Wi-Fi as a Sensor: Capabilities, Challenges, and Defenses

Survey on Security and Privacy Defenses in Wi-Fi Sensing

Sofia Dimieva
Supervisor(s): Arash Asadi , Fabian Portner
[1]EEMCS, Delft University of Technology, The Netherlands

A Thesis Submitted to EEMCS Faculty Delft University of Technology,
In Partial Fulfilment of the Requirements
For the Bachelor of Computer Science and Engineering

Name of the student: Sofia Dimieva Final project course: CSE3000 Research Project
Thesis committee: Arash Asadi ,Fabian Portner

Abstract

This paper addresses the growing privacy concerns associated with Wi-Fi sensing, a technology that uses existing wireless infrastructure to extract sensitive information such as human presence, motion, breathing, and typing. Although Wi-Fi sensing enables cost-effective and scalable applications in smart homes, healthcare, and security, its passive nature allows adversaries to eavesdrop on Wi-Fi transmissions and analyze signal patterns for private data extraction. While researchers have proposed various defense mechanisms to counter these threats, a unified framework for categorizing and comparing these approaches, particularly in terms of their trade-offs and practicality, has been notably absent. This paper aims to fill this gap by providing a structured categorization of existing defense mechanisms against Wi-Fi sensing, including a detailed analysis of their limitations, and informing the design of more robust and privacy-preserving solutions. The key contributions include a comprehensive taxonomy of defense mechanisms based on their core strategy and operational layer, a comparative analysis of their effectiveness, feasibility, and trade-offs, and an identification of open challenges for future research.

1 Introduction

Wi-Fi sensing is a rapidly advancing technology that uses passive wireless communication signals to infer human presence, motion, and even fine-grained activities such as breathing or typing. Unlike traditional sensing systems, Wi-Fi sensing leverages existing wireless infrastructure, making it both cost-effective and highly scalable. As Wi-Fi devices become increasingly widespread and research in this field continues to advance rapidly, Wi-Fi sensing is gaining significant traction for applications in smart homes, healthcare, security, and human-computer interaction. However, the same capabilities that make Wi-Fi sensing powerful also raise serious privacy concerns, as passive adversary can simply eavesdrop on WiFi transmissions and analyze the signal patterns to sense private information

In response to these concerns, researchers have proposed various defense mechanisms intended to disrupt or limit the ability of adversaries to use Wi-Fi for sensing. These defenses range from signal obfuscation and MAC address randomization to physical-layer interference and environmental manipulation. However, the literature lacks a unified framework to categorize these approaches based on how they function and what trade-offs they introduce. Most existing work focuses on proposing new defenses or enhancing attack strategies, while relatively little attention has been paid to systematically comparing these defenses in terms of their practicality and effectiveness.

This paper aims to address that gap by providing a structured categorization of existing defense mechanisms against Wi-Fi sensing and offering a detailed analysis of their limitations. Through this analysis, we aim to inform the design of more robust, practical, and privacy-preserving solutions to mitigate the risks associated with Wi-Fi sensing technologies.

The key contributions of this paper are as follows:

- We present a comprehensive taxonomy of defense mechanisms against Wi-Fi sensing, organizing them by their core strategy and operational layer.
- We conduct a comparative analysis of these defenses in terms of effectiveness, feasibility, and trade-offs.
- We identify and discuss open challenges and limitations in existing defense strategies.

- We provide guidance for future research aimed at balancing privacy protection with Wi-Fi performance and usability.

The remainder of this paper is structured as follows: In Section 2, we review existing work and prior surveys on the topic. Section 3 provides an overview of Wi-Fi sensing and the associated privacy threats. In Section 4, we describe our methodology for comparing the tables. Section 5 introduces our taxonomy, outlines the main categories of defenses, and analyzes their limitations through comparative tables. Finally, Section 6 summarizes the key findings of this review and outlines potential avenues for future research.

2 Related Work

Prior research on Wi-Fi sensing has focused heavily on enabling accurate sensing applications, while relatively less attention has been paid to defending against passive sensing threats. Recent surveys by Liu et al. [14] and Geng et al. [11] categorize sensing threats and countermeasures, offering broad overviews but limited comparative analysis of defense limitations. Furthermore, these surveys tend to explain defenses primarily in the context of particular attacks, rather than providing a comprehensive categorization of distinct defense methodologies themselves.

This paper addresses that gap through a taxonomy and critical analysis of existing defenses against passive Wi-Fi sensing.

3 Background

Wi-Fi sensing leverages physical layer characteristics of wireless communication to extract environmental and contextual information without requiring dedicated sensors. As wireless signals propagate through space, they interact with objects and human bodies, undergoing attenuation, reflection, scattering, and multipath effects. These interactions introduce measurable perturbations that can be used to infer presence, motion, and fine-grained activities.

Modern Wi-Fi systems extensively utilize **Multiple-Input Multiple-Output (MIMO)** technology. MIMO employs multiple transmit and receive antennas at both the transmitter and receiver, enabling the creation of multiple spatial streams. A key technique employed with MIMO systems is **beamforming**. Beamforming dynamically adjusts the phase and amplitude of signals transmitted from each antenna to constructively interfere at the intended receiver and destructively interfere elsewhere.

