

# Federated Learning in Dentistry: Chances and Challenges

Journal of Dental Research  
2022, Vol. 101(11) 1269–1273  
© International Association for Dental Research and American Association for Dental, Oral, and Craniofacial Research 2022



Article reuse guidelines:  
sagepub.com/journals-permissions  
DOI: 10.1177/00220345221108953  
journals.sagepub.com/home/jdr

R. Rischke<sup>1</sup>, L. Schneider<sup>2,3</sup> , K. Müller<sup>1</sup>, W. Samek<sup>1</sup> ,  
F. Schwendicke<sup>2,3</sup> , and J. Krois<sup>2,3</sup> 

## Abstract

Building performant and robust artificial intelligence (AI)-based applications for dentistry requires large and high-quality data sets, which usually reside in distributed data silos from multiple sources (e.g., different clinical institutes). Collaborative efforts are limited as privacy constraints forbid direct sharing across the borders of these data silos. Federated learning is a scalable and privacy-preserving framework for collaborative training of AI models without data sharing, where instead the knowledge is exchanged in form of wisdom learned from the data. This article aims at introducing the established concept of federated learning together with chances and challenges to foster collaboration on AI-based applications within the dental research community.

**Keywords:** computer vision/convolutional neural networks, artificial intelligence, deep learning/machine learning, privacy, dental informatics/bioinformatics, mathematical modeling

## Introduction

The potential of artificial intelligence (AI) to transform health care is vast. AI-based applications in dentistry may help in research, prevention, diagnostics, decision support, and automating routine tasks to facilitate treatment at low cost for more people, eventually allowing for personalized, predictive, preventive, and participatory dentistry (Schwendicke et al. 2020). However, particularly in health care, aspects related to data privacy and data sharing have been identified as hampering factors (Rieke et al. 2020). Especially, dentistry is affected by this, as patients can be identified from anonymized radiographs due to individual structures of tissues, tooth anatomy, and restoration status.

Hence, scalable methods for building AI that respect privacy constraints are required to unlock the full potential of AI-based applications in dentistry. A promising direction is the paradigm shift from centralized data-pooling for the development of AI to federated learning (FL) approaches. FL is closely related to deep learning, which is an AI subfield that aims at training neural networks to extract statistical patterns in given data to eventually make predictions on unseen data. During the so-called training phase, the neural network is iteratively and repeatedly exposed to training data, which consist of data points (e.g., images) with associated expert-based labels (e.g., “healthy” or “caries”). By minimizing the prediction error, the model learns, for instance, to distinguish healthy and decayed teeth in bite-wing radiographs. A common approach for this training process is to first pool data and then train a model on these data, known as centralized learning. Such an approach, however, often lacks generalizability as the data pool stems from one or very few data sources (e.g., 1 or 2 contributing hospitals or research institutes). *Federated learning*, in contrast, enables multiple participants to collaboratively train AI models. FL widens the access

to knowledge from many more and diverse data sources without sharing them directly. Instead, knowledge is exchanged in the form of trained AI models or their outputs (Rieke et al. 2020; Kairouz et al. 2021). This enables generalizability, while keeping the training data private, which is a well-known restriction for medical data analysis tasks. A central server (e.g., a trusted service provider) usually orchestrates the whole FL training process based on defined protocols.

One popular protocol for FL is *federated averaging (FedAvg)*, introduced by McMahan et al. (2017), where all participants agree on an AI model (e.g., a certain neural network architecture), and each participant (e.g., research institute) trains its model on its local training data (e.g., collected bitewing radiographs with labels of “healthy” or “caries”). The updated models are then sent by the participants to the central server, which aggregates all received models to a new global model by averaging all model parameters. This global model carries knowledge originating from data of all participants and is broadcasted to all participants for the next round of this iterative training procedure, unless a certain stopping criterion is met. Figure 1 gives an overview of this communication protocol.

