Sadalītā mācīšanās ar pielietojumiem laikrindas datiem.

Distributed learning with applications for time series data.

Autors: Audris Arzovs aa17083
Promocijas darba vadītājs: LU docents Kārlis Freivalds

Dalītā mācīšanās (Distributed learning)

ML modeļu decentralizēta trenēšana nodrošinot ML algoritmu darbību uz robežierīcēm.

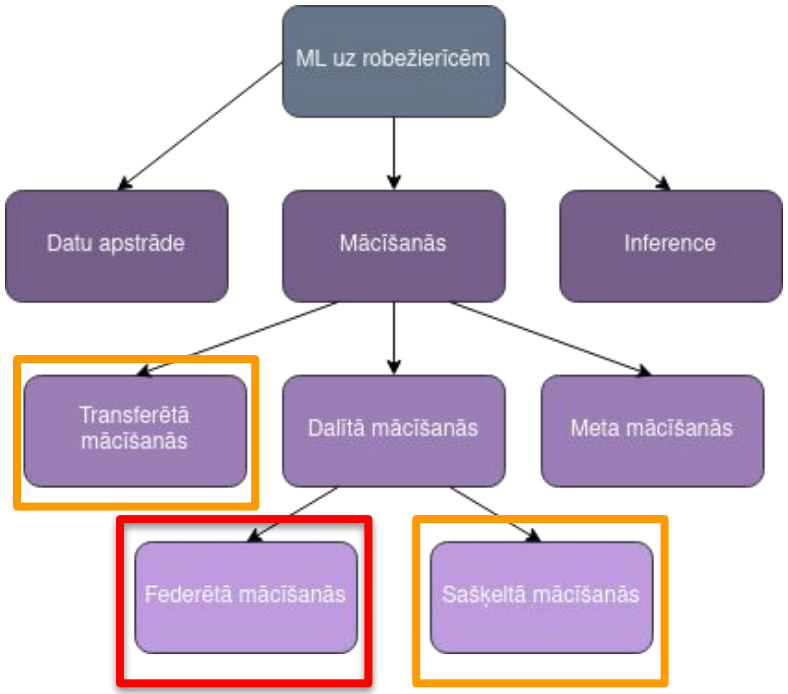

W. Li, H. Hacid, E. Almazrouei, and M. Debbah, “A review and a taxonomy of edge machine learning: Requirements, paradigms, and techniques,” arXiv preprint arXiv:2302.08571, 2023.

Laikrindas datu izmantošana

- Divi izmantošanas virzieni
  - Atlikušā dzīves ilguma paredzēšana dažādām energoelektronikas komponentēm, piemēram, tranzistoriem (**power electronics**);
  - Kritiskās infrastruktūras raksturojošo datu analīze (**IoT-Edge-Cloud continuum**).

Goals of Power electronics

- Remaining useful life analysis for:
  - Batteries (e.g. electric vehicle);
  - MOSFETs etc.;
- Failure detection for:
  - Smart grids;
  - Specific electronic components.
- Maintaining privacy for different manufacturers while providing global value with federated learning.

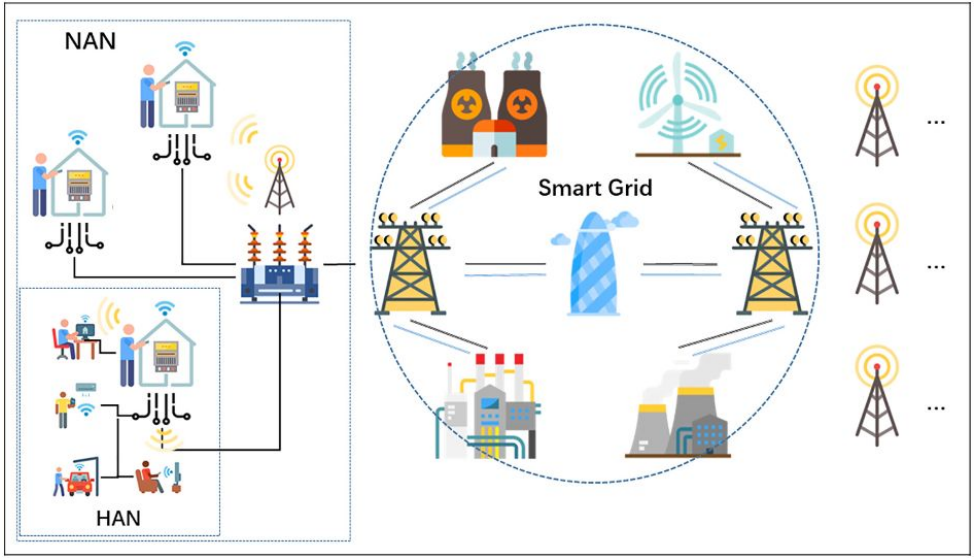

H. Cao, S. Liu, R. Zhao, and X. Xiong, “Ifed: A novel federated learning framework for local differential privacy in power internet of things,” International Journal of Distributed Sensor Networks, vol. 16, no. 5, p. 1550147720919698, 2020.

Goals of IoT-Edge-Cloud Continuum

Maintaining privacy of machine learning models, e.g., that predict certain characteristics of the critical infrastructure related to water quality assessment facilities and building structural integrity with Differential privacy and Federated learning.

