Analysis Study of Participant Selection Methods in Federated Learning

Wafa Bouras[1, 2], Mohamed Benouis[1,2, 4], Kamel Eddine Heraguemi[1, 2, 3], Bouderah Brahim[1, 2], and Samir Akhrouf[1, 2]

[1]LIM Laboratory, SoukAhras university
[2]University of M'sila, PO Box166 Ichebilia, 28000 M'sila, Algeria
[3]National school of artificial intelligence, Sidi AbdAllah, Algiers, Algeria
[4]Chair for Human-Centered AI, University of Augsburg, Universitätsstraße 6a, 86159 Augsburg, Germany

***Abstract*—To the best of current knowledge, the performance of federated learning predominantly depends on the efficiency of the aggregation server scheme utilized to consolidate model parameters received from distributed local devices. However, in practical scenarios, the global server often faces single-point failures due to four major issues: 1) variations in data distribution settings, such as independent identical distribution (IID) or non-independent identical distribution; 2) communication overhead; 3) limitations in hardware and resource storage availability; and 4) diverse participant participation behaviors. To address the latter concern, limited research has endeavored to establish a correlation between these heterogeneous settings and federated learning performance by analyzing different aspects of participant behavior. Inspired by the absence of a definitive verdict regarding the relationship between the global server and participant behavior, this paper investigates the aspect of participant selection methods and conducts a detailed comparative study among various participant selection methods.**

***Index Terms*—security, federated learning, participant selection, machine learning**

I. Introduction

Federated Learning (FL) facilitates collaborative model creation across diverse participants, from individual devices to organizational entities, under the central administration of a server. FL's decentralization of training data, distributed across participant devices, is a key feature substantiated in prior research [1]. At FL's outset, participant selection is a pivotal milestone where the server curates a subset for model training in each iteration. Among strategies, the Fed-avg algorithm [2] is notable for its randomized selection, yet it may pose accuracy challenges. Recognizing participant selection's critical role, various methodological approaches have emerged, aiming to optimize selection efficacy [3]. These methods consider factors like hardware heterogeneity and data characteristics. However, there's a need for a comprehensive overview paper on FL participant selection for quick comprehension in this field. Although these efforts have made some progress in participant selection methods, to the best of our knowledge, there is no comprehensive and comparative study on the recent participant selection methods applied to FL. This work studies three recently proposed selection methods tested and validated on WESAD datasets. This work aims to study the adaptability and performance of such methods on FL performance without any specialization in the original dataset case study. We conduct preprocessing and a set of handcuffed feature methods to clean and process all the data. In this paper, we make the following contributions:

1) We present a comprehensive study on three popular participant selection methods in the FL context.
2) We conduct a comparative analysis of the methods by using a recurrent neural network (RNN).
3) To facilitate quick reproduction and further research on the topic of the participant selection method, we release the code for this work, [1], which contains the implementations of all the methods in this study.

The remainder of this paper is organized as follows: In the next section, we describe the three participant selection methods in detail, namely oort, pieces, and active-federated learning. We then describe the experiments, including the datasets, a comparison of the study methods' performances, and the parameters' impact. Finally, we present the concluding remarks.

II. Background and related works

A. Related Works

Federated learning has gained significant traction in the realm of healthcare systems. A critical stage within this framework is participant selection, where the process of choosing suitable participants is central to the success of the federated learning endeavor. Numerous studies and research efforts have been dedicated to participant selection in federated learning, particularly in its application to healthcare settings. Initially, the primary focus of these methods was to tackle the inherent challenge of device heterogeneity. This challenge arises due to the diverse range of devices and equipment used within healthcare systems, and addressing it is crucial to ensuring the effectiveness and reliability of federated learning in this domain. [4] was earlier research on participant selection. This research methodology primarily revolved around the optimization of participant selection to maximize the utilization of available resources for model training while adhering to predefined deadlines. in [5], A novel and impartial clustered sampling strategy for the selection of participants is introduced, aimed at reducing the variance in weights for participant aggregation while also enhancing the likelihood of selecting participants

