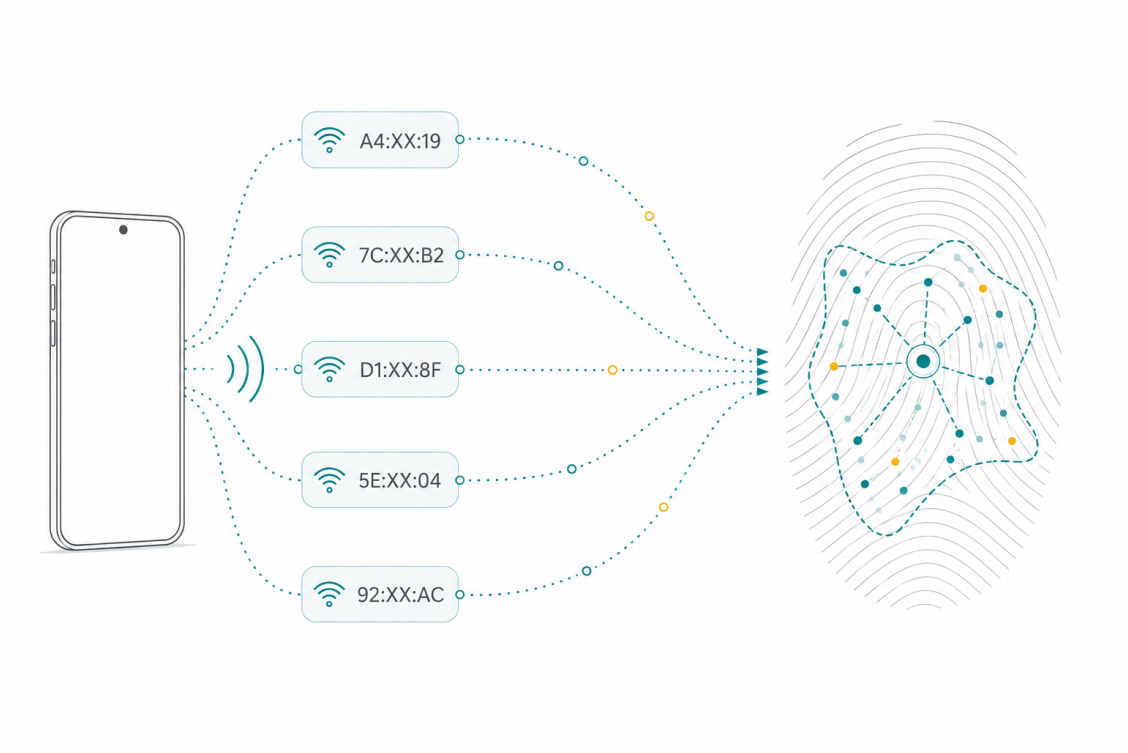

Device Fingerprinting from Passive WiFi Traffic: Can Devices Be Uniquely Identified?

A Study on Device Fingerprinting from Passive WiFi Traffic

Department of Computer Science and Engineering

Miranda Bergström, Ludvig Dahl, Måns Detterfelt,
Rima Jiang Safady, Monika Statelova, Seyyid Osman Uzel

Department of Computer Science and Engineering

CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026
www.chalmers.se

Bachelor project report 2026

Device Fingerprinting from Passive WiFi Traffic: Can Devices Be Uniquely Identified?

Department of Computer Science and Engineering

Miranda Bergström, Ludvig Dahl, Måns Detterfelt,
Rima Jiang Safady, Monika Statelova, Seyyid Osman Uzel

UNIVERSITY OF
GOTHENBURG

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Computer Science and Engineering
Chalmers University of Technology
University of Gothenburg
Gothenburg, Sweden 2026

Device Fingerprinting from Passive WiFi Traffic: Can Devices Be Uniquely Identified?
A Study on Device Fingerprinting from Passive WiFi Traffic
Miranda Bergström, Ludvig Dahl, Måns Detterfelt,
Rima Jiang Safady, Monika Statelova, Seyyid Osman Uzel

Supervisor: Mohamed Hashim Changrampadi, Department of Computer Science and Engineering

Graded by teacher: Magnus Almgren, Department of Computer Science and Engineering
Examiner: Arne Linde, Department of Computer Science and Engineering

Bachelor's Thesis 2026
Department of Computer Science and Engineering
Chalmers University of Technology
SE-412 96 Gothenburg
Sweden
Telephone +46 31 772 1000

Cover: The image was generated using OpenAI's ChatGPT Images through ChatGPT in 2026, based on an author-written prompt describing passive WiFi device fingerprinting.

Typeset in LaTeX
Gothenburg, Sweden 2026

Device Fingerprinting from Passive WiFi Traffic: Can Devices Be Uniquely Identified?
A Study on Device Fingerprinting from Passive WiFi Traffic
Miranda Bergström, Ludvig Dahl, Måns Detterfelt,
Rima Jiang Safady, Monika Statelova, Seyyid Osman Uzel
Department of Computer Science and Engineering
Chalmers University of Technology
University of Gothenburg

Abstract

Wireless devices are widely used in homes, workplaces, and public spaces, making WiFi communication a pervasive part of everyday digital infrastructure. Even when payload encryption and Media-Access Control (MAC) address randomization are enabled, wireless devices can still expose metadata that may support device identification. Such identification could enable social engineering attacks, surveillance, and compromises of personal privacy. To examine this risk, this study investigated whether WiFi devices can be fingerprinted using only passively observed traffic, without decrypting payload content or actively interacting with the network. Data was collected through passive monitoring in a controlled home environment and analyzed using statistical methods, DBSCAN clustering, Uniform Manifold Approximation and Projection (UMAP) dimensionality reduction, and a cluster-based prediction model. The extracted features included probe request structure, vendor-specific metadata, timing behavior, RSSI, and power-management indicators. The results show that combinations of passively observable metadata can reveal distinguishable patterns, particularly at the vendor and device-type level. Probe request metadata was the most informative feature category, while RSSI and sleep/wake behavior mainly provided contextual information. However, reliable identification of individual devices remained limited because similar devices often produced overlapping metadata patterns. The findings show that MAC address randomization and encryption reduce direct tracking, but do not fully prevent metadata-based classification or partial device re-identification.

Keywords: WiFi, device fingerprinting, security, metadata, passive monitoring, privacy.

Acronyms

AP
Access Point. xi, 12, V

DBSCAN
Density-based spatial clustering of applications with noise. iv, viii, x–xii, 9, 11, 17, 23–27, 31, 33, VI–IX

IE
Information Element. 7, 12, 14, 16, 17, 23–25, 31, 32, 34, 37

IEEE
Institute of Electrical and Electronics Engineers. 1, 4–7, 11, 15

MAC
Media-Access Control. iv, vii, viii, xii, 1–3, 5–9, 12, 15–17, 19, 23–28, 30–38

OUIs
Organizationally Unique Identifiers. 14, 16, 25

PCA
Principal Component Analysis. 17

RQ
Research Question. 10, 11

RSSI
Received Signal Strength Indicator. iv, vii, viii, x, xi, 1, 3, 4, 6, 10–13, 16, 19–22, 28, 30, 32–35, 37, I–IV

t-SNE
t-distributed Stochastic Neighbour Embedding. 17

UMAP
Uniform Manifold Approximation and Projection. iv, viii, x, xi, 17, 23, 26, 27, 31, 33, VI–IX

Contents

List of Figures

List of Tables

1

Introduction

Wireless communication has become an essential component of everyday life. Smartphones, laptops, tablets, wearables, smart televisions, and Internet of Things devices, such as smart home sensors, thermostats, or voice assistants, are often continuously connected to WiFi networks in both domestic and public environments. This makes WiFi a central technology for modern connectivity, but it also creates important privacy and security challenges.

Unlike wired communication, wireless transmissions are broadcast over radio waves and can therefore be observed by any receiver within range. To protect users, modern wireless systems include several privacy mechanisms designed to limit device tracking. Link-layer encryption protects the content of transmitted data, thereby preventing direct inspection of user communications [1]. In addition, many operating systems implement MAC address randomization, which periodically changes the transmitter address to reduce the possibility of persistent linkage between traffic and a fixed hardware identifier [2], [3]. From this perspective, passive monitoring, defined here as observing IEEE 802.11 transmissions without joining a network association, decrypting traffic, or injecting frames, is often assumed to provide limited capability for long-term device identification [2].

However, despite these protections, certain information must remain visible for the network to function correctly. Management and control frames are transmitted without encryption, and encrypted data frames carry side-channel information through protocol headers, timing information and physical-layer indicators accessible to a passive receiver [1]. Such metadata includes probe request structure, frame timing, power-management behavior, and Received Signal Strength Indicator (RSSI).

This thesis investigates whether such passively observable WiFi metadata can be used for device fingerprinting. In this context, device fingerprinting refers to recognizing a device based on observable characteristics of its communication behavior rather than relying on an explicit identifier.

1.1 Motivation and Project Objectives

The motivation for this study lies in the distinction between concealing explicit identifiers and preventing device recognition altogether. Existing privacy mechanisms, such as encryption and MAC address randomization, reduce direct forms of

tracking, but they do not necessarily remove all observable patterns from wireless communication. Even without a stable identifier, a device may still communicate in ways that are partly shaped by its hardware, software, configuration, and current usage state [4].

This creates a privacy concern because device recognition does not always depend on one strong identifying feature. Instead, several weak and indirect observations may be combined into a more informative profile. Individually, such characteristics may be noisy or unreliable, but combined they may reveal patterns that persist across observations.

The practical relevance of this problem is especially clear in domestic environments. Devices in a home network are often closely connected to individual users and daily routines. If a passive observer can infer that the same device is present across different points in time, or distinguish between devices without relying on their MAC addresses, then metadata-based observation may undermine the privacy protection that users expect from encryption and address randomization. Even partial or probabilistic recognition may therefore have privacy implications, particularly when observations are collected over longer periods.

At the same time, it is not clear in advance which metadata features are useful for fingerprinting. Wireless communication is affected by the physical environment, device state, user activity, network conditions, and implementation differences between devices.

The primary objective of this thesis is therefore to investigate what can realistically be inferred from passively observed WiFi metadata under controlled conditions. The study first explores several categories of observable metadata and then evaluates which features remain sufficiently stable and informative to contribute to device fingerprinting. In doing so, it examines whether combinations of structural, temporal, signal-based, and behavioral characteristics can support device distinction or re-identification despite MAC address randomization, without analyzing the encrypted payload content.

As a secondary objective, the project develops a visualization dashboard that presents observed networks, devices, and selected metadata features in real time. The dashboard also provides an interface for initiating device analysis, in which newly captured metadata is processed by a cluster-based fingerprinting model. The model compares observed feature representations with previously trained device centroids and produces a classification based on device type, behavioral similarity, or vendor-related information, together with a confidence indication.

1.2 Research Question

The core research question investigated in this thesis is formulated as follows:

> *Can devices be fingerprinted using passively observed wireless traffic alone, despite MAC randomization and without analyzing encrypted payload*

content?

To address this question, the study is guided by the following sub-questions:

1. Which categories of WiFi metadata can be observed and extracted under passive monitoring conditions?
2. Which of the explored metadata features remain sufficiently stable and informative to contribute to device fingerprinting?
3. Can combinations of weak metadata features provide more reliable device distinction than individual features considered in isolation?
4. What do the findings imply for the effectiveness of MAC address randomization as a privacy mechanism?

Addressing these questions requires identifying which metadata features are observable under passive monitoring conditions, evaluating their stability over time, and determining whether combinations of such features enable reliable cross-session matching.

1.3 Contributions

This thesis makes three main contributions.

First, the study explores which passively observable WiFi metadata features can be used for device fingerprinting. This includes evaluating RSSI, probe request metadata, and power-management behavior as possible sources of device-specific information. The analysis shows that these features differ in reliability and interpretability: RSSI provides contextual information about proximity and movement, probe request metadata reveals structural and behavioral patterns related to vendors and device types, and power-management behavior provides additional context during live observation.

Second, the study explores machine learning-based analysis of passively observed metadata. Unsupervised clustering is used to investigate whether extracted features naturally form meaningful groups without predefined device labels. In addition, a cluster-based fingerprinting model is implemented to classify newly observed devices based on similarity to previously trained device representations. Together, these methods provide a practical way to evaluate whether the selected metadata features can support automated fingerprinting beyond manual inspection.

Third, the study implements a visualization dashboard for passive WiFi metadata. The dashboard presents observed networks and devices in real time, including selected features such as RSSI, frame activity, timestamps, and sleep/wake status. It also includes fingerprinting-related fields, such as the predicted fingerprint, confidence score, classification tier, and vendor hint, and provides an interface for initiating device analysis using the implemented fingerprinting model. This makes the dashboard both a tool for inspecting passive observations and a practical interface for evaluating fingerprinting-relevant patterns.

1.4 Scope and Limitations

The scope of this project was limited to passive monitoring of WiFi traffic in a controlled home-network environment. Public WiFi networks typically generate high volumes of traffic from many unrelated devices, making controlled analysis significantly more complex. By restricting the study to a home setting, it was possible to maintain awareness of which devices were present and active during experiments. This made it easier to relate observed traffic patterns to known devices, which was necessary when evaluating metadata-based fingerprinting. However, this also limits the generalizability of the findings, since device behavior in a home network may differ from behavior in dense public environments.

All traffic monitoring was conducted passively. The study did not use active interference or manipulation, such as frame injection, or man-in-the-middle techniques. This was a deliberate methodological limitation because the project specifically investigated what could be inferred under passive observation conditions, without privileged access or intrusive techniques. Excluding active methods ensured that any observed fingerprinting capability arose from naturally exposed metadata.

The study did not perform any decryption of encrypted traffic. This decision was motivated both by ethical and technical considerations. From an ethical perspective, decryption could expose sensitive information associated with individuals using the network, which fell outside of the intended scope of the project. From a technical perspective, implementing decryption mechanisms would have substantially increased the technical complexity of the project and shifted the focus away from metadata analysis. Consequently, the study was limited to information observable without access to payload content.

Monitoring and analysis were performed in real time using two wireless network interfaces configured in monitor mode. In this mode, an interface can passively capture nearby IEEE 802.11 frames without associating with an access point or transmitting traffic as part of the network. The use of two interfaces reflected a realistic resource-limited observation scenario, since a typical laptop can only accommodate a limited number of external wireless adapters. However, this also restricted monitoring range and channel coverage. As a result, some channel activity or short-duration transmissions may not have been captured. Wireless measurements were also affected by environmental factors such as distance, physical obstruction, interference, device orientation, and multipath propagation [5]. These factors are especially important when interpreting signal-based measurements such as RSSI.

