SAND2021-8517C

The Center for Cyber Defenders
Expanding computer security knowledge

Privacy-Preserving AutoML

Sandia National Laboratories/NM, U.S. Department of Energy

July 27, 2021

Alycia N. Carey | University of Arkansas | Ph.D. Computer Science | May 2024

Susan Gardener | Nicholas Pattengale | 5629

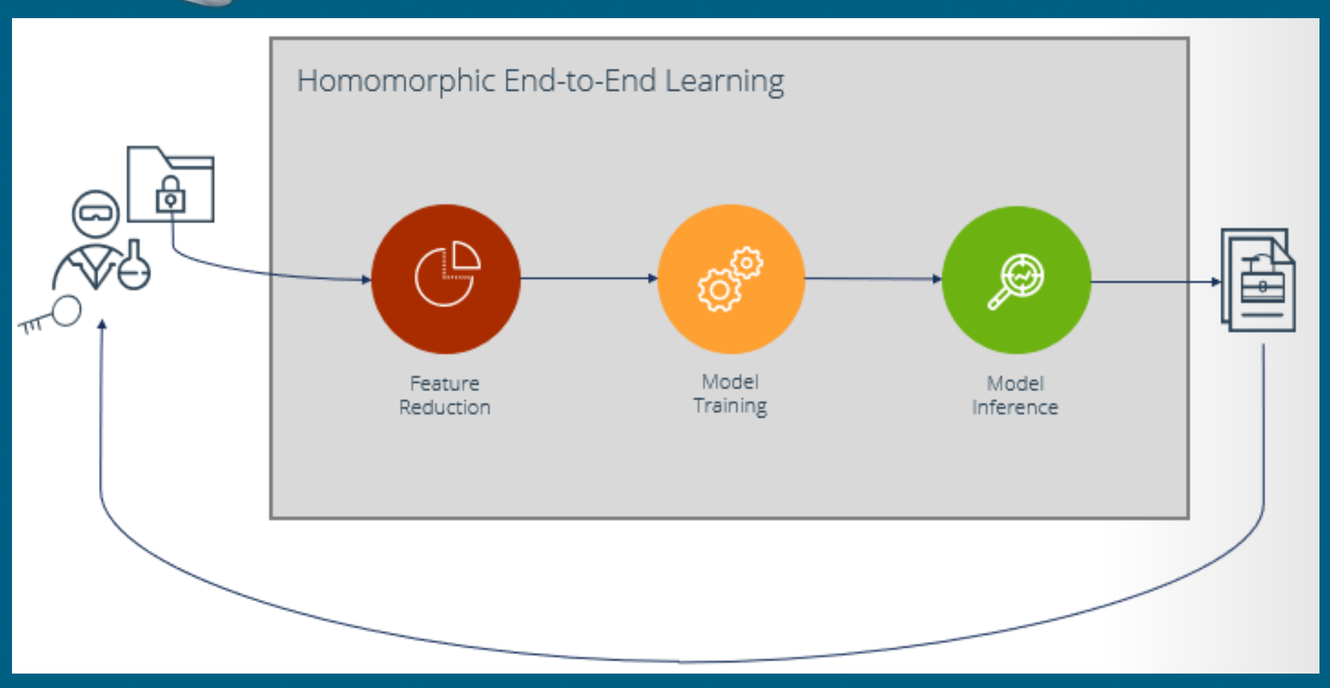

Abstract: Recent years have yielded breakthrough progress on performing machine learning training and inference under fully homomorphic encryption (FHE). While this recent progress signals a promising future for encrypted computation, it will be many years before FHE is adopted in to the standard practice of computing. In attempt to understand the feasibility of short-circuiting incremental FHE progress, we formulate a vision for privacy-preserving automated machine learning (AutoML, which automates all parts of the machine learning process, including feature selection, model training, parameter selection, and model inference). Our envisioned capability, PPautoML, has the potential to bootstrap the machine learning community with a leap forward security posture. We are executing two tasks to explore this goal: surveying and demonstrating. We also examine national security use cases where a PPautoML pipeline might be useful. We have completed the survey process and are beginning to assess the feasibility of demonstrating PPautoML using already developed components.

