big data and cognitive computing

MDPI

Article

Federated Learning to Safeguard Patients Data: A Medical Image Retrieval Case

Gurtaj Singh [1], Vincenzo Violi [1,2] and Marco Fisichella [3,*]

[1] Laboratory ARTS, University Mediterranea of Reggio Calabria, 89124 Reggio Calabria, Italy
[2] Department of Information Engineering, University of Pisa, 56122 Pisa, Italy
[3] L3S Research Center, Leibniz University Hannover, 30167 Hannover, Germany
* Correspondence: mfisichella@L3S.de

Abstract: Healthcare data are distributed and confidential, making it difficult to use centralized automatic diagnostic techniques. For example, different hospitals hold the electronic health records (EHRs) of different patient populations; however, transferring this data between hospitals is difficult due to the sensitive nature of the information. This presents a significant obstacle to the development of efficient and generalizable analytical methods that require a large amount of diverse Big Data. Federated learning allows multiple institutions to work together to develop a machine learning algorithm without sharing their data. We conducted a systematic study to analyze the current state of FL in the healthcare industry and explore both the limitations of this technology and its potential. Organizations share the parameters of their models with each other. This allows them to reap the benefits of a model developed with a richer data set while protecting the confidentiality of their data. Standard methods for large-scale machine learning, distributed optimization, and privacy-friendly data analytics need to be fundamentally rethought to address the new problems posed by training on diverse networks that may contain large amounts of data. In this article, we discuss the particular qualities and difficulties of federated learning, provide a comprehensive overview of current approaches, and outline several directions for future work that are relevant to a variety of research communities. These issues are important to many different research communities.

Keywords: federated learning; health; privacy preserving

check for updates

Citation: Singh, G.; Violi, V.; Fisichella, M. Federated Learning to Safeguard Patients Data: A Medical Image Retrieval Case. *Big Data Cogn. Comput.* **2023**, *7*, 18. https://doi.org/10.3390/bdcc7010018

Academic Editor: Moulay A. Akhloufi

Received: 8 December 2022
Revised: 9 January 2023
Accepted: 16 January 2023
Published: 18 January 2023

Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

Traditional machine learning models require centralized access to complete training data. Such conventional learning architectures collect local data from sensors and devices (i.e., clients) and transmit it to the central server for processing; later, the learned model weights are transmitted back to the local clients. This round trip limits the ability of a model to learn in real time [1–3].

Federated learning (FL), on the other hand, does not require centralized access to training data; instead, it trains the local model using local training data. The clients then transmit these locally trained models back to the central server, where they are aggregated (i.e., weights are averaged) before being returned to the clients as a single optimized global model. More specifically, FL allows machine learning algorithms to be adapted to learn over multiple iterations from different datasets distributed across multiple clients [4].

In order to demonstrate how federated learning is creating new possibilities, let us examine a few examples; medical privacy laws prevent research institutions from combining patient data to create more potent machine learning models, even though they rely on machine learning to detect cancer on MRI pictures. Using federated learning, they can better identify cancer patients so they can provide life-saving medicines by training highly accurate machine learning models on vast, heterogeneous, worldwide datasets [5].

All machine learning algorithms benefit significantly from accessing data that approximates the global distribution. FL elegantly addresses the challenges connected with

the egress of sensitive medical data by allowing different parties to train collaboratively without the need to communicate or centralize data sets. This reduces the need for centralized data storage. As a result, it has the potential to open up new avenues for innovative research and entrepreneurship, as well as to improve medical care around the world. The main contribution of our paper is the following:

- Privacy: through the use of FL, we try to provide a distributed learning solution that is privacy oriented, as clients do not forward their information to a central server.
- Data quality and robustness: we carried out tests on the variation of the distribution of datasets held by individual clients, and we demonstrated that the proposed solution is robust as it does not depend on the quality of the data held by individual clients.
- Scalability: we carried out a study based on scalability by evaluating the performance of our solution as the number of FL clients varies.
- Medical domain: where data are often distributed across different hospitals, clinics, and research organizations. We proposed a way in which FL can be used in the medical domain. In this particular case, our goal was to define a system based on a neural network capable of recognizing a case of COVID-19 from other pathologies through the use of X-ray images.

The X-ray image dataset is open-source, as are all the tools used to perform the proposed work.

1.1. General Data Protection Regulation

Due to the General Data Protection Regulation (GDPR) [6] implementation in 2018, EU residents now have a legal right to protect their data. This presents a barrier to conducting medical research, especially when more than one hospital is involved. It is possible to train decentralized machine learning models without sharing data between institutions, thus preserving users' right to confidentiality. Researchers in the healthcare industry have extensively used FL to analyze medical datasets.

1.2. Challenges of Fl in Healthcare

The use of FL in the healthcare industry raises several complicated issues. Hospitals and clinics may collect data in very different ways, and FL must understand a lot of textual and unstructured data. As a result, FL for health frequently calls for some data preprocessing on the part of the medical partner. Additionally, due to the resource-intensive nature of training machine learning models, FL systems are not commonly employed in hospitals and doctor's offices. Computers with GPU capabilities and gradient calculation capabilities are required in many hospitals. With the exception of IoT applications, this hinders development. FL is still a possibility, though, because there is a rapid model training option, and it is acceptable to take more time to produce a good model [7].

2. Related Work

The concept of Federated Learning (FL), proposed in 2017 [8], has attracted the research community's interest. It allows multiple clients to train global models without revealing private data. This training mode protects its users' privacy without violating owner authority and combines different data sources to maximize their utility. However, the data samples on each FL participating device are typically not independently and identically distributed (IID), which poses significant challenges to FL in terms of statistical heterogeneity [9]. In this paper, we explore the problem with non-IIID data. We also present some challenges that this problem may pose for FL. In recent years, many efforts have been made to develop algorithms for federated learning and data exploitation. Many use cases are related to FL. In [10,11], FL is applied to mobile data users, successfully using federated learning models to protect privacy and achieve maximum precision.

Mobile devices typically have limited data plans and slow network connections to the central server where the global model is managed, as seen in [12,13]. Therefore, a widely used application of federated learning is to reduce network link congestion to minimize

communication overhead as much as possible. Another application is highlighted in [14], where the goal is to secure resources from constrained devices and consider the problem of learning model parameters from data distributed across multiple edge nodes without sending raw data to a central location. It is also possible to solve other problems, such as reducing communication costs, as shown in [15], where a two-stream model with a MMD (Maximum Mean Discrepancy) constraint is trained instead of the single model on devices in default federated learning settings.

A critical issue is the use of medical data. They are kept in separate data silos, and access to this data is restricted due to patient privacy concerns, so the learning model cannot fully utilize these data. However, when enough data are available, machine learning can reach its full potential and transition from research to clinical practice. In particular, in our study, federated learning techniques are used to analyze a case of medical data analysis to overcome this type of challenge, especially in the processing of chest X-ray images. Deep learning is commonly utilized in FL to diagnose diseases based on chest X-ray images [16,17].