2

Theoretical Background

This chapter provides the technical background needed to understand passive WiFi monitoring, MAC address randomization, device fingerprinting, and the use of machine learning methods for analyzing wireless metadata. WiFi networks are based on the IEEE 802.11 family of standards, which define wireless communication at the physical layer and the Media-Access Control (MAC) layer [1]. At the MAC layer, communication is organized into units called frames rather than packets. These frames are commonly divided into management, control, and data frames. Management frames support network discovery and connection establishment [2], control frames coordinate access to the wireless medium, and data frames carry higher-layer traffic once communication has been established [1].

Before a device can communicate through a WiFi network, it must discover available access points. This can occur through passive scanning, where the device listens for beacon frames transmitted by access points, or through active scanning, where the device sends probe requests and receives probe responses [2]. After discovering a network, the device typically performs authentication and association with an access point. In protected networks, association is followed by key generation and distribution, including the IEEE 802.11i four-way handshake, which establishes session keys used for protected data communication [6]. Although encryption protects the contents of data frames, it does not hide all wireless metadata. Frame timing, frame type, addressing fields, signal strength, and management-frame behavior may still be observable to a passive monitor [1], [6].

These properties of IEEE 802.11 communication are important for this thesis because they determine which information remains visible without decrypting traffic or actively participating in the network. The following sections therefore explain how passively observable WiFi metadata relates to MAC address randomization, device fingerprinting, re-identification, and unsupervised learning methods used to analyze recurring patterns in wireless traffic.

2.1 Passive Observability and WiFi Metadata

The widespread adoption of WiFi has increased its relevance from a security and privacy perspective. Unlike wired communication, wireless transmissions are broadcast over radio signals that can be received by devices within range [1]. This characteris-

tic makes wireless networks inherently observable and relevant for both adversarial tracking scenarios and academic research in areas such as network security, privacy, and traffic analysis [2], [4]. A passive observer does not need to join the network in order to capture nearby IEEE 802.11 frames, provided that the wireless interface supports monitor mode. Monitor mode is a mode which can allow any device, that has the ability to wirelessly connect to a network, to monitor and capture WiFi traffic.

Although encryption mechanisms protect the contents of transmitted data, certain information must remain visible for the network to function correctly [2]. This includes frame headers, management-frame fields, timing information, channel information, and physical-layer indicators. Such information may reveal communication patterns, device behavior, and network activity even when payload data is encrypted [1].

In this thesis, passive monitoring refers to observing IEEE 802.11 traffic without authenticating to the network, decrypting protected payloads, or injecting frames. The observable information is therefore limited mainly to physical-layer characteristics and MAC-layer frame information. In this context, metadata refers to information about a wireless transmission other than the protected payload itself, such as frame type, transmission timing, addressing fields, management-frame contents, channel information, and Received Signal Strength Indicator (RSSI) [1], [7]. These observable properties provide the basis for the fingerprinting techniques discussed in the following sections.

2.2 MAC Addresses and Address Randomization

A Media-Access Control (MAC) address is a 48-bit link-layer identifier assigned to a network interface [8]. For globally administered addresses, this identifier is intended to be unique to a specific device interface [8]. A MAC address is commonly written as six hexadecimal octets separated by colons, for example `A4:5E:60:12:34:56`. In IEEE 802.11 networks, MAC addresses appear in frame header address fields and are used to indicate entities such as the transmitter, receiver, station, and access point involved in a wireless transmission [1], [9].

Historically, wireless devices often used static hardware addresses assigned to their network interfaces [8]. Since these identifiers could remain constant across networks and over long periods of time, they enabled long-term tracking of devices and, indirectly, their users across different locations [8]. This created a privacy concern in WiFi communication, especially during network discovery, where devices may transmit management frames before they are associated with a network.

To mitigate this privacy risk, modern operating systems and wireless standards have introduced MAC address randomization [2]. Instead of always transmitting the factory-assigned address, a device may use a randomly generated locally administered address in certain WiFi contexts [8]. This prevents a passive observer from relying only on a stable MAC address to link observations across time or across

different networks [3].

The duration for which a randomized MAC address remains constant depends on the operating system, device configuration, and communication context. During network discovery, devices may use randomized addresses in probe requests, and these addresses may change between scans, sessions, or probe-request bursts [9]. However, randomization is not necessarily synchronized with every probing round, and the same randomized address may appear in consecutive rounds. When a device connects to a known network, the address may remain stable for longer periods, for example as a persistent per-network randomized address. Therefore, MAC address randomization should not be understood as one uniform behavior, but as a set of implementation-dependent mechanisms intended to reduce linkability.

While MAC address randomization reduces the ability to track devices through explicit identifiers, it does not eliminate all identifying characteristics from wireless communication [2]. Other observable properties, such as probe request behavior, frame timing, supported capabilities, sequence behavior, and signal-related characteristics, may still provide information that can help distinguish or re-identify devices [2], [8]. For this reason, MAC address randomization motivates the use of device fingerprinting techniques that rely on patterns in observable metadata rather than on a fixed MAC address.

2.3 Device Fingerprinting and Re-Identification

Fingerprinting is the process of distinguishing or recognizing an entity based on a combination of observable characteristics rather than a single explicit identifier. In this thesis, device fingerprinting refers to applying this idea to WiFi devices by analyzing observable characteristics of their communication behavior instead of relying on identifiers such as MAC addresses. A fingerprint is therefore not one feature by itself, but a collection of features that together form a characteristic pattern. In WiFi communication, such features may include probe request behavior, frame timing, frame type distributions, supported capabilities, sequence behavior, signal-strength characteristics, and other metadata observable in IEEE 802.11 traffic [2], [10], [11].

In practice, device fingerprinting is performed by collecting observations from one or more devices, extracting measurable features from those observations, and comparing the resulting feature sets across time or across sessions [2], [8]. These comparisons may be performed using direct similarity measures, statistical methods, or machine learning techniques. If observations collected at different times produce sufficiently similar fingerprints, they may be linked as likely originating from the same device.

Previous work has shown that WiFi devices can expose identifying information even when explicit identifiers are unavailable or unreliable. Studies such as [2] and [8] have investigated fingerprinting based on probe request contents, Information Element (IE)s, sequence-number behavior, timing patterns, and signal-related features. These approaches show that device-specific or implementation-specific behavior may

remain observable at the WiFi and MAC layers, even when MAC address randomization is used [2], [9], [11].

From a security perspective, this means that hiding or changing a direct identifier does not necessarily prevent recognition. Device fingerprinting may allow an observer to link observations across different communication sessions, and therefore support re-identification. In this context, re-identification refers to recognizing that separate observations likely correspond to the same device or device category, even when explicit identifiers such as randomized MAC addresses are unavailable or have changed [8]. Over time, such observations can reveal patterns of presence, movement, or behavior that may support profiling or targeted attacks.

2.4 Machine Learning Approaches for Device Fingerprinting

Device fingerprinting can be performed using different types of comparison methods. In simple cases, fingerprints may be compared using manually defined rules or similarity measures. However, when fingerprints consist of several features, such as timing behavior, frame type distributions, signal-strength statistics, and management-frame characteristics, machine learning methods can be used to analyze patterns that may be difficult to describe manually [2].

In this context, an observation refers to a collected sample of WiFi traffic or metadata from a device during a specific measurement period. From each observation, measurable features can be extracted and represented as a feature vector. A feature vector is a numerical representation of the observation, where each value describes one property of the observed traffic, such as probe request count, frame timing, supported capabilities, vendor-specific metadata, or received signal strength [2].

Machine learning approaches for fingerprinting can be supervised or unsupervised [2]. Unsupervised learning analyzes observations without predefined labels and attempts to find structure or groupings based on similarity between feature vectors [12]. Supervised learning uses observations with known labels, such as the device identity, device type, or operating system, to classify new observations. Both approaches are relevant in this thesis: unsupervised clustering is used to explore structure in metadata, while a trained similarity-based prediction model is used to compare new observations with previously learned device profiles.

2.5 Unsupervised Learning for Fingerprinting

Unsupervised learning refers to methods that analyze data without predefined class labels. In this thesis, observations are compared based on their extracted WiFi metadata features rather than being grouped in advance by a known device identity. The goal is to examine whether the data contains natural groupings that may indicate recurring device-specific, vendor-level, or implementation-level behavior.

Clustering is a common unsupervised learning approach [2]. It groups observations so that observations within the same cluster are more similar to each other than to observations in other clusters [12]. In passive WiFi fingerprinting, clustering can therefore be used as an exploratory tool to identify recurring similarity patterns among feature vectors [2], [12].

One clustering method relevant to this thesis is Density-based spatial clustering of applications with noise (DBSCAN) [4]. DBSCAN identifies clusters as dense regions of points in a feature space, separated by areas of lower point density [13]. A point represents one observation after it has been converted into a feature vector. A key advantage of DBSCAN is that it does not require the number of clusters to be specified in advance. It can also classify points that do not belong to any dense region as noise, which is useful when passive WiFi observations contain outliers, incomplete measurements, or environmental variation [4].

In the context of MAC address randomization, unsupervised learning is relevant because the explicit identifier may not be reliable or persistent [2]. In this thesis, DBSCAN is not treated as a direct identifier by itself. Instead, it is used to explore whether combinations of observed WiFi metadata features produce meaningful groupings that may support later fingerprinting and re-identification analysis.

2.6 Supervised Learning for Fingerprinting

Supervised learning refers to methods where a model is trained using observations with known labels [12]. In device fingerprinting, these labels may represent a specific device, device type, vendor, or operating system. The model learns relationships between extracted feature vectors and these labels, and can then assign a label to a new observation based on similarity to the training data [12].

For WiFi fingerprinting, supervised learning is useful when ground-truth information is available about which device produced each observation. However, its reliability depends on how representative the training data is. A model trained on a limited set of devices may perform well on known or similar devices, but may struggle with unseen devices, changed device behavior, different environments, or overlapping feature patterns.

In this thesis, the supervised component is represented by the trained cluster-based prediction model. Known-device captures are used to create device-specific feature representations, and new observations are classified by comparing them with these learned profiles. The outputs should therefore be interpreted as similarity-based classification rather than deterministic identification.

3

Method

This chapter describes the experimental design, data collection, and analytical methodology used to investigate whether passive WiFi metadata could support device fingerprinting under controlled conditions. The methodology simulated a resource-limited passive observer and evaluated whether stable structural and behavioral signatures could be extracted without active interference, privileged access, or traffic decryption.

The methodological workflow consisted of a passive monitoring setup, a set of feature-evaluation experiments, a cluster-based prediction model, and a visualization dashboard. To make the experimental structure easier to follow, the main experiments and evaluation components are numbered in Table 3.1. Experiments E1–E6 evaluate individual categories of passively observable metadata, while E7 evaluates the implemented fingerprinting model and demonstrates the visualization dashboard.

Table 3.1: Overview of experimental and evaluation components

No.	**Experiment**	**Research question**	**Purpose**	**Main comparison**
E1	Short-range RSSI measurements	RQ1, RQ2	Evaluate whether RSSI reflects relative proximity under controlled short-range conditions.	RSSI at 1, 2, and 4 meters.
E2	Long-range RSSI measurements	RQ1, RQ2	Examine whether the distance–RSSI relationship remains visible over larger distances.	RSSI at 2, 6, and 10 meters.
E3	Movement-based RSSI measurements	RQ2	Assess how movement affects short-term RSSI stability.	Packet-level RSSI while the target device moved relative to the monitoring point.

No.	Experiment	Research question	Purpose	Main comparison
E4	Obstruction-based RSSI measurements	RQ2	Evaluate the effect of a physical barrier on signal strength and variability.	RSSI before and after wall obstruction.
E5	Probe request structure and scanning behavior	RQ1, RQ2, RQ3, RQ4	Determine whether probe request metadata contains stable device-, vendor-, or implementation-level patterns.	Information Element structure, supported rates, vendor-specific fields, burst timing, and DBSCAN clusters.
E6	Sleep/wake and power-management behavior	RQ1, RQ2, RQ3	Assess whether power-management indicators provide useful behavioral context.	Active and low-activity device states during longer captures.
E7	Cluster-based prediction model and visualization dashboard	RQ2, RQ3, RQ4	Develop and evaluate a cluster-based prediction model using the features extracted from E1–E6, and visualize the resulting metadata, clusters, and prediction outputs in a dashboard.	Known-device captures compared with new observation windows, with dashboard visualization of metadata, clusters, and predictions.

3.1 Passive Monitoring Setup

The experimental setup consisted of one access point, a set of target devices, and one or more monitoring devices equipped with external wireless adapters operating in monitor mode. Monitor mode allowed nearby IEEE 802.11 frames to be captured without associating with the observed network or interacting with the observed devices. This setup made it possible to collect passively observable metadata from management, control, and data frames while avoiding active interference with normal network operation. The main components of the experimental setup are presented in Table 3.2.

Table 3.2: Experimental setup parameters

Parameter	Configuration
Access Point	Home WiFi access point and mobile phone hotspot
Monitoring device	Laptop running Linux-based packet capture tools
Wireless adapter(s)	External WiFi adapters supporting monitor mode
Capture mode	Passive monitor mode
Frequency band	2.4 GHz and/or 5 GHz, depending on the capture
Channels monitored	Channel selected by access point or selected channel-hopping setup
Target devices	Two iPhones, one laptop and one Android phone
Capture format	PCAP/PCAPNG and extracted CSV/structured data
Physical environment	Controlled home-network environment
Observer position	Stationary unless otherwise specified

3.1.1 Capture Tools and Data Processing

Wireless traffic was collected using the Aircrack-ng suite, specifically the `airodump-ng` tool, in combination with `tcpdump` and `tshark` for low-level packet capture. These tools were used to collect raw traffic and metadata such as MAC addresses, RSSI values, channel information, frame statistics, timestamps, and device presence over time [14], [15], [16]. The raw captures also enabled extraction of structural frame characteristics, including Information Element composition and ordering, supported capabilities, and vendor-specific fields. Wireshark was used during the exploratory and validation phases to inspect selected packets, verify parsing and filtering logic, and confirm which metadata fields remained observable despite encrypted payloads [17]. Windows command-line network utilities were used for supplementary inspection during controlled experiments.

Custom Python scripts based on Scapy were developed to automate packet parsing, feature extraction, aggregation, and real-time processing. These scripts were used to extract probe request timing, burst behavior, Information Element structure, supported rates, capability fields, RSSI statistics, vendor-specific metadata, and power-management indicators.

