





Personalized and Resilient Distributed Learning Through Opinion Dynamics

Luca Ballotta , *Member, IEEE*, Nicola Bastianello , *Member, IEEE*,
Riccardo M. G. Ferrari , *Senior Member, IEEE*, and Karl H. Johansson , *Fellow, IEEE*

Abstract—In this article, we address two practical challenges of distributed learning in multiagent network systems, namely, personalization and resilience. Personalization is the need of heterogeneous agents to learn local models tailored to their own data and tasks, while still generalizing well; on the other hand, the learning process must be resilient to cyberattacks or anomalous training data to avoid disruption. Motivated by a conceptual affinity between these two requirements, we devise a distributed learning algorithm that combines distributed gradient descent and the Friedkin–Johnsen model of opinion dynamics to fulfill both of them. We quantify its convergence speed and the neighborhood that contains the final learned models, which can be easily controlled by tuning the algorithm parameters to enforce a more personalized/resilient behavior. We numerically showcase the effectiveness of our algorithm on synthetic and real-world distributed learning tasks, where it achieves high global accuracy both for personalized models and with malicious agents compared to standard strategies.

Index Terms—Distributed learning, Friedkin–Johnsen (FJ) model, personalized learning, resilient distributed learning.

I. INTRODUCTION

DISTRIBUTED learning is the natural extension of machine learning to multiagent systems where devices collaboratively train models. Applications include multirobot systems [1], connected vehicles [2], smart grids [3], wind power forecasting [4], and the Internet of Things [5]. Distributed learning is

more robust than federated learning (FL) whereby the central aggregator may bias the global model toward some agents and is vulnerable to cyberattacks [6]. Yet, other problems are present. In this article, we focus on two key requirements in multiagent network systems: 1) personalization reflects the need of heterogeneous agents to prioritize their own local data and tasks and 2) resilience is the capability of agents to deter unknown intruders trying to disrupt the learning process.

A. Literature Review

Distributed learning is based on recent developments in distributed optimization [7], [8]. Among distributed optimization algorithms, distributed gradient descent (DGD), gradient tracking (GT), and dual methods (especially alternating direction method of multipliers (ADMM)) are widespread [8]. Algorithms based on GT and ADMM achieve exact convergence to an optimal solution using fixed parameters, while DGD converges inexactly unless a vanishing step size, which slows down convergence, is employed. These algorithms are leveraged to solve distributed learning problems [9], [10].

Most works assume that all agents are honest. However, this may not be the case as an adversary can deploy attacks through the wireless network [6], [11], [12]. Attacks can target either data, to infer or corrupt the agents' datasets, or models, to worsen their accuracy or bias. In this article, we focus on model attacks, which requires securing the training algorithm. A growing literature is devoted to resilient distributed algorithms; see [12] for a survey. A core component of distributed algorithms is consensus averaging, during which each agent exchanges models and possibly gradients with neighbors and updates its own local model based on received information. This step is vulnerable because attackers can transmit malicious data and pollute the agents' updates. While actively discriminating malicious agents to disregard their model updates is effective, it requires cybersecurity mechanisms not suited for low-power devices with limited hardware or fast algorithm execution. Hence, designing a passively resilient distributed algorithm with low computational footprint is of utmost importance. To this aim, resilient distributed algorithms replace the consensus step with a robust protocol. A common strategy uses trimmed means without the largest and smallest received values [13], [14], which requires dense interagent connectivity and a shared bound on the number of adversaries. Shang [15] replaces the average with the median, but still requires dense connectivity. A related

Received 1 October 2025; revised 27 November 2025; accepted 11 December 2025. Date of publication 29 December 2025; date of current version 19 March 2026. This work was supported in part by the EU Horizon-RIA Projects TWAIN under Grant 101122194 and ULTIMATE under Grant 101070162 and in part by the Swedish Research Council Distinguished Professor Grant 2017-01078 (Knut and Alice Wallenberg Foundation Wallenberg Scholar Grant). Recommended by Associate Editor J. A. Taylor. (Luca Ballotta and Nicola Bastianello contributed equally to this work.) (Corresponding author: Luca Ballotta.)

Luca Ballotta was with the Delft University of Technology, 2628 CD Delft, The Netherlands. He is now with the Department of Information Engineering, University of Padova, 35131 Padova, Italy (e-mail: luca.ballotta@unipd.it).

Nicola Bastianello and Karl H. Johansson are with the School of Electrical Engineering and Computer Science, and Digital Futures, KTH Royal Institute of Technology, 10044 Stockholm, Sweden (e-mail: nicolba@kth.se; kallej@kth.se).

Riccardo M. G. Ferrari is with the Delft Center for Systems and Control, Delft University of Technology, 2628 CD Delft, The Netherlands (e-mail: r.ferrari@tudelft.nl).

Digital Object Identifier 10.1109/TCNS.2025.3649108

algorithm in [16] uses the centerpoint, a generalization of the median to higher dimensions, but assumes that all agents have the same optimizer. Abbas et al. [17] leverage “trusted” agents, which must be secured from cyberattacks. A line of works uses “trust” obtained from physical information channels to filter out malicious messages, whose performance depends on the statistics of trust observations [18], [19], [20]. By contrast, in this work, we focus effort on a solution that does not require dense connectivity among agents, shared parameters, or extra resources and provides guarantees under standard assumptions.

Another challenge of distributed learning is the statistical heterogeneity of the cooperating agents [21]. The agents’ local data are often drawn from different distributions, which reflect their peculiar perspectives (e.g., different information sources, sensors, or plant specifications). As a consequence, each agent may wish to learn a model that is highly accurate on its own local distribution, possibly trading this need with lower accuracy on other agents’ distributions. This concept is referred to as *personalization*. Personalized learning has been widely explored in FL, wherein it relies on the central coordinator [21], [22], [23]. Also, it has been sparsely addressed in distributed learning [24], [25], but these approaches transform the original problem to a computationally burdensome bilevel optimization, especially challenging with compute-intensive models such as in deep learning. On the contrary, here, we are interested in embedding personalization into the algorithm itself rather than the problem formulation to avoid an excessive computational burden.

The synergy of personalization with other objectives has been explored in FL. Kundu et al. [26] integrate personalization with robustness to outliers, that is, agents with a significantly different local distribution. Han et al. [27] discuss how personalization and generalization can both be achieved via tailored algorithm design. Bietti et al. [28] explore the impact of personalized learning for privacy preservation (i.e., robustness against attacks on data). Li et al. [29] discuss the potential of personalization to achieve robustness to model attacks and fairness, which ensures uniform performance of the trained models across agents. However, all these works do not apply to a fully distributed setup and rely on computationally intensive reformulations of the problem to achieve personalization.

B. Contribution

We propose a novel distributed learning algorithm to achieve personalization and resilience. The idea is to purposely bias, in a controlled way, each agent toward its own local cost. In collaborative but heterogeneous settings, this strategy improves local accuracy achieved by the agents, enhancing personalization; with malicious agents, it makes training resilient by reducing influence of those on the learned models. In light of this connection, we develop one algorithm to tackle both problems. We draw inspiration from the resilient consensus approach in [30] and combine the Friedkin–Johnsen (FJ) opinion dynamics model, originally conceived to capture disagreement, with DGD. The key ingredient of our algorithm is a scalar parameter $\lambda \in [0, 1]$, modeling opinion “stubbornness” in the original FJ model, that allows agents to smoothly transition from collaborative training

($\lambda = 0$), which targets high global accuracy irrespectively of heterogeneous agents or attacks, to local training ($\lambda = 1$), which achieves high local accuracy but cannot generalize well. We characterize the geometric convergence rate of our algorithm and the distance between learned models and optimum of the nominal distributed learning problem as functions of design parameters. We conduct an extensive numerical campaign with synthetic and real-world classifications problems. Our algorithm achieves superior personalization compared to DGD and enhances resilience with over 10% of malicious agents scattered across a sparse network, improving accuracy of DGD by up to 78%. To the best of our knowledge, this is the first approach for personalized distributed learning that does not require a computationally intensive reformulation of the original problem. Furthermore, it does not require dense communication or extra resources such as secured agents or trust information, and enjoys formal guarantees under standard assumptions on the cost functions. In the context of resilience, our algorithm may be used as a first defense mechanism to let agents partially collaborate from the start till the adversaries are detected by more sophisticated but slower strategies, such as standard network security or the trust-based algorithms in [18], [19], and [20].

Organization: The rest of this article is organized as follows. We introduce the distributed learning setup in Section II, describing in detail the concepts of personalization (see Section II-A) and resilience (see Section II-B) and discussing their affinity (see Section II-C). In Section III, we develop our distributed learning algorithm by combining DGD and the FJ model. In Section IV, we analyze its fixed point and convergence speed in both cases with fully collaborative agents and with malicious agents sending bounded values. In Section V, we test our algorithm on synthetic and real-world classification tasks, where it outperforms DGD and local training in achieving personalization and resilience. We discuss current limitations and potential directions of improvements in Section VI. Finally, Section VII concludes this article.

