UNECE Expert Meeting on Statistical Data Confidentiality 2023

UNECE

INSIGHTS INTO PRIVACY-PRESERVING FEDERATED MACHINE LEARNING FROM THE PERSPECTIVE OF A NATIONAL STATISTICAL OFFICE

- Benjamin Santos, Julian Templeton, Rafik Chemli, Saeid Molladavoudi (Statistics Canada)
- Francesco Pugliese, Erika Cerasti, Massimo De Cubellis (ISTAT)
- Matjaz Jug (Statistics Netherlands)
- Statistics Canada, Statistics Netherlands, ISTAT

Wiesbaden (Germany)

28° September 2023

Introduction

- National Statistical Offices (NSOs) have a wealth of data but are limited in some ways on what can be collected
  Sensitive topics, legally protected data, organizational data, …

- To get insights from data that cannot be collected, Federated Learning is a potential solution

- To expand the research being done by NSOs, ISTAT, Statistics Canada, and Statistics Netherlands have expanded previous research on Federated Learning and its potential utility for NSOs

Introduction

- This work explores applying Federated Learning (FL) on a Human Activity Recognition dataset by testing the following

  ✓ Different federated aggregation strategies

  ✓ Using Differential Privacy to better protect the locally trained Machine Learning (ML) models

  ✓ Homomorphically encrypting model weights to hide their values from the central authority (aggregator)

What is Federated Learning?

- FL allows a centralized ML model to be trained on data residing on distributed client devices, with the locally updated models then being aggregated

- This allows analytics to be derived from data sources that cannot be collected

- The performance of the trained models can reach similar performance of centralized approaches, but a careful selection of the hyperparameters and the aggregation method is important

What is Federated Learning?

- While this allows previously impossible analytics to be possible to generate, the approach on its own does not remove all privacy risks

  ✓ Locally trained client models can still be attacked

- Other Privacy Enhancing Technologies (PETs) can be used in conjunction with FL to defend against these concerns

  ✓ Differential Privacy (trade-off of performance vs privacy)

  ✓ Homomorphic Encryption (adds more computational complexity)

Main Federated Learning Strategies

- After training the model with data locally, a client will send the weights or gradients back to the server to be aggregated

- Within this work we test the following federated aggregation methods:

  ✓ Federated Averaging (FedAvg)
  ✓ Federated Adaptive Gradient (FedAdagrad)
  ✓ FedAdam (Federated Adam)
  ✓ FedYogi (Federated Yogi)

Main Federated Learning Strategies

- ✓ Federated Averaging (FedAvg): it is a federated learning algorithm that aims to train a global model by aggregating the local model updates from multiple clients by calculating the average of the model parameters.

- ✓ Federated Adaptive Gradient (FedAdagrad): it is a variant algorithm that exploits the adaptive gradient descent method called Adagrad. It adapts the learning rate for each model parameter based on its historical gradients, allowing the model to converge faster and achieve better performance.

Main Federated Learning Strategies

- ✓ FedAdam (Federated Adam): it is another federated learning algorithm that combines the advantages of the Adam optimizer with the federated learning setting. It employs adaptive learning rates and momentum to efficiently update the global model using the local updates from clients. The gradients computed locally by the devices are aggregated in the central server.

- ✓ FedYogi (Federated Yogi): it is a federated learning algorithm inspired by the Yogi optimizer. It incorporates elements of both adaptive learning rates and momentum to handle non-convex optimization problems in federated learning scenarios.

What is Differential Privacy?

- At a high level, DP is a PET which injects noise into data or statistics such that the results are the same whether any single datapoint within a database is or is not present within the database

- Injects noise during the training to control the amount of increased privacy with a corresponding drop in performance

- The privacy budget ϵ determines the amount of privacy to be added, where a lower ϵ adds more privacy

Experiment 1 – Aggregation Methods and Differential Privacy

- We compare different aggregation methods with and without DP being used
- Different ϵ values are used to observe the privacy/performance tradeoff
- Using DP with FL helps protect the privacy of the model weights or gradients being sent to the central authority

