The Human Side of Service Engineering, Vol. 108, 2023, 244–253
https://doi.org/10.54941/ahfe1003126

Cross-Silo Federated Learning in Enterprise Networks With Cooperative and Competing Actors

Kristina Müller and Freimut Bodendorf

Friedrich-Alexander-University of Erlangen-Nürnberg, Nuremberg, Germany

ABSTRACT

The performance and generalizability of AI-based enterprise applications depends on the quantity, quality, and diversity of training data. However, data usually exists in the form of data silos at individual sites. With *cross-silo federated learning* knowledge extracted from data silos that are distributed across multiple enterprises can be combined to improve the predictive performance of AI models without sharing and centralizing potentially sensitive raw data. The decentralized learning approach thus offers new privacy-preserving opportunities for cross-company collaboration, knowledge management, and the development of intelligent applications and services in federated enterprise networks. Since federated learning enables collaboration between both cooperating and competing companies, this literature review of application-based papers analyzes the differences in the design and strategic management of federated enterprise networks as a function of the actors' relationships.

Keywords: Cross-silo federated learning, Federated networks, Enterprise collaboration

INTRODUCTION

Artificial Intelligence (AI) is one of the key drivers for optimizing business process efficiency, supporting human decision making, and creating new intelligent products or services. The accuracy and robustness of AI models depends on the availability of large amounts of high-quality and diverse training data. However, data usually exists in form of homogeneous data silos at individual sites, which leads to overfitting and thus a poor generalizability of the AI models (Lyu et al., 2020).

To address these AI training data requirements, multiple organizations can collaborate in a data sharing ecosystem and train more robust AI models on a centralized corpus of multi-organizational data. However, the strategy of *collaborative data sharing* is often not feasible due to confidentiality concerns and privacy regulations (Sheller et al., 2020).

Although originally introduced for collaborative yet decentralized training of AI models on edge devices (*cross-device federated learning*) (McMahan et al., 2017) the benefits of federated learning can also be unleashed in an enterprise context (*cross-silo federated learning*) (Kairouz et al., 2021). By combining the knowledge of multiple organizations without sharing raw data, cross-silo federated learning offers a solution to situations where AI models of single organizations fail but data sharing is not possible. In this *data-private collaborative learning method*, multiple data owners (also referred to as actors) jointly contribute to an overall AI model based on their decentralized data (Sheller et al., 2020). To this end, each actor trains a local model on its own data, which is then aggregated with the local models of the other participants in the federated network (Li et al., 2021). Repeating the training process several times results in a joint AI model with better accuracy and generalizability than those of the respective individual models (Sim et al., 2020).

Since sensitive raw data remains decentralized with its owners, thus preserving privacy and data protection, cross-silo federated learning facilitates AI alliances between (1) several internal units of a company (Röder et al., 2022), (2) a company and its customers (Zhang et al., 2021), (3) within supply chains (Liu et al., 2022), and between (4) competitors (Deng et al., 2022). This diversity of collaboration opportunities in federated enterprise networks requires differentiation of their respective design and strategic management requirements.

Therefore, this paper aims to answer the following research question: *What are the differences in the design and strategic management of cross-silo federated networks depending on the relationship of the actors involved?*

The remainder of the paper is organized as follows. First, the methodological approach of this study is described. Next, the findings for the respective actor relationships are presented and finally summarized in the conclusion.

METHODOLOGY

Following the guidelines of Brocke et al. (2009) and Webster and Watson (2002), a literature review of application-based papers is conducted to identify and synthesize the different requirements for the design and strategic management of cross-silo federated enterprise networks as a function of the actors involved.

Scopus, IEEE Xplore, and ACM Digital Library are queried with the search term “federated learning” AND (“cross-silo” OR “enterprise” OR “institution”) to search for papers containing the term in their title, abstract, or keywords. The database search takes place between late December 2022 and mid-January 2023 and yields a total of 688 hits, which are screened by title and, in case of ambiguity, also by abstract.

In the first step, all research papers that do not describe an application of federated learning to a specific use case are excluded. By applying the exclusion criteria, the search results are narrowed down to 33 journal and conference papers, which are reviewed based on their full text and complemented by a backward and forward search. Based on full-text reading, 14 papers are identified that describe specific requirements arising from the relationship of the actors involved in a cross-silo federated network. These papers are the basis of this study.

