Combining differential privacy and homomorphic encryption for privacy-preserving collaborative machine learning

June, 12^{th}, 2023

Arnaud GRIVET SEBERT

Supervisors: Renaud SIRDEY and Cédric GOUY-PAILLER

Reviewers: Melek ÖNEN and Jan RAMON

Examiners : Rachid GUERRAOUI and David POINTCHEVAL

Outline

1. Introduction

Alice and Bob's dream

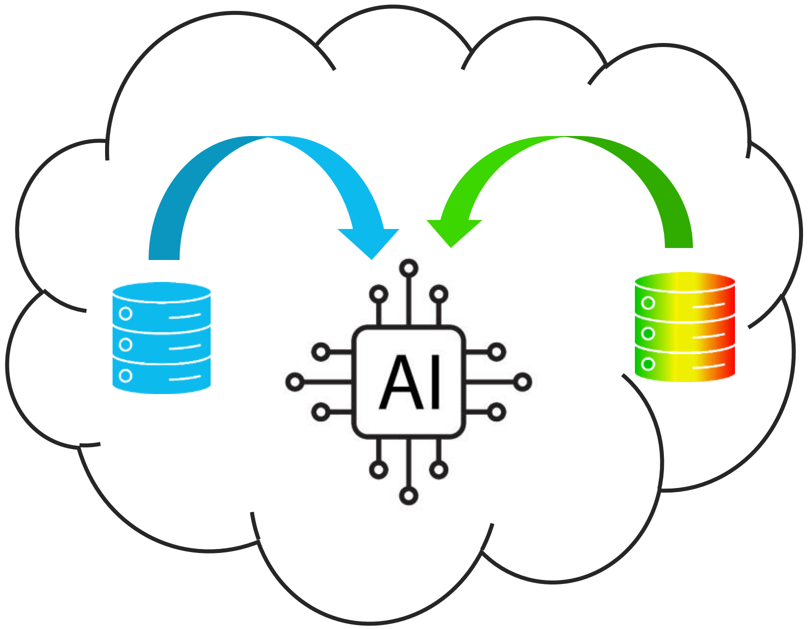

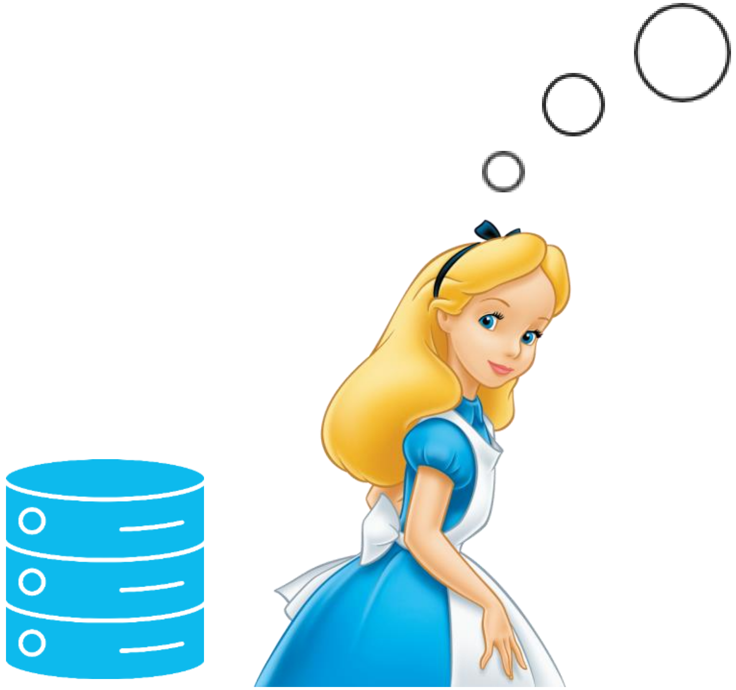

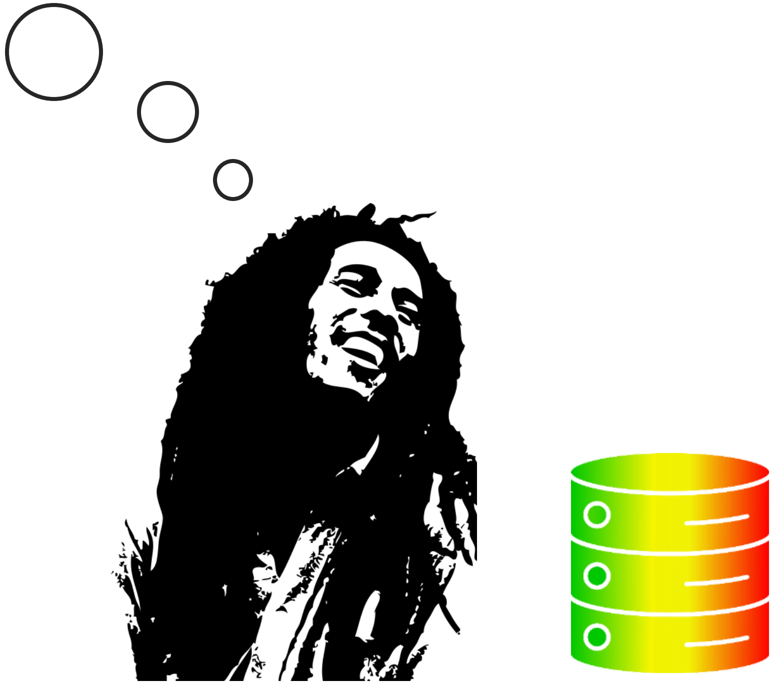

Alice and Bob both have data. They dream of building a public ML model that could learn from both their knowledge.

E.g. medical data from patients.

Part of the image comes from vecteezy.com

Alice and Bob's dream

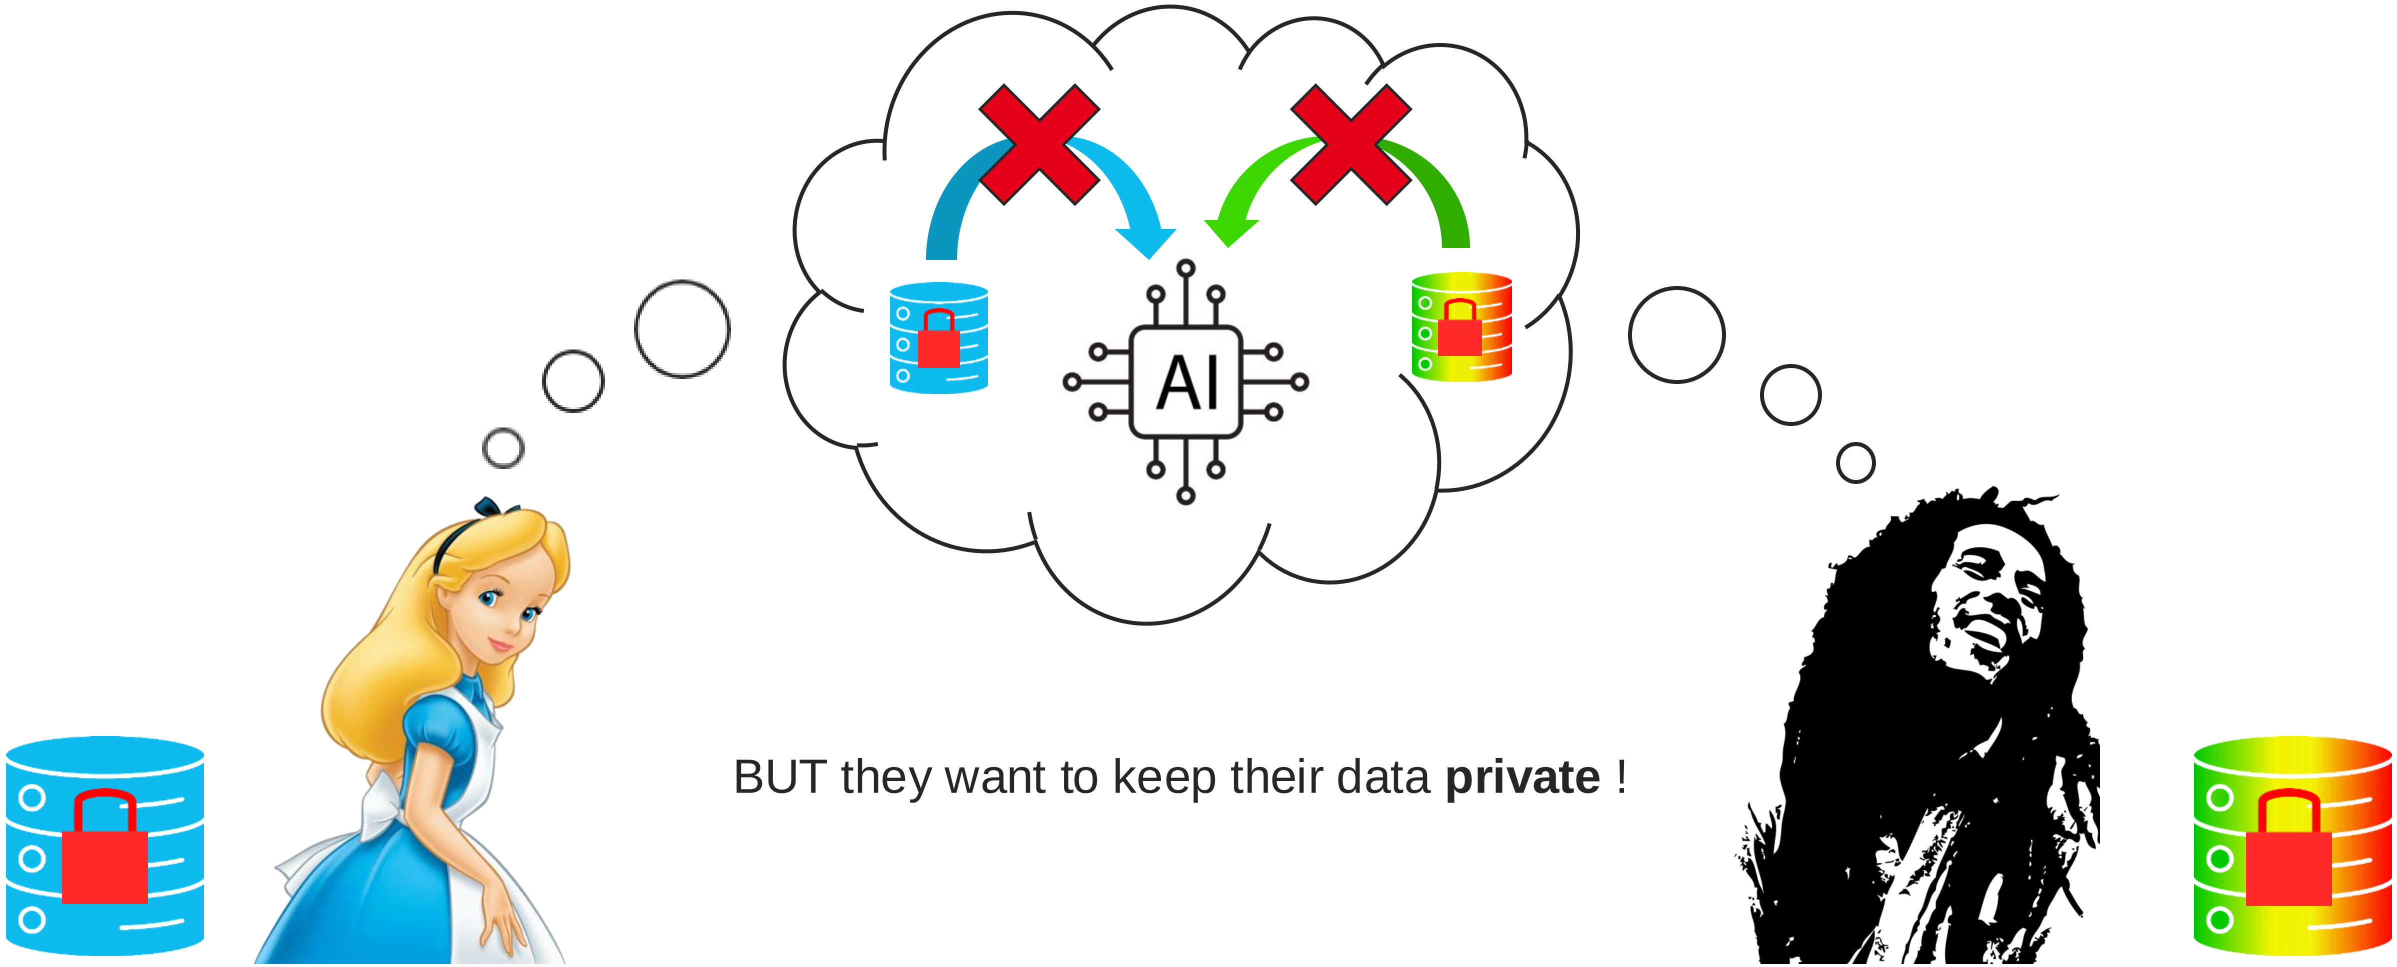

BUT they want to keep their data **private** !