<sup>1</sup>Department of Artificial Intelligence, Fraunhofer Heinrich Hertz Institute, Berlin, Germany

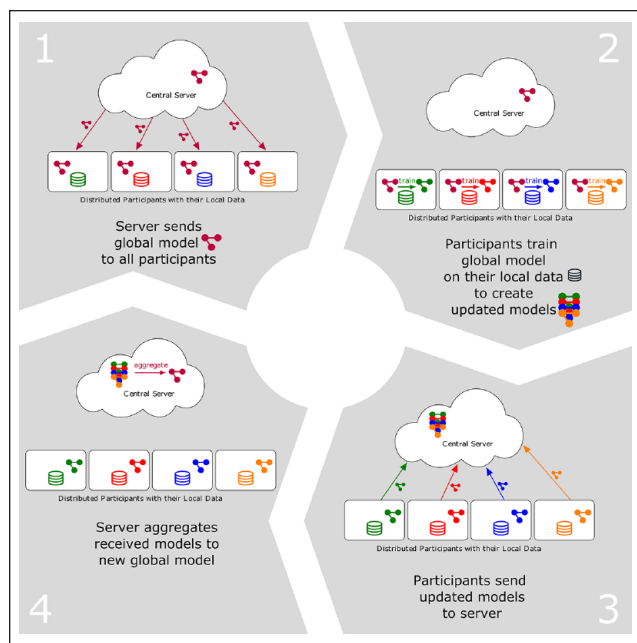
<sup>2</sup>Department of Oral Diagnostics, Digital Health and Health Services Research, Charité-Universitätsmedizin, Berlin, Germany

<sup>3</sup>ITU/WHO Focus Group on AI for Health, Topic Group Dental Diagnostics and Digital Dentistry, Geneva, Switzerland

## Corresponding Author:

L. Schneider, Department of Oral Diagnostics, Digital Health and Health Services Research, Charité-Universitätsmedizin, Aßmannshauer Str. 4-6, Berlin, 14197, Germany.

Email: schneider.lisa@charite.de



**Figure 1.** Algorithmic overview of the iterative federated averaging protocol with 4 steps: 1) server broadcasts artificial intelligence model to all participants; 2) each participant individually trains model on its local data to create an updated model; 3) participants send model updates back to central server; 4) server averages all model updates to aggregate them to a new global model for the next round.

Another established algorithmic paradigm for FL is *federated distillation (FD)* (Jeong et al. 2018; Li and Wang 2019; Sattler, Korjakow, et al. 2021; Sattler, Marban, et al. 2021), where instead of model parameters (as in FedAvg), only model outputs for an unlabeled public data set are exchanged between participants and the central server. This approach allows participants to train heterogeneous model architectures but requires the existence of an appropriate public data set having similar distribution as the training data sets. Silva et al. (2018) published a public data set with 1,500 panoramic images, which may be used for employing FD in dentistry.

FL has already successfully fostered collaborative AI training without data sharing in different health care settings, especially for tasks in medical imaging (Roth et al. 2020; Sheller et al. 2020; Kaissis et al. 2021) and specifically also for international COVID-19 research (Raisaro et al. 2020; Dayan et al. 2021; Yang et al. 2021). The findings of these studies illustrate the value proposition of FL:

1. The performance of all local AI models regarding their predictive quality and generalizability can be improved by collaborative training (Roth et al. 2020; Dayan et al. 2021; Yang et al. 2021).
2. FL has the potential to achieve results of centralized learning (Sheller et al. 2020; Kaissis et al. 2021).
3. There are security mechanisms for preserving the privacy within FL (Raisaro et al. 2020; Kaissis et al. 2021).

