

TurboSVM-FL: Boosting Federated Learning through SVM Aggregation for Lazy Clients

Mengdi Wang*, Anna Bodonheli*, Efe Bozkir, Enkelejda Kasneci

Chair for Human-Centered Technologies for Learning, Technical University of Munich, Munich, Bavaria, Germany
 {mengdi.wang, anna.bodonheli, efe.bozkir, enkelejda.kasneci}@tum.de

Abstract

Federated learning is a distributed collaborative machine learning paradigm that has gained strong momentum in recent years. In federated learning, a central server periodically coordinates models with clients and aggregates the models trained locally by clients without necessitating access to local data. Despite its potential, the implementation of federated learning continues to encounter several challenges, predominantly the slow convergence that is largely due to data heterogeneity. The slow convergence becomes particularly problematic in cross-device federated learning scenarios where clients may be strongly limited by computing power and storage space, and hence counteracting methods that induce additional computation or memory cost on the client side such as auxiliary objective terms and larger training iterations can be impractical. In this paper, we propose a novel federated aggregation strategy, **TurboSVM-FL**, that poses no additional computation burden on the client side and can significantly accelerate convergence for federated classification task, especially when clients are “lazy” and train their models solely for few epochs for next global aggregation. TurboSVM-FL extensively utilizes support vector machine to conduct selective aggregation and max-margin spread-out regularization on class embeddings. We evaluate TurboSVM-FL on multiple datasets including FEMNIST, CelebA, and Shakespeare using user-independent validation with non-iid data distribution. Our results show that TurboSVM-FL can significantly outperform existing popular algorithms on convergence rate and reduce communication rounds while delivering better test metrics including accuracy, F1 score, and MCC.

Introduction

With the increasing importance of data privacy, a giant stride in distributed machine learning has been observed in recent years. As one promising distributed machine learning paradigm, federated learning (FL) has been growing at an astounding rate after its introduction (McMahan et al. 2017). In the common FL settings, data is distributed over numerous end clients, while the central server possesses no data by itself. After the server initiates a model and sends the model to clients, each client trains the model locally using its own data. The server periodically aggregates the locally trained

models and synchronizes local models of clients with the latest aggregated one. With such a process, FL provides a primary privacy guarantee to a large extent since the server does not require data sharing and is hence preferred in many privacy-preserving scenarios where sensitive data is utilized.

Based on the characteristics of participating entities, FL can be further categorized into cross-silo FL and cross-device FL (Kairouz et al. 2021). In cross-silo FL, the target clients are often large-scale institutions such as hospitals, data centers, educational organizations, and high-tech companies. Such stakeholders commonly possess decent resource for computing, storage, and internet connection, while the number of attending institutions is relatively low. Therefore, the probability that each client takes part in all aggregation rounds is high. In contrast to cross-silo FL, cross-device FL focuses more on training on end-user devices like smartphones and personal computers using user data. The scale of participating clients in cross-device FL can be fairly large, while each client may be strongly limited by its computing power and connectivity. As a consequence, only a (small) portion of clients could share their models during a global aggregation. An additional critical aspect of cross-device FL is that each client might only contain data collected from a single user, which exacerbates data heterogeneity across clients.

Despite recent advancements in FL, its implementation in practice is still facing some challenges. Among these, the slow convergence is a primary concern: a significantly greater number of aggregation rounds are often needed to reach convergence compared to non-FL setups. Several factors contribute to this inefficiency according to (Wu et al. 2023), such as client drift caused by data heterogeneity (Karimireddy et al. 2020), lack of adaptive optimization (Reddi et al. 2020), and the increase in model complexity and data size. One trivial solution proposed in (McMahan et al. 2017) is to increase client computation load by a larger number of local training iterations. Although this solution vastly speeds up the convergence, it multiplies the computation load on the client side, which can be problematic when clients are constrained by computing power, like in cross-device FL case. Other existing solutions mainly target at either client drift or adaptive optimization. The former often requires to optimize additional objective functions on the client side, which in turn also increases the client com-

*These authors contributed equally.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