Part of the image comes from vecteezy.com

Federated learning (FL)

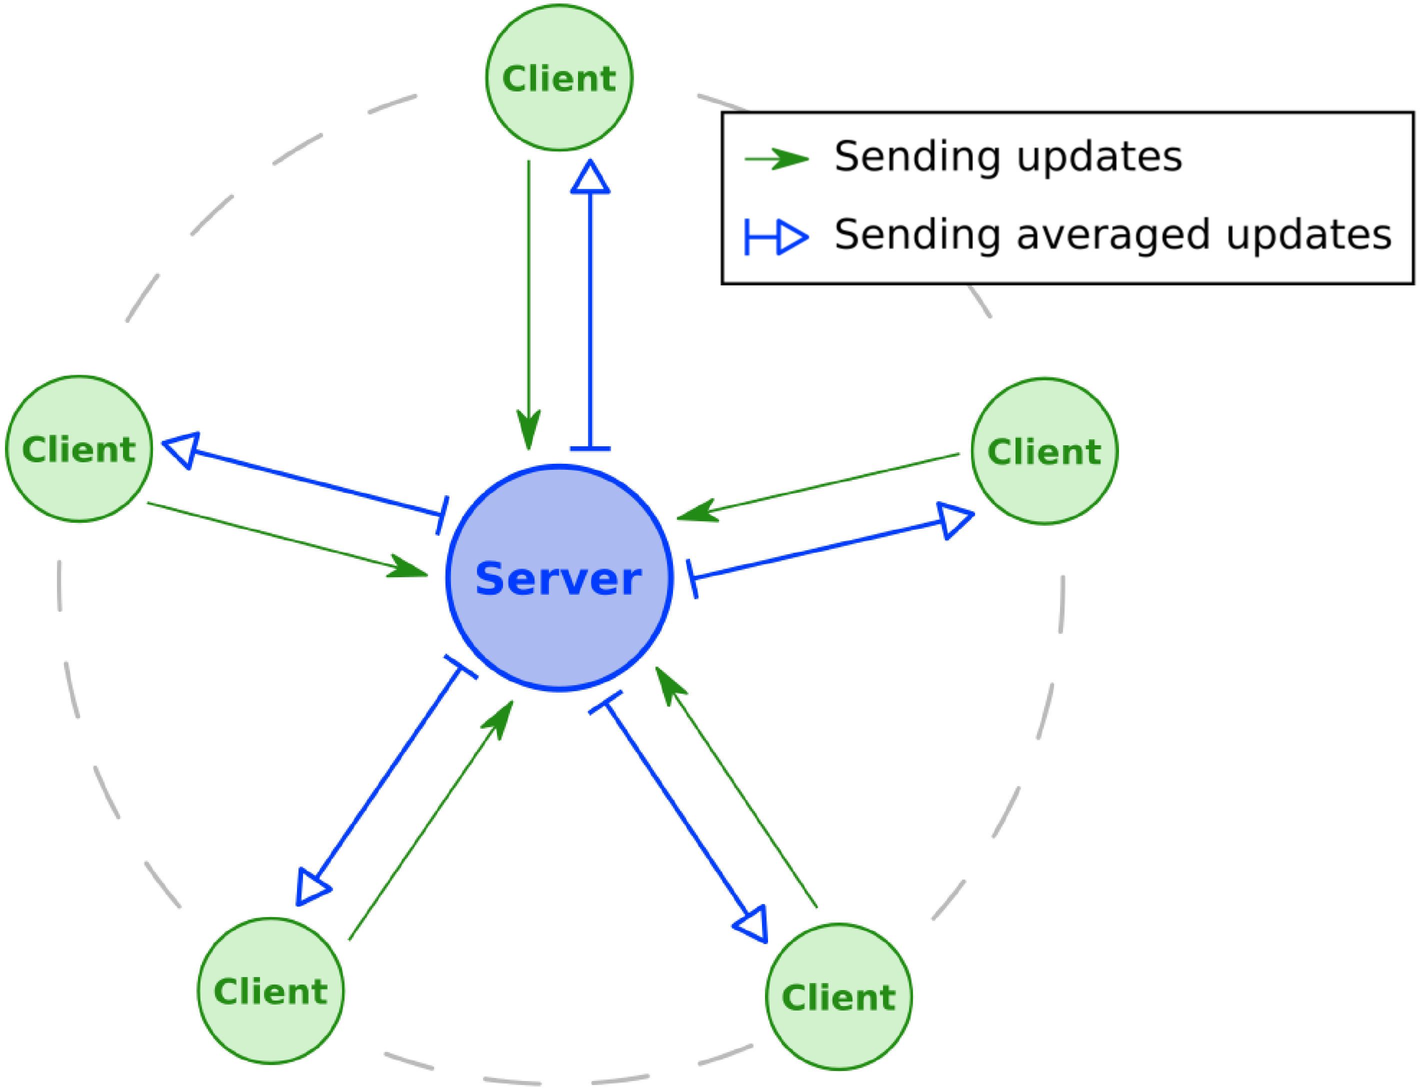

- *Federated learning of deep networks using model averaging*, McMahan et al. (2016)
- *Communication-efficient learning of deep networks from decentralized data*, McMahan et al. (2017)

At each iteration:

- The **server** sends the current model to a subset of the **clients**
- The clients compute updates of the model using their local data.
- The clients send their updates to the server.
- The server **averages** the clients' updates.

→ **The data are not outsourced.**

FL is not a panacea

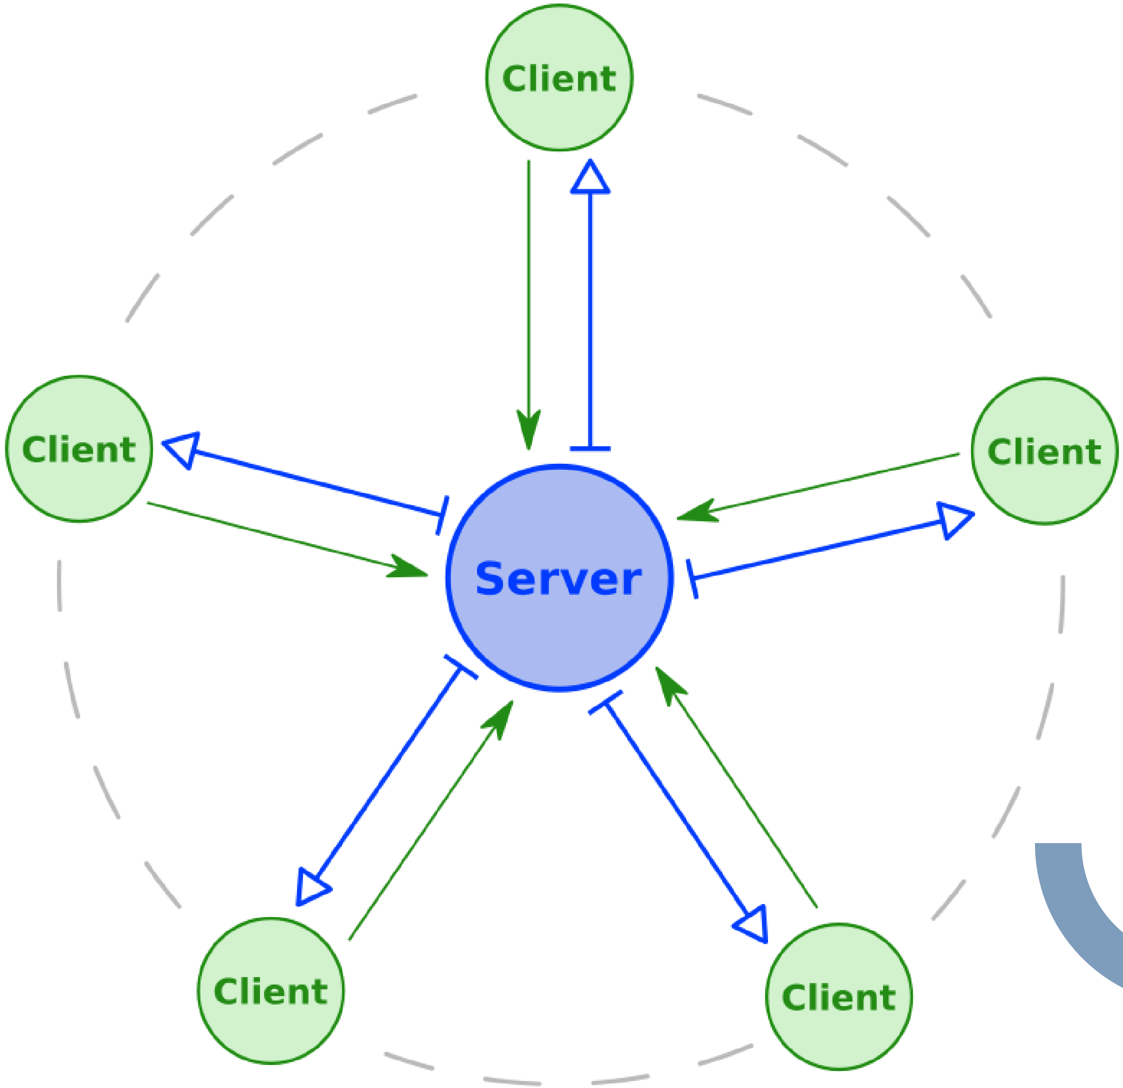

Threats may come from:

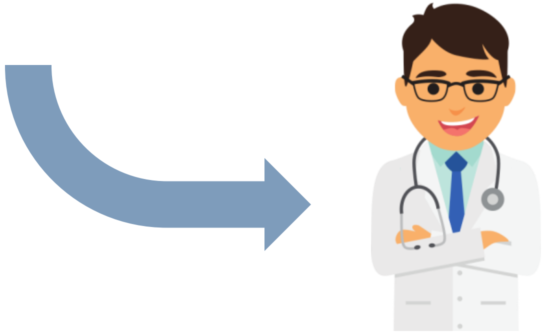

End-user

Part of the image comes from vecteezy.com

FL is not a panacea

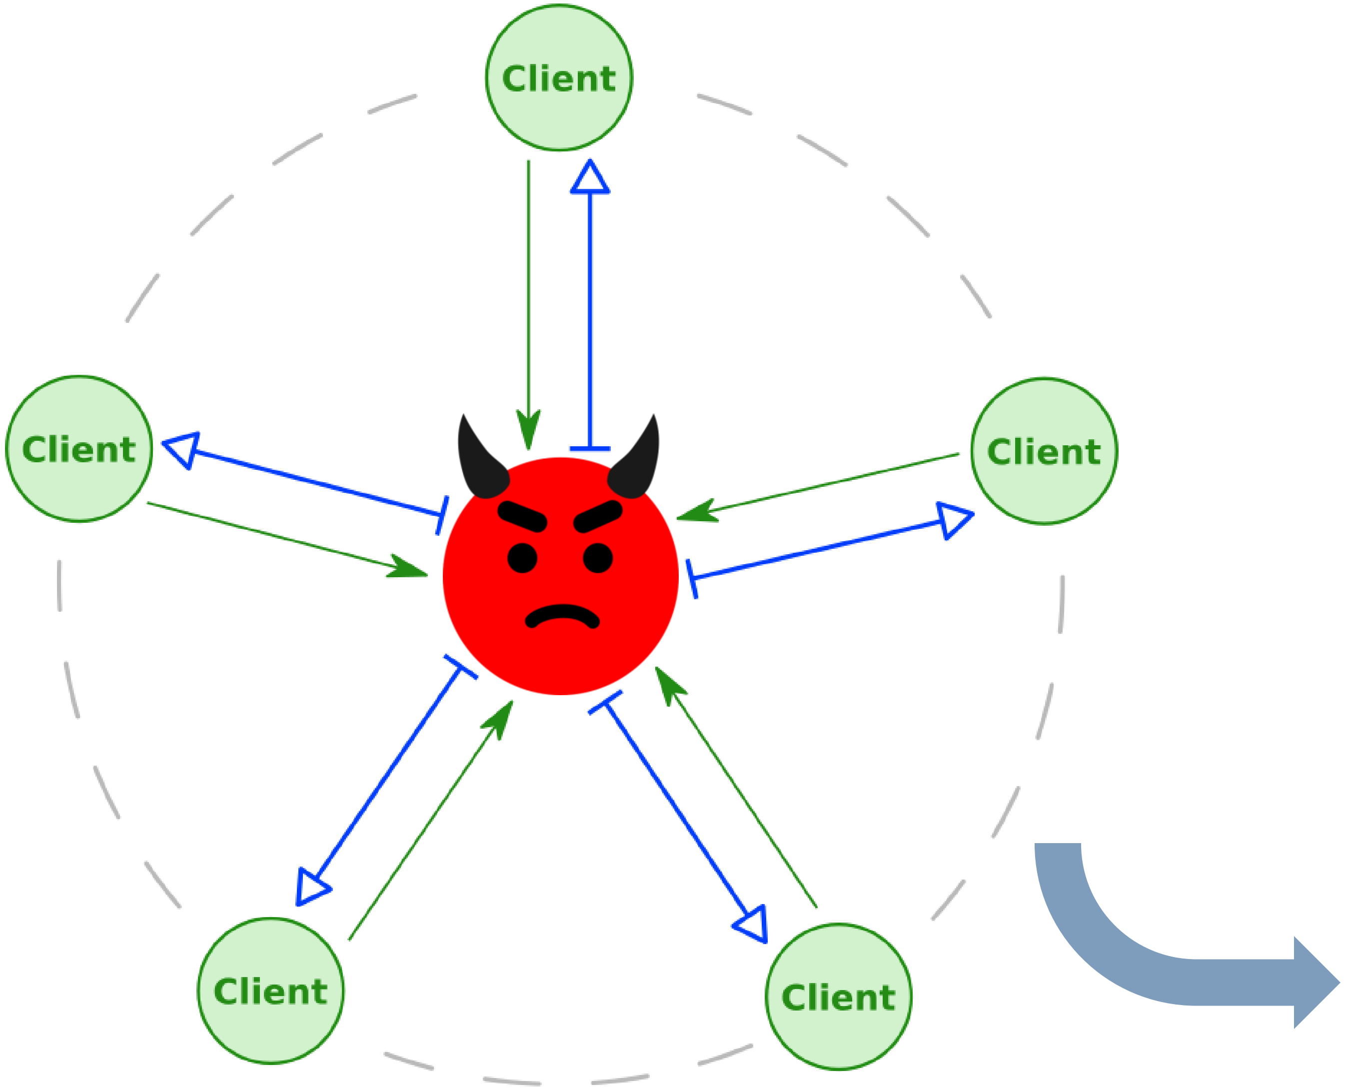

End-user

Threats may come from:

- The **server** who sees the individual updates

Part of the image comes from vecteezy.com

FL is not a panacea

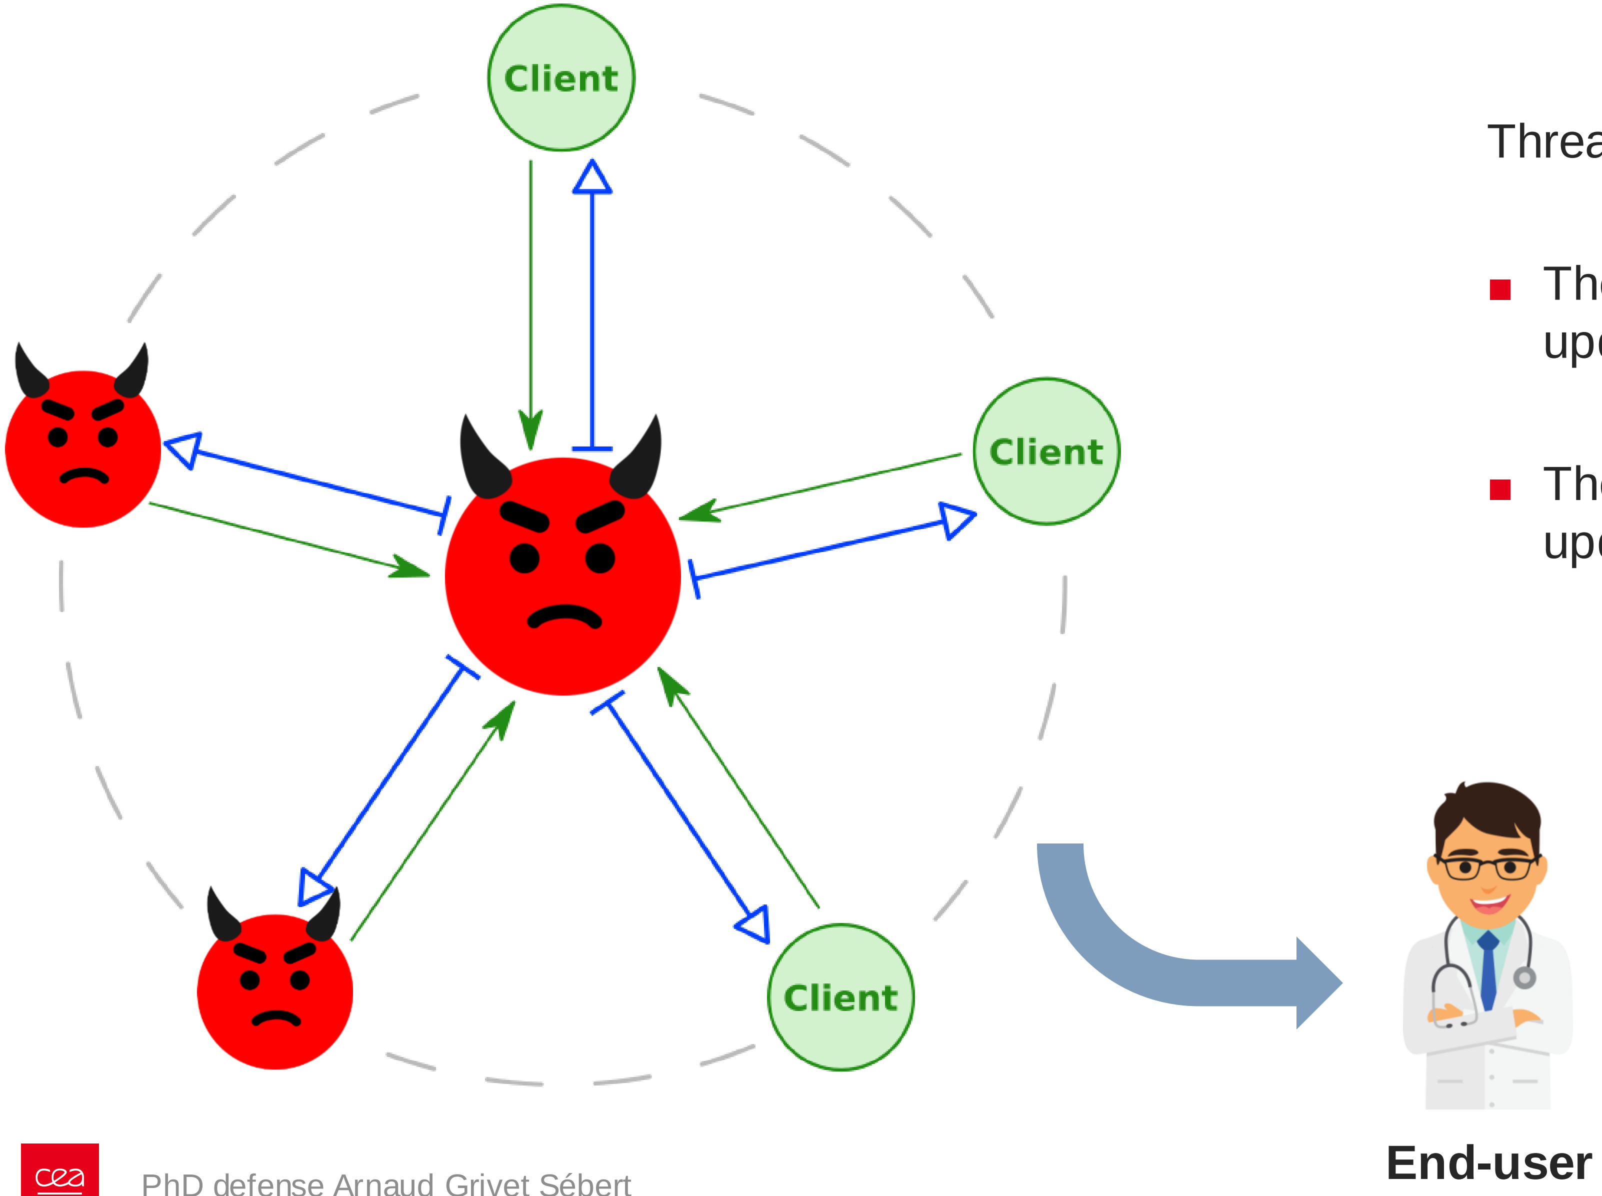

Threats may come from:

- The **server** who sees the individual updates
- The **other clients** who see the aggregated updates