3.1.2 Capture Configuration and Controlled Conditions

Before each capture session, the wireless adapter was configured in monitor mode and assigned either to a specific channel or to a relevant channel-hopping strategy. For experiments involving an associated target device, the monitored channel was selected based on the channel used by the access point.

Raw captures were stored before feature extraction so that original observations could be reprocessed and validated when necessary. During controlled experiments, variables such as device placement, observer position, distance, movement, obstruc-

tion, and activity state were either kept constant or changed deliberately according to the purpose of the experiment. This made it possible to compare capture sessions under controlled conditions and reduce the risk of attributing environmental variation to device-specific behavior.

3.2 Feature Evaluation Experiments

The feature evaluation consisted of six experiments, denoted E1–E6 in Table 3.1. These experiments evaluated three categories of passively observable WiFi metadata: RSSI variation, probe request structure and behavior, and power-management patterns. Experiments E1–E4 evaluated RSSI under controlled distance, movement, and obstruction conditions. Experiment E5 evaluated probe request structure and scanning behavior, while Experiment E6 evaluated power-management behavior. Together, these experiments were used to assess which metadata features were sufficiently stable, interpretable, and informative to contribute to device fingerprinting.

3.2.1 Proximity Estimation Using Received Signal Strength Indicator (RSSI)

Experiments E1–E4 evaluated the reliability, variability, and contextual sensitivity of RSSI measurements, as well as whether they could contribute to fingerprinting. The target device was placed at predefined distances from the monitoring device while the monitoring position and device orientation were kept stable where possible.

The relationship between physical distance and RSSI was examined using both short-range and long-range measurements. For short-range measurements, the device was positioned at distances of 1, 2, and 4 meters, while for long-range observations, distances of 2, 6, and 10 meters were used. This enabled analysis of how signal strength decreased with distance and established a baseline for expected RSSI variation.

Controlled movement was then introduced by moving the target device past the monitoring point at a controlled speed. This made it possible to observe short-term RSSI variation and sensitivity to motion. Finally, an obstruction scenario was introduced by placing a wall between the transmitter and receiver in order to evaluate the impact of obstructions and multipath effects on signal stability.

The extracted RSSI features included mean signal strength, variance, short-term fluctuation patterns, and stability over time. Graphical representations of the measurements were generated to visualize trends, assess variability, and identify systematic relationships between distance, movement, and environmental conditions.

3.2.2 Evaluation of Device-Specific Patterns Using Probe Request Structure and Scanning Behavior

Experiment E5 analyzed probe request behavior to assess its suitability for device fingerprinting. The analysis considered both temporal behavior of active scanning

and structural properties of probe request frames, including Information Element structure, capability fields, and vendor-specific metadata.

To isolate probe request activity, target devices were placed within range of the monitoring setup while remaining unassociated with any access point. This ensured that observed traffic reflected active scanning behavior rather than communication from an established connection.

Captured probe request frames were processed to extract a set of temporal features to characterize scanning behavior. These included inter-frame timing, mean probe interval, timing variability, and the total number of transmitted probe requests. Probe requests were also grouped into bursts using a predefined inter-frame timing threshold of 0.3 seconds, allowing rapid intra-burst transmissions to be separated from longer inter-burst gaps. From these bursts, features such as burst size, burst position, inter-burst timing, and estimated scan-cycle periodicity were derived.

Structural features were extracted to capture device-specific properties of frame composition. These included the ordered sequence of Information Elements, the number of elements per frame, supported and extended data rates, aggregate rate statistics such as maximum and mean supported rate, capability fields, HT and VHT capability indicators, extended capabilities, and vendor-specific elements. Since these characteristics are influenced by hardware, firmware, driver, and operating-system implementations, they were expected to remain relatively stable across captures.

Vendor-related features were derived from Organizationally Unique Identifiers (OUIs) and vendor-specific Information Elements. These included the number of vendor-specific elements, the diversity of OUIs, and the size of vendor-specific data fields. Such features provided insight into manufacturer-level implementation patterns and protocol extensions.

The analysis evaluated whether probe request metadata combined stable structural characteristics with distinguishable scanning behavior. Temporal features were assessed through statistical measures and distributions, while structural and vendor-related features were evaluated based on their consistency across captures and their ability to separate devices, vendors, or device types.

3.2.3 Evaluation of Device-Specific Patterns Using Power-Management Behavior

Experiment E6 used long-duration captures to evaluate whether sleep/wake cycles and power-management behavior were consistent enough to contribute to device-specific fingerprinting. The target device was connected to the access point and observed under two activity conditions. In the active condition, the device generated sustained traffic, such as video streaming or other bandwidth-intensive activity, while transmissions were passively captured for three to four hours. In the low-activity condition, the device remained connected but was not actively used, apart from background processes such as notifications or synchronization.

The analysis focused on power-management indicators, including QoS null frames,

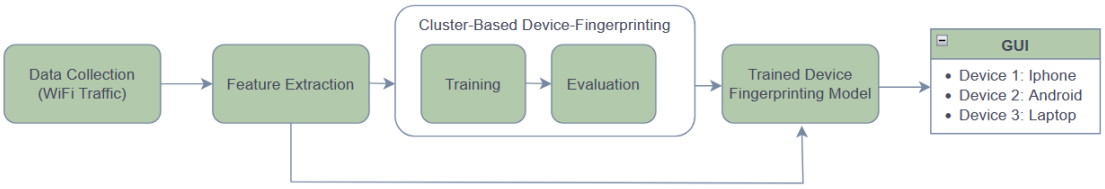

Figure 3.1: Overview of the passive WiFi fingerprinting pipeline from traffic capture to GUI-based fingerprint visualization.

null data frames, and power-management flags, since these could indicate transitions between active and power-saving states. From these observations, sleep/wake timing patterns and mean cycle durations were estimated.

The purpose of this experiment was to determine whether sleep/wake behavior produced identifiable and stable patterns, or whether it was too dependent on activity state and environmental conditions to serve as a reliable fingerprinting feature.

3.3 Fingerprinting System Architecture

This section describes the pipeline used to transform passively captured WiFi traffic into fingerprinting results. As shown in Figure 3.1, raw traffic was converted into structured feature representations, which were then used for clustering-based training and evaluation. The resulting fingerprinting model produced device labels, confidence levels, and supporting metadata that were presented in the visualization dashboard.

3.3.1 Features Used for Fingerprinting

The feature set was selected based on two criteria: whether the feature could be observed passively without decrypting traffic, and whether it could contribute to distinguishing devices, device types, or vendors. Since MAC addresses were not treated as persistent identifiers, the analysis focused on metadata available in IEEE 802.11 frames. The feature categories and their intended roles in the analysis are summarized in Table 3.3.

Table 3.3: Feature categories used in the fingerprinting analysis.

Feature category	**Extracted features**	**Role in fingerprinting**
Probe request behavior	Probe intervals, timing variability, probe count, burst size, burst position, inter-burst timing, scan-cycle periodicity	Captured active scanning behavior, but was less reliable than structural probe request features for fingerprinting.
Information Element structure	IE ordering, IE count, supported and extended rates, HT/VHT capabilities, extended capabilities, capability fields	Captured implementation-specific frame structure and supported device/vendor-level separation.
Vendor-related metadata	OUIs, vendor-specific IE count, OUI diversity, vendor data length	Provided manufacturer- or implementation-level information, but did not uniquely identify devices.
Received Signal Strength Indicator	Mean RSSI, variance, short-term fluctuation, temporal stability	Used as contextual information related to proximity, movement, and environmental conditions.
QoS and power-management behavior	QoS Null frames, null data frames, power-management flags, sleep/wake timing, mean cycle duration	Evaluated as a supplementary behavioral feature for active and idle state interpretation.
MAC address	Source and transmitter addresses observed in frames	Used for analysis and validation where relevant, but excluded as a persistent identifier due to MAC randomization.

3.3.2 Feature Extraction and Preprocessing

Feature extraction was implemented as a processing pipeline that converted raw packet captures into the feature categories summarized in Table 3.3. The raw data was first filtered to remove invalid or malformed frames, ordered by timestamp, checked for duplicate records, and processed to handle missing values. Potential measurement artifacts, such as packet loss and channel-switching gaps, were considered because they could affect timing-based, signal-based, and channel-related measurements.

For each relevant frame or observation window, temporal features were computed from probe request timing, burst structure, and inter-frame intervals. Structural features were extracted from Information Element ordering, element counts, supported rates, and capability fields. Vendor-related features were derived from OUI values and vendor-specific Information Elements, while signal-based and behavioral features were computed from RSSI statistics, QoS Null frames, null data frames, power-management flags, and estimated sleep/wake timing.

Categorical features, including vendor-related fields and Information Element presence, were encoded using binary or multi-label representations. Numerical features were normalized to make dimensions comparable before clustering, visualization, and prediction. In addition to packet-level features, observations were aggregated over time windows and device-level sessions to create more stable feature representations for clustering, visualization, and model evaluation.

3.3.3 Unsupervised Clustering and Pattern Discovery

Density-based clustering (DBSCAN) was used to examine whether the extracted features produced meaningful groupings without assuming the number of devices in advance. This was suitable for passive metadata because the observations could contain noise, incomplete measurements, and clusters of irregular shape.

Clustering was performed using combined temporal, structural, and vendor-related features. The resulting clusters were analyzed to determine whether observations from the same device, vendor, or device type were grouped together. Cluster-level summaries included packet counts, the number of unique source MAC addresses, dominant IE structures, and dominant vendor-related attributes. Cluster purity was also calculated with respect to source MAC addresses and inferred device labels, providing a quantitative measure of how well the clusters reflected underlying device-level, vendor-level, or implementation-level behavior.

3.3.4 Dimensionality Reduction and Visualization

Dimensionality reduction was used to inspect the structure of the high-dimensional feature space. Principal Component Analysis (PCA), t-distributed Stochastic Neighbour Embedding (t-SNE) and Uniform Manifold Approximation and Projection (UMAP) were applied to project feature vectors into two-dimensional representations. PCA was used to examine dominant variance patterns, while t-SNE and UMAP were used to visualize local neighborhood structure and possible cluster separation.

The resulting visualizations were used to support the interpretation of DBSCAN results, examine separation between devices or vendors, and assess how different feature categories contributed to the observed structure in the data.

3.3.5 Cluster-Based Device Prediction Model

Experiment E7 evaluated whether known devices could be identified from aggregated passive WiFi metadata using the cluster-based prediction model. Packet-level observations were grouped into fixed-duration or hybrid windows, where each window was represented as a feature vector. These aggregated windows were used because they provided a more stable representation of device behavior than individual packets.

During the analysis, it was observed that the same device did not always form a single compact group in the feature space. Instead, some devices produced multiple recurring feature patterns. These patterns may have reflected changes in activity

level, power-management behavior, environmental variation, or other factors affecting the captured metadata. As a result, a single centroid was not always sufficient to represent one device. Each device was therefore allowed to form multiple clusters, with the number of clusters determined by the variation observed in that device's training data.

The model was trained on data from four target devices, including both different device types and devices of the same general category. This made it possible to evaluate whether the extracted features could distinguish clearly different devices and whether similar devices produced overlapping feature representations. For each device, K-means clustering was applied to its training samples in the normalized feature space. This produced one or more device-specific clusters, with some devices represented by two clusters and others by three clusters depending on their observed behavioral variation.

During prediction, each new feature window was processed using the same feature extraction, encoding, and normalization steps as the training data. The model then calculated the distance between the observed feature vector and the trained clusters for each known device. The predicted label was assigned based on the closest matching device cluster, provided that the distance was within the corresponding acceptance threshold.

To account for uncertainty, the model used three prediction tiers based on the distance between an observed feature window and the trained device clusters. The tier depended on whether the observation was within the device-specific acceptance threshold, within an extended near-match margin of 1.25 times the threshold, or outside both conditions.

Table 3.4: Prediction tiers used by the cluster-based model.

Tier	**Condition**	**Output**
Accepted	Distance below the device-specific threshold.	Classified as corresponding known device with high confidence.
Near match	Distance within an extended margin with a clear gap between the nearest cluster and the competing clusters.	Labeled with a `_like` suffix, indicating partial similarity without full confidence.
Unknown / fallback	Outside accepted and near-match criteria. Vendor metadata used as fallback information.	*Unknown*, or assigned a coarse manufacturer or operating-system family hint when possible.

The model was evaluated using controlled single-device captures and multi-device scenarios. This allowed assessment of classification accuracy for known devices, robustness under realistic capture conditions, and the extent to which the extracted features supported device-level identification rather than only broader device-type or vendor-level separation. The cluster-based approach also made it possible to account for intra-device variation observed across different capture conditions, rather than assuming one stable behavioral representation per device.

3.4 Visualization Dashboard

The visualization dashboard formed part of Experiment E7 and was developed as a practical demonstration of the passive WiFi fingerprinting approach. The dashboard integrated the main findings of the study into an interactive system for observing networks, tracking devices, and displaying fingerprinting results in real time.

The dashboard presented detected wireless networks together with devices observed in relation to the selected network. For each device, the interface displayed metadata and fingerprinting outputs such as the observed MAC address, predicted fingerprint label, confidence level, classification tier, vendor hint, RSSI, movement estimate, mapped frame count, estimated distance, and sleep/wake status. These outputs connected the extracted metadata features to the higher-level interpretations developed in the analysis.

The dashboard also distinguished between active and disconnected devices. Active devices remained visible while recent activity was observed, whereas inactive devices were moved to a disconnected-device section after a defined period without observed transmissions. This made it possible to follow device presence and changes in observable behavior over time.

Overall, the dashboard functioned as a practical fingerprinting tool and proof-of-concept implementation. It demonstrated how passive WiFi metadata could be transformed into real-time device-level interpretations, including device classification, confidence assessment, vendor hints, proximity estimation, movement status, and power-management indicators.

4

Results

This chapter presents the results obtained from the experiments described in *Chapter 3*. The results are organized according to the main experimental phases. For each experiment, visualizations are used to illustrate observed patterns and to support interpretation of feature behavior.

4.1 RSSI Measurements Under Controlled Conditions

This section presents the results from Experiments E1–E4, which evaluated RSSI under controlled distance, movement, and obstruction conditions. In all figures, the horizontal axis represents time in seconds, while the vertical axis represents RSSI in dBm. Additional results are provided in Appendix A.

4.1.1 Short-Range Measurements

Measurements recorded at 1, 2, and 4 meters were categorized as short-range measurements. The measurements in Figure 4.1 show a consistent decrease in signal strength as distance increases. At approximately 1 meter, RSSI values are generally observed in the range of -35 to -40 dBm. At the distances of 2 meters and 4 meters, the values decrease to approximately -45 to -50 dBm and lower.

