Can Machine Learning Break Wi-Fi Privacy? A Study on MAC Address Randomization

Marta Puig
Universitat Pompeu Fabra
Barcelona, Spain
martazheng.puig@upf.edu

Costas Michaelides
Universitat Pompeu Fabra
Barcelona, Spain
costas.michaelides@upf.edu

Lucia Pintor
University of Cagliari
Cagliari, Italy
lucia.pintor@unica.it

Boris Bellalta
Universitat Pompeu Fabra
Barcelona, Spain
boris.bellalta@upf.edu

Francesc Wilhelmi
Universitat Pompeu Fabra
Barcelona, Spain
francisco.wilhelmi@upf.edu

***Abstract*—Although IEEE 802.11 Medium Access Control (MAC) address randomization mitigates passive tracking, Machine Learning (ML)-based fingerprinting still exposes device identities. By combining unencrypted Probe Frame hardware specs, Inter-Probe Frame Arrival Time (IFAT), and Simulated Received Signal Strength Indication (SRSSI), we introduce a bitwise decomposition of the High Throughput (HT) capabilities information field to boost identification accuracy. Evaluating K-Means, DBSCAN, and OPTICS across 22 devices from six manufacturers, DBSCAN achieves up to 89.6% global accuracy using decomposed HT features and three SRSSI measurements. This demonstrates that existing MAC randomization is insufficient, highlighting the need for stronger Wi-Fi privacy standards.**

***Index Terms*—IEEE 802.11, MAC address randomization, privacy, machine learning, Wi-Fi**

I. Introduction

Wi-Fi clients, or Stations (STAs), discover available networks using two methods: passive scanning, where they listen for Beacon frames, or active scanning, which involves transmitting Probe Request frames and receiving Probe Response frames from nearby Access Points (APs). A critical vulnerability of active discovery is that its management frames are transmitted entirely unencrypted, making them a primary target for passive eavesdropping and device tracking [1].

To mitigate third-party tracking, vendors such as Apple (starting from iOS 8) and Microsoft (starting from Windows 10) introduced their own Medium Access Control (MAC) address randomization mechanisms before formal standardization. Later, to address interoperability issues, the IEEE formally addressed MAC address randomization in the 802.11aq amendment for Probe Requests during the pre-association stage [1]. Despite randomization, Probe Requests still contain static header fields that can be used to track a given user. Moreover, temporal information can be extracted from Inter-Probe Frame Arrival Time (IFAT) and spatial information can be derived from the Received Signal Strength Indication (RSSI), which further contribute to passively fingerprinting users even when the MAC address is randomized [2].

Unlike prior work focused on a single or a couple of features (Table I), this paper explores the robustness of MAC address randomization by combining Information Elements (IEs), IFAT, and Simulated Received Signal Strength Indication (SRSSI), which are made available to an eavesdropper using unsupervised clustering. We omit Sequence Number (SN) due to unpredictable increments and extend the dataset in [5], [6], testing 22 devices from six manufacturers across four feature configurations and four density levels (5, 10, 15, and 22). Our findings show current MAC randomization is insufficient, specifically: *i)* Bitwise decomposition of High Throughput (HT) capabilities information enhances device separation, facilitating tracking, *ii)* IFAT introduces variance that degrades clustering performance, *iii)* Spatial features (SRSSI) improve de-randomization accuracy over header-only data.

II. Proposed Methodology

A. Data Description

We consider the dataset introduced in [5], which consists of 20-minute captures from 22 different devices under six distinct operational modes. Among all possible configurations, only captures in "Mode S" (screen inactive, Wi-Fi enabled, power-saving disabled) were considered. Analyzing the dataset more deeply, we observed a significant variance in the number of frames transmitted. Despite all devices being captured under the same conditions, certain devices broadcast considerably more frames within those 20 minutes, ranging from 5 frames to over 1,200 frames. The information considered for MAC address de-randomization is as follows:

- **Decomposed HT Capabilities:** Unlike prior work, where all 28 bytes of the HT capabilities field (IE 45) are used as a single decimal value [6], we take only the *HT capabilities information*, the *A-MPDU parameters*, and the *HT extended capabilities* fields.

This work was supported by the following projects: TRUE Wi-Fi PID2024-155470NB-I00 (MICIU/AEI/10,13039/501100011033/FEDER,UE), ICREA Academia 2024 (00077 AGAUR), MdM CEX2021-001195-M (MICIU/AEI/10.13039/501100011033), and by NextGenerationEU (Italian NRRP, Mission 4, Component 2, Investment 1.2, CUP F23C25000440006).