To perform efficient beamforming, Wi-Fi devices exchange **Beamforming Feedback Information (BFI)**, which consists of measurements reported by the receiver back to the transmitter [22]. These reports help the transmitter select optimal beamforming parameters based on the current channel conditions. BFI is crucial for maintaining high data rates but also exposes sensitive channel characteristics that can be exploited for sensing.

Several signal metrics are commonly exploited in sensing systems:

- **Received Signal Strength Indicator (RSSI)** captures the power level of a received signal. *What it reveals:* coarse motion or presence changes (e.g., someone walking between the AP and client), proximity estimation, and large-scale environmental dynamics.
- **Channel State Information (CSI)** represents how wireless signals propagate from the transmitter to the receiver, capturing both amplitude attenuation and phase shift

across the sub-carriers. A typical MIMO-OFDM link models CSI as the complex matrix

$$H = \begin{bmatrix} h_{11} & \cdots & h_{1N} \\ \vdots & \ddots & \vdots \\ h_{M1} & \cdots & h_{MN} \end{bmatrix},$$

where each entry h_{mn} is the channel response for the n-th sub-carrier of the m-th Tx-Rx antenna pair. *What it reveals:* fine-grained human activity such as breathing, gestures, or keystrokes; detailed room or object mapping through multipath analysis.

- **Time-of-Flight (ToF)** and **Angle-of-Arrival (AoA)**, derived from CSI, estimate propagation delay and incident angle. *What they reveal:* centimeter-level ranging and direction finding-enabling indoor localization, trajectory tracking, and motion direction inference.

As illustrated in Figure 1, a passive eavesdropper positioned near a Wi-Fi link can intercept over-the-air transmissions between an access point and a client device. Beyond RSSI, AoA, and CSI, adversaries can exploit additional signal features embedded in protocol headers and physical-layer signatures. Unencrypted MAC frame headers expose device addresses, frame types, and session metadata. Inter-frame timing and traffic bursts reveal behavioral routines and user activity. Radiometric fingerprints - unique distortions caused by hardware imperfections - support persistent device tracking even when MAC addresses are randomized. Traffic metadata, such as packet size and transmission frequency, can further aid in app usage inference and environmental profiling. Together, these features constitute a rich sensing surface accessible to passive adversaries.

To aid in understanding how different defenses mitigate these threats, later sections of this paper analyze which of these signal features are protected by each defense.

Specifically, our taxonomy (Section 5) and accompanying feature mapping table (Table 5) highlight how each method addresses different aspects of this sensing surface.

4 Methodology

We conducted a structured literature review to identify and analyze defense mechanisms specifically aimed at mitigating passive Wi-Fi sensing. The review focused on papers published between 2014 and 2025, sourced from IEEE Xplore, ACM Digital Library, and Google Scholar.

4.1 Search and Selection Criteria

Search queries combined keywords such as *"Wi-Fi sensing"*, *"privacy"*, *"passive sensing"*, *"channel obfuscation"*, and *"defense mechanisms"*. A paper was considered eligible if it met the following criteria:

- It addressed defenses against **passive** Wi-Fi sensing-i.e., attacks that do not require the target's participation.
- It provided sufficient technical detail to enable classification and comparison of the proposed defense.
- It was published in a peer-reviewed journal or conference between 2014 and 2025.

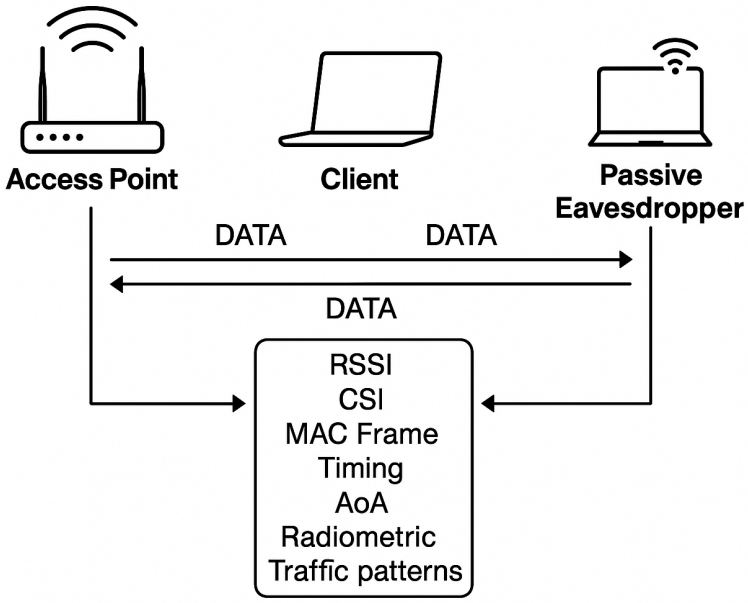

Figure 1: Passive Wi-Fi sensing architecture.

The initial search yielded approximately 113 papers. After screening titles and abstracts, we selected 42 papers for full-text analysis. From these, 26 papers met our inclusion criteria and were used to develop the taxonomy and comparative analysis.

4.2 Taxonomy Development

To structure our findings, we created a taxonomy based on the operational layer (e.g., physical, MAC, application) and the defense's core mechanism (e.g., signal distortion, feedback obfuscation, ML-based spoofing). These categories were derived inductively by analyzing recurring patterns and techniques across the reviewed literature.

