ProxyFL: Decentralized Federated Learning through Proxy Model Sharing

Shivam Kalra

University of Waterloo <https://orcid.org/0000-0002-8415-8627>

Junfeng Wen

Layer6

Jesse Cresswell

Layer6

Maksims Volkovs

Layer6

Hamid Tizhoosh (✉ tizhoosh.hamid@mayo.edu)

Mayo Clinic

Article

Keywords:

Posted Date: December 16th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-1168002/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

ProxyFL: Decentralized Federated Learning through Proxy Model Sharing

Shivam Kalra^{1,2,3} Junfeng Wen¹
Jesse C. Cresswell¹ Maksims Volkovs¹ Hamid R. Tizhoosh^{*2,3,4}

¹Layer 6 AI ²Kimia Lab, University of Waterloo ³Vector Institute ⁴Mayo Clinic

shivam.kalra@uwaterloo.ca, {junfeng, jesse, maks}@layer6.ai, tizhoosh@uwaterloo.ca

Abstract

Institutions in highly regulated domains such as finance and healthcare often have restrictive rules around data sharing. Federated learning is a distributed learning framework that enables multi-institutional collaborations on decentralized data with improved protection for each collaborator’s data privacy. In this paper, we propose a communication-efficient scheme for decentralized federated learning called *ProxyFL*, or proxy-based federated learning. Each participant in ProxyFL maintains two models, a private model, and a publicly shared proxy model designed to protect the participant’s privacy. Proxy models allow efficient information exchange among participants using the PushSum method without the need of a centralized server. The proposed method eliminates a significant limitation of canonical federated learning by allowing model heterogeneity; each participant can have a private model with any architecture. Furthermore, our protocol for communication by proxy leads to stronger privacy guarantees using differential privacy analysis. Experiments on popular image datasets, and a pan-cancer diagnostic problem using over 30,000 high-quality gigapixel histology whole slide images, show that ProxyFL can outperform existing alternatives with much less communication overhead and stronger privacy.

1 Introduction

Access to large-scale datasets is a primary driver of advancement in machine learning, with well-known datasets such as ImageNet [Deng et al., 2009] in computer vision, or SQuAD [Rajpurkar et al., 2016] in natural language processing leading to remarkable achievements. Other domains such as healthcare and finance face restrictions on sharing data, due to regulations and privacy concerns. It is impossible for institutions in these domains to pool and disseminate their data, which limits the progress of research and model development. The ability to share information between institutions while respecting the data privacy of individuals would lead to more robust and accurate models. In this paper we explore the setting of multi-institutional collaboration in highly regulated domains, and propose a method for decentralized model training that preserves data privacy.

In the healthcare domain, for example, histopathology has seen widespread adoption of digitization, offering unique opportunities to increase objectivity and accuracy of diagnostic interpretations through machine learning [Tizhoosh and Pantanowitz, 2018]. Digital images of tissue specimens exhibit significant heterogeneity from the preparation, fixation and staining protocols used at the preparation site, among other factors. Because of this variety, the integration of medical data across multiple institutions is essential. However, centralization of medical data faces regulatory obstacles, as well as workflow and technical challenges including managing and distributing the data. The latter is particularly relevant for digital pathology since each histopathology image is generally a gigapixel file, often one or more gigabytes in size. Distributed machine learning on decentralized data could be a solution to overcome these challenges, and promote the adoption of machine learning in healthcare and similar highly regulated domains.

Federated learning (FL) is a distributed learning framework that was designed to train a model on data that could not be centralized [McMahan et al., 2017]. It trains a model in a distributed manner directly on client devices where data is generated, and gradient updates are communicated back to the centralized

Figure 1: *ProxyFL* is a communication-efficient, decentralized federated learning method where each client (e.g, hospital) maintains a private model, a proxy model, and private data. During distributed training, the client communicates with others only by exchanging their proxy model which enables data and model autonomy. After training, a client’s private model can be used for inference.

server for aggregation. However, the canonical FL setting is not suited to the multi-institutional collaboration problem, as it involves a centralized third party that controls a single model. Considering a collaboration between hospitals, creating one central model may be undesirable. Each hospital may seek autonomy over its own model for regulatory compliance and tailoring to its own specialty.

While it is often claimed that FL provides improved privacy since raw data never leaves the client’s device [McMahan et al., 2017], it does not provide the guarantee of security that regulated institutions require. FL involves each client sending unaudited gradient updates to the central server, which is problematic since deep neural networks are capable of memorizing individual training examples, which may completely breach the client’s privacy [Carlini et al., 2019].

In contrast, meaningful and quantitative guarantees of privacy are provided by the differential privacy (DP) framework [Dwork et al., 2006b]. In DP, access to a database is only permitted through randomized queries in a way that obscures the presence of individual data points. More formally, let $\mathcal{D}$ represent a set of data points, and M a probabilistic function, or *mechanism*, acting on databases. We say that the mechanism is (ϵ, δ) -differentially private if for all subsets of possible outputs $\mathcal{S} \subset \text{Range}(M)$, and for all pairs of databases $\mathcal{D}$ and $\mathcal{D}'$ that differ by one element,

$$\Pr[M(\mathcal{D}) \in \mathcal{S}] \leq \exp(\epsilon) \Pr[M(\mathcal{D}') \in \mathcal{S}] + \delta. \quad (1)$$

The spirit of this definition is that when one individual’s data is added or removed from the database, the outcomes of a private mechanism should be largely unchanged in distribution. This will hold when ϵ and δ are small positive numbers. In this case an adversary would not be able to learn about the individual’s data by observing the mechanism’s output, hence, privacy is preserved. DP mechanisms satisfy several useful properties, including strong guarantees of privacy under composition, and post-processing [Dwork and Roth, 2014, Dwork et al., 2010]. These properties make DP a suitable solution for ensuring data privacy in a collaborative FL setting.

In this paper, we propose proxy-based federated learning, or *ProxyFL*, for decentralized collaboration between institutions which enables training of high-performance and robust models, without sacrificing data privacy or communication efficiency. Our contributions are: (i) a method for decentralized FL in multi-institutional collaborations that is adapted to heterogeneous data sources, and preserves model autonomy for each participant; (ii) incorporation of DP for rigorous privacy guarantees; (iii) analysis and improvement of the communication overhead required to collaborate.

