RESEARCH-ARTICLE

Let It Beam: Enabling Selective and Secure CSI-based Sensing via Beamforming

GIOVANNI ANGELO ALGHISI, University of Brescia, Brescia, BS, Italy

GIOVANNI PERIN, University of Brescia, Brescia, BS, Italy

FRANCESCA MENEGHELLO, University of Padua, Padua, PD, Italy

FRANCESCO GRINGOLI, University of Brescia, Brescia, BS, Italy

Open Access Support provided by:

University of Padua

University of Brescia

Published: 04 November 2025

Citation in BibTeX format

WiNTECH '25: ACM Workshop on Wireless Network Testbeds, Experimental evaluation & Characterization
November 4 - 8, 2025
Hong Kong, China

Conference Sponsors:
SIGMOBILE

Let It Beam: Enabling Selective and Secure CSI-based Sensing via Beamforming

Giovanni Angelo Alghisi
University of Brescia & CNIT
Brescia, Italy
giovanni.alghisi@unibs.it

Giovanni Perin*
University of Brescia & CNIT
Brescia, Italy
giovanni.perin@unibs.it

Francesca Meneghello†
University of Padova
Padova, Italy
francesca.meneghello.1@unipd.it

Francesco Gringoli
University of Brescia & CNIT
Brescia, Italy
francesco.gringoli@unibs.it

Abstract

Integrated sensing and communication (ISAC) strategies are key components of next-generation wireless networks, including Wi-Fi systems, enabling new services and supporting network management operations. The main idea behind this integration is that the channel state information (CSI), which is continuously estimated for communication purposes, can be leveraged for environmental sensing, especially with the aid of artificial intelligence (AI) algorithms. However, this capability comes with privacy concerns, as passive eavesdroppers, even when using commercial off-the-shelf (COTS) devices, can estimate the CSI and potentially infer sensitive information. In this paper, we propose a new technique to mitigate this risk, leveraging the potentialities of multiple-input multiple-output (MIMO) systems. We design an obfuscation and de-obfuscation system that conceals the real CSI from eavesdroppers, while enabling trusted devices to reverse the distortion and perform sensing normally. We implemented a prototype of our system through software-defined radios (SDRs) and evaluated the effectiveness of our proposed approach considering a device-free localization task. The results show that the obfuscation makes it unfeasible to perform sensing at unauthorized devices (the accuracy drops to about 20%) while the de-obfuscation at legitimate devices allows reaching almost 100% in sensing accuracy. To support reproducibility, we make the dataset and code publicly available.

CCS Concepts

• **Networks** → **Network experimentation**; **Network measurement**; • **Hardware** → **Wireless devices**.

Keywords

Integrated sensing and communication (ISAC), Wi-Fi sensing, channel state information (CSI), MIMO systems, beamforming, physical layer security, device-free localization.

ACM Reference Format:
Giovanni Angelo Alghisi, Giovanni Perin, Francesca Meneghello, and Francesco Gringoli. 2025. Let It Beam: Enabling Selective and Secure CSI-based Sensing via Beamforming. In *ACM Workshop on Wireless Network Testbeds, Experimental evaluation & Characterization (WiNTECH '25), November 4–8, 2025, Hong Kong, China.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3737895.3768295

1 Introduction

In recent years, the integrated sensing and communications (ISAC) paradigm has gained momentum, and it is expected to play a key role in next-generation wireless systems for improving network functionalities and offer new services to the users [8]. This approach enables simultaneous communication and sensing, allowing inference of dynamic environmental changes without additional sensors. Current state-of-the-art solutions are promising, enabling applications such as intrusion detection, activity and gesture recognition, vital signs monitoring, indoor localization, and tracking [8, 7].

The channel state information (CSI) is the key enabler for this dual functionality. In orthogonal frequency-division multiplexing (OFDM)-based networks, such as Wi-Fi, it represents the channel frequency response (CFR) between the transmitter and the receiver devices over all the operational sub-carriers. This information is continuously estimated at

*Also with University of Padova, Padova, Italy.
†Also with Northeastern University, Boston, USA.

WiNTECH '25, Hong Kong, China

ACM ISBN 979-8-4007-1972-1/2025/11
https://doi.org/10.1145/3737895.3768295

the receiver side thanks to known predefined symbols, called long training symbols (LTSs), included in the preamble of every transmitted frame. The CSI is traditionally used to equalize the signals at the receiver and properly decode the data. However, since the CFR depends on the environment, it is possible to infer information about the surroundings from the CSI itself, in most cases using machine learning (ML) models to learn the underlying patterns [11].

This raises a serious privacy concern: Current wireless communication standards allow any receiver to estimate the CSI. This means even a free-riding attacker can collect CSI data and train models to perform sensing [9], and this can be done using commercial off-the-shelf (COTS) devices [6]. As an example of privacy violation, imagine someone is at home with a nearby unauthorized party sniffing the Wi-Fi packets that in-home devices are exchanging. By extracting the CSI, the latter could extrapolate sensitive information about the former without them even knowing. What makes this threat even more dangerous is that it is completely passive [9].