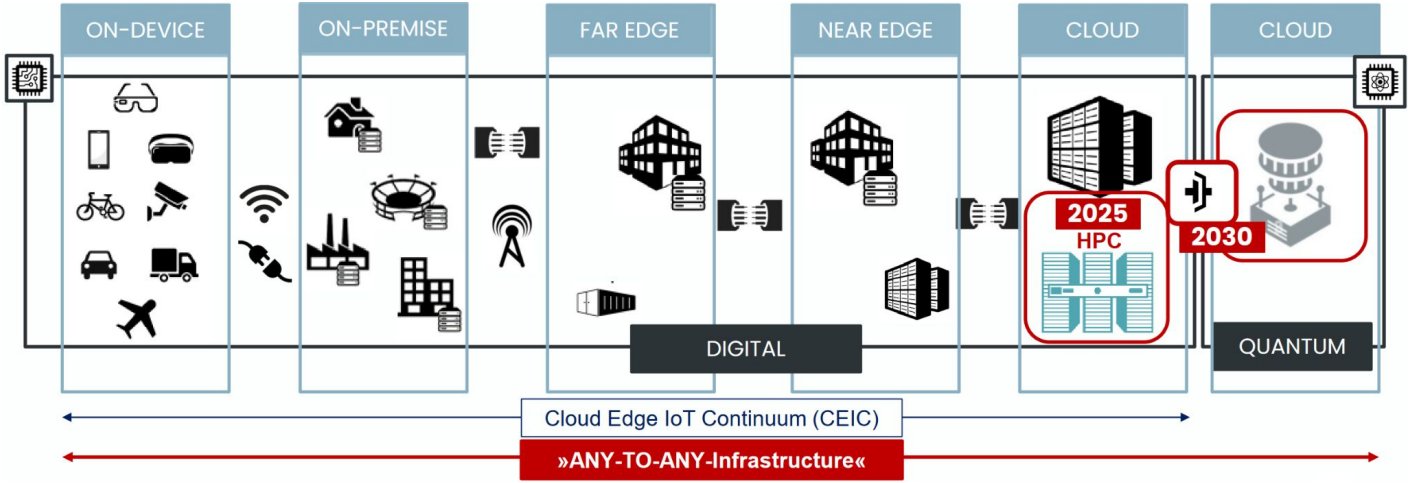

Fritz, M. General Challenges for a computing continuum. Available online: https://eucloudedgeiot.eu/wp-content/uploads/2023/05/AIOps_merged.pdf (accessed on 13.06.2023), 2023.
Arzovs, A.; Judvaitis, J.; Nesenbergs, K.; Selavo, L. Distributed Learning in the IoT–Edge–Cloud Continuum. Mach. Learn. Knowl. Extr. 2024, 6, 283-315. https://doi.org/10.3390/make6010015

Federated learning

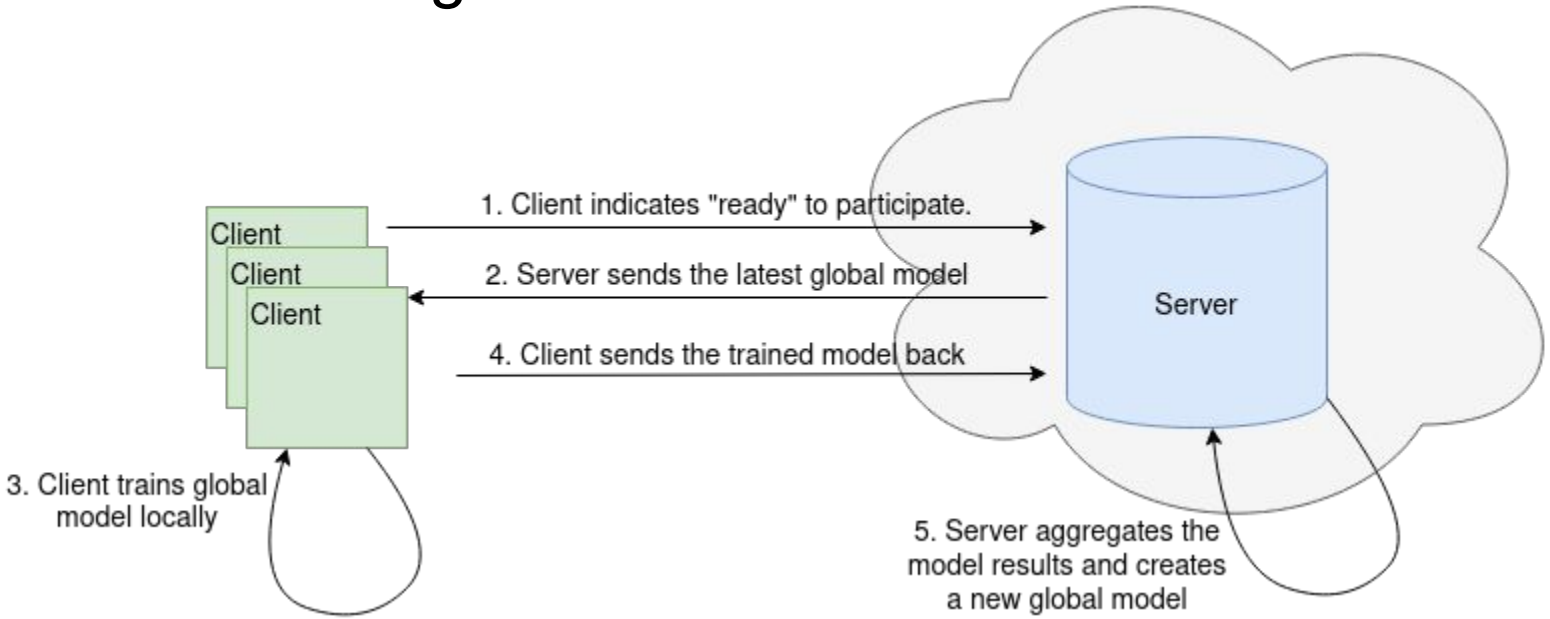

M. Rabbat, “Meta FL research presentation.” [web]. Accessible: https://semla.polymtl.ca/wp-content/uploads/2022/11/Rabbat-AsyncFL-SEMLA22.pdf, 2022. [accessed 15.04.2023].

Attack vectors against Federated learning

- Membership inference attacks;
  - Data reconstruction;
  - Property inference;
  - Attribute inference.
- Model inversion attacks;
- Poisoning attacks.
  - Model poisoning;
  - Data poisoning;
  - Backdoor attacks.

Initial input data

Data reconstruction from gradients

M. Rabbat, “Meta FL research presentation.” [web]. Accessible: https://semla.polymtl.ca/wp-content/uploads/2022/11/Rabbat-AsyncFL-SEMLA22.pdf, 2022. [accessed 15.04.2023].
Geiping, J., Bauermeister, H., Dröge, H. and Moeller, M., “Inverting gradients-how easy is it to break privacy in federated learning?”. Advances in Neural Information Processing Systems, 33, pp.16937-16947, 2020.

