Secure and Private AI: Differential Privacy and Federated Learning

(A course by Facebook AI offered via Udacity)

Seyed Iman Mirzadeh

Washington State University

seyediman.mirzadeh@wsu.edu

August 2019

Overview

1. Differential Privacy

2. Federated Learning

About The Course

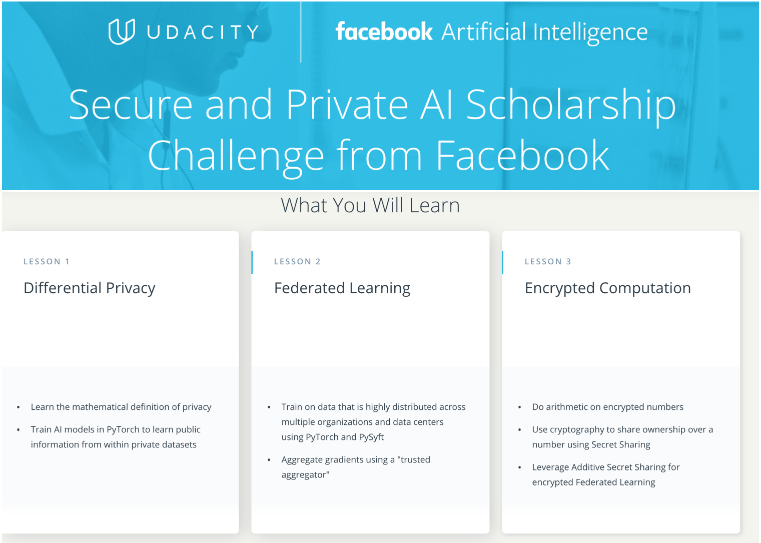

UDACITY | facebook Artificial Intelligence

Secure and Private AI Scholarship Challenge from Facebook

What You Will Learn

LESSON 1

Differential Privacy

- Learn the mathematical definition of privacy
- Train AI models in PyTorch to learn public information from within private datasets

LESSON 2

Federated Learning

- Train on data that is highly distributed across multiple organizations and data centers using PyTorch and PySyft
- Aggregate gradients using a "trusted aggregator"

LESSON 3

Encrypted Computation

- Do arithmetic on encrypted numbers
- Use cryptography to share ownership over a number using Secret Sharing
- Leverage Additive Secret Sharing for encrypted Federated Learning

What is Differential Privacy?

Let's start by the goal of DP:

DP Goal

* We have a dataset(or database) and we run some analysis(e.g., statistical analysis)
* BUT, we want to make sure our analysis doesn't compromise the privacy of any particular individual within that dataset.

Therefore, in order to accomplish our goal, we need:

- A robust definition of privacy!
- Then, definition of compromising/preserving that privacy

Definition of Differential Privacy

‘Differential privacy‘ describes a promise, made by a data holder to a data subject: ”You will not be affected, adversely or otherwise, by allowing your data to be used in any study or analysis, no matter what other studies, data sets, or information sources, are available.”

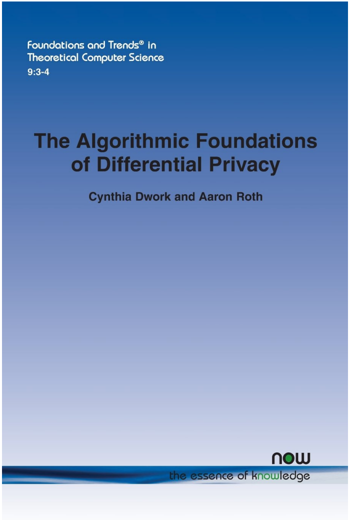

Foundations and Trends® in
Theoretical Computer Science
9:3-4

The Algorithmic Foundations of Differential Privacy

Cynthia Dwork and Aaron Roth

now
the essence of knowledge

Differential Privacy by Example

- Suppose we want have a dataset of people responding "Yes/No" to the question "Have you ever driven under the influence?"

Name	DUI?
John	Yes
Jack	No
Jennifer	Yes
James	No

Table: Dataset

Differential Privacy by Example

- Suppose we want have a dataset of people responding "Yes/No" to the question "Have you ever driven under the influence?"

Name	DUI?
John	Yes
Jack	No
Jennifer	Yes
James	No

Table: Dataset

- Naturally, if we publish the real dataset, we'll be in a big trouble!

Differential Privacy by Example

- Suppose we want have a dataset of people responding "Yes/No" to the question "Have you ever driven under the influence?"

Name	DUI?
John	Yes
Jack	No
Jennifer	Yes
James	No

Table: Dataset

- Naturally, if we publish the real dataset, we'll be in a big trouble!
- However, one technique that is used in DP is to add noise. Here's how we secure the data for each row:

Differential Privacy by Example

- Suppose we want have a dataset of people responding "Yes/No" to the question "Have you ever driven under the influence?"