II. DISTRIBUTED LEARNING SETUP

Consider N agents that exchange information over a wireless network modeled as an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} \doteq \{1, \dots, N\}$ and $(i, j) \in \mathcal{E}$ means that there exists a direct communication link between the two agents labeled i and j . Agent i is equipped with a decision variable $x_i \in \mathbb{R}^n$ and cost function $f_i : \mathbb{R}^n \rightarrow \mathbb{R}$. Notation $\{x_i\}_{i \in \mathcal{V}}$ denotes the set of all agents’ variables. Each agent directly manipulates only its variable x_i but can both access its cost function f_i and exchange information with neighbors. In distributed optimization, the agents cooperatively solve the following problem:

$$\begin{aligned} & \underset{\{x_i\}_{i \in \mathcal{V}}}{\text{minimize}} && \sum_{i \in \mathcal{V}} f_i(x_i) \\ & \text{subject to} && x_i = x_j \quad \forall i, j \in \mathcal{V}. \end{aligned} \quad (1)$$

We turn to the specific case of distributed learning, which is the focus of this article. We first recap the underlying probabilistic optimization framework. Here, agent i samples data from a local probability distribution \mathcal{P}_i that reflects the agent’s perspective

on, or accessed portion of, the global phenomenon being observed by the network. Each agent i aims to learn a model, parameterized by x_i , able to describe (or “explain”) both its local distribution \mathcal{P}_i and those of the other agents $\{\mathcal{P}_j\}_{j \neq i}$.¹ Hence, the agents cooperate to solve the risk-minimization problem

$$\begin{aligned} & \underset{\{x_i\}_{i \in \mathcal{V}}}{\text{minimize}} && \sum_{i \in \mathcal{V}} \mathbb{E}_{d \sim \mathcal{P}_i} [\ell(x_i; d)] \\ & \text{subject to} && x_i = x_j \quad \forall i, j \in \mathcal{V}. \end{aligned} \quad (2)$$

In problem (2), the loss function $\ell : \mathbb{R}^n \times \mathbb{R}^p \rightarrow \mathbb{R}$ quantifies how accurately the model x_i describes a data point d sampled from \mathcal{P}_i . A classic example is image classification, where the data are image–class pairs $d = (a, b)$ and the model matches an image a to a semantic class b such as an alphanumeric character or an object. In this case, x_i minimizes $\ell(x_i; (a, b))$ if the model outputs b given input a . However, problem (2) can be solved only if all distributions $\{\mathcal{P}_i\}_{i \in \mathcal{V}}$ are known, which is not the case in practice. Thus, the agents address a deterministic approximation of (1) called empirical risk minimization and characterized as follows. Agent i owns a private dataset $\mathcal{D}_i = \{d_{i,j}\}_{j \in \mathcal{I}_i}$, where each data point is assumed sampled from the local distribution, namely, $d_{i,j} \sim \mathcal{P}_i$. This allows the agent i to approximate its risk function $\mathbb{E}_{d \sim \mathcal{P}_i} [\ell(x_i; d)]$ with the empirical risk associated with \mathcal{D}_i and to seek a learning-based model parameterized by x_i that suitably describes the data in \mathcal{D}_i .² The local cost f_i is consequently instantiated as

$$f_i(x_i) = \frac{1}{|\mathcal{I}_i|} \sum_{j \in \mathcal{I}_i} \ell(x_i; d_{i,j}) + \gamma \rho(x_i). \quad (3)$$

The regularization function ρ in (3) is designed to prevent overfitting the training dataset \mathcal{D}_i and usually chosen as a norm with weight $\gamma > 0$. In words, minimization of the (regularized) empirical risk (3) aims to make agent i learn a parameter x_i that describes \mathcal{D}_i and also generalizes to the distribution \mathcal{P}_i , namely, it explains samples $d \sim \mathcal{P}_i$ not present in \mathcal{D}_i .

Problem (1) can be solved by training one model on the aggregated dataset $\{\mathcal{D}_i\}_{i \in \mathcal{V}}$. However, privacy concerns and communication constraints prevent the agents from sharing data, ruling out this option. At the same time, if each agent trains a model only on its own dataset, it achieves poor generalization. These concurrent issues make distributed optimization algorithms suited to the distributed learning specialization of (1) where the local cost functions f_i are the regularized losses defined in (3). The agents solve (1) without exchanging the data \mathcal{D}_i but only the local parameters x_i or the gradients $\nabla f_i(x_i)$.

To derive formal guarantees of convergence, we require standard assumptions on convexity and smoothness of the local losses [9], [31], [32]. In the following, $\|\cdot\|$ denotes the 2-norm.

¹ This is effectively possible only if the distributions $\{\mathcal{P}_i\}_{i \in \mathcal{V}}$ are “similar” enough.

² Since we assume that all agents learn the same model (e.g., a linear model or a neural network model) and only differ in the model parameters x_i , in the following, we refer to x_i just as “parameter(s)” for the sake of simplicity.

Definition 1: Function f is μ -strongly convex if there exists $\mu > 0$ such that, for all x and y , it holds

$$f(y) \geq f(x) + \nabla f(x)^\top (y - x) + \frac{\mu}{2} \|y - x\|^2. \quad (4)$$

Definition 2: Function f is L -smooth if ∇f is globally Lipschitz with Lipschitz constant $L > 0$.

Assumption 1: The local loss f_i in (3) is μ -strongly convex and L -smooth for all $i \in \mathcal{V}$.

Connection with FL: The FL setup is the special case where each agent is connected to only one external agent, called the aggregator. This receives local models from agents, combines them into a *global model*, and sends the latter back to the agents, which refine it with local data, and the process repeats itself. While the communication topology of FL is a simple star network centered at the aggregator, this corresponds to a complete graph for (1) whereby the aggregator implements all communication links.

A. What Is Personalization in Distributed Learning?

The distributed learning problem previously introduced treats all agents as if they are homogeneous, i.e., sample data from the same underlying distribution. Indeed, problem (1) forces all agents to learn a unique model under the implicit assumption that this will generalize to the distribution \mathcal{P}_i of each agent i . However, in practical cases, the agents both have heterogeneous [non independent identically distributed (i.i.d.)] data and may want to prioritize their own local distributions. For instance, if industrial partners collaboratively train models for data-driven predictive maintenance, each partner may want a model tailored to its own machines. This necessity calls for *personalization* of the parameter x_i such that the model learned by agent i does not just generalize to all \mathcal{P}_j ’s—due to the collaboration with other agents—but provides high accuracy specifically on the distribution \mathcal{P}_i . Of course, there is a tradeoff between personalization and generalization that depends on interagent heterogeneity.

1) Personalized FL: In FL, the global model broadcast by the aggregator can serve as a reference for personalization. If the global model differs significantly from a local model, that agent may prioritize its local loss. This can be accommodated as proposed in [22]. Denoting the global parameters by x , problem (1) is formally redefined as

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \quad \sum_{i \in \mathcal{V}} F_i(x) \quad (5a)$$

where the local loss of agent i is replaced by

$$F_i(x) = \min_y \left\{ f_i(y) + \frac{\pi_i}{2} \|y - x\|_2^2 \right\} \quad (5b)$$

and the parameter π_i quantifies the importance agent i gives to the global model. The reformulation (5) allows agent i to personalize its parameter x_i for its own dataset \mathcal{D}_i while controlling how different its local FL model and the global one are through the coefficient π_i .

2) Personalized Distributed Learning: The personalization-oriented reformulation (5) relies on the hierarchical structure of FL, whereby the aggregator computes and broadcasts the

global model x to all agents. However, an analogous formulation cannot be constructed for the fully distributed setup we consider, where agents access only information sent by a few neighbors.

Problem 1 (Personalized distributed learning): Given the distributed learning problem (1), how to increase the local accuracy of each agent while maintaining high global accuracy? Formally, how to optimally trade between minimization of the local loss $f_i(x_i)$ of each agent i and the total loss $\sum_{j \in \mathcal{V}} f_j(x_j)$?

B. What Is Resilience in Distributed Learning?

The distributed learning setup (1) assumes that all agents are “honest” and correctly run the distributed algorithm. However, wireless communication allows external malicious agents to interfere and degrade the learned models. This can be done by attacking (or acting as) an agent and spread misleading information during training, e.g., by sharing noisy parameters x_i to lower accuracy. Even without malicious attackers, agents with many outliers or very different data distributions may degrade the distributed training of local models.

Standard approaches for privacy preservation include computationally expensive security mechanisms, e.g., encryption, or ad hoc solutions such as perturbing the transmitted parameters. These approaches do little against disrupting model attacks. Proactive cybersecurity mechanisms leverage the communication protocol to detect adversaries from transmitted packets, but this can take quite some time, neglects the content of messages, and is not robust to normal agents with outliers in the training data. Therefore, there is a need to make distributed training not just secure but also *resilient*, that is, insensitive to disruptions introduced by adversaries or agents with poor training data without degrading the models learned by normal agents.

Problem 2 (Resilient distributed learning): Given the distributed learning problem (1), how to make the honest agents learn globally accurate parameters x_i with unknown malicious agents? Formally, let $\mathcal{H} \subset \mathcal{V}$ denote honest agents; then, we are interested in solving the problem

$$\begin{aligned} & \underset{\{x_i\}_{i \in \mathcal{H}}}{\text{minimize}} && \sum_{i \in \mathcal{H}} f_i(x_i) \\ & \text{subject to} && x_i = x_j \quad \forall i, j \in \mathcal{H}. \end{aligned} \quad (6)$$

C. Affinity Between Personalization and Resilience

As motivated in the previous sections, personalization and resilience respond to two different needs. On the one hand, an agent personalizes its own model to achieve high local accuracy, possibly at the cost of lower accuracy on other agents’ distributions. On the other hand, model attacks urge the honest agents to implement resilient mechanisms that prevent disruption of the cooperative training.