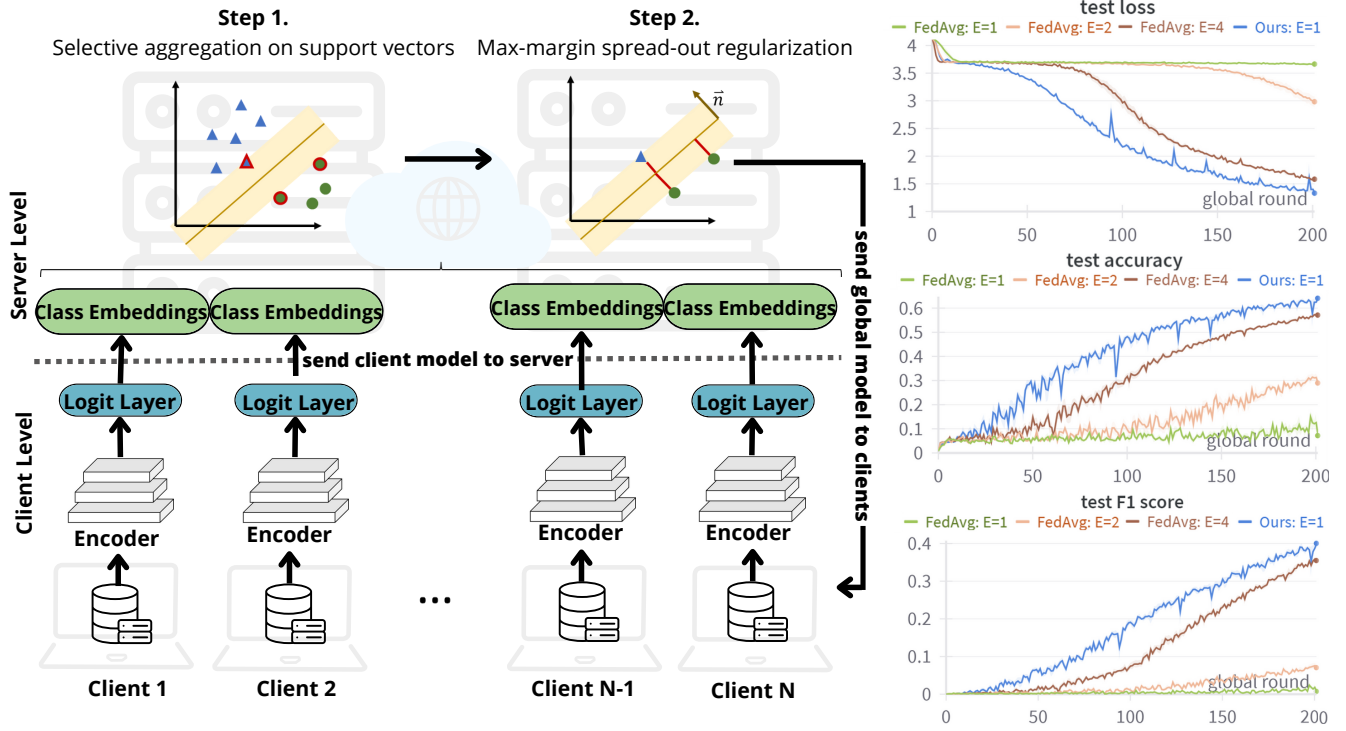


Figure 1: Left: pipeline of TurboSVM-FL. Right: test performance of TurboSVM-FL against FedAvg. E indicates the number of client local training epochs. The results were obtained on FEMNIST dataset using a suboptimal client learning rate.

putation load and even the need for storage, while the latter can be particularly hard to tune because it is often necessary to decide the choice of optimizers and learning rates jointly between server and clients. There also exist solutions that require additional data on the server side, which may increase the risk of privacy leaks if not handled properly.

In this paper, we focus on embedding-based neural network classifiers, which means the input samples are encoded into the same space as class representations and their similarity determines the probability of the sample belonging to the class (Yu et al. 2020). We propose a novel federated aggregation strategy for classification tasks called TurboSVM-FL. TurboSVM-FL induces no additional computation cost or storage consumption on the client side compared to the vanilla FL algorithm, and shows great potential in reducing communication rounds, especially when clients are “lazy” and train their local models for very few iterations. TurboSVM-FL extensively exploits the property of support vector machine (SVM) and consists of three steps. Firstly, TurboSVM-FL reformalizes classification while treating client models in a model-as-sample way, and fits SVMs using latent class embeddings. Then, it conducts selective aggregation on latent class features that form support vectors. Lastly, TurboSVM-FL applies max-margin spread-out regularization on aggregated class representations upon SVM hyperplanes. By adopting adaptive methods in max-margin spread-out regularization, TurboSVM-FL can also benefit from adaptive optimization. The main contributions of this work can be summarized as follows:

- We introduce a novel perspective to interpret classification in FL setting and a model-as-sample strategy, which lay the foundation for further FL improvements such as selective aggregation and outlier detection.
- We propose a novel federated aggregation algorithm named TurboSVM-FL that vastly boosts convergence for federated classification tasks using deep neural networks. The proposed method extensively exploits support vector machine (SVM) to conduct selective model aggregations and max-margin spread-out regularization.
- We conduct experiments on various benchmarks in user-independent validation and show the potential of TurboSVM-FL in reducing communication cost. The benchmarks contain three different tasks covering image classification and natural language processing with non-iid data distribution over clients.

Related Work

Federated Learning

The concept of FL was originally introduced in (McMahan et al. 2017). Unlike centralized learning where the goal is to fit an optimal model on a collection of data, FL aims to train a model that delivers superior performance across data subsets. In the remaining of this work, we narrow down our focus to federated classification tasks. In a federated classification task with K classes and N clients, denote the local dataset of each client as D_1, \dots, D_N with $D_n = \{(x, y)\}, n \in [N]$, where $(x, y) \in \mathbb{R}^P \times [K]$ is

a sample point of class $y \in [K]$. Let $D_G = \bigcup_{n=1}^N D_n$ with $|D_G| = \sum_{n=1}^N |D_n|$ describe the collection of local datasets and $\ell(x, y, \theta)$ be the objective function measured on the sample (x, y) with model θ . Then, the goal of centralized learning is to find an optimal model θ^* that satisfies:

$$\theta^* = \underset{\theta}{\operatorname{argmin}} \mathbb{E}_{(x,y) \sim D_G} [\ell(x, y, \theta)] \quad (1)$$

In contrast, FL aims to fit a model that performs optimally across clients:

$$\theta^* = \underset{\theta}{\operatorname{argmin}} \sum_{n=1}^N \frac{|D_n|}{|D_G|} \mathbb{E}_{(x,y) \sim D_n} [\ell(x, y, \theta)] \quad (2)$$

