Federated Learning in Medical Imaging

Darzidehkalani, Erfan; Ghasemi-Rad, Mohammad; van Ooijen, P M A

Published in:
Journal of the american college of radiology

DOI:
10.1016/j.jacr.2022.03.016

Publication date:
2022

Federated Learning in Medical Imaging: Part II: Methods, Challenges, and Considerations

Erfan Darzidehkalani, MS[a,b], Mohammad Ghasemi-rad, MD[c], P. M. A. van Ooijen, PhD[a,b]

Abstract

Federated learning is a machine learning method that allows decentralized training of deep neural networks among multiple clients while preserving the privacy of each client's data. Federated learning is instrumental in medical imaging because of the privacy considerations of medical data. Setting up federated networks in hospitals comes with unique challenges, primarily because medical imaging data and federated learning algorithms each have their own set of distinct characteristics. This article introduces federated learning algorithms in medical imaging and discusses technical challenges and considerations of real-world implementation of them.

Key Words: Federated learning, medical imaging, privacy-preserving machine learning

J Am Coll Radiol 2022;19:975-982.

INTRODUCTION

With the rapid developments in machine learning in health care and computer-aided diagnosis, access to medical data has become a matter of interest. Clinicians, computer scientists, and medical technologists require access to more data to enable machine learning-based projects. However, it is always a challenging task to balance between building more powerful machines to be used in the health care industry and the limitations of accessing large amount of data under the privacy considerations. Generally, sharing data requires hospitals to deal with General Data Protection Regulation restrictions and have approval from the institutional review board. An institutional review board or ethical committee determines to which degree a hospital can share information with other hospitals and ensures that hospitals are compliant with the General Data Protection Regulation restrictions. As a result, data centers in hospitals do not usually have large and diverse data sets required to train deep neural networks.

Federated learning (FL) [1] is a machine learning concept proposed by McMahan et al to tackle this problem. In this concept, a neural network is trained with data sets from multiple hospitals, and the whole training process is managed through a central server. At each round, hospitals train a neural network on their local data and share the updated models with the central server. The server gathers all the updated models and aggregates them into an updated global model. In the next round, the updated global model is sent back to the hospitals. This way of training enables the researchers to use the data from multiple clients, while ensuring that the sensitive data are kept locally.

There exist several FL algorithms. McMahan et al [1] proposed federated averaging method (FedAvg) to minimize parameter change among hospitals. The algorithm is straightforward: A subset of the clients is selected each round. Training is distributed among multiple clients. Each client will compute an updated model on its own local data set. All model instances on the clients should start with the same random initialization to achieve convergence. Clients communicate with the central server once their local training has been finished. Finally, the central server gathers the updates of the respective clients. An immediate effect of local training can be seen at this stage. The updated global model can

[a]Department of Radiation Oncology, University Medical Center Groningen, University of Groningen, Groningen, the Netherlands.
[b]Machine Learning Lab, Data Science Center in Health, University Medical Center Groningen, University of Groningen, the Netherlands.
[c]Department of Interventional Radiology, Baylor College of Medicine, Houston, Texas.
Corresponding author and reprints: Erfan Darzidehkalani, Universitair Medisch Centrum Groningen, Department of Radiation Oncology, Hanzeplien1, Groningen, The Netherlands; e-mail: e.darzidehkalani@umcg.nl.

The authors state that they have no conflict of interest related to the material discussed in this article. The authors are non-partner/non-partnership track/ employees.

1546-1440/22/$36.00 ■ https://doi.org/10.1016/j.jacr.2022.03.016

be tested against a test data set, and comparing its performance with the previous round can give insight into how much improvement was achieved during the last round of training. An illustration of this step is shown in Figure 1. Blockchain-based technologies can also be used in the aggregation stage. In a blockchain network, local clients (miners) replace the central server and distribute the aggregation process among themselves. In this case, the whole process will be decentralized. Blockchain networks can be valuable since they prevent failure if the central server or clients fail [2].