Experiment 1 – Aggregation Methods and Differential Privacy

Accuracy comparison over different strategies and epsilon values

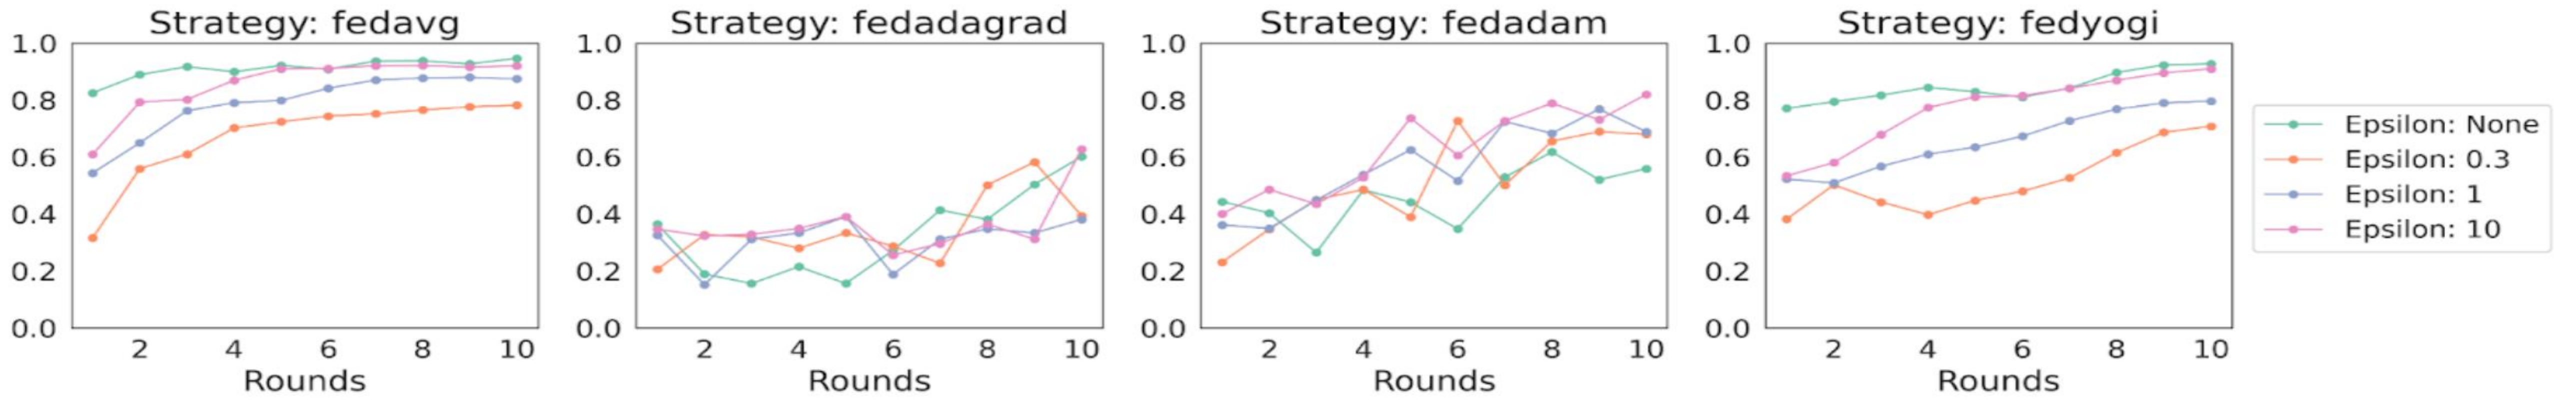

