ABSTRACT

PRIVACY-PRESERVING DEEP NEURAL NETWORK USING HYBRID FHE SCHEMES

By:

Anas Banta Seutia

21/475041/PA/20517

Privacy concerns in cloud-based machine learning necessitate cryptographic solutions that enable computation on encrypted data. This thesis presents a hybrid Fully Homomorphic Encryption (FHE) framework for deep neural networks that leverages scheme switching between CKKS (Cheon-Kim-Kim-Song) and TFHE (Fast Fully Homomorphic Encryption scheme over the Torus) to eliminate input-dependent preprocessing requirements. Unlike existing approaches that rely on polynomial approximations requiring predetermined input ranges and client-side precomputation, this implementation achieves exact nonlinear activation evaluation through strategic scheme transitions.

The system performs CKKS SIMD optimised encrypted convolutions using toeplitz packing and matrix multiplications for arithmetic efficiency, then switches to TFHE scheme for exact ReLU computation via functional bootstrapping, enabling domain-unlimited operation over $(-\infty, \infty)$ without range calibration.

Evaluation of ReLU using scheme switching demonstrates 12.4 bits of effective precision. The evaluation of the MNIST test data shows 15.65s reduction in preprocessing overhead compared to approximation-based methods, while maintaining comparable total inference time of 215.26 seconds. The unbounded operational capability proves essential for encrypted training scenarios where activation distributions evolve unpredictably. This work validates that hybrid FHE schemes can practically support privacy-preserving deep learning through modular architectures that leverage complementary strengths of different cryptographic schemes, establishing a foundation for confidential cloud-based neural network services in healthcare, finance, and other privacy-critical domains.

Keywords: Fully Homomorphic Encryption, Deep Neural Networks, Scheme Switching, Privacy-Preserving Machine Learning, CKKS, TFHE