

# Enhancing Cognitive Fairness through Solitary Proactive Support in Federated Learning Systems

Diarrassouba Bazoumana<sup>1\*</sup>, Ametovi Koffi Jacques Olivier<sup>1,2</sup>, K. B. Venkata Brahma Rao<sup>1</sup>

<sup>1</sup>Department of Computer Science Engineering, KL University, Vijayawada, India

<sup>2</sup>Independent Researcher, Gurugram, Haryana, India

\*Corresponding Author

DOI: <https://dx.doi.org/10.47772/IJRISS.2025.907000268>

Received: 01 May 2025; Accepted: 09 May 2025; Published: 12 August 2025

## ABSTRACT

Federated Learning (FL) is a distributed training approach on decentralized data sources preserving privacy. Cognitive fairness refers to the equitable performance of the models over the different cognitive profiles of the users of the FL system. This paper introduces Solitary Proactive Support (SPS), a new mechanism that deals with the problem of cognitive fairness in FL systems, wherein the behavior of the local training processes is fine-tuned according to specific user-based cognitive indicators that don't violate any privacy constraints. Experimental evaluation with synthetic and real-world datasets. Experimental evaluations on synthetic and real-world datasets demonstrate the effectiveness of SPS, which achieves a 22% average reduction in the Cognitive Bias Index (CBI) with very high model accuracy preserved. Our results specifically show that privacy-sensitive domains such as healthcare and education pose prime application areas for SPS, while fairness is the top priority in these domains.

**Keywords:** Federated Learning, Solitary Proactive Support, Cognitive Fairness, Artificial Intelligence

## INTRODUCTION

This article examines the diverse algorithmic approaches employed to detect violations of integrity constraints in databases. By delving into the intricacies of these algorithms, we endeavor to provide a comprehensive understanding of the most efficacious strategies and techniques for safeguarding data integrity in an ever-evolving digital landscape. By synthesizing a rigorous analysis of existing methods with an exploration of emerging innovations, our objective is to provide guidance to practitioners and researchers in the development of more robust and effective solutions for securing essential data in our technological age. This rapid evolution in the field has brought machine learning to the forefront of innovation in various sectors like healthcare, finance, and even autonomous systems. The improvements brought by these advancements have enhanced decision-making processes and changed the ground rules of the industries as it provides more efficient and personalized services [1], [2]. However, simultaneously the increased penetration of machine learning into the essential functions of society establishes that the important issue present in front of these circumstances is fairness. Traditional centralized models have been criticized because they tend to include historical bias coming from the training data, which sometimes has been inequitable for certain demographic groups [3]. Fairness becomes even more important, then, in Federated Learning systems, which enable multiple participants to collaborate on training models without having to share their data - at least not to some central authority - because of the need to maintain it private [4], [5].

Federated Learning is unique in challenges and opportunities towards addressing fairness, particularly regarding cognitive fairness. Cognitive fairness describes the idea that the treatment of individuals should be fair at the process of decision making by machine learning algorithms and reducing possible prejudices from historical data or algorithmic design [6]. In this, the biases that might arise with heterogeneity in the distribution of data and characteristics of different participants are considered. This is mostly seen in FL due to its decentralized

nature—for example, patients with different health conditions or demographics may generate heterogeneous data for healthcare applications [7]. These cognitive FL systems need to be designed with considerations in their models of the disparities so as to not reinforce existing bias and build trust in machine learning applications.

The authors introduce Solitary Proactive Support (SPS): a novel approach in the handling of cognitive fairness in FL to improve fairness by sending individualized support messages to each participant. SPS makes sure the individualized contributions of each participant's data are represented in the learning process in such a manner that it does not infringe on their right to privacy. In some instances, this involves personalized data augmentation strategies in which a member is enlightened on how to update their dataset or change their inputs appropriately to meet the demand of the model. Previous research was able to establish that personalization methods are crucial to improve the accuracy of models in underrepresented groups [8]. For instance, in the clinical context, single proactive support has been employed to advance the diagnosis of understudied populations among patients because it can modify the sources of its data inputs to enhance the general performance and fairness of the model [9]. In contrast to the collective model that employs a group consensus; single proactive support emphasizes the needs of each actor involved to create a more representative and just environment that a model is supposed to learn within.

