FedFuzzyGA: A Federated Evolutionary Framework for Learning Compact Fuzzy Rule-Based Classifiers

Asier Urio-Larrea*†, Javier Andreu-Perez‡, Graçaliz Pereira Dimuro§ and Humberto Bustince*†
*Departamento de Estadística, Informática y Matemáticas
Universidad Pública de Navarra, Pamplona, Navarra
Email: {asier.urio,bustince}@unavarra.es
†Institute For Smart Cities, Universidad Pública de Navarra
‡University of Essex, Colchester, United Kingdom
Email:j.andreu-perez@essex.co.uk
§Universidade Federal do Rio Grande do Sur, Rio Grande, Brazil
Email:graçaliz.dimuro@furg

***Abstract*—This paper presents FedFuzzyGA, a new federated learning framework designed to train interpretable Fuzzy Rule-Based Classifiers using decentralized and non-IID data. The primary challenge we address is that existing methods for federated fuzzy learning often result in excessively complex rule bases that are difficult to interpret. To address this problem, we adapted a centralized Genetic Algorithm into a privacy-preserving collaborative process through a method we call cyclical population synchronization. In this setup, clients perform local evolutionary searches on their private data and periodically exchange only their elite rule sets with a central server. This allows the system to share knowledge through a client-driven fusion strategy without ever compromising raw data privacy. Our experimental results show that FedFuzzyGA not only enables effective collaborative learning but often yields accuracy that meets or exceeds what a client could achieve trough training in isolation. Most significantly, the framework ensures model transparency by generating compact, readable rule sets. This makes federated evolutionary learning of fuzzy rule-based classifiers a viable method for high-stakes applications that require human-understandable AI.**

I. Introduction

The increased adoption of artificial intelligence (AI) models in high-stakes domains such as healthcare, finance, and legal services has brought the need for trustworthy systems [1]. Trustworthiness in AI, as defined by the EU's guidelines [2], is a framework that extends beyond predictive performance to include ethical and technical principles, including human oversight, technical robustness, privacy, transparency, fairness, societal well-being, and accountability. Two of these pillars are explainability, the ability to understand and justify a model's decisions, and privacy, the protection of sensitive data used in training and operation. These requirements are driven by both ethical imperatives and regulatory frameworks like the General Data Protection Regulation (GDPR) [3], which mandate transparency and control over personal data.

To address these requirements, Fuzzy Rule-Based Classifiers (FRBCs) offer a solution for inherent explainability through human-readable 'IF-THEN' rules [4], while Federated Learning (FL) [5] provides a framework for privacy by enabling collaborative training without data sharing. While initial approaches to combining FL and FRBCs exist [6], [7], [8], they frequently produce models with large and complex rule sets, which undermine their capacity for explainability.

The challenge, therefore, is to create a FL framework designed for the optimization of FRBCs. Genetic Algorithms (GAs) [9] offer a suitable foundation for this task. They are population-based optimization techniques inspired by natural selection, effective for navigating the complex, non-differentiable search spaces typical of FRBC design. Indeed, since the seminal work of Ishibuchi et al. [10], GAs have been widely employed for automatically generating, selecting, and tuning fuzzy rules and membership functions, as exemplified by the FARCHD algorithm [11]. This established paradigm of Evolutionary Fuzzy Systems [12] has proven successful for learning compact and accurate rule sets in centralized settings. While recent research has begun to explore Federated Genetic Algorithms for various optimization tasks [13], [14], applying these techniques to learn FRBCs within a federated context, where the evolutionary population is distributed across private clients, remains under-explored. This gap presents the core opportunity for our work.

To address this challenge, we propose FedFuzzyGA, a novel Federated Genetic Algorithm framework for learning Fuzzy Rule-Based Classifiers. Our method adapts the evolutionary process to the FL setting by distributing the population across clients, ensuring data privacy while promoting global convergence and, crucially, the preservation of compact, interpretable rule structures. This paper presents a foundational investigation into this paradigm. We introduce the FedFuzzyGA framework and conduct a detailed, controlled empirical analysis to answer core questions about its federated behavior, inherent trade-offs, and practical viability, using well-established benchmarks to ensure the clarity and interpretability of the results.

