

How to Evaluate Privacy in WiFi Sensing Systems?

An Assessment of Metrics and Their Application to ISAC Use Cases

Óscar G. Martins^{*†}, Marco Gomes[†], João P. Vilela^{*}

^{*}CRACS/INESCTEC, CISUC and Dep. of Computer Science, Faculty of Sciences, University of Porto, Portugal

[†]Instituto de Telecomunicações and Department of Electrical and Computer Engineering, University of Coimbra, Portugal

Abstract—Integrated Sensing and Communication (ISAC) reuses wireless signals for both communication and sensing, improving spectrum efficiency and enabling new applications, but also creating privacy risks. Adversaries can exploit standard communication signals to infer sensitive information without users’ consent or awareness. This article addresses the lack of clear, structured metrics for evaluating privacy in ISAC by proposing a taxonomy of metrics and demonstrating their application to representative sensing use cases.

Our results indicate that data-similarity metrics can help reveal privacy leakage, but they may not fully capture the risks of unauthorized communication-based sensing. Feature extraction choices strongly influence these metrics and their interpretation. By contrast, adversarial-success metrics more directly measure an attacker’s practical ability to infer sensitive information, although they remain dependent on the assumed attack model, sensing setup, and learning algorithms.

Index Terms—WiFi-based Sensing, Privacy-Aware ISAC, Privacy Metrics

I. INTRODUCTION

Integrating sensing into wireless communication infrastructure has enabled Integrated Sensing and Communication (ISAC). This improves hardware and spectrum efficiency [1] and supports new applications by using sensing data to enhance communication and communication waveforms to enable ubiquitous sensing [2].

In WiFi systems, communication data can be used to train Machine Learning (ML) models to understand and classify events occurring in the propagation medium. While this enables applications in smart homes, healthcare, and industrial automation, it also raises privacy concerns. By repurposing standard signals, adversaries can “see” through walls and identify human presence, gestures, and even biometric signatures without consent [3, 4]. Because WiFi networks are so widespread, such pervasive sensing can remain unnoticed by targets, posing severe privacy risks.

National Funds finance this work through the Portuguese funding agency, FCT - Fundação para a Ciência e a Tecnologia, within project LA/P/0063/2020. This article/publication is based on work from COST Action 6GPYSEC (CA22168), supported by COST (European Cooperation in Science and Technology). This work is funded by national funds through FCT – Fundação para a Ciência e a Tecnologia, I.P., under the support UID/50014/2025 (<https://doi.org/10.54499/UID/50014/2025>) and, when eligible, co-funded by EU funds under project/support UID/50008/2025 – Instituto de Telecomunicações, with DOI identifier <https://doi.org/10.54499/UID/50008/2025>. M. Gomes also thanks P. Birungi (WorldTalents) for partially funding his research group’s activities. Óscar Martins wishes to acknowledge financial support from the Portuguese funding institution Fundação para a Ciência e a Tecnologia (FCT) under the grant 2024.06005.BD.

Most ISAC research focuses on improving sensing performance, while fewer works address security and privacy [5, 6]. These studies often use the same decision models as legitimate sensing systems and evaluate privacy only through model performance, without applying dedicated privacy metrics. As a result, they obscure both the privacy risks of ISAC and the effectiveness of privacy-preserving mechanisms (PPMs).

The reasoning behind this is two-fold. First, privacy has no universal metric and depends on the data and the attacker’s goal [7]. Second, privacy leakage and legitimate sensing are often measured with the same metrics, typically accuracy [8], differing only in perspective. In ISAC, this leaves a need for metrics that capture what sensitive information can actually be inferred, rather than only utility-oriented performance. For example, a Channel State Information (CSI) obfuscation method for unauthorized sensing in device-free localization is proposed [6]; however, the attacker performance is only evaluated with regard to accuracy, without looking at the distance error between the predicted and actual locations, which can provide better insights into the adversary’s performance.

This article addresses the lack of structured privacy evaluation metrics for ISAC systems by presenting: (1) a taxonomy of metrics; (2) their application to selected ISAC sensing use cases spanning spatiotemporal and identification-related aspects; and (3) an evaluation of the privacy implications of unauthorized communication-based sensing in ISAC systems using the defined metrics, threat model, and objectives.