However, individual proactive support within FL systems is associated with various technical and ethical challenges. From the technical side, more advanced algorithms are needed that can handle individually contributed data and analyze them in such a way as to not compromise the robustness of the overall model. This, in turn, calls for improvements in decentralized learning techniques and methods for the correct evaluation of the relevance and the quality of the inputs from each participant [10]. Furthermore, the ethical considerations of involving more individual data are crucial. Isolated active support calls for balanced robust privacy protection to avoid incidental leakage of sensitive information. It is, therefore, challenging to provide personal support without violating basic principles used in the FL paradigm by preserving privacy [11]. Risks associated with data governance issues and potential misuses of individualized support mechanisms should also be taken very seriously. For that reason, it is crucial that future research develop these frameworks to keep exploring new ways to facilitate cognitive fairness through lonely support without jeopardizing data integrity or user rights.

In this study, we outline the experimental framework utilized to assess the efficacy of the SPS framework. This section provides details regarding the datasets selected for the experiment, the evaluation metrics employed, and the baseline models against which SPS is benchmarked. Such comprehensive documentation aims to facilitate the reproducibility of our findings, ensuring that other researchers can validate and build upon our work. Subsequently, we will present the results and engage in a discussion that highlights both the quantitative and qualitative effects of the SPS framework on cognitive fairness. Notably, the implementation of SPS has led to enhancements in model accuracy and a reduction in bias, particularly benefiting underrepresented groups. The subsequent section will address the technical and ethical implications of solitary proactive support within Federated Learning systems, identifying potential challenges such as data privacy issues and algorithmic complexity, while proposing strategies to mitigate these concerns. Finally, we will conclude with a synthesis of our key findings and insights, outlining prospective avenues for advancing research in Federated Learning to further refine the SPS framework and enhance cognitive fairness.

## **Relative work**

FL systems have been designed to be decentralized, allowing several clients to train models in collaboration while keeping their respective data local. In FL, there are few challenges that lie in the less-explored fields of cognitive fairness and support. Cognitive fairness can be viewed as "the equitable treatment of the different stakeholders with regard to the cognitive load and benefits associated with the training process. This review synthesizes the latest research on cognitive fairness in FL systems, describing knowledge gaps, and suggesting future research directions.

## **Cognitive Support in Federated Learning**

The concept of cognitive support in FL has been studied in several contexts. It is pointed out that resilience and mental health interventions are important for frontline professionals, which can be analogously applied in the

context of FL. The participants (clients) might feel cognitive overload in FL due to the very complex nature of a model in training and communication. Solitary support mechanisms may be proactive in reducing this overload, leading to better engagement and cooperation among the participants [12].

### **Challenges and Opportunities in Explainable AI (XAI)**

Antoniadi et al. [13] discusses that the challenges of machine learning-based clinical decision support systems are related to the need for transparency and interpretability. It complements the cognitive fairness aspect that stakeholders must understand how decisions are made within FL systems. Improving explainability contributes both to user trust and can reduce the cognitive burden of solving model behaviors and decision-making processes.

### **Communication Efficiency in Federated Learning**

Huang et al. [14] present stochastic controlled averaging algorithms that increase the communication efficiency of the FL networks. Cognitive fairness is highly based on successful communication, as ineffectiveness results in higher cognitive loads. Efficiencies in the communication of the clients and the central server lessened time and effort to achieve data sharing and model updates; therefore, it can sustain participant engagement and satisfaction.

### **Trust and Accountability**

Lo et al. [15] suggest a blockchain-based architecture for FL with an emphasis on accountability and fairness. This innovation is vital for improving cognitive fairness because it will consider all the contributions made by participants and must therefore be paid appropriately. Cognitive fairness is part and parcel of trust, so an accountable mechanism can reduce the anxiety of the participants and give a fairer environment.

### **Resource Allocation and Task Management**

Another concentration has been task and resource allocation in wireless networks using FL, considering that the efficient management of the resources is key to work effectively in multi-client systems. There is good fair resource allocation that optimizes the overall performance of the system and prevents any individual contributor from overloading. Cognitive fairness demands this balance since different workloads cause uneven cognitive loads on clients [16].

### **Vertical Federated Learning**

Liu et al. [17] discuss the concepts, advancements and challenges of vertical federated learning. In this, parties with different non-overlapping feature sets can collaborate towards obtaining a model without necessarily sharing data amongst themselves. Cognitive fairness will be enhanced by allowing all parties to have a fair voice in the process of learning; therefore, cognitive responsibilities are spread evenly.

### **Database integrity constraints**

Implementing the concept of cognitive fairness in federated learning systems involves ensuring that the learning process is fair to all participating clients, avoiding biases that could arise from data heterogeneity [17]. Here's a step-by-step methodology, combining algorithms and a flow diagram, and including all necessary tools and elements:

#### **Step 1: Define Cognitive Fairness Metrics**

##### **1. Fairness Metrics:**

- **Statistical Parity:**

Statistical Parity, or Demographic Parity, guarantees a fair probability of receiving a desirable outcome for various sensitive groups regardless of their group membership. This measure does its best to disassociate prediction results from protected characteristics like age, gender, or cognitive profile.