The remainder of this paper is organized as follows: Section II provides background information on FRBCs, GAs and FL. In Section III, we describe the design and implementation details of our proposed federated GA for learning FRBCs. The experimental setup is described in Section IV and the results are presented in Section V. Finally, Section VI discusses our findings and suggests potential future directions for research.

Code and Data Availability: To facilitate reproducibility, the implementation of the federated genetic algorithm and the scripts for dataset generation will be published in a public repository upon paper acceptance.

II. Background

A. Fuzzy Rule-Based Classifiers

A FRBC is a structured system built upon two core components: a knowledge base and an inference mechanism. The knowledge base encapsulates the specialized information needed for classification tasks, and is itself divided into a database and a rule base.

The database defines the fuzzy sets and their associated membership functions, which map crisp input values to degrees of membership in linguistic terms (e.g., "low," "medium," "high"). The rule base is a set of fuzzy `IF-THEN` rules that model the relationships between input features and output classes, effectively encoding the system's decision logic. A single fuzzy rule R_i is formally expressed as:

$$R_i : \text{If } x_1 \text{ is } A_i^1 \text{ and } x_2 \text{ is } A_i^2 \text{ and } \dots \text{ and } x_n \text{ is } A_i^n, \\ \text{then Class is } C_i \text{ with weight } RW_i$$

Here, $x_1, x_2, \dots, x_n$ denote the input features; A_i^j is the fuzzy linguistic label for the j-th feature in the i-th rule; C_i is the consequent class label; and $RW_i \in [0, 1]$ is a rule weight that quantifies the rule's confidence or importance in the final classification decision [15].

The inference procedure classifies an input $\mathbf{x}$ by first calculating the membership degree $\mu_{A_{ij}}(x_j)$ of each feature to a rule's fuzzy sets. The activation strength $\alpha_i(\mathbf{x})$ of a rule R_i is computed via a t-norm over these memberships: $\alpha_i(\mathbf{x}) = T(\mu_{A_{i1}}(x_1), \dots, \mu_{A_{in}}(x_n))$. This strength is then weighted by the rule's confidence to produce an association degree $\beta_i^k(\mathbf{x}) = \alpha_i(\mathbf{x}) \cdot RW_i$ for class k. For each class, these degrees are aggregated into a class soundness $Y_k = S_{j:C_j=k}(\beta_j^k(\mathbf{x}))$. The final classification is determined by selecting the class with the maximum soundness: $\text{Class}(\mathbf{x}) = \arg\max_{k=1,\dots,N} Y_k$.

B. Genetic Algorithms-based Fuzzy Rule Learning

GAs are metaheuristics inspired by the principles of natural selection; they solve complex optimization problems by evolving a population of candidate solutions over successive generations. The canonical procedure, as illustrated in the dashed box in Figure 1, involves a cyclical process: a selection operator favors individuals with higher fitness (e.g., better classification accuracy) for reproduction; a crossover operator recombines genetic material from parent chromosomes to create offspring; and a mutation operator introduces random alterations to maintain population diversity. Through this iterative cycle of evaluation and variation, the population converges toward increasingly optimal solutions.

When applied to FRBC design [12], the GA's optimize the classifier's knowledge base. In this paradigm, an individual chromosome typically encodes a complete candidate rule set or a subset of a rule base. The GA-based optimization is applied to three tasks in FRBC design, as exemplified by the FARCHD algorithm [11]. First, the GA searches the combinatorial space of possible IF-THEN rules, using its genes to encode antecedent conditions and consequents. Second, it filters the rules, where each gene representing the inclusion or exclusion of a rule from the rule candidate pool, enabling the evolution of a compact and high-performance rule subset. Third, it refines the continuous parameters of predefined fuzzy membership functions, optimizing the classifier's decision boundaries for better accuracy.