Nonetheless, personalization and resilience share similarities. They both reduce the emphasis on collaboration and focus on each individual agent to either enhance local accuracy (personalization) or reduce uncertainty in the source of received information (resilience). In both scenarios, each agent regards all other agents as potentially harmful to its own learning goal: on the one hand, the other agents may have different data

distributions, introducing undesired variance to the model; on the other hand, the identity of malicious agents is unknown, forcing honest agents to be suspicious of their neighbors.

This discussion motivates the quest for a common mechanism that can accommodate both objectives, which we propose and evaluate in the next sections.

III. ALGORITHM DESIGN

In this section, we devise our distributed learning algorithm tailored to Problems 1 and 2. We consider gradient-based algorithms consistently with a large portion of the machine learning literature. Motivated by the discussion in Section II-C, our design aims to optimally trade local accuracy (of each agent) for global accuracy (of all agents) to achieve personalized and resilient models. We first compare GT and DGD in Section III-A. DGD achieves higher local accuracy than GT and performs better with noisy training. However, this behavior of DGD naturally emerges as a by-product of its intrinsic suboptimality and cannot be tuned. Therefore, in Section III-B, we combine DGD with the FJ model, which captures stubbornness in opinion dynamics. This model introduces an additional scalar parameter, which we use to tune the learned models toward either local or global accuracy and account for variable requirements.

A. DGD: Advantages of Inaccuracy

Consider the standard iterate of DGD

$$\mathbf{x}_{k+1} = W \mathbf{x}_k - \alpha \nabla f(\mathbf{x}_k) \quad (\text{DGD})$$

where the bold symbol $\mathbf{x} \in \mathbb{R}^{Nn}$ stacks the parameters of all agents and $f(\mathbf{x}_k)$ stacks all local losses, each evaluated at the corresponding agent’s parameter. The matrix W is doubly stochastic and can be easily built in practice, e.g., with Metropolis weights.

It is well known that (DGD) does not solve problem (1) but achieves a neighborhood of the globally optimal model. The size of this neighborhood depends on the heterogeneity of the local losses and is larger with more varied training data. The next result quantifies the distance between the learned parameters and the optimum of (1).

Proposition 1: Let $x^* = \operatorname{argmin}_{x \in \mathbb{R}^n} \sum_{i=1}^N f_i(x)$ and \bar{x} be the fixed point of (DGD). Then, the following bound holds:

$$\|\bar{x} - \mathbf{1} \otimes x^*\|^2 \leq \mathcal{O}(c + \alpha D) \quad (7)$$

where c is a constant offset and

$$D = \sqrt{2L \sum_{i=1}^N (f_i(x^*) - f_i(x_i^*))} \quad (8)$$

with $x_i^* \doteq \operatorname{argmin}_{x \in \mathbb{R}^n} f_i(x)$.

Proof: By [32, Thm. 7], we know that (DGD) converges to the point \bar{x} . In addition, we have that

$$\|\bar{x} - \mathbf{1} \otimes x^*\|^2 \leq \mathcal{O}(\alpha D')$$

TABLE I

ACCURACY WITH GRADIENT TRACKING (ED) AND DISTRIBUTED GRADIENT DESCENT (DGD AND ATC)

	Dataset	Min	Mean \pm Std	Max
ED	global	0.934	0.934 \pm 0.000	0.934
	local	0.640	0.934 \pm 0.110	1.000
DGD	global	0.848	0.876 \pm 0.026	0.918
	local	0.840	0.974 \pm 0.054	1.000
ATC	global	0.868	0.891 \pm 0.021	0.928
	local	0.840	0.966 \pm 0.063	1.000

TABLE II

ACCURACY WITH GRADIENT TRACKING (ED) AND DISTRIBUTED GRADIENT DESCENT (DGD AND ATC) WITH NOISY UPDATES

	Dataset	Min	Mean \pm Std	Max
ED	global	0.502	0.506 \pm 0.003	0.512
	local	0.020	0.502 \pm 0.274	1.000
DGD	global	0.604	0.674 \pm 0.052	0.758
	local	0.400	0.686 \pm 0.194	1.000
ATC	global	0.520	0.627 \pm 0.051	0.704
	local	0.000	0.676 \pm 0.320	1.000

where

$$D' = \sqrt{2L \sum_{i=1}^N (f_i(0) - f_i(x_i^*))}.$$

Summing and subtracting $f_i(x_i^*)$, rearranging, and using $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ if $a, b \geq 0$, we get

$$D' \leq c + D$$

with $c \doteq \sqrt{2L \sum_{i=1}^N (f_i(0) - f_i(x_i^*))}$. \square

The constant c in the bound (7) does not depend on the cost functions. On the other hand, if the agents' losses are very heterogeneous, the global minimizer x^* and each local minimizer x_i^* differ widely, causing D to increase.

This behavior of (DGD) is undesired in the nominal case where one aims to solve (1) and achieve global optimality. It turns useful to improve on local accuracy of each agent instead, because the updates (DGD) converge near both the global optimum x^* and the minimizers of the local losses f_i .

Numerical evaluation: To concretely assess potential advantages of DGD methods, we consider a binary classification task on synthetic data.³ We compare (DGD) with exact diffusion (ED), a GT algorithm proposed in [33] that solves (1), and adapt-then-combine (ATC), whose k th iterate is

$$\mathbf{x}_{k+1} = W(\mathbf{x}_k - \alpha \nabla f(\mathbf{x}_k)). \quad (\text{ATC})$$

The classification accuracy of model x on dataset \mathcal{D} is defined as

$$\text{accuracy} \doteq 1 - \frac{|\{(a, b) \in \mathcal{D} : \hat{b}_x(a) \neq b\}|}{|\mathcal{D}|} \quad (9)$$

where $\hat{b}_x(a)$ is the label output by model x given features a . Table I reports accuracy statistics across the N agents. For agent i , local accuracy is computed on the test dataset $\mathcal{D}_i^{\text{test}}$ and global accuracy on the combined dataset $\cup_{i \in \mathcal{V}} \mathcal{D}_i^{\text{test}}$. The values show that (DGD) outperforms ED on local accuracy, because the latter converges to the minimizer of the global loss and cannot accommodate for heterogeneous datasets. Algorithm (ATC) places between (DGD) and ED in terms of average accuracy.

Notably, GT is sensitive to noise. In distributed learning, this clashes against noisy gradients (e.g., in stochastic gradient

descent), wireless channel erasure, or messages sent by malicious agents. Table II reports accuracy scores computed as in Table I but with algorithm updates additively perturbed by i.i.d. Gaussian noise drawn from $\mathcal{N}(0, 2)$. The accuracy of ED drops significantly, whereas (DGD) and (ATC) are more robust as shown by the smaller decreases in average accuracy.

We conclude that DGD methods improve over exact GT both on local accuracy and with noisy updates, suggesting opportunities for enhancing personalization and resilience in light of the discussion in Section II-C. However, neither (ATC) nor (DGD) provide ways to compare local models with a “global” consensus-based model, as opposed to GT, and offer no easy way to tune local or global accuracy reached by each agent. We next address this issue through an opinion dynamics model that structurally accounts for heterogeneous agents.

B. FJ-DGD: HARNESSING STUBBORNESS FOR LOCAL ACCURACY

We draw inspiration from the FJ model to tune between global and local accuracy achieved by learned models and enhance personalization and resilience. The original FJ model tracks dynamic evolution of opinions \mathbf{x} starting from an initial condition \mathbf{x}_0 , assuming that the agents are partially stubborn and retain their initial opinions throughout [34], [35]

$$\mathbf{x}_{k+1} = (I - \Lambda)W\mathbf{x}_k + \Lambda\mathbf{x}_0. \quad (10)$$

Matrix Λ is diagonal and its i th diagonal element $\lambda_i \in [0, 1]$ represents the stubbornness of agent i . The case with $\lambda_i \equiv 0$ reduces (10) to the consensus algorithm, whereas arbitrary stubbornness parameters $\lambda_i > 0$ prevent a consensus apart from trivial cases. In words, agents behaving according to (10) embed the others' opinions (through the consensus term $W\mathbf{x}_k$) but always mediate with their own initial opinion. The closer λ_i is to 1, the less the final opinion x_i is affected by the others.