TABLE I
MAC DE-RANDOMIZATION PRIOR ART. FEATURE COLUMNS: SN = SEQ. NUMBER, IE = INFORMATION ELEMENT, T = TIMING, R = RSSI.

Work	SN	IE	T	R	Environment & Dataset	Algorithm	Result
Freudiger [3]	✓	✓	×	×	Laboratory (3 channels). Labeled (True): Specific devices with known OS.	Linear	Proven SN leak
Praharenka et al. [4]	×	×	✓	×	Isolated (3 channels). Labeled (True): Manual labeling in a controlled environment.	NB/RF/SVM	94.5% / 93.2% / 90.4% Accuracy
Pintor et al. [5], [6]	×	✓	×	×	Isolated (3 channels). Labeled (True): Single-device captures merged into 12 scenarios.	DBSCAN	7.5% Average Error
Cifuentes-Urtubey et al. [7]	✓	✓	×	×	Conference venue (1 channel). Unlabeled: Real-world MobiCom 2023 traffic.	Heuristics	45% reduction in MAC IDs
Pérez-Hernández et al. [8]	×	✓	×	✓	Building (1 channel). Unlabeled: Ground truth via manual counting.	K-Means	98% Counting Accuracy
This work	×	✓	✓	✓	**Isolated (1 channel). Labeled (True): Single-device captures merged into 3 scenarios.**	**DBSCAN**	**89.6% Global Accuracy**

- **IFAT:** We compute the IFAT as the time difference between two consecutive Probe Requests from the same MAC address. Frames are first grouped into bursts, which are consecutive frames that share the same source MAC address and for which the elapsed time between frames is below 1 s. Once the bursts are identified, the mean IFAT is computed and assigned uniformly to all frames, including the first one.
- **SRSSI:** We simulate the RSSI perceived by an attacker using Log-Normal Shadowing Model (LNSM) with dynamic variance that depends on the distance from the STA to a sniffer [9].

B. Clustering Algorithms

Three clustering algorithms are evaluated: K-Means, DBSCAN, and OPTICS. K-Means has a lower computational complexity, $O(knT)$, where k is the number of clusters, n is the number of samples, and T is the number of iterations. However, it uses a predefined number of clusters. This is addressed by DBSCAN, which identifies them automatically and does not force all samples into clusters (some can be treated as outliers), but assumes that clusters have similar densities. Since devices may generate different numbers of Probe Frames, OPTICS is considered, as it better handles clusters with varying densities. Both DBSCAN and OPTICS have a worst-case computational complexity of $O(n^2)$ [10].

C. Evaluation Metrics

Unsupervised clustering performance is assessed using four metrics: Global Accuracy, Precision, Recall, and Individual Device Accuracy. Although the clustering algorithms are unsupervised, we evaluate their performance using the ground-truth labels in the dataset. To map the predicted clusters to the corresponding ground-truth labels in the dataset, we apply the Hungarian algorithm, which efficiently finds the assignment that maximizes the number of correct matches without evaluating all the possible permutations [11].

III. EXPERIMENTAL SETUP

A. Feature Importance Analysis

Random Forest (RF) feature importance analysis confirms that decomposed HT subfields and IFAT are highly critical (with IFAT ranking in the top three), while spatial RSSI consistently achieves the highest importance when present. Zero-importance features (`channel`, `rates`, `ext_rates`, and `vht_caps`) were excluded. `ssid` was retained because non-wildcard Service Set Identifiers (SSIDs) help differentiate devices with identical HT specs. We also derived `oui_from_mac` (first 24 bits of resolved MAC) as a vendor identifier alongside `oui`. The final feature set comprises `time`, `time_diff` (IFAT), decomposed HT subfields (or raw `ht_caps`), `htex_caps`, `ampdu`, `oui`, `oui_from_mac`, and SRSSI.

B. Scenarios and Configurations

We consider three scenarios to reflect different levels of an eavesdropper's infrastructure and tracking capabilities:

- *Sce. 1* (Without RSSI): Passive attack relying solely on Frame headers and IFAT timing.
- *Sce. 2* (Single Sniffer): One monitoring node providing a single SRSSI per Frame.
- *Sce. 3* (Multi-Sniffer): Three nodes in a triangular topology providing three SRSSI values per Frame.

