Ensembling Personalized Quantum Models with Local Differential Privacy via Meta-Learning

Flavjo Xhelollari, Juntao Chen

Fordham University

fx1 [at] fordham [dot] edu

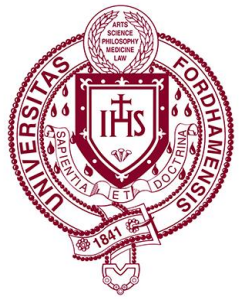

Motivation

Quantum Machine Learning (QML) holds promises for faster, more efficient Machine Learning.

Privacy concerns arise as user data is required for model training.

Utilizing a de-centralized approach for building a model.

Problem Statement

- Q: How to train user-specific QML models that ensure privacy and maintain high accuracy?

- Proposed solution : Hybrid QC framework with ε_i –LDP at the user level, and meta-learner to combine individual models.

Background

QML combines QC and ML, and makes use of VQCs.

LDP ensures data privacy by perturbing information at the user level, before sharing. It also balances privacy and utility.

Related Works

Hybrid QC Architectures :

- Focus on individual models, not ensembles.

[Hirce et al., 2023 ; Kashefi 2023 ; Gong et al., 2020]

DP and QML :

- Rely on centralized mechanisms (e.g. Federated Learning).

[Nunez et al., 2022 ; Kashefi 2023 ; Watkins et al., 2023]

<u>Our contribution:</u>

- Customizable LDP integrated in QML training.
- Meta-Learning ensemble for multiple VQC-based models.

Methodology

Phase 1 - User-specific training:

- Users apply LDP locally
- Data is encoded to quantum states via R-gates
- Train QML models via Parameter-Shift rule

Phase 2 - Meta-Learner:

- Combine predictions from multiple models
- Aggregate them via Neural Networks
- Train the ensemble for convergence

Framework

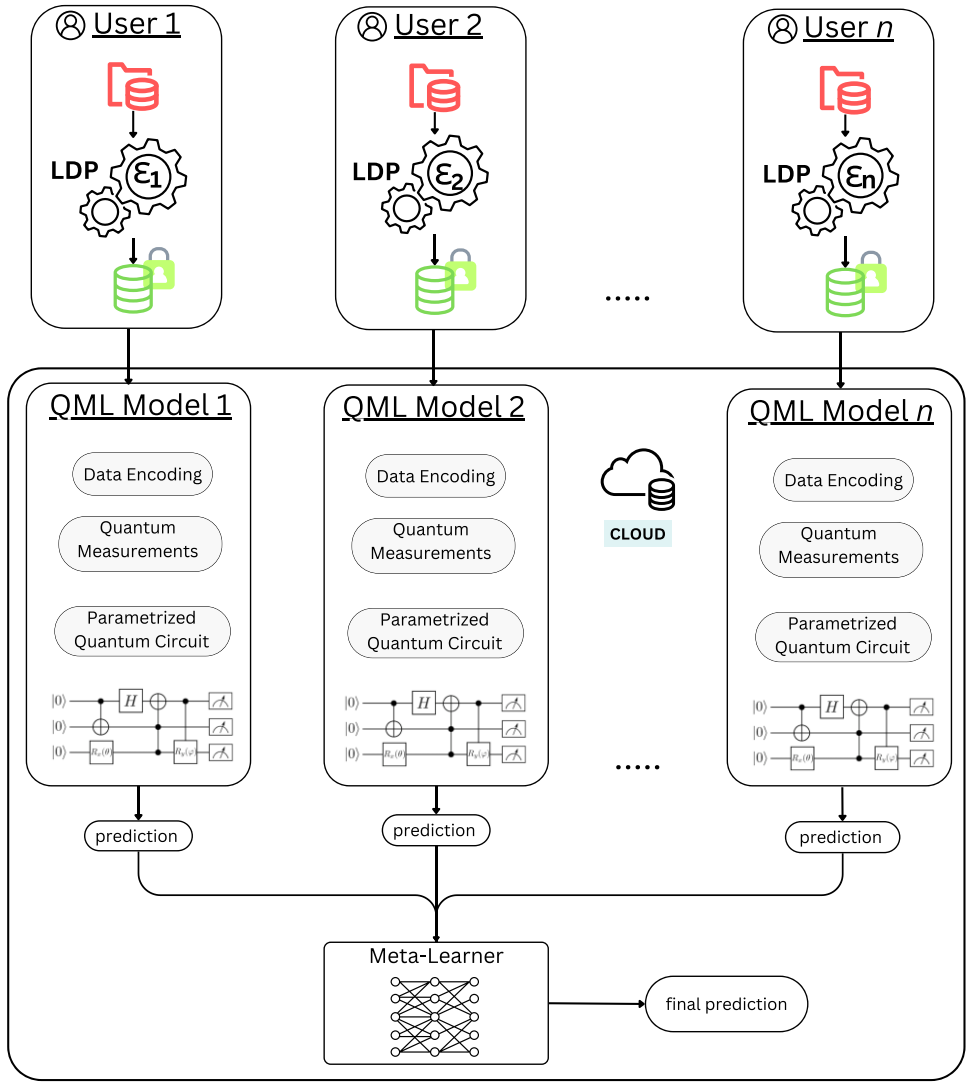

Figure 1. Framework Overview

Experimental Setup

01

Datasets : MNIST, IRIS for binary classification.

02

Privacy budgets ranging from *0.1* to *5.0*.

03

Quantum Models built via PennyLane, using 8-qubits with strongly entangling layers.

04

Metrics : Accuracy, loss, and privacy-performance tradeoffs.