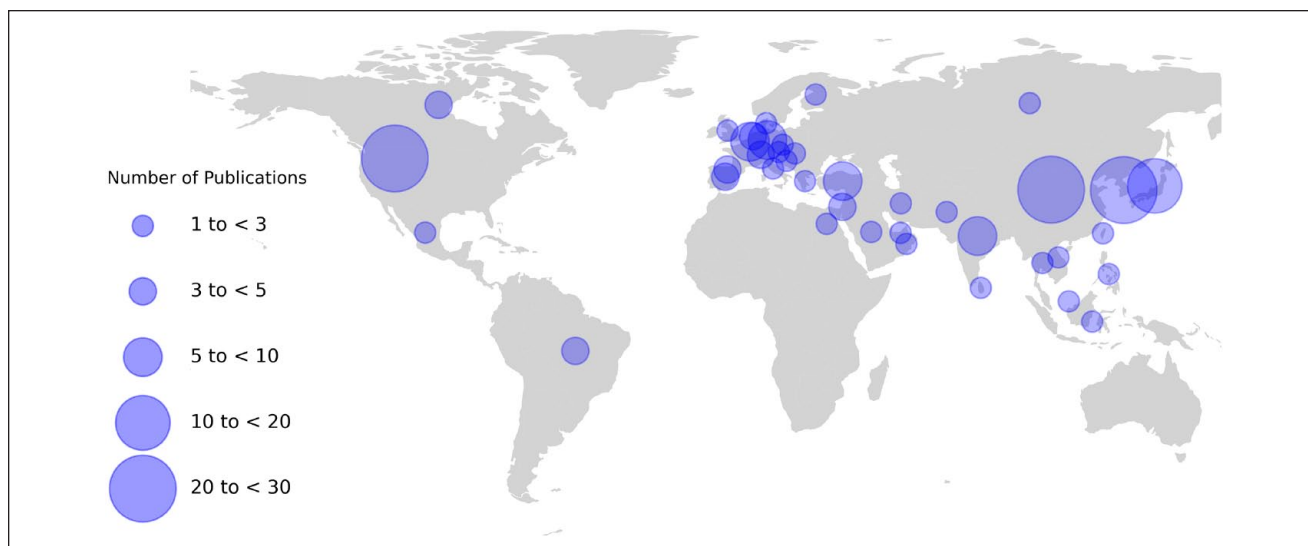
As for dentistry, the above findings and insights apply likewise. Hence, we aim to introduce the FL concept, chances, current challenges, and potential to the dental research community. Embracing FL may foster collaboration and cultivate knowledge exchange, while respecting data privacy concerns, and overall improves AI-based applications in dentistry.

## Chances of Federated Learning in Dentistry

AI is an emerging field in dentistry. In 2020, over 240 AI-related publications were listed on PubMed, which among others predict the occurrence of caries lesions (Lee et al. 2018), periodontal bone loss (Krois et al. 2019), periodontally compromised teeth (Thanathornwong and Suebnukarn 2020), and apical lesions (Ekert et al. 2019). The number of institutes taking active part in this endeavor and their widespread geographic locations are represented in Figure 2. Training such AI-based systems requires large dental data sets, which ideally capture all possible anatomical structures and pathologies. More available data are likely to improve the predictive performance of models, and with more diversity of samples, models are likely to generalize better across different data sources. However, the procurement of such data is difficult as it takes a considerable amount of time, effort, and financial resources to collect and store the data as well as to establish a ground truth. Especially, the latter aspect is often a limiting factor for the size of data sets in dentistry, as most use-cases are based on segmentation or detection tasks, which are generally more time-consuming in ground truth generation than classification tasks. As this holds for many applications, FL is also better understood for classification tasks than segmentation or detection, which necessitates further research in this direction.

In addition, dental data are considered highly sensitive, and data sharing is regulated to preserve the privacy of patients, which restricts the publication of data sets (Van Panhuis et al. 2014). To address this, several techniques for deidentification have been proposed, such as removing patient-specific information (i.e., age or date of birth). Yet, it has been shown that this is not sufficient to protect patients' privacy (Rocher et al. 2019). Especially dental images and radiographs bear a high risk for reidentification. Structures of hard and soft tissue as well as teeth anatomy and dental restorations are unique to individuals, which make radiographs highly relevant for forensic odontology. In this field, reidentification is highly desired: a person, often an unknown deceased individual, may be identified by comparing ante- and postmortem dental radiographs. For example, after the tsunami disaster in 2004, 79% of deceased individuals were identified based on intraoral radiography alone (James et al. 2005). For general data sharing between institutes, these reidentification capabilities of forensic odontology highlight the risk of privacy sensitivity and possible privacy breaches.

To enable joint data-driven research between dental institutes, FL addresses these privacy concerns with its privacy-by-design approach of avoiding raw data exchanges and instead



**Figure 2.** World map with artificial intelligence (AI)-related publications in dentistry (based on the first author’s affiliation) grouped by country. The data stem from a systematic review (unpublished) that screened the online repositories of Medline, IEEE, and arXiv for publications related to AI and dentistry that were published between 2015 and May 2021.

sharing the wisdom learnt from the data, as described in the previous section. A typical setting for dentistry is *cross-silo* FL (Kairouz et al. 2021). Thereby, a few participants (e.g., 2–100 hospitals) have access to large data silos that they cannot share directly with each other due to the described privacy concerns.

