

Towards Reliable Participation in UAV-Enabled Federated Edge Learning on Non-IID Data

YOUSRA CHERIGUENE ^{1,2}, WAEEL JAAFAR ³ (Senior Member, IEEE),
HALIM YANIKOMEROGU ² (Fellow, IEEE), AND CHAKER ABDELAZIZ KERRACHE ¹

¹Laboratoire d'Informatique et de Mathématiques, Université Amar Telidji de Laghouat, Laghouat 03000, Algeria

²Non-Terrestrial Network (NTN) Lab, Department of Systems and Computer Engineering, Carleton University, Ottawa, ON K1S 5B6, Canada

³Department of Software and Information Technology Engineering, École de Technologie Supérieure, Montreal, QC H3C 1K3, Canada

CORRESPONDING AUTHOR: YOUSRA CHERIGUENE (e-mail: y.cheriguene@lagh-univ.dz).

This work was supported by International Study in Canada Scholarship (SiCS).

ABSTRACT Federated Learning (FL) is a decentralized machine learning (ML) technique that allows a number of participants to train an ML model collaboratively without having to share their private local datasets with others. When participants are unmanned aerial vehicles (UAVs), UAV-enabled FL would experience heterogeneity due to the majorly skewed (non-independent and identically distributed -IID) collected data. In addition, UAVs may demonstrate unintentional misbehavior in which the latter may fail to send updates to the FL server due, for instance, to UAVs' disconnectivity from the FL system caused by high mobility, unavailability, or battery depletion. Such challenges may significantly affect the convergence of the FL model. A recent way to tackle these challenges is client selection, based on customized criteria that consider UAV computing power and energy consumption. However, most existing client selection schemes neglected the participants' reliability. Indeed, FL can be targeted by poisoning attacks, in which malicious UAVs upload poisonous local models to the FL server, by either providing targeted false predictions for specifically chosen inputs or by compromising the global model's accuracy through tampering with the local model. Hence, we propose in this article a novel client selection scheme that enhances convergence by prioritizing fast UAVs with high-reliability scores, while eliminating malicious UAVs from training. Through experiments, we assess the effectiveness of our scheme in resisting different attack scenarios, in terms of convergence and achieved model accuracy. Finally, we demonstrate the performance superiority of the proposed approach compared to baseline methods.

INDEX TERMS Dropouts, edge computing, federated learning, model poisoning, stragglers, UAV, unmanned aerial vehicle.

I. INTRODUCTION

Machine Learning (ML)-assisted techniques are garnering an increasing interest in various research disciplines, including unmanned aerial vehicle (UAV)-enabled wireless networks, as a result of the booming traffic data, the rising deployment of UAVs, and reduced costs [1], [2], [3]. Specifically, UAVs are extensively employed for data collection and offloading [1], [2], which may be complex due to environment and user quality-of-service (QoS) requirements' unpredictable changes. Thus, the development of intelligent management schemes for related applications, e.g., based on ML, is needed. Conventional ML algorithms are

cloud-centric, which makes them inappropriate for stringent QoS and UAV-enabled wireless networks. Indeed, relying on the cloud may expose sensitive information, such as the UAVs' locations and IDs, to attacks, besides increasing the UAVs' energy and bandwidth consumption needed for data transmissions. To bypass this risk, federated edge learning (FEEL) has been recently introduced [4], [5], where UAVs distributively train models on their local datasets without exposing their raw data through transmissions. However, UAV-based FEEL faces several challenges, such as statistical heterogeneity and system heterogeneity. Indeed, due to the different locations of UAVs, collected data is

significantly skewed, i.e., non-independent and identically distributed (non-IID). Moreover, the differences in UAVs' storage, computing, and power consumption, affect the latter's data collection and model training quality. Hence, system heterogeneity results in the appearance of FEEL stragglers, forcing the federated learning (FL) server to wait for the slowest UAVs prior to model aggregation, thus resulting in a slower FL convergence.

In addition, UAV participants can demonstrate unintentional misbehavior during training due to poor connectivity or system faults. For instance, a UAV participant may fail to respond to the FL server due to its high mobility and/or low battery level, resulting in a small number of FL updates. In this case, the FL server would generate a biased global model and/or need to relaunch the FL round from the beginning. Another unexplored issue is the potential intentional malicious behavior of UAV participants [6], [7]. Indeed, such behavior may lead to denial-of-service (DoS) attacks on the FL server (untargeted attacks), or produce false predictions for FL targeted inputs (poisoning targeted attacks). To tackle such problems, defense mechanisms have to rely on specific assumptions related to data distributions, e.g., data IIDness [7], and the number of attackers [6], [8].

The combination of the aforementioned FEEL risks would inevitably degrade the performances of FL, e.g., slow or even cancel model convergence, as well as increase training time, energy consumption, and bandwidth consumption. To the best of our knowledge, no prior work has jointly addressed the intentional and unintentional misbehavior of FL participants in UAV-enabled systems and its impact on FL performance. Consequently, we aim here to investigate the UAV participants' selection problem in heterogeneous data and system settings, in the presence of unreliable UAV participants categorized into one of the following profiles: Adversaries, stragglers, or dropouts. The classification of UAVs into these profiles is realized by analyzing their history and response behavior during training. In the absence of malicious UAVs, i.e., adversaries, our proposed solution acts as an FL client selection scheme that prioritizes reliable UAVs. In contrast, in the presence of malicious UAVs, we target detecting and eliminating the latter in the early stages of FL training. The main contributions of this article can be summarized as follows:

- 1) The diversity of locally available resources has an influence on overall learning performance, therefore, we propose a novel participant selection policy based on favoring highly reliable participants in order to eliminate training instability and reach the best FL accuracy performances.
- 2) In the presence of malicious participants performing untargeted or targeted attacks, we propose a line of defense that aims to identify and filter malicious updates/participants from training without requiring raw data exchanges.
- 3) Under the assumption of non-IID data distributions among UAV participants, we demonstrate through experiments that our proposed selection scheme efficiently

mitigates unreliable participants and achieves superior performances, in terms of FL model accuracy, training time, dropout ratio, and minimal attack success rate, compared to baseline methods.

The remainder of the article is organized as follows. Section II presents the related works. In Section III, the system model is introduced. Subsequently, the problem formulation and the proposed scheme are detailed in Section IV. In Section V, experiments and performance evaluation are conducted. Finally, Section VI concludes the article.

II. RELATED WORKS

A. CLIENT SELECTION IN FL

Participants' resources can be extremely heterogeneous, e.g., different communication and computation capabilities, which may affect FL training [9], [10], thus causing the appearance of stragglers. A straggler can be defined as an FL participant who spends a high time training its data and uploading its related local model to the aggregator or FL edge computing (EC) server. This slow training can be due to several factors such as the low computing power, the degraded communication link with the EC server, and the large size of the dataset. Indeed, the EC server should be receiving updates from participants in a reasonable time. On the other hand, a participant who leaves the FL training process before finishing its given tasks is referred to as a dropout. Dropping out from the FL system, i.e., quitting during the training or the aggregation phase in a given round, can significantly impact the FL performance. In fact, in such a situation, the EC server shall interrupt the current training round and relaunch it, thus leading to lost data and wasted energy consumption. The dropping out event can be caused by intentional or unintentional quitting from the FL system. The latter can be due to the participant being hijacked, or disconnected due to traveling away from the EC server.

To tackle the aforementioned issues, client selection has been developing in recent years. In particular, for the stragglers' issue, client selection policies are set up to improve the overall training latency. For instance, the proposed client selection technique of [11], called Oort, favored the participation of clients with the most useful data for both model accuracy enhancement and faster training. To mitigate the effect of unreliable participants, Oort removed clients that participated successfully in a given number of FL rounds. Nevertheless, this defense mechanism treated malicious and non-malicious participants equally, i.e., without any distinction.

Subsequently, client clustering has been involved in the selection process to improve FL performances. For example, authors of [12] presented the TiFL algorithm, which divided participants into different tiers according to their training response latency. Then, participants belonging to the same tier would be selected at each training round. TiFL used an adaptive tier selection technique that adjusted clustering (or tiering) to the observed training performances over time. In [13], the authors proposed FedRN, a robust FL method

that exploits k -reliable neighboring participants with high data similarity to train the model only with clean samples over reliable participants.

Due to the heterogeneity of computing and communication resources at FL participants, the authors of [14] tackled the stragglers' problem by proposing FedCS. FedCS is a greedy FL mechanism that favors incorporating more participants into its deadline-constrained training process, regardless of the datasets choice, and communication and computing capabilities.

Given the complexity of FL systems, intelligent participant selection methods based on reinforcement learning (RL) have been proposed. For instance, a deep RL control method has been presented in [15] to accomplish the desired trade-off between local updating and global aggregation, in order to minimize the loss function while adhering to a resource budget limitation. In the context of FL for wireless networks, authors in [16] examined the client scheduling problem to achieve low FL latency and rapid convergence, while taking into account the full or partial knowledge of channel state information (CSI) in communication links. The scheduling problem has been solved using a multi-armed bandit (MAB) solution. Specifically, MAB was proposed to learn statistical data online without prior CSI knowledge and use it to build the scheduling strategy. Similarly, federated reinforcement learning (FRL) was proposed in [17] aiming to improve local driving models of connected and autonomous vehicles (CAVs). Specifically, CAV selection for FRL has been investigated, while taking into account CAVs' reputation, quality of driving local models, and training time overhead.

To automate quality-aware FL participant selection, authors of [18] proposed a novel scheme named AUCTION. The latter used reinforcement learning (RL) to progressively enhance the participants' selection policy through interaction with the FL platform after encoding the selection policy into an attention-based neural network. Although the developed method considered the quantity and quality of training data and cost of training, it ignored the effect of communication and computing latency, i.e., the stragglers' effect. Alternatively, the latter can be reduced using asynchronous FL where participants upload their local models asynchronously to the server [19], [20], [21].

On the other hand, the dropout problem was investigated by only a few in the literature. For instance, [10] proposed a multi-criteria-based client selection technique that combines for each client the computing power, storage, and energy consumption to forecast the latter's potential to carry out FL training rounds and handle the least number of discarded rounds caused by client dropouts. When client dropouts are caused by highly dynamic environmental conditions, authors of [22] proposed decentralized FL to reduce mandatory communications with the central FL server and rely instead on cooperative aggregation with nearby participants. In [23], a mobile edge computing (MEC) framework is proposed for a multi-layer federated learning protocol, named HybridFL. This approach implemented model aggregation over two

levels, namely edge and cloud levels, using several aggregation techniques. Moreover, it relied on a probabilistic method for participant selection at the edge layer to mitigate both stragglers and dropouts, while ignoring the state of participants (whose reliability is agnostic). Through simulations, HybridFL is shown to shorten FL rounds' length and accelerate convergence.

In the aforementioned works, although participants may struggle to be performing well in FL, they have been considered benign or non-malicious, i.e., they do not misbehave intentionally against the FL process.