[1]https://github.com/wafabou/participantSelectioninFederatedLearningAComparativeStudy

with distinctive data distributions. The research authors introduce two distinct clustered sampling methodologies: clustered sampling grounded in sample size considerations and clustered sampling predicated on similarity metrics. Through a series of rigorous experiments, it has been empirically demonstrated that the adoption of clustered sampling techniques results in superior and expedited convergence in the context of model training. [6]introduced the problem of vertically federated participant selection; the researchers conducted an investigation into the process of participant selection within the framework of vertically federated learning. In this endeavor, they framed the problem from a mutual information-based perspective. To address this challenge, they introduced an innovative mutual information estimator named VF-MINE, which leverages Fagin's algorithm. To facilitate participant selection, the researchers devised a group testing-based framework, VF-PS, built upon the foundation of VF-MINE. Furthermore, the authors conducted a comprehensive examination of the security properties inherent in their proposed methodologies, demonstrating that their methods yielded remarkable improvements across orders of magnitude in terms of efficacy and performance. In [7], a novel method enhances federated learning in mobile settings with a reputation-aware user selection scheme, significantly improving aggregated model accuracy. In [8], a new participant selection scheme for Federated Reinforcement Learning (FRL) is introduced based on rewards. This scheme accelerates learning and reduces agent requirements, marking a significant advancement in learning paradigm efficiency. In [9] tackled the critical concern of worker selection in the context of federated learning within mobile networks, with the overarching objective of enhancing the reliability of the federated learning process. To achieve this, they devised a reputation-based scheme aimed at the discerning selection of dependable and trustworthy workers. In order to establish an efficient and secure mechanism for reputation management, the researchers employed a multi-weight subjective logic model to compute workers' reputations. Furthermore, to ensure tamper resistance and non-repudiation while maintaining a decentralized framework, they leveraged a consortium blockchain for the management of reputations.

[10] presents a novel approach for joint sampling and device-to-device (D2D) offloading optimization in Federated Learning (FedL), with new convergence bounds and a Graph Convolutional Network (GCN)-based algorithm for sampling strategy determination. while [11] focuses on minimizing wall-clock convergence time in Federated Learning (FL) by developing an optimal participant sampling strategy. It introduces a novel convergence bound, a non-convex optimization problem formulation, and efficient algorithms for parameter estimation and solution approximation. [12], the authors propose FedPNS, a node selection strategy for Federated Learning (FL) with non-i.i.d. datasets. It prioritizes nodes based on a probabilistic approach guided by optimal aggregation results. This approach improves model convergence by excluding detrimental local updates and comparing local gradients with the global gradient. The authors conducted a theoretical analysis showing FedPNS's superior convergence rate over conventional FedAvg.

In the context of federated learning, the process of aggregating local updates from participating participants is iterative and continues until a predefined accuracy threshold is achieved. However, various factors related to device capabilities and data characteristics can influence this iterative convergence process.

B. Background

1) Participant Selection Definition: Participant selection constitutes a cornerstone of federated learning methodology, enabling the deliberate identification of participants for model training. This process encompasses random assignment or the evaluation of specific participant attributes to optimize training effectiveness and precision, often referred to as sampling. In the realm of federated learning, participants denoted as workers or users are devices operating under the supervision of a central server and actively participating in the training phase to refine the model. Various resources, including those cited such as [7], [13], and [14], leverage coordinators to execute participant selection as an auxiliary function alongside server operations.

2) Characteristics Considered in Participant Selection:

- **Device Heterogeneity**: Device heterogeneity encompasses variations in hardware capabilities, such as battery life, storage capacity, and network connectivity. Devices with limited resources, often referred to as "stragglers," can negatively impact federated learning by degrading model accuracy, impeding training efficiency, and increasing the dropout rate. To address these challenges, several algorithms focus on mitigating system-related issues and assessing individual device capabilities.
- **Data Heterogeneity**: Data heterogeneity pertains to the diversity in datasets residing on participant devices, categorized as follows:
  - *Independently and Identically Distributed (IID) Data*: In federated learning, IID data assumes that each device's data follows an independent and identical distribution, albeit with allowances for slight deviations from complete independence and identical distribution.
  - *Non-IID Data*: Non-IID data deviates from the IID characteristics and can lead to biased training, causing over-fitting to specific data sources. Ensuring IID-like characteristics on participant devices is crucial in federated learning, and techniques such as differential privacy and specialized algorithms are employed to address non-IID data and achieve fairness and accuracy.
- **Dynamic Behavior**: Dynamic behavior in federated learning refers to the ability of participants to join or leave the network, even replacing other participants in the event of dropouts. This feature enhances system scalability and adaptability, ensuring a more robust and flexible training process.

