Powerful
data science & analytics
with stronger privacy

How to do AI on personal data while respecting privacy

How to do AI on personal data while respecting privacy: anonymize it

Anonymizing data: a decades old problem

Census administrations have tried to make data public while protecting privacy for decades. With the evolution of privacy regulations, now data protection becomes core to all data processing, making anonymization techniques attractive.

Two main themes have emerged:

- **Truncation**: removing identifying parts from a data point to publish an anonymous dataset
- **Aggregation**: aggregate information so that no individual can be singled out

The anonymization game

Let's try to anonymize a dataset of company wages so that we can learn general properties without being impeded by the sensitiveness of the data (eg: gender pay gap, salary trends

Name	Level	G	Salary
Mark Zuck	CEO	M	$22,554,543
Mike Schroepfer	CTO	M	$19,757,363
Alissa Wang	SW Eng	F	$190,000
Vincent Peters	SW Eng	M	$140,000

⋮

Klay Donald	SW Eng	M	$150,000
Ella Robson	SW Eng	F	$210,000

Anonymizing data: field deletion

Delete the field that lead to re-identification

Name	Level	G	Salary
Mark Zuck	CEO	M	$22,554,543
Mike Schroepfer	CTO	M	$19,757,363
Alissa Wang	SW Eng	F	$190,000
Vincent Peters	SW Eng	M	$140,000

⋮

Klay Donald	SW Eng	M	$150,000
Ella Robson	SW Eng	F	$210,000

Anonymizing data: field deletion

Delete the field that lead to re-identification

Name	Level	G	Salary
Mark Zuck	CEO	M	$22,554,543
Mike Schroepfer	CTO	M	$19,757,363
Alissa Wang	SW Eng	F	$190,000
Vincent Peters	SW Eng	M	$140,000
⋮			
Klay Donald	SW Eng	M	$150,000
Ella Robson	SW Eng	F	$210,000

Name	Level	G	Salary
*	CEO	M	$22,554,543
*	CTO	M	$19,757,363
*	SW Eng	F	$190,000
*	SW Eng	M	$140,000
⋮			
*	SW Eng	M	$150,000
*	SW Eng	F	$210,000

Anonymizing data: field deletion

Delete the field that lead to re-identification

Name	Level	G	Salary
Mark Zuck	CEO	M	$22,554,543
Mike Schroepfer	CTO	M	$19,757,363
Alissa Wang	SW Eng	F	$190,000
Vincent Peters	SW Eng	M	$140,000
⋮			
Klay Donald	SW Eng	M	$150,000
Ella Robson	SW Eng	F	$210,000

Name	Level	G	Salary
*	C-level	M	$22,554,543
*	C-level	M	$19,757,363
*	SW Eng	F	$190,000
*	SW Eng	M	$140,000
⋮			
*	SW Eng	M	$150,000
*	SW Eng	F	$210,000

Anonymizing data: field deletion

Delete the field that lead to re-identification

Name	Level	G	Salary
Mark Zuck	CEO	M	$22,554,543
Mike Schroepfer	CTO	M	$19,757,363
Alissa Wang	SW Eng	F	$190,000
Vincent Peters	SW Eng	M	$140,000
⋮			
Klay Donald	SW Eng	M	$150,000
Ella Robson	SW Eng	F	$210,000

on 1/31	Level	G	Salary
*	SW Eng	F	$190,000
*	SW Eng	M	$140,000

on 2/1	Level	G	Salary
*	SW Eng	F	$194,000
*	SW Eng	M	$140,000

Anonymizing data: field deletion

Delete the field that lead to re-identification

Name	Level	G	Salary
Mark Zuck	CEO	M	$22,554,543
Mike Schroepfer	CTO	M	$19,757,363
Alissa Wang	SW Eng	F	$190,000
Vincent Peters	SW Eng	M	$140,000
⋮			
Klay Donald	SW Eng	M	$150,000
Ella Robson	SW Eng	F	$210,000

on 1/31	Level	G	Salary
*	SW Eng	F	$190,000
*	SW Eng	M	$140,000

on 2/1	Level	G	Salary
*	SW Eng	F	$194,000
*	SW Eng	M	$140,000

What else can we delete?

Anonymizing data: field deletion

Also, you would probably want to use a deeper dataset.

Name	Level	G	Salary	Citizenship	City	DoB	Degree	Performance	Last promo. date
Mark Zuck	CEO	M	$22,554,543	...					
Mike Schroepfer	CTO	M	$19,757,363						
Alissa Wang	SW Eng	F	$190,000						
Vincent Peters	SW Eng	M	$140,000						

Netflix $1m contest on “anonymized” data

Anonymizing data: aggregation

K-anonymity: guarantee that each combination of features has at least k records (achieved by generalizing, grouping features or even removing features)

Name	Level	G	Salary
Mark Zuck	CEO	M	$22,554,543
Mike Schroepfer	CTO	M	$19,757,363
Alissa Wang	SW Eng	F	$190,000
Vincent Peters	SW Eng	M	$140,000
⋮			
Klay Donald	SW Eng	M	$150,000
Ella Robson	SW Eng	F	$210,000

Anonymizing data: aggregation

K-anonymity: guarantee that each combination of features has at least k records (achieved by generalizing, grouping features or even removing features)

Name	Level	G	Salary
Mark Zuck	CEO	M	$22,554,543
Mike Schroepfer	CTO	M	$19,757,363
Alissa Wang	SW Eng	F	$190,000
Vincent Peters	SW Eng	M	$140,000
⋮			
Klay Donald	SW Eng	M	$150,000
Ella Robson	SW Eng	F	$210,000

Level	G	Count	Salary
C-level	M	5	$15,000,000
SW eng	F	120	$180,000
SW eng	M	140	$170,000

Anonymizing data: aggregation

K-anonymity: guarantee that each combination of features has at least k records (achieved by generalizing, grouping features or even removing features)

Name	Level	G	Salary
Mark Zuck	CEO	M	$22,554,543
Mike Schroepfer	CTO	M	$19,757,363
Alissa Wang	SW Eng	F	$190,000
Vincent Peters	SW Eng	M	$140,000
⋮			
Klay Donald	SW Eng	M	$150,000
Ella Robson	SW Eng	F	$210,000

Level 1/31	G	Count	Salary
C-level	M	5	$15,000,000
SW eng	F	120	$180,000
SW eng	M	140	$170,000

Level 2/1	G	Count	Salary
C-level	M	5	$15,000,000
SW eng	F	121	$181,000
SW eng	M	140	$170,000