Thus, this paper provides the following contributions:

- We present a taxonomy of metrics for privacy analysis in ISAC systems;
- We define a comprehensive threat model with passive and active adversarial scenarios, covering different levels of access and capability for unauthorized sensing in ISAC systems, as well as the use of PPMs to mitigate the privacy implications;
- We evaluate the proposed metrics on selected ISAC sensing use cases to quantify privacy leakage in spatiotemporal and identification tasks.

The remainder of this paper is organized as follows. Section II reviews WiFi-based sensing, privacy threats, and evaluation approaches. Section III presents our privacy metric taxonomy for ISAC. Section IV-A describes the system and threat models, datasets, and sensing tasks. Section V evaluates the privacy impact of unauthorized communication-based sensing and mitigation mechanisms. Section VI concludes the paper.

II. WiFi SENSING, PRIVACY THREATS AND PRIVACY EVALUATION

Integrated Sensing and Communication (ISAC) blurs the line between communication and sensing for efficient spectrum usage [1]. WiFi-based sensing uses standard communication signals for human-centric applications, including localization [5], Human Activity Recognition (HAR) [9, 10], and Person Identification (PI) [11, 12].

Although limited by the generalization of the underlying ML models, WiFi-based sensing raises privacy concerns, as adversaries can exploit the same signals to infer sensitive information without users' consent or awareness [4].

Assessing these privacy implications requires a clear view of the risks and attack vectors. Prior work has studied malicious sensing, such as inferring user presence and activities through walls [9], and mitigation mechanisms such as signal obfuscation and access control [3, 6]. In [13], the authors propose transmitter-side IQ symbol manipulation to distort location information, evaluating privacy leakage through localization accuracy and a location reliability metric, $\mathcal{L}_R$, based on the Euclidean distance between the estimated and true location. In [3], only location accuracy is used to evaluate privacy-preserving capability, while [14] evaluates selective sensing, i.e., only allowed devices can perform sensing, using location accuracy only, without considering errors introduced by legitimate or illegitimate sensing agents.

For privacy protection in other WiFi-based sensing applications, such as Human Activity Recognition and Person Identification, evaluation also relies on accuracy metrics, such as the classification accuracy of a HAR or PI model. In [15], the authors evaluate a CSI fuzzing obfuscation mechanism for HAR using only the accuracy of a Gated Recurrent Unit (GRU) model, without stating whether the model is trained on clean or obfuscated data, while in [6] the attacker model is trained on obfuscated data, but only accuracy is used to evaluate privacy leakage.

While accuracy is useful for model performance, it does not capture the full privacy impact. It overlooks leakage that can persist under obfuscation and adversarial adaptation through new training strategies. It also misses exploitable cues such as prediction confidence, which adversaries could use to target easier samples or train a stronger model.

III. PRIVACY ASSESSMENT METRICS FOR ISAC SYSTEMS

To comprehensively evaluate the privacy implications of unauthorized communication-based sensing in ISAC systems, we define a taxonomy of privacy assessment metrics that captures both theoretical data leakage and practical adversarial success. The proposed metrics are classified into two main categories: Data Leakage Metrics and Adversarial Success Metrics.

A. Data Leakage Metrics

Data Leakage Metrics quantify how much information can be inferred from communication signals, independent of the adversarial model or ML technique used. They capture the

inherent privacy risk in the signals and the sensitive information they may reveal, even when PPMs such as obfuscation or other anonymization methods are applied. In practice, these metrics measure the relationship between sensing data and private attributes, as well as the similarity and separability of data across sensitive property values, which indicate potential privacy leakage.

