Article

Not peer-reviewed version

# STellar-FL: A Decentralized Federated Learning Architecture for Scalable Cross-Institution AI Under Network-Constrained Environments

Hsu Ouyang , Shao-Yu Chang , Chao-Chun Chuang , Chia-Hung Yang , Yu-Chieh Lin *

Posted Date: 23 March 2026

doi: 10.20944/preprints202603.1776.v1

Keywords: decentralized federated learning; privacy-preserving AI; cross-institution collaboration

*Article*

# STellar-FL: A Decentralized Federated Learning Architecture for Scalable Cross-Institution AI Under Network-Constrained Environments

**Hsu Ouyang [1], Shao-Yu Chang [1], Chao-Chun Chuang [2], Chia-Hung Yang [1] and Yu-Chieh Lin [3,*]**

[1] ST Medical Technology INC
[2] National Center of High-Performance Computing
[3] Digital Medicine Center, Translational Health Research Institute, Vilnius University
* Correspondence: yu-chieh.lin@mf.vu.lt

**Abstract**

Federated learning (FL) enables collaborative model training without sharing raw data and has become an important paradigm for privacy-sensitive applications such as healthcare and other regulated domains. However, most existing federated learning frameworks rely on centralized coordination servers, fixed network configurations, and complex infrastructure requirements, which limit their deployment in real-world institutional environments with strict cybersecurity and data governance constraints. In this work, we propose STellar-FL, a decentralized federated learning architecture designed for scalable cross-institution model training under network-constrained environments. The proposed framework adopts a microservice-based design consisting of a federated training orchestration module, a distributed communication layer, and federated execution nodes. STellar-FL enables secure model exchange through relay-assisted peer connectivity, eliminates the need for centralized servers with public IP exposure, and provides a unified workflow for model development, deployment, and validation. Compared with conventional federated learning frameworks, STellar-FL reduces deployment complexity, improves system robustness by removing single points of failure, and supports flexible collaboration across heterogeneous institutional infrastructures. The proposed architecture provides a practical foundation for real-world privacy-preserving AI deployment in healthcare and other data-sensitive domains.

**Keywords:** decentralized federated learning; privacy-preserving ai; cross-institution collaboration

## 1. Introduction

The rapid adoption of artificial intelligence (AI) has created increasing demand for large-scale collaborative model training across institutions[1-3]. However, in highly data-sensitive domains such as healthcare and military, data are often distributed across institutions due to privacy regulations, data sovereignty requirements, and institutional governance policies. These constraints make centralized data aggregation infeasible and limit the development of robust machine learning models[4-6].

Federated learning (FL) has emerged as a promising paradigm for collaborative model training without sharing raw data[7-10]. By enabling local model updates to be aggregated across distributed nodes, FL allows institutions to jointly train models while preserving data privacy. As a result, FL has gained significant attention in data-sensitive applications, including medical imaging analysis, clinical prediction, and digital pathology.

Despite its promise, existing federated learning frameworks remain difficult to deploy in real-world institutional environments[11]. Most current FL systems rely on centralized coordination servers, fixed network configurations, and predefined communication endpoints. These requirements introduce several practical limitations. First, centralized FL servers typically require public IP

 doi:10.20944/preprints202603.1776.v1



addresses and port forwarding privileges, which are often restricted in hospital networks due to cybersecurity policies. Second, virtual private network (VPN) configurations are frequently required to establish communication channels[12], introducing additional setup complexity and security concerns. Third, deployment and maintenance of heterogeneous runtime environments, such as container-based execution pipelines, create substantial operational overhead. These constraints significantly hinder the adoption of federated learning in real-world cross-institution collaborations.

Recent research has explored decentralized federated learning to reduce reliance on centralized coordination[13-14]. However, existing approaches primarily focus on algorithmic optimization or peer-to-peer model aggregation and often lack system-level solutions for deployment challenges such as firewall constraints, distributed resource management, and scalable infrastructure integration. In particular, current frameworks rarely address the need for storage-aware model synchronization, flexible network connectivity, and unified training workflows in heterogeneous institutional settings.

To address these limitations, we propose **STellar-FL**, a decentralized federated learning architecture designed for practical deployment in privacy-sensitive environments. The proposed framework integrates distributed storage, peer-to-peer communication, and modular microservice components to enable scalable cross-institution model training without requiring centralized infrastructure or fixed network endpoints. The architecture introduces a unified training workflow, firewall-compatible communication mechanisms, and distributed resource coordination to simplify federated learning deployment across heterogeneous environments.

