

Privacy Enhancing & Privacy Preserving Technologies

An Introduction

Yogesh Kumar



Disclaimer

The views and opinions expressed in this presentation, including any options discussed, are the personal views and opinions of the presenters and do not necessarily reflect the views or positions of any entities they represent, including their current or former employers. The presenters is solely responsible for the accuracy of the information presented and any opinions expressed.

This disclaimer is intended to clarify that the presenters are speaking on their own behalf and that their opinions do not necessarily reflect the views of their employers. It is important to note that this disclaimer is not a substitute for legal advice, and you should consult with an attorney if you have any questions about your specific situation.

Copyright Disclaimer

All copyrighted material cited in this work is the property of the respective copyright holders. The author of this work has made all reasonable efforts to identify and credit all copyright holders, but if any copyright holders have not been properly identified or credited, the author apologizes for the oversight.

About me

- Delhi College of Engineering- BE, Rutgers University Executive MBA
- Currently: Senior Manager @ Walmart
- Previously: Amazon, AT&T, Nokia
- Product Development, R&D, Program Management, Business Development
- Now in East Bay. Seattle, Atlanta, Beijing, Shanghai, Tampere/Finland, New Delhi
- Rock Climbing, Mountaineering, Hiking
- Complex Systems Theory, History of Ideas, Evolution of Technology, Philosophy





LinkedIn



Agenda

- Introduction
- Key Data Privacy Issues addressed by PETs/PPTs
- PETs What do various PETs solve and how?
 - Homomorphic Encryption
 - Confidential Computing
 - Differential Privacy
 - Data Clean Rooms
 - Disclosure Avoidance System
- Summary

Purpose of PETs & PPTs



- Solutions aimed at **protecting individual privacy**
- Variety of methods aimed at different aspects of data handling
 - Enhancing *minimize the amount of personal* data collection, processing, and sharing.
 - Preserving data processing and analysis in a way that does not compromise the privacy of individuals.
- Make regulatory compliance cheaper and easier
 - 'data protection by design and by default' approach
- Safeguard personal data from unauthorized disclosure
- Enable data sharing without revealing sensitive information
- Reduce risk of re-identification

Key Issues addressed by PETs

Individual Privacy Protection

Inference / Difference Attacks

- Re-identification risk when sharing aggregate data
 - Possibility of identifying or inferring sensitive individual information
 - background knowledge and auxiliary information to re-identify individuals and extract sensitive information.
- ✓ Ensure that the inclusion or exclusion of a single individual's data does not significantly affect the outcome of the analysis,
 - ✓ making it difficult to infer any specific individual's data.





Balancing Privacy and Utility

Privacy Protection and Legal Requirements
 Data Sharing Restrictions
 Reduced accuracy and utility

- ✓ Rigorous / mathematical frameworks that allow for a tunable balance between privacy protection and data utility,
- Enable the use of data for meaningful analysis while protecting individual privacy.



Data Sharing and Collaboration

Sharing data between organizations or with the public can <u>risk</u>
<u>exposing sensitive information</u>.

- Data sharing and collaboration can introduce risks such as <u>data</u> <u>breaches, misuse, or unauthorized access</u>.
- Enables secure data sharing by ensuring that shared data does not compromise individual privacy, fostering collaboration while maintaining confidentiality.

 Secure environment that mitigates these risks by implementing strict access controls, auditing capabilities, and data
 processing restrictions.



Regulatory Compliance

Privacy regulations

- Require robust protection of personal data.
- Impose strict controls on how personal data is handled, shared, and processed.

Need for rigorous, quantifiable, and verifiable approach to privacy protection

Data Integration from Multiple Sources



- ➢Combining data from different sources (e.g., first-party, secondparty, third-party data) for analysis can be complex and risky in terms of privacy.
- Need to enhance data analytics capabilities by leveraging external data, but sharing detailed data can compromise privacy.
- ✓ Facilitate the integration and analysis of data from multiple sources without exposing the underlying data, allowing for richer insights and more comprehensive analyses.



Data Monetization

- Monetize the data assets while protecting customer privacy and maintaining data security.
- ➢Offer insights and aggregated data analysis services without directly selling or exposing raw data, thereby monetizing data while protecting privacy.
- ✓Allow multiple parties to run analyses on combined datasets without directly sharing the raw data, thus maintaining the privacy and security of each party's data.