2 Related Work

Decentralized FL for highly regulated domains. Unlike centralized FL [McMahan et al., 2017, Li et al., 2020a] where federated clients coordinate to train a centralized model that can be utilized by everyone as a service, decentralized FL is more suitable for multi-institutional collaborations due to regulatory constraints. The main challenge of decentralized FL is to develop a protocol that allows information passing

in a peer-to-peer manner. Gossip protocols [Kempe et al., 2003] can be used for efficient communication and information sharing [Nedić and Olshevsky, 2016, Nedić et al., 2018]. There are different forms of information being exchanged in the literature, including model weights [Li et al., 2021, Huang et al., 2021], knowledge representations [Wittkopp and Acker, 2020] or model outputs [Lin et al., 2020, Ma et al., 2021]. However, unlike our method, none of these protocols provides a guarantee of privacy for participants, and therefore cannot be used safely in highly regulated domains.

Mutual learning. Each client in our ProxyFL has two models that serve different purposes. They are trained using a DP variant of deep mutual learning (DML) [Zhang et al., 2018] which is an approach for mutual knowledge transfer. DML compares favourably to knowledge distillation between a pre-trained teacher and a typically smaller student [Hinton et al., 2015] since it allows training both models simultaneously from scratch, and provides beneficial information to both models. Federated Mutual Learning (FML) [Shen et al., 2020] introduces a meme model that resembles our proxy model, which is also trained mutually with each client’s private model, but is aggregated at a central server. However, FML is not well-suited to the multi-institutional collaboration setting as it is centralized and provides no privacy guarantee to clients.

Differential privacy in FL. Although raw data never leaves client devices, FL is still susceptible to breaches of privacy [Melis et al., 2019, Bhowmick et al., 2018]. DP has been combined with FL to train centralized models with a guarantee of privacy for all clients that participate [McMahan et al., 2018]. By ensuring that gradient updates are not overly reliant on the information in any single training example, gradients can be aggregated centrally with a DP guarantee [Abadi et al., 2016]. We take inspiration from these ideas for ProxyFL.

Computational pathology. The main application domain considered in this work is computational pathology. Various articles have emphasized the need for privacy-preserving FL when facing large-scale computational pathology workloads. Li et al. [2019] and Ke et al. [2021] used FL for medical image augmentation and segmentation. Their method used a centralized server to aggregate selective weight updates that were treated in a DP framework, but they did not account for the total privacy budget expended over the training procedure. Li et al. [2020b] and Lu et al. [2020] built medical image classification models with FL, and added noise to model weights for privacy. However, model weights have unbounded sensitivity, so no meaningful DP guarantee is achieved with these techniques.

3 Method - ProxyFL

ProxyFL, or proxy-based federated learning is our proposed approach for decentralized federated learning. It is designed for multi-institutional collaborations in highly-regulated domains, and as such incorporates quantitative privacy guarantees with efficient communication.

3.1 Problem Formulation & Overview

We consider the decentralized FL setting involving a set of clients $\mathcal{K}$, each with a local data distribution $\mathcal{D}_k, \forall k \in \mathcal{K}$. Every client maintains a *private model* $f_{\phi_k} : \mathcal{X} \rightarrow \mathcal{Y}$ with parameters ϕ_k , where $\mathcal{X}, \mathcal{Y}$ are the input/output spaces respectively. In this work, we assume that all private models have the same input/output specifications, but may have different structures¹. The goal is to train the private models collectively so that each generalizes well on the joint data distribution.

There are three major challenges in this setting: (i) The clients may not want to reveal their private model’s structure and parameters to others. Revealing model structure can expose proprietary information, increase the risk of adversarial attacks [Fredrikson et al., 2015], and can leak private information about the local datasets [Truex et al., 2019]. (ii) In addition to model heterogeneity, the clients may not want to rely on a third party to manage a shared model, which precludes centralized model averaging schemes. (iii) Information sharing must be efficient, robust, and peer-to-peer. To address the above challenges, we introduce an additional *proxy model* $h_{\theta_k} : \mathcal{X} \rightarrow \mathcal{Y}$ for each client with parameters θ_k . It serves as an interface between the client and the outside world. As part of the communication protocol, all clients agree on a common proxy model architecture for compatibility.

In every round of ProxyFL, each client trains its private and proxy models jointly so that they can benefit from one another. With differentially private training, the proxy can extract useful information

¹This can be further relaxed by including client-specific input/output adaptation layers.

124 from private data, ready to be shared with other clients without violating privacy constraints. Then, each
 125 client sends its proxy to its out-neighbors and receives new proxies from its in-neighbors according to a
 126 communication graph, specified by an adjacency matrix P and de-biasing weights $\mathbf{w}$. Finally, each client
 127 aggregates the proxies they received, and *replaces* their current proxy. The overall procedure is shown in
 Fig. 1 and Algorithm 1. We discuss each step in detail in the subsequent subsections.

Algorithm 1 ProxyFL

Require: Proxy parameters $\theta_k^{(0)}$, private parameters $\phi_k^{(0)}$, de-biasing weight $w_k^{(0)}$ for client k , DML weights $\alpha, \beta \in (0, 1)$, learning rate $\eta > 0$, adjacency matrix $P^{(t)}$

```

1: for each round  $t = 0, \dots, T - 1$  at client  $k \in \mathcal{K}$  do
2:   for each local optimization step do
3:     Sample mini-batch  $\mathcal{B}_k = \{(\mathbf{x}_i, y_i)\}_{i=1}^B$  from  $\mathcal{D}_k$ 
4:     Update local proxy and private models:
            $\theta_k^{(t)} \leftarrow \theta_k^{(t)} - \eta \tilde{\nabla} \hat{\mathcal{L}}_{\theta_k}(\mathcal{B}_k)$  # DP update
            $\phi_k^{(t)} \leftarrow \phi_k^{(t)} - \eta \nabla \hat{\mathcal{L}}_{\phi_k}(\mathcal{B}_k)$  # non-DP update
5:   end for
6:    $\phi_k^{(t+1)} \leftarrow \phi_k^{(t)}$ 
7:   Send  $(P_{k',k}^{(t)} \theta_k^{(t)}, P_{k',k}^{(t)} w_k^{(t)})$  to out-neighbors;
      receive  $(P_{k,k'}^{(t)} \theta_{k'}^{(t)}, P_{k,k'}^{(t)} w_{k'}^{(t)})$  from in-neighbors
8:   Update local proxy  $\theta_k^{(t+1)} \leftarrow \sum_{k'} P_{k,k'}^{(t)} \theta_{k'}^{(t)}$ 
9:   Update de-bias weight  $w_k^{(t+1)} \leftarrow \sum_{k'} P_{k,k'}^{(t)} w_{k'}^{(t)}$ 
10:  De-bias  $\theta_k^{(t+1)} \leftarrow \theta_k^{(t+1)} / w_k^{(t+1)}$ 
11: end for
12: return  $\theta_k^{(T)}, \phi_k^{(T)}$ 

