

Barriers to the Practical Adoption of Federated Machine Learning in Cross-company Collaborations

Tobias Müller^{1,2}, Nadine Gärtner¹, Nemrude Verzano¹ and Florian Matthes²

¹SAP SE, Dietmar-Hopp-Allee 16, 69190 Walldorf, Germany

²Chair for Software Engineering for Business Information Systems, Technical University of Munich, Boltzmannstrasse 3, 85748 Garching bei München, Germany

Keywords: Big Data, Anonymization, Encryption, Data Markets, Privacy-enhancing Techniques, Federated Learning.

Abstract: Research in federated machine learning and privacy-enhancing technologies has spiked recently. These technologies could enable cross-company collaboration, which yields the potential of overcoming the persistent bottleneck of insufficient training data. Despite vast research efforts and potentially large benefits, these technologies are only applied rarely in practice and for specific use cases within a single company. Among other things, this little and specific utilization can be attributed to a small amount of libraries for a rich variety of privacy-enhancing methods, cumbersome design of end-to-end privacy-enhancing pipelines and unwieldy customizability to needed requirements. Hence, we identify the need for an easy-to-use privacy-enhancing tool to support and enable cross-company machine learning, suitable for varying scenarios and easily adjustable to the desired corresponding privacy-utility desiderata. This position paper presents the starting point for our future work aiming at the development of the described application.

1 PROBLEM STATEMENT

The ever-increasing trend of data-driven decision making, automation of traditional manufacturing and industrial practices requires companies to collect data and generate useful information from the collected data. This data-driven automation can be achieved by implementing artificial intelligence (AI) systems which are trained to identify patterns based on the collected data and then act according to some predefined rules. These systems allow companies to design better products, optimize processes and enhance resilience along the digital supply chain without the need of human intervention (Ivanov et al., 2019).

However, the complexity of AI models tends to increase as problems get more complicated. The common ground of training complex AI systems is the requirement for large and diverse datasets. The requirement for large datasets mostly tends to be the bottleneck of these AI systems. Since a large amount of high-quality training data is key, this can be overcome by voluntary data sharing by data owners and multi-institutional dataset aggregation. Hence, the need for cross-company collaboration emerges. Data markets provide a possible approach for enabling cross-company collaboration (Ohrimenko et al., 2019).

Representatives from business, politics, and science have recognized the added value and potential economic benefit of said data markets and therefore recently supported major multi-national projects that provide a common digital infrastructure for data-sharing. Initiatives like International Data Spaces Association (IDSA), GAIA-X or the automotive-related Catena-X aim to enable cross-company collaboration in a shared environment of trust. This work takes place in the context of the Catena-X project.

Even though data markets promote collaboration across companies and consequently foster innovation, these markets bear the risk of revealing sensitive business information or even personal data. Besides, trust and confidence in the protection of privacy and intellectual property is a critical barrier in the users' willingness to participate in collaborative information sharing (Schomakers et al., 2020). In addition to preventing privacy breaches, these systems also have to ensure legal compliance e.g. with GDPR regulations.

It is usual practice to anonymize or pseudo-anonymize sensitive data at the originating institution, before transmitting the data to a central location, where it is stored, analyzed and used for model training (Sheller et al., 2020). However as seen in the infamous Netflix prize (Narayanan and Shmatikov, 2008),

using only anonymization has proven to be insufficient against re-identification attacks. Consequently, a multitude of privacy-enhancing techniques (PETs) recently emerged from this problem, each with a varying scope and for differing privacy concerns. For an extensive overview of PETs in IoT data markets we refer the reader to the work of (Garrido et al., 2021).

Moreover, the described government funded projects mainly deal with data exchange and less with collective data processing. A widely used collaborative data processing framework is federated learning, in which a machine learning model is jointly trained among participants. Each data owner receives a copy of the machine learning model from the model owner and trains it locally. After local training, the resulting update gradients are shared with the model owner, which then updates the centralized model through the accumulated gradients (Konečný et al., 2015; Hard et al., 2018). By this, sensitive data never leaves the users' device. However, further privacy risks like model stealing or reverse-engineering of training data arise when applying federated learning (Gonçalves et al., 2021). Hence, we need further precautions and techniques to enhance privacy of sensitive data in federated machine learning processes.

