

FEDERATED LEARNING FROM SMALL DATASETS

Michael Kamp

Institute for AI in medicine (IKIM)
University Hospital Essen, Essen Germany, and
Ruhr-University Bochum, Bochum Germany, and
Monash University, Melbourne, Australia
michael.kamp@uk-essen.de

Jonas Fischer

Harvard T.H. Chan School of Public Health
Department of Biostatistics
Boston, MA, United States
jfisher@hsph.harvard.edu

Jilles Vreeken

CISPA Helmholtz Center for Information Security
Saarbrücken, Germany
vreeken@cispa.de

ABSTRACT

Federated learning allows multiple parties to collaboratively train a joint model without having to share any local data. It enables applications of machine learning in settings where data is inherently distributed and undisclosable, such as in the medical domain. Joint training is usually achieved by aggregating local models. When local datasets are small, locally trained models can vary greatly from a globally good model. Bad local models can arbitrarily deteriorate the aggregate model quality, causing federating learning to fail in these settings. We propose a novel approach that avoids this problem by interleaving model *aggregation* and *permutation* steps. During a permutation step we redistribute local models across clients through the server, while preserving data privacy, to allow each local model to train on a daisy chain of local datasets. This enables successful training in data-sparse domains. Combined with model aggregation, this approach enables effective learning even if the local datasets are extremely small, while retaining the privacy benefits of federated learning.

1 INTRODUCTION

How can we learn *high quality* models when data is *inherently distributed* across sites and cannot be shared or pooled? In federated learning, the solution is to iteratively train models locally at each site and share these models with the server to be aggregated to a global model. As only models are shared, data usually remains undisclosed. This process, however, requires sufficient data to be available at each site in order for the locally trained models to achieve a minimum quality—even a single bad model can render aggregation arbitrarily bad (Shamir and Srebro, 2014). In many relevant applications this requirement is not met: In healthcare settings we often have as little as a few dozens of samples (Granlund et al., 2020; Su et al., 2021; Painter et al., 2020). Also in domains where deep learning is generally regarded as highly successful, such as natural language processing and object detection, applications often suffer from a lack of data (Liu et al., 2020; Kang et al., 2019).

To tackle this problem, we propose a new building block called *daisy-chaining* for federated learning in which models are trained on one local dataset after another, much like a daisy chain. In a nutshell, at each client a model is trained locally, sent to the server, and then—instead of aggregating local models—sent to a random other client as is (see Fig. 1). This way, each local model is exposed to a daisy chain of clients and their local datasets. This allows us to learn from small, distributed datasets simply by consecutively training the model with the data available at each site. Daisy-chaining alone, however, violates privacy, since a client can infer from a model upon the data of the client it received it from (Shokri et al., 2017). Moreover, performing daisy-chaining naively would lead to overfitting which can cause learning to diverge (Haddadpour and Mahdavi, 2019). In this paper, we propose to combine daisy-chaining of local datasets with aggregation of models, both orchestrated by the server, and term this method *federated daisy-chaining* (FEDDC).




Figure 1: *Federated learning settings*. A standard federated learning setting with training of local models at clients (middle) with aggregation phases where models are communicated to the server, aggregated, and sent back to each client (left). We propose to add daisy chaining (right), where local models are sent to the server and then redistributed to a random permutation of clients as is.

We show that our simple, yet effective approach maintains privacy of local datasets, while it provably converges and guarantees improvement of model quality in convex problems with a suitable aggregation method. Formally, we show convergence for FEDDC on non-convex problems. We then show for convex problems that FEDDC succeeds on small datasets where standard federated learning fails. For that, we analyze FEDDC combined with aggregation via the Radon point from a PAC-learning perspective. We substantiate this theoretical analysis for convex problems by showing that FEDDC in practice matches the accuracy of a model trained on the full data of the SUSY binary classification dataset with only 2 samples per client, outperforming standard federated learning by a wide margin. For non-convex settings, we provide an extensive empirical evaluation, showing that FEDDC outperforms naive daisy-chaining, vanilla federated learning FEDAVG (McMahan et al., 2017), FEDPROX (Li et al., 2020a), FEDADAGRAD, FEDADAM, and FEDYOGI (Reddi et al., 2020) on low-sample CIFAR10 (Krizhevsky, 2009), including non-iid settings, and, more importantly, on two real-world medical imaging datasets. Not only does FEDDC provide a wide margin of improvement over existing federated methods, but it comes close to the performance of a gold-standard (centralized) neural network of the same architecture trained on the pooled data. To achieve that, it requires a small communication overhead compared to standard federated learning for the additional daisy-chaining rounds. As often found in healthcare, we consider a cross-SILO scenario where such small communication overhead is negligible. Moreover we show that with equal communication, standard federated averaging still underperforms in our considered settings.

In summary, our contributions are (i) FEDDC, a novel approach to federated learning from small datasets via a combination of model permutations across clients and aggregation, (ii) a formal proof of convergence for FEDDC, (iii) a theoretical guarantee that FEDDC improves models in terms of ϵ, δ -guarantees which standard federated learning can not, (iv) a discussion of the privacy aspects and mitigations suitable for FEDDC, including an empirical evaluation of differentially private FEDDC, and (v) an extensive set of experiments showing that FEDDC substantially improves model quality on small datasets compared to standard federated learning approaches.

2 RELATED WORK

Learning from small datasets is a well studied problem in machine learning. In the literature, we find both general solutions, such as using simpler models and transfer learning (Torrey and Shavlik, 2010), and more specialized ones, such as data augmentation (Ibrahim et al., 2021) and few-shot learning (Vinyals et al., 2016; Prabhu et al., 2019). In our scenario overall data is abundant, but the problem is that data is distributed into small local datasets at each site, which we are not allowed to pool. Hao et al. (2021) propose local data augmentation for federated learning, but their method requires a sufficient quality of the local model for augmentation which is the opposite of the scenario we are considering. Huang et al. (2021) provide generalization bounds for federated averaging via the NTK-framework, but requires one-layer infinite-width NNs and infinitesimal learning rates.

Federated learning and its variants have been shown to learn from incomplete local data sources, e.g., non-iid label distributions (Li et al., 2020a; Wang et al., 2019) and differing feature distributions (Li et al., 2020b; Reisizadeh et al., 2020a), but fail in case of large gradient diversity (Haddadpour and Mahdavi, 2019) and strongly dissimilar label distribution (Marfoq et al., 2021). For small

datasets, local empirical distributions may vary greatly from the global distribution: the difference of empirical to true distribution decreases exponentially with the sample size (e.g., according to the Dvoretzky–Kiefer–Wolfowitz inequality), but for small samples the difference can be substantial, in particular if the distribution differs from a Normal distribution (Kwak and Kim, 2017). Shamir and Srebro (2014) have shown the adverse effect of bad local models on averaging, proving that even due to a single bad model averaging can be arbitrarily bad.

A different approach to dealing with biased local data is by learning personalized models at each client. Such personalized FL (Li et al., 2021) can reduce sample complexity, e.g., by using shared representations (Collins et al., 2021) for client-specific models, e.g., in the medical domain (Yang et al., 2021), or by training sample-efficient personalized Bayesian methods (Achituve et al., 2021). It is not applicable, however, to settings where you are not allowed to learn the biases or batch effects of local clients, e.g., in many medical applications where this would expose sensitive client information. Kiss and Horvath (2021) propose a decentralized and communication-efficient variant of federated learning that migrates models over a decentralized network, storing incoming models locally at each client until sufficiently many models are collected on each client for an averaging step, similar to Gossip federated learning (Jelasi et al., 2005). The variant without averaging is similar to simple daisy-chaining which we compare to in Section 7. FEDDC is compatible with any aggregation operator, including the Radon machine (Kamp et al., 2017), the geometric median (Pillutla et al., 2022), or neuron-clustering (Yurochkin et al., 2019), and can be straightforwardly combined with approaches to improve communication-efficiency, such as dynamic averaging (Kamp et al., 2018), and model quantization (Reisizadeh et al., 2020b). We combine FEDDC with averaging, the Radon machine, and FedProx (Li et al., 2020a) in Sec. 7.

3 PRELIMINARIES

We assume iterative learning algorithms (cf. Chp. 2.1.4 Kamp, 2019) $\mathcal{A} : \mathcal{X} \times \mathcal{Y} \times \mathcal{H} \rightarrow \mathcal{H}$ that update a model $h \in \mathcal{H}$ using a dataset $D \subset \mathcal{X} \times \mathcal{Y}$ from an input space \mathcal{X} and output space \mathcal{Y} , i.e., $h_{t+1} = \mathcal{A}(D, h_t)$. Given a set of $m \in \mathbb{N}$ clients with local datasets $D^1, \dots, D^m \subset \mathcal{X} \times \mathcal{Y}$ drawn iid from a data distribution \mathcal{D} and a loss function $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$, the goal is to find a single model $h^* \in \mathcal{H}$ that minimizes the risk $\varepsilon(h) = \mathbb{E}_{(x,y) \sim \mathcal{D}}[\ell(h(x), y)]$. In *centralized learning*, datasets are pooled as $D = \bigcup_{i \in [m]} D^i$ and \mathcal{A} is applied to D until convergence. Note that applying \mathcal{A} on D can be the application to any random subset, e.g., as in mini-batch training, and convergence is measured in terms of low training loss, small gradient, or small deviation from previous iterate. In standard *federated learning* (McMahan et al., 2017), \mathcal{A} is applied in parallel for $b \in \mathbb{N}$ rounds on each client locally to produce local models h^1, \dots, h^m . These models are then centralized and aggregated using an aggregation operator $\text{agg} : \mathcal{H}^m \rightarrow \mathcal{H}$, i.e., $\bar{h} = \text{agg}(h^1, \dots, h^m)$. The aggregated model \bar{h} is then redistributed to local clients which perform another b rounds of training using \bar{h} as a starting point. This is iterated until convergence of \bar{h} . When aggregating by averaging, this method is known as federated averaging (FEDAVG). Next, we describe FEDDC.