Although the overall trend is clear, the measurements also show variation within each distance segment. Fluctuations of several dB occur even when the device remains at a fixed distance. This indicates that RSSI can reflect relative proximity under controlled conditions, but that its precision is limited by signal variability.

4.1.2 Long-Range Measurements

Measurements recorded at 2, 6, and 10 meters were categorized as long-range measurements and used to evaluate signal behavior over a wider range. The results (Figure 4.2) confirm the negative relationship between distance and RSSI. Compared with the short-range measurements, the variability is more pronounced at longer distances, and the spread of measured values increases. This suggests that environmental factors have a stronger effect as the distance between the device and the monitoring point increases.

Figure 4.1: Short-range RSSI measurements using an iPhone 15 Pro as the target device. Dashed vertical lines indicate transitions between distances, while regression lines illustrate the average signal level within each segment. Shaded regions represent the variability of the measurements.

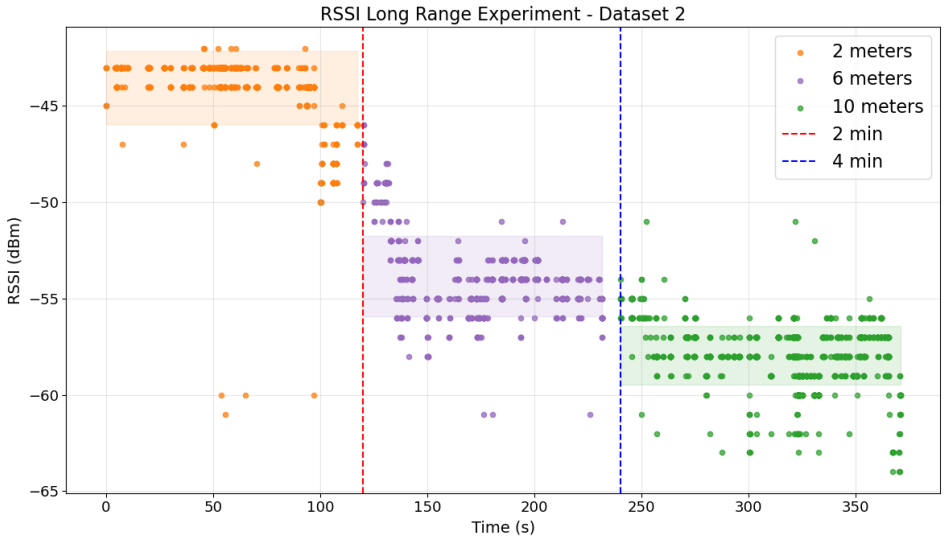

Figure 4.2: Long-range RSSI measurements using an iPhone 15 Pro as the target device. Dashed vertical lines indicate transitions between distances, while regression lines illustrate the average signal level within each segment. Shaded regions represent the variability (standard deviation) of the measurements.

4.1.3 Movement-Based Measurements

The movement-based experiment (Figure 4.3) examined how RSSI varies dynamically when a device moves relative to the monitoring point. The observed signal follows an overall trend corresponding to changes in distance over time, and the fit-

ted curve captures the general variation as the device moves away from and towards the receiver.

However, the individual packet-level measurements show substantial short-term fluctuations around this trend. This demonstrates that movement introduces additional variability, making RSSI sensitive not only to distance, but also to orientation and surrounding environmental conditions.

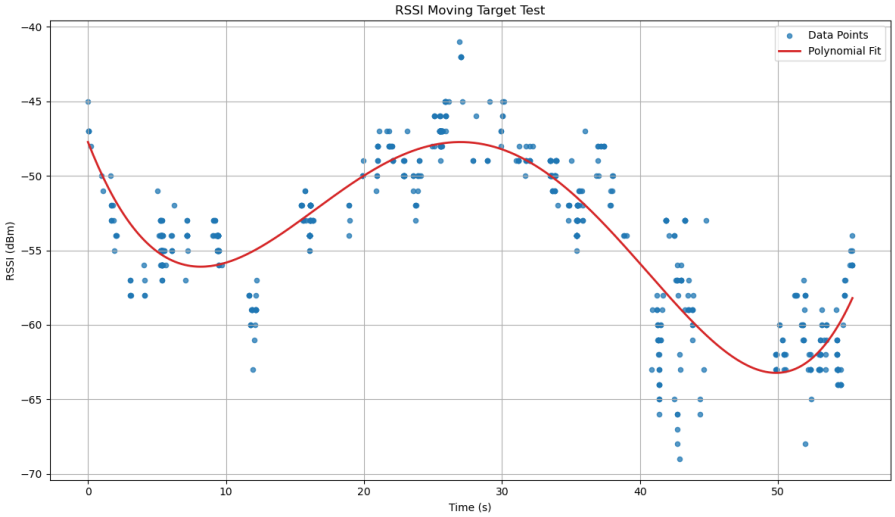

Figure 4.3: RSSI measurements for a moving target device passing the capture point. Each point represents an individual packet measurement, while the polynomial curve provides a smoothed approximation of the signal strength trend as the distance between the device and the capture point varies.

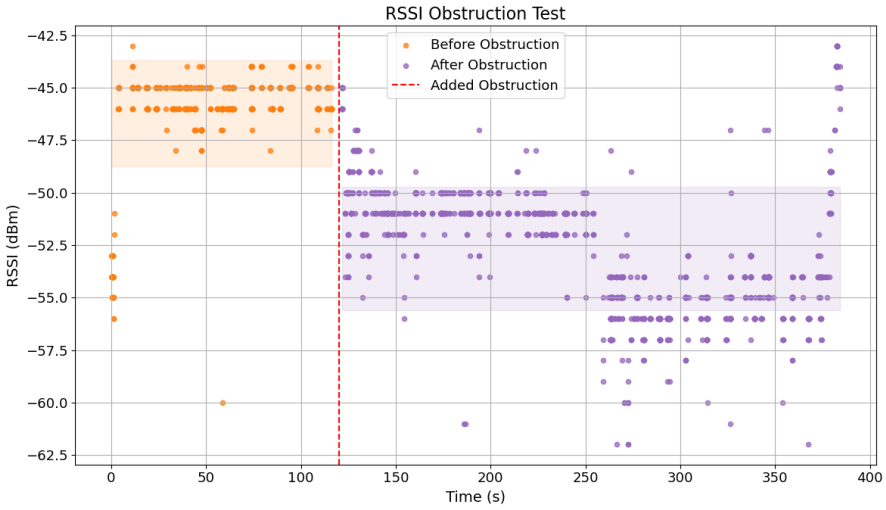

Figure 4.4: RSSI measurements using an iPhone 15 Pro as the target device. The dashed vertical line indicates the point at which the device is moved behind a wall. Regression lines illustrate the average signal strength before and after obstruction, while shaded regions represent measurement variability.

4.1.4 Obstruction-Based Measurements

The obstruction experiment evaluated the effect of introducing a physical barrier between the target device and the monitoring point. The results (Figure 4.4) show a clear decrease in average signal strength after the wall was introduced. The spread of the measurements also increases, indicating that obstruction adds further variability to the received signal.

4.2 Probe Request Clustering Analysis

This section presents the results from Experiment E5, which evaluated the packet-level analysis of probe request frames. Each packet was represented using features based on Information Element (IE) structure, supported rate characteristics, capability-related fields, vendor-specific information, and burst-timing behavior. The resulting feature matrix was standardized and clustered using DBSCAN. The results are reported through MAC-level summaries, cluster composition, cluster-purity measures, and two-dimensional Uniform Manifold Approximation and Projection (UMAP) projections. Additional cluster-level statistics and UMAP projections are provided in Appendix C.

4.2.1 MAC-Level Summary of Probe Request Behavior

The ten most frequently observed source MAC addresses (Table 4.1) show limited structural variation in their probe request frames. Most of the addresses use only one unique Information Element ordering, suggesting that probe request structure remains relatively stable within each observed source address.

At the same time, timing behavior differs between addresses, particularly in mean inter-packet delay and mean burst size. This indicates that devices with similar probe request structure may still differ in temporal transmission behavior.

Table 4.1: Top 10 source MAC addresses ranked by observed packet count.

Source MAC	Packet count	Vendor group	Unique IE orders	Typical IE order	Mean dt (s)	Mean burst size (packets/burst)
14:ac:60:a3:da:15	593	Wi-Fi Alliance	1	0-1-50-3-127-221	1.516	6.153
4c:d5:77:0c:bb:b7	208	Wi-Fi Alliance	1	0-1-50-3-127-221	2.629	18.31
bc:6e:e2:ff:a7:1d	199	Microsoft	1	0-1-50-3-45-221	4.525	3.975
88:f4:da:33:77:fa	161	Microsoft	1	0-1-50-3-45-221	5.496	3.646
5c:3a:45:d8:ae:8f	105	Unknown	2	0-1-50-45	8.500	1.952
8a:c6:05:bf:8d:38	92	Apple	1	0-1-50-3-45-127-221-221-221	7.212	5.435
ca:6e:5c:fa:6d:b8	78	Apple	1	0-1-50-3-45-127-221-221-221	10.287	5.821
7a:25:34:4b:de:f8	71	Apple	1	0-1-50-3-45-127-221-221-221	12.560	5.338
3e:22:8f:a1:d6:d3	49	Apple	1	0-1-50-3-45-127-221-221-221	17.411	4.673
c2:62:8c:c5:4a:a4	48	Unknown	1	0-1-50-3-45	6.058	4.708

4.2.2 Cluster-Level Summary: Cluster Composition by Source MAC Address

To support interpretation of the DBSCAN results, one representative MAC address was selected for each major non-noise cluster (Table 4.2). The representative address was defined as the source MAC contributing the largest number of packets to the cluster.

Some clusters are dominated by a small number of source addresses. For example, Cluster 8 and Cluster 16 have relatively high MAC-level dominance. In contrast, Clusters 1, 2, 3, and 4, contain packets from many different source addresses. These clusters therefore appear to represent broader shared probe request patterns rather than traffic from a single device.

Table 4.2: Representative source MAC addresses for major non-noise clusters.

Cluster	Cluster packets	Vendor group	Representative MAC	Representative MAC packets	MAC purity	Unique MACs	Typical IE order
8	857	Wi-Fi Alliance	14:ac:60:a3:da:15	593	0.692	6	0-1-50-3-127-221
2	633	Apple	8a:c6:05:bf:8d:38	91	0.144	149	0-1-50-3-45-127-221-221-221
1	493	Unknown	06:11:dd:a4:ce:b9	2	0.004	351	0-1-50-3-45-127
16	361	Microsoft	bc:6e:e2:ff:a7:1d	199	0.551	3	0-1-50-3-45-221
3	262	Unknown	a6:17:95:99:11:9d	2	0.008	261	0-1-50-3-45-127
4	170	Apple	06:0a:ce:e7:9f:02	1	0.006	170	0-1-50-3-45-221-127

4.2.3 Cluster-Level Summary: Cluster Composition by DBSCAN

The cluster-level results (Table 4.3) show that most non-noise clusters have a stable common Information Element order, whereas the noise cluster (Cluster -1) contains greater structural variation. Several clusters also align with vendor-related labels, including Apple, Microsoft, Wi-Fi Alliance, and Unknown. This indicates that the extracted probe request features capture vendor- and implementation-related structure in the data.

4.2.4 Cluster-Level Summary: Cluster Purity Analysis

The vendor-level purity results (Table 4.4) show that all non-noise clusters have a purity value of 1.000. This means that each non-noise cluster contains packets associated with only one collapsed vendor category. The noise cluster is more heterogeneous, containing five vendor labels and a purity of 0.346.

Table 4.3: DBSCAN cluster composition and dominant labels.

Cluster	Packets	Unique MACs	Unique IE orders	Common IE order	Top vendor OUIs	Top vendor name	Top vendor type	MACs per packet
-1	280	75	21	0-1-50-3-45-127-221-221-221	00:50:f2	Microsoft	22	0.268
0	10	8	1	0-1-50-45-127-221	50:6f:9a	Wi-Fi Alliance	22	0.800
1	493	351	1	0-1-50-3-45-127	none	Unknown	Unknown	0.712
2	633	149	2	0-1-50-3-45-127-221-221-221	00:17:f2	Apple	10	0.235
3	262	261	1	0-1-50-3-45-127	none	Unknown	Unknown	0.996
4	170	170	1	0-1-50-3-45-221-127	00:17:f2	Apple	10	1.000
5	110	15	1	0-1-50-3-45-127-221-221	50:6f:9a	Wi-Fi Alliance	22	0.136
6	144	36	1	0-1-50-3-45-191-221-127-221-221	00:50:f2	Microsoft	8	0.250
7	42	30	1	0-1-50-3-45-127-221-221-221-221	00:17:f2	Apple	10	0.714
8	857	6	1	0-1-50-3-127-221	50:6f:9a	Wi-Fi Alliance	22	0.007
9	12	8	1	0-1-50-3-221	00:50:f2	Microsoft	8	0.667
10	161	7	2	0-1-50-3-45	none	Unknown	Unknown	0.043
11	26	6	1	0-1-50-3-45-127-191-221	00:50:f2	Microsoft	8	0.231
12	42	7	2	0-1-50-45-127-221	50:6f:9a	Wi-Fi Alliance	22	0.167
13	19	2	1	0-1-50-3-45-127-221-221-221	00:90:4c	Epigram	51	0.105
14	17	6	1	0-1-50-3-45-127-191-221	50:6f:9a	Wi-Fi Alliance	22	0.353
15	52	1	1	0-1-45-191	none	Unknown	Unknown	0.019
16	361	3	1	0-1-50-3-45-221	00:50:f2	Microsoft	8	0.008

Table 4.4: Cluster purity by collapsed vendor label.

Cluster	Packets	Unique values	Dominant value	Dominant count	Purity
8	857	1	Wi-Fi Alliance	857	1.000
2	633	1	Apple	633	1.000
1	493	1	Unknown	493	1.000
16	361	1	Microsoft	361	1.000
3	262	1	Unknown	262	1.000
4	170	1	Apple	170	1.000
10	161	1	Unknown	161	1.000
6	144	1	Microsoft	144	1.000
5	110	1	Wi-Fi Alliance	110	1.000
15	52	1	Unknown	52	1.000
7	42	1	Apple	42	1.000
12	42	1	Wi-Fi Alliance	42	1.000
11	26	1	Microsoft	26	1.000
13	19	1	Epigram	19	1.000
14	17	1	Wi-Fi Alliance	17	1.000
9	12	1	Microsoft	12	1.000
0	10	1	Wi-Fi Alliance	10	1.000
-1	280	5	Microsoft	97	0.346

The MAC-level purity results (Table 4.5) show a different pattern. Purity with respect to source MAC addresses varies considerably across clusters. Cluster 15 is fully pure, while Cluster 8 and Cluster 16 show moderate dominance by one source address. Several larger clusters, including Cluster 1, 3, and 4, have very low MAC-

level purity.