DESIGN AND STRATEGIC MANAGEMENT OF CROSS-SILO FEDERATED NETWORKS

The decentralized training process of cross-silo federated learning has already been simulated for various benchmark data sets and tasks in the last two years (e.g., prediction of the remaining machine lifetime (Ranathunga et al., 2022), detection of fraudulent transactions (Myalil et al., 2021), or forecasting the length of patient stay in hospitals (Rahman et al., 2022)). Thereby, most studies focus on demonstrating comparable model performance to traditional training on the centralized version of the data corpus, without considering the specific design and strategic management requirements resulting from different enterprise relationships.

Moreover, the majority of the papers deals with the application of cross-silo federated learning in competitive relationships. To date, only a few papers implement the decentralized learning method among cooperating enterprises.

In federated networks, the problem of statistical heterogeneity is prevalent because the distribution of data between actors is often inconsistent (e.g., different label distribution at each actor) (Myalil et al., 2021). Consequently, the assumption of an independent and identical distribution (IID) among actors is not met, which can severely degrade the global model performance for certain actors. As countermeasure, several of the studies reviewed describe various personalization techniques that are important to leverage the benefits of federated learning (Myalil et al. (2021), Durrant et al. (2022), Kanani et al. (2021), Deng et al. (2022), Zhang et al. (2021), and Yu and Huang (2022)). However, their implementation is independent of the relationship between the actors involved and is therefore not further described here.

Although local data is not exposed in the federated learning process, model parameters can still leak sensitive information about the raw data. Therefore, some of the studies reviewed combine federated learning with additional privacy protection measures to avoid the reconstruction of raw data. The most common techniques in this context are differential privacy (Schreyer et al., 2022), homomorphic encryption (Ranathunga et al., 2022), and secure multi-party computation (Deng et al., 2022). The choice of one of the techniques is rarely justified by conceivable attacker scenarios (malicious central server, insider or outsider attacker) in the particular use case. However, this is important since privacy protection measures are associated with a degradation of model performance or increased computation and communication costs (Abdulrahman et al., 2021).

Furthermore, some papers combine federated learning with blockchain technology to ensure credibility and traceability of model updates and prevent model poisoning attacks (Ranathunga et al., 2022).

In the following, the focus is on the design and strategic management aspects, which so far are justified from the relationships of the actors involved. The four previously defined types of cross-silo federated enterprise networks are distinguished in terms of the goal of collaboration, the data partitioning among actors, the communication architecture, and the implicit incentives of participation or the design of explicit incentives through rewards. Table 1 summarizes the results.

Table 1. Characteristics of different cross-silo federated networks.

	Internal Business Unit	Customer Relationship	Supply Chain	Competitors
Relationship type	Cooperative (internal)	Cooperative (external)	Cooperative (external)	Competitive
Objective	Indirect consolidation of internal, physically dispersed data sets	Delivering enhanced intelligent customer services	Combining data features of a joint sample space	Increasing the amount and diversity of training data
Data partitioning	Horizontal	Horizontal	Vertical	Horizontal
Communication architecture	Centralized	Centralized	Centralized	Decentralized or centralized
Participation incentives	Cost reduction in AI development (implicit)	Monetary remuneration of customers depending on their contribution to the training process (explicit)	Increasing efficiency and resilience in hierarchical actor structures (implicit)	Model performance corresponding to the value of each actor's local updates to the global model (explicit)
Literature Sources	Röder et al. (2022)	Zhang et al. (2021), Mohr et al. (2021), Schreyer et al. (2022)	Wang et al. (2022), Liu et al. (2022), Che et al. (2022)	Myalil et al. (2021), Durrant et al. (2022), Ranathunga et al. (2022), Nguyen et al. (2022), Kanani et al. (2021), Deng et al. (2022), Yu and Huang (2022)

Internal Business Units

Overall, some groups or individual companies have large and diverse data sets internally that are potentially suitable for training AI models, but are dispersed across different subsidiaries, business units, or locations. Such internal data silos can also arise as part of mergers and acquisitions when data inventories remain in separated databases of the former individual companies (Brundyn et al., 2022). For the training of AI models, the dispersed data is traditionally consolidated on a central server (Zhang et al., 2021). However, internal data transfer is associated with high transmission and storage costs and, in the case of personal data, is complicated by legal requirements such as the European General Data Protection Regulation (GDPR) (Brundyn et al., 2022). In this case, federated learning is suitable for connecting internal, physically dispersed data silos and training AI models on the entire data available within the group or enterprise.