The research presented in [18] has developed an algorithm capable of detecting pneumonia on chest radiographs at any time. This level of performance is superior to that of practicing radiologists using a dense convolutional network.

In [19], a comparison of the performance of four popular models (MobileNetv2, ResNet18, ResNeXt, and COVID-Net) with the federated learning framework and without the framework is proposed to detect COVID-19 from chest X-ray images. In [20], a federated learning system with a dynamic focus for COVID-19 recognition on CXR images, called FedFocus, is developed. To improve the training efficiency and accuracy, the training loss of each model is used as the basis for parameter aggregation weighting. As the depth of the training layer increases, a continuously updated dynamic factor is developed to stabilize the aggregation process.

3. Federated Learning

McMahan first used FL in an edge server architecture while updating language models on mobile devices [8]. There are a large number of mobile edge devices, and each one stores private information.

Researchers at Google developed a federated learning system to frequently update a collective model to update predictive models in the Gboard system, which is Google's keyboard auto-completion system [21]. This was performed to improve the accuracy of Gboard's predictions. Users of the Gboard system receive recommended search terms, and the system remembers whether the terms were clicked or not. The word prediction algorithm used by Gboard is constantly being improved through a process known as federated averaging. This method takes into account data from all mobile devices, not just one (FedAvg). Federated averaging does not require data from each edge device to be transmitted to a central location. In federated learning, each mobile device, such as a smartphone or tablet, encrypts the model before uploading it to the cloud.

All encrypted models are combined into a single global model that is also encrypted. This is performed to prevent the cloud server from reading the data from individual devices. Then, the new encrypted model is downloaded to each device at the edge of the cloud system. The federated learning system that Google has developed is an excellent example of business-to-consumer (B2C) interaction in the context of developing a secure and decentralized learning environment for B2C applications [22]. In a business-to-consumer environment, federated learning can increase speed while protecting user privacy. This is because information can be transferred faster between edge devices and the central server.

In addition to the business-to-consumer paradigm, the business-to-business paradigm (often referred to as B2B) can also benefit from federated learning. As a fundamental change to the algorithmic architecture, federated learning replaces the communication of model parameters with the traditional practice of exchanging data from one site to another. This prevents unauthorized third parties from "guessing" the content of data belonging to others.

We categorize federated learning by the process by which the data is distributed among the many participants.

3.1. Definition

The goal of federated learning is to build a collaborative machine-learning model (ML) using data from multiple sources. Model training and inference are the two distinct processes that occur in federated learning. Only information about the neural network weights can be passed back and forth between parties while model training is being performed. Due to the nature of the transaction, neither confidential nor private information is revealed on either side [8]. A party can keep the trained model for itself or share it with multiple parties. The model is applied to new information when it is time to finalize the data. For example, in a business-to-business situation, a federated medical imaging system may accept a new patient even if the patient's diagnosis came from another facility. In conclusion, it is necessary to establish a method for equitably sharing the profits generated by collaboration. The viability of the association should be considered in the design of the processes. A model is a function that associates a data instance at a party with an outcome. Federated Learning is an algorithmic framework for building machine learning (ML) models [23]. A model is a function that associates a data instance at a party with an outcome.

- At least two different groups have shown interest in constructing a machine learning model together, and each group has data that it would want to utilize to train the model.
- During the process of training a model, each partner is responsible for keeping all of the data.
- The model might be encrypted and partly shared between parties, preventing third parties from re-engineering the data from a particular party. This would be accomplished by employing a method for encryption.
- The performance of the finished model is equivalent to that of an ideal model that was built with all data submitted to a single party throughout the construction process.

The FL paradigm is characterized by many properties and is applicable to many scenarios. In particular, two possible implementations of FL exist and are reported hereafter: horizontal FL and vertical FL. Regarding its properties, one of the most important is related to privacy. Some advantages due to this feature have already emerged, but in the following subsections, we will discuss them better. Given our focus on the medical domain, we will discuss how FL can be applied in this scenario.

3.2. Horizontal Federated Learning

Horizontal federated learning, also known as sample-based federated learning, is implemented whenever data sets have the same feature space but distinct samples. This kind of federated learning is also known as sample-based federated learning. For instance, the intersection set of customers at two regional banks may be a very small group consisting of significantly different user groups hailing from various places. The highlight spaces of both companies are the same due to the similarities between their operations. A cooperative deep-learning technique is suggested in the reference. This strategy calls for participants to train individually and only share sections of parameter changes with one another. In 2017, Google proposed a horizontal federated learning strategy to improve Android phone models [8]. Using this framework, a single Android phone user may adjust model parameters locally and then upload them to the Android cloud. This helps train the centralized model in coordination with other data owners. They also provide a safe aggregation solution as part of their federated learning architecture to guarantee the privacy of aggregated user updates [24]. Additively homomorphic encryption is used for model parameter aggregation in [25], which provides protection from the central server.

The work in [26] advocates using a multi-task federated learning system to enable several places to carry out autonomous tasks while simultaneously sharing information and keeping the system secure. In addition to this, the multi-task learning technique that they give may help solve problems with fault tolerance, laggards, and costly communication.

The authors proposed the development of a secure client-server architecture in [27], which would allow models created on client devices to collaborate with models created at the server site to form a global federated model. This would be possible through an encrypted connection between the client and the server. The development of the model guarantees that no sensitive information will be divulged to any unauthorized parties. The authors of [28] also offered several techniques to cut down on communication costs so that it would be easier to train centralized models utilizing data that was spread among mobile clients.

3.3. Vertical Federated Learning

For data that has been vertically partitioned, many machine-learning strategies have been proposed to protect users' privacy. These strategies include classification [29], gradient descent [30], secure linear regression [31,32], association rule mining [33], and cooperative statistical learning [34]. A vertical federated learning strategy for generating a logistic regression model while protecting privacy was recently published in references [35,36]. The authors investigated the effect of entity resolution on learning performance and used Taylor approximation for the loss and gradient functions to make it possible to use homomorphic encryption for computations that protect privacy. This was performed in order to enable the use of homomorphic encryption. Vertical federated learning, also known as feature-based federated learning, is appropriate for situations where two data sets share the same sample ID space but have different feature spaces. Vertical federated learning can also be viewed as a vertical version of horizontal federated learning.

Consider two local firms, one of which is an online retailer and the other a financial institution located in the same city. The bulk of the locals is likely to be included in their user sets, which results in a substantial amount of overlap between their user spaces. E-commerce maintains the browsing and purchase histories of the user, but banks record the user's income and spending habits in addition to their credit rating; as a result, the feature spaces of e-commerce and banks are quite different from one another. Let us imagine that we want a prediction model for a product purchase based on data from both the user and the product and that this data are shared between the user and the company. Vertically federated learning combines these variables and calculates the training loss and gradients to protect users' privacy from developing a model that collectively uses data from both sides. This model can then be used to make predictions. Because it allows all participants to build a plan for "shared wealth" while maintaining their unique identities and rankings, this form of federal mechanism is known as "federated learning".

