Neural Network Training on Encrypted Data with TFHE

Luis Montero, Jordan Frery, Celia Kherfallah, Roman Bredehoft, Andrei Stoian

Abstract

We show that learning gradient based classifiers with **Stochastic Gradient Descent** using **TFHE** and **6-bit quantization** yields accuracies and runtimes comparable with the state of the art. We showcase this approach on **Logistic Regression** and **Multi-Layer-Perceptrons**.

Code available on Github

Rounding and Quantization

$$x_q = \lfloor (\sigma_x x_f - \mu_x) \rceil$$

$$\sigma'_x = 2^p \sigma_x \implies x_r = \lfloor \frac{x_q}{2^p} \rceil 2^p$$

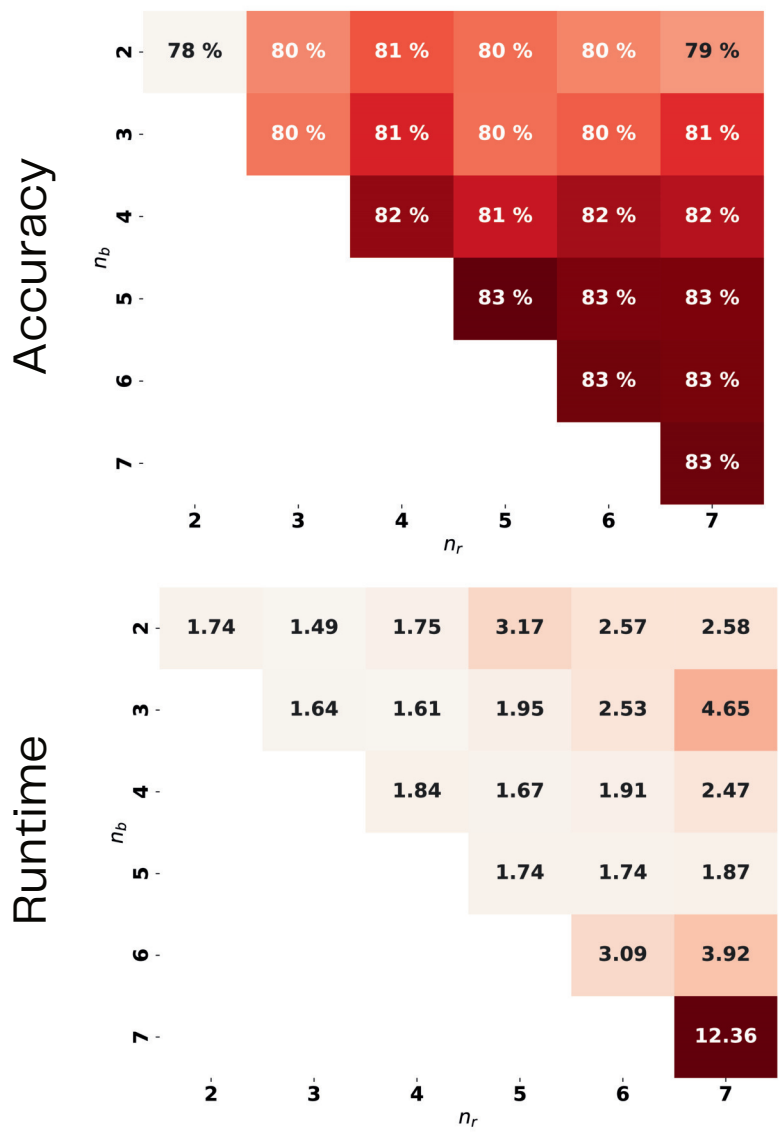

FHE vs FP32

Dataset	Model/# params	Quantization bits	Best FHE accuracy	Best FP32 accuracy	Batch latency
Breast cancer	Logistic / 30	6b	98.25%	99.12%	11.8s
Breast cancer	MLP (1 hidden) / 930	4b	98.25%	99.12%	149s
Mortality	Logistic / 10	6b	90.09%	90.47%	7.2s
Mortality	MLP (1 hidden) / 165	4b	87.25%	90.44%	45s
Heart disease	Logistic / 13	6b	88.52%	89.47%	8.02s