Sales forecasting based on historical sales data from multiple stores of a retail company (Röder et al., 2022) or fraud detection based on credit card transactions from multiple stores (Brundyn et al., 2022) are just two examples of the in-plant use of federated learning. In the two examples, each store owns a portion of the total sample space of the group or company with the same features of interest for AI model training (horizontal data partitioning).

The collaborative training process is initiated and controlled by the corporate headquarters, i.e., the participating subsidiaries or branches send their locally trained models to a central server. Besides the task of aggregating the local models, the headquarters can also contribute its own model. Subsequently, the final global model is available to both the headquarters and the participating subsidiaries or branches and can also be shared with other internal business units.

Customer Relationship

A company can learn from the data of multiple customers in federated networks without having to access their raw data centrally. The functionality of the global model, which is the aggregation of the local customer models, is leveraged by the company to provide enhanced intelligent customer services aimed at strengthening customer loyalty and attracting new customers (Zhang et al., 2021).

As described by Mohr et al. (2021), predictive maintenance services are an example of this federated learning use case. Here, the aim is to reduce unplanned production downtime by identifying and replacing the parts of a machine that need repair before they fail. Due to the relatively rare occurrence of machine failures, individual customers may not be able to train a high-quality prediction model using their own data. Based on the more accurate prediction results of the model trained on the data of multiple customers owning the same machine (horizontal data partitioning), the machine manufacturer can offer its customers a maintenance service tailored to their individual needs.

In a federated network formed by a company and its customers, tasks and resources are distributed differently among the actors. While the customers are responsible for data collection and local model training the company orchestrates the federated learning process and provides the corresponding services. Accordingly, this use case of federated learning employs a centralized communication architecture with a central server as coordinator and aggregator, where customers interact directly only with the central server, but not with other customers (Schreyer et al., 2022).

To be able to offer its intelligent services, the company depends on the willingness of its customers to participate in the learning process. Customers can be encouraged to contribute their data and computing resources through suitable incentive mechanisms (Zhang et al., 2021). On the one hand, customers should be rewarded according to their contribution (Zhang et al., 2021). On the other hand, each customer should receive the same service quality, i.e., comparable predictive performance of the model for all customers (Li et al., 2020). Consequently, monetary rewards that do not affect model performance are suitable incentives in this federated learning use case, e.g., in form of discounts for services or the purchase of new products from the company (Zhang et al., 2021).

Supply Chain

Multiple enterprises or institutions may hold different data characteristics of an overlapping set of sample IDs that, when combined, provide a more comprehensive picture of the joint sample space (Che et al., 2022). In this case, data is vertically distributed. Although there are few implementations to date (e.g., Wang et al. (2022), Liu et al. (2022) or Che et al. (2022)), vertical federated learning lends itself to training AI models based on all available data features without the need to share data.

For instance, by combining different information from suppliers, manufacturers, distributors, and retailers, vertical federated learning offers great potential for predicting supply chain risks, such as delivery delays or demand changes, to minimize disruptions that can significantly impact the performance and efficiency of the entire supply chain and the individual actors involved (Validi et al., 2018). In addition to traditional supply chains, federated learning can also be applied to develop diagnosis and treatment models based on scattered data collected from the same patients by different healthcare institutions such as general practitioners, specialists, hospitals, etc., thereby avoiding multiple data collection (Che et al., 2022).

Since data is partitioned by features among the actors and usually one of them owns the label of interest, a different training procedure is employed in vertical federated learning compared to the horizonal case, where actors share the same feature and label space with each actor owning the respective labels for its individual samples (Xu et al., 2022). The vertical training process consists of two phases. The first step is to identify the overlap of data samples based on identifiers such as universal or predefined identification numbers. This is followed by collaborative model training on the common sample space. Coordinated by a central server of a trusted third party, actors exchange intermediate model updates that assist the other actors in gradient calculation, while being protected from raw data retrieval by homomorphic encryption (Wang et al. (2022) and Liu et al. (2022)).