Defending against privacy leakage - obfuscation

(ε, δ)-Differential privacy

Creates accuracy and privacy trade-off.

$$Pr(M(x) = t) <= e^{\varepsilon} Pr(M(x') = t) + \delta$$

x, x' - neighboring datasets (differ in at most 1 record);
M - mechanism that adds noise to the data from a probability density function e.g. Laplacian or Gaussian;
ε - privacy budget;
δ - relaxation parameter for ML use cases.

C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," Journal of Privacy and Confidentiality, vol. 7, no. 3, pp. 17–51, 2016.
N. Ponomareva, H. Hazimeh, A. Kurakin, Z. Xu, C. Denison, H. B. McMahan, S. Vassilvitskii, S. Chien, and A. Thakurta, "How to dp-fy ml: A practical guide to machine learning with differential privacy," arXiv preprint arXiv:2303.00654, 2023.

Difficulty of applying Differential privacy

- Privacy and utility trade-off;
- Clear algorithm doesn’t exist yet;
- Unclear specification of privacy budget;
- Requires empirical investigation of each implementation;
- Hyperparameter configuration guidelines are scarce;
- Auditing of DP application is suggested;
- There exist many tangents from the initial definition.
  - For machine learning and other directions.
- The further from the initial data source (even at the prediction level), the better results are achieved.

N. Ponomareva, H. Hazimeh, A. Kurakin, Z. Xu, C. Denison, H. B. McMahan, S. Vassilvitskii, S. Chien, and A. Thakurta, “How to dp-fy ml: A practical guide to machine learning with differential privacy,” arXiv preprint arXiv:2303.00654, 2023.

Defending against poisoning attacks - denying access

- Homomorphic encryption allows computation on encrypted data from parties without disclosing their data to other parties;
- Requires a third party to do the computation;
- High computational cost;
- Widely researched but isn't frequently available in FL tools.

$$\forall_{m1,m2} \epsilon M, E(m1)\, o\, E(m2) = E(m1\, o\, m2)$$

Acar, A.; Aksu, H.; Uluagac, A.S.; Conti, M. A survey on homomorphic encryption schemes: Theory and implementation. ACM Computing Surveys (Csur) 2018, 51, 1–35.
Evans, D.; Kolesnikov, V.; Rosulek, M.; et al. A pragmatic introduction to secure multi-party computation. Foundations and Trends® in Privacy and Security 2018, 2, 70–246.

Defending against poisoning attacks with transparency

Because denying access is not enough. We need to establish trust between the actors.

“There is no root of trust in existing federated learning methods, i.e., from the service provider’s perspective, every client could be malicious”

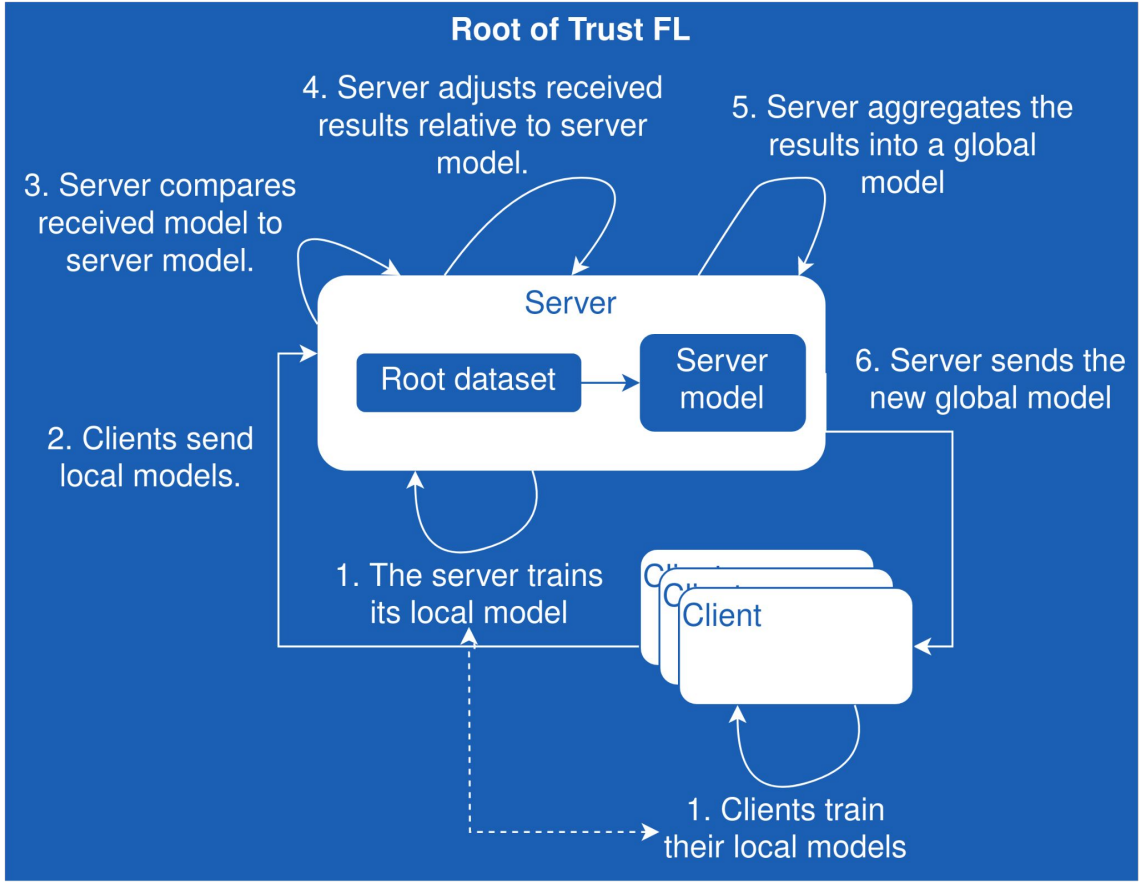

X. Cao, M. Fang, J. Liu, and N. Z. Gong, “FItrust: Byzantine-robust federated learning via trust bootstrapping,” arXiv preprint arXiv:2012.13995, 2020.

Current developments for Power electronics

- Existing electric vehicle battery remaining useful life prediction model ported to Federated learning (FL);
  - NASA randomized battery cycling dataset as data source;
  - Autoencoder for reducing the size of cycles;
  - LSTM and CNN networks for predictions.