The typical workflow of FL can be broken down into three steps. First, a server initializes a model and broadcasts this model to all clients. Then, each client trains the received model on its own dataset for E epochs and sends the trained model back to the server. In the next step, the server aggregates locally trained models into a new global model and synchronizes clients with the latest global model. The last two steps are repeated for multiple rounds until convergence. The first federated aggregation algorithm, FedAvg, was introduced in (McMahan et al. 2017) and applies weighted average over client models:

$$\theta_G = \sum_{n=1}^N \frac{|D_n|}{|D_G|} \theta_n, \theta_n = \underset{\theta}{\operatorname{argmin}} \mathbb{E}_{(x,y) \sim D_n} [\ell(x, y, \theta)] \quad (3)$$

where θ_G and θ_n denote the aggregated global model and the local model of n -th client, respectively.

One of those challenges that FL is facing is the large amount of aggregation rounds needed to approach convergence. While increasing local training iterations E can significantly advance convergence, it also vastly increases the computation load on the client side, which can be extremely problematic in cross-device FL. Many follow-up works aim to speed up FL convergence, and they can be mainly categorized into two groups. The first group endeavor to address client drift (Karimireddy et al. 2020) caused by data heterogeneity, while the other group attempt to benefit FL with adaptive learning methods, which we describe as follows.

Client Drift in Federated Learning

Client drift (Karimireddy et al. 2020) describes the phenomenon that client local models approach individual local optima rather than global optima and their average is drifted away from global optima as well, which is caused by data heterogeneity across client local datasets and can dramatically impact convergence behavior. Various recent works attempt to solve client drift on the client side. SCAFOLD (Karimireddy et al. 2020) introduces a control variate term to stabilize gradient update. FedProx (Li et al. 2020) proposes an additional loss term based on L2 distance between global model θ_G and client model θ_n during local training. MOON (Li, He, and Song 2021) and FedProc (Mu et al. 2023) suggest the use of contrastive learning to combat data heterogeneity. The former introduces an objective based on latent features extracted respectively by the global model,

current client model, and previous client model, while the latter penalizes the dissimilarity between latent features and class representations. A common drawback of the aforementioned methods lies in that they increase either computation burden or memory consumption or even both on the client side, which can be quite a challenge for end-user devices like smartphones and tablets.

There also exists works that aim to solve client drift on the server side. For instance, FedAwS (Yu et al. 2020) addresses an extreme data distribution case where clients may have only data from a single class. FedAwS utilizes spread-out regularizer (Zhang et al. 2017) and raises a penalty term based on cosine similarities among class embeddings on the server side.

Adaptive Federated Learning

In centralized learning, advanced adaptive and momentum-based optimization techniques such as AdaGrad (Duchi, Hazan, and Singer 2011), Adam (Kingma and Ba 2014), and Yogi (Zaheer et al. 2018) have shown great success in convergence acceleration. In contrast, in vanilla FedAvg, client models are trained with stochastic gradient descent (Robbins and Monro 1951; Kiefer and Wolfowitz 1952) and server aggregation is (weighted) averaging. Numerous works have been devoted to benefiting FL with advanced server-side optimizers. As a forerunner in this field, (Reddi et al. 2020) proposed a family of adaptive aggregation methods called FedOpt. Different from weighted average in Equation 3, FedOpt computes pseudo-gradient (Chen and Huo 2016; Nichol, Achiam, and Schulman 2018) from client models and updates the global model with a chosen optimizer:

$$\Delta \leftarrow \sum_{n=1}^N \frac{|D_n|}{|D_G|} \theta_n - \theta_G \quad (4)$$

$$\theta_G \leftarrow \text{server_optimizer}(\theta_G, -\Delta, \eta) \quad (5)$$

where η indicates the learning rate. Depending on the choice of optimizer, FedOpt can be derivated into multiple variants. For instance, in (Reddi et al. 2020), the researchers introduced FedAdaGrad, FedAdam, and FedYogi with their names indicating the choice of optimization technique. FedAMS (Wang, Lin, and Chen 2022) suggests the use of AMSGrad (Reddi, Kale, and Kumar 2019) optimizer on the server side, which is an improved version of Adam. According to (Wang et al. 2021a), it can be hard to tune FedOpt-family methods due to the additional implementation of optimizer on the server side, and it is often necessary to search for optimal learning rates jointly for client optimizer and server optimizer.

Compared to server-level adaptive learning, adaptive optimization on the client side is less studied. Client adaptivity poses its own challenges, particularly due to the potential for the states of client optimizers to significantly diverge from each other as a result of data heterogeneity. To address this challenge, (Wang et al. 2021b) proposes to reset client optimizer status in each global round, while LocalAMSGrad (Chen, Li, and Li 2020) and FAFED (Wu et al. 2023) suggest the sharing and aggregation of client optimizer states similarly to client models. These methods

can be meaningless if clients are limited by computation resource and can only train their local models for few epochs.