Name	DUI?
John	Yes
Jack	No
Jennifer	Yes
James	No

Table: Dataset

- Naturally, if we publish the real dataset, we'll be in a big trouble!
- However, one technique that is used in DP is to add noise. Here's how we secure the data for each row:
  - Flip a coin.

Differential Privacy by Example

- Suppose we want have a dataset of people responding "Yes/No" to the question "Have you ever driven under the influence?"

Name	DUI?
John	Yes
Jack	No
Jennifer	Yes
James	No

Table: Dataset

- Naturally, if we publish the real dataset, we'll be in a big trouble!
- However, one technique that is used in DP is to add noise. Here's how we secure the data for each row:
  - Flip a coin.
  - If tails, don't change the real value

Differential Privacy by Example

- Suppose we want have a dataset of people responding "Yes/No" to the question "Have you ever driven under the influence?"

Name	DUI?
John	Yes
Jack	No
Jennifer	Yes
James	No

Table: Dataset

- Naturally, if we publish the real dataset, we'll be in a big trouble!
- However, one technique that is used in DP is to add noise. Here's how we secure the data for each row:
  - Flip a coin.
  - If tails, don't change the real value
  - If heads, then flip a second coin and respond Yes if heads and No if tails.

Differential Privacy by Example(2)

Name	DUI?
John	Yes
Jack	No
Jennifer	Yes
James	No

Table: Dataset

- Each individual has a degree of "plausible deniability": If we have in our dataset that "John" has driven under the influence, we can't be sure that he really did this. Because there's a chance that our first coin came "heads" and thus this row is random.

Differential Privacy by Example(2)

Name	DUI?
John	Yes
Jack	No
Jennifer	Yes
James	No

Table: Dataset

- Each individual has a degree of "plausible deniability": If we have in our dataset that "John" has driven under the influence, we can't be sure that he really did this. Because there's a chance that our first coin came "heads" and thus this row is random.
- Basically, 50% of data is real and the other 50% is random.

Differential Privacy by Example(2)

Name	DUI?
John	Yes
Jack	No
Jennifer	Yes
James	No

Table: Dataset

- But we can calculate "number of actual people who committed DUI" without actually knowing whether or not any individual committed DUI. If p is the true fraction of people who committed DUI:

$$\text{Num "Yes" answers} = \frac{1}{4}(1-p) + \frac{3}{4}p$$

Differential Privacy by Example(2)

Name	DUI?
John	Yes
Jack	No
Jennifer	Yes
James	No

Table: Dataset

- But we can calculate "number of actual people who committed DUI" without actually knowing whether or not any individual committed DUI. If p is the true fraction of people who committed DUI:

$$\text{Num "Yes" answers} = \frac{1}{4}(1-p) + \frac{3}{4}p$$

- Randomization is essential; more precisely, any non-trivial privacy guarantee that holds regardless of all present or even future sources of auxiliary information, requires randomization.

Probability Simplex and Randomized Algorithm

Definition 2.1 (Probability Simplex). Given a discrete set B, the *probability simplex* over B, denoted $\Delta(B)$ is defined to be:

$$\Delta(B) = \left\{ x \in \mathbb{R}^{|B|} : x_i \geq 0 \text{ for all } i \text{ and } \sum_{i=1}^{|B|} x_i = 1 \right\}$$

Definition 2.2 (Randomized Algorithm). A randomized algorithm $\mathcal{M}$ with domain A and discrete range B is associated with a mapping $M : A \to \Delta(B)$. On input $a \in A$, the algorithm $\mathcal{M}$ outputs $\mathcal{M}(a) = b$ with probability $(M(a))_b$ for each $b \in B$. The probability space is over the coin flips of the algorithm $\mathcal{M}$.

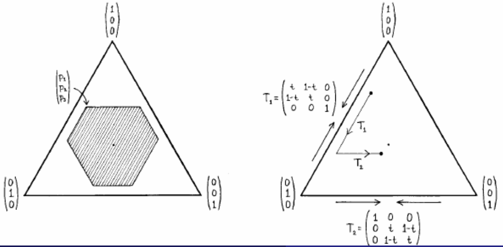
$\begin{pmatrix}1\\0\\0\end{pmatrix}$
$\begin{pmatrix}p_1\\p_2\\p_3\end{pmatrix}$
$\begin{pmatrix}0\\1\\0\end{pmatrix}$
$\begin{pmatrix}0\\0\\1\end{pmatrix}$
$\begin{pmatrix}1\\0\\0\end{pmatrix}$
$T_1 = \begin{pmatrix}t & 1-t & 0\\1-t & t & 0\\0 & 0 & 1\end{pmatrix}$
T_1
T_2
$\begin{pmatrix}0\\1\\0\end{pmatrix}$
$\begin{pmatrix}0\\0\\1\end{pmatrix}$
$T_2 = \begin{pmatrix}1 & 0 & 0\\0 & t & 1-t\\0 & 1-t & t\end{pmatrix}$

