Fully-Decentralized Consensus-Based Federated Learning for Cell Outage Detection in Cellular Networks

Andrea WRONA*
Dept. of Computer, Control, and Management Engineering
Sapienza University of Rome
Via Ariosto 25, 00185 Rome, Lazio, Italy

Simone GENTILE
1) Dept. of Computer, Control, and Management Eng.
Sapienza University of Rome
Via Ariosto 25, 00185 Rome, Lazio, Italy
2) Dept. of Electrical and Information Eng.
Polytechnic of Bari
Via Re David 200 70125 Bari, Puglia, Italy

Emanuele DE SANTIS
Dept. of Computer, Control, and Management Engineering
Sapienza University of Rome
Via Ariosto 25, 00185 Rome, Lazio, Italy

***Abstract*—Cell Outage Detection (COD) mechanisms in 5G and beyond cellular networks play an increasingly important role in ensuring uninterrupted services to end users by promptly identifying possible outages at the radio and cell levels. Traditionally, COD algorithms have used aggregated data at the core network level to detect anomalies, but there have been scalability and data confidentiality issues. This work proposes a novel fully-decentralized consensus-based Federated Learning approach. This approach utilizes Random Trees and federated feature removals to identify anomalies at the cell level. It is based only on data available locally at the Base Station (BS), but relies on knowledge acquired by all BSs participating in the federation. The approach is fully decentralized in the sense that it does not involve a central entity responsible for aggregating the knowledge of the learning agents. A set of simulations based on a dataset with real cell data has been employed to demonstrate the effectiveness of the proposed approach in comparison to other baseline approaches, even in the presence of malicious agents attempting to disrupt the learning process.**

I. INTRODUCTION

Fifth-generation and beyond (5G/B5G) networks offer data rates in the order of gigabits per second, very low latency at both the radio and core levels, the ability to connect thousands of User Equipments (UEs) simultaneously, thus enabling entirely new use cases with widespread service availability and high performance [1]. However, these system performances are strictly dependent on the efficient functioning of the cellular network, with degradation or system unavailability in case of equipment failures, malicious attacks, capacity saturation at the base station (BS) level, etc. This has important consequences for users in the affected areas, with a reduction in the perceived quality of service [2], [3].

Cell Outage Detection (COD) mechanisms aim to mitigate such problems by providing timely identification of anomalies occurring in the cellular network, thus allowing the network operator to take appropriate countermeasures as soon as possible to mitigate the problems for the affected users.

Several approaches to COD have been proposed in the literature. Among the others, some algorithms, like the one proposed in [4], make some geographically based correlation between the users, some make use of handover metrics and statistics [5], other works use reports from neighboring cells [6] or Channel Quality Indicators (CQIs) of the UEs [7] to detect faulty cells. Some works have also applied traditional machine learning methods to COD: [8] used a Naive Bayes classifier to cluster cells and identify faulty ones in a UMTS network, using real and simulated data of the network for training, while [9] uses Hidden Markov Models for COD, predicting the status of a 4G eNodeB by training on data from healthy and anomalous cells. Other approaches use unsupervised learning for COD as in [10], [11] to cluster healthy cells and cells with anomalies, while the work in [12] uses unsupervised learning to identify the location of the faults. Other approaches also use unsupervised deep learning for COD, such as [13], [14], which implements an autoencoder-based architecture for anomaly detection based on RSRP and RSRQ measures.

As for supervised deep learning approaches, several works have been produced for COD, and in particular [15], [16], which distribute the classification process among the edge nodes of the network, making the approach scalable with the number of base stations and cells to be monitored, but suffer from a lack of knowledge sharing among the training agents, which train only on their own datasets. This leads to a potential imbalance in classification performance due to the different number of datasets in each cell.

In order to take advantage of the availability of additional data belonging to different learning agents, the federated learning methodology is gaining popularity due to its ability to share knowledge between agents without sharing actual data, thus preserving the privacy and confidentiality of the data. Federated learning techniques have been used in the context of COD by [17], which detects faults and optimizes resources to recover users falling in the faulty cell using neural networks in combination with standard constrained optimization techniques, or by [18], which implements a federated unsupervised learning strategy for detecting antenna tilt anomalies.