4 FEDERATED DAISY-CHAINING

We propose federated daisy chaining as an extension to federated learning in a setup with m clients and one designated sever.¹ We provide the pseudocode of our approach as Algorithm 1.

The client: Each client trains its local model in each round on local data (line 4), and sends its model to the server every b rounds for aggregation, where b is the aggregation period, and every d rounds for daisy chaining, where d is the daisy-chaining period (line 6). This re-distribution of models results in each individual model conceptually following a daisy chain of clients, training on each local dataset. Such a daisy chain is interrupted by each aggregation round.

The server: Upon receiving models, in a daisy-chaining round (line 9) the server draws a random permutation π of clients (line 10) and re-distributes the model of client i to client $\pi(i)$ (line 11), while in an aggregation round (line 12), the server instead aggregates all local models and re-distributes the aggregate to all clients (line 13-14).

¹This star-topology can be extended to hierarchical networks in a straightforward manner. Federated learning can also be performed in a decentralized network via gossip algorithms (Jelasi et al., 2005).

Algorithm 1: Federated Daisy-Chaining FEDDC

Input: daisy-chaining period d , aggregation period b , learning algorithm \mathcal{A} , aggregation operator agg , m clients with local datasets D^1, \dots, D^m , total number of rounds T

Output: final model aggregate \bar{h}_T

```

1 initialize local models  $h_0^1, \dots, h_0^m$ 
2 Locally at client  $i$  at time  $t$  do
3   sample  $S$  from  $D^i$ 
4    $h_t^i \leftarrow \mathcal{A}(S, h_{t-1}^i)$ 
5   if  $t \bmod d = d - 1$  or  $t \bmod b = b - 1$  then
6     send  $h_t^i$  to server
7     receive new  $h_t^i$  from server // receives either aggregate  $\bar{h}_t$  or some  $h_t^j$ 
8 At server at time  $t$  do
9   if  $t \bmod d = d - 1$  then // daisy chaining
10    draw permutation  $\pi$  of  $[1, m]$  at random
11    for all  $i \in [m]$  send model  $h_t^i$  to client  $\pi(i)$ 
12  else if  $t \bmod b = b - 1$  then // aggregation
13     $\bar{h}_t \leftarrow \text{agg}(h_t^1, \dots, h_t^m)$ 
14    send  $\bar{h}_t$  to all clients

```

Communication complexity: Note that we consider cross-SILO settings, such as healthcare, where communication is not a bottleneck and, hence, restrict ourselves to a brief discussion in the interest of space. Communication between clients and server happens in $O(\frac{T}{d} + \frac{T}{b})$ many rounds, where T is the overall number of rounds. Since FEDDC communicates every d th and b th round, the amount of communication rounds is similar to FEDAVG with averaging period $b_{\text{FedAvg}} = \min\{d, b\}$. That is, FEDDC increases communication over FEDAVG by a constant factor depending on the setting of b and d . The amount of communication per communication round is linear in the number of clients and model size, similar to federated averaging. We investigate the performance of FEDAVG provided with the same communication capacity as FEDDC in our experiments and in App. A.3.6.


5 THEORETICAL GUARANTEES

In this section, we formally show that FEDDC converges for averaging. We, further, provide theoretical bounds on the model quality in convex settings, showing that FEDDC has favorable generalization error in low sample settings compared to standard federated learning. More formally, we first show that under standard assumptions on the empirical risk, it follows from a result of [Yu et al. \(2019\)](#) that FEDDC converges when using averaging as aggregation and SGD for learning—a standard setting in, e.g., federated learning of neural networks. We provide all proofs in the appendix.

Corollary 1. *Let the empirical risks $\mathcal{E}_{\text{emp}}^i(h) = \sum_{(x,y) \in D^i} \ell(h_i(x), y)$ at each client $i \in [m]$ be L -smooth with σ^2 -bounded gradient variance and G^2 -bounded second moments, then FEDDC with averaging and SGD has a convergence rate of $\mathcal{O}(1/\sqrt{mT})$, where T is the number of local updates.*

Since model quality in terms of generalization error does not necessarily depend on convergence of training ([Haddadpour and Mahdavi, 2019](#); [Kamp et al., 2018](#)), we additionally analyze model quality in terms of probabilistic worst-case guarantees on the generalization error ([Shalev-Shwartz and Ben-David, 2014](#)). The average of local models can yield as bad a generalization error as the worst local model, hence, using averaging as aggregation scheme in standard federated learning can yield arbitrarily bad results (cf. [Shamir and Srebro, 2014](#)). As the probability of bad local models starkly increases with smaller sample sizes, this trivial bound often carries over to our considered practical settings. The Radon machine ([Kamp et al., 2017](#)) is a federated learning approach that overcomes these issues for a wide range of learning algorithms and allows us to analyze (non-trivial) quality bounds of aggregated models under the assumption of convexity. Next, we show that FEDDC can improve model quality for small local datasets where standard federated learning fails to do so.

A Radon point ([Radon, 1921](#)) of a set of points S from a space \mathcal{X} is—similar to the geometric median—a point in the convex hull of S with a high centrality (i.e., a Tukey depth ([Tukey, 1975](#);



(a) FEDDC with Radon point with $d = 1, b = 50$. (b) Federated learning with Radon point with $b = 1$. (c) Federated learning with Radon point with $b = 50$.

Figure 2: *Results on SUSY*. We visualize results in terms of train (green) and test error (orange) for (a) FEDDC ($d = 1, b = 50$) and standard federated learning using Radon points for aggregation with (b) $b = 1$, i.e., the same amount of communication as FEDDC, and (c) $b = 50$, i.e., the same aggregation period as FEDDC. The network has 441 clients with 2 data points per client. The performance of a central model trained on all data is indicated by the dashed line.

(Gilad-Bachrach et al., 2004) of at least 2). For a Radon point to exist, $S \subset \mathcal{X}$ has to have a minimum size $r \in \mathbb{N}$ called the Radon number of \mathcal{X} . For $\mathcal{X} \subseteq \mathbb{R}^d$ the radon number is $d + 2$. Here, the set of points S are the local models, or more precisely their parameter vectors. We make the following standard assumption (Von Luxburg and Schölkopf, 2011) on the local learning algorithm \mathcal{A} .

Assumption 2 ((ϵ, δ) -guarantees). *The learning algorithm \mathcal{A} applied on a dataset drawn iid from \mathcal{D} of size $n \geq n_0 \in \mathbb{N}$ produces a model $h \in \mathcal{H}$ s.t. with probability $\delta \in (0, 1]$ it holds for $\epsilon > 0$ that $\mathbb{P}(\varepsilon(h) > \epsilon) < \delta$. The sample size n_0 is monotonically decreasing in δ and ϵ (note that typically n_0 is a polynomial in ϵ^{-1} and $\log(\delta^{-1})$).*

Here $\varepsilon(h)$ is the risk defined in Sec. 3. Now let $r \in \mathbb{N}$ be the Radon number of \mathcal{H} , \mathcal{A} be a learning algorithm as in assumption 2, and risk ε be convex. Assume $m \geq r^h$ many clients with $h \in \mathbb{N}$. For $\epsilon > 0, \delta \in (0, 1]$ assume local datasets D_1, \dots, D_m of size larger than $n_0(\epsilon, \delta)$ drawn iid from \mathcal{D} , and h_1, \dots, h_m be local models trained on them using \mathcal{A} . Let \mathbf{t}_h be the iterated Radon point (Clarkson et al., 1996) with h iterations computed on the local models (for details, see App. A.2). Then it follows from Theorem 3 in Kamp et al. (2017) that for all $i \in [m]$ it holds that

$$\mathbb{P}(\varepsilon(\mathbf{t}_h) > \epsilon) \leq (r \mathbb{P}(\varepsilon(h_i) > \epsilon))^{2^h} \quad (1)$$

where the probability is over the random draws of local datasets. That is, the probability that the aggregate \mathbf{t}_h is bad is doubly-exponentially smaller than the probability that a local model is bad. Note that in PAC-learning, the error bound and the probability of the bound to hold are typically linked, so that improving one can be translated to improving the other (Von Luxburg and Schölkopf, 2011). Eq. 1 implies that the iterated Radon point only improves the guarantee on the confidence compared to that for local models if $\delta < r^{-1}$, i.e. $\mathbb{P}(\varepsilon(\mathbf{t}_h) > \epsilon) \leq (r \mathbb{P}(\varepsilon(h_i) > \epsilon))^{2^h} < (r\delta)^{2^h} < 1$ only holds for $r\delta < 1$. Consequently, local models need to achieve a minimum quality for the federated learning system to improve model quality.

Corollary 3. *Let \mathcal{H} be a model space with Radon number $r \in \mathbb{N}$, ε a convex risk, and \mathcal{A} a learning algorithm with sample size $n_0(\epsilon, \delta)$. Given $\epsilon > 0$ and any $h \in \mathbb{N}$, if local datasets D_1, \dots, D_m with $m \geq r^h$ are smaller than $n_0(\epsilon, r^{-1})$, then federated learning using the Radon point does not improve model quality in terms of (ϵ, δ) -guarantees.*

In other words, when using aggregation by Radon points alone, an improvement in terms of (ϵ, δ) -guarantees is strongly dependent on large enough local datasets. Furthermore, given $\delta > r^{-1}$, the guarantee can become arbitrarily bad by increasing the number of aggregation rounds.

