

Federated Learning For Everyone: An accessible ecosystem for clinical research

Ashkan Pirmani, Martijn Oldenhof, Liesbet M. Peeters, Edward De Brouwer, Yves Moreau

Submitted to: JMIR Formative Research
on: December 15, 2023

Disclaimer: © The authors. All rights reserved. This is a privileged document currently under peer-review/community review. Authors have provided JMIR Publications with an exclusive license to publish this preprint on its website for review purposes only. While the final peer-reviewed paper may be licensed under a CC BY license on publication, at this stage authors and publisher expressly prohibit redistribution of this draft paper other than for review purposes.

Table of Contents

Original Manuscript..... 4

Supplementary Files..... 28

 Figures 29

 Figure 1..... 30

 Figure 2..... 31

 Figure 4..... 32

 Multimedia Appendixes 33

 Multimedia Appendix 0..... 34

Federated Learning For Everyone: An accessible ecosystem for clinical research

Ashkan Pirmani^{1, 2, 3, 4} MSc; Martijn Oldenhof² MSc; Liesbet M. Peeters^{1, 3, 4} PhD; Edward De Brouwer^{2*} PhD; Yves Moreau^{2*} PhD

¹Data Science Institute Hasselt University Diepenbeek BE

²ESAT-STADIUS KU Leuven Leuven BE

³University Multiple Sclerosis Center Hasselt University Diepenbeek BE

⁴Biomedical Research Institute Hasselt University Diepenbeek BE

*these authors contributed equally

Corresponding Author:

Yves Moreau PhD

ESAT-STADIUS

KU Leuven

Kasteelpark Arenberg 10

Leuven

BE

Abstract

The integrity and reliability of clinical research outcomes hinge on substantial amounts of data. However, their fragmented distribution across multiple contributors, coupled with ethical and regulatory difficulties, amplifies the challenges of accessing the most pertinent data. Despite federated learning's potential to harness insights from fragmented data, its adoption is hindered by implementation complexities, scalability issues, and inclusivity challenges. To navigate these, this paper presents the Federated Learning For Everyone (FL4E) framework, a versatile and accessible ecosystem that promotes multi-stakeholder clinical research collaboration. FL4E simplifies the intricacies of federated learning with a user-friendly, modular structure that underpins an ecosystem-based design. An innovative 'degree of federation' feature strikes a balance between centralized and federated learning, paving the way for a tailored application in various healthcare settings. Through rigorous testing on real-world healthcare datasets, we have demonstrated the efficacy of FL4E. Our findings show that the degree of federation's hybrid model achieves performance comparable to fully federated models without the associated significant overhead. Such results underscore the effectiveness of federation and suggest that FL4E has the potential to revolutionize collaborative clinical research by bridging the gap between centralized and federated approaches. The detailed implementation of FL4E, along with Docker configurations and related analyses, can be found on the corresponding GitHub page.

(JMIR Preprints 15/12/2023:55496)

DOI: <https://doi.org/10.2196/preprints.55496>

Preprint Settings

1) Would you like to publish your submitted manuscript as preprint?

Please make my preprint PDF available to anyone at any time (recommended).

Please make my preprint PDF available only to logged-in users; I understand that my title and abstract will remain visible to all users.

Only make the preprint title and abstract visible.

✓ **No, I do not wish to publish my submitted manuscript as a preprint.**

2) If accepted for publication in a JMIR journal, would you like the PDF to be visible to the public?

✓ **Yes, please make my accepted manuscript PDF available to anyone at any time (Recommended).**

Yes, but please make my accepted manuscript PDF available only to logged-in users; I understand that the title and abstract will remain visible to all users.

Yes, but only make the title and abstract visible (see Important note, above). I understand that if I later pay to participate in http://www.jmir.org/preprint/55496

Original Manuscript

Original Paper

Federated Learning For Everyone: An accessible ecosystem for clinical research

Ashkan Pirmani^{1,2,3,4}, Martijn Oldenhof¹, Liesbet M Peeters^{2,3,4}, Edward De Brouwer^{*1,5}, Yves Moreau^{*1}

¹KU Leuven, ESAT-STADIUS, Kasteelpark Arenberg 10, 3001 Leuven, Belgium.

²Data Science Institute, U Hasselt, Hasselt, Belgium.

³Biomedical Research Institute, U Hasselt, Hasselt, Belgium.

⁴University MS Center, U Hasselt, Hasselt, Belgium.

⁵Department of Genetics, Yale University, New Haven, USA.

yves.moreau@esat.kuleuven.be

*These authors contributed equally.

Abstract

Background: The integrity and reliability of clinical research outcomes rely heavily on access to vast amounts of data. However, the fragmented distribution of this data across multiple institutions, along with ethical and regulatory barriers, presents significant challenges to accessing relevant data. While federated learning offers a promising solution to leverage insights from fragmented datasets, its adoption faces hurdles due to implementation complexities, scalability issues, and inclusivity challenges.

Objective: This paper introduces the Federated Learning For Everyone (FL4E), an accessible framework facilitating multi-stakeholder collaboration in clinical research. It focuses on simplifying federated learning through an innovative ecosystem-based approach.

Methods: The ‘degree of federation’ is a fundamental concept of FL4E, allowing for flexible integration of federated and centralized learning models. This feature provides a customizable solution by enabling users to choose the level of data decentralization based on specific healthcare settings or project needs, making federated learning more adaptable and efficient. By using an ecosystem-based collaborative learning strategy, FL4E encourages a comprehensive platform for managing real-world data, enhancing collaboration and knowledge sharing among its stakeholders.

Results: Evaluating FL4E’s effectiveness using real-world healthcare datasets has highlighted its ecosystem-oriented and inclusive design. By applying hybrid models to two distinct analytical tasks—classification and survival analysis—within real-world settings, we have effectively measured the ‘degree of federation’ across various contexts. These evaluations show that FL4E’s hybrid models not only match the performance of fully federated but also avoid the substantial overhead usually linked with these models. Achieving this balance greatly enhances collaborative initiatives and broadens the scope of analytical possibilities within the ecosystem.

Conclusion: FL4E represents a significant step forward in collaborative clinical research by merging the benefits of centralized and federated learning. Its modular ecosystem-based design and the ‘degree of federation’ feature make it an inclusive, customizable framework suitable for a wide array of clinical research scenarios, promising to revolutionize the field through improved collaboration and data utilization. Detailed implementation and analyses are available at <https://github.com/ashkan-pirmani/FL4E-Analysis>.

Keywords: Federated Learning; Multi-stakeholder Collaboration; Real-World Data.

Introduction:

Real-world Data in Healthcare: Towards Federated Learning Solutions

Real-world data (RWD), which encompasses clinical and health data, is a fundamental asset in healthcare, providing invaluable insights for research and clinical applications [1], [2]. However, for privacy, legal, collaborative, or practical reasons, RWD is typically *fragmented* into different data sources, such as electronic health records, patient registries, and other clinical repositories. This fragmentation, characterized by the dispersion of RWD across numerous healthcare systems and geographic locations, hampers the Findability, Accessibility, Interoperability, and Reusability (FAIRness) of the RWD [3]. Thus, the capacity to harness these fragmented data sources for evidence-based healthcare decision-making is hindered and must be improved [4], [5].

While limited data sharing impedes the clinical research progress in general, the impact is particularly significant for research into low-prevalence diseases, where data scarcity already poses a challenge [6]. This issue becomes especially critical in urgent situations, like the COVID-19 pandemic [7], [8], where rapid data sharing is vital to quickly understand the impact of a novel disease.

Nevertheless, many challenges stemming from RWD fragmentation are systemic and can be seen as constraints that are difficult to eliminate. Legal data protection frameworks, ethical considerations, and organizational policies of different countries often limit the dissemination of such data both within and across countries [9], [10], [11], [12]. These challenges are further complicated by different data registries using different data formats, standards, and infrastructures to reflect their specific needs [13]. This context raises the need for a global data sharing framework to facilitate the secure and efficient sharing of RWD, while accommodating most of the complexities of the clinical data ecosystem.

A key component of such a global data sharing framework is Federated learning (FL), which offers a solution by training machine learning models on distributed datasets without centralization [14], [15]. Nevertheless, current FL packages crucial for implementing this system fail to cater to the multifaceted requirements of healthcare stakeholders.

Related Works

The FL framework ecosystem is characterized by its heterogeneity, prompted by contributions from both the open-source and industry communities, with each framework casted to address specific facets of FL's inherent challenges. From enhancing security and privacy to ensuring scalability and promoting framework agnosticism to facilitating ease of use and customization, these frameworks collectively foster a conducive environment for the flourishing of FL applications.

Industry solutions like IBM FL [16], the Clara Training Framework [17], and Sherpa.ai [18] offer specialized features but suffer from restricted accessibility due to their proprietary nature, limiting widespread application and innovation. Open-source FL frameworks, despite their appeal of broader accessibility, grapple with challenges unique to their design and intended use cases. Syft also known as PySyft [19] in the context of offering advanced privacy-enhancing technologies, focuses on secure and private data analysis across various domains. Despite it emphasizes structured transparency systems, the complex setup process associated with Syft could impede its broad adoption. As a pioneering framework in the realm of privacy technology, Syft has undergone substantial revisions across its iterations. While aimed at improvement, these continuous overhauls may pose challenges

for early adopters in terms of staying abreast with the framework's evolution.

Frameworks such as FATE [20] and OpenFL [21] are distinguished by their strong emphasis on security and privacy, which are critical in sector like healthcare, where data confidentiality is essential. FATE is noted for its robust security and privacy features. However, it requires substantial infrastructural investment, and a steep learning curve, as there are some discrepancies in documentation clarity and availability across different parts of the framework. OpenFL also offers Intel-enabled trusted execution environment known as Safe Guard and mutual TLS for all communications to ensure maximum privacy. However, implementing this framework on a broader scale is hindered and may necessitate specific hardware, software, or the activation of certain services [22]. TensorFlow Federated [23], while offering a specialized toolkit for TensorFlow models, limits its utility to those committed to a singular modeling framework, potentially sidelining a broader audience that operates across diverse machine learning frameworks.

Flower [24] and FedLab [24] are recognized for their framework agnosticism and scalability. Flower emerges as a user-friendly and lightweight option, albeit still progressing through its developmental stages with real-world applications yet to be fully demonstrated. This reflects a common developmental trajectory among open-source FL frameworks, where achieving a balance between accessibility, ease of use, and broad applicability poses a significant challenge.

FedML [25] distinguishes itself through its versatility and recent architectural improvements, positioning it alongside frameworks like Plato [26], known for supporting a diverse array of computing paradigms. This dual focus on facilitating FL research and the development of practical applications underscores its strengths. However, integrating significant updates may necessitate substantial maintenance for applications developed on its platform. Additionally, developers may need to reacquaint themselves with the framework's updated features, presenting a potential challenge in adaptability and ongoing engagement.