Together, these results show that the clustering separates packets more clearly by vendor- or implementation-level characteristics rather than by individual source MAC address.

Table 4.5: Cluster purity by source MAC addresses.

Cluster	Packets	Unique values	Dominant value	Dominant count	Purity
15	52	1	5c:3a:45:d8:ae:8f	52	1.000
8	857	6	14:ac:60:a3:da:15	593	0.692
13	19	2	88:a2:9e:42:56:12	11	0.579
16	361	3	bc:6e:e2:ff:a7:1d	199	0.551
12	42	7	20:c1:9b:80:ad:9a	19	0.452
14	17	6	2e:6f:01:16:94:7c	6	0.353
10	161	7	5c:3a:45:d8:ae:8f	53	0.329
9	12	8	82:30:df:e6:25:89	3	0.250
11	26	6	fa:f1:6f:c8:ce:c1	6	0.231
6	144	36	c4:71:0f:61:fb:61	30	0.208
0	10	8	00:d4:9e:00:c6:6b	2	0.200
2	633	149	8a:c6:05:bf:8d:38	91	0.144
5	110	15	46:f3:53:50:b9:dd	12	0.109
-1	280	75	2c:cf:67:56:62:84	28	0.100
7	42	30	0a:c8:a9:fe:8e:4e	2	0.048
3	262	261	a6:17:95:99:11:9d	2	0.008
4	170	170	06:0a:ce:e7:9f:02	1	0.006
1	493	351	06:11:dd:a4:ce:b9	2	0.004

4.2.5 UMAP Projection After Noise Removal

The UMAP plots (Figure 4.5) present two views of the packet-level feature space after DBSCAN noise points have been removed. The left plot is colored by cluster assignment, while the right plot is colored by collapsed vendor label.

After noise removal, the remaining packets form clearer regions in the UMAP embedding. Large clusters, including Cluster 2, 8 and 16 remain visually distinct. The vendor-colored projection also shows that Apple, Microsoft, Wi-Fi Alliance, and Unknown occupy largely separate areas of the embedding.

These visual results are consistent with the vendor-purity analysis and support the finding that probe request metadata contains strong vendor- and implementation-level patterns. However, because several clusters contain packets from many different source MAC addresses, the results do not demonstrate reliable individual-device identification at packet-level.

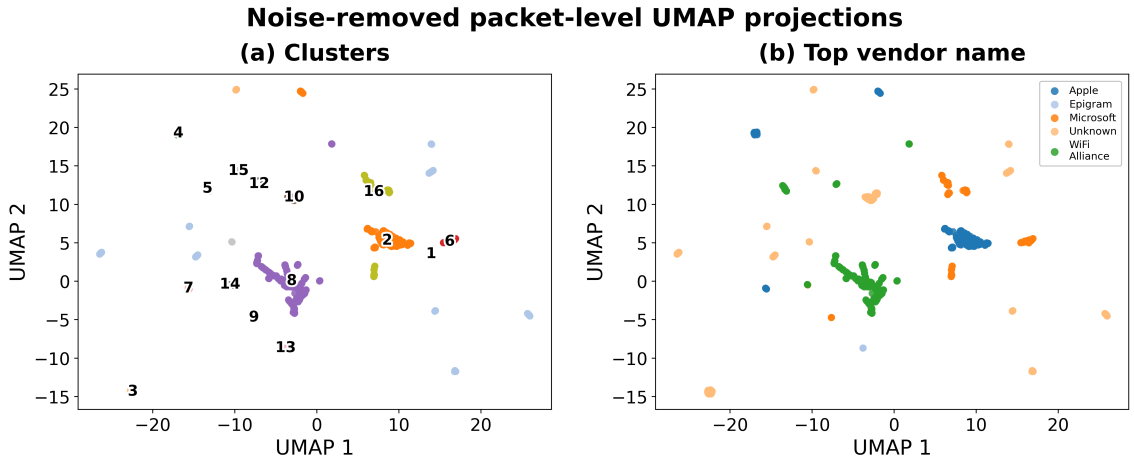

Figure 4.5: Two-dimensional UMAP projections of the packet-level feature space after removal of DBSCAN noise points. The left plot shows DBSCAN cluster assignments, and the right plot shows collapsed vendor labels.

4.3 Sleep/Wake Cycles

This section presents the results from Experiment E6, which evaluated power-management flags. These flags are located in the Frame Control Field of a WiFi frame and indicate whether a device is entering or leaving power-saving mode, referred to here as "sleep" [1, p. 664].

The captures used a laptop as the main target device and a mobile phone acting as the access point. The first capture also included a second mobile phone as a target device, while the second capture included only the laptop. The measurements used for this analysis (Table 4.6) were taken at 5 and 10 minutes into each capture. The supporting sleep/wake screenshots are provided in Appendix B.

The measured sleep/wake cycle times vary both between capture sessions and within the same capture. For the laptop, the mean cycle duration at the 5-minute measurement point differs from 11.15 seconds in Capture 1 to 5.69 seconds in Capture 2. This variation indicates that sleep/wake timing is not stable across sessions.

Table 4.6: Mean sleep/wake cycle duration at 5 and 10 minutes.

Capture	Device	MAC address	Mean S/W cycle at 5 min (s)	Mean S/W cycle at 10 min (s)
1	Mobile phone	56:A1:CB:39:D2:2C	0.98	2.56
1	Laptop	74:4C:A1:83:76:B9	11.15	14.34
2	Laptop	74:4C:A1:83:76:B9	5.69	9.82

These results suggest that sleep/wake timing is not sufficiently consistent to serve as a reliable fingerprinting feature on its own. However, it may still provide contextual

information about device state and activity level during passive observation.

4.4 Evaluation of the Cluster-Based Prediction Model

This section presents the results from Experiment E7. The model correctly classified all 12 observed source addresses at the station level after aggregation (Table 4.7). At the window level, it classified 46 out of 58 windows correctly, corresponding to an overall accuracy of 79.3%. Station-level correctness refers to the aggregated prediction per source address, while window-level accuracy refers to the proportion of correctly classified aggregation windows. The strongest results were obtained for iPhone and laptop devices, whereas Android phone captures produced more unknown windows. This indicates that the model was able to recognize known device categories, but that its confidence varied between device types and capture conditions.

Table 4.7: Cluster-based prediction results by device category.

Device category	Devices	Windows	Accepted	Like	Unknown	Correct	Accuracy
iPhone	6	6	5	1	0	6	100.0%
Android phone	3	38	25	2	11	27	71.1%
Laptop	3	14	12	1	1	13	92.9%
Total	12	58	42	4	12	46	79.3%

4.5 Visualization Dashboard

As part of Experiment E7, the visualization dashboard was implemented to display passively captured WiFi metadata and cluster-based prediction outputs in real time. It provided an interface for inspecting detected networks, observed devices, and selected fingerprinting-related outputs.

The Networks view (Figure 4.6) displayed detected networks together with device-level metadata. For each selected network, the dashboard showed the observed MAC address, RSSI, estimated distance, movement state, frame counts, timestamp of the most recent observation, sleep/wake status, and mean sleep cycle. It also displayed fingerprinting-related fields, including predicted fingerprint, confidence score, classification tier, and vendor hint. This allowed raw observations and model outputs to be inspected in the same interface.

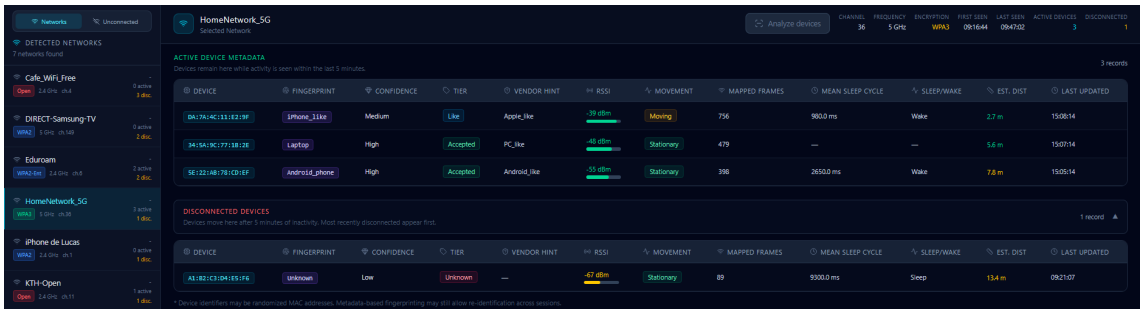

Figure 4.6: Networks view of the visualization dashboard.

The dashboard also included an Unconnected Devices view for devices observed without direct association to a detected access point. Since this view was not used as the main basis for the fingerprinting evaluation, it is included in Appendix E.

5

Discussion

This chapter discusses the results presented in Chapter 4 in relation to the main research question: whether devices can be fingerprinted using passively observed WiFi traffic alone, despite MAC address randomization and without inspecting encrypted payload content. The results show that passive metadata does not provide a deterministic identifier equivalent to a stable MAC address, but it can still reveal recurring patterns that support probabilistic device distinction. The strongest evidence comes from probe request metadata, while RSSI and sleep/wake behavior mainly provide contextual information during live observation.

5.1 RSSI and Sleep/Wake Behavior as Contextual Features

The RSSI experiments showed a clear relationship between signal strength, distance, movement, and obstruction under controlled conditions. This makes RSSI useful for estimating proximity or movement during a specific observation period. However, it is strongly affected by environmental factors such as device orientation, obstacles, reflections, interference, and the position of the monitoring device. Therefore, it should not be treated as a stable long-term identifier.

Sleep/wake behavior showed a similar limitation. Power-management indicators can be observed passively and may provide useful information about whether a device is active, idle, or entering a power-saving state. However, the measured sleep/wake cycle durations varied across captures, which limits their reliability as persistent fingerprinting features. This behavior is likely influenced by user activity, background processes, device state, and operating-system power management.

The main value of these features is therefore contextual rather than identifying. They can help interpret what is happening during a live capture and may support short-term differentiation between devices that otherwise appear to be similar. For example, two devices with similar probe request structures may still differ in signal strength, movement pattern, or current activity state. However, these features are not stable enough to replace structural metadata as the main basis for fingerprinting.

5.2 Interpretation of Probe Request Metadata

Probe request metadata provided the clearest fingerprinting-relevant patterns in the study. The clustering results show that features such as IE ordering, supported rates, capability fields, and vendor-specific elements capture implementation-level behavior. This means that passive observations can reveal how devices construct probe request frames, even when payload content is encrypted and MAC addresses are randomized. This is consistent with previous work, such as [2] and [8], showing that these features can remain informative even when explicit identifiers are randomized.

The strongest separation appeared at the vendor and device-type level. The non-noise clusters were highly consistent with collapsed vendors (Table 4.3), and the Uniform Manifold Approximation and Projection (UMAP) projection (Figure 4.5) showed that vendor-related groups occupied mostly distinct regions in the feature space. This indicates that probe request metadata is useful for distinguishing broad implementation groups, such as Apple, Microsoft, Wi-Fi Alliance, and Unknown categories. However, this should not be interpreted as reliable identification of individual physical devices. Several clusters contained many different source MAC addresses, which suggests that similar devices or implementations can produce similar probe request structures.

Therefore, a fingerprint in this study should not be understood as a unique identifier in the same sense as a fixed hardware address. Instead, it is better understood as a behavioral and structural profile. Such a profile can reduce the anonymity set of an observed device by assigning it to a likely vendor, device family, or behavior group, but it cannot always determine which exact physical device produced the traffic.

The results also show that structural stability alone is not sufficient for unique identification. As shown in Table 4.1, many frequently observed source addresses used a stable IE order, but the same or similar structures also appeared across multiple devices. Therefore, IE ordering is useful, but only as part of a broader feature set. Its value increases when combined with capability fields, supported rates, and vendor-specific metadata. This supports the idea that passive fingerprinting is strongest when several weak features are combined rather than when one feature is used in isolation.

Noise and unknown clusters further illustrate the limits of the approach. The Density-based spatial clustering of applications with noise (DBSCAN) noise cluster (Cluster -1) contained more heterogeneous observations, which may represent rare behavior, incomplete captures, transitional states, or patterns that were not frequent enough to form dense clusters. Similarly, an Unknown label does not necessarily mean that the traffic lacks structure. It may only mean that the available metadata could not be confidently linked to a known vendor category. For this reason, the cluster labels should be treated as interpretive categories rather than verified physical-device identities.

5.3 Combined Fingerprinting Model

The combined fingerprinting model demonstrates how the most stable extracted metadata can be used in a practical classification workflow. Based on the feature evaluation, the model focused on structural and vendor-related probe request metadata, such as Information Element structure, supported capabilities, and vendor-specific fields. More context-dependent features, including temporal behavior, RSSI, and sleep/wake cycles, were excluded from the core model inputs because they varied more strongly across capture conditions.

By using the more stable metadata features, the model could compare new observations with previously learned device profiles and assign predictions, confidence scores, and classification tiers. The model should therefore be interpreted as a similarity-based classification system rather than a deterministic identification system. A high-confidence prediction means that the observed structural and vendor-related metadata closely matches a learned profile. A near-match prediction still indicates strong similarity, since the observation remains close to the learned centroid, but it does not satisfy the stricter acceptance criteria. In contrast, a low-confidence prediction indicates greater uncertainty.

The evaluation results support this interpretation. After aggregation, the model correctly classified all 12 observed source addresses at the station level, while the window-level evaluation reached 79.3% accuracy. This suggests that the model can recognize known device categories when stable structural and vendor-related features are used. However, the difference between device categories also shows that the model's confidence depends on capture conditions and device behavior. The stronger results for iPhone and laptop devices, compared with the larger number of unknown Android phone windows, indicate that the model is not equally reliable across all device types. Therefore, the model should be interpreted as a proof-of-concept similarity-based classifier rather than a general real-world identification system.

5.4 Implications for MAC Randomization

The results are consistent with previous studies showing that MAC address randomization improves privacy by reducing direct address-based tracking, but it does not remove all metadata-based linkability. Although randomization prevents simple tracking based on a fixed address, other observable fields and behavioral patterns, such as those explored in this study, may remain stable enough to support classification or partial re-identification.