3.4. Privacy

The protection of one's privacy is an essential part of federated learning. It is necessary to use security models and conduct analysis to give tangible privacy guarantees. This part will analyze and contrast some different privacy solutions for federated learning. Additionally, we will outline ways to reduce indirect leakage and significant difficulties that may arise [37].

3.4.1. Data Anonymization

Data anonymization, also known as de-identification, is concealing (through methods such as hashing, for example) or removing sensitive characteristics. For instance, PersonallyIdentifying Information (PII) in order to make the altered dataset (also known as the anonymous dataset) unrecognizably anonymous.

The process of data anonymization has to find a balance between protecting individuals' privacy and the value of the data collected. Hiding or omitting information may make the dataset less usable. A data subject may also be re-identified when linked with auxiliary data from other anonymous datasets, rendering them vulnerable to a linkage attack. This makes it possible to re-identify a data subject. Numerous methods, such as those based on k-anonymity, l-diversity, and closeness, preserve the distribution of sensitive attributes in a dataset to reduce the risk of re-identification of a data subject within the

same quasi-identifier group, have been proposed as ways to thwart linkage attacks [37]. These methods preserve the distribution of sensitive attributes in a dataset. Sadly, such privacy-protecting methods cannot defend against linkage assaults with opponents who are aware of the sensitive qualities. The procedures based on k-anonymity need to be improved, which will need new strategies, such as differential privacy, that provide a comprehensive assurance of confidentiality.

3.4.2. Secure Multi-Party Computation (SMC)

SMC security models inherently include several parties, and they provide total zero knowledge by providing security proof inside a well-defined simulation framework. This means that each party knows nothing but its own input and output. The attribute of having zero knowledge is very desirable; nevertheless, accomplishing this goal often calls for convoluted computational processes, and there is no guarantee that it can be performed quickly. If enough safeguards are in place for the situation, it is possible that the revelation of some but not all of one's information might be deemed appropriate. It is feasible to construct a security model using SMC while requiring a lesser level of security in order to achieve higher levels of efficiency [29]. In recent research [38], the SMC framework has been utilized for training machine learning models using two servers and semi-honest assumptions. The multi-party computation (MPC) methods are used for model training and verification in [39], which prevents users from disclosing sensitive data. Sharemind [40] is one of the most advanced SMC frameworks currently available. A 3PC model with an honest majority and security considerations for both semi-honest and malevolent assumptions was provided in reference [41], which may be found in references [42–44]. These efforts need the data of participants to be covertly transferred across servers that are not collaborating with one another.

3.4.3. Differential Privacy

Techniques such as Differential Privacy [45] and k-anonymity are used to protect users' data privacy [46–49]. Adding noise to the data or using generalization methods to obscure certain sensitive attributes until the third party cannot distinguish the individual is the main goal of differential privacy, k-anonymity, and diversification [50] methods. This ensures that the data cannot be restored, hence leading to privacy guarantees. However, at their core, these approaches still need the transmission of the data to another location, and the execution of these methods often necessitates a compromise between accuracy and privacy. The authors of paper [51] presented an adaptation of the differential privacy method in a federated learning environment to protect clients' privacy rights. This was accomplished by concealing client contributions while the system was being trained.

3.5. Federated Learning and Healthcare

In the past couple of years, we have experienced a great deal of transition due to a pandemic. During this historical period, it was clear that healthcare service providers did not have sufficient resources at their disposal. Medical practitioners must have access to trustworthy technology to provide excellent patient treatment. However, substantial and varied datasets are required to train an algorithm for therapeutic applications. It becomes more difficult to communicate sensitive information when rigorous limitations such as Health Insurance Portability and Accountability Act (HIPAA) [52] are in place. FL was founded to respond to problems such as these. Participating organizations use their internal data pool to train the same algorithm. These establishments might profit from these trained algorithms. They could avoid compliance with some restrictions while maintaining access to others. In addition, it creates a database from which they may glean information and expertise.

For the field of machine learning, FL is an innovative new idea and method. This has the potential to bring significant progress in the healthcare sector. FL does not want to play the role of medical authority at this time. It is designed to free them up so they may focus on providing better patient care [5].

4. Approach

We will discuss the methods used to successfully complete the examinations in the following paragraphs.

The workflow proposed by us is shown in Figure 1. This technique provides an illustration of the processes that are carried out on both the client side and the central server.

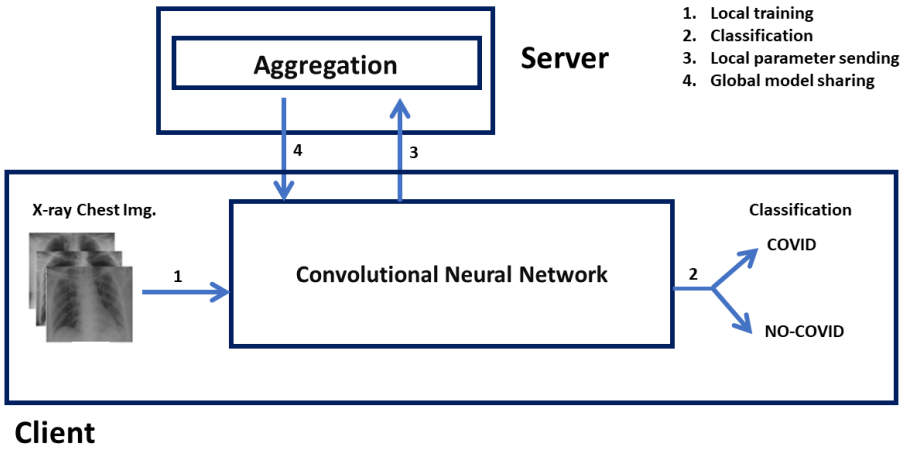

Figure 1. Local training and global aggregation workflow.

The client is responsible for maintaining the local dataset and, in this instance, includes pictures of the patient's radiographs. In the next phase, indicated by *Local Training*, all clients use their datasets to train the local neural network. Clients learn to classify inputs during the training phase indicated by *Classification*. That is, they learn to recognize images related to COVID-19. The approach used in this case study is based on supervised learning, where the samples of the dataset are accompanied by a label indicating the type of image in question.

In the third step, indicated with *Local Parameter Sending*, all the clients of the scenario send the weights of the locally trained neural network to the FL server in question. As mentioned in the previous sections, the server has the task of aggregating and averaging the weights from all clients to update the global neural network. This global neural network will, therefore, have a higher generalization capacity than the local one precisely because it is built with the help of all clients. The next step is to distribute the global learner weights to all clients. This process will be repeated until the desired metrics are achieved.

The neural network model's structure is designed to detect even the most minute details. The implementation details of the proposed model are reported in Table 1. The CNN [53] structure comprises 17 Conv2d layers, a flattening layer, and a linear layer.

Table 1. CNN structure.