```

128

129 3.2 Training Objectives

130 For concreteness, we consider classification tasks. To train the private and proxy models at the start of each
 131 round of training, we apply a variant of DML [Zhang et al., 2018]. Specifically, when training the private
 132 model for client k , in addition to the cross-entropy loss (CE)

$$\mathcal{L}_{\text{CE}}(f_{\phi_k}) := \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_k} \text{CE}[f_{\phi_k}(\mathbf{x}) \| y], \quad (2)$$

133 DML adds a KL divergence loss (KL)

$$\mathcal{L}_{\text{KL}}(f_{\phi_k}; h_{\theta_k}) := \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_k} \text{KL}[f_{\phi_k}(\mathbf{x}) \| h_{\theta_k}(\mathbf{x})], \quad (3)$$

134 so that the private model can also learn from the current proxy model. The objective for learning the private
 135 model is given by

$$\mathcal{L}_{\phi_k} := (1 - \alpha) \cdot \mathcal{L}_{\text{CE}}(f_{\phi_k}) + \alpha \cdot \mathcal{L}_{\text{KL}}(f_{\phi_k}; h_{\theta_k}), \quad (4)$$

136 where $\alpha \in (0, 1)$ balances between the two losses. The objective for the proxy model is similarly defined
 137 as

$$\mathcal{L}_{\theta_k} := (1 - \beta) \cdot \mathcal{L}_{\text{CE}}(h_{\theta_k}) + \beta \cdot \mathcal{L}_{\text{KL}}(h_{\theta_k}; f_{\phi_k}). \quad (5)$$

138 where $\beta \in (0, 1)$. As in DML, we alternate stochastic gradient steps between the private and proxy models.

139 In our context, mini-batches are sampled from the client’s private dataset. Releasing the proxy model
 140 to other clients risks revealing that private information. Therefore, each client uses differentially pri-
 141 vate stochastic gradient descent (DP-SGD) [Abadi et al., 2016] when training the proxy (but not the pri-
 142 vate model). Let $\mathcal{B}_k = \{(\mathbf{x}_i, y_i)\}_{i=1}^B$ denote a mini-batch sampled from $\mathcal{D}_k$. The stochastic gradient is

143 $\nabla \hat{\mathcal{L}}_{\phi_k}(\mathcal{B}_k) := \frac{1}{B} \sum_{i=1}^B \mathbf{g}_{\phi_k}^{(i)}$ where

$$\mathbf{g}_{\phi_k}^{(i)} := (1 - \alpha) \nabla_{\phi_k} \text{CE}[f_{\phi_k}(\mathbf{x}_i) \| y_i] + \alpha \nabla_{\phi_k} \text{KL}[f_{\phi_k}(\mathbf{x}_i) \| h_{\theta_k}(\mathbf{x}_i)]. \quad (6)$$

144 $\nabla \hat{\mathcal{L}}_{\theta_k}(\mathcal{B}_k)$ and $\mathbf{g}_{\theta_k}^{(i)}$ are similarly defined for the proxy. To perform DP training for the proxy, the per-
 145 example gradient is clipped, then aggregated over the mini-batch, and finally Gaussian noise is added [Abadi
 146 et al., 2016]:

$$\begin{aligned} \bar{\mathbf{g}}_{\theta_k}^{(i)} &:= \mathbf{g}_{\theta_k}^{(i)} / \max\left(1, \|\mathbf{g}_{\theta_k}^{(i)}\|_2 / C\right), \\ \tilde{\nabla} \hat{\mathcal{L}}_{\theta_k}(\mathcal{B}_k) &:= \frac{1}{B} \left(\sum_{i=1}^B \bar{\mathbf{g}}_{\theta_k}^{(i)} + \mathcal{N}(0, \sigma^2 C^2 I) \right), \end{aligned} \quad (7)$$

147 where $C > 0$ is the clipping threshold and $\sigma > 0$ is the noise level (see Lines 2–5 in Algorithm 1).

148 3.3 Privacy Guarantee

149 The proxy model is the only entity that a client reveals, so each client must ensure this sharing does not
 150 compromise the privacy of their data. Since arbitrary post-processing on a DP-mechanism does not weaken
 151 its (ϵ, δ) guarantee [Dwork and Roth, 2014], it is safe to release the proxy as long as it was trained via a
 152 DP-mechanism. DP-SGD as defined in Eq. (7) is based on the Gaussian mechanism [Dwork et al., 2006a]
 153 which meets the requirement of Eq. (1) by adding Gaussian noise to the outputs of a function f with
 154 bounded sensitivity C in L_2 norm,

$$M(x) = f(x) + \mathcal{N}(0, \sigma^2 C^2 I). \quad (8)$$

155 DP-SGD simply takes $f(x)$ to be the stochastic gradient update, with clipping to ensure bounded sensitiv-
 156 ity.

157 Every application of the DP-SGD step incurs a privacy cost related to the clipping threshold C , and
 158 noise level σ . A strong bound on the total privacy cost over many applications of DP-SGD is obtained by
 159 using the framework of Rényi differential privacy [Mironov, 2017, Mironov et al., 2019] to track privacy
 160 under compositions of DP-SGD, then convert the result to the language of (ϵ, δ) -DP as in Eq. (1) [Balle
 161 et al., 2020].

162 Finally, privacy guarantees are tracked on a per-client basis. In a multi-institutional collaboration, every
 163 client has an obligation to protect the privacy of the data it has collected. Hence, each client individually
 164 tracks the parameters (ϵ, δ) for its own proxy model training, and can drop out of the protocol when its
 165 prespecified privacy budget is reached. Throughout the paper we specify δ based on the dataset size, and
 166 compute ϵ .

167 3.4 Communication Efficiency & Robustness

168 The proxies serve as interfaces for information transfer and must be locally aggregated in a way that facil-
 169 itates efficient learning among clients. One may use a central parameter server to compute the average of
 170 the proxies, similar to Shen et al. [2020]. However, this will incur a communication cost that grows linearly
 171 in the number of clients, and is not decentralized. We propose to apply the PushSum scheme [Kempe et al.,
 172 2003, Nedić et al., 2018] to exchange proxies among clients.