Overview of PETs / PPTs

Key Technologies and Methods

Differential Privacy	Secure Multi-party Computation
Data Clean Rooms	Zero-Knowledge Proofs (ZKPs)
Federated Learning	Confidential Computing
	Homomorphic Encryption

Each technology works on specific set of problems / scenarios Heavy on customization – tailored solutions Lack of off-the-shelf solutions Consume Compute Resources – cost prohibitive Combine various methods to meet specific goals Must be considered during product / service design Require commitment and upfront investment with sometimes difficult to quantify immediate benefits

Homomorphic Encryption



Allows computations analytical functions to be performed on encrypted data without needing to decrypt it first. Protects against theft & surveillance.. Eliminates need for local copies.



Fully Homomorphic Encryption Libraries



- HElib Optimized for performance C++ library that implements various schemes, including BGV. HElib GitHub
- PALISADE Comprehensive open-source library for lattice cryptography, including homomorphic encryption. Supports several schemes such as BGV, BFV, and CKKS. <u>PALISADE GitHub</u>
- SEAL (Microsoft SEAL) Easy-to-use, efficient library. Supports the BFV and CKKS schemes and is designed to be highly accessible for developers with extensive documentation and examples. <u>Microsoft SEAL GitHub</u>
- TFHE Fast fully homomorphic encryption over the torus. Focuses on gate bootstrapping techniques enabling practical FHE operations at a bit level. <u>TFHE GitHub</u>
- OpenFHE Relatively recently development that builds on the work of PALISADE, HElib, and SEAL. Flexible and extensible platform for FHE and other lattice-based cryptographic schemes. <u>OpenFHE GitHub</u>



Homomorphic Encryption

Implementation Steps

- Select Homomorphic Encryption Library to support required operations
- Prepare and Test Environment
 - Encryption and Decryption
 - Handling of Keys
 - Analytic Operations
 - Sensitivity to Noise \rightarrow accumulation of noise as operations get layered
 - Noise Management
 - Performance guardrails \rightarrow results available within acceptable window
- Implement Client Side Data Encryption; Server Side Analytics
- Setup monitoring & quality controls



Duality's FHE Example

Step 1

Configure model parameters and data and send to Hub

Step 2

Distribute the parameters and data requirements to Data Owners

Step 3

Pre-process and prepare the data for encrypted computation

Step 4

Encrypt intermediate results and send it to compute

Step 5

Encrypted aggregation of intermediate results

Step 6

Results are decrypted and actioned



Homomorphic Encryption in Machine Learning

Opportunities & Challenges

Ref: FB's Cheetah Paper

Cheetah: Optimizing and Accelerating Homomorphic Encryption for Private Inference

Brandon Reagen^{*1,2}, Wooseok Choi^{*3}, Yeongil Ko⁴, Vincent T. Lee⁵ Hsien-Hsin S. Lee², Gu-Yeon Wei⁴, David Brooks⁴

*Equal contribution New York University¹, Facebook AI Research², Seoul National University³ Harvard University⁴, Facebook Reality Labs Research⁵

Abstract— As the application of deep learning continues to grow, so does the amount of data used to make predictions. While traditionally, big-data deep learning was constrained by computing performance and off-chip memory bandwidth, a new constraint has emerged: privacy. One solution is homomorphic encryption (HE). Applying HE to the client-cloud model allows cloud services to perform inference directly on the client's encrypted data. While HE can meet privacy constraints, it introduces enormous computational challenges and remains impractically slow in current systems.

This paper introduces Cheetah, a set of algorithmic and hardware optimizations for server-side HE DNN inference to approach real-time speeds. Cheetah proposes HE-parameter tuning optimization and operator scheduling optimizations, which together deliver $79 \times$ speedup over state-of-the-art. However, this still falls short of real-time inference speeds by almost four orders of magnitude. Cheetah further proposes an accelerator architecture, when combined with the algorithmic optimizations, to bridge the remaining performance gap. We evaluate several DNNs and show that privacy-preserving HE inference for ResNet50 can be done at near real-time performance with an accelerator dissipating 30W and 545mm² in 5nm.

TABLE I: Generalization of	privacy-preserving	techniques
----------------------------	--------------------	------------

Solution	Security	Limitation
Local	System	Edge performance; leaks model
TEE	System	Performance; side-channels
DP	Statistical	Applications; utility-privacy tradeoff
MPC	Cryptographic	Communication bandwidth
HE	Cryptographic	Compute

(DP) [4], [13], [18], [21], (ii) secure multi-party compute (MPC) [32], [36], [48], [49], and (iii) homomorphic encryption (HE) [7], [28], [31], [53]. Table I summarizes the techniques and limitations associated with each with respect to wide-scale deployment.

