Presenter Name: David Zagardo

Talk Title: "Block-wise Gradient Aggregation Enables Privacy Preserving Deep Learning"

Presenter Short Bio: David Zagardo earned his Master's in Privacy Engineering from Carnegie Mellon University in 2022. After graduation, he worked at Dynamo AI and built their MultiGPU differentially private fine tuning SDK for distributed large language model training. He now splits his time between scanning the web for privacy violations at WebXRay.ai and dreaming up new differential privacy algorithms for deep learning.

Short Talk Abstract: Traditional Differentially Private Stochastic Gradient Descent (DP-SGD) introduces statistical noise on top of gradients drawn from a Gaussian distribution to ensure privacy. This talk introduces the novel Mutual Information Differentially Private Aggregation (MI-Aggregate) algorithm for deep learning. MI-Aggregate builds off of existing private deep learning literature, but makes a definitive shift by taking a block-wise aggregation approach to its privacy mechanism modeled after information theoretic privacy analyses. The theoretical results presented in this talk show that the combination of block-wise aggregation, parameter-specific block size selection, block-level clipping, and gradient accumulation allows MI-Aggregate to achieve space complexity identical to that of non-private training while maintaining statistically similar privacy guarantees to DP-SGD. MI-Aggregate is found to be significantly more performant with respect to perplexity than DP-SGD. The theoretical results are validated by the experimental findings.