Another approach is averaging the outputs of the local models trained on the clients individually (ensemble single client models). A general definition for ensemble learning is different machine learning algorithms doing the same task combined into one algorithm. Each algorithm extracts information or features from the input data, and the resulting information will be fused using various mechanisms, such as averaging, and voting. Generally, ensembles consistently outperform each of their constituting algorithms alone. In the federated setting of ensemble learning, neither the models nor the data will be shared among clients in the training cycle. All the clients will be assigned a similar model with random initial values. Each client will train its model. Their outputs for the same task will be averaged in the deployment phase, resulting in an accumulated knowledge from multiple models.

A third algorithm is the single weight transfer (SWT). In this algorithm, a deep learning model is trained at a single client up to a particular time and then transferred to the next client. There are numerous options to decide when to finish a local training and pass its model to the next client. Standard criteria are the number of epochs per client and validation loss or accuracy depending on the problem. For example, Chang et al [3] chose to reach the plateau of validation loss as a sign of moving to the next client. Cyclic weight transfer (CWT) is another algorithm in which a model is trained at each client for a predetermined number of epochs, then transferred to the next client. In this algorithm, the model visits each client than once.

The functionality of models and tasks in an FL scenario differs depending on the FL algorithm. Algorithms that transfer models are more versatile and adaptable than other algorithms. Deep learning models' performance in a federated environment can also vary from model to model. Models' adaptability can determine the overall performance of an FL network. For example, research has shown that some deep neural network components (such as batch normalization layers) cause performance issues and are harder to adjust in a federated setup. In contrast, components like convolutional layers could be easily averaged, averaging their results in a proper global model. As a result, deep learning models that have more suitable components are a better choice for FL. Research is ongoing to develop specific models that perform better in a federated environment [4].

Comparison of the FL Algorithms

We may categorize the algorithms based on what is exchanged between the server and the client to compare federated algorithms. Techniques such as FedAvg, SWT, and CWT transfer the model between the server and the clients. Approaches like split learning [5] transfer middle layer outputs of a neural network. The middle layer outputs can be regarded as a distorted form of the input data. In other words, as the neural network processes the input data, it undergoes numerous modifications that distort the input. Methods such as ensemble methods share their models' final output and broadcast it to a central server.

The amount of data transferred is relatively tiny in methods in which the model is moved to the central server and is independent of the amount of training data at each site. It is solely determined by the size of the deep learning model. The majority of popular deep learning algorithms are tens of megabytes in size. However, an FL algorithm that transfers a model does not necessarily have a low overall communication overhead. The overall amount of exchanged data also depends on the number of communication rounds

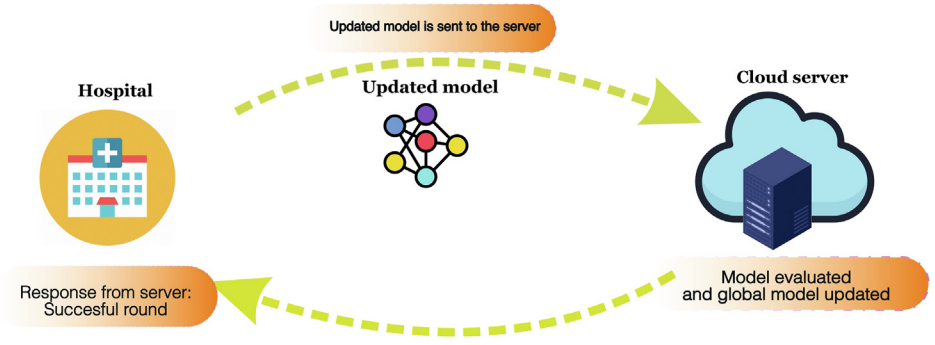

Fig. 1. Cloud server gathers the locally updated model from clients.

between clients and servers. The hyperparameters can determine the number of communication rounds, and communication overhead could be high if there is too much exchange between clients.