F1-Score comparison over different strategies and epsilon values

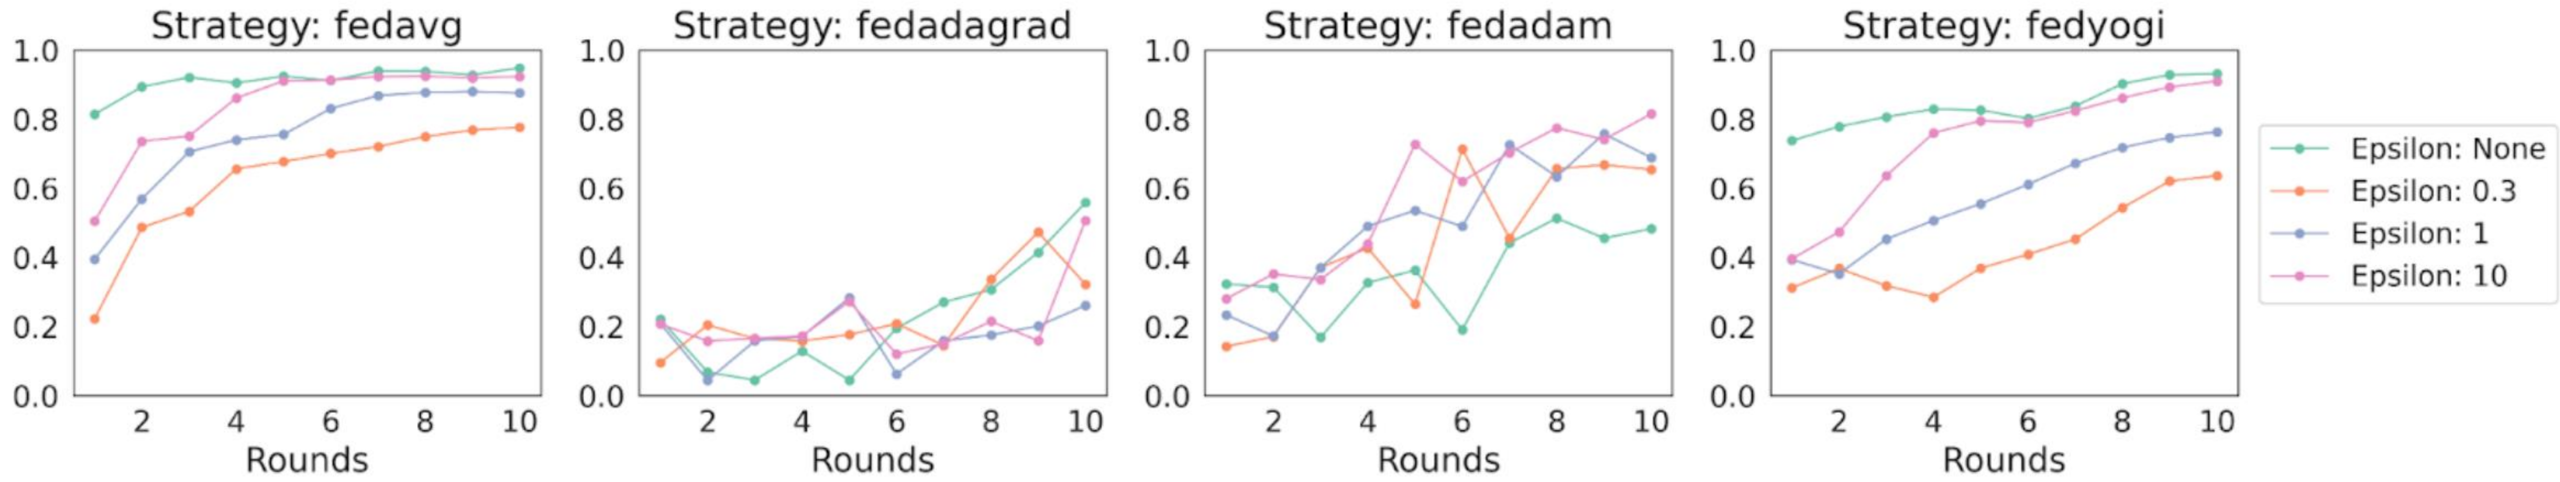

Experiment 1 – Aggregation Methods and Differential Privacy

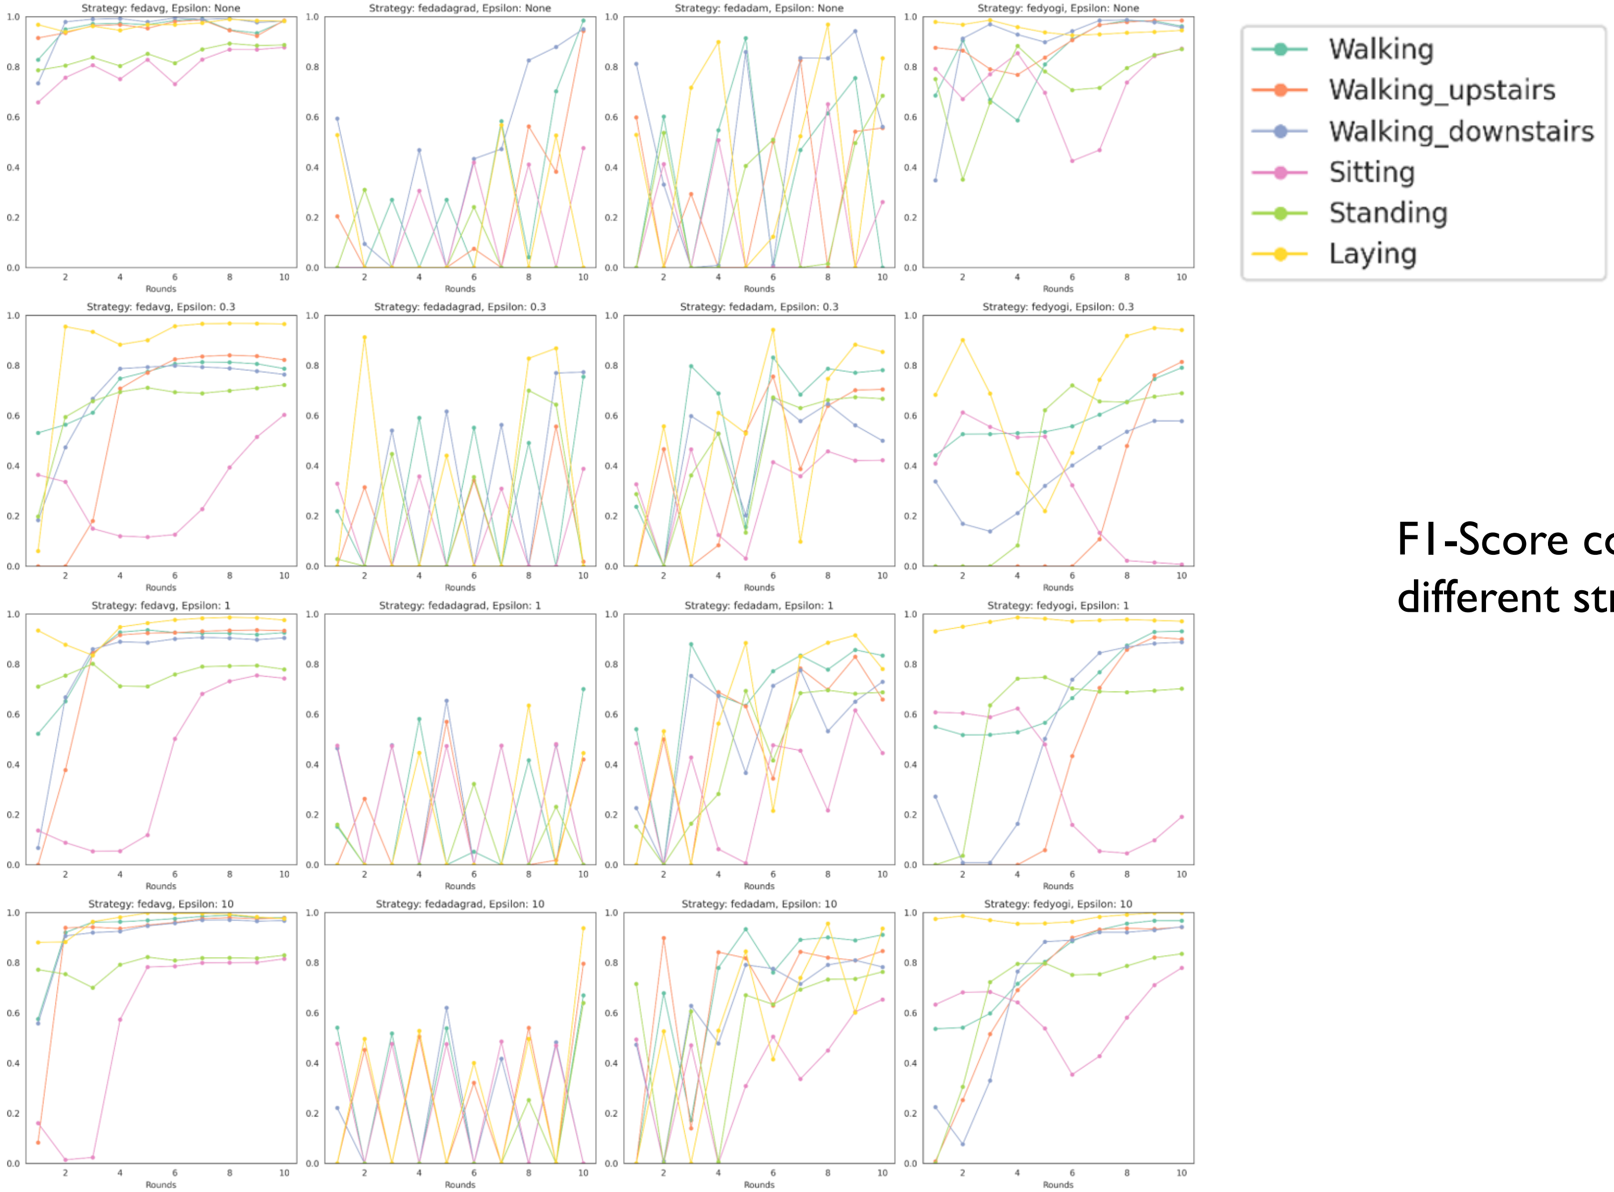