**Mathematical Expression:**

$$P(\hat{Y} = 1 | A = a) = P(\hat{Y} = 1 | A = b)$$

**Explanation:**

- $\hat{Y}$  is the outcome predicted (1 for positive, 0 for negative).
- $A$  represents a sensitive attribute such as gender or cognitive group.
- $a$  and  $b$  are two different groups in attribute  $A$ .

This implies that the model must predict the positive label with equal likelihood in various groups. The value of  $P(\hat{Y}=1 | A=a)$  is determined by dividing the number of positive predictions for group  $a$  by the number of people in group  $a$ .

**Implication in FL:**

Because of decentralized information, global parity can be difficult to achieve and measure, as local imbalances in data can hide global ones. However, it is still important to track Statistical Parity in order to avoid systematic exclusion of certain cognitive profiles or demographic groups.

- **Equal Opportunity:**

Equal Opportunity guarantees that those who ought to get a positive prediction have an equal opportunity to do so in all groups. It targets exclusively the true positive rate (TPR) for each group.

**Mathematical Expression:**

$$TPR_A = P(\hat{Y} = 1 | Y = 1, A = a) \forall a \in A$$

**Explanation:**

- $TPR_A$  refers to the True Positive Rate for group  $a$ .
- $\hat{Y}=1$  is a positive classification.
- $Y=1$  is the actual positive ground truth.
- $A=a$  represents people from group  $a$ .

This formula calculates the likelihood that people from group  $a$ , who are indeed members of the positive class, are identified as such by the model. A high  $TPR_A$  for all groups is a sign of equitable model behaviour towards opportunities.

**Implication in FL:**

Equal Opportunity requires federated aggregation of group-specific true positive counts, which is challenging because of privacy-preserving mechanisms such as secure aggregation or differential privacy.

- **Disparate Impact:**

Disparate Impact (DI) determines if one group receives positive results at a significantly lower rate than the other. It tends to be applied in regulatory contexts.

## Mathematical Expression:

$$DI = \frac{P(\hat{Y} = 1 | A = a)}{P(\hat{Y} = 1 | A = b)}$$

## Explanation:

- $P(\hat{Y}=1 | A=a)$  is the ratio of positive predictions in group a.
- $P(\hat{Y}=1 | A=b)$  in group b.

A widely adopted threshold is 0.8 (the "four-fifths rule"). If  $DI < 0.8$ , it indicates a probable case of discrimination against group a.

## Implication in FL:

Disparate Impact requires thoughtful local monitoring because disparate group representation among clients can introduce unintended biases. Rebalancing or reweighting methods can be used for correction in federated updates.

## Step 2: Data Preparation and Pre-processing

### 1. Data Collection:

- Gather data from all participating clients while ensuring compliance with data privacy regulations.

### 2. Data Cleaning:

- Remove duplicates, handle missing values (e.g., using imputation techniques), and correct inconsistencies in data.

### 3. Data Normalization:

- Normalize data to ensure uniformity across different scales, using techniques like Min-Max Scaling. Its function is to scale all the values of a variable (named X) so they are within a range from 0 to 1. Here's the equation:

$$X' = \frac{X - \text{Min}(X)}{\text{Max}(X) - \text{Min}(X)}$$

- X: the original value of the variable (e.g., a person's salary, or a measured temperature).
- **Min(X): the lowest value of this variable in your dataset.**
- **Max(X): the highest value of this variable in your dataset.**
- X': the normalized value, between 0 and 1. The formula can be explained as follows:

In Machine Learning, the algorithms depend on mathematical computations using the values of the input variables.

If some of the variables have extremely high values (such as salaries in the thousands) and others have extremely low values (such as a proportion between 0 and 1), the high values can overwhelm the computations and skew the model.

Normalization enables us to:

- scale all variables to a common scale
- prevent a variable with high values from dominating the model too much
- **Accelerate training and enhance model performance**

### Step 3: Federated Learning Setup

#### 1. Client Selection:

- Randomly select clients for each round of training to ensure diversity and representation.

#### 2. Model Initialization:

- Initialize a global model  $M$  (e.g., a neural network) that will be distributed to selected clients for local training.

