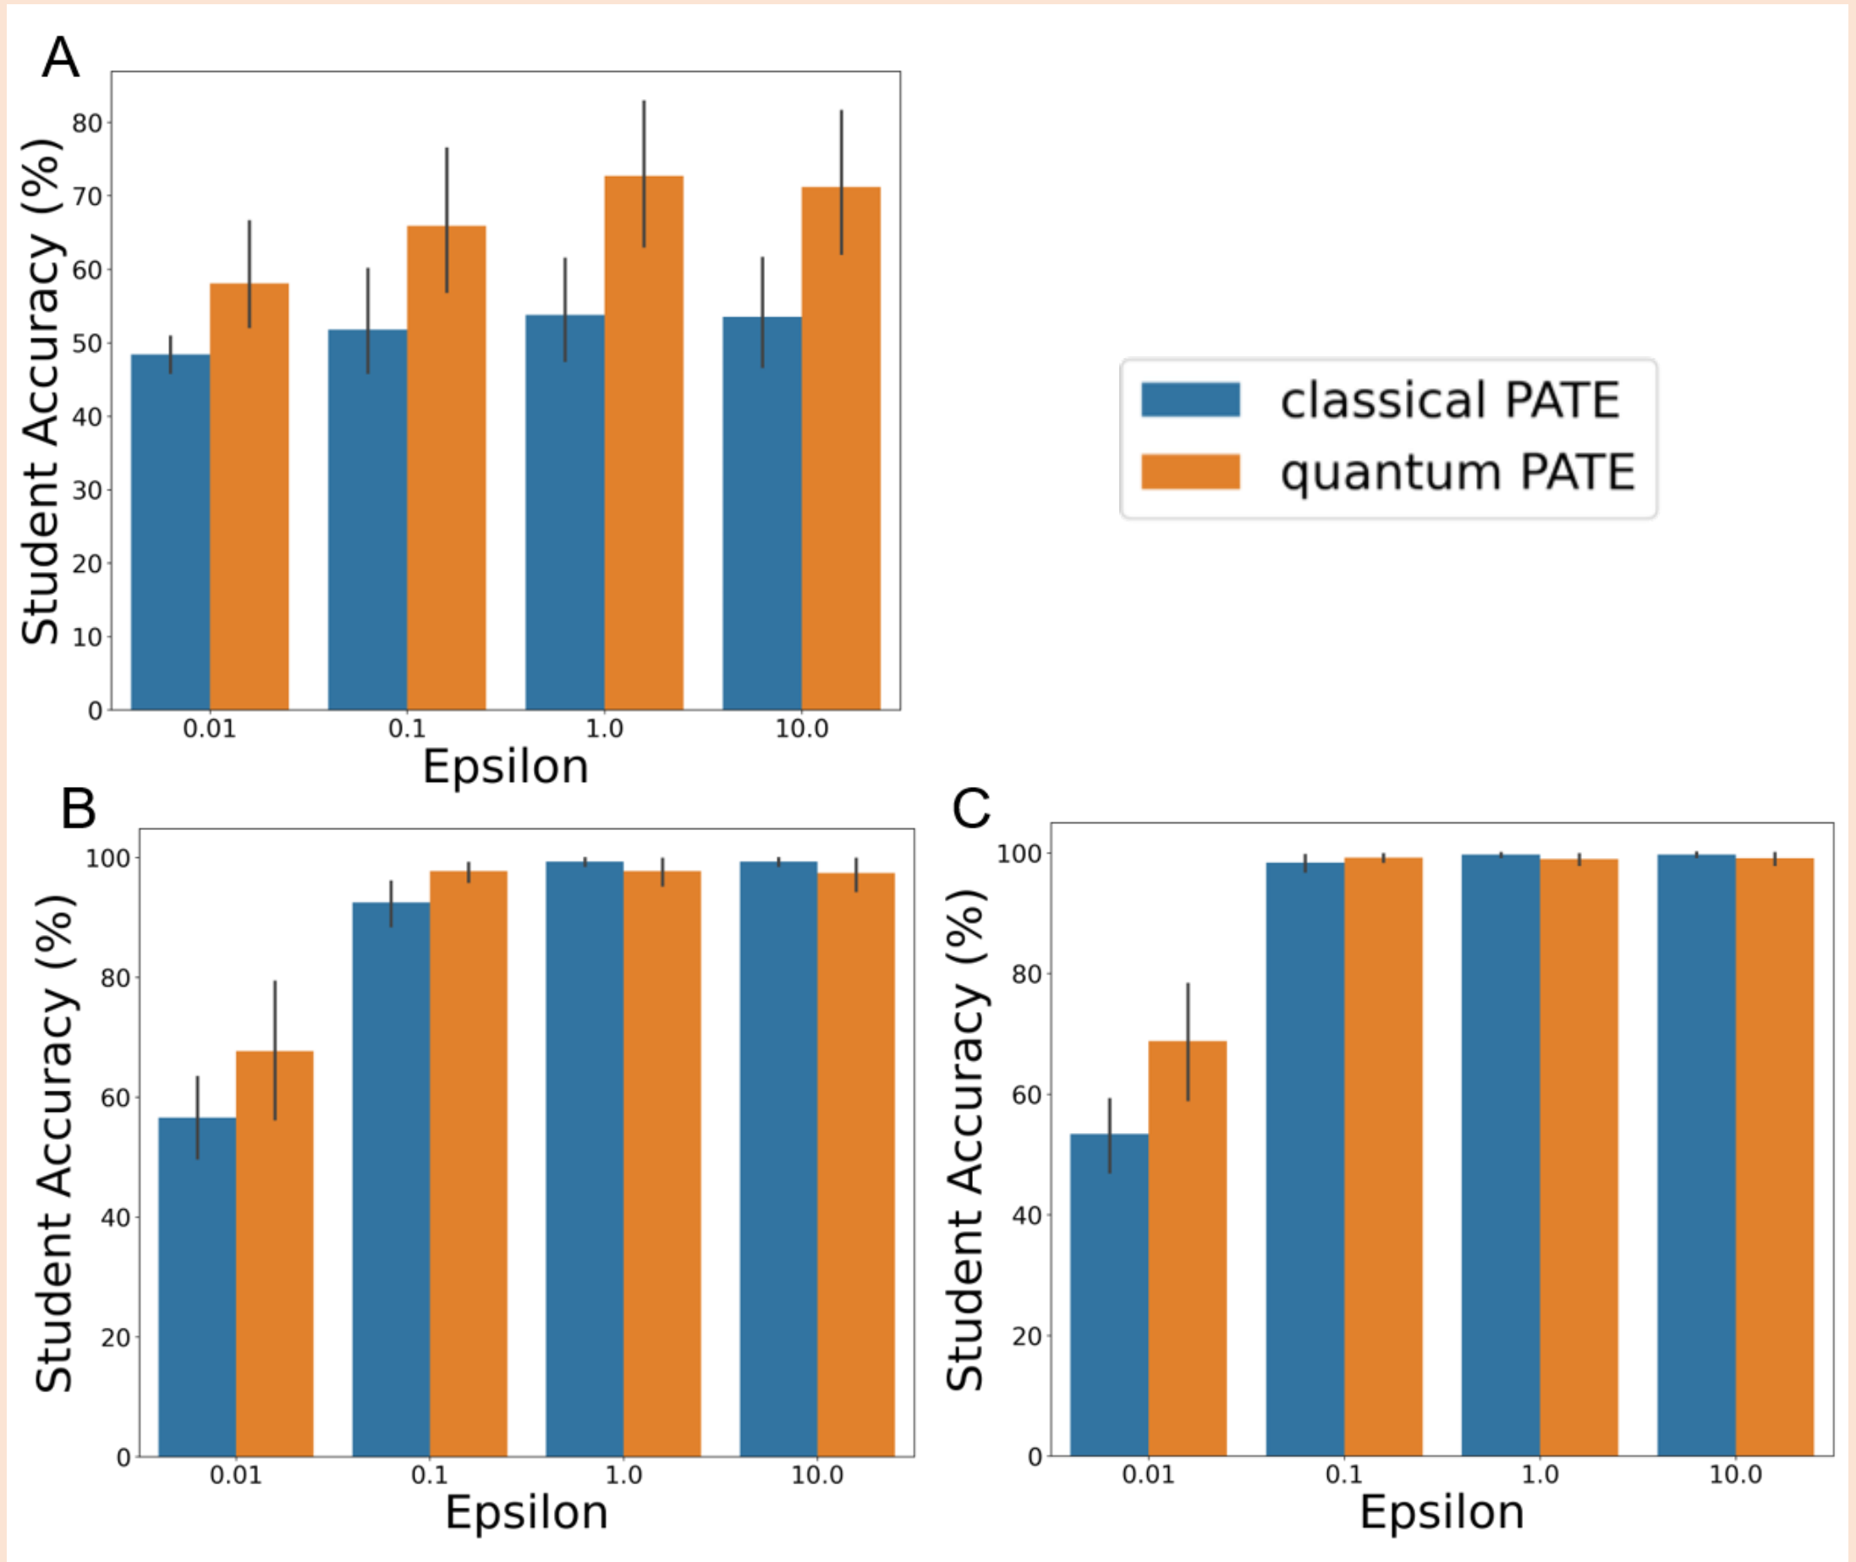

Figure 5. Accuracy vs. Epsilon for 4 teachers in classical PATE and quantum PATE. We averaged the results from 10 experiments, and the error bar denotes the standard deviation. (A), (B), and (C) respectively show the result of 1, 10, and 20 epoch training.

Discussion

Impacts

- We demonstrated the potential of hybrid quantum-classical framework for accurate and privacy-preserving machine learning
- QPATE shows a challenge of balancing accuracy and privacy (ε values)
- Hybrid approach improves prediction accuracy at low ε values compared to classical DNNs

Limitations

- Trade-off between accuracy and privacy not investigated based on number of teachers
- Further research needed to establish quantum advantage in differential privacy
- Potential of VQC in PATE evaluated with limited subcircuits and qubits; scalability requires exploration
- Hybrid quantum-classical classifiers evaluated in simplified settings; needs more complex tasks