Innovative approaches are evident in PaddleFL [27] and Galaxy FL [28], with Galaxy FL pioneering blockchain technology integration for improved data privacy and model confidentiality. This innovation broadens the FL framework landscape, introducing novel data protection and utilization aspects. However, blending FL with blockchain remains a debated topic within the literature, especially concerning the significant computation and communication overheads [29].

Lastly, FLSim [30] specializes in FL simulations, providing a comprehensive framework that incorporates differential privacy and secure aggregation scenarios. Similarly, Leaf [31] focuses primarily on simulation studies. While this specialization in simulation offers valuable insights, it may also narrow their applicability, potentially restricting their direct translation into real-world deployments.

For a more in-depth discussion of the FL frameworks, we guide readers to the appendix. There, Table Multimedia Appendix 1 offers a comprehensive summary of the highlights and key features of various FL frameworks. Additionally, Table Multimedia Appendix 2 details the developmental progress and recent advancements within these frameworks, providing a clear snapshot of their evolution and current state.

Navigating the Path Forward: the need for an accessible ecosystem for clinical research

Beyond the challenges already covered, a recent review of FL in the biomedical domain reveals that only some research studies have provided reusable or reproducible source code alongside their research [32]. This highlights a common issue: the tendency to start a unique project for each separate data sharing initiative. This practice not only consumes valuable time and resources but also hinders the evolution of a universal ecosystem. Furthermore, the opportunity to build upon and contribute to an existing initiative is often overlooked.

Despite recent progress in the field, there is a critical demand for a more inclusive, user-friendly, and holistic framework in healthcare collaborative research [33], [34]. A framework that empowers stakeholders from diverse backgrounds to fully engage in collaborative learning opportunities is essential. This requires FL solutions characterized by intuitive interfaces, streamlined workflows, comprehensive documentation, and a design process that actively seeks and incorporates feedback from a wide range of stakeholders. By prioritizing usability and inclusivity in developing FL solutions, we can significantly lower technical barriers. This approach enables wider participation and amplifies the contributions of collaborative healthcare research.

In this work, we present Federated Learning for Everyone (FL4E)—a dynamic, data-centric framework designed with inclusivity and adaptability at its core. FL4E is dedicated to efficiently converting RWD into actionable evidence, facilitating engagement from a broad spectrum of stakeholders, and enhancing collaborative learning experiences. Its hierarchical architecture offers a customizable data sharing ecosystem, facilitating a tunable balance between centralization and federation. With its modular approach, FL4E is general and provides an ecosystem-based environment that can support a wide range of collaborative research projects. This modularity further allows all stakeholders to contribute to the project by selectively engaging in the relevant part of the data sharing pipeline.

Key contributions: (1) We propose a unique FL framework designed to cater to the diverse needs of healthcare stakeholders and to adapt to various study designs. This design is further enhanced by offering a tunable balance between centralization and federation. (2) We formalize this balance by introducing the concept of ‘*degree of federation*’. (3) Our evaluation of the framework, using two real-world datasets, underscores the value of federation. Specifically, hybrid experiments conducted at different degrees of federation deliver performance comparable to fully federated setups.

Methods

In the methods section of this study, we rigorously investigate the fundamental principles underlying FL4E, systematically articulating its core concepts. This entails a meticulous examination of the pivotal modules and essential components that constitute FL4E, with a particular focus on their functionalities and interrelationships. Moreover, this section highlights the architectural nuances of FL4E, comprehensively detailing its structural configuration and the dynamics of its operational framework.

Key concepts of FL4E:

1) The concept of ‘degree of federation’ within the FL4E framework introduces a novel, adaptable, and pragmatic approach to collaborative data analysis designed to meet stakeholders’ varied needs and limitations. This feature enables participants to customize their level of involvement and data sharing based on their specific infrastructural capacities, regulatory requirements, and privacy considerations.

FL4E acknowledges that while federated models offer significant privacy advantages, their practicality may vary due to different constraints. Therefore, FL4E is designed to accommodate various levels of participation for stakeholders, ensuring flexibility if one data sharing stream proves unfeasible.

Therefore, the choice between a centralized or federated model within FL4E is presented not as a

dichotomy but as a continuum. On the one hand, centralization may simplify operations and enhance efficiency, but it could compromise privacy to some degree. Conversely, full federation prioritizes privacy, possibly adding operational and analytical complexity. FL4E's design empowers stakeholders to find a balance that best suits their needs, governance frameworks, and regulatory environments.

This adaptable strategy has proven effective in accommodating diverse data sharing needs, as demonstrated by the Global Data Sharing Initiative for Multiple Sclerosis and COVID-19 [35], [36]. The deployment of a versatile hybrid data acquisition system, merging traditional data sharing with federated methods, has not only broadened inclusivity and flexibility but also led to a notable 45.47% increase in data collection volume compared to initial efforts [37]. These results underscore the 'degree of federation' practical value, showcasing the advantages of integrating such adaptability into data acquisition and sharing frameworks. Different configurations within the degree of federation are depicted in Figure 1.

However, it is essential to emphasize that the framework fosters regulatory compliance across all participation methods. Whether stakeholders prefer the simplicity and direct oversight of a more centralized model or the enhanced privacy of a federated approach, stakeholders are reminded of their collective obligation to adhere to prevailing data protection regulations.

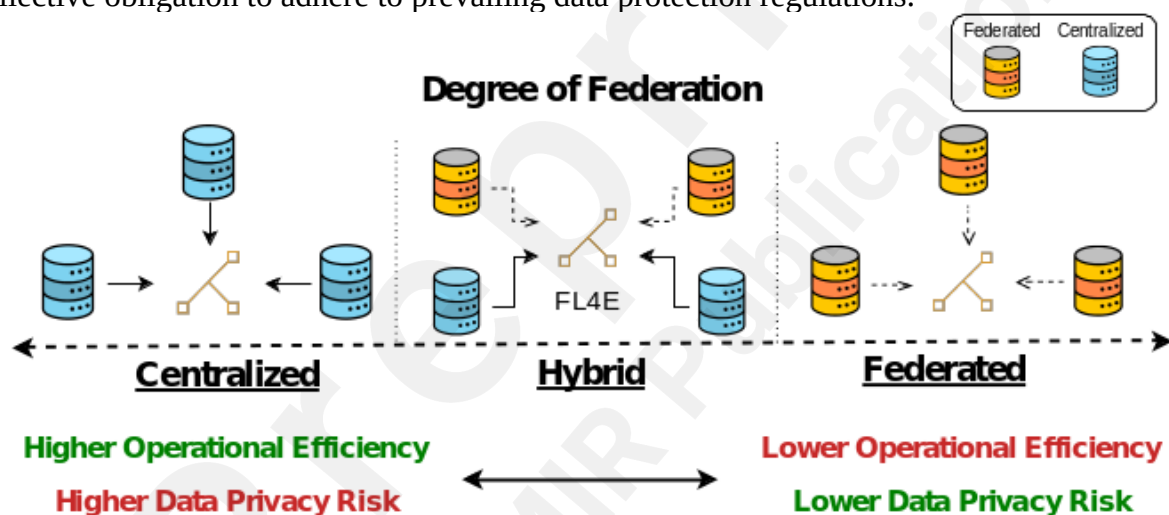


Figure 1 - The degree of federation characterizes the balance between fully centralized and fully federated setups, leading to hybrid solutions where some of the stakeholders centralize their data while others prefer a federated approach

2) Ecosystem-based collaborative learning: FL4E framework introduces a paradigm shift from conventional project-specific FL to an 'ecosystem-based' approach, underscoring a strategic pivot towards creating a unified, adaptable platform tailored for conducting a various analyses centered on specific diseases or health conditions. At the heart of the 'ecosystem-based' methodology is its holistic perspective, aiming to formulate a comprehensive and synergistic environment for data sharing and collaborative research.

An ecosystem, in this context, is envisioned not just as a collection of data or series of disjointed projects but as a dynamic, collaborative network of stakeholders. This includes researchers, healthcare professionals, patients, and policy-makers, all united by a common goal: understanding and addressing a particular health issue. By fostering such an ecosystem, FL4E encourages the pooling of diverse resources, expertise, and perspectives. This collaborative efficiency breaks down traditional silos, allowing for more efficient utilization of resources through shared data and findings.

Moreover, encompassing a wide array of stakeholders and research questions enables the ecosystem to provide a more nuanced and comprehensive understanding of the health condition under

investigation. The diversity of perspectives and expertise enriches the research outcomes, offering holistic insights far beyond what could be achieved in isolation.

This approach is inherently adaptable and capable of evolving in response to new research findings, emerging health challenges, and shifts in the regulatory landscape. Such flexibility ensures the long-term sustainability and relevance of the ecosystem, enabling it to continue generating valuable insights over time.

Regarding its structure, FL4E features a nested design. Within each ecosystem, several studies are carried out, each focusing on specific research questions related to the overarching theme. These studies, while operating under their unique regulatory and consent protocols, incorporate the degree of federation. They can include a dataset, data quality scripts, and analyses tailored to their specific research queries.

FL4E's structure also incorporates simplicity and accessibility in its design. The intuitive user interface encourages contributions from stakeholders, fostering a collaborative environment that is integral to the ecosystem's success. This systematic approach to collaboration and data sharing is visually depicted in Figure 2.

The FL4E framework builds upon the foundational concepts of the Flower framework [38]. Flower's flexibility, support for multi-node execution, and agnosticism to both ML frameworks and programming languages significantly contribute to the overall adaptability and versatility of FL4E.

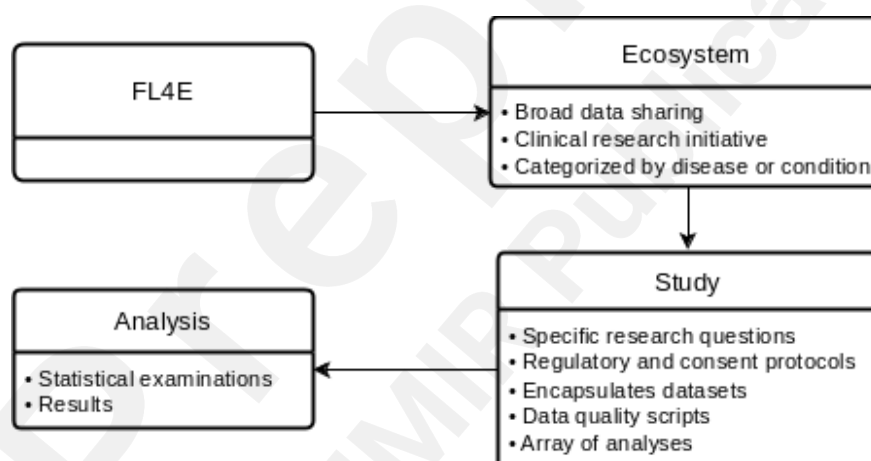


Figure 2 - FL4E's Nested Approach Schema: The schema presents the ecosystem approach of the FL4E, facilitating the incorporation of diverse studies framed by different research questions. It further highlights the support for a wide range of analyses, leading to the execution of statistical experiments.