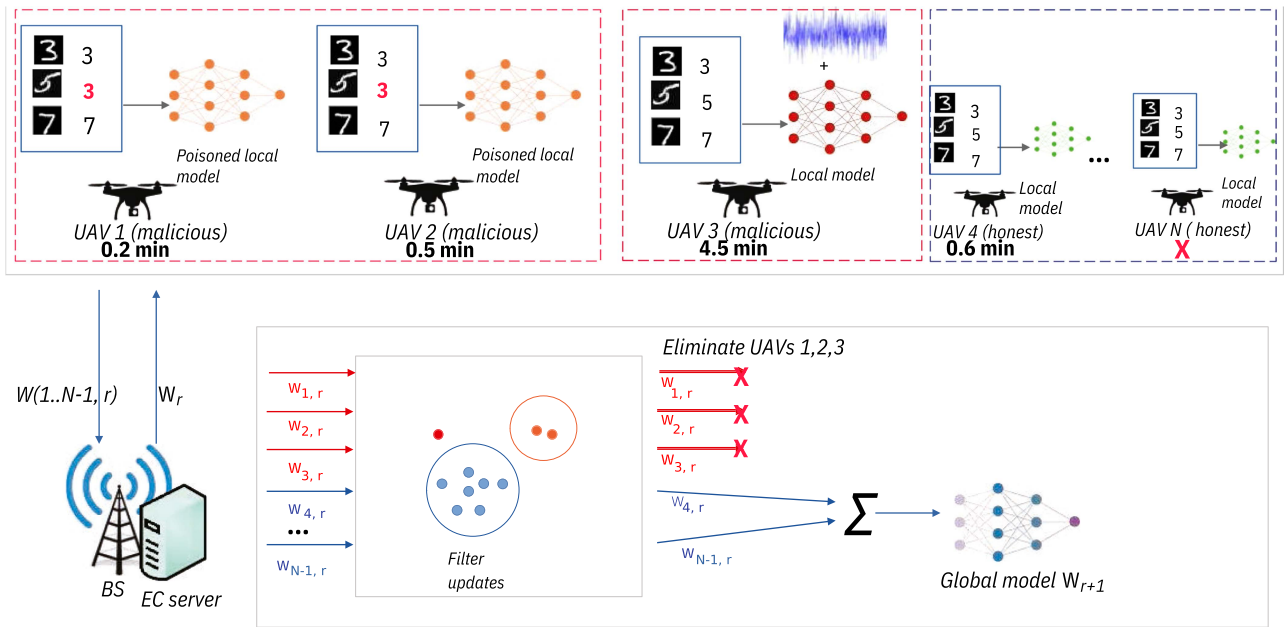

B. ATTACKS ON FL AND DEFENSE MECHANISMS

Despite FL's successful deployment in many fields, this ML technique is prone to attacks. Indeed, given that participants have complete authority over the local training process, the latter becomes an attack surface through which an attacker, e.g., a malicious client, can access the learning process through the client system's vulnerabilities. Subsequently, a poisoning (targeted) attack can be conducted, in which the ML model is altered during training, thus inducing the global model into producing erroneous outputs that correspond to the attacker's action on certain inputs, e.g., the attacker can conduct a data poisoning attack where it alters an image classifier to change an attacker-targeted label of local data samples with specific attributes. An example of a targeted attack is illustrated in Fig. 1 for UAVs 1 and 2 [24]. Another type of attack is the untargeted attack, where the attacker aims to disrupt the training process by producing incorrect outputs on all inputs and thus slowing down FL convergence. We illustrate this case in Fig. 1 for UAV 3 [25].

For the sake of clarity, we detail the targeted and untargeted attack objectives in Fig. 2. In the case of a targeted attack, attackers with a specific target or objective (red arrows) work together in each training round to advance towards the defined objective, by gradually shifting their malicious inputs. In an untargeted attack (orange arrows), attackers demonstrate different and even divergent behaviors, to achieve different local objectives.

Since untargeted attacks have a direct and significant effect on the FL accuracy performance, they can be detected and mitigated through robust aggregation schemes performed at the server level [6], [8]. Similarly, a method called attestedFL was proposed in [25] where each node's training history is monitored via a fine-grained assessment technique to detect malicious participants. However, the latter was proven inefficient in highly non-IID datasets.

For targeted attacks, authors of [26] proposed a defense mechanism, called CONTRA, which uses a cosine-similarity-based metric and builds reputation by dynamically rewarding or punishing participants depending on their historical contributions to the FL global model. In [7], authors presented Justinian's GAAvernor (GAA), a defense method that relies on RL to mitigate malicious attacks. Specifically, GAA generated reward signals for policy learning using a quasi-validation set, i.e., a combination of a small dataset that


FIGURE 1. System model.

FIGURE 2. Model attacks' objectives vs. true objective (Red and orange arrows point to the directions of the attacks' objectives, and blue arrows to the true objective).

exhibits a distribution similar to the actual sample distribution, and historical interactions data with the participants, as experience. However, this method was applied only for IID data samples, which is not realistic in the context of UAV-enabled systems. In [8], [27], model updates were clustered together, and the small-sized clusters were identified as malicious, thus eliminated from the FL process. Despite its simplicity, this approach would inevitably delete relevant data from non-malicious participants, due to clustering inaccuracy.

Other methods relied on reputation design to avoid the negative impact of the participants' malicious behavior on FL performance. For instance, authors of [28], [29], [30] proposed reputation-based participants' selection, where a reputation score is updated either according to the state and channel quality of each participant [28], or using reputation-aware clustering [29], or using an incentive mechanism to favor the frequent selection of highly reliable participants [30].

According to the above discussion, there is a need to mitigate unreliable participants in FL heterogeneous systems.

Table 1 below summarizes the key differences between related works and our contribution. To the best of our knowledge, existing works are not effective in defending against unreliable participants demonstrating different behaviors. Furthermore, most works can not be applied directly to UAV-based systems due to the latter's heterogeneity.

III. SYSTEM MODEL

The considered system model is depicted in Fig. 1. It is composed of a base station (BS) that is paired to an edge computing (EC) server to satisfy ground users' and N_U UAVs' FL tasks. Ground users and N_U UAVs are deployed in the BS's 3D coverage area. The UAVs act as edge nodes with lower capabilities than the BS, but sufficiently high to perform local ML training on their local datasets, denoted as DS_i , $\forall i = \{1, \dots, N_U\}$ ¹. For the sake of simplicity, we assume here that the FL process is set up among the BS as the aggregator and only the UAVs as participants, where the model's local update parameters, denoted by weight vector $w_{i,r}$ for UAV i in FL round r , are uploaded from the UAVs to the BS. Moreover, we assume that the EC server has four engines: A participant filtering engine, a participant selection engine, a security engine, and an aggregation engine. The participant filtering engine identifies and eliminates the stragglers in the participants' pool. The participants' selection engine chooses the UAVs that will train and upload their model parameters in each FL round. The security engine is used to eliminate unreliable participants, while the aggregation engine performs model aggregation on the local updates of selected UAVs into

¹We assume here that UAVs have sufficient power to execute ML while flying. This can be justified by the advancements in powerful miniaturized computing devices with low-power consumption, e.g., tiny ML [32], [33].

TABLE 1. Comparison of Related Works

Reference	Dataset Setting	Features				Defense strategy
		Dropouts	Stragglers	Client Selection	Attack type	
[11]	Non-IID	–	✓	✓	–	Pre-defined maximum number of participation rounds per client.
[12], [15], [16], [14], [17]	Non-IID	–	✓	✓	–	–
[10], [22], [23]	Non-IID	✓	✓	✓	–	–
[13]	Non-IID	–	–	✓	–	–
[18]	Non-IID	–	–	✓	–	–
[19]–[21]	Non-IID	–	✓	✓	–	–
[26], [27]	IID	–	–	Random	Untargeted	Robust aggregation based on median statistics of received gradients.
[8]	IID	–	–	Random	Untargeted	The gradient with the closest distance to its neighboring gradients is selected as the global gradient.
[25]	IID	–	–	Random	Untargeted	Use of cosine similarity between models sent by the same participant + use of quasi-validation dataset prior to aggregation.
[7]	IID	–	–	Random	Targeted	RL technique that uses participants' interactions as experience and a quasi-validation dataset for reward signals.
[28]	Non-IID	–	–	✓	Targeted	Using a reputation metric, the global model rewards or punishes particular participants depending on their past contributions.
[30]	Non-IID	–	–	✓	Targeted	Use of a reputation model based on beta distribution function to determine the credibility of local participants.
[31]	Non-IID	–	–	✓	Targeted	Use of a reputation-aware FL participant selection method based on stochastic integer programming.
[32]	Non-IID	–	–	✓	Targeted	The combination of a multi-weight subjective logic model with a reputation-based participant selection approach for reliable FL.
[33]	IID	–	–	Random	Targeted	Use of dimensionality reduction techniques + K-means clustering of received updates.
This work	Non-IID	✓	✓	✓	Targeted + untargeted	IQR-based straggler filtering + reliability-based dropouts filtering + DBSCAN clustering of cosine similarity distances between received updates for malicious UAVs filtering.

a global model, as shown in Fig. 1. Furthermore, for ease of reading, notations are described in Table 2.

A. FEDERATED LEARNING MODEL

Initially, the EC server creates a generic model for a specific FL task, then selects a set of participants with whom to communicate model parameters. The minimum number of participants selected is $P_r = N_U \times f_r$, where $f_r \in [0, 1]$ is the fraction of participants involved in round r . Once selected, UAV client i downloads the most recent global model w_{r-1} from the EC server for synchronization. Then, it executes multiple iterations of stochastic gradient descent (SGD) on mini-batches $M_{i,r}$ from its dataset DS_i , to update its model weight $w_{i,r}$, such that

$$w_{i,r} = SGD(w_{r-1}, M_{i,r}) - w_{r-1}, \forall i = 1, \dots, P_r. \quad (1)$$

Periodically, the resulting weight update is uploaded to the EC server. Subsequently, the EC server aggregates the collected local models' parameters to produce the global model w_r as follows:

$$w_r = w_{r-1} + \sum_{i=1}^{P_r} \frac{|M_{i,r}|}{|M_r|} w_{i,r}, \quad (2)$$

where $M_r = \cup_{i=1}^{P_r} M_{i,r}$ denotes the number of data samples from P_r participants in round r . If an insufficient number of model updates is received, i.e., a high dropout ratio, the round is canceled by the EC server.

B. FL LATENCY MODEL

Based on the hardware configuration information and size of the local training dataset of the i^{th} UAV, the EC server can estimate the local training time as [34]

$$t_{i,r} = T_e \frac{\kappa |M_{i,r}|}{\gamma_i}, \forall i = 1, \dots, P_r, \quad (3)$$

where κ is the number of CPU-cycles needed to process one data sample, T_e is the number of local iterations, and γ_i is the available CPU-frequency of UAV client i .

At the end of T_e local training iterations, the neural networks' (NN) updated parameters are transmitted through the wireless channel to the EC server, via the BS.