This offers a simple workaround to improve local accuracy by letting each agent retain local information at all times. Intuitively, the FJ model can purposely bias the local model x_i of agent i toward the optimal one for local data \mathcal{D}_i , whereby the closer λ_i to 1, the stronger such as bias. A straightforward implementation of this intuition yields the following algorithm:

$$\mathbf{x}_{k+1} = \Lambda\mathbf{x}^* + (I - \Lambda)(W\mathbf{x}_k + \alpha \nabla f(\mathbf{x}_k)) \quad (11)$$

³ Datasets, losses, and models are described in Section V, Example 1.

where $\mathbf{x}^* \doteq [(x_1^*)^\top, \dots, (x_N^*)^\top]^\top$ and $x_i^* = \operatorname{argmin}_x f_i(x)$ is the local optimizer of agent i . This strategy requires the agents to spend extra time to precompute the local optimal parameters x_i^* before collaborative training, with the added complication of synchronizing the start. To avoid this issue, we let the agents concurrently update the local models during collaborative training. This results in the following algorithm, whereby agent i tracks its local minimizer x_i^* with variable y_i

$$\begin{aligned} \mathbf{y}_{k+1} &= \mathbf{y}_k - \alpha \nabla f(\mathbf{y}_k), & \mathbf{y}_0 &= \mathbf{x}_0 \\ \mathbf{x}_{k+1} &= \Lambda \mathbf{y}_{k+1} + (I - \Lambda)(W \mathbf{x}_k - \alpha \nabla f(\mathbf{x}_k)). \end{aligned} \quad (\text{FJ-DGD-1})$$

Since \mathbf{y}_k converges to \mathbf{x}^* , both (11) and (FJ-DGD-1) have the same fixed point. Algorithm (FJ-DGD-1) combines the consensus-based aggregation of DGD and the stubbornness of the FJ model to learn models x_i , which are influenced by both behaviors throughout the training. Alternatively, the stubborn and consensus-based models may be separately computed and combined at a second stage. This yields the following algorithm, whereby each agent i uses two extra variables y_i and z_i

$$\begin{aligned} \mathbf{y}_{k+1} &= \mathbf{y}_k - \alpha \nabla f(\mathbf{y}_k), & \mathbf{y}_0 &= \mathbf{x}_0 \\ \mathbf{z}_{k+1} &= W \mathbf{z}_k - \alpha \nabla f(\mathbf{z}_k), & \mathbf{z}_0 &= \mathbf{x}_0 \\ \mathbf{x}_{k+1} &= \Lambda \mathbf{y}_{k+1} + (I - \Lambda) \mathbf{z}_{k+1}. \end{aligned} \quad (\text{FJ-DGD-2})$$

In Section IV, we analyze convergence of (FJ-DGD-2). This choice is due to the higher flexibility of (FJ-DGD-2), whose \mathbf{y} - and \mathbf{z} -updates are independent. A similar analysis can be carried out for (FJ-DGD-1).

C. FJ-DGD WITH CORRUPTED UPDATES

The previous section considers a nominal collaborative scenario where all agents truthfully obey the designed algorithm. We now extend this case to the scenario where some agents do not exactly follow the designed update rule. In particular, this model can capture malicious agents that transmit noisy or deceiving models to disrupt the collaborative training carried out by the other agents. In light of the affinity between personalization and resilience discussed in Section II-C, FJ-DGD can be tuned to achieve good global accuracy in this case. Intuitively, if the agents reduce collaboration in a controlled manner, they mitigate the effect of malicious or low-quality information they receive. In the following, we focus on (FJ-DGD-2) for the sake of conciseness.

Let \mathcal{H} and \mathcal{M} be the subsets of honest and malicious agents, respectively, such that $\mathcal{H} \cap \mathcal{M} = \emptyset$ and $\mathcal{V} = \mathcal{H} \cup \mathcal{M}$. We assume that malicious agents participating in the learning process freely choose \mathbf{z} in the updates of (FJ-DGD-2) to degrade performance of honest agents' local models. Define the vector \mathbf{e}_k such that

$$[\mathbf{e}_k]_i = \begin{cases} 0, & \text{if } i \in \mathcal{H} \\ e_{i,k}, & \text{if } i \in \mathcal{M}. \end{cases} \quad (12)$$

Error vector \mathbf{e}_k represents the deviation from nominal updates caused by malicious behavior. From the perspective of agent $i \in \mathcal{H}$ (which ignores whether $e_{j,k} = 0$ or $e_{j,k} \neq 0$ for any neighbor

j), the updates read

$$\begin{aligned} z_{i,k+1} &= w_{ii} z_{i,k} + \sum_{j \in \mathcal{N}_i \cap \mathcal{H}} w_{ij} z_{j,k} \\ &+ \sum_{j \in \mathcal{N}_i \cap \mathcal{M}} w_{ij} (z_{j,k} + e_{j,k}) - \alpha \nabla f_i(z_{i,k}) \end{aligned} \quad (13)$$

and replacing the \mathbf{z} -update of (FJ-DGD-2) yields

$$\begin{aligned} \mathbf{y}_{k+1} &= \mathbf{y}_k - \alpha \nabla f(\mathbf{y}_k), & \mathbf{y}_0 &= \mathbf{x}_0 \\ \mathbf{z}_{k+1} &= W \mathbf{z}_k - \alpha \nabla f(\mathbf{z}_k) + W \mathbf{e}_k, & \mathbf{z}_0 &= \mathbf{x}_0 \\ \mathbf{x}_{k+1} &= \Lambda \mathbf{y}_{k+1} + (I - \Lambda) \mathbf{z}_{k+1}. \end{aligned} \quad (\text{FJ-DGD-N})$$

This modified version behaves worse because of the unmodeled disturbances in \mathbf{e}_k , which may not relate to the algorithm itself. The following section is dedicated to the performance analysis of (FJ-DGD-2) and of its corrupted version (FJ-DGD-N).

IV. CONVERGENCE ANALYSIS

We propose two main results. The first one quantifies the final models learned by (FJ-DGD-2) within a fully collaborative setup, along with the speed of convergence. This gives an indication of how long training is required. Recall that we use the notation $\mathbf{x}^* \doteq [(x_1^*)^\top, \dots, (x_N^*)^\top]^\top$, with $x_i^* = \operatorname{argmin}_x f_i(x)$ the local optimizer of agent i .

Theorem 1: Let $\bar{\mathbf{x}}$ be the fixed point of (DGD). Algorithm (FJ-DGD-2) converges linearly to $\Lambda \mathbf{x}^* + (I - \Lambda) \bar{\mathbf{x}}$ at a rate of

$$\zeta = \max \{ |1 - \alpha\mu|, |1 - \alpha L|, |\lambda_{\min}(W) - \alpha L| \} \quad (14)$$

where $\lambda_{\min}(W)$ is the smallest eigenvalue of W .

Proof: The result is a consequence of the following facts.

- 1) $\mathbf{y}_{k+1} = \mathbf{y}_k - \alpha \nabla f(\mathbf{y}_k)$ converges linearly to \mathbf{x}^* , at a rate $\zeta' = \max \{ |1 - \alpha\mu|, |1 - \alpha L| \}$.
- 2) (DGD) converges linearly to $\bar{\mathbf{x}}$, at a rate $\zeta'' = \max \{ |1 - \alpha\mu|, |\lambda_{\min}(W) - \alpha L| \}$.

Combining these items yields the rate in (14). \square

Theorem 1 shows that the algorithm (FJ-DGD-2) enjoys a geometric convergence rate. The final models are convex combinations between the global optimum \mathbf{x}^* and the models $\bar{\mathbf{x}}_i$ learned with (DGD). This allows a designer to easily tune the behavior of (FJ-DGD-2) based on how much local accuracy is preferred over global accuracy. In Section V, we show how to use this to enforce personalization or help resilience. We also remark that the choice of Λ here does not affect the convergence rate of the algorithm, but only its fixed point. Indeed, the convergence rate is fully characterized by the separate convergence rates of the \mathbf{y} - and \mathbf{z} -updates, which do not depend on Λ . The fixed point though does depend on it, as it is characterized as the convex combination $\Lambda \mathbf{x}^* + (I - \Lambda) \bar{\mathbf{x}}$.

The next result evaluates the convergence of (FJ-DGD-N) with noisy terms $e_{i,k}$ injected by malicious agents. We assume that these cannot transmit arbitrary messages and bound their magnitude. This assumption is not restrictive because, if honest agents receive inappropriately large model weights, they can easily detect malicious agents and exclude them from training.

Theorem 2: Let $\|e_k\| \leq \tau$. The algorithm (FJ-DGD-N) converges to a neighborhood of $\hat{\mathbf{x}} = \Lambda \mathbf{x}^* + (I - \Lambda)\bar{\mathbf{x}}$ characterized by

$$\|\mathbf{x}_k - \hat{\mathbf{x}}\| \leq \zeta^k \|\mathbf{x}_0 - \hat{\mathbf{x}}\| + (1 - \min_i \lambda_i) \tau \sum_{h=0}^{k-1} \zeta^{k-h-1} \quad (15)$$

at a rate of ζ , defined in (14).

Proof: The goal is to provide a bound to the distance $\|\mathbf{x}_k - \hat{\mathbf{x}}\|$, with $\Lambda \mathbf{x}^* + (I - \Lambda)\bar{\mathbf{x}}$ being the fixed point of (FJ-DGD-2); see Theorem 1. Using the characterization of the update (FJ-DGD-N) and the triangle inequality yields

$$\begin{aligned} \|\mathbf{x}_{k+1} - \hat{\mathbf{x}}\| &\leq \|\Lambda \mathbf{y}_{k+1} + (I - \Lambda)(W \mathbf{z}_k - \alpha \nabla f(\mathbf{z}_k)) - \hat{\mathbf{x}}\| \\ &\quad + \|(I - \Lambda)W e_k\| \\ &\stackrel{(i)}{\leq} \zeta \|\mathbf{x}_k - \hat{\mathbf{x}}\| + \|(I - \Lambda)W e_k\| \end{aligned}$$

where (i) holds by contractiveness of (FJ-DGD-2). Finally, by submultiplicativity of the norm, we have $\|(I - \Lambda)W e_k\| \leq \|(I - \Lambda)\| \|W\| \|e_k\|$ and using $\|I - \Lambda\| = 1 - \min_i \lambda_i$, $\|W\| = 1$, and $\|e_k\| \leq \tau$, we have

$$\|\mathbf{x}_{k+1} - \hat{\mathbf{x}}\| \leq \zeta \|\mathbf{x}_k - \hat{\mathbf{x}}\| + (1 - \min_i \lambda_i) \tau \quad (16)$$

and iterating (16) over time yields the result. \square

Theorem 2 proves that (FJ-DGD-N) converges inexactly, due to the noise injected by the malicious agents, to a neighborhood whose radius is upper bounded by $(1 - \min_i \lambda_i) \tau / (1 - \zeta)$. Honest agents can regulate their λ_i s to attenuate the impact of malicious agents, with the extreme choice of $\lambda_i = 1$ leading to all agents isolating and performing local training only.