## 2. Related Work

### *2.1. Federated Learning Frameworks*

Federated learning was first introduced as a distributed training paradigm that enables collaborative model optimization without centralized data sharing. The standard federated learning setting typically relies on a centralized server that coordinates model aggregation across participating clients using iterative optimization schemes such as Federated Averaging (FedAvg)[15].

Several production-grade federated learning frameworks have been developed to support practical deployment. NVFlare (NVIDIA), OpenFL (Intel), and Flower provide infrastructure for distributed training, model synchronization, and secure communication across nodes[16-18]. These frameworks commonly adopt a client–server architecture in which a central coordinator manages training rounds and aggregation procedures. While effective in controlled environments, such centralized designs introduce deployment challenges in institutional settings where network accessibility and infrastructure configuration are restricted.

Moreover, many existing frameworks require fixed communication endpoints, public network accessibility, or virtual private network configurations to establish secure connections. These requirements limit scalability and hinder adoption in environments with strict cybersecurity policies, such as hospital networks.

### *2.2. Decentralized and Peer-to-Peer Federated Learning*

To reduce dependence on centralized coordination, decentralized federated learning approaches have been proposed[13-14, 19-21]. These methods typically adopt peer-to-peer communication schemes or distributed optimization strategies, allowing nodes to exchange model updates without a central server. Prior work has explored gossip-based aggregation, decentralized stochastic gradient descent, and blockchain-based federated learning architectures.

Although these approaches improve system robustness and remove single points of failure, most existing decentralized FL methods primarily focus on optimization algorithms rather than deployment infrastructure. Practical considerations such as firewall compatibility, distributed storage management, and system-level orchestration remain underexplored. In addition, many decentralized approaches assume stable network connectivity or predefined peer relationships, which may not hold in real-world institutional environments.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 23 March 2026 doi:10.20944/preprints202603.1776.v1



### 2.3. Secure and Privacy-Preserving Learning

Security and privacy preservation are central concerns in federated learning. Prior work has proposed secure aggregation protocols, differential privacy mechanisms, and cryptographic techniques to prevent information leakage from model updates[22-24]. These approaches aim to protect data confidentiality and mitigate adversarial threats during collaborative training.

While these techniques enhance privacy guarantees, they do not address infrastructure-level challenges related to deployment complexity, network configuration, and resource coordination across heterogeneous environments. Practical adoption of federated learning requires both algorithmic security mechanisms and system-level architectural design.

## 3. STellar-FL Architecture

### 3.1. System Overview

In contrast to existing federated learning approaches, STellar-FL is a decentralized architecture designed to enable scalable and secure cross-institution model training without relying on centralized coordination infrastructure. An overview of the proposed STellar-FL architecture is illustrated in Figure 1. The system adopts a modular microservice-based design that integrates distributed communication and unified training orchestration, enabling deployment across heterogeneous institutional environments.


**Figure 1.** Overview of STellar-FL system architecture.

Unlike conventional federated learning frameworks that depend on centralized servers with fixed network configurations, STellar-FL eliminates the requirement for public IP addresses, port forwarding privileges, and virtual private network (VPN) configurations. Instead, the system employs distributed relay and proxy mechanisms to enable firewall-compatible communication between participating nodes.

The overall architecture consists of three major components:

1. **Federated Training Orchestration(STellar-FL and STefan-FL)** – the STellar-FL service orchestrates training tasks, device discovery, and deployment of FL clients across institutions;

 doi:10.20944/preprints202603.1776.v1



STefan-FL provides the unified FL workflow, framework adapters (NVFlare, OpenFL, Flower), and simulation.

2. **Distributed Communication Layer (STellar-Go)** – provides peer-to-peer connectivity, relay services, and distributed resource coordination.
3. **Federated Execution Nodes** – institutional compute nodes that perform local model training and participate in distributed optimization. Model parameters and training state are synchronized through the communication layer and the FL orchestration workflow.

Each participating institution operates an independent node that maintains local data and performs local model updates. Model parameters and training states are exchanged through distributed communication channels rather than a centralized aggregation server.

*3.2. System Architecture and Network Topology*

STellar-FL adopts a decentralized network topology in which participating institutions operate autonomous training nodes that collaboratively perform distributed model optimization. The architecture supports dynamic node participation and does not require persistent direct connections between all peers, allowing flexible deployment across heterogeneous network environments. The overall network topology is illustrated in Figure 2.