Assuming that communications between the BS and UAVs follow the free-space air-to-ground path loss model, that BS and UAVs are equipped with a single-antenna, and that simultaneous communications from UAVs to the BS are orthogonal in the frequency domain, then the communication link between UAV i and the BS can be represented as

$$h_i = \sqrt{\beta_0} d_i^{-\frac{\alpha_i}{2}}, \forall i = 1, \dots, N_U, \quad (4)$$

where d_i denotes the 3D distance between UAV i and the BS, β_0 is the reference channel gain, and α_i is the path-loss component expressed as

$$\alpha_i = \frac{a_1}{1 + a_4 \exp(a_3(\theta_i - a_4))} + a_2, \forall i = 1, \dots, N_U, \quad (5)$$

TABLE 2. Notations

Notation	Description
N_U	Number of UAVs
r_0	Number of global rounds
$w_{i,r}$	Weight vector of client i in round r
w_r	Aggregated global model in round r
$M_{r,i}$	Number of data samples of client i
$t_{i,r}$	Local training time of client i in round r
$t_{i,BS}$	Upload time of local updates of client i
$t_{BS,i}$	Download time of global model to client i
t_r	Total time for round r
t_{ag}	Aggregation time on weight updates
t_r^i	Training and uploading time of client i in round r
T_e	Number of local iterations
κ	Number of CPU-cycles per data sample
γ_i	CPU frequency of client i
d_i	3D distance between UAV i and the BS
β_0	The reference channel gain
h_i	Communication link between UAV i and the BS
$R_{i,BS}$	Uplink data rates between UAV i and the BS
$R_{BS,i}$	Downlink data rates between UAV i and the BS
B_i	Bandwidth reserved to UAV i by the BS
P_i	Transmission power of UAV i
P_{BS}	Transmission power of the BS
V	The number of NN model parameters
z	The size of a model parameter
σ^2	The unitary power of the additive white Gaussian noise
α_i	Path-loss component
$A_{i,r}$	Accuracy of local model of client i in round r
A_g	Global accuracy after r_0 rounds of training
$\Lambda_{i,r}$	Reliability score of client i in round r
ρ	Min threshold value for the reliability score
ρ_0	Max threshold of the reliability score
W_r	Received weight updates in round r
$\Gamma_{i,r}$	Score of client i in round r
$\kappa_{i,j}$	Cosine similarity between weight updates i and j
$\hat{w}_{i,r}$	Weight update drift of client i in round r
\mathcal{P}_r	Participating UAVs in round r
\mathcal{P}_r^h	Set of honest UAVs in round r
\mathcal{P}_r^s	Set of straggler UAVs in round r
\mathcal{P}_r^d	Set of dropout UAVs in round r
\mathcal{P}_r^m	Set of malicious UAVs in round r
$\tilde{\mathcal{P}}_r$	Participating UAVs in aggregation in round r
ν	Upper bound scaling parameter
f	Attackers ratio in the UAV network
χ_d	Dropout ratio in the UAV network
f_r	Minimum UAV participants fraction
τ_a	Average FL round time

where θ_i is the elevation angle between the BS and UAV i , and $\{a_1, \dots, a_n\}$ are parametric constants related to the urban environment [35].² Subsequently, the achievable uplink and downlink data rates for FL data exchange can be, respectively, given by

$$R_{i,BS} = B_i \log_2 \left(1 + \frac{P_i |h_i|^2}{B_i \sigma^2} \right), \quad \forall i = 1, \dots, N_U, \quad (6)$$

and

$$R_{BS,i} = B \log_2 \left(1 + \frac{P_{BS} |h_i|^2}{B \sigma^2} \right), \quad \forall i = 1, \dots, N_U, \quad (7)$$

²Different air-to-ground channel models can be considered. However, their impact is expected to be similar to that of the considered channel model since variations in channel quality would lead to the same effect from the FL perspective, e.g., slower FL convergence, lower FL accuracy, and higher UAV power consumption.

where B_i (resp. B) is the bandwidth allocated to the transmission of UAV i (resp. the BS), P_i (resp. P_{BS}) is the transmit power of UAV i (resp. the BS), and σ^2 is the unitary power of the additive white Gaussian noise (AWGN). Hence, the required times for FL updates' upload over the uplink, and download over the downlink, can be respectively expressed by

$$t_{i,BS} = Vz/R_{i,BS}, \quad (8)$$

and

$$t_{BS,i} = Vz/R_{BS,i}, \quad (9)$$

where V is the number of NN model parameters and z is the size of a single model parameter, e.g., a NN weight. Consequently, the overall duration of global FL round r at the EC server is expressed as follows:

$$t_r = \max_{i=1, \dots, P_r} (t_{i,r} + t_{i,BS}) + t_{ag}, \quad (10)$$

where t_{ag} is the time taken by the EC server to perform aggregation on the participants' weight updates. Moreover, since $t_{BS,i}$ is negligible, thus it is eliminated from t_r .

C. EC SERVER AND UAVS' SECURITY PROFILES

In our system, we assume that the EC server is benign, i.e., is not prone to cyberattacks and do respond correctly to FL operations. Regarding the UAV participants, we consider two security profiles, namely reliable UAVs and unreliable UAVs. Expressly, we assume that:

- A reliable UAV i is an honest participant that accurately reports its model parameter values, i.e., it uploads non-tampered vector weights $w_{i,r}$ to the EC server in round r .
- An unreliable UAV would demonstrate unreliable behavior in different manners. The first manner consists of dropping out, intentionally or not, from FL training after it started. The second manner is to submit poisonous model parameters to the EC server. This situation happens when an attacker takes control of one or several UAVs, makes them malicious, and obliges them to alter their own model parameters, either trained or transmitted ones. To align with practical systems, we assume that an attacker cannot corrupt more than 50% of the available UAV participants.

IV. PROBLEM FORMULATION AND PROPOSED SOLUTION

A. PROBLEM FORMULATION

In order to design a resilient FL system, we need to address the negative impact of statistical and system heterogeneity on the accuracy metric, denoted $A_{i,r}$, of each UAV participant i in round r . Then, we address the security and reliability of the FL process in the presence of unreliable UAV participants. Formally, let $\mathbf{x}_r = [x_{1,r}, \dots, x_{N_U,r}]$ be the vector of binary variables that indicate the participation or not of UAVs in FL round r , while ζ is a deadline for local training and update weights transmission. We assume that among the participating

N_U UAVs, there exist four sets, namely \mathcal{P}_r^h , \mathcal{P}_r^s , \mathcal{P}_r^d , and \mathcal{P}_r^m , of sizes P_r^h , P_r^s , P_r^d , and P_r^m , respectively, such that $N_U = P_r^h + P_r^s + P_r^d + P_r^m$. These sets correspond respectively to honest (i.e., non-straggler, non-dropout, and non-malicious), stragglers, dropouts, and malicious participants.³ Our objective consists of maximizing the accuracy by aggregating the update weights from \mathcal{P}_r^h only. Let $\tilde{\mathcal{P}}_r^h$ be the decided set of participating UAVs in the aggregation process, based on the \mathbf{x}_r strategy, i.e., $|\tilde{\mathcal{P}}_r^h| = \sum_{i=1}^{N_U} x_{i,r}$. Subsequently, the related problem can be formulated as follows:

$$\max_{\mathbf{x}_r, \zeta} \sum_{i=1}^{N_U} x_{i,r} A_{i,r} \quad (\text{P1})$$

$$\text{s.t. } t_{i,r} + t_{i,BS} \leq \zeta, \forall i \in \tilde{\mathcal{P}}_r^h, \quad (\text{P1.a})$$

$$0 \leq \sum_{i=1}^{N_U} x_{i,r} \leq N_U, \quad (\text{P1.b})$$

$$x_{i,r} \in \{0, 1\}, \forall i = 1, \dots, N_U, \quad (\text{P1.c})$$

where constraint (P1.a) indicates that local training and update weights transmission should not exceed a deadline ζ . If ζ is exceeded, the concerned UAV participant is considered as a dropout. Constraint (P1.b) ensures that the maximal number of participants does not exceed the number of available UAVs, while (P1.c) reflects the binary nature of the participant selection variable.

Problem (P1) can be proven NP-hard. Indeed, for a given ζ , it aims to maximize a sum objective, while respecting a resource constraint (maximal number of UAV participants), (P1) can be reduced to a knapsack problem, which is already proven to be NP-hard [36]. Since it is intractable to directly find the global optimum of (P1), we propose to solve it using a cascaded solution, where the first sub-solution eliminates the stragglers and sets ζ , the second focuses on selecting non-dropouts from the remaining non-stragglers UAVs, while the third sub-solution filters out malicious participants from the resulting set of the second sub-solution. These sub-solutions are detailed in the subsections below.

B. PROPOSED NON-STRAGGLERS PARTICIPANT SELECTION

Most research in the literature has focused on achieving a higher convergence rate for client scheduling by minimizing the number of communication rounds required to reach a specific level of global accuracy. However, this approach does not minimize the elapsed round time during training. Indeed, training time, $t_{i,r}$, is constrained by stragglers who would extend the training round time, due for instance to limited CPU frequency speed γ_i .⁴ In addition, participants may be

³For the sake of simplicity we assume that \mathcal{P}_r^h , \mathcal{P}_r^s , \mathcal{P}_r^d , and \mathcal{P}_r^m are distinct sets.

⁴To be noted that the EC server can estimate the UAVs' CPU information through the processing of a standardized training task on them and tracking the time required to complete it.




FIGURE 3. Distribution of training round time for different UAVs' CPU frequencies.

training different numbers or sizes of data samples, due to the statistical heterogeneity, thus affecting the training round time. In Fig. 3, we illustrate an example of the system heterogeneity's effect, in terms of $t_{i,r}$, and as a function of UAVs' CPU frequency ranges of γ_i . As it can be seen, having UAVs with a large range of CPU capabilities, e.g., $\gamma_i \in [10^5, 10^7]$ Hz, resulted in the highest number of outliers or stragglers (black diamonds).

Stragglers can be eliminated from the participants' set \mathcal{P}_r by applying the interquartile range (IQR) [37], which is a popular statistical method that detects outliers. Our choice of IQR value is motivated by its high robustness in skewed data. IQR value is calculated as the difference between quartiles $Q3$ (i.e., third quartile of the data or upper 75% of data), and $Q1$ (i.e., lower 25% of data). Differently from the classic IQR method, we define our decision range as the upper bound of the accepted data points, and we do not consider the lower bound. Hence, any data point that has $t_{i,r}$ larger than the upper bound, defined as

$$IQR_{up} = Q3 + \nu \times IQR, \quad (11)$$

where ν is a scaling parameter, where $\nu > 0$. Thus, the participants' filtering engine eliminates stragglers from \mathcal{P}_r and generates a novel set of UAVs, denoted \mathcal{P}_r^{ms} . Subsequently, we define by

$$\zeta = \frac{2}{|\mathcal{P}_r^{ms}|} \sum_{i \in \mathcal{P}_r^{ms}} t_{i,r}, \quad (12)$$

the deadline for local training and update weights transmission. For the sake of simplicity, the value of ζ is chosen conservatively (twice the average local training time of non-stragglers), since transmission time is typically one or several orders of magnitudes lower than training time.