Results

IRIS :

- Stronger privacy results in slightly lower accuracy compared to the no-LDP.
- Smaller dataset sizes result in smoother convergence patterns.

MNIST :

- LDP accuracy improves with higher Ɛ.
- Both LDP and no-LDP versions converge, but LDP ensemble shows higher variance due to the noise.

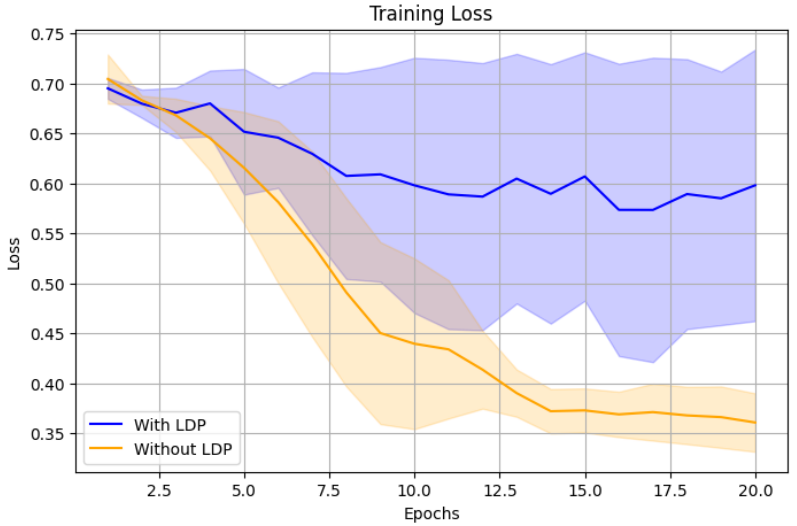

Figure 2. Training loss of the ensembles

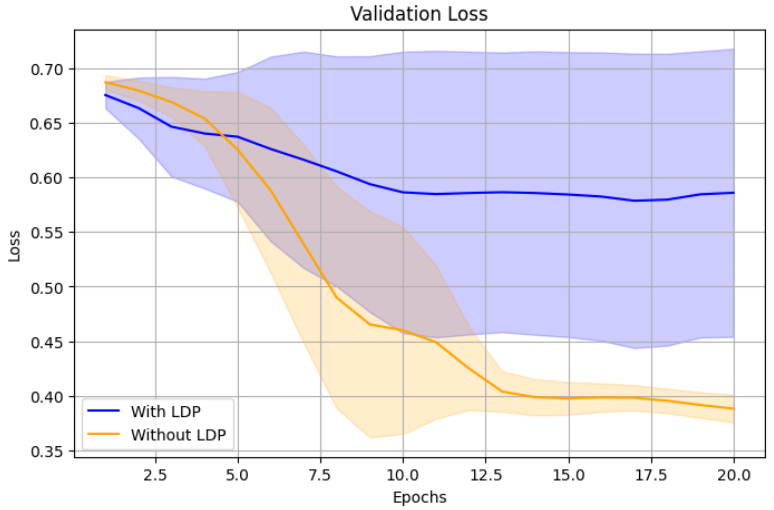

Figure 3. Validation loss of the ensembles

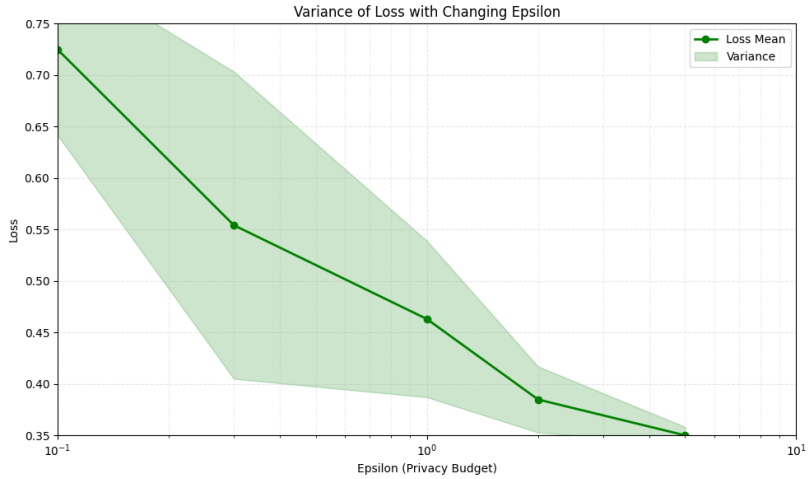

Figure 4. Variance loss w.r.t. $\mathcal{E}$

Conclusion

THE PROPOSED FRAMEWORK ENSURES STRONG PRIVACY GUARANTEES WITH MINIMAL PERFORMANCE LOSS.

META-LEARNING ENSEMBLES MITIGATE THE TRADE-OFF BETWEEN PRIVACY AND ACCURACY.

EXPERIMENTAL RESULTS VALIDATE THE FRAMEWORK'S EFFECTIVENESS ACROSS DIFFERENT SCENARIOS.

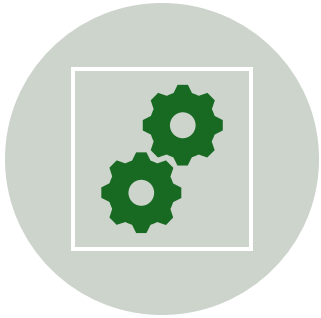

FUTURE WORK:

EXPLORE ADVANCED ARCHITECTURES FOR LARGER DATASETS, EXTEND THE FRAMEWORK TO MULTI-MODAL CLASSIFICATION TASKS, INVESTIGATE IN REAL-WORLD APPLICATIONS (HEALTHCARE, FINANCE..)

Q&A

THANK YOU