A promising defense against passive attacks in single-input single-output (SISO) systems is *transmitter-side obfuscation* [2]. The strategy here is to introduce pre-distortions to the signals before their transmission. These distortions can be induced by applying an artificial channel response or by multiplying the modulated Wi-Fi subcarriers by time-varying weights in the frequency domain. When the distortions are applied consistently across an entire frame, the receiver can still decode the data; however, the CSI appears to vary randomly over time, making it difficult for an attacker to distinguish between environmental and artificially-crafted changes. This approach is especially appealing for ISAC systems, as a legitimate receiver, knowing the distortion pattern, can recover the real CSI and preserve sensing capabilities [1].

This approach can be further enhanced by considering the multi-antenna transmissions characterizing multiple-input multiple-output (MIMO) systems. In these scenarios, the CSI observed at each receiving antenna results from the linear combination of the CFRs of all transmitting antennas. By dynamically changing the precoding weights over time, the resulting CSI becomes highly variable, increasing the efforts required by a hypothetical attacker to perform sensing. This idea has been shown to be effective in [10] and [3]. Nonetheless, both works have limitations: [10] breaks communication compatibility with other Wi-Fi devices, whereas [3] hinders sensing for both legitimate receivers and attackers, making the solution unsuitable for ISAC systems.

In this work, we build upon the foundation established in [3] and extend it to enable secure and selective sensing. Specifically, we propose a system that:

(1) employs time-varying beamforming precoding weights to introduce controlled virtual distortions into the CSI, effectively *obfuscating* it from eavesdroppers and preventing unauthorized access to sensitive information;
(2) allows legitimate receivers with knowledge of the distortion pattern to recover the real CFR, i.e., *de-obfuscate* the CSI, and perform sensing tasks;
(3) is designed for seamless integration into future Wi-Fi standards.

The proposed method is validated through experimental analysis in a device-free localization scenario (see Sect. 5.3). The goal is to protect user privacy by concealing their position from an eavesdropper, while still allowing authorized devices to localize the user accurately.

2 Related Work

In recent years, there has been a growing interest in CSI-based sensing solutions [8]. However, awareness of threats targeting the physical layer remains limited.

One of the first countermeasures has been proposed in [12], where the authors introduce a system that uses a relay node to generate time-varying multipath components, thereby disrupting sensing capabilities. A more advanced solution is [13], where a reflective intelligent surface (RIS) introduces time-varying reflections to hide the presence of a person within a room. While these techniques can theoretically prevent both passive and active attacks, i.e., scenarios where the attacker actively transmits and thus controls data collection, they require additional hardware, such as one reflective node or a RIS, increasing the complexity of the system.

A simpler and more affordable alternative for SISO systems is transmitter-side obfuscation [1, 2], and a natural extension to MIMO systems is BeamDancer [3], which protects the CSI from eavesdropping by using multi-antenna transmission and continuously changing beamforming vectors. Remarkably, it can run on COTS devices without modifying the Wi-Fi protocol. Another relevant system is mimoCrypt [10], which prevents unauthorized hand gesture recognition using multi-antenna transmissions. Though similar, it differs significantly from both BeamDancer and our method. In mimoCrypt, the authors use beamforming to *encrypt* (obfuscate) the LTS only instead of all symbols. This design limits compatibility with legacy systems, as the estimated CSI does not allow for correct equalization. To overcome this, the authors assume that legitimate receivers know the beamforming weights used for the LTS in each transmission and can *decrypt* (de-obfuscate) the CSI, thus decoding the data. In practice, this approach would require significant changes to the standard, which limits its deployability. In contrast, our approach preserves communication compatibility and, as detailed in Sect. 4.3, requires a few additional header fields solely for legitimate sensing. As a result, it maintains low complexity and deployment costs.

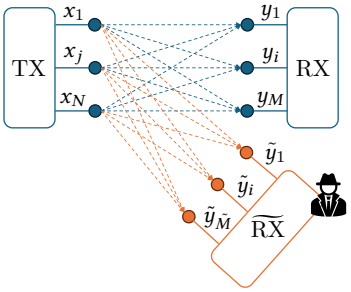

Figure 1: Diagram of the MIMO system showing links to the authorized and unauthorized receivers.

Moreover, mimoCrypt tries to recover the CFR from the encrypted LTSs using a method similar to what is called here *de-obfuscation*. While the authors acknowledge that precoding impacts the quality of the recovered CSI, their approach employs a brute-force optimization to generate precoding weights that compensate for residual distortion. In contrast, an analytical framework for generating precoding weights is presented in this work (Sect. 4.2), along with a dedicated algorithm providing an optimal solution to estimate the real CSI from obfuscated measurements (Sect. 4.4).

3 System Model

We consider the MIMO communication scheme shown in Fig. 1 and focus on the channel between the transmitter and the legitimate receiver. The following reasoning naturally extends to the attacker channel.

Let M and N denote the number of receiving and transmitting antennas, respectively. We define $\boldsymbol{x}(f,t) \in \mathbb{C}^N$ and $\boldsymbol{y}(f,t) \in \mathbb{C}^M$ as the transmitted and received signal vectors at subcarrier f and time t, and $\boldsymbol{H}(f,t) \in \mathbb{C}^{M\times N}$ as the channel matrix. The MIMO system can then be modeled in the frequency domain as

$$\boldsymbol{y}(f,t) = \boldsymbol{H}(f,t)\,\boldsymbol{x}(f,t) + \boldsymbol{e}(f,t), \tag{1}$$

where $\boldsymbol{e}(f,t) \in \mathbb{C}^M$ represents additive noise.