Introduction: AutoML provides methods and processes to make machine learning available for non-machine learning experts or researchers. It aims to improve the efficiency of machine learning as well as accelerate the speed of research. Unfortunately, some data domains cannot be processed in an unencrypted manner by a third-party AutoML system due to security constraints. However, the capability to perform automated machine learning in an encrypted domain does not exist today. The creation of a privacy-preserving AutoML pipeline could significantly reduce the risk of data compromise and the burden of data protection in the application of machine learning. It is our goal to not only compile an overview of the state of the field of fully homomorphic encryption in machine learning, but also provide a realistic look into the implementation of a privacy-preserving AutoML pipeline.

	Papers	Implementation
Total	198	24
Feature Reduction	14	1
Model Training	68	8
Model Inference	127	21
Application	6	-
Background	32	-

Methods:

1. **Survey**: Using papers published on IEEE, ArXiV, IACR, and ACM
   a. Catalog published techniques
   b. Note which have open source implementations
   c. If an implementation exists, grade the maturity and quality
   d. Note which machine learning techniques should have privacy preserving implementations if none currently exist
2. **Demonstrate**: Using only open source data and modules
   a. Demo/articulate utility of technique mashup under:
      a. Same FHE schemes
      b. Mixed FHE schemes using implementation tools like CHIMERA or following a multi-encrypt/decrypt process