In contrast, in algorithms that transfer some type of actual data, whether distorted input data (eg, split learning [5]) or output data (eg, ensemble models), the size of sent data can vary greatly. However, because medical imaging data are enormous, the amount of communicated information is usually more significant than with methods that transfer the model. CDS also falls into this category because it requires actual data transfer to a central server. These two groups differ significantly in terms of communication burden as well as privacy level. Because input and output data are not sent in any format, methods that transfer models are more secure since retrieving patient data from deep learning models is difficult.

In ensemble models, the ensembling process is done locally, and outputs of the models are sent to the global server instead of model parameters. As a result, heavy server-side computations are avoided and a federated network can be set up easily. Because ensemble models are proven to perform well in various areas of medical imaging, using ensembles can help improve the accuracy, generalizability, and stability of a federated network. However, ensemble methods impose some challenges. First, the risk of data leakage is serious in this setting. Some sort of output data like segmentation masks is very likely to reveal patients' identities. Second, contrary to model size, outputs can vary quite a lot in their size. Outputs in image format require too much communication load. In addition, ensemble models are design dependent. Models that do not necessarily share the same objective function can be combined into one ensemble. This leads to one complex multiobjective model having disparate optimization goals. This is not necessarily harmful, but there is a lack of research on the theoretical analysis of ensembles, and the result of an ensemble remains almost always unclear, making ensemble methods unreliable.

Besides, there is always a compromise between training time, model complexity, performance, and generalizability. Although these measures have been thoroughly investigated in single machine learning models, the literature on their relationship in a complex ensemble is still not explored much.

Another aspect of comparing FL models is that FL algorithms, in which a model is transferred, can consistently be averaged by the central server, regardless of the task they are performing. Deep neural networks performing classification, segmentation, regression, or other tasks could be averaged as long as there is a proper deep learning model for that. All the mentioned tasks have been demonstrated and proved to work in a federated manner. However, averaging the output from many sources is not always feasible for other federated earning algorithms. For example, if the task is multiclass classification, an ensemble approach cannot simply average the class output of distinct clients. The ensemble approach is thus limited in the jobs it can tackle.

Several research publications have been published that compare FL implementations. Nilsson et al [6] compared various FL methods in practice. They demonstrated that FedAvg is the best FL algorithm. Despite having slightly lower performance than CDS, it is practically comparable in their comparative performance analysis to a nonfederated architecture. There are numerous variants of the FedAvg algorithm and other FL approaches. However, the original FedAvg method remains one of the top methods in comparison studies. Chang et al [3] investigated several FL algorithms in the radiology area. According to this study, FedAvg does not impose any bias compared with other algorithms because it considers all clients equally and does not arrange them in any particular order. As shown in Figure 2, with algorithms such as SWT and CWT, clients are placed in a sequence and trained one after the other. As a result of catastrophic forgetting, the model is more representative of the most recent clients it observed and less of the earlier clients [7]. As a result, there is a bias favoring the most recent institution in models with sequential training. Although CWT can mitigate this effect by running the model through institutions multiple times in a cyclic fashion, bias remains. Table 1 shows the essential characteristics of FL algorithms. There is also a sample of use cases of these algorithms in the medical domain.

Pan et al [8] investigated the impact of the model ensemble for automatic bone age estimation based on imaging data. The results showed that combining heterogenous, uncorrelated models lead to more robust ensembles. Conversely, naively combining top models does not necessarily ensure top-notch performance. The researchers were able to demonstrate how data FL can aid in identifying comparable patients while protecting their privacy.

CHALLENGES AND CONSIDERATIONS

FL still has a long way to go in radiology. There are numerous challenges in both the theoretical formulation and practical implementation. FL algorithms could be divided into fully decentralized, peer-to-peer methods requiring a trusted central server. Each category comes up with its challenges. Generally speaking, methods with a central server offer more flexibility and better performance, and decentralized methods are more reliable and secure.