Within (11), the scaling factor ν is a constant multiplier that is applied to the IQR in order to determine the outliers' threshold. Actually, the value of ν to be used significantly depends on the application type and datasets properties. Nevertheless, most applications in the literature opted for a scaling factor of

1.5, as it is agreed to be a typical acceptable value [38], [39]. In any case, if the training time showcases extreme values or high variability, a higher scaling factor would be preferable. In contrast, a small scaling factor would be suitable to prevent false positives if the training times are homogeneous. In addition, using a large ν would lead to an aggressive straggler elimination strategy, which might affect the training accuracy, whereas a small ν value leads to a more conservative straggler elimination, and probably to longer training times.

C. PROPOSED NON-DROPOUTS PARTICIPANT SELECTION

Given the high mobility and fragility of UAV systems, unsteady communications or equipment failure may occur, which will lead to losing the client during the training round, i.e., dropping out of the FL process. As a result, the convergence and performance of the FL under partial participation mechanism is significantly impacted by the quantity of aggregated models. According to [40], increasing the number of aggregated local models improves the accuracy performance. However, in a realistic system, not all selected participants would succeed in training and uploading their updates within the fixed deadline. Consequently, the reliability of participants is vital to ensure convergence to an enhanced accuracy.

To mitigate dropouts and maintain resilient operations, we define here a reliability score, denoted by $\Lambda_{i,r}$ for UAV $i \in \mathcal{P}_r^{ns}$, which provides an assessment of the participant's behavior within the FL process. Specifically, a UAV uploading its model parameters within ζ is rewarded, while those failing, i.e., dropping out of the training round, are penalized. The reliability score is expressed by

$$\Lambda_{i,r} = \Lambda_{i,r-1} + g, \quad \forall i \in \mathcal{P}_r^{ns}, \quad (13)$$

where $\Lambda_{i,r-1}$ is the obtained reliability score for the previous round $r - 1$, and

$$g = \begin{cases} 0, & \text{if UAV } i \text{ is asked to participate in round } r \\ & \text{but did not respond, or started responding} \\ & \text{but got disconnected before completing,} \\ 1, & \text{if UAV } i \text{ participated and (P1.a) is satisfied} \\ -1, & \text{otherwise.} \end{cases} \quad (14)$$

Assuming that a ‘‘reliable’’ (non-dropout) participant is defined by a value $\Lambda_{i,r}$ equal or greater than a reliability threshold ρ , then the set of potential participants can be reduced to $\mathcal{P}_r^{nd} = \{i \in \mathcal{P}_r^{ns} \mid \Lambda_{i,r} \geq \rho\}$. Furthermore, in order to avoid bias in the global model and fast battery exhaustion of participants, we establish a client over-selection avoidance mechanism, where if $\Lambda_{i,r}$ reaches a pre-fixed maximum value ρ_0 , it is reset to 0.

In the particular case of a portion or all scores being lower than ρ , Based on the calculated $\Lambda_{i,r}$ values, the participant selection engine prioritizes the inclusion of $P_r - |\mathcal{P}_r^{nd}|$ high-score UAVs to \mathcal{P}_r^{nd} , which have $\Lambda_{i,r} < \rho$.

D. PROPOSED NON-MALICIOUS AGGREGATION

Following the previous steps to obtain a non-straggler non-dropout participants set \mathcal{P}_r^{nd} , we focus here on achieving non-malicious aggregation of update weights at the EC server.

As an input for our problem, we use the update weights inputs, denoted w_r^{nd} , of the participants' set \mathcal{P}_r^{nd} . The EC server evaluates $w_{i,r}^{nd}$, $\forall i \in \mathcal{P}_r^{nd}$ to classify them as malicious (i.e., targeted or untargeted attack) or non-malicious. Then, malicious update weights are eliminated from FL aggregation.

To do so, we start by characterizing the malicious attacks as follows. First, targeted attacks are easy to perform, e.g., erroneously classifying samples in a dataset, but difficult to detect since they subtly effect the global model. if the number of attackers is large, they typically demonstrate similar behaviours that are different from those with non-malicious intentions. Second, untargeted attacks demonstrate behaviours dissimilar from any other participant, thus their weight updates can be considered as outliers. A successful attack would eventually deviate the FL convergence direction towards a different one. By measuring the pairwise angular distance (PAD) between the participants' update weights, the desired convergence direction can be estimated. The use of PAD is preferred over other measures, such as the Euclidean distance, since malicious participants cannot alter the direction of a gradient without reducing their attack's efficacy. Hence, we propose here to calculate the PAD using the cosine similarity.

Unlike [41] that calculated cosine similarity PAD on mini-batch gradients to estimate the direction of the true gradient, authors of [42] demonstrated that an approximate estimation of the true gradient's direction can be obtained by applying cosine similarity PAD on the FL's weight updates, under sufficiently smooth FL loss function and low learning rate conditions. Consequently, we opt here for a similar approach as [42], where the cosine similarities between weight updates, defined by $\kappa_{i,j}$, $\forall (i, j) \in \mathcal{P}_r^{nd} \times \mathcal{P}_r^{nd}$, are calculated as

$$\kappa_{i,j} = \frac{\langle w_{i,r}, w_{j,r} \rangle}{\|w_{i,r}\| \|w_{j,r}\|}, \quad (15)$$

where $\langle \cdot, \cdot \rangle$ is the Euclidean dot product and $\|\cdot\|$ is the magnitude's operator.

With the cosine similarity factors in hand, we transpose this information into a graph \mathcal{G} having the elements of \mathcal{P}_r^{nd} as vertices, and $\kappa_{i,j}$ as edges. According to the previous discussion, targeted attackers would behave similarly, while untargeted attackers can be seen as outliers. Backed with these assumptions, in addition to the ratio of attackers (f) being typically below half of participants in \mathcal{P}_r^{nd} , i.e., $f < 50\%$, we can adopt a clustering approach to detect targeted and untargeted attackers in \mathcal{G} . For the sake of simplicity, we opt for the density-based spatial clustering of applications with noise (DBSCAN) to execute this task. The expected result is several clusters, where the largest one would correspond to the non-malicious participants, while smaller ones correspond to targeted attackers, and outliers to untargeted attackers. Subsequently, the EC server aggregates $w_{i,r}^{nd}$ from the non-malicious (largest) cluster only, denoted as $\tilde{\mathcal{P}}_r^h$.

Algorithm 1: UAV Reliable Participation.

Input: number of UAVs N_U , minimum number of participants P_r , number of global rounds r_0 ;

Steps:

EC server initializes w_0 ;

Randomly assign $\Lambda_{i,1} \in [0, 9], \forall UAV_i \in \{1, \dots, N_U\}$;

```

for round  $r = 1 \dots r_0$  do
  /* Filter 1: Remove stragglers */
  Calculate  $IQR_{up}$  using eq.(12) ;
  for each  $UAV_i$  in  $\{1, \dots, N_U\}$  do
    if  $t_{i,r} < IQR_{up}$  then
      |  $\mathcal{P}_r^{ns} \leftarrow \mathcal{P}_r^{ns} \cup \{UAV_i\}$ ;
    end
  end
  Calculate  $\zeta$  using eq.(13) ;
  /* Filter 2: Remove dropouts */
  for each  $UAV_i$  in  $\mathcal{P}_r^{ns}$  do
    if  $\Lambda_{i,r} \geq \rho_0$  then
      |  $\Lambda_{i,r} \leftarrow 0$ ;
    end
    if  $\Lambda_{i,r} \geq \rho$  then
      |  $\mathcal{P}_r^{nd} \leftarrow \mathcal{P}_r^{nd} \cup \{UAV_i\}$  ;
    end
  end
  if  $|\mathcal{P}_r^{nd}| < P_r$  then
    Sort  $\mathcal{P}_r^{ns} \setminus \mathcal{P}_r^{nd}$  in descending order w.r.t.  $\Lambda_{i,r}$  ;
    Add the first  $P_r - |\mathcal{P}_r^{nd}|$  UAVs from  $\mathcal{P}_r^{ns} \setminus \mathcal{P}_r^{nd}$  to  $\mathcal{P}_r^{nd}$ 
    (with highest  $\Lambda_{i,r}$ ) ;
  end
  /* Filter 3: Remove malicious participants */
  for each  $UAV_i$  in  $\mathcal{P}_r^{nd}$  do
    Update local model using eq.(1) for  $T_e$  local iterations;
    for each  $UAV_j$  in  $\mathcal{P}_r^{nd}$  where  $j \geq i + 1$  do
      Calculate and store cosine similarity  $\kappa_{i,j}$  using
      eq.(16) ;
    end
  end
  Apply DBSCAN and obtain clusters sorted from the largest
  to the smallest as  $\mathcal{C}_r = \{\mathcal{C}_1, \dots, \mathcal{C}_{|\mathcal{C}_r|}\}$  ;
  Set  $\tilde{\mathcal{P}}_r^h = \mathcal{C}_1$  ;
  Update  $w_r$  on  $\tilde{\mathcal{P}}_r^h$  using eq.(2) ;
  Update  $\Lambda_{i,r+1}$  using eq.(14),  $\forall UAV_i \in \{1, \dots, N_U\}$  ;
end
Output:  $w_{r_0}$  global model after  $r_0$  rounds of training ;

```

As it can be seen in the above subsections IV-B to IV-D, our solution consists of three cascaded sub-solutions, where the output of one sub-solution is the input of the next one. We summarize its operation in Algorithm 1.

V. EXPERIMENTS AND RESULTS

A. EXPERIMENT SETUP

We consider an FL system with a single BS/EC server and $N_U = 50$ UAVs within its coverage area. The UAVs' computing frequencies $\gamma_i, \forall i = 1, \dots, N_U$, are randomly sampled from the range $[10^6, 10^8]$ Hz. To filter out stragglers, we set $\nu = 1.5$ [38], [39]. Also, ρ is set to -5 and ρ_0 to 10. These values were determined numerically through experimentation. We assume that $\kappa = 7 \times 10^4$ CPU cycles, $P_i = 0.28$ Watt, $\forall i = 1, \dots, N_U$. Additionally, we have set up a network configuration wherein the model weights are transmitted from the

FL client to the server with a total delay ranging between 20 ms and 200 ms, using a total bandwidth of 10 MHz. We test our solution's effectiveness on an image classification FL task, where the dataset is split into a training set and a testing set at ratios 80% and 20%, respectively. We run our defense approach on two popular datasets, namely MNIST and CIFAR-10, where MNIST consists of black and white handwritten digits of 60,000 training samples with size 28×28 pixels [43], while CIFAR-10 is composed of colored pictures with size 32×32 pixels, including 50,000 training images and 10,000 testing images [44]. Both datasets are divided into 10 classes, for numbers from "0" to "9" in MNIST, and for object and animal classes in CIFAR-10. We assume that each UAV can have a maximum number of $|M_r| = 1300$ samples. To take into account dataset non-IIDness in our experiment, we use different distribution methods. To build the first data distribution, called "Distribution 1", we use random sampling with the notion of " $k\%$ non-IID", where k measures the degree of non-IIDness in percentage. Accordingly, by setting $k = 80$, 80% of the dataset is dispersed randomly between the UAV FL clients while the remaining 20% is equally distributed among them. This data distribution strategy leads to unbalanced local datasets. For the second distribution, called "Distribution 2", we intentionally spread the data such that each UAV has data issued from $N_c = 2$ classes only. To be noted that both distributions 1 and 2 reflect two different strategies of data non-IIDness. This distributions design task is executed only once, and the resulting local datasets will be used in the remaining of this section.