There has been large research and investment in powerful privacy-enhancing measures. Even though there is a lot of theoretical work, there seems to be little practical application. When applied, PETs are only utilized for specific scenarios and with no general, generic adaption. The factor, that PETs are unwieldy for non-experts is the reason for this observation (Renaud et al., 2014; Gerber et al., 2019). Most methods lack a reusable library or the required combination of anonymization and secure computation (Garrido et al., 2021). Therefore, multiple approaches have to be reasonably combined when data of multi-institutional sources have to be processed. Designing such an end-to-end privacy-enhancing architecture is cumbersome, especially for non-experts.

Additionally, there is always a trade-off between the degree of privacy and data utility. From a user perspective, some use cases might demand precise accuracy, while other scenarios prefer high data privacy over high utility. To make informed trade-off decisions, users need to know the specifics about the underlying architecture and implemented PETs. Again, this might be cumbersome, especially for non-experts.

Lastly, the composition of data sources, model owners and stakeholders varies with each use case. Each scenario also brings different data structures, which results in the need for interchangeable machine learning models depending on the use case.

This position paper presents the starting point for our future work aiming to lower the hurdle for users to practically deploy privacy-enhancing techniques to enhance trust and accordingly support collaborative, cross-company machine learning.

Thus we can summarize the challenges and the goals of a possible application as follows:

- PETs are currently mainly usable by experts. The application should enable non-experts to easily enforce privacy-enhancing techniques.
- There is always a trade-off between utility and privacy. The application has to enable non-experts to make informed trade-off decisions.
- The solution should be generically usable for different scenarios with varying model owners, data sources and machine learning models.

After an initial technology, we seek to evaluate and compare existing PETs and group them according to performance indicators like the achievable degree of privacy, addressed privacy concern, computational overhead and adaptability. We also plan to benchmark existing libraries in the process and derive capabilities and limitations of existing technologies, solutions and libraries. With the help of the resulting findings, we want to construct multiple privacy-enhancing architectures, which cover varying levels of privacy-utility requirements. Finally, we pursue the development of an easy-to-use privacy-enhancing application for multi-institutional federated machine learning, suitable for varying scenarios and easily adjustable to the corresponding utility-privacy need. The user should be able to integrate own machine learning models. Since our solution mainly aims to provide privacy on the data processing layer, the overall application should be compatible with trustworthy distributed data infrastructures, such as Gaia-X, which provide additional privacy on the storage and communication layer.

In section 2, we will give a short introduction to the terms security, privacy, federated machine learning and the accompanying privacy concerns. This is followed by section 3 with an overview of trustworthy distributed data infrastructures (Gaia-X and IDSA) and existing solutions. Here, we provide an initial assessment about the advantages and limitations of the most widely-used PETs. Afterwards, we present related work (section 4) before we finally conclude in section 5.

2 PRELIMINARIES

The following will introduce the terms **security**, **privacy** and the concept of **federated machine learning**. Additionally, we will further debate the necessity of privacy-enhancing techniques for cross-company data sharing and argue that contractual protection and incentivization is not enough to motivate users to participate in collaborative data processing.

Security and Privacy. Even though the terms data security and data privacy are mostly used interchangeably, it is important to differentiate both concepts as well as define their interrelation. For this, we adapt the definition as given by Jain et al. (2016):

Data **security** is specified as the practice of defending the confidentiality, integrity and availability of data. Security aims to prevent data compromise through the use of technology, processes and training.

Data **privacy**, on the other hand, is concerned with the control over the collection and usage of personal information. To preserve data privacy, an individual should have the capability to stop personal information from becoming known to unauthorized or undesired people. The focus lies on the use and governance of individual data.

“While security is fundamental for protecting data, it’s not sufficient for addressing privacy.” (Jain et al., 2016).

Federated Machine Learning. For collaborative data processing, we will rely on federated machine learning as the main framework, which works as follows (Konečný et al., 2015):

1. The model owner provides a centralized machine learning model.
2. Each participant downloads the current model and trains it on local data.
3. Each participant summarizes the changes as a small focused update, which then is sent back to the model owner.
4. The updates are aggregated, merged through an averaging scheme and then used to update the central model.
5. Repeat the process.

Since the data stays on the users’ device, this process yields a strong privacy and security benefit. The model also resides on the device, which makes real-time prediction possible, even without internet connection. However, federated learning only preserves privacy to a certain degree, communication could become a potential bottleneck and the framework may

need to anticipate low levels of participation as well as statistical and systems heterogeneity (Li et al., 2020).