Differential Privacy, a definition of an anonymous response

Definition

https://www.cis.upenn.edu/~aaroth/Papers/privacybook.pdf

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

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

Differential Privacy

Let's compute a mean

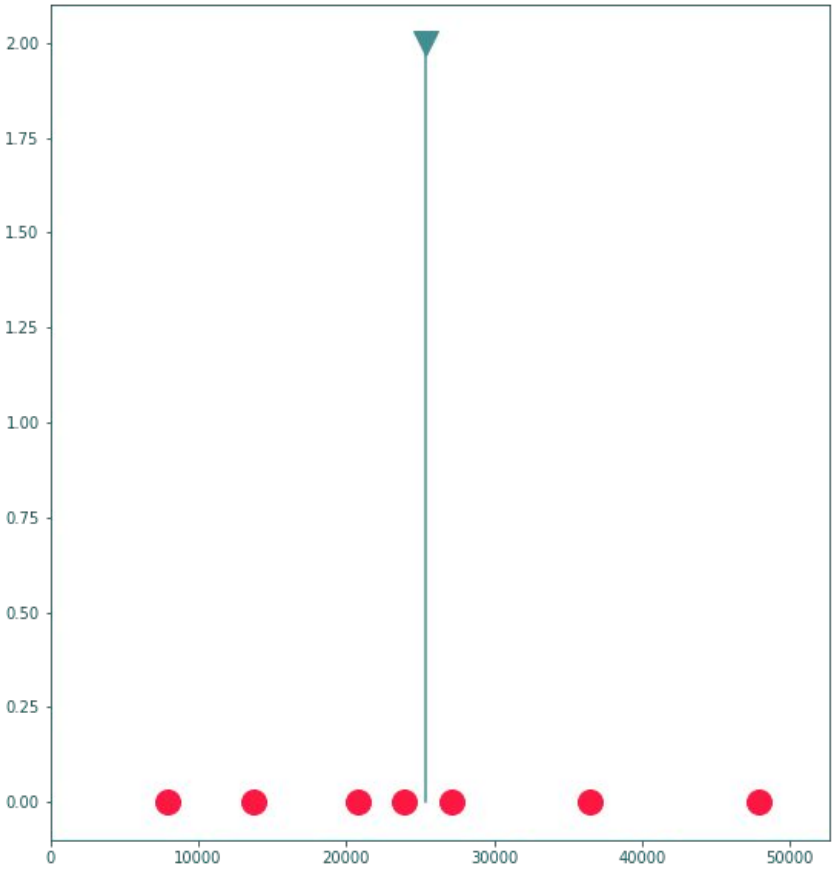

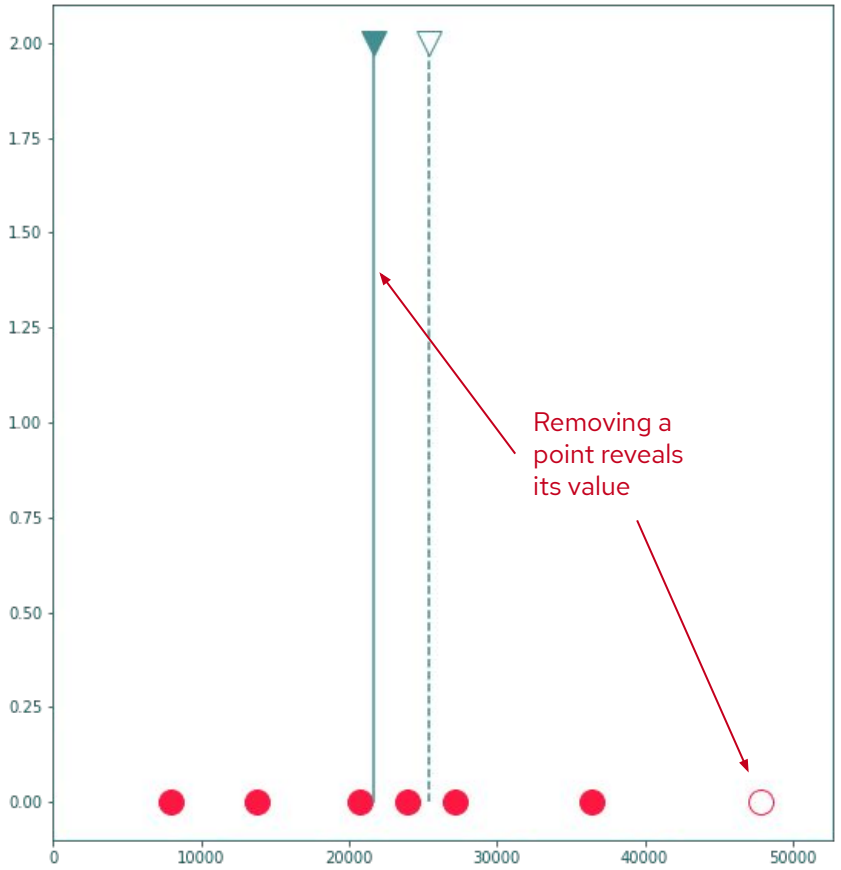

Differential Privacy

Add some noise to prevent inference

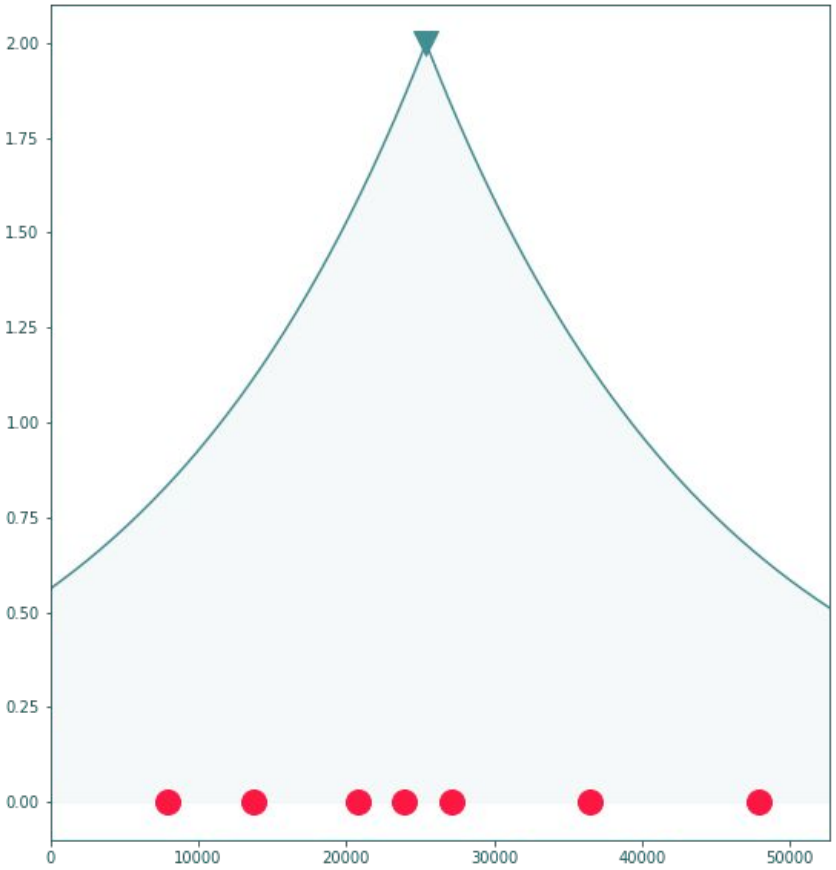

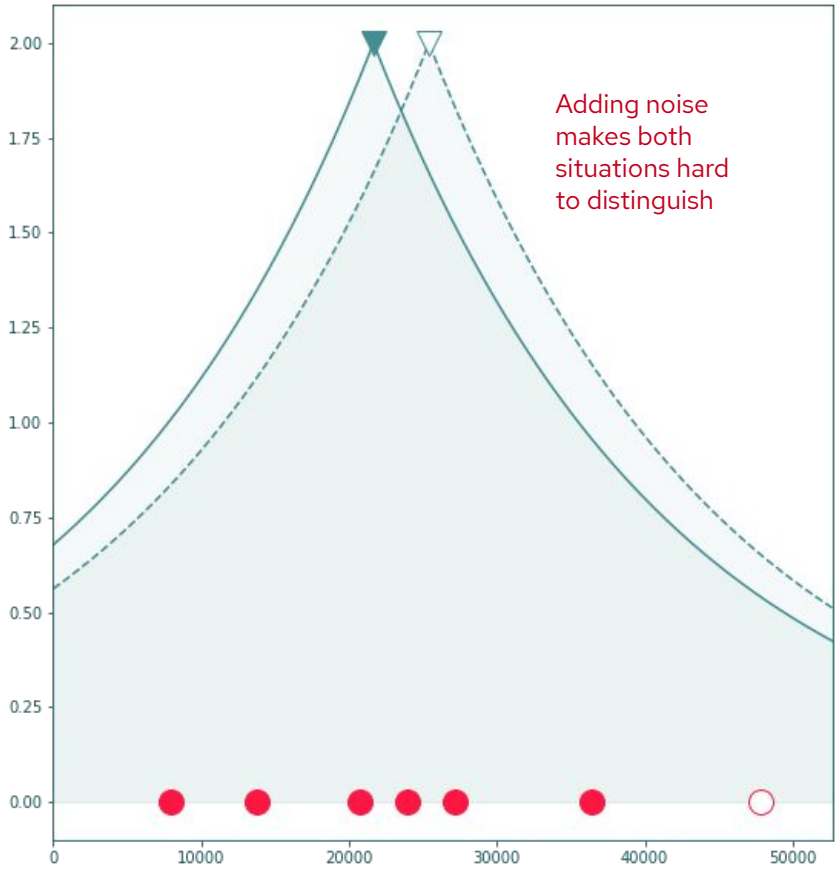

Differential Privacy

$$\frac{P(\bar{A}|X=x)}{P(A|X=x)} = \frac{P(X=x|\bar{A})P(\bar{A})}{P(X=x|A)P(A)} = \frac{P(X=x|\bar{A})}{P(X=x|A)} = \text{privacy loss}$$

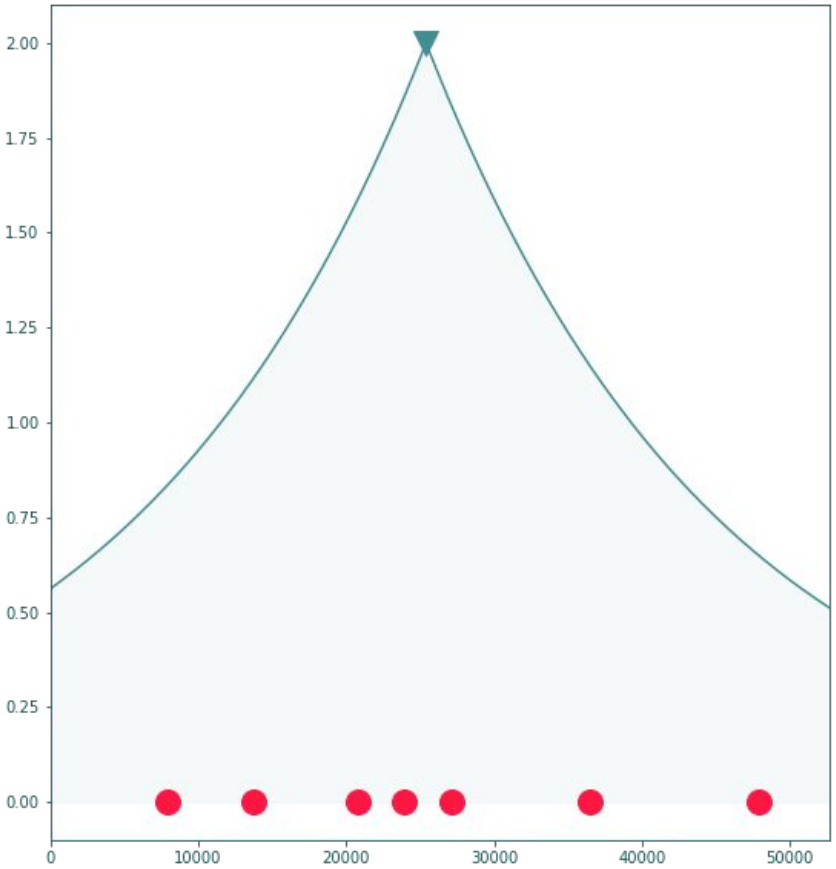

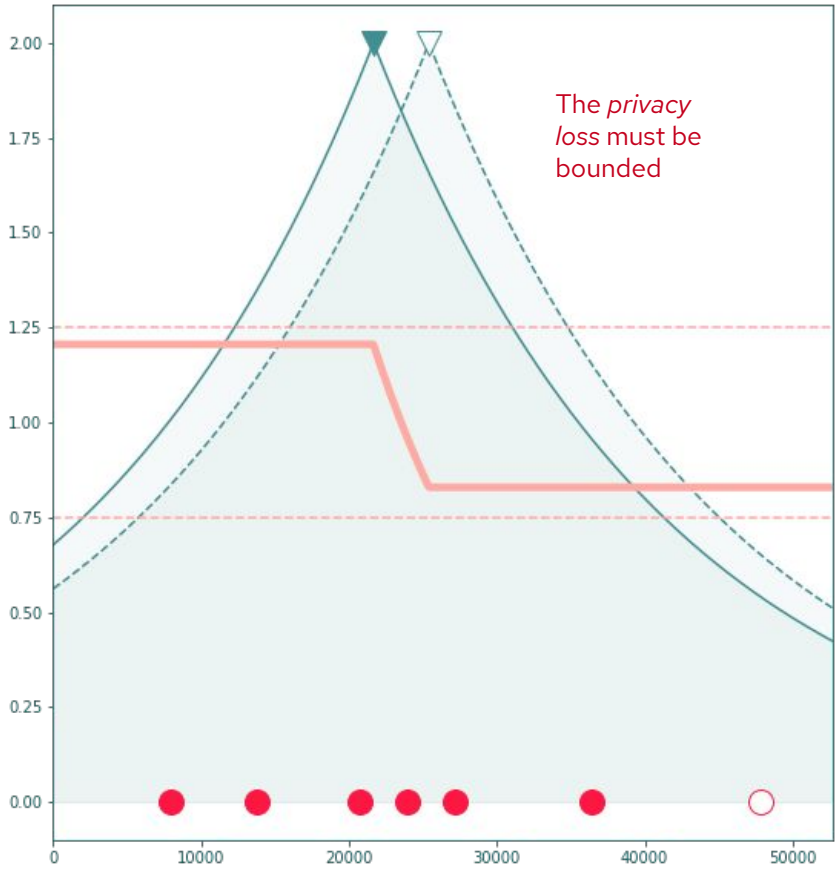

So let's anonymize our data with DP!

