

EXPLAINABLE CROSS-LAYER MACHINE LEARNING FOR ENTERPRISE WLAN AUDITING

Viktor Stoynov

Faculty of Telecommunications, Technical University of Sofia, Sofia
Bulgaria

* Corresponding Author, e-mail: vstoynov@tu-sofia.bg

Abstract: This paper presents an explainable cross-layer machine learning method for auditing enterprise and campus wireless local area networks. It combines passive radio and Medium Access Control evidence with service, mobility, security and configuration indicators to identify anomalous windows, infer root causes and generate audit scores for operator review. The learning process uses unsupervised anomaly admission and supervised root-cause classification, while the reporting layer produces evidence-supported recommendations. Evaluation is conducted with CrossAudit-WLANSim, a Python-based discrete-time simulator configured for three enterprise sites, twelve access points and 69120 audit windows. Under site drift, the drift-aware configuration reaches an anomaly F1-score of 0.770 and an eight-state macro-F1 of 0.855, outperforming a network-only baseline and reducing false-positive alarms.

Key words: wlan audit, explainable machine learning, cross-layer telemetry, concept drift.

1. INTRODUCTION

Enterprise and campus wireless local area networks (WLANs) have become an essential access technology for institutional communication, education, administration and everyday digital services. Their operation is governed by the IEEE 802.11 family of standards, but their practical assessment cannot be reduced only to protocol conformance or nominal access point (AP) capacity [1]. In dense deployments, service quality depends on radio conditions, contention, retransmissions, client density, roaming behaviour, service-plane delays, security events and configuration consistency. These requirements make WLAN assurance a cross-layer problem rather than a simple monitoring task.

A central difficulty is that the same user-visible symptom can originate from several technical causes. Low throughput may be caused by co-channel interference, excessive channel occupation, AP overload, poor roaming, Dynamic Host Configuration Protocol (DHCP) delay, Domain Name System (DNS) delay or management-frame attack activity. The contention behaviour of IEEE 802.11 further shows that medium access conditions can strongly affect perceived performance even when the radio link is

available [2]. Recent work confirms that machine learning (ML) can support IEEE 802.11 performance analysis and WLAN state detection, but many studies address narrower objectives such as throughput prediction, attack detection or change-point identification [3,4].

The distinction between monitoring and auditing is important for this paper. Monitoring describes the collection of counters, alarms and logs. Auditing evaluates whether the observed state is acceptable with respect to performance, security and configuration compliance. Modern network-monitoring studies show that richer telemetry and multi-indicator representations are more informative than isolated counters for anomaly detection and root-cause reasoning [5,6]. However, an audit result must also be defensible, comparable across sites and linked to evidence that can be inspected by a network operator. In this sense, a WLAN audit should not only report that an anomaly exists; it should also explain which indicators support the decision and how the decision affects operational risk.

The increasing volume of controller telemetry and the heterogeneity of WLAN indicators motivate data-driven audit methods. However, direct application of ML to operational telemetry is not sufficient. A model trained on one site can produce excessive false positives when applied to another site with different client density, utilisation profile or radio environment. In addition, a high classification score alone does not provide the type of evidence normally required in an audit report. Explainable ML is therefore relevant because the model output has to be interpreted by engineers, reviewers or security personnel rather than consumed as an opaque label [7-9]. These limitations require a design in which evidence construction, anomaly admission, root-cause inference, explanation and audit scoring are treated as separate but connected functions.

This paper proposes CrossAudit-WLAN as an explainable ML method for WLAN audit. The method combines passive radio and Medium Access Control (MAC) evidence with service-plane, mobility, security and configuration indicators. It admits anomalous audit windows, assigns the most likely operational cause, and produces a risk-aware audit score together with an explanation bundle. The objective is not to replace specialised controller diagnostics, but to provide a compact and reproducible assessment layer that can support technical reports and comparative evaluation of WLAN sites.

The main contribution is threefold. First, a six-stage audit pipeline is defined in which evidence acquisition and feature construction are separated from the learning stages. Second, an audit scoring mechanism is introduced to combine performance, security and configuration-compliance dimensions without forcing all decisions into a single classifier output. Third, the method is evaluated under site drift using a repeatable simulation environment, feature ablation, window-size sensitivity checks and explanation-fidelity analysis. The evaluation follows conservative splitting principles motivated by concept drift, leakage-aware wireless experiments and systematic anomaly-detection testing in time series [10-12].