- **Mutual Information (MI):** Mutual Information quantifies how much a communication signal reveals about a sensitive property, such as location or activity. Higher MI indicates greater privacy leakage, since more information can be inferred from the signal. It can be used to assess theoretical data leakage and compare the effectiveness of different PPMs. If we denote X as the communication signal and Y as the sensitive property, the mutual information can be calculated as:

$$I(X; Y) = H(Y) - H(Y|X) \\ = \sum_{x \in X} \sum_{y \in Y} P(x, y) \log \left(\frac{P(x, y)}{P(x)P(y)} \right), \quad (1)$$

where H and P are the entropy and probability distribution function, respectively. Since MI depends on the number of possible outcomes, it can be normalized by the entropy of the sensitive property, $H(Y)$, to obtain a normalized mutual information (NMI) in the range $[0, 1]$. Thus, NMI is calculated as:

$$NMI(X; Y) = \frac{I(X; Y)}{H(Y)}. \quad (2)$$

- **Statistical Distance:** Statistical distance measures the separability between probability distributions, though it does not directly quantify the amount of information that can be inferred about the sensitive property. A common measure is the Kullback-Leibler divergence (KL), which can be used to compare the distributions corresponding to different values of a sensitive property, where a larger KL suggests a greater potential for privacy leakage as the distributions associated with different property values are more distinct.

If " Q " and " P " are probability distributions representing two subsets of X corresponding each to a class of the a private attribute Y , then, the KL can be calculated as:

$$D_{KL}(P||Q) = \sum_{x \in X} P(x) \log \left(\frac{P(x)}{Q(x)} \right). \quad (3)$$

Since KL is not symmetric, it can be useful to obtain a symmetric measure of distributional separability, such as the Jensen-Shannon Divergence (JSD). The JSD can be calculated from KL as:

$$JSD(P||Q) = \frac{1}{2} D_{KL}(P||M) + \frac{1}{2} D_{KL}(Q||M), \quad (4)$$

where $M = \frac{1}{2}(P+Q)$ is the average distribution of P and Q . The JSD ranges between 0 and 1, where 0 indicates that the two distributions are identical and 1 indicates that they are completely different.

B. Adversarial Success Metrics

Adversarial Success Metrics evaluate how effectively an adversary can exploit leaked information to achieve malicious goals. In ISAC systems, these goals may include sensing tasks such as localization, HAR, or PI. These metrics are important because they reflect the adversary's practical ability to infer sensitive information from communication signals. However, they depend strongly on the adversarial model, sensing framework, and ML methods, and therefore may not fully capture privacy risk in general.

- **Adversarial Success Rate (ASR):** This metric measures the percentage of successful inferences made by the adversary under the conditions defined in the threat model. ASR depends on the specific sensing task and on how the adversarial model defines a successful inference. It aggregates privacy breaches across all samples and provides an overall view of the adversary's performance. Thus, it can be computed as:

$$ASR = \frac{N_{\text{Successful Inferences}}}{N_{\text{Total Inferences}}} \times 100\% \quad (5)$$

- **Accuracy:** Accuracy represents the percentage of correct classifications made by the adversary; however, it does not account for other characteristics that can be exploited by adversaries, such as the confidence of the model in its predictions, and restrictions on the adversary's objectives.

$$Accuracy = \frac{N_{\text{Correct}}}{N_{\text{Total}}} \times 100\% \quad (6)$$

- **Top- k Identification Accuracy:** Percentage of samples whose true class appears among the model's top- k predictions. It captures the adversary's ability to narrow the candidate set, even without a correct top-1 prediction.

$$Acc_{top-k} = \frac{N_{\text{correct in top-}k}}{N_{\text{Total Number of Samples}}} \times 100\% \quad (7)$$

- **Error:** In addition to accuracy, measuring errors reveals limitations of the adversary's inferences. For example, for localization, average distance error is more informative than accuracy because it indicates how close incorrect predictions are to the true location. The error metric depends on the task and requires a distance function d (e.g., Euclidean or cosine distance) [7]. We define error as the mean distance between inference and the true value:

$$Error = \frac{1}{N} \sum_{i=1}^N d(\text{Adversary's Inference}_i, \text{True Value}_i) \quad (8)$$

These metrics can be used to form a concise framework to evaluate privacy risks from unauthorized communication-based sensing in ISAC systems, covering both theoretical data leakage and practical adversarial success.