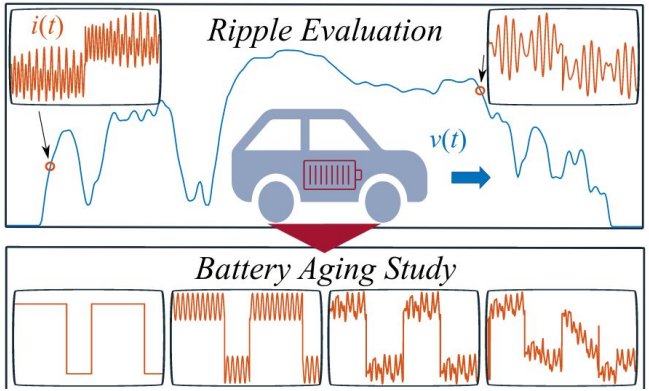

M. Bosello, C. Falcomer, C. Rossi, and G. Pau, “To charge or to sell? ev pack useful life estimation via lstms, cnns, and autoencoders,” Energies, vol. 16, no. 6, p. 2837, 2023.
E. Goldammer, M. Gentejohann, M. Schlüter, D. Weber, W. Wondrak, S. Dieckerhoff, C. Gühmann, and J. Kowal, “The impact of an overlaid ripple current on battery aging: the development of the sicwell dataset,” Batteries, vol. 8, no. 2, p. 11, 2022.
FL for battery RUL with Flower framework. Available online: https://github.com/Audris-A/FL-for-battery-RUL-with-Flower-framework (accessed on 22.01.2024).

Current developments for Power electronics

- Federated learning implemented in Flower framework;
- Added Homomorphic encryption (CKKS) using Pyfhel library;
  - Clients share keys;
  - Server and clients share context;
- Analyzed differential privacy and homomorphic encryption impact on model utility and performance.
  - With the assumption that the noise level is added with the sequential DP mechanism relative to the rounds.

FL for battery RUL with Flower framework. Available online: https://github.com/Audris-A/FL-for-battery-RUL-with-Flower-framework (accessed on 22.01.2024).
Xiong, X.; Liu, S.; Li, D.; Cai, Z.; Niu, X. A comprehensive survey on local differential privacy. Secur. Commun. Netw. 2020, 2020, 1–29.

ELEKTRONIKAS UN DATORZINĀTNU INSTITŪTS INSTITUTE OF ELECTRONICS AND COMPUTER SCIENCE
MOTE
Powerized
Digitalization of Power Electronic Applications within Key Technology Value Chains
Chips JU
Co-funded by the European Union

Current developments for Power electronics

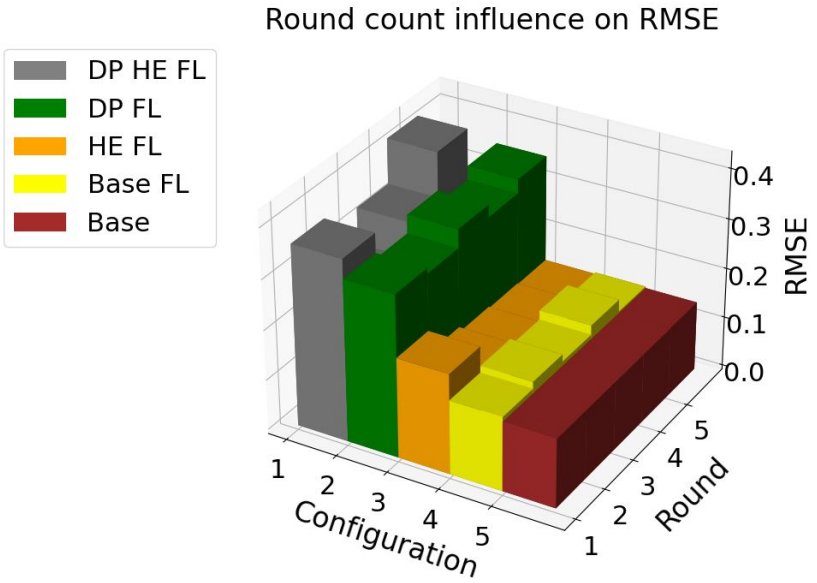

Differential privacy increased the RMSE by 3x in comparison to base configurations.

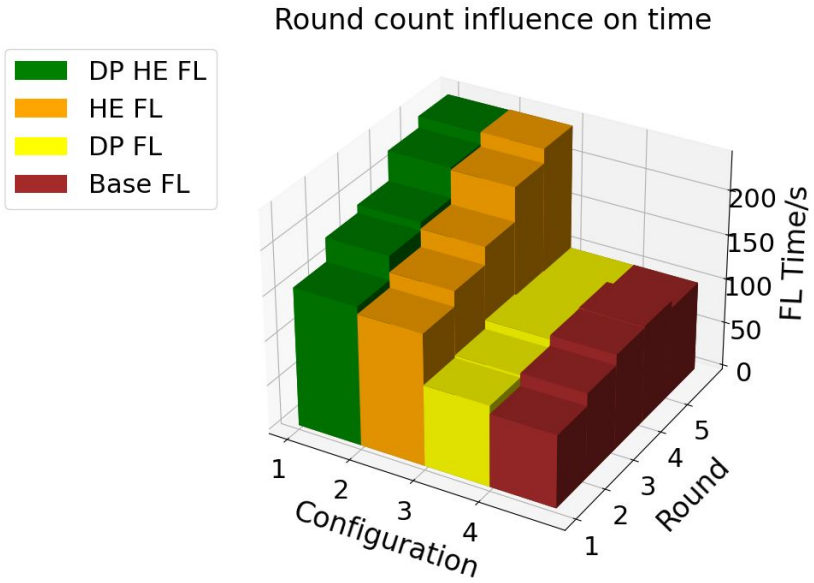

Homomorphic encryption prolonged the training process by approximately 2 times.

ELEKTRONIKAS UN DATORZINĀTNU INSTITŪTS INSTITUTE OF ELECTRONICS AND COMPUTER SCIENCE
MOTE
Powerized
Digitalization of Power Electronic Applications within Key Technology Value Chains
ChipsJU
Co-funded by the European Union

Future goal for Power electronics - Heterogeneous Federated learning

- Heterogeneity types
  - Models or parts from it;
  - Data (Independent and identically distributed (IID) and non-IID).
- DP noise explosion;
- Personalization;
- Knowledge distillation.

