
Federated Learning and AI Regulation in the European Union: Who is Responsible? – An Interdisciplinary Analysis

Herbert Woisetschläger^{*1} Simon Mertel^{*2} Christoph Krönke³ Ruben Mayer² Hans-Arno Jacobsen⁴

Abstract

The European Union Artificial Intelligence Act mandates clear stakeholder responsibilities in developing and deploying machine learning applications to avoid substantial fines, prioritizing private and secure data processing with data remaining at its origin. Federated Learning (FL) enables the training of generative AI Models across data siloes, sharing only model parameters while improving data security. Since FL is a cooperative learning paradigm, clients and servers naturally share legal responsibility in the FL pipeline. Our work contributes to clarifying the roles of both parties, explains strategies for shifting responsibilities to the server operator, and points out open technical challenges that we must solve to improve FL’s practical applicability under the EU AI Act.

1. Introduction

With the introduction of the European Union Artificial Intelligence Act (AI Act) (Council of the European Union, 2021) and other international regulations being on the horizon, e.g., in the United States (The White House, 2023) and Canada (House Of Commons of Canada, 2022), everyone concerned with the development and deployment of AI has to adapt to new game rules. This entails data governance, robustness against adversarial scenarios, and energy consid-

erations (Woisetschläger et al., 2024a). The AI Act puts the *service provider* into the spotlight, who has to assume responsibility for model development and deployment within the meaning of Article 3. Especially regarding data governance, the AI Act instantiates extensive rules for high-risk and general-purpose AI applications (GPAI, Article 52) that cater to data privacy and system security. The majority of generative AI applications fall under the GPAI definition in Article 3.

Federated Learning (FL) presents a privacy-enhancing and data-protecting machine learning technique (McMahan et al., 2017) that has recently received increased attention for enabling access to data silos for generative AI applications (Woisetschläger et al., 2024b). In FL, a server operator provides an ML model sent to several clients and then trained on the clients’ local data, which collaboratively train a global model via a central server, aggregating their local model updates. Private and secure computing techniques like Differential Privacy or Trusted Execution Environments help improve data privacy and system security (Bonawitz et al., 2017; Andrew et al., 2021). FL’s data locality removes the key challenge of monitoring data lineage and simplifies accounting for user consent. Specifically, we study the FL workflow in alignment with related work (Li et al., 2020; Hard et al., 2018; McMahan et al., 2017) to touch up on the following:

Data Acquisition. The server operator can only employ a variety of client sampling strategies (Malinovsky et al., 2023; Wang & Ji, 2022; McMahan et al., 2017) for an FL training round, without the ability to directly investigate client data or process integrity.

Data Storage. Similarly, the clients decide how, where, and when to store data. This has implications on data availability, which directly touches upon the AI Act data governance requirements (Article 10)¹.

Data Preprocessing. While the server operator can provide instructions on how to preprocess data so that the data is compatible with the ML model, the clients have the freedom to run additional preprocessing steps. Since the server oper-

^{*}Equal contribution ¹School of Computation, Information and Technology, Technical University of Munich, Germany ²Department of Computer Science, University of Bayreuth, Germany ³Department of Law & Economics, University of Bayreuth, Germany ⁴Department of Electrical and Computer Engineering, University of Toronto, Canada. Correspondence to: Herbert Woisetschläger <herbert.woisetschlaeger@tum.de>, Simon Mertel <simon.mertel@uni-bayreuth.de>, Christoph Krönke <christoph.kroenke@uni-bayreuth.de>, Ruben Mayer <ruben.mayer@uni-bayreuth.de>, Hans-Arno Jacobsen <jacobsen@eecg.toronto.edu>.

Presented at the GenLaw ’24 Workshop at the 41st International Conference on Machine Learning, Vienna, Austria. PMLR 235, 2024. Copyright 2024 by the author(s).

¹In the following, the term Article refers to articles in the AI Act if not specified otherwise

ator has no direct data access, verifying data integrity before training is challenging. For FL applications, there are numerous approaches to improve data integrity (Sánchez Sánchez et al., 2024; Roy Chowdhury et al., 2022).

Model Aggregation. While acting as the FL training orchestrator, the server operator handles the model integrity control mechanism when aggregating model updates. Thus, FL appears to be a well-suited solution to open up data silos and provide access to additional data. This would significantly benefit the training or fine-tuning of generative models due to their sheer appetite for ever-increasing amounts of data (Zhou et al., 2023).