4.3 Evaluation Criteria

Each defense was evaluated along four practical dimensions that reflect key design trade-offs in real-world deployments:

- **Effectiveness** - How well the defense reduces or degrades sensing accuracy.
- **Residual Leakage** - The extent of information that may still be inferred by an adversary.
- **Resource Dependence** - The level of additional hardware, firmware, or protocol modifications required.
- **Usability Impact** - The performance cost introduced by the defense, such as latency, throughput degradation, or reduced device compatibility.

These dimensions were selected based on trends observed during preliminary reviews and refined through iterative analysis.

5 Wi-Fi Sensing Defenses and Their Limitations

This section presents a comprehensive taxonomy of defenses against Wi-Fi sensing, organized by their operational layer and core defensive strategy. We classify existing approaches into two main categories: *RF Manipulation*, which directly alters the wireless propagation environment, and *Higher-Level Defenses*, which operate at the protocol or system layer. For each defense category, we provide detailed descriptions of representative techniques, analyze their practical limitations, and summarize comparative trade-offs in structured tables. Figure 2 visualizes the overall taxonomy.

5.1 RF Manipulation

RF Manipulation techniques directly interfere with or alter the wireless signals to prevent accurate sensing. This category is further divided into two subcategories:

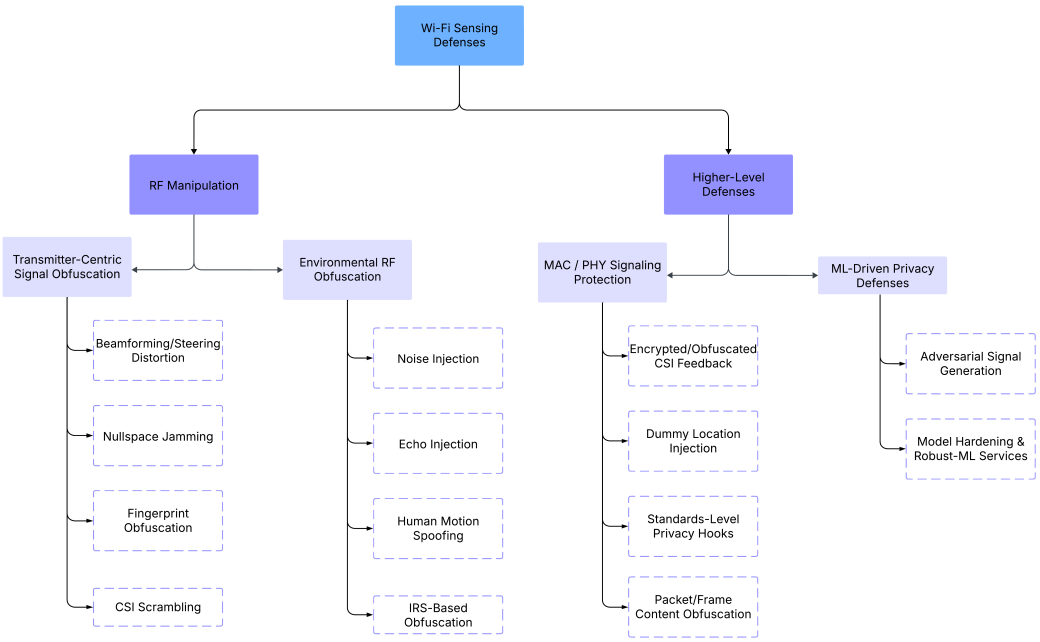

Figure 2: Taxonomy of Wi-Fi Sensing Defenses.

- **Transmitter-Centric Signal Obfuscation**: These defenses involve modifications at the transmitter to distort or obfuscate the emitted Wi-Fi signals.
  - **Beamforming/Steering Distortion**: This technique manipulates the directional properties of Wi-Fi signals at the transmitter to mislead sensing systems, particularly those relying on AoA for localization. A naive approach involves

beamforming a null towards the Access Point (AP) to hide the direct signal path. While this obfuscates location, it significantly degrades the Signal-to-Noise Ratio (SNR) and thus communication throughput. MIRAGE [3] offers an advanced method by intentionally adding a delay to the direct path signal, making it appear to have traveled a longer distance than reflected paths. This invalidates the sensing system's heuristic that the direct path is the shortest, thereby obfuscating the true AoA without compromising communication throughput or SNR.