Name	Level	G	Salary
Mark Zuck	CEO	M	$22,554,543
Mike Schroepfer	CTO	F	$19,757,363
Alissa Wang	SW Eng	F	$190,000
Vincent Peters	SW Eng	M	$140,000
⋮			
Klay Donald	SW Eng	M	$150,000
Ella Robson	SW Eng	F	$210,000

A WAY OUT: Data scientists work on original data without accessing it.

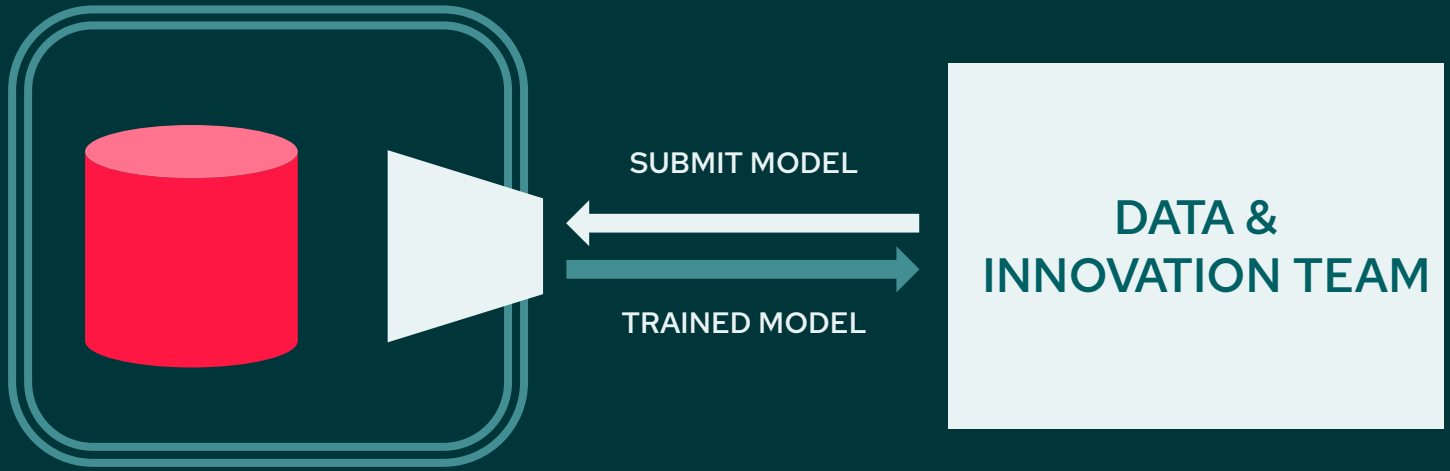

A WAY OUT: Data scientists work on original data without accessing it.

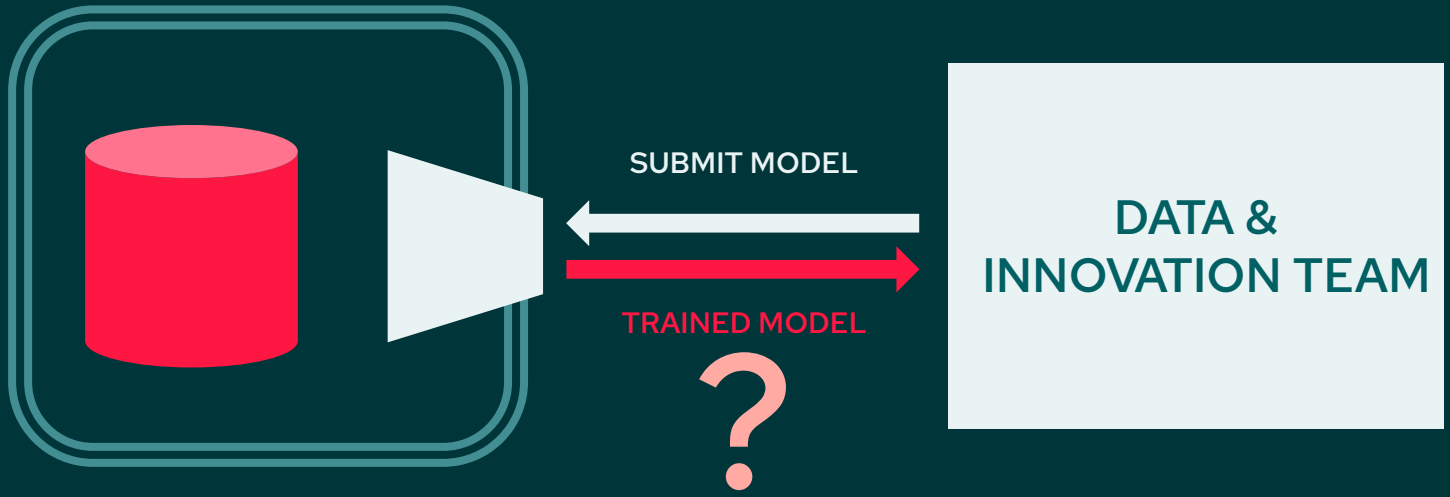

Membership Inference attacks

Name	Date of birth	Comorbidities	Profession	Zip code	Serious covid
XXX	5/7/1962	Cancer	Agriculture	14230	Yes
XXX	8/13/2004	None	Student	75005	No
...	...	...	...	..	...

TRAIN

$$Risk(birthdate, comorbidity, profession, zipcode)$$

How can we ensure that the trained model does not reveal membership of individuals?

With sparse data (which is the case in high dimension training data) and large dimensionality of models, it is likely that the model will “spike” near points from he training set.

Figure 1: An image recovered using a new model inversion attack (left) and a training set image of the victim (right). The attacker is given only the person’s name and access to a facial recognition system that returns a class confidence score.

THE SARUS WAY: Data scientists work on original data without accessing it, results are protected by differential privacy

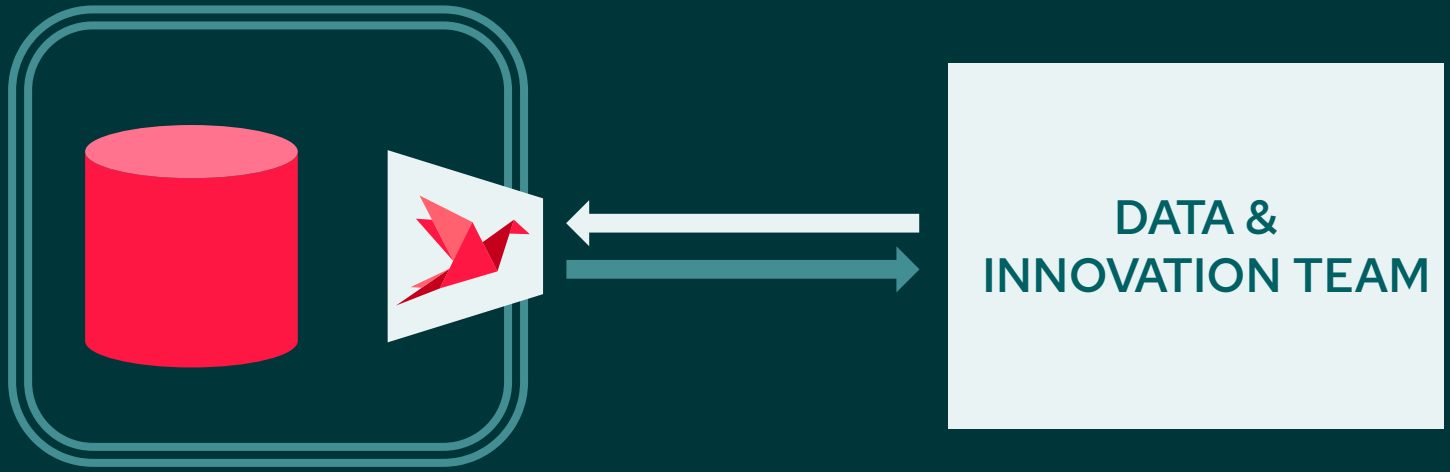

Private data science workflow

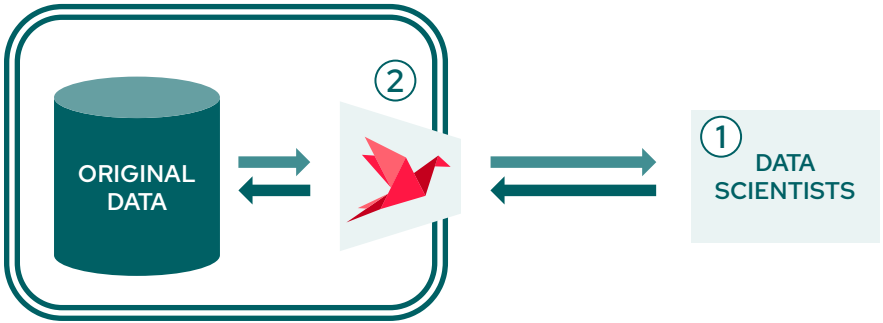

1. Prepare models
2. Submit models for training on original data

AI under DP constraint

The problems considered are non-convex *Empirical Risk Minimization* (ERM) and in particular *Deep Neural Networks* training.

$$\min_{\theta} \sum_{(x,y)\in S} L\left(f_\theta\left(x\right), y\right)$$

Let the risk be:

$$R(\theta) = \mathbb{E}\left[L\left(f_\theta\left(x\right), y\right)\right]$$

and empirical risk on sample S:

$$R_S(\theta) = \frac{1}{|S|} \sum_{(x,y)\in S} L\left(f_\theta\left(x\right), y\right)$$

AI under DP constraint

- Large dimensionality, width, make second-order methods (Newton, Hessian based) impractical,
- Large number of observations, height, make non-stochastic variants, non practical.

The standard way of solving these problems is to use *Stochastic Gradient Descent*.

$$\theta_{t+1} = \theta_t - \rho \sum_{(x,y)\in S_t} \nabla L\left(f_\theta\left(x\right), y\right)$$

$$\theta_{t+1} = \theta_t - \rho \nabla R_{S_t}(\theta)$$

SGD and its variants [Ruder, 2016] is very well suited to ERM on very large datasets.

[Ruder, 2016] Ruder, S. (2016). An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*.