Support Vector Machine

Support vector machine (SVM) (Cortes and Vapnik 1995) is a widely-used and robust supervised learning model that can be used for both regression and classification tasks. Unlike traditional linear classifiers, where the decision boundary is a linear combination of all data points, the separating hyperplane of SVM is a combination of selected samples, which are also called support vectors and lie the closest to the decision boundary. While common classifiers minimize solely the classification objective, SVM struggles to control the trade-off between discriminative error minimization and margin maximization while allowing some misclassifications. The margin refers to the distance between the support vectors of different classes and the decision boundary. The primal problem of SVM can be formalized as:

$$\operatorname{argmin}_{w, \zeta_1, \dots, \zeta_m} \frac{1}{2} \|w\|^2 + \lambda \sum_{i=1}^m \zeta_i, \text{ s.t. } y_i w^\top x_i \geq 1 - \zeta_i, \zeta_i \geq 0 \quad (6)$$

where ζ_i defines the distance of a misclassified sample to its correct margin plane. The coefficient λ controls the magnitude of regularization. A smaller λ prioritizes larger margins and may result in a greater number of support vectors.

It is important to note that there are several prior works that integrate FL and SVM, such as (Bakopoulou, Tillman, and Markopoulou 2021; Navia-Vázquez, Díaz-Morales, and Fernández-Díaz 2022; Bemani and Björnell 2022). Our approach is distinctly different from them in the sense that our algorithm leverages SVM to improve global aggregation and offers a solution to the problem “how to FL”. In contrast, in previous works, SVM serves as the core model to be trained in FL and thus addresses the question of “what to FL”.

Methodology

In this paper, we propose a novel federated aggregation algorithm for classification task called TurboSVM-FL that is able to boost convergence significantly. TurboSVM-FL sophisticatedly leverages SVM to conduct selective aggregation and max-margin spread-out regularization. By adopting an adaptive optimizer, TurboSVM-FL can also benefit from adaptivity. Compared to vanilla FedAvg (McMahan et al. 2017), TurboSVM-FL requires no additional computation, storage, or communication cost on the client side. A pseudocode for TurboSVM-FL is given in Algorithm 3, and a graphical illustration is depicted Figure 1.

In the following, we present our algorithm in detail, starting by reformalizing the classification task as “finding nearest neighbor”. We reduce our discussion to embedding-based deep learning networks and ignore the logit activation, which means the model θ can be divided into two parts: g and W , with $f_\theta(x) = Wg(x)$, where $g: \mathbb{R}^P \rightarrow \mathbb{R}^d$ is the feature encoder that maps input $x \in \mathbb{R}^P$ to a latent representation in \mathbb{R}^d and W is the last projection layer containing class embeddings. In a classification task with K classes, W will be of shape $R^{K \times d}$ and is also called logit layer.

Then, the class inference \hat{y} of a sample (x, y) is indifferent from finding the nearest neighbor to $g(x)$ in w_1, \dots, w_K with $w_k \in \mathbb{R}^d$ being the k -th row in W indicating the class embedding of class k . Implicitly, the metric used to measure distance is vector inner product, and the choice of nearest neighbor is regardless of activation function and loss function. For simplicity, we ignore bias terms, and the class inference can then be represented as:

$$\hat{y} = \operatorname{argmax}_{k \in [K]} \operatorname{sim}(w_k, g(x)), \operatorname{sim}(w_k, g(x)) = w_k^\top \cdot g(x) \quad (7)$$

Hence, training the last layer in classification can be regarded as encouraging the correct class embedding to approach instance embedding while discouraging all other class embeddings to be close.

Next, TurboSVM-FL treats client models as sample points for a secondary classification task at higher level, and fits SVM using these samples. The SVM is constructed as a multi-class classification among K classes, and the SVM training samples are exactly the collected class embeddings, i.e., $\{(w_k^n, k) | k \in [K], n \in [N]\}$. In other words, for each class k , there are N sample points $\{(w_k^1, k), \dots, (w_k^N, k)\}$, and each sample point is the k -th row of the weight matrix of the logit layer from a client model.

In vanilla FL, the class embeddings in the logit layer of the global model are obtained by averaging client models, in other words, $w_k^G = \sum_{n=1}^N \frac{|D_n|}{|D_G|} w_k^n$. Due to data heterogeneity among clients, the class embeddings of some clients can be drifted away from global optima and hence seriously disturb the aggregation. TurboSVM-FL addresses this problem with the help of support vectors during global update. SVM aims at a margin-maximization decision boundary that is a linear combination of selected samples, which are also called support vectors. The support vectors can be regarded as the most informative samples of each class and function similarly to contrastive anchors. In other words, fitting SVM is to some extent equivalent to selective aggregation over samples. TurboSVM-FL brings this property to federated aggregation by averaging only class embeddings that form support vectors, as depicted in Algorithm 1.

Moreover, TurboSVM-FL employs spread-out regularization across projected global class embeddings to maintain the margin-maximization property. This is crucial for two reasons: first, although support vectors are the most informative data points, they are close to decision boundary and can be misclassified; second, the weights used during FL ag-

Algorithm 1: TurboSVM-FL part 1: selective aggregation

Input: fitted SVM, sizes of local datasets $|D_1|, \dots, |D_N|$.

for $k \in [K]$ **do**
 retrieve support vectors $\{v_k^m\}$ for class k from SVM
 $w_k^G \leftarrow \frac{\sum_m |D_m| v_k^m}{\sum_m |D_m|} \triangleright m$: index to client model
end for
return global class embeddings w_1^G, \dots, w_K^G

Algorithm 2: TurboSVM-FL part 2: max-margin spread-out regularization

Input: fitted SVM, global class embeddings $w_k^G, k \in [K]$, server learning rate η_G .