Within each scenario, four distinct setups are tested to evaluate the contribution of specific features:

- Raw/Decomposed Feature Representation: Comparing the use of the raw Hexadecimal `ht_caps` string against the decomposed binary subfields.
- With/Without Temporal Information: Evaluating the models both with and without the IFAT.

Each combination (12 in total) is evaluated across four device densities (5, 10, 15, and 22 devices). For the 5, 10, and 15 devices, five random subsets of devices are generated without replacement, and results are averaged across these iterations. For the 22-device case, the full dataset is used, so only a single result is obtained, and no averaging is required.

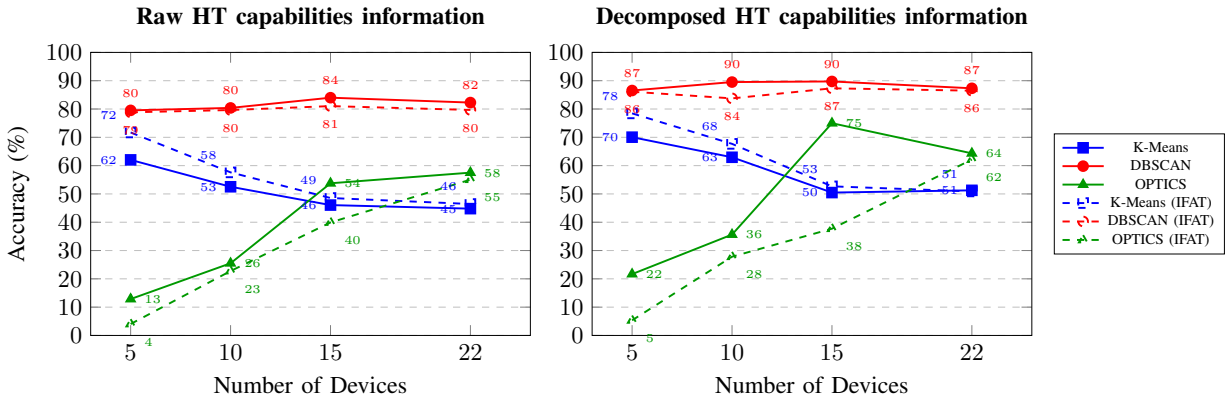


Fig. 1. Clustering accuracy vs. device count for Scenario 1 (without RSSI). Left: raw hexadecimal HT capabilities information. Right: decomposed HT capabilities information. Solid lines: without IFAT; dashed lines: with IFAT.

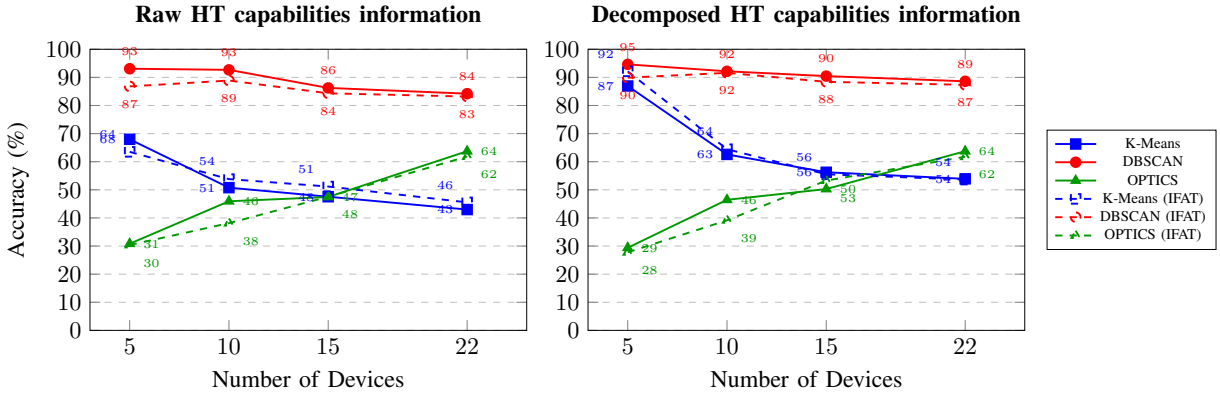


Fig. 2. Clustering accuracy vs. device count for Scenario 2 (single-sniffer RSSI). Left: raw hexadecimal HT capabilities information. Right: decomposed HT capabilities information. Solid lines: without IFAT; dashed lines: with IFAT.