The main problem with federated learning architectures is that they typically require a central entity that coordinates the

* Corresponding author. Email: wrona@diag.uniroma1.it

learning process and aggregates the knowledge received by the training agents. This is a single point of failure in the training process, since if the central entity stops its orchestration and aggregation tasks for any reason (e.g., attacks, malfunctions, overload, etc.), the entire learning process is blocked. To overcome this problem, fully decentralized methods for federated learning have recently been studied in the literature, for example by [19], [20], [21].

This paper, based on a previous work by the authors [22], introduces a fully distributed consensus-based approach for a federated Random Forest classifier of cell anomalies. As in the previous work, the learning agents' Random Forest classifiers progressively remove features selected by the federation as less important. However, the newly proposed fully decentralized approach takes advantage of the federated learning framework while mitigating possible risks due to attacks and/or malfunction of the central entity.

The main original contributions of this thesis are summarized as follows:

- a consensus-based mechanism for solving the COD problem to aggregate the importance scores of the features learned by each agent to all agents in the federation; this makes the approach fully decentralized, since no central entity is needed for such aggregation;
- a modified consensus law to make the decentralized learning agents robust against possible other malicious agents (attackers) that want to destroy or degrade the learning performance of the other agents in the federation.

The remainder of the paper is as follows: Section II shows the considered mobile network scenario for cell anomaly detection problem; Section III details the proposed fully-decentralized federated learning approach, together with some preliminary notions on consensus theory; Section IV presents some simulations of the proposed approach in a realistic cellular network scenario, thus showing the effectiveness of our work with respect to baseline approaches; finally Section V summarizes the contribution of the paper and analyzes future works.

II. Considered Cellular Network Scenario

Fig. 1 shows the considered cellular network scenario, as composed by a set of N Base Stations (BSs). Each Base Station i serves a set C_i of c_i cells (hexagons in the figures). Even if usually each Base Station hosts 3 cells, in this work c_i can be considered equal to any number of cells, even different Base Station per Base Station (i.e., $c_i \neq c_j, j \neq i$).

Some cells $k \in C_i$ may be faulty, e.g., due to a malfunctioning, due to an attack, or due to severe overloading (represented in red), while the others are normal cells (represented in yellow, green, light blue and orange).

Each Base Station i has access to an Edge Server (ES) located in its neighborhood, having some computation capability and being able to connect with a low-latency communication to a certain number of neighboring Edge Servers, belonging to neighbor BSs N_i. In Fig. 1, each BS i is considered, without loss of generality, to be connected to its two nearest BSs in a ring topology.

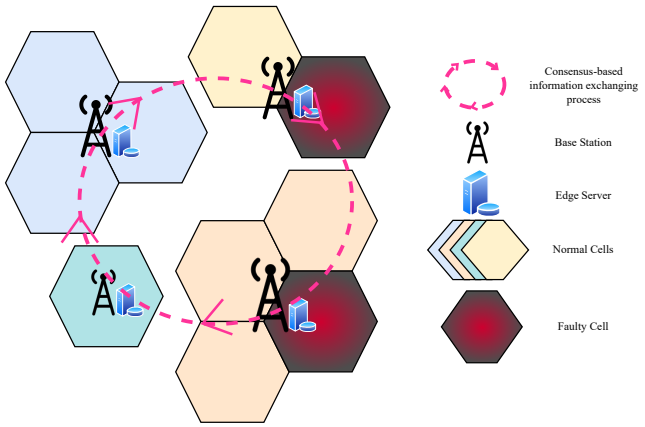

Fig. 1. Cellular Network Scenario.

In the presented scenario, a number $N_a < N$ of BSs may be considered compromised (e.g., by a cyber attack to their respective ES) The role of these compromised BSs is to degrade or destroy as much as possible the fedeated fully-distributed learning process of the remaining $N - N_a$ non-compromised BSs.