Importance of Privacy. As already stated, PETs introduce more computational overhead and reduce the data utility to a certain extent. So the question remains, if adding further PETs on top of federated machine learning is really necessary to support participation in collaborative data processing systems or if legal restrictions and (financial) incentives might suffice. A series of three studies from Schomakers et al. (2020) have shown, that the relevant barriers regarding data sharing are privacy concerns with the protection of sensitive data as the most essential condition. The level of anonymity is important for all user groups, which even outweighs monetary benefits for users with high privacy concerns (Schomakers et al., 2020). We assume that competing companies belong to the group of users with higher privacy concerns which therefore prioritize data security. Further studies confirm that secure data sharing is the most important factor in fostering sharing-centered collaboration (Panahifar et al., 2018; Woldaregay et al., 2020).

Studying the usage pattern of PETs, Coopamootoo (2020) found that non-technology methods and technologies which are integrated into services, therefore easily usable, are the most utilized PETs. Methods which are cumbersome to apply are less likely to be used (Coopamootoo, 2020). If and how this observation fits with industrial practice has to be determined. But we argue, that these observation also applies to industrial machine learning practices which results in a need for generically and easily usable privacy-enhancing framework for collaborative data processing.

3 EXISTING APPROACHES

In this section, we will present an initial technology overview for privacy-enhancing data processing, including a first discussion about their respective benefits, limitations and assessment of utility for our project. We group the methods into technologies based on anonymization and cryptography.

3.1 Anonymization Approaches

In contrast to cryptography-based techniques, anonymization does not try to hide data from unauthorized parties, but to help prevent identification or re-identification by protecting direct identifiers, quasi-identifiers and sensitive attributes (Majeed and

Lee, 2020). Most anonymization techniques fundamentally rely on non-perturbative and perturbative masking based on statistics, probability theory and heuristics. Several approaches have been proposed in the past (Salas and Domingo-Ferrer, 2018):

Differential Privacy. Strictly speaking, **Differential Privacy (DP)** (Dwork, 2006) is not a method, but a mathematically provable promise of privacy in a dataset. More specifically, DP guarantees that an algorithm \mathcal{M} is (ϵ, δ) -differentially private if for any neighboring datasets \mathcal{D} and \mathcal{D}' differing on at most one element (neighboring datasets) and any set of possible outputs $S \in \text{Range}(\mathcal{M})$:

$$\Pr[\mathcal{M}(\mathcal{D}) \in S] \leq e^\epsilon \times \Pr[\mathcal{M}(\mathcal{D}') \in S] + \delta \quad (1)$$

with ϵ as the likeliness of finding individuals and δ as the possibility, by which outputs differ for different datasets. Hence through differential privacy, an adversary will essentially get the same inference about any individual's private information, which makes the outputs "differentially" indistinguishable (Wood et al., 2018). This promise is mostly achieved by adding noise randomly sampled from a probability density function.

Differential Privacy is widely applicable (Hassan et al., 2020) and has also been presented to obfuscate the resulting update gradients of training deep neural networks. This extension to traditional neural networks is one possibility to counteract reverse-engineering of training data (Abadi et al., 2016). Adding noise to data or to update gradients lowers the degree of information and therefore the utility. However, the amount of noise can be easily customized, which shifts the privacy-utility trade-off. Due to the easy adaptability and broad applicability, we expect DP to be a valuable technique for our project.

K-Anonymity. As the name suggests, a dataset provides k -anonymity (Sweeney, 2002) if the identifying information of each individual is indistinguishable from at least $k-1$ other individuals. This can be realized by clustering a set of sensitive attribute values into equivalence classes of size k . A higher value for k represents higher anonymity and consequently a lower probability of correctly linking sensitive attributes to an associated individual. However, finding an optimal value for k with minimum information loss is a NP-hard problem (Meyerson and Williams, 2004; Liang and Samavi, 2020) and k -anonymity is not secure against homogeneity attacks or background knowledge attacks (Maheshwarkar et al., 2011).