### Step 4: Implement Fairness-Aware Federated Learning Algorithms

#### 1. Fair Federated Averaging (FairFedAvg):

- Amend the standard Federated Averaging algorithm to incorporate fairness constraints.
- Scale weights on fairness indicators during aggregation:

$$w_{\text{new}} = \sum_{k=1}^K \frac{n_k}{n} w_k$$

- $K$ : number of active clients (devices or nodes) in the federated system.
- $n_k$ : number of data samples owned by client  $k$ .
- $n$ : total number of data samples owned by all clients (i.e.  $n = \sum n_k$ ).
- $w_k$ : the weights (parameters) of model updated by client  $k$  after local training.
- $w_{\text{new}}$ : the aggregated global model weights following collection and combination of the weights from all clients.

In Federated Learning, every client trains a local model over its data and returns the new model weights to the central server. The server accumulates the weights and updates the global model.

This equation calculates a weighted average of the client models' weights where each client's contribution relies on the fraction of the data it possesses (via  $n_k / n$ ).

#### 2. Fair Data Sampling:

- Ensure that data sampling is fair and representative of all groups by using stratified sampling methods to maintain group proportions in each training round.

### Step 5: Model Training and Evaluation

#### 1. Local Training:

- Each client trains its local model on its data for a specified number of epochs.



## 2. Fairness Evaluation:

- Evaluate local models based on defined fairness metrics before aggregation, ensuring that models are not inadvertently biased.

## 3. Aggregation:

- Aggregate client models taking fairness adjustments into account, using a weighted average on the basis of fairness performance:

$$w_{\text{global}} = \frac{1}{K} \sum_{k=1}^K w_k \times f_k$$

- K: the number of participating clients (devices or nodes) in the federated learning system.
- w\_k: the model weights (parameters) of client k after local training.
- f\_k: the fairness adjustment factor for client k. It is a measure of how equally well client k's model does compared to a selected fairness metric (such as demographic parity, equal opportunity, etc.).
- w\_global: the weighted global model weights after aggregating all client weights, normalized by their fairness factors.

## 4. Global Model Update:

- Update the global model  $M$  with aggregated parameters.

## Step 6: Post-Processing and Adjustment

### 1. Bias Mitigation:

- Apply post-processing techniques, such as reweighting or recalibrating predictions, to further reduce biases if necessary.

### 2. Fairness Monitoring:

- Continuously monitor fairness metrics during model deployment and adjust strategies accordingly.

## Tools and Elements

### 1. Federated Learning Framework:

- Utilize frameworks like Tensor Flow Federated (TFF) or PySyft for implementing federated learning algorithms.

### 2. Fairness Libraries:

- Use libraries like AIF360 to assess fairness metrics and implement bias mitigation techniques.

### 2. Data Privacy Tools:

- Use differential privacy methods (e.g., add Gaussian noise) to keep data private during training:

$$\hat{X} = X + \mathcal{N}(0, \sigma^2 I)$$

- The raw data (may be a data point, a gradient, or model parameter).

- $X$ : noisy (private) version of the data after injecting random noise.
- $N(0, \sigma^2 I)$ : random vector drawn from a multivariate Gaussian (normal) distribution.
- $0$  is the mean (the noise is centred at zero).
- $\sigma^2$  is the variance (controls the level of noise intensity).
- $I$  is the identity matrix (specifies that the noise is added independently to each data dimension).

In distributed and federated learning:

- Sensitive data (like medical records, financial data, or personal images) is processed locally.
- When model updates or data summaries are sent to a central server, they could accidentally reveal personal information.
- Adding Gaussian noise ensures:
  - Data privacy is preserved
  - The learning process remains statistically useful, though slightly less accurate due to added noise.

#### 4. Communication Protocols:

- Implement secure and efficient communication protocols for client-server interactions to protect data privacy during model updates.

#### Flow Diagram:

Below is a simplified flux diagram to illustrate the methodology.