In addition to the privacy aspects, FL also offers a solution to data and model ownership concerns present in centralized learning, as in FL data never leave their source and the AI model is trained by all participants. This can be a security advantage, as in FL, there is not only a single institute that an attacker can concentrate on to tamper the training process or get access to private information.

Overall, FL offers great potential to accelerate advances in AI for dentistry. There are already many distinguished gravitational centers of AI research in dentistry (see Fig. 2), so that the provision of a technology that connects them may result in a tremendous development boost for the community. Each research group most likely has access to large amounts of siloed data of varying modalities and may be willing to contribute to collaborative training of dental AI. Dentistry by nature offers beyond the already existing AI research centers many potential FL collaborators, as it is common to store data from small dental offices at large research institutions, which potentially have the required resources to overcome initial investments to join a federation.

### Coping with Challenges for FL in Dentistry

Despite all the chances of FL for dentistry, there are also challenges that require appropriate solutions to unlock the potential of FL. First, FL requires AI preknowledge of the participants, since they agree on a certain model architecture (in FedAvg) or

train on their individual architecture (in FD). Performing an architecture search (Elsken et al. 2019) within FL is computationally very expensive and thus should happen before the joint training phase on the local data to find an appropriate starting model for the considered learning task. Similarly, the tuning process of parameters as well as monitoring and debugging of FL systems are more difficult without direct access to data. Within FL, all participants should have appropriate local computational resources to join the federation, which may be a high initial investment for new participants. Furthermore, these learning resources are costly in maintenance and require higher coordination and effort for deployment than centralized systems.

Second, the participants’ data silos are often statistically heterogeneous by nature due to different medical standards as well as social and economic determinants, also known as data set shift (Quionero-Candela et al. 2009). More formally, the whole training data are usually not distributed in an *iid* (independent and identically distributed) fashion among the participants, which would be fulfilled in an ideal FL setting. Data heterogeneity (“non-*iid*-ness”), however, usually has negative effects on the final model performance in FL and is also a cause of the performance gap to centralized learning (Li et al. 2018; Zhao et al. 2018; Hsu et al. 2019; Kairouz et al. 2021; Sattler, Korjakow, et al. 2021; Sattler, Müller, and Samek 2021). Data heterogeneity in dentistry may be caused by a covariate shift (e.g., different x-ray machines used for data acquisition), prior probability shift (e.g., more implants in developed countries), or unbalancedness (e.g., some regions collect more data), as characterized by Kairouz et al. (2021).

FL applications usually must cope with a mixture of these types, especially when applied at an international scale. However, tasks in dentistry are usually less driven by personal preferences but rather by defined standards (e.g., when

considering caries classification). This is crucial for a feasible FL task, as shown by other success stories in medical imaging outlined in the first section.

There are different strategies for coping with data heterogeneity in FL (see, e.g., Kairouz et al. 2021). General approaches are data set augmentation by publicly available data or sharing some data between participants (Zhao et al. 2018) to make the data silos statistically more homogeneous as well as modifying existing FL algorithms together with their hyperparameters for more effective and efficient training in the presence of data heterogeneity (Li et al. 2018; Li and Wang 2019; Karimireddy et al. 2020; Sattler, Korjakow, et al. 2021). If the assumption that the participants share a similar probability distribution is not fulfilled, meaning that training a single model satisfying all participants is not possible, then more specific strategies are required such as personalization of the global model (Wang et al. 2019), multitask learning (Smith et al. 2017), participant clustering (Sattler, Müller, and Samek 2021), and meta-learning (Khodak et al. 2019).

Third, although FL is a privacy-by-design approach and has, due to its decentralized nature, a lower security risk than centralized learning, vanilla FL protocols without further protection mechanisms have been shown to be vulnerable to certain privacy attacks (Kairouz et al. 2021; Mothukuri et al. 2021). While attacks in centralized learning usually focus on compromising the central server to leak data, privacy attacks in FL rather aim at the reconstruction of the participants' local training data or tampering the training process to leak private information through model inversion attacks (Fredrikson et al. 2015), membership inference attacks (Shokri et al. 2017), or Generative Adversarial Network (GAN) attacks (Hitaj et al. 2017). Thus, privacy within FL systems often also depends on the system's security. Hence, setting up an FL system that is based on trust and secured by cryptography (e.g., secure multiparty computation or homomorphic encryption) is a common approach to preserve privacy (Kairouz et al. 2021; Kaissis et al. 2021; Mothukuri et al. 2021). Another established concept for privacy preservation within FL is differential privacy (Dwork and Roth 2014), where carefully chosen noise is added to the model parameters or model outputs to obfuscate them before release so that it is hard to reconstruct individual training samples. Differential privacy provides a mathematically provable guarantee of privacy but often creates a privacy versus performance trade-off.