Part of the image comes from vecteezy.com

FL is not a panacea

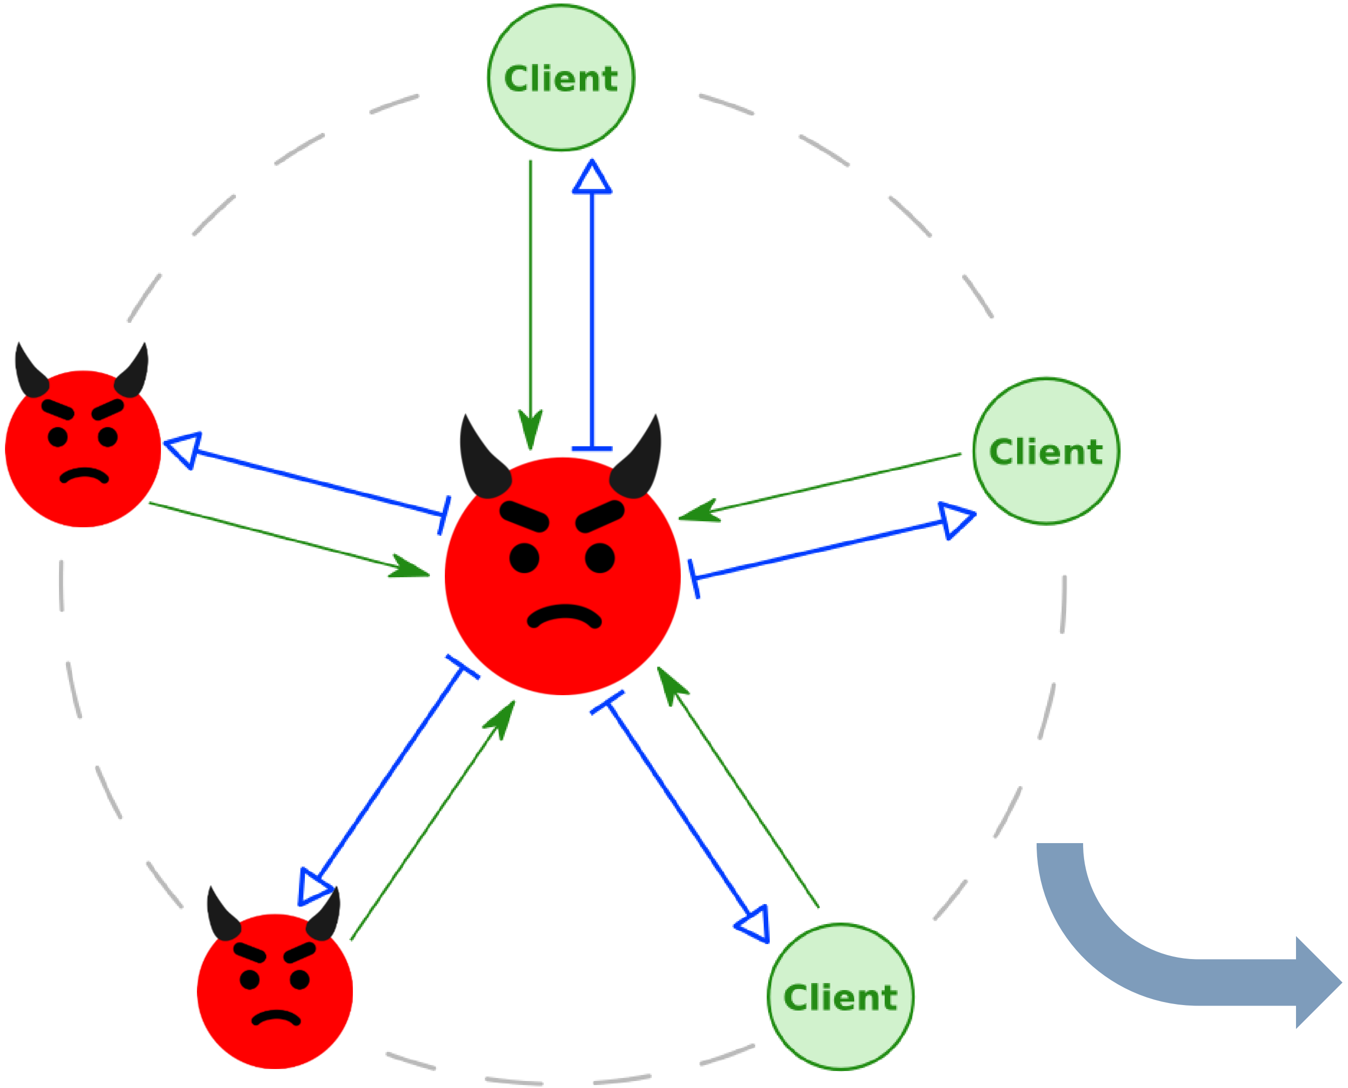

Threats may come from:

- The **server** who sees the individual updates
- The **other clients** who see the aggregated updates
- The **end-users** who see the final model

End-user

Part of the image comes from vecteezy.com

Model inversion and similar attacks

The end-users may « retro-engineer » the trained model to get information about the training data.

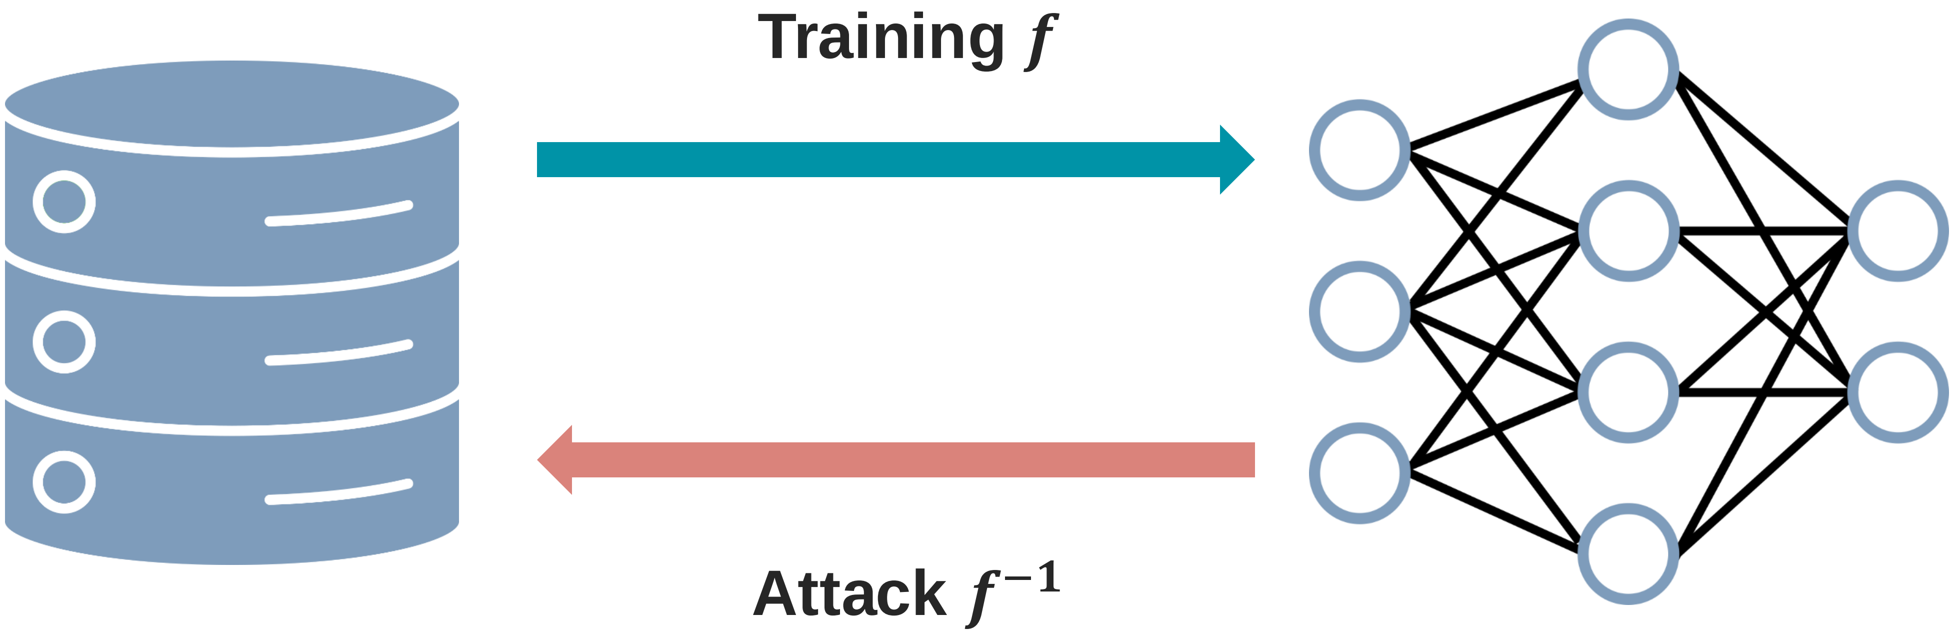

Part of the image comes from vecteezy.com

2. Privacy tools

Differential privacy (DP)

- *Calibrating noise to sensitivity in private data analysis*, Dwork et al. (2006)
- *The algorithmic foundations of differential privacy*, Dwork and Roth (2014)

- Binary relation of **adjacency** on the databases (generally, differing by one individual)
- A **probabilistic** mechanism satisfies **DP** when an adversary cannot infer from the output which of two adjacent databases he hesitates on was the input.

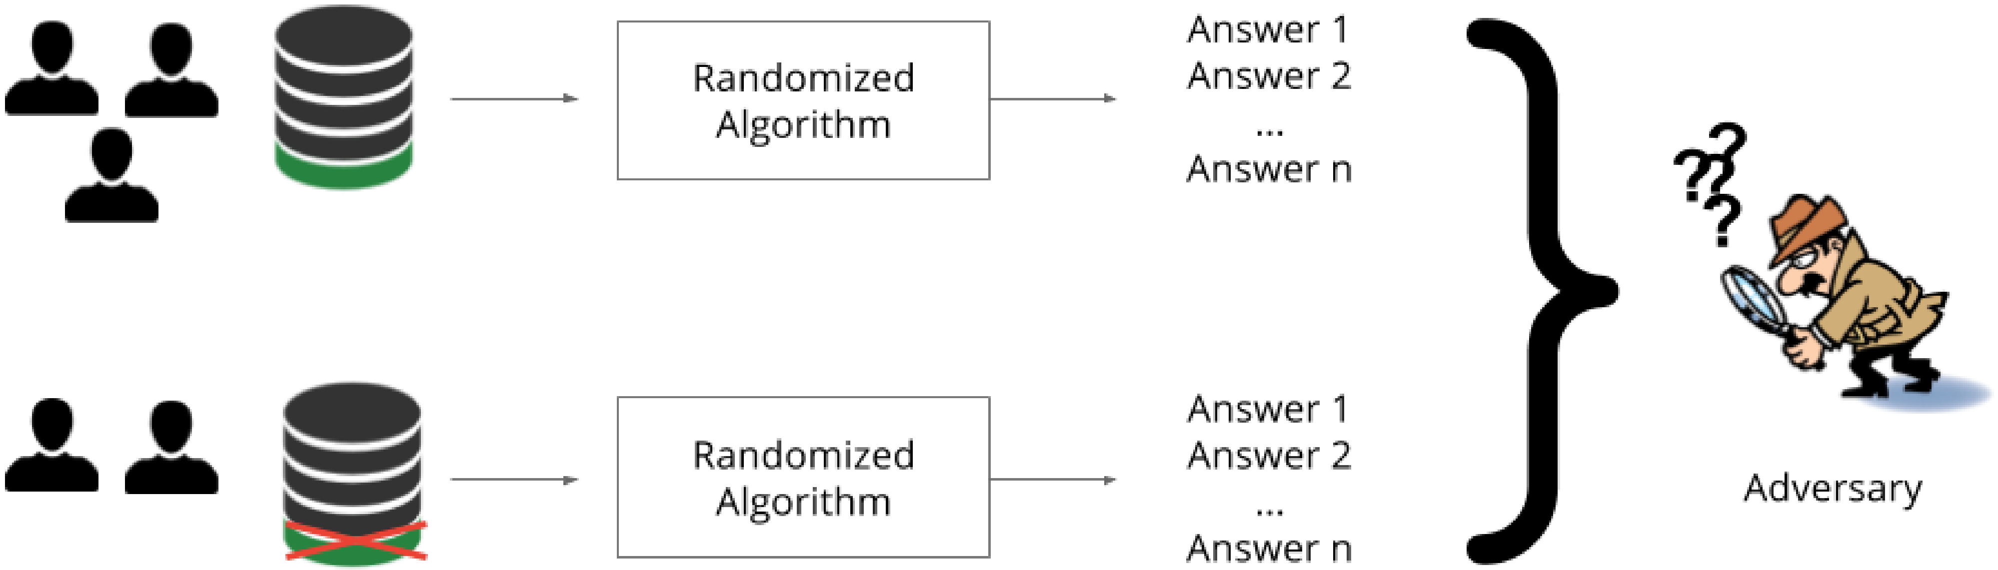

Privacy and machine learning: two unexpected allies ?, Papernot et Goodfellow (blog post, 2018)

Differential privacy (DP)

- The **indistinguishability** of two adjacent databases is **quantified** by the proximity of the associated output distributions.

- By computing the maximum of this quantity over all pairs of adjacent databases, one gets the **privacy cost** of the mechanism.