V. NUMERICAL EXPERIMENTS

We now test our algorithm on distributed learning tasks where we wish to impose personalization and resilience. While these two needs share similarities as previously discussed, they correspond to two different scenarios—fully collaborative heterogeneous agents versus some noncollaborative or adversarial agents whose identity is unknown. Therefore, we perform two different numerical tests. Also, this allows us to better identify the effects of our algorithm for personalization and resilience.

A. Evaluating Personalization

We use the following distributed learning task to conveniently model heterogeneous agents and assess personalization ability.

Example 1 (Binary classification): Consider a classification problem with local loss function

$$f_i(x_i) = \frac{1}{|\mathcal{I}_i|} \sum_{j \in \mathcal{I}_i} \log(1 + \exp(-b_{i,j} a_{i,j}^\top x_i)) + \gamma \|x_i\|^2 \quad (17)$$

where $d_{i,j} = (a_{i,j}, b_{i,j})$, with $a_{i,j} \in \mathbb{R}^n$ being the feature vector and $b_{i,j} \in \{-1, 1\}$ the classification label. Problem (1) with the costs (17) is convex. We set the size of model parameters $n = 15$. Each agent i has access to $|\mathcal{D}_i| = 450$ training samples and $|\mathcal{D}_i^{\text{test}}| = 50$ test samples such that $\mathcal{D}_i^{\text{test}} \cap$

$\mathcal{D}_i = \emptyset$. Following [36], we simulate heterogeneous agents by generating data points $(a_{i,j}, b_{i,j})$ according to $b_{i,j} = \text{argmax}(\text{softmax}(W_i a_{i,j} + c_i))$ where each element of W_i and c_i is drawn from $\mathcal{N}(\mu_i, 1)$ with $\mu_i \sim \mathcal{N}(0, 1)$, and $a_{i,j} \sim \mathcal{N}(\nu_i, \Sigma_i)$ where each element of ν_i is drawn from $\mathcal{N}(B_i, 1)$ with $B_i \sim \mathcal{N}(0, 1)$ and Σ_i is diagonal with $[\Sigma_i]_{jj} = j^{-1.2}$. Agent i classifies the feature vector $a_{i,j}$ as

$$\hat{b}_{x_i}(a_{i,j}) = \begin{cases} 1, & \text{if } (1 + \exp(-a_{i,j}^\top x_i))^{-1} > 0.5 \\ -1, & \text{otherwise.} \end{cases} \quad (18)$$

We set $N = 10$ agents connected through the ring topology. \square

We train with the setup in Example 1 and $\gamma = 0.01$. We choose $\Lambda = \lambda I$ to more easily interpret the experiment through the scalar parameter $\lambda \in [0, 1]$. The results are reported in Figs. 1 and 2, where the label ‘‘Local GD’’ refers to each agent independently training on its own dataset without any collaboration. Note that Local GD is the instantiation of (FJ-DGD-1) and (FJ-DGD-2) with $\lambda = 1$, whereas (DGD) corresponds to setting $\lambda = 0$.

The test accuracy is very similar to the one obtained from training data, suggesting that the learned models generalize well. Compared to (DGD), both FJ-inspired modifications exhibit superior personalization performance. In particular, Fig. 1(a) and (c) reveals a graceful improvement in personalization, along with a controlled degradation of global accuracy in Fig. 1(b) and (d), by suitably tuning the parameter λ . As this increases, the agents turn more ‘‘stubborn’’ and consistently bias the local models toward their respective local optimizers, improving local accuracy. Notably, both (FJ-DGD-1) and (FJ-DGD-2) with $\lambda = 0.75$ slightly outperform ‘‘Local GD’’ on the local test accuracy (mean accuracy 0.93 versus 0.926), suggesting that collaboration with other agents is helpful even for personalization purposes, possibly because agents with similar distributions benefit from each other. On the other hand, the global accuracy is much higher for (FJ-DGD-1)–(FJ-DGD-2) than ‘‘Local GD,’’ which highlights the need for (partial) collaboration for generalizing to other agents’ distributions.

Fig. 2 spotlights the lowest local accuracy across all agents with the four compared strategies on the test sets. Note that (DGD) and ‘‘Local GD’’ feature straight lines because they do not depend on λ . DGD performs the worst, as expected since it values the contribution of all agents equal. Less intuitive is the fairly poor performance of ‘‘Local GD’’ that is even worse than the least personalized FJ-based algorithm. This also hints at partial collaboration as useful to achieve consistent effective personalization for all agents.

Overall, the best performance is provided by (FJ-DGD-2), which scores just slightly higher than (FJ-DGD-1) in terms of local accuracy—and hence provides similar personalization—but feature significantly higher (6%–16% higher mean) global accuracy. This suggests that a two-stage cascade comprising computation of local and global models and subsequent convex combination of the two is an effective strategy to personalize the local models without excessively compromising global performance. On the other hand, (FJ-DGD-1) personalizes the models as well and requires only two-thirds of the memory used

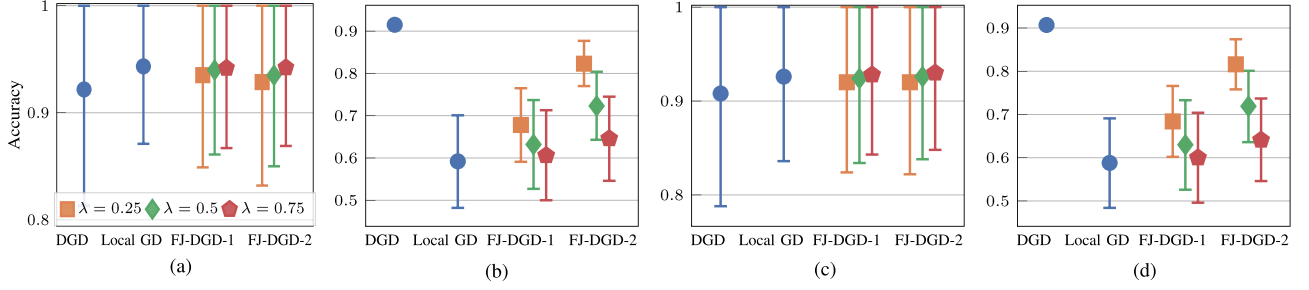


Fig. 1. Accuracy with DGD and its FJ-based variants for the task in Example 1. Marks show the mean and bars one standard deviation intervals across all agents. (a) Local training accuracy. (b) Global training accuracy. (c) Local test accuracy. (d) Global test accuracy.




Fig. 2. Minimal local test accuracy achieved for the task in Example 1.

by (FJ-DGD-2), which is especially useful to train large models as compared to the storage capacity of the agents.

We have demonstrated the superiority of FJ-inspired algorithms for personalization. We now explore how the algorithms under study behave as the interagent heterogeneity varies.

Increasing heterogeneity: We propose a study to isolate the effect heterogeneity plays in trading personalization for global accuracy and to assess how well the different algorithms personalize the local models. To this aim, we consider the following simplified version of Example 1 that allows us to easily tune the heterogeneity among agents' local distributions and to visually compare the learned models.

Example 2 (Binary classification with 2-D features): The distribution \mathcal{P}_i of agent i produces samples $d_{i,j} = (a_{i,j}, b_{i,j})$, where $a_{i,j} \in \mathbb{R}^2$ and the corresponding label is generated according to the linear model

$$b_{i,j} = \begin{cases} -1, & \text{if } w_i^\top a_{i,j} + v_{i,j} \geq 0 \\ 1, & \text{otherwise} \end{cases} \quad (19)$$

with noise $v_{i,j} \sim \mathcal{N}(0, 0.01)$. Given a parameter $\theta > 0$, we construct the ground-truth vectors $\{w_i\}_{i \in \mathcal{V}}$ as $w_i = [1 \ \theta_i]^\top$, where the parameters $\{\theta_i\}_{i \in \mathcal{V}}$ are evenly spaced between $-\theta$ and θ (included). In words, a larger value of θ amplifies the differences among the slopes of vectors w_i and makes the local agents' distributions more heterogeneous. Two cases are depicted in Fig. 3 with $\theta \in \{0.1, 1\}$. Agent i trains its parameter x_i to learn w_i using the same loss of Example 1. The $N = 10$ agents communicate according to a circulant graph where each agent has four neighbors. We set $\gamma = 10^{-5}$ and train for 1000 iterations with 500 training samples per agent. \square




Fig. 3. Samples of and model learned by agent 1 along with true classifiers w_i of Example 2. Vector w_1 is solid; the other vectors $w_i, i \neq 1$, dotted. (a) Model with (DGD), $\theta = 0.1$. (b) Model with (FJ-DGD-1), $\theta = 0.1$. (c) Model with (DGD), $\theta = 1$. (d) Model with (FJ-DGD-1), $\theta = 1$.