F1-Score comparison per class over different strategies and epsilon values

What is Homomorphic Encryption (HE)?

- Allows to perform arithmetic operations on encrypted data
  HE is a public-key cryptographic scheme.

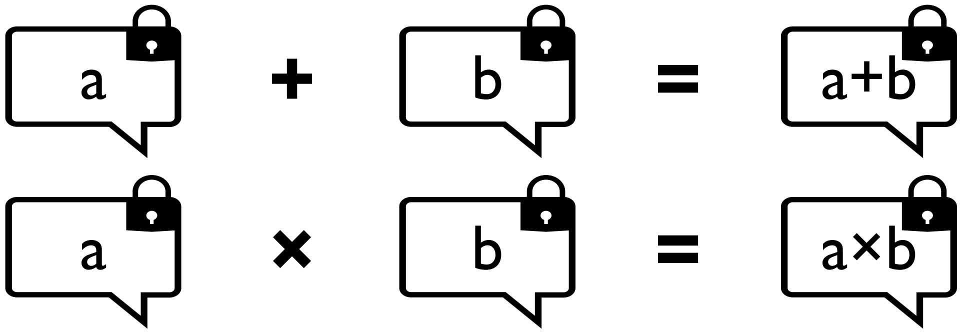

- **Application:** delegated computing! Unparalleled cryptographic security at the cost of higher computational and storage requirements.