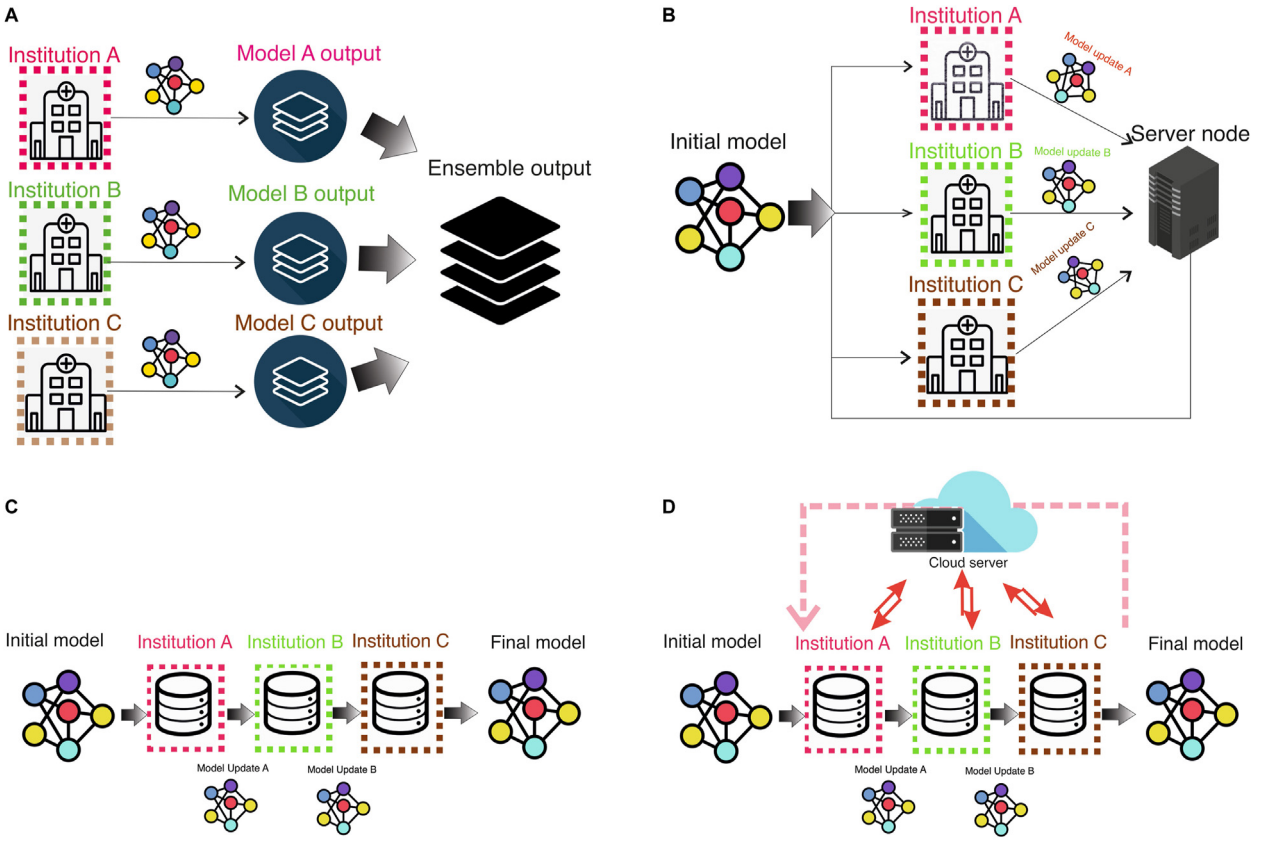

Fig. 2. Different decentralized learning methods. (a) Ensemble methods: Clients train local models on their own data set; model outputs of different clients are averaged. (b) Federated averaging method: An initial model is sent to the clients; each train the model on their own data and the resulting local models are averaged in a central server. (c) Single weight transfer: An initial model is sequentially passed through clients and visit each client once. Final model is the model trained on the latest client. (d) Cyclic weight transfer: similar to single weight transfer, but model is passed through institutions multiple times.

However, there is still some risk associated with the FL infrastructure [9]. An adversary could reconstruct private data from the local model updates [10]. Hospitals can do additional security measures to prevent adversaries from accessing the exchanged data between the server and clients.