- **Nullspace Jamming / Channel Encryption**: This defense physically encrypts Wi-Fi channels using MIMO capabilities to thwart unauthorized sensing while preserving communication and legitimate sensing capabilities for authorized users [15]. The core idea is to apply a secret pre-coding matrix (W) to the transmitted signals at the user's device [15]. For a legitimate receiver possessing the secret W, the effective channel can be properly decoded. However, for an unauthorized eavesdropper, the channel impulse response (H) is transformed into a random-looking effective channel ($H_{eff} = HW$) due to the unknown and constantly changing W, thereby preventing accurate sensing extraction [15]. This method effectively jams the channel for unauthorized parties by rendering the Channel State Information (CSI) unintelligible, without degrading the legitimate communication or sensing performance.
- **Fingerprint Obfuscation**: This defense aims to hide unique radiometric fingerprints inherent to Wi-Fi transceivers, which arise from manufacturing imperfections (e.g., Carrier Frequency Offset (CFO), In-phase/Quadrature (IQ) imbalance) [1]. RF-Veil [1] is an example of such a method that proposes to obfuscate these fingerprints for unauthorized eavesdroppers. While these fingerprints can be used for device identification and tracking, RF-Veil provides a robust solution to protect user privacy by masking these unique signatures without impacting communication performance or interoperability with standard Wi-Fi devices [1].
- **CSI Scrambling via Preprocessing**: This defense operates by modifying the transmitted Wi-Fi signal at the physical layer to embed false or altered channel information, thereby scrambling the CSI observed by unauthorized sensing devices. Two examples demonstrate different approaches to achieve this:
  - **OpenWiFi CSI Fuzzer** [13]: This method imposes an artificial, pre-known channel response onto the transmitted signal before it leaves the device [13]. As a result, the CSI perceived by any receiver is a combination of the true channel response and this artificial response. Only authorized receivers, privy to the secret artificial response, can subtract it to reconstruct the actual channel and perform sensing; unauthorized parties receive scrambled, unusable CSI [13].
  - **AntiSense** [7]: This defense involves an active device that continuously monitors the wireless channel. Upon detecting an activity, it superimposes a slightly modified copy of the current Wi-Fi frame onto the same channel [7]. This superimposed signal, differing subtly in phase, is designed to render the CSI information meaningless for any unauthorized attacker, while still allowing the original communication to proceed without disruption [7].

For each transmitter-centric signal obfuscation method discussed above, we provide a comparative overview in Table 1, evaluating their limitations, residual leakage, resource dependence, and usability impact.

Table 1: Limitations of Transmitter-Centric Signal Obfuscation Subcategories

Subcategory	**Limitations**	**Residual leakage**	**Resource dependence**	**Usability impact**	**Ref**
Beamforming/ Steering Distortion	Requires multi-antenna hardware and precise feedback; susceptible to RSSI-only attacks	Low-Moderate: RSSI/ time-of-flight still usable by attacker	MIMO-capable radios, firmware changes	Slight phase/timing variation; minimal throughput effect	[3]
Null-Space/ Artificial-Noise Jamming	Needs accurate nullspace estimation; can reduce spatial reuse	Low: broad CSI patterns remain	Multiple antennas, power amplifier overhead	Negligible if nullspace correctly computed	[15], [21]
Fingerprint Randomization	Requires key management and sync; potential standard-compliance issues	Moderate: coarse fingerprinting from other RF impairments	Firmware or driver updates	No impact on throughput	[1], [12]
CSI Scrambling/ Precoder Encryption	Custom firmware/SDR required; key distribution complexity	Moderate: adversary sees scrambled CSI but may exploit side-channels	Software/ hardware updates; key-management	Minimal latency if descrambling efficient	[13], [10], [7]

- **Environmental RF Obfuscation**: These defenses involve altering the wireless environment or introducing external factors to disrupt sensing.
  - **Controlled Interference (Noise Injection)**: This defense involves injecting carefully designed noise or interference into the wireless environment to perturb CSI and disrupt Wi-Fi sensing, while ensuring minimal impact on ongoing legitimate communication. Aegis [24] does this by generating and injecting customized, low-power random noise into the wireless channel, making eavesdroppers' sensing output unusable without harming legitimate Wi-Fi communication [24]. Similarly, PhyCloak [16] obfuscates human activity recognition by injecting spatially adaptive and time-varying signals that make the channel response appear constant to sensing applications, even during motion, thus distorting physical signatures without interfering with communication [16].
  - **Synthetic Multipath Injection (Echo Injection)**: This defense strategy involves intentionally generating and injecting artificial echoes or synthetic multipath components into the wireless environment. The goal is to distort the real CSI observed by sensing systems, making it difficult for attackers to accurately infer information like location or activity. Wi-Pi [5] exemplifies this approach by using an active device to superimpose a copy of the legitimate Wi-Fi frame onto the channel [5]. This artificially created signal, designed to confuse unauthorized sensing, causes the malicious receiver to perceive additional, false multipath components in the CSI, thereby scrambling the true channel characteristics without significantly degrading the original communication quality [5]. The resulting

received signal can be modeled as:

$$y(t) = s(t) + \alpha \cdot s(t - \Delta)$$

where $s(t)$ is the original transmitted signal, α is the attenuation factor of the injected echo, and Δ is the artificial delay. This delayed replica mimics a new multipath component, misleading the adversary's sensing algorithms.