C. Federated Learning

FL was introduced to address the challenges of data privacy and the computational resources required for the training of machine learning models. This paradigm enables the training of global models through the collaboration of multiple devices without the need to share raw data. Each device performs its own local training and shares the model parameters with a central server, which aggregates them to update the global model. Additionally, federated learning can function in a decentralized manner, where devices share updates directly with each other, eliminating the need for a central server. This approach ensures privacy and security while leveraging distributed data sources for collaborative model training.

One of the most important aspects of federated learning is the aggregation of information from individual devices. Originally, McMahan's proposal, FedAvg, created the global model by computing a weighted average of the neural network weights. However, subsequent works have explored alternative aggregation methods under different conditions to improve performance and robustness [16].

Another challenge in federated learning is the potential difference in data distribution across clients. In traditional centralized learning, datasets are often assumed to be independently and identically distributed (IID), meaning that each data point is drawn from the same underlying distribution and is independent of the others. However, in federated learning, data originates from diverse sources, such as different devices, users, or organizations. This means that the distribution of data across clients may vary significantly, reflecting differences in user behavior, environmental conditions, or other factors [17]. As a result, the assumption of IID data no longer holds, which must be taken into account to ensure good model performance, generalization, and stable convergence during the training of the models.

D. Federated Genetic Algorithms

In an FGA framework, the optimization process occurs hierarchically [18]. Clients execute genetic operations locally, evaluating candidate solutions against their private data to derive fitness values or train local models. Instead of transmitting raw data, clients upload only model parameters, gradients, or elite individuals to a central orchestrator. The server aggregates these contributions, often using strategies like weighted

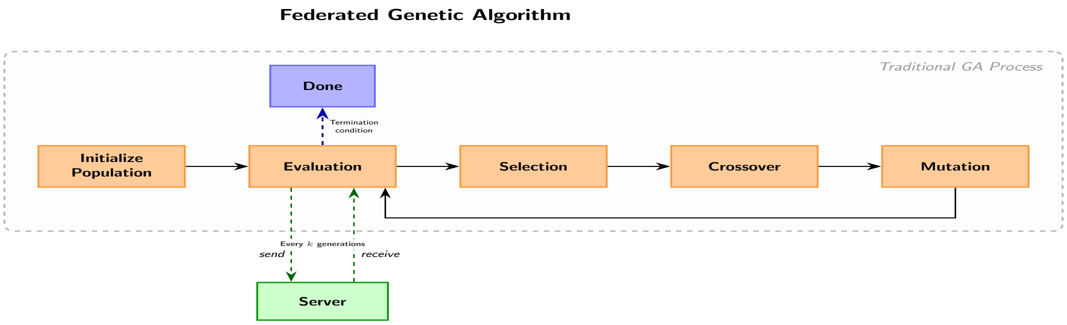

Fig. 1: Architectural overview of the FedFuzzyGA. The dashed boundary encapsulates the standard local GA workflow, while the external green channel illustrates the periodic synchronization with the central server for global aggregation.

averaging or elitist selection [14], to construct a global model or population, which is subsequently broadcast back to the clients [19]. This mechanism ensures that sensitive information remains on local devices while enabling the system to solve complex optimization problems across heterogeneous environments.

III. The FedFuzzyGA Framework

In this section, we introduce our proposed Federated Genetic Algorithms-based Fuzzy Rule Learning system. First, we will introduce the base GA upon we will build our FedFuzzyGA system. Secondly, we present the federated implementation, and finally detail the different configurations admited by our algorithm.

A. Genetic Algorithm for Fuzzy Rule Learning

The FedFuzzyGA framework is built upon a centralized GA for FRBC learning, specifically the implementation provided by the *ex-fuzzy* library [20]. A key advantage of this library is that it is designed to evolve compact, interpretable fuzzy rule sets by actively promoting a low number of rules and antecedents per rule. This library, which utilizes the *pymoo* framework [21] for optimization, handles the chromosome representation and evolutionary operators, allowing our work to focus on the novel federated coordination mechanism without sacrificing the primary goal of explainability.

