Loki: Physical-World Adversarial Attacks on Wireless Indoor Localization via Differentiable Object Placement

Xueqiang Han, *Student Member, IEEE*, Jinyang Huang, *Member, IEEE*, Meng Li, *Senior Member, IEEE*, Chao Cai, *Member, IEEE*, Tianyue Zheng✉, *Member, IEEE*

***Abstract*—As a cornerstone for numerous sensing applications, wireless indoor localization has been a pivotal area of research over the last two decades. While techniques such as jamming, spoofing, and adversarial perturbation have been exploited to compromise wireless indoor localization, existing attacks face challenges in accessibility to wireless systems and stealthiness. To address these limitations, we introduce *Loki*, a novel physical-world attack on wireless indoor localization via differentiable object placement. Specifically, we develop a differentiable wireless ray-tracing technique that allows us to optimize object placement in the scene. By repositioning an existing object in the scene by just a few centimeters, *Loki* fools existing wireless indoor localization systems into generating erroneous localization results. We also show via experiments that the object placement generated by *Loki* aligns with wireless sensing theory (e.g., the forward scattering region and Fresnel zone), confirming its explainability. Additionally, *Loki* proves effective across various localization models and scenarios, highlighting its generalizability.**

***Index Terms*—Wireless indoor localization, differentiable ray-tracing, adversarial attack, physical attack.**

I. Introduction

The increasing demand for context-based services has driven the development of localization technologies. Wireless indoor localization, serving as a complementary solution to outdoor GNSS localization, has emerged as a promising field attracting significant attention from both academia [1], [2] and industry [3]–[5]. The wide-scale proliferation of smartphones and other wireless devices has provided the necessary infrastructure for practical and efficient indoor localization [1], making it a pivotal technology for a wide range of applications. These include enhancing healthcare [6], strengthening security and surveillance [7], enabling smart home [8], and streamlining asset tracking [9]. With its vast potential, wireless indoor localization has transitioned from research to real-life deployments. Commercial successes like Apple iBeacon [3], Google Maps for Indoors [4], and Cisco DNA Spaces [5] underscore its growing integration into our daily lives.

However, the widespread adoption of wireless indoor localization has raised significant security concerns, particularly in safety-critical applications like healthcare [6] and surveillance [7]. In these sectors, efficient and accurate positioning is of utmost importance. However, the inherent broadcast nature of wireless signals exacerbates these risks, creating opportunities for adversaries to eavesdrop on and manipulate wireless transmissions [10]. Such adversarial actions can disrupt or even disable indoor localization systems. As a result, investigating the security vulnerabilities and understanding the potential attack vectors against wireless indoor localization systems are of critical importance.

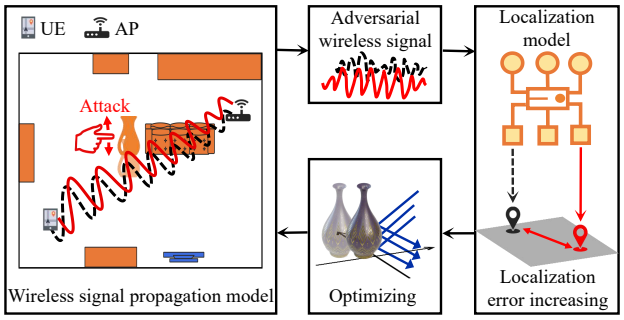

Fig. 1. The attack scenario of *Loki*.

Conventional attacks on wireless indoor localization consist of jamming [11] and spoofing [12]. Jamming primarily affects location estimation by emitting strong interference, while spoofing aims to intentionally mislead the localization systems using artificially fabricated signals. However, orchestrating these attacks is challenging and may disrupt existing communication and sensing services, making them highly conspicuous [13]. More recent research has shifted towards adversarial attacks on wireless indoor localization [14]–[17], where the focus is on crafting adversarial samples in the *signal space* using gradient-based methods. These methods iteratively alter the values of the captured wireless data to maximize location estimation errors. However, they assume that wireless signals can be directly modified in *signal space*, which is unrealistic because it necessitates the laborious hacking of low-level firmware or hardware, thereby limiting accessibility.