Federated Daisy-Chaining as given in Alg. 1 permutes local models at random, which is in theory equivalent to permuting local datasets. Since the permutation is drawn at random, the amount of permutation rounds T necessary for each model to observe a minimum number of distinct datasets k with probability $1 - \rho$ can be given with high probability via a variation of the coupon collector problem as $T \geq d \frac{m}{\rho^{\frac{1}{m}}} (H_m - H_{m-k})$, where H_m is the m -th harmonic number—see Lm. 5 in

App. A.5 for details. It follows that when we perform daisy-chaining with m clients and local datasets of size n for at least $dm\rho^{-\frac{1}{m}}(H_m - H_{m-k})$ rounds, then each local model will with probability at least $1 - \rho$ be trained on at least kn distinct samples. For an ϵ, δ -guarantee, we thus need to set b large enough so that $kn \geq n_0(\epsilon, \sqrt{\delta})$ with probability at least $1 - \sqrt{\delta}$. This way, the failure probability is the product of not all clients observing k distinct datasets and the model having a risk larger than ϵ , which is $\sqrt{\delta}\sqrt{\delta} = \delta$.

Proposition 4. *Let \mathcal{H} be a model space with Radon number $r \in \mathbb{N}$, ϵ a convex risk, and \mathcal{A} a learning algorithm with sample size $n_0(\epsilon, \delta)$. Given $\epsilon > 0$, $\delta \in (0, r^{-1})$ and any $h \in \mathbb{N}$, and local datasets D_1, \dots, D_m of size $n \in \mathbb{N}$ with $m \geq r^h$, then Alg. 1 using the Radon point with aggr. period*

$$b \geq d \frac{m}{\delta^{\frac{1}{2m}}} \left(H_m - H_{m - \lceil n^{-1} n_0(\epsilon, \sqrt{\delta}) \rceil} \right) \quad (2)$$

improves model quality in terms of (ϵ, δ) -guarantees.

This result implies that if enough daisy-chaining rounds are performed in-between aggregation rounds, federated learning via the iterated Radon point improves model quality in terms of (ϵ, δ) -guarantees: the resulting model has generalization error smaller than ϵ with probability at least $1 - \delta$. Note that the aggregation period cannot be arbitrarily increased without harming convergence. To illustrate the interplay between these variables, we provide a numerical analysis of Prop. 4 in App. A.5.1.

This theoretical result is also evident in practice, as we show in Fig. 2. There, we compare FEDDC with standard federated learning and equip both with the iterated Radon point on the SUSY binary classification dataset (Baldi et al., 2014). We train a linear model on 441 clients with only 2 samples per client. After 500 rounds FEDDC daisy-chaining every round ($d = 1$) and aggregating every fifty rounds ($b = 50$) reached the test accuracy of a gold-standard model that has been trained on the centralized dataset (ACC=0.77). Standard federated learning with the same communication complexity using $b = 1$ is outperformed by a large margin (ACC=0.68). We additionally provide results of standard federated learning with $b = 50$ (ACC=0.64), which shows that while the aggregated models perform reasonable, the standard approach heavily overfits on local datasets if not pulled to a global average in every round. More details on this experiment can be found in App. A.3.2. In Sec. 7 we show that the empirical results for averaging as aggregation operator are similar to those for the Radon machine. First, we discuss the privacy-aspects of FEDDC.

6 DATA PRIVACY

A major advantage of federated over centralized learning is that local data remains undisclosed to anyone but the local client, only model parameters are exchanged. This provides a natural benefit to data privacy, which is the main concern in applications such as healthcare. However, an attacker can make inferences about local data from model parameters (Ma et al., 2020) and model updates or gradients (Zhu and Han, 2020). In the daisy-chaining rounds of FEDDC clients receive a model that was directly trained on the local data of another client, instead of a model aggregate, potentially facilitating membership inference attacks (Shokri et al., 2017)—reconstruction attacks (Zhu and Han, 2020) remain difficult because model updates cannot be inferred since the server randomly permutes the order of clients in daisy-chaining rounds.

Should a malicious client obtain model updates through additional attacks, a common defense is applying appropriate clipping and noise before sending models. This guarantees ϵ, δ -differential privacy for local data (Wei et al., 2020) at the cost of a slight-to-moderate loss in model quality. This technique is also proven to defend against backdoor and poisoning attacks (Sun et al., 2019). Moreover, FEDDC is compatible with standard defenses against such attacks, such as noisy or robust aggregation (Liu et al., 2022)—FEDDC with the Radon machine is an example of robust aggregation. We illustrate the effectiveness of FEDDC




Figure 3: *Differential privacy results.* Comparison of FEDDC (top solid line) to FEDDC with clipped parameter updates and Gaussian noise (dashed lines) on CIFAR10 with 250 clients.




Figure 4: *Synthetic data results.* Comparison of FEDDC (a), FEDAVG with same communication (b) and same averaging period (c) for training fully connected NNs on synthetic data. We report mean and confidence accuracy per client in color and accuracy of central learning as dashed black line.

with differential privacy in the following experiment. We train a small ResNet on 250 clients using FEDDC with $d = 2$ and $b = 10$, postponing the details on the experimental setup to App. A.1.1 and A.1.2. Differential privacy is achieved by clipping local model updates and adding Gaussian noise as proposed by Geyer et al. (2017). The results as shown in Figure 3 indicate that the standard trade-off between model quality and privacy holds for FEDDC as well. Moreover, for mild privacy settings the model quality does not decrease. That is, FEDDC is able to robustly predict even under differential privacy. We provide an extended discussion on the privacy aspects of FEDDC in App. A.7.

7 EXPERIMENTS ON DEEP LEARNING

Our approach FEDDC, both provably and empirically, improves model quality when using Radon points as aggregation which, however, require convex problems. For non-convex problems, in particular deep learning, averaging is the state-of-the-art aggregation operator. We, hence, evaluate FEDDC with averaging against the state of the art in federated learning on synthetic and real world data using neural networks. As baselines, we consider federated averaging (FEDAVG) (McMahan et al., 2017) with optimal communication, FEDAVG with equal communication as FEDDC, and simple daisy-chaining without aggregation. We further consider the 4 state-of-the-art methods FEDPROX (Li et al., 2020a), FEDADAGRAD, FEDYOGI, and FEDADAM (Reddi et al., 2020). As datasets we consider a synthetic classification dataset, image classification in CIFAR10 (Krizhevsky, 2009), and two real medical datasets: MRI scans for brain tumors² and chest X-rays for pneumonia³. We provide additional results on MNIST in App. A.3.8. Details on the experimental setup are in App. A.1.1, A.1.2, code is publicly available at <https://github.com/kampmichael/FedDC>.

Synthetic Data: We first investigate the potential of FEDDC on a synthetic binary classification dataset generated by the sklearn (Pedregosa et al., 2011) `make_classification` function with 100 features. On this dataset, we train a simple fully connected neural network with 3 hidden layers on $m = 50$ clients with $n = 10$ samples per client. We compare FEDDC with daisy-chaining period $d = 1$ and aggregation period $b = 200$ to FEDAVG with the same amount of communication $b = 1$ and the same averaging period $b = 200$. The results presented in Fig. 4 show that FEDDC achieves a test accuracy of 0.89. This is comparable to centralized training on all data which achieves a test accuracy of 0.88. It substantially outperforms both FEDAVG setups, which result in an accuracy of 0.80 and 0.76. Investigating the training of local models between aggregation periods reveals that the main issue of FEDAVG is overfitting of local clients, where FEDAVG *train* accuracy reaches 1.0 quickly after each averaging step. With these promising results on vanilla neural networks, we next turn to real-world image classification problems typically solved with CNNs.

CIFAR10: As a first challenge for image classification, we consider the well-known CIFAR10 image benchmark. We first investigate the effect of the aggregation period b on FEDDC and FEDAVG, separately optimizing for an optimal period for both methods. We use a setting of 250 clients with

²kaggle.com/navoneel/brain-mri-images-for-brain-tumor-detection

³kaggle.com/praveengovi/coronahack-chest-xraydataset

a small version of ResNet, and 64 local samples each, which simulates our small sample setting, drawn at random without replacement (details in App. A.1.2). We report the results in Figure 5 and set the period for FEDDC to $b = 10$, and consider federated averaging with periods of both $b = 1$ (equivalent communication to FEDDC with $d = 1, b = 10$) and $b = 10$ (less communication than FEDDC by a factor of 10) for all subsequent experiments.

Next, we consider a subset of 9600 samples spread across 150 clients (i.e. 64 samples per client), which corresponds to our small sample setting. Now, each client is equipped with a larger, untrained ResNet18.⁴ Note that the combined amount of examples is only one fifth of the original training data, hence we cannot expect typical CIFAR10 performance. To obtain a gold standard for comparison, we run centralized learning CENTRAL, separately optimizing its hyperparameters, yielding an accuracy of around 0.65. All results are reported in Table 1, where we report FEDAVG with $b = 1$ and $b = 10$, as these were the best performing settings and $b = 1$ corresponds to equal amounts

of communication as FEDDC. We use a daisy chaining period of $d = 1$ for FEDDC throughout all experiments for consistency, and provide results for larger daisy chaining periods in App. A.3.5, which, depending on the data distribution, might be favorable. We observe that FEDDC achieves substantially higher accuracy over the baseline set by federated averaging. In App. A.3.7 we show that this holds also for client subsampling. Upon further inspection, we see that FEDAVG drastically overfits, achieving training accuracies of 0.97 (App. A.3.1), a similar trend as on the synthetic data before. Daisy-chaining alone, apart from privacy issues, also performs worse than FEDDC. Intriguingly, also the state of the art shows similar trends. FEDPROX, run with optimal $b = 10$ and $\mu = 0.1$, only achieves an accuracy of 0.51 and FEDADAGRAD, FEDYOGI, and FEDADAM show even worse performance of around 0.22, 0.31, and 0.34, respectively. While applied successfully on large-scale data, these methods seem to have shortcomings when it comes to small sample regimes.