173 Let $\Theta^{(t)} \in \mathbb{R}^{|\mathcal{K}| \times d_\theta}$ represent the stacked proxies at round t , where the rows are the proxy parameters
 174 $\theta_k^{(t)}, \forall k \in \mathcal{K}$. We use $P^{(t)} \in \mathbb{R}^{|\mathcal{K}| \times |\mathcal{K}|}$ to denote the weighted adjacency matrix representing the graph
 175 topology at round t , where $P_{k,k'}^{(t)} \neq 0$ indicates that client k receives the proxy from client k' . Note
 176 that $P^{(t)}$ needs to be column-stochastic, but need not be symmetric (bidirectional communication) nor
 177 time-invariant (across rounds). Such a $P^{(t)}$ will ensure efficient communication when it is sparse. The
 178 communication can also handle asymmetrical connections such as different upload/download speeds, and
 179 can adapt to clients joining or dropping out since it is time-varying.

180 With these notations, every round of communication can be concisely written as $\Theta^{(t+1)} = P^{(t)} \Theta^{(t)}$.
 181 Under certain mixing conditions [Seneta, 2006], it can be shown that $\lim_{T \rightarrow \infty} \prod_{t=0}^T P^{(t)} = \pi \mathbf{1}^\top$, where
 182 π is the limiting distribution of the Markov chain and $\mathbf{1}$ is a vector of all ones. Suppose for now that there
 183 is no training for the proxies between rounds, i.e., updates to the proxies are due to communication and

Figure 2: An illustration of the directed exponential graphs used in our experiments. Dark arrows indicate the communication path of the bottom node at each round, while light arrows show the communication path of others. Solid nodes indicate clients who have received information from the bottom client. After only $\lceil \log_2(|\mathcal{K}|) \rceil$ rounds, all nodes have access to that information.

Figure 3: Test performance with DP training, reporting mean and standard deviation over clients and runs.

184 replacement only. In the limit, θ_k will converge to $\theta_k^{(\infty)} = \pi_k \sum_{k' \in \mathcal{K}} \theta_{k'}^{(0)}$. To mimic model averaging (i.e.,
185 computing $\frac{1}{|\mathcal{K}|} \sum_{k' \in \mathcal{K}} \theta_{k'}^{(0)}$), the bias introduced by π_k must be corrected. This can be achieved by having
186 the clients maintain another set of weights $\mathbf{w} \in \mathbb{R}^{|\mathcal{K}|}$ with initial values $\mathbf{w}^{(0)} = \mathbf{1}$. By communicating
187 $\mathbf{w}^{(t+1)} = P^{(t)} \mathbf{w}^{(t)}$, we can see that $\mathbf{w}^{(\infty)} = \pi \mathbf{1}^\top \mathbf{w}^{(0)} = |\mathcal{K}| \pi$. As a result, the de-biased average is
188 given by $\theta_k^{(\infty)} / w_k^{(\infty)} = \frac{1}{|\mathcal{K}|} \sum_{k' \in \mathcal{K}} \theta_{k'}^{(0)}$.

189 Finally, recall that the proxies are trained locally in each round. Instead of running the communica-
190 tion to convergence for proxy averaging, we alternate between training (Lines 2–5 in Algorithm 1) and
191 communicating (Lines 7–10) proxies, similar to Assran et al. [2019].

192 4 Empirical Evaluation

193 In this section, we demonstrate the effectiveness and efficiency of ProxyFL on popular image classification
194 datasets as well as a pan-cancer diagnostic problem. In our experiments, we use the exponential com-
195 munication protocol of Assran et al. [2019], as illustrated in Fig. 2. The communication graph $P^{(t)}$ is a
196 permutation matrix so that each client only receives and sends one proxy per communication round. Each
197 client communicates with its peer that is $2^0, 2^1, \dots, 2^{\lceil \log_2(|\mathcal{K}|-1) \rceil}$ steps away periodically. The communi-
198 cation and synchronization is handled via the OpenMPI library [Graham et al., 2005]. Our code is available
199 at <https://github.com/layer6ai-labs/ProxyFL>.

200 4.1 Image Classification

201 **Datasets & settings.** We conducted experiments on MNIST [LeCun et al., 1998], Fashion-MNIST (FaM-
202 NIST) [Xiao et al., 2017] and CIFAR-10 [Krizhevsky, 2009]. Fa/MNIST has 60k training images of size
203 28x28, while CIFAR-10 has 50k RGB training images of size 32x32. Each dataset has 10k test images,
204 which are used to evaluate the model performance. Experiments were conducted on a server with 8 V100
205 GPUs, which correspond to 8 clients. In each run, every client had 1k (Fa/MNIST) or 3k (CIFAR-10)
206 non-overlapping private images sampled from the training set. To test robustness on non-IID data, clients
207 were given a skewed private data distribution. For each client, a randomly chosen class was assigned and
208 a fraction p_{major} (0.8 for Fa/MNIST; 0.3 for CIFAR-10) of that client’s private data was drawn from that

Figure 4: Communication time. Figure 5: Degree of non-IID skew. Figure 6: Heterogeneous models.

Figure 7: The effect of DP-SGD on different methods.

class. The remaining data was randomly drawn from all other classes in an IID manner. Hence, clients must learn from collaborators to generalize well on the IID test set.

Baselines. We compare our method to FedAvg [McMahan et al., 2017], FML [Shen et al., 2020], AvgPush, Regular, and Joint training. FedAvg and FML are centralized schemes that average models with identical structure. FML is similar to ProxyFL in that every client has two models, except FML does centralized averaging and originally did not incorporate DP training. AvgPush is a decentralized version of FedAvg using PushSum for aggregation. Regular training uses the local private datasets without any collaboration. Joint training mimics a scenario without constraints on data centralization by combining data from all clients and training a single model. Regular, Joint, FedAvg, and AvgPush use DP-SGD for training their models, while ProxyFL and FML use it for their proxies.

Implementation details. Following Shen et al. [2020], the private/proxy models are LeNet5/MLP for Fa/MNIST, and CNN2/CNN1 for CIFAR10. All methods use the Adam optimizer [Kingma and Ba, 2014] with learning rate of 0.001, weight decay of $1e-4$, mini-batch size of 250, clipping $C = 1.0$ and noise level $\sigma = 1.0$. Each round of local training takes a number of gradient steps equivalent to one epoch over the private data. For proper DP accounting, mini-batches are sampled from the training set independently with replacement by including each training example with a fixed probability [Yu et al., 2019]. The DML parameters α, β are set at 0.5 for FML and ProxyFL. We report mean and standard deviations based on 5 random seeds. Additional details and results can be found in Appendix A and B.