3.1 Beamforming

Beamforming is a technique that improves performance in multi-antenna systems by directing the signal energy toward a specific receiver. This is especially useful in environments where transmitted signals reflect off walls and objects, creating multiple propagation paths. These multipath components may interfere constructively or destructively at the receiver, affecting the overall signal quality. To address this, beamforming leverages knowledge of the MIMO channel to adjust the phase and amplitude of the signals transmitted from each antenna, ensuring they combine constructively at the intended receiver. This improves the received signal-to-noise ratio (SNR) and reduces interference.

Let $s(f,t)$ be the complex symbol to transmit on subcarrier f at time t, which may represent either data or known pilots like the LTSs. Instead of sending it identically from each antenna, the transmitter applies a complex-valued precoding vector $\boldsymbol{w}(f,t) \in \mathbb{C}^N$, forming the transmit signal:

$$\boldsymbol{x}(f,t) = \boldsymbol{w}(f,t)\,s(f,t). \tag{2}$$

From Eq. (1), the received signal is:

$$\boldsymbol{y}(f,t) = \boldsymbol{H}(f,t)\,\boldsymbol{w}(f,t)\,s(f,t), \tag{3}$$

where noise is omitted for clarity.

At the receiver, channel estimation proceeds as defined by the standard using the LTSs. Note that the effective channel estimated for decoding is the beamformed channel, i.e., the product $\boldsymbol{H}(f,t)\,\boldsymbol{w}(f,t)$. This estimate allows a proper combination of the signals collected at the different receiver antennas to reconstruct the transmitted symbol.

Although beamforming can also support more advanced use cases, such as serving multiple users in multi-user MIMO (MU-MIMO) (spatial multiplexing), this work focuses on the simpler single-user, single-stream case, which already captures the essential behavior for the proposed system.

4 Bridging Privacy and Sensing

The CSI can be obtained by any Wi-Fi-enabled device in the surroundings of the transmitter as the LTSs are defined in the standard and, in turn, universally known. Having this information, any device can execute sensing algorithms and gain knowledge about the environment. Notably, as discussed in Sect. 3.1, beamforming can induce changes in the CSI extracted by the receiver, thus possibly preventing sensing. However, conventional beamforming, which aims to improve the received SNR, does not bring adequate obfuscation. Indeed, the optimal precoding for beamforming—obtained from the singular value decomposition (SVD) of $\boldsymbol{H}$ [4]—still yields $\boldsymbol{H}(f,t)\,\boldsymbol{w}(f,t)$ as a fingerprint of the environment, allowing an attacker to perform sensing. To address this limitation, we propose a beamforming-based technique that differs from the standard approach by introducing *temporal randomness* into the precoding vectors. This ensures that $\boldsymbol{H}(f,t)\,\boldsymbol{w}(f,t)$ no longer serves as a stable environmental fingerprint.

Next, we present our obfuscation strategy together with the de-obfuscation to be applied at the legitimate receiver to perform sensing while maintaining adequate communication performance.

4.1 Obfuscation

The obfuscation objective is to make an attacker, unaware of the precoding weights used during transmission, unable to distinguish whether variations in the CSI are caused by the environment or by the precoding itself. To do this, our

intuition is that introducing *temporal randomness* in the precoding vectors makes it harder for the attacker to extract meaningful information from the CSI without knowledge of the precoding weights, even if it collects many measurements and has considerable computational power. However, changes in the channel caused by user movement occur relatively slowly. If the masking distortions vary much faster, attackers can simply average the CSI over time to remove these rapid variations, filtering them out and recovering the underlying channel information. Moreover, previous studies on transmitter-side obfuscation in SISO systems show that these techniques can negatively impact communication performance [2]. Strong virtual distortions act similarly to poor physical channels, often reducing the packet delivery ratio (PDR). In the MIMO case, this issue becomes even more critical: not only strong distortions may degrade performance, but an unlucky choice of the precoding vector can also lead to destructive interference between the transmitted signals, further lowering the SNR. Finally, to be effective, obfuscation should be applied consistently across all transmissions, from beacons to data frames, or an attacker can obtain sensing information on the not-beamformed parts of the frame.

Naturally, using beamforming for obfuscation is not always feasible. For instance, a single-antenna device cannot apply beamforming, making this approach impractical in such cases. However, we believe that beamforming-based obfuscation is a promising strategy that can be integrated into future Wi-Fi standards as part of a broader set of technologies aimed at enhancing physical-layer privacy.

Attack feasibility. Let us now consider the system from the attacker's perspective. Assume that the attacker is equipped with $\tilde{M}$ receiving antennas and observes T transmissions of a known LTS symbol $s(f,t)$, each time crafted with a different precoding. Furthermore, let us also assume that the surrounding environment is static over these T transmissions, and $\tilde{\boldsymbol{H}}(f,t) = \tilde{\boldsymbol{H}}(f), \forall\, t = 1,\dots,T$. The received signals, precoding vectors, and noise terms can be packed into matrices $\tilde{\boldsymbol{Y}}(f) = \left[\tilde{\boldsymbol{y}}(f,1),\dots,\tilde{\boldsymbol{y}}(f,T)\right], \boldsymbol{W}(f) = \left[\boldsymbol{w}(f,1),\dots,\boldsymbol{w}(f,T)\right]$, and $\tilde{\boldsymbol{E}}(f) = \left[\tilde{\boldsymbol{e}}(f,1),\dots,\tilde{\boldsymbol{e}}(f,T)\right]$. This yields the attacker's received signal to be formulated in the frequency domain as

$$\tilde{\boldsymbol{Y}}(f) = \tilde{\boldsymbol{H}}(f)\,\boldsymbol{W}(f)\,s(f) + \tilde{\boldsymbol{E}}(f), \tag{4}$$

with $s(f,t) = s(f)$ the known LTS.