- Chromosome Encoding: Each individual encodes a complete rule set. The antecedents are represented via two lists of length nAnts × nRules: a binary vector for feature selection and an integer vector for linguistic label assignment. The consequents are stored in a separate integer vector of length nRules specifying the output class for each rule.
- Evolutionary Process: The algorithm follows a standard single-objective GA workflow:
  - Initialization: The population is generated using integer random sampling, ensuring each gene is correctly initialized within its valid domain.
  - Variation: Offspring are created through Simulated Binary Crossover (SBX) [22] and Polynomial Mutation (PM) [23], with results rounded to maintain integer representation.
  - Selection: A Tournament Selection mechanism drives the population toward higher fitness, where fitness is evaluated as the classifier's accuracy (optionally combined with a complexity penalty) on the local dataset.

TABLE I: Structure of a gen representing a FRBC

Antecedent Mask	Linguistic Labels	Consequents
$nAnts \times nRules$ 0/1	$nAnts \times nRules$ $0..nTerms - 1$	$nRules$ $0..nClasses - 1$

Following the evolutionary process, a rule pruning step is applied to each candidate classifier: rules whose computed confidence or weight falls below a predefined threshold are discarded from the final rule base. Thus, enforcing the compactness and interpretability of the evolved models

B. Federated Genetic Algorithm for Fuzzy Rule Learning

Our contribution lies in the transformation of the centralized GA for FRBC learning into a collaborative framework. FedFuzzyGA introduces a cyclical population synchronization protocol which, instead of the widely used averaging of model parameters, acts as an exchange of diverse genetic material between clients. This process enables participants to collaboratively search for an optimal, compact FRBC by leveraging global evolutionary progress without ever compromising their private data.

The federated workflow, illustrated in Figure 1, operates in synchronized cycles across N clients and proceeds as follows:

1) Local Evolutionary Phase: Each client independently runs the local GA (described in section III-A) for a fixed number of k generations. All fitness evaluations are performed exclusively on the client's private dataset.
2) Synchronization Phase: Every k generations, a global synchronization round is triggered:
   - Client-to-Server: Each client selects its top performing individuals based on local fitness and transmits only these individuals to a central server.

- Server Aggregation: The server collects the migrant individuals, forming a global elite pool.
- Server-to-Client: The server broadcasts this set of global representatives back to all clients.
- Population Integration Phase: Each client integrates the global representatives into its local population. The specific integration logic, defining which individuals are replaced, is a design choice described in the final fusion strategy (see Section III-C).

3) Cycle Continuation: The process repeats, with the included new individuals, until a terminal condition is met (e.g., a maximum number of generations reached).

This mechanism allows the search for an optimal FRBC to leverage insights learned from all clients' data distributions, while the final rule pruning step (applied locally post-evolution) ensures the interpretability of the collaboratively evolved model is preserved.

C. Variations

The cyclical synchronization protocol provides the basic federated framework. To investigate its effective configuration for FRBC learning, we formalize and implement three key control variants that govern the intensity of exchange, the final model derivation, and the potential for local specialization. These parameters define a family of FedFuzzyGA strategies.

- Number of Communication Rounds (*R*): This parameter sets how many times the population synchronization protocol is executed. It controls the balance between local computation and global collaboration. Given a number of T local generations per client, the number of generations performed between each synchronization event is $k = T/R$
- Migration Size (*M*): This parameter controls the communication volume and selection pressure during synchronization. A small M exchanges only the very best individuals, promoting convergence but potentially reducing genetic diversity. A larger M preserves more client-specific evolutionary paths at the cost of increased communication.
- Final Fusion Strategy: This variant determines how the final global model is produced after the last synchronization round. It addresses the trade-off between consensus and preserving local refinements.
  - Merge-Then-Select (MS) : After the final broadcast, all clients integrate the global representatives into their populations. The single best-performing individual across all final client populations is selected as the global model.
  - Select-Only (SO): Each client select the best individual from the global representatives, without considering its previous population. This strategy prioritizes a consensus model based on aggregated elite knowledge.
- Post-Fusion Fine-Tuning: This enables client-specific personalization after the final model is selected. After creating the population by the union between the local population and the received server individuals, if enabled, each client performs an additional run of k generations of the *ex-fuzzy* GA. This produces a personalized model adapted to the local distribution while being initialized from the collaboratively learned, interpretable rule base. If disabled, the best individual is selected from the combined population.

