Fancy-garbling

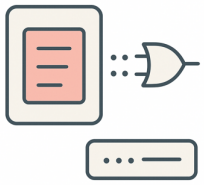

Machine learning (ML) models, particularly neural networks (NNs), are extensively used in tasks such as image recognition, data classification, and analytics. Often, the primary privacy concern lies in protecting sensitive input data rather than the ML model itself. Indeed, models can frequently be extracted or learned, even when kept secret. Practical examples of scenarios with focus on input privacy include:

- A client outsourcing sensitive data classification tasks to a cloud service without revealing their inputs.
- Public classifiers evaluating data secretly shared among multiple parties.
- Proving claims about private inputs without disclosing underlying data.

In these contexts, the garbled circuit (GC) approach is particularly suitable, as it efficiently ensures data privacy during the evaluation of public ML models. The goal of this thesis is to explore the application of garbled circuits to privacy-preserving neural network evaluations, specifically building upon recent cryptographic advances.

The thesis will focus on implementing and experimentally evaluating optimized GC techniques tailored specifically for ML applications. Neural networks can be securely evaluated by encoding their computations into garbled circuits. While general GC implementations exist, recent studies have demonstrated significant efficiency improvements through specialized garbled operations designed explicitly for neural network components, including sign function, approximate computations, multiplication improvements.

The student can approach this project through the following structured phases:

1. Reviewing and understanding GC techniques, in particular reviewing recent improvements in GC specifically targeting ML applications (e.g., approximate activation functions and optimized arithmetic operations).
2. Implementing improved GC operations for ML components such as ReLU, max pooling, and sign functions. and experimenting with approximate versions of these functions to explore trade-offs between accuracy and efficiency (preliminaries implementations and public libraries are available).
3. Implement and evaluate two privacy scenarios, with public and private weights

Practicalities

Promoter:	Emmanuela Orsini emmanuela.orsini@unibocconi.it Office 2-C2-03
Nature of the work:	40% literature, 20% theoretical work, 40% software
Number of students:	Max 2
Prerequisites:	A background in cryptography and programming skills are required