$\ell_{sp} \leftarrow 0$

for $k \in [K]$ **do**

for $k' \in [K]$ with $k' > k$ **do**

 retrieve hyperplane $h_{k,k'}$ from SVM

$\ell_{sp} \leftarrow \ell_{sp} + \exp(-\frac{(w_k^{G\tau} \cdot h_{k,k'} - w_{k'}^{G\tau} \cdot h_{k,k'})^2}{2||h_{k,k'}||^2})$

end for

end for

for $k \in [K]$ **do**

$w_k^G \leftarrow \text{server_optimizer}(w_k^G, -\nabla_{w_k^G} \ell_{sp}, \eta_G)$

end for

return global class embeddings w_1^G, \dots, w_K^G

gregation may differ from the coefficients assigned to support vectors during SVM fitting, which could undermine the SVM property. Spread-out regularization like in (Yu et al. 2020) offers the potential to distinguish class embeddings. Nevertheless, we propose that omnidirectional regularization is not the most efficient method. Instead, we leverage once again the SVM property, namely we project the aggregated embeddings back onto the SVM decision boundaries, and penalize the similarities among projected embeddings:

$$\ell_{sp}(w_1, \dots, w_K) = \sum_{k \in [K]} \sum_{k' \neq k} \text{sim}(\frac{w_k^\tau \cdot h_{k,k'}}{||h_{k,k'}||}, \frac{w_{k'}^\tau \cdot h_{k,k'}}{||h_{k,k'}||}) \quad (8)$$

where $h_{k,k'}$ is the normal of the separating hyperplane for classes k and k' retrieved from fitted SVM.

In (Yu et al. 2020), the authors proved that the classification error can be upper-bounded by the separation of class embeddings. We extend their analysis for TurboSVM-FL by showing that by applying selective aggregation and max-margin spread-out regularization, TurboSVM-FL effectively enlarges the difference between projected logits. For simplicity, we narrow down to binary classification and denote the distance relaxation terms for embeddings of each class as ζ_n^+ and ζ_n^- , $n \in [N]$, respectively. Let h be the decision boundary of the fitted SVM. Then, under further simplification that all class embeddings serve as support vectors and all clients have same amount of samples, given a new positive sample x^* , the difference between the positive and negative logits when projected on h can be bounded as follows:

$$\begin{aligned} & \text{proj}(\text{logit}^+(x^*), h) - \text{proj}(\text{logit}^-(x^*), h) \\ & \geq \frac{[2N - \sum_{n=1}^N (\zeta_n^+ + \zeta_n^-)](1 - \zeta^*)}{N||h||^2} \end{aligned} \quad (9)$$

where ζ^* is the SVM relaxation term for x^* . By averaging support vectors and applying max-margin spread-out regularization, TurboSVM-FL reduces $||h||$ and ζ_n^\pm in essence according to SVM theory, and the term above that bounds logit distance from below is hence increased. A more detailed analysis of this is given in the Appendix.

Algorithm 3: The TurboSVM-FL Framework

Input: clients $n \in [N]$, client local datasets D_1, \dots, D_N , $|D_G| = |D_1| + \dots + |D_N|$, number of global epochs T , number of client epochs E , number of classes K , server learning rate η_G , client learning rate η , mini-batch size B

ServerUpdate:

initialize global model $\theta_0^G = (g_0^G, W_0^G)$

for $t = 0, 1, \dots, T - 1$ **do**

for client $n \in [N]$ **do**

$\theta_{t+1}^n \leftarrow \text{ClientUpdate}(n, \theta_t^G)$

end for

$g_{t+1}^G \leftarrow \sum_{n=1}^N \frac{|D_n|}{|D_G|} g_{t+1}^n$

 fit SVM using samples $\{(w_k^n, k) | k \in [K], n \in [N]\}$

$W_{t+1}^G \leftarrow \text{Algorithm 1 TurboSVM-FL part 1}$

$W_{t+1}^G \leftarrow \text{Algorithm 2 TurboSVM-FL part 2}$

end for

return $\theta_T^G = (g_T^G, W_T^G)$

ClientUpdate(n, θ_t^G): \triangleright Run on client n with model θ_t^G

$\theta_{t+1}^n \leftarrow \theta_t^G$

for $e = 0, 1, \dots, E - 1$ **do**

for mini-batch \mathcal{B} of size B in D_n **do**

$\theta_{t+1}^n \leftarrow \text{client_optimizer}(\theta_{t+1}^n, -\nabla_{\theta_{t+1}^n} \ell(\mathcal{B}, \eta))$

end for

end for

return θ_{t+1}^n

Since projections onto the same axis are always co-linear, the common cosine similarity between them is always either 1 or -1 and thus not meaningful. We hence use Gaussian function as a similarity measurement because of its outstanding capability (Yang et al. 2021) as similarity kernel:

$$\begin{aligned} & \text{sim}(\frac{w_k^\tau \cdot h_{k,k'}}{||h_{k,k'}||}, \frac{w_{k'}^\tau \cdot h_{k,k'}}{||h_{k,k'}||}) = \\ & \exp(-\frac{(w_k^\tau \cdot h_{k,k'} - w_{k'}^\tau \cdot h_{k,k'})^2}{2||h_{k,k'}||^2}) \end{aligned} \quad (10)$$

TurboSVM-FL then optimizes class embeddings regarding the objective ℓ_{sp} and can benefit from adaptivity and momentum with a proper choice of optimizer. The max-margin spread-out regularization part of TurboSVM-FL is illustrated in Algorithm 2. The pseudocode for the whole algorithm is given in Algorithm 3.