- **Human Motion Spoofing**: This defense strategy involves generating and injecting artificial human motion signatures into the wireless environment to mislead Wi-Fi sensing systems. The goal is to create "fake humans" or trajectories that deceive sensing applications, thereby protecting the privacy of actual individuals. RF-Protect [19] is a notable example that enables privacy by injecting fake human motion into sensed data [19]. It achieves this through a novel hardware reflector design capable of modifying radio waves to create reflections at arbitrary locations, coupled with a generative mechanism that produces realistic human trajectories. This allows RF-Protect to create synthetic reflections that mimic human movement, thereby obfuscating real activities without requiring high bandwidth hardware or physical motion from the user [19].
- **IRS-Based Randomization (IRS-Based Obfuscation / Micro-Doppler Obfuscation with Passive Radar)**: This defense leverages Intelligent Reflecting Surfaces (IRS) to dynamically and randomly manipulate the wireless propagation environment, thereby confusing Wi-Fi sensing systems without actively transmitting signals. IRShield [20] is an example that implements this by applying random reflection profiles to a passive IRS [20]. The IRS, controlled by software, constantly updates its reflection profile (e.g., every 5 ms) to randomly reflect parts of the wireless signals [20]. This continuous and random manipulation of the signal's propagation path introduces noise and distortion into the CSI observed by adversaries, significantly reducing their ability to detect human motion or accurately localize individuals [20].

For each RF obfuscation method discussed above, we provide a comparative overview in Table 2, evaluating their limitations, residual leakage, resource dependence, and usability impact.

5.2 Higher-Level Defenses

Higher-Level Defenses operate at the protocol or system layers, intervening in how Wi-Fi devices communicate and interact rather than directly manipulating RF signals. These methods aim to prevent sensing by altering communication patterns, metadata, or user interaction.

- **MAC / PHY Signaling Protection**: These mechanisms aim to protect specific information exchanged within Wi-Fi protocols or at the physical layer.
  - **Encrypted/Obfuscated CSI Feedback**: This category of defenses is designed to protect the CSI and beamforming feedback reports from unauthorized access and exploitation by Wi-Fi sensing systems [22]. Defenses in this category operate by applying transformations or encryption to the feedback information itself before it is transmitted. The fundamental principle is to render the transmitted CSI

Table 2: Limitations of Environmental RF Obfuscation Subcategories

Subcategory	**Limitations**	**Residual leakage**	**Resource dependence**	**Usability impact**	**Ref**
Controlled Noise Injection	Requires dedicated hardware, sensitive to room layout, careful calibration	Moderate: coarse presence/motion still detectable	External noise emitter(s); power and installation	Low throughput impact; moderate setup overhead	[24], [16]
Echo/ Multipath Injection	Precise timing/alignment needed; risk of self-interference if misaligned	Low-Moderate: static channel features persist	External relay/reflector device; power	Minimal if well-calibrated; adds device deployment	[8], [9]
Human-Motion Spoofing	Complex hardware and control logic; phantom signals may confuse occupants	High: adversary sees ghost reflections but may infer some real motion	Reflective surfaces or active spoofers	Low after deployment; user confusion risk	[19], [25]
IRS-Based Obfuscation	Expensive metasurface, complex calibration, geometry-sensitive	Moderate: attacker may correlate IRS settings over time	Reconfigurable Intelligent Surface hardware	Negligible to link throughput; visual /budget overhead	[20], [2]

or feedback unintelligible or misleading to any unauthorized eavesdropper. This is achieved by manipulating the data in such a way that only legitimate, authorized receivers, possessing a shared secret or a pre-defined recovery mechanism, can accurately reconstruct the original, usable information. This ensures that communication integrity is maintained while preventing privacy breaches through passive sensing [14]. A prime example of this defense type is the **LeakyBeam mitigation** [22]. This defense applies a temporally varying random transformation to the original BFI before transmission [22]. If V represents the original BFI and Q is a randomly generated, time-varying complex matrix or vector (known to the legitimate AP), the obfuscated BFI, V_{obf}, is generated through an element-wise product:

$$V_{obf} = Q \odot V$$

where $\odot$ denotes element-wise multiplication [22]. For an unauthorized eavesdropper, Q is unknown and constantly changing, making V_{obf} appear as random noise, thus rendering the CSI information within it unintelligible and useless for sensing purposes [22]. Conversely, a legitimate AP, which possesses the same Q (or can derive it), can easily recover the original BFI, V, by performing the inverse operation:

$$V = V_{obf} \oslash Q$$

where $\oslash$ denotes element-wise division [22]. This allows the AP to correctly utilize the BFI for its intended communication purposes (e.g., beamforming ad-

justments) without any degradation in performance. This approach typically requires firmware or protocol updates on Wi-Fi devices [22].