Fourth, the value propositions of FL in the first section might already motivate owners of data silos to join FL efforts. Nonetheless, FL requires an investment in local computing resources (including maintenance and coordination) and is based on trust among the participants. Both investments should pay off fairly to motivate participation. Designing mechanisms for FL that incentivize truthful participation and fairness is an interesting open research question (Donahue and Kleinberg 2021; Shi et al. 2021; Witt et al. 2021).

## Call for Action for Dental Researchers

The authors encourage researchers to join their forces to improve AI-based applications in dentistry through collaborative FL

initiatives. To get started, interested parties should get familiar with the FL life cycle and training process as, for instance, described in Kairouz et al. (2021). Based on a solid understanding of the theoretical concepts, one may decide on the data modality, the task to be solved, and potential collaborators for the FL initiative. The topic group "Dental Diagnostics and Digital Dentistry" (TG-Dental) within ITU/WHO Focus Group Artificial Intelligence for Health (FG-AI4H) provides a point of contact to find collaborators. Once the general conditions are defined, all participants must agree on certain settings of the FL process. This includes a decision on the collaboratively trained model architecture. A consensus might be reached by simulations of the FL process, which is not necessary if a well-performing model architecture is already known. Alternatively, participants may use the FD protocol, allowing each participant to train individual architectures. As discussed in the previous section, setting up an FL system can be challenging in practice given the constraints in health care. Thus, the discussed aspects of data heterogeneity, privacy, security, and incentives should be considered already during the initial design phase. Finally, the implementation of the FL training pipelines may be supported by frameworks such as Flower (<https://flower.dev/>), PySyft (<https://github.com/OpenMined/PySyft>), NVIDIA Clara (<https://developer.nvidia.com/clara-medical-imaging>), or TensorFlow Federated (<https://www.tensorflow.org/federated>).

## Conclusion

FL is an established, scalable, and privacy-preserving concept for collaborative AI training without data sharing. Its value propositions have been successfully verified in different medical domains. Although challenges remain to be solved, dentistry should be among the early adopters to use the potential that lies in this concept. This in turn may foster digitalization and standardization efforts in dentistry to address some of the discussed challenges.

## Author Contributions

R. Rischke, L. Schneider, contributed to conception, design, data analysis, and interpretation, drafted and critically revised the manuscript; K. Müller, W. Samek, contributed to conception and design, critically revised the manuscript; F. Schwendicke, contributed to conception, design, data analysis, and interpretation, critically revised the manuscript; J. Krois, contributed to conception, design and data analysis, drafted and critically revised the manuscript. All authors gave final approval and agree to be accountable for all aspects of the work.

## Declaration of Conflicting Interests

The authors declared the following potential conflicts of interest with respect to the research, authorship, and/or publication of this article: F. Schwendicke and J. Krois are cofounders of the dentalXrai Ltd., a startup. dentalXrai Ltd. did not have any role in conceiving, conducting or reporting this study.

## Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

## ORCID iDs

L. Schneider  <https://orcid.org/0000-0002-4431-2669>

W. Samek  <https://orcid.org/0000-0002-6283-3265>

F. Schwendicke  <https://orcid.org/0000-0003-1223-1669>

J. Krois  <https://orcid.org/0000-0002-6010-8940>

## References

- Dayan I, Roth HR, Zhong A, Harouni A, Gentili A, Abidin AZ, Liu A, Costa AB, Wood BJ, Tsai CS, et al. 2021. Federated learning for predicting clinical outcomes in patients with COVID-19. *Nat Med.* 27(10):1735–1743.
- Donahue K, Kleinberg J. 2021. Models of fairness in federated learning. arXiv preprint arXiv:2112.00818. <https://arxiv.org/abs/2112.00818>.
- Dwork C, Roth A. 2014. The algorithmic foundations of differential privacy. *Found Trends Theor Comput Sci.* 9(3–4):211–407.
- Ekert T, Krois J, Meinhold L, Elhennawy K, Emará R, Golla T, Schwendicke F. 2019. Deep learning for the radiographic detection of apical lesions. *J Endod.* 45(7):917–922.
- Elsken T, Metzen JH, Hutter F. 2019. Neural architecture search: a survey. *J Mach Learn Res.* 20(55):1–21.
- Fredrikson M, Jha S, Ristenpart T. 2015. Model inversion attacks that exploit confidence information and basic countermeasures. In: *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security (CCS)*. New York (NY): Association for Computing Machinery. p. 1322–1333.
- Hitaj B, Ateniese G, Perez-Cruz F. 2017. Deep models under the GAN: information leakage from collaborative deep learning. In: *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS)*. New York (NY): Association for Computing Machinery. p. 603–618.
- Hsu TMH, Qi H, Brown M. 2019. Measuring the effects of non-identical data distribution for federated visual classification. arXiv preprint arXiv:1909.06335. <https://arxiv.org/abs/1909.06335>.
- James H. 2005. Thai tsunami victim identification overview to date. *J Forensic Odontostomatol.* 23(1):1–18.
- Jeong E, Oh S, Kim H, Park J, Bennis M, Kim S. 2018. Communication-efficient on-device machine learning: federated distillation and augmentation under non-iid private data. arXiv preprint arXiv:1811.11479. <https://arxiv.org/abs/1811.11479>.
- Kairouz P, McMahan HB, Avent B, Bellet A, Bennis M, Bhagoji AN, Bonawitz K, Charles Z, Cormode G, Cummings R, et al. 2021. Advances and open problems in federated learning. *Found Trends Mach Learn.* 14(1–2):1–210.
- Kaissis G, Ziller A, Passerat-Palmbach J, Ryffel T, Trask A, Lima I Jr, Mancuso J, Jungmann F, Steinborn MM, Saleh A, et al. 2021. End-to-end privacy preserving deep learning on multi-institutional medical imaging. *Nat Mach Intell.* 3(6):473–484.
- Karimireddy SP, Kale S, Mohri M, Reddi S, Stich S, Suresh AT. 2020. SCAFFOLD: Stochastic controlled averaging for federated learning. In: *Proceedings of the 37th International Conference on Machine Learning (ICML)*. Proc Mach Learn Res. 119:5132–5143.
- Khodak M, Balcan MF, Talwalkar A. 2019. Adaptive gradient-based meta-learning methods. *Adv Neural Inf Process Syst.* 32:5915–5926.
- Krois J, Ekert T, Meinhold L, Golla T, Kharbot B, Witte-meier A, Dörfer C, Schwendicke F. 2019. Deep learning for the radiographic detection of periodontal bone loss. *Sci Rep.* 9(1):8495.
- Lee JH, Kim DH, Jeong SN, Choi SH. 2018. Detection and diagnosis of dental caries using a deep learning-based convolutional neural network algorithm. *J Dent.* 77:106–111.
- Li D, Wang J. 2019. Fedmd: heterogeneous federated learning via model distillation. arXiv preprint arXiv:1910.03581. <https://arxiv.org/abs/1910.03581>.
- Li T, Sahu AK, Zaheer M, Sanjabi M, Talwalkar A, Smith V. 2018. Federated optimization in heterogeneous networks. arXiv preprint arXiv:1812.06127. <https://arxiv.org/abs/1812.06127>.
- McMahan B, Moore E, Ramage D, Hampson S, y Arcas BA. 2017. Communication-efficient learning of deep networks from decentralized data. In: *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*. Proc Mach Learn Res. 54:1273–1282.
- Mothukuri V, Parizi RM, Pouriye S, Huang Y, Dehghantanha A, Srivastava G. 2021. A survey on security and privacy of federated learning. *Future Gener Comput Syst.* 115:619–640. doi:10.1016/j.future.2020.10.007