Fig. 4. Test accuracy with DGD and its FJ-based variations under increasing interagent heterogeneity for the binary classification task in Example 2. (a) Local accuracy. (b) Global accuracy.

We set $\lambda = 0.5$ for both (FJ-DGD-1) and (FJ-DGD-2). Fig. 4 summarizes the performance achieved with the four algorithms previously compared. Each mark displays the average accuracy as the location on the y -axis and the standard deviation across agents as the size.⁴ As θ increases, the local distributions become more and more different and (DGD) struggles to balance global for local accuracy. The latter degrades both in average and deviation, with some agents hitting low scores. If each agent independently trains (“Local GD”), the local accuracy barely changes with θ but the global accuracy decreases. On the other hand, the FJ-based variants of DGD gracefully mediate between global accuracy and personalization, achieving significantly higher local accuracy than (DGD) with smaller standard deviation. Interestingly, both (FJ-DGD-1) and (FJ-DGD-2) perform almost identical for this task, especially with high heterogeneity (large θ).

A visual explanation of enhanced personalization is given in Fig. 3 focusing on agent 1. When θ is small, all ground-truth local classifiers are similar to each other, and both (DGD) and (FJ-DGD-1) yield similar models x_1 . However, when θ is large and the local models differ significantly, the model x_1 learned by (FJ-DGD-1) is closer to the true vector w_1 compared to the one learned by (DGD), enhancing personalization while retaining high global accuracy. This behavior can be easily tuned through the parameter λ in (FJ-DGD-1) or (FJ-DGD-2) and makes these algorithms flexible to various needs.

The experiments of Examples 1 and 2 suggest that the FJ-inspired adaptations of DGD are effective for learning personalized models while also retaining good global accuracy according to the nominal problem (1). Their simplicity makes them attractive both for practical implementation and for interpretation of the models, whose degree of personalization can be easily tuned through the parameter λ . In the next section, we consider the scenario with malicious agents.

B. Evaluating Resilience

For this set of experiments, we consider the multiclass classification task described next.

Example 3 (Image classification): We study classification on the MNIST dataset with samples $d = (a, b)$, where a is the flattened (one-dimensional) photo of a handwritten digit and b is the corresponding numerical value from 0 to 9. Each agent i learns the parameter $x_i \in \mathbb{R}^{p \times c}$ of a multiclass logistic classifier, where $p = 784$ is the image size (28×28 pixels) and $c = 10$ is the number of classes (digits). Denoting the ℓ th column of x_i by $[x_i]_\ell$, the multiclass logistic loss is

$$f_i(x_i) = \frac{1}{|I_i|} \sum_{j \in I_i} \left(\log \left(\sum_{c'=1}^c \exp(a_{i,j}^\top [x_i]_{c'}) \right) - a_{i,j}^\top [x_i]_{b_{i,j}} \right) + \gamma \|x_i\|^2 \tag{20}$$

⁴ We scale all standard deviations by 100 to make the marks visible.




Fig. 5. Accuracy on MNIST “Hom” in Example 3 with malicious agents after 300 learning iterations.




Fig. 6. Accuracy on MNIST “Het-2” in Example 3 with malicious agents. (a) After 300 learning iterations. (b) After 1000 learning iterations.




Fig. 7. Local losses, global losses, and local test accuracy for all agents on MNIST “Het-2” in Example 3 with malicious agents. (a) DGD. (b) (FJ-DGD-2) with $\lambda = 0.25$.

with regularization weight $\gamma = 0.1$. Labels are then assigned according to the softmax policy as

$$\hat{b} = \arg \max_{\ell \in \{1, \dots, c\}} \frac{\exp(a^\top [x_i]_\ell)}{\sum_{c'=1}^c \exp(a^\top [x_i]_{c'})} - 1. \tag{21}$$

We simulate $N = 100$ total agents communicating over a random geometric network on the unitary square in \mathbb{R}^2 with communication radius $\rho = 0.25$. We generate three datasets that make collaborative learning increasingly difficult.

- 1) *Homogeneous (Hom)*: We randomly assign 554 samples of MNIST to each agent. Samples are randomly split between local training set (443 samples) and test set (111

samples) for each agent. In this case, all agents most likely have representatives of all classes in both training and test data.

- 2) *Heterogeneous-2 (Het-2)*: We modify the “Hom” datasets and randomly remove two classes from each agent’s local train and test data. Thus, each agent’s local datasets contain at most eight out of the ten digit classes.
- 3) *Heterogeneous-5 (Het-5)*: We randomly remove five classes from each agent’s local “Hom” train and test data.

In all experiments, we show only the global accuracy since in this case we are not interested in personalized models. \square

Motivated by the higher performance of (FJ-DGD-2) in the previous section, we focus on it for the next experiments. First, we set the baseline evaluating classification accuracy in the ideal case with no malicious agents, running 1000 learning iterations; see Fig. 5. We use $\lambda = 0.5$ for this test. As expected, the algorithm (DGD) performs best, followed by (FJ-DGD-2) and finally by local training, and the accuracy decreases as more classes are removed from each local dataset. Note that the steady-state accuracy on the dataset “Hom” is expected to be almost equal with the three algorithms because all agents qualitatively have the same information in this case. Nonetheless, due to different speeds of convergence and augmented training data compared to local datasets, both (DGD) and (FJ-DGD-2) achieve higher accuracy than “Local GD” even after 1000 learning iterations. Such a faster convergence can be a further argument in favor of collaborative training.

We then randomly select ten malicious agents across the network. This yields a ratio between malicious and honest agents of over 11%. We set stealthy attacks whereby each malicious agent m trains its local parameter x_m with (DGD) but, at each iteration k , communicates to its neighbors purposely corrupted parameters \tilde{x}_m computed as follows:

$$\tilde{x}_{m,k} = x_{m,k} + v_k, \quad v_k \sim \mathcal{N}(0, \text{diag}(\min\{\eta|x_{m,k}|, \kappa\})) \quad (22)$$

where the minimization is elementwise.

1) Dataset “Hom”: We first run a shorter training on dataset “Homogeneous” with $\eta = \kappa = 5$. The accuracy is reported in Fig. 6, where we make explicit that $\lambda = 0$ and $\lambda = 1$ are, respectively, (DGD) and “Local GD” (“L. GD”). Our approach tailored to personalization enhances resilience as well, outperforming the benchmarks. The algorithm (DGD) treats all agents, malicious ones included, equally. “Local GD” is insensitive to attacks but does not leverage collaboration with honest agents. Our algorithm (FJ-DGD-2) with $\lambda \in (0, 1)$ stands in the middle; agents do not fully rely on the others but retain benefits of mutual collaboration. The U-shaped curve confirms the intuition that partially reducing collaboration is beneficial as it mitigates the influence of malicious agents, but excessively doing so degrades performance because updates are too conservative. The same behavior was analytically characterized in [30] for resilient consensus. The rapid rise of the accuracy suggests that (FJ-DGD-2) is sensitive to small values of λ , which already improve resilience significantly.

2) Dataset “Het-2”: We then train the agents on “Heterogeneous-2” with the choice $\eta = 10$ and $\kappa = 3$ in (22)

The achieved accuracy is shown in Fig. 7. In addition, Fig. 8 illustrates the behavior of two different algorithms during training, where the solid lines show the average values and the gray filled area delimits one standard deviation interval across all agents at each iteration. For the sake of computational time, the global loss is computed at each iteration on a random subset (chosen beforehand and fixed through the iterations) of the global training set. The influence of malicious agents causes the loss produced by (DGD) to steadily increase in the leftmost and center panels of Fig. 8(a), which quickly settles the accuracy to a suboptimal value (rightmost panel). On the other hand, setting just $\lambda = 0.25$ in (FJ-DGD-2) provides a good level of resilience as particularly evident by the global loss in the center panel of Fig. 8(b) that initially increases but recovers a decreasing trend after about 100 iterations. In this case, the gap between accuracy achieved with our approach and (DGD) local training after 300 iterations is even sharper, as Fig. 7(a) highlights.

Early stopping: Contrary to the experiment without malicious agents in Fig. 5, running (FJ-DGD-2) with $\lambda \in \{0.25, 0.5\}$ for 1000 iterations causes several agents to overfit the training data with malicious agents. To overcome this issue, we implement an early stopping policy during training. At every iteration, each honest agent i computes a moving average (MA) of the accuracy on its local training set over a sliding window of length W , storing the maximal (so smoothed) accuracy achieved so far along with the corresponding local parameter x_i . If the maximal MA of the accuracy does not increase, i.e., no improvement is done, for more than W_{imp} consecutive iterations, the agent resets its local parameter to the one corresponding to the maximal accuracy MA and stops local training, but keeps transmitting its local parameter to neighbors. We set $W = W_{\text{imp}} = 20$ after trial and error.

Fig. 7(b) illustrates the outcome with early stopping. The latter makes (DGD) slightly degrade performance, possibly due to more randomness in the updates; thus, in Fig. 7(b), we report the accuracy (DGD) without early stopping for fairness. Nonetheless, our approach still outperforms both (DGD) and local training with $\lambda = 0.25$, some agents performing particularly well (global accuracy over 80%), and achieves accuracy comparable to local training with $\lambda \in \{0.5, 0.75\}$.