Modules and functionality:

This section provides an in-depth overview of the core modules within the FL4E framework: the Study Center, Repository Center, Model Center, and Data Center. Each module's role and anticipated functionality are detailed, showcasing how the division into distinct centers contributes to effectively deploying the overarching ecosystem. Figure 3 complements this by visually representing the architectural arrangement of server and client components. By detailing the anticipated interactions among stakeholders and illustrating the operational flow within the framework, the architecture provides a solid foundation for understanding the potential and practical application of FL4E in advancing healthcare data-driven research and insights.

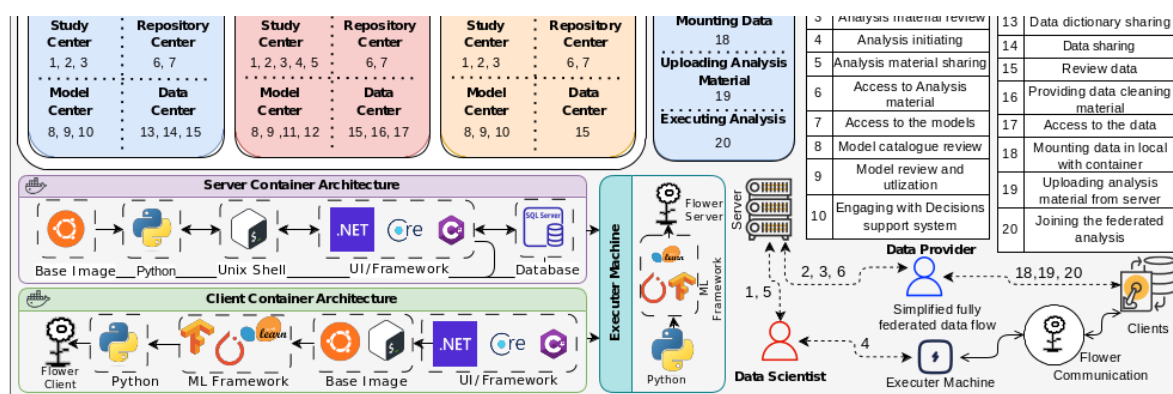


Figure 3 - The High-level Architecture of FL4E: This figure showcases a comprehensive framework designed to accommodate the nuanced interactions of diverse stakeholders. The architecture includes the expected user stories across three primary categories of participants: data providers, data scientists, and downstream users. Each stakeholder group engages with FL4E platform through distinct pathways. The architecture diagram features a table outlining the unique interactions of the stakeholders, mapping their respective roles and activities within the FL4E framework. This detailed mapping clearly explains how each stakeholder contributes to and benefits from the FL4E, highlighting the platform's versatility and user-centric design. At the core of our architecture lie three fundamental components: the server, client, and executor machine, each vital to the execution of FL tasks. The diagram we provide elucidates their interconnected roles, showcasing the seamless flow of data, scripts, and analytical results across the system. The server acts as the orchestrator for FL tasks, hosting the primary web application within a Docker container as an ASP.NET application. It securely manages the platform's data, housed on an SQL server in a dedicated hosting environment. This component is crucial for coordinating tasks and employs a Python environment to manage secure data sharing and preprocessing of 'data center' module of the framework. On the client side, implemented as a Docker-based image, it runs on the data contributor's machine. This component is essential for integrating RWD into the FL process. Developed using Python and ASP.NET for web applications, the client-side component establishes a connection to the executor machine. On the other hand, the executor machine plays a crucial role in conducting the analysis. It is designed to receive clients updates. This adaptable component configuration allows data scientists to tailor it according to their specific analytical needs and preferences.

Study Center:

The Study Center is a critical piece in the ecosystem design of the framework. It coordinates research activities via the Study Catalogue and the Analysis Center, highlighting the interconnectivity and synergy within the system.

The Study Catalogue functions as a cooperative repository where various stakeholders can share and access overarching information or metadata about different studies. Collaborations can be initiated by a data scientist with an intriguing research question but needing more data or by a data provider possessing the required data but lacking the analytical tools. Downstream users can further enrich the catalogue by contributing scientific articles or proposing research questions relevant to other stakeholders. Any new entry to the catalogue should include the title, a brief description, the study lead or author, and the current status of the entry (*i.e.*, active or not). If the entry involves any data, this should also be noted, as the data would be subsequently managed within the Data Center.

In contrast, the Analysis Center is primarily designed for data scientists. Here, they can create an analysis entry derived from a study catalogue record. While doing so, they must furnish comprehensive details, including the exact analysis title, a descriptive summary, and the privacy status of the analysis (private or public). In a private setting, data scientists can selectively grant access to specific users (data providers). This means only authorized users can access these entries.

In the FL4E framework, data scientists play a crucial role in orchestrating the analysis by providing essential scripts for client-side execution. Given that the FL4E backend is built upon the Flower framework, these scripts must align with this framework's requirements. The ML-agnostic nature of this foundation offers a substantial flexibility to data scientists, allowing them to craft these scripts

within the ambit of their expertise and the specific needs of the analysis. This flexibility is further augmented by Flower's support for secure channels, enabling the addition of privacy-preserving configurations to the analysis execution. The integration of Flower's capabilities within FL4E not only simplifies the technical implementation of FL projects but also empowers data scientists to innovate and experiment with various analytical models and privacy configurations. This approach ensures that FL4E remains a versatile and secure platform for FL, accommodating a wide range of research objectives and data privacy standards.

To ensure transparency, the process of sharing scripts among the stakeholders is not automated. Clients may review and, if agreed, run these scripts locally using the FL4E-client component. In keeping with the principles of modularity and customization, the server does not include a built-in compiler to execute the analysis. Data scientists should disclose the IP and port address of the designated machine to data providers for analysis execution, enabling FL process participation. It is vital to note that no analysis should proceed in the Analysis Center without explicit consent from participating centers. If a particular study is associated with any data on central server, data scientists should incorporate the data as a centralized client into the federated analysis, effectively acting as an additional presumptive federated client in the data scientist's machine.

Repository Center:

The Repository Center serves as a warehouse for all scripts and models essential for stakeholders' research and collaboration. It serves both as a repository to disseminate the scripts and models to the relevant stakeholders, and as an archive that stores past analyses and provides examples facilitating the creation of new studies. It seamlessly integrates with the Study Center and Model Center, enabling direct and organized access to all shared scripts and models across these centers.

Model Center:

The Model Center is designed to facilitate the coordination and sharing of trained models. Similar to the Study Center, the Model Center is structured around two main elements: the Model Catalogue and the Model Repository.

The Model Catalogue serves as a database for high-level information about various trained models. This can include the title, description, type of model, study lead, and status. The Model Catalogue is particularly useful for downstream users and data scientists who come across an interesting model and wish to share it within the ecosystem, or for those who want to share the outcome of an analysis conducted within the Study Center.

Furthermore, the Model Repository is where detailed information about each trained model is stored. This includes the relationship of the model to a specific Model Catalogue entry or Analysis Center entry. Like the Study Center, the Model Center also adopts authorization layers to ensure that only authorized users can access certain models. A typical scenario is a trained model that would only be shared among the participants that contributed with data.

However, if needed, the Model Center can also be used to disseminate a model more broadly. Indeed, users can attach trained models to the results of an analysis. This means that a user who may not have been able to join an analysis can still benefit from the trained models, download them, and test them on their local data. This is especially useful for data providers who want to validate models against their datasets or for data scientists who wish to use them as pre-trained models in a similar domain.

Furthermore, the Model Center allows data scientists to provide a decision support system as a result of a trained model. They can inject an interactive HTML web page into the Model Repository entry,

serving as an interactive dashboard. This user-friendly interface enables downstream users, even those without expert knowledge, to engage with the system effectively.

Data Center:

The Data Center plays a crucial role in actualizing the concept of the “degree of federation” within the FL4E framework. Beyond providing data sharing capabilities, it also enables data integration and aggregation, meaning this module stores the data in a central server. Data providers and data scientists are central to the sharing process within the Data Center, employing a meticulous three-step routine to ensure data quality, along with effective data sharing and aggregation. The first step is (1) the raw data sharing scheme, followed by (2) the data cleaning and enhancement, and (3) the data sharing and aggregation.

Raw Data Sharing Scheme: The process initiates when the data provider shares a raw data scheme via a data dictionary, a static file describing the data structure. This dictionary must be attached to an existing study catalogue entry along with details like title, description, and data type (public or private). For private data, the provider must list authorized users.

Data Cleaning and Enhancement Script: Once the data scheme is established, data scientists must create a Python script for data cleaning and enhancement. This script ensures data quality by handling duplicates, missing values, and validating value ranges based on information of the data schema. This script should be uploaded to the Data Center and linked to the data dictionary.

Data Sharing and Aggregation: After setting up the data dictionary and cleaning script, data sharing can occur. Each data share triggers the cleaning script execution, with the outcomes appended to a core dictionary and each upload saved as a distinct file. This setup permits various data providers to share data in line with the data dictionary, with the script guaranteeing aggregation into a single file. Data scientists can then access this aggregated file for analytical use if access is granted to them.

Implementation of the framework:

The FL4E framework is divided into server-side and client-side implementations. Server-side, the framework is a microservice-based web application, designed using ASP.NET Core [39]. It adheres to the N-tier architecture and employs the Model-View-Controller design, along with a repository pattern, leading to a scalable and maintainable application. The application is segmented into four layers: Application, Model, Utility, and Data Access. Each layer performs distinct tasks, such as managing user authorization, defining data structures, handling requests, and interacting with the database. ASP.NET Core Identity module is utilized for authentication and authorization, while an admin control panel supervises user registrations in the interface layer of the framework [39].

Client-side, the framework aims for simplicity and ease of use, also leveraging the ASP.NET Core framework. It incorporates three primary interfaces for data mounting, script uploading, and script execution. The data mounting interface facilitates user selection and mounting of data files, the script uploading interface enables users to upload scripts acquired from the server, and the script execution interface allows for the running of these scripts, using the Flower framework to secure the communication between the client and server.

Ensuring the seamless operation and accessibility of the FL4E platform, both the client and server components are encapsulated as Docker containers. This strategic choice underpins the platform's commitment to ease of deployment, consistency across different environments, and portability, thereby significantly enhancing the user experience for participants in FL projects.