- **Security and Privacy**: Security and privacy considerations are paramount in federated learning. The selection of malicious participants poses significant risks, potentially leading to vulnerabilities and manipulation. Various algorithms focus on identifying and excluding malicious participants from the learning process to mitigate these risks.
- **Fairness**: Fairness mandates that all participants should have an equal opportunity to participate in the training process, with no participants being unfairly excluded. Several algorithms, such as random sampling or active learning-based selection, are designed to ensure fairness in participant selection

3) The Objective of Participant Selection: The primary objective of participant selection in the federated learning context is to establish an optimal process that ensures high accuracy while minimizing training time. The selection strategy aims to achieve various goals, including:

1) Maximizing the number of participants in each iteration to enhance data diversity.
2) Prioritizing the selection of trustworthy machines to safeguard the integrity of the process.
3) Efficiently reducing training time for faster convergence
4) Balancing privacy preservation with learning accuracy
5) Mitigating the impact of non-IID data.
6) Enhancing the overall reliability of performance.

Efficient participant selection strategies play a vital role in achieving these objectives, thereby ensuring the success of federated learning while upholding data privacy and model accuracy.

4) Participant Selection Categories: Participant selection strategies in federated learning are categorized based on their approach and objectives [15]:

- **Random Selection**: Participants are chosen randomly in each iteration. While simple, this approach may result in the frequent selection of straggling participants, impacting accuracy and training time and introducing bias into models.
- **Performance-Based Selection**: Relies on past performance metrics to select participants. The goal is to choose participants according to their resource capabilities while excluding those who have not completed tasks promptly. Reputation-based algorithms often implement this strategy.
- **Data-Based Selection**: Addresses data heterogeneity issues to prevent training on non-IID data. While fairness may not always apply, other techniques should handle non-IID data effectively.
- **Security-Based Selection**: Focuses on preventing the inclusion of malicious participants, as their presence can lead to attacks and manipulate accuracy. Reputation-based mechanisms are commonly used.
- **Group-Based Selection**: Participants are divided into groups based on criteria like geographic location or data type. Selection is based on group characteristics.
- **Characteristics-Based Selection**: Involves selecting participants based on specific features that influence results, like weight divergence or probability allocation.

Alternative categorization methods include:

- **Synchronous Selection**: Concurrent or coordinated selection of individuals for model training.
- **Asynchronous Selection**: Independent selection of participants occurring at varying intervals without strict coordination.
- **Hybrid Selection (Synchronous and Asynchronous)**: Integrates elements of both synchronous and asynchronous methodologies in participant selection.

Each category addresses different aspects of federated learning, aiming to optimize the training process while preserving data privacy, ensuring accuracy, and mitigating vulnerabilities. The choice of strategy depends on specific system requirements and goals.

III. Experimental Study

A. Approaches to Participant Selection Employed in Our Study

1) Synchronous method: Oort: The primary aim of this method is to optimize processing time and accuracy by integrating an Oort executor for participant selection in federated learning (FL). Oort operates as the FL coordinator, efficiently identifying and selecting participants based on predefined criteria. It interfaces with the FL driver for seamless participant management. [13]This operational sequence involves job submission, participant selection, execution, and aggregation. The coordinator communicates selection criteria to Oort, who selects participants who meet these criteria. Participants compute results independently, and the coordinator aggregates their updates. After aggregating updates from sufficient participants, the coordinator triggers the next training round, repeating these steps iteratively. Periodic federated testing assesses accuracy attainment.

2) Asynchronous method: Pisces: Pisces: The primary objective is to maximize participant involvement and effectively utilize struggling participants to enhance accuracy within a shorter timeframe. This is achieved through a coordinated orchestration involving three main components: the participant manager, coordinator, and executor. The configuration phase involves participant manager initialization based on the training plan, while the registration phase records participant meta-information as participants join. In the orchestration phase, the coordinator interacts with the participant manager and executor. The background communication aspect encompasses server-participant communication, local update storage, and model validation. During each iteration, the coordinator consults the participant manager for model aggregation necessity and, if required, delegates the task to the executor. The coordinator also seeks guidance from the participant manager to identify idle participants who respond with a "no" or provide an eligible participant selection plan. [14]