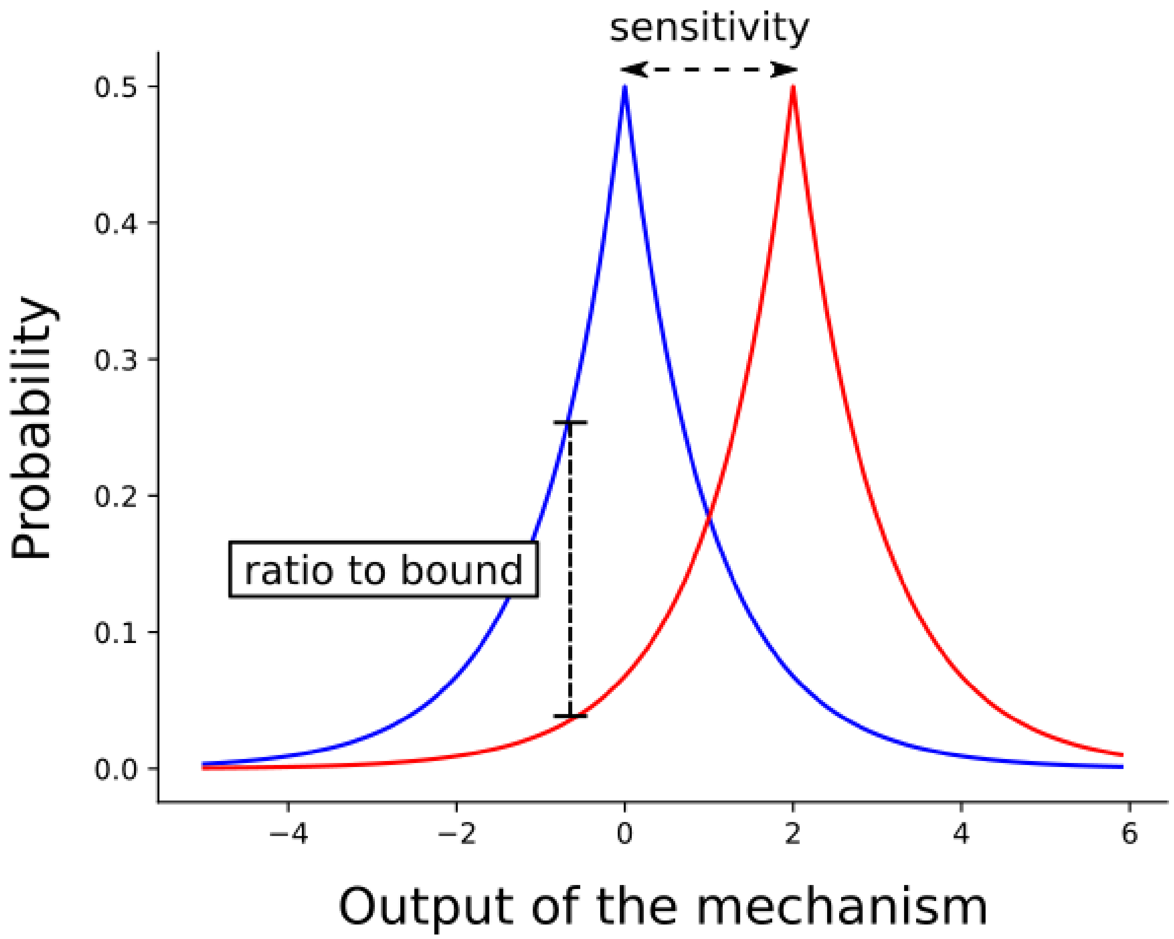

- Pure ε-DP : $$\mathbb{P}\left[\mathcal{A}(d) \in S\right] \leq e^{\epsilon}\mathbb{P}\left[\mathcal{A}(d') \in S\right]$$

Differential privacy (DP)

- The **indistiguishability** of two adjacent databases is **quantified** by the proximity of the associated output distributions.

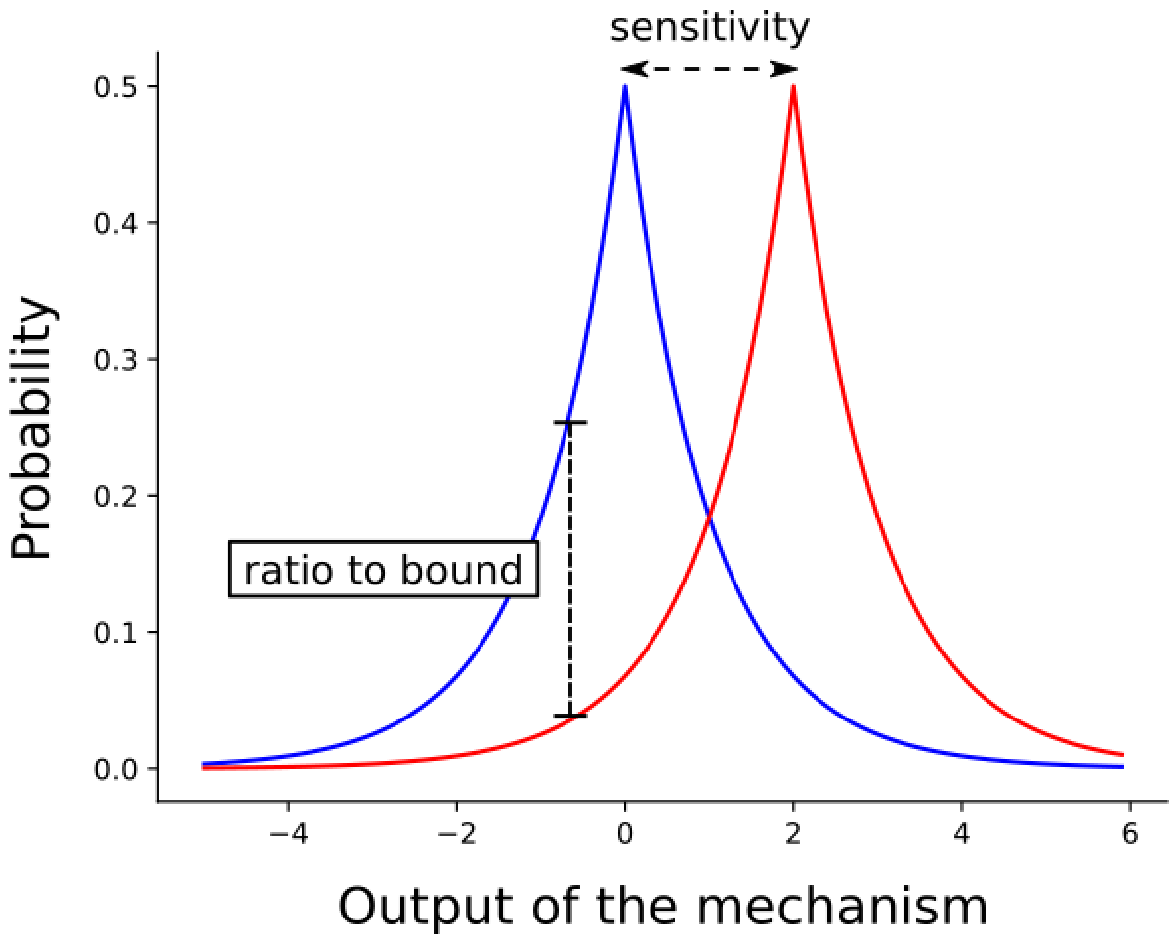

- By computing the maximum of this quantity over all pairs of adjacent databases, one gets the **privacy cost** of the mechanism.

- Approximate (ε, δ)-DP : $\mathbb{P}\left[\mathcal{A}(d) \in S\right] \leq e^{\epsilon}\mathbb{P}\left[\mathcal{A}(d') \in S\right] + \delta$

Properties of differential privacy

- A DP mechanism is **necessarily probabilistic**: most often, random noise is added to a deterministic mechanism (usually Laplacian or Gaussian).

- DP protects **individual** information, not statistical information.

- **Composition**: keeping track of the privacy cost associated to **several queries** (adaptive or not)
  - The more queries one asks, the easier it is to rebuild the underlying distribution.

- **Immunity to post-processing**: The privacy cost of f o M is (at most) the same as M.

Homomorphic encryption (HE)

Cryptographic paradigm allowing to compute operations in the encrypted domain without access to the data nor the result in clear

$$\mathsf{Enc}(m_1) \oplus \mathsf{Enc}(m_2) = \mathsf{Enc}(m_1 + m_2) \in \Omega.$$
$$\mathsf{Enc}(m_1) \otimes \mathsf{Enc}(m_2) = \mathsf{Enc}(m_1 m_2) \in \Omega.$$

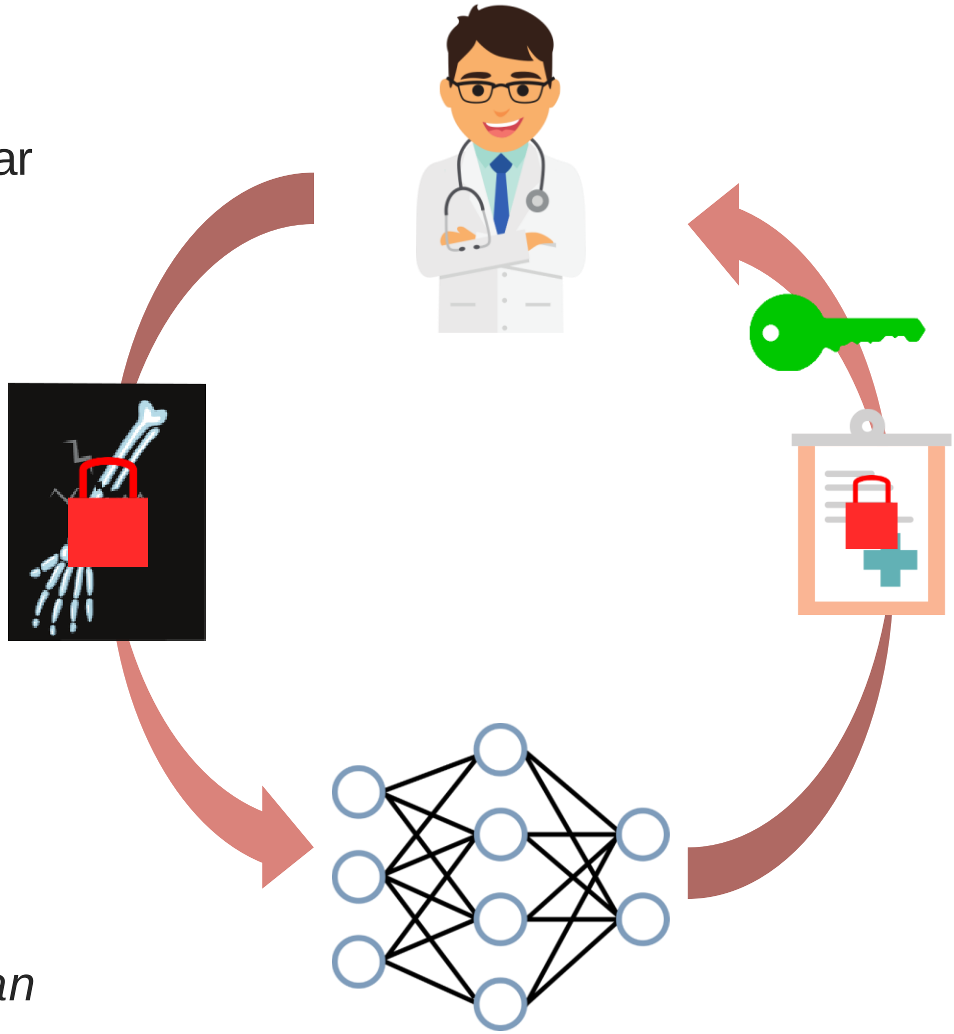

- *Somewhat practical fully homomorphic encryption*, Fan and Vercauteren (2012)
- *(Leveled) fully homomorphic encryption without bootstrapping*, Brakerski et al. (2014)
- *Faster fully homomorphic encryption: Bootstrapping in less than 0.1 seconds*, Chilotti et al. (2016)

Part of the image comes from vecteezy.com

Homomorphic encryption (HE)

- There is not one HE but many HEs:
  - Additive, multiplicative cryptosystems, fully HE (**FHE**)
  - Different FHE cryptosystems: BFV, BGV, TFHE, CKKS

- FHE is not an off-the-shelf tool, you need to tweak the parameters (plaintext and ciphertext modulus, polynomial degree, programmable bootstrapping)

- Limitations of FHE:
  - **Computationally intensive** (especially some operations like comparisons, divisions)
  - Only handles **discrete values** (Z_p, $Z_p[X]$)

- Enhancing features:
  - **Batching** (BFV, BGV, CKKS) allows to encapsulate **several cleartexts** into **one single ciphertext** $\rightarrow$ greatly accelerate the computations
  - **Multi-key** FHE, threshold FHE: several users are needed to decrypt

Comparison of DP and HE

Differential privacy	Homomorphic encryption
Information-theoretic security	Computational security
+ **Statistical** information accessible - Vulnerable to statistical attacks (e.g. reconstruction of average data)	Totally **blinds** the adversary
- What is a good **privacy cost** in practice is unclear	+ Cryptographic standards of **security parameters**
- The mechanism is noisy → **trade-off** between **accuracy** and privacy	+ The decrypted information is (almost) **intact**
+ Almost free in terms of computation	- **Computationally intensive**
Usually uses **continuous** noise	Only works with **discrete** values