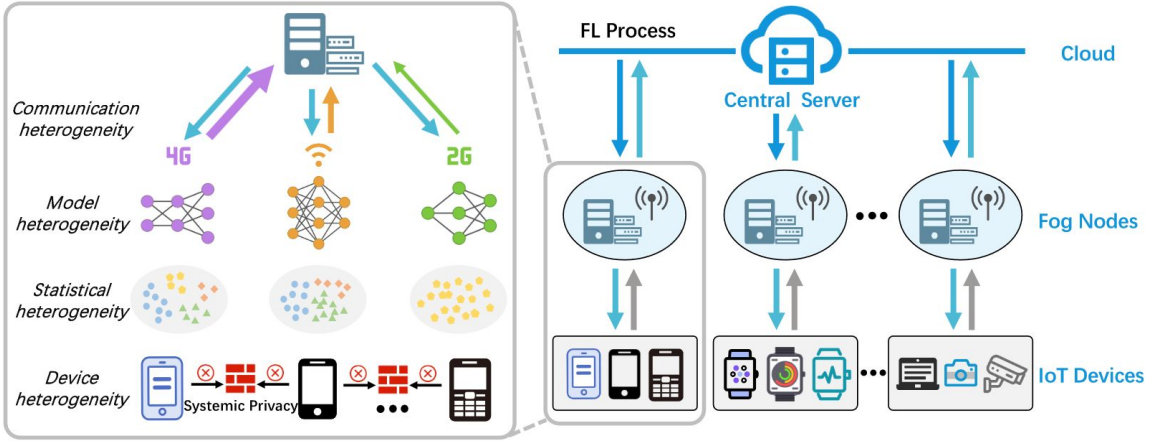

Ye, M.; Fang, X.; Du, B.; Yuen, P.C.; Tao, D. Heterogeneous federated learning: State-of-the-art and research challenges. ACM Computing Surveys 2023, 56, 1–44.

Future goal for Power electronics - Physics-informed neural networks

- Two component optimization - NN loss + PDEs;
- Lacks research in DP and FL application;
- Combines neural-networks and the laws of physics;

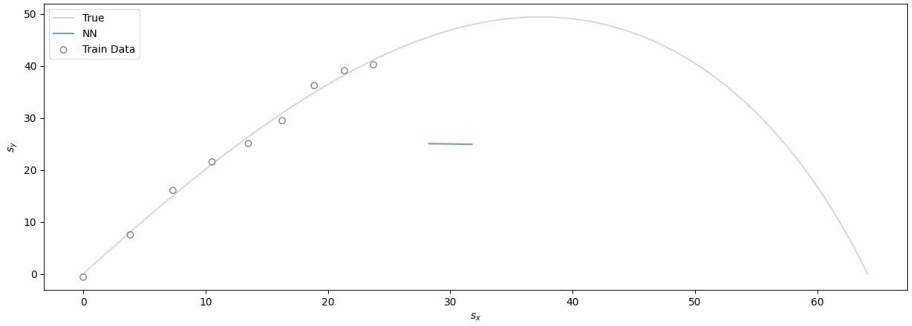

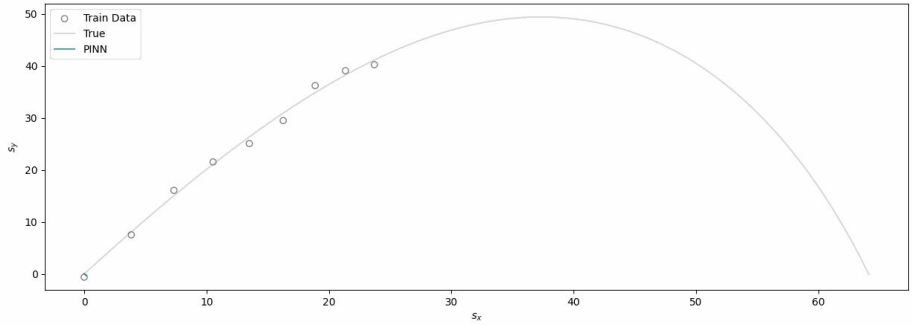

Physics-informed neural networks. Available online: https://towardsdatascience.com/physics-informed-neural-networks-pinns-an-intuitive-guide-fff138069563 (accessed on 01.02.2024).

Current developments for IoT-Edge-Cloud Continuum

100 average sensor value analysis relative to the applied differential privacy level for obfuscated data further application in ML model training.

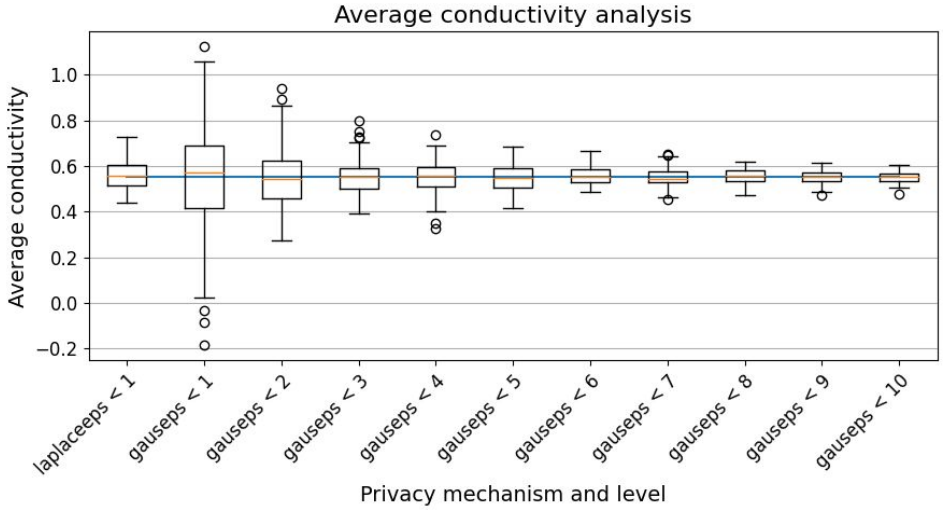

Proof of concept water quality classification with differential privacy using a simple neural network model.