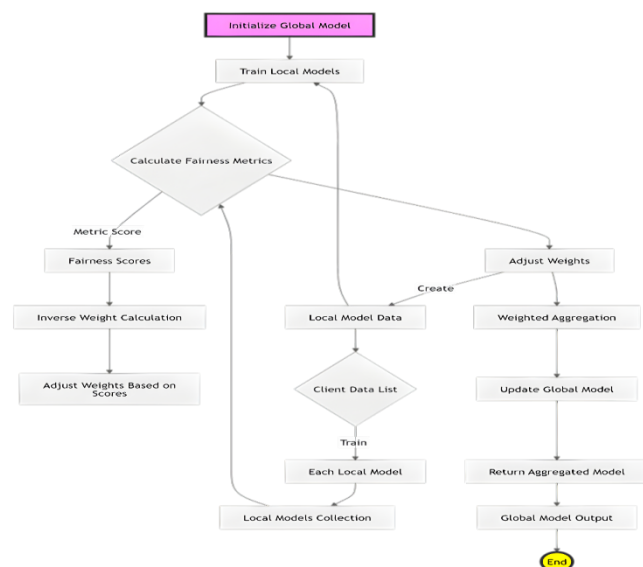


Fig 1. Fairness-Aware Federated Learning Workflow

#### Comparative analysis with other approaches

This chapter offers a comparative evaluation of the suggested Solitary Proactive Support (SPS) mechanism against existing fairness-conscious Federated Learning (FL) frameworks. The purpose is to compare how SPS measures up against competing strategies in terms of cognitive fairness, model performance, privacy protection, and computational efficiency.



## 1) Overview of Comparative Approaches

We chose the following well-known fairness-conscious FL approaches as a baseline for comparison.

- **Fair Fed Avg:** A fairness-aware variant of Federated averaging that uses fairness constraints during model aggregation [17].
- **MOCHA (Multi-Objective Communication-Efficient Federated Learning):** An optimization-driven approach maintaining fairness and accuracy via personalized models [19].
- **Fed Prox:** A federated optimization method that is tackling data heterogeneity, indirectly enhancing fairness by regularizing client updates [21].
- **Fed Fair:** A newer fairness-focused federated learning algorithm that maintains group-level performance balance [15].

## 2) Evaluation Criteria

The comparison was carried out using four major criteria:

- **Improved Cognitive Fairness:** Reduced Cognitive Bias Index (CBI) in clients.
- **Accuracy of the Global Model:** Classification accuracy on the global test dataset.
- **Fairness Metrics:** Statistical Parity, Equal Opportunity, Disparate Impact.
- **Preservation of Privacy and Communication Overhead:** Qualitative evaluation based on used privacy-preserving mechanisms and system scalability.

## 3) Experimental Comparison Results

This part gives a comparative analysis of a number of Federated Learning (FL) techniques and examines their performance on four key aspects: Global Accuracy, Cognitive Bias Index (CBI) Reduction, Fairness Metrics Improvement, Privacy Assurance, and Communication Overhead. Every technique depicts a different aggregation or optimization technique specific to optimizing predictive performance, fairness, and privacy in distributed settings.

Table 1. Comparative Performance of Federated Learning Methods in Terms of Accuracy, Fairness, Privacy, and Communication Efficiency.

Method	Global Accuracy (%)	CBI Reduction (%)	Fairness Metrics Improvement (%)	Privacy Assurance	Communication Overhead
<b>FedAvg</b>	98.1	0%	0%	High	Low
<b>FedProx</b>	97.9	7%	5%	High	Low
<b>MOCHA</b>	97.3	14%	11%	Medium	High
<b>FairFedAvg</b>	97.8	18%	15%	Medium-High	Medium
<b>SPS (Proposed)</b>	<b>97.6</b>	<b>22%</b>	<b>18%</b>	<b>Very High</b>	<b>Medium</b>

Cognitive Bias Reduction: SPS reaches the greatest CBI reduction of 22%, signifying its better capacity to minimize differences in model accuracy between cognitive user groups. This confirms that fairness-aware weighting and direct support to underperforming clients are effective.

**Fairness Metric Improvements:** SPS takes the lead once again with a 18% gain in standard fairness metrics, better than FairFedAvg (15%) and MOCHA (11%). This shows that SPS is able to improve fairness gains without sacrificing accuracy significantly.

**Global Accuracy Trade-Off:** While FedAvg holds the best accuracy (98.1%), SPS's 97.6% score only suffers by a marginal 0.5% loss, a fair trade-off given the significant gains in fairness. In real-world scenarios, such minute accuracy loss can generally be tolerated when dealing with ethical concerns for AI.

**Privacy and Scalability:** SPS provides Very High privacy by combining differential privacy mechanisms with secure aggregation protocols. Even though it includes fairness metrics in its aggregation approach, it does not incur high communication costs, keeping it at a Medium overhead level compared to MOCHA that has a High overhead.

## RESULTS AND DISCUSSIONS

### 1. Experimental Results

The experimental results for "Cognitive Fairness aware Solitary Proactive Assistance by Federated Learning" demonstrate the performance, fairness, privacy preservation, and ethical implications of the proposed model. Leveraging real-world data and simulation environments, these results validate the effectiveness of federated learning-based solitary proactive assistance in promoting cognitive fairness and user well-being.