This does not mean that MAC randomization is ineffective. The limited source-MAC purity in several clusters shows that individual-device identification becomes more difficult when the address is not reliable. However, the results also show that changing the address alone is not enough to eliminate all fingerprinting risk. Probe request structure, capability information, supported rates, and vendor-specific elements may still reveal implementation-level patterns.

A broader privacy protection strategy would therefore need to consider more than address randomization. Reducing linkability may require standardizing, minimizing, or randomizing other observable metadata in management frames. Otherwise, passive observers may still be able to classify or compare devices based on the remaining structural and behavioral information.

5.5 Limitations and Future Work

Several limitations affect how the findings should be interpreted. First, the dataset was limited in size, device diversity, and capture environment. The experiments were conducted in a controlled home-network setting, which made the analysis manageable but limits generalizability. A larger dataset with more devices, vendors, operating systems, chipsets, and environments would therefore be needed to evaluate how stable the observed patterns are across broader real-world conditions.

Second, some results use the observed source MAC address to interpret clustering behavior (Table 4.5). This is useful for evaluating how the addresses are grouped, but it should not be treated as definitive proof of physical-device identity. Under MAC randomization, the same physical device may appear under different source addresses over time, especially during active scanning. Therefore, source-MAC purity reflects separation of observed addresses rather than necessarily separation of individual physical devices. Future work could use longer observation periods and repeated captures with stronger ground truth to evaluate whether fingerprints remain stable across randomized addresses, device states, environments, and normal user behavior.

Third, several features were context-sensitive. RSSI was affected by distance, movement, and obstruction, while sleep/wake behavior varied across captures. These features can support live interpretation, but they are not reliable enough to serve as long-term identifiers on their own.

Fourth, DBSCAN and UMAPintroduce limitations related to the exploratory analysis. DBSCAN depends on feature selection, normalization, and parameter choices, so different settings could change the resulting clusters or noise assignments. UMAPalso simplifies the high-dimensional feature space into two dimensions, meaning that the visual separation should be treated as supporting evidence rather than proof of separability. Therefore, the DBSCAN clusters and UMAPplots should be interpreted as exploratory indications of metadata structure, not as final evidence of device identification. Future work could improve this by performing a sensitivity analysis across different DBSCAN parameters, feature sets, and normalization strategies.

Finally, the cluster-based prediction model has different limitations. It compares new observations with learned device-specific profiles based on aggregated feature windows, K-means clusters, and distance thresholds. Its predictions therefore depend on how representative the training captures are and how well the learned centroids capture normal device variation. Although the model performed well at the station level in the proof-of-concept evaluation, the evaluation used a limited number of known devices and should not be treated as a general estimate of real-world

classification performance. The model may still struggle when a device behaves differently from the training data or when two devices have very similar structural and vendor-related metadata. Future work could evaluate the model using longer and more diverse captures, repeated sessions, unseen devices, and alternative classification methods.

5.6 Sustainability Considerations

The environmental impact of the study was limited because the experiments used passive traffic capture with two wireless adapters and one computer, without generating additional network traffic. The main material impact came from the hardware used for the measurements. Therefore, the project's broader sustainability relevance is mainly social rather than environmental. The findings relate to privacy, user control, and the potential misuse of passive metadata for tracking, profiling, or commercial use, which are discussed in Chapter 6.

5.7 Answer to the Research Question

The main research question asked whether devices can be fingerprinted using passively observed wireless traffic alone, despite MAC randomization and without inspecting encrypted payload content. Based on the results, the answer is yes, but only with important limitations. Passive WiFi metadata can support probabilistic fingerprinting and device distinction, especially at the vendor, device-type, or behavioral-profile level. However, the results do not support reliable deterministic identification of every individual physical device.

Table 5.1: Summary of answers to the research sub-questions.

Sub-question	Summary of findings
Which metadata can be observed passively?	RSSI, probe request structure, packet timing, burst behavior, capability fields, vendor-specific information, and sleep/wake behavior were observable.
Which features are stable and informative?	Probe request structural metadata was the most stable and informative, especially for vendor- and device-type classification. RSSI and sleep/wake behavior were mainly contextual, and temporal features were not as stable as structural features.
Do feature combinations improve distinction?	Yes. Combining IE structure, capability-related, and vendor-related metadata provided a stronger basis for fingerprinting than using individual features alone.
What does this imply for MAC randomization?	MAC randomization improves privacy by reducing address-based tracking, but it does not remove all metadata-based linkability.

6

Societal and Ethical Aspects

This thesis examined whether WiFi devices could be passively fingerprinted using metadata, despite MAC randomization and encrypted payloads. The study did not attempt to identify individuals directly. However, the results showed that passively observable metadata could distinguish devices at broader levels, such as vendor and device type, and in some cases supported partial re-identification across observations. Since personal devices are often closely associated with individuals, these findings raise ethical and societal concerns related to privacy, consent, and potential misuse.

6.1 Privacy, Consent and Data Sensitivity

A central ethical concern is that MAC address randomization may give users a stronger sense of privacy than it actually provides. In this study, the passively collected WiFi metadata did not contain direct personal information such as names, account details, message contents, or other payload data. Nevertheless, the results showed that combinations of metadata features could reveal distinguishable patterns, particularly at the vendor and device-type level. In practical settings, devices such as smartphones and laptops are often closely linked to individual users. As a result, even partial device re-identification may indirectly support user tracking.

The privacy risk is not limited to identifying technical properties of a device. Activity-dependent features, such as traffic timing and sleep/wake behavior, may provide indications of how and when a device is used. Similarly, RSSI measurements collected over time may provide indirect information about proximity, movement, and relative position within an environment. Although these features are noisy and context-dependent, their aggregation may reveal patterns such as device presence, activity periods, movement, or absence from a monitored area. In a domestic environment, such inferences may expose aspects of daily life that users reasonably expect to remain private.

These findings challenge common assumptions about privacy in wireless communication. Encryption protects payload content, and MAC randomization conceals static identifiers, but neither mechanism removes all stable structural or behavioral characteristics from wireless transmissions. The results therefore suggest that current privacy mechanisms mainly protect against direct identification, while metadata-based classification and re-identification remain possible. This makes WiFi metadata a

meaningful and often underestimated source of privacy leakage.

To account for these privacy implications, the data collection was limited to controlled home environments and mainly involved devices belonging to the researchers or informed participants. This reduced the risk of collecting unrelated personal data. However, passive WiFi monitoring created an inherent consent challenge because wireless transmissions can be captured from nearby devices that are not explicitly participating in the study. In practice, it may not be possible to inform every individual whose device is within capture range. This means that some risk of incidental data collection remained, even though the study was conducted in a restricted and controlled setting. To reduce this risk, the study avoided linking observations to real-world identities, stored data locally, restricted access to the research group, and removed the collected data after the project. These measures reduced the ethical risk of the study, but they did not fully eliminate the broader consent problem associated with passive wireless monitoring.

6.2 Dual-use and Societal Value

The techniques examined in this thesis have clear dual-use potential. On one hand, they can help researchers, developers, and standardization bodies understand weaknesses in existing privacy mechanisms. On the other hand, similar methods could be used for unauthorized tracking, surveillance, or commercial profiling. Even classification at the vendor or device-type level may reveal useful information about a person's devices, habits, or presence in a specific environment. The results should therefore be interpreted as evidence of a privacy weakness rather than as a method for identifying or monitoring individuals.

Despite these risks, the study offers significant societal value. By showing that passively observable metadata can support vendor-level and device-type classification, and in some cases partial re-identification, the study contributes to a more realistic understanding of the limitations of current WiFi privacy protections. This knowledge can support the development of stronger privacy-preserving mechanisms, inform future standards, and help users and policymakers better understand what encryption and MAC address randomization can and cannot protect against.

Overall, the study was ethically justifiable because the data collection was limited in scope, the analysis avoided personal identification, and mitigation measures were applied to reduce potential harm. At the same time, the findings highlight the need for continued discussion about passive monitoring, metadata privacy, and the limits of current privacy mechanisms in everyday wireless environments.

7

Conclusion

This thesis investigated whether WiFi devices can be fingerprinted using passively observed wireless traffic alone, despite MAC address randomization and without inspecting encrypted payload content. The results show that passively observable WiFi metadata can support device fingerprinting, but not in the form of reliable unique identification of every individual physical device. Instead, the fingerprinting capability demonstrated in this study is best understood as probabilistic device distinction based on structural and behavioral similarity.

Several categories of metadata were observable under passive monitoring conditions, including RSSI, probe request structure, packet timing, burst behavior, capability fields, vendor-specific information, and sleep/wake behavior. Among these, probe request metadata was the most informative feature category. Features such as Information Element structure, supported capabilities, and vendor-specific fields revealed stable implementation-level patterns that supported vendor- and device-type classification. RSSI and sleep/wake behavior were more dependent on environmental conditions and device state, making them more useful as contextual features during live observation.

The clustering results showed that devices from the same vendor or implementation family often produced similar probe request structures. This allowed broader groups, such as Apple, Microsoft, Wi-Fi Alliance, and Unknown categories, to be separated in the feature space. However, this separation did not translate into reliable unique identification of individual physical devices. Several clusters contained traffic from many different source MAC addresses, showing that multiple devices can share similar metadata patterns.

The cluster-based prediction model demonstrated how the most stable structural and vendor-related metadata can be used in a practical fingerprinting workflow. By comparing new observations with learned device profiles, the model could assign predictions, confidence scores, and classification tiers. In the proof-of-concept evaluation, it correctly classified all 12 observed source addresses after aggregation and achieved 79.3% accuracy at the window level. However, because the model was evaluated on a limited number of known devices and relies on similarity to learned profiles, it may struggle when devices have similar hardware, operating systems, or wireless implementations. The final fingerprint should therefore be interpreted as a similarity-based profile, not as a replacement for a persistent hardware address.

Overall, the answer to the research question is yes, but with important limitations. Passive WiFi metadata can support classification, comparison, and partial re-identification of devices, especially at the vendor, device-type, or behavioral-profile level. However, the results do not support reliable deterministic identification of every individual physical device. The study therefore shows that MAC address randomization and encryption reduce direct device tracking, but do not fully eliminate the risk of metadata-based device recognition.

References

[1] IEEE, "IEEE Standard for Information Technology–Telecommunications and Information Exchange between Systems Local and Metropolitan Area Networks–Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," *IEEE Std 802.11-2024 (Revision of IEEE Std 802.11-2020)*, pp. 1–5956, 2025.

[2] A. Pérez-Hernández, M. N. Barreras-Martín, J. A. Becerra, M. J. Madero-Ayora, and P. Aguilera, "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.

[3] Mathy Vanhoef and Célestin Matte and Mathieu Cunche and Leonardo S. Cardoso and Frank Piessens, "Why MAC Address Randomization is Not Enough: An Analysis of Wi-Fi Network Discovery Mechanisms," in *Proceedings of the 11th ACM Asia Conference on Computer and Communications Security (ASIACCS)*, 2016, pp. 413–424.

[4] P. Xiao, D. Lin, and M. Wang, "A RF Fingerprint Clustering Method Based on Automatic Feature Extractor," in *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, ser. LNICST. Cham: Springer, 2023.

[5] S. Ahmadi, "Chapter 9 - The IEEE 802.16m Physical Layer (Part I)," in *Mobile WiMAX*. Academic Press, 2011, pp. 335–487.

[6] S. Frankel, B. Eydt, L. Owens, and K. Scarfone, "Establishing Wireless Robust Security Networks: A Guide to IEEE 802.11i," National Institute of Standards and Technology, Tech. Rep. Special Publication 800-97, 2007.

[7] H. Lin, Y. Chen, S. Li, and W. Peng, "A Dual-Mode Framework for Indoor Localization via Temporal Learning and Knowledge Distillation," *Ad Hoc Networks*, vol. 181, p. 104089, 2026.

[8] N. Bogahawatta, Y. S. Karunanayaka, S. Seneviratne, K. Thilakarathna, and R. Masood, "Device Type Classification Using WiFi Probe Requests: From Signals to Insights," *IEEE Transactions on Mobile Computing*, 2026.

[9] J. Martin, T. Mayberry, C. Donahue, L. Foppe, L. Brown, C. Riggins, E. C. Rye, and D. Brown, "A Study of MAC Address Randomization in Mobile Devices

and When it Fails," *Proceedings on Privacy Enhancing Technologies*, vol. 2017, no. 4, pp. 365–383, 2017.

[10] A. Kobusińska, J. Brzeziński, and K. Pawulczuk, "Device Fingerprinting: Analysis of Chosen Fingerprinting Methods," in *Proceedings of the 2nd International Conference on Internet of Things, Big Data and Security - IoTBDS*, INSTICC. SciTePress, 2017, pp. 167–177.

[11] J. Kim and D. Han, "Passive WiFi Fingerprinting Method," in *Proceedings of the 9th International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE, 2018.

[12] A. Aksoy and S. Varma, "AI-Driven Genetic Algorithms for Enhanced Numeric Feature Quantization in IoT Device Fingerprinting for Threat Detection," in *Proceedings of the 2024 IEEE International Conference on Big Data (BigData)*. IEEE, 2024, pp. 2568–2574.

[13] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise," in *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, 1996, pp. 226–231.

[14] Aircrack-ng Project, "airodump-ng Documentation: What's the meaning of the fields displayed by airodump-ng?" https://www.aircrack-ng.org/doku.php?id=airodump-ng#what_s_the_meaning_of_the_fields_displayed_by_airodump-ng, 2022, accessed: 2026-04-16.

[15] The Tcpdump Group, "tcpdump Manual Page," https://www.tcpdump.org/manpages/tcpdump.1.html, 2026, accessed: 2026-04-16.

[16] Wireshark Foundation. (n.d.) TShark Manual Page. [Online]. Available: https://www.wireshark.org/docs/man-pages/tshark.html

[17] ——, "Wireshark User's Guide," https://www.wireshark.org/docs/wsug_html/, n.d., accessed: 2026-04-16.

[18] H. Hermansson and S. O. Hansson, "A Three-Party Model Tool for Ethical Risk Analysis," *Risk Management*, vol. 9, no. 3, pp. 129–144, 7 2007.