The remainder of the paper is organized as follows. Section 2 discusses related work and motivates the audit-oriented design. Section 3 presents the proposed cross-layer audit method, its evidence model, learning stages and audit scoring mechanism. Section 4 describes the simulation environment and experimental protocol. Section 5 analyses

the obtained results, including ablation and interpretability checks. Section 6 concludes the paper and outlines future validation steps.

2. RELATED WORK AND MOTIVATION

Machine learning has been extensively investigated for improving IEEE 802.11 operation. Existing studies cover channel selection, rate adaptation, performance prediction, interference recognition, mobility optimisation and anomaly detection. The survey in [3] shows that ML-based WLAN optimisation can exploit radio and MAC indicators that are difficult to interpret by fixed threshold rules. Change-point detection approaches, including random-forest-based methods, further show that deviations in WLAN behaviour can be detected from temporally structured observations [4]. These results are important, but they do not directly solve the audit problem, because audit requires an explicit connection between a decision, its supporting evidence and its operational meaning.

Network telemetry research has moved from isolated counters to richer representations of network state. Programmable monitoring and in-band telemetry studies demonstrate that combining multiple indicators can improve anomaly detection and root-cause analysis [5,6]. For WLANs, an analogous logic is needed even when the telemetry is not collected through programmable data-plane mechanisms. Radio quality, retry behaviour, channel utilisation, client density, service delays, management events and configuration warnings all describe different aspects of the same operational state. Treating these indicators separately often leads to incomplete conclusions.

Recent work on multiclass IoT anomaly classification has shown that resampling strategies and ensemble classifiers can improve predictive performance under imbalanced conditions [13]. This is relevant for WLAN audit, where normal windows dominate the data and anomalous operational states must still be separated into meaningful root-cause classes. For an audit-oriented system, the transparency of the decision path is as important as the numerical accuracy of the model.

Explainable artificial intelligence (XAI) has become a necessary component when ML outputs are used in operational or security-sensitive environments. In networking, explanations can be used to identify which features contributed to an anomaly decision and to provide evidence for further verification [7]. Broader XAI surveys emphasise that explanation techniques should be selected according to the decision context, the user of the explanation and the required level of accountability [8,9]. For WLAN audit, the primary user is usually an operator, engineer or reviewer who needs compact evidence rather than a full model description.

Security and configuration aspects require separate attention. Wi-Fi attack detection studies demonstrate that management-frame abuse and authentication anomalies can be recognised by data-driven methods [14]. Misconfigured bring-your-own-device environments and configuration inconsistencies also create measurable degradation or risk [15]. Related work on misconfiguration in open radio access networks shows that AI/ML components may themselves be affected by configuration errors and policy mismatches [16]. These observations justify the inclusion of security and configuration evidence as first-class audit inputs rather than as optional annotations.

Concept drift is another important consideration. WLAN behaviour changes with the site, time of day, device mix, user density and local radio conditions. Models trained on temporally adjacent or randomly split samples can appear accurate while actually benefiting from leakage between training and test observations. Drift and leakage concerns have been discussed in the broader ML and wireless sensing literature [10,11], and systematic anomaly-detection evaluation in time series also stresses the importance of realistic data separation [12]. For this reason, the experiments in this paper use both site separation and chronological separation.

From a methodological perspective, Isolation Forest provides an efficient unsupervised mechanism for anomaly detection in tabular telemetry [17]. Tree ensembles are suitable for heterogeneous numerical and Boolean features [18], while SHAP-based local feature attributions can support operator-facing explanations [19]. Recent studies on Wi-Fi performance prediction and AI-driven telemetry processing further motivate the adopted cross-layer feature representation [20,21].

The proposed method is positioned between monitoring, intrusion detection and assurance reporting. It does not aim to replace low-level WLAN simulators or vendor-specific management tools. Instead, it formalises the audit process as a chain of evidence construction, anomaly admission, cause interpretation, explanation and scoring. This positioning is useful in enterprise and campus environments, where the same infrastructure must be assessed from performance, security and compliance perspectives.