IV. SYSTEM AND EXPERIMENTAL MODEL

The proposed privacy assessment metrics aim to evaluate the privacy implications of unauthorized communication-based

sensing across scenarios and tasks. Data Leakage Metrics capture the theoretical risk posed by different signals, while Adversarial Success Metrics require a defined attacker model. Thus, we introduce the following threat model for ISAC systems.

A. System and Threat Model

We assume a WLAN with a multi-antenna AP (Tx) operating at 2.4/5 GHz and multiple UEs (Rx) in the environment that periodically collect CSI for sensing. The adversary (Eve) is an external eavesdropper aiming to obtain CSI for unauthorized sensing while avoiding detection. We consider two data-collection scenarios for Eve:

- **Scenario 1 (passive):** Eve either compromises legitimate UEs or connects her own UEs to an available WLAN from outside the sensing area to extract CSI and perform sensing.
- **Scenario 2 (active):** Eve actively probes the environment with her own transmitter to obtain CSI. This attack allows it to bypass transmitter-side obfuscation as it uses clean CSI estimation pilots.

To perform sensing, the attacker needs a trained inference model; thus, we assume a strong attacker with access to a suitable model. Model architecture is outside this paper's scope, as our goal is not state-of-the-art accuracy, but evaluation of the proposed privacy metrics.

B. Sensing Environments and Datasets

To evaluate the performance of the proposed metrics in WiFi-based sensing scenarios, we utilize two distinct datasets and sensing tasks: localization and gait-based person identification (PI).

For the first task, we use the dataset from [5], which contains CSI data collected in a controlled environment (see Fig. 1) with multiple positions and passive/active setups at 802.11ac/5 GHz. This addresses both Scenarios 1 and 2. The dataset includes clean and obfuscated recordings using two privacy-preserving mechanisms (PPMs) proposed by the authors: "random walk" and "swipe" obfuscation techniques. These techniques are implemented by a specialized repeater that injects variable time delays to scramble CSI and hinder unauthorized sensing. In the "swipe" technique, delays sweep linearly between maximum and minimum values, whereas in the "random walk" technique, delays change stepwise following a random Markov process with uniform step distribution [5]. The obfuscation level is controlled by a parameter ϵ , where higher values correspond to stronger obfuscation.

For the second task, we use the dataset from [11], with the layout shown in Fig. 2. The CSI data is collected from 6 receivers positioned in the environment, operating in the 802.11n standard at 2.4 GHz, which addresses Scenario 1, the passive attacker, of our threat model. The dataset includes multiple individuals, allowing for the evaluation of PI tasks based on the gait patterns of the individuals walking in the environment in the 4 different paths.



Fig. 1: Layout for the dataset in [5]. “RX” represent receiver devices, “TX_P” and “TX_A” are passive and active attack transmitters, stars “A”, “B”, and “C” are possible obfuscator positions. P_1 to P_8 denote the 8 prospective user locations.




Fig. 2: Layout for the dataset in [11].

These datasets provide CSI data and private attribute labels, i.e., user positions for the localization task and user identities for the PI gait-based task. Notice that CSI is complex-valued, and thus, it can be represented in terms of its amplitude and phase, or in terms of its real and imaginary parts. In this work, we focus on the CSI Phase data, as it has been shown to be more effective for sensing tasks than the amplitude data [8], and thus, it is more relevant for the evaluation of privacy implications in ISAC systems.

V. PRIVACY METRICS EVALUATION

In this section, we apply the proposed metrics to evaluate the privacy implications of unauthorized communication-based sensing in ISAC systems under the defined threat model and using the selected datasets.

A. Data Leakage

To assess the dependence between the communication signals and the sensitive property, we compute the NMI between the data and the private labels (e.g. possible positions or users), and the JSD divergence between the corresponding data distributions for each pair of classes. Both metrics range from 0 to 1, with higher values indicating greater potential privacy leakage.

While these metrics can be computed directly from CSI Phase data, the high dimensionality of the data leads to the need for dimensionality reduction to better compute them.