Moreover, Edge Server of BS i contains a dataset $\mathcal{D}_i$ composed by history records of the cells C_i served by Base Station i. This data include information on (i) date and time, (ii) Physical Resource Blocks (PRBs) usage, (iii) number of connected User Equipments (UEs), (iv) other user data, as for example Call Detail Records [15]. It is important to notice that each Base Station i relies only upon data coming from its cells C_i for its local training (performed by its associated ES), while it leverages on the knowledge (not the data) of the other BS j by the means of the federated algorithm detailed in the next section.

III. Consensus-based Feature Removal Process

A significant challenge in the COD process is determining the optimal set of features that ensures high classification performance. The selection or removal of features is essential when handling large volumes of data, particularly when it is noisy and originates from multiple users.

As mentioned in Section I, this study presents a collaborative and privacy-preserving method for feature selection utilizing Random Forests. This ensemble learning method builds a collection of decision trees, hierarchical structures composed of nodes, where each leaf node represents a prediction and each internal node corresponds to a decision based on a specific feature. Each decision tree is trained on a randomly sampled subset of the training data, using the bagging technique.

For classification tasks such as anomaly detection, each tree independently predicts a class based on the given input, and the final prediction is determined by the class with the highest number of occurrences.

One of the advantages of using a Random Forest classifier is its ability to calculate the importance of a feature by

considering how often it is selected for splitting the trees' nodes in the ensemble. This results in an importance score for each feature, which helps determine the significance of each feature within the dataset.

In this work, importance scores calculated at the local level are combined through a consensus-based approach to enhance the overall anomaly detection performance. Consensus theory determines how autonomous agents can reach a common agreement through local interactions, without relying on a central authority.

The dynamics of each node (or agent) under the consensus protocol are described by:

$$\dot{x}_i = \sum_{j \in N_i} a_{ij}(x_j - x_i), \tag{1}$$

where:

- x_i represents the state of agent i ;
- N_i is the set of neighbors of agent i;
- a_{ij} is the weight representing the connectivity between agents.

In matrix form, the system evolves as

$$\dot{x} = -Lx, \tag{2}$$

where L is the Laplacian matrix of the graph, defined as:

$$L = D - A, \tag{3}$$

where D is the degree matrix, a diagonal matrix where D_{ii} represents the degree of node i, and A is the adjacency matrix of the graph, defined such that $A_{ij} = 1$ if there is an edge between nodes i and j, and $A_{ij} = 0$ otherwise.

Global consensus is achieved if the graph is connected, which is ensured when the second smallest eigenvalue λ_2 of the Laplacian matrix is strictly positive [23], [24].

The integration of consensus theory into the AI-based COD framework eliminates the reliance on a central authority in charge of averaging data coming from local nodes. This prevents vulnerabilities associated with single points of failure, and promotes scalability in large-scale cellular networks.

A. Robust Consensus Algorithm

The consensus-based FedRF algorithm is detailed in the pseudocode of Alg. 1.

In the initial phase, all ESs are part of the federation, where they independently train on their local datasets and calculate feature importance scores.

At each iteration, nodes exchange their feature importance scores only with their direct neighbors. Using the consensus theory, each node updates its scores based on the information received, leading to a global agreement on feature importance across the network. This can be expressed mathematically in the following way. Node i updates each of its importance scores x for the metrics m according to the following dynamical system:

$$\dot{x}_i^m(t) = \sum_{j=1}^{N_i} w_{j,i}(t)\left(x_j^m(t) - x_i^m(t)\right) \tag{4}$$

where N_i is the number of neighbors of node i, and $w_{j,i}$ is the reliability weight of node j according to node i, computed as:

$$w_{j,i}(t) = \begin{cases} 1, & \text{if } \left|x_i^m(t) - x_j^m(t)\right| \leq \tau, \\ 0, & \text{otherwise} \end{cases} \tag{5}$$

with τ being an appropriate tolerance threshold. This reliability weight $w_{j,i}$ aims to exclude from the consensus protocol all agents who have a state $x_j^m(t)$ that is too far from the other agents one ($x_i^m(t)$, $i \in N_j$ neighbors of node j). Having such a different state is very unlikely to be caused by differences in the datasets and in the training process, thus indicates that the agents with outlier states may be attackers. This situation may happen, for example, in the case of Byzantine attackers that manipulate their reported states $x_j^m(t)$ to try to steer all the other agents towards a (very) different shared state once the consensus is reached, thus leading to poor performances of the whole federated training process.