Competitors

Since the distributed training process of federated learning enables collaboration despite simultaneous competition, federated networks can also be established by enterprises in the same industry with similar business models and AI use cases. All participating data owners aim for a robust AI model with high predictive performance for their local tasks by training on a larger and more diverse dataset, overcoming their individual data bottlenecks.

Several research papers simulate the implementation of cross-silo federated learning in competitive relationships for a variety of use cases in different industries. So far, most of them develop diagnosis and treatment models for healthcare (e.g., Nguyen et al. (2022), Kanani et al. (2021) or Yu and Huang (2022)), with a few papers investigating use cases in finance (e.g., Myalil et al. (2021)), manufacturing (e.g., Ranathunga et al. (2022)), and agriculture (e.g., Durrant et al. (2022)). In such cases, the actors share or have previously aligned the same feature and label space, while each actor owns a subset of the sample space (horizontal data partitioning).

The competitive relationship generally implies that the actors involved treat each other with distrust and their actions are driven by maximizing their own profit. Against this background, a centralized implementation of federated networks, where one of the involved actors hosts the central server for model aggregation and thus gains control over the training process as well as exclusive access to the potentially sensitive model updates of all other actors, is impractical in this setting. Instead, Nguyen et al. (2022), Ranathunga et al. (2022), and Yu and Huang (2022) replace the central server with a peer-to-peer network structure. In this decentralized communication architecture, model updates are exchanged between neighboring actors on a communication graph and the received models are aggregated locally by each actor (Yu and Huang, 2022). As an alternative to circumvent the high communication costs of decentralized communication architectures, a trusted third party can be employed as central authority (Durrant et al., 2022). In addition, Durrant et al. (2022) and Kanani et al. (2021) apply local differential privacy to the centralized architecture to protect model updates from raw data disclosure on the central server, but at the cost of degrading model performance.

Typically, enterprises possess different amounts and quality of local data. Data owners with a small amount of data have a strong incentive to participate in the federated learning process and benefit from a much more accurate and robust model compared to their individual model (Kanani et al., 2021). To ensure fairness between competitors with different resources and incentivize participation of enterprises with high quality and large data sets, actors should be rewarded according to their contribution to the collaborative model training (Ranathunga et al., 2022). Considering that all participants intend to use the final model for their internal task, model rewards are appropriate for this purpose, i.e., the local accuracy of the aggregated model reflects the value of each participant's contribution (Sim et al. (2020) and Ranathunga et al. (2022)). In a decentralized federated network, model rewards can be provided by the hierarchical arrangement of actors (Ranathunga et al., 2022). Based on their contribution, actors are placed in an ascending order. As the training process proceeds sequentially from bottom up, actors with higher contribution at the top of the hierarchy obtain aggregated models containing the local models of several predecessors and are expected to provide more accurate and robust predictions. For realizing desired model rewards in the case of a central third party, Sim et al. (2020) propose to train a separate model for each actor on the central server, adding varying amounts of noise to the aggregated data of the other participants depending on the contribution of the respective actor.

CONCLUSION

Federated learning offers a new way for enterprises to collaborate on AI development and maximize the value of their distributed data silos. Through the decentralized training process, knowledge can be shared and aggregated internally between different locations of a company, but also externally between a company and its customers, within a supply chain, as well as between competitors, without centralizing and disclosing raw data.

This literature review provides an overview of the specific design and strategic management requirements as a function of the actors involved in a federated enterprise network. These differ in terms of the goal of collaboration, the data partitioning among the actors, the design of the communication architecture, and the incentives for participation. The overview serves as a starting point for managing federated enterprise networks in practice, as well as for further research investigating the implementation of federated learning in different actor relationships.

It should be noted that the application-based papers reviewed in this work simulate the federated learning process on a single server and therefore do not face the challenges of real-world implementation. Consequently, most of the papers focus only on selected aspects of design and strategic management arising from the relationships of the actors. Other aspects, such as the choice of privacy measures or the implementation of blockchain, are rarely justified by the requirements of the federated network in question. Further research needs to differentiate these aspects in terms of the actors involved in the training process.