To model different data distributions across clients that could occur in for example our healthcare setting, we ran further experiments on simulated non-iid data, gradually increasing the locally available data, as well as on non-privacy preserving decentralized learning. We investigate the effect of non-iid data on FEDDC by studying the “pathological non-IID partition of the data” (McMahan et al., 2017). Here, each client only sees examples from 2 out of the 10 classes of CIFAR10. We again use a subset of the dataset. The results in Tab. 2 show that FEDDC outperforms FEDAVG by a wide margin. It also outperforms FEDPROX, a method specialized on heterogeneous datasets in our considered small sample setting. For a similar training setup as before, we show results for gradually increasing local datasets in App. A.3.4. Most notably, FEDDC outperforms FEDAVG even with 150 samples locally. Only when the full CIFAR10 dataset is distributed across the clients, FEDAVG is on par with FEDDC (see App. Fig. 7). We also compare with distributed training through gradient sharing (App. A.3.3), which discards any privacy concerns, implemented by mini-batch SGD with parameter settings corresponding to our federated setup as well as a separately optimized version. The results show that such an approach is outperformed by both FEDAVG as well as FEDDC, which is in line with previous findings and emphasize the importance of model aggregation.

As a final experiment on CIFAR10, we consider daisy-chaining with different combinations of aggregation methods, and hence its ability to serve as a building block that can be combined with other federated learning approaches. In particular, we consider the same setting as before and combine FEDPROX with daisy chaining. The results, reported in Tab. 2, show that this combination is not only successful, but also outperforms all others in terms of accuracy.

Medical image data: Finally, we consider two real medical image datasets representing actual health related machine learning tasks, which are naturally of small sample size. For the brain MRI scans, we simulate 25 clients (e.g., hospitals) with 8 samples each. Each client is equipped with a CNN




Figure 5: *Averaging periods on CIFAR10.* For 150 clients with small ResNets and 64 samples per client, we visualize the test accuracy (higher is better) of FEDDC and FEDAVG for different aggregation periods b .

⁴Due to hardware restrictions we are limited to training 150 ResNets, hence 9600 samples across 150 clients.

	CIFAR10	MRI	Pneumonia
FEDDC (ours)	62.9 \pm 0.02	78.4 \pm 0.61	83.2 \pm 0.84
DC (baseline)	58.4 \pm 0.85	57.7 \pm 1.57	79.8 \pm 0.99
FEDAVG (b=1)	55.8 \pm 0.78	74.1 \pm 1.68	80.1 \pm 1.53
FEDAVG (b=10)	48.7 \pm 0.87	75.6 \pm 1.18	79.4 \pm 1.11
FEDPROX	51.1 \pm 0.80	76.5 \pm 0.50	80.0 \pm 0.36
FEDADAGRAD	21.8 \pm 0.01	45.7 \pm 1.25	62.5 \pm 0.01
FEDYOGI	31.4 \pm 4.37	71.3 \pm 1.62	77.6 \pm 0.64
FEDADAM	34.0 \pm 0.23	73.8 \pm 1.98	73.5 \pm 0.36
CENTRAL	65.1 \pm 1.44	82.1 \pm 1.00	84.1 \pm 3.31

Table 1: Results on image data, reported is the average test accuracy of the final model over three runs (\pm denotes maximum deviation from the average).

	CIFAR10
FEDDC	62.9 \pm 0.02
FEDDC +FEDPROX	63.2 \pm 0.38
Non-IID	
FEDDC	34.2 \pm 0.61
FEDAVG (b=1)	30.2 \pm 2.11
FEDAVG (b=10)	24.9 \pm 1.95
FEDPROX	32.8 \pm 0.00
FEDADAGRAD	11.7 \pm 0.00
FEDADAM	13.0 \pm 0.00
FEDYOGI	12.5 \pm 0.04

Table 2: Combination of FEDDC with FEDAVG and FEDPROX and non-iid results on CIFAR10.

(see App. A.1.1). The results for brain tumor prediction evaluated on a test set of 53 of these scans are reported in Table 1. Overall, FEDDC performs best among the federated learning approaches and is close to the centralized model. Whereas FEDPROX performed comparably poorly on CIFAR10, it now outperforms FEDAVG. Similar to before, we observe a considerable margin between all competing methods and FEDDC. To investigate the effect of skewed distributions of sample sizes across clients, such as smaller hospitals having less data than larger ones, we provide additional experiments in App. A.3.5. The key insight is that also in these settings, FEDDC outperforms FEDAVG considerably, and is close to its performance on the unskewed datasets.

For the pneumonia dataset, we simulate 150 clients training ResNet18 (see App. A.1.1) with 8 samples per client, the hold out test set are 624 images. The results, reported in Table 1, show similar trends as for the other datasets, with FEDDC outperforming all baselines and the state of the art, and being within the performance of the centrally trained model. Moreover it highlights that FEDDC enables us to train a ResNet18 to high accuracy with as little as 8 samples per client.

8 DISCUSSION AND CONCLUSION

We propose to combine daisy-chaining and aggregation to effectively learn high quality models in a federated setting where only little data is available locally. We formally prove convergence of our approach FEDDC, and for convex settings provide PAC-like generalization guarantees when aggregating by iterated Radon points. Empirical results on the SUSY benchmark underline these theoretical guarantees, with FEDDC matching the performance of centralized learning. Extensive empirical evaluation shows that the proposed combination of daisy-chaining and aggregation enables federated learning from small datasets in practice. When using averaging, we improve upon the state of the art for federated deep learning by a large margin for the considered small sample settings. Last but not least, we show that daisy-chaining is not restricted to FEDDC, but can be straight-forwardly included in FEDAVG, Radon machines, and FEDPROX as a building block, too.

FEDDC permits differential privacy mechanisms that introduce noise on model parameters, offering protection against membership inference, poisoning and backdoor attacks. Through the random permutations in daisy-chaining rounds, FEDDC is also robust against reconstruction attacks. Through the daisy-chaining rounds, we see a linear increase in communication. As we are primarily interested in healthcare applications, where communication is not a bottleneck, such an increase in communication is negligible. Importantly, FEDDC outperforms FEDAVG in practice also when both use the same amount of communication. Improving the communication efficiency considering settings where bandwidth is limited, e.g., model training on mobile devices, would make for engaging future work.

We conclude that daisy-chaining lends itself as a simple, yet effective building block to improve federated learning, complementing existing work to extend to settings where little data is available per client. FEDDC, thus, might offer a solution to the open problem of federated learning in healthcare, where very few, undisclosable samples are available at each site.

ACKNOWLEDGMENTS

The authors thank Sebastian U. Stich for his detailed comments on an earlier draft. Michael Kamp received support from the Cancer Research Center Cologne Essen (CCCE). Jonas Fischer is supported by a grant from the US National Cancer Institute (R35CA220523).

REFERENCES

- Idan Achituve, Aviv Shamsian, Aviv Navon, Gal Chechik, and Ethan Fetaya. Personalized federated learning with gaussian processes. In *Advances in Neural Information Processing Systems*, volume 34. Curran Associates, Inc., 2021. 3
- Giuseppe Ateniese, Luigi V Mancini, Angelo Spognardi, Antonio Villani, Domenico Vitali, and Giovanni Felici. Hacking smart machines with smarter ones: How to extract meaningful data from machine learning classifiers. *International Journal of Security and Networks*, 10(3):137–150, 2015. 22
- Pierre Baldi, Peter Sadowski, and Daniel Whiteson. Searching for exotic particles in high-energy physics with deep learning. *Nature communications*, 5(1):1–9, 2014. 6, 15
- Arjun Nitin Bhagoji, Supriyo Chakraborty, Prateek Mittal, and Seraphin Calo. Analyzing federated learning through an adversarial lens. In *International Conference on Machine Learning*, pages 634–643. PMLR, 2019. 22
- Kenneth L Clarkson, David Eppstein, Gary L Miller, Carl Sturtivant, and Shang-Hua Teng. Approximating center points with iterative radon points. *International Journal of Computational Geometry & Applications*, 6(03):357–377, 1996. 5, 15
- Liam Collins, Hamed Hassani, Aryan Mokhtari, and Sanjay Shakkottai. Exploiting shared representations for personalized federated learning. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 2089–2099. PMLR, 18–24 Jul 2021. 3
- Robin C Geyer, Tassilo Klein, and Moin Nabi. Differentially private federated learning: A client level perspective. *arXiv preprint arXiv:1712.07557*, 2017. 7
- Ran Gilad-Bachrach, Amir Navot, and Naftali Tishby. Bayes and tukey meet at the center point. In *International Conference on Computational Learning Theory*, pages 549–563. Springer, 2004. 5
- Kristin L Granlund, Sui-Seng Tee, Hebert A Vargas, Serge K Lyashchenko, Ed Reznik, Samson Fine, Vincent Laudone, James A Eastham, Karim A Touijer, Victor E Reuter, et al. Hyperpolarized mri of human prostate cancer reveals increased lactate with tumor grade driven by monocarboxylate transporter 1. *Cell metabolism*, 31(1):105–114, 2020. 1
- Farzin Haddadpour and Mehrdad Mahdavi. On the convergence of local descent methods in federated learning. *arXiv preprint arXiv:1910.14425*, 2019. 1, 2, 4
- Weituo Hao, Mostafa El-Khamy, Jungwon Lee, Jianyi Zhang, Kevin J Liang, Changyou Chen, and Lawrence Carin Duke. Towards fair federated learning with zero-shot data augmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3310–3319, 2021. 2
- Baihe Huang, Xiaoxiao Li, Zhao Song, and Xin Yang. FI-ntk: A neural tangent kernel-based framework for federated learning analysis. In *International Conference on Machine Learning*, pages 4423–4434. PMLR, 2021. 2
- Marwa Ibrahim, Mohammad Wedyan, Ryan Alturki, Muazzam A Khan, and Adel Al-Jumaily. Augmentation in healthcare: Augmented biosignal using deep learning and tensor representation. *Journal of Healthcare Engineering*, 2021, 2021. 2
- Márk Jelasity, Alberto Montresor, and Ozalp Babaoglu. Gossip-based aggregation in large dynamic networks. *ACM Transactions on Computer Systems (TOCS)*, 23(3):219–252, 2005. 3