One can think that the server operator is automatically also the *service provider*. Yet, FL is a cooperative ML training technique where a central entity typically provides the ML model, and clients can decide when to participate and what data to use for training. As such, we see that the server (model) and the clients (data) control parts of the FL lifecycle, rendering them both legally responsible for their respective parts. Thus, this opens up the question:

Who is the service provider at what point in the FL workflow, and how can each party assume adequate responsibility?

Our paper studies technical and legal requirements that need to be established so that the FL server operator can assume responsibility as a service provider. This requires future technical work on auditability, verifiability, integrity, and privacy. Further, we need to establish regulatory references for the terms and services of FL applications.

2. Technical Solutions Need to Focus on Transferring Responsibility to the Server Operator

For practical FL applications, the server operator must assume the role of the service provider by employing appropriate technical solutions.

When establishing an FL system that could potentially entail thousands of clients at a time, managing responsibilities is likely to become a key challenge. Thus, we require solutions that provide for auditability, verifiability, integrity, and privacy.

Auditability & verifiability. There is a natural trade-off between privacy and data audits. The core paradigm of FL is to not share data beyond a client’s area of control. Data is strictly inaccessible for everybody but the data owner, and even in-house restricted to authorized personnel. Thus, we face a challenge when aiming to audit all steps that happen on a client device or data server. For instance, a work by (Liu et al., 2023) uses a Bayesian Nash equilibrium and a market mechanism to incentivize truthful client behavior,

i.e., submission of useful model updates. While this approach significantly reduces the risk of adversarial attacks, it does not meet the requirements for auditing in the context of the AI Act, which are well-defined. Quintessentially, any data that is being captured, processed, and used in a training process must be evaluated for potential bias or adversarial information. To achieve this, numerous works combining FL with blockchain technology explore auditing the data processing steps and the training itself (Nguyen et al., 2021; Ma et al., 2020). What remains open is to develop solutions against data tampering.

Integrity & privacy. Particularly, we have to rethink the obligations of the provider concerning data integrity and protection (Articles 8–10), such that responsibility is transferred to the FL server. To account for the asymmetry of access and control-by-design in FL systems, we must develop data integrity measures that capture the nature of client data at the time of collection, while preprocessing the data, and immediately before starting the training process. Peer-based verification schemes of model updates are a promising direction to identify adversarial clients (Roy Chowdhury et al., 2022). Extending such schemes from client models to client data without infringing privacy would be interesting. At the same time, technical solutions must be in line with the requirements set out in the GDPR, which are not (necessarily) aligned with the concepts and rules of the AI Act.

3. Regulatory Implementations Need to Foster Integrity and Verifiability

We need FL server operators to assume full responsibility; clients are technically and legally obligated to comply.

Service Provider. The GDPR (Council of the European Union, 2016) defines the term *data controller*. Complementary, the AI Act defines the *service provider* of AI systems. For data protection assessment when processing personal information at first, we need to clarify who is the data controller responsible and accountable for each distinct phase of the data processing and must demonstrate compliance with the requirements of the GDPR (Article 5). The AI Act does not have a differentiated allocation of roles for separate processing phases and focuses on one central “provider” of a (compliant) AI system, defined in Article 3, with the obligations arising from Article 8.

While both the FL server and clients could be considered providers under the AI Act, since the AI Act (unlike the GDPR) focuses less on responsibilities for individual, definable data processing phases and more on secure system design as a whole, the provider concept has to be teleologically limited to the FL server. Thus, the server acts as the fully responsible service provider under the AI Act (Article 8), especially concerning data governance (Article 10) and General-Purpose AI service (Article 52).

General Terms and Conditions for AI Systems. While in Article 4 of the GDPR, the controller is the person who, alone or jointly with others, decides on the purposes ("why") and means ("how") of the processing of personal data, the AI Act focuses on the (traditionally single and) central provider of an AI system. However, since clients are autonomously in control of their data while the server is in control of the model, we see an inconsistency between what is controllable by the service provider and what he is responsible for. To close this gap, we need a two-pronged approach – technical and legal ("how" & "why"). Responsibility in FL should depend on the server's physical, technical, and legal ability to influence decentralized model training and configuration. This poses challenges. Unlike Article 26 and 28 GDPR, Article 8 et seq. of the AI Act do not provide details on governance in networked processing environments like FL systems.