REFERENCES

Abdulrahman, S., Tout, H., Ould-Slimane, H., Mourad, A., Talhi, C, & Guizani, M. (2021). A Survey on Federated Learning: The Journey From Centralized to Distributed On-Site Learning and Beyond. IEEE Internet of Things Journal, 8(7), 5476–5497. https://doi.org/10.1109/JIOT.2020.3030072

Brocke, J., Simons, A., Niehaves, B., Niehaves, B., Reimer, K., Plattfaut, R., & Cleven, A. (2009). RECONSTRUCTING THE GIANT: ON THE IMPORTANCE OF RIGOUR IN DOCUMENTING THE LITERATURE SEARCH PROCESS. ECIS 2009 Proceedings, Article 161. https://aisel.aisnet.org/ecis2009/161

Brundyn, A., Scoullos, E., & Williams, D. (2022, August 16). Using Federated Learning to Bridge Data Silos in Financial Services. NVIDIA. https://developer.nvidia.com/blog/using-federated-learning-to-bridge-data-silos-in-financial-services/

Che, S., Kong, Z., Peng, H., Sun, L., Leow, A., Chen, Y., & He, L. (2022). Federated Multi-view Learning for Private Medical Data Integration and Analysis. ACM Transactions on Intelligent Systems and Technology, 13(4), Article 61. https://doi.org/10.1145/3501816

Deng, T., Li, Y., Liu, X., & Wang, L. (2022). Federated learning-based collaborative manufacturing for complex parts. Journal of Intelligent Manufacturing. https://doi.org/10.1007/s10845-022-01968-3

Durrant, A., Markovic, M., Matthews, D., May, D., Enright, J., & Leontidis, G. (2022). The role of cross-silo federated learning in facilitating data sharing in the agri-food sector. Computers and Electronics in Agriculture, 193, Artikle 106648. https://doi.org/10.1016/j.compag.2021.106648

Kairouz, P., McMahan, H. B., Avent, B., et al. (2021). Advances and Open Problems in Federated Learning. Foundations and Trends in Machine Learning, 14(1-2). https://doi.org/10.1561/2200000083

Kanani, P., Marathe, V. J., Peterson, D., Harpaz, R., & Bright, S. (2021). Private Cross-Silo Federated Learning for Extracting Vaccine Adverse Event Mentions. In Kamp, M., et al. (Eds.) Machine Learning and Principles and Practice of Knowledge Discovery in Databases, Communications in Computer and Information Science 1525 (pp. 490–505). Springer. https://doi.org/10.1007/978-3-030-93733-1_37

Li, Q., Wen, Z., Wu, Z., Hu, S., Wang, N., Li, Y., Liu, X., & He, B. (2021). A Survey on Federated Learning Systems: Vision, Hype and Reality for Data Privacy and Protection. IEEE Transactions on Knowledge and Data Engineering. https://doi.org/10.1109/TKDE.2021.3124599

Li, T., Sanjabi, M., Beirami, A., &. Smith, V. (2020). Fair Resource Allocation in Federated Learning. International Conference on Learning Representations.

Liu, H., Ye, B., Qin, Z., & Zhang, J. (2022). The High-Performance Solution with Federated Learning in Supply Chain System. Proceedings of the 8th International Conference on Computing and Artificial Intelligence, 241–245. https://doi.org/10.1145/3532213.3532249

Lyu, L., Yu, J., Nandakumar, K., Li, Y., Ma, X., Jin, J., Yu, H., & Ng, K. S. (2020). Towards Fair and Privacy-Preserving Federated Deep Models. IEEE Transactions on Parallel and Distributed Systems, 31(11), 2524-2541. https://doi.org/10.1109/TPDS.2020.2996273

McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Agüera y Arcas, B. (2017). Communication-Efficient Learning of Deep Networks from Decentralized Data. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, 54.