Experiments and Results

We benchmarked TurboSVM-FL on three datasets against six FL algorithms, including FedAvg (McMahan et al. 2017), FedAdam (Reddi et al. 2020), FedAMS (Wang, Lin, and Chen 2022), FedProx (Li et al. 2020), MOON (Li, He, and Song 2021), and FedAwS (Yu et al. 2020). In this section, we provide task descriptions and results. For more details such as reproducibility and model structures, we redirect readers to the Appendix and our GitHub repository¹.

¹ <https://github.com/Kasneci-Lab/TurboSVM-FL>.

Dataset	Source	Task	# Classes	# Users	Type & Dim	Mean/Std per User
FEMNIST	(LeCun 1998) (Cohen et al. 2017)	image classification	62	3550	BW image 28×28	226.8 / 88.9
CelebA	(Liu et al. 2015)	smile detection	2	9343	RGB image 84×84	21.4 / 7.6
Shakespeare	(Shakespeare 2014) (McMahan et al. 2017)	next char prediction	80	1129	string 80	3743.2 / 6212.3

Table 1: Overview of used datasets.




Figure 2: Test metrics on FEMNIST dataset.

Tasks

We benchmarked TurboSVM-FL on three different datasets covering data types of both image and nature language, namely FEMNIST (LeCun 1998; Cohen et al. 2017), CelebA (Liu et al. 2015), and Shakespeare (Shakespeare 2014; McMahan et al. 2017) (Table 1). The task in FEMNIST dataset is handwritten digit and letter classification using grayscale images. The number of classes in FEMNIST is 62 (10 digits, 26 lowercase letters, and 26 uppercase letters) and the resolution of images is 28×28 . The CelebA dataset contains 84×84 RGB images of faces of celebrities, and the task is binary classification between smiling and non-smiling faces. The Shakespeare dataset consists of scripts of different roles from Shakespeare’s works, and the task is next-character prediction given an input string of length 80. The number of unique characters and symbols in this dataset is 80. All three datasets can be acquired on LEAF (Caldas et al. 2018). For the two image classification tasks, CNN models were implemented, while for the language task LSTM model was utilized. We adopted the model structures as given in LEAF. Details about data distributions and models can be found in the Appendix.

We also adopted the data split given in LEAF. More specifically, we conducted 90% – 10% train-test-split in a user-independent way, which means we had a held-out set of clients for validation rather than a fraction of validation data on each client (Wang et al. 2021a). The main reason for conducting user-independent validation is that such a test is a more valid approach for unseen data and, thus, more representative for real-world applications. Moreover, it is more challenging to fit a model in a user-independent setting compared to a user-dependent data split.

Algorithm	FEMNIST	CelebA	Shakespr.
FedAvg	144.4±4.6	91.6±18.2	51.0±5.0
FedAdam	110.8±16.2	>200	54.8±3.2
FedAwS	81.4±2.2	84.2±24.4	45.0±2.3
FedProx	145.4±3.4	94.4±18.4	157.2±5.3
FedAMS	116.4±19.6	>200	51.6±2.1
MOON	145.6±3.7	94.2±18.8	52.4±3.1
TurboSVM-FL	54.6±1.6	46.4±9.4	43.4±2.9

Table 2: Number of communication rounds needed to reach certain test accuracy on FEMNIST (70%), CelebA (70%), and Shakespeare (50%) datasets. Smaller is better.

Results

To compare the convergence rate of different FL algorithms, we reported two groups of metrics: number of global aggregation rounds needed to reach certain validation accuracy (70% for FEMNIST and CelebA, 50% for Shakespeare), and the achieved F1 score, accuracy, and MCC (Matthews Correlation Coefficient) after 100 aggregation rounds. The results are given in Tables 2 and 3, and also visualized in Figures 2, 4 and 5 (Appendix) for each task respectively.

Our results clearly indicate that on FEMNIST and CelebA datasets TurboSVM-FL yields a significantly faster convergence in contrast to other FL methods, while on the Shakespeare dataset TurboSVM-FL slightly outperforms others. Compared to the baseline FedAvg, TurboSVM-FL successfully reduces the number of global rounds by 62.2%, 49.3%, and 14.9% to reach the same test accuracy as given in Table 2. When all FL methods are run for the same rounds, TurboSVM-FL yields in much better test metrics on both image classification tasks in comparison to other meth-