Number of Layer	Layer Type	Output Shape	Number of Trainable Parameters
1	Conv2d	[8, 256, 256]	216
2	Conv2d	[16, 128, 128]	1152
3	Conv2d	[32, 64, 64]	4608
4	Conv2d	[16, 66, 66]	512
5	Conv2d	[32, 256, 256]	4608
6	Conv2d	[64, 33, 33]	18,432

Table 1. *Cont.*

Number of Layer	Layer Type	Output Shape	Number of Trainable Parameters
7	Conv2d	[32, 35, 35]	2048
8	Conv2d	[64, 35, 35]	18,432
9	Conv2d	[128, 17, 17]	73,728
10	Conv2d	[74, 19, 19]	8192
11	Conv2d	[128, 19, 19]	73,728
12	Conv2d	[256, 9, 9]	294,912
13	Conv2d	[128, 11, 11]	32,768
14	Conv2d	[256, 11, 11]	294,912

4.1. Dataset

A database of X-ray images of patients with COVID-19 and other conditions, such as pneumonia, was used during this specific examination [54]. This dataset is obtained by merging some datasets used in previous works and is freely available on GitHub.

Table **??** shows the attributes associated with each image. Compared to the starting dataset, we carried out a preprocessing phase. In particular, we have identified all cases of COVID-19, and we have defined all the rest of the cases as NO-COVID since our goal is to train a neural network that can recognize COVID-19 cases from all the rest. We also filtered the attributes associated with each image as many defined in [54] were not helpful for the training process.

Table 2. Image attributes.

Attribute	Description
Patient ID	Internal identifier
Offset	Number of days since the start of symptoms or hospitalization for each image. If a report indicates "after a few days", then 5 days is assumed
Sex	Male (M), Female (F), or blank
Age	Age of the patient in years
Finding	Type of pneumonia
Survival	Yes or No
View	Posteroanterior (PA), Anteroposterior (AP), AP Supine (APS), or Lateral (L) for X-rays; Axial or Coronal for CT scans
Modality	CT, X-ray, or something else

Number of Layer	Layer Type	Output Shape	Number of Trainable Parameters
15	Conv2d	[128, 13, 13]	256
16	Conv2d	[256, 13, 13]	294,912
17	Conv2d	[2, 13, 13]	4608
18	Flatten	[338]	0
19	Linear	[2]	678

4.2. Dataset Distribution

Regarding the dataset's distribution, we have reported two scenarios: *(i)* an IID scenario where data are equally distributed in a stratified fashion among the clients. *(ii)* a non-IID scenario where the data has been distributed unevenly, following a Gaussian distribution with $N(0, 4)$.

The choice of $\sigma^2 = 4$ is dictated to have a smooth distribution of data in contrast to values such as 1 that introduce a major skew of the slope. For instance, with five clients, we ended up having 36% of data for the first client, 27% for the second client, 20%, 11%, and 6% for the remaining.

5. Results

Before discussing the tests carried out, it is necessary to define that each subset of data associated with each individual client has been partitioned to avoid the overfitting phenomenon; in particular, the following subdivision has been carried out: 60% of the data are reserved for the training procedure, 20% of the data are reserved for the validation set useful to avoid overfitting, and the remaining 20% are reserved for the test set to evaluate the performance of the trained network.

Each measurement made was repeated five times, and the average value of these tests is reported in the results that will be discussed in this section.

Simulations have been run to illustrate the applicability of federated learning regarding the application of the example presented before. All experiments were run on a machine with 16 GB of RAM, an Intel core i7 9750H processor, and an NVIDIA Geforce GTX 1050 Ti graphics card. According to the results of our study, the performance of federated learning is extremely comparable to that of the centralized learning scheme.

The final purpose of the paper is to highlight the evident variation of the IID and non-IID cases of accuracy. In detail, the term "Identically Distributed" refers to a situation in which there are no overarching trends, the distribution does not vary, and all the items in the sample come from the same probability distribution. The term "independent" refers to the fact that each sample item is a separate occurrence. In other words, they do not have any relationship with one another at all. The results demonstrate that the data distribution parameter can affect the accuracy. It has been established that even if the data are disseminated identically, the number of data distribution clients still substantially impacts the ultimate accuracy and processing of the data.

The results obtained from our case for the centralized system have excellent accuracy, as can bee seen in Figure 2. For example, in the best case, achieving an accuracy peak of about 98% after 200 epochs is possible. This is because the central server, which performs the training together, has a more comprehensive and complete overview and can thus achieve better accuracy for each case.

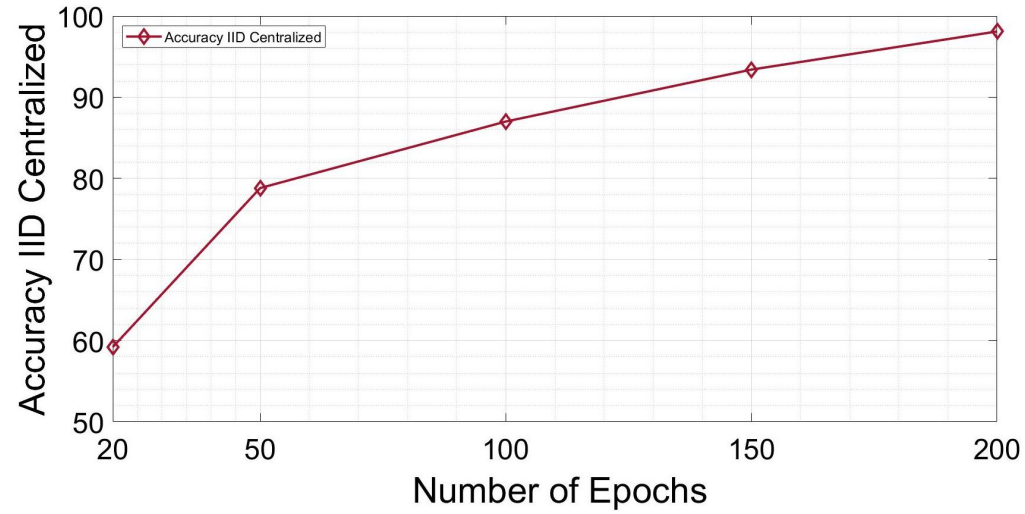

Figure 2. Accuracy in Centralized IID case of use.

However, as explained in the previous sections, this proposed methodology has some limitations. Several tests have been conducted with the federated case, IID, and non-IID to overcome these problems, with satisfactory results. Looking at Figure 3, we notice that, in the IID case, setting the number of rounds to 200, we have an accuracy of 96% for the scenario with 5 clients and an accuracy of about 93% for the scenario with 15 clients.