The data protection assessment of the FL lifecycle may be impacted if the FL server sets requirements for the clients, which may lead to the server being classified as the controller under the GDPR. At first glance and from a strictly technical perspective, both the FL server and the FL clients fall under the provider concept of Article 3 GDPR, yet a "joint providership" (based on "joint controllership" under the GDPR, Article 26) does not exist under the AI Act. In fulfillment of its obligations under the AI Act, in particular Article 10, the FL server can set far-reaching requirements for the FL clients concerning the training of models and handling of training data. These requirements could lead to the FL server being classified as a controller under data protection law within the means of Article 4 GDPR, while the FL client is classified as a mere processor within the scope of Article 28 GDPR.

Thus, a key legal instrument for ensuring compliance with the AI Act and GDPR (Articles 26 & 28) is likely the development of specific General Terms and Conditions binding for both FL server and clients. The server operator, as the service provider, has to oblige clients to provide sufficient reporting compliant with the AI Act. This can be supported by cryptographic tools that minimize the need for trust among entities (Nguyen et al., 2021).

4. Considerations on Major Federated Learning Architectures

Cross-silo FL may allow for more flexibility in system design and responsibility distribution between clients and server than cross-device FL.

While Section 2 and Section 3 are generally applicable to any FL application, there are two major system architectures that create further opportunities to organize responsibilities: cross-silo and cross-device training.

4.1. Cross-Device Federated Learning

Cross-device FL typically entails a large number of devices (> 1,000). In such a setup, FL clients are characterized by having a very small number of local data samples and little participation time in the federated training process (Hard et al., 2018). This is a major challenge regarding client accountability and, ultimately, becomes problematic when a client should assume responsibility as a service provider. Hence, for practical considerations, all responsibility has to be assumed by the server in the cross-device setting and there is practically no room for client-side responsibility and a strong need for tools and methods that allow the FL server to cover all compliance criteria. This implies that the runtime environment on clients must be as encapsulated as possible, coupled with strict terms of service agreements.

4.2. Cross-Silo Federated Learning

In contrast, in cross-silo settings, individual clients hold a significant amount of data and participate in multiple training rounds, and usually come with higher computational capabilities than in cross-device FL. Typically, cross-silo FL can involve large institutions such as hospitals (Huang et al., 2022), which themselves have a high commitment to regulatory compliance and take strong precautions regarding security and privacy protection. As such, it is an open research direction to explore the synergies between established institutional processes (e.g., medical record keeping) and the AI Act requirements (e.g., on data transparency). The terms and conditions must be balanced between ensuring appropriate regulatory compliance and practical utility such that clients are incentivized to participate in training, and we can assume partial responsibility on the side of clients. Such synergies could help better balance the service provider responsibilities and reduce costs for clients and the server, not only improving the economic viability of FL but also its ecological footprint.

5. Conclusion

We study the FL life-cycle responsibilities under the AI Act. We find client-side responsibility for numerous steps, which practically limits the applicability of FL to open up additional data silos that would benefit the training of foundation models. Yet, there are promising directions that deserve increased attention such that a server operator can become the *service provider* without clients being required to assume extensive liability. With this, one can drive the adoption of FL and help decrease data bias by directly relying on user data. Further clarifying the outlined service provider question directly responds to the EU AI Office's call for contributions to help implement the AI Act (Nature, 2024).

Acknowledgements

This work is partially funded by the Bavarian Ministry of Economic Affairs, Regional Development and Energy (Grant: DIK0446/01).