A fully functional prototype of the FL4E server component has been successfully deployed to

Microsoft Azure, showcasing the platform's robustness [40]. This deployment not only demonstrates FL4E's capability to leverage cloud infrastructure but also its readiness to support FL projects.

To further acquaint potential users and stakeholders with FL4E's functionalities and user interface, a comprehensive video demonstration has been made available [41]. This visual guide serves as an invaluable resource for understanding the platform's operational flow, features, and the simplicity of initiating and participating in FL studies.

For those interested in a deeper exploration or customization of FL4E, the source code for both the server and client components is readily accessible [42]. This openness ensures transparency, fosters a community of collaboration, and enables continuous improvement of the platform by allowing researchers, developers, and data scientists to contribute to its development.

Ethical considerations:

This study employs secondary analyses of openly licensed repository datasets for demonstration purposes in the results section. No sensitive data were used for the co-creation of the framework. All dataset sources are cited and adhere to open licensing agreements. For more detailed information about the datasets, please refer to the original sources.

Results

Datasets:

We evaluated our framework using two real-world clinical datasets from Flamby [43]: Fed-Heart-Disease and Fed-Tcga-Brca. The Fed-Heart-Disease dataset contains 740 records from four centers, detailing 13 clinical features and a binary heart disease indicator for each patient. We employed a logistic regression model for prediction. The Fed-Tcga-Brca dataset, sourced from The Cancer Genome Atlas's Genomics Data Commons portal, has data from 1,088 breast cancer patients across six centers. With 39 clinical features and each patient's time of death, we used a Cox model to predict the risk of death.

Additionally, the Wisconsin Breast Cancer dataset was incorporated into the GitHub repository as an extra use case [42], [44]. It serves as a set of foundational scripts that can be generalized and act as guidelines for hybrid experiment settings in the FL4E framework.

Degree of federation scenarios:

To demonstrate how FL4E adapts to various research questions, and how it modulates the degree of federation, we conducted three experiments using these datasets:

Fully federated experiment: In this scenario, every client participates in the FL process with each one contributing to the global model using their respective local datasets. This scenario aligns with the conventional FL setup. The data flow for this experiment is depicted in Figure 3.

Hybrid experiment: In the hybrid setup, while all clients are part of the training process, some of them contribute their data centrally to the server. For instance, with the Fed-Heart-Disease dataset, Clients 1 and 2 were randomly selected to centrally contribute their data, while the rest participated as federated clients. For the Fed-Tcga-Brca dataset, Clients 2, 3, and 5 joined as centralized clients.

Centralized experiment: For this experiment, all clients send their data to the central server, and the model is trained centrally. This design aims to provide a benchmark for comparing the performance

of the federated and hybrid approaches with the centralized model.

Experimental Findings

We evaluated our FL models using Area Under the Receiver Operating Characteristic Curve (ROC-AUC) and accuracy metrics for the Fed-Heart-Disease dataset and the concordance index (C-index) for the Fed-Tcga-Brca dataset. These metrics were chosen according to their ability to assess model performance in binary classification and survival analysis task, while accounting for potential class imbalance. For federated experiments, we implemented multiple FL strategies: FedAVG [14], FedOpt [45] (including FedAdagrad and FedYogi), and FedProx [46]. Hyperparameter tuning was performed using grid search on a centralized setting and applied to each experiment of the datasets. To ensure consistency, each experiment was repeated five times. The results of these experiments are summarized in Table 1.

Table 1 - Experiments Carried Out Using FL4E With the Fed-Heart-Disease And Fed-Tcga-Brca Datasets, Demonstrating Adaptability To Varying Degrees Of Federation.

		Fed-Heart-Disease		Fed-Tcga-Brca
		ROC-AUC	Accuracy	C-Index
Fully Federated	FedAvg	0.846 ± 0.002	0.733 ± 0.007	0.732 ± 0.030
	FedAdagrad	0.841 ± 0.029	0.726 ± 0.05	0.748 ± 0.016
	FedYogi	0.803 ± 0.051	0.715 ± 0.032	0.745 ± 0.037
	FedProx	0.846 ± 0.003	0.741 ± 0.006	0.725 ± 0.007
Hybrid Experiment	FedAvg	0.825 ± 0.004	0.740 ± 0.054	0.656 ± 0.06
	FedAdagrad	0.821 ± 0.008	0.741 ± 0.005	0.776 ± 0.036
	FedYogi	0.794 ± 0.013	0.710 ± 0.039	0.726 ± 0.041
	FedProx	0.822 ± 0.004	0.737 ± 0.012	0.439 ± 0.227
Local Experiment	Client 0	0.842 ± 0.009	0.753 ± 0.011	0.668 ± 0.064
	Client 1	0.882 ± 0.007	0.800 ± 0.013	0.445 ± 0.237
	Client 2	0.546 ± 0.271	0.550 ± 0.199	0.635 ± 0.166
	Client 3	0.542 ± 0.054	0.559 ± 0.096	0.570 ± 0.140
	Client 4	n/a		0.851 ± 0.078
	Client 5	n/a		0.666 ± 0.001
	Client 1,2	0.819 ± 0.003	0.752 ± 0.008	n/a
	Client 2,3 & 5	n/a		0.578 ± 0.125
Centralized		0.812 ± 0.003	0.753 ± 0.007	1.6090.207

In our analysis of the Fed-Heart-Disease dataset, as illustrated in Figure 4 and detailed in Table 1, notable behavior patterns emerges from Clients 2 and 3. These patterns indicate the effects of class imbalance and quantity skew, suggesting that smaller clients with imbalanced datasets may benefit more from participating in an FL environment than operating in isolation. This insight is harmonious with the foundational principles of the FL paradigm, which aspires to harness the potential of diverse data sources to enhance model robustness and generalizability.

To clarify the evaluation process, we introduce the notation where K represents the total number of clients participating in the FL process, and n_i denotes the size of the dataset for the i -th client, where i ranges from 1 to K . Furthermore, E_i signifies the evaluation metric (such as accuracy or ROC-AUC) obtained from assessing the global model by the i -th client. The sum of the dataset sizes across all clients is represented by N , calculated as $N = \sum_{i=1}^K n_i$. Therefore, the overall evaluation metric

E for the global FL model is calculated using the formula $E = \frac{1}{N} \sum_{i=1}^K (n_i E_i)$, offering a comprehensive measure of performance across the federated network.

This approach significantly impacts the evaluation process by prioritizing the characteristics of individual clients over the generalizability of the FL model. This provides a more comprehensive understanding of model performance across diverse environments. However, the potential for skewness in federated evaluation, often stemming from imbalanced clients, arises from its reliance on a weighted average of individual client metrics. This introduces challenges, particularly when directly comparing it to centralized performance.

It is important to note that for a direct comparison between federated and centralized evaluations, the ideal scenario involves assessing the FL model using a separate global test set on the server side. While this approach aligns more closely with the methodology used in centralized evaluations, it is not always feasible, especially in real-world settings where maintaining a separate test set on the server can be impractical due to data governance and privacy concerns.

Contrary to initial expectations that centralized models would outperform others, our analysis unveiled that federated and hybrid models exhibit superior performance based on the ROC-AUC measure, even though the accuracy metrics might suggested a different narrative. To ensure the reliability and consistency of our findings, we have conducted rigorous comparisons with the original benchmark [43]. These comparisons demonstrate that our observations, especially concerning the performance metrics of centralized and federated models such as accuracy, are in line with existing research. This alignment reinforces the validity of our analysis despite the inherent challenges in conducting direct comparisons between federated and centralized performance evaluations [43].




Figure 4 - Delineating Client Heterogeneity in Real-World Applications: Exploration of class imbalance in the Fed-Heart-Disease analysis.

Similarly, the Fed-Tcga-BrcA study substantiates the advantage of FL. The heterogeneity of clients in the dataset analysis, as highlighted by Flamby [43], indicates that this diversity significantly influences training performance, underscoring the potential benefits of FL.

Notably, the hybrid experiment yields a crucial insight: it offers performance that matches fully federated models while significantly outperforming local settings for smaller clients. Importantly, this is achieved with considerably less complexity and overhead associated with managing federation across multiple data centers. This setup enables centralized management of in-house data, providing analysts with direct access (within regulatory boundaries) and flexibility for advanced analyses. By having data information at their disposal, analysts can develop more effective scripts, enabling advanced analyses with greater flexibility.

Discussion:

FL4E presents a novel approach to collaborative learning in a healthcare context, by innovatively balancing centralization and federation. This allows healthcare organizations to engage in FL according to their infrastructure and policy capabilities. Composed of the Study Center, Repository Center, Model Center, and Data Center, FL4E's modular, ecosystem-oriented design caters to a wide variety of stakeholders, thereby facilitating multi-party clinical research collaborations.

However, FL4E is not alone in promoting deployment ease and inclusivity in FL environments. FeatureCloud [47] advances FL by offering a platform that simplifies deployment across healthcare and research sectors. Aimed at democratizing multi-institutional analyses, FeatureCloud ensures that advanced programming skills or complex setups do not impede participation. The AI Store module, facilitating algorithm sharing and access, highlights FeatureCloud's commitment to broadening FL adoption.

In parallel, the emergence of UniFed [48] underscores a similar commitment to simplifying FL's application by unifying diverse open-source frameworks. UniFed's approach standardizes FL experimentation and deployment, offering a configuration-based, schema-enforced task specification that eases the distributed execution of FL projects. By supporting various FL frameworks, UniFed addresses the challenges of workflow diversity, data format discrepancies, and interpretation variations, thereby streamlining the FL experimentation process.

What distinguishes FL4E in the evolving landscape is its focus on the 'degree of federation' and the adaptability it brings, ensuring the framework is inclusive, adaptable and can host hybrid experiments. Unlike FeatureCloud, which was developed from scratch, FL4E builds on the well-regarded Flower framework, providing a robust, accessible foundation and utilizing this base, an adaptive, ecosystem-based model, designed to cater to diverse stakeholder preferences and requirements. FL4E extends the democratization vision of FL, incorporating a modular, customizable design that enhances engagement and ensures versatility across various applications. Additionally, FL4E adheres to key coding principles such as N-tier architecture and the repository pattern. These design choices promote ease of use, flexibility for modifications, and future-proofing the framework, further distinguishing FL4E by making it more accessible, maintainable, and adaptable to advancements in FL.

As the FL field continues to mature, platforms like FeatureCloud and UniFed play crucial roles in lowering technical barriers and fostering an environment conducive to collaborative innovation. Their contributions, along with FL4E's nuanced understanding of stakeholder engagement, collectively propel the FL ecosystem towards a future where data-driven insights are accessible to a broader spectrum of the healthcare and research community, thereby amplifying the impact of FL in advancing healthcare outcomes and research endeavors.