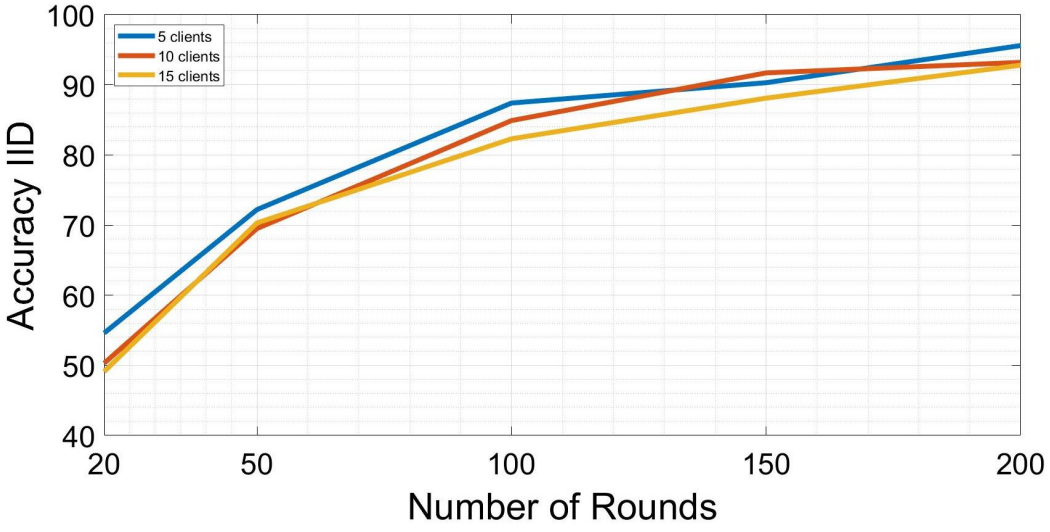

Figure 3. Accuracy in the scenario with IID data when varying the number of clients.

The case with NO-IID is reported in Figure 4. When the number of rounds varies, the accuracy with 5 clients can reach peaks of 90% accuracy and 87% with 15 clients.

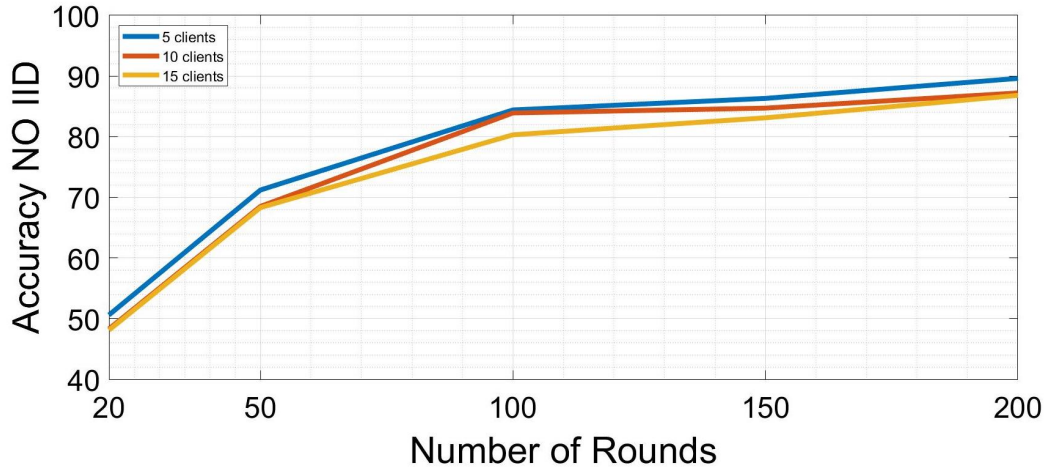

Figure 4. Accuracy in the scenario with NO-IID data when varying the number of clients.

For this type of study, both the IID and the NO-IID distribution cases exhibit similar accuracy trends, as shown in Figures 3 and 4. In our scenario, we observed that the accuracy improves with increasing rounds, reaching peaks of 90% after 200 rounds and 87% in the worst case (where the distribution of 10 and 15 clients coincides).

The comparison in Figure 5 shows that IID data give better results for 5 clients, with a difference of around 6% between IID and NO-IID after 200 rounds. With a split of 10 or 15 clients, the results are comparable with that of 5 clients. Moreover, the centralized case is slightly better than the federated one, but at the expense of data protection.

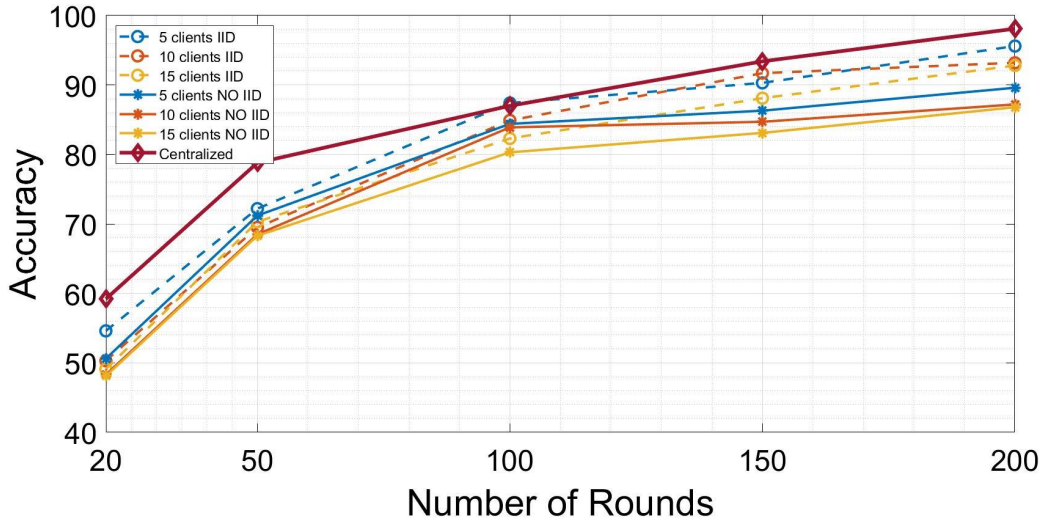

Figure 5. Accuracy comparison between the IID, NO IID, and centralized IID cases.

Figure 6 also shows the confusion matrices for the IID case with varying clients and for the centralized case, while Figure 7 shows the confusion matrices for the NO-IID case with varying clients. The confusion matrices shown have the ground truth values as rows and the predictions as columns. Each element of each matrix represents the values *true positive* (first row and first column), *false positive* (second row and first column), *false negative* (first row and second column), and *true negative* (second row and second column).

Furthermore, since the matrices are built with the parameters just mentioned, functions such as precision and recall are easily obtainable.