State of the art on privacy-preserving FL…

- FL with cryptographic primitives: **Ariann: *Low-interaction privacy-preserving deep learning via function secret sharing***, Ryffel et al. (2020), ***Batchcrypt: Efficient homomorphic encryption for cross-silo federated learning***, Zhang et al. (2020)
  - Communication between the clients and no DP

- FL with secure multi-party computation and DP : ***A generic framework for privacy preserving deep learning***, Ryffel et al. (2018), ***Comprehensive comparison of multiparty secure additions with differential privacy***, Goryczka and Xiong (2015), ***Distributed differential privacy via shuffling***, Cheu et al. (2019) and ***Amplification by shuffling: From local to central differential privacy via anonymity***, Erlingsson et al. (2019)
  - Communication between the clients and, in some cases, trusted entity needed
  - With secure shuffling, DP guarantees not as good as with secure aggregation

- FL with HE and DP: ***Efficient and privacy-enhanced federated learning for industrial artificial intelligence***, Hao et al. (2019) and ***A hybrid approach to privacy-preserving federated learning***, Truex et al. (2019)
  - Do not take into account the quantisation

… and alternatives to FL

- Decentralised federated learning: ***Distributed differentially private averaging with improved utility and robustness to malicious parties***, Sabater et al. (2020), ***Privacy amplification by decentralization***, Cyffers and Bellet (2022) and ***An accurate, scalable and verifiable protocol for federated differentially private averaging***, Sabater et al. (2022)
  - Communication between the clients
  - In the second paper, vulnerable to colluding clients and eavesdroppers
  - In the third paper, they use a central coordinator → no full decentralisation

- Private Aggregation of Teacher Ensembles (PATE): ***Semi-supervised knowledge transfer for deep learning from private training data***, Papernot et al. (2016) and **Scalable private learning with PATE**, Papernot et al. (2018)
  - Requires a trusted aggregator

3. Bridging DP and FHE via Poisson quantisation

Arnaud GRIVET SEBERT, Marina CHECRI, Renaud SIRDEY, Oana STAN, Cédric GOUY-PAILLER, **Combining homomorphic encryption and differential privacy in federated learning**, *https://arxiv.org/abs/2205.04330*

Private FL with centralised noise

Combining DP and HE

- **DP** against **end-users** and **other clients**
- **HE** against the **server**

Two problems

- The DP guarantees do not apply to the server
  → The server cannot have access to the final model in clear
- **Verifiable computing** cannot be implemented

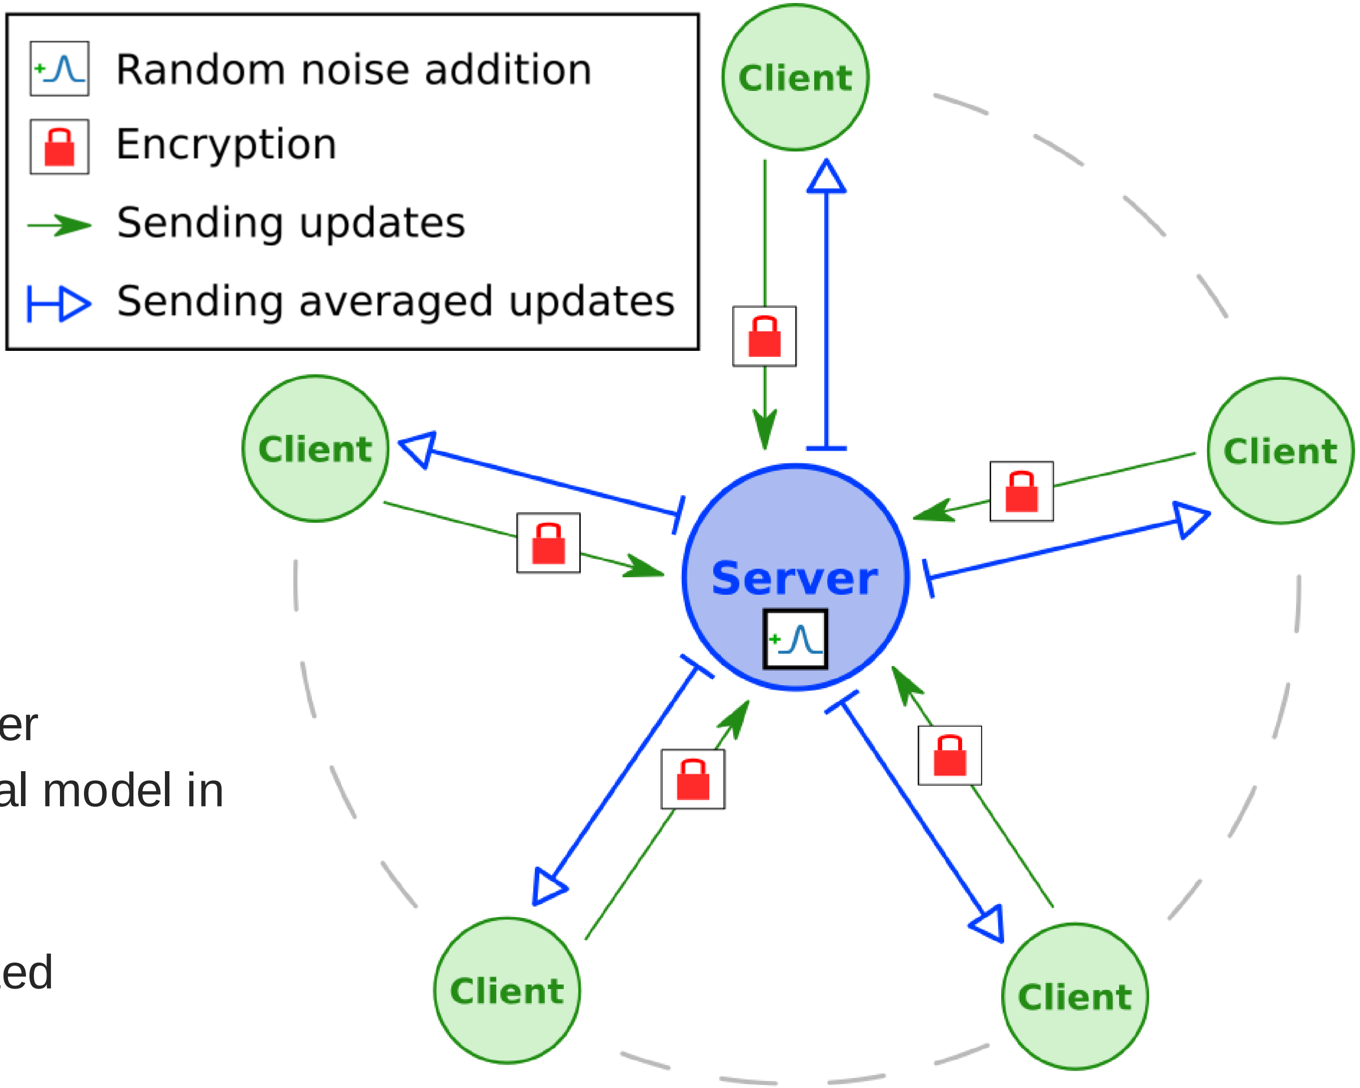

Private FL with distributed noise

- The clients add a **Gaussian noise** to their updates.

  → the average of the individual noises would still be Gaussian

- Then, the clients **encrypt the noised updates**.

Problem

The noised updates have to be **quantised** and **bounded**, which modifies the aggregated noise and make it hard to analyze.

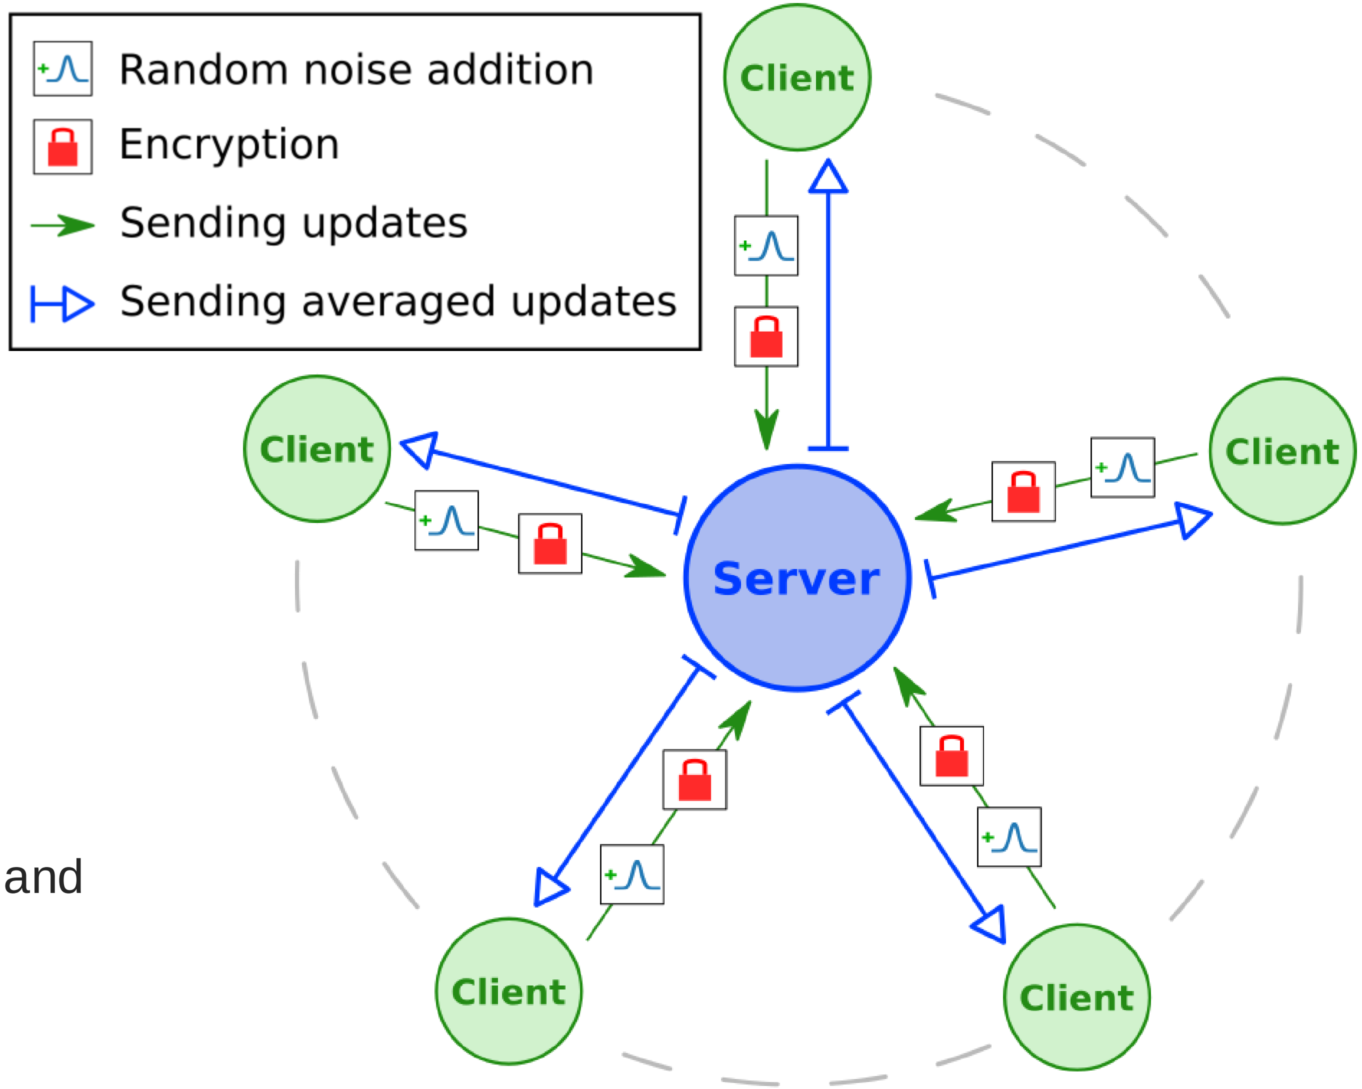

State of the art on FL with discrete distributed DP

- ***CPSGD: Communication-efficient and differentially-private distributed SGD***, Agarwal et al., 2018, ***Tight differential privacy for discrete-valued mechanisms and for the subsampled Gaussian mechanism using FFT***, Koskela et al., 2021 and ***The Poisson binomial mechanism for unbiased federated learning with secure aggregation***, Chen et al., 2022
  - DP guarantees for multidimensional binomial mechanism only in specific cases

