Efficiently Computing Similarities to Private Datasets

Arturs Backurs
Microsoft Research

Zinan Lin
Microsoft Research

Sepideh Mahabadi
Microsoft Research

Sandeep Silwal
MIT -> Wisconsin Madison

Jakub Tarnawski
Microsoft Research

Hiring students and postdocs!

Motivation

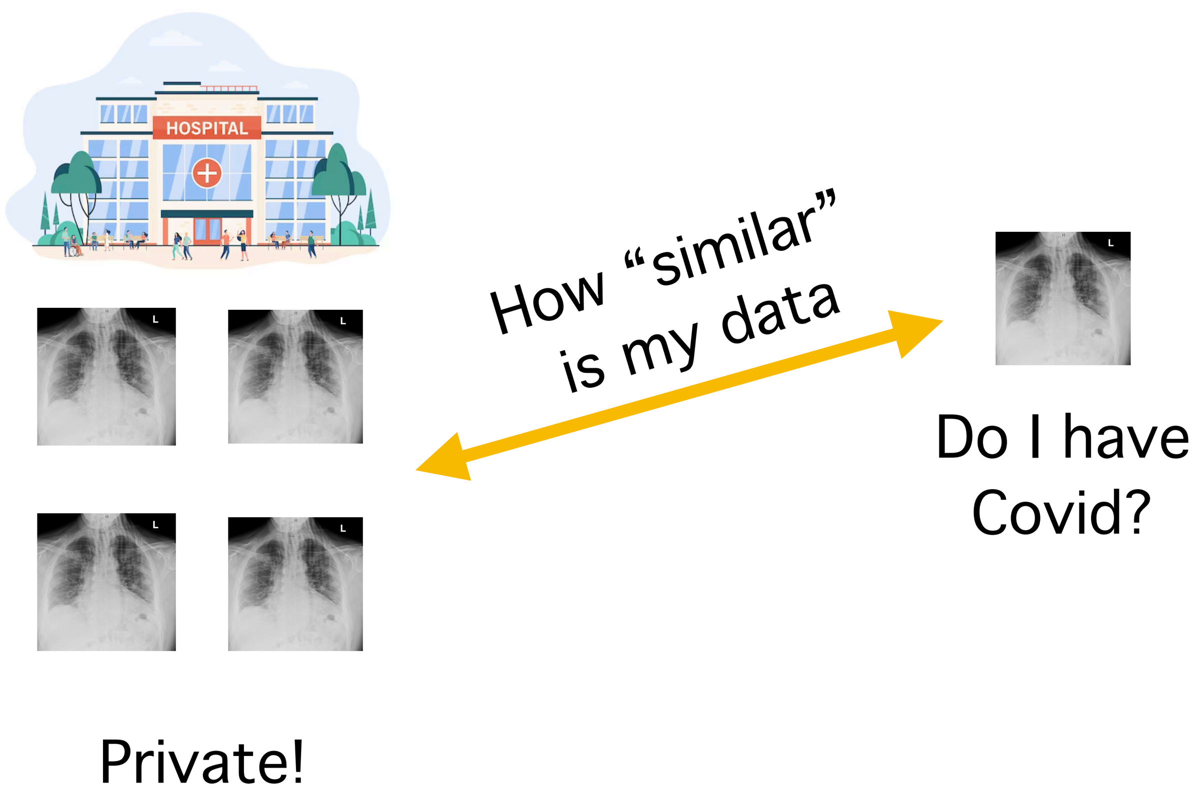

Requirements:

- Preserve 'privacy' of existing patients
- Output an accurate answer

Formal Definition

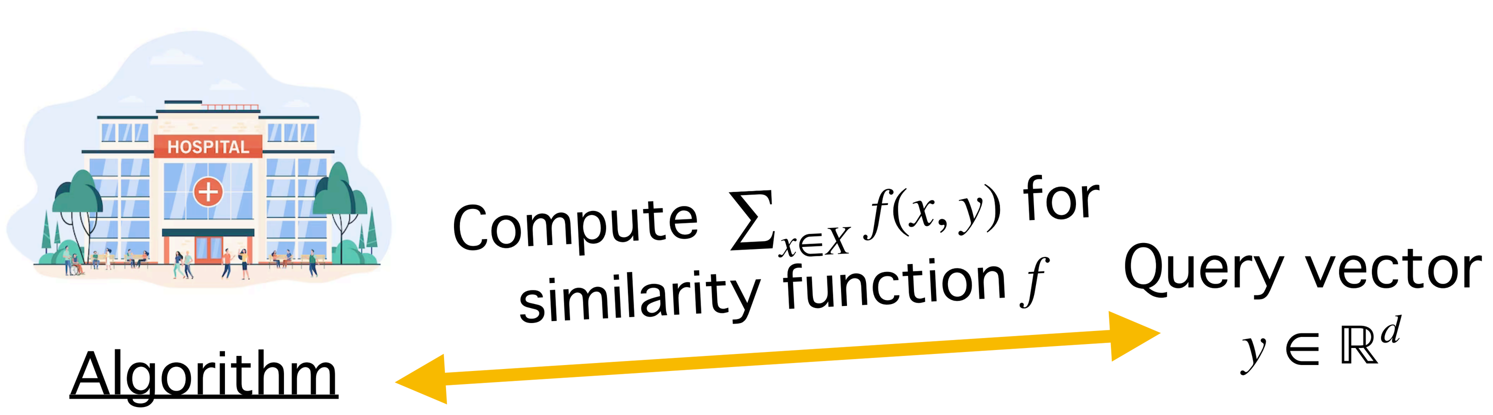

- Takes in input dataset X
- Alg outputs a data structure $\mathcal{D}$
- $\mathcal{D}$ is differential private wrt X
- On any query y, $\mathcal{D}$ approximates sum

f can be a distance function e.g. $\|x-y\|_2$
f can be a kernel function e.g. $\exp(-\|x-y\|_2)$

Our results

What are the tradeoffs between accuracy and privacy?
Privacy measured with respect to (ε, δ)-Differential Privacy

n = Dataset Size
d = Dataset dimension

(M, A) means $\mathbb{E}[\,|\,\text{True} - \mathcal{D}(y)\,|\,] \le M \cdot \text{True} + A$
Hiding ε and log terms

$f(x,y)$	Our Error	Prior Error	Ref.
$\|x-y\|_1$	$\alpha,\ d^{1.5}/\sqrt{\alpha}$	$0,\ \text{poly}(n,d)$	Bounded data [Huang, Roth '14]
$\|x-y\|_2$	$\alpha,\ 1/\alpha^{1.5}$	$0,\ \text{poly}(n,d)$	Bounded data [Huang, Roth '14]
$e^{-\|x-y\|_2}$	$0,\ \alpha$	$0,\ \alpha$	Prior algo slower [Wagner et al. '23]
$\frac{1}{1+\|x-y\|_2}$	$0,\ \alpha$	-	

Sneak Preview of Algorithms

- $\|x-y\|_1 = \sum_i |x_i - y_i| \implies$ Sufficient to solve 1D case
- $e^{-\|x-y\|_2}$: Prove novel dimensionality reduction result which preserves kernel sums
- $\frac{1}{1+\|x-y\|_2}$: Reduce this kernel to the case of $e^{-\|x-y\|_2}$
  - Use function approx. theory to write $\frac{1}{1+z}$ as a sublinear number of terms that look like e^{-z}

Experiments

Private Image Classification

- Use private similarity data structures to assign labels
- Similar measured on embeddings of images
- Embeddings curated from a large public model