- **Dummy Location Injection**: This defense strategy aims to protect user privacy by actively introducing false or misleading location information into the wireless environment or into the sensing system's data processing pipeline, thereby preventing accurate inference of the user's true position or movement [14]. This can be achieved through various means, such as transmitting fabricated Wi-Fi frames that carry deceptive location metadata, or by manipulating existing signals to create the illusion of a user being present at a dummy location. The mechanism could also involve a device actively generating false Wi-Fi fingerprints (e.g., RSSI or CSI patterns) that, when observed by a sensing system, lead to the computation of an erroneous location. The overarching principle is to confuse an unauthorized sensing system into calculating an inaccurate or randomized location for the user, while ideally allowing authorized systems to filter out or disregard the injected dummy data [14].
- **Standards-Level Privacy Hooks**: These defenses propose integrating privacy-enhancing features directly into Wi-Fi communication standards (e.g., IEEE 802.11), aiming to provide protection against Wi-Fi sensing at the protocol layer. This approach involves modifying the fundamental operations of Wi-Fi devices to inherently obscure or protect sensitive information from unauthorized sensing, rather than relying on external mechanisms or post-processing. These methods involve embedding privacy features directly into physical (PHY) or Media Access Control (MAC) layer signaling. This can include mechanisms like built-in CSI scrambling or directional nulling, which are designed to operate transparently to legitimate communication while rendering sensing data unusable for unauthorized parties. For instance, a built-in CSI scrambling scheme might use randomly generated scrambling vectors applied to the transmitted signal [4]. This process ensures that the CSI observed by eavesdroppers appears as random noise, effectively preventing them from extracting meaningful physical activity data [4]. Crucially, such schemes are designed to allow legitimate Wi-Fi receivers, equipped with a pre-shared secret key or a specific decoding mechanism, to successfully recover the true channel state and maintain normal communication performance [4]. A prime example of such a defense is **CSS (Built-in Channel State Scrambling)** [4] which implements a built-in channel state scrambling mechanism that uses randomly generated scrambling vectors to emulate human activities in the observed CSI [4]. This allows all transmitted frames to be scrambled, preventing eavesdroppers from recovering human physical activity from the channel state. Despite this obfuscation, CSS ensures that legacy receivers can still successfully decode the frames, and legitimate Wi-Fi sensors, possessing the shared secret key, can recover the true activity with high success rates [4]. This method offers an "always-on" protection for CSI without sacrificing legitimate Wi-Fi link performance [4].

For each MAC/PHY signaling protection method discussed above, we provide a comparative overview in Table 3, evaluating their limitations, residual leakage, resource dependence, and usability impact.

- **ML-Driven Privacy Defenses**: This defense category leverages the power of machine learning (ML) to either create deceptive signals that mislead Wi-Fi sensing sys-

Table 3: Limitations of MAC/PHY Signaling Protection Subcategories

Subcategory	**Limitations**	**Residual leakage**	**Resource dependence**	**Usability impact**	**Ref**
Encrypted/ Obfuscated CSI & Beam Feedback	Requires protocol extensions and key-management; interoperability hurdles	Low: RF side-channels (RSSI, AoA) remain	Firmware/driver upgrades, PKI support	Minor latency/ processing overhead	[22], [6]
Standards-Compliant Scrambling Hooks	Backwards-compatibility issues; limited scrambling granularity	Low: attacker may exploit alternate control fields	Firmware or driver updates	Minimal throughput/latency impact	[10], [7]
Packet/ Frame-Content Obfuscation	Adds padding and delays, increasing latency and framing overhead	Moderate: timing patterns still partially inferable	Software or driver-level support	Increased latency and jitter	[4]

tems or to fortify existing ML models against privacy-exploiting adversarial attacks.

- **Adversarial Signal Generation**: These defenses generate "adversarial examples" in the physical layer, which appear legitimate but cause sensing models to infer incorrect or random information, thereby protecting actual user privacy. This is often achieved by training deep learning models, particularly sequence models, to generate CSI patterns that mimic specific activities or states while an actual different activity is occurring [18]. This generated adversarial CSI, when introduced into the wireless environment or directly into the sensing data stream, is designed to perturb the input of a target sensing system's ML model, leading to misclassification of the sensed activity or an inability to detect any activity at all. The challenge lies in ensuring these generated signals are effective against a range of sensing models and do not disrupt legitimate Wi-Fi communication. A prime example is **Wi-Spoof** [18], which utilizes deep learning techniques to create adversarial CSI sequences that, when processed by a HAR system, lead to the misidentification of human activities. For instance, it can make a system perceive "no activity" when a person is actually moving, effectively obscuring their presence and actions without physically altering the environment or directly jamming the signal [18].
- **Model Hardening and Robust-ML Services**: This defense category focuses on making the machine learning models used in Wi-Fi sensing inherently more robust and resistant to adversarial attacks. Rather than manipulating the physical signals, these methods modify the learning algorithms, model architectures, or training procedures to enhance their resilience against input perturbations that could compromise privacy by using adversarial training, where models are trained on both clean and adversarially perturbed data to improve their generalization and robustness against attacks [23]. Other strategies include input regularization, defensive distillation, or incorporating provable robustness guarantees through formal verification techniques. The goal is to ensure that even

if an adversary introduces subtle changes to the input data (e.g., through minor signal alterations), the sensing model's output remains consistent and does not reveal private information or fall victim to deception. **SecureSense** [23] is an example of a learning framework designed for model hardening in device-free human activity recognition [23].

For each adversarial machine learning-based defense discussed above, we provide a comparative overview in Table 4, evaluating their limitations, residual leakage, resource dependence, and usability impact.

