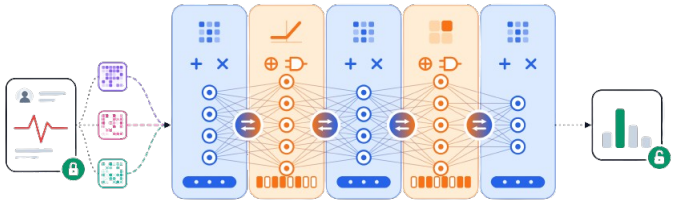
The right share at the right time: Secret-Sharing Conversions for Privacy-Preserving Machine Learning



Privacy-preserving machine learning (PPML) allows machine-learning models to be evaluated on sensitive data without requiring the underlying data to be revealed. Many PPML protocols rely on secret sharing (SS) to represent intermediate values, with arithmetic and Boolean SS being two commonly used representations. While arithmetic sharing is efficient for addition and multiplication, Boolean sharing or Garbled Circuits (GC) can be advantageous for operations such as comparisons and nonlinear functions (ReLU, maxpool, truncation). Switching between these representations, however, requires dedicated conversion protocols that introduce additional computation and communication overhead.

In this master's thesis, we want to investigate the trade-offs involved in SS conversions at different stages of PPML protocols, with particular attention to non-linear layers and comparisons. Several conversion techniques exist, ranging from direct A2B (arithmetic-to-boolean) and B2A (boolean-to-arithmetic) conversions to mixed-protocol approaches such as ABY and ABY3, and preprocessing-based methods like daBits and edaBits. Which conversion technique is preferable may vary depending on the non-linear layer being evaluated, the number of parties, the ring or field the computation lives in, and the latency and bandwidth of the deployment. The goal is to determine how and where the choice of conversion technique affects the performance of complete PPML workloads.

The student will, after a careful literature review on secure multi-party computation (MPC) and secret-sharing conversion techniques, implement and benchmark several conversions inside an existing MPC framework and integrate them into representative PPML computations. The comparison will cover the common non-linear layers in isolation as well as complete inference on a small neural network, evaluating computation time, communication overhead, and round complexity under varying network conditions. This will result in a clear picture of the trade-offs and, ideally, practical guidance on which conversion to select at each stage of the pipeline.

An introduction to mixed-protocol MPC can be found in the ABY paper [1]; ABY3[2] extends this to the three-party machine learning setting, and the edaBits[3] paper is a good starting point for state-of-the-art conversion techniques.

Practicalities

Promotor:	Bart Preneel		
Daily supervisor:	Aysajan Abidin	aysajan.abidin@esat.kuleuven.be	office 01.13
	Neus Vayà	nvayaviv@esat.kuleuven.be	office B01.18
Nature of the work:	30% literature, 70% software		
Number of students:	1		

[1] https://www.ndss-symposium.org/wp-content/uploads/2017/09/08_2_1.pdf
[2] https://dl.acm.org/doi/10.1145/3243734.3243760
[3] https://eprint.iacr.org/2020/338