In our MNIST experiment, we adopt at each FL participant the LeNet-5 convolutional neural network (CNN) architecture [45], while we adopt the VGG11 architecture for CIFAR-10-related system [46]. The latter is composed of eight CNN layers and three fully connected layers. We set the mini-batch size to 64, the number of local iterations for each round $T_e = 10$, and we use the vanilla stochastic gradient descent optimizer. Moreover, we select the minimum number of participants $P_r = 5, r_0 = 200$ as the number of FL rounds, 0.01 as the learning rate, and the cross-entropy loss as the loss function. Finally, the hyperparameter ϵ for DBSCAN clustering is set to $\epsilon = 0.02$ in the MNIST-related system, and to $\epsilon = 0.06$ for the CIFAR-10 related system.

Here, we evaluate our solution's performance in two different settings. In the first setting, we assume statistical and system heterogeneous environments using different data distributions, where stragglers and dropouts can exist, but no attackers, i.e., the ratio of attackers is $f = 0\%$. In the second setting, we consider $f > 0\%$, with different targeted and untargeted attack scenarios. The experimented targeted and untargeted attacks are chosen as the label flipping attack and additive noise attack, respectively, which are detailed as follows:

- *Label flipping attack:* It attempts to modify the behavior of an ML model on a subset of data samples while retaining the primary model performance over the whole testing dataset. Specifically, attackers attempt to compel the model to identify samples with the label "5"

as images with label “3”, as shown in Fig. 1 (Top-left rectangle).

- *Additive noise attack*: During this attack, a malicious participant contributes with Gaussian noise in its local model updates in order to degrade the global model’s accuracy and prevent its convergence, as illustrated in Fig. 1 (Top-middle rectangle).

The performance metrics evaluated in our system are the FL accuracy in the absence of malicious participants ($f = 0\%$), while under attacks ($f > 0\%$), we evaluate the attack success rate, denoted as ASR, and the false positive (FP) and false negative (FN) ratios of detected attackers. In a targeted attack scenario, ASR_{ta} is defined as the ratio of the number of altered dataset samples (i.e., those with their label flipped from ‘5’ to ‘3’) incorrectly classified (i.e., classified to anything but ‘5’) to the total number of samples within the dataset. It can be written as

$$ASR_{ta} = |S_{ta}|/|S|, \quad (16)$$

where $|S_{ta}|$ (resp. $|S|$) is the size of incorrectly classified data from the attacked class dataset (resp. size of dataset). Similarly, ASR_{ua} for an untargeted attack can be expressed by

$$ASR_{ua} = |A_1 - A_2|/A_1, \quad (17)$$

where A_1 and A_2 are the measured accuracy values under the scenarios of no attacks and an untargeted attack, respectively. Moreover, FP denotes the ratio of honest participants wrongfully identified as attackers, while FN refers to the ratio of non-detected attackers.


B. RESULTS AND DISCUSSION

1) SETTING 1 ($f = 0$)

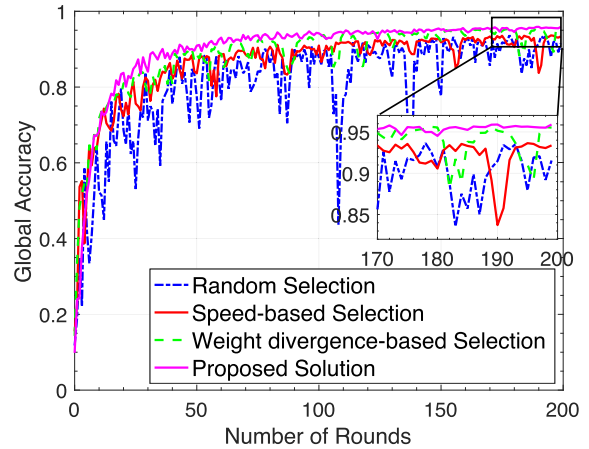
We evaluate the global accuracy, denoted A_g , of our solution, and compare it to those of four baseline methods. The first benchmark is the centralized learning approach, which is a model training technique characterized by the collection and processing of all data in a single server. The second baseline is the random selection scheme of participants, denoted “Random Selection”, which operates conventionally with FL. The third baseline, named “Speed-based Selection”, chooses the fastest participants in terms of data processing time. The fourth approach called “Weight divergence-based Selection”, defines the weight divergence metric as the distance between the weights of the local and global models, calculated as

$$\hat{w}_{i,r} = \left| \frac{w_{i,r} - w_{r-1}}{w_{r-1}} \right|, \quad \forall i = 1, \dots, N_U, \quad (18)$$

and describes the update drift in non-IID datasets conditions [16], [47]. Specifically, a large weight divergence value implies that the client’s data has a high non-IIDness degree compared to other participants, thus it should take part in more upcoming rounds to reduce this gap and maximize its accuracy. Moreover, selecting participants with larger weight divergence values would prevent bias, which is a result of the global model “drifting” to a client’s local optimizer [48].



(a) Distribution 1



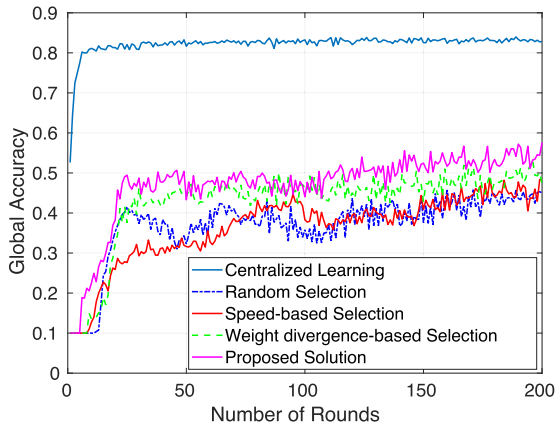
(b) Distribution 2

FIGURE 4. Global accuracy vs. FL rounds under different data distributions (MNIST dataset).

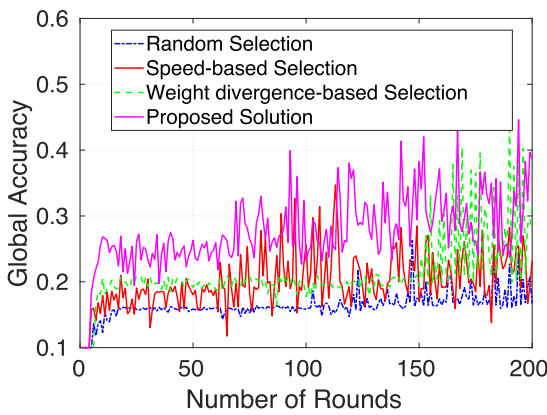
In Figs. 4(a) and 5(a), we present the global accuracy of the aforementioned participants selection methods, under Distribution 1 scenario, and using the MNIST and CIFAR-10 datasets, respectively. Accordingly, centralized learning achieves the best accuracy performance of 98.9% for the MNIST dataset and 82.4% for the CIFAR-10 dataset after 50 epochs.⁵ In contrast, the remaining techniques achieve performances ranging between 90% (random selection) and 96% (proposed solution) for MNIST, while for CIFAR-10 their performances are between 45.03% (random selection) and 57.5% (proposed solution). The findings indicate that the convergence and performances of FL, even when optimized, are below those of centralized learning, as a result of data non-IIDness.⁶ In any case, the proposed solution demonstrates its superiority over the FL baselines.

⁵Here, the number of rounds represents the number of epochs for centralized learning.

⁶Given that centralized learning is computationally heavy and due to its very high performance gap with FL techniques, we decide to drop it in the remaining experiments.



(a) Distribution 1



(b) Distribution 2

FIGURE 5. Global accuracy vs. FL rounds under different data distributions (CIFAR-10 dataset).

TABLE 3. Comparison of Participants Selection Methods in Terms of Accuracy, Dropout Ratio, and Average Round Time (MNIST Dataset)

Distribution	Selection Approach	A_g	τ_a	χ_d
Distribution 1	Weight divergence-based	95.90 %	37.13	0.15
	Speed-based	95.28 %	1.82	0.11
	Random	93.18%	41.85	0.15
	Proposed solution	97.02%	7.31	0.05
Distribution 2	Weight divergence-based	94.12 %	36.82	0.17
	Speed-based	93.37 %	2.10	0.10
	Random	91.54%	38.35	0.18
	Proposed solution	95.52%	5.52	0.037

The superior performance of our solution is mainly due to the careful participant selection that guarantees a small number of dropouts and a small round time, which is not the case with most baselines. Indeed, through the results of Tables 3 and 4 (Distribution 1), we can see that our solution has the lowest dropout ratio χ_d , and the second lowest average round time τ_a . “Speed-based Selection” achieves the best τ_a performance since it always selects participants with low processing times, regardless of their dropout risk. The proposed solution offers a good trade-off between reliability and convergence speed while maximizing global accuracy.

TABLE 4. Comparison of Participants Selection Methods in Terms of Accuracy, Dropout Ratio, and Average Round Time (CIFAR-10 Dataset)

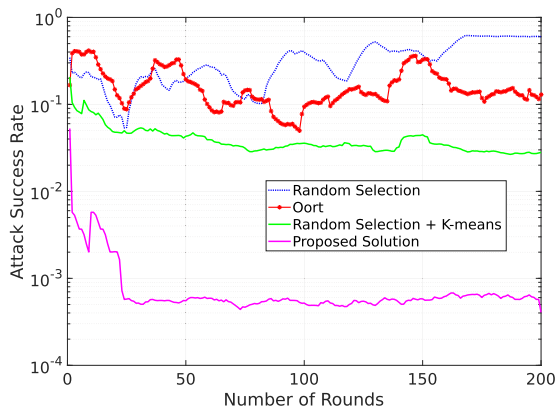


Distribution	Selection Approach	A_g	τ_a	χ_d
Distribution 1	Weight divergence-based	48.16%	32.70	0.14
	Speed-based	46.34%	2.32	0.1
	Random	45.03%	39.89	0.17
	Proposed solution	57.5%	6.55	0.06
Distribution 2	Weight divergence-based	31.71%	39.4	0.17
	Speed-based	23.17%	1.70	0.12
	Random	16.70%	40.34	0.16
	Proposed solution	38.80%	7.65	0.05

In Figs. 4(b) and 5(b), we extend the performance evaluation to Distribution 2. As noticed, the proposed solution is robust to highly unbalanced data distribution and achieves high accuracy, compared to baselines. Also, it converges the fastest. Similarly to Distribution 1, the dropout ratio and average round time are the lowest and second lowest ones as shown in Tables 3 and 4 (Distribution 2).