Thus, multiple extensions have with further privacy requirements have been proposed. For instance,

l -diversity (Machanavajjhala et al., 2007) proposed to have least l -different values in sensitive attributes or t -closeness (Li et al., 2007) with the additional property that the distance t between the distribution of sensitive values within each class is less than or equal to the overall dataset distribution of the attribute. Even though, we also identified less prominent approaches like β -likeness (Cao and Karras, 2012), δ -presence (Nergiz et al., 2007) or δ -disclosure privacy (Brickell and Shmatikov, 2008), we initially focus on k -anonymity, l -diversity and t -closeness.

The clustering of sensitive attribute values into k equivalence classes can be done by replacing a value with less specific but semantically consistent value (**generalization**), by removing selected data points (**suppression**) or by changing data to something else, which can be reverted back with the help of the original data (**distortion**).

These methods are mainly used to obfuscate the input and output and therefore enhance privacy of the input data and computational output. Of course, this small list of privacy-preserving anonymization techniques is far from complete. We refer to the study of Majeed and Lee (2020) for a more detailed overview.

3.2 Cryptographic Approaches

In contrast to anonymization, cryptography-based approaches hide sensitive data from unauthorised access by converting sensitive information (plaintext) into unintelligible form (ciphertext). Naturally, a multitude of cryptography-based methods have been proposed not only to transfer data, but also for secure data processing:

Homomorphic Encryption. Homomorphic Encryption (HE) is an encryption scheme, which allows the computation of encrypted data, where the output can be decrypted using the corresponding secret key (Ogburn et al., 2013). Homomorphic encryption schemes can be classified into fully homomorphic encryption (FHE), where both - addition and multiplication - is supported and partially homomorphic encryption (PHE), where only one operation is possible. Any schema in-between is classified as somewhat homomorphic encryption (Kogos et al., 2017). Most recently, a FHE extension for deep learning has been proposed, which achieved nearly identical results compared to training on non-encrypted data. Evaluated on the CIFAR-10 dataset (Krizhevsky, 2012), the model performed extremely well with a high security level and classification accuracy. The proposed model is still limited by computational overhead and the runtime, which is about 4

hours to infer a single image (Lee et al., 2021). Consequently, this technology is still impractical to use, but might be promising after further research.

Oblivious RAM. An oblivious Random Access Memory (ORAM) (Goldreich, 1987) aims to prevent leakage of sensitive information, which can be inferred by the behavior of the user rather than data content itself. More specifically, ORAMs conceal the access pattern to a remote storage by continuously shuffling and re-encrypting data as it is accessed, while preserving the input-output behavior of the original data. The access to the physical storage location can be observed, but the ORAM algorithm obfuscates the access pattern such that one can not infer the true (logical) access pattern from the observation. By this, an adversary can not obtain nontrivial information about the execution of a program or the nature of the data, which usually can be leaked even though data values are all encrypted. As an extension to the original concept, PathORAM (Stefanov et al., 2018) with enhanced significantly less bandwidth cost.

Although ORAMs do not specifically enhance privacy in the data processing layer, we still need to consider the possibility, that the communication and data access patterns might indirectly leak sensitive information.

Secure Multi-party Computation (SMPC). is a protocol, which enables multiple parties to perform and evaluate computations without revealing any of the private data held by each party, without the necessity of a trusted third party. Several peers perform a shared computation without sharing it amongst themselves. Given M participants and N computing parties, each SMPC protocol follows three steps (Torkzadehmahani et al., 2020):

1. Each participant shares separate and different secrets to each of the N computing parties according to a selected secret sharing method.
2. Each computing party computes intermediate results and shares the results with the other computing parties.
3. Each computing party aggregates the results from all computing parties and computes a global result, which is then sent back to all participants.

Whereby additive secret sharing (Vaidya and Clifton, 2003) and Shamir secret sharing (Shamir, 1979) are the most prominent secret sharing schemes.

In contrast to traditional cryptography, SMPC only conceals partial information about the data while performing computation with the data and usually

entail multiple rounds of interactive communication, which leads to a sensitivity to network latency. Since machine learning frameworks can be integrated (Knott et al., 2021) and because these SMPC protocols target the problem of keeping the input private and of ensuring correctness, SMPC might be a valuable technique for our project.

Zero Knowledge Proofs. Generally, Zero Knowledge Proofs (ZKP) protocols allows one party (verifier) to verify the authenticity of a given computation conducted by another party (prover), without having any knowledge about the prover, the computation or the underlying data (Feige and Shamir, 1990). ZKPs can be distinguished into interactive ZKPs where a sequential message exchange is required and non-interactive ZKPs with no sequential message exchange.