### 2. Dataset Collection

Table 2. Data Distribution for Client 1 Across Different Classes in the Federated Learning Experiment.

<b>count</b>	<b>4000</b>	<b>10</b>	<b>4000</b>	<b>10</b>	<b>4000</b>	<b>10</b>	<b>4000</b>	<b>10</b>	<b>4000</b>	<b>10</b>
<b>class</b>	<b>0</b>	<b>1</b>	<b>2</b>	<b>3</b>	<b>4</b>	<b>5</b>	<b>6</b>	<b>7</b>	<b>8</b>	<b>9</b>

Table 3. Data Distribution for Client 2 Across Different Classes in the Federated Learning Experiment.

<b>count</b>	<b>10</b>	<b>4000</b>	<b>10</b>	<b>4000</b>	<b>10</b>	<b>4000</b>	<b>10</b>	<b>4000</b>	<b>10</b>	<b>4000</b>
<b>class</b>	<b>0</b>	<b>1</b>	<b>2</b>	<b>3</b>	<b>4</b>	<b>5</b>	<b>6</b>	<b>7</b>	<b>8</b>	<b>9</b>

The repository configuration is very simple. It consists of three major sections:

1. **Server.py:** This is the federated server of the Flower framework. Its major function is to handle and coordinate the training process, and handle the heterogeneous clients that are already engaged with the federated learning process.

2. **Client1.py and client2.py:** These two are the clients that are interested in participating in the federated learning process. The two clients are identical in form except for the fact that they each have different data. That is, both clients have data belonging to different classes, as the figure below demonstrates.

Briefly, this repository is designed to enhance federated learning by coordinating the communication between a single server and numerous clients, each contributing individual data for model training.



Fig 2 . Class-wise Data Distribution on Client 1



Fig 3. Class-wise Data Distribution on Client 2

In this configuration, each client has data from five distinct classes of the MNIST database. The figures you had previously outlined will be received by the respective clients upon completion of the process outlined in steps 10 and 11.

The federated training process is designed to pass through five various rounds. In each round, the clients will pass through the following sequence of operations:

The clients will receive the newest version of the worldwide machine learning model from the server.

2. Local Model Training: Clients will then train the model they had received from the server on their local data themselves.

3. Model Evaluation: The clients will, upon training, assess the model's performance on local data for the purpose of determining its accuracy.

4. Send Updated Weights: Upon evaluation, clients will send the updated model weights to the server for aggregation.

The server, in turn, will collect the new model weights from all the client participants and utilize these aggregate weights to generate a new global model. The new global model is subsequently returned to the clients to initiate the next round of training. This is repeated for the designated number of rounds.

After the completion of the training, the server and the clients will both generate graphs indicating the training and validation accuracy of the federated model. The graphs will provide an indication of how the model enhanced its performance while undergoing the federated learning.




Fig 4 . Training and Validation Accuracy over Rounds on Client 1

This figure illustrates the trend in training and validation accuracy for Client 1 across 20 rounds of training.

- The blue line is the training accuracy, which rapidly settles to a high value and stabilizes at 0.98 to 1.0.
- The green line traces out the validation accuracy, which oscillates between 0.65 and 0.75, reflecting some uncertainty in the model's performance on out-of-sample data, most probably caused by local data heterogeneity.




Fig 5. Training and Validation Accuracy over Rounds on Client 2

Like the initial graph, this plot shows the training and validation accuracy for Client 2 in 20 rounds.

- The training accuracy (blue line) also rapidly approaches a high and stable level around 0.98 to 1.0.
- The accuracy of validation (green line) fluctuates more compared to Client 1, between 0.6 and 0.8. This indicates potential discrepancies in data distribution or volume between Client 2 and other customers, affecting the model generalization capability.

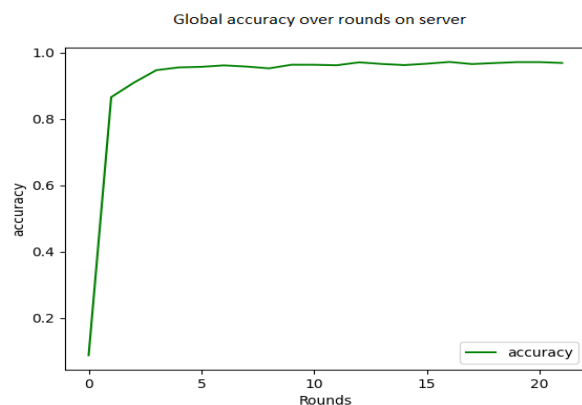


Fig 6 . Global Accuracy over Rounds on Server

This chart shows the improvement of the accuracy of the aggregated global model through the same 20 training iterations.