Each of the above solutions have differing limitations. Local execution offers individual users improved security, but there is risk of sensitive information leaking or being stolen through the model, plus model-privacy concerns for service providers [60]. TEEs have been shown to be vulnerable to side-channel attacks [15]. DP offers statistical privacy levels



Homomorphic Encryption

Data Privacy in Computation	Processing by third-party services, creating risks of exposure or breaches. Information remains secure and private throughout the processing
Secure Data Outsourcing	Data storage and computation using cloud services Processing without exposing the raw data to the service providers
Regulatory Compliance	Strict protection of personal data - including protection during processing Data remains encrypted during processing
Collaboration on Sensitive Data	Joint data analysis between organizations or departments can risk exposing sensitive information Collaborative computations on shared data without revealing the underlying sensitive information to any party involved
Data Monetization with Privacy	Monetize data assets without compromising user privacy

Limitations: Works best with structured data; Computationally expensive



Confidential Computing

- Focus on the security and privacy of **data in use** (not just at-rest and at-transit)
- Protection of sensitive computing and data elements from
 - Own operators and software
 - Cloud Providers and their system software
- Model training with private data
- Distrusting entities to pool their data to train models or use
- Trusted Execution Environment (TEE):
 - Secure area of a processor that guarantees confidentiality and integrity for code and data
 - Data are encrypted and only accessible from within the protected environments.
- Secure Enclaves:
 - A form of TEE that provides an isolated execution environment to protect the data and code from being accessed or tampered with by other parts of the system, including the operating system, hypervisors, and administrators.



Source: NVIDIA



Confidential Computing Technologies

- Intel SGX (Software Guard Extensions):
 - Provides hardware-based memory encryption to isolate specific application code and data in memory.
- AMD SEV (Secure Encrypted Virtualization)
 - Encrypts virtual machine memory to protect against unauthorized access from the hypervisor.
- Microsoft Azure
 - Virtual machines and services using Intel SGX to protect data during processing in the cloud.
- Google Cloud
 - Confidential VMs using AMD SEV to encrypt data in use
- AWS Nitro
 - Nitro EC2 Instances specialized h/w, chips
 - AWS Nitro Enclaves isolated compute environment for multi-party compute
- NVIDIA Hopper / H100
 - Supports GPU assignments to VMs in various ways

Zero-Knowledge Proofs (ZKPs)



- Cryptographic technique that enables a prover to assure a verifier that a statement is true without revealing any additional information
- Enable a person (the prover) to assure a verifier that it meets contractual, legal, or other requirements without revealing underlying information such as contract terms, financial data, and protected information
- Improves compliance with data laws and data agreements enabling parties to interact and transact without sharing protected or confidential information
- Use Cases
 - Confidential Data Sharing
 - Private Data Queries
 - Authentication without Revealing Identity
 - Secure Multi-Party Computation (SMPC)

Federated Learning

- Eliminates need for centralized collection of training data
- Solves for regulatory & legal restrictions on sharing of raw sensitive information
- Central server sends a copy of the partially ٠ trained model to each participating organization and collects model updates instead of raw data
- Some risk of data leak from targeted attacks on ٠ model updates & trained models - relatively smaller risk compared to central collection of training data.
- Open-source frameworks for federated learning
 - FATE, Flower, PySyft, OpenFL, FedML, NVFlare, and ٠ **Tensorflow Federated**
- Mobile and IoT devices make great candidates •
- Famous implementations Google Keyboard, ٠ Speech, Apple

Randomly selected 100s / 1000s out of millions of devices for each round of model training







Secure Multi-party Computation

- Motivations for SMPC:
 - Each participant's data must remain confidential not shared with other participants.
 - Computed result is correct and verifiable.
 - Each participant inputs their data independently, without knowing the inputs of the other participants.
 - No Collusion even if some participants collude, they cannot deduce the inputs of the other honest participants.
- Use Cases
 - Privacy-Preserving Data Analysis
 - Secure Voting Systems
 - Private Auctions
 - Joint Financial Computations