Data Heterogeneity

The FedAvg algorithm authors claim that their proposed method can handle heterogeneous data. However, the decentralized structure of the data makes data processing challenging to verify the completeness and quality of their findings. Further investigations revealed that this claim is not always valid [11]. In almost all cases, heterogenous data deteriorate the accuracy of the FL model. The degree of divergence depends on how heterogenous the data are. Local models are trained on data with different patient profiles, resulting in a global model that could not represent all the profiles. In some cases, heterogenous data prevent model convergence.

Data homogeneity significantly impacts the version of the federated model to be chosen to train the model. The difference between CDS and FL might vary from similar to CDS better depending on the data. One rule of thumb would be that if data have very different distribution in different data centers, simply averaging each client's data in every round might affect the performance negatively.

Zhao et al [12] examined the effect of bias that distributing data can have on final performance of FL algorithms. According to their study, difference in data distribution can have negative effect on the model accuracy up to 55%. Another difficulty is that data heterogeneity might lead to a situation in which a best global model might be a poor model for some clients or a best global model might work quite well on some clients and perform poorly on other clients. Consequently, all participants should agree on the concept of optimum model training in advance of the training. Further technical studies should be carried out to find the

Table 1. Comparison of FL methods

Methods	Summary	Transferred Data	Communication Load	Advantages	Disadvantages	Use Cases
FedAvg	In each round, every client trains the global model on local data. Then models are averaged.	Model	Low	Easily converged	Weak robustness with imbalanced client distribution	COVID-19 CT scans [24], lung nodule detection [25]
SWT	Model is passed through clients sequentially; visits each client once.	Model	Very low	Low communication load	Highly biased toward the latest institution	Diabetic retinotherapy [3], mammography [3]
CWT	Model is passed through clients sequentially; the sequence is repeated multiple times.	Model	Low	High performance	Needs many rounds to converge	Breast cancer data [26], EHR [26]
Ensemble methods	All the computations are done locally; the outputs are averaged.	Output	High	Easy to deploy	High possibility of data leakage, high communication load	Patient health records [27]
CDS	Data are moved from clients into a central data center.	Data	-	High performance	No privacy	MRI reconstruction [28], dermoscopy image synthesis [29]

CDS = centralized data sharing; COVID-19 = coronavirus disease 2019; CWT = cyclic weight transfer; EHR = electronic health record; FedAvg = federated averaging; FL = federated learning; SWAT = single weight transfer.

optimum technique for updating the central model with heterogeneous data. FedAvg is the standard method for accumulating the data from clients. Still, other distributed optimization methods that can tackle distribution differences are a subject of research.

Bias

Bias is a prevalent issue in distributed networks. Bias is a state that a neural network is inclined toward the distribution of a client more than other clients. It results in the model performing well on that client by compromising the performance on the other clients. The cause of bias could be the difference in the size or the distribution of clients' data. Also, the FL algorithm itself could be a source of bias.

Sheller et al [7] showed that CWT is a less biased algorithm than SWT. The degree of bias could vary, depending on which client was trained last. They favored FedAvg over SWT and CWT. FedAvg conducted the FL more fairly. For algorithms like SWT and CWT, there is always a bias toward the latest clients on whom they were trained. In FedAvg, however, the results of local training are aggregated every round, avoiding bias. In SWT, the global model changes after visiting each client, and succeeding clients mitigate the model's bias toward the preceding institutions. However, there is no mitigation for the latest institution the model is trained on.

The global aggregation method (ie, server algorithm) should be designed to minimize bias. It also should be robust to local variations, as well as perturbations added by security measures. Reducing bias and designing models that capture diversity is possible by calculating the level of bias arising from each client, then modifying the algorithm to address the difference in the distributions.

However, if the distribution difference is taken into account properly, bias might still emerge later in training. Some features, as well as general data distribution, might vary over time. For example, the number of patients with a particular disease in a particular hospital might change for several reasons. This can cause a domain shift: a change in a client's data distribution. There could be more work on data domain shifts and somehow explicitly address the alterations in sex, patient profile, age, and disease among different institutes or one institute. Models could also be further developed to consider economic or racial status into model training and modify a model to handle diversity in images [13].