3) Dataset “Het-5”: Finally, we train on the dataset “Heterogeneous-5” for 1000 iterations using early stopping. Fig. 9 reports the achieved accuracy, which shows the same qualitative behavior observed with “Heterogeneous-2.” In this case, the gap between our approach and the baselines is ever wider. Setting $\lambda = 0.25$ sizably outperforms both (DGD) and “Local GD,” while $\lambda \in \{0.5, 0.75\}$ yields slightly higher accuracy than “Local GD.”

VI. LIMITATIONS AND FUTURE RESEARCH

This work provides formal guarantees and study the performance achieved with a globally assigned parameter λ . More realistic scenarios might require that each agent i locally sets its own parameter λ_i . This raises both technical and practical questions, such as if and how convergence can be characterized and how should the agents meaningfully set local parameters

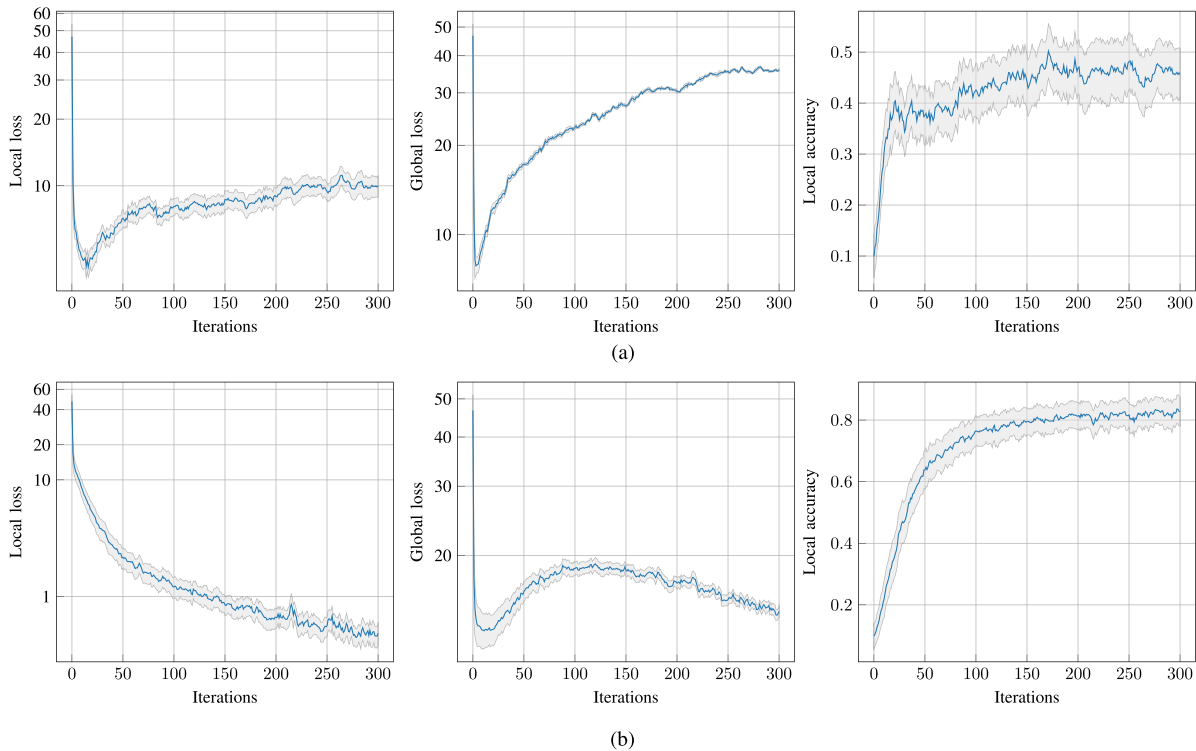


Fig. 8. Accuracy on the MNIST dataset in Example 3 without malicious agents. Marks show the mean and bars the 75% percentile interval across agents. (a) Dataset “Hom.” (b) Dataset “Het-2.” (c) Dataset “Het-5.”




Fig. 9. Accuracy on MNIST “Het-5” in Example 3 with malicious agents after 1000 learning iterations.

λ_i s. It may be particularly challenging with malicious agents, whereby a poor choice of λ_i may degrade resilience and yield low performance. Literature on resilient consensus offers interesting options to dynamically adjust weights, such as heuristic metrics of dissimilarity with neighbors [37], [38] or inspired by the Hegselmann–Krause model [39]. These simple mechanisms may be effective to locally tune competition parameters λ_i s in a decentralized manner during training.

Related to the previous discussion, it is interesting to pair robust training with detection of adversaries. For example, the algorithms in [18] and [19] use trust observations obtained from the wireless channel to filter out suspicious messages at each iteration, and identify all adversaries in finite time almost surely. More broadly, if active and/or time-consuming security

mechanisms are available, the approach proposed in the present work can make the early training rounds resilient and be replaced by a standard distributed learning algorithm after adversaries have been identified.

VII. CONCLUSION

In this article, we have proposed a lightweight distributed learning algorithm, which accommodates for personalization and resilience by combining DGD and the FJ model. We quantified both its geometric convergence rate and the worst case distance from the nominal global optimum in the presence of attacks and experimentally showed its effectiveness by testing it on various classification tasks. In particular, we showcased its ability to learn personalized models that retain good generalization and demonstrated how it enhances resilience against attacks aimed at disrupting the training.

REFERENCES

- [1] J. Yu, J. A. Vincent, and M. Schwager, “DiNNO: Distributed neural network optimization for multi-robot collaborative learning,” *IEEE Robot. Autom. Lett.*, vol. 7, no. 2, pp. 1896–1903, Apr. 2022.
- [2] X. Ma, J. Zhao, and Y. Gong, “Joint scheduling and resource allocation for efficiency-oriented distributed learning over vehicle platooning networks,” *IEEE Trans. Veh. Technol.*, vol. 70, no. 10, pp. 10894–10908, Oct. 2021.
- [3] H. Huang, H. Xu, Y. Cai, R. S. Khalid, and H. Yu, “Distributed machine learning on smart-gateway network toward real-time smart-grid energy management with behavior cognition,” *ACM Trans. Des. Autom. Electron. Syst.*, vol. 23, no. 5, 2018, Art. no. 56.
- [4] B. Sommer, P. Pinson, J. W. Messner, and D. Obst, “Online distributed learning in wind power forecasting,” *Int. J. Forecasting*, vol. 37, no. 1, pp. 205–223, 2021.