3. PROPOSED CROSS-LAYER AUDIT METHOD

3.1. Audit window and evidence model

CrossAudit-WLAN operates on audit windows collected from APs, controller logs and optional packet-summary taps. For AP a and time window t , the input vector $x(a,t)$ contains radio, MAC, service, mobility, security and configuration evidence. The output is $o(a,t) = (y, c, P, S, C, Q, E, R)$, where y is the anomaly state, c is the inferred root cause, P , S and C are performance, security and configuration-compliance scores, Q is the global audit score, E is the explanation bundle and R is the remediation set.

The audit-window abstraction is deliberately practical. WLAN controllers and APs commonly expose counters and event summaries over fixed time intervals. Such intervals are also suitable for reporting, because an operator can compare the decision with logs, configuration changes and incident tickets for the same period. The proposed method therefore treats each window as a compact evidence object. It does not require packet capture of all traffic, although packet summaries can enrich the service and security indicators when available.

The evidence groups serve different diagnostic roles. Radio and MAC features describe signal quality, contention and retransmission behaviour. Service features describe the user-facing outcome, including throughput, latency, jitter and name-resolution or address-assignment delays. Mobility features represent roaming behaviour, while security features capture authentication failures, deauthentication bursts and scanning activity. Configuration features describe conditions such as channel overlap, disabled Protected Management Frames (PMF), Quality of Service (QoS) mismatch,

abnormal power variance and controller warnings. A single feature rarely proves a cause; the diagnostic value comes from the pattern across groups.

3.2. Pipeline organization and learning stages

Figure 1 summarizes the six-stage pipeline. Stage A acquires passive evidence from radio, MAC, service, mobility, security and management-related sources. Stage B converts counters into rates, deltas, rolling medians and AP-relative deviations. These two stages are not learning stages; they do not train or apply an ML model but construct the auditable feature space used by the subsequent methods. This separation is important because audit evidence should remain interpretable even when the learning model is changed.




Figure 1. CrossAudit-WLAN staged audit pipeline and operator-facing outputs.

Stage C is the first explicit ML stage and performs unsupervised anomaly admission. In the experiments, Stage C uses Isolation Forest for unsupervised anomaly admission. The purpose of this stage is not to provide a final cause, but to decide whether a window deserves further audit attention. This design reduces the number of windows passed to the more detailed root-cause stage and reflects the operational reality that most WLAN windows are normal or near-normal.

Stage D applies supervised ML to admitted anomalous windows and assigns a primary operational class, such as congestion, interference, AP overload, roaming instability, misconfiguration, security event or mixed condition. The primary class represents the dominant interpretation of the audit window, while secondary indicators remain available through the explanation bundle and the audit score. The mixed-condition class accounts for cases in which several degradation factors occur simultaneously. A tree-ensemble classifier is used because it can process heterogeneous numerical and Boolean telemetry and provide local feature attributions.

Stage E packages the decision into an explanation bundle using local feature attributions and a compact operator-facing evidence statement. The explanation bundle is intended to be readable by an operator: it lists the most influential indicators, their direction of deviation and the associated audit interpretation. Stage F implements the audit scoring and remediation layer. This final stage converts the model output into a reporting form, because an audit decision must state not only what was detected, but also how severe the detected state is and which evidence supports the recommended action.

3.3. Risk-aware audit scoring

The audit scoring mechanism is not a classifier but a reporting layer:

$$Q(a, t) = 100 \left(w_p P(a, t) + w_s S(a, t) + w_c C(a, t) \right) - \lambda A(a, t) - \mu r(c(a, t)) \quad (1)$$

where all terms are evaluated for AP a and time window t , $w_p + w_s + w_c = 1$, A is anomaly confidence, $r(c)$ is a class-risk penalty, and λ and μ control the anomaly and class-risk penalties. Thus, an organization can emphasize performance, security or compliance without changing the ML models. For example, an educational campus may prioritise service continuity during examinations, whereas an enterprise handling sensitive information may assign a stronger penalty to security-related events and PMF misconfiguration.