By tuning these parameters, we can produce different federated behaviors, from tightly synchronized global learning to more independent local adaptation. The experimental section (Section IV) will analyze their impact.

IV. Experimental Setup

In this section, we will describe the experimental design, with a particular focus on the generation of the datasets and the models used for evaluation.

Datasets

For the experimental validation, we selected three benchmark datasets: Iris, Balance, and Dermatology [24]. We used these standard benchmarks to guarantee reproducibility and, for our foundational research, to make a detailed ablation study computationally practical. The Iris dataset served as the primary testbed for this ablation study, ensuring the analysis remained tractable and the results clearly interpretable. The Balance and Dermatology datasets were then added to evaluate the comparative performance and generalization capability of the final FedFuzzyGA configuration against baseline methods in more scenarios.

To simulate a realistic federated scenario, we introduce feature distribution skew across clients, creating a non-IID data environment. This is achieved using the FedArtML tool [25], which partitions the dataset among $N = 5$ clients based on the Dirichlet distribution $DirN(\alpha)$. We set the concentration parameter $\alpha = 0.5$ to induce a high degree of feature heterogeneity. This process ensures each client receives a unique and biased subset of the original data.

Baseline Methods

We compare FedFuzzyGA against two baselines:

- **Local Learning**: In this case, each agent has access only to its own data. They train independently, and there is no collaboration among the different agents.
- **Federated Chi**: The federated Chi algorithm described in [8], wherein clients locally execute Chi's algorithm, transmit their resulting rule bases to a server for aggregation, and subsequently receive a combined federated rule base

Evaluation Metrics and FedFuzzyGA Configuration

We assess performance using: Global Model Accuracy (primary) and Average Rule Count. All methods use the same *ex-fuzzy* hyperparameters: a population size of 50, maximum number of generations T=100 generations, and a rule pruning threshold of 0.1. For FedFuzzyGA, we perform an ablation study over the grid: $R \in \{1, 5, 10, 20\}$, $M \in \{1, 5, 10\}$, Fusion

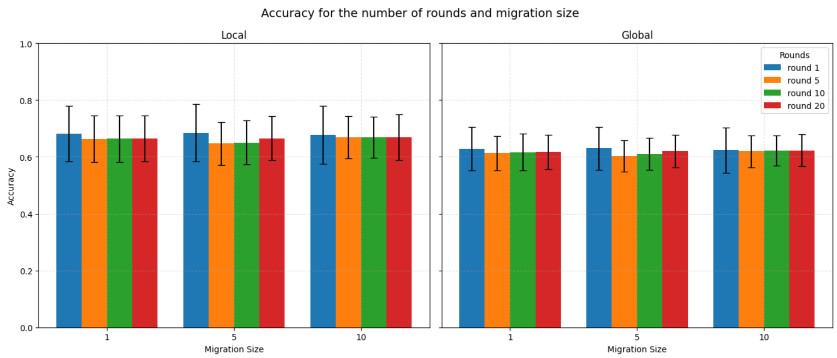

(a) Rounds and migration size

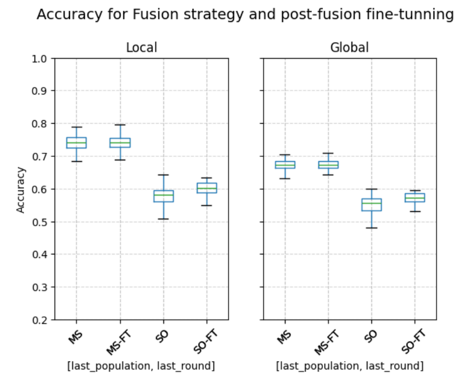

(b) Last round strategies

Fig. 2: Ablation study results

TABLE II: Performance comparison across datasets and methods.