Once consensus is reached, the feature with the lowest agreed importance score is identified as the least significant. Each node then removes this feature and retrains its Random Forest model accordingly. The ESs that show a decline in performance leave the federation while maintaining the identified feature, whereas the remaining ones proceed with a new iteration. The process repeats iteratively until the federation no longer contains any agents.

Algorithm 1 Consensus-based Feature Removal in Federated Random Forest (FedRF)

Require: Individual datasets $\mathcal{D}_1, \mathcal{D}_2, \ldots, \mathcal{D}_N$ from N base stations
Require: Federation: $\mathfrak{F} = \langle \text{BS}_1, \ldots, \text{BS}_N \rangle$
Ensure: Final selected features $\mathcal{F}_{\text{final}}$

```
Initialize F_final with all features
for each client i in 𝔉 do
    Train Random Forest on D_i with features F_final
    Compute feature importance scores (I_j^i)
end for
repeat
    I_j^cons ← ConsensusUpdate(I_j),  ∀ feature j
    j* ← arg min_j(I_j^cons)
    Remove least significant feature j* from F_final
    for each client i in 𝔉 do
        Train Random Forest on D_i with features F_final
        Compute feature importance scores (I_j^i)
        if Client i worsens its performance then
            Client i rejects group's decision and exits 𝔉
        end if
    end for
until 𝔉 ≠ ∅
return F_final
```

The following section will evaluate the proposed consensus-based federated approach using a dataset suitable for solving the COD problem.

IV. Simulation Results

A. Dataset and Metrics

The dataset used to carry out anomaly detection at radio level has been taken from [25]. It contains various radio

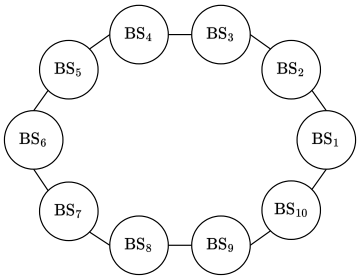

Fig. 2. Ring Network Topology

and connectivity metrics, collected every 15 minutes for two weeks, from a set of $N = 10$ BSs, each one handling a different number of cells. Tab. I contains an overview of the features available for classifying the cell anomalies. Based on the *CellName* feature, the dataset has been split among the 10 ESs. Specifically, 80% of each private dataset $\mathcal{D}_i$ has been allocated for training purposes, serving as the main input for the model optimization process. Subsequently, the remaining 20% has been earmarked for testing, thereby facilitating the rigorous evaluation of the model's performance and generalization capabilities on unseen data.

In order to realize the consensus procedure, the 10 Base Stations have been linked together so that each one of them is connected to other two, realizing the ring structure shown in Fig. 2. Anyway, our approach seamlessly works with different network topologies, with the condition that the network graph is connected (i.e., there are no separate node subsets) for every federation round. This may require additional neighbor discovery process in case in a certain round one or more nodes become disconnected from the network graph.

In what follows performance are evaluated using the F1-score, a performance metric that balances precision and recall using their armonic mean. Formally, let:

- TP be the number of correctly predicted anomalies.
- FP be the number of incorrectly predicted anomalies
- FN be the number of incorrectly predicted normal samples

The precision P is defined by

$$P = \frac{\text{TP}}{\text{TP} + \text{FP}}, \tag{6}$$

whereas the recall R by

$$R = \frac{\text{TP}}{\text{TP} + \text{FN}}. \tag{7}$$

The F1-score is defined as

$$F_1 = 2\frac{PR}{P + R}. \tag{8}$$

The F1-score ranges from 0 to 1, where a higher value indicates a better balance between precision and recall.