3) Synchronous and Asynchronous Methods: Active Federated Learning: Active Federated Learning (AFL) represents a pioneering approach to user cohort selection in the realm of Federated Learning (FL), characterized by its dynamic adaptation to both the model's state and the data residing on individual participants. AFL distinguishes itself by incorporating reliability considerations into the selection process, thus mitigating bias by intensifying training attempts on less reliable users. Notably, the loss value function employed in AFL is versatile, rendering it applicable to a wide range of supervised problems. Investigating the application of AFL in conjunction with more intricate models presents an intriguing avenue for future research. [16]

B. Experiments

1) Dataset Description: The WESAD dataset, detailed in [17], is designed for affective state recognition. It includes physiological signals from various sources like blood volume pulse, electrocardiogram, etc. We focus on binary stress classification, dividing signals into 700-sample segments with 50% overlap. We extract features like mean, variance, and nonlinear characteristics from 121,813 segments to maximize inter-subject correlation and minimize intra-subject variance.

2) Classifier architecture: Given the dataset's affiliation with the time series domain, we employed recurrent neural networks (RNN) as a classifier. This RNN architecture comprises 70 memory cells, a dropout layer, and a subsequent flattened layer utilized for the generation of binary output classes (i.e., stress versus non-stress). The training process employs the cross-entropy loss function and utilizes a stochastic gradient descent (SGD) optimizer with a learning rate of β=0.005. A comprehensive summary of the hyperparameters of the RNN model is presented in Table I. The model is specifically configured to focus on training with the WESAD dataset, emphasizing binary class recognition for stress versus no-stress scenarios.

We set up various FL settings, including participant number communication rounds, each participant's data site, and the RNN model's hyperparameter, performed over the WESAD database. As the main purpose of this study is to study the effect of the participant selection method in the FL context, we provide a study of the influence of data distribution on the global performance model. In IID, the data holds the same label destruction among all participant participants; however, in the case of no-IID, which is mostly the particle case in real-world scenarios, the labels are not distributed equally among participants.

The parameters that work with this experiment are represented in **??**

models	RNN
optimizer	SGD
batch size	32
Learning rate	0.005
Momentum	0.9
epochs	1

TABLE I: parameters on the experiments.

3) Discussion: The visual representation under scrutiny presents a line chart that portrays the comparative performance of three participant selection algorithms—namely, "Oort," "Pisces," and "Active Federated Learning." This analysis was conducted within the context of Plato [18], an open-source Federated Learning (FL) platform renowned for its ability to abstract away the complexities of the underlying machine learning infrastructure by providing user-friendly APIs.

To assess the efficacy of these selection algorithms, we executed a comprehensive experiment using Plato. This framework has been widely acknowledged for its utility in facilitating the development and evaluation of FL methodologies. Through Plato's seamless integration of various FL components and its convenient API, we could design and conduct a rigorous analysis of the aforementioned algorithms.

The dataset employed for this study was WESAD, which was used for this comparative study evaluation. The chart effectively communicates how these three participant selection methods performed concerning the given performance metric. Notably, the chart reveals a convergence of all three methods to a common performance level of 77.22%, as depicted in Figure 1.

Intriguingly, "Pisces" exhibits a noteworthy early advantage in the initial rounds of the experiment but ultimately converges with the other two methods as the experiment progresses, a trend that becomes apparent after approximately 10 rounds. Additionally, "Oort" displays some intriguing irregularities or fluctuations within its performance line, suggesting potential variability or distinct behavior relative to the other methods. A more in-depth analysis would be necessary to elucidate the underlying causes of these fluctuations.

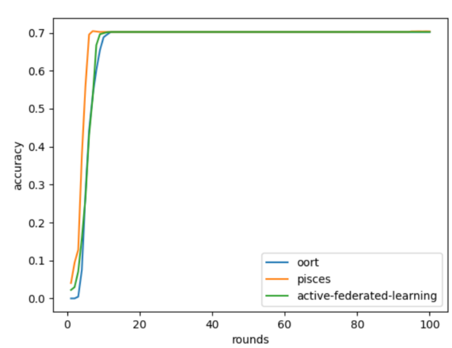

Fig. 1: Accuracy performance versus round communication

Figure 2 visually represents the temporal aspects of participant selection methods in the Federated Learning (FL) context. Among these methods, "Active Federated Learning" emerges as the swiftest regarding the time required for participant selection. It is closely followed by "Pisces," which demonstrates temporal performance nearly equivalent to that of "Active Federated Learning."