- ***The discrete Gaussian for differential privacy***, Cannone et al., 2020 and ***The distributed discrete Gaussian mechanism for federated learning with secure aggregation***, Kairouz et al., 2021
  - The discrete Gaussian distribution is **not stable by addition** → not ideal for distributed DP

- ***The Skellam mechanism for differentially private federated learning***, Agarwal et al., 2021

Beyond the state of the art

Common shortcomings of the state of the art

- Do not reach the Gaussian mechanism's DP guarantees
- Needs the **quantisation scale** to tend to zero to approach the Gaussian mechanism's DP guarantees
  → the communication cost approaches infinity
- Involved mathematical analysis

Our promise

+ **Very same** DP guarantees as the Gaussian mechanism
+ DP guarantees **uncoupled** from the quantisation scale
+ Straightforward analysis

Poisson quantisation

New **stochastic quantisation operator** based on Poisson distribution:

$$Q_{s,\mu} : x \in]\mu; +\infty[\mapsto sY + \mu$$

where $\mu \in s\mathbb{Z}$ and $Y \sim \mathcal{P}\left(\frac{x-\mu}{s}\right)$

- μ: common lower bound to the values to quantise
- s: quantisation scale

$Q_{s,\mu}(x) \in s\mathbb{Z}$ and $\mathbb{E}\left[Q_{s,\mu}(x)\right] = x \rightarrow$ actual unbiased quantisation operator

Poisson quantisation as a post-processing

Commutes with the sum (the sum of the quantised values has the same distribution as the quantised sum)

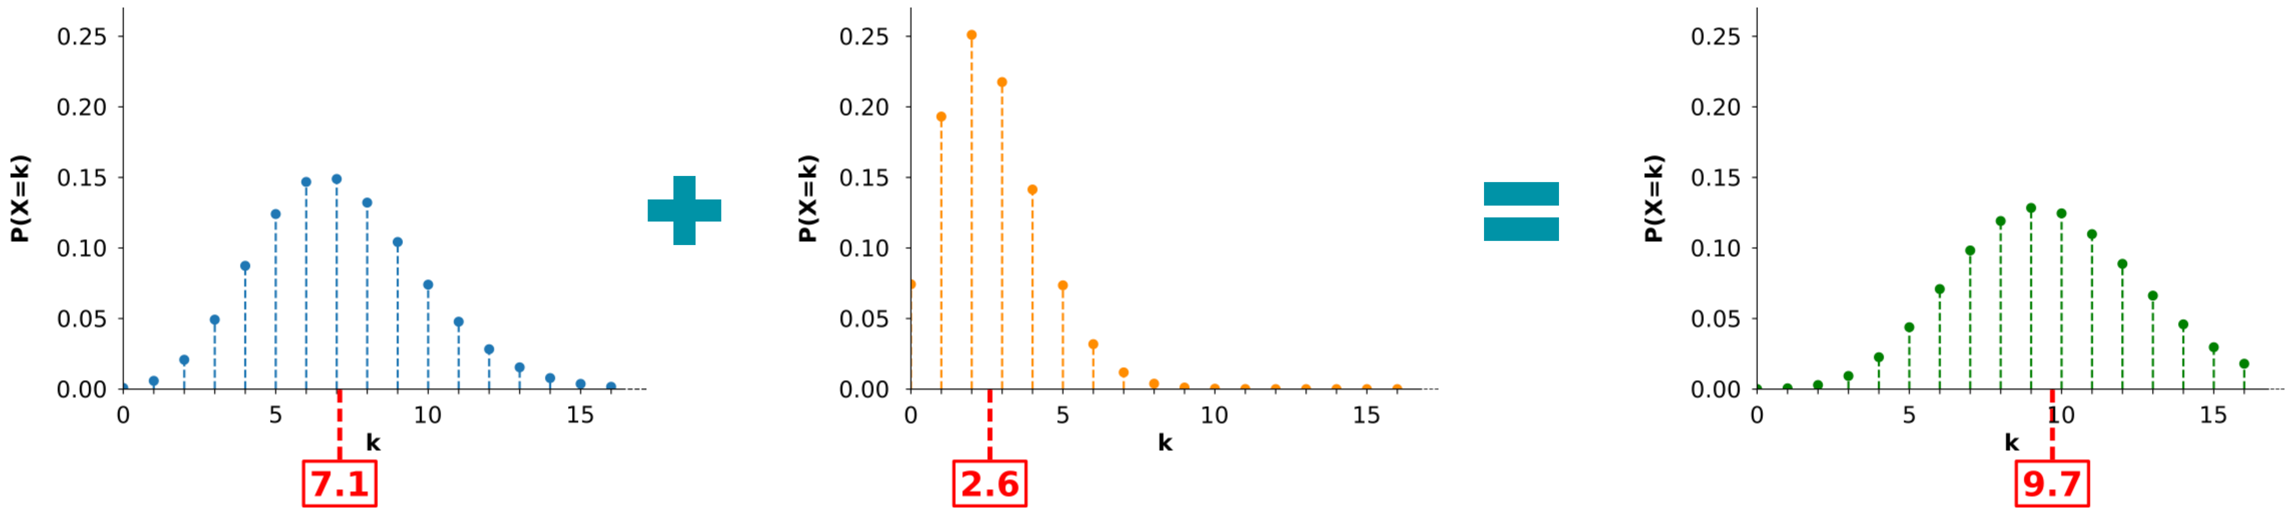

→ quantisation can be viewed as a **post-processing** → no impact on the DP guarantees

Isolating the sum out

$$\sum_{i=1}^{K} \underbrace{Q_{s,\mu}\left(x_i + \mathcal{N}\left(0,\sigma^2\right)\right)}_{\text{individual value sent to the server}} \overset{d}{\underset{\text{stability of the Poisson distribution}}{=}} Q_{s,K\mu}\left(\sum_{i=1}^{K}\left(x_i + \mathcal{N}\left(0,\sigma^2\right)\right)\right) \overset{d}{\underset{\text{stability of the normal distribution}}{=}} \overbrace{Q_{s,K\mu}}^{\text{post-processing}} \overbrace{\left(\underbrace{\sum_{i=1}^{K} x_i}_{\text{information of interest}} + \mathcal{N}\left(0,K\sigma^2\right)\right)}^{\text{information protected by DP}}$$

- Actually implemented process
  - Account for HE
  - Unclear DP analysis
- Virtual equivalent process
  - Not compatible with HE
  - Straightforward DP analysis

Additional issues

- Poisson quantisation requires a **common lower bound** for the noised updates.
  - Lower bound for the unnoised updates already ensured by clipping
  - Lower bound of the Gaussian noise « thanks to » the imperfection of the sampling algorithms (Box-Muller in Cartesian and polar forms, ziggurat)

- Poisson distribution is **not bounded** but the encryption automatically applies a modulo operation
  - The modulo operation does not affect the DP guarantees (**post-processing**):

$$\sum_{i=1}^{K}(x_i \mod N) \mod N = \sum_{i=1}^{K} x_i \mod N$$

  - The encrypted values **very rarely** (probability ~ 10^{-5} or lower) go beyond the 26-bit modulus we use and, in practice, it does not affect the model accuracy either.

Impact of the quantisation scale

Trade-off between communication/computation cost and accuracy

- The quantisation scale s impacts the **model accuracy** in two ways:
  - **Precision** of the updates
  - Variance of the **Poisson noise** = s(x-μ)

- s influences the **size of the ciphertexts** and then the choice of the **plaintext modulus** (because the values are multiplied by 1/s before encryption to have only integers)

Uncoupled from privacy

- The quantisation scale s does not impact the **DP guarantees** (because the quantisation can be seen as a post-processing)

Single/multi-key HE scheme

Single-key set-up

- There is only one decryption key.
- Anyone owning this key can decrypt.

k-out-of-n threshold set-up

- There are n shares of the decryption key distributed among the n clients.
- At least k shares are needed for decryption.

Single key set-up	k-out-of-n threshold set-up
+ **No communication** needed between the clients.	- Interclient **communication** needed for decryption
- One client may **collude** with the server and share the decryption key with it.	+ **Robust** to up to k-1 clients that collude with the server.

Differential privacy analysis

Same DP guarantees as the **Gaussian mechanism** without additional analysis:

- Privacy amplification by **subsampling**

- DP guarantee from the **point of view of a colluding participant**
  - $(1 - \chi) K\sigma^2$ instead of $K\sigma^2$, where K is the number of participants and χ is the ratio of colluding participants

- DP guarantee from the **point of view of non-colluding participant:** $(K-1)\sigma^2$ instead of $K\sigma^2$

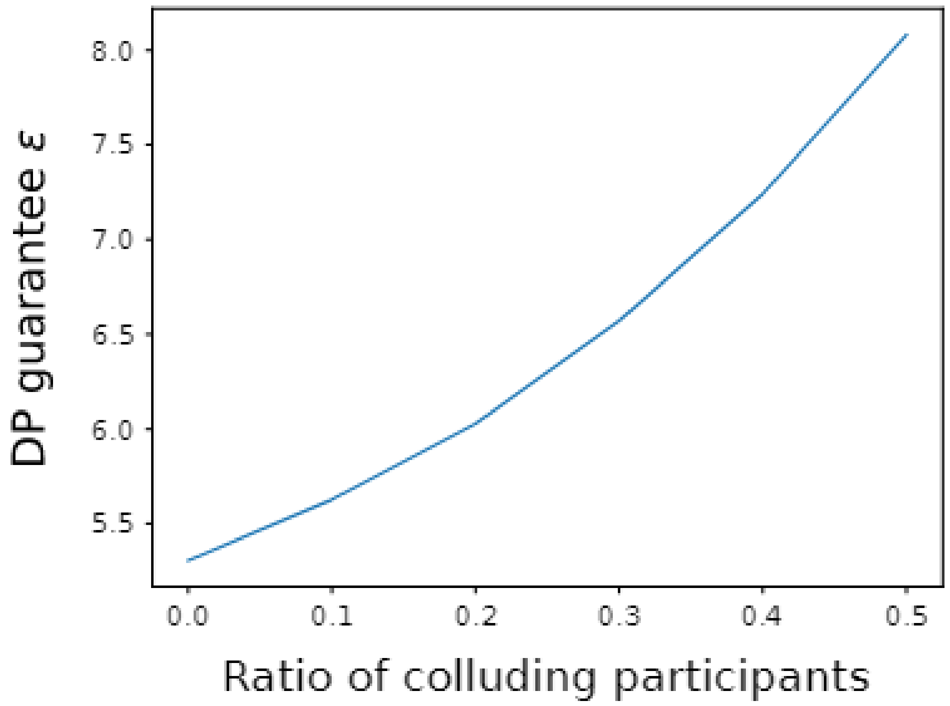

- With $\boldsymbol{\delta = 10^{-5}}$, we get **ε = 5.306** for an end-user, **ε = 5.313** for a non-colluding participant

- The noise induced by the Poisson quantisation may also help to sanitise the data (cf. Chen et al., 2022) but we only consider it as a post-processing

Experimental results (accuracy)

- Experiments on FEMNIST (Federated Extended MNIST)
- M = 3596 clients, K = 1000 participants per round, T = 100 rounds, quantization scale s = 10^{-4}, clipping bound S = 1
- The DP guarantees increase whith the number of rounds (queries) → need to decrease it
- Role of M and K:
  - When K/M decreases, the **probability for a client to be selected** decreases and so does the privacy cost
  - With K/M fixed, when K increases, the individual information is more **diluted** so the privacy cost decreases.