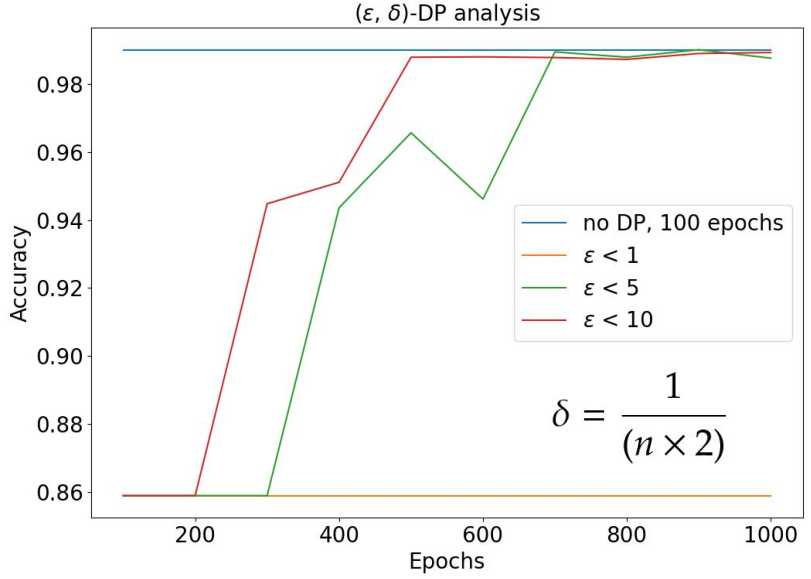

Future goal for IECC

- Use differential privacy at the input level and use the resulting obfuscated dataset for neural network training.

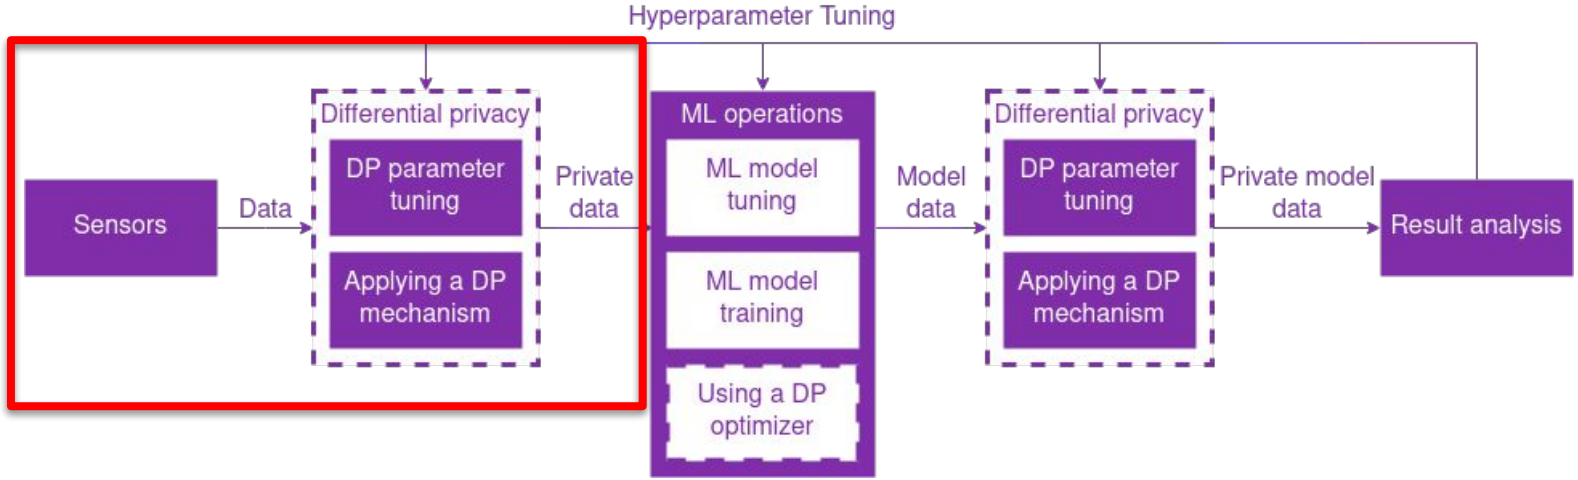

ELEKTRONIKAS UN DATORZINĀTNU INSTITŪTS INSTITUTE OF ELECTRONICS AND COMPUTER SCIENCE
MOTE
Powerized
Digitalization of Power Electronic Applications within Key Technology Value Chains
ChipsJU
Co-funded by the European Union

Publikācijas

- Arzovs A, Judvaitis J, Nesenbergs K, Selavo L. Distributed Learning in the IoT–Edge–Cloud Continuum. Machine Learning and Knowledge Extraction. 2024 Feb 1;6(1):283-315. (MDPI MAKE - Q1 žurnāls);
- Šobrīd ir plāns veidot vēl divas.

Dalība konferencēs un citos pasākumos

- Abstrakta prezentēšana konferencē IWoEDI’2023, Rīgā;
- PowerizeD General Assembly, Vienna, Austria;
- PowerizeD review meeting, Leuven, Belgium;
- Joint Estonian-Latvian Theory Days, Randivälja, Estonia;
- LU 82. starptautiskā konference, tiešsaistē.

Summary

- Federated learning requires additional privacy and security methods;
- Differential privacy (DP) implementation complexity;
- DP privacy-utility cost;
- Security impact on performance;
- Federated learning with DP and homomorphic encryption has been applied to an existing electric vehicle battery RUL prediction CNN model;
- A neural network was trained using a DP optimizer for a water quality prediction task;
- DP has been applied to initial water quality dataset for further ML model training;

Acknowledgement

- This research is funded by “Digitalization of Power Electronic Applications within Key Technology Value Chains” (PowerizeD), which has received funding from the Chips Joint Undertaking (Chips-JU) under grant agreement No 101096387. The Chips-JU receives support from the European Union’s Horizon Europe research and innovation programme and Germany, Finland, Spain, Netherlands, Sweden, Belgium, Austria, Italy, Hungary, Switzerland, Greece and Latvia. Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or Chips Joint Undertaking. Neither the European Union nor the Chips Joint Undertaking can be held responsible for them;
- This research is funded by the Latvian Council of Science, project “Smart Materials, Photonics, Technologies and Engineering Ecosystem” No VPP-EM-FOTONIKA-2022/1-0001.