- Michael Kamp. *Black-Box Parallelization for Machine Learning*. PhD thesis, Rheinische Friedrich-Wilhelms-Universität Bonn, Universitäts-und Landesbibliothek Bonn, 2019. 3
- Michael Kamp, Mario Boley, Olana Missura, and Thomas Gärtner. Effective parallelisation for machine learning. In *Advances in Neural Information Processing Systems*, volume 30, pages 6480–6491. Curran Associates, Inc., 2017. 3, 4, 5, 14, 15
- Michael Kamp, Linara Adilova, Joachim Sicking, Fabian Hüger, Peter Schlicht, Tim Wirtz, and Stefan Wrobel. Efficient decentralized deep learning by dynamic model averaging. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 393–409. Springer, 2018. 3, 4
- Bingyi Kang, Zhuang Liu, Xin Wang, Fisher Yu, Jiashi Feng, and Trevor Darrell. Few-shot object detection via feature reweighting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8420–8429, 2019. 1
- Péter Kiss and Tomas Horvath. Migrating models: A decentralized view on federated learning. In *Proceedings of the Workshop on Parallel, Distributed, and Federated Learning*. Springer, 2021. 3
- Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, Toronto, 2009. 2, 7
- Sang Gyu Kwak and Jong Hae Kim. Central limit theorem: the cornerstone of modern statistics. *Korean journal of anesthesiology*, 70(2):144, 2017. 3
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 18
- Qinbin Li, Bingsheng He, and Dawn Song. Model-contrastive federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10713–10722, 2021. 3
- Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. In *Conference on Machine Learning and Systems, 2020a*, 2020a. 2, 3, 7
- Xiaoxiao Li, Meirui Jiang, Xiaofei Zhang, Michael Kamp, and Qi Dou. Fedbn: Federated learning on non-iid features via local batch normalization. In *International Conference on Learning Representations*, 2020b. 2
- Pei Liu, Xuemin Wang, Chao Xiang, and Weiye Meng. A survey of text data augmentation. In *2020 International Conference on Computer Communication and Network Security (CCNS)*, pages 191–195. IEEE, 2020. 1
- Pengrui Liu, Xiangrui Xu, and Wei Wang. Threats, attacks and defenses to federated learning: issues, taxonomy and perspectives. *Cybersecurity*, 5(1):4, 2022. 6, 22
- Chuan Ma, Jun Li, Ming Ding, Howard H Yang, Feng Shu, Tony QS Quek, and H Vincent Poor. On safeguarding privacy and security in the framework of federated learning. *IEEE network*, 34(4): 242–248, 2020. 6, 22
- Othmane Marfoq, Giovanni Neglia, Aurélien Bellet, Laetitia Kameni, and Richard Vidal. Federated multi-task learning under a mixture of distributions. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2021. 2
- Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*, pages 1273–1282, 2017. 2, 3, 7, 8
- Peter Neal. The generalised coupon collector problem. *Journal of Applied Probability*, 45(3): 621–629, 2008. 20

- Corrie A Painter, Esha Jain, Brett N Tomson, Michael Dunphy, Rachel E Stoddard, Beena S Thomas, Alyssa L Damon, Shahrayz Shah, Dewey Kim, Jorge Gómez Tejeda Zañudo, et al. The angiosarcoma project: enabling genomic and clinical discoveries in a rare cancer through patient-partnered research. *Nature medicine*, 26(2):181–187, 2020. 1
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011. 7, 14
- Krishna Pillutla, Sham M Kakade, and Zaid Harchaoui. Robust aggregation for federated learning. *IEEE Transactions on Signal Processing*, 70:1142–1154, 2022. 3
- Viraj Prabhu, Anitha Kannan, Murali Ravuri, Manish Chaplain, David Sontag, and Xavier Amatriain. Few-shot learning for dermatological disease diagnosis. In *Machine Learning for Healthcare Conference*, pages 532–552. PMLR, 2019. 2
- Johann Radon. Mengen konvexer Körper, die einen gemeinsamen Punkt enthalten. *Mathematische Annalen*, 83(1):113–115, 1921. 4, 15
- Sashank J Reddi, Zachary Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konečný, Sanjiv Kumar, and Hugh Brendan McMahan. Adaptive federated optimization. In *International Conference on Learning Representations*, 2020. 2, 7, 14
- Amirhossein Reisizadeh, Farzan Farnia, Ramtin Pedarsani, and Ali Jadbabaie. Robust federated learning: The case of affine distribution shifts. In *Advances in Neural Information Processing Systems*, volume 33, pages 21554–21565. Curran Associates, Inc., 2020a. 2
- Amirhossein Reisizadeh, Aryan Mokhtari, Hamed Hassani, Ali Jadbabaie, and Ramtin Pedarsani. Fedpaq: A communication-efficient federated learning method with periodic averaging and quantization. In *International Conference on Artificial Intelligence and Statistics*, pages 2021–2031. PMLR, 2020b. 3
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014. 4
- Ohad Shamir and Nathan Srebro. Distributed stochastic optimization and learning. In *2014 52nd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 850–857. IEEE, 2014. 1, 3, 4, 16
- Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 3–18. IEEE, 2017. 1, 6, 22
- Xiaoping Su, Xiaofan Lu, Sehrish Khan Bazai, Eva Compérat, Roger Mouawad, Hui Yao, Morgan Rouprêt, Jean-Philippe Spano, David Khayat, Irwin Davidson, et al. Comprehensive integrative profiling of upper tract urothelial carcinomas. *Genome biology*, 22(1):1–25, 2021. 1
- Ziteng Sun, Peter Kairouz, Ananda Theertha Suresh, and H Brendan McMahan. Can you really backdoor federated learning? *arXiv preprint arXiv:1911.07963*, 2019. 6, 22
- Lisa Torrey and Jude Shavlik. Transfer learning. In *Handbook of research on machine learning applications and trends: algorithms, methods, and techniques*, pages 242–264. IGI global, 2010. 2
- John W Tukey. Mathematics and picturing data. In *Proceedings of the International Congress of Mathematics*, volume 2, pages 523–531, 1975. 4
- Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *Advances in neural information processing systems*, volume 29, pages 3630–3638. Curran Associates, Inc., 2016. 2
- Ulrike Von Luxburg and Bernhard Schölkopf. Statistical learning theory: Models, concepts, and results. In *Handbook of the History of Logic*, volume 10, pages 651–706. Elsevier, 2011. 5

- Hongyi Wang, Mikhail Yurochkin, Yuekai Sun, Dimitris Papailiopoulos, and Yasaman Khazaeni. Federated learning with matched averaging. In *International Conference on Learning Representations*, 2019. 2
- Kang Wei, Jun Li, Ming Ding, Chuan Ma, Howard H Yang, Farhad Farokhi, Shi Jin, Tony QS Quek, and H Vincent Poor. Federated learning with differential privacy: Algorithms and performance analysis. *IEEE Transactions on Information Forensics and Security*, 15:3454–3469, 2020. 6, 22
- Qian Yang, Jianyi Zhang, Weituo Hao, Gregory P. Spell, and Lawrence Carin. Flop: Federated learning on medical datasets using partial networks. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, page 3845–3853. Association for Computing Machinery, 2021. 3
- Hao Yu, Sen Yang, and Shenghuo Zhu. Parallel restarted sgd with faster convergence and less communication: Demystifying why model averaging works for deep learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5693–5700, 2019. 4, 19
- Mikhail Yurochkin, Mayank Agarwal, Soumya Ghosh, Kristjan Greenewald, Nghia Hoang, and Yasaman Khazaeni. Bayesian nonparametric federated learning of neural networks. In *International Conference on Machine Learning*, pages 7252–7261. PMLR, 2019. 3
- Chengliang Zhang, Suyi Li, Junzhe Xia, Wei Wang, Feng Yan, and Yang Liu. Batchcrypt: Efficient homomorphic encryption for cross-silo federated learning. In *USENIX Annual Technical Conference*, pages 493–506, 2020. 22
- Ligeng Zhu and Song Han. Deep leakage from gradients. In *Federated learning*, pages 17–31. Springer, 2020. 6

A APPENDIX

A.1 DETAILS ON EXPERIMENTAL SETUP

In this section we provide all details to reproduce the empirical results presented in this paper. Furthermore, the implementation provided at <https://github.com/kampmichael/FedDC> allows to directly reproduce the result. Experiments were conducted on an NVIDIA DGX with six A6000 GPUs.

A.1.1 NETWORK ARCHITECTURES

Here, we detail network architectures considered in our empirical evaluation.

MLP for Synthetic Data A standard multilayer perceptron (MLP) with ReLU activations and three linear layers of size 100,50,20.

Averaging round experiment For this set of experiments we use smaller versions of ResNet architectures with 3 blocks, where the blocks use 16, 32, 64 filters, respectively. In essence, these are smaller versions of the original ResNet18 to keep training of 250 networks feasible.

CIFAR10 & Pneumonia For CIFAR10, we consider a standard ResNet18 architecture, where weights are initialized by a Kaiming Normal and biases are zero-initialized. Each client constructs and initializes a ResNet network separately. For pneumonia, X-ray images are resized to (224, 224).

MRI For the MRI data, we train a small convolutional network of architecture Conv(32)-Batchnorm-ReLU-MaxPool-Conv(64)-Batchnorm-ReLU-MaxPool-Linear, where Conv(x) are convolutional layers with x filters of kernel size 3. The pooling layer uses a stride of 2 and kernel size of 2. The Linear layer is of size 2 matching the number of output classes. All scan images are resized to (150, 150).

A.1.2 TRAINING SETUP

In this section, we give additional information for the training setup for each individual experiment of our empirical evaluation.