Dataset	Method	Global Acc.	Local Acc.	Avg. # Rules
Iris	Federated Chi	**0.87** $\pm$ **0.12**	**0.85** $\pm$ **0.14**	18 $\pm$ 1.19
	FedFuzzyGA (Ours)	0.79 $\pm$ 0.05	0.72 $\pm$ 0.07	**4.5** $\pm$ **1.19**
	Local Learning	0.76 $\pm$ 0.09	0.68 $\pm$ 0.06	4.6 $\pm$ 1.1
Dermatology	Federated Chi	0.29 $\pm$ 0.04	0.27 $\pm$ 0.13	262.8 $\pm$ 1.48
	FedFuzzyGA (Ours)	**0.57** $\pm$ **0.03**	**0.53** $\pm$ **0.07**	3.5 $\pm$ 1.4
	Local Learning	0.52 $\pm$ 0.07	0.51 $\pm$ 0.05	**3.38** $\pm$ **1.3**
Balance	Federated Chi	0.36 $\pm$ 0.09	0.31 $\pm$ 0.12	90.2 $\pm$ 2.5
	FedFuzzyGA (Ours)	**0.61** $\pm$ **0.04**	**0.59** $\pm$ **0.02**	**6.24** $\pm$ **2.17**
	Local Learning	0.61 $\pm$ 0.05	0.59 $\pm$ 0.04	7.1 $\pm$ 1.6

Strategy $\in$ {Merge-Then-Select, Select-Only}, and the post-fusion fine-tunning which can be applied or not. All results are averaged over 10 independent runs.

The models are evaluated to assess different capabilities. Local accuracy measures a model's performance on the specific data distribution of its originating client. Global accuracy evaluates the generalization capability of the consensus model by testing it on the combined dataset from all clients.

V. Results

A. Ablation Study

An ablation study was conducted to asses the impact of FedFuzzyGA's control parameters: the number of communication rounds (R), the migration size (M), the final fusion strategy, and the post-fusion fine-tuning. The results in Figure 2 reveal a robust to most configuration choices.

- The number of communication rounds (R) and migration size (M) showed no significant impact on accuracy (Figure 2a). This indicates the basic mechanism of intermittent elite exchange is effective and simplifies deployment.
- The Merge-Then-Select (MS) approach, which allows a final client-side evaluation of global elites, outperformed the server-only Select-Only (SO) strategy (Figure 2 b). This result shows that local validation of collaborative solutions improves results.
- Post-fusion fine-tuning provided no significant gain for learning a global model (Figure 2b), suggesting the main evolutionary learning occurs during synchronized rounds.

FedFuzzyGA's efficacy stems from the synergy of elite diffusion and local ratification. The system is not sensitive to communication granularity, but requires the final selection pressure to be applied directly to client data (MTS strategy). This finding clarifies the core mechanism and provides a clear guideline for application: prioritize the MTS fusion strategy.

B. Comparative results

The results in Table II delineate a clear trade-off between model interpretability and federated performance. The prior federated method (Federated Chi) achieves high accuracy only on the simplest dataset (Iris) and fails on more complex problems (Dermatology, Balance), while consistently generating complex models with hundreds of rules, rendering them unsuitable for explainable AI.

Our proposed FedFuzzyGA mantains the intrinsic interpretability of fuzzy systems within a federated setting, producing compact rule sets (3-6 rules) comparable to those of models trained on isolated local data.