The design is inherently cross-layer. Low throughput, high retry rate and poor signal quality indicate radio degradation, whereas low throughput combined with many clients and high channel-busy time is more consistent with congestion or AP overload. Deauthentication bursts, authentication failures and increased scanning activity represent security-related indicators, while a QoS mismatch together with channel-overlap flags points to configuration risk. Therefore, the role of the ML method is to identify relationships among symptoms observed across different protocol layers, rather than to perform simple packet-level classification.

The scoring layer also avoids a common weakness of purely classification-based assurance. A class may be easy to identify but operationally severe, or difficult to identify but moderate in its expected impact. CrossAudit-WLAN therefore separates recognition quality from audit acceptability. This allows the method to report, for example, that a security event has high classification confidence and low audit acceptability, while a moderate congestion interval may receive a lower risk penalty even when it is detected with similar confidence.

4. EXPERIMENTAL SIMULATOR AND CONFIGURATION

The experiments were conducted with CrossAudit-WLANSim, a Python-based discrete-time simulator developed specifically for this study to support repeatable WLAN audit experiments. It does not reproduce all IEEE 802.11 physical-layer (PHY) details. Instead, it generates audit-level telemetry windows linking radio state, MAC contention, service delay, mobility, management events and configuration state. This abstraction is appropriate because the objective is to assess the audit and inference chain rather than to reproduce a complete PHY model.

The simulator consists of four modules. The traffic module generates daily WLAN activity with diurnal variation, changing client presence and site-specific baselines. The scenario module introduces congestion, interference, AP overload, roaming instability, misconfiguration, security events and mixed conditions. The drift module shifts utilisation, retry rate, client diversity and service timing between training and test sites. Finally, the labelling module assigns labels and audit scores to each 15-min window. Random noise and overlapping intervals are added to prevent deterministic class separation.

The simulated environment comprises three managed WLAN sites, each with four APs and heterogeneous clients: laptops, smartphones and Internet of Things (IoT) devices. The 60-day observation period is divided into 15-min windows, yielding 69120 samples. Normal operation follows representative ranges for client density, channel utilisation, retry rate and service timing. Deviations introduce the seven operational

conditions. The ML method uses Isolation Forest for anomaly admission, a tree-ensemble root-cause classifier and feature attribution for explanations.

The training protocol used both site and chronological separation. Sites A and B were used for training on days 1-40 and for validation on days 41-60. Site C was kept outside the training process and was used only for calibration on days 41-44 and final testing on days 45-60. Random splitting was avoided to reduce leakage from repeated AP observations and temporally adjacent windows. Drift-aware recalibration updates only the robust-scaling parameters and the anomaly threshold, leaving all final test windows unseen.

The feature vector contains thirty-one variables grouped into radio/MAC, service, mobility, security and configuration categories. The radio/MAC group covers signal quality, retry rate, frame-error rate, channel utilisation and client count, while the service group covers throughput, latency, jitter, DHCP delay and DNS delay. The remaining groups describe roaming delay, authentication failures, deauthentication rate, scan rate, channel-overlap flag, PMF disabled flag, QoS mismatch, power variance and controller warnings. The selected variables provide a cross-layer representation of audit-level WLAN conditions.

Three design choices should be noted. First, the simulator intentionally introduces overlapping symptoms, because real WLAN incidents rarely affect one indicator only. Second, site C is not used as a conventional training site, which makes the test more conservative than a random split. Third, recalibration is limited to scaling and anomaly-threshold adjustment. The root-cause classifier is not retrained on final test windows, so the reported results represent transfer to a shifted site rather than memorisation of the test distribution.

5. ANALYSIS OF EXPERIMENTAL RESULTS

Table 1 reports the main detection and classification results. Dash symbols mark metrics that are not applicable to the corresponding evaluation level or were not reported for that experiment. For binary anomaly admission, the relevant measures are precision, recall, F1-score and false-positive rate (FPR). For the eight-state decision, macro-F1 and accuracy provide the main multiclass summary. The static-threshold version of CrossAudit-WLAN preserved high recall under site drift, but it also produced a high FPR. Drift-aware recalibration reduced FPR from 0.230 to 0.093 and improved anomaly F1-score from 0.724 to 0.770. In comparison, the network-only baseline achieved 0.648 F1-score and 0.279 FPR, which shows that service indicators alone are insufficient for reliable WLAN audit.