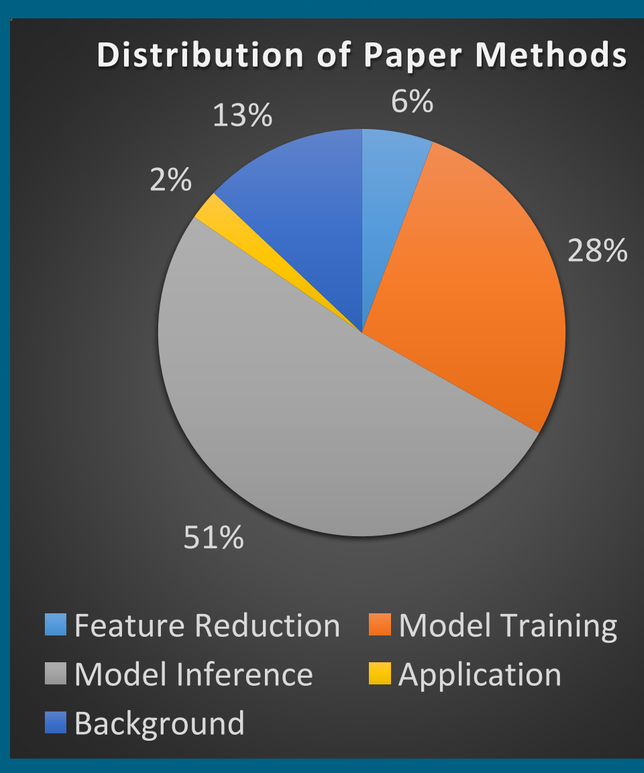

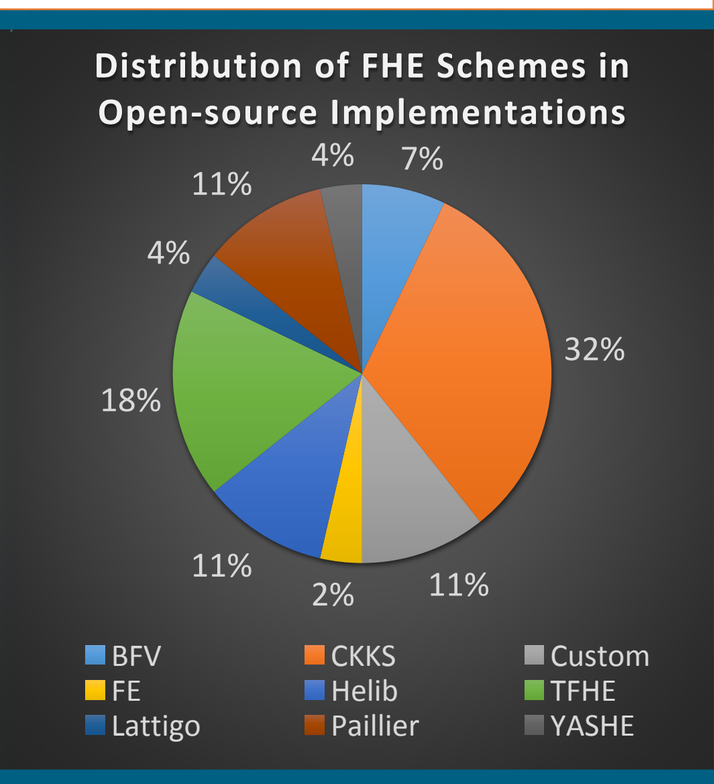

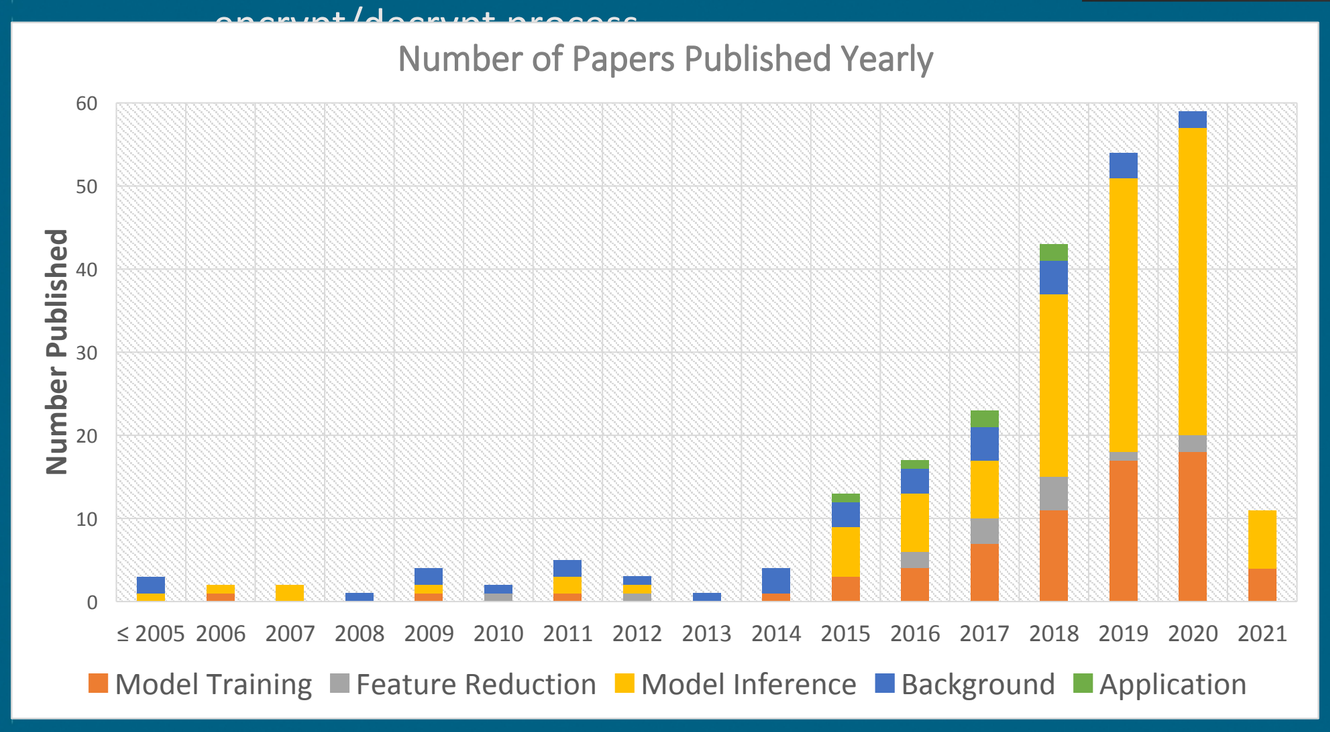

Results and Discussion: We were able to collect approximately 200 papers pertaining to privacy preserving AutoML processes such as feature reduction, model training, and model inference. After the completion of the survey phase, it was found that most papers that have no open-source implementation utilize custom FHE schemes, most open-source implementations do not use the same FHE, and there is a massive lack of feature reduction papers and open-source feature reduction implementations. The next phase, demonstration, will reveal if creating a privacy-preserving AutoML pipeline using open-source implementations is as feasible as originally thought. The major hurdle to realizing this goal is the disparate utilization of FHE schemes in the currently available open-source methods.