Acurracy per batch on Breast Cancer Dataset:

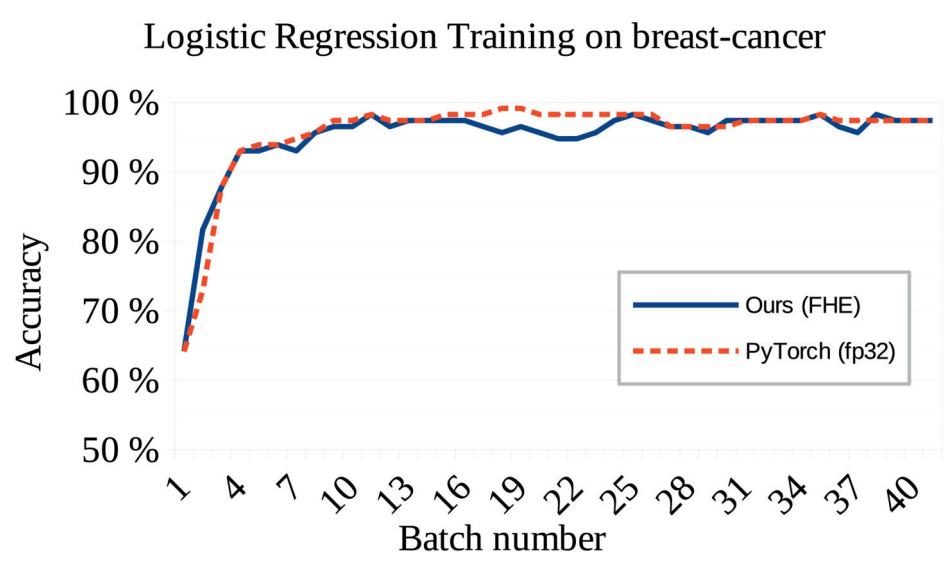

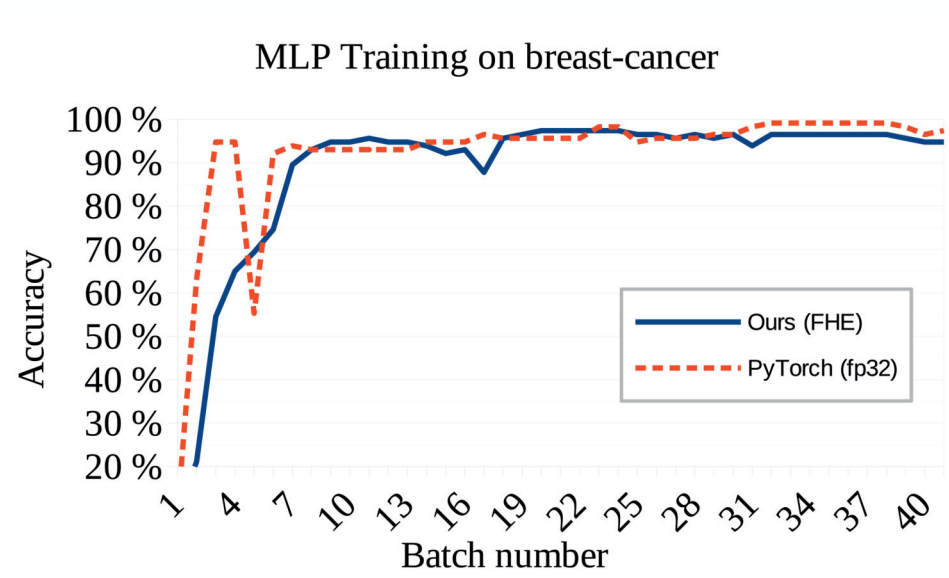

ZAMA

Resources
zama.ai/blog
github.com/zama-ai/concrete-ml