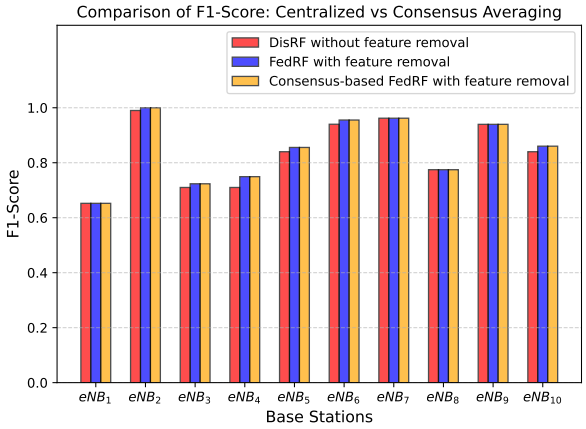

Fig. 3. Centralized vs Consensus. Nothing changes.

B. Equivalence between Centralized Averaging and Consensus

The histogram presented in Fig. 3 compares the COD F1-scores across the 10 BSs with three different methods: (i) no cooperation among base stations (DisRF - red), (ii) standard mathematical averaging, i.e. the approach proposed in our previous work [22] (FedRF - blue), and (iii) consensus-based FedRF (orange), namely the procedure described in Section III. These results have been obtained without any attack from malicious nodes. As expected, the F1-scores for consensus-based averaging and standard mathematical averaging are identical across all base stations. This equivalence highlights that consensus averaging effectively matches the performance of traditional mathematical averaging by a central authority, confirming its validity and reliability as a distributed computation approach.

Conversely, the F1-scores in the no-cooperation scenario (red bars) are consistently lower for certain base stations. This reduction underscores the critical role of cooperation and information sharing among base stations in improving fault detection performance.

C. Attack Policy

Now it is supposed that some nodes may be malicious ones, injecting false data about feature importance scores throughout the federated learning process. The attack policy considered in this work is as follows. Suppose F features are used at a specific round of the federated process. The attacker computes the feature importance vector f_{score} on its dataset. Then, he picks

$$i^* = \text{argmax} f_{\text{score}}$$

The attack vector f_{attack} is set in such a way that its i-th component is equal to:

$$f^i_{\text{attack}} = \begin{cases} \beta/F, & \text{if } i = i^* \\ \dfrac{1 - \beta/F}{F - 1}, & \text{otherwise.} \end{cases} \tag{9}$$

where β is a free attack parameter available to the malicious node that must be chosen in such a way that the importance assigned to i^* is lower than all the other ones. By simple calculations, this is obtained enforcing

$$\beta < \frac{F}{F - 2}. \tag{10}$$

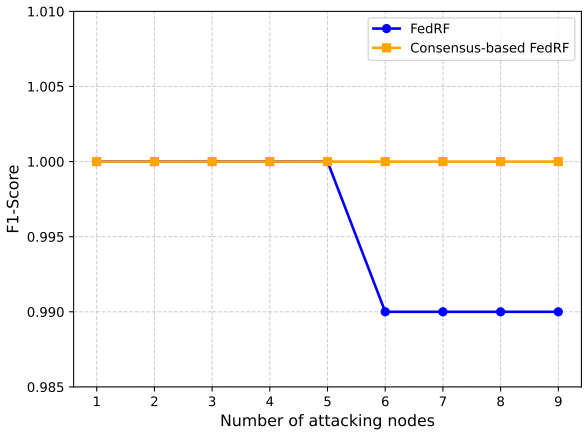

Fig. 4. Variation of F1-score on BS_2 against increasing number of attackers.

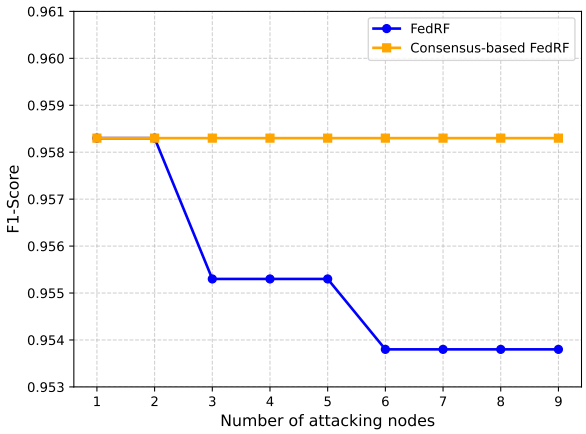

Fig. 5. Variation of F1-score on BS_6 against increasing number of attackers.