- Quionero-Candela J, Sugiyama M, Schwaighofer A, Lawrence ND. 2009. *Dataset shift in machine learning*. Cambridge (MA): The MIT Press.
- Raisaro JL, Marino F, Troncoso-Pastoriza J, Beau-Lejdstrom R, Bellazzi R, Murphy R, Bernstam EV, Wang H, Bucalo M, Chen Y, et al. 2020. SCOR: a secure international informatics infrastructure to investigate COVID-19. *J Am Med Inform Assoc.* 27(11):1721–1726.
- Rieke N, Hancox J, Li W, Milletari F, Roth HR, Albarqouni S, Bakas S, Galtier MN, Landman BA, Maier-Hein K, et al. 2020. The future of digital health with federated learning. *NPJ Digit Med.* 3:119.
- Rocher L, Hendrickx JM, De Montjoye YA. 2019. Estimating the success of reidentifications in incomplete datasets using generative models. *Nat Commun.* 10(1):3069.
- Roth HR, Chang K, Singh P, Neumark N, Li W, Gupta V, Gupta S, Qu L, Ihsani A, Bizzo B, et al. 2020. Federated learning for breast density classification: a real-world implementation. In: Albarqouni S, Bakas S, Kamnitsas K, Cardoso MJ, Landman B, Li W, Milletari F, Rieke N, Roth H, Xu D, et al., editors. *Domain adaptation and representation transfer and distributed and collaborative learning*. Cham (Switzerland): Springer International Publishing. p. 181–191. doi:10.1007/978-3-030-60548-3\_18
- Sattler F, Korjakow T, Rischke R, Samek W. 2021. Fedaux: leveraging unlabeled auxiliary data in federated learning. *IEEE Trans Neural Netw Learn Syst* [epub ahead of print 1 Dec 2021]. doi:10.1109/TNNLS.2021.3129371
- Sattler F, Marban A, Rischke R, Samek W. 2022. CFD: communication-efficient federated distillation via soft-label quantization and delta coding. *IEEE Trans Netw Sci Eng.* 9(4):2025–2038.
- Sattler F, Müller KR, Samek W. 2021. Clustered federated learning: model-agnostic distributed multitask optimization under privacy constraints. *IEEE Trans Neural Netw Learn Syst.* 32(8):3710–3722.
- Schwendicke F, Samek W, Krois J. 2020. Artificial intelligence in dentistry: chances and challenges. *J Dent Res.* 99(7):769–774.
- Sheller MJ, Edwards B, Reina GA, Martin J, Pati S, Kotrotsou A, Milchenko M, Xu W, Marcus D, Coleen RR, et al. 2020. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Sci Rep.* 10(1):12598.
- Shi Y, Yu H, Leung C. 2021. A survey of fairness-aware federated learning. arXiv preprint arXiv:2111.01872. <https://arxiv.org/abs/2111.01872>.
- Shokri R, Stronati M, Song C, Shmatikov V. 2017. Membership inference attacks against machine learning models. arXiv preprint arXiv:1610.05820. <https://arxiv.org/abs/1610.05820>.
- Silva G, Oliveira L, Pithon M. 2018. Automatic segmenting teeth in x-ray images: trends, a novel data set, benchmarking and future perspectives. *Expert Syst Appl.* 107:15–31. doi:10.1016/j.eswa.2018.04.001
- Smith V, Chiang C, Sanjabi M, Talwalkar AS. 2017. Federated multi-task learning. *Adv Neural Inf Process Syst.* 30:4424–4434.
- Thanathornwong B, Suebnukarn S. 2020. Automatic detection of periodontal compromised teeth in digital panoramic radiographs using faster regional convolutional neural networks. *Imaging Sci Dent.* 50(2):169–174.
- Van Panhuis WG, Paul P, Emerson C, Grefenstette J, Wilder R, Herbst AJ, Heymann D, Burke DS. 2014. A systematic review of barriers to data sharing in public health. *BMC Public Health.* 14:1144.
- Wang K, Mathews R, Kiddon C, Eichner H, Beaufays F, Ramage D. 2019. Federated evaluation of on-device personalization. arXiv preprint arXiv:1910.10252. <https://arxiv.org/abs/1910.10252>.
- Witt L, Zafar U, Shen K, Sattler F, Li D, Samek W. 2021. Reward-based 1-bit compressed federated distillation on blockchain. arXiv preprint arXiv:2106.14265. <https://arxiv.org/abs/2106.14265>.
- Yang D, Xu Z, Li W, Myronenko A, Roth HR, Harmon S, Xu S, Turkbey B, Turkbey E, Wang X, et al. 2021. Federated semi-supervised learning for COVID region segmentation in chest CT using multi-national data from China, Italy, Japan. *Med Image Anal.* 70:101992.
- Zhao Y, Li M, Lai L, Suda N, Civin D, Chandra V. 2018. Federated learning with non-iid data. arXiv preprint arXiv:1806.00582. <https://arxiv.org/abs/1806.00582>.