Algorithm	[%]	FEMN.	CelebA	Shakespr.
FedAvg	F_1	33.1±1.3	70.5±2.2	17.4±0.5
	A	59.3±1.8	70.6±2.1	52.9±0.8
	M	57.9±1.9	41.6±4.2	48.9±0.8
FedAdam	F_1	48.1±4.9	34.2±0.3	18.1±0.2
	A	66.6±3.9	51.6±0.1	53.8±0.9
	M	65.4±4.1	1.0±2.3	49.8±0.9
FedAwS	F_1	54.8±0.8	72.2±3.0	18.3±0.6
	A	73.1±0.6	72.3±2.8	53.3±0.7
	M	72.2±0.6	44.9±5.6	49.3±0.7
FedProx	F_1	32.8±1.3	70.4±2.1	12.5±0.6
	A	59.2±1.8	70.6±2.1	46.0±1.2
	M	57.7±1.9	41.4±4.1	41.2±1.2
FedAMS	F_1	46.8±6.5	36.3±5.0	18.2±0.4
	A	65.8±5.0	52.6±2.3	54.0±0.5
	M	64.6±5.2	4.7±10.5	50.0±0.6
MOON	F_1	33.4±3.7	70.7±2.3	17.6±0.6
	A	58.2±3.0	70.9±2.2	52.8±1.0
	M	56.8±3.1	41.8±4.4	48.8±1.0
<i>TurboSVM-FL</i>	F_1	61.9±0.7	77.2±1.3	19.2±0.3
	A	76.6±1.0	77.3±1.2	53.7±0.2
	M	75.8±1.0	55.0±2.4	49.7±0.2

Table 3: Achieved test F1 score, accuracy (A), and MCC score (M) after 100 aggregation rounds. Greater is better.

ods, while its performance is a fair match to the adaptive algorithms on the next-character prediction task. Moreover, we show in Figure 4 that while adaptive FL methods like FedAdam and FedAMS are not stable, TurboSVM-FL can still robustly benefit from adaptivity on the server side.

Impact of Embedding Size on Selectivity

We further explored the impact of embedding size on the number of client models that form support vectors. To approach this, we ran experiments on the FEMNIST dataset with varying embedding size and number of participating clients C . We recorded the number of support vectors for class 1 at 200th round in Table 4. It is clear that a higher embedding size is associated with better performance and fewer support vectors when there exist enough participating clients, while a scarcity of clients can lead to full use of class embeddings as support vectors. Furthermore, larger embedding size also leads to higher complexity and burden, which needs to be balanced off depending on the specific tasks.

		Embedding Dimension				
C		4	16	64	256	1024
8	#SV	8	8	8	8	8
	F_1 [%]	6.1	14.2	58.9	63.9	67.4
64	#SV	62	63	58	52	37
	F_1 [%]	16.3	43.2	60.5	65.9	68.3
512	#SV	197	150	147	100	77
	F_1 [%]	16.7	43.2	59.7	64.5	67.0

Table 4: Influence of embedding size on support vectors.

Discussion

As TurboSVM-FL focuses on class representations on the server side while other parts of the global model are still aggregated with average and client training is done with vanilla SGD, the use of TurboSVM-FL in combination with other FL algorithms is promising. For instance, on the client side FedProx can be applied to counteract data heterogeneity, while on the server side, class embeddings are aggregated with TurboSVM-FL. Another example is the use of adaptive FL methods like FedAdam for encoder aggregation while logit layers are aggregated with TurboSVM-FL. Moreover the idea of model-as-sample can be further explored, for example, for anomaly client detection and client clustering.

TurboSVM-FL is particularly suitable for cross-device FL where edge devices are often constrained by computation and storage resources, but its improvement is also not excluded from cross-silo case. A typical application scenario of TurboSVM-FL is federated transfer learning, where a pre-trained model like VGG16 and Resnet50 is adopted, and all of the layers except the last few ones are frozen. In this case, each client only needs to train and share the last few layers, which makes TurboSVM-FL extremely efficient. The capability of TurboSVM-FL is also not constrained to single-output tasks. For multi-output tasks such as multi-label classification and multi-task learning, TurboSVM-FL can also be applied. To approach this, separate classification heads for different tasks should be implemented where the backbone encoder shares its weights among tasks, and then TurboSVM-FL should be applied to each head in parallel.

One improvement direction of TurboSVM-FL is to relax the implicit assumption about linear separability of class embeddings with kernelization during SVM fitting and class inference. A piloting ablation study is included in the Appendix in this regard. Furthermore, while posing no additional computation cost on the client side, TurboSVM-FL requires the server to be powerful such that it can fit SVMs efficiently, especially when the SVMs are in one-vs-one (OVO) pattern. We chose OVO instead of OVR (one-vs-rest) mainly for two reasons: 1. in general, OVO performs better than OVR; 2. for TurboSVM-FL, OVO never suffers from class imbalance while OVR always does, since the numbers of samples for each class are always the same. Although OVO imposes more computation on the server side, we think that to approach FL, a powerful server is a must-have, and OVO is no burden for such a server. In case the number of classes is large, the computation burden can be further resolved by sampling a proportion of classes on which SVMs are fitted.

Conclusion

In this work, we proposed a novel federated aggregation strategy called TurboSVM-FL, which extensively exploits SVM to conduct selective aggregation and max-margin spread-out regularization for class embeddings and can vastly reduce communication rounds. We tested our approach on three publicly available datasets, and our results show that TurboSVM-FL outperforms existing FL methods largely on convergence rate regarding various metrics.

Acknowledgements

We acknowledge the funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - Project number KA 4539/5-1.