Mohr, M., Becker, C., Möller, R., & Richter, M. (2021). Towards Collaborative Predictive Maintenance Leveraging Private Cross-Company Data. In Reussner, R. H., Koziolek, A., & Heinrich, R. (Eds.), INFORMATIK 2020 (pp.427–432). Gesellschaft für Informatik. https://doi.org/10.18420/inf2020_39

Myalil, D., Raja, M. A., Apte, M., & Lodha, S. (2021). Robust Collaborative Fraudulent Transaction Detection using Federated Learning. Proceedings of the 20th IEEE International Conference on Machine Learning and Applications, 373–378. https://doi.org/10.1109/ICMLA52953.2021.00064

Nguyen, T. V., Dakka, M. A., Diakiw, S. M., VerMilyea, M. D., Perugini, M., Hall, J. M. M., & Perugini, D. (2022). A novel decentralized federated learning approach to train on globally distributed, poor quality, and protected private medical data. Scientific Reports, 12, Article 8888. https://doi.org/10.1038/s41598-022-12833-x

Rahman, M. M., Kundu, D., Suha, S. A., Siddiqi, U. R., and Dey, S. K. (2022). Hospital patients' length of stay prediction: A federated learning approach. Journal of King Saud University - Computer and Information Sciences, 34(10), 7874–7884. https://doi.org/10.1016/j.jksuci.2022.07.006

Ranathunga, T., McGibney, A., Rea, S., & Bharti, S. (2022). Blockchain based Decentralised Model Aggregation for Cross-Silo Federated Learning in Industry 4.0. IEEE Internet of Things Journal. https://doi.org/10.1109/JIOT.2022.3218704

Röder, M., Kowalczyk, P. & Thiesse, F. (2022). TRACING DOWN THE VALUE OF CO-CREATION IN FEDERATED AI ECOSYSTEMS. ECIS 2022 Research Papers, Article 155. https://aisel.aisnet.org/ecis2022_rp/155

Schreyer, M., Sattarov, T., & Borth, D. (2022). Federated and Privacy-Preserving Learning of Accounting Data in Financial Statement Audits. Proceedings of the Third ACM International Conference on AI in Finance, 105–113. https://doi.org/10.1145/3533271.3561674

Sheller, M. J., Edwards, B., Reina, G. A., Martin, J., Pati, S., Kotrotsou, A., Milchenko, M., Xu, W., Marcus, D., Colen, R. R., & Bakas, S. (2020). Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. Scientific Reports, 10, Article 12598. https://doi.org/10.1038/s41598-020-69250-1

Sim, R. H. L., Zhang, Y., Chan, M. C., & Low, B. K. H. (2020). Collaborative machine learning with incentive-aware model rewards. Proceedings of the 37th International Conference on Machine Learning, 8927–8936.

Validi, S., Dani, S., & Antoniou, G. (2018). Supply chain risk management and artificial intelligence: state of the art and future research directions. International Journal of Production Research, 57(7), 2179–2202. https://doi.org/10.1080/00207543.2018.1530476

Wang, H., Xie, F., Duan, Q., & Li, J. (2022). Federated Learning for Supply Chain Demand Forecasting. Mathematical Problems in Engineering. https://doi.org/10.1155/2022/4109070

Webster, J. and Watson, R. T. (2002). Analyzing the Past to Prepare for the Future: Writing a Literature Review. MIS Quarterly, 26(2), xiii-xxiii. https://www.jstor.org/stable/4132319

Xu, R., Baracaldo, N., Zhou, Y., Abay, A., & Anwar, A. (2022). Privacy-Preserving Vertical Federated Learning. In Ludwig, H. and Baracaldo, N. (Eds.) Federated Learning (pp. 417–438). Springer. https://doi.org/10.1007/978-3-030-96896-0_18

Yu, L. and Huang, J. (2022). Cyclic Federated Learning Method Based on Distribution Information Sharing and Knowledge Distillation for Medical Data. Electronics, 11(23), Article 4039. https://doi.org/10.3390/electronics11234039

Zhang, W., Lu, Q., Yu, Q., Li, Z., Liu, Y., Lo, S. K., Chen, S., Xu, X., & Zhu, L. (2021). Blockchain-Based Federated Learning for Device Failure Detection in Industrial IoT. IEEE Internet of Things Journal, 8(7), 5926–5937. https://doi.org./10.1109/JIOT.2020.3032544