To recover the channel, the attacker must solve a *nonlinear* system where both $\tilde{\boldsymbol{H}}$ and $\boldsymbol{W}$ are unknowns, for each subcarrier f. The condition on the number of antennas $\tilde{M}$ and the number of observed transmissions T to find a solution is

$$\tilde{M}\cdot T \geq \tilde{M}\cdot N + N\cdot T. \tag{5}$$

Rearranging, this yields the condition

$$\tilde{M} \geq \frac{T\cdot N}{T-N}, \quad \text{with } T > N. \tag{6}$$

This shows that the attacker must (1) have more receiving antennas than the number of transmit antennas (i.e., $\tilde{M} > N$), and (2) observe a sufficient number of transmissions to retrieve a solution ($T = N$ for $\tilde{M} \to \infty$).

Even when this condition is satisfied, the attacker faces a *nonlinear inverse problem* with a high number of unknowns. Solving it would require complex optimization techniques, potentially large computational resources, and may still fail due to the ambiguity between channel effects and the hidden precoding. In practice, this makes it highly challenging for an attacker to recover meaningful CSI and perform unauthorized sensing without knowledge of the precoding vectors.

4.2 De-obfuscation

For the legitimate receiver, we can adapt the system from Eq. (4), but considering the legitimate channel $\boldsymbol{H}(f)$ as

$$\boldsymbol{Y}(f) = \boldsymbol{H}(f)\,\boldsymbol{W}(f)\,s(f) + \boldsymbol{E}(f). \tag{7}$$

As $\boldsymbol{W}(f)$ is known to the legitimate receiver, and is a square[1] and invertible matrix, i.e., $T = N$ frames are transmitted using N linearly independent precoding vectors, the channel can be estimated by

$$\hat{\boldsymbol{H}}(f) = \boldsymbol{Y}(f)\,[\boldsymbol{W}(f)\,s(f)]^{-1}. \tag{8}$$

Unfortunately, this closed-form solution is highly sensitive to noise, particularly when $\boldsymbol{W}(f)$ is ill-conditioned. In such cases, the inversion step amplifies the noise component, yielding larger estimation errors:

$$\hat{\boldsymbol{H}}(f) - \boldsymbol{H}(f) = \boldsymbol{E}(f)\,[\boldsymbol{W}(f)\,s(f)]^{-1}. \tag{9}$$

This amplification is measured by the *condition number* of $\boldsymbol{W}(f)$, which is the ratio between the largest and the smallest singular values [5]. Therefore, to ensure robust de-obfuscation and minimize noise amplification, the precoding vectors should be chosen so that $\boldsymbol{W}(f)$ is well-conditioned, ideally by maintaining near-orthogonality (which implies singular values with unitary modulus). This final observation highlights the importance of carefully designing the obfuscation pattern to find a trade-off between three aspects: (1) obscuring the CSI from attackers, (2) guaranteeing reliable communication performance, and (3) maintaining high quality of the reconstructed CSI for authorized sensing.

4.3 Towards a practical protocol

The following approach is based on generating a sequence of precoding matrices $\boldsymbol{W}(f,q)$, where the index q uniquely identifies each matrix in the sequence. Transmissions are

[1] When $\boldsymbol{W}(f)$ is not square, the pseudoinverse would be required instead.

organized in *cycles*, with each cycle consisting of one frame transmitted for each column of $\boldsymbol{W}(f,q)$.

The overall protocol proceeds in the following 4 phases:

Phase 1 **TX–RX association.** Nodes agree, through secured cryptographic exchange, on a shared random seed for pseudo-random number generators. Set $q \leftarrow 1$.

Phase 2 **Precoding generation.** Using the shared seed, nodes generate the precoding matrices $\boldsymbol{W}(f,q)$.

Phase 3 **Transmission.** For the current index q, multiple cycles occur until Phase 4 is triggered.

Phase 4 **Precoding update.** After a period τ, increment $q \leftarrow q+1$ and return to Phase 2.

Assuming the channel remains static during a cycle (i.e., transmissions are faster than channel variations), the scheduling in Phase 3 allows the receiver to estimate the full channel matrix $\boldsymbol{H}(f)$ in a single cycle.

As detailed, to perform de-obfuscation successfully, the receiver must know the precoding matrix $\boldsymbol{W}(f,q)$ used during each cycle. For this, we propose embedding special fields within the headers of Wi-Fi frames to indicate the current cycle and the precoding vector, i.e., which column of $\boldsymbol{W}(f,q)$ was applied. Embedding this information is crucial because frame drops can otherwise disrupt the de-obfuscation process. In addition, by securing these bits, we further complicate the attacker's efforts to gather side information that could simplify the attack. We stress that $\boldsymbol{W}(f,q)$ is known at the receiver thanks to the agreement on the seed in Phase 1.