Results & discussions. Fig. 3 shows the performance on the test datasets. There are a few observations: (i) The private models of ProxyFL achieve the best overall performance on all datasets, even better than the centralized counterpart FML. The improvements of ProxyFL-private over other methods are statistically significant for every dataset (p -value $< 1e-5$). Note that the Joint method serves as an upper bound of the problem when private datasets are combined. (ii) As a decentralized scheme, ProxyFL has a much lower communication cost compared to FML, as shown in Fig. 4. The exponential protocol has a constant time complexity per round regardless of the number of clients, which makes ProxyFL much more scalable. (iii) Decentralized schemes seem to be more robust to DP training, as AvgPush outperforms FedAvg and ProxyFL outperforms FML consistently.

4.2 Ablation

We ablated ProxyFL to see how different factors affect its performance on the MNIST dataset. Unless specified otherwise, all models, including the private ones in FML and ProxyFL, have the same MLP

Client	# Slides		
	Training	Test	Total
C1: University of Pittsburgh	1,310	562	1,872
C2: Indivumed	1,004	431	1,435
C3: Asterand	818	351	1,169
C3: MSKCC	798	342	1,140
Total			5,616

Table 1: Distribution of WSIs across 4 different participating clients. The total of 5,616 WSIs accounts for around 6 TB of imaging data.

structure. Additional ablation results can be found in the Appendix B.

IID versus non-IID. Fig. 5 shows the performance of the methods with different levels of non-IID dataset skew. $p_{\text{major}} = 0.1$ is the IID setting because there are 10 classes. We can see that as the setting deviates from IID, most methods degrade except for Joint training since it unifies the datasets. The private model of ProxyFL is the most robust to the degree of non-IID dataset skew. Note that the proxy model of ProxyFL achieves similar performance to the private model of FML, which is trained *without* DP guarantees. This indicates that ProxyFL is a scheme that is robust to distribution shifts among the clients.

Different private architectures. One important aspect of ProxyFL is that the private models can be heterogeneous—customized to meet the special needs of the individual clients. On the MNIST task we use all four model architectures, one for every two clients². Results are shown in Fig. 6. We see that different models can achieve very diverse and sub-optimal performances with individual Regular training, while ProxyFL can improve all architectures’ performance. The improvements for weaker models are more significant than for stronger models.

DP versus non-DP. Next we investigate the effect of DP-SGD on the algorithms. Fig. 7a shows the test accuracies of the different training methods with and without DP-SGD’s gradient clipping and noise addition. Clearly, all methods can outperform Regular training when there is no privacy constraint. However, with DP-SGD, centralized methods like FedAvg and FML-proxy perform poorly, even worse than Regular training. ProxyFL-private shows the smallest decrease in performance when DP-SGD is included and remains closest to the upper bound of Joint training. Fig. 7b shows how efficient each training method is with its privacy budget³.

4.3 Pan-Cancer Analysis

Dataset. In this experiment, we evaluated ProxyFL on a multi-origin real-world dataset. We considered the largest public archive of whole-slide images (WSIs), namely The Cancer Genome Atlas (TCGA) [Weinstein et al., 2013]. TCGA provides about 30,000 H&E stained WSIs originating from various institutions, distributed across multiple primary diagnoses. The client data for this study was derived from TCGA by splitting it across four major institutions: i) University of Pittsburgh, ii) Indivumed, iii) Asterand, and iv) Memorial Sloan Kettering Cancer Center (MSKCC). The data splits for each participating client are described in Table 1, and the anatomic site distributions in Fig. 10. The total data size is around 6 TB for all hospitals.

WSI pre-processing & model setup. Each WSI is an extremely large image (more than 50,000 x 50,000 pixels with a size often much larger than several hundred MBs), and cannot be directly processed by a CNN. In order to classify a WSI, we divided it into a small number of representative patches called a mosaic, using the techniques from Kalra et al. [2020a]. The mosaic patches were then converted into feature vectors using a pre-trained DenseNet [Huang et al., 2017]. Each WSI corresponds to a set of features; these sets are then used for training a classifier based on the DeepSet architecture [Zaheer et al., 2017]. In the context of ProxyFL, both the private and proxy models are DeepSet-based.

Experimental setup. The experiments were conducted using four V100 GPUs. Three FL methods were compared: ProxyFL, FML, and FedAvg. In each scenario, training was conducted for 50 rounds with a mini-batch size of 16. All methods were tested with two DP settings, one with strong privacy $\sigma = 1.4$,

²CNN1 and CNN2 are slightly adapted to fit MNIST images.

³For illustration, we report ϵ for $\delta = 10^{-3}$. The DP guarantees are evaluated on a per-client basis.

Figure 8: Performance of ProxyFL, FML, and FedAvg on the histopathology dataset involving four hospitals. The mean accuracy and standard deviation of clients on both internal and external data is recorded at the end of each round for two DP settings, and is presented in (a) and (b) respectively. Three random seeds were used. As expected, stronger privacy results in the lower overall accuracy for the internal dataset, but ProxyFL and FML show commensurate changes. Privacy guarantees for each method are listed in (c), computed based on the training set sizes in Table 1. The communication time per client for 150 rounds of training is shown in (d). FedAvg has less efficient communication because it exchanges the larger private model whereas ProxyFL and FML exchange the lightweight proxy models.

and the other with comparatively weak privacy $\sigma = 0.7$, both with $C = 0.7$. The client-level privacy guarantees for the two DP settings are provided in Fig. 8c, computed based on the training set sizes in Table 1. FedAvg and the proxy models used the DP-SGD optimizer with learning rate 0.01, whereas the private models used the Adam with learning rate 0.001. For ProxyFL and FML the private models are used to compute the accuracy values, whereas the central model is used in the case of FedAvg.

Performance was computed based on two test datasets – internal and external. Both datasets are local to the clients. Internal test data is sampled from the same distribution as the client’s private training data, whereas external test data comes from other clients involved in the federated training, and hence a different institution entirely⁴. The 32 unique primary diagnoses in the dataset can be further grouped into 13 tumor types⁵. The tumor type of a WSI is generally known at inference time, so the objective is to predict the cancer sub-type. We evaluated our method by its accuracy of classifying a cancer sub-type (primary diagnosis) of a WSI given that its tumor type is already known.