These concepts can also be used to prove the results, correctness, consistency and achieved accuracy of each participants' locally trained model without leaking any information about the data (Zhang et al., 2020; Liu et al., 2021). ZKPs appear to be an efficient solution to address authentication problems and to confirm given predictions without compromising data utility and accuracy of the machine learning model itself.

Naturally, this list is also far from complete. We refer to the study of (Kaaniche et al., 2020) for a more comprehensive overview.

3.3 Trusted Execution Environments

Moreover, we want to present trusted execution environments (TEE), which is a combination of hardware and software components and allow users to define secure areas of memory (enclaves) that enhance confidentiality, as well as the data and computation integrity. These processing environments are isolated and thus impede other programs outside the enclave to act on the data (Sabt et al., 2015). These environments should be secure against sophisticated attacks like probing external memories, measuring execution time and against attacks aiming to retrieve cryptographic key material (Shepherd et al., 2016). Physical access to the hardware is the only way of compromising TEEs. One has to bypass remote attestation and the sealed storage by manipulating the system to provide false certifications (Colman et al., 2019).

However, there is no standardization and most open-source proposals lack maturity. Hardware and a trusted third party might be needed, which introduces a single point of failure (Shepherd et al., 2016; Busch et al., 2020) and reduces the usefulness for our in-

tentions. Proprietary solutions like the Intel Software Guard Extension (SGX) (Anati et al., 2013), where Intel is the trusted third party and therefore single point of failure have been mostly used due to the lack of maturity of open-source protocols (Garrido et al., 2021). As a fallback, TEEs could help enable secure centralized machine learning, in case it turns out that a federated learning approach is not feasible.

4 RELATED WORK

From recommender systems for videos (Duan et al., 2020), payments in smart finance (Liu et al., 2020) to predicting the energy demand of electric vehicle networks (Saputra et al., 2019), federated learning has found its way into a broad range of industrial domains (Zhou et al., 2021). However, little work has been focused on federated learning across companies, which introduces further privacy concerns and therefore privacy requirements compared to industrial applications within a single company.

Most work on multi-institutional machine learning has been made in medical research (Sheller et al., 2020; Sarma et al., 2021; Guo et al., 2021). Notably, Kaissis et al. (2021) recently presented PriMIA, an end-to-end privacy-preserving deep learning framework for medical imaging on multi-institutional data. Combining federated learning, secure multi-party computation and differential privacy, the authors have been able to prevent model inversion attacks with comparable accuracy and reasonably longer (1.5-2.91 times longer) inference time (Kaissis et al., 2021).

In the context of industrial IoT, some authors proposed blockchain-based federated learning systems for data sharing in industrial IoT (Lu et al., 2020; Mohr et al., 2021). Although all these inspiring approaches are promising, they focus on specific use cases with certain privacy concerns. Hence, we identify a lack of a generally usable privacy-preserving framework for collaborative processing of industrial data.

5 CONCLUSION

We observe a reluctance of participating in collaborative, cross-company machine learning due to high privacy concerns. Since wary users value data privacy and data security over monetary benefits and contractual agreements, we argue that the barriers to multi-institutional machine learning persist due to un-intuitive usability of privacy-enhancing technologies,

small amount of available libraries and the momentarily inevitable need for expert knowledge to adjust these technologies to the desired needs given by each use case.

Hence, this position paper represents the starting point for our future work aiming towards a framework, which enables non-experts to easily apply desired privacy-enhancing technologies for collaborative data processing in trustworthy distributed data infrastructures.

As a first step towards this application, we presented an initial technology overview with a first discussion about the practicality and usability of the corresponding methods. Naturally, this list is far from complete and the approaches have yet to be thoroughly evaluated. The results of this evaluation will provide the necessary basis for the application architecture and integration concept into distributed data infrastructures.

ACKNOWLEDGEMENTS

The authors would like SAP SE and the Catena-X Automotive Network for supporting this work.