- [5] M. Le, T. Huynh-The, T. Do-Duy, T.-H. Vu, W.-J. Hwang, and Q.-V. Pham, "Applications of distributed machine learning for the Internet-of-Things: A comprehensive survey," *IEEE Commun. Surv. Tuts.*, vol. 27, no. 2, pp. 1053–1100, Apr. 2025.
- [6] C. Ma et al., "Trusted AI in multiagent systems: An overview of privacy and security for distributed learning," *Proc. IEEE*, vol. 111, no. 9, pp. 1097–1132, Sep. 2023.
- [7] A. Nedić, A. Olshevsky, and M. G. Rabbat, "Network topology and communication-computation tradeoffs in decentralized optimization," *Proc. IEEE*, vol. 106, no. 5, pp. 953–976, May 2018.
- [8] G. Notarstefano, I. Notarnicola, and A. Camisa, "Distributed optimization for smart cyber-physical networks," *Found. Trends Syst. Control.*, vol. 7, no. 3, pp. 253–383, 2019.
- [9] R. Xin, S. Kar, and U. A. Khan, "Decentralized stochastic optimization and machine learning: A unified variance-reduction framework for robust performance and fast convergence," *IEEE Signal Process. Mag.*, vol. 37, no. 3, pp. 102–113, May 2020.
- [10] T.-H. Chang, M. Hong, H.-T. Wai, X. Zhang, and S. Lu, "Distributed learning in the nonconvex world: From batch data to streaming and beyond," *IEEE Signal Process. Mag.*, vol. 37, no. 3, pp. 26–38, May 2020.
- [11] L. Lyu et al., "Privacy and robustness in federated learning: Attacks and defenses," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 7, pp. 8726–8746, Jul. 2024.
- [12] C. Liu, N. Bastianello, W. Huo, Y. Shi, and K. H. Johansson, "A survey on secure decentralized optimization and learning," 2024, *arXiv:2408.08628*.
- [13] S. Sundaram and B. Ghahserifard, "Distributed optimization under adversarial nodes," *IEEE Trans. Autom. Control*, vol. 64, no. 3, pp. 1063–1076, Mar. 2019.
- [14] C. Fang, Z. Yang, and W. U. Bajwa, "BRIDGE: Byzantine-resilient decentralized gradient descent," *IEEE Trans. Signal Inf. Process. Netw.*, vol. 8, pp. 610–626, 2022.
- [15] Y. Shang, "Median-based resilient consensus over time-varying random networks," *IEEE Trans. Circuits Syst. II*, vol. 69, no. 3, pp. 1203–1207, Mar. 2022.
- [16] J. Li, W. Abbas, M. Shabbir, and X. Koutsoukos, "Byzantine resilient distributed learning in multirobot systems," *IEEE Trans. Robot.*, vol. 38, no. 6, pp. 3550–3563, Dec. 2022.
- [17] W. Abbas, A. Laszka, and X. Koutsoukos, "Improving network connectivity and robustness using trusted nodes with application to resilient consensus," *IEEE Trans. Control Netw. Syst.*, vol. 5, no. 4, pp. 2036–2048, Dec. 2018.
- [18] M. Yemini, A. Nedić, A. J. Goldsmith, and S. Gil, "Resilient distributed optimization for multi-agent cyberphysical systems," *IEEE Trans. Autom. Control*, vol. 70, no. 6, pp. 3952–3967, Jun. 2025.
- [19] L. Ballotta and M. Yemini, "The role of confidence for trust-based resilient consensus," in *Proc. Amer. Control Conf.*, 2024, pp. 2822–2829.
- [20] C. N. Hadjicostis and A. D. Domínguez-García, "Trustworthy distributed average consensus," in *Proc. IEEE Conf. Decis. Control*, 2022, pp. 7403–7408.
- [21] A. Fallah, A. Mokhtari, and A. Ozdaglar, "Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2020, pp. 3557–3568.
- [22] C. T. Dinh, N. Tran, and J. Nguyen, "Personalized federated learning with Moreau envelopes," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2020, pp. 21394–21405.
- [23] F. Hanzely, B. Zhao, and M. Kolar, "Personalized federated learning: A unified framework and universal optimization techniques," *Trans. Mach. Learn. Res.*, 2022. [Online]. Available: <https://jmlr.org/tmlr/index.html>
- [24] A. Sadiet et al., "Decentralized personalized federated learning: Lower bounds and optimal algorithm for all personalization modes," *Eur. J. Comput. Optim.*, vol. 10, 2022, Art. no. 100041.
- [25] M. T. Toghiani, S. Lee, and C. A. Uribe, "PARS-Push: Personalized, asynchronous and robust decentralized optimization," *IEEE Control Syst. Lett.*, vol. 7, pp. 361–366, 2023.
- [26] A. Kundu, P. Yu, L. Wynter, and S. H. Lim, "Robustness and personalization in federated learning: A unified approach via regularization," in *Proc. IEEE Int. Conf. Edge Comput. Commun.*, 2022, pp. 1–11.
- [27] D.-J. Han, D.-Y. Kim, M. Choi, C. G. Brinton, and J. Moon, "SplitGP: Achieving both generalization and personalization in federated learning," in *Proc. IEEE Conf. Comput. Commun.*, 2023, pp. 1–10.
- [28] A. Bietti, C.-Y. Wei, M. Dudik, J. Langford, and S. Wu, "Personalization improves privacy-accuracy tradeoffs in federated learning," in *Proc. Int. Conf. Mach. Learn.*, 2022, pp. 1945–1962.
- [29] T. Li, S. Hu, A. Beirami, and V. Smith, "Ditto: Fair and robust federated learning through personalization," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 6357–6368.
- [30] L. Ballotta, G. Como, J. S. Shamma, and L. Schenato, "Can competition outperform collaboration? The role of misbehaving agents," *IEEE Trans. Autom. Control*, vol. 69, no. 4, pp. 2308–2323, Apr. 2024.
- [31] N. Bastianello and E. Dall'Anese, "Distributed and inexact proximal gradient method for online convex optimization," in *Proc. Eur. Control Conf.*, 2021, pp. 2432–2437.
- [32] K. Yuan, Q. Ling, and W. Yin, "On the convergence of decentralized gradient descent," *SIAM J. Optim.*, vol. 26, no. 3, pp. 1835–1854, 2016.
- [33] K. Yuan, B. Ying, X. Zhao, and A. H. Sayed, "Exact diffusion for distributed optimization and learning—Part I: Algorithm development," *IEEE Trans. Signal Process.*, vol. 67, no. 3, pp. 708–723, Feb. 2019.
- [34] N. E. Friedkin and E. C. Johnsen, "Social influence and opinions," *J. Math. Sociol.*, vol. 15, nos. 3/4, pp. 193–206, 1990.
- [35] A. V. Proskurnikov and R. Tempo, "A tutorial on modeling and analysis of dynamic social networks. Part I," *Annu. Rev. Control*, vol. 43, pp. 65–79, 2017.
- [36] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," *Proc. Mach. Learn. Syst.*, vol. 2, pp. 429–450, 2020.
- [37] J. S. Baras and X. Liu, "Trust is the cure to distributed consensus with adversaries," in *Proc. Mediterranean Conf. Control Autom.*, Akko, Israel, 2019, pp. 195–202.
- [38] V. Bonagura, C. Fioravanti, G. Oliva, and S. Panzneri, "Resilient consensus based on evidence theory and weight correction," in *Proc. Amer. Control Conf.*, 2023, pp. 393–398.
- [39] R. Hegselmann and U. Krause, "Opinion dynamics and bounded confidence: Models, analysis and simulation," *J. Artif. Soc. Soc. Simul.*, vol. 5, no. 3, 2002. [Online]. Available: <https://jasss.soc.surrey.ac.uk/5/3/2.html>



Luca Ballotta (Member, IEEE) received the Ph.D. degree in information engineering from the University of Padova, Padova, Italy, in 2023.

From 2020 to 2022, he was a Visiting Student with the Massachusetts Institute of Technology, Cambridge, MA, USA. From 2023 to 2025, he was a Postdoctoral Researcher with the Delft University of Technology, Delft, The Netherlands. He is currently an Assistant Professor with the Department of Information Engineering, University of Padova. His research

interests include resource allocation in multiagent and network control systems, resilient distributed control and learning, and safe control.

Dr. Ballotta was the recipient of the Young Author Prize at 2020 IFAC World Congress and the finalist of the 2024 EECI Ph.D. Award.



Nicola Bastianello (Member, IEEE) received the bachelor's degree in information engineering and the master's degree in automation engineering in 2015 and 2018, respectively, and the Ph.D. degree in information engineering in 2021, all from the University of Padova, Padova, Italy.

From 2021 to 2022, he was a Postdoctoral Researcher with the Department of Information Engineering, University of Padova. During the Ph.D. degree, he was a Visiting Student with the

Department of Electrical, Computer, and Energy Engineering, University of Colorado Boulder, Boulder, CO, USA. He is currently a Postdoctoral Researcher with the School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology, Stockholm, Sweden. His research interests include the intersection of optimization and learning, with a focus on multiagent systems.

Dr. Bastianello is on the Editorial Board of the IEEE Control Systems Society and the European Control Association Conference.



Riccardo M. G. Ferrari (Senior Member, IEEE) received the Laurea degree (cum laude and printing honors) in electronic engineering and the Ph.D. degree in information engineering from the University of Trieste, Trieste, Italy, in 2004 and 2009, respectively.

He has held both academic and industrial research and development positions, in particular as a Researcher in the field of process instrumentation and control for the steel-making sector. He is a Marie Curie alumnus and currently an Associate Professor with the Delft Center for Systems and Control, Delft University of Technology, Delft, The Netherlands. His research interests include wind power fault-tolerant control and fault diagnosis and attack detection in large-scale cyber-physical systems, with applications to electric vehicles, cooperative autonomous vehicles, and industrial control systems.

Dr. Ferrari was the recipient of the 2005 Giacomini Award of the Italian Acoustic Society. He obtained the second place in the Competition on Fault Detection and Fault-Tolerant Control for Wind Turbines at 2011 IFAC World Congress. Furthermore, he was awarded an Honorable Mention for the Pauk M. Frank Award at 2018 IFAC Symposium on Fault Detection, Supervision and Safety for Technical Processes. He also received an Airbus Award at 2020 IFAC World Congress for the best contribution to the competition on aerospace industrial fault detection.



Karl H. Johansson (Fellow, IEEE) received the M.Sc. degree in electrical engineering and the Ph.D. degree in automatic control from Lund University, Lund, Sweden, in 1992 and 1997, respectively.

He is currently a Swedish Research Council Distinguished Professor in Electrical Engineering and Computer Science with KTH Royal Institute of Technology, Stockholm, Sweden, and the Founding Director of Digital Futures. He has held visiting positions at the University of California, Berkeley, CA, USA; California Institute of Technology, Pasadena, CA; Nanyang Technological University, Singapore, and other prestigious institutions. His research interests include networked control systems and cyber-physical systems with applications in transportation, energy, and automation networks.

Dr. Johansson received numerous best paper awards and various distinctions from the IEEE, International Federation of Automatic Control (IFAC), and other organizations for his scientific contributions. He has been awarded Distinguished Professor by the Swedish Research Council, a Wallenberg Scholar by the Knut and Alice Wallenberg Foundation, and a Future Research Leader by the Swedish Foundation for Strategic Research. He has also received the triennial IFAC Young Author Prize. He is an IEEE Control Systems Society (CSS) Distinguished Lecturer. He is the recipient of the 2024 IEEE CSS Hendrik W. Bode Lecture Prize. His extensive service to the academic community includes being the President of the European Control Association, IEEE CSS Vice President Diversity, Outreach and Development, and a Member of the IEEE CSS Board of Governors and the IFAC Council. He has served on the Editorial Board of *Automatica*, IEEE TRANSACTIONS ON AUTOMATIC CONTROL, IEEE TRANSACTIONS ON CONTROL OF NETWORK SYSTEMS, and many other journals. He is also a Member of the Swedish Scientific Council for Natural Sciences and Engineering Sciences. He is a Fellow of the Royal Swedish Academy of Engineering Sciences.