Privacy-Preserving Inference of Deep Differential Logic-Gate Networks under TFHE

Mahmoud Y. M. Yassin, Mahmoud AbdelHafeez Sayed, Mostafa Taha
Systems and Computer Engineering, Carleton University
Correspondence: mahmoudyassin3@cmail.carleton.ca

AI/CRV 2026

SCI-LAB Carleton University

Problem Statement

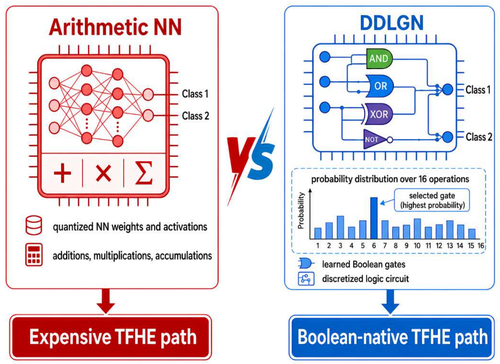

- Cloud AI inference can expose sensitive user data.
- TFHE protects privacy, but arithmetic neural networks are costly under encrypted evaluation.
- DDLGNs offer a Boolean-native path: neural models discretized into learned two-input logic-gate circuits.

Summary

- DDLGNs align naturally with TFHE because inference becomes Boolean circuit evaluation.
- Encrypted latency scales mainly with the number of evaluated logic gates.
- PP-DDLGN is a promising path toward efficient, privacy-preserving AI inference.

References

1. F. Petersen, C. Borgelt, H. Kuehne, and O. Deussen, "Deep differentiable logic gate networks," *NeurIPS*, vol. 35, pp. 2006–2018, 2022.
2. I. Chillotti, N. Gama, M. Georgieva, and M. Izabachène, "TFHE: Fast fully homomorphic encryption over the torus," *Journal of Cryptology*, vol. 33, no. 1, pp. 34–91, 2020.
3. A. Stoian et al., "Deep neural networks for encrypted inference with TFHE," in *CSCML*. Springer, 2023, pp. 493–500.

Want to know more?
Let's connect on LinkedIn.
Scan QR

Proposed Encrypted Inference Pipeline

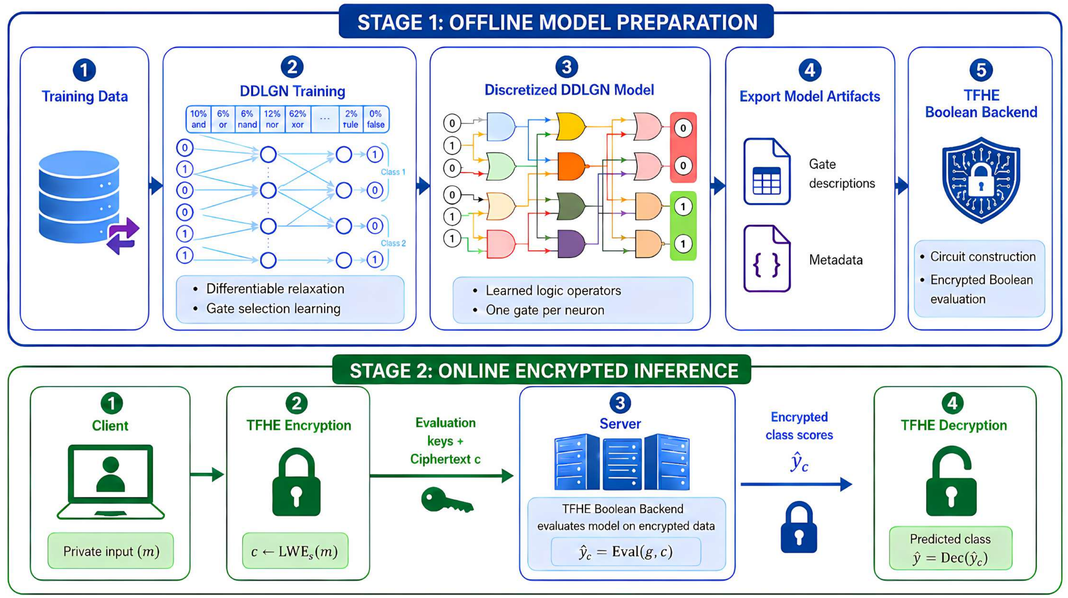

- The server trains a DDLGN and discretizes it into a learned Boolean circuit.
- The client encrypts the input using TFHE, sends it to the server, and keeps the secret key locally.
- The server evaluates the circuit directly over encrypted bits, without observing the plaintext input.
- The client decrypts returned class scores to obtain the final prediction.

Key Findings

- We experiment with three PP-DDLGN sizes on MNIST: 8k, 24k, and 48k logic gates.
- Larger Boolean circuits improve accuracy, but increase latency.
- Homomorphic circuit evaluation dominates the end-to-end runtime; encryption and decryption are comparatively small.
- Compared with Quantization-Aware Training Fully Connected Neural Network (QAT-FCNN) baselines, the proposed PP-DDLGN achieves higher accuracy with lower encrypted inference latency.

Experimental Results