Nonetheless, FL4E still presents some limitations. Version control can pose challenges due to the Flower package dependency, and significant version changes may require internal updates. Moving

from technical concerns to regulatory issues, ensuring adherence to regional data privacy regulations necessitates diligence from initiators and study leads, who must establish proper agreements for data sharing. Additionally, the framework's dependence on manual communication among stakeholders could potentially impede its universal accessibility. This last point is particularly critical, as the system's effective functioning heavily relies on the presence of a pre-established network of collaborators where members have a certain level of familiarity or knowledge of each other.

Although FL4E and other FL systems offer potential technical solutions to collaborative learning in healthcare, the influence of societal barriers should not be underestimated. Hence, continuous education, stakeholder engagement, and proof-of-concept demonstrations within controlled environments are crucial for successful deployment.

Future development path to FL4E, the demand for advanced privacy measures becomes increasingly apparent. While FL introduces a novel approach to leveraging patient-level data across decentralized sources, it is essential to recognize that this method is not without its vulnerabilities [49], [50], [51]. The process of sharing impersonal statistics during FL can inadvertently encode sensitive information, thus introducing potential risks and challenges to data privacy [49], [52], [53].

The architecture of FL, which is designed to enhance privacy by eliminating the need for direct data sharing, is nevertheless susceptible to sophisticated attacks and privacy breaches. These vulnerabilities include inference attacks, where malicious entities may attempt to reconstruct personal data from model updates, as well as direct cyber threats that target the integrity and confidentiality of the FL process [49]. The body of literature on FL privacy highlights the critical importance of acknowledging these risks, with studies identifying various attack vectors that could potentially compromise the privacy guarantees of federated model sharing [53].

Therefore, despite FL's significant reduction in the risk of direct data exposure, it introduces new complexities in safeguarding against indirect data inference and cyber-attacks. This nuanced understanding necessitates a cautious approach in the implementation of FL, particularly when handling sensitive healthcare information. The design of FL workflows, including the development of infrastructure and study plans, must incorporate robust security measures and privacy-preserving techniques.

To this end, it is crucial to leverage the latest advancements in secure computing, differential privacy, and encryption technologies to strengthen the FL process against potential vulnerabilities [54], [55], [56]. Acknowledging these risks and adequately preparing for them is paramount, ensuring that the advantages of utilizing decentralized data through FL do not come at the expense of patient privacy or data security. The implementation of industry-scale privacy measures is critically important in FL applications, as emphasized by recent studies [55]. However, these measures must be carefully balanced to ensure that data protection does not compromise the computational efficiency of FL analysis.

To ease deployment, recent research in practical FL points towards promising techniques, such as soft computing's fusion with FL [57], and the use of FL in cloud-edge collaborative architectures [58]. Such strategies can address concerns over data privacy, high communication costs, and data management difficulties, offering a promising direction for future investigations.

While the manual steps (for instance, script exchanges) in FL4E are designed for transparency, automating these steps could improve the framework's adaptability and usability. Additionally, transitioning from the current ASP.Net to Python could amplify the framework's appeal, especially given Python's strong foothold in the data science realm. As healthcare data expands in volume and complexity, frameworks like FL4E, with further advancements and community involvement, stand to

be vital tools in catalyzing transformative progress and collaboration in the domain.

Acknowledgements

YM, MO and AP are funded by (1) VLAIO PM: Augmenting Therapeutic Effectiveness through Novel Analytics (HBC.2019.2528); Research Council KU Leuven: Symbiosis 4 (C14/22/125), Symbiosis3 (C14/18/092; CELSA - Active Learning (CELSA/21/019). YM, MO and AP are affiliated to Leuven.AI and received funding from the Flemish Government (AI Research Program, FWO SBO (S003422N), and ELIXIR Belgium (I002819N)). EDB was funded by a FWO-SB grant.

Data Availability

The datasets used for the analysis purposes of this study are available at the FLamby repository on GitHub [59]. FLamby serves as a benchmark for cross-silo datasets for FL applications, with a focus on healthcare application [43].

Abbreviations

RWD: Real-world data

FAIR : Findable, Accessible, Interoperable, and Reusable

FL4E : Federated Learning for Everyone

FL: Federated Learning

C-index: concordance index

ROC-AUC: Area Under the Receiver Operating Characteristic Curve

References

- [1] R. E. Sherman *et al.*, “Real-World Evidence - What Is It and What Can It Tell Us?,” *N Engl J Med*, vol. 375, no. 23, pp. 2293–2297, Dec. 2016, doi: 10.1056/NEJMSB1609216.
- [2] J. Wu, J. Roy, and W. F. Stewart, “Prediction modeling using EHR data: challenges, strategies, and a comparison of machine learning approaches,” *Med Care*, vol. 48, no. 6 Suppl, Jun. 2010, doi: 10.1097/MLR.0B013E3181DE9E17.
- [3] M. D. Wilkinson *et al.*, “The FAIR Guiding Principles for scientific data management and stewardship,” *Scientific Data* 3:1, vol. 3, no. 1, pp. 1–9, Mar. 2016, doi: 10.1038/sdata.2016.18.
- [4] P. B. Jensen, L. J. Jensen, and S. Brunak, “Mining electronic health records: towards better research applications and clinical care,” *Nat Rev Genet*, vol. 13, no. 6, pp. 395–405, Jun. 2012, doi: 10.1038/NRG3208.
- [5] N. G. Weiskopf and C. Weng, “Methods and dimensions of electronic health record data quality assessment: enabling reuse for clinical research,” *J Am Med Inform Assoc*, vol. 20, no. 1, p. 144, 2013, doi: 10.1136/AMIAJNL-2011-000681.
- [6] Y. Kodra *et al.*, “Recommendations for Improving the Quality of Rare Disease Registries,” *Int J Environ Res Public Health*, vol. 15, no. 8, p. 1644, Aug. 2018, doi: 10.3390/IJERPH15081644.
- [7] V. Moorthy, A. M. H. Restrepo, M. P. Preziosi, and S. Swaminathan, “Data sharing for novel coronavirus (COVID-19),” *Bull World Health Organ*, vol. 98, no. 3, p. 150, Mar. 2020, doi: 10.2471/BLT.20.251561.
- [8] L. M. Peeters *et al.*, “COVID-19 in people with multiple sclerosis: A global data sharing initiative,” *Multiple Sclerosis Journal*, vol. 26, no. 10, pp. 1157–1162, Sep. 2020, doi: 10.1177/1352458520941485/ASSET/IMAGES/LARGE/10.1177_1352458520941485-FIG1.JPEG.
- [9] M. Aitken, J. De St Jorre, C. Pagliari, R. Jepson, and S. Cunningham-Burley, “Public responses to the sharing and linkage of health data for research purposes: a systematic review and thematic

- synthesis of qualitative studies,” *BMC Med Ethics*, vol. 17, no. 1, pp. 1–24, Nov. 2016, doi: 10.1186/S12910-016-0153-X/TABLES/4.
- [10] S. Holm and T. Ploug, “Big Data and Health Research—The Governance Challenges in a Mixed Data Economy,” *J Bioeth Inq*, vol. 14, no. 4, pp. 515–525, Dec. 2017, doi: 10.1007/S11673-017-9810-0/METRICS.
- [11] D. Taruscio, E. Mollo, S. Gainotti, M. P. de la Paz, F. Bianchi, and L. Vittozzi, “The EPIRARE proposal of a set of indicators and common data elements for the European platform for rare disease registration,” *Arch Public Health*, vol. 72, no. 1, 2014, doi: 10.1186/2049-3258-72-35.
- [12] N. A. Dreyer *et al.*, “Direct-to-Patient Research: Piloting a New Approach to Understanding Drug Safety During Pregnancy,” *JMIR Public Health Surveill*, vol. 1, no. 2, p. e22, Dec. 2015, doi: 10.2196/publichealth.4939.
- [13] M. G. Kahn *et al.*, “A Harmonized Data Quality Assessment Terminology and Framework for the Secondary Use of Electronic Health Record Data,” *eGEMS*, vol. 4, no. 1, p. 18, Sep. 2016, doi: 10.13063/2327-9214.1244.
- [14] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, “Communication-Efficient Learning of Deep Networks from Decentralized Data.” PMLR, pp. 1273–1282, Apr. 10, 2017. Accessed: Dec. 15, 2023. [Online]. Available: <https://proceedings.mlr.press/v54/mcmahan17a.html>
- [15] J. Konečný, H. Brendan McMahan, F. X. Yu, A. Theertha Suresh, D. Bacon Google, and P. Richtárik, “Federated Learning: Strategies for Improving Communication Efficiency,” Oct. 2016, Accessed: Dec. 15, 2023. [Online]. Available: <https://arxiv.org/abs/1610.05492v2>
- [16] H. Ludwig *et al.*, “IBM Federated Learning: an Enterprise Framework White Paper V0.1,” Jul. 2020, Accessed: Dec. 15, 2023. [Online]. Available: <https://arxiv.org/abs/2007.10987v1>
- [17] “Federated Learning powered by NVIDIA Clara | NVIDIA Technical Blog.” Accessed: Dec. 15, 2023. [Online]. Available: <https://developer.nvidia.com/blog/federated-learning-clara/>
- [18] N. Rodríguez-Barroso *et al.*, “Federated Learning and Differential Privacy: Software tools analysis, the Sherpa.ai FL framework and methodological guidelines for preserving data privacy,” *Information Fusion*, vol. 64, pp. 270–292, Dec. 2020, doi: 10.1016/J.INFFUS.2020.07.009.
- [19] A. Ziller *et al.*, “PySyft: A Library for Easy Federated Learning,” *Studies in Computational Intelligence*, vol. 965, pp. 111–139, 2021, doi: 10.1007/978-3-030-70604-3_5/FIGURES/7.
- [20] LiuYang, FanTao, ChenTianjian, XuQian, and YangQiang, “FATE,” *The Journal of Machine Learning Research*, vol. 22, pp. 1–6, Jan. 2021, doi: 10.5555/3546258.3546484.
- [21] G. A. Reina *et al.*, “OpenFL: An open-source framework for Federated Learning,” *Phys Med Biol*, vol. 67, no. 21, May 2021, doi: 10.1088/1361-6560/ac97d9.
- [22] “Try Federated Learning with OpenFL.” Accessed: Mar. 18, 2024. [Online]. Available: <https://www.intel.com/content/www/us/en/developer/articles/technical/try-federated-learning-with-openfl.html>
- [23] “TensorFlow Federated.” Accessed: Dec. 15, 2023. [Online]. Available: <https://www.tensorflow.org/federated>
- [24] D. Zeng, S. Liang, X. Hu, H. Wang, and Z. Xu, “FedLab: A Flexible Federated Learning Framework,” Jul. 2021, Accessed: Mar. 18, 2024. [Online]. Available: <https://arxiv.org/abs/2107.11621v4>
- [25] C. He *et al.*, “FedML: A Research Library and Benchmark for Federated Machine Learning,” Jul. 2020, Accessed: Mar. 18, 2024. [Online]. Available: <https://arxiv.org/abs/2007.13518v4>
- [26] B. Li, N. Su, C. Ying, and F. Wang, “Plato: An Open-Source Research Framework for Production Federated Learning,” *Proceedings of ACM Turing Award Celebration Conference, CHINA 2023*, pp. 1–2, Jul. 2023, doi: 10.1145/3603165.3607364.
- [27] “GitHub - PaddlePaddle/PaddleFL: Federated Deep Learning in PaddlePaddle.” Accessed: Mar. 18, 2024. [Online]. Available: <https://github.com/PaddlePaddle/PaddleFL>
- [28] L. Liu, C. Wu, and J. Xiao, “Blockchain-Based platform for Distribution AI,” Feb. 2019, doi: 10.29007/4BZ1.