Results. The sub-type classification results for internal and external data on two different DP settings (strong and weak privacy) for each method are reported in Fig. 8a and Fig. 8b. ProxyFL achieves overall higher accuracy compared to FML and FedAvg on the internal test data for both privacy settings. For the external test data, all three methods perform similar to each other with FedAvg slightly ahead when using stronger privacy. ProxyFL has noticeably better convergence compared to FML as shown by the lower variance in both privacy settings. When strong privacy is used, the FedAvg central model has converged by around the 25th round showing no improvement in the performance across both test datasets. Both ProxyFL and FML are more communication efficient than FedAvg because they exchange lightweight proxy models rather than the larger private models (Fig. 8d), but ProxyFL has the lowest communication overhead due to using fewer model exchanges.

5 Conclusion and Future Work

The recent surge in digitization of highly-regulated areas, such as healthcare and finance, is creating a large amount of data useful for training intelligent models. However, access to this data is limited due to regulatory and privacy concerns. In this paper, we proposed a novel decentralized federated learning scheme,

⁴For the external test data, we only use examples with a primary diagnosis present in the client’s local data. If access to this type of external data is not possible due to privacy concerns, model performance could be validated on public external data.

⁵Tumor types are from Tables 3 and 4 of Kalra et al. [2020b].

ProxyFL, for multi-institutional collaborations without revealing participants’ private data. ProxyFL preserves data and model privacy, and provides a decentralized and communication-efficient mechanism for distributed training. Experiments suggest that ProxyFL is competitive compared to other baselines in terms of model accuracy, communication efficiency, and privacy preservation. Furthermore, we evaluated our approach on a real-world scenario consisting of four medical institutions collaborating to train models for pan-cancer classification.

In contrast to ProxyFL, some FL methods for private knowledge transfer require a public dataset [Papernot et al., 2017, Sun and Lyu, 2021]. This distinction is important, as the core issue in highly regulated domains is that large-scale public datasets are not readily available. Access to a public dataset greatly increases the variety of techniques that could be used, and is an avenue for future research. This work focused on the *privacy* aspects of the ProxyFL protocol, but not on its *security*. We have assumed that all clients collaborate in *good faith*. Handling malicious participants is another important direction of future research, and the ideas from Bhagoji et al. [2019] are relevant.

Acknowledgments

SK and HRT have been supported by The Natural Sciences and Engineering Research Council of Canada (NSERC). SK is also supported by a Vector Institute internship.

Author contributions

SK conceptualized the idea of federated learning through proxy model sharing. SK and JW have the equal contributions in the experimental analysis and paper writing. JCC was involved in the discussion, contributed to the differential privacy analysis of the method, reviewed the code, and provided important edits to the paper. MV and HRT were involved in the discussions of the approach, and provided critical feedback to the paper. HRT guided the external validations with TCGAimages.

Data availability

The publicly available dataset of 30,072 WSIs from TCGA (Weinstein et al. [2013]) (Genomic Data Commons GDC) was used for conducting this study.

References

- Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 308–318, 2016.
- Mahmoud Assran, Nicolas Loizou, Nicolas Ballas, and Mike Rabbat. Stochastic gradient push for distributed deep learning. In *International Conference on Machine Learning*, pages 344–353. PMLR, 2019.
- Borja Balle, Gilles Barthe, Marco Gaboardi, Justin Hsu, and Tetsuya Sato. Hypothesis testing interpretations and renyi differential privacy. In Silvia Chiappa and Roberto Calandra, editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 2496–2506. PMLR, 26–28 Aug 2020. URL <https://proceedings.mlr.press/v108/balle20a.html>.
- 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.
- Abhishek Bhowmick, John Duchi, Julien Freudiger, Gaurav Kapoor, and Ryan Rogers. Protection against reconstruction and its applications in private federated learning. *arXiv preprint arXiv:1812.00984*, 2018.

- 346 Nicholas Carlini, Chang Liu, Úlfar Erlingsson, Jernej Kos, and Dawn Song. The Secret Sharer: Evalu-
 347 ating and Testing Unintended Memorization in Neural Networks. In *Proceedings of the 28th USENIX*
 348 *Conference on Security Symposium, SEC'19*, page 267–284, USA, 2019. USENIX Association. ISBN
 349 9781939133069.
- 350 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical
 351 image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255.
 352 Ieee, 2009.
- 353 Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Found. Trends Theor.*
 354 *Comput. Sci.*, 9(3–4):211–407, August 2014. ISSN 1551-305X. doi: 10.1561/04000000042. URL
 355 <https://doi.org/10.1561/04000000042>.
- 356 Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our
 357 data, ourselves: Privacy via distributed noise generation. In *Advances in Cryptology (EURO-*
 358 *CRYPT 2006)*, volume 4004 of *Lecture Notes in Computer Science*, pages 486–503. Springer Ver-
 359 lag, May 2006a. URL [https://www.microsoft.com/en-us/research/publication/](https://www.microsoft.com/en-us/research/publication/our-data-ourselves-privacy-via-distributed-noise-generation/)
 360 [our-data-ourselves-privacy-via-distributed-noise-generation/](https://www.microsoft.com/en-us/research/publication/our-data-ourselves-privacy-via-distributed-noise-generation/).
- 361 Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in
 362 private data analysis. In Shai Halevi and Tal Rabin, editors, *Theory of Cryptography*, pages 265–284,
 363 Berlin, Heidelberg, 2006b. Springer Berlin Heidelberg. ISBN 978-3-540-32732-5.
- 364 Cynthia Dwork, Guy N Rothblum, and Salil Vadhan. Boosting and differential privacy. In *2010 IEEE 51st*
 365 *Annual Symposium on Foundations of Computer Science*, pages 51–60. IEEE, 2010.
- 366 Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. Model inversion attacks that exploit confidence
 367 information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC Conference on*
 368 *Computer and Communications Security, CCS '15*, page 1322–1333, New York, NY, USA, 2015. As-
 369 sociation for Computing Machinery. ISBN 9781450338325. doi: 10.1145/2810103.2813677. URL
 370 <https://doi.org/10.1145/2810103.2813677>.
- 371 Richard L Graham, Timothy S Woodall, and Jeffrey M Squyres. Open mpi: A flexible high performance
 372 mpi. In *International Conference on Parallel Processing and Applied Mathematics*, pages 228–239.
 373 Springer, 2005.
- 374 Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint*
 375 *arXiv:1503.02531*, 2015.
- 376 Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convo-
 377 lutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*,
 378 pages 4700–4708, 2017.
- 379 Yutao Huang, Lingyang Chu, Zirui Zhou, Lanjun Wang, Jiangchuan Liu, Jian Pei, and Yong Zhang. Per-
 380 sonalized cross-silo federated learning on non-iid data. In *Proceedings of the AAAI Conference on*
 381 *Artificial Intelligence*, volume 35-9, pages 7865–7873, 2021.
- 382 Shivam Kalra, Hamid R Tizhoosh, Charles Choi, Sultaan Shah, Phedias Diamandis, Clinton JV Campbell,
 383 and Liron Pantanowitz. Yottixel—an image search engine for large archives of histopathology whole slide
 384 images. *Medical Image Analysis*, 65:101757, 2020a.
- 385 Shivam Kalra, Hamid R Tizhoosh, Sultaan Shah, Charles Choi, Savvas Damaskinos, Amir Safarpour,
 386 Sobhan Shafiei, Morteza Babaie, Phedias Diamandis, Clinton JV Campbell, et al. Pan-cancer diagnostic
 387 consensus through searching archival histopathology images using artificial intelligence. *NPJ digital*
 388 *medicine*, 3(1):1–15, 2020b.
- 389 Jing Ke, Yiqing Shen, and Yizhou Lu. Style normalization in histology with federated learning. In *2021*
 390 *IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pages 953–956. IEEE, 2021.
- 391 David Kempe, Alin Dobra, and Johannes Gehrke. Gossip-based computation of aggregate information. In
 392 *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings.*, pages 482–
 393 491. IEEE, 2003.