REFERENCES

- 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, CCS '16*, page 308–318, New York, NY, USA. Association for Computing Machinery.
- Anati, I., Gueron, S., Johnson, S. P., and Scarlata, V. R. (2013). Innovative technology for cpu based attestation and sealing.
- Brickell, J. and Shmatikov, V. (2008). The cost of privacy: Destruction of data-mining utility in anonymized data publishing. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '08*, page 70–78, New York, NY, USA. Association for Computing Machinery.
- Busch, M., Westphal, J., and Mueller, T. (2020). Unearthing the trustedcore: A critical review on huawei's trusted execution environment. In *14th USENIX Workshop on Offensive Technologies (WOOT 20)*. USENIX Association.
- Cao, J. and Karras, P. (2012). Publishing microdata with a robust privacy guarantee. *Proc. VLDB Endow.*, 5(11):1388–1399.
- Colman, A., Chowdhury, M. J. M., and Baruwal Chhetri, M. (2019). Towards a trusted marketplace for wearable data. In *2019 IEEE 5th International Conference on*

- Collaboration and Internet Computing (CIC)*, pages 314–321.
- Coopamootoo, K. P. (2020). Usage patterns of privacy-enhancing technologies. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security, CCS '20*, page 1371–1390, New York, NY, USA. Association for Computing Machinery.
- Duan, S., Zhang, D., Wang, Y., Li, L., and Zhang, Y. (2020). Jointrec: A deep-learning-based joint cloud video recommendation framework for mobile iot. *IEEE Internet of Things Journal*, 7(3):1655–1666.
- Dwork, C. (2006). Differential privacy. In Bugliesi, M., Preneel, B., Sassone, V., and Wegener, I., editors, *Automata, Languages and Programming*, pages 1–12, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Feige, U. and Shamir, A. (1990). Zero knowledge proofs of knowledge in two rounds. In Brassard, G., editor, *Advances in Cryptology — CRYPTO' 89 Proceedings*, pages 526–544, New York, NY. Springer New York.
- Garrido, G. M., Sedlmeir, J., Uludag, Ö., Alaoui, I. S., Luckow, A., and Matthes, F. (2021). Revealing the landscape of privacy-enhancing technologies in the context of data markets for the iot: A systematic literature review. *CoRR*, abs/2107.11905.
- Gerber, N., Zimmermann, V., and Volkamer, M. (2019). Why johnny fails to protect his privacy. In *2019 IEEE European Symposium on Security and Privacy Workshops (EuroS PW)*, pages 109–118.
- Goldreich, O. (1987). Towards a theory of software protection and simulation by oblivious rams. In *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing, STOC '87*, page 182–194, New York, NY, USA. Association for Computing Machinery.
- Gonçalves, C., Bessa, R. J., and Pinson, P. (2021). A critical overview of privacy-preserving approaches for collaborative forecasting. *International Journal of Forecasting*, 37(1):322–342.
- Guo, P., Wang, P., Zhou, J., Jiang, S., and Patel, V. M. (2021). Multi-institutional collaborations for improving deep learning-based magnetic resonance image reconstruction using federated learning.
- Hard, A., Kiddon, C. M., Ramage, D., Beaufays, F., Eichner, H., Rao, K., Mathews, R., and Augenstein, S. (2018). Federated learning for mobile keyboard prediction.