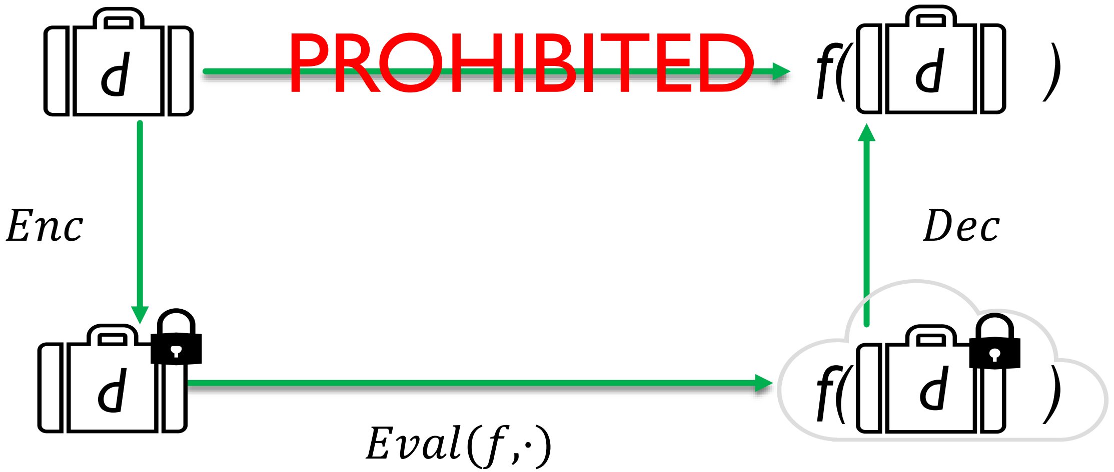

HE aggregation in a FL setting

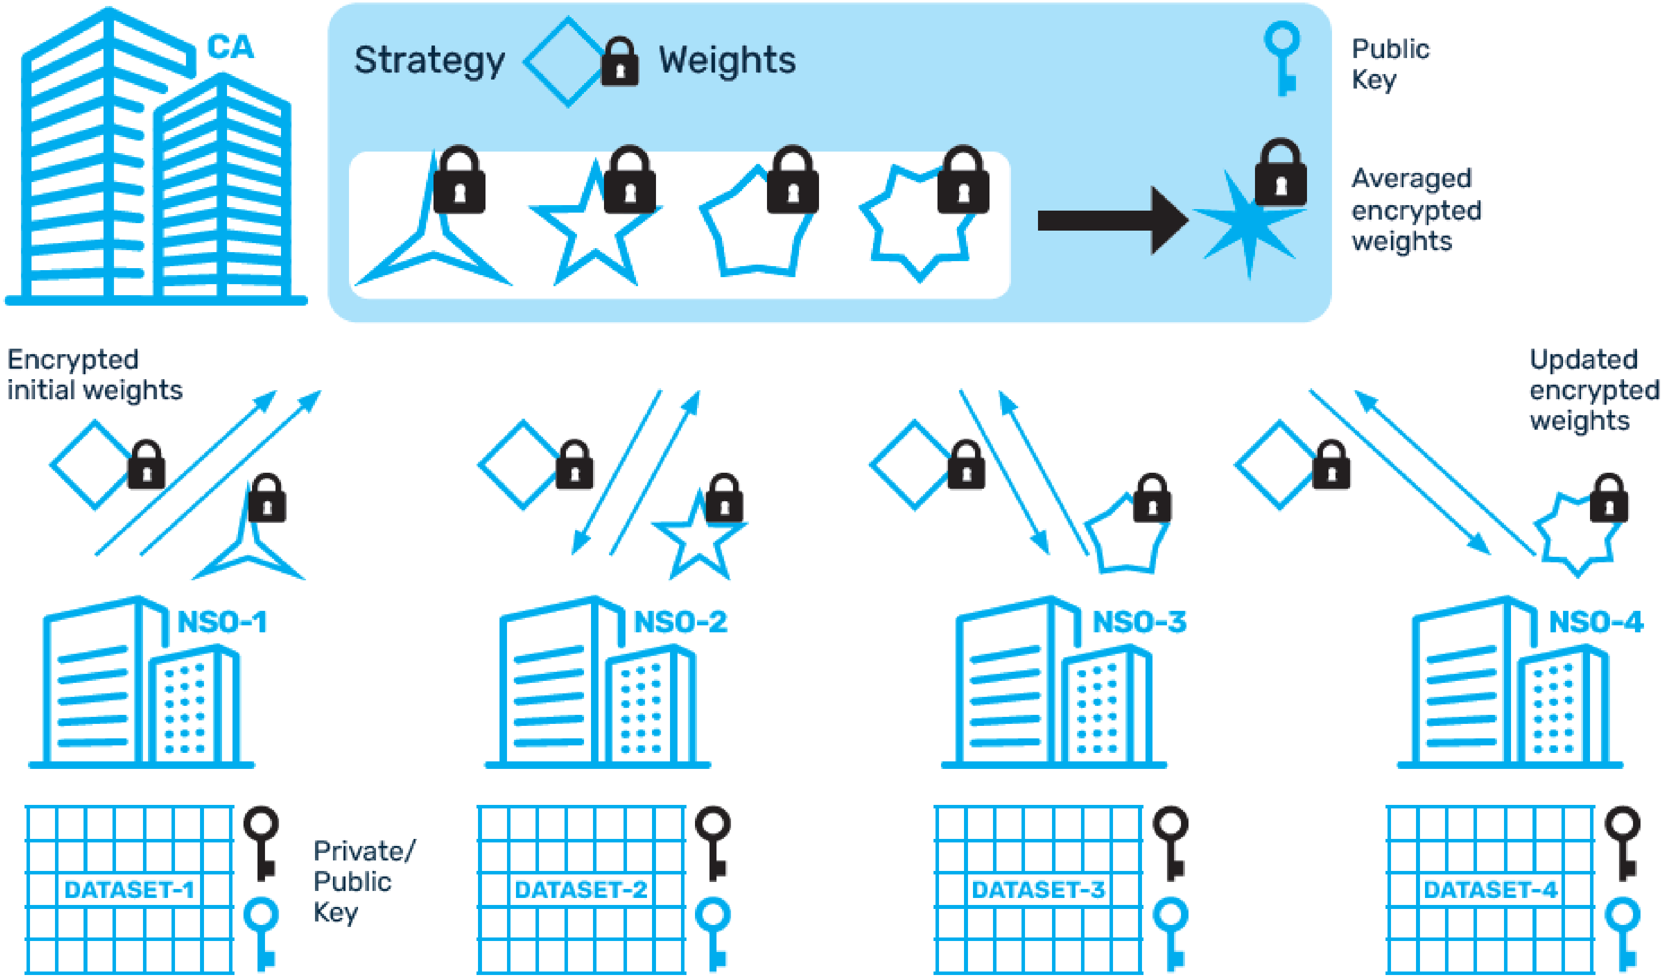

**UNITED NATIONS, 2023, UNITED NATIONS GUIDE ON PRIVACY-ENHANCING TECHNOLOGIES FOR OFFICIAL STATISTICS*

Experiment 2 – Homomorphic Encryption Aggregation Results

Strategy	**Relative training time**	**Relative RAM**	**Relative model size**	**Relative Encryption - Serialization Time**	**Relative Deserialization - Decryption Time**
FedAvg	1	1	1	1	1
eFedAvg (1)	1.04	1	8.78	23	631
eFedAvg (2)	597	1.08	4001	1237	2532

- FedAvg: model’s last layer encryption
- FedAvg (1): model’s last two layers encryption
- FedAvg (2): model’s last three layers encryption

Conclusions

- FedAvg and FedYogi perform the best in this experiment when unoptimized
- DP's effect on the performance will vary depending on the aggregation strategy and ϵ must carefully be selected
- Homomorphic Encryption can add significant time and communication costs, scaling with the amount of encrypted weights/gradients
- Overall, FL is a feasible approach to be considered by NSOs when data cannot be collected

References

Beutel, D. J., T. Topal, A. Mathur, X. Qiu, J. Fernandez-Marques, Y. Gao, L. Sani, K. H. Li, T. Parcollet, P. P. B. de Gusmão, and N. D. Lane (2022). Flower: A friendly federated learning research framework.

Cheon, J., A. Kim, M. Kim, and Y. Song (2017). Homomorphic encryption for arithmetic of approximate numbers. In Advances in Cryptology-ASIACRYPT 2017, pp. 409–437. Springer.

Dugdale, C., S. Molladavoudi, B. Santos, and J. Templeton (2022). Privacy enhancing technologies at statistics canada. In Proceedings of the Annual Meeting, Statistical Society of Canada.

Dwork, C. (2006). Differential privacy. In M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener (Eds.), Automata, Languages and Programming, Berlin, Heidelberg, pp. 1–12. Springer Berlin Heidelberg.

Gurbuzbalaban, M., U. Simsekli, and L. Zhu (2021). The heavy-tail phenomenon in sgd. In International Conference on Machine Learning. PMLR.