In terms of accuracy, FedFuzzyGA shows a nuanced profile: it maintains performance parity with the Local ex-fuzzy baseline on the Balance dataset, provides a modest but consistent improvement on the more complex Dermatology dataset, and shows a slight decrease on the simple Iris dataset where the Federated Chi method excels. This indicates that the benefit of our federated collaboration is most pronounced in challenging scenarios, while the core value lies in its ability to enable collaborative learning without sacrificing the explainability that is lost in other federated approaches.

Our experiments confirm that FedFuzzyGA learns more general models than those trained locally. As shown in Table II, the global model's accuracy, tested on aggregated data from all clients, exceeds the average accuracy of the local models evaluated on their respective private test sets. This demonstrates a collaborative gain: the global model is not merely an average of local solutions, but a superior model that integrates the knowledge of all individual clients. Consequently, it performs better on the complete data distribution than any single client's specialized model performs on its own subset. This suggests that our cyclical population synchronization enables effective knowledge fusion, creating a model that is more robust and generalizable than any model learned in isolation.

VI. Conclusions

This work presents FedFuzzyGA, a novel method for learning FRBCs by federating the GA responsible for rule discovery. Our framework adapts the evolutionary process to a distributed setting through cyclical population synchronization, enabling collaborative learning while preserving data privacy.

The experimental results demonstrate that FedFuzzyGA successfully learns accurate and interpretable models in non-IID federated scenarios. The model achieves this while producing compact rule sets, an essential property for explainability that is lost in other federated approaches like the federated version of Chi's algorithm. While the accuracy gains over isolated local training are context-dependent, our method consistently enables effective collaboration without sacrificing model interpretability, establishing a new trade-off point in the federated learning landscape.

This work establishes FedFuzzyGA as a novel and viable paradigm for explainable FL. While evaluated on standard benchmarks to facilitate a clear ablation study and direct comparison, the demonstrated principles provide a strong foundation for future application to larger, real-world datasets

Future work will focus on several directions: 1) Extending the federated protocol to other components of the evolutionary cycle, such as different server aggregation strategies, distributed fitness evaluation or mutation operators; 2) Conducting more experiments on real-world distributed datasets; and 3) Formalizing convergence guarantees for the federated evolutionary process to strengthen its theoretical foundation.

Acknowledgments

This work is supported by FAPERGS (23/2551-0001865-9, 23/2551-0000126-8), CNPq (304118/2023-0, 407206/2023-0), MCIN/AEI/10.13039/50100011033/FEDER,UE (PID2022-136627NB-I00).

References

[1] S. Thiebes, S. Lins, and A. Sunyaev, "Trustworthy artificial intelligence," *Electronic Markets*, vol. 31, no. 2, pp. 447–464, 2021.

[2] High-Level Expert Group on Artificial Intelligence (AI HLEG), "Ethics guidelines for trustworthy AI." Technical Report, 2019. Directorate-General for Communications Networks, Content and Technology.

[3] "Regulation (eu) 2016/679 of the european parliament and of the council of 27 april 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing directive 95/46/ec (general data protection regulation)."

[4] A. Moral, C. Castiello, L. Magdalena, and C. Mencar, *Explainable fuzzy systems*. Springer, 2021.

[5] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial intelligence and statistics*, pp. 1273–1282, PMLR, 2017.

[6] A. Wilbik and P. Grefen, "Towards a Federated Fuzzy Learning System," in *2021 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE)*, pp. 1–6, July 2021. ISSN: 1558-4739.

[7] Y.-C. C. L. Zhang, Y. Shi and C.-T. Lin, "Federated fuzzy neural network with evolutionary rule learning," *Transactions on Fuzzy Systems*, vol. 31, no. 5, p. 1653–1664, 2023.

[8] M. Daole, P. Ducange, F. Marcelloni, and A. Renda, "Trustworthy AI in Heterogeneous Settings: Federated Learning of Explainable Classifiers," in *2024 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE)*, pp. 1–9, June 2024. ISSN: 1558-4739.