SMPC Workflow and Protocols

- Implementing SMPC
 - Secret Sharing: Chunking and sharing in such a way that no individual share reveals any information about the original input.
 - Local Computation: Participants perform computations on their shares locally.
 - Recombination: Result without revealing the individual inputs.
 - Data Clean Room
- SMPC Protocols:
 - Yao's Garbled Circuits
 - Secret Sharing Schemes
 - Homomorphic Encryption
 - Oblivious Transfer
 - Private Set Intersection

Differential Privacy – Key Concepts



- Ensures that the inclusion or exclusion of a single individual's data does not significantly affect the outcome of the analysis
- Makes it difficult to infer any specific individual's data from disclosed aggregated results

Formal definition: Query q is ε-differentially private if, for any two databases d1 and d2 that differ in a single row, and any set of outputs R :

$\Pr[q(d1) \in R] \leq e^{\varepsilon} \cdot \Pr[q(d2) \in R] + \delta$

- q query, algorithm, function etc. d dataset. R results
- ε is privacy loss parameter aka privacy budget quantifies the allowable deviation between data disclosures and the most accurate real-world scenarios, effectively measuring the impact of an individual's data on the analysis's outcome.
- Smaller ε value indicates stronger privacy protection, reducing the risk to individual privacy,
- Larger ε suggests weaker protection, increasing potential privacy risks.
- ε of 0 would mean perfect privacy, indicating no increase in privacy risk for any individual from the analysis.
- δ The probability that any row in the result fails to be epsilon-differentially private0

* Compare & contrast w/ Synthetic Data



Differential Privacy Libraries

Google's Differential Privacy Library	Easy integration with various data analysis tools, customizable privacy parameters. Suitable for large-scale data processing tasks.
PyDP (Python Differential Privacy)	Python wrapper for Google's differential privacy library.
PySyft	Extends PyTorch and TensorFlow to enable differential privacy, federated learning, and multi-party computation.
IBM Differential Privacy Library	Supports various differential privacy mechanisms, easy integration with data science workflows. Useful for enterprise-level applications requiring robust privacy guarantees.
OpenDP & SmartNoise	A collaboration between Microsoft and Harvard, SmartNoise provides tools to enable privacy-preserving data analytics.
TensorFlow Privacy	TensorFlow Privacy extends TensorFlow to enable training of machine learning models with differential privacy.



Data Clean Rooms



- Available from Cloud Providers as product grade service
- May offer existing integrations and useful data ready for consumption
- Open-source or proprietary implementation of Confidential Computing frameworks
- Leading providers of Data Clean Rooms: AWS, Snowflake, LiveRamp, Google Ads Data Hub, InfoSum, Aqilliz



Disclosure Avoidance System

- Techniques and processes designed to protect individual privacy when releasing data collected from surveys, censuses, or other data processing efforts
- Critical for organizations that publish information that cannot be used to identify specific individuals
- Maintaining data utility with privacy protection. Ensuring compliance.
- Key Components of a DAS
 - Data Anonymization, Perturbation, Suppression, Aggregation
 - Qualified, Controlled Access
 - Limited / Restricted Data Transfer API/ Query-based mechanisms
 - Privacy-preserving Analytics



Opportunities & Challenges

- Each technology works on specific set of problems / scenarios
- Heavy on customization tailored solutions
- Lack of off-the-shelf solutions
- Consume Compute Resources cost prohibitive
- Combine various methods to meet specific goals
- Must be considered during product / service design
- Require commitment and upfront investment with sometimes difficult to quantify immediate benefits





References

- Intro to Homomorphic Encryption ; How would you explain homomorphic encryption?
- autodp: Automating differential privacy computation
- Differentially Private Analytics at Scale
- PRIVIC: A privacy-preserving method for incremental collection of location data
- <u>UK ICO PET Recommendations</u>
- FLOWER: A Friendly Federated Learning Framework
- <u>Google Differential Privacy Library</u>
- Analyzing Leakage of Personally Identifiable Information in Language Models
- EzPC: Easy Secure Multiparty Computation
- <u>A Pragmatic Introduction to Secure Multi-Party Computation</u>
- Zero Knowledge Proofs: An illustrated primer
- <u>Google Differential Privacy Library</u>
- <u>PySyft DP Library</u>
- <u>Designing Access with Differential Privacy</u> MIT Handbook of Using Admin Data for Research
- Apple's Differential Privacy ; Apple's Learning with Privacy at Scale
- Google's Big Query Differential Privacy; Difference Privacy Clause example
- <u>Secure multi-party computation</u>



Thank you!