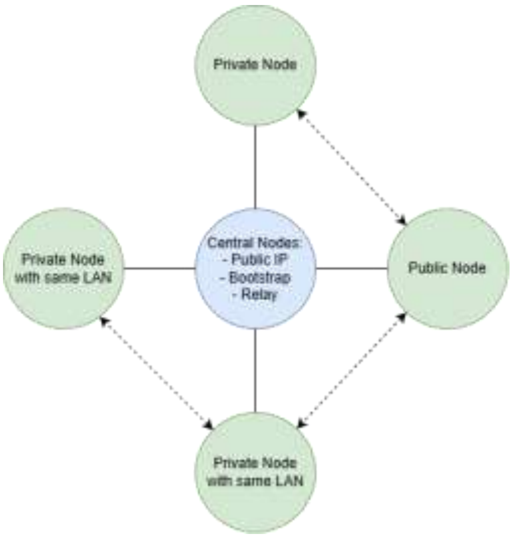


**Figure 2.** Network topology of the STellar-FL decentralized communication architecture.

As shown in Figure 2, the system consists of multiple private nodes, which represent institutional training environments that maintain local data and perform model updates, and publicly reachable nodes that provide communication support. A central node provides bootstrap and relay services to facilitate node discovery and communication initialization. These nodes do not store sensitive data or perform centralized aggregation; instead, they serve as communication intermediaries that establish secure peer-to-peer channels between training nodes.

To enable connectivity under restricted network conditions, STellar-FL relies on relay-assisted communication and outbound connections from private institutional networks. This design removes the requirement for direct inbound access to institutional infrastructure and enables deployment in environments with strict firewall policies.

The communication layer provides the following capabilities:

- dynamic node discovery,
- proxy-based message routing,
- distributed task coordination,
- secure model exchange.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 23 March 2026 doi:10.20944/preprints202603.1776.v1



By decoupling communication infrastructure from training execution, the proposed architecture improves deployment flexibility and supports scalable cross-institution collaboration in network-constrained environments.

### *3.3. Distributed Communication and Connectivity Mechanism*

A key design objective of STellar-FL is to enable federated learning deployment in environments where direct network access is restricted. To address this challenge, STellar-FL introduces a distributed communication layer that supports proxy-based connectivity, secure message routing, and relay-assisted peer communication. The architecture of the communication layer is illustrated in Figure 3.

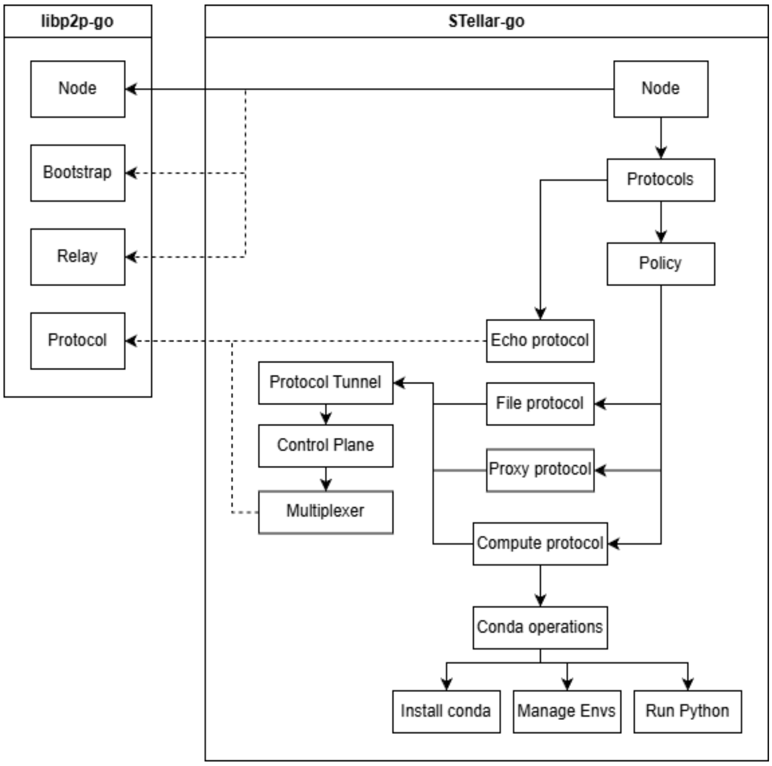


**Figure 3.** Distributed communication architecture of STellar-FL.

As shown in Figure 3, the communication layer adopts a protocol-driven design built upon a peer-to-peer networking framework. The system provides a protocol tunnel and control plane that manage node communication, task coordination, and resource access across distributed environments. Multiple communication protocols are supported, including proxy-based routing, file transfer, and compute execution protocols, enabling flexible interaction between participating nodes.