Thus, we employ Linear Discriminant Analysis (LDA) [16], a supervised technique to find a linear combination of discriminant components that better separate classes (e.g., P_1 and P_2 in Fig. 1) in the raw data, providing insights into privacy leakage for the specific sensing task. Alternatively, unsupervised techniques can be applied to evaluate data leakage more generally, though resulting feature distributions are likely less task-relevant. To address this latter case, we employ a k-means clustering approach, where features are grouped together, reducing the dimensionality of the data to k categorical groups.

Thus, following Eq. (2), NMI is computed between the features, i.e., full CSI Phase or reduced-dimensionality features (LDA or k-means), as the random variable X , and the private labels, i.e., user positions for the localization task and user identities for the PI gait-based task, as the random variable Y . For mean divergence, we compute JSD between each pair of distributions Q and P , i.e., different user positions or identities (classes), as defined in Eq. (4). Thus, the mean JSD is computed as

$$\text{Mean JSD} = \frac{1}{C} \sum_{i=1}^{N-1} \sum_{j=i+1}^N \text{JSD}(P_i || P_j) \quad (9)$$

where N is the number of classes of private attribute Y , P_i and P_j are the probability distributions for classes i and j , respectively, and $C = N(N-1)/2$ is the number of class pairs combinations.

To assess the impact of feature extraction techniques on data leakage, we use Gait Identification[11] as an example, with results shown in Table I. The value for k in the k-means clustering was chosen as the number of user classes, i.e., 10, to allow for a better comparison with the LDA features.

TABLE I: Mean NMI and JSD for Gait Identification using the dataset in [11].

Features	Mean NMI	Mean JSD
Full CSI	0.0378	1
Continuous LDA	0.0389	0.7725
k-means	0.0378	0.0636

While NMI remains similar across feature extraction methods, JSD varies notably, with LDA producing higher values than k-means. This suggests that LDA features are more separable and may reveal more sensitive information, highlighting the impact of feature extraction on leakage assessment.

For the localization task [5], the NMI between LDA-extracted features and private labels, along with the interclass JSD, is shown in Table II. Due to data availability in this dataset, we focus on evaluating the impact of PPMs (obfuscation) on data leakage.

In the clean scenario, both NMI and JSD show relatively low values, suggesting theoretically low data leakage. However, the results do not show a clear correlation between reduced data leakage and obfuscation techniques. While NMI and Mean JSD are higher for clean data in the active scenario,

TABLE II: Mean NMI and JSD for the localization task using clean and obfuscated data with random walk (rw) and swipe techniques from [5].

Obfuscation	Parameter	Active		Passive	
		NMI	JSD	NMI	JSD
clean		0.1829	0.0219	0.0003	0.0119
rw	0.01	0.0023	0.0146	0.0014	0.0130
rw	0.05	0.0025	0.0124	0.0010	0.0124
rw	0.1	0.0010	0.0124	0.0015	0.0129
rw	0.2	0.0019	0.0127	0.0017	0.0127
swipe	0.5	0.0620	0.0162	0.0015	0.0127

this does not hold for the passive scenario, where NMI is actually higher for some obfuscation techniques. JSD shows no significant differences across obfuscation techniques in the passive scenario, misaligned with expected behavior.

As we will see in the next subsection, the adversary’s success metrics do not necessarily reflect this low theoretical data leakage. Thus, statistical similarity metrics can be informative but insufficient to capture privacy risks in ISAC sensing. Table I further shows that choice of features strongly affects leakage metrics, underscoring the need for careful feature selection when interpreting privacy outcomes.

B. Adversarial Success

Adversarial Success Metrics require a specific adversarial model and attacker methodology to be defined. In this work, we assume a strong adversary that trains a simple 1D Convolutional Neural Network (CNN) model to perform localization based on CSI data and a ResNet18-based model [8] for Gait Identification using single receiver CSI phase visualizations.

For the Gait Identification task, we evaluate top-5 accuracy, which is shown in Table III. The adversary model achieves a base accuracy of approximately 82%, which is significantly higher than the expected random classification performance of 10% for this task. Likewise, top- k accuracy demonstrates that the adversary’s model is able to narrow down to the four most likely targets with an accuracy above 96%, indicating that the model successfully extracts private information.