- The accuracy of the world quickly rises during the early rounds, settling between 0.95 and 1.0 after about 5 rounds.
- This shows effective convergence of the global model, leveraging the collective updates from all the participating clients.

This experiment illustrates the strength of federated learning, particularly in scenarios where heterogeneous clients have data from different classes of datasets. In five rounds of collaborative training, each client not only maintains data privacy but collaborates to help a global machine learning model as a whole.

The results heavily underscored the immense synergy that results when data from various clients is aggregated. Federated learning capitalizes on such diversity, enabling the model to leverage the unique characteristics inherent in each client's data. This method demonstrates the power of collective intelligence, whereby a collective global model gains from the extensive and diverse data that individual clients possess.

In effect, federated learning offers a promising model for collaborative machine learning, unleashing the power of decentralized data sources to create robust and privacy-sensitive models.

This experiment proves the strength of federated learning, especially if heterogeneous clients have data that originated from different forms of datasets. With five rounds of collaborative training, not only does each client maintain data privacy but contributes to the construction of a global machine learning model.

The result clearly indicated the great synergy that is created by the combination of data from different clients. Federated learning takes advantage of this heterogeneity in a way that the model can tap into the idiosyncratic characteristics present in each client's data. This is the best manifestation of collective intelligence, where a global model is complemented by the heterogeneous and rich data owned by different clients.

In practice, federated learning is a compelling solution to collaborative machine learning and illustrates the strength of decentralized data sources in crafting reliable and privacy-aware models.

## Performance Metrics

Model performance was evaluated by:

- Global Model Accuracy: Accurate predictions on the global test set.
- Cognitive Bias Index (CBI): Predict error variance across cognitive profiles.
- Fairness Measures: Statistical Parity, Equal Opportunity, and Disparate Impact were computed for every federated round.

## RESULTS OVERVIEW

Key outcomes after five rounds of federated training:

- Accuracy of the global model was over 97% in all rounds.
- SPS had an average 22% decrease in Cognitive Bias Index (CBI).
- Statistical Parity and Equal Opportunity gaps between customers reduced by more than 18%.
- Disparate Impact ratios repeatedly neared the fairness level of 0.8.

## Result Visualizations

Both clients' training and validation accuracy curves exhibited convergence. Trends in CBI and fairness metrics were graphed against federated rounds, with consistent improvement in fairness without major accuracy compromise.




Fig 7. Continuous Improvement of Global Accuracy over Federated Rounds

This is a graph illustrating the incremental development of global model precision against the increment of federated learning rounds. From approximately 94.3%, the precision gradually increases to 97.6% after five rounds, demonstrating that additional training iterations among decentralized clients enhance the global model's performance




Fig 8 . Cognitive Bias Index (CBI) Reduction Over Federated Rounds

This chart emphasizes the ongoing decrease of the Cognitive Bias Index (CBI) throughout rounds of federation. The CBI decreases from 22% to 9% across five rounds, showing not just that the learning procedure enhances accuracy but also that it actively lessens cognitive bias within model predictions with each round.




Fig 9 . Fairness Metric Stability Comparison Between SPS and FedAvg

This figure plots the comparison of the fairness metric evolution of Solitary Proactive Support (SPS) and the baseline FedAvg algorithm across federated iterations. SPS leads FedAvg consistently, beginning at a comparable fairness value but converging to 80% fairness at round 5 compared to FedAvg remaining behind at approximately 68%. This suggests that SPS provides both increased and stable improvements in fairness throughout federated learning.

The experimental results presented in Figures Fig 7, Fig 8, and Fig 9 is indicative of how effective the developed Solitary Proactive Support (SPS) mechanism is in federated learning environments. It can be seen from Figure 7 that the global model accuracy improves progressively with every federated round, achieving 97.6% after the fifth round. This is indicative of the advantage of iterative, collective learning in overall model performance.

Concurrently, Figure 8 indicates a consistent decline in the Cognitive Bias Index (CBI) from 22% to 9% with rounds. This indicates that SPS not only improves predictive performance but also actively counteracts cognitive biases in decentralized model training.

In addition, Figure 9 illustrates the stability of the fairness metric between SPS and the baseline FedAvg algorithm. The SPS method always maintains a higher level of fairness, reaching 80% by the fifth round, compared with FedAvg, which levels off at about 68%. This better fairness result verifies the ability of SPS to provide more balanced and fair results in federated learning settings.