In this way, the attacker assigns the lowest score to the most important feature upon his dataset. Note that, by Eq. (9), the sum of the attack vector remains equal to 1, as if its values represent real importance scores.

D. Resilience against attacks through consensus-based COD

In this section we evaluate the performance of BS_2 and BS_6, in case of the presence of attacking nodes. However, similar considerations could be made referring to any of the other BSs. The simulation results depicted in Fig. 4 illustrate the degradation in system performance of BS_2 under an increasing number of attacking nodes. Performance is evaluated again using the F1-score in two scenarios: (i) FedRF with standard averaging [22] and (ii) consensus-based FedRF. The F1 score has been chosen because it is the most representative metric in binary decision problems. However, comparable results have been obtained with metrics such as precision and recall..

As the number of attacking nodes increases from 1 to 9, the F1-score decreases from 6 attackers on using centralized averaging, whereas performance is not affected even with 9 attackers in the case of the proposed consensus algorithm. Hence, the centralized approach exhibits a pronounced decline in F1-score, reflecting its inherent vulnerability to malicious interference. In contrast, the consensus-based averaging method demonstrates maximum resilience, keeping the F1-score invariant with respect to the number of attacking nodes.

Similar patterns are observed for BS_6, as shown in Fig. 5. In this case, relying on a centralized averaging approach, the F1-score begins to drop with 2 attackers, whereas again the consensus averaging strategy maintains full stability, reinforcing its suitability for adversarial scenarios across multiple base stations.

Eventually, Fig. 6 demonstrates the working principle of the consensus algorithm in relation with the feature maxUE_UL and in presence of 3 attacking nodes and 7 normal ones. It is possible to appreciate how the proposed algorithm lets the normal nodes reach consensus avoiding the false injection issued by the attackers.

V. Conclusion

In this work, we proposed a consensus-based feature removal mechanism to enhance cell outage detection in cellular networks by addressing the vulnerabilities of traditional centralized approaches. Our method improves upon standard centralized averaging by eliminating the single point of failure and incorporating a robust consensus process that mitigates the impact of malicious nodes. Specifically, our approach prevents adversarial participants from manipulating feature importance

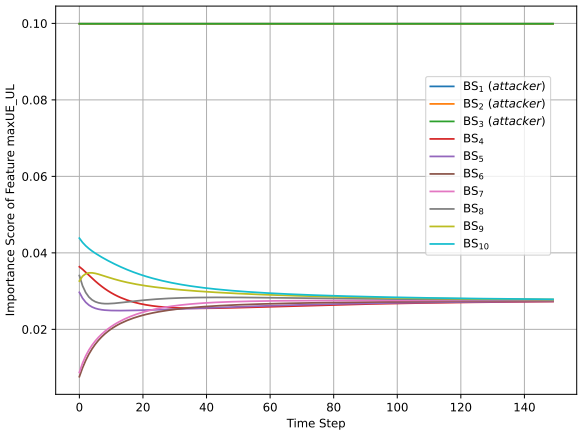

Fig. 6. Time evolution of the feature maxUE_UL during the consensus procedure with 3 attackers and 7 normal nodes.

TABLE I. Features of the COD dataset.

Feature	Description	Measurement Unit
Time	Hour of the day	(hh:mm)
CellName	Unique identifier for BS and cell	Text String
PRBUsageUL	Resource utilization in uplink	%
PRBUsageDL	Resource utilization in downlink	%
meanThr_UL	Average carried traffic in uplink	Mbps
meanThr_DL	Average carried traffic in downlink	Mbps
maxThr_UL	Maximum carried traffic in uplink	Mbps
maxThr_DL	Maximum carried traffic in downlink	Mbps
meanUE_UL	Average number of active UEs in uplink	-
meanUE_DL	Average number of active UEs in downlink	-
maxUE_UL	Maximum number of active UEs in uplink	-
maxUE_DL	Maximum number of active UEs in downlink	-
maxUE_UL+DL	Maximum number of total active UEs	-
Anomaly	Label for supervised learning	0 or 1

scores to mislead other nodes into removing critical features, thereby preserving the integrity of the anomaly detection process. Through extensive simulations, we demonstrated that our mechanism outperforms centralized strategies in terms of resilience, robustness, and accuracy. By leveraging a decentralized consensus approach, our method ensures that feature removal are made collectively within the cellular network, reducing the risk of biased or compromised outcomes. The results highlight that our system can effectively withstand adversarial influence while maintaining a high detection rate for cell anomalies.