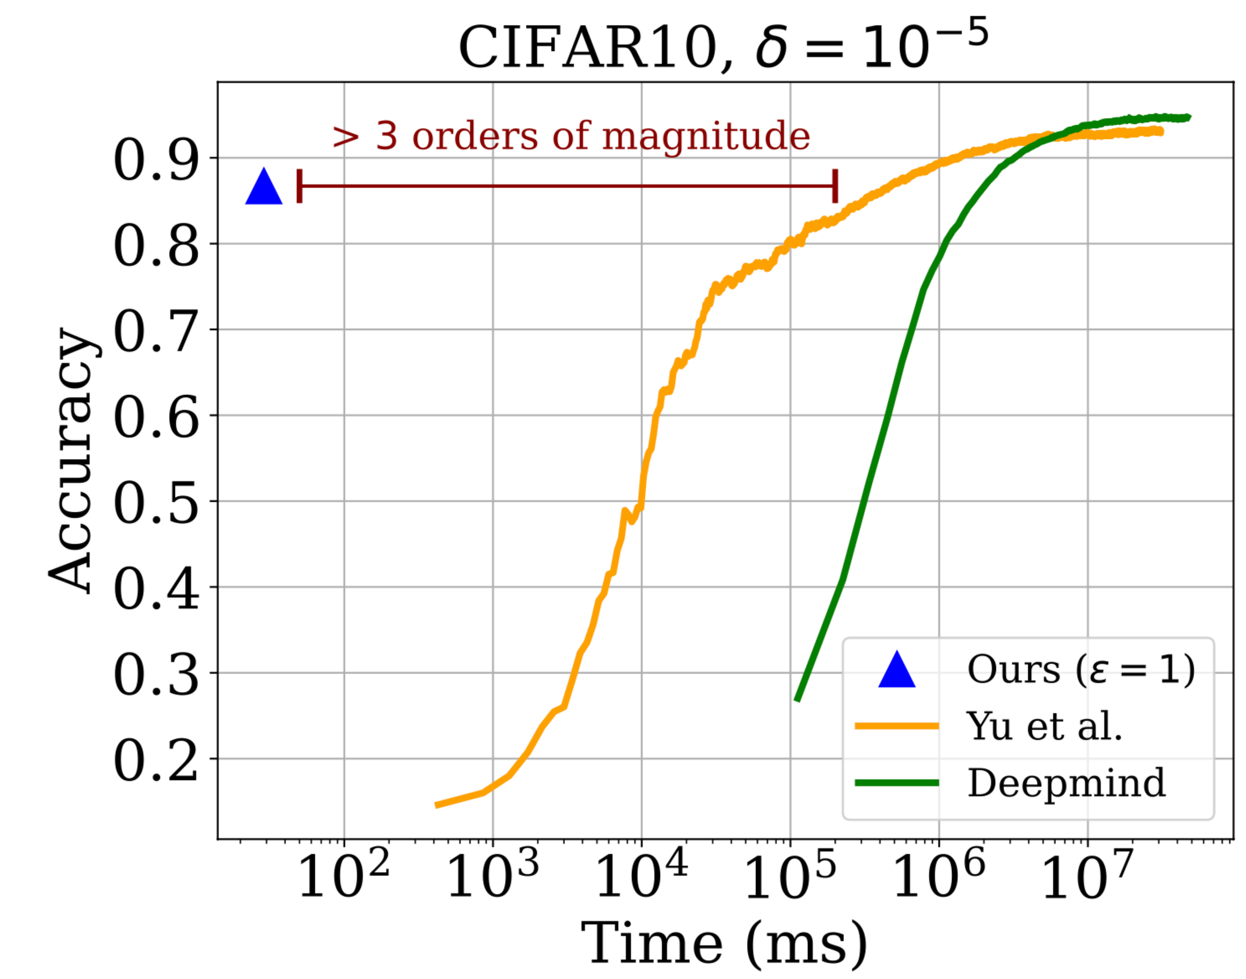

No deep learning required!

$> 10^3$x faster than SOTA (which uses deep learning magic) for comparable acc.