TABLE III: Top- k Accuracy for Gait Identification using the dataset in [8].

k	1	2	3	4	5
Accuracy %	82.78	92.82	94.89	96.30	97.24

Accuracy and Top-3 accuracy for the localization task are shown in Fig. 3. In this scenario, these metrics show that obfuscation heavily degrades the adversary’s accuracy performance in both active and passive scenarios, as accuracy drops from nearly 100% and 80% (respectively) to under 20% (near the 12.5% random baseline). The top-3 accuracy drops from nearly 100% to nearly 40% in the obfuscated scenarios (slightly above the random-classification baseline of 37.5%). Thus, even though the model’s top-1 accuracy is only about 80%, it can still narrow the true location to the three most




Fig. 3: Averaged accuracy, top-3 accuracy, and mean distance error for the localization task under the passive (solid line) and active attacks (dashed line), with and without obfuscation techniques presented in [5].

likely candidates. This demonstrates the value of top- k metrics for assessing adversary performance beyond top-1 accuracy.

The mean distance error (right scale of Fig. 3) increases substantially with obfuscation, reaching ~ 2.5 m in both scenarios (close to the average inter-location distance of ~ 2.72 m for the locations in Fig. 1). These results indicate the proposed PPM [5] drives the model toward random-like outputs in both accuracy and error, implying the adversary cannot reliably extract location information when PPM is applied.

TABLE IV: Mean Adversarial Success Rate (ASR) based on confidence threshold (th) for active (A.) and passive (P.) sensing with/without obfuscation, based on the data in [5].

Scenario	Confidence ASR Threshold (th)								
	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9
P. Clean	80.73	79.99	79.18	78.29	77.61	76.85	75.80	74.54	73.04
P. Obf.	15.01	14.66	14.32	13.97	13.59	13.19	12.74	12.21	11.52
A. Clean	98.64	98.56	98.48	98.32	98.08	97.82	97.59	97.39	96.93
A. Obf.	15.99	15.57	15.15	14.71	14.26	13.78	13.26	12.67	11.92

Table IV shows the variation of ASR with higher confidence thresholds, enabling evaluation in higher-risk scenarios. While passive attacks exhibit lower-confidence predictions, the decrease seen from $th = 0.5$ to $th = 0.9$ indicates most correct predictions have high confidence. With obfuscation, despite ASR remaining extremely low (below random chance), the model is likely frequently confidently wrong.

Fig. 4 uses mean distance-error thresholds (th_e) to assess how PPMs affect adversarial error. In the clean scenario, performance appears saturated, with many errors likely far from the true position and outside the plotted range. In obfuscated scenarios, attackers still retain partial location information. We can also study the behavior of different PPM parameters. “rw0.2” gives the lowest ASR at $th = 1.5$ m but degrades as th increases, while “rw0.1” provides better overall privacy by keeping ASR low for both small (see Fig. 3) and larger thresholds.

As shown in this section, Adversarial Success Metrics offer a direct and practical view of adversary performance and are easier to interpret for privacy analysis. However, they




Fig. 4: ASR based on mean distance error threshold th_e for the localization task under the passive scenario in [5].

depend on the chosen ML methods and may not capture all privacy risks. In contrast, Data Leakage Metrics quantify statistical links between communication signals and sensitive properties, but they are harder to interpret and do not always align with adversarial performance. Therefore, both metric families are needed for a comprehensive privacy assessment of unauthorized communication-based sensing in ISAC systems, but a better understanding of data-leakage metrics, including how to properly assess them and when they are more relevant, is required.

VI. CONCLUSION

In this work, we proposed a comprehensive set of privacy assessment metrics to evaluate the implications of unauthorized communication-based sensing in ISAC systems. We defined both *Data Leakage Metrics*, which aim to quantify the theoretical amount of information that can be inferred from communication signals, and *Adversarial Success Metrics*, which evaluate how effectively an adversary can exploit leaked information to achieve specific malicious objectives. We applied these metrics under a defined threat model using two datasets for localization and gait identification tasks, considering both active and passive attack scenarios as well as PPMs.