- Hassan, M. U., Rehmani, M. H., and Chen, J. (2020). Differential privacy techniques for cyber physical systems: A survey. *IEEE Communications Surveys Tutorials*, 22(1):746–789.
- Ivanov, D., Dolgui, A., Das, A., and Sokolov, B. (2019). *Digital Supply Chain Twins: Managing the Ripple Effect, Resilience, and Disruption Risks by Data-Driven Optimization, Simulation, and Visibility*, pages 309–332. Springer International Publishing, Cham.
- Jain, P., Gyanchandani, M., and Khare, N. (2016). Big data privacy: a technological perspective and review. *Journal of Big Data*, 3.
- Kaaniche, N., Laurent, M., and Belguith, S. (2020). Privacy enhancing technologies for solving the privacy-personalization paradox: Taxonomy and survey. *Journal of Network and Computer Applications*, 171:102807.
- Kaissis, G., Ziller, A., Passerat-Palmbach, J., Ryffel, T., Usynin, D., Trask, A., Lima, I., Mancuso, J., Jungmann, F., Steinborn, M.-M., Saleh, A., Makowski, M., Rueckert, D., and Braren, R. (2021). End-to-end privacy preserving deep learning on multi-institutional medical imaging. *Nature Machine Intelligence*, 3:1–12.
- Knott, B., Venkataraman, S., Hannun, A. Y., Sengupta, S., Ibrahim, M., and van der Maaten, L. (2021). Crypten: Secure multi-party computation meets machine learning. *ArXiv*, abs/2109.00984.
- Kogos, K. G., Filippova, K. S., and Epishkina, A. V. (2017). Fully homomorphic encryption schemes: The state of the art. In *2017 IEEE Conference of Russian Young Researchers in Electrical and Electronic Engineering (EICoN Rus)*, pages 463–466.
- Konečný, J., McMahan, B., and Ramage, D. (2015). Federated optimization: distributed optimization beyond the datacenter.
- Krizhevsky, A. (2012). Learning multiple layers of features from tiny images. *University of Toronto*.
- Lee, J.-W., Kang, H., Lee, Y., Choi, W., Eom, J., Deryabin, M., Lee, E., Lee, J., Yoo, D., Kim, Y.-S., and No, J.-S. (2021). Privacy-preserving machine learning with fully homomorphic encryption for deep neural network. *Future Internet*.
- Li, N., Li, T., and Venkatasubramanian, S. (2007). t-closeness: Privacy beyond k-anonymity and l-diversity. In *2007 IEEE 23rd International Conference on Data Engineering*, pages 106–115.
- Li, T., Sahu, A. K., Talwalkar, A., and Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60.
- Liang, Y. and Samavi, R. (2020). Optimization-based k-anonymity algorithms. *Computers & Security*, 93:101753.
- Liu, T., Xie, X., and Zhang, Y. (2021). zkcnn: Zero knowledge proofs for convolutional neural network predictions and accuracy. *Cryptology ePrint Archive*, Report 2021/673.
- Liu, Y., Ai, Z., Sun, S., Zhang, S., Liu, Z., and Yu, H. (2020). *FedCoin: A Peer-to-Peer Payment System for Federated Learning*, pages 125–138.
- Lu, Y., Huang, X., Dai, Y., Maharjan, S., and Zhang, Y. (2020). Blockchain and federated learning for privacy-preserved data sharing in industrial iot. *IEEE Transactions on Industrial Informatics*, 16(6):4177–4186.
- Machanavajjhala, A., Kifer, D., Gehrke, J., and Venkatasubramanian, M. (2007). L-diversity: Privacy beyond k-anonymity. *ACM Trans. Knowl. Discov. Data*, 1(1):3–es.
- Maheshwarkar, N., Pathak, K., and Chourey, V. (2011). Privacy issues for k-anonymity model.
- Majeed, A. and Lee, S. (2020). Anonymization techniques