In contrast, "Oort" exhibits a relatively prolonged selection duration when compared to the aforementioned methods. This finding underscores a notable temporal divergence between "Oort" and the other two methods, "Active Federated Learning" and "Pisces." Such variations in time consumption for participant selection can have significant implications for the overall efficiency and responsiveness of the FL framework.

It is important to note that these temporal observations are instrumental in understanding the practical implications of the participant selection methods within the FL framework, offering valuable insights into the trade-offs between time efficiency and other performance metrics. Further analysis and contextualization of these temporal disparities are warranted to comprehensively assess the suitability of these methods in specific FL scenarios.

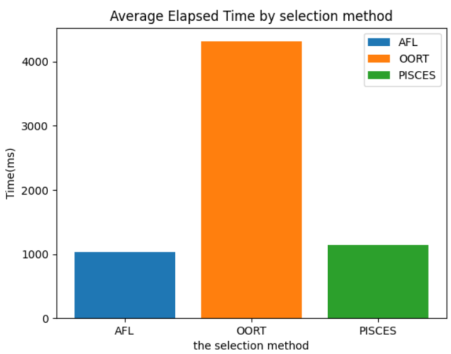

Fig. 2: Accuracy performance versus time

Figure 3 illuminates the pivotal relationship between the number of participants and the resultant accuracy within the context of our study. Three distinct participant selection methods—namely, "Active Federated Learning," "Pisces," and "Oort"—are under scrutiny in this analysis. The impact of varying participant counts on these methods' accuracy is particularly interesting.

"Active Federated Learning" exhibits remarkable steadiness in its accuracy performance, regardless of the fluctuations in the number of participants. This consistent performance suggests a degree of robustness within the method, making it a reliable choice in scenarios characterized by varying participant participation.

In contrast, "Pisces" presents an interesting pattern whereby accuracy fluctuates in response to changes in the number of participants. This oscillatory behavior indicates a certain level of sensitivity to participant count variations, which merits further investigation. "Oort," on the other hand, displays pronounced fluctuations in accuracy as the number of participants varies. These substantial changes in accuracy highlight a significant degree of volatility within the "Oort" method when confronted with shifts in participant participation.

Furthermore, the accuracy values reveal distinct performance disparities among the three methods. "Active Federated Learning" consistently maintains an accuracy level of 70.22%, indicating its reliability in achieving a predetermined performance threshold. Conversely, "Pisces" demonstrates a notably higher accuracy rate, while "Oort" exhibits a considerably lower accuracy level.

These observations underline the intricate interplay between the number of participants and the accuracy achieved by the participant selection methods. The findings underscore the importance of methodological considerations when selecting an appropriate approach within the Federated Learning framework, as accuracy fluctuations can have significant implications for the overall success of FL deployments in real-world scenarios. Further exploration and contextualization of these accuracy trends are warranted to gain a more comprehensive understanding of the behavior and suitability of these methods in varying FL contexts.

Overall, our exploration of participant selection methods in the context of federated learning underscores the pivotal role these methods play in shaping the efficacy and performance of FL systems. The prominence of "Active Federated Learning" as the method of choice highlights the need for careful consideration when selecting participant selection strategies, thereby enhancing the prospects for the successful deployment of FL in various real-world applications.

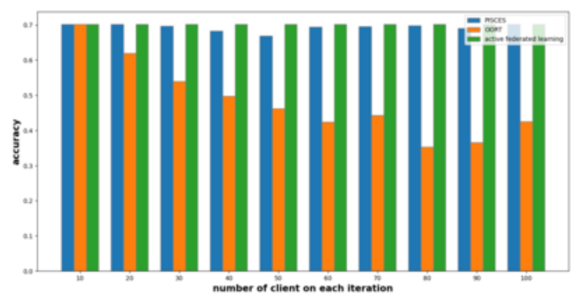

Fig. 3: Accuracy performance versus number of participants

Table 2 illustrates a comparative analysis of similarities and distinctions evident in the experiences associated with participant selection methods.