  → Increase M and K
- An unweighted mean is easier to sanitise (lower **sensitivity**)

	Accuracy
State of the art	84.6%
Decrease the number of learning rounds T	83.58%
Increase M and K	81.04%
Assign same coefficients	80.21%
Clipping of the updates	79.26%
Quantisation	79.03%
Modulo operation	79.07%
Adding random noise	**76.84%**

Influence of successive adaptations on accuracy

- The model trained with Gaussian noise but without quantisation or modulo operation has the same accuracy 76.84% → The quantisation comes at **no cost** !

Experimental results (computational overhead due to HE)

- FedAvg only needs homomorphic addition
  - 20h for training / 100 rounds = 12 min per FL round
  - Up to 8 seconds of HE latency (encryption, evaluation and decryption)

  → Only **1.1%** of time overhead due to HE

- Security level: 128 bits

- BFV cryptosystem enables massive batching

  → 60 ciphertexts of 8192 slots contain the 486,654 updates

- HE used to protect the data but, as a **bonus effect**, the model parameters are also protected from the server

Users (keys)	1	1000	3596
Participants (additions to perform)	1000	1000	1000
Context generation	0,0036	0,0073	0,0073
Key generation	0,0019	1,5545	5,5885
Encoding	0,0062	0,0072	0,0072
Encryption	0,1509	0,2358	0,2355
Evaluation	1,2761	3,0368	3,0544
Total decryption latency	0,0279	1,4376	4,6242

Computation time (in seconds) of HE operations with a 26-bit modulus for the full 486654 weights model

4. Conclusion and perspectives

Contribution summary

- **A secure federated learning framework using homomorphic encryption and verifiable computing** (joint work with Abbass Madi, Oana Stan, Aurélien Mayoue, Cédric Gouy-Pailler and Renaud Sirdey)
  - Uses **verifiable computing** to address the computation errors from the server.
  - Lacks DP but initiated reflections about private FL.

- **Combining homomorphic encryption and differential privacy in federated learning** (joint work with Marina Checri, Renaud Sirdey, Oana Stan and Cédric Gouy-Pailler)
  - Seamlessly combines HE and DP to protect **FL** thanks to a novel **quantisation operator.**

- **SPEED: Secure, PrivatE, and Efficient Deep learning** (joint work with Rafaël Pinot, Martin Zuber, Cédric Gouy-Pailler and Renaud Sirdey)
  - Extends the scope of threats of PATE framework to the **honest-but-curious server** by adding a **HE layer** on the server side $\rightarrow$ homomorphic argmax
  - Thorough analysis of the privacy cost as a function of the ratio of colluding data owners

- **When approximate design for fast homomorphic computation provides differential privacy guarantees** (joint work with Martin Zuber, Renaud Sirdey, Oana Stan and Cédric Gouy-Pailler)
  - Proposes **SHIELD** (Secure and Homomorphic Imperfect Election via Lightweight Design), a homomorphic argmax operator whose approximate behaviour lighten homomorphic computation while ensuring DP
  - Integrates SHIELD in SPEED workflow.

Delights and hardships of a PhD student

- Started with COVID → writing of SPEED article and first submission in full remote work

- Research is about meeting people and collaborating: Renaud and Cédric of course, Martin, Rafaël, Oana, Pierre-Emmanuel, Aymen, Marina, Abbass, Aurélien

- Many tries that fail and some that success (e.g. transfer learning and quantisation)

Conclusion

- Training a model collaboratively and without harming the data privacy is **challenging** and has a **cost** (computation, accuracy)

- DP and HE are **complementary**: they address threats coming from **different actors**

- HE is very demanding in terms of computation (and time). It thus has to be used with **parsimony**, and in **tailored protocols** that reduce the homomorphic computations to the minimum.

- Ready-to-use technology to some extent

Perspectives

- Thoroughly and formally study SHIELD algorithm and its extensions
- Extend the threat model beyond honest-but-curious server via **verifiable computing**
- Use the noise induced by HE to ensure DP
- Aggregation robust to Byzantine clients
- Propose **systematic approaches** to choose the required DP guarantees ε and δ
- Broader reflections on DP (using ignorance as additional noise, metric-based DP, data-dependent vs data-independent DP)

Publications and talks

- **Published articles and preprints:**
  - Grivet Sébert, A., Pinot, R., Zuber, M., Gouy-Pailler, C., Sirdey, R. (2021). **SPEED: Secure, PrivatE, and Efficient Deep learning**. *Machine Learning*, 110(4), 675-694, presented at ECML-PKDD 2020
  - Madi, A., Stan, O., Mayoue, A., Grivet-Sébert, A., Gouy-Pailler, C., Sirdey, R. (2021). **A secure federated learning framework using homomorphic encryption and verifiable computing**. In *2021 Reconciling Data Analytics, Automation, Privacy, and Security: A Big Data Challenge (RDAAPS)* (pp. 1-8). IEEE
  - Grivet Sébert, A., Checri, M., Sirdey, R., Stan, O., Gouy-Pailler, C. (2022). **Combining homomorphic encyption and differential privacy in federated learning**. *arXiv preprint arXiv:2205.04330* (submitted)
  - Grivet Sébert, A., Zuber, M., Sirdey, R., Stan, O., Gouy-Pailler, C. (2023). **When approximate design for fast homomorphic computation provides differential privacy guarantees**. *arXiv preprint arXiv:2304.02959* (to be submitted)

- **Popular science paper:**
  - Sirdey, R., Grivet Sébert, A., Gouy-Pailler, C. (2022). **[Cahier technique] Cryptographie homomorphe: l'art de partager sans divulguer** (in French). *Industrie et technologies 1054, juin 2022*

- **Talks:**
  - Talk about **Protecting Data from all Parties: Combining FHE and DP in Federated Learning** at Paris Privacy Preserving AI Meetup, June, 8th, 2022
  - **Machine learning without jeopardising the training data** at Principles of Distributed Learning (PODL) workshop in ACM Principles of Distributed Computing (PODC) 2022, Salerno, Italy, July, 25th, 2022
  - **Privacy in collaborative learning: differential privacy meets homomorphic encryption** at CNIL Privacy Research day, Paris France, June, 14th, 2023 (to be presented)

Some references (1)

- Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). **Deep learning with differential privacy**. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*
- Agarwal, N., Suresh, A. T., Yu, F. X. X., Kumar, S., & McMahan, B. **CPSGD: Communication-efficient and differentially-private distributed SGD**. *NeurIPS*, 31:7564–7575, 2018
- Agarwal, N., Kairouz, P., & Liu, Z. (2021). **The Skellam mechanism for differentially private federated learning**. *Advances in Neural Information Processing Systems*, *34*
- Brakerski, Z., Gentry, C., & Vaikuntanathan, V. (2014). (Leveled) fully homomorphic encryption without bootstrapping. *ACM Transactions on Computation Theory (TOCT)*
- Chakraborty, O. & Zuber, M (2022). **Efficient and accurate homomorphic comparisons**. In *WAHC'22*, pages 35–46. Association for Computing Machinery, 2022
- Chatzikokolakis, K., Andrés, M. E., Bordenabe, N. E., & Palamidessi, C. (2013). **Broadening the scope of differential privacy using metrics**. In *PETS 2013*
- Chen, W.-N., Ozgur, A., & Kairouz, P. (2022). **The Poisson binomial mechanism for unbiased federated learning with secure aggregation**, in *ICML*, 2022
- Cheu, A., Smith, A., Ullman, J., Zeber, D., & Zhilyaev, M. (2019). **Distributed differential privacy via shuffling**. In *Advances in Cryptology–EUROCRYPT 2019: 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques*
- Chillotti, I., Gama, N., Georgieva, M., & Izabachene, M. (2016). **Faster fully homomorphic encryption: Bootstrapping in less than 0.1 seconds**. In *Advances in Cryptology–ASIACRYPT 2016: 22nd International Conference on the Theory and Application of Cryptology and Information Security*
- Cyffers, E., & Bellet, A. (2022). **Privacy amplification by decentralization**. In *International Conference on Artificial Intelligence and Statistics* (pp. 5334-5353). PMLR
- Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2006). **Calibrating noise to sensitivity in private data analysis**. In *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006*

Some references (2)

- Dwork, C., & Roth, A. (2014). **The algorithmic foundations of differential privacy**. *Foundations and Trends® in Theoretical Computer Science*
- Erlingsson, Ú., Feldman, V., Mironov, I., Raghunathan, A., Talwar, K., & Thakurta, A. (2019). **Amplification by shuffling: From local to central differential privacy via anonymity**. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*
- Fan, J., & Vercauteren, F. (2012). **Somewhat practical fully homomorphic encryption**. *Cryptology ePrint Archive*.
- Gentry, C. (2009). **Fully homomorphic encryption using ideal lattices**. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*
- Goryczka, S. & Xiong, L. (2015) **A comprehensive comparison of multiparty secure additions with differential privacy**. *IEEE transactions on dependable and secure computing*
- Hao, M., Li, H., Luo, X., Xu, G., Yang, H., & Liu, S. (2019). **Efficient and privacy-enhanced federated learning for industrial artificial intelligence**. *IEEE Transactions on Industrial Informatics*, *16*(10), 6532-6542.
- Iliashenko, I. & Zucca, V. (2021). **Faster homomorphic comparison operations for BGV and BFV**. In *Proceedings on Privacy Enhancing Technologies*, 2021
- Kairouz, P., Liu, Z., & Steinke, T. (2021, July). **The distributed discrete Gaussian mechanism for federated learning with secure aggregation**. In *International Conference on Machine Learning* (pp. 5201-5212). *PMLR*
- Koskela, A., Jälkö, J., Prediger, L., & Honkela, A. (2021). **Tight differential privacy for discrete-valued mechanisms and for the subsampled Gaussian mechanism using FFT**. In *International Conference on Artificial Intelligence and Statistics*, pages 3358–3366. *PMLR*, 2021
- McMahan, H. B., Moore, E., Ramage, D., & y Arcas, B. A. (2016). **Federated learning of deep networks using model averaging**. *arXiv preprint arXiv:1602.05629*, *2*.

Some references (3)

- McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). **Communication-efficient learning of deep networks from decentralized data**. In *Artificial intelligence and statistics* (pp. 1273-1282). PMLR.
- Papernot, N., Abadi, M., Erlingsson, U., Goodfellow, I., & Talwar, K. (2016). **Semi-supervised knowledge transfer for deep learning from private training data**. *ICLR*, 2016
- Papernot, N., Song, S., Mironov, I., Raghunathan, A., Talwar, K., and Erlingsson, U. (2018) **Scalable private learning with PATE**. *ICLR*, 2018
- Paillier, P. (1999). **Public-key cryptosystems based on composite degree residuosity classes**. In *Advances in Cryptology—EUROCRYPT'99: International Conference on the Theory and Application of Cryptographic Techniques*
- Ryffel, T., Trask, A., Dahl, M., Wagner, B., Mancuso, J., Rueckert, D., & Passerat-Palmbach, J. (2018). **A generic framework for privacy preserving deep learning**. *arXiv preprint arXiv:1811.04017*.
- Ryffel, T., Tholoniat, P., Pointcheval, D., & Bach, F. (2022). **Ariann: Low-interaction privacy-preserving deep learning via function secret sharing**. *Proceedings on Privacy Enhancing Technologies*, *2022*(1), 291-316.
- Sabater, C., Bellet, A., & Ramon, J. (2020). **Distributed differentially private averaging with improved utility and robustness to malicious parties**. *arXiv preprint arXiv:2006.07218*
- Stevens, T., Skalka, C., Vincent, C., Ring, J., Clark, S., & Near, J. (2022). **Efficient differentially private secure aggregation for federated learning via hardness of learning with errors**. In *31st USENIX Security Symposium (USENIX Security 22)*
- Truex, S., Baracaldo, N., Anwar, A., Steinke, T., Ludwig, H., Zhang, R., & Zhou, Y. (2019, November). **A hybrid approach to privacy-preserving federated learning**. In *Proceedings of the 12th ACM workshop on artificial intelligence and security* (pp. 1-11).
- Wang, J., Schuster, R., Shumailov, I., Lie, D., & Papernot, N. (2022). **In differential privacy, there is truth: on vote-histogram leakage in ensemble private learning**. *Advances in Neural Information Processing Systems*, *35*