IV. Results and Discussion

A. Clustering Performance

1) Sce. 1 (Without RSSI): Scenario 1 evaluates baseline performance using only Probe Frame headers and temporal data, leaving it vulnerable to hardware collisions among devices from the same manufacturer (Fig. 1). Decomposing HT capabilities information into individual subfields consistently improves clustering across all algorithms over raw hexadecimal values. This decomposition yields the largest accuracy gains for K-Means, though its performance still degrades as device density increases. DBSCAN exhibits superior robustness, reaching an average global accuracy of 88.3%, while OPTICS underperforms both models. Finally, IFAT aids K-Means cluster separation but degrades DBSCAN by introducing temporal variance that density-based clustering interprets as noise.

2) Sce. 2 (Single Sniffer): Scenario 2 adds spatial context via a single SRSSI value to estimate device distance (Fig. 2), improving performance across all algorithms compared to Sce. 1. Decomposing HT capabilities information remains essential despite the addition of spatial data. Under the decomposed configuration, DBSCAN demonstrates superior robustness, achieving 88.6% accuracy across 22 devices. K-Means benefits from SRSSI in lower device densities but fails to scale, while OPTICS shows only marginal improvements. Finally, IFAT remains detrimental to DBSCAN, introducing noise that compensates for the benefits of spatial separation.

3) Sce. 3 (Multi-Sniffer): Scenario 3 evaluates a multi-sniffer architecture using three SRSSI signals to improve physical localization (Fig. 3), yielding the highest accuracy across all scenarios. These spatial features enable effective clustering regardless of MAC randomization or identical hardware specs. Nonetheless, decomposing HT capabilities information into subfields remains critical; in the 22-device environment, DBSCAN reaches 89.6% accuracy with decomposed features—a 22.9% increase over raw hexadecimal values. While K-Means and OPTICS achieve their best performance in this scenario, K-Means still misclusters approximately 40% of frames with 22 devices, and OPTICS remains significantly inferior to DBSCAN.

B. Individual Device Analysis

While the last setup (Sce. 3, Decomposed subfields, w/o IFAT) achieved a high global accuracy of 89.6% using DBSCAN, this result can be misleading. Because devices do

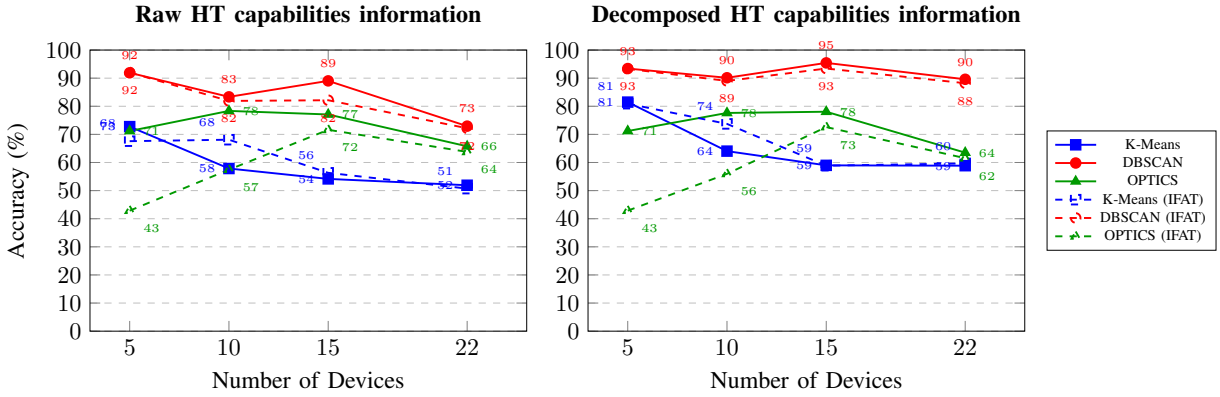


Fig. 3. Clustering accuracy vs. device count for Scenario 3 (multi-sniffer RSSI). Left: raw hexadecimal HT capabilities information. Right: decomposed HT capabilities information. Solid lines: without IFAT; dashed lines: with IFAT.

TABLE II
PER-DEVICE ACCURACY (DECOMPOSED SUBFIELDS, W/O IFAT, DBSCAN). SHADING: ≥80% , 50–79% , 1–49% , 0% . THE THIRD COLUMN FOR EACH DEVICE INDICATES THE NUMBER OF MATCHED/TOTAL PROBE REQUESTS.