SUSY experiments SUSY is a binary classification dataset with 18 features. We train linear models with stochastic gradient descent (learning rate 0.0001, found by grid-search on an independent part of the dataset) on 441 clients, aggregating every 50 rounds via the iterated Radon point (Kamp et al., 2017) with $h = 2$ iterations. FEDDC performs daisy-chaining with period $d = 1$. The test accuracy is evaluated on a test set with 1 000 000 samples drawn iid at random.

Synthetic Data The synthetic binary classification dataset is generated by the sklearn (Pedregosa et al., 2011) `make_classification` function with 100 features of which 20 are informative, 60 are redundant, and 5 are repeated. We generate 3 clusters per class with a class separation of 1.0, a shift of 1.0 and a scale of 3.0. Class labels are randomly flipped with probability 0.02.

Averaging rounds parameter optimization To find a suitable number when averaging should be carried out, we explore $b \in \{1, 10, 20, 50, 100, 200, 500, \infty\}$ on CIFAR10 using 250 clients each equipped with a small ResNet. We assign 64 samples to each client drawn at random (without replacement) from the CIFAR10 training data and use a batch size of 64. For each parameter, we train for $10k$ rounds with SGD using cross entropy loss and initial learning rate of 0.1, multiplying the rate by a factor of .5 every 2500 rounds.

FedAdam, FedAdagrad, and FedYogi We use the standard values for β_1 and β_2 , i.e., $\beta_1 = 0.9$, $\beta_2 = 0.999$, as suggested in Reddi et al. (2020). We optimized learning rate η_l and global learning rate η from the set $\{0.001, 0.01, 0.1, 1.0, 2.0\}$ yielding optimal parameters $\eta_l = 0.1$ and $\eta = 1.0$.

CIFAR10 differential privacy and main experiments We keep the same experimental setup as for hyperparameter tuning, but now use 100 clients each equipped with a ResNet18.

A.2 ITERATED RADON POINTS AND THE RADON MACHINE

The Radon machine (Kamp et al., 2017) aggregates a set $S = h^1, \dots, h^m$ of local models $h^i \in \mathcal{H}$ via the iterated Radon point algorithm (Clarkson et al., 1996). For models with $d \in \mathbb{N}$ parameters, $r = d + 2$ many models are required to compute a single Radon point, where r is called the Radon number of \mathcal{H} . Let $m = r^h$ for some $h \in \mathbb{N}$, then the iterated Radon point aggregates models in h iterations. In each iteration, the set S is partitioned into subsets of size r and the Radon point of each subset is calculated. The final step of each iteration is to replace the set S of models by the set of Radon points. After h iterations, a single Radon point τ_h is obtained as the aggregate.

Radon points can be obtained by solving a system of linear equations of size $r \times r$ (Kamp et al., 2017): In his main theorem, Radon (1921) gives the following construction of a Radon point for a set $S = \{s_1, \dots, s_r\} \subseteq \mathbb{R}^d$. Find a non-zero solution $\lambda \in \mathbb{R}^{|S|}$ for the following linear equations.

$$\sum_{i=1}^r \lambda_i s_i = (0, \dots, 0) \quad , \quad \sum_{i=1}^r \lambda_i = 0$$

Such a solution exists, since $|S| > d + 1$ implies that S is linearly dependent. Then, let I, J be index sets such that for all $i \in I$: $\lambda_i \geq 0$ and for all $j \in J$: $\lambda_j < 0$. Then a Radon point is defined by

$$\tau(\lambda) = \sum_{i \in I} \frac{\lambda_i}{\Lambda} s_i = \sum_{j \in J} \frac{\lambda_j}{\Lambda} s_j \quad ,$$

where $\Lambda = \sum_{i \in I} \lambda_i = -\sum_{j \in J} \lambda_j$. Any solution to this linear system of equations is a Radon point. The equation system can be solved in time r^3 . By setting the first element of λ to one, we obtain a unique solution of the system of linear equations. Using this solution λ , we define the Radon point of a set S as $\tau(S) = \tau(\lambda)$ in order to resolve ambiguity.

A.3 ADDITIONAL EMPIRICAL RESULTS

In Sec. 7 we have shown that FEDDC performs well on benchmark and real-world datasets. In the following we provide additional empirical results, both to investigate the main results more closely, as well as to further investigate the properties of FEDDC. For the CIFAR10 experiment, we investigate training accuracy (App. A.3.1) and present results for distributed mini-batch SGD (App. A.3.3). For the SUSY experiment, we compare to FEDAVG (App. A.3.2). As additional experiments, we investigate the impact of local dataset size (App. A.3.4) and skewed dataset size distributions (App. A.3.5), and analyze the communication-efficiency of FEDDC (App. A.3.6). Finally, we present results on MNIST where FEDDC achieves state-of-the-art accuracy (App. A.3.8).

A.3.1 TRAIN AND TEST ACCURACIES ON CIFAR10:


In Table 3 we provide the accuracies on the entire training set for the final model, together with test accuracies, on CIFAR10. The high training accuracies of FEDAVG (≈ 0.97)—and to a lesser degree FEDPROX (0.96)—indicate overfitting on local data sets. The poor training performance of FEDADAGRAD, FEDYOGI, and FEDADAM hint at insufficient model updates. A possible explanation is that the adaptive learning rate parameter (which is proportional to the sum of past model updates) becomes large quickly, essentially stopping the training process. The likely reason is that due to large differences in local data distributions, model updates after each aggregation round are large.

A.3.2 ADDITIONAL RESULTS ON SUSY

In Sec. 5 we compared FEDDC to federated learning with the iterated Radon point. For completeness, we compare it to FEDAVG as well, i.e., federated learning using averaging on the same SUSY binary classification dataset (Baldi et al., 2014). The results shown in Fig. 6 are in line with the findings in Sec. 5: FEDDC with the iterated Radon point outperforms FEDAVG both with the same amount of communication ($b = 1$) and the same aggregation period ($b = 50$). The results for $b = 50$ show that

	Test	Train
FEDDC (ours)	62.9 \pm 0.02	94.7 \pm 0.52
DC (baseline)	58.4 \pm 0.85	94.1 \pm 2.31
FEDAVG ($b=1$)	55.8 \pm 0.78	97.2 \pm 0.87
FEDAVG ($b=10$)	48.7 \pm 0.87	97.4 \pm 0.23
FEDPROX	51.1 \pm 0.80	95.9 \pm 0.42
FEDADAGRAD	21.8 \pm 0.01	31.7 \pm 0.25
FEDYOGI	31.4 \pm 4.37	72.4 \pm 0.90
FEDADAM	34.0 \pm 0.23	73.9 \pm 0.89

Table 3: Train and test accuracy on CIFAR10 of the final model over three runs (\pm denotes maximum deviation from the average).



(a) FEDDC with Radon point with $d = 1, b = 50$. (b) Federated learning with averaging with $b = 1$. (c) Federated learning with averaging with $b = 50$.

Figure 6: *Results on SUSY*. We visualize results in terms of train (green) and test error (orange) for FEDDC with the iterated Radon point (a) and FEDAVG with $b = 1$ (b) as well as FEDAVG with $b = 50$ (c). The network has 441 clients with 2 data points per client. The performance of a central model trained on all data is indicated by the dashed line.

FEDAVG exhibits the same behavior on local training sets that indicates overfitting. Overall, FEDAVG performs comparably to federated learning with the iterated Radon point. That is, FEDAVG ($b = 1$ and $b = 50$) has an accuracy of 0.67, resp. 0.64, compared to 0.68, resp. 0.64 for federated learning with the iterated Radon point.

A.3.3 COMPARISON WITH DISTRIBUTED MINI-BATCH SGD ON CIFAR10

We compare to distributed mini-batch SGD, i.e., central updates where gradients are computed distributedly on CIFAR10. We use the same setup as for the other experiments, i.e., $m = 150$ clients and a mini-batch size of $B = 64$, so that the effective mini-batch size for each update is $mB = 9600$, the optimal learning rate is $\lambda = 0.01$. Here, mini-batch SGD achieves a test accuracy of 19.47 ± 0.68 . Since a plausible explanation for the poor performance is the large mini-batch size, we compare it to the setting with $B = 1$ to achieve the minimum effective mini-batch size of $B = 150$. The results are substantially improved to an accuracy of 50.14 ± 0.63 , underlining the negative effect of the large batch size in line with the theoretical analysis of Shamir and Srebro (2014). Running it 64 times the number of rounds that FedDC uses improves the accuracy just slightly to 54.3. Thus, even with optimal $B = 1$ and a 64-times slower convergence, mini-batch SGD is outperformed by both FedAvg and FedDC, since it cannot use the more favorable mini-batch size of $B = 64$ on $m = 150$ clients.

A.3.4 LOCAL DATASET SIZE

In our experiments, we used dataset sizes common in the medical domain, e.g., for radiological images. To further investigate the impact of local dataset sizes on the performance of FEDDC wrt. FEDAVG, we evaluate the performance for local dataset sizes ranging from 2 to 256 (given the size of CIFAR10, 256 is the maximum without creating overlap between local datasets). The results in

Fig. 7 show that FEDDC outperforms all baselines for smaller local datasets. Only for as much as 256 examples FEDAVG performs as good as FEDDC.

These results further confirm that FEDDC is capable of handling heterogeneous data: for $n < 10$ the clients only see a subset of labels due to the size of their local datasets (with $n=2$, each client can at most observe two classes). We find this a more natural non-iid setting. These results indicate that the shuffle mechanism indeed mitigates data heterogeneity well. A further study of the impact of non-iid data, a comparison with personalized FL, and potential improvements to the shuffling scheme are interesting directions for future work




Figure 7: Test accuracy wrt. local dataset size on CIFAR10 with 150 clients ($n = 2^i$ for $i \in \{1, \dots, 8\}$) with linear (left) and logarithmic (right) x-axis.