Lack of Standard Data

Standardizing data prevents irrelevant information from being considered meaningful in neural networks. It removes the variability between institutions. Electronic data management is the norm in medical imaging and medical communications, and DICOM is the globally recognized image data format and the near-global care standard for electronic file storage. However, not all the available data in the medical imaging sector are standardized. Many institutions still lack the infrastructure to handle their imaging data according to current management standards. One factor is the lack of a universal method to organize and manage patient records. Data management is expensive [14]. Not all hospitals have advanced data management facilities. This issue leads to the preselection of hospitals participating in research, which is a source of bias.

Medical data are very diverse because of the diversity of modalities, dimensionalities, and features and because of variables such as variances in the acquisition, medical equipment brand, or local demographics within a specific protocol. There is still no uniform data standardization method. As a result, health care federated networks are likely to have clients with disparate data quality and distribution. Methods like FedAvg are generally likely to fail under these circumstances. One way to avoid bias is to harmonize data and make each client data type similar, following a similar preprocessing. This also might require sharing metadata among institutions to find a general method of harmonization in data that suit all the institutions. However, this could be tricky given the restrictions of individual institutions. Hence, one way of further development for FL systems is that the clinicians and computer scientists collaborate to standardize handling data among multiple institutions concerning the privacy restrictions and considerations.

Privacy and Security

Data breach is a major concern, and medical data must be safeguarded in compliance with accepted confidentiality procedures. FL has proven to be effective in protecting patients' privacy and anonymity by keeping data locally. However, there are some privacy-related challenges associated with FL. Despite many attempts to deidentify personal data from DICOM images, patient information could still be re-identified [15,16]. Recent studies have successfully rebuilt a patient's face from the MRI data. Furthermore, adversaries could steal the data or access the algorithm for nonencrypted networks.

Furthermore, deep learning models still have some sensitive information in the weights they carry. On a decentralized network, it is feasible to reconstitute portions of patients' information having only the local model from one client [17-19]. Adversaries can decrypt deep learning models and reveal patients' information with a very high accuracy [20]. Malicious parties can distort the deep learning models. False outputs produced by such models

can have severe consequences if used in practice. As a result, it should be ensured that models are secure and that adversaries cannot breach models to be employed in the real-world setting [21].

There are specific measures to improve privacy. Particular countermeasures, such as model encryptions, differential privacy (DP) [22], adversarial defense against malicious clients [19], and increased communication security, can be done. DP refers to the practice of keeping a data set's global statistical distribution while minimizing individually identifiable information. DP can be done by adding perturbation to each sample. Adding noise to a data set to reduce the chance of private data being revealed is based on the argument that one can preserve general data distribution while individual samples are changed by randomly altering a data set. Adding systematic noise helps machine models to learn the whole distribution of training data while keeping each sample anonymous.

However, such countermeasures complicate the training algorithms and can affect training performance. Sometimes much longer training times are required, or accuracy will be dramatically decreased. This can impose an additional cost on the whole network. Hence, it is quintessential to consider whether deploying a countermeasure is necessary. The cost efficiency of implementing them depends mainly on the level of trust among involved parties and the project scale. If clients do not trust each other, then DP is a necessity. This is because federated clients have regular communication, and critical information can be exchanged in the interactions. So each client's data should be safeguarded from other clients. This shows how important it is to clarify the level of trust among clients. This argument holds true in fully decentralized algorithms, in which no central node is involved, and also in algorithms including a central server, in which the client-server trust is also essential. Total image anonymization is still a problem. In the absence of encryption, attackers may acquire private information from local datacenters or intercept the communication pathways and rob the passing data.