Future works may extend the COD problem considering multi-class anomaly detection, thus moving towards the problem of anomaly diagnosis. Other research directions may focus on extending our framework to accommodate more sophisticated adversarial models, improving the efficiency of the consensus process, and integrating adaptive weighting strategies to further enhance the robustness against varying levels of adversarial behavior.

ACKNOWLEDGMENT

This work has been supported by the EU in the scope of the NANCY project, funded by Smart Networks and Services Joint Undertaking (SNS JU) under the European Union's Horizon Europe research and innovation programme, Grant Agreement No 101096456. This work has been also co-funded by the European Union - Next Generation Eu - under the National Recovery and Resilience Plan (NRRP), Mission 4 Component 1 Investment 4.1 - Decree No. 118 (2nd March 2023) of Italian Ministry of University and Research - Concession Decree No. 2333 (22nd December 2023) of the Italian Ministry of University and Research, Project code D93C23000450005, within the Italian National Program PhD Programme in Autonomous Systems (DAuSy).

REFERENCES

[1] V. S. Pana, O. P. Babalola, and V. Balyan, "5g radio access networks: A survey," *Array*, vol. 14, p. 100170, 2022.

[2] D. Mulvey, C. H. Foh, M. A. Imran, and R. Tafazolli, "Cell fault management using machine learning techniques," *IEEE Access*, vol. 7, pp. 124 514–124 539, 2019.

[3] A. K. Sangaiah, S. Rezaei, A. Javadpour, F. Miri, W. Zhang, and D. Wang, "Automatic fault detection and diagnosis in cellular networks and beyond 5g: Intelligent network management," *Algorithms*, vol. 15, no. 11, p. 432, 2022.

[4] W. Wang, Q. Liao, and Q. Zhang, "Cod: A cooperative cell outage detection architecture for self-organizing femtocell networks," *IEEE Transactions on Wireless Communications*, vol. 13, no. 11, pp. 6007–6014, 2014.

[5] I. de-la Bandera, R. Barco, P. Munoz, and I. Serrano, "Cell outage detection based on handover statistics," *IEEE Communications Letters*, vol. 19, no. 7, pp. 1189–1192, 2015.

[6] C. M. Mueller, M. Kaschub, C. Blankenhorn, and S. Wanke, "A cell outage detection algorithm using neighbor cell list reports," in *Self-Organizing Systems: Third International Workshop, IWSOS 2008, Vienna, Austria, December 10-12, 2008. Proceedings 3*. Springer, 2008, pp. 218–229.

[7] Q. Liao, M. Wiczanowski, and S. Stańczak, "Toward cell outage detection with composite hypothesis testing," in *2012 IEEE international conference on communications (ICC)*. IEEE, 2012, pp. 4883–4887.

[8] R. M. Khanafer, B. Solana, J. Triola, R. Barco, L. Moltsen, Z. Altman, and P. Lazaro, "Automated diagnosis for umts networks using bayesian network approach," *IEEE Transactions on vehicular technology*, vol. 57, no. 4, pp. 2451–2461, 2008.

[9] M. Alias, N. Saxena, and A. Roy, "Efficient cell outage detection in 5g hetnets using hidden markov model," *IEEE Communications Letters*, vol. 20, no. 3, pp. 562–565, 2016.

[10] O. Onireti, A. Zoha, J. Moysen, A. Imran, L. Giupponi, M. A. Imran, and A. Abu-Dayya, "A cell outage management framework for dense heterogeneous networks," *IEEE Transactions on Vehicular Technology*, vol. 65, no. 4, pp. 2097–2113, 2015.