- [29] Y. Qu, M. P. Uddin, C. Gan, Y. Xiang, L. Gao, and J. Yearwood, "Blockchain-enabled Federated Learning: A Survey," *ACM Comput Surv*, vol. 55, no. 4, p. 70, Nov. 2022, doi: 10.1145/3524104.
- [30] L. Li, J. Wang, and C. Z. Xu, "FLSim: An Extensible and Reusable Simulation Framework for Federated Learning," *Lecture Notes of the Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering, LNICST*, vol. 369, pp. 350–369, 2021, doi: 10.1007/978-3-030-72792-5_30/TABLES/8.
- [31] S. Caldas *et al.*, "LEAF: A Benchmark for Federated Settings," Dec. 2018, Accessed: Dec. 15, 2023. [Online]. Available: <https://arxiv.org/abs/1812.01097v3>
- [32] M. G. Crowsonid *et al.*, "A systematic review of federated learning applications for biomedical data," *PLOS Digital Health*, vol. 1, no. 5, p. e0000033, May 2022, doi: 10.1371/JOURNAL.PDIG.0000033.
- [33] V. A. Rudrapatna and A. J. Butte, "Opportunities and challenges in using real-world data for health care," *J Clin Invest*, vol. 130, no. 2, pp. 565–574, 2020, doi: 10.1172/JCI129197.
- [34] K. Hiramatsu, A. Barrett, and Y. Miyata, "Current Status, Challenges, and Future Perspectives of Real-World Data and Real-World Evidence in Japan," *Drugs Real World Outcomes*, vol. 8, no. 4, pp. 459–480, Dec. 2021, doi: 10.1007/S40801-021-00266-3/TABLES/8.
- [35] S. Simpson-Yap *et al.*, "Associations of Disease-Modifying Therapies With COVID-19 Severity in Multiple Sclerosis," *Neurology*, vol. 97, no. 19, pp. E1870–E1885, Nov. 2021, doi: 10.1212/WNL.00000000000012753.
- [36] S. Simpson-Yap *et al.*, "Updated Results of the COVID-19 in MS Global Data Sharing Initiative: Anti-CD20 and Other Risk Factors Associated With COVID-19 Severity," *Neurol Neuroimmunol Neuroinflamm*, vol. 9, no. 6, Nov. 2022, doi: 10.1212/NXI.0000000000200021/ASSET/54BBC39A-4F08-43E1-8192-463361034550/ASSETS/GRAPHIC/NXI.0000000000200021TU1.GIF.
- [37] A. Pirmani, E. De Brouwer, L. Geys, T. Parciak, Y. Moreau, and L. M. Peeters, "The Journey of Data Within a Global Data Sharing Initiative: A Federated 3-Layer Data Analysis Pipeline to Scale Up Multiple Sclerosis Research.," *JMIR Med Inform*, vol. 11, no. 1, p. e48030, Nov. 2023, doi: 10.2196/48030.
- [38] D. J. Beutel *et al.*, "Flower: A Friendly Federated Learning Research Framework," Jul. 2020, Accessed: Dec. 15, 2023. [Online]. Available: <https://arxiv.org/abs/2007.14390v5>
- [39] "ASP.NET Core | Open-source web framework for .NET." Accessed: Dec. 15, 2023. [Online]. Available: <https://dotnet.microsoft.com/en-us/apps/aspnet>
- [40] "FL4E-Working Prototype." Accessed: Dec. 15, 2023. [Online]. Available: <https://fl-server.azurewebsites.net/>
- [41] "Federated Learning For Everyone Demo - YouTube." Accessed: May 14, 2024. [Online]. Available: https://www.youtube.com/watch?v=K8n7kffpvQ0&ab_channel=AshkanPirmani
- [42] "GitHub - ashkan-pirmani/FL4E-Analysis: FL4E-Analysis." Accessed: Dec. 15, 2023. [Online]. Available: <https://github.com/ashkan-pirmani/FL4E-Analysis>
- [43] J. O. du Terrail *et al.*, "FLamby: Datasets and Benchmarks for Cross-Silo Federated Learning in Realistic Healthcare Settings," *Adv Neural Inf Process Syst*, vol. 35, Oct. 2022, Accessed: Dec. 15, 2023. [Online]. Available: <https://arxiv.org/abs/2210.04620v3>
- [44] "Breast Cancer Wisconsin (Prognostic) - UCI Machine Learning Repository." Accessed: Dec. 15, 2023. [Online]. Available: <https://archive.ics.uci.edu/dataset/16/breast+cancer+wisconsin+prognostic>
- [45] S. J. Reddi *et al.*, "Adaptive Federated Optimization," *ICLR 2021 - 9th International Conference on Learning Representations*, Feb. 2020, Accessed: Dec. 15, 2023. [Online]. Available: <https://arxiv.org/abs/2003.00295v5>
- [46] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated Optimization in Heterogeneous Networks," Dec. 2018, Accessed: Dec. 15, 2023. [Online]. Available: <https://arxiv.org/abs/1812.06127v5>
- [47] J. Matschinske *et al.*, "The FeatureCloud Platform for Federated Learning in Biomedicine: Unified Approach," *J Med Internet Res*, vol. 25, no. 1, p. e42621, Jul. 2023, doi: 10.2196/42621.

- [48] X. Liu *et al.*, “UniFed: All-In-One Federated Learning Platform to Unify Open-Source Frameworks,” Jul. 2022, Accessed: Mar. 07, 2024. [Online]. Available: <https://arxiv.org/abs/2207.10308v3>
- [49] P. Kairouz *et al.*, “Advances and open problems in federated learning,” *Foundations and Trends in Machine Learning*, vol. 14, no. 1–2, pp. 1–210, Jun. 2021, doi: 10.1561/22000000083.
- [50] R. Nasirigerdeh, R. Torkzadehmahani, J. Baumbach, and D. B. Blumenthal, “On the Privacy of Federated Pipelines,” *SIGIR 2021 - Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 1975–1979, Jul. 2021, doi: 10.1145/3404835.3462996.
- [51] A. Hartebrodt, R. Röttger, and D. B. Blumenthal, “Federated singular value decomposition for high-dimensional data,” *Data Min Knowl Discov*, pp. 1–38, Oct. 2023, doi: 10.1007/S10618-023-00983-Z/TABLES/6.
- [52] V. Tolpegin, S. Truex, M. E. Gursoy, and L. Liu, “Data poisoning attacks against federated learning systems,” *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 12308 LNCS, pp. 480–501, 2020, doi: 10.1007/978-3-030-58951-6_24/TABLES/6.
- [53] E. Bagdasaryan, A. Veit, Y. Hua, D. Estrin, and V. Shmatikov, “How To Backdoor Federated Learning,” *Proc Mach Learn Res*, vol. 108, pp. 2938–2948, Jul. 2018, Accessed: Mar. 07, 2024. [Online]. Available: <https://arxiv.org/abs/1807.00459v3>
- [54] C. Dwork, “Differential privacy,” *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 4052 LNCS, pp. 1–12, 2006, doi: 10.1007/11787006_1/COVER.
- [55] R. Lozi, W. Puech, J. Park, and H. Lim, “Privacy-Preserving Federated Learning Using Homomorphic Encryption,” *Applied Sciences 2022, Vol. 12, Page 734*, vol. 12, no. 2, p. 734, Jan. 2022, doi: 10.3390/APP12020734.
- [56] A. Wood, K. Najarian, and D. Kahrobaei, “Homomorphic Encryption for Machine Learning in Medicine and Bioinformatics,” *ACM Computing Surveys (CSUR)*, vol. 53, no. 4, Aug. 2020, doi: 10.1145/3394658.
- [57] Y. Supriya and T. R. Gadekallu, “A Survey on Soft Computing Techniques for Federated Learning-Applications, Challenges and Future Directions,” *ACM J Data Inf Qual*, vol. 15, no. 2, Jun. 2023, doi: 10.1145/3575810.
- [58] Y. Huang *et al.*, “A High-Precision Method for 100-Day-Old Classification of Chickens in Edge Computing Scenarios Based on Federated Computing,” *Animals 2022, Vol. 12, Page 3450*, vol. 12, no. 24, p. 3450, Dec. 2022, doi: 10.3390/ANI12243450.
- [59] “GitHub - owkin/FLamby: Cross-silo Federated Learning playground in Python. Discover 7 real-world federated datasets to test your new FL strategies and try to beat the leaderboard.” Accessed: May 14, 2024. [Online]. Available: <https://github.com/owkin/FLamby>
- [60] “PySyft - OpenMined Blog.” Accessed: Mar. 18, 2024. [Online]. Available: <https://blog.openmined.org/tag/pysyft/>
- [61] A. Ziller *et al.*, “PySyft: A Library for Easy Federated Learning,” *Studies in Computational Intelligence*, vol. 965, pp. 111–139, 2021, doi: 10.1007/978-3-030-70604-3_5/FIGURES/7.
- [62] “Fate.” Accessed: Mar. 18, 2024. [Online]. Available: <https://fate.fedai.org/>
- [63] “FATE-Community/FATE_Project_Technical_Charter.pdf at master · FederatedAI/FATE-Community · GitHub.” Accessed: Mar. 18, 2024. [Online]. Available: https://github.com/FederatedAI/FATE-Community/blob/master/FATE_Project_Technical_Charter.pdf
- [64] “GitHub - FederatedAI/FATE: An Industrial Grade Federated Learning Framework.” Accessed: Mar. 18, 2024. [Online]. Available: <https://github.com/FederatedAI/FATE>
- [65] “FedML: A Research Library and Benchmark for Federated Machine Learning — MIT Media Lab.” Accessed: Mar. 18, 2024. [Online]. Available: <https://www.media.mit.edu/publications/fedml-a-research-library-and-benchmark-for-federated-machine-learning/>
- [66] “FEDML - Your Generative AI Platform at Scale.” Accessed: Mar. 18, 2024. [Online]. Available:

<https://www.fedml.ai/home>

- [67] D. J. Beutel *et al.*, “Flower: A Friendly Federated Learning Research Framework,” Jul. 2020, Accessed: Mar. 18, 2024. [Online]. Available: <https://arxiv.org/abs/2007.14390v5>
- [68] “Flower: A Friendly Federated Learning Framework.” Accessed: Mar. 18, 2024. [Online]. Available: <https://flower.ai/>
- [69] Y. Xie *et al.*, “FederatedScope: A Flexible Federated Learning Platform for Heterogeneity”, Accessed: Apr. 24, 2024. [Online]. Available: <https://github.com/alibaba/FederatedScope>
- [70] “OpenFL – Linux Foundation Project.” Accessed: Mar. 18, 2024. [Online]. Available: <https://openfl.io/>
- [71] “GitHub - SMILELab-FL/FedLab: A flexible Federated Learning Framework based on PyTorch, simplifying your Federated Learning research.” Accessed: Mar. 18, 2024. [Online]. Available: <https://github.com/SMILELab-FL/FedLab>
- [72] D. Stripelis, C. Anastasiou, P. Toral, A. Asghar, and J. L. Ambite, “MetisFL: An Embarrassingly Parallelized Controller for Scalable & Efficient Federated Learning Workflows KEYWORDS federated learning; distributed machine learning; machine learning systems; data management”, doi: 10.1145/3630048.3630186.
- [73] H. R. Roth *et al.*, “NVIDIA FLARE: Federated Learning from Simulation to Real-World,” Oct. 2022, doi: 10.48550/arXiv.2210.13291.
- [74] “PaddleFL — PaddleFL 0.1.0.beta documentation.” Accessed: Mar. 18, 2024. [Online]. Available: <https://paddlefl.readthedocs.io/en/latest/>
- [75] F. Lai *et al.*, “FedScale: Benchmarking Model and System Performance of Federated Learning at Scale,” *Proc Mach Learn Res*, vol. 162, pp. 11814–11827, May 2021, Accessed: Apr. 24, 2024. [Online]. Available: <https://arxiv.org/abs/2105.11367v5>
- [76] “GitHub - TL-System/plato: A federated learning framework to support scalable and reproducible research.” Accessed: Mar. 18, 2024. [Online]. Available: <https://github.com/TL-System/plato>
- [77] W. Zhuang, X. Gan, Y. Wen, and S. Zhang, “EasyFL: A Low-code Federated Learning Platform For Dummies,” *IEEE Internet Things J*, vol. 9, no. 15, pp. 13740–13754, May 2021, doi: 10.1109/JIOT.2022.3143842.
- [78] “GitHub - GalaxyLearning/GFL: Galaxy Federated Learning Framework.” Accessed: Mar. 18, 2024. [Online]. Available: <https://github.com/GalaxyLearning/GFL>
- [79] “GitHub - facebookresearch/FLSim: Federated Learning Simulator (FLSim) is a flexible, standalone core library that simulates FL settings with a minimal, easy-to-use API. FLSim is domain-agnostic and accommodates many use cases such as vision and text.” Accessed: Mar. 18, 2024. [Online]. Available: <https://github.com/facebookresearch/FLSim>

Multimedia Appendix A

A rich diversity marks the landscape of FL frameworks, each offering a unique integration of features tailored to address the multifaceted challenges of distributed, privacy-preserving machine learning. A methodical investigation uncovers a range of design principles, varying from those placing a high priority on security and privacy to frameworks that emphasize scalability, platform neutrality, and user-friendliness. This analysis endeavors to cover this landscape, highlighting the intrinsic attributes, potential limitations, and the distinct application domains of these frameworks, thereby providing a comprehensive overview.

Multimedia Appendix 1, offers an insightful comparison across various FL frameworks, shedding light on their distinctive features, limitations, key application areas, and their status as either open-source or proprietary tools. This comparative analysis reveals the rich diversity within the FL ecosystem, demonstrating how each framework is uniquely positioned to address specific needs and

challenges.

This compilation highlights the inherent trade-offs involved when selecting an FL framework, such as balancing ease of use against computational demands or gauging the level of community support available.

Multimedia Appendix 1 - Comparison of Federated Learning Frameworks: An Overview of Features, Limitations, and Application Domains

Framework	Core/Key Features	Limitations/Remarks	Specific Use Cases/Applications
PySyft [60], [61]	- Enhances privacy in data processing, combining PETs for structured transparency systems. - Supports homomorphically encrypted operations and split neural networks for inference.	- Aims to enable secure and private data analysis across domains, introducing Duet for privacy through encryption and controlled data access.	Secure and private data analysis across various domains, supporting structured transparency in data flows.
FATE [62], [63], [64]	- Industrial-grade, supports various FL scenarios and algorithms. - Modular, scalable, with a strong emphasis on security and privacy through homomorphic encryption and MPC. - Backed by a robust community and ecosystem.	- Complexity and steep learning curve for new users. - Deployment challenges, especially in large-scale environments. - Significant resource requirements.	Suitable for industries where data security and privacy are paramount, such as finance and healthcare.
FedML [25], [65], [66]	- Versatile, supporting a wide range of domains including NLP, CV, and IoT. - Flexible and customizable architecture with comprehensive benchmarks. - Promotes open and collaborative development.	- Complexity for beginners. - Resource intensity for full utilization. - Dependent on active community engagement.	Broad spectrum of FL research and applications, from edge devices to cloud-based environments.
Flower [67], [68]	- Highly customizable and extendable. Framework agnostic, supporting integration with popular ML libraries. - Scalable, tested with over 10,000 clients. - Focuses on both academic research and practical application.	- Potential integration complexities and scalability challenges in very large setups not specifically mentioned.	Projects requiring integration with various ML frameworks and scalability, from mobile to edge devices.
TensorFlow Federated[23]	- Two-layer architecture for custom federated algorithms. - Encourages research and experimentation. - Support for non-learning computations for federated analytics.	- Requires familiarity with TensorFlow for optimal use. - Design choices may limit flexibility with non-TensorFlow technologies.	Developing privacy-preserving models like mobile keyboard prediction without central data collection.
FederatedScope[69]	- Modular and extensible design, supporting a broad range of FL applications. - Unique event-driven architecture for flexible and powerful asynchronous federated training. - Comprehensive documentation and tutorials.	- Ease of use for new users enhanced by extensive documentation, though specific challenges are not mentioned.	Varied FL applications, from CV and NLP to graph FL.
OpenFL [21], [22], [70]	- Scalable, capable of managing large federations. - High emphasis on security, employing Intel-enabled TEEs. - Designed for ease of use with support for popular ML frameworks.	- Setup and familiarity with security features required for optimal use.	Expansive FL projects requiring high security, such as healthcare data analysis.
FedLab [24], [71]	- Modular design, supporting diverse data distributions and federated algorithms. - Comprehensive baselines and data partitions for research.	- Effectiveness in highly heterogeneous environments or under specific operational constraints could vary.	Research focusing on FL under realistic conditions, addressing IID and non-IID settings.

Framework	Core/Key Features	Limitations/Remarks	Specific Use Cases/Applications
MetisFL[72]	<ul style="list-style-type: none"> - Adheres to principles of modularity, extensibility, and configurability. - Supports TensorFlow and PyTorch, extendable to other backends. - Application domains span CV, NLP, and more, reflecting its general-purpose, domain-agnostic nature. 	<ul style="list-style-type: none"> - Challenges with very specific or niche requirements might arise, not explicitly mentioned. 	General-purpose applications across various domains such as CV and NLP.
NVFlare[73]	<ul style="list-style-type: none"> - Focuses on privacy preservation. - Supports popular ML/DL frameworks. - Provides tools for secure provisioning, orchestration, and monitoring. 	<ul style="list-style-type: none"> - Implementation challenges in domain-specific applications might be encountered. 	Secure, privacy-preserving collaborations in AI research fields across distributed setups.
IBM FL[16]	<ul style="list-style-type: none"> - Highly configurable and extensible, supporting a wide range of ML techniques. - Includes a comprehensive library of fusion algorithms. 	<ul style="list-style-type: none"> - Significant setup and configuration might be required. - Broad applicability with potential for complex deployments, though specific limitations are not discussed. 	Enterprise environments with needs for data privacy in supervised, unsupervised, and reinforcement learning projects.
PaddleFL [27], [74]	<ul style="list-style-type: none"> - Supports horizontal and vertical FL strategies. - Leverages PaddlePaddle for distributed training and Kubernetes for deployment. 	<ul style="list-style-type: none"> - Preparing to extend support for vertical FL and deployment schemes, indicating ongoing development and optimization. 	Applications in CV, NLP, and recommendation systems, addressing data isolation and secure data knowledge sharing.
FedScale[75]	<ul style="list-style-type: none"> - Offers realistic FL tasks and scalable runtime with diverse datasets. - Focuses on benchmarking model and system performance at scale. 	<ul style="list-style-type: none"> - Aims to address client heterogeneity and large-scale FL deployment challenges, reflecting a comprehensive approach to FL benchmarking. 	Facilitates reproducible FL research across domains like image classification and speech recognition, aiming for heterogeneity-aware optimizations.
Plato [26], [76]	<ul style="list-style-type: none"> - Designed for scalable and reproducible FL research. - Supports a wide range of FL algorithms, scalable, and extensible. 	<ul style="list-style-type: none"> - Focuses on scalability, extensibility, real-world system implementations, and reproducibility, suggesting broad applicability. 	Scalable and reproducible FL research, supporting deployment in various environments.
EasyFL[77]	<ul style="list-style-type: none"> - Low-code platform for FL experimentation and prototyping. - Supports heterogeneity simulation, training flow abstraction, and distributed training optimization. 	<ul style="list-style-type: none"> - Emphasizes ease of use, deployability, and efficiency in FL experimentation and development, suitable for users with varying expertise levels. 	User-friendly experimentation and prototyping in FL, suitable for research and industrial application prototypes.
Galaxy Learning [28], Federated [78]	<ul style="list-style-type: none"> - Decentralized FL framework based on blockchain technology. - Designed to train models with distributed data while maintaining data ownership and model interest. 	<ul style="list-style-type: none"> - Aims at improving network bandwidth utilization and maintaining security under malicious node attacks, suggesting a novel approach to FL. 	Addressing data privacy, high training costs, and ownership issues in AI development through a decentralized approach.
FLSim [30], [79]	<ul style="list-style-type: none"> - Domain-agnostic, scalable, and supports differential privacy, secure aggregation, and compression techniques. - Designed for simulating FL environments efficiently. 	<ul style="list-style-type: none"> - Provides a flexible and comprehensive simulation framework for exploring various aspects of FL. 	Efficiently building simulators for FL, suitable for research and development in cross-device FL.