Table 4: Limitations of Adversarial ML-Centric Defense Subcategories

Subcategory	**Limitations**	**Residual leakage**	**Resource dependence**	**Usability impact**	**Ref**
Adversarial Signal Generation/ Spoofing	Requires knowledge of target model; specialized full-duplex hardware	High: robust/retrained attackers can adapt	Compute resources; specialized radios	Negligible to throughput; complex to deploy	[18], [25]
Robust-Model/ Defensive ML	Continuous retraining needed; arms-race with attacker	Moderate: adaptive adversaries can find new exploits	High compute for training/inference	None to Wi-Fi link; high development cost	[23], [17]

In summary, **Transmitter-Centric Signal Obfuscation** demands multi-antenna or firmware upgrades and precise channel knowledge, but residual RSSI or timing cues often persist. **MAC/PHY Signaling Protection** can hide control plane information with minimal throughput cost, but RF side channels remain exploitable and backward compatibility issues frequently arise. Finally, **Adversarial ML-Centric Defenses** impose heavy computing and model maintenance burdens and tend to fall short against adaptive attackers who can retrain or circumvent the countermeasures. Together, these comparisons underscore that no single defense fully eliminates information leakage or deployment overhead. Finally, to conclude the taxonomy and clarify the effectiveness of each defense mechanism, Table 5 presents a feature-level mapping. It identifies which specific signal features are protected or remain exposed under each defense. This summary helps illustrate the gaps in coverage across defenses and emphasizes the difficulty in achieving comprehensive sensing obfuscation that covers them all.

6 Conclusions and Future Work

This paper provides a concise taxonomy and detailed analysis of defense mechanisms against passive Wi-Fi sensing. Our survey categorized defenses into RF Manipulation and Higher-Level Defenses. RF Manipulation techniques, including Transmitter-Centric Signal Obfuscation (e.g., Beamforming/Steering Distortion, Nullspace Jamming, Fingerprint Obfuscation, CSI Scrambling) and Environmental RF Obfuscation (e.g., Controlled Interference, Synthetic Multipath Injection, Human Motion Spoofing, IRS-Based Obfuscation) , often require specialized hardware or firmware, and struggle with residual leakage. Higher-Level Defenses include MAC/PHY Signaling Protection (e.g., Encrypted/Obfuscated CSI Feedback, Dummy Location Injection, Standards-Level Privacy Hooks) and ML-Driven Privacy

Table 5: Mapping of Wi-Fi Sensing Defenses to Protected Signal Features

Defense	Category	RSSI	CSI	MAC	Time	AoA	RadioID	Traffic
Aegis [24]	Controlled Interference	No	Yes	No	Yes	Yes	No	No
MIRAGE [3]	Beamforming Distortion	No	Yes	No	Yes	Yes	No	No
RF-Veil [1]	Fingerprint Obfuscation	No	No	No	No	No	Yes	No
OpenWiFi CSI Fuzzer [13]	CSI Scrambling	No	Yes	No	No	No	No	No
AntiSense [7]	CSI Overwrite	No	Yes	No	Yes	No	No	No
Wi-Pi [5]	Synthetic Multipath Injection	No	Yes	No	No	Yes	No	No
RF-Protect [19]	Human Motion Spoofing	No	Yes	No	No	Yes	No	No
IRShield [20]	IRS-Based Obfuscation	No	Yes	No	No	Yes	No	No
LeakyBeam [22]	Encrypted BFI Feedback	No	Yes	Yes	Yes	Yes	No	No
CSS [4]	Standards-Level CSI Scrambling	No	Yes	No	Yes	No	No	No
Wi-Spoof [18]	Adversarial CSI Generation	No	Yes	No	No	No	No	No
SecureSense [23]	Model Hardening	No	Yes	No	No	No	No	No

Defenses (e.g., Adversarial Signal Generation, Model Hardening). These approaches often involve protocol extensions, robust key management, or computational intensity, and face challenges like backward compatibility and ineffectiveness against adaptive adversaries.

Our analysis consistently shows that no single defense fully eliminates information leakage or deployment overhead. The feature-level mapping highlights varying protection levels for different signal features, underscoring the difficulty of comprehensive sensing obfuscation.

Future work can focus on developing simulations or proof-of-concept implementations to validate key techniques or explore using multiple techniques simultaneously.

7 Responsible Research

To ensure that this paper follows commonly adopted standards of academic integrity, we have followed the principles recommended by the Netherlands Code of Conduct for Research Integrity focusing on reproducibility and transparency of this research. We have detailed the methods employed for gathering and selecting papers in Section 4. All papers used throughout the research are cited and can be found in the References section, allowing for full traceability and verification of our sources.

Large language models (LLMs) were employed in this study strictly for the purpose of

sentence paraphrasing and rephrasing to enhance clarity and conciseness of the text. Their use was limited to stylistic improvements and did not involve content generation, analysis, or interpretation of research findings.

References

[1] Luis Fernando Abanto-Leon, Andreas Bäuml, Gek Hong Sim, Matthias Hollick, and Arash Asadi. Stay connected, leave no trace: Enhancing security and privacy in wifi via obfuscating radiometric fingerprints. *Abstract Proceedings of the 2021 ACM SIGMETRICS / International Conference on Measurement and Modeling of Computer Systems*, 2020.

[2] Antonios Argyriou. Obfuscation of human micro-doppler signatures in passive wireless radar. *IEEE Access*, 11:40121–40127, 2023.

[3] Roshan Sai Ayyalasomayajula, Aditya Arun, Wei Sun, and Dinesh Bharadia. Users are closer than they appear: Protecting user location from wifi aps. *Proceedings of the 24th International Workshop on Mobile Computing Systems and Applications*, 2022.

[4] Alexander Bienstock, Paul Rösler, and Yi Tang. Asmesh: Anonymous and secure messaging in mesh networks using stronger, anonymous double ratchet. *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, 2023.