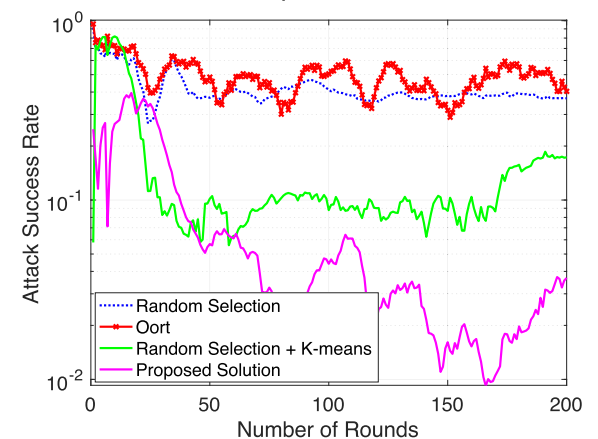
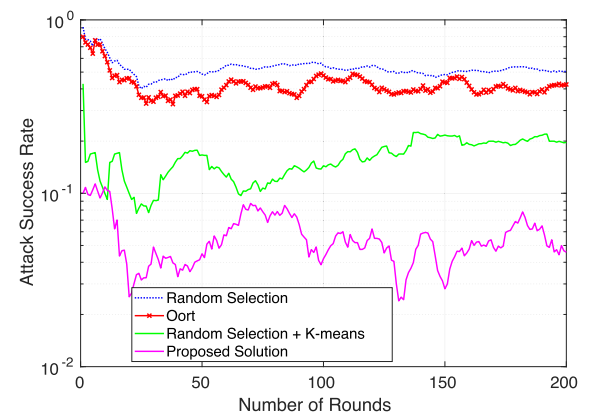
2) SETTING 2 ($f > 0$)

In this setting, and under the data assumption of Distribution 1, we compare our solution to three baselines. First, we use the previously defined “Random Selection” where malicious attackers filtering is in place. Then, we define “Random Selection + K-means”, which adds to “Random Selection” a K-means [49] based clustering defense mechanism to detect attackers when (15) is applied. The number of clusters is set to $C = 2$, in order to distinguish between the non-malicious and malicious UAVs. Finally, the third baseline is “Oort” [11], which relies on controlling participation to FL by limiting the number of participation rounds per UAV to 20. Experiments have been conducted with attackers ratios (among available $N_U = 50$ UAVs) $f \in \{10\%, 20\%, 33\%\}$.

Performance under untargeted attacks: Under additive noise attacks (untargeted attacks), we evaluate in Figs. 6 and 7 the ASR_{ua} performances of the proposed solution and baselines, given different f ratios, and for the MNIST and CIFAR-10 systems, respectively. For any f , the proposed solution (pink lines) achieves the lowest, and thus the best, performance in terms of attack success rate. Looking into the subfigures in Fig. 6 (resp. Fig. 7), we notice that our solution succeeds in keeping the ASR_{ua} below 1% (resp. 10%) for most of the time with MNIST (resp. CIFAR-10). In contrast, all baselines’ performances are worse than that of our method, for any f . Specifically, “Random Selection” (blue lines) and “Oort” (red lines) achieve the worst ASR_{ua} . Indeed, since “Oort” eliminates participants through rounds, at some point, it can improve the ASR performance. However, since this method does not clearly distinguish between non-malicious and malicious UAVs, over time, the UAV pool from which participants are selected is shrinking, thus leading to degraded performances. In the meanwhile, “Random Selection + K-means” demonstrates better results than the previous baselines, but still worse than our solution. Moreover, as f increases, this method demonstrates less robustness. The poor performance of this method is due to the high FP


 (a) $f=10\%$

 (b) $f=20\%$

 (c) $f=33\%$
FIGURE 6. Attack success rate vs. FL rounds under untargeted attack (MNIST dataset).

and FN of attackers' detection using K-means, thus causing the non-detection of malicious UAVs and the elimination of non-malicious UAVs during training. In Tables 5 and 6, we see that as f increases, FP and FN deteriorate for the benchmarks, hence leading to worse results. Finally, the obtained ASR_{ua} results for MNIST are better than those for CIFAR-10, for any method. Indeed, the CIFAR-10-related task is more complex, thus a stronger data non-IIDness is observed, leading to the cluster-based selective aggregation being less efficient.


 (a) $f=10\%$

 (b) $f=20\%$

 (c) $f=33\%$
FIGURE 7. Attack success rate vs. FL rounds under untargeted attack (CIFAR-10 dataset).

Performance under targeted attacks: Under a targeted attack such as flipping labels of a target class, the accuracy for that specific class is affected, especially when f is high. In Figs. 8 and 9, we illustrate the performance, in terms of ASR_{ta} , of the different participant selection methods, for the MNIST and CIFAR-10-related datasets, respectively. For the

TABLE 5. Comparison of Participant Selection Methods in Terms of FN, FP, and global/class Accuracy (MNIST Dataset)

Attack Type	f (%)	Selection Method	FN (%)	FP (%)	Acc. (%)	
No attack	0	Rand. Sel. + K-means	0	66	65.65	
		Oort	0	98	96.03	
		Proposed solution	0	2	97.02	
Untargeted Attack	10	Rand. Sel. + K-means	0	30	77.65	
		Oort	92	96	10.01	
		Proposed solution	0	0.5	97.01	
	20	Rand. Sel. + K-means	50	10	55.32	
		Oort	88	97	9.74	
		Proposed solution	0	0.1	97	
	33	Rand. Sel. + K-means	70	33	32.82	
		Oort	90	95	5.66	
		Proposed solution	0	3	96.8	
Targeted Attack	10	Rand. Sel. + K-means	0	50	68.67 92.15	
		Oort	91	94	10.78 85.02	
		Proposed solution	0	0.1	97 98.27	
	20	Rand. Sel. + K-means	11	33	60.01 85.75	
		Oort	90	92	9.35 64.22	
		Proposed solution	0.2	0.5	96.97 98.22	
	33	Rand. Sel. + K-means	37	30	55.37 83.56	
		Oort	96	98	6.78 45.85	
		Proposed solution	0.2	1	96.86 98.19	
	Targeted + Untargeted Attack	10	Rand. Sel. + K-means	0	40	66.20 94.25
			Oort	90	95	9.22 66.80
			Proposed solution	0	0.1	97 98.29
20		Rand. Sel. + K-means	30	10	58.90 89.89	
		Oort	95	93	8.07 66.23	
		Proposed solution	0.1	0.4	96.98 98.25	
33		Rand. Sel. + K-means	50	30	53.89 90.24	
		Oort	98	95	5.87 50.20	
		Proposed solution	0.5	2	96.84 98.21	

TABLE 6. Comparison of Participant Selection Methods in Terms of FN, FP, and global/class Accuracy (CIFAR-10 Dataset)

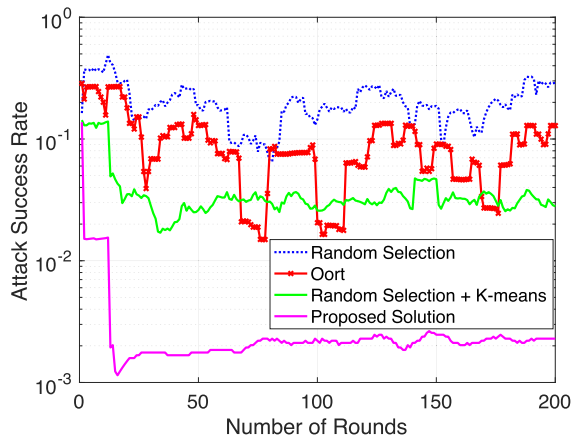
Attack Type	f (%)	Selection Method	FN (%)	FP (%)	Acc. (%)
No attack	0	Rand. Sel. + K-means	0	71	45.03
		Oort	0	97.3	55.17
		Proposed solution	0	5	57.5
Untargeted Attack	10	Rand. Sel. + K-means	45	33	47.60
		Oort	93	95	8.01
		Proposed solution	0.2	0.1	56.82
	20	Rand. Sel. + K-means	48	36	42.32
		Oort	85	98	6.23
		Proposed solution	0.5	1	54.90
	33	Rand. Sel. + K-means	70	33	39.82
		Oort	91	96	5.22
		Proposed solution	2	3	54.51
Targeted Attack	10	Rand. Sel. + K-means	3	48	45.67 89.02
		Oort	90	95	9.57 57.3
		Proposed solution	0.1	0.3	55.98 90.59
	20	Rand. Sel. + K-means	33	10	43.53 86.5
		Oort	93	96	7.75 53.65
		Proposed solution	2	1.1	55.45 90.32
	33	Rand. Sel. + K-means	52	22	40.65 83.62
		Oort	95	98	6.56 45.22
		Proposed solution	3	2	54.89 90.21
Targeted + Untargeted Attack	10	Rand. Sel. + K-means	2	43	45.01 88.02
		Oort	89	97	6.33 53.63
		Proposed solution	0.1	0.2	55.80 91.32
	20	Rand. Sel. + K-means	39	20	43.20 87.89
		Oort	95	93	5.88 51.20
		Proposed solution	0.4	1.3	55.22 90.75
	33	Rand. Sel. + K-means	65	22	40.01 85.24
		Oort	99	92	5.55 48.51
		Proposed solution	2	3	54.65 90.33

MNIST dataset, our method demonstrates a stable (i.e., robust) behavior for any f value at performance around 0.2%. Meanwhile, the benchmarks' performances are above 2%. For the CIFAR-10 dataset, the proposed algorithm achieves ASR_{ta} below 1% for $f \leq 10\%$ and around 7% for $f > 10\%$. Meanwhile, the other approaches realize ASR_{ta} above 10% for any f .


Unlike the benchmarks, the proposed solution is able to realize correct UAV clustering and efficient malicious participants elimination, even for high f and complex datasets such as the CIFAR-10. On the other hand, "Random Selection + K-means" demonstrated better performances than the other benchmarks but less than the proposed method and with instability through rounds. This is due to the randomness in

K-means initialization that affects FP and eliminates non-malicious participants, as shown in Tables 5 and 6. Finally, "Oort" demonstrates the worst results, similar to "Random Selection", since it does not effectively eliminate malicious UAVs.