Defending against poisoning attacks - denying access

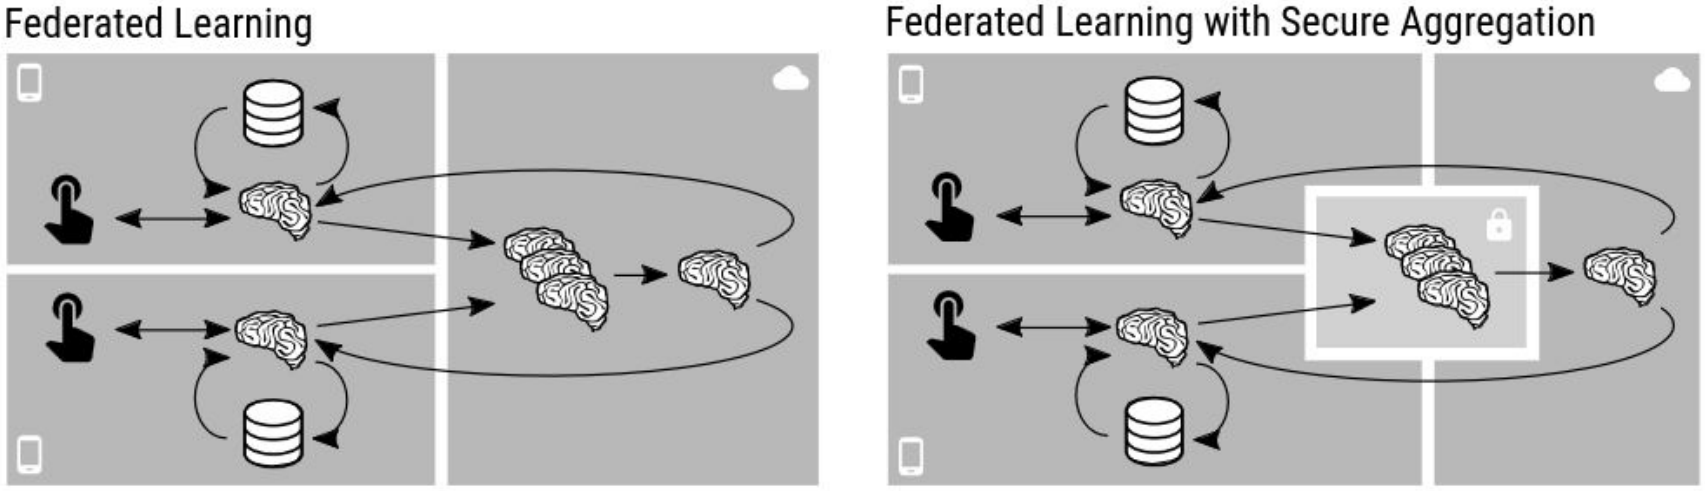

K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth, “Practical secure aggregation for privacy-preserving machine learning,” in proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, pp. 1175–1191, 2017.

Differential privacy application stages

The further from the initial data source (even at the prediction level), the better results are achieved.

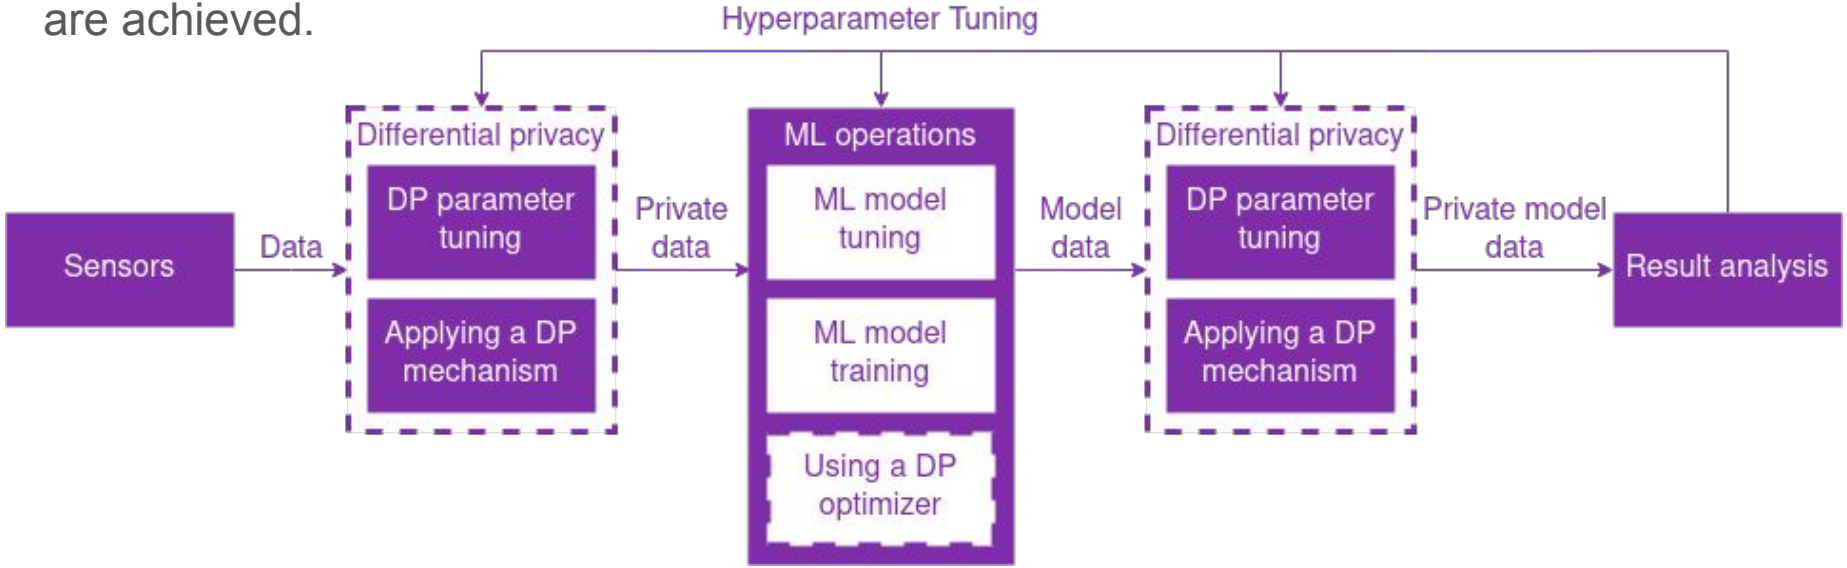

Ponomareva, N.; Hazimeh, H.; Kurakin, A.; Xu, Z.; Denison, C.; McMahan, H.B.; Vassilvitskii, S.; Chien, S.; Thakurta, A.G. How to dp-fy ml: A practical guide to machine learning with differential privacy. Journal of Artificial Intelligence Research 2023, 77, 1113–1201.