[5] Poh Chan, Ju-Chin Chao, and Ruey-Beei Wu. A wi-fi-based passive indoor positioning system via entropy-enhanced deployment of wi-fi sniffers. *Sensors*, 23:1376, 01 2023.

[6] Renato Lo Cigno, Francesco Gringoli, Marco Cominelli, and Lorenzo Ghiro. Integrating csi sensing in wireless networks: Challenges to privacy and countermeasures. *IEEE Network*, 36:174–180, 2022.

[7] Marco Cominelli et al. Antisense: Standard-compliant csi obfuscation against unauthorized wi-fi sensing. *Computer Communications*, 185:92–103, 2022.

[8] Marco Cominelli, Francesco Gringoli, and Renato Lo Cigno. Non intrusive wi-pi csi obfuscation against active localization attacks. *2021 16th Annual Conference on Wireless On-demand Network Systems and Services Conference (WONS)*, pages 1–8, 2021.

[9] Marco Cominelli, Francesco Gringoli, and Renato Lo Cigno. On the properties of device-free multi-point csi localization and its obfuscation. *Comput. Commun.*, 189:67–78, 2022.

[10] Xiao Deng, Dongyu Xia, Xun Wang, Shuyu Shi, and Wei Wang. Css: Built-in channel state scrambling for secure wi-fi based sensing. *2024 IEEE 44th International Conference on Distributed Computing Systems (ICDCS)*, pages 1119–1130, 2024.

[11] Ruixu Geng et al. A survey of wireless sensing security from a role-based view: Victim, weapon, and shield. *arXiv preprint arXiv:2412.03064*, 2024.

[12] Hadi Givehchian, Nishant Bhaskar, Alexander Redding, Han Zhao, Aaron Schulman, and Dinesh Bharadia. Practical obfuscation of ble physical-layer fingerprints on mobile devices. *2024 IEEE Symposium on Security and Privacy (SP)*, pages 2867–2885, 2024.

[13] Xianjun Jiao, Michael Tetemke Mehari, Wei Liu, Muhammad Aslam, and Ingrid Moerman. openwifi csi fuzzer for authorized sensing and covert channels. *Proceedings of the 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, 2021.

[14] Xingyu Liu, Xin Meng, Hancong Duan, Ze Hu, and Min Wang. A survey on secure wifi sensing technology: Attacks and defenses. *Sensors*, 25(6), 2025.

[15] Jun Luo, Hangcheng Cao, Hongbo Jiang, Yanbing Yang, and Zhe Chen. Mimocrypt: Multi-user privacy-preserving wi-fi sensing via mimo encryption. *2024 IEEE Symposium on Security and Privacy (SP)*, pages 2812–2830, 2023.

[16] Yue Qiao, Ouyang Zhang, Wenjie Zhou, Kannan Srinivasan, and Anish Arora. Phycloak: Obfuscating sensing from communication signals. In *USENIX Annual Technical Conference*, 2016.

[17] Yamini Shankar and Ayon Chakraborty. Practical defense against adversarial wifi sensing. *2024 IEEE International Conference on Advanced Networks and Telecommunications Systems (ANTS)*, pages 1–6, 2024.

[18] Aryan Sharma, Deepak Mishra, Sanjay Jha, and Aruna Seneviratne. Wi-spoof: Generating adversarial wireless signals to deceive wi-fi sensing systems. *J. Inf. Secur. Appl.*, 91:104052, 2025.

[19] Jayanth Shenoy, Zikun Liu, Bill Tao, Zachary Kabelac, and Deepak Vasisht. Rf-protect: privacy against device-free human tracking. *Proceedings of the ACM SIGCOMM 2022 Conference*, 2022.

[20] Paul Staat, S R Mulzer, Stefan Roth, Veelasha Moonsamy, Aydin Sezgin, and Christof Paar. Irshield: A countermeasure against adversarial physical-layer wireless sensing. *2022 IEEE Symposium on Security and Privacy (SP)*, pages 1705–1721, 2021.

[21] Yuwei Wang, Li Sun, and Qinghe Du. Multi-antenna signal masking and round-trip transmission for privacy-preserving wireless sensing. *IEEE Transactions on Information Forensics and Security*, 19:6305–6320, 2024.

[22] Rui Xiao, Xiankai Chen, Yinghui He, Jun Han, and Jinsong Han. Lend me your beam: Privacy implications of plaintext beamforming feedback in wifi. *Proceedings 2025 Network and Distributed System Security Symposium*, 2025.

[23] Jianfei Yang, Han Zou, and Lihua Xie. Securesense: Defending adversarial attack for secure device-free human activity recognition. *IEEE Transactions on Mobile Computing*, 23:823–834, 2022.

[24] Yao Yao, Yan Li, Xin Liu, Zicheng Chi, Wei Wang, Tiantian Xie, and Ting Zhu. Aegis: An interference-negligible rf sensing shield. *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications*, pages 1718–1726, 2018.

[25] Siwang Zhou, Wei Zhang, Dan Peng, Yonghe Liu, Xing-Yu Liao, and Hongbo Jiang. Adversarial wifi sensing for privacy preservation of human behaviors. *IEEE Communications Letters*, 24:259–263, 2020.