- 394 Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Confer-*
395 *ence on Learning Representations*, 2014.
- 396 Alex Krizhevsky. Learning Multiple Layers of Features from Tiny Images. *Technical Report, University*
397 *of Toronto, Toronto*, 2009.
- 398 Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to docu-
399 ment recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- 400 Chengxi Li, Gang Li, and Pramod K Varshney. Decentralized federated learning via mutual knowledge
401 transfer. *IEEE Internet of Things Journal*, 2021.
- 402 Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. Federated learning: Challenges, methods,
403 and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60, 2020a.
- 404 Wenqi Li, Fausto Milletari, Daguang Xu, Nicola Rieke, Jonny Hancox, Wentao Zhu, Maximilian Baust,
405 Yan Cheng, Sébastien Ourselin, M Jorge Cardoso, et al. Privacy-preserving federated brain tumour
406 segmentation. In *International workshop on machine learning in medical imaging*, pages 133–141.
407 Springer, 2019.
- 408 Xiaoxiao Li, Yufeng Gu, Nicha Dvornek, Lawrence H. Staib, Pamela Ventola, and James S. Duncan.
409 Multi-site fmri analysis using privacy-preserving federated learning and domain adaptation: Abide re-
410 sults. *Medical Image Analysis*, 65:101765, 2020b. ISSN 1361-8415. doi: [https://doi.org/10.1016/](https://doi.org/10.1016/j.media.2020.101765)
411 [j.media.2020.101765](https://doi.org/10.1016/j.media.2020.101765). URL [https://www.sciencedirect.com/science/article/pii/](https://www.sciencedirect.com/science/article/pii/S1361841520301298)
412 [S1361841520301298](https://www.sciencedirect.com/science/article/pii/S1361841520301298).
- 413 Tao Lin, Lingjing Kong, Sebastian U Stich, and Martin Jaggi. Ensemble distillation for robust model
414 fusion in federated learning. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors,
415 *Advances in Neural Information Processing Systems*, volume 33, pages 2351–2363. Curran Associates,
416 Inc., 2020.
- 417 Ming Y Lu, Dehan Kong, Jana Lipkova, Richard J Chen, Rajendra Singh, Drew FK Williamson, Tiffany Y
418 Chen, and Faisal Mahmood. Federated learning for computational pathology on gigapixel whole slide
419 images. *arXiv preprint arXiv:2009.10190*, 2020.
- 420 Jiaxin Ma, Ryo Yonetani, and Zahid Iqbal. Adaptive distillation for decentralized learning from heteroge-
421 neous clients. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 7486–7492.
422 IEEE, 2021.
- 423 H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas.
424 Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proceedings of*
425 *the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2017. URL
426 <http://arxiv.org/abs/1602.05629>.
- 427 H. Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. Learning differentially private recur-
428 rent language models. In *International Conference on Learning Representations*, 2018.
- 429 Luca Melis, Congzheng Song, Emiliano De Cristofaro, and Vitaly Shmatikov. Exploiting unintended
430 feature leakage in collaborative learning. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages
431 691–706. IEEE, 2019.
- 432 Ilya Mironov. Rényi differential privacy. *2017 IEEE 30th Computer Security Foundations Symposium*
433 *(CSF)*, Aug 2017. doi: 10.1109/csf.2017.11. URL [http://dx.doi.org/10.1109/CSF.2017.](http://dx.doi.org/10.1109/CSF.2017.11)
434 [11](http://dx.doi.org/10.1109/CSF.2017.11).
- 435 Ilya Mironov, Kunal Talwar, and Li Zhang. Rényi Differential Privacy of the Sampled Gaussian Mecha-
436 nism. *arXiv preprint arXiv:1908.10530*, 2019.
- 437 Angelia Nedić and Alex Olshevsky. Stochastic gradient-push for strongly convex functions on time-varying
438 directed graphs. *IEEE Transactions on Automatic Control*, 61(12):3936–3947, 2016.