Finally, while an attacker might attempt to extract the CSI during the association phase, the protocol is designed to protect physical-layer information at any stage. Specifically, if beamformed signals are transmitted during the association process before a secure link is established, an attacker would be unable to estimate the channel. Crucially, legitimate receivers do not need knowledge of the precoding matrix to decode the transmitted data, allowing the CSI to remain protected from the outset. Accurate channel estimation can be performed later, once the association is complete.

4.4 Precoding Matrix Design

Based on the obfuscation and de-obfuscation requirements described above, we design the precoding matrix as

$$\boldsymbol{W}(f,q) = \boldsymbol{A}(f,q)\,\boldsymbol{E}(f,q)\,\boldsymbol{G}(f,q), \tag{10}$$

where $\boldsymbol{A}(f,q) = \mathrm{diag}\,(a_1(f,q),\ldots,a_N(f,q))$ is a real matrix encoding the amplitude scaling per antenna stream; $\boldsymbol{E}(f,q) = \mathrm{diag}\left(e^{j\varphi_1(f,q)},\ldots,e^{j\varphi_N(f,q)}\right)$ is a unitary matrix applying per-stream phase rotations; and $\boldsymbol{G}(f,q)$ is a real orthogonal matrix constructed as a sequence of Givens rotations in an N-dimensional space. A Givens rotation is a transformation that performs a rotation in the 2D plane spanned by two coordinate axes, leaving all other coordinates unchanged [5]. To construct $\boldsymbol{G}(f,q)$, we first apply a rotation in the $(1,2)$ plane by an angle $\theta_1(f,q)$, then in the $(2,3)$ plane by $\theta_2(f,q)$, and so on till $(N-1,N)$, finally rotating in the $(N,1)$ plane to complete the cycle. The matrix $\boldsymbol{G}(f,q)$ represents the overall linear transformation resulting from all of these successive Givens rotations. This structured sequence produces smooth, yet diverse, spatial transformations across antennas, enhancing the obfuscation effect while preserving column orthogonality.

For each variable (a, φ, and θ), we define a two-dimensional random surface over the time-frequency domain identifying the range of possible values. This is done by initializing a grid of i.i.d. Gaussian random values, followed by smoothing through a two-dimensional Gaussian filter. To avoid boundary artifacts, the random grid is generated with extended support along the frequency axis (i.e., with more subcarriers than needed), and the central region is cropped to match the required dimensions. Let $\bar{z}(f,q)$ represent the raw random grid and $g(f,q)$ denote the Gaussian kernel. The smoothed surface writes as

$$z(f,q) = (\bar{z} * g)(f,q), \tag{11}$$

where $*$ denotes the 2D convolution. Here, $z(f,q)$ represents a generic variable surface that can correspond to any of $a_k(f,q)$, $\varphi_k(f,q)$, or $\theta_k(f,q)$. At each obfuscation cycle, for each variable, we extract a frequency slice at index q and normalize it to a target interval $[z_{\min}, z_{\max}]$, specific to the type of the variable. The resulting values for $a_k(f,q)$, $\varphi_k(f,q)$, and $\theta_k(f,q)$ are then used to generate the precoding matrix $\boldsymbol{W}(f,q)$ following Eq. (10). This process ensures that the parameters $z(f,q)$ vary smoothly across frequency and time, thereby generating orthogonal precoding matrices, temporally and spectrally smooth, able to inject plausible, physically consistent distortions into the CSI.

In this initial investigation, as detailed in Sect. 5, we do not consider variations in the amplitude terms a_k and instead fix $\boldsymbol{A}(f,q) = \boldsymbol{I}$ for all (f,q). This choice is done to evaluate the de-obfuscation process under favorable numerical conditions, i.e., making $\boldsymbol{W}(f,q)$ orthogonal for all (f,q) to contain amplification of measurement errors. Moreover, allowing $\boldsymbol{A}(f,q)$ to vary would change the norm of the precoding vectors, leading to deviations from unity and causing fluctuations in transmission power. While this variability could strengthen obfuscation against adversarial inference, increasing the power may violate regulatory limits and increase energy consumption, whereas reducing it could lower the SNR, thus lowering estimation accuracy and PDR. Future work will systematically explore these trade-offs.

5 Experimental Evaluation

To evaluate the proposed approach, we consider a CSI-based device-free localization task, where the goal is to infer a person's position among eight predefined spots using CSI

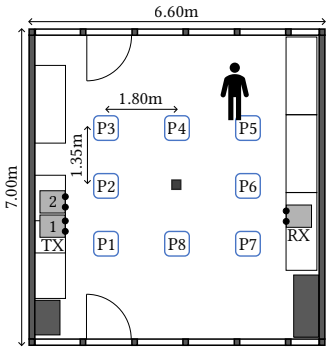

Figure 2: Layout of the testbed. The RX device acts as both the legitimate receiver and the attacker's one.

data.[2] Experiments are conducted using a Wi-Fi 5 (IEEE 802.11ac) network deployed in the advanced networking systems (ANS) laboratory at the University of Brescia (Fig. 2).[3] To isolate the effect of the obfuscation from the devices' placement, we use the same receiver for both attacker and legitimate roles.