References

- Andrew, G., Thakkar, O., McMahan, B., and Ramaswamy, S. Differentially private learning with adaptive clipping. *Advances in Neural Information Processing Systems*, 34: 17455–17466, 2021.
- Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., and Seth, K. Practical Secure Aggregation for Privacy-Preserving Machine Learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS '17*. ACM, October 2017. doi: 10.1145/3133956.3133982. URL <http://dx.doi.org/10.1145/3133956.3133982>.
- Council of the European Union. General data protection regulation (GDPR), apr 2016. URL <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A32016R0679>. Document 32016R0679.
- Council of the European Union. Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL - LAYING DOWN HARMONISED RULES ON ARTIFICIAL INTELLIGENCE (ARTIFICIAL INTELLIGENCE ACT) AND AMENDING CERTAIN UNION LEGISLATIVE ACTS, apr 2021. URL <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206>. Document 52021PC0206.
- Hard, A., Kiddon, C. M., Ramage, D., Beaufays, F., Eichner, H., Rao, K., Mathews, R., and Augenstein, S. Federated learning for mobile keyboard prediction, 2018. URL <https://arxiv.org/abs/1811.03604>.
- House Of Commons of Canada. An Act to enact the Consumer Privacy Protection Act, the Personal Information and Data Protection Tribunal Act and the Artificial Intelligence and Data Act and to make consequential and related amendments to other Acts, 6 2022. URL <https://www.parl.ca/DocumentViewer/en/44-1/bill/C-27/first-reading>.
- Huang, C., Huang, J., and Liu, X. Cross-silo federated learning: Challenges and opportunities, 2022. URL <https://arxiv.org/abs/2206.12949>.
- Li, T., Sahu, A. K., Talwalkar, A., and Smith, V. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60, May 2020. ISSN 1558-0792. doi: 10.1109/msp.2020.2975749. URL <http://dx.doi.org/10.1109/MSP.2020.2975749>.
- Liu, Y., Lou, R., and Wei, J. Auditing for federated learning: A model elicitation approach. In *The Fifth International Conference on Distributed Artificial Intelligence, DAI '23*. ACM, November 2023. doi: 10.1145/3627676.3627683. URL <http://dx.doi.org/10.1145/3627676.3627683>.
- Ma, C., Li, J., Ding, M., Shi, L., Wang, T., Han, Z., and Poor, H. V. When federated learning meets blockchain: A new distributed learning paradigm, 2020. URL <https://arxiv.org/abs/2009.09338>.
- Malinovsky, G., Horváth, S., Burlachenko, K. P., and Richtárik, P. Federated learning with regularized client participation. In *Federated Learning and Analytics in Practice: Algorithms, Systems, Applications, and Opportunities*, 2023. URL <https://openreview.net/forum?id=6CDBpf7kNG>.
- McMahan, B., Moore, E., et al. Communication-Efficient Learning of Deep Networks from Decentralized Data. In Singh, A. and Zhu, J. (eds.), *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pp. 1273–1282. PMLR, 20–22 Apr 2017. URL <https://proceedings.mlr.press/v54/mcmahan17a.html>.
- Nature. There are holes in Europe’s AI Act — and researchers can help to fill them. *Nature*, 625(7994): 216–216, January 2024. ISSN 1476-4687. doi: 10.1038/d41586-024-00029-4. URL <http://dx.doi.org/10.1038/d41586-024-00029-4>.
- Nguyen, D. C., Ding, M., Pham, Q.-V., Pathirana, P. N., Le, L. B., Seneviratne, A., Li, J., Niyato, D., and Poor, H. V. Federated learning meets blockchain in edge computing: Opportunities and challenges. *IEEE Internet of Things Journal*, 8(16):12806–12825, August 2021. ISSN 2372-2541. doi: 10.1109/jiot.2021.3072611. URL <http://dx.doi.org/10.1109/JIOT.2021.3072611>.
- Roy Chowdhury, A., Guo, C., Jha, S., and van der Maaten, L. Eiffel: Ensuring integrity for federated learning. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, CCS '22*. ACM, November 2022. doi: 10.1145/3548606.3560611. URL <http://dx.doi.org/10.1145/3548606.3560611>.
- Sánchez Sánchez, P. M., Huertas Celdrán, A., Xie, N., Bovet, G., Martínez Pérez, G., and Stiller, B. Federatedtrust: A solution for trustworthy federated learning.

Future Generation Computer Systems, 152:83–98, March 2024. ISSN 0167-739X. doi: 10.1016/j.future.2023.10.013. URL <http://dx.doi.org/10.1016/j.future.2023.10.013>.

The White House. Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, 10 2023. URL <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>.

Wang, S. and Ji, M. A unified analysis of federated learning with arbitrary client participation. In Oh, A. H., Agarwal, A., Belgrave, D., and Cho, K. (eds.), *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=qSs7C7c4G8D>.

Woisetschläger, H., Erben, A., Marino, B., Wang, S., Lane, N. D., Mayer, R., and Jacobsen, H.-A. Federated learning priorities under the european union artificial intelligence act, 2024a. URL <https://arxiv.org/abs/2402.05968>.

Woisetschläger, H., Isenko, A., Wang, S., Mayer, R., and Jacobsen, H.-A. A survey on efficient federated learning methods for foundation model training, 2024b. URL <https://arxiv.org/abs/2401.04472>.

Zhou, C., Li, Q., Li, C., Yu, J., Liu, Y., Wang, G., Zhang, K., Ji, C., Yan, Q., He, L., Peng, H., Li, J., Wu, J., Liu, Z., Xie, P., Xiong, C., Pei, J., Yu, P. S., and Sun, L. A comprehensive survey on pretrained foundation models: A history from bert to chatgpt, 2023. URL <https://arxiv.org/abs/2302.09419>.