References

- Bakopoulou, E.; Tillman, B.; and Markopoulou, A. 2021. Fedpacket: A federated learning approach to mobile packet classification. *IEEE Transactions on Mobile Computing*, 21(10): 3609–3628.
- Bemani, A.; and Björssell, N. 2022. Aggregation Strategy on Federated Machine Learning Algorithm for Collaborative Predictive Maintenance. *Sensors*, 22(16): 6252.
- Caldas, S.; Duddu, S. M. K.; Wu, P.; Li, T.; Konečný, J.; McMahan, H. B.; Smith, V.; and Talwalkar, A. 2018. Leaf: A benchmark for federated settings. *arXiv preprint arXiv:1812.01097*.
- Chen, K.; and Huo, Q. 2016. Scalable training of deep learning machines by incremental block training with intra-block parallel optimization and blockwise model-update filtering. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 5880–5884. IEEE.
- Chen, X.; Li, X.; and Li, P. 2020. Toward communication efficient adaptive gradient method. In *Proceedings of the 2020 ACM-IMS on Foundations of Data Science Conference*, 119–128.
- Cohen, G.; Afshar, S.; Tapson, J.; and Van Schaik, A. 2017. EMNIST: Extending MNIST to handwritten letters. In *2017 International Joint Conference on Neural Networks (IJCNN)*, 2921–2926. IEEE.
- Cortes, C.; and Vapnik, V. 1995. Support-vector networks. *Machine learning*, 20: 273–297.
- Duchi, J.; Hazan, E.; and Singer, Y. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7).
- Kairouz, P.; McMahan, H. B.; Avent, B.; Bellet, A.; Bennis, M.; Bhagoji, A. N.; Bonawitz, K.; Charles, Z.; Cormode, G.; Cummings, R.; et al. 2021. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2): 1–210.
- Karimireddy, S. P.; Kale, S.; Mohri, M.; Reddi, S.; Stich, S.; and Suresh, A. T. 2020. Scaffold: Stochastic controlled averaging for federated learning. In *International Conference on Machine Learning*, 5132–5143. PMLR.
- Kiefer, J.; and Wolfowitz, J. 1952. Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics*, 462–466.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- LeCun, Y. 1998. The MNIST database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>.
- Li, Q.; He, B.; and Song, D. 2021. Model-contrastive federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10713–10722.
- Li, T.; Sahu, A. K.; Zaheer, M.; Sanjabi, M.; Talwalkar, A.; and Smith, V. 2020. Federated optimization in heterogeneous networks. *Proceedings of Machine learning and systems*, 2: 429–450.
- Liu, Z.; Luo, P.; Wang, X.; and Tang, X. 2015. Deep learning face attributes in the wild. In *Proceedings of the IEEE international conference on computer vision*, 3730–3738.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, 1273–1282. PMLR.
- Mu, X.; Shen, Y.; Cheng, K.; Geng, X.; Fu, J.; Zhang, T.; and Zhang, Z. 2023. Fedproc: Prototypical contrastive federated learning on non-iid data. *Future Generation Computer Systems*, 143: 93–104.
- Navia-Vázquez, A.; Díaz-Morales, R.; and Fernández-Díaz, M. 2022. Budget Distributed Support Vector Machine for Non-ID Federated Learning Scenarios. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 13(6): 1–25.
- Nichol, A.; Achiam, J.; and Schulman, J. 2018. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*.
- Reddi, S.; Charles, Z.; Zaheer, M.; Garrett, Z.; Rush, K.; Konečný, J.; Kumar, S.; and McMahan, H. B. 2020. Adaptive federated optimization. *arXiv preprint arXiv:2003.00295*.
- Reddi, S. J.; Kale, S.; and Kumar, S. 2019. On the convergence of adam and beyond. *arXiv preprint arXiv:1904.09237*.
- Robbins, H.; and Monro, S. 1951. A stochastic approximation method. *The annals of mathematical statistics*, 400–407.
- Shakespeare, W. 2014. *The complete works of William Shakespeare*. Race Point Publishing.
- Wang, J.; Charles, Z.; Xu, Z.; Joshi, G.; McMahan, H. B.; Al-Shedivat, M.; Andrew, G.; Avestimehr, S.; Daly, K.; Data, D.; et al. 2021a. A field guide to federated optimization. *arXiv preprint arXiv:2107.06917*.
- Wang, J.; Xu, Z.; Garrett, Z.; Charles, Z.; Liu, L.; and Joshi, G. 2021b. Local adaptivity in federated learning: Convergence and consistency. *arXiv preprint arXiv:2106.02305*.
- Wang, Y.; Lin, L.; and Chen, J. 2022. Communication-efficient adaptive federated learning. In *International Conference on Machine Learning*, 22802–22838. PMLR.
- Wu, X.; Huang, F.; Hu, Z.; and Huang, H. 2023. Faster adaptive federated learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 10379–10387.
- Yang, J.; Wu, Z.; Peng, K.; Okolo, P. N.; Zhang, W.; Zhao, H.; and Sun, J. 2021. Parameter selection of Gaussian kernel SVM based on local density of training set. *Inverse Problems in Science and Engineering*, 29(4): 536–548.
- Yu, F.; Rawat, A. S.; Menon, A.; and Kumar, S. 2020. Federated learning with only positive labels. In *International Conference on Machine Learning*, 10946–10956. PMLR.

- Zaheer, M.; Reddi, S.; Sachan, D.; Kale, S.; and Kumar, S. 2018. Adaptive methods for nonconvex optimization. *Advances in neural information processing systems*, 31.
- Zhang, X.; Yu, F. X.; Kumar, S.; and Chang, S.-F. 2017. Learning spread-out local feature descriptors. In *Proceedings of the IEEE international conference on computer vision*, 4595–4603.