A convolutional neural network (CNN) is trained to perform position classification, and localization accuracy is used as the performance metric. To preserve user privacy, we adopt the transmission scheme described in Sect. 4.3. We assume TX and RX successfully exchange seeds during association to generate the precoding matrices for sensing obfuscation. Each transmitted frame carries information about the current precoding matrix and the specific precoding vector used, enabling proper de-obfuscation. Localization accuracy is compared under three scenarios:

Clear: Optimal baseline with original CFR.
Obfuscated: At the attacker, no precoding knowledge.
De-obfuscated: At the authorized receiver, which reconstructs the original channel for sensing.

In each experiment, a person stands at each of the eight locations in Fig. 2 (P1 to P8), while alternating transmissions of one null data packet (NDP), to estimate the full clear channel, and one cycle of obfuscated cycle of data frames.[4] At each position, 450 of such alternations are recorded. For the precoding matrix, we use the parameters summarized in Tab. 1. It is updated approximately every $\tau = 100$ ms, during which we collect on average 10 clear and 10 obfuscated CSI snapshots per spot. Data are collected over two loops (P1 to P8): the first for training, the second for testing.

[2]The dataset and source code used for the experiments are publicly available at https://doi.org/10.5281/zenodo.17098131 and https://github.com/ansresearch/let-it-beam.
[3]https://ans.unibs.it
[4]Note that this NDP is not part of the protocol described in Sect. 4.3; it is used only during the experiments to obtain a clear channel estimate. In practice, no NDP should be transmitted while obfuscation is active, as it would allow the attacker to estimate the CSI.

Table 1: Parameters for generating precoding matrices.

Component	Parameter	Value
$a_k(f,q)$	$[a_{\min}, a_{\max}]$	$[1, 1]$
$\varphi_k(f,q)$	$[\varphi_{\min}, \varphi_{\max}]$	$[0, 2\pi]$
$\theta_k(f,q)$	$[\theta_{\min}, \theta_{\max}]$	$[0, 2\pi]$
Gaussian filter	Kernel shape	2D symmetric
	Std-dev σ	21
	Size $K = 2\lceil 3\sigma \rceil + 1$	67×67

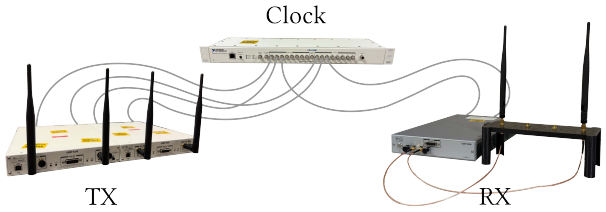

Figure 3: Experimental setup with synchronized SDRs for MIMO transmissions.

Finally, we assess the system's communication performance in an empty room by comparing the PDR for different modulation and coding schemes (MCSs) in two scenarios: (1) standard beamforming with a single spatial stream (SS), using the optimal precoding vector obtained via SVD of the channel matrix to maximize receiver SNR, and (2) the proposed obfuscation scheme. For each MCS, we consider 2000 updates of the precoding matrix.

5.1 Experimental System Setup

The experimental setup, shown in Fig. 3, consists of a 4×2 MIMO network (4 TX and 2 RX antennas).

Using Matlab on a workstation, we generate IEEE 802.11ac 80 MHz frames with beamforming applied. Per-antenna IQ samples are then sent to the transmitter, which consists of two synchronized Ettus USRP X310 software-defined radios (SDRs) (each with 2 antennas). These transmit the frames over the air to the receiver, implemented using an Ettus USRP N300, and IQ samples are then forwarded to a second workstation, where Matlab is used to extract the CSI and decode the data. All SDRs are synchronized using an Ettus Octoclock. Synchronization ensures coherent transmission between the two transmitting SDRs and helps avoid phase errors in the CSI estimation at the receiver due to the carrier frequency offset (CFO), which could affect de-obfuscation. While perfect synchronization between TX and RX is unrealistic in practical deployments, we deliberately enforce tight synchronization in this preliminary study to eliminate such non-idealities. This allows us to isolate and validate the core

functionality of the proposed de-obfuscation mechanism under ideal conditions. The analysis of how these impairments affect de-obfuscation and the design of mitigation strategies, such as those described in [14], is left for future work.

5.2 System Evaluation Methodology

For the localization task, a small size 1D CNN is adopted. Specifically, it accounts for two convolutional layers with 32 kernels of size 5 and a stride of 2 steps. After each convolutional layer, activated through the rectified linear unit (ReLU) function, max-pooling of size 2 is added. The resulting feature map is classified using a fully-connected layer with a softmax activation function to differentiate among the 8 possible positions.

Training and testing of the CNN are conducted under different conditions to evaluate the impact of obfuscation:

Clear (*Clear*): The network is fed with the CSI along the temporal dimension, and each input channel maps to a specific TX-RX antenna pair at the same time instant.

Obfuscated (*Obf.*): The attacker extracts the per-subcarrier CSI from each received frame as a 2×1 matrix. Over one cycle of 4 frames, this yields $2 \times 1 \times 4 = 8$ elements. Assuming the wireless channel remains static throughout the cycle, each element can be treated as an independent input channel, ensuring a fair comparison with other cases.