AI under DP constraint: DP-SGD

The privacy losses incurred by each gradient updates can be composed

To provide some privacy guarantees a popular approach is to add some noise to the gradient descent.

$$\theta_{t+1} = \theta_t - \rho \frac{N}{L} \nabla R_{S_t}(\theta) + \sqrt{2\rho\sigma}\varepsilon_t \tag{2}$$

In their famous paper [Abadi et al., 2016], Abadi et al. follow this approach. The update process is the following:

[Abadi et al., 2016] Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., and Zhang, L. (2016). Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 308–318.

AI under DP constraint: DP-SGD

DP-SGD is tested on MNIST and CIFAR-10 with reasonable results

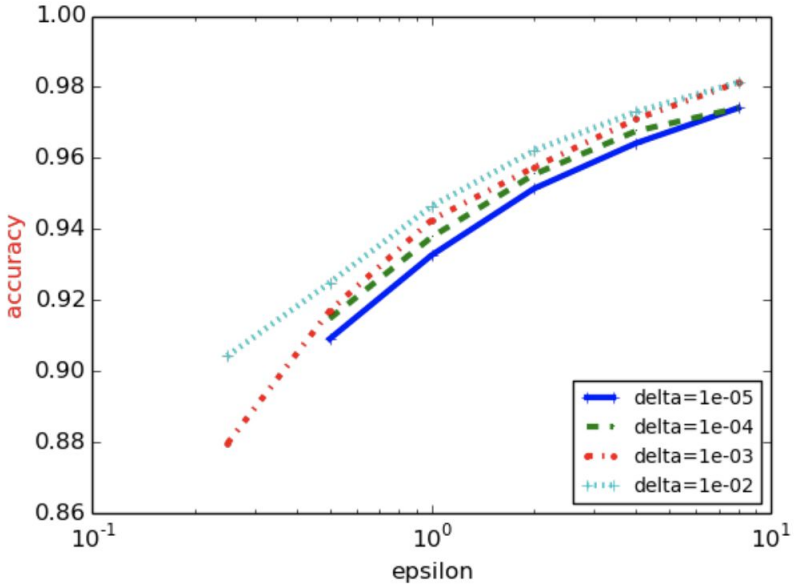

Figure 4: Accuracy of various (ε, δ) privacy values on the MNIST dataset. Each curve corresponds to a different δ value.

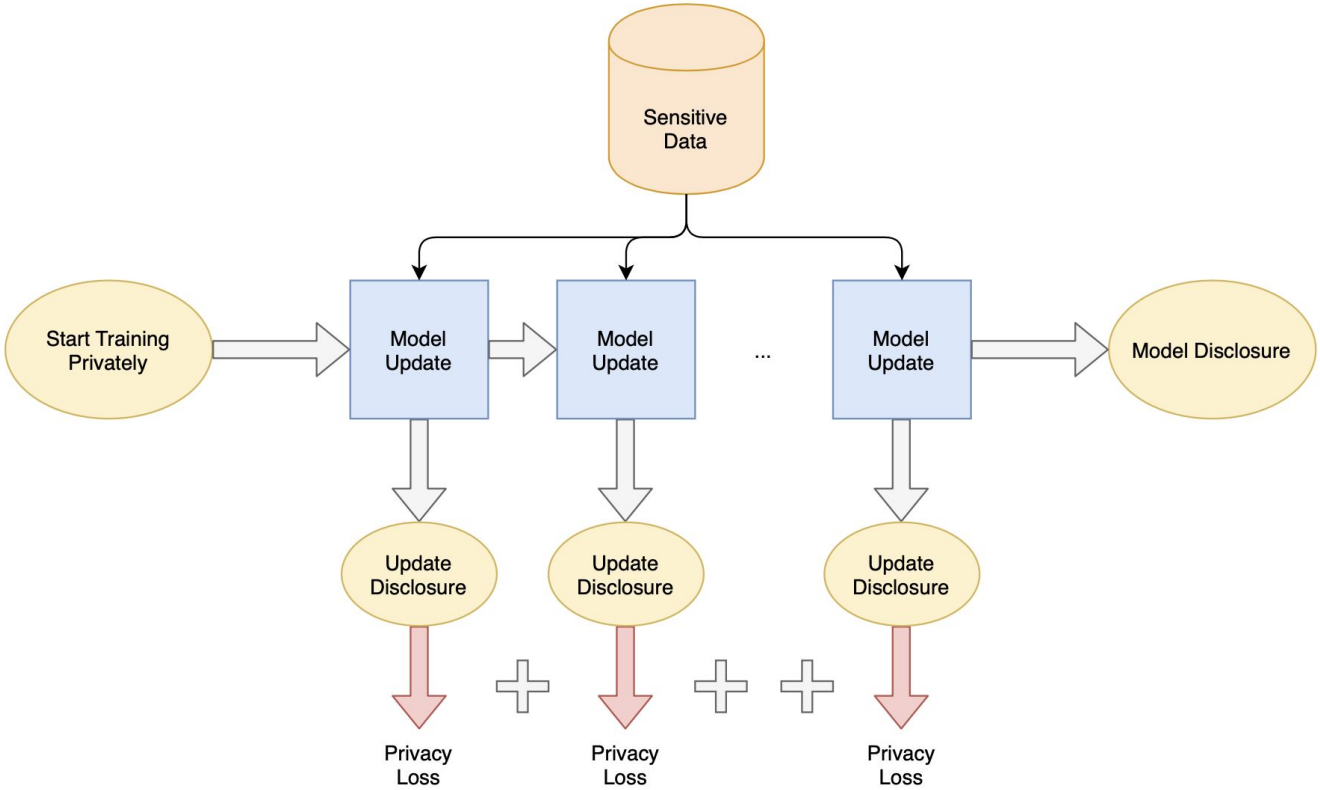

Figure 2: For DP-SGD, each model update is a private query to sensitive data that could be disclosed. Each query depends on the previous queries since they update the same set of parameters. The resulting model is simply the aggregation of all those updates.

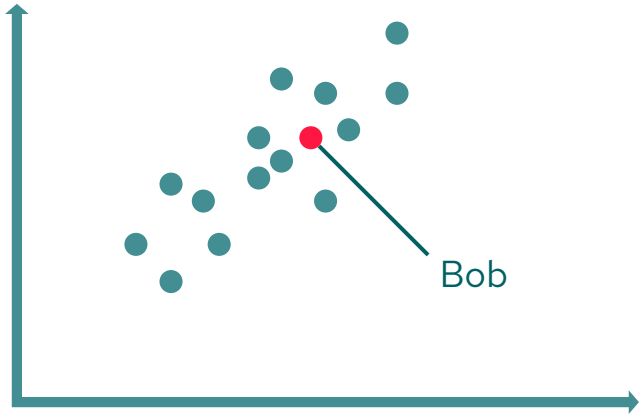

Noisy representation

Noisy representation

Shortcomings of DP-SGD

What if we don't need to publish gradients?

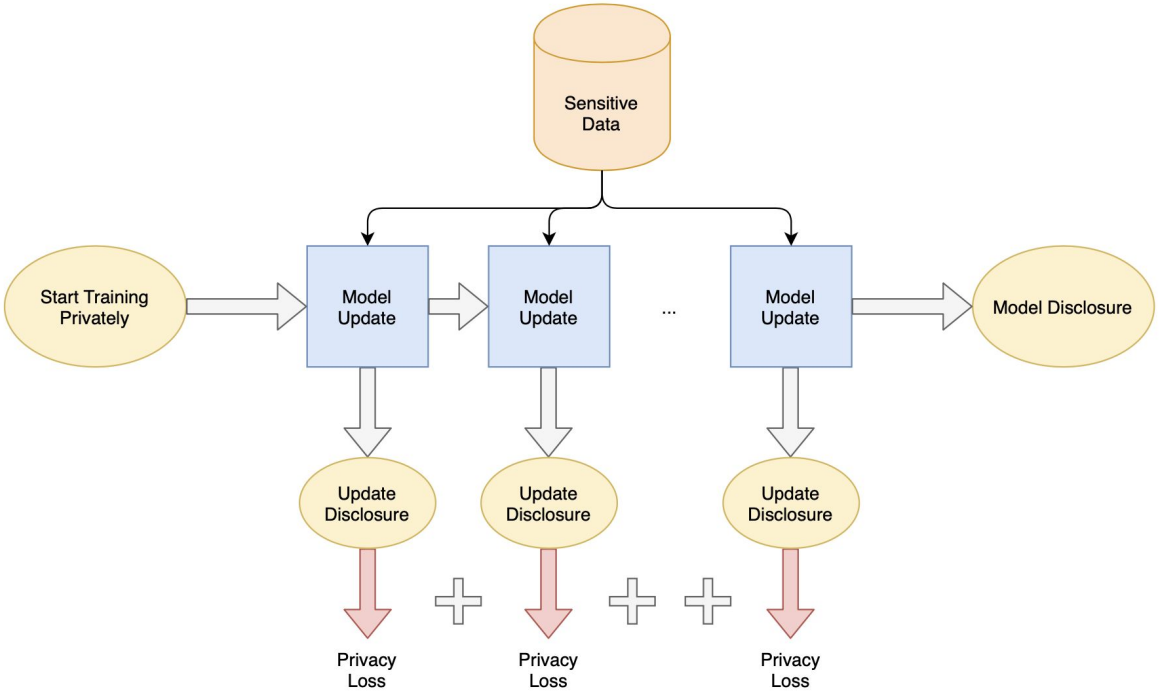

Figure 2: For DP-SGD, each model update is a private query to sensitive data that could be disclosed. Each query depends on the previous queries since they update the same set of parameters. The resulting model is simply the aggregation of all those updates.

Shortcomings of DP-SGD

What if we don't need to publish gradients?