The end-to-end eight-state decision also improved. The full cross-layer model achieved a macro-F1 of 0.855 and an overall accuracy of 0.883 on the held-out site and time split, whereas the network-only baseline reached a macro-F1 of 0.562. This improvement is particularly relevant for states with similar service-level symptoms but different cross-layer patterns. Congestion, interference and AP overload all degrade service quality, but they differ in channel-busy time, retry behaviour, signal quality, client density and controller warnings.

Table 1. Main anomaly-admission and eight-state audit performance on the held-out site and time split.

Method	Evaluation level	Precision	Recall	F1 / macro-F1	FPR	Accuracy
Network-only baseline	Anomaly admission	0.508	0.895	0.648	0.279	-
CrossAudit-WLAN, static threshold	Anomaly admission	0.576	0.973	0.724	0.230	-
CrossAudit-WLAN, drift-aware threshold	Anomaly admission	0.736	0.808	0.770	0.093	-
Network-only baseline	Eight-state decision	-	-	0.562	-	-
CrossAudit-WLAN, full cross-layer model	Eight-state decision	-	-	0.855	-	0.883

Figure 2 presents two complementary views. The left part reports scenario-level root-cause F1-score. Misconfiguration, security events and mixed conditions are recognized most reliably, as their signatures are abrupt and concentrated in security or configuration indicators. Roaming instability is also well separated due to distinctive mobility-delay and roaming-failure patterns. Congestion and AP overload remain more difficult, since both involve many clients, increased latency and high utilisation.




Figure 2. Scenario-level root-cause F1-score and risk-aware audit score in the held-out test set.

The right part of Fig. 2 reports the audit scores, which reflect operational acceptability rather than classification quality. Thus, a security event may obtain a high F1-score because its pattern is easy to recognize, but a low audit score because it represents a severe assurance violation. A mixed condition follows the same logic: it is highly separable yet receives the lowest score because several assurance pillars degrade at once. Therefore, Fig. 2 separates model performance on the left from WLAN risk assessment on the right. The Spearman correlation between scenario severity and audit score was 0.898.

Table 2 summarizes the additional experimental checks. Removing radio and MAC features reduced macro-F1 to 0.731, mainly degrading interference recognition. Removing security and configuration features reduced macro-F1 to 0.764 and made misconfiguration and management-plane events harder to separate from generic degradation. A 5-min window improved responsiveness but increased false positives, whereas a 30-min window delayed admission. The 15-min window therefore remains a practical audit compromise.

Table 2. Additional experimental findings on the held-out site C.

<i>Experiment</i>	<i>Main quantitative result</i>	<i>Interpretation</i>
<i>Full CrossAudit-WLAN, drift-aware</i>	<i>Anomaly F1 = 0.770; macro-F1 = 0.855; FPR = 0.093</i>	<i>Reference configuration; balanced sensitivity and false-positive control</i>
<i>Static threshold without recalibration</i>	<i>Recall = 0.973; FPR = 0.230</i>	<i>Sensitive but assigns too many anomalies to shifted site-C windows</i>
<i>Network-only baseline</i>	<i>Anomaly F1 = 0.648; macro-F1 = 0.562</i>	<i>Service counters alone do not separate radio, security and configuration causes</i>
<i>Radio/MAC feature ablation</i>	<i>Macro-F1 = 0.731</i>	<i>Interference and congestion become less distinguishable</i>
<i>Security/configuration ablation</i>	<i>Macro-F1 = 0.764</i>	<i>Misconfiguration and management-plane events lose their diagnostic pattern</i>

The attribution analysis confirms the interpretability of the audit output. In sampled anomalous windows, 73% of cases had at least two of the three top-ranked features corresponding to the injected cause. Security events were mainly associated with deauthentication rate, authentication failures and scan rate. Misconfigurations were characterized by channel-overlap flag, QoS mismatch and power variance, whereas roaming instability was captured by roaming delay, roaming-failure ratio and DHCP delay. These patterns do not demonstrate causality but provide evidence verifiable against logs or configuration records.