De-obfuscated training and testing (*De-obf. Tr&Te*): The network is trained and tested on the reconstructed 8-channel CSI, as in the clear case, assessing whether meaningful localization information is preserved.

De-obfuscated testing only (*De-obf. TeO*): The network trained on clear CSI data is used to predict positions from de-obfuscated CSI samples. This tests the similarity between clear and de-obfuscated CSI.

In the first three cases, for each of the ten experiments, we train and test eight models, each initialized with different random weights. This repeated training strategy reduces the impact of initialization bias and provides a more robust assessment of performance. In the last case, predictions use the pre-trained clear models.

5.3 Experimental Results

Fig. 4 shows the localization results as box plots showing the median, the 25th and 75th percentiles (box), outliers as individual points computed using the inter-quartile range, and whiskers extending to the most extreme non-outlier data points. Statistics are computed over the results of all 10 experiments, each performed with 8 models. In the *clear* case, accuracy approaches 100 % with minimal deviation. Under *obfuscation*, accuracy drops drastically, with median values around 22 %, close to random guessing, confirming the effectiveness of our obfuscation strategy. Some models still extract meaningful information from obfuscated data, achieving up to 46 % accuracy, indicating that obfuscation generally hinders training, but its effect is highly variable. Both *de-obfuscated scenarios* recover accuracy close to the *clear* case, demonstrating that the restored CSI preserves meaningful localization information (*De-obf. Tr&Te*) and closely resembles the clear one (*De-obf. TeO*). For completeness, the confusion matrices of the localization results for the four cases are reported in Fig. 5.

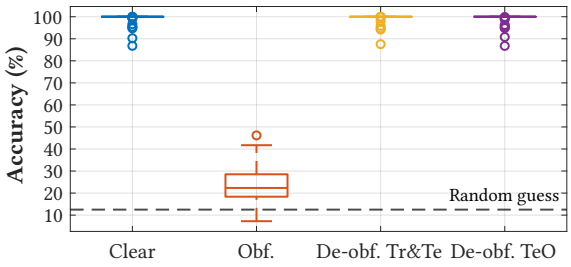

Figure 4: Box plots comparing localization accuracy in clear, obfuscated, and de-obfuscated cases.

Finally, Fig. 6 compares the PDR under optimal beamforming and obfuscation across MCSs 0–9. Optimal beamforming maintains 100 % PDR consistently, while obfuscation achieves 100 % only up to MCS 4 and then steadily declines to near zero at the highest MCS, illustrating a *trade-off between privacy and communication* reliability at higher data rates.

6 Conclusion

Building on the foundation laid by [3], this work advances privacy protection in Wi-Fi sensing by introducing a selective obfuscation mechanism that leverages time-varying beamforming precoding weights to conceal the CFR from unauthorized sensing. Unlike previous approaches, our framework allows trusted devices to retain full sensing capabilities while preserving backward compatibility for communication with legacy receivers. We validate the effectiveness of our system through real experiments, using device-free localization as a representative sensing task. The results show that the proposed obfuscation scheme successfully conceals user location from unauthorized receivers, while legitimate devices experience no degradation in localization performance, achieving accuracy comparable to the clear case.

As part of future work, we will explore how real-world impairments, such as the CFO, impact the robustness of the de-obfuscation process. Additionally, as with most obfuscation strategies, our technique introduces a trade-off: although communication remains reliable at low to moderate data rates, higher modulation schemes experience a drop in PDR. This calls for a deeper exploration of the design space to balance privacy protection, communication efficiency, and sensing accuracy.

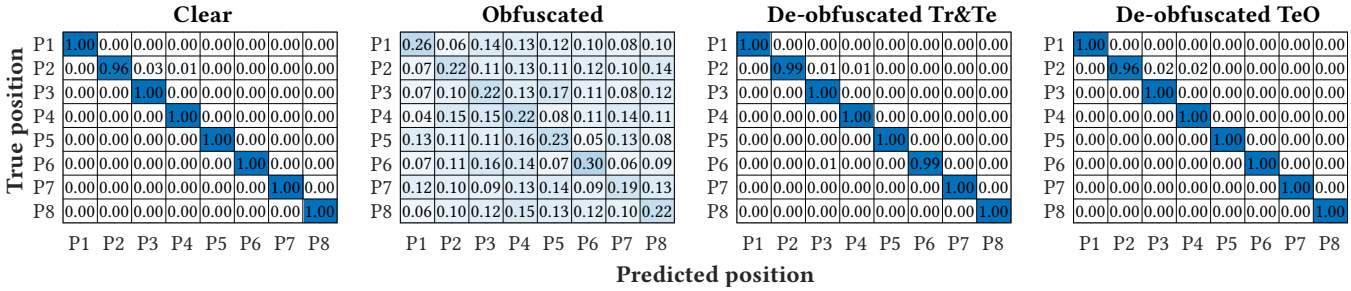

Figure 5: Confusion matrices showing localization performance in clear, obfuscated, and de-obfuscated cases. Rows are normalized to highlight how accurately each true position is classified.