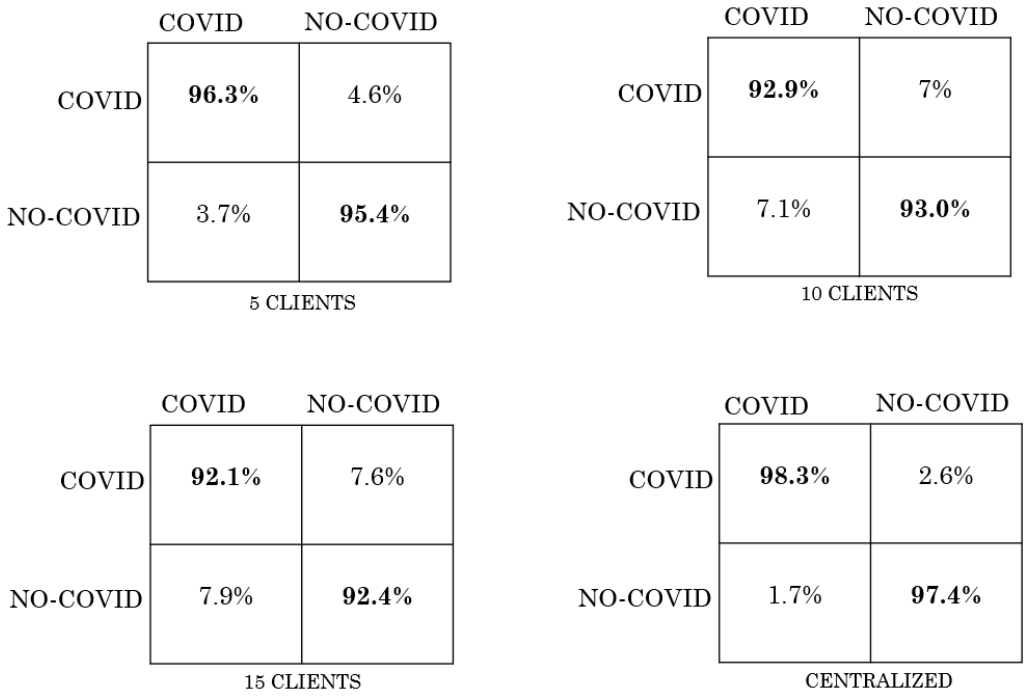

Figure 6. Confusion matrices for the IID and centralized scenarios.

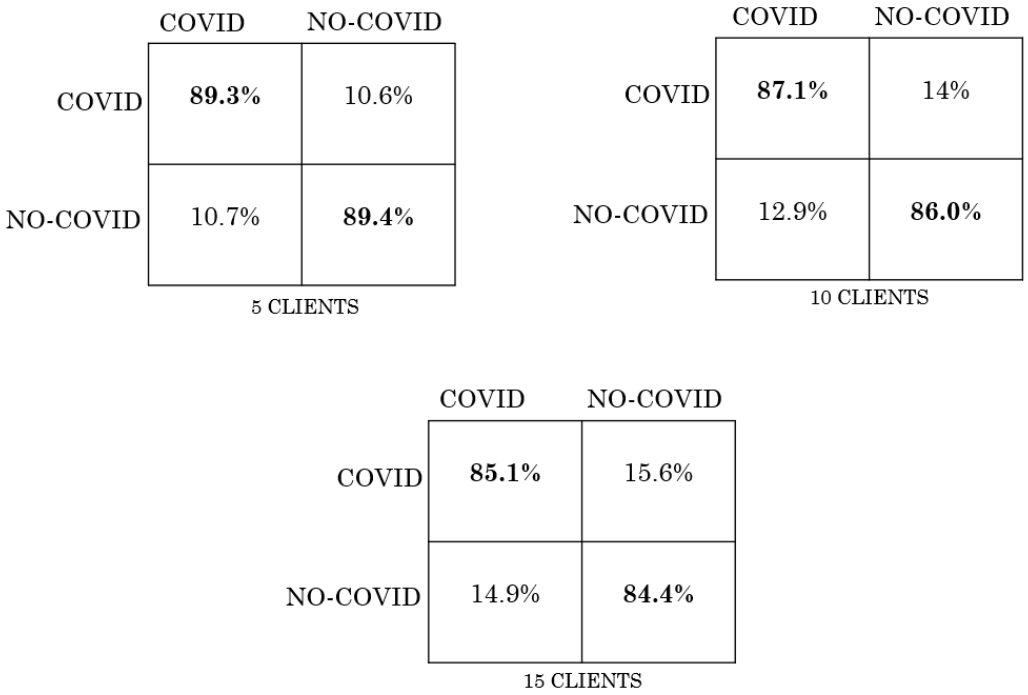

Figure 7. Confusion matrices for the NO-IID scenario.

Table 3 shows the *precision* values obtained for each scenario.

Table 3. Precision achieved for each scenario.

Scenario	IID	NO-IID
FL with 5 clients	0.97	0.93
FL with 10 clients	0.9	0.87
FL with 15 clients	0.93	0.85
Centralized	0.98	0.98

By carefully observing the confusion matrices and the results obtained in terms of precision, we can effectively state that the centralized case has superior performance, but despite this, the various FL approaches that vary in terms of scalability and in terms of dataset distribution still have excellent performance. In particular, as seen in the accuracy measures, we observe that the results obtained are in agreement with the previous ones, as the scenario with a distribution of IID data performs slightly better than the NO-IID scenario.

6. Conclusions

This work allows us to expand our use cases in many areas of engineering, healthcare, and computer science and summarize overviews of federated learning. We discussed several medical distribution examples that show how federated learning can overcome the problems of the centralized case mentioned above and provide a valid and efficient alternative, especially in terms of privacy and algorithm efficiency.

In summary, in the federated settings, the performance results are slightly better when the distribution is identical and independent, i.e., IID versus non-IID. Even if a centralized system could achieve slightly better results in terms of accuracy, federated learning is a good way to replace the centralized algorithm while preserving data privacy and achieving excellent results. Indeed, the use of FL, for example, in the health sector, often requires some pre-treatment of the data on the part of the medical partner. This is something that needs to be taken into account, as mentioned earlier. In addition, because of the extensive use of computing resources required to train machine learning models, FL systems are not widely used in healthcare settings such as hospitals and physician offices. Many hospitals need computers equipped with GPU capabilities and cannot perform gradient calculations. Nevertheless, despite all of its drawbacks, FL demonstrates capabilities that enable good results.

Author Contributions: G.S.: Conceptualization, Methodology, Software, Investigation, Writing—original draft, Visualization. V.V.: Conceptualization, Methodology, Software, Investigation, Writing—original draft, Visualization. M.F.: Supervision, Conceptualization, Validation, Resources, Writing—original draft, Writing—review & editing, Funding acquisition. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by L3S Research Center of Leibniz University of Hannover, Germany.

Institutional Review Board Statement: The dataset and project in [54] were approved by the University of Montreal's Ethics Committee #CERSES-20-058-D. The dataset in [17] is an open-source medical database extracted from the National Institutes of Health Clinical Center's clinical database PACS with ethical approval from the Intramural Research Programs of the NIH Clinical Center and National Library of Medicine. Note that the original radiology reports (related to these chest x-ray examinations) are not for public use for many reasons. For retrospective data collection, written informed consent was obtained from participating patients before enrollment in the study.

Data Availability Statement: The data presented in this study are available in [54].