- for privacy preserving data publishing: A comprehensive survey. *IEEE Access*, PP:1–1.
- Meyerson, A. and Williams, R. (2004). On the complexity of optimal k-anonymity. *Proceedings of the ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, 23.
- Mohr, M., Becker, C., Möller, R., and Richter, M. (2021). Towards collaborative predictive maintenance leveraging private cross-company data. In Reussner, R. H., Koziol, A., and Heinrich, R., editors, *INFORMATIK 2020*, pages 427–432. Gesellschaft für Informatik, Bonn.
- Narayanan, A. and Shmatikov, V. (2008). Robust de-anonymization of large sparse datasets. In *2008 IEEE Symposium on Security and Privacy (sp 2008)*, pages 111–125.
- Nergiz, M. E., Atzori, M., and Clifton, C. (2007). Hiding the presence of individuals from shared databases. In *Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data, SIGMOD '07*, page 665–676, New York, NY, USA. Association for Computing Machinery.
- Ogburn, M., Turner, C., and Dahal, P. (2013). Homomorphic encryption. *Procedia Computer Science*, 20:502–509. Complex Adaptive Systems.
- Ohrimenko, O., Tople, S., and Tschitschek, S. (2019). Collaborative machine learning markets with data-replication-robust payments.
- Panahifar, F., Byrne, P., Salam, M., and Heavey, C. (2018). Supply chain collaboration and firm's performance: The critical role of information sharing and trust. *Journal of Enterprise Information Management*, 31:00–00.
- Renaud, K., Volkamer, M., and Renkema-Padmos, A. (2014). Why doesn't jane protect her privacy? In De Cristofaro, E. and Murdoch, S. J., editors, *Privacy Enhancing Technologies*, pages 244–262, Cham. Springer International Publishing.
- Sabt, M., Achemlal, M., and Bouabdallah, A. (2015). Trusted execution environment: What it is, and what it is not. In *2015 IEEE Trustcom/BigDataSE/ISPA*, volume 1, pages 57–64.
- Salas, J. and Domingo-Ferrer, J. (2018). Some basics on privacy techniques, anonymization and their big data challenges. *Mathematics in Computer Science*, 12.
- Saputra, Y. M., Hoang, D. T., Nguyen, D. N., Dutkiewicz, E., Mueck, M. D., and Srikanthswara, S. (2019). Energy demand prediction with federated learning for electric vehicle networks.
- Sarma, K. V., Harmon, S., Sanford, T., Roth, H. R., Xu, Z., Tetreault, J., Xu, D., Flores, M. G., Raman, A. G., Kulkarni, R., Wood, B. J., Choyke, P. L., Priester, A. M., Marks, L. S., Raman, S. S., Enzmann, D., Turkbey, B., Speier, W., and Arnold, C. W. (2021). Federated learning improves site performance in multicenter deep learning without data sharing. *Journal of the American Medical Informatics Association*, 28(6):1259–1264.
- Schomakers, E.-M., Lidynia, C., and Zieffle, M. (2020). All of me? users' preferences for privacy-preserving data markets and the importance of anonymity. *Electronic Markets*, 30.
- Shamir, A. (1979). How to share a secret. *Commun. ACM*, 22(11):612–613.
- Sheller, M., Edwards, B., Reina, G., Martin, J., Pati, S., Kotrotsou, A., Milchenko, M., Xu, W., Marcus, D., Colen, R., and Bakas, S. (2020). Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Scientific Reports*, 10.
- Shepherd, C., Arfaoui, G., Gurulian, I., Lee, R. P., Markantonakis, K., Akram, R. N., Sauveron, D., and Conchon, E. (2016). Secure and trusted execution: Past, present, and future - a critical review in the context of the internet of things and cyber-physical systems. In *2016 IEEE Trustcom/BigDataSE/ISPA*, pages 168–177.
- Stefanov, E., Dijk, M. V., Shi, E., Chan, T.-H. H., Fletcher, C., Ren, L., Yu, X., and Devadas, S. (2018). Path oram: An extremely simple oblivious ram protocol. *J. ACM*, 65(4).
- Sweeney, L. (2002). K-anonymity: A model for protecting privacy. *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.*, 10(5):557–570.
- Torkzadehmahani, R., Nasirigerdeh, R., Blumenthal, D. B., Kacprowski, T., List, M., Matschinske, J., Späth, J., Wenke, N. K., Bihari, B., Frisch, T., Hartebrodt, A., Hausschild, A.-C., Heider, D., Holzinger, A., Hötendorfer, W., Kastelitz, M., Mayer, R., Nogales, C., Pustozero, A., Röttger, R., Schmidt, H. H. H. W., Schwalber, A., Tschohl, C., Wohner, A., and Baumbach, J. (2020). Privacy-preserving artificial intelligence techniques in biomedicine.
- Vaidya, J. and Clifton, C. (2003). Privacy-preserving k-means clustering over vertically partitioned data. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '03*, page 206–215, New York, NY, USA. Association for Computing Machinery.
- Woldaregay, A., Henriksen, A., Issom, D.-Z., Pfuhl, G., Sato, K., Richard, A., Lovis, C., Arsan, E., Rochat, J., and Hartvigsen, G. (2020). User expectations and willingness to share self-collected health data. *Studies in health technology and informatics*, 270.
- Wood, A., Altman, M., Bembenek, A., Bun, M., Gaboardi, M., Honaker, J., Nissim, K., O'Brien, D. R., Steinke, T., and Vadhan, S. (2018). Differential privacy: A primer for a non-technical audience. *Vanderbilt Journal of Entertainment & Technology Law*, 21(1):209–275.
- Zhang, J., Fang, Z., Zhang, Y., and Song, D. (2020). Zero knowledge proofs for decision tree predictions and accuracy. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security, CCS '20*, page 2039–2053, New York, NY, USA. Association for Computing Machinery.
- Zhou, J., Zhang, S., Lu, Q., Dai, W., Chen, M., Liu, X., Pirtikangas, S., Shi, Y., Zhang, W., and Herrera-Viedma, E. (2021). A survey on federated learning and its applications for accelerating industrial internet of things.