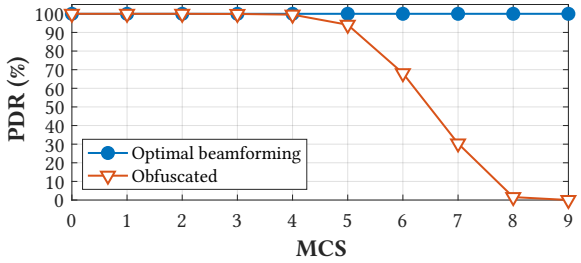

Figure 6: Packet delivery ratio comparison between optimal beamforming and obfuscation.

Acknowledgments

This work was partially supported by the European commission at University of Brescia under the Italian National Recovery and Resilience Plan (NRRP) of NextGenerationEU through projects ISP5G+ (CUP D33C22001300002) and SCARPHASE (CUP C89J24000580008) part of the SERICS program (PE00000014) and EMBRACE (CUP E63C22002070006) part of the RESTART program (PE00000001); and by project 6G-INTENSE (Intent-driven NaTive AI architecturE supporting Compute-Network abstraction and Sensing at the Deep Edge), Horizon Europe SNS JU Grant N. 101139266; at University of Padova through the project ROBUST-6G, Horizon Europe SNS JU, Grant N. 101139068, and CAMELIA (CUP C93C24004880002) under the Italian NRRP of NextGenerationEU.

References

[1] Giovanni Angelo Alghisi, Francesco Gringoli, Marco Cominelli, Shabbir Raza, and Renato Lo Cigno. 2025. For your eyes only: bridging privacy and sensing in Wi-Fi networks through CSI obfuscation. In *2025 23rd Mediterranean Communication and Computer Networking Conference (MedComNet 2025)*. Cagliari, Italy, (June 2025).

[2] Marco Cominelli, Felix Kosterhon, Francesco Gringoli, Renato Lo Cigno, and Arash Asadi. 2021. IEEE 802.11 CSI randomization to preserve location privacy: an empirical evaluation in different scenarios. *Computer Networks*, 191, 107970. doi:10.1016/j.comnet.2021.107970.

[3] Marco Cominelli, Shaghayegh Shahcheraghi, Jakob Link, Matthias Hollick, Federico Cerutti, Francesco Gringoli, and Arash Asadi. 2024. Physical-layer privacy via randomized beamforming against adversarial wi-fi sensing: analysis, implementation, and evaluation. *IEEE Transactions on Wireless Communications*, 23, 12, 19603–19617. doi:10.1109/TWC.2024.3485477.

[4] Andrea Goldsmith. 2005. *Wireless Communications*. Cambridge University Press, USA. ISBN: 0521837162.

[5] Gene H. Golub and Charles F. Van Loan. 2013. *Matrix Computations*. (4th ed.). Johns Hopkins University Press, Philadelphia, PA. doi:10.1137/1.9781421407944.

[6] Francesco Gringoli, Matthias Schulz, Jakob Link, and Matthias Hollick. 2019. Free your CSI: a channel state information extraction platform for modern Wi-Fi chipsets. In *Proceedings of the 13th International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization* (WiNTECH '19). Association for Computing Machinery, Los Cabos, Mexico, 21–28. ISBN: 9781450369312. doi:10.1145/3349623.3355477.

[7] Huawei Technologies. 2024. Integrated sensing and communication (ISAC)—from concept to practice. (2024). Retrieved Aug. 12, 2025 from https://www.huawei.com/en/huaweitech/future-technologies/integrated-sensing-communication-concept-practice.

[8] Jian Liu, Hongbo Liu, Yingying Chen, Yan Wang, and Chen Wang. 2019. Wireless sensing for human activity: A survey. *IEEE Communications Surveys & Tutorials*, 22, 3, 1629–1645.

[9] Renato Lo Cigno, Francesco Gringoli, Marco Cominelli, and Lorenzo Ghiro. 2022. Integrating CSI Sensing in Wireless Networks: Challenges to Privacy and Countermeasures. *IEEE Network*, 36, 4.

[10] Jun Luo, Hangcheng Cao, Hongbo Jiang, Yanbing Yang, and Zhe Chen. 2024. MIMOCrypt: multi-user privacy-preserving Wi-Fi sensing via MIMO encryption. In *2024 IEEE Symposium on Security and Privacy (SP)*, 2812–2830. doi:10.1109/SP54263.2024.00025.

[11] Yongsen Ma, Gang Zhou, and Shuangquan Wang. 2019. WiFi sensing with channel state information: a survey. *ACM Computing Surveys*, 52, 3.

[12] Yue Qiao, Ouyang Zhang, Wenjie Zhou, Kannan Srinivasan, and Anish Arora. 2016. PhyCloak: obfuscating sensing from communication signals. In *Proceedings of the 13th Usenix Conference on Networked Systems Design and Implementation* (NSDI'16). USENIX Association, Santa Clara, CA, 685–699.

[13] Paul Staat, Simon Mulzer, Stefan Roth, Veelasha Moonsamy, Markus Heinrichs, Rainer Kronberger, Aydin Sezgin, and Christof Paar. 2022. IRShield: a countermeasure against adversarial physical-layer wireless sensing. In *IEEE Symposium on Security and Privacy (SP)*. doi:10.1109/SP46214.2022.9833676.

[14] Hongzi Zhu, Yiwei Zhuo, Qinghao Liu, and Shan Chang. 2018. π-splicer: Perceiving accurate CSI phases with commodity WiFi devices. *IEEE Transactions on Mobile Computing*, 17, 9.