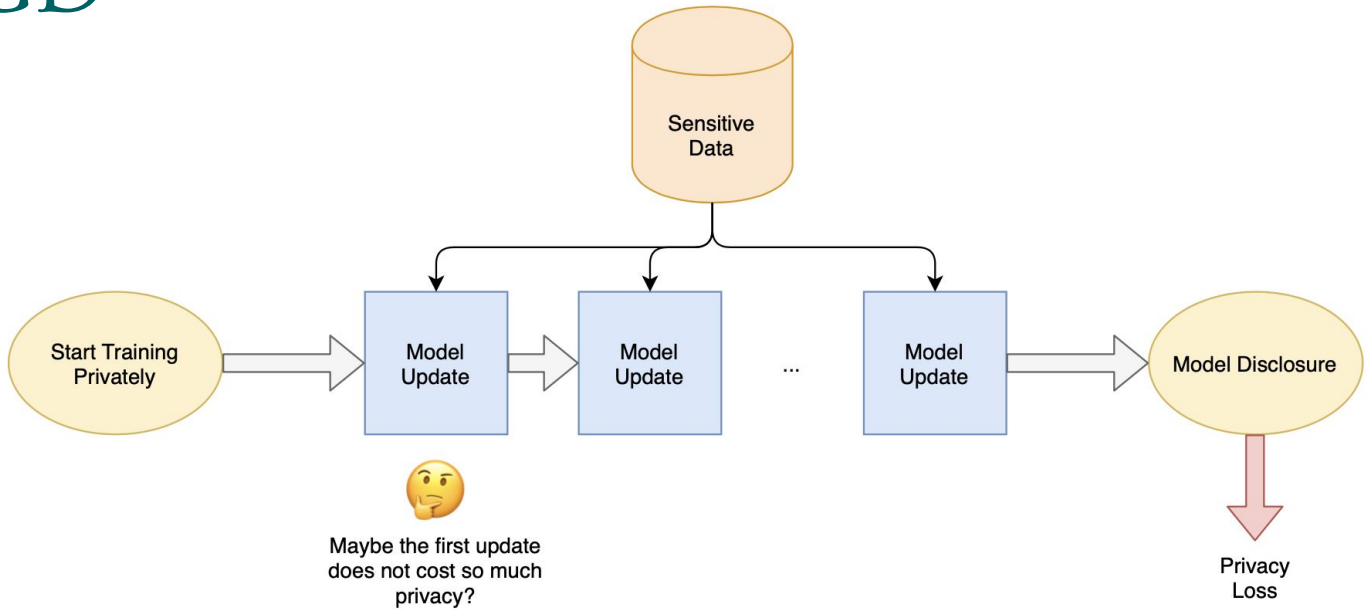

Figure 3: Because only the resulting model is disclosed, there is no need to consider each model update is disclosed. Maybe the first updates have far less impact than the last ones.

DP-SGLD

The noisy gradient descent from DP-SGD in [Abadi et al., 2016], can be looked at as a Markov Diffusion Process, as described in [Neal et al.,].

$$\theta_{t+1} = \theta_t - \rho \nabla R_S(\theta) + \sqrt{2\rho\sigma}\varepsilon_t$$

In [Welling and Teh, 2011] Welling and Teh, look at the stochastic version of this markov chain, where the gradient is computed on a sub-sample of the full dataset.

For $S_t \subset S$, where $|S_t| = L$ and $|S| = N$.

$$\theta_{t+1} = \theta_t - \rho \frac{N}{L} \nabla R_{S_t}(\theta) + \sqrt{2\rho\sigma}\varepsilon_t$$

This procedure, named *Stochastic Gradient Langevin Dynamics* (SGLD), can be seen as the discretisation of the following diffusion process.

$$d\theta_t = -\rho \nabla R_S(\theta) dt + \sqrt{2\rho\sigma} dW_t$$

and that, as $\rho \to 0$, the discretization error[3] goes to zero and the stochasticity in the gradients averages out.

DP-SGLD

The density $f(\theta, t)$ of θ along the SGLD diffusion process is decribed by the Fokker-Planck equation as derived in [Pavliotis, 2014]:

$$\frac{\partial f}{\partial t} = -\rho \cdot \Delta R_S(\theta) f - \rho \nabla R_S(\theta) \nabla f + \rho \sigma \Delta f$$

The stationnary distribution of the SGLD diffusion $\theta_\infty \sim f(\theta)$ solves:

$$\sigma \Delta f = \Delta R_S \cdot f + \nabla R_S \cdot \nabla f$$

It can be shown to be:

$$f(\theta) = \frac{e^{-\frac{1}{\sigma} R_S(\theta)}}{\int_\theta e^{-\frac{1}{\sigma} R_S(\theta)}}$$

The exponential mechanism $\mathcal{M}_E(D, u, \Theta)$ selects and outputs an element $\theta \in \Theta$ with probability proportional to:

$$\exp\left(\frac{\varepsilon u(D, \theta)}{2\Delta u}\right).$$

$\mathcal{M}_E(D, u, \Theta)$ is $(\varepsilon, 0)$-differentially private.

DP-SGLD

DP-SGLD is promising
Guarantees of convergence are still missing in the non-convex case

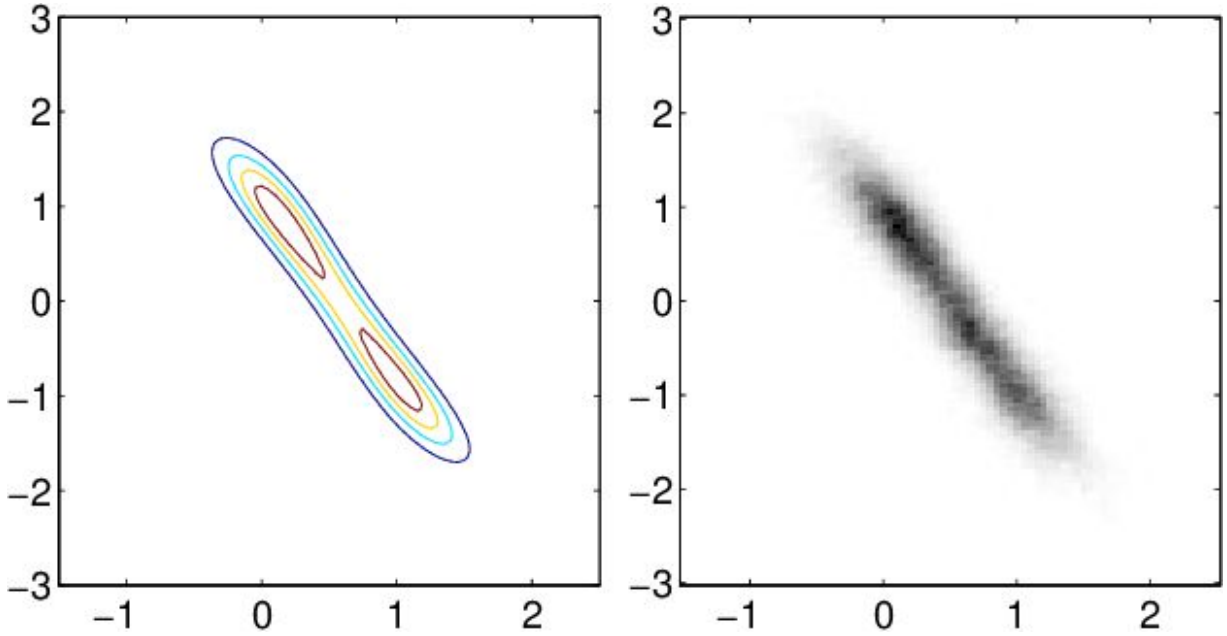

Private data science workflow

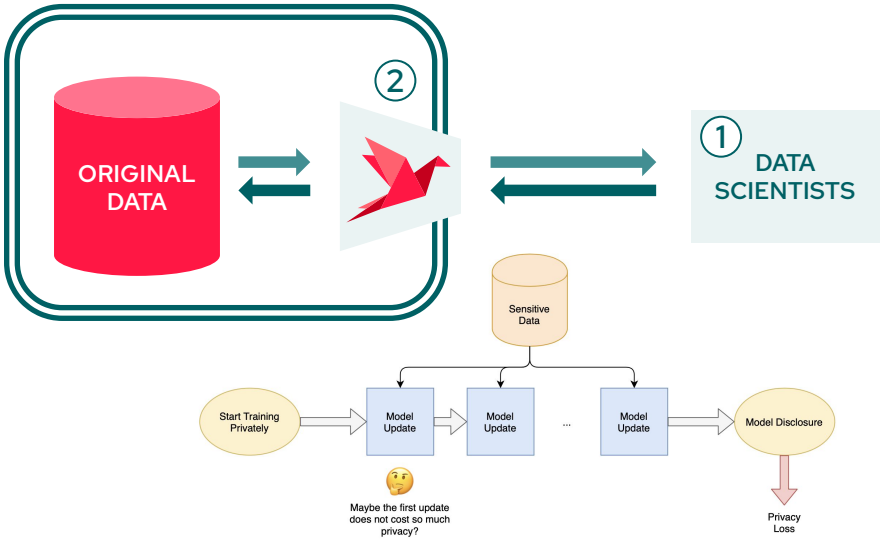

1. Prepare models
2. Submit models for training on original data

Practical difficulties synthetic DB

Today we have an autoregressive approach inspired by:

[1] Bengio, Y. & Bengio, Samy. (2001). *Modeling High-Dimensional Discrete Data with Multi-Layer Neural Networks*.

[2] Benigno Uria, Marc-Alexandre Côté, Karol Gregor, Iain Murray, and Hugo Larochelle. *Neural autoregressive distribution estimation*.