Defending against poisoning attacks with transparency

- Use the blockchain primitives in the Federated learning network;
- Transactions - Data to be stored on-chain, sent from one of the nodes in the network (sending the model update);
- Shared ledger - accounting mechanism of all verified transactions in the blockchain network (transparency about all model updates);
- Consensus mechanism - In order to verify a transaction, a network consensus has to be reached where network nodes agree upon whether or not to accept a transaction (evaluate each update);
- Peer-to-peer networking - Decentralized communication mechanism (escaping centralized bottlenecks);
- On-chain and off-chain storage - on-chain data is smaller in size (e.g. metadata). Off-chain storage (e.g. The InterPlanetary File System (IPFS)) for models.

Witt, L.; Heyer, M.; Toyoda, K.; Samek, W.; Li, D. Decentral and incentivized federated learning frameworks: A systematic literature review. IEEE Internet of Things Journal 2022.
Ali, M.; Karimipour, H.; Tariq, M. Integration of blockchain and federated learning for Internet of Things: Recent advances and future challenges. Computers & Security 2021, 108, 102355.

Future goal for Power electronics - Heterogeneous Federated learning

Methods for heterogeneous devices

Split learning

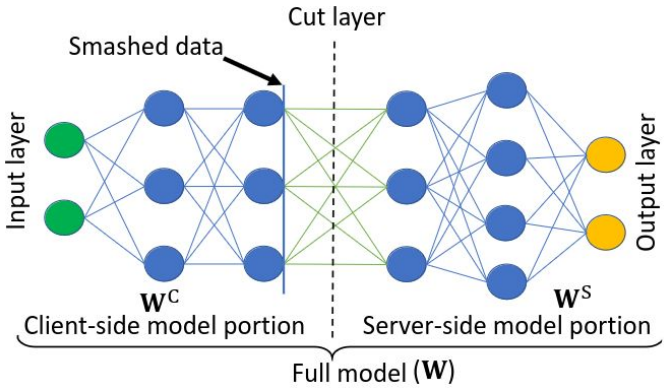

Transfer learning

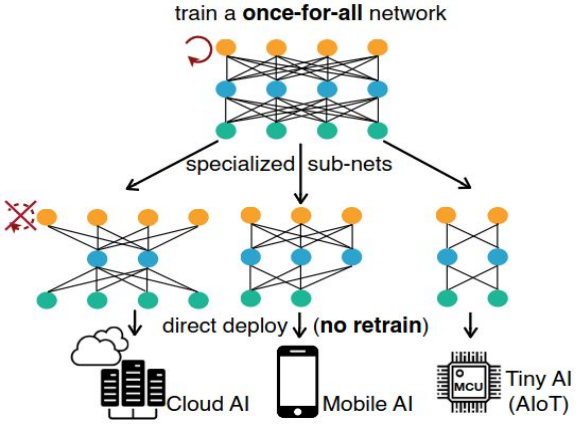

Thapa, C.; Chamikara, M.; Camtepe, S.A. Advancements of federated learning towards privacy preservation: from federated learning to split learning. arXiv preprint arXiv:2011.14818 2020.
Cai, H.; Gan, C.; Wang, T.; Zhang, Z.; Han, S. Once-for-all: Train one network and specialize it for efficient deployment. arXiv preprint arXiv:1908.09791 2019.

Outline

- The goals of each research direction - Power electronics and IoT-Edge-Cloud continuum, and how federated learning can help to reach these goals;
- Federated learning, its drawbacks - why we need more methods for privacy and security related problems;
- How and why Federated learning can benefit from using differential privacy and homomorphic encryption.
- This talk focuses on differential privacy for neural network models and not databases (See the 2019 theory days talk “Taming epsilon of differential privacy” from Alisa Pankova);
- Description of current developments regarding battery remaining useful life prediction and water quality prediction.

Main goal

- Private and secure (perhaps also robust) iterative heterogeneous machine learning model training and knowledge aggregation without sharing any input data.

ELEKTRONIKAS UN DATORZINĀTNU INSTITŪTS INSTITUTE OF ELECTRONICS AND COMPUTER SCIENCE
MOTE
Powerized
Digitalization of Power Electronic Applications within Key Technology Value Chains
ChipsJU
Co-funded by the European Union

Defending against privacy leakage

- Anonymity methods - obsolete;
  - k-anonymity;
  - l-diversity;
  - t-closeness.
- L2 regularization.
- Defense against:
  - Linkage attacks (weakness of anonymity methods);
  - Inference attacks;
  - Backdoor attacks (Weak-DP).
- Creating data privacy.
  - Obfuscation of initial data, while still providing value.

(ε, δ)-Differential privacy

Creates accuracy and privacy trade-off.

$$Pr(M(x) = t) <= e^{\varepsilon} Pr(M(x') = t) + \delta$$

x, x' - neighboring datasets (differ in at most 1 record);
M - mechanism that adds noise to the data from a probability density function e.g. Laplacian or Gaussian;
ε - privacy budget;
δ - relaxation parameter for ML use cases.

C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," Journal of Privacy and Confidentiality, vol. 7, no. 3, pp. 17–51, 2016.
N. Ponomareva, H. Hazimeh, A. Kurakin, Z. Xu, C. Denison, H. B. McMahan, S. Vassilvitskii, S. Chien, and A. Thakurta, "How to dp-fy ml: A practical guide to machine learning with differential privacy," arXiv preprint arXiv:2303.00654, 2023.