A.3.5 REALISTIC DATASET SIZE DISTRIBUTION

In medical applications, a common scenario is that some hospitals, e.g., university clinics, hold larger datasets, while small clinics, or local doctors’ offices only hold very small datasets. To simulate such a scenario, we draw local dataset sizes for a dataset of size n so that a fraction c of the clients hold only a minimum number of samples n_{min} (the local doctor’s offices), and the other clients have an increasing local dataset size starting from n_{min} until all data is distributed. That is, for clients $i = [1, \dots, m - \lfloor cm \rfloor]$ the dataset sizes are given by $n_{min} + ai$ with

$$a = \frac{2(n - (\lfloor cm \rfloor)n_{min})}{\lfloor (1-c)m \rfloor (\lfloor (1-c)m \rfloor - 1)}$$

We use the MRI brainscan dataset with $c = 0.3$ and $n_{min} = 2$. The results presented in Tab. 4 show that FEDDC performs well in that setting. FEDDC ($d = 4, b = 10$) outperforms all other methods with an accuracy of around 0.81, is similar to FEDDC on equally distributed data (0.79), and is even close to the centralized gold-standard (0.82).

MRI	
FEDDC (d=1, b=10)	74.7 ± 0.61
FEDDC (d=2, b=10)	74.4 ± 1.15
FEDDC (d=4, b=10)	80.6 ± 0.66
FEDDC (d=5, b=10)	77.8 ± 1.42
DC (baseline)	53.9 ± 0.25
FEDAVG (b=1)	70.1 ± 2.59
FEDAVG (b=10)	75.3 ± 2.37
central	79.9 ± 6.23

Table 4: Results for realistic dataset size distribution on MRI, reported is the average test accuracy of the final model over three runs (\pm denotes maximum deviation from the average).

A.3.6 COMMUNICATION EFFICIENCY OF FEDDC

Although communication is not a concern in cross-silo applications, such as healthcare, the communication efficiency of FEDDC is important in classical federated learning applications. We therefore compare FEDDC with varying amounts of communication to FEDAVG on CIFAR10 in Tab. 5. The results show that FEDDC ($d=2$) outperforms FEDAVG ($b=1$), thus outperforming just using half the amount of communication, and FEDDC ($d=5$) performs similar to FEDAVG ($b=1$), thus outperforming using five times less communication. FEDDC with $d = 10$ and $b = 10$ significantly outperforms FEDAVG ($b=10$), which corresponds to the same amount of communication in this low-sample setting.

CIFAR10	
FEDDC ($d=1, b=10$)	62.9 \pm 0.02
FEDDC ($d=2, b=10$)	60.8 \pm 0.65
FEDDC ($d=5, b=10$)	55.4 \pm 0.11
FEDDC ($d=10, b=20$)	53.8 \pm 0.47
FEDAVG ($b=1$)	55.8 \pm 0.78
FEDAVG ($b=10$)	48.7 \pm 0.87
central	65.1 \pm 1.44

Table 5: Communication efficiency of FEDDC compared to FEDAVG., where FEDDC ($d=1, b=10$) and FEDAVG ($b=1$), respectively FEDDC ($d=10, b=20$) and FEDAVG ($b=10$) have the same amount of communication.

A.3.7 CLIENT SUBSAMPLING

A widely used technique to improve communication-efficiency in federated learning is to subsample clients in each communication round. For example, instead of averaging the models of all clients in vanilla FedAvg, only a subset of clients sends their models and receives the average of this subset. By randomly sampling this subset, eventually all clients will participate in the process. The fraction $C \in (0, 1]$ of clients sampled is a hyperparameter. In cross-SILO applications, such as healthcare, communication-efficiency is not relevant and client participation is assumed to be $C = 1.0$. Client subsampling can naturally be used in FEDDC by sampling clients both in daisy-chaining and aggregation rounds. We conducted an experiment on CIFAR10 where we compare FEDDC using subsampling to FEDAVG ($b = 10$). The results in Table 6 show that FEDDC indeed works well with client subsampling and outperforms FEDAVG with $C = 0.2$, similar to full client participation ($C=1.0$). However, due to the restricted flow of information, the training process is slowed. By prolonging training from 10000 to 30000 rounds, FEDDC with $C = 0.2$ reaches virtually the same performance as in full client participation, but with higher variance. The same holds true for FEDAVG.

	$T = 10\,000$		$T = 30\,000$
	$C = 1.0$	$C = 0.2$	$C = 0.2$
FEDDC ($d=1, b=10$)	62.9 \pm 0.02	53.2 \pm 3.42	61.0 \pm 1.07
FEDAVG ($b=10$)	48.7 \pm 0.87	45.9 \pm 6.94	49.3 \pm 5.23

Table 6: Client subsampling of FEDDC compared to FEDAVG on CIFAR10.

A.3.8 ADDITIONAL RESULTS ON MNIST

In order to further demonstrate the efficiency of FEDDC on clients that achieve state-of-the-art performance we perform experiments on the MNIST (LeCun et al., 1998) dataset. We use a CNN with two convolutional layers with max-pooling, followed by two linear layers with 1024, resp. 100 neurons. Centralized training on all 60 000 training samples of MNIST achieves a test accuracy of 0.994 which is similar to the state-of-the-art. The results for $m = 50$ clients in Tab. 7 show that

FEDDC outperforms FEDAVG both with the same amount of communication, i.e., FEDAVG ($b = 1$) and FEDAVG ($b = 10$). In line with the results on CIFAR10 (cf. Fig. 7), the advantage of FEDDC shrinks with increasing local dataset size. Using $n = 1200$, i.e., the full training set distributed over $m = 50$ clients, results in virtually the same performance of FEDDC and FEDAVG, both reaching a test accuracy of around 0.96.

	FEDDC (d=1,b=10)	FEDAVG (b=1)	FEDAVG (b=10)
$n = 8$	87.6	84.4	84.9
$n = 1200$	96.7	96.3	96.5

Table 7: Performance of FEDDC and FEDAVG on MNIST for varying local dataset sizes n .

A.4 PROOF OF CONVERGENCE

Corollary. *Let the empirical risks $\mathcal{E}_{emp}^i(h) = \sum_{(x,y) \in D^i} \ell(h_i(x), y)$ at each client be L -smooth with σ^2 -bounded gradient variance and G^2 -bounded second moments, then FEDDC with averaging and SGD as learning algorithm has a convergence rate of $\mathcal{O}(1/\sqrt{mT})$, where $T \in \mathbb{N}$ is the number of local updates.*

Proof. We assume that each client $i \in [m]$ uses SGD as learning algorithm. In each iteration $t \in [T]$, a client i computes the gradient $\mathbf{G}_t^i = \nabla \ell(h_t^i(x), y)$ with $x, y \in D^i$ drawn randomly and updates the local model $h_{t+1}^i = h_t^i - \gamma G_t^i$, where $\gamma > 0$ denotes the learning rate. FEDDC with daisy-chaining period d and averaging period b , run for T local iterations, computes T/b local gradients at each of the m clients before averaging. Each local gradient is computed on an iid sample from \mathcal{D} , independent of whether local models are permuted. Therefore, FEDDC with averaging and SGD is equivalent to parallel restarted SGD (PR-SGD) (Yu et al., 2019) with b/d times larger local datasets. Yu et al. (2019) analyze the convergence of PR-SGD with respect to the average \bar{h}_t of local models in round t . Since $\mathcal{E}_{emp}^i(h) = \sum_{(x,y) \in D^i} \ell(h_i(x), y)$ at each client be L -smooth with σ^2 -bounded gradient variance and G^2 -bounded second moments, Theorem 1 in Yu et al. (2019) is applicable. It then follows from Corollary 1 (Yu et al., 2019) that for $\gamma = \sqrt{m}/(L\sqrt{T})$ and $b \leq T^{\frac{1}{4}}/m^{\frac{3}{4}}$ it holds that

$$\begin{aligned} \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_{(x,y) \sim \mathcal{D}} \left[\mathcal{E}_{emp}^i(\bar{h}_t(x), y) \right] \right] &\leq \frac{2L}{\sqrt{mT}} \left(\mathbb{E}_{(x,y) \sim \mathcal{D}} \left[\mathcal{E}_{emp}^i(\bar{h}_0(x), y) \right] \right. \\ &\quad \left. - \mathbb{E}_{(x,y) \sim \mathcal{D}} \left[\mathcal{E}_{emp}^i(\bar{h}^*(x), y) \right] \right) \\ &\quad + \frac{1}{\sqrt{mT}} (4G^2 + \sigma^2) \in \mathcal{O} \left(\frac{1}{\sqrt{mT}} \right). \end{aligned}$$

Here, the first expectation is over the draw of local datasets and \bar{h}^* is given by

$$\bar{h}^* = \arg \min_{h \in \mathcal{H}} \mathbb{E}_{(x,y) \sim \mathcal{D}} \left[\mathcal{E}_{emp}(\bar{h}(x), y) \right].$$

Thus, FEDDC with averaging and SGD converges in $\mathcal{O}(1/\sqrt{mT})$. \square

A.5 PROOF OF MODEL QUALITY IMPROVEMENT BY FEDDC

In order to proof Prop. 4, we first need the following Lemma.

Lemma 5. *Given $\delta \in (0, 1]$, $m \in \mathbb{N}$ clients, and $k \in [m]$, if Algorithm 1 with daisy chaining period $d \in \mathbb{N}$ is run for $T \in \mathbb{N}$ rounds with*

$$T \geq d \frac{m}{\rho^{\frac{1}{m}}} (H_m - H_{m-k})$$

where H_m is the m -th harmonic number, then each local model has seen at least k distinct datasets with probability $1 - \rho$.

Note that $H_m \approx \log m + \gamma + \frac{1}{2} + \mathcal{O}\left(\frac{1}{m}\right)$ where $\gamma \approx 0.5772156649$ denotes the Euler-Mascheroni-constant.