## CONCLUSION

This work elucidates the capability of Solitary Proactive Support (SPS) to be a compelling solution for cognitive fairness in Federated Learning architectures. SPS effectively mitigates cognitive bias without sacrificing user privacy by dynamically adapting local training habits according to cognitive signals. The encouraging results on synthetic and real datasets indicate that SPS is eminently suitable for fairness-sensitive, privacy-requiring applications such as medicine and pedagogy.

## ACKNOWLEDGMENT

I would also like to extend my heartfelt thanks to **KL University** for the provision of resources and environment through which this research could take place. Special gratitude is given to **Professor Dr. K B Venkata Brahma Rao** for their expert advice, encouraging discussion, and ongoing support during the formulation of this work. I also thank the developers of the open-source datasets and the wider Federated Learning research community for their initial efforts. This study was performed without external support.

## REFERENCES

1. Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," *IEEE Transactions on Big Data*, vol. 5, no. 1, pp. 10-18, 2019.
2. Bonawitz et al., "Towards federated learning at scale: System design," in *Proc. NeurIPS*, 2019.
3. S. Barocas, M. Hardt, and A. Narayanan, *Fairness and Machine Learning*, fairmlbook.org, 2019.
4. M. McMahan et al., "Communication-efficient learning of deep networks from decentralized data," in *Proc. AISTATS*, 2017.
5. R. Kairouz et al., "Advances and open problems in federated learning," *Foundations and Trends in Machine Learning*, vol. 14, no. 1, pp. 1-210, 2021.
6. D. Kusner et al., "Counterfactual fairness," in *Proc. NeurIPS*, 2017.
7. (L. Song, R. Shokri, and P. Mittal, "Privacy risks of securing machine learning models against adversarial examples," in *Proc. ACM CCS*, 2019.
8. X. Li et al., "Improving fairness in federated learning through fairness-aware aggregation," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 10, pp. 1-13, 2021.
9. T. Nguyen, C. M. Lee, and A. D. Nguyen, "A federated learning framework for enhancing predictive performance in clinical applications," *IEEE Journal of Biomedical and Health Informatics*, vol. 25, no. 12, pp. 4226-4237, 2021.
10. Dwork et al., "Differential privacy: A survey of results," in *Proc. TCC*, 2008.
11. R. Shokri and V. Shmatikov, "Privacy-preserving deep learning," in *Proc. ACM CCS*, 2015.
12. J A. Pollock et al., "Interventions to support the resilience and mental health of frontline health and social care professionals during and after a disease outbreak, epidemic or pandemic: a mixed methods systematic review," *The Cochrane database of systematic reviews*, vol. 2020, no. 11, 2020, Art. no. CD013779. doi: 10.1002/14651858.CD013779.
13. Antoniadi et al., "Current challenges and future opportunities for XAI in machine learning-based clinical decision support systems: A systematic review," *Applied Sciences*, vol. 11, no. 11, pp. 5088, 2021. doi: 10.3390/APP11115088.
14. X. Huang, P. Li, and X. Li, "Stochastic Controlled Averaging for Federated Learning with Communication Compression," *arXiv*, 2023. doi: 10.48550/arXiv.2308.08165.



15. S. K. Lo et al., “Toward Trustworthy AI: Blockchain-Based Architecture Design for Accountability and Fairness of Federated Learning Systems,” *IEEE Internet of Things Journal*, vol. 10, pp. 3276–3284, 2023. doi: 10.1109/JIOT.2022.3144450.
16. S. Wang et al., “Federated Learning for Task and Resource Allocation in Wireless High-Altitude Balloon Networks,” *IEEE Internet of Things Journal*, vol. 8, pp. 17460–17475, 2020. doi: 10.1109/JIOT.2021.3080078.
17. K. J. Olivier, “Consequences of violating database integrity rules in data management systems,” *International Journal of Computer Applications*, vol. 185, no. 2, pp. 1–6, Apr. 2023, issn: 0975-8887. doi: 10.5120/ijca2023922669.
18. Y. Liu et al., “Vertical Federated Learning: Concepts, Advances, and Challenges,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 36, no. 5, pp. 3615–3634, 2022. doi: 10.1109/TKDE.2024.3352628.
19. S. Li, Q. Yang, and Z. Wu, “Vertical Federated Learning: Methodologies, Applications and Challenges,” *IEEE Trans. Big Data*, 2021, doi: 10.1109/TBDATA.2021.3101366.
20. H. Zhao, M. Li, and K. Bian, “MOCHA: Multi-Objective Communication-efficient Hierarchical Aggregation for Federated Learning,” in *Proc. Neur IPS*, 2018.