The communication mechanism provides the following capabilities:

- firewall-compatible communication channels,
- relay-assisted peer connectivity,
- encrypted message transmission,
- distributed task synchronization.

Instead of requiring publicly accessible training servers, institutional nodes establish outbound connections to relay services that facilitate bidirectional communication between peers. This design removes the need for inbound network exposure, significantly simplifies deployment in restricted institutional environments, and reduces cybersecurity risks associated with open network ports.

By decoupling network connectivity from training execution, the proposed communication mechanism enables robust cross-institution collaboration under heterogeneous network conditions.

### 3.4. Unified Federated Training Orchestration

STellar-FL provides a unified federated learning workflow that standardizes model implementation, data loading, and training execution across heterogeneous environments. The architecture of the unified orchestration module is illustrated in Figure 4.

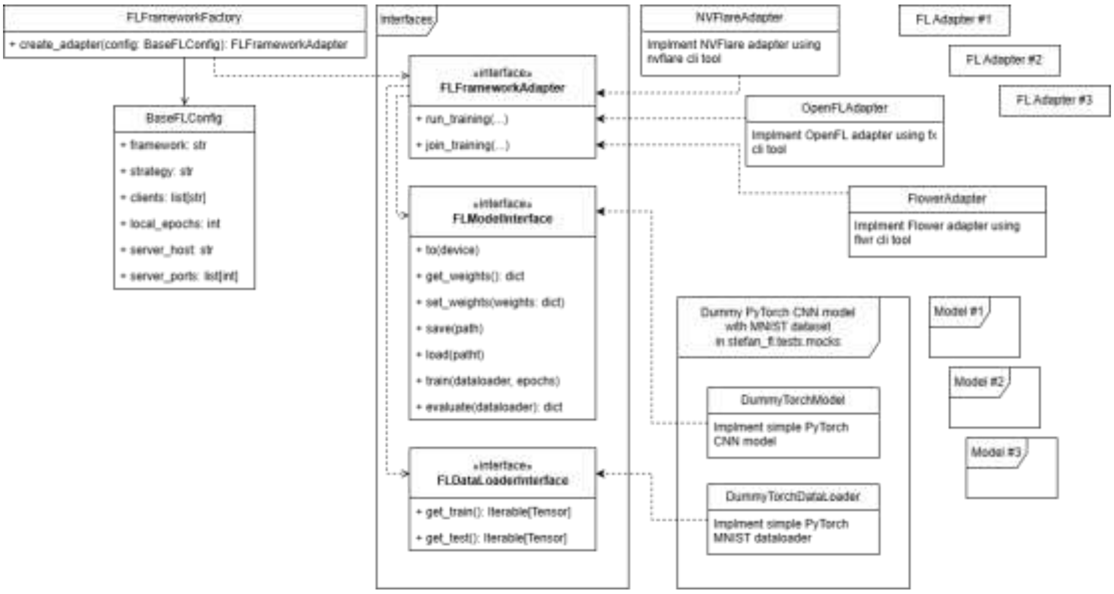


**Figure 4.** Unified federated training orchestration architecture in STellar-FL.

As shown in Figure 4, the orchestration module introduces a set of standardized interfaces that decouple model implementation and data processing from the underlying federated learning infrastructure. The system adopts an adapter-based design that abstracts different federated learning backends and provides consistent interfaces for model training and evaluation.

Through this abstraction layer, STellar-FL supports integration with multiple federated learning frameworks, including NVFlare, OpenFL, and Flower, while maintaining a unified execution workflow. By separating model implementation from execution infrastructure, the framework enables flexible deployment across heterogeneous environments and improves compatibility with existing federated learning ecosystems.

The unified workflow reduces engineering overhead, simplifies system configuration, and improves reproducibility in collaborative research settings. In addition, STellar-FL provides automated environment configuration and runtime management mechanisms that ensure consistent dependency setup and standardized execution pipelines across different institutional infrastructures. This design reduces operational complexity associated with container-based deployment strategies and improves compatibility across diverse operating systems and hardware configurations.