[19] Ben Wolford, "How does the GDPR affect email? - GDPR.eu," accessed: 2026-04-17. [Online]. Available: https://gdpr.eu/email-encryption/

A

RSSI Measurement Results

A.1 Short-Range RSSI Measurements

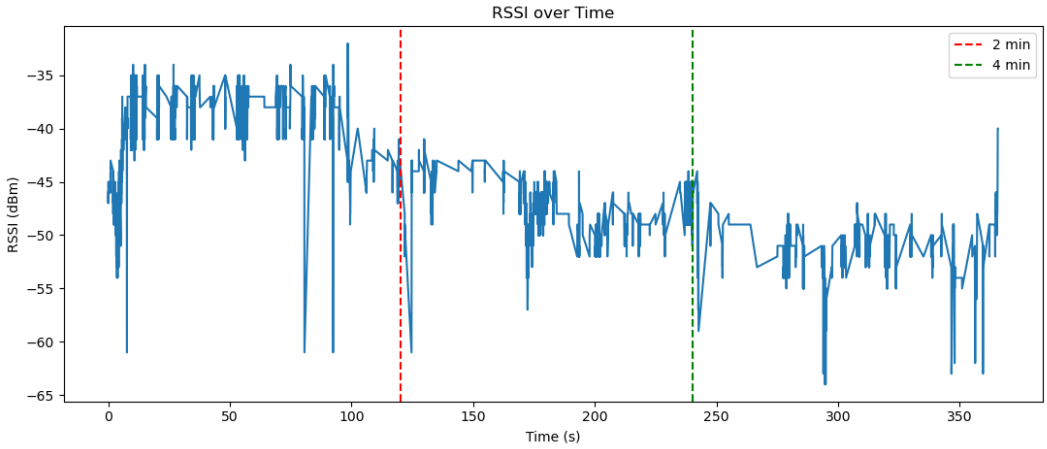

Figure A.1: Short-range RSSI measurements during initial capture. Dashed vertical lines indicate transitions between distances.

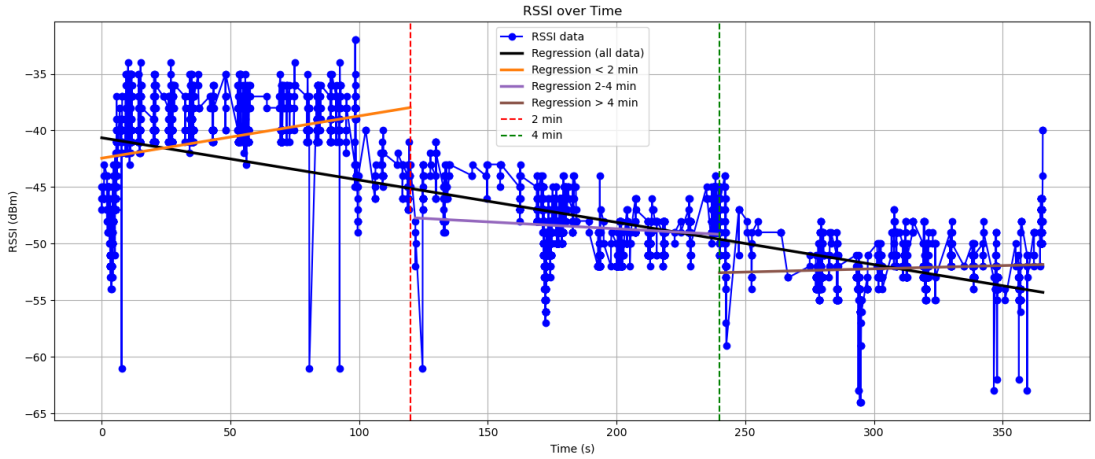

Figure A.2: Short-range RSSI measurements from initial capture with fitted regression lines for each distance segment. Each point corresponds to an individual captured packet. Dashed vertical lines indicate transitions between distances.

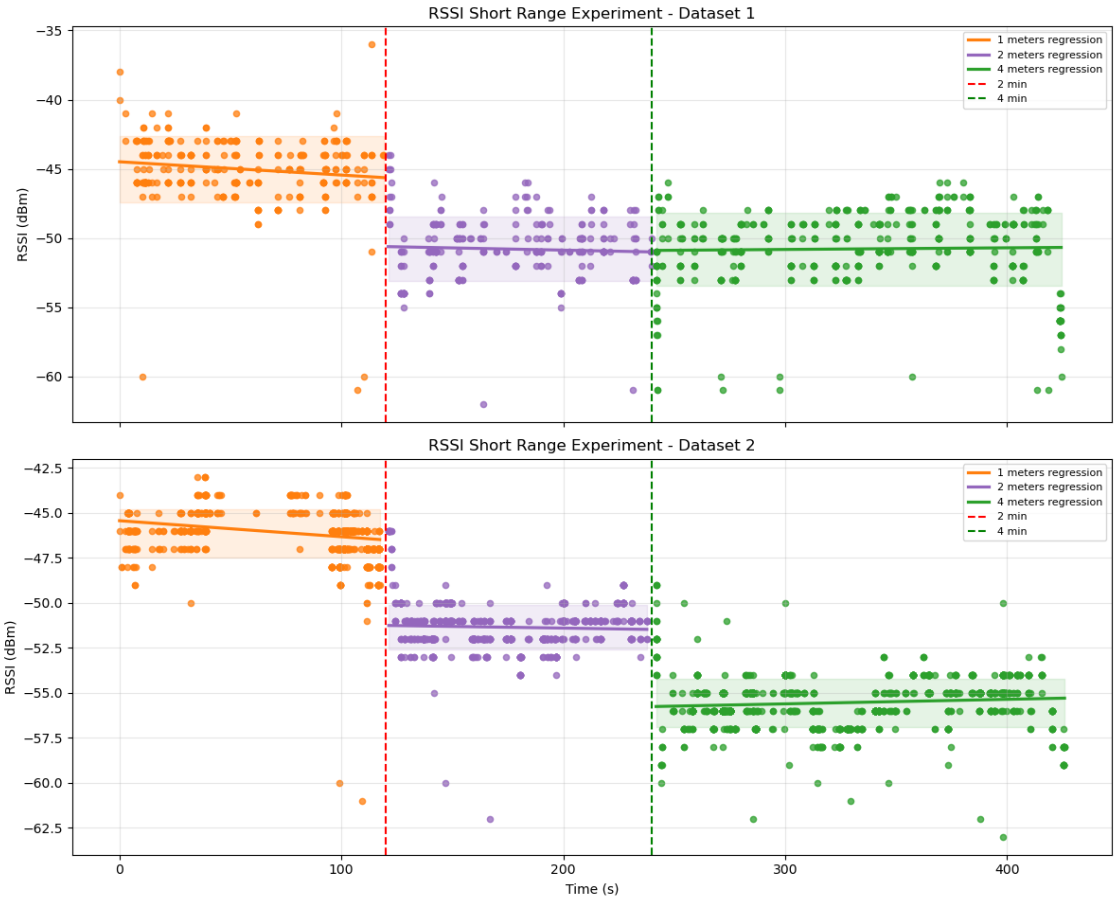

Figure A.3: Short-range RSSI measurements using a Nothing Phone 2 and an iPhone 15 Pro as the target devices. Dataset 1 is the Nothing Phone and Dataset 2 is the iPhone. Dashed vertical lines indicate transitions between distances, while regression lines illustrate the average signal level within each segment. Shaded regions represent the variability (standard deviation) of the measurements.

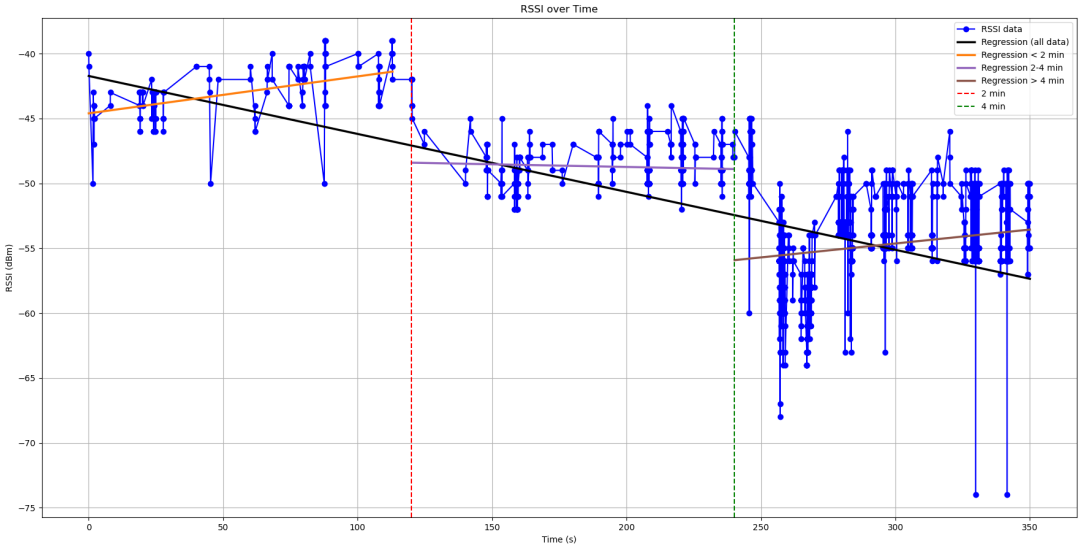

Figure A.4: Long-range RSSI measurements from initial capture with fitted regression lines for each distance segment. Each point corresponds to an individual captured packet. Dashed vertical lines indicate transitions between distances.

A.2 Long-Range RSSI Measurements and Obstruction Experiment

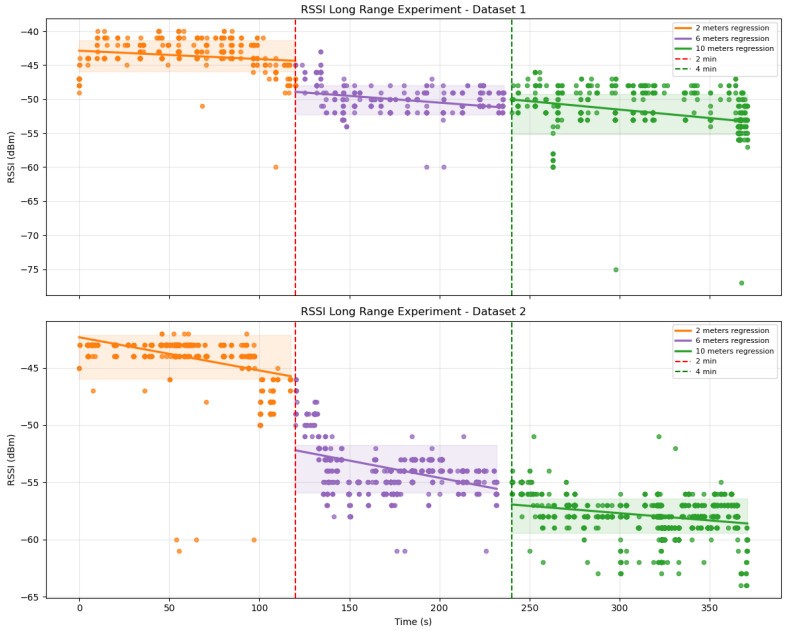

(a) Long-range RSSI measurements.

(b) Obstruction experiment RSSI measurements.

Figure A.5: RSSI measurements using a Nothing Phone 2 and an iPhone 15 Pro as target devices. Dataset 1 corresponds to the Nothing Phone 2, while Dataset 2 corresponds to the iPhone 15 Pro. Subfigure (a) presents the long-range measurements with distance transitions marked by dashed vertical lines, while subfigure (b) presents the obstruction experiment, where the dashed vertical line indicates when the device was moved behind a wall. Regression lines show average signal trends, and shaded regions represent measurement variability.

B

Sleep/Wake Cycle Screenshots

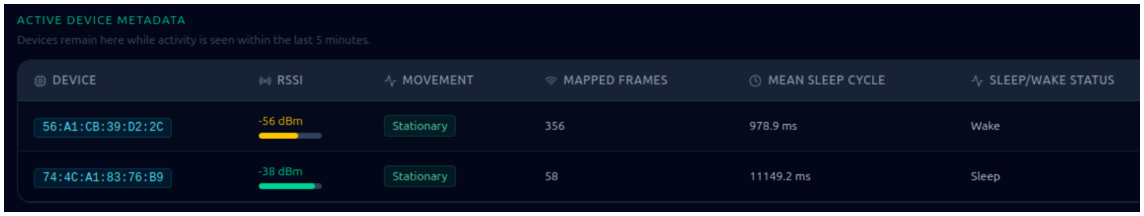

Figure B.1: Screenshot of the GUI depicting two devices (one phone 56:A1:***, and one laptop 74:4C:***) and their respective Sleep/Wake cycles and current status.

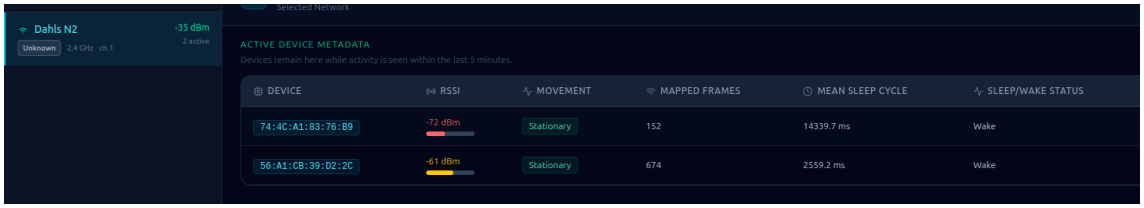

Figure B.2: Same capture session as Figure B.1 after five additional minutes.

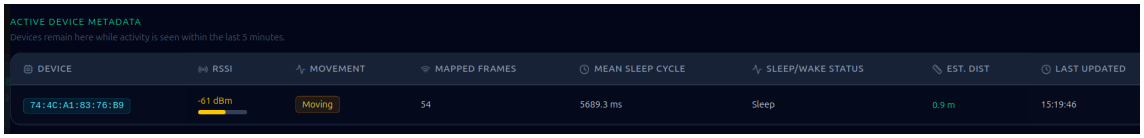

Figure B.3: Screenshot of another capture with the same laptop and AP at 5 minutes of capturing.

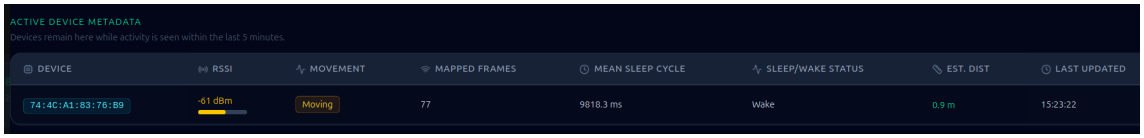

Figure B.4: Same capture session as Figure B.3 after five additional minutes.

C

DBSCAN Cluster and UMAP Results

The full cluster-level results are presented as tables for readability. The following tables report structural and timing-related statistics, rate and capability-related statistics, and vendor-specific feature statistics.