[9] M. Mitchell, *An Introduction to Genetic Algorithms*. The MIT Press, 02 1996.

[10] H. Ishibuchi, K. Nozaki, N. Yamamoto, and H. Tanaka, "Construction of fuzzy classification systems with rectangular fuzzy rules using genetic algorithms," *Fuzzy Sets and Systems*, vol. 65, no. 2, pp. 237–253, 1994. Fuzzy Methods for Computer Vision and Pattern Recognition.

[11] J. Alcalá-Fdez, R. Alcala, and F. Herrera, "A fuzzy association rule-based classification model for high-dimensional problems with genetic rule selection and lateral tuning," *IEEE Transactions on Fuzzy systems*, vol. 19, no. 5, pp. 857–872, 2011.

[12] A. Fernández, V. López, M. J. del Jesus, and F. Herrera, "Revisiting evolutionary fuzzy systems: Taxonomy, applications, new trends and challenges," *Knowledge-Based Systems*, vol. 80, pp. 109–121, 2015. 25th anniversary of Knowledge-Based Systems.

[13] R. Marin Machado de Souza, A. Holm, M. Biczyk, and L. N. de Castro, "A systematic literature review on the use of federated learning and bioinspired computing," *Electronics*, vol. 13, no. 16, 2024.

[14] Y.-F. Ge, H. Wang, J. Cao, Y. Zhang, and G. Kambourakis, "Federated genetic algorithm: Two-layer privacy-preserving trajectory data publishing," in *Proceedings of the Genetic and Evolutionary Computation Conference*, GECCO '24, (New York, NY, USA), p. 749–758, Association for Computing Machinery, 2024.

[15] H. Ishibuchi and T. Nakashima, "Effect of rule weights in fuzzy rule-based classification systems," *IEEE Transactions on Fuzzy Systems*, vol. 9, no. 4, pp. 506–515, 2001.

[16] K. Pillutla, S. M. Kakade, and Z. Harchaoui, "Robust aggregation for federated learning," *IEEE Transactions on Signal Processing*, vol. 70, pp. 1142–1154, 2022.

[17] Y. Zhao, M. Li, L. Lai, N. Suda, D. Civin, and V. Chandra, "Federated learning with non-iid data," *arXiv preprint arXiv:1806.00582*, 2018.

[18] J. Xu, Y. Jin, W. Du, and S. Gu, "A federated data-driven evolutionary algorithm," *Knowledge-Based Systems*, vol. 233, p. 107532, 2021.

[19] S. Khatua, A. Mukherjee, and D. De, "Fedgen: Federated learning-based green edge computing for optimal route selection using genetic algorithm in internet of vehicular things," *Vehicular Communications*, vol. 49, p. 100812, 2024.

[20] J. Fumanal-Idocin and J. Andreu-Perez, "Ex-fuzzy: A library for symbolic explainable ai through fuzzy logic programming," *Neurocomputing*, vol. 599, p. 128048, 2024.

[21] J. Blank and K. Deb, "pymoo: Multi-objective optimization in python," *IEEE Access*, vol. 8, pp. 89497–89509, 2020.

[22] K. Deb, K. Sindhya, and T. Okabe, "Self-adaptive simulated binary crossover for real-parameter optimization," in *Proceedings of the 9th Annual Conference on Genetic and Evolutionary Computation*, GECCO '07, (New York, NY, USA), p. 1187–1194, Association for Computing Machinery, 2007.

[23] K. Deb and D. Deb, "Analysing mutation schemes for real-parameter genetic algorithms," *International Journal of Artificial Intelligence and Soft Computing*, vol. 4, no. 1, pp. 1–28, 2014. PMID: 59280.

[24] D. Dua and C. Graff, "UCI machine learning repository,"

[25] G. D. M. Jimenez, A. Anagnostopoulos, I. Chatzigiannakis, and A. Vitaletti, "Fedartml: A tool to facilitate the generation of non-iid datasets in a controlled way to support federated learning research," *IEEE Access*, 2024.