Formalized Definition of Differential Privacy

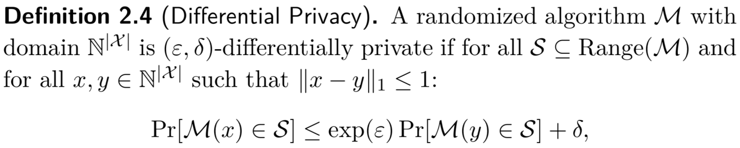

Definition 2.4 (Differential Privacy)**.** A randomized algorithm $\mathcal{M}$ with domain $\mathbb{N}^{|\mathcal{X}|}$ is (ε, δ)-differentially private if for all $\mathcal{S} \subseteq \mathrm{Range}(\mathcal{M})$ and for all $x, y \in \mathbb{N}^{|\mathcal{X}|}$ such that $\|x - y\|_1 \leq 1$:

$$\Pr[\mathcal{M}(x) \in \mathcal{S}] \leq \exp(\varepsilon) \Pr[\mathcal{M}(y) \in \mathcal{S}] + \delta,$$

It means, if we remove individuals, the answer distribution will remain the same by a factor of e^{ϵ}

The smaller the ϵ, the more accurate our analysis will be. But, the more data we also leak.

The ultimate goal of Differential Privacy algorithms is to find better randomized methods to reduce ϵ as much as possible while harming the accuracy as little as possible.

Federated Learning

- Federated learning is a response to the question: can a model be trained without the need to move and store the training data to a central location?
- It enables mobile phones to collaboratively learn a shared prediction model while keeping all the training data on device, decoupling the ability to do machine learning from the need to store the data in the cloud

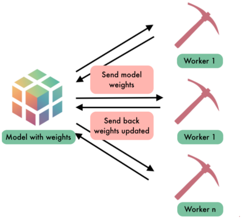

-

Federated Learning General Mechanism

A typical round of learning consists of the following sequence.

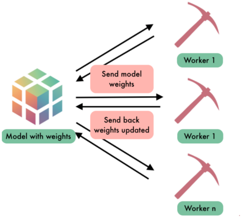

Federated Learning General Mechanism

A typical round of learning consists of the following sequence.

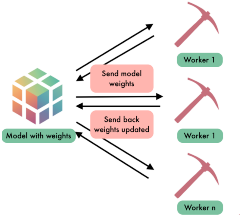

1. A random subset of members of the Federation (known as clients) is selected to receive the global model synchronously from the server.

Federated Learning General Mechanism

A typical round of learning consists of the following sequence.

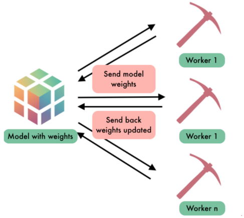

1. A random subset of members of the Federation (known as clients) is selected to receive the global model synchronously from the server.
2. Each selected client computes an updated model using its local data.

Federated Learning General Mechanism

A typical round of learning consists of the following sequence.

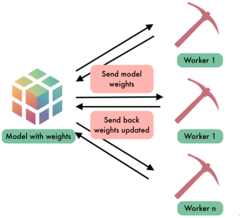

1. A random subset of members of the Federation (known as clients) is selected to receive the global model synchronously from the server.
2. Each selected client computes an updated model using its local data.
3. The model updates are sent from the selected clients to the server.

Federated Learning General Mechanism

A typical round of learning consists of the following sequence.

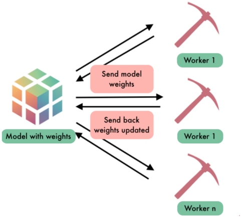

1. A random subset of members of the Federation (known as clients) is selected to receive the global model synchronously from the server.
2. Each selected client computes an updated model using its local data.
3. The model updates are sent from the selected clients to the server.
4. The server aggregates these models (typically by averaging) to construct an improved global model.

Why Differential Privacy and Federated Learning?

1. With not collecting clients' data secure, more clients will trust and more data can be gathered. (FL)

Why Differential Privacy and Federated Learning?

1. With not collecting clients' data secure, more clients will trust and more data can be gathered. (FL)
2. It will be possible to share many private databases after securing them and more data will be there to be analyzed. (DP)

Why Differential Privacy and Federated Learning?

1. With not collecting clients' data secure, more clients will trust and more data can be gathered. (FL)
2. It will be possible to share many private databases after securing them and more data will be there to be analyzed. (DP)
3. By doing the computation on edge devices, the load from servers will reduce and we can run online machine algorithms with low cost. (FL)

Thank You