$$P(Z_1 \dots Z_n) = \prod_{i=1}^{n} P(Z_i | Z_1 \dots Z_{i-1}). \quad (1)$$

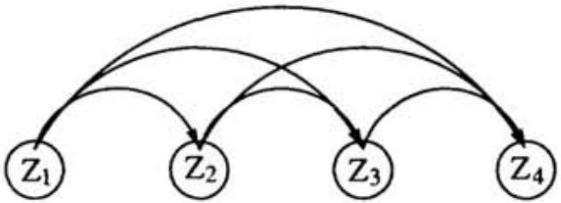

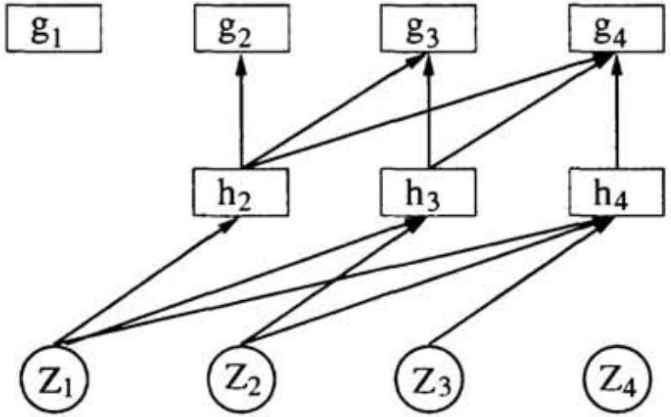

Private data science workflow

$$P(Z_1 \dots Z_n) = \prod_{i=1}^{n} P(Z_i | Z_1 \dots Z_{i-1}). \quad (1)$$

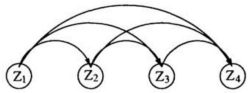

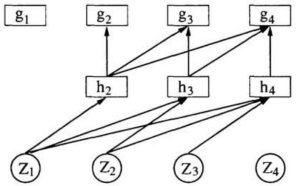

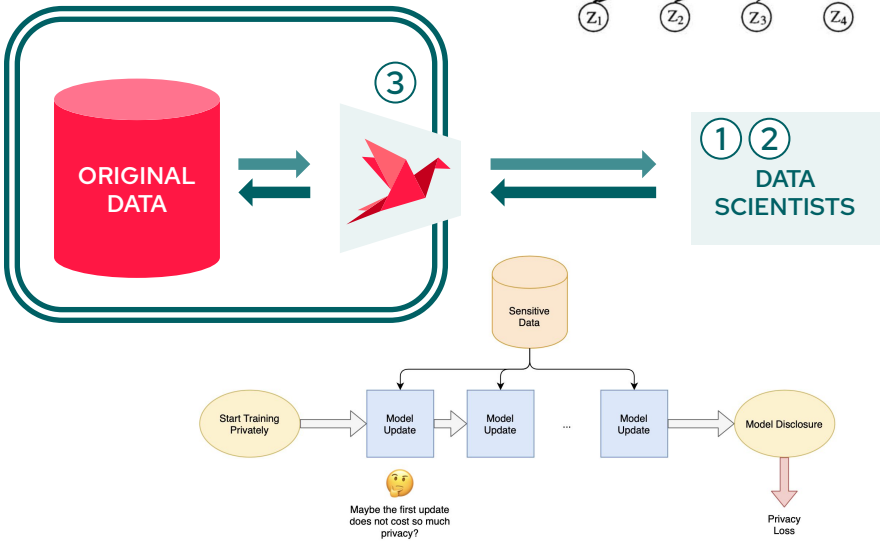

1. Download synthetic data
2. Prepare models on synthetic data
3. Submit models for training on original data

Sarus: more opportunities with better data protections

Faster iterations

Test new ideas and potential of your data in minutes.

Get results without going through lengthy compliance processes for every new application.

Reduce data leakage

Data privacy is always protected with the highest level: differential privacy.

No data leaves your secure infrastructure, there won't be any copy of anything sensitive out in the wild

New Opportunities

Leverage the full potential of data that was altered before, or not even accessible.

Make datasets available to more internal teams or external partners.

Appendix

How companies use private data for innovation

TRADITIONAL PRIVATE DATA WORKFLOW

Data scientists are given access to data to build models. Data may be personal or belong to another business unit, or external partners.

DATA REQUIREMENTS

- Some data to be made accessible to data scientist
- This data must be anonymized to protect privacy
- Stringent security applies to data engineering and access control

=> Compliance process and approval specific to each project and dataset

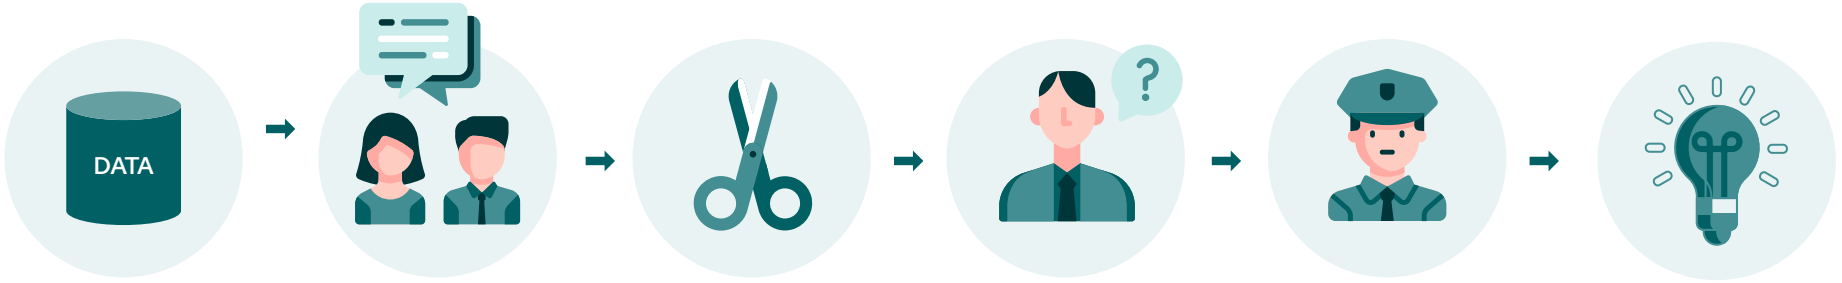

Innovation suffers

INHIBITS INNOVATION POTENTIALLY KILLING PROJECTS

- Takes too long to prove new ideas
- Difficult to bring in external partners

INFERIOR RESULTS

- Anonymization techniques alter data, which can destroy value
- Some data is unavailable because it is too personal

DATA LEAKAGE RISK

- Weak privacy properties
- Copies of the data may be at risk in the wild

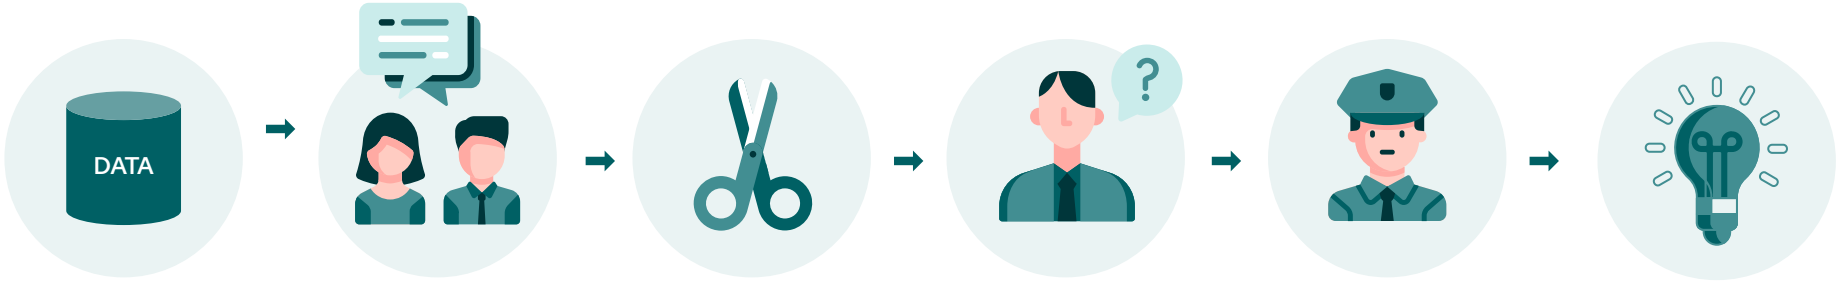

2+ MONTHS FOR EACH ITERATION

THE SARUS WAY: Data scientists and analysts work on original data without accessing it.

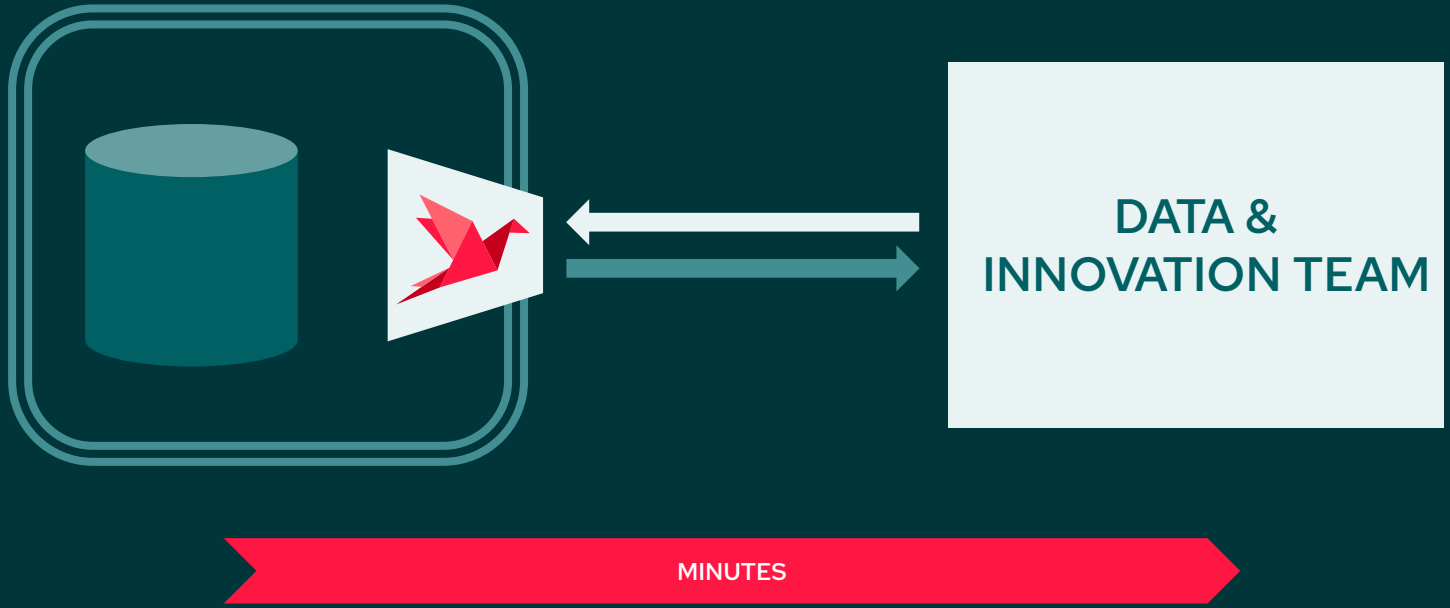

With Sarus, practitioners work on data without accessing it

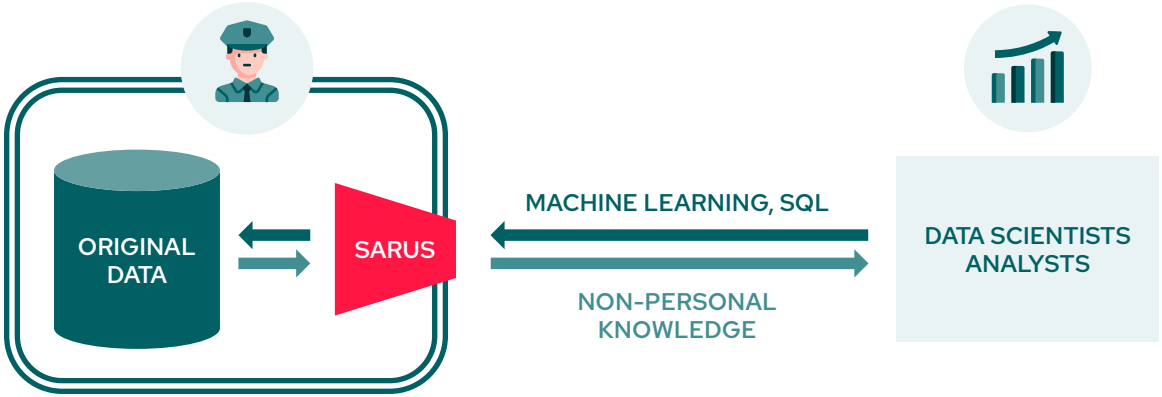

Differential privacy ensures that personal information never leaves the secure bubble.
for faster