Runtime remained low, with a mean inference time of 0.195 ms per audit window. With the two-of-three window rule, median admission latency was 15 min and mean latency was 21.6 min. Overall, 93.3% of injected incidents were admitted before the scenario ended. These values indicate that the proposed method is suitable for periodic audit processing and near-real-time operator review, although production deployment would still require integration with a controller, log collector or WLAN assurance platform.

The results are simulation-based evidence, not full production validation. CrossAudit-WLANSim makes the experiments repeatable and allows controlled scenario variation. However, it cannot represent all controller behaviour, firmware differences or site-specific policies in real WLAN deployments. The contribution is methodological. Results show that cross-layer evidence, drift-aware thresholding and the

audit scoring mechanism can form a coherent WLAN audit method. Future work will extend validation with campus WLAN traces, operator-confirmed labels and remediation rules adapted to specific controller platforms.

6. CONCLUSION

This paper formulated the operational assessment of enterprise WLANs as a cross-layer audit problem. CrossAudit-WLAN converts passive telemetry into anomaly decisions, root-cause labels, audit scores and remediation evidence. Its design separates evidence construction from the ML stages, which makes the audit process easier to inspect and adapt. The method also separates model recognition from audit acceptability, so that the operational severity of an event can be reported independently from the confidence of the classifier.

The simulation-based evaluation demonstrates the value of combining radio/MAC, service, mobility, security and configuration evidence. Under site drift, the drift-aware configuration reduced the false-positive rate and improved anomaly F1-score in comparison with the static-threshold variant. The full cross-layer model also achieved a substantially higher eight-state macro-F1 than the network-only baseline. Ablation checks confirmed that radio/MAC evidence is important for distinguishing interference and congestion, while security and configuration features are necessary for recognising management-plane events and misconfiguration.

The main limitation is the remaining ambiguity between congestion and AP overload, which share several service-level and utilisation-related symptoms. Another limitation is that the current evaluation is based on a controlled simulator, although the protocol includes site drift and leakage-aware splitting. Future work will therefore focus on validation with real campus WLAN traces, operator-confirmed incident labels and controller-specific remediation rules. Further development will also consider adaptive scoring policies, additional explanation formats and integration with WLAN assurance dashboards used in operational environments.

ACKNOWLEDGEMENT

This work was supported by the research project KII-06-KOCT/13/22.05.2024 "COST Action CA22168" of the Bulgarian National Science Fund.

REFERENCES

- [1] IEEE. IEEE Standard for Information Technology - Wireless LAN MAC and PHY Specifications. IEEE Std 802.11-2020, IEEE, New York, USA, 2021, 4379 p. DOI: 10.1109/IEEESTD.2021.9363693.
- [2] Bianchi, G. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE Journal on Selected Areas in Communications*, vol.18, no.3, 2000, pp. 535-547. DOI: 10.1109/49.840210.
- [3] Szott, S., Kosek-Szott, K., Gawlowicz, P., Gozdecki, J., Loziak, K., Natkaniec, M., Prasnal, L., Szott, M., Wagrowski, M. Wi-Fi meets ML: A survey on improving IEEE 802.11 performance with machine learning. *IEEE Communications Surveys and Tutorials*, vol.24, no.3, 2022, pp. 1843-1893. DOI: 10.1109/COMST.2022.3177584.