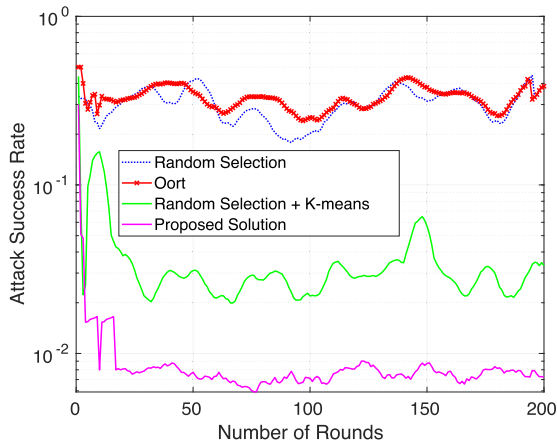
In Tables 5 and 6, we compare the results of the above solutions in terms of FN, FP, and accuracy, under different attack scenarios, where targeted and untargeted attacks refer to the events where the ratio of UAVs f is totally dedicated to that specific attack, while the targeted+untargeted scenario means that only $f/2$ are dedicated for untargeted attacks and the other $f/2$ to targeted ones. To be noted that the accuracy refers to the global accuracy when no or only untargeted attacks are performed, while in scenarios of targeted attack and targeted+untargeted attack, two accuracy values are presented



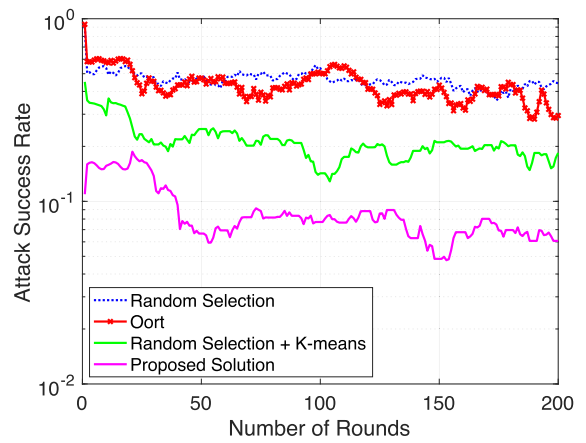
(a) $f=10\%$



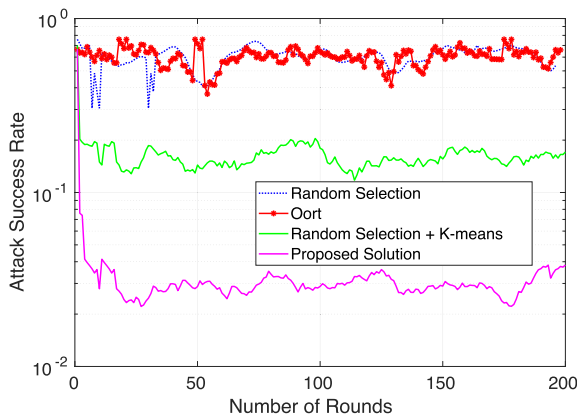
(a) $f=10\%$



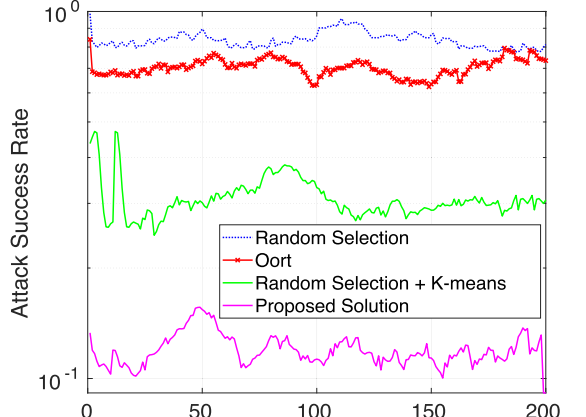
(b) $f=20\%$



(b) $f=20\%$



(c) $f=33\%$



(c) $f=33\%$

FIGURE 8. Attack success rate vs. FL rounds under targeted attack (MNIST dataset).

FIGURE 9. Attack success rate vs. FL rounds under targeted attack (CIFAR-10 dataset).

for each f value, where the first (top) refers to the global accuracy, while the second (bottom) to the accuracy of the targeted class with the label flipping attack only, i.e., the classes of the number “5” and “cat” in MNIST and CIFAR-10 experiments, respectively. Accordingly, the use of DBSCAN in our solution is more efficient in identifying clusters of non-malicious and malicious participants than “Random selection

+ K-means” and “Oort” in any scenario. Indeed, due to the K-means’ sensitivity to outliers and noise, it tends to provide high FP and FN ratios, in particular at high f values, and thus the system’s accuracy is impacted. In contrast, DBSCAN demonstrates stable behavior regardless of the type of attackers, since it is a density-based algorithm. “Oort” provides




FIGURE 10. Global accuracy vs. ρ and ρ_0 (MNIST dataset, Distribution 1).

the worst results as it is not using an efficient mechanism to mitigate attacks. The impact of both untargeted and targeted attacks have a small impact on the global accuracy of the proposed solution, while “Random Selection + K-means” performance significantly degrades with f , compared to the no attack scenario. Indeed, this method is inefficient when no attackers are present, thus underperforming, i.e., it is mis-clustering the UAVs. However, when $f > 0$, K-means operates well and thus is able to achieve better results. In contrast, “Oort” performs very well when no attacks occur, but its accuracy drastically drops as soon as $f > 0$. This is mainly due to its incapability to correctly assess the maliciousness of participants.

C. DISCUSSION ON HYPERPARAMETERS SENSITIVITY

To enforce a criterion for the minimal contribution from the participants, we established the reputation threshold ρ and its maximal value ρ_0 as settable coefficients. Participants having reputations below ρ are not considered in subsequent FL rounds. In order to understand the impact of the values of ρ and ρ_0 , we illustrate in Fig. 10 the achieved global accuracy performance as a function of ρ and ρ_0 , and under the distribution 1 setting. We notice that the choice of ρ and ρ_0 values have a significant impact on the global accuracy A_g . For instance, a small ρ degrades A_g for any ρ_0 , which is the result of a very slow process of unreliable UAVs elimination. Another negative impact on A_g is the choice of close ρ and ρ_0 values, e.g., $\rho = 9$ and $\rho_0 = 10$ respectively. In this case, several UAVs will be eliminated, seen as with very low reliability scores or falsely being identified as dropouts. In contrast, a large ρ_0 value may lead to clients’ over-selection. For instance, if a UAV begins as a reliable client before turning unreliable after a certain number of rounds, it will still have a high reliability score, which will have a detrimental effect on the accuracy performance. Therefore, a tradeoff between ρ and ρ_0 is necessary to improve A_g .

VI. CONCLUSION

In this article, we have investigated the problem of participant selection in FL under the presence of misbehaving UAVs, by considering straggler, dropout, and malicious participants’ profiles. To reach efficient and reliable FL, we proposed a novel three-step selection method that sequentially eliminates misbehaving participants, starting with stragglers, then dropouts, and finally malicious ones. Through experiments, we demonstrate the superiority of our solution in terms of accuracy, convergence, average round time, and attack success rate, compared to baseline methods, and considering different data distribution among UAVs and under the absence or presence of malicious nodes. As a future work, we will extend our investigation on FL to other types of attacks and heterogeneous systems, where participants are of different types and operate under varying constraints.

REFERENCES

- [1] B. Brik, A. Ksentini, and M. Bouaziz, “Federated learning for UAVs-enabled wireless networks: Use cases, challenges, and open problems,” *IEEE Access*, vol. 8, pp. 53841–53849, 2020.
- [2] O. Ghdiri, W. Jaafar, S. Alfattani, J. B. Abderrazak, and H. Yanikomeroglu, “Offline and online UAV-enabled data collection in time-constrained IoT networks,” *IEEE Trans. Green Commun. Netw.*, vol. 5, no. 4, pp. 1918–1933, Dec. 2021.
- [3] J. Liu, Z. Wu, J. Liu, and X. Tu, “Distributed location-aware task offloading in multi-UAVs enabled edge computing,” *IEEE Access*, vol. 10, pp. 72416–72428, 2022.
- [4] G. Zhu, D. Liu, Y. Du, C. You, J. Zhang, and K. Huang, “Toward an intelligent edge: Wireless communication meets machine learning,” *IEEE Commun. Mag.*, vol. 58, no. 1, pp. 19–25, Jan. 2020.
- [5] W. Wen, Z. Chen, H. H. Yang, W. Xia, and T. Q. S. Quek, “Joint scheduling and resource allocation for hierarchical federated edge learning,” *IEEE Trans. Wireless Commun.*, vol. 21, no. 8, pp. 5857–5872, Aug. 2022.
- [6] D. Yin, Y. Chen, R. Kannan, and P. Bartlett, “Byzantine-robust distributed learning: Towards optimal statistical rates,” in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 5650–5659.
- [7] X. Pan, M. Zhang, D. Wu, Q. Xiao, S. Ji, and M. Yang, “Justinian’s GAVERNOR: Robust distributed learning with gradient aggregation agent,” in *Proc. USENIX Conf. Secur. Symp.*, 2020, Art. no. 93.
- [8] P. Blanchard, E. M. El Mhamdi, R. Guerraoui, and J. Stainer, “Machine learning with adversaries: Byzantine tolerant gradient descent,” in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 118–128.
- [9] Y. Lu, X. Huang, Y. Dai, S. Maharjan, and Y. Zhang, “Blockchain and federated learning for privacy-preserved data sharing in industrial IoT,” *IEEE Trans. Ind. Informat.*, vol. 16, no. 6, pp. 4177–4186, Jun. 2020.
- [10] S. Abdulrahman, H. Tout, A. Mourad, and C. Talhi, “FedMCCS: Multi-criteria client selection model for optimal IoT federated learning,” *IEEE Internet Things J.*, vol. 8, no. 6, pp. 4723–4735, Mar. 2021.
- [11] F. Lai, X. Zhu, H. V. Madhyastha, and M. Chowdhury, “Oort: Efficient federated learning via guided participant selection,” in *Proc. USENIX Symp. Operating Syst. Des. Implementation*, 2021, pp. 19–35.
- [12] Z. Chai et al., “TiFL: A tier-based federated learning system,” in *Proc. Int. Symp. High-Perform. Parallel Distrib. Comput.*, 2020, pp. 125–136.
- [13] S. Kim, W. Shin, S. Jang, H. Song, and S.-Y. Yun, “FedRN: Exploiting k-reliable neighbors towards robust federated learning,” in *Proc. ACM Int. Conf. Inf. Knowl. Manage.*, 2022, pp. 972–981.
- [14] T. Nishio and R. Yonetani, “Client selection for federated learning with heterogeneous resources in mobile edge,” in *Proc. IEEE Int. Conf. Commun.*, 2019, pp. 1–7.
- [15] S. D. Apostolidis, P. C. Kapoutsis, A. C. Kapoutsis, and E. B. Kosmatopoulos, “Cooperative multi-UAV coverage mission planning platform for remote sensing applications,” *Auton. Robots*, vol. 46, no. 2, pp. 373–400, 2022.
- [16] W. Xia, W. Wen, K.-K. Wong, T. Q. Quek, J. Zhang, and H. Zhu, “Federated-learning-based client scheduling for low-latency wireless communications,” *IEEE Wireless Commun.*, vol. 28, no. 2, pp. 32–38, Apr. 2021.