Works with structured and unstructured data for maximum utility

Compatible with data scientists' tools and ways of working

Data remains in the original infrastructure for maximum security

Application: Learning without moving BU data

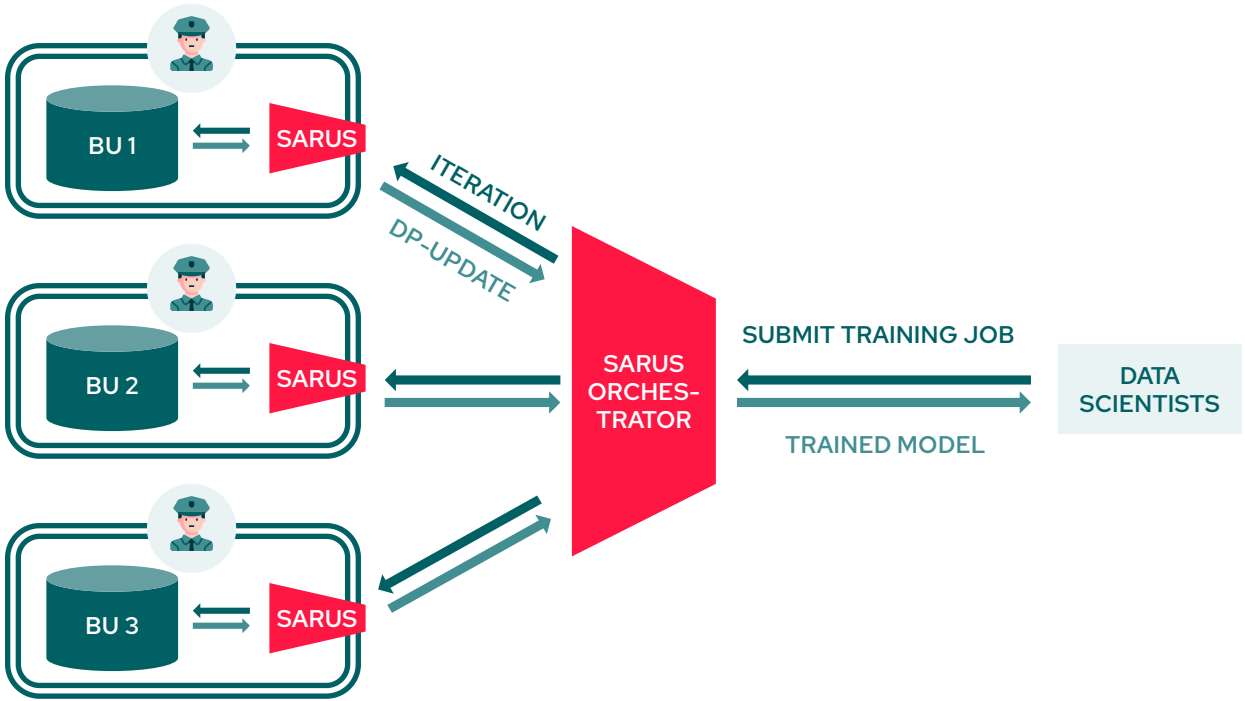

Problem Moving BU data to central repository for innovation faces business/compliance roadblocks.

Solution: Bank installs Sarus on participating BU's infrastructure. Central data science team work remotely on sensitive data.

Results: No data has been moved outside of BU's IT or across borders. Yet, it was 100% accessible for innovation.

Application: Open innovation Validate vendor added-value

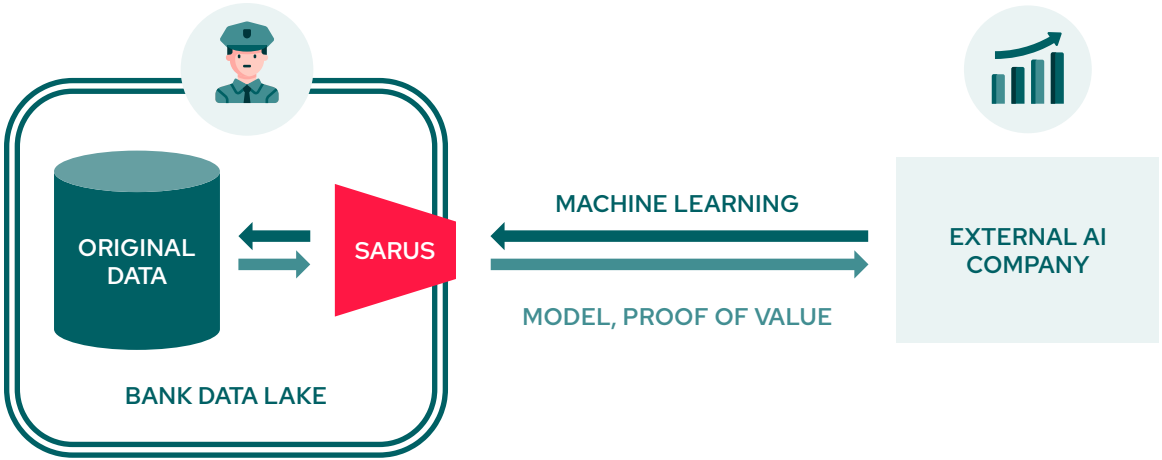

Problem: Bank wants to validate vendor's anti-money laundering algorithms but cannot share data with them.

Solution: Bank installs Sarus on its data lake and the vendor can prove the value of their solution. No personal data leaks out.

Results: The bank quickly validates upside of working with vendor before entering deeper partnership.

Application: learn on imaging center medical data

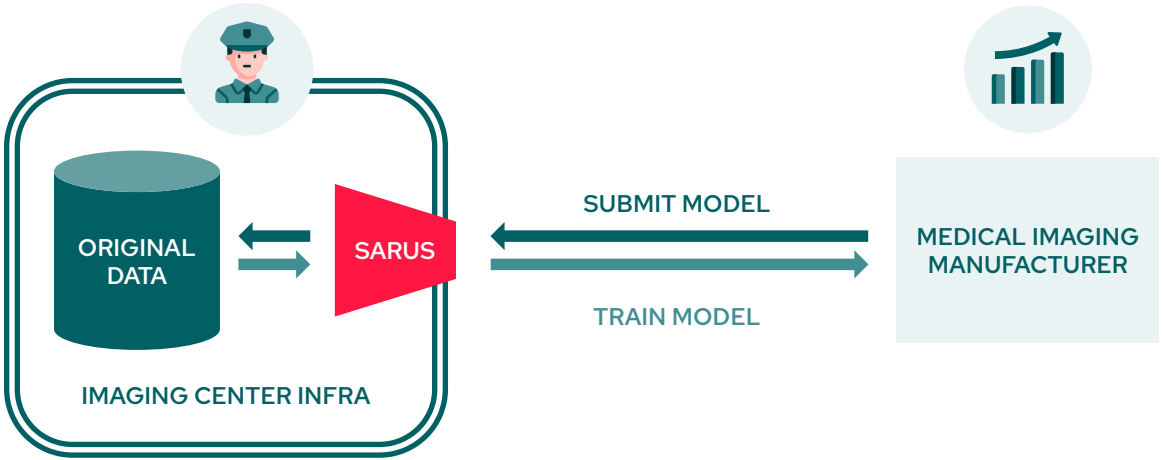

Problem: Medical imagery requires data that does not belong to the manufacturer

Solution: Sarus is installed on imaging systems and is used to train non-personal models

Results: Manufacturer can improve models and propose AI-based decision tools on a full range of equipments without jeopardizing privacy

How it compares to:

Data masking

Privacy is protected in a generic way, going beyond rules-based DM in versatility and protection.

It combines static DM (synthetic data can be used freely) and dynamic DM (end-result is generated dynamically).

It applies to unstructured data while never releasing this data, hence more secure than USR techniques.

Synthetic data

Sarus provides synthetic data for practical reasons but in the end, the model or calculation is always performed on the original data.

No pattern is lost to data generation and the end result is provably valid, which is more convincing to the data scientist or external authority (FDA).

Also, synthetic data is fragile for complex data structures.

Homomorphic encryption

In HE, the computation on the data set is protected but nothing guarantees that the result does not reveal personal information. With Sarus, we prove the result of the calculation does not reveal personal information. Implicitly we assume the data can be processed in clear but it could combine with HE in the cases of this assumption does not hold.

Streamline compliance with mathematical privacy guarantees

“De-identified data isn’t”

Crowds don’t protect individuals: average wages becomes *personal data* when a new employee joins. Most approaches are subject to re-identification attacks. Even concepts such as k-anonymity or l-diversity fail to prevent re-identification.

Enter Differential Privacy

DP was proposed by Cynthia Dwork in 2006. The idea is that whether a person is in a dataset or not should not matter to the result. It implies randomization.

Sarus implements the latest DP research throughout the workflow to accelerate compliance processes

Differential privacy is at the core of Sarus technology. It is implemented in all interaction with the data practitioner so that the data owner can be confident personal information is protected regardless of data type and end-use. Guarantees even extend to the final models or analyses that are protected against inference attacks.

Leverage the full depth of data for maximum utility

Anonymization techniques target simple data structures
The more complex the data the more ways there are to identify users. Traditional anonymization techniques rely on aggregation or field deletion to make re-identification harder. It cannot work on rich data

Sarus makes this step irrelevant by never sharing data
Instead of trying to suppress any way re-identification can happen, which is bound to fail on rich data, Sarus lets users work on the full data set and focuses on making sure the results of the computation will not leak personal information.

Natively applicable to all data types, including unstructured data
By doing away with ad hoc anonymization techniques, Sarus provides a universal solution to work with all data type, including longitudinal data, free text, images, audio files, or patient health records.