Device	Acc.	Matched/Total	Device	Acc.	Matched/Total
A	88.9%	40/45	M	90.5%	19/21
B	100%	129/129	N	100%	27/27
C	100%	1236/1236	O	14.5%	9/62
D	100%	6/6	Q	100%	98/98
E	0%	0/115	R	100%	287/287
G	100%	617/617	S	100%	64/64
H	100%	12/12	T	70.6%	12/17
I	100%	28/28	U	0%	0/5
J	44.2%	73/165	V	0%	0/27
K	82.4%	28/34	W	100%	539/539
L	0%	0/54	X	0%	0/7

not transmit the same number of Probe Frames, this class imbalance has an impact on individual identification accuracy (Table II). Devices with very few samples are frequently misclustered; for example, sparse transmitters like devices O and T are often absorbed by dominant devices like Q. Consequently, extreme minority classes (such as U, V, and X) exhibit 0% individual accuracy. Due to their sparse data, these Probe Frames are easily absorbed into the closest dominant neighboring clusters. Devices L and V overlap and cluster with device B, while device U is absorbed by device W. Beyond sparse data, the clustering task is further complicated when a single device broadcasts two different HT capabilities information values. This is seen with devices J and M, where the use of multiple signatures makes it nearly impossible to achieve 100% individual accuracy. By splitting their Probe Frames across different configurations, their maximum accuracies were limited to 44.2% and 90.5%, respectively.

V. CONCLUSIONS AND FUTURE WORK

This study evaluated an eavesdropper's ability to bypass Wi-Fi MAC address randomization using Machine Learning (ML)-based fingerprinting across static, temporal, and spatial features. Decomposing the 16-bit HT capabilities information into individual subfields proved critical for device separation, while multi-sniffer RSSI integration significantly boosted performance. Overall, DBSCAN emerged as the most robust model, achieving 89.6% global accuracy across 22 devices. However, tracking performance remains constrained by low-frequency transmitters failing to form dense clusters, missing HT information forcing reliance on weaker secondary features, and devices dynamically alternating HT signatures over time. Future work will extend this evaluation to newer standards like Extremely High Throughput (EHT) and explore temporal models (e.g., Long Short-Term Memories (LSTMs)) to capture user mobility in real-world environments.

REFERENCES

[1] J. Henry *et al.*, *Wi-Fi 7 In Depth: Your Guide to Mastering Wi-Fi 7, the 802.11be Protocol, and Their Deployment*. Cisco Press, 2024.

[2] M. Vanhoef *et al.*, "Why MAC Address Randomization is not Enough: An Analysis of Wi-Fi Network Discovery Mechanisms," in *Proc. ACM ASIA CCS*, 2016, pp. 413–424.

[3] J. Freudiger, "How talkative is your mobile device? an experimental study of Wi-Fi probe requests," in *Proc. ACM WiSec*, 2015.

[4] W. Praharenka and I. Nikolaidis, "Identifying device type from cross channel probe request behavior," in *Proc. ACM WiSec*, 2021, pp. 392–394.

[5] L. Pintor and L. Atzori, "A dataset of labelled device Wi-Fi probe requests for MAC address de-randomization," *Comput. Netw.*, vol. 205, p. 108783, 2022.

[6] L. Pintor, "Advancements in Wi-Fi-Based Passenger Counting and Crowd Monitoring: Techniques and Applications," Ph.D. dissertation, Univ. degli Studi di Cagliari, 2024.

[7] F. Cifuentes-Urtubey and R. Kravets, "Poster: Can You Find Me?: Linking Devices Despite Wi-Fi MAC Randomization at MobiCom 2023," in *Proc. ACM MobiCom*, 2024, pp. 1668–1670.

[8] A. Pérez-Hernández *et al.*, "De-Randomization of MAC Addresses Using Fingerprints and RSSI With ML for Wi-Fi Analytics," *IEEE Access*, vol. 12, pp. 150 857–150 868, 2024.

[9] J. Xu *et al.*, "Distance measurement model based on RSSI in WSN," *Wireless Sens. Netw.*, vol. 2, pp. 606–611, 2010.

[10] scikit-learn, "Clustering," https://scikit-learn.org/stable/modules/clustering.html, 2026.

[11] J. Gil-Aluja, *The Hungarian assignment algorithm*. Springer, 1998, pp. 148–158.