ICO (2022). Privacy-enhancing technologies (pets). hiips://ico.org.uk/media/about-the-ico/consultations/4021464/chapter-5-anonymisation-pets.pdf.

Kairouz, P., H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings, et al. (2021). Advances and open problems in federated learning. Foundations and Trends® in Machine Learning 14(1–2), 1–210.

9

Konečný, J., H. B. McMahan, F. X. Yu, P. Richtarik, A. T. Suresh, and D. Bacon (2016). Federated learning: Strategies for improving communication efficiency. In NIPS Workshop on Private Multi-Party Machine Learning.

References

Konečný, J., H. B. McMahan, F. X. Yu, P. Richtarik, A. T. Suresh, and D. Bacon (2016). Federated learning: Strategies for improving communication efficiency. In NIPS Workshop on Private Multi-Party Machine Learning.

McMahan, B., E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas (2017). Communication-efficient learning of deep networks from decentralized data. In Artificial intelligence and statistics, pp. 1273–1282. PMLR.

Nguyen, H., L. Phan, H. Warrier, and Y. Gupta (2022). Federated learning for non-iid data via client variance reduction and adaptive server update. arXiv preprint arXiv:2207.08391.

Nilsson, A., S. Smith, G. Ulm, E. Gustavsson, and M. Jirstrand (2018). A performance evaluation of federated learning algorithms. In Proceedings of the Second Workshop on Distributed Infrastructures for Deep Learning, DIDL '18, New York, NY, USA, pp. 1–8. Association for Computing Machinery.

Paillier, P. (1999). Public-key cryptosystems based on composite degree residuosity classes. In Advances in Cryptology -EUROCRYPT '99. Springer.

Reddi, S., Z. Charles, M. Zaheer, Z. Garrett, K. Rush, J. Konečn`y, S. Kumar, and H. B. McMahan (2020). Adaptive federated optimization. arXiv preprint arXiv:2003.00295.

Shokri, R., M. Stronati, C. Song, and V. Shmatikov (2017). Membership inference attacks against machine learning models. In 2017 IEEE symposium on security and privacy (SP), pp. 3–18. IEEE.

References

Shokri, R., M. Stronati, C. Song, and V. Shmatikov (2017). Membership inference attacks against machine learning models. In 2017 IEEE symposium on security and privacy (SP), pp. 3–18. IEEE.

Tolpegin, V., S. Truex, M. E. Gursoy, and L. Liu (2020). Data poisoning attacks against federated learning systems. In L. Chen, N. Li, K. Liang, and S. Schneider (Eds.), Computer Security – ESORICS 2020, Cham, pp. 480–501. Springer International Publishing.

UN (2023). United nations guide on privacy-enhancing technologies for official statistics: case study 15. Technical report, United Nations Committee of Experts on Big Data and Data Science for Official Statistics,
New York.

UNECE (2023). Input-privacy preservation report. Technical report, High-Level Group for the Modernisation
of Official Statistics, Brussels, Belgium.

Varno, F., M. Saghayi, L. Rafiee Sevyeri, S. Gupta, S. Matwin, and M. Havaei (2022). Adabest: Minimizing client drift in federated learning via adaptive bias estimation. In S. Avidan, G. Brostow, M. Cissé, G. M. Farinella, and T. Hassner (Eds.), Computer Vision – ECCV 2022, Cham, pp. 710–726. Springer Nature Switzerland.

References

Yousefpour, A., I. Shilov, A. Sablayrolles, D. Testuggine, K. Prasad, M. Malek, J. Nguyen, S. Ghosh, A. Bharadwaj, J. Zhao, G. Cormode, and I. Mironov (2022). Opacus: User-friendly differential privacy library in pytorch.

Zaheer, M., S. Reddi, D. Sachan, S. Kale, and S. Kumar (2018). Adaptive methods for nonconvex optimization. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.), Advances in Neural Information Processing Systems, Volume 31. Curran Associates, Inc.

Zanussi, Z., B. Santos, and S. Molladavoudi (2021). Supervised text classification with leveled homomorphic encryption. In Proceedings 63rd ISI World Statistics Congress, Volume 11, pp. 16

Acknowledgements

Thank You!

Questions?