- [17] Y. Fu, C. Li, F. R. Yu, T. H. Luan, and Y. Zhang, "A selective federated reinforcement learning strategy for autonomous driving," *IEEE Trans. Intell. Transp. Syst.*, vol. 24, no. 2, pp. 1655–1668, Feb. 2023.
- [18] Y. Deng et al., "AUCTION: Automated and quality-aware client selection framework for efficient federated learning," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 8, pp. 1996–2009, Aug. 2022.
- [19] Y. Chen, Y. Ning, M. Slawski, and H. Rangwala, "Asynchronous online federated learning for edge devices with non-IID data," in *Proc. IEEE Int. Conf. Big Data*, 2020, pp. 15–24.
- [20] W. Wu, L. He, W. Lin, R. Mao, C. Maple, and S. Jarvis, "SAFA: A semi-asynchronous protocol for fast federated learning with low overhead," *IEEE Trans. Comput.*, vol. 70, no. 5, pp. 655–668, May 2021.
- [21] A. Imteaj and M. H. Amini, "FedAR: Activity and resource-aware federated learning model for distributed mobile robots," in *Proc. IEEE Int. Conf. Mach. Learn. Appl.*, 2020, pp. 1153–1160.
- [22] L. Barbieri, S. Savazzi, M. Brambilla, and M. Nicoli, "Decentralized federated learning for extended sensing in 6G connected vehicles," *Veh. Commun.*, vol. 33, 2022, Art. no. 100396.
- [23] W. Wu, L. He, W. Lin, and R. Mao, "Accelerating federated learning over reliability-agnostic clients in mobile edge computing systems," *IEEE Trans. Parallel Distrib. Syst.*, vol. 32, no. 7, pp. 1539–1551, Jul. 2021.
- [24] E. Bagdasaryan, A. Veit, Y. Hua, D. Estrin, and V. Shmatikov, "How to backdoor federated learning," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2020, pp. 2938–2948.
- [25] R. A. Mallah, D. López, G. Badu-Marfo, and B. Farooq, "Untargeted poisoning attack detection in federated learning via behavior attestation," *IEEE Access*, vol. 11, pp. 125064–125079, 2023.
- [26] S. Awan, B. Luo, and F. Li, "CONTRA: Defending against poisoning attacks in federated learning," in *Proc. Eur. Symp. Res. Comput. Secur.*, 2021, pp. 455–475.
- [27] S. Shen, S. Tople, and P. Saxena, "Auror: Defending against poisoning attacks in collaborative deep learning systems," in *Proc. ACM Annu. Conf. Comput. Secur. Appl.*, 2016, pp. 508–519.
- [28] Z. Song, H. Sun, H. H. Yang, X. Wang, Y. Zhang, and T. Q. S. Quek, "Reputation-based federated learning for secure wireless networks," *IEEE Internet Things J.*, vol. 9, no. 2, pp. 1212–1226, Jan. 2022.
- [29] X. Tan, W. C. Ng, W. Y. B. Lim, Z. Xiong, D. Niyato, and H. Yu, "Reputation-aware federated learning client selection based on stochastic integer programming," *IEEE Trans. Big Data*, pp. 1–12, 2022, doi: [10.1109/TBDATA.2022.3191332](https://doi.org/10.1109/TBDATA.2022.3191332).
- [30] J. Kang, Z. Xiong, D. Niyato, S. Xie, and J. Zhang, "Incentive mechanism for reliable federated learning: A joint optimization approach to combining reputation and contract theory," *IEEE Internet Things J.*, vol. 6, no. 6, pp. 10700–10714, Dec. 2019.
- [31] D. Li, W. E. Wong, W. Wang, Y. Yao, and M. Chau, "Detection and mitigation of label-flipping attacks in federated learning systems with KPCA and K-means," in *Proc. IEEE Int. Conf. Depend. Syst. Appl.*, 2021, pp. 551–559.
- [32] V. Parmar, S. S. Sarwar, Z. Li, H.-H. S. Lee, B. D. Salvo, and M. Suri, "Memory-oriented design-space exploration of edge-AI hardware for XR applications," in *IEEE Micro*, vol. 43, no. 6, pp. 40–49, Nov./Dec. 2023, doi: [10.1109/MM.2023.3321249](https://doi.org/10.1109/MM.2023.3321249).
- [33] S. Ray, X. Sun, N. Tremelling, M. Gordiyenko, and P. Kinget, "How tiny can analog filterbank features be made for ultra-low-power on-device keyword spotting?," 2023, *arXiv:2304.08541*.
- [34] L. Li, H. Xiong, Z. Guo, J. Wang, and C.-Z. Xu, "SmartPC: Hierarchical pace control in real-time federated learning system," in *Proc. IEEE Real-Time Syst. Symp.*, 2019, pp. 406–418.
- [35] A. Al-Hourani and K. Gomez, "Modeling cellular-to-UAV path-loss for suburban environments," *IEEE Wireless Commun. Lett.*, vol. 7, no. 1, pp. 82–85, Feb. 2018.
- [36] T. Dantzig, *Number: The Language of Science*. New York, NY, USA: Plume Book, 2007.
- [37] H. Wu et al., "Probabilistic automatic outlier detection for surface air quality measurements from the China national environmental monitoring network," *Adv. Atmospheric Sci.*, vol. 35, Dec. pp. 1522–1532, 2018.
- [38] R. McGill, J. W. Tukey, and W. A. Larsen, "Variations of box plots," *Amer. Statistician*, vol. 32, no. 1, pp. 12–16, 1978.
- [39] K. D. Buch, "Decision based non-linear filtering using interquartile range estimator for Gaussian signals," in *Proc. Annu. IEEE India Conf.*, 2014, pp. 1–5.
- [40] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. Y. Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2017, pp. 1273–1282.
- [41] Y. Lin, S. Han, H. Mao, Y. Wang, and W. J. Dally, "Deep gradient compression: Reducing the communication bandwidth for distributed training," 2018, *arXiv:1712.01887*.
- [42] F. Sattler, K. Müller, and W. Samek, "Clustered federated learning: Model-agnostic distributed multi-task optimization under privacy constraints," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 32, no. 8, pp. 3710–3722, Aug. 2021.
- [43] L. Deng, "The MNIST database of handwritten digit images for machine learning research," *IEEE Signal Process. Mag.*, vol. 29, no. 6, pp. 141–142, Nov. 2012.
- [44] A. Krizhevsky, "Learning multiple layers of features from tiny images," pp. 32–33, 2009. [Online]. Available: <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>
- [45] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.
- [46] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. 3rd Int. Conf. Learn. Representations*, 2015, pp. 1–14.
- [47] M. Kamp et al., "Efficient decentralized deep learning by dynamic model averaging," in *Proc. Eur. Conf. Mach. Learn. Knowl. Discov. Databases*, 2019, pp. 393–409.
- [48] T. Huang, W. Lin, W. Wu, L. He, K. Li, and A. Y. Zomaya, "An efficiency-boosting client selection scheme for federated learning with fairness guarantee," *IEEE Trans. Parallel Distrib. Syst.*, vol. 32, no. 7, pp. 1552–1564, Jul. 2021.
- [49] S. Lloyd, "Least squares quantization in PCM," *IEEE Trans. Inf. Theory*, vol. 28, no. 2, pp. 129–137, Mar. 1982.



YOUSRA CHERIGUENE received the M.Sc degree in computer science from University of Laghouat, Algeria, in 2020. She is currently a Ph.D student with the Informatics and Mathematics Laboratory (LIM), the University of Laghouat. Her research interests include mobile edge computing, unmanned aerial vehicles, and federated learning.



WAEEL JAAFAR (Senior Member, IEEE) received the master's and Ph.D. degrees in electrical engineering from Polytechnique Montreal, QC, Canada, in 2014. He has been an Associate Professor with the Software and IT Engineering Department, École de Technologie Supérieure, University of Quebec, Montreal, since September 2022. Between 2019 and 2022, he was with the Systems and Computer Engineering Department, Carleton University, Ottawa, ON, Canada, as an NSERC Postdoctoral Fellow. From 2014 to 2018, he has

pursued a career in the telecommunications industry, where he has been involved in designing telecom solutions for projects across Canada and abroad. During his career, he was a Visiting Researcher within Khalifa University, Abu Dhabi, UAE in 2019, Keio University, Tokyo, Japan, in 2013, and UQAM, Montreal, in 2007. His research interests include wireless communications, integrated terrestrial and non-terrestrial networks, resource allocation, edge caching and computing, and machine learning for communications and networks. He was the recipient of several prestigious grants including NSERC Alexandre-Graham Bell scholarship, FRQNT internship scholarship, and best paper awards at IEEE ICC 2021 and ISCC 2023.



HALIM YANIKOMEROGLU (Fellow, IEEE) received the B.Sc. degree in electrical and electronics engineering from Middle East Technical University, Ankara, Turkey, in 1990 and the M.A.Sc. degree in electrical engineering and the Ph.D. degree in electrical and computer engineering from the University of Toronto, Toronto, ON, Canada, in 1992 and 1998, respectively. Since 1998, he has been with the Department of Systems and Computer Engineering, Carleton University, Ottawa, ON,

where he is currently a Chancellor's Professor. He has given more than 110 invited seminars, keynotes, panel talks, and tutorials in the last five years. He has supervised or hosted more than 150 Postgraduate researchers with the laboratory, Carleton University. His extensive collaborative research with industry resulted in 39 granted patents. His research interests include wireless communications and networks, with a special emphasis on non-terrestrial networks (NTN) in recent years. Dr. Yanikomeroglu is a Fellow of the Engineering Institute of Canada (EIC) and the Canadian Academy of Engineering (CAE). He is also a Member of the IEEE ComSoc Governance Council, IEEE ComSoc GIMS, IEEE ComSoc Conference Council, and IEEE PIMRC Steering Committee. He was the recipient of the several awards for research, teaching, and service, including the IEEE ComSoc Satellite and Space Communications TC Recognition Award, in 2023, the IEEE ComSoc Fred W. Ellersick Prize, in 2021, the IEEE VTS Stuart Meyer Memorial Award, in 2020, and the IEEE ComSoc Wireless Communications TC Recognition Award, in 2018. He was also the recipient of the Best Paper Award from the IEEE Competition on Non-Terrestrial Networks for 5G and 6G, in 2022 (grand prize), IEEE ICC 2021, and IEEE WISEE 2021 and 2022. He was the General Chair and the Technical Program Chair for several IEEE conferences. He is the Chair of the Steering Committee of IEEE's Flagship Wireless Event and the Wireless Communications and Networking Conference (WCNC). He was also on the editorial boards for various IEEE periodicals. He is a Distinguished Speaker of the IEEE Communications Society and the IEEE Vehicular Technology Society, and an Expert Panelist of the Council of Canadian Academies (CCA|CAC). He has also served on editorial boards of various IEEE periodicals.



CHAKER ABDELAZIZ KERRACHE received the M.Sc. degree in computer science degree from the University of Laghouat, Laghouat, Algeria, in 2012 and the Ph.D. degree in computer science degree from the University of Laghouat, in 2017. He is currently an Associate Professor with the Department of Computer Science and the Head with Informatics and Mathematics Laboratory (LIM), the University of Laghouat. His research interests include trust and risk management, secure multi-hop communications, vehicular networks, named

data networking, and UAVs.