Our results show that while Data Leakage Metrics can provide insights into the statistical relationship between communication signals and sensitive properties, which helps to understand the inherent privacy risks without relying on specific adversarial models, they require careful interpretation and do not necessarily correlate with the practical performance of adversaries. On the other hand, Adversarial Success Metrics provide a more direct evaluation of performance and are easier to interpret in terms of privacy implications, but they depend on the specific use cases and corresponding ML techniques employed and may fail to capture the full extent of privacy risks.

Overall, our findings highlight the need for a more in-depth evaluation of statistical and theoretical methods concerning the relationship between communication signals and the private information they contain, including computational methods and their practical implications. Furthermore, they underscore

the necessity of a comprehensive and standardized evaluation framework that combines both Data Leakage Metrics and Adversarial Success Metrics to fully understand the privacy implications of unauthorized ISAC sensing and to effectively assess the performance of different PPMs in mitigating these risks. Finally, privacy should be assessed jointly with both legitimate sensing and communications performance, and metrics used should take into account the privacy-utility trade-off that proposed PPMs bring.

REFERENCES

- [1] K. Wu et al., *Joint Communications and Sensing*. Wiley, 2022.
- [2] J. A. Zhang et al., “Enabling joint communication and radar sensing in mobile networks—a survey,” *IEEE Communications Surveys & Tutorials*, vol. 24, no. 1, pp. 306–345, 2022.
- [3] R. L. Cigno et al., “Integrating csi sensing in wireless networks: Challenges to privacy and countermeasures,” *IEEE Network*, vol. 36, no. 4, pp. 174–180, 2022.
- [4] Ó. G. Martins et al., “Delving into security and privacy of joint communication and sensing: A survey,” *IEEE Open Journal of the Communication Society*, vol. 6, pp. 4978–5004, 2025.
- [5] M. Cominelli et al., “Antisense: Standard-compliant csi obfuscation against unauthorized wi-fi sensing,” *Computer Communications*, vol. 185, pp. 92–103, 2022.
- [6] Z. Chu et al., “Privacy-preserving wifi sensing in wsns via csi obfuscation,” *Computers & Security*, vol. 157, p. 104594, 2025.
- [7] I. Wagner et al., “Technical privacy metrics: A systematic survey,” *ACM Comput. Surv.*, vol. 51, no. 3, Jun. 2018.
- [8] Y. Gu et al., “WiGRUNT: WiFi-enabled gesture recognition using dual-attention network,” *IEEE Transactions on Human-Machine Systems*, vol. 52, no. 4, pp. 736–746, 2022.
- [9] F. Meneghello et al., “SHARP: Environment and person independent activity recognition with commodity IEEE 802.11 access points,” *IEEE Transactions on Mobile Computing*, 2022.
- [10] Y. Zhang et al., “Widar3.0: Zero-effort cross-domain gesture recognition with Wi-Fi,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 11, pp. 8671–8688, 2022.
- [11] J. Zhang et al., “Gate-ID: WiFi-based human identification irrespective of walking directions in smart home,” *IEEE Internet of Things Journal*, vol. 8, no. 9, pp. 7610–7624, 2021.
- [12] C. Li et al., “WiHF: Gesture and user recognition with Wi-Fi,” *IEEE Transactions on Mobile Computing*, vol. 21, no. 2, pp. 757–768, 2022.
- [13] M. Cominelli et al., “IEEE 802.11 CSI randomization to preserve location privacy: An empirical evaluation in different scenarios,” *Computer Networks*, vol. 191, p. 107970, 2021.
- [14] G. A. Alghisi et al., “For your eyes only: Bridging privacy and sensing in wi-fi networks through csi obfuscation,” in *23rd Mediterranean Communication and Computer Networking Conference*, 2025, pp. 1–6.
- [15] T. Zhang et al., “Privacy protection in wifi sensing via csi fuzzing,” in *2024 IEEE/ACM Symposium on Edge Computing (SEC)*, 2024, pp. 410–416.
- [16] P. Belhumeur et al., “Eigenfaces vs. fisherfaces: Recognition using class specific linear projection,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 19, no. 7, pp. 711–720, 1997.