System Architecture

Medical data in federated networks requires on-premise or cloud-based data storage. Hospitals might need private or cloud-based computing power, as well as software for preprocessing and standardization of data, such as PACS. To allow the local model training hardware (graphics processing unit), connections and data centers should be set up in local centers. These bring their challenges, such as high computational power, to ensure harmony with other clients and high-performance bandwidth and connection between different centers, which is not always feasible in medical centers. Many hospitals still lack computing resources and strong Internet connections [23]. Besides, for the whole network to work correctly, redundant computational facilities and data centers should be devised to prevent data loss. If one computational client fails, the network could continue its training, which imposes an additional challenge. Robustness of the network is also critical; federated models should be structured so that adding or removing clients and increasing or decreasing the amount of data in a center does not negatively impact patient data or model privacy.

CONCLUSION

This article introduces the main FL algorithms used in radiology and compared their characteristics. A federated setting faces myriad challenges; designing algorithms to address them results in various algorithms with distinct optimization objectives. In general, developments focus on privacy, communication load, data heterogeneity, and model performance as their objective. This article discusses and compares FL algorithms based on these objectives. We start by introducing FL and its paramount role in medical imaging research. Then we present the most popular FL algorithms and discuss their challenges and considerations These challenges are current lines of research and require extra attention in implementing FL networks.

TAKE-HOME POINTS

- Implementation of FL pipelines for medical imaging can mitigate privacy concerns to a large extent. However, distinct characteristics of medical images and health care institutions can cause specific obstacles that differ significantly from those encountered with other data types.
- Health care institutions generally lack cloud-based or on-premise computational facilities, which are critical for establishing federated networks. They might also need to prepare data management and standardization pipelines and have robust network connections.
- Main functional challenges include bias toward one hospital, data heterogeneity, local model performance, and security issues.
- Several FL algorithms were designed to address those issues. Some promising results enhance privacy, communication load, data heterogeneity, and model performance. Research is ongoing, and a universal solution is yet to come.

ACKNOWLEDGMENTS

This research is supported by KWF Kankerbestrijding and the Netherlands Organisation for Scientific Research Domain AES, under the project number 17924, AI in

Medical Imaging for novel Cancer User Support, as part of their joint strategic research programme: Technology for Oncology IL. The collaboration project is cofunded by the PPP allowance made available by Health Holland, Top Sector Life Sciences Health, to stimulate public-private partnerships.