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Ceroni, A.; Gadiraju, U.; Matschke, J.; Wingert, S.; Fisichella, M. Where the Event Lies: Predicting Event Occurrence in Textual Documents. In Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2016, Pisa, Italy, 17–21 July 2016; Perego, R., Sebastiani, F., Aslam, J.A., Ruthven, I., Zobel, J., Eds.; ACM: New York, NY, USA, 2016; pp. 1157–1160. [CrossRef]
2. Ceroni, A.; Gadiraju, U.K.; Fisichella, M. Improving Event Detection by Automatically Assessing Validity of Event Occurrence in Text. In Proceedings of the 24th ACM International Conference on Information and Knowledge Management, CIKM 2015, Melbourne, VIC, Australia, 19–23 October 2015; Bailey, J., Moffat, A., Aggarwal, C.C., de Rijke, M., Kumar, R., Murdock, V., Sellis, T.K., Yu, J.X., Eds.; ACM: New York, NY, USA, 2015; pp. 1815–1818. [CrossRef]
3. Banabilah, S.; Aloqaily, M.; Alsayed, E.; Malik, N.; Jararweh, Y. Federated learning review: Fundamentals, enabling technologies, and future applications. *Inf. Process. Manag.* **2022**, *59*, 103061. [CrossRef]
4. Nguyen, D.C.; Ding, M.; Pathirana, P.N.; Seneviratne, A.; Li, J.; Poor, H.V. Federated learning for internet of things: A comprehensive survey. *IEEE Commun. Surv. Tutorials* **2021**, *23*, 1622–1658. [CrossRef]
5. Nguyen, D.C.; Pham, Q.V.; Pathirana, P.N.; Ding, M.; Seneviratne, A.; Lin, Z.; Dobre, O.; Hwang, W.J. Federated learning for smart healthcare: A survey. *ACM Comput. Surv. (CSUR)* **2022**, *55*, 1–37. [CrossRef]
6. Viorescu, R. 2018 reform of EU data protection rules. *Eur. J. Law Public Adm.* **2017**, *4*, 27–39. [CrossRef]
7. Pfitzner, B.; Steckhan, N.; Arnrich, B. Federated learning in a medical context: A systematic literature review. *ACM Trans. Internet Technol.* **2021**, *21*, 1–31. [CrossRef]
8. McMahan, H.B.; Moore, E.; Ramage, D.; Hampson, S.; Arcas, B.A.Y. Communication-Efficient Learning of Deep Networks from Decentralized Data. *arXiv* **2016**, arXiv:1602.05629. https://doi.org/10.48550/ARXIV.1602.05629.
9. Younis, R.; Fisichella, M. FLY-SMOTE: Re-Balancing the Non-IID IoT Edge Devices Data in Federated Learning System. *IEEE Access* **2022**, *10*, 65092–65102. [CrossRef]
10. Liu, Q.; Huang, S.; Opadere, J.; Han, T. An edge network orchestrator for mobile augmented reality. In Proceedings of the IEEE INFOCOM 2018-IEEE Conference on Computer Communications, Honolulu, HI, USA, 16–19 April 2018; pp. 756–764.
11. Nilsson, A.; Smith, S.; Ulm, G.; Gustavsson, E.; Jirstrand, M. A performance evaluation of federated learning algorithms. In Proceedings of the Second Workshop on Distributed Infrastructures for Deep Learning, Rennes, France, 10–11 December 2018; pp. 1–8.
12. Luping, W.; Wei, W.; Bo, L. CMFL: Mitigating communication overhead for federated learning. In Proceedings of the 2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS), Dallas, TX, USA, 7–10 July 2019; pp. 954–964.
13. Spirovska, K.; Didona, D.; Zwaenepoel, W. Paris: Causally consistent transactions with non-blocking reads and partial replication. In Proceedings of the 2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS), Dallas, TX, USA, 7–10 July 2019; pp. 304–316.
14. Wang, S.; Tuor, T.; Salonidis, T.; Leung, K.K.; Makaya, C.; He, T.; Chan, K. Adaptive federated learning in resource constrained edge computing systems. *IEEE J. Sel. Areas Commun.* **2019**, *37*, 1205–1221. [CrossRef]
15. Yao, X.; Huang, C.; Sun, L. Two-stream federated learning: Reduce the communication costs. In Proceedings of the 2018 IEEE Visual Communications and Image Processing (VCIP), Taichung, Taiwan, 9–12 December 2018; pp. 1–4.
16. Feki, I.; Ammar, S.; Kessentini, Y.; Muhammad, K. Federated learning for COVID-19 screening from Chest X-ray images. *Appl. Soft Comput.* **2021**, *106*, 107330. [CrossRef]
17. Wang, X.; Peng, Y.; Lu, L.; Lu, Z.; Bagheri, M.; Summers, R.M. Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21–26 July 2017; pp. 2097–2106.
18. Lerzynski, G. Ethical Implications of Digitalization in Healthcare. In *Digitalization in Healthcare*; Springer: Berlin/Heidelberg, Germany, 2021; pp. 163–170.
19. Liu, B.; Yan, B.; Zhou, Y.; Yang, Y.; Zhang, Y. Experiments of federated learning for covid-19 chest x-ray images. *arXiv* **2020**, arXiv:2007.05592.
20. Li, Z.; Xu, X.; Cao, X.; Liu, W.; Zhang, Y.; Chen, D.; Dai, H. Integrated CNN and federated learning for COVID-19 detection on chest X-ray images. *IEEE/ACM Trans. Comput. Biol. Bioinform.* **2022**, *early access*. [CrossRef]
21. Yang, Q.; Liu, Y.; Cheng, Y.; Kang, Y.; Chen, T.; Yu, H. Federated learning. *Synth. Lect. Artif. Intell. Mach. Learn.* **2019**, *13*, 1–207.
22. Yang, Q.; Fan, L.; Yu, H. *Federated Learning: Privacy and Incentive*; Springer: Berlin/Heidelberg, Germany, 2020; Volume 12500.
23. Lim, W.Y.B.; Luong, N.C.; Hoang, D.T.; Jiao, Y.; Liang, Y.C.; Yang, Q.; Niyato, D.; Miao, C. Federated learning in mobile edge networks: A comprehensive survey. *IEEE Commun. Surv. Tutorials* **2020**, *22*, 2031–2063. [CrossRef]
24. Segal, A.; Marcedone, A.; Kreuter, B.; Ramage, D.; McMahan, H.B.; Seth, K.; Bonawitz, K.A.; Patel, S.; Ivanov, V. Practical Secure Aggregation for Privacy-Preserving Machine Learning. In Proceedings of the CCS, Dallas, TX, USA, 30 October–3 November 2017.
25. Phong, L.T.; Aono, Y.; Hayashi, T.; Wang, L.; Moriai, S. Privacy-Preserving Deep Learning via Additively Homomorphic Encryption. *IEEE Trans. Inf. Forensics Secur.* **2018**, *13*, 1333–1345. [CrossRef]
26. Smith, V.; Chiang, C.K.; Sanjabi, M.; Talwalkar, A.S. Federated multi-task learning. *Adv. Neural Inf. Process. Syst.* **2017**, 4427–4437. [CrossRef]

27. McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; Arcas, B.A.Y. Communication-efficient learning of deep networks from decentralized data. In Proceedings of the Artificial Intelligence and Statistics (PMLR), Fort Lauderdale, FL, USA, 20–22 April 2017; pp. 1273–1282.
28. Konečnỳ, J.; McMahan, H.B.; Ramage, D.; Richtárik, P. Federated optimization: Distributed machine learning for on-device intelligence. *arXiv* **2016**, arXiv:1610.02527.
29. Du, W.; Han, Y.S.; Chen, S. Privacy-preserving multivariate statistical analysis: Linear regression and classification. In Proceedings of the 2004 SIAM International Conference on Data Mining (SIAM), Lake Buena Vista, FL, USA, 22–24 April 2004; pp. 222–233.
30. Wan, L.; Ng, W.K.; Han, S.; Lee, V.C.S. *Privacy-Preservation for Gradient Descent Methods*; Association for Computing Machinery: New York, NY, USA, 2007. [CrossRef]
31. Gascón, A.; Schoppmann, P.; Balle, B.; Raykova, M.; Doerner, J.; Zahur, S.; Evans, D. Secure Linear Regression on Vertically Partitioned Datasets. *IACR Cryptol. ePrint Arch.* **2016**, *2016*, 892.
32. Sanil, A.P.; Karr, A.F.; Lin, X.; Reiter, J.P. Privacy preserving regression modelling via distributed computation. In Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Seattle, DC, USA, 22–25 August 2004; pp. 677–682.
33. Vaidya, J.; Clifton, C. Privacy Preserving Association Rule Mining in Vertically Partitioned Data. In Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '02, Edmonton, AB, Canada, 23–26 July 2002; Association for Computing Machinery: New York, NY, USA, 2002 ; pp. 639–644. [CrossRef]
34. Du, W.; Atallah, M. Privacy-Preserving Cooperative Statistical Analysis. In Proceedings of the Seventeenth Annual Computer Security Applications Conference, New Orleans, LA, USA, 10–14 December 2001.
35. Hardy, S.; Henecka, W.; Ivey-Law, H.; Nock, R.; Patrini, G.; Smith, G.; Thorne, B. Private federated learning on vertically partitioned data via entity resolution and additively homomorphic encryption. *arXiv* **2017**, arXiv:1711.10677.
36. Nock, R.; Hardy, S.; Henecka, W.; Ivey-Law, H.; Patrini, G.; Smith, G.; Thorne, B. Entity resolution and federated learning get a federated resolution. *arXiv* **2018**, arXiv:1803.04035.
37. Fisichella, M.; Lax, G.; Russo, A. Partially-federated learning: A new approach to achieving privacy and effectiveness. *Inf. Sci.* **2022**, *614*, 534–547. .: 10.1016/j.ins.2022.10.082. [CrossRef]
38. Mohassel, P.; Zhang, Y. Secureml: A system for scalable privacy-preserving machine learning. In Proceedings of the 2017 IEEE Symposium on Security and Privacy (SP), San Jose, CA, USA, 22–26 May 2017; pp. 19–38.
39. Kilbertus, N.; Gascón, A.; Kusner, M.; Veale, M.; Gummadi, K.; Weller, A. Blind justice: Fairness with encrypted sensitive attributes. In Proceedings of the International Conference on Machine Learning (PMLR), Stockholm, Sweden, 10–15 July 2018; pp. 2630–2639.
40. Agrawal, D.; Aggarwal, C.C. On the design and quantification of privacy preserving data mining algorithms. In Proceedings of the 20th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, Santa Barbara, CA, USA, 21–23 May 2001; pp. 247–255.
41. Mohassel, P.; Rindal, P. ABY3: A Mixed Protocol Framework for Machine Learning. In Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS '18, Toronto, Canada, 15–19 October 2018; Association for Computing Machinery: New York, NY, USA, 2018; pp. 35–52. [CrossRef]
42. Araki, T.; Furukawa, J.; Lindell, Y.; Nof, A.; Ohara, K. High-Throughput Semi-Honest Secure Three-Party Computation with an Honest Majority. In Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS'16, Vienna, Austria, 24–28 October 2016; Association for Computing Machinery: New York, NY, USA, 2016; pp. 805–817. [CrossRef]
43. Furukawa, J.; Lindell, Y.; Nof, A.; Weinstein, O. High-throughput secure three-party computation for malicious adversaries and an honest majority. In Proceedings of the Annual international conference on the theory and applications of cryptographic techniques, Paris, France, 30 April–4 May 2017; pp. 225–255.
44. Mohassel, P.; Rosulek, M.; Zhang, Y. Fast and Secure Three-Party Computation: The Garbled Circuit Approach. In Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, CCS'15, Denver, CO, USA, 12–16 October 2015; Association for Computing Machinery: New York, NY, USA, 2015; pp. 591–602. [CrossRef]
45. Dwork, C. Differential Privacy: A Survey of Results. In Proceedings of the 5th International Conference on Theory and Applications of Models of Computation, TAMC'08, Xi'an, China, 25–29 April 2008; pp. 1–19.
46. Abadi, M.; Chu, A.; Goodfellow, I.; McMahan, H.B.; Mironov, I.; Talwar, K.; Zhang, L. Deep Learning with Differential Privacy. In Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS'16, Vienna, Austria, 24–28 October 2016; Association for Computing Machinery: New York, NY, USA, 2016; pp. 308–318. [CrossRef]
47. Chaudhuri, K.; Monteleoni, C. Privacy-preserving logistic regression. In Proceedings of the NIPS, Vancouver, BC, Canada, 8–11 December 2008.
48. McMahan, H.B.; Ramage, D.; Talwar, K.; Zhang, L. Learning Differentially Private Recurrent Language Models. *arXiv* **2017**, arXiv:1710.06963. https://doi.org/10.48550/ARXIV.1710.06963.
49. Stochastic gradient descent with differentially private updates. In Proceedings of the2013 IEEE Global Conference on Signal and Information Processing, Austin, TX, USA, 3–5 December 2013. [CrossRef]
50. Agrawal, R.; Srikant, R. Privacy-Preserving Data Mining. *SIGMOD Rec.* **2000**, *29*, 439–450. [CrossRef]
51. Geyer, R.C.; Klein, T.; Nabi, M. Differentially Private Federated Learning: A Client Level Perspective. *arXiv* **2017**, arXiv:1712.07557. https://doi.org/10.48550/ARXIV.1712.07557.

52. Mbonihankuye, S.; Nkunzimana, A.; Ndagijimana, A. Healthcare data security technology: HIPAA compliance. *Wirel. Commun. Mob. Comput.* **2019**, *2019*, 1927495. [CrossRef]
53. Ozturk, T.; Talo, M.; Yildirim, E.A.; Baloglu, U.B.; Yildirim, O.; Acharya, U.R. Automated detection of COVID-19 cases using deep neural networks with X-ray images. *Comput. Biol. Med.* **2020**, *121*, 103792. [CrossRef]
54. Cohen, J.P.; Morrison, P.; Dao, L. COVID-19 image data collection. *arXiv* **2020**, arXiv:2003.11597.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.