### 3.5. Comparison with Conventional Federated Learning Frameworks

To better position the proposed architecture, we compare STellar-FL with representative conventional federated learning frameworks. Existing systems typically adopt a centralized client–server design in which a central coordinator manages training rounds and model aggregation. While effective in controlled environments, such architectures introduce deployment challenges in institutional settings due to networking constraints, infrastructure requirements, and scalability limitations.

In contrast, STellar-FL adopts a decentralized architecture with relay-assisted communication and unified training orchestration. The design eliminates centralized infrastructure dependency and enables deployment in heterogeneous environments with strict network restrictions. The key architectural differences are summarized in Table 1.

 doi:10.20944/preprints202603.1776.v1



**Table 1.** Architectural comparison between conventional federated learning frameworks and STellar-FL.

| Aspect | Conventional Federated Learning | STellar-FL (Proposed) |
|---|---|---|
| **Architecture** | Centralized client–server | Decentralized |
| **Central coordinator** | Required | Not required |
| **Public IP / port configuration** | Typically required | Not required |
| **Network connectivity** | Fixed endpoints / VPN dependent | Relay-assisted peer connectivity |
| **Deployment complexity** | High | Low |
| **Scalability** | Limited by central server | Distributed scaling |
| **Single point of failure** | Yes | No |

The comparison highlights the fundamental architectural shift introduced by STellar-FL. Conventional federated learning frameworks rely on centralized coordination and fixed network configurations, which complicate deployment in real-world institutional environments. In contrast, STellar-FL removes the need for centralized infrastructure and supports dynamic peer connectivity through relay-assisted communication.

The decentralized design improves system robustness by eliminating single points of failure and enables scalable cross-institution collaboration without requiring complex network configuration, making the framework more suitable for privacy-sensitive and infrastructure-constrained environments.

## 4. LLM-Assisted Development and Extensibility

The unified interface and adapter-based design of the federated training orchestration layer (STefan-FL) create a natural extension point for large language model (LLM)–assisted development. Because model implementation and data loading are decoupled from the underlying FL infrastructure through well-defined contracts (e.g., model and dataloader callables), LLMs can generate or compose FL-capable Python code that conforms to these contracts without needing to handle framework-specific APIs or deployment details. Generated code can target training or inference workflows and can be validated in simulation mode before deployment, reducing the risk of runtime failures in distributed settings.

This extensibility is a direct consequence of the design choice to standardize the FL workflow: a single abstraction layer supports multiple backends (NVFlare, OpenFL, Flower) and exposes a consistent entry point for model and dataloader implementations. By keeping the surface area small and deterministic, the architecture lowers the barrier for automated or human-in-the-loop code generation. Domain experts and developers can describe desired behavior in natural language or high-level specifications, and LLM-generated scripts can be checked quickly via local or simulated FL runs before being used in cross-institution training.

The approach offers several advantages: faster prototyping of federated experiments, reduced boilerplate and integration effort, and the ability to reuse or adapt generated code across different FL backends and institutions. It also opens a path toward low-code and no-code federated learning. In a low-code setting, users might assemble training or inference pipelines from predefined blocks (model templates, data loaders, aggregation strategies) with minimal hand-written code. In a no-code setting, natural language or visual workflows could drive the generation of FL-capable artifacts that are then executed through the same unified orchestration and communication stack. Such directions could broaden participation in federated learning—including from clinical or operational staff who may not write code—while preserving the flexibility and control needed for research and production deployment.

 doi:10.20944/preprints202603.1776.v1



## 5. Discussion and Limitations

### *5.1. Practical Implications for Real-World Deployment*

The proposed STellar-FL architecture addresses key barriers that hinder the adoption of federated learning in real-world institutional environments. By eliminating the requirement for centralized coordination infrastructure, public network exposure, and complex network configuration, the framework enables practical deployment in settings with strict data governance and cybersecurity constraints. This is particularly relevant for healthcare applications, where data sovereignty, regulatory compliance, and institutional autonomy are critical considerations.

The decentralized communication mechanism and unified training workflow significantly reduce deployment complexity and enable scalable cross-institution collaboration. These characteristics make STellar-FL particularly suitable for multi-center medical AI development, where heterogeneous infrastructure and restricted network accessibility are common challenges.