[11] L. Bodrog, M. Kajo, S. Kocsis, and B. Schultz, "A robust algorithm for anomaly detection in mobile networks," in *2016 IEEE 27th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC)*. IEEE, 2016, pp. 1–6.

[12] W. Xue, M. Peng, Y. Ma, and H. Zhang, "Classification-based approach for cell outage detection in self-healing heterogeneous networks," in *2014 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 2014, pp. 2822–2826.

[13] P.-C. Lin, "Large-scale and high-dimensional cell outage detection in 5g self-organizing networks," in *2019 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2019, pp. 8–12.

[14] Y.-H. Ping and P.-C. Lin, "Cell outage detection using deep convolutional autoencoder in mobile communication networks," in *2020 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2020, pp. 1557–1560.

[15] B. Hussain, Q. Du, S. Zhang, A. Imran, and M. A. Imran, "Mobile edge computing-based data-driven deep learning framework for anomaly detection," *IEEE Access*, vol. 7, pp. 137 656–137 667, 2019.

[16] W. H. L. Pinaya, S. Vieira, R. Garcia-Dias, and A. Mechelli, "Convolutional neural networks," in *Machine learning*. Elsevier, 2020, pp. 173–191.

[17] M. H. Mahmoud, A. Albaseer, M. Abdallah, and N. Al-Dhahir, "Federated learning resource optimization and client selection for total energy minimization under outage, latency, and bandwidth constraints with partial or no csi," *IEEE Open Journal of the Communications Society*, vol. 4, p. 936–953, 2023. [Online]. Available: http://dx.doi.org/10.1109/OJCOMS.2023.3263962

[18] D. Mulvey, C. H. Foh, M. A. Imran, and R. Tafazolli, "Cellular network antenna tilt anomaly detection using federated unsupervised learning," in *ICC 2023 - IEEE International Conference on Communications*. IEEE, May 2023, p. 3048–3053. [Online]. Available: http://dx.doi.org/10.1109/ICC45041.2023.10279460

[19] E. T. Martínez Beltrán, M. Q. Pérez, P. M. S. Sánchez, S. L. Bernal, G. Bovet, M. G. Pérez, G. M. Pérez, and A. H. Celdrán, "Decentralized federated learning: Fundamentals, state of the art, frameworks, trends, and challenges," *IEEE Communications Surveys & Tutorials*, vol. 25, no. 4, p. 2983–3013, 2023. [Online]. Available: http://dx.doi.org/10.1109/COMST.2023.3315746

[20] H. Ye, L. Liang, and G. Y. Li, "Decentralized federated learning with unreliable communications," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 3, p. 487–500, Apr. 2022. [Online]. Available: http://dx.doi.org/10.1109/JSTSP.2022.3152445

[21] Y. Li, C. Chen, N. Liu, H. Huang, Z. Zheng, and Q. Yan, "A blockchain-based decentralized federated learning framework with committee consensus," *IEEE Network*, vol. 35, no. 1, p. 234–241, Jan. 2021. [Online]. Available: http://dx.doi.org/10.1109/MNET.011.2000263

[22] A. Wrona, S. Gentile, E. De Santis, A. Giuseppi, A. Pietrabissa, and F. Delli Priscoli, "A cooperative feature removal mechanism for cell outage detection in wireless telecommunication networks," in *International Conference on Critical Information Infrastructures Security*. Springer, 2024, pp. 84–95.

[23] F. Cacace, M. Mattioni, S. Monaco, and L. R. Celsi, "Output containment via multiconsensus for heterogeneous linear systems on digraphs," *IEEE Transactions on Control of Network Systems*, vol. 11, no. 2, pp. 1012–1023, 2023.

[24] F. Cacace, M. Mattioni, S. Monaco, and D. Normand-Cyrot, "Consensus and multi-consensus for discrete-time lti systems," *Automatica*, vol. 166, p. 111718, 2024.

[25] J. Vidal, "Anomaly detection in cellular networks," 2020. [Online]. Available: https://kaggle.com/competitions/anomaly-detection-in-4g-cellular-networks