Compatible with all tools and habits that matter to data scientists

Data scientists use standard libraries
Sarus wants to make the data scientist experience identical whether data is accessible directly or via Sarus. They can use the same libraries and submit their own models for training on the remote data.

Differential privacy never gets in the way
Protecting individuals is a shared objective between the data owner (who wants to protect their data) and the data scientist (who wants their models to generalize beyond the training set). Applying the right level of differential privacy helps them both achieve their targets.

Synthetic data samples for prototyping and testing
Because testing code is more efficient when done locally, data scientists can download sample data via the Sarus gateway. This data is synthetic data that shares the format and statistical properties of the original dataset but never embeds personal information. When ready, data scientists can train their model on the original data.

Data remains in the original infrastructure for maximum security

Data stay where it is
The number one enemy of data security is moving copies of data around. Sarus software is deployed inside our customers' infrastructure, being a gateway between sensitive data and end-users. Data is never exposed or exported.

No change to security practices
Security policies remain the same. Requires a single one open port for API; traffic is fully monitored and logged.

Compatible with any environment
Docker-based solution can be installed on any system, from clouds (AWS, GCP, Azure) to on-premise, on Linux or Windows

Compatible with major data stores
Original data can sit on SQL databases, storage such a S3, data lake such as HADOOP, etc.

Membership Inference attacks references

Privacy in Pharmacogenetics: An End-to-End Case Study of Personalized Warfarin Dosing

Performing an in-depth case study on privacy in personalized warfarin dosing, we show that suggested models carry privacy risks, in particular because attackers can perform what we call model inversion: an attacker, given the model and some demographic information about a patient, can predict the patient's genetic markers.

Carlini, N., Liu, C., Erlingsson, Ú., Kos, J., & Song, D. (2019). The secret sharer: Evaluating and testing unintended memorization in neural networks

Specifically, for models trained without consideration of memorization, we describe new, efficient procedures that can extract unique, secret sequences, such as credit card numbers

Fredrikson, M., Jha, S., & Ristenpart, T. (2015, October). Model inversion attacks that exploit confidence information and basic countermeasures. Example of MIA

Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017, May). Membership inference attacks against machine learning models

Using realistic datasets and classification tasks, including a hospital discharge dataset whose membership is sensitive from the privacy perspective, we show that these models can be vulnerable to membership inference attacks.

Hayes, J., Melis, L., Danezis, G., & De Cristofaro, E. (2019). LOGAN: Membership inference attacks against generative models.

Present attacks based on both white-box and black-box access to the target model, against several state-of-the-art generative models, over datasets of complex representations of faces (LFW), objects (CIFAR-10), and medical images (Diabetic Retinopathy).

Private data science workflow

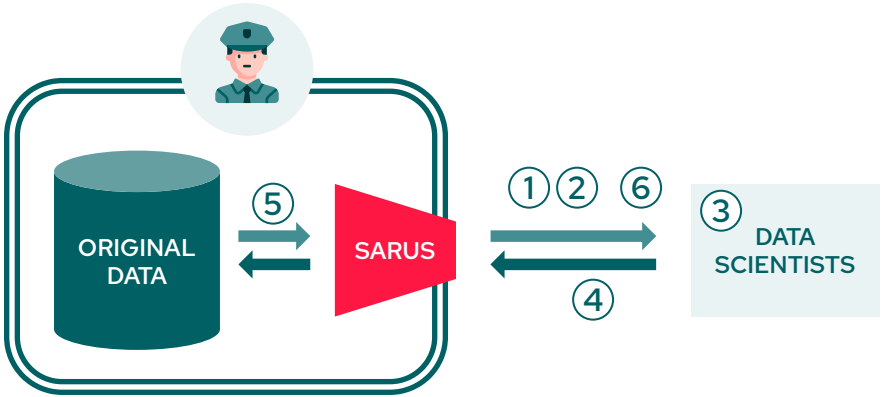

1. List data, access metadata
2. Download synthetic data
3. Prepare models, test locally on synthetic data
4. Submit models for training
5. Train on original data
6. Download differentially-private models

Without Sarus

- BROWSE AVAILABLE DATASETS
- DOWNLOAD SAMPLE
- ANALYZE SAMPLE, SPEC OUT MODEL
- WRITE MODEL, TEST ON SAMPLE
- TRAIN ON DISTRIBUTED INFRA
- ANALYZE PERFORMANCE, ITERATE
- PUSH MODEL TO PRODUCTION

With Sarus

- BROWSE AVAILABLE DATASETS
- DOWNLOAD **SYNTHETIC** SAMPLE
- ANALYZE **SYNTHETIC** SAMPLE, SPEC OUT MODEL
- WRITE MODEL, TEST ON **SYNTHETIC** SAMPLE
- TRAIN ON **REMOTE &** DISTRIBUTED INFRA
- ANALYZE PERFORMANCE, ITERATE
- PUSH MODEL TO PRODUCTION

Anomaly detection on synthetic data

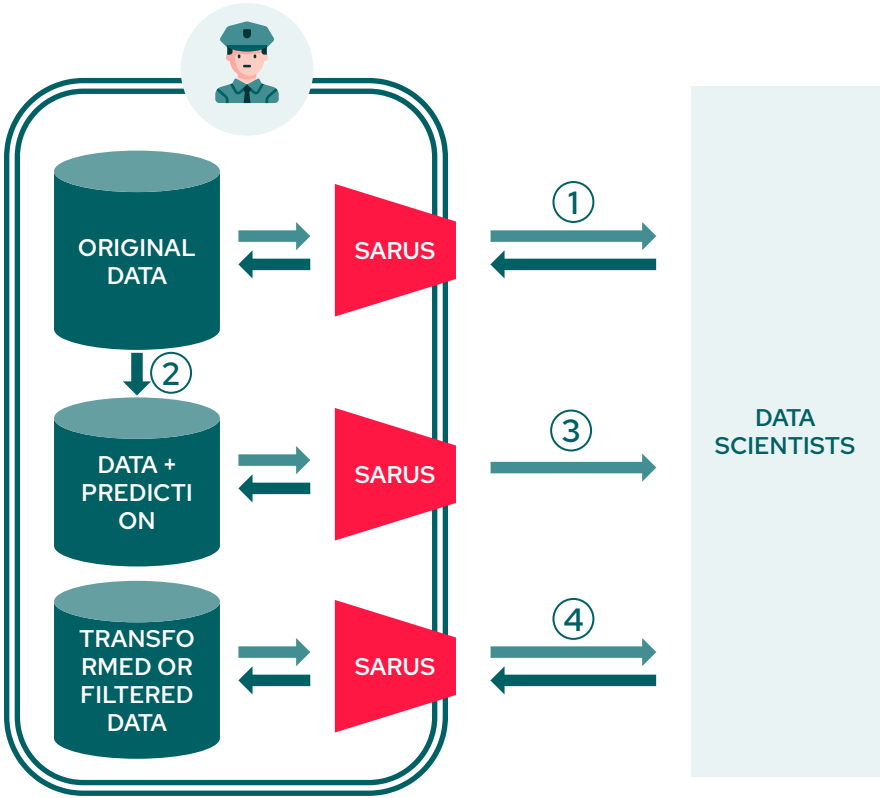

Step 1: Train model on full data set

If performance below par
Step 2: Create secondary data set with model prediction as a feature (or filter by model prediction). Sarus automatically generates synthetic data

Step 3: Inspect synthetic data and look for error inducing patterns

Step 4: Fix data or train with new transform functions

Annotation on hidden data 👨‍🔬

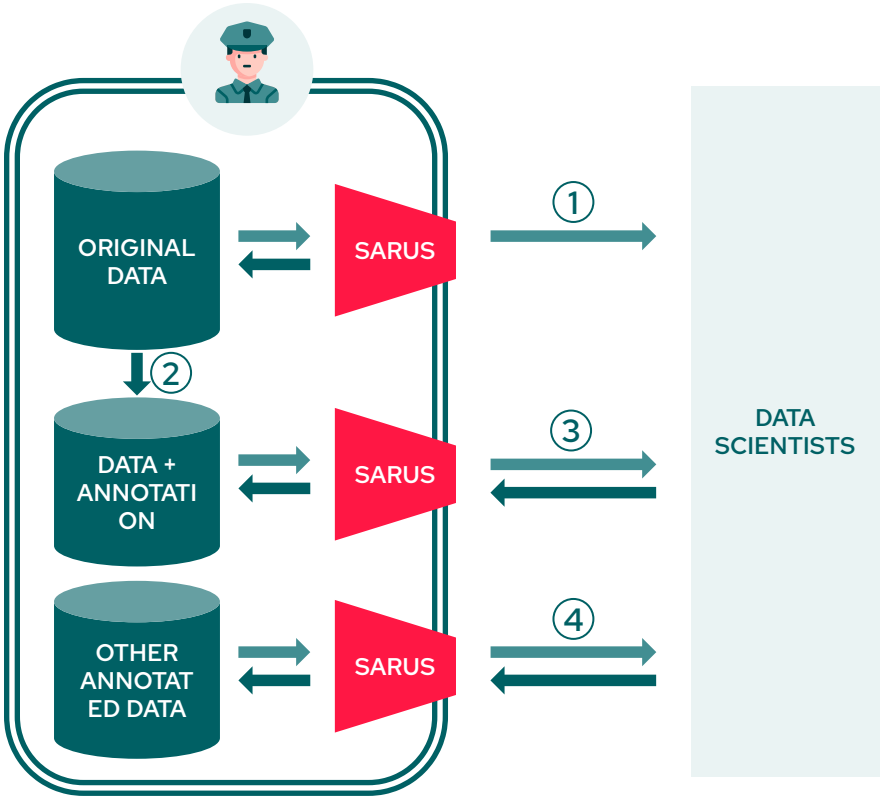

Step 1: Train annotation model on synthetic data

Step 2: Create augmented data set with annotations from step-1 model

Step 3: Supervised training on annotated data

Step 4 [optional]: Validate on separate dataset with manual annotation

ε,δ - differential privacy

Definition
https://www.cis.upenn.edu/~aaroth/Papers/privacybook.pdf

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

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

Team

Maxime Agostini, CEO

Ecole Polytechnique graduate,
worked in finance and tech.
Cofounder-CEO of AlephD

Nicolas Grislain, Chief of Science

Graduated from ENS,
worked in finance, data startups.
Cofounder-Chief of Science
of AlephD

Vincent Lepage, CTO

Ecole Polytechnique graduate,
worked in finance, data and health
startups (ex-CPO OWKIN).
Cofounder-CTO of AlephD