Proof. For a single local model, it follows from the coupon collector problem (Neal, 2008) that the expected number of daisy-chaining rounds R required to see at least k out of m clients is at least $m(H_m - H_{m-k})$, where

$$H_m = \sum_{i=1}^m \frac{1}{i}$$

is the m -th harmonic number. To see that, consider that for the first distinct client the chance to pick it is $\frac{m}{m-1}$, for the second $\frac{m}{m-2}$ and for the k -th it is $\frac{m}{m-k+1}$, which sums up to $m(H_m - H_{m-k})$.

Applying the Markov inequality yields

$$P\left(R \geq \frac{1}{\rho} m(H_m - H_{m-k})\right) \leq \rho.$$

The probability for all local models to have seen at least k clients then is at most ρ^m . Thus, if we perform at least

$$R \geq \frac{m}{\rho^{\frac{1}{m}}} m(H_m - H_{m-k})$$

daisy-chaining rounds, then the probability that each local model has not seen at least k distinct datasets is smaller than ρ . The result follows from the fact that the number of daisy-chaining rules is $R = T/d$. \square

Note that H_k can be approximated as

$$H_m \approx \log m + \gamma + \frac{1}{2} + \mathcal{O}\left(\frac{1}{m}\right)$$

where $\gamma = \lim_{m \rightarrow \infty} (H_m - \ln m) \approx 0.5772156649$ denotes the Euler-Mascheroni-constant. From this it follows that

$$H_m - H_{m-k} \approx \ln \frac{m}{m-k} + \mathcal{O}\left(\frac{1}{m} - \frac{1}{m-k}\right)$$

With this, we can now proof Prop. 4 that we restate here for convenience.

Proposition. *Let \mathcal{H} be a model space with Radon number $r \in \mathbb{N}$, ε a convex risk, and \mathcal{A} a learning algorithm with sample size $n_0(\varepsilon, \delta)$. Given $\varepsilon > 0$, $\delta \in (0, r^{-1})$ and any $h \in \mathbb{N}$, and local datasets D_1, \dots, D_m of size $n \in \mathbb{N}$ with $m \geq r^h$, then Alg. 1 using the Radon point with aggr. period*

$$b \geq d \frac{m}{\delta^{\frac{1}{2m}}} \left(H_m - H_{m - \lceil n^{-1} n_0(\varepsilon, \sqrt{\delta}) \rceil} \right)$$

improves model quality in terms of (ε, δ) -guarantees.

Proof. For

$$b \geq d \frac{m}{\delta^{\frac{1}{2m}}} \left(H_m - H_{m - \lceil n^{-1} n_0(\varepsilon, \sqrt{\delta}) \rceil} \right)$$

it follows from Lemma 5 with $k = \lceil n^{-1} n_0(\varepsilon, \sqrt{\delta}) \rceil$ that with probability $1 - \sqrt{\delta}$ all local models are trained on at least $kn = n_0(\varepsilon, \sqrt{\delta})$ samples. Thus an $(\varepsilon, \sqrt{\delta})$ -guarantee holds for each model with probability $1 - \sqrt{\delta}$. It follows from Eq. 1 that the probability that the risk is higher than ε is

$$\mathbb{P}(\varepsilon(\mathbf{t}_h) > \varepsilon) < \left(r \sqrt{\delta} \sqrt{\delta} \right)^{2^h} = (r\delta)^{2^h}.$$

The result follows from $\delta < r^{-1}$ and Eq. (1). \square

A.5.1 NUMERICAL ANALYSIS OF PROPOSITION 4

The lower bound on the aggregation period


$$b \geq d \frac{m}{\delta^{\frac{1}{2m}}} \left(H_m - H_{m - \lceil n^{-1} n_0(\epsilon, \sqrt{\delta}) \rceil} \right)$$

grows linearly with the daisy-chaining period and a factor depending on the number of clients m , the error probability δ , and the required number of hops $\lceil n^{-1} n_0(\epsilon, \sqrt{\delta}) \rceil$. Applying the bound to the experiment using the Radon point on SUSY in Sec. 5 with $m = 441$ clients, daisy-chaining period $d = 1$, and local dataset size of $n = 2$ for local learners achieving an $(\epsilon = 0.05, \delta = 0.01)$ -guarantee requires $b \geq 49.9$ to improve model quality according to Prop. 4. For $b = 50$ like in our experiments, Prop. 4 thus predicts that model quality is improved, under the assumption that


$$n_0(\epsilon, \delta) = \frac{1}{\epsilon} \log \frac{1}{\delta}.$$

Even though, the experiments on CIFAR10 in Section 7 are non-convex, and thus Prop. 4 does not apply, we can still evaluate the predicted required aggregation period: with $m = 150$ clients, daisy-chaining period $d = 1$, local dataset size of $n = 64$, and local learners achieving an $(\epsilon = 0.01, \delta = 0.01)$ -guarantee requires $b \geq 8.81$.

We now analyze the scaling behavior with the error probability δ for various local dataset sizes in Fig. 8a. The lower the error probability, the larger the required aggregation period b , in particular for small local datasets. If local datasets are sufficiently large, the aggregation period can be chosen very small. In Fig. 8b we investigate the required aggregation period b for local learners achieving an $(\epsilon = 0.01, \delta = 0.01)$ -guarantee in relation to the local dataset size n . Indeed, the smaller the local dataset size, the larger the required aggregation period. We also see that for smaller numbers of clients, more aggregation rounds are required, since the chance of a model visiting the same client multiple times is larger.



(a) Aggregation period b for error probability δ for varying local dataset sizes with $m = 150$ clients and $\epsilon = 0.01$.



(b) Aggregation period b for local dataset size n for varying numbers of clients m , with $\epsilon = 0.01$ and $\delta = 0.01$.

A.6 NOTE ON COMMUNICATION COMPLEXITY

In our analysis of the communication complexity, we assume the total amount of communication to be linear in the number of communication rounds. For some communication systems the aggregated model can be broadcasted to individual clients which is not possible in daisy-chaining rounds, reducing the communication complexity in FedAvg. In most scenarios, fiber or GSM networks are used where each model has to be sent individually, so there is no substantial difference between broadcasting a model to all clients and sending an individual model to each client. Therefore, also in this case the amount of communication rounds determines the overall amount of communication.

A.7 EXTENDED DISCUSSION ON PRIVACY

Federated learning only exchanges model parameters, and no local data. It is, however, possible to infer upon local data given the model parameters, model updates or gradients (Ma et al., 2020). In classical federated learning there are two types of attacks that would allow such inference: (i) an attacker intercepting the communication of a client with the server, obtaining models and model updates to infer upon the clients data, and (ii) a malicious server obtaining models to infer upon the data of each client. A malicious client cannot learn about a specific other client’s data, since it only obtains the average of all local models. In federated daisy-chaining there is a third possible attack: (iii) a malicious client obtaining model updates from another client to infer upon its data.

In the following, we discuss potential defenses against these three types of attacks in more detail. Note that we limit the discussion on attacks that aim at inferring upon local data, thus breaching data privacy. Poisoning (Bhagoji et al., 2019) and backdoor (Sun et al., 2019) attacks are an additional threat in federated learning, but are of less importance for our main setting in healthcare: there is no obvious incentive for a hospital to poison a prediction. It is possible that FEDDC presents a novel risk surface for those attacks, but such attack strategies are non-obvious. Robust aggregation, such as the Radon point, are suitable defenses against such attacks (Liu et al., 2022). Moreover, the standard mechanisms that guarantee differential privacy also defend against backdoor and poisoning attacks (Sun et al., 2019).

A general and wide-spread approach to tackle all three possible attack types is to add noise to the model parameters before sending. Using appropriate clipping and noise, this guarantees ϵ, δ -differential privacy for local data (Wei et al., 2020) at the cost of a slight-to-moderate loss in model quality. We empirically demonstrated that FEDDC performs well under such noise in Sec. 6.

Another approach to tackle an attack on communication (i) is to use encrypted communication. One can furthermore protect against a malicious server (ii) by using homomorphic encryption that allows the server to average models without decrypting them (Zhang et al., 2020). This, however, only works for particular aggregation operators and does not allow to perform daisy-chaining. Secure daisy-chaining in the presence of a malicious server (ii) can, however, be performed using asymmetric encryption. Assume each client creates a public-private key pair and shares the public key with the server. To avoid the malicious server to send clients its own public key and act as a man in the middle, public keys have to be announced (e.g., by broadcast). While this allows sending clients to identify the recipient of their model, no receiving client can identify the sender. Thus, inference on the origin of a model remains impossible. For a daisy-chaining round the server sends the public key of the receiving client to the sending client, the sending client checks the validity of the key and sends an encrypted model to the server which forwards it to the receiving client. Since only the receiving client can decrypt the model, the communication is secure.

In standard federated learning, a malicious client cannot infer specifically upon the data of another client from model updates, since it only receives the aggregate of all local models. In federated daisy-chaining, it receives the model from a random, unknown client in each daisy-chaining round. Now, the malicious client can infer upon the membership of a particular data point in the local dataset of the client the model originated from, i.e., through a membership inference attack (Shokri et al., 2017). Similarly, the malicious client can infer upon the presence of data points with certain attributes in the dataset (Ateniese et al., 2015). The malicious client, however, does not know the client the model was trained on, i.e., it does not know the origin of the dataset. Using a random scheduling of daisy-chaining and aggregation rounds at the server, the malicious client cannot even distinguish between a model from another client or the average of all models. Nonetheless, daisy-chaining opens up new potential attack vectors (e.g., clustering received models to potentially determine their origins). These potential attack vectors can be tackled in the same way as in standard federated learning, i.e., by adding noise to model parameters as discussed above, since “[d]ifferentially private models are, by construction, secure against membership inference attacks” (Shokri et al., 2017).