- 439 Angelia Nedić, Alex Olshevsky, and Michael G Rabbat. Network topology and communication-
440 computation tradeoffs in decentralized optimization. *Proceedings of the IEEE*, 106(5):953–976, 2018.
- 441 Nicolas Papernot, Martín Abadi, Úlfar Erlingsson, Ian Goodfellow, and Kunal Talwar. Semi-supervised
442 knowledge transfer for deep learning from private training data. In *Proceedings of the International
443 Conference on Learning Representations*, 2017. URL <https://arxiv.org/abs/1610.05755>.
- 444 Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100, 000+ questions for
445 machine comprehension of text. In *EMNLP*, pages 2383–2392, 2016. URL <http://aclweb.org/anthology/D/D16/D16-1264.pdf>.
- 447 E. Seneta. *Non-negative Matrices and Markov Chains*. Springer Series in Statistics. Springer New York,
448 2006. ISBN 9780387327921.
- 449 Tao Shen, Jie Zhang, Xinkang Jia, Fengda Zhang, Gang Huang, Pan Zhou, Kun Kuang, Fei Wu, and Chao
450 Wu. Federated mutual learning. *arXiv preprint arXiv:2006.16765*, 2020.
- 451 Lichao Sun and Lingjuan Lyu. Federated model distillation with noise-free differential privacy. In *Pro-
452 ceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 1563–
453 1570. International Joint Conferences on Artificial Intelligence Organization, 8 2021. Main Track.
- 454 Hamid Reza Tizhoosh and Liron Pantanowitz. Artificial intelligence and digital pathology: challenges and
455 opportunities. *Journal of pathology informatics*, 9, 2018.
- 456 Stacey Truex, Ling Liu, Mehmet Emre Gursoy, Lei Yu, and Wenqi Wei. Demystifying membership infer-
457 ence attacks in machine learning as a service. *IEEE Transactions on Services Computing*, pages 1–1,
458 2019. doi: 10.1109/TSC.2019.2897554.
- 459 John N Weinstein, Eric A Collisson, Gordon B Mills, Kenna R Mills Shaw, Brad A Ozenberger, Kyle
460 Ellrott, Ilya Shmulevich, Chris Sander, and Joshua M Stuart. The cancer genome atlas pan-cancer
461 analysis project. *Nature genetics*, 45(10):1113–1120, 2013.
- 462 Thorsten Wittkopp and Alexander Acker. Decentralized federated learning preserves model and data pri-
463 vacy. In *International Conference on Service-Oriented Computing*, pages 176–187. Springer, 2020.
- 464 Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking
465 machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- 466 Lei Yu, Ling Liu, Calton Pu, Mehmet Emre Gursoy, and Stacey Truex. Differentially private model pub-
467 lishing for deep learning. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 332–349. IEEE,
468 2019.
- 469 Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan Salakhutdinov, and Alexan-
470 der Smola. Deep sets. *arXiv preprint arXiv:1703.06114*, 2017.
- 471 Ying Zhang, Tao Xiang, Timothy M Hospedales, and Huchuan Lu. Deep mutual learning. In *Proceedings
472 of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4320–4328, 2018.

Appendix

A Experiment Details

A.1 Image Classification

We used the model architectures from Shen et al. [2020] for direct comparison. Specifically, the MLP consists of two hidden layers with 200 units each and ReLU activations. The CNN1 model consists of (conv(3x3, 6), ReLU, maxpool(2x2), conv(3x3, 16), ReLU, maxpool(2x2), fc(64), ReLU, fc(10)) where conv(3x3, 6) means a 3x3 convolution layer with 6 channels, maxpool(2x2) is a max pooling layer with 2x2 kernel, and fc(64) means a fully connected layer with 64 hidden units, etc. The CNN2 model is (conv(3x3, 128), ReLU, maxpool(2x2), conv(3x3, 128), ReLU, maxpool(2x2), fc(10)).

A.2 WSI Distribution and Mosaic

A WSI, a gigapixel image, is first converted into a *mosaic*, or set of representative patches based on Kalra et al. [2020a]. A sample mosaic is illustrated in Fig. 9. These patches are then transformed into a bag of features, $X = \{x_1, \dots, x_n\}$, where x_i is the feature vector of the i^{th} patch in the mosaic. The ordering of elements inside X should not affect the classification of the corresponding WSI. Therefore, a WSI classification can be modelled as a set learning problem, or multi-instance learning (MIL).

Figure 9: A sample mosaic of a WSI diagnosed with lung cancer. It contains 70 patches, each of size 1000×1000 .

To mimic a real-world scenario, WSIs were selected from TCGA and assigned to clients based on the hospital of origin. The distribution of WSIs by client for training and testing is summarized in Table 1, while Fig. 10 shows the class distribution per client. This distribution illustrates that the data across clients is naturally non-IID.

Figure 10: WSI distributions among the top-5 anatomic sites. Note that the most common site is different for each client.

A.3 DeepSet Architecture

For a set classification problem, applying a permutation φ to the items in the set X , giving $\{x_{\varphi_1}, \dots, x_{\varphi_n}\}$, should not affect the class label. Hence, permutation invariant models are appropriate for classification on sets. We utilized the DeepSet architecture for WSI classification [Zaheer et al., 2017]. DeepSet introduces a universal set approximation as a neural architecture. It is composed of an element-wise operation f , followed by a symmetric function ϕ , where f is a neural network, and ϕ is a pooling operation, such as mean or max. For the histopathology experiments, both private and proxy models are DeepSet-based. For the **private model**, the element-wise function $f_{private}$ consists of (fc(512), ReLU), and for the **proxy model**, f_{proxy} consists of (fc(256), ReLU). Both $f_{private}$ and f_{proxy} are applied element-wise on an input set (features of patches in a mosaic). The mean is used as the pooling function ϕ , and following this the output is classified into 32 cancer diagnosis types using a MLP consisting of (fc(128), ReLU, fc(64), ReLU, fc(32), Softmax).

B Additional Results

Figure 11: MLP for both private and proxy.

Figure 12: Effect of batch size.

Figure 13: Effect of DML weight.

Here we show additional results for the ablation studies on MNIST. In the main text the private models for ProxyFL and FML were different from the proxy/centralized models. Fig. 11 shows the training curves when all models used the same MLP architecture. We see that ProxyFL-private can achieve the best performance regardless of the private model structure, as compared with Fig. 3 in the main text.

Increasing mini-batch size is beneficial for the stability of stochastic gradient updates and for computation speed. However, smaller batch sizes lead to stronger (ϵ, δ) privacy guarantees when using DP-SGD [Abadi et al., 2016]. In the main text, for computation efficiency we used $B = 250$ images per DP-SGD step, Eq. (7). Fig. 12 shows how batch size can affect the privacy guarantee. Smaller batch sizes do not have a significant effect on the final accuracy of ProxyFL-private, while the privacy guarantee dramatically improves. Therefore, we can have a strong privacy guarantee with ProxyFL by using small batch sizes at the cost of longer training time.

Finally, Fig. 13 shows the performance with different DML weights α, β . Note that this only affects ProxyFL and FML. We fix $\alpha = \beta$ in the experiments. We can see that ProxyFL is very robust to different choices of α , unlike its centralized counterpart FML. With a larger α , the performance gap between the private and proxy models gets smaller, which is reasonable as they tend to mimic each other.