C.1 Cluster-Level Statistics

Table C.1: Structural and timing statistics for DBSCAN clusters.

Cluster	Mean frame length	Mean dt	Mean burst size
-1	186.171	68.348	5.021
0	180.400	0.008	1.400
1	155.621	0.005	1.576
2	195.371	5.435	4.115
3	129.714	0.001	1.008
4	143.000	0.000	1.000
5	175.000	0.045	9.218
6	208.049	0.282	5.542
7	229.000	0.006	1.571
8	107.712	1.127	9.319
9	95.083	1.666	1.500
10	114.317	2.587	3.621
11	184.385	0.045	4.923
12	169.690	3.574	2.952
13	190.789	0.043	6.211
14	190.294	1.671	3.176
15	125.000	14.591	1.942
16	162.873	4.946	3.820

Table C.2: Rate and capability-related statistics for DBSCAN clusters.

Cluster	Mean num rates	Mean num ext rates	Mean max rate	Mean rate avg	Mean HT cap	Mean VHT cap	Mean extcap len
-1	5.286	6.629	54.000	19.025	4559.514	162924519.364	8.707
0	8.000	4.000	54.000	18.875	2535.000	0.000	13.000
1	4.000	8.000	54.000	18.875	16429.000	0.000	9.363
2	4.000	8.000	54.000	18.875	16429.000	0.000	9.861
3	8.000	4.000	54.000	18.875	50923.355	0.000	12.626
4	8.000	4.000	54.000	18.875	4207.000	0.000	8.000
5	8.000	4.000	54.000	18.875	2349.000	0.000	10.000
6	4.000	8.000	54.000	18.875	2350.639	865189186.000	9.986
7	4.000	8.000	54.000	18.875	16429.000	0.000	10.000
8	8.000	4.000	54.000	18.875	0.000	0.000	8.986
9	4.000	8.000	54.000	18.875	0.000	0.000	0.000
10	8.000	4.000	54.000	18.875	6351.031	0.000	0.000
11	4.000	8.000	54.000	18.875	2532.846	865204626.000	9.000
12	8.000	4.000	54.000	18.875	2535.190	0.000	9.024
13	4.000	8.000	54.000	18.875	33.000	0.000	8.000
14	4.000	8.000	54.000	18.875	1845.000	865180772.824	9.941
15	8.000	0.000	54.000	25.875	6639.000	63992242.000	0.000
16	8.000	4.000	54.000	18.875	6639.000	0.000	0.000

Table C.3: Vendor-specific feature statistics for DBSCAN clusters.

Cluster	Mean vendor tag count	Mean vendor OUI count	Mean vendor OUI unique	Mean vendor type count	Mean vendor type unique	Mean vendor data count	Mean vendor data len
-1	2.021	2.021	1.836	2.021	2.021	0.446	4.707
0	1.000	1.000	1.000	1.000	1.000	0.000	0.000
1	0.000	0.000	0.000	0.000	0.000	0.000	0.000
2	3.000	3.000	3.000	3.000	3.000	1.000	12.000
3	0.000	0.000	0.000	0.000	0.000	0.000	0.000
4	1.000	1.000	1.000	1.000	1.000	0.000	0.000
5	2.000	2.000	2.000	2.000	2.000	0.000	0.000
6	3.000	3.000	3.000	3.000	3.000	1.000	8.000
7	4.000	4.000	4.000	4.000	4.000	1.000	12.000
8	1.000	1.000	1.000	1.000	1.000	0.000	0.000
9	1.000	1.000	1.000	1.000	1.000	0.000	0.000
10	0.000	0.000	0.000	0.000	0.000	0.000	0.000
11	1.000	1.000	1.000	1.000	1.000	0.000	0.000
12	1.000	1.000	1.000	1.000	1.000	0.000	0.000
13	3.000	3.000	3.000	3.000	3.000	1.000	10.000
14	1.000	1.000	1.000	1.000	1.000	0.000	0.000
15	0.000	0.000	0.000	0.000	0.000	0.000	0.000
16	1.000	1.000	1.000	1.000	1.000	0.000	0.000

C.2 UMAP Projections

The figures show the UMAP projection of the packet-level feature space for all sampled packets, colored by DBSCAN cluster assignment and collapsed vendor label.

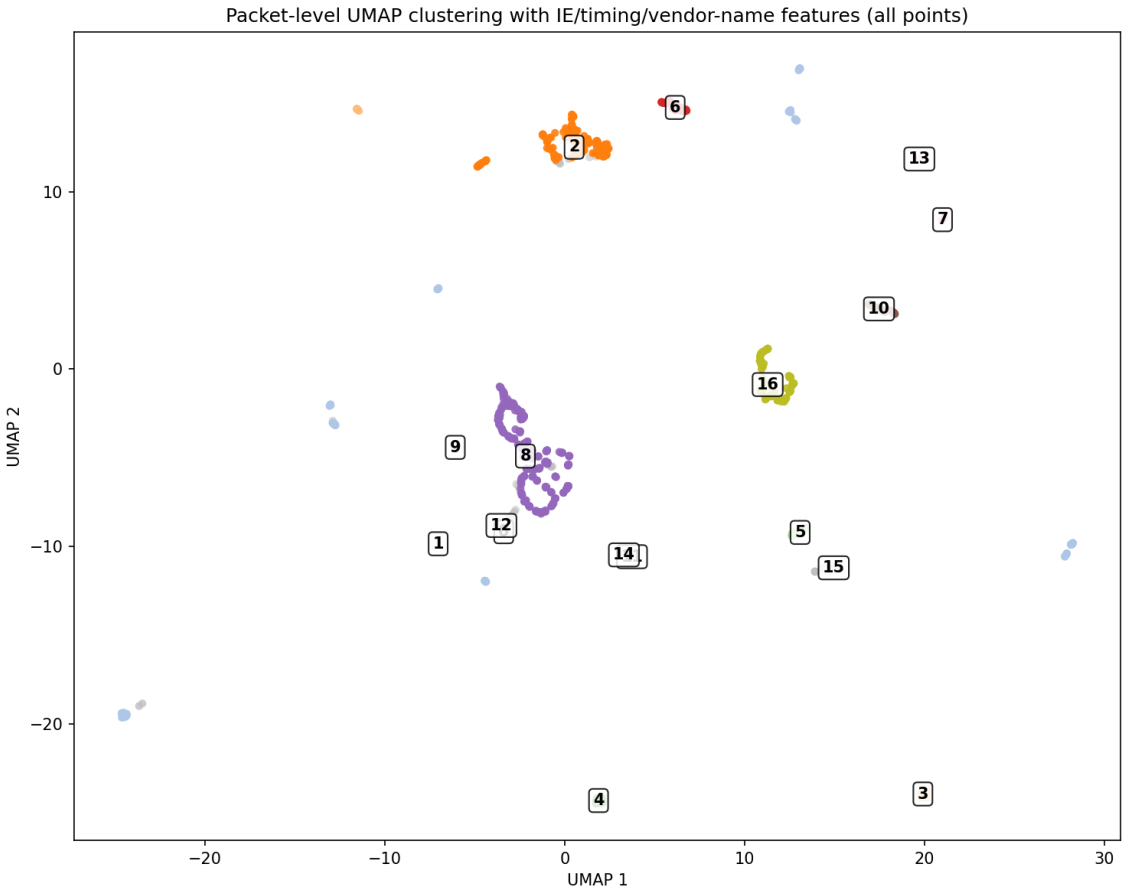

Figure C.1: Two-dimensional UMAP projection of the packet-level feature space for all sampled packets, colored by DBSCAN cluster assignment.

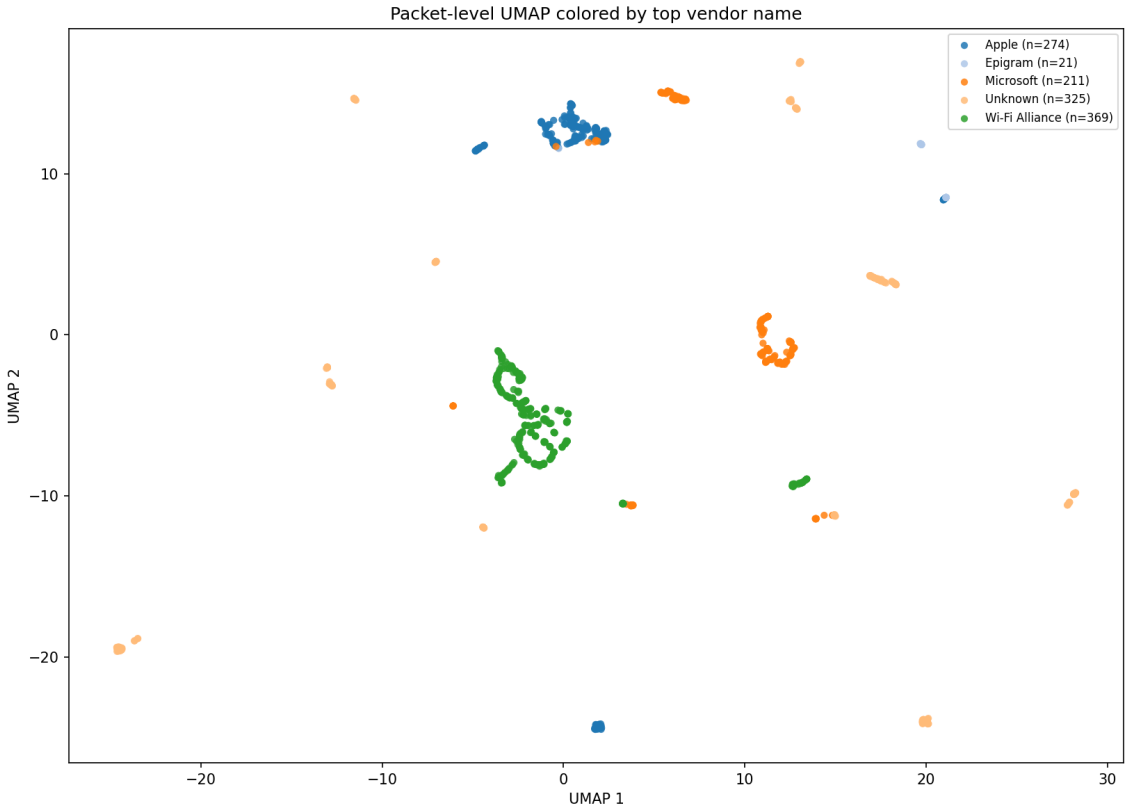

Figure C.2: Two-dimensional UMAP projection of the packet-level feature space for all sampled packets, colored by collapsed vendor label.

D

Ethical Analysis Using the Three-Party Model

To further assess the ethical implications of the study, the framework proposed by Hermansson and Hansson is applied. This model evaluates ethical risk by considering three key roles [18]:

- **Beneficiaries**, who gain from the activity
- **Risk-exposed individuals**, who may be affected by it
- **Decision-makers**, who design and control the process

In this thesis, the beneficiaries include the research community and society at large, as the work aims to improve understanding of privacy risks. The decision-makers are the research group and authors of the thesis, responsible for data collection, analysis, and interpretation. The risk-exposed group consists of individuals whose devices may be incidentally captured during passive monitoring.

In the following subsections, the assessment of ethical risk and responsibility is discussed, following the guiding questions proposed in the three-party model [18, 133–139]:

- "To what extent do the risk-exposed benefit from the risk exposure?" [18, p. 133]
- "Is the distribution of risks and benefits fair?" [18, p. 134]
- "Can the distribution of risks and benefits be made less unfair by redistribution or compensation?" [18, p. 135]
- "To what extent is the risk exposure decided by those who run the risk?" [18, p. 136]
- "Do the risk-exposed have access to all relevant information about the risk?" [18, p. 137]
- "Are there risk-exposed persons who cannot be informed or included in the decision process?" [18, p. 138]
- "Does the decision-maker benefit from other people's risk exposure?" [18, p.

139]

D.1 Assessment of Risks and Benefits

The risk-exposed individuals do not benefit directly from participation, as data collection occurs passively and without explicit involvement. However, indirect benefits may arise if the findings contribute to improved privacy protections in the future.

The distribution of risks and benefits is therefore asymmetrical, with risks primarily affecting nearby individuals while benefits are more broadly distributed. This imbalance is partially mitigated by limiting data collection to controlled environments and minimizing data retention.

D.2 Control and Awareness of Risk Exposure

Risk exposure is not actively controlled by the affected individuals, as they are generally unaware of the data collection process. Consequently, they cannot meaningfully consent to or evaluate the associated risks.

In addition, not all relevant information about the data collection process is accessible to potentially affected individuals. This includes details such as when data is captured, what features are extracted, and how the data is stored and used. These limitations are inherent to passive monitoring and represent a key ethical challenge.

D.3 Feasibility of Inclusion and Mitigation

In practice, it is not feasible to inform or include all potentially affected individuals in the data collection process. Physical constraints, such as multi-unit housing environments, make comprehensive notification difficult.

For example, posting notices at the experimental site may not reach individuals located in adjacent apartments or on different floors, and some individuals may not notice or understand such communication. Similarly, sending advance notifications, such as emails to all potentially affected individuals, could improve awareness, but introduces additional challenges related to feasibility, data access, and compliance with data protection regulations such as GDPR [19].

As a result, the study relies on minimizing harm rather than achieving complete inclusion. This includes limiting data scope, avoiding identity leakage, and ensuring secure handling of collected data.

D.4 Researcher Benefit and Responsibility

The researchers benefit academically from the collected data, as it enables the completion of the study. This raises an ethical consideration regarding the use of data obtained without consent.

However, the intent of the research is to investigate systemic privacy limitations rather than to exploit individual data. When conducted responsibly and transparently, such work can be justified by its potential contribution to improving privacy protections.

D.5 Overall Ethical Assessment

Applying the three-party model indicates that the project involves limited ethical risks. These risks are primarily indirect and arise from incidental data collection rather than targeted observation.

Given the mitigation measures implemented, the restricted scope of the study, and the broader societal benefits, the project can be considered ethically justifiable. However, the findings also highlight the need for continued discussion on how passive monitoring techniques should be regulated and communicated in practice.

The analysis therefore supports the conclusion that the project is ethically acceptable within the framework of Hermansson and Hansson's three-party model.

E

Unconnected Devices View

This appendix shows the implemented dashboard view for unconnected devices. The view lists detected devices that are not currently associated with the monitored access point.

Figure E.1: Implemented unconnected devices view of the visualization dashboard.

DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

UNIVERSITY OF
GOTHENBURG

CHALMERS
UNIVERSITY OF TECHNOLOGY