Furthermore, the modular microservice-based design allows flexible integration with existing federated learning frameworks and computational environments, facilitating interoperability and future system extensions.

### *5.2. System-Level Trade-Offs*

Despite its advantages, the proposed architecture introduces several trade-offs. The relay-assisted communication mechanism may introduce additional routing overhead compared with direct client–server communication. While this overhead is generally modest, it may affect training efficiency in environments with extremely limited network bandwidth.

In addition, decentralized coordination may introduce increased system complexity in scheduling and synchronization compared with centralized aggregation. Careful system configuration and resource management are therefore required to ensure optimal performance in large-scale deployments.

### *5.3. Limitations*

This study focuses primarily on system architecture and deployment feasibility rather than algorithmic optimization or model performance improvement. The evaluation emphasizes infrastructure-level characteristics such as communication mechanisms and deployment flexibility, and does not provide extensive benchmarking across diverse machine learning tasks.

Furthermore, although the proposed architecture supports distributed model synchronization, the current implementation assumes reliable node participation during training. Future work may explore enhanced fault tolerance mechanisms and adaptive node management strategies for highly dynamic environments.

Finally, while the framework is motivated by healthcare applications, large-scale clinical validation and long-term operational evaluation remain future directions.

## 6. Conclusion and Future Directions

This work presents STellar-FL, a decentralized federated learning architecture designed for scalable and secure cross-institution model training. STellar-FL adopts a microservice-based design consisting of a federated training orchestration module, a distributed communication layer, and federated execution nodes. The framework enables secure model exchange through relay and proxy nodes, supports flexible environment management, and provides simulation and automated testing workflows for model development. These features collectively reduce deployment complexity while improving system robustness and scalability in heterogeneous institutional environments.

The main contributions of this work are summarized as follows:

- We propose STellar-FL, a decentralized federated learning architecture that enables cross-institution model training through distributed communication and peer-to-peer connectivity.

- We design a firewall-compatible communication mechanism that eliminates the need for centralized servers with public IP addresses or VPN-based connectivity.
- We introduce a unified federated learning workflow with modular system components to simplify deployment and environment management.
- We discuss how this design enables LLM-assisted generation of FL-capable training and inference code and outlines a path toward low-code and no-code federated learning.
- We demonstrate the feasibility of the proposed architecture in multi-institution medical AI scenarios.

The proposed framework provides a practical infrastructure foundation for real-world privacy-preserving AI deployment in healthcare and other regulated domains. By addressing key infrastructure barriers in federated learning deployment, STellar-FL enables scalable and flexible collaboration across institutions with heterogeneous computational and network environments.

Future work will explore communication optimization strategies, enhanced fault tolerance mechanisms, and large-scale deployment in multi-institution healthcare settings. Further research may also investigate integration with advanced privacy-preserving techniques, adaptive resource management and LLM-assisted as well as low-code and no-code federated learning workflows to improve system efficiency, usability and robustness in dynamic environments.