- [4] Huet, A., Krolkowski, J., Navarro, J.M., Chen, F., Rossi, D. Change point detection in WLANs with Random AP Forests. *Companion of the 19th International Conference on Emerging Networking Experiments and Technologies*, Paris, France, December 2023, pp. 29-36. DOI: 10.1145/3624354.3630587.
- [5] Kravets, O. Ja., Aksenov, I.A., Redkin, Yu.V., Rahman, P.A., Kochegarov, M.V., Gorshkov, A.V., Sorokin, S.A. Modelling of neural network monitoring agent to predict traffic spikes and agent training. *International Journal on Information Technologies and Security*, vol.16, no.3, 2024, pp. 49-56. DOI: 10.59035/VDTR2117.
- [6] Duan, Y., Li, C., Bai, G., Chen, G., Zhou, F., Chen, J., Gao, Z., Zhang, C. MFGAD-INT: In-band network telemetry data-driven anomaly detection using multi-feature fusion graph deep learning. *Journal of Cloud Computing: Advances, Systems and Applications*, vol.12, no.1, 2023, 126. DOI: 10.1186/s13677-023-00492-w.
- [7] Marantis, P., Ramantas, K., Verikoukis, C. Anomaly detection with quantified explanations for multivariate time series networking data. *2024 IEEE Global Communications Conference (GLOBECOM)*, Cape Town, South Africa, December 2024, pp. 4214-4219. DOI: 10.1109/GLOBECOM52923.2024.10901585.
- [8] Mersha, M., Lam, K., Wood, J. et al. Explainable artificial intelligence: A survey of needs, techniques, applications, and future direction. *Neurocomputing*, vol.599, 2024, 128111. DOI: 10.1016/j.neucom.2024.128111.
- [9] Sharma, A., Rani, S., Shabaz, M. A comprehensive review of explainable AI in cybersecurity: Decoding the black box. *ICT Express*, vol.11, no.6, 2025, pp. 1200-1219. DOI: 10.1016/j.icte.2025.10.004.
- [10] Hinder, F., Vaquet, V., Hammer, B. One or two things we know about concept drift: Detecting concept drift. *Frontiers in Artificial Intelligence*, vol.7, 2024, 1330257. DOI: 10.3389/frai.2024.1330257.
- [11] Varga, D. Mitigating data leakage in a WiFi CSI benchmark for human action recognition. *Sensors*, vol.24, no.24, 2024, 8201. DOI: 10.3390/s24248201.
- [12] Zhang, A., Deng, S., Cui, D., Yuan, Y., Wang, G. An experimental evaluation of anomaly detection in time series. *Proceedings of the VLDB Endowment*, vol.17, no.3, 2023, pp. 483-496. DOI: 10.14778/3632093.3632110.
- [13] Alharbi, F. A comparative study of SMOTE and ADASYN for multiclass classification of IoT anomalies. *International Journal on Information Technologies and Security*, vol.17, no.2, 2025, pp. 15-24. DOI: 10.59035/QEFU7977.
- [14] Moharam, M.H., Ashraf, K., Alaa, H. et al. Real-time detection of Wi-Fi attacks using hybrid deep learning models on NodeMCU. *Scientific Reports*, vol.15, no.1, 2025, 32544. DOI: 10.1038/s41598-025-18947-2.
- [15] Choi, J., Kim, T., Bahk, S. Detection of misconfigured BYOD devices in Wi-Fi networks. *Applied Sciences*, vol.10, no.20, 2020, 7203. DOI: 10.3390/app10207203.
- [16] Yungaicela-Naula, N.M., Pastor, A., Contreras, L.M. Misconfiguration in O-RAN: Analysis of the impact of AI/ML. *Computer Networks*, vol.252, 2024, 110455. DOI: 10.1016/j.comnet.2024.110455.

- [17] Liu, F.T., Ting, K.M., Zhou, Z.-H. Isolation Forest. *2008 Eighth IEEE International Conference on Data Mining (ICDM)*, Pisa, Italy, December 2008, pp. 413-422. DOI: 10.1109/ICDM.2008.17.
- [18] Breiman, L. Random forests. *Machine Learning*, vol.45, 2001, pp. 5-32. DOI: 10.1023/A:1010933404324.
- [19] Lundberg, S.M., Lee, S.-I. A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, vol.30, 2017, pp. 4765-4774.
- [20] Formis, G., Garlisi, D., Gringoli, F. Improving Wi-Fi network performance prediction with deep learning models. *arXiv preprint arXiv:2507.11168*, 2025. URL: <https://arxiv.org/abs/2507.11168> (Visited on 09.07.2026).
- [21] Foroughi, P., Brockners, F., Rougier, J.-L. ADT: AI-driven network telemetry processing on routers. *Computer Networks*, vol.220, 2023, 109474. DOI: 10.1016/j.comnet.2022.109474.

Information about the authors:

Viktor Stoynov is a Senior Assistant Professor at the Faculty of Telecommunications, Technical University of Sofia, Bulgaria. His research interests include wireless and mobile networks, 5G/6G communication systems, network robustness, physical-layer security, IoT and machine-learning-based network assurance.

Manuscript received on 09 July 2026