Based on the Table Multimedia Appendix 2, a systematic analysis reveals insightful patterns regarding the popularity, development activity, and maintenance status of the existing FL frameworks. This summary is grounded in the provided metrics: stars (indicating popularity), forks (suggesting community involvement and usage), commits (reflecting development activity), releases (showing versioning and updates), contributors (indicating collaborative effort), and maintenance status.

PySyft emerges as the most popular framework with 9.2K stars, attributed to its comprehensive approach to privacy-preserving machine learning. FATE and FedML also exhibit strong popularity, with 5.4K and 4K stars, respectively, indicating their significant impact and recognition within the community.

Regarding community engagement, PySyft leads with 2K forks, followed closely by FATE with 1.5K forks. This suggests a high level of interest and collaboration in these projects. Notably, TensorFlow Federated and FedML also show considerable engagement with 572 and 754 forks, respectively.

In terms of development activity, PySyft stands out with a remarkable 27,664 commits, indicating a highly active and evolving project. FATE follows with 13,775 commits, underscoring its robust development effort.

The contributor metric further illustrates the community and collaborative aspects of these projects. PySyft leads significantly with 421 contributors, reflecting a broad and diverse development community. Flower and TensorFlow Federated also show strong collaboration, with 120 and 106 contributors, respectively.

The Apache 2.0 license predominates among the frameworks, favored for its permissive nature, encouraging broad adoption and contribution. Exceptions include MetisFL and FLSim, which are under the BSD License, and IBM FL, which utilizes a specific license.

Maintenance status reveals a mixed landscape. While most projects are actively maintained, a few, such as FederatedScope, IBM FL, PaddleFL, and GalaxyFederatedLearning, are not recently maintained, raising considerations for potential adopters regarding long-term support and updates.

Multimedia Appendix 2 - Development Landscape of Federated Learning Frameworks

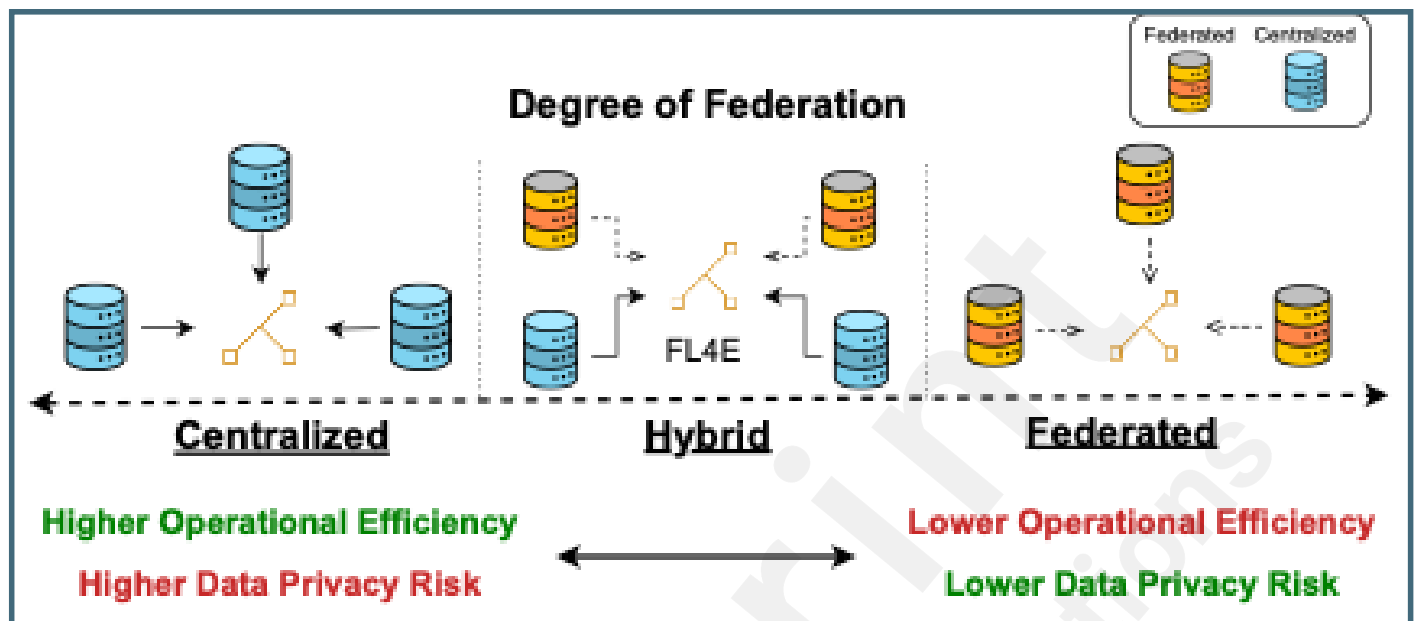
Name	License	Developed By	Stars	Forks	Commits	Releases	Contributors	Maintenance
PySyft	Apache 2.0	OpenMined	9.2K	2K	27,664	111	421	Recently Maintained
FATE	Apache 2.0	Webank	5.4K	1.5K	13,775	48	86	Recently Maintained
FedML	Apache 2.0	FedML	4K	754	11,726	6	66	Recently Maintained
Flower	Apache 2.0	Adap	3.9K	714	1,986	17	120	Recently Maintained
TensorFlow Federated	Apache 2.0	Google	2.3K	572	5,191	77	106	Recently Maintained
FederatedScope	Apache 2.0	Alibaba	1.2K	192	544	3	18	Not Recently Maintained
OpenFL	Apache 2.0	Intel	644	172	606	9	77	Recently Maintained
FedLab	Apache 2.0	Smilelab	633	120	1,422	9	13	Recently Maintained
MetisFL	BSD License	BioInt	530	49	380	0	9	Recently Maintained
NVFlare	Apache 2.0	Nvidia	508	133	1,532	43	34	Recently Maintained
IBM FL	Specific License	IBM	473	135	178	10	11	Not Recently Maintained
PaddleFL	Apache 2.0	PaddlePaddle	488	122	735	5	25	Not Recently Maintained
FedScale	Apache 2.0	SymbioTech	359	114	707	1	21	Recently Maintained
Plato	Apache 2.0	Theory Lab	304	71	1,280	20	24	Recently Maintained

Name	License	Developed By	Stars	Forks	Commits	Releases	Contributors	Maintenance
EasyFL	Apache 2.0	FLGO	409	78	1,015	8	9	Recently Maintained
Galaxy Federated Learning	Apache 2.0	GalaxyLearning	235	58	118	0	3	Not Recently Maintained
FLSim	BSD Clause License	Facebook Research	229	47	175	3	24	Recently Maintained

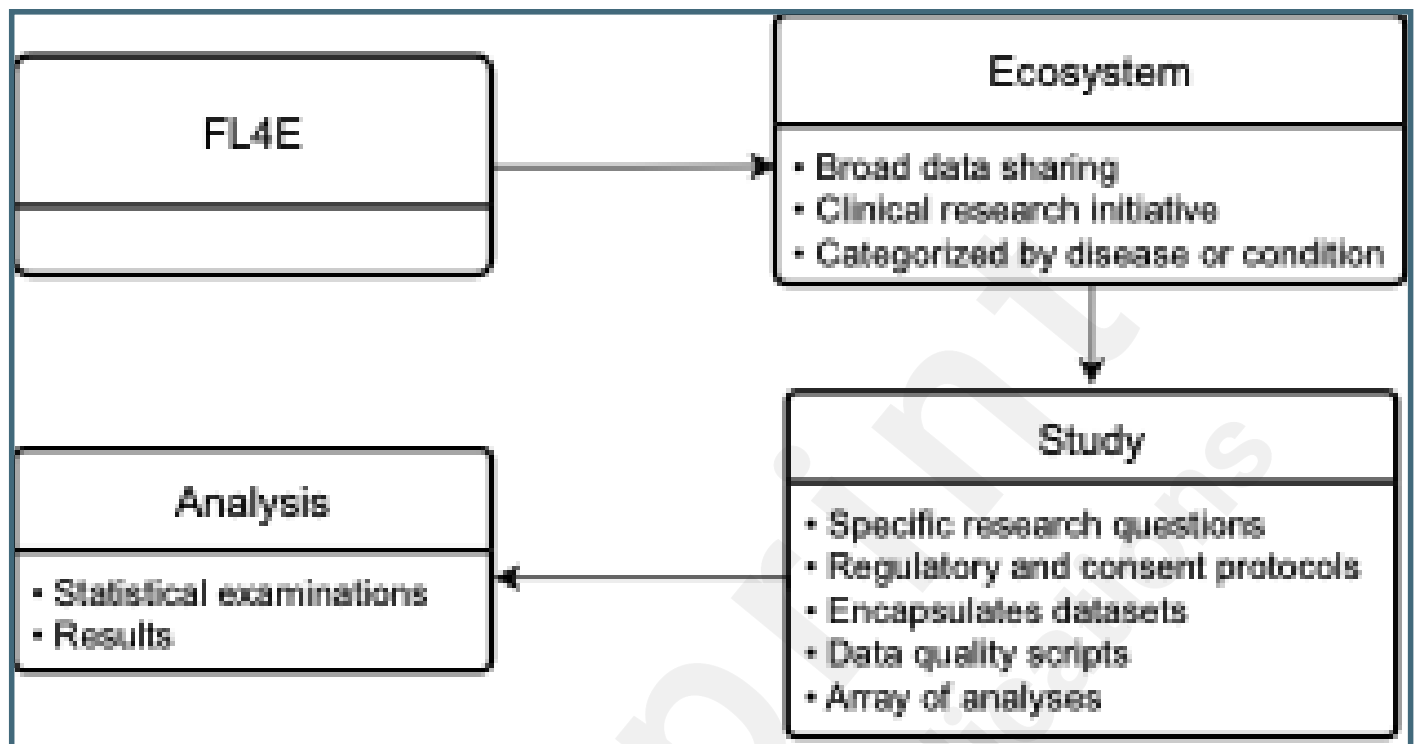
Supplementary Files

Figures


The degree of federation characterizes the balance between fully centralized and fully federated setups, leading to hybrid solutions where some of the stakeholders centralize their data while others prefer a federated approach.



FL4E's Nested Approach Schema: The schema presents the ecosystem approach of the FL4E, facilitating the incorporation of diverse studies framed by different research questions. It further highlights the support for a wide range of analyses, leading to the execution of statistical experiments.



Delineating Client Heterogeneity in Real-World Applications: Exploration of class imbalance in the Fed-Heart-Disease analysis.



Multimedia Appendixes

Responses for the reviewers.

URL: <http://asset.jmir.pub/assets/04568fdeb162e84dfdadd431974fc8f1.docx>