## Reference


1. Hunt, E. B. (2014). *Artificial intelligence*. Academic Press.
2. Al Kuwaiti, A., Nazer, K., Al-Reedy, A., Al-Shehri, S., Al-Muhanna, A., Subbarayalu, A. V., … Al-Muhanna, F. A. (2023). A review of the role of artificial intelligence in healthcare. *Journal of Personalized Medicine, 13*(6), 951. https://doi.org/10.3390/jpm13060951
3. Jaboob, A., Durrah, O., & Chakir, A. (2024). Artificial intelligence: An overview. *Engineering Applications of Artificial Intelligence*. https://doi.org/10.1007/978-3-031-50300-9_1
4. Oniani, D., Hilsman, J., Peng, Y., Poropatich, R. K., Pamplin, C. O. L., Legault, L. T. C., & Wang, Y. (2023). From military to healthcare: Adopting and expanding ethical principles for generative artificial intelligence. *arXiv*. https://arxiv.org/abs/2308.02448
5. Rashid, A. B., Kausik, A. K., Al Hassan Sunny, A., & Bappy, M. H. (2023). Artificial intelligence in the military: An overview of the capabilities, applications, and challenges. *International Journal of Intelligent Systems, 2023*, 8676366. https://doi.org/10.1155/2023/8676366
6. Yadav, N., Pandey, S., Gupta, A., Dudani, P., Gupta, S., & Rangarajan, K. (2023). Data privacy in healthcare: In the era of artificial intelligence. *Indian Dermatology Online Journal, 14*(6), 788–792. https://doi.org/10.4103/idoj.idoj_123_23
7. Li, L., Fan, Y., Tse, M., & Lin, K. Y. (2020). A review of applications in federated learning. *Computers & Industrial Engineering, 149*, 106854. https://doi.org/10.1016/j.cie.2020.106854
8. Zhang, C., Xie, Y., Bai, H., Yu, B., Li, W., & Gao, Y. (2021). A survey on federated learning. *Knowledge-Based Systems, 216*, 106775. https://doi.org/10.1016/j.knosys.2021.106775
9. Wen, J., Zhang, Z., Lan, Y., Cui, Z., Cai, J., & Zhang, W. (2023). A survey on federated learning: Challenges and applications. *International Journal of Machine Learning and Cybernetics, 14*(2), 513–535. https://doi.org/10.1007/s13042-022-01647-1
10. Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., … Cardoso, M. J. (2020). Federated learning for medical imaging: A review of the state of the art. *npj Digital Medicine, 3*, 1–12. https://doi.org/10.1038/s41746-020-00323-1
11. Kairouz, P., McMahan, H. B., Avent, B., et al. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*. https://doi.org/10.1561/2200000083
12. Moncada-Torres, A., Martin, F., Sieswerda, M., Van Soest, J., & Geleijnse, G. (2021). VANTAGE6: An open source privacy-preserving federated learning infrastructure for secure insight exchange. In *AMIA Annual Symposium Proceedings* (Vol. 2020, p. 870).
13. Beltrán, E. T. M., Pérez, M. Q., Sánchez, P. M. S., Bernal, S. L., Bovet, G., Pérez, M. G., … Celdrán, A. H. (2023). Decentralized federated learning: Fundamentals, state of the art, frameworks, trends, and

challenges. *IEEE Communications Surveys & Tutorials*, *25*(4), 2983–3013. https://doi.org/10.1109/COMST.2023.3292088
14. Yuan, L., Wang, Z., Sun, L., Yu, P. S., & Brinton, C. G. (2024). Decentralized federated learning: A survey and perspective. *IEEE Internet of Things Journal*, *11*(21), 34617–34638. https://doi.org/10.1109/JIOT.2024.3351234
15. McMahan, B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. Y. (2017). Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics* (pp. 1273–1282).
16. Roth, H. R., Cheng, Y., Wen, Y., Yang, I., Xu, Z., Hsieh, Y. T., … Feng, A. (2022). Nvidia FLARE: Federated learning from simulation to real-world. *arXiv*. https://arxiv.org/abs/2210.13291
17. Reina, G. A., Gruzdev, A., Foley, P., Perepelkina, O., Sharma, M., Davidyuk, I., … Bakas, S. (2021). OpenFL: An open-source framework for federated learning. *arXiv*. https://arxiv.org/abs/2105.06413
18. Beutel, D. J., Topal, T., Mathur, A., Qiu, X., Fernandez-Marques, J., Gao, Y., … Lane, N. D. (2020). Flower: A friendly federated learning research framework. *arXiv*. https://arxiv.org/abs/2007.14390
19. Roy, A. G., Siddiqui, S., Pölsterl, S., Navab, N., & Wachinger, C. (2019). BrainTorrent: A peer-to-peer environment for decentralized federated learning. *arXiv*. https://arxiv.org/abs/1905.06731
20. Wink, T., & Nochta, Z. (2021). An approach for peer-to-peer federated learning. In *2021 IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W)* (pp. 150–157). https://doi.org/10.1109/DSN-W52839.2021.00039
21. Qin, Z., Yan, X., Zhou, M., & Deng, S. (2024). BlockDFL: A blockchain-based fully decentralized peer-to-peer federated learning framework. In *Proceedings of the ACM Web Conference 2024* (pp. 2914–2925). https://doi.org/10.1145/3589334.3645477
22. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., & Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS)* (pp. 1175–1191). https://doi.org/10.1145/3133956.3133982
23. McMahan, H. B., Ramage, D., Talwar, K., & Zhang, L. (2018). Learning differentially private recurrent language models. In *International Conference on Learning Representations (ICLR)*.
24. Geyer, R. C., Klein, T., & Nabi, M. (2017). Differentially private federated learning: A client-level perspective. *arXiv*. https://arxiv.org/abs/1712.07557