REFERENCES

1. McMahan B, Moore E, Ramage D, Hampson S, y Arcas BA. Communication-efficient learning of deep networks from decentralized data. In: Proceedings of the 20th international conference on artificial intelligence and statistics, Fort Lauderdale, FL, USA, 54. PMLR; 2017: 1273-82.
2. Wang Z, Hu Q. Blockchain-based federated learning: a comprehensive survey. Preprint. arXiv 2021; 2110.02182.
3. Chang K, Balachandar N, Lam C, et al. Distributed deep learning networks among institutions for medical imaging. J Am Med Inform Assoc 2018;25:945-54.
4. Li X, Jiang M, Zhang X, Kamp M, Dou Q. Fedbn: Federated learning on non-iid features via local batch normalization. Preprint. arXiv 2021; 2102.07623.
5. Poirot MG, Vepakomma P, Chang K, Kalpathy-Cramer J, Gupta R, Raskar R. Split learning for collaborative deep learning in healthcare. Preprint. arXiv 2019; 1912.12115.
6. 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. New York, NY, United States: Association for Computing Machinery; 2018:1-8.
7. Sheller MJ, Edwards B, Reina GA, et al. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. Sci Rep 2020;10:1-12.
8. Pan I, Thodberg HH, Halabi SS, Kalpathy-Cramer J, Larson DB. Improving automated pediatric bone age estimation using ensembles of models from the 2017 rsna machine learning challenge. Radiol Artif Intell 2019;1:e190053.
9. Yin H, Mallya A, Vahdat A, Alvarez JM, Kautz J, Molchanov P. See through gradients: image batch recovery via gradinversion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition,Nashville, TN, USA. IEEE; 2021:16337-46.
10. Wang Z, Song M, Zhang Z, Song Y, Wang Q, Qi H. Beyond inferring class representatives: user-level privacy leakage from federated learning. In: IEEE INFOCOM 2019-IEEE Conference on Computer Communications. IEEE; 2019:2512-20.
11. Li T, Sahu AK, Zaheer M, Sanjabi M, Talwalkar A, Smith V. Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems*, 2. Austin, TX, USA: IEEE; 2020:429-50.
12. Zhao Y, Li M, Lai L, Suda N, Civin D, Chandra V. Federated learning with non-iid data. Preprint. arXiv 2018; 1806.00582.
13. Li X, Gu Y, Dvornek N, Staib LH, Ventola P, Duncan JS. Multi-site fmri analysis using privacy preserving federated learning and domain adaptation: abide results. Med Image Anal 2020;65:101765.
14. Wang SJ, Middleton B, Prosser LA, et al. A cost-benefit analysis of electronic medical records in primary care. Am J Med 2003;114:397-403.
15. Aryanto YE, Oudkerk M, Van Ooijen P. Free DICOM de-identification tools in clinical research: functioning and safety of patient privacy. Eur Radiol 2015;25:3685-95.
16. Monteiro E, Costa C, Oliveira JL. A machine learning methodology for medical imaging anonymization. In: 2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). Milan, Italy: IEEE; 2015:1381-4.
17. Carlini N, Liu C, Erlingsson Ú, Kos J, Song D. The secret sharer: Evaluating and testing unintended memorization in neural networks. In: 28th {USENIX} Security Symposium ({USENIX} Security 19). Santa Clara, CA, USA: USENIX; 2019:267-84.
18. Zhang Y, Jia R, Pei H, Wang W, Li B, Song D. The secret revealer: Generative model-inversion attacks against deep neural networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. Seattle, WA, USA: IEEE; 2020:253-61.
19. Hitaj B, Ateniese B, Perez-Cruz F. Deep models under the gan: information leakage from collaborative deep learning. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security. Dallas, TX, USA: ACM; 2017:603-18.
20. Fredrikson M, Jha S, Ristenpart T. Model inversion attacks that exploit confidence information and basic countermeasures. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security. Denver, CO, USA: ACM; 2015:1322-33.
21. Tomsett R, Chan K, Chakraborty S. Model poisoning attacks against distributed machine learning systems. In: . Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications. International Society for Optics and Photonics; 2019;11006:110061D.
22. Abadi M, Chu A, Goodfellow I, et al. Deep learning with differential privacy. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. New York, NY, United States: Association for Computing Machinery; 2016:308-18.
23. Li M, Yu S, Ren K, Lou W. Securing personal health records in cloud computing: patient-centric and fine-grained data access control in multi-owner settings. In: International Conference on Security and Privacy in Communication Systems. Washington, DC, USA: Springer; 2010:89-106.
24. Dou Q, So TY, Jiang M, et al. Federated deep learning for detecting COVID-19 lung abnormalities in CT: a privacy-preserving multinational validation study. NPJ Digit Med 2021;4:1-11.
25. Baheti P, Sikka M, Arya K, Rajesh R. Federated learning on distributed medical records for detection of lung nodules. In: International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications, VISIGRAPP (4: VISAPP). Valletta, Malta: European Association for Computer Graphics; 2020:445-51.
26. Beaulieu-Jones BK, Yuan W, Finlayson SG, Wu ZS. Privacy-preserving distributed deep learning for clinical data. Preprint 2018; arXiv: 01484.
27. Li Y, Bai C, Reddy CK. A distributed ensemble approach for mining healthcare data under privacy constraints. Inform Sci 2016;330:245-59.
28. Quan TM, Nguyen-Duc T, Jeong W-K. Compressed sensing MRI reconstruction using a generative adversarial network with a cyclic loss. IEEE Trans Med Imaging 2018;37:1488-97.
29. Yi X, Walia E, Babyn PS. Unsupervised and semi-supervised learning with categorical generative adversarial networks assisted by wasserstein distance for dermoscopy image classification. Preprint 2018;03700. arXiv:1804.