IV. Conclusion

In conclusion, this study has provided a comprehensive analysis of participant selection methods in the context of Federated Learning (FL). Through the utilization of the Plato framework and the "WESAD" dataset, the study elucidated the intricate dynamics governing these methods within FL frameworks. Notably, "Active Federated Learning" emerged as the most effective participant selection method, demonstrating superiority over alternatives such as "Pisces" and highlighting the limitations of the "Oort" approach. The findings underscore the critical role that participant selection methods play in shaping the performance, efficiency, and security of FL systems. Moving forward, it is imperative for future research to prioritize fortifying the security aspects of these methods and optimizing them to foster a more streamlined and efficient FL ecosystem. This study contributes valuable insights that will inform ongoing efforts to advance the field of federated learning and its practical applications.

TABLE II: Comparison between participant selection methods,

	Oort	***Pisces***	***Active federated learning***
Approach	Centralised	Centralised	Dynamic
Privacy-Preserving	no	no	no
Heterogeneity Consideration	yes	yes	yes
Dynamic Selection	no	no	yes
Fairness and Bias Considerations	yes	yes	yes
Communication Efficiency	yes	yes	yes
Security Considerations	no	no	no
Aggregation Strategy Alignment	no	no	
Adaptive Strategies	yes	yes	yes
Compatibility with Learning Tasks	yes	yes	yes
Accuracy with IID	77%	77%	77%
Accuracy with NONIID	55%	71%	77%

REFERENCES

[1] Kairouz, Peter, et al. ”Advances and open problems in federated learning.” Foundations and trends® in machine learning 14.1–2 (2021): 1-210.

[2] McMahan, Brendan, et al. ”Communication-efficient learning of deep networks from decentralized data.” Artificial intelligence and statistics. PMLR, 2017.

[3] Fu, Lei, et al. ”Client selection in federated learning: Principles, challenges, and opportunities.” IEEE Internet of Things Journal (2023).

[4] Nishio, Takayuki, and Ryo Yonetani. ”Client selection for federated learning with heterogeneous resources in mobile edge.” ICC 2019-2019 IEEE international conference on communications (ICC). IEEE, 2019.

[5] Fraboni, Yann, et al. ”Clustered sampling: Low-variance and improved representativity for clients selection in federated learning.” International Conference on Machine Learning. PMLR, 2021.

[6] Jiang, Jiawei, et al. ”Vf-ps: How to select important participants in vertical federated learning, efficiently and securely?.” Advances in Neural Information Processing Systems 35 (2022): 2088-2101.

[7] Wang, Yuwei, and Burak Kantarci. ”A novel reputation-aware client selection scheme for federated learning within mobile environments.” 2020 IEEE 25th International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD). IEEE, 2020.

[8] Lee, Woonghee. ”Reward-based participant selection for improving federated reinforcement learning.” ICT Express 9.5 (2023): 803-808.

[9] Kang, Jiawen, et al. ”Reliable federated learning for mobile networks.” IEEE Wireless Communications 27.2 (2020): 72-80.

[10] Wang, Su, et al. ”Device sampling for heterogeneous federated learning: Theory, algorithms, and implementation.” IEEE INFOCOM 2021-IEEE Conference on Computer Communications. IEEE, 2021.

[11] Luo, Bing, et al. ”Tackling system and statistical heterogeneity for federated learning with adaptive client sampling.” IEEE INFOCOM 2022-IEEE conference on computer communications. IEEE, 2022.

[12] Wu, Hongda, and Ping Wang. ”Node selection toward faster convergence for federated learning on non-iid data.” IEEE Transactions on Network Science and Engineering 9.5 (2022): 3099-3111.

[13] Lai, Fan, et al. ”Oort: Efficient federated learning via guided participant selection.” 15th USENIX Symposium on Operating Systems Design and Implementation (OSDI 21). 2021.

[14] Jiang, Zhifeng, et al. ”Pisces: Efficient federated learning via guided asynchronous training.” Proceedings of the 13th Symposium on Cloud Computing. 2022.

[15] Soltani, Behnaz, et al. ”A survey on participant selection for federated learning in mobile networks.” Proceedings of the 17th ACM Workshop on Mobility in the Evolving Internet Architecture. 2022.

[16] Goetz, Jack, et al. ”Active federated learning.” arXiv preprint arXiv:1909.12641 (2019).

[17] Schmidt, Philip, et al. ”Introducing wesad, a multimodal dataset for wearable stress and affect detection.” Proceedings of the 20th ACM international conference on multimodal interaction. 2018.

[18] Li, Baochun, et al. ”Plato: An open-source research framework for production federated learning.” Proceedings of the ACM Turing Award Celebration Conference-China 2023. 2023.