To address the limitations, researchers have explored performing adversarial attacks directly in the *physical world* rather than in the *signal space* [18]. A notable example is Phy-Adv [19], which seeks to mislead WiFi-based localization by employing adversarial shields around access points (APs) to manipulate the received signal strength. However, shielding APs with artificial enclosures is not a common practice in daily settings, rendering such methods conspicuous and readily detectable. Furthermore, Phy-Adv's dependence on signal strength manipulation confines its effectiveness to systems utilizing received signal strength indication (RSSI) for indoor localization. This narrows its applicability, as the majority

X. Han and T. Zheng are with the Department of Computer Science and Engineering, Southern University of Science and Technology, China.
E-mail: hanxq@mail.sustech.edu.cn, zhengty@sustech.edu.cn
J. Huang and M. Li are with the School of Computer and Information, Hefei University of Technology, China.
E-mail: {hjy, mengli}@hfut.edu.cn
C. Cai is with the College of Life Science and Technology, Huazhong University of Science and Technology, China.
E-mail: chriscai@hust.edu.cn
✉ Corresponding author: Tianyue Zheng.

of current solutions have transitioned to using channel state information (CSI) for more accurate indoor localization [20].

Facing the aforementioned issues, we ask a critical question: is it possible to develop an attack on wireless indoor localization that achieves the goals of *effectiveness*, *accessibility*, and *inconspicuity* while also being *generalizable* to a wide range of environments? Correspondingly, we envision an attack on wireless indoor localization operating in the physical world, where we can perform the attack by strategically moving existing objects without introducing new ones. However, designing such attacks faces three significant challenges. Firstly, it is impractical to manipulate object placements exhaustively to obtain the optimal attack configuration in the real world. To solve this issue, we must establish a wireless signal propagation model that enables efficient optimization of object placement, thereby guiding strategic adjustments without exhaustive physical trials. Secondly, bridging the gap between a model and the real world is a complex task, as discrepancies between the two are inevitable. To ensure the effectiveness of our attack, we must develop an efficient approach that is resilient to such discrepancies, allowing for successful execution in real-world scenarios. Last but not least, our ultimate goal is not merely to attack wireless indoor localization but to reveal vulnerabilities and raise awareness. To achieve this, we must conduct theoretical analysis and gain valuable physical insights and rules (beyond how to move objects) that can contribute to the development of more secure and robust wireless indoor localization systems.

To address these challenges, we introduce *Loki*, a novel adversarial attack on wireless indoor localization through strategic, differentiable object placement. As illustrated in Fig. 1, a wireless indoor localization model attempts to pinpoint the location of user equipment (UE), such as a smartphone, using signals captured from a WiFi AP. *Loki* manipulates localization by subtly repositioning an existing object (such as a common vase) within a few centimeters, thereby altering wireless signals and leading the system to produce erroneous location results. *Loki* begins by constructing a model of the environment, including its 3-D geometry and physical attributes (such as dielectric properties and reflectivity). We then apply fully differentiable ray tracing (RT) [21], [22] to create a model that associates wireless signals with object placement. Utilizing this model, we execute adversarial attacks through a gradient-based optimization method. To enhance the model's practicality, we introduce a novel robustness augmentation technique designed to compensate for discrepancies between the model and real-world conditions, ensuring that *Loki* is robust and effective in real-world applications. Additionally, we conduct a theoretical analysis of *Loki*'s outcomes, uncovering physical insights consistent with wireless sensing theory, such as the Fresnel zone and forward scattering region. Our extensive experiments, both simulated and in real-world settings, further validate *Loki*'s performance across various configurations. The primary contributions of our work are summarized as follows:

- We present an innovative adversarial attack in the physical world that targets wireless indoor localization systems through *differentiable object placement*. To our knowledge, *Loki* is the first attack of its kind that is effective, accessible, and inconspicuous.
- To support this attack, we develop a wireless RT engine capable of *fully differentiable object placement*. It facilitates the use of gradient-based methods to effectively maximize localization errors.
- We develop a novel robustness augmentation method that integrates hardware, material, and placement variations, achieving resiliency to real-world variations.
- We analyze the object placements generated by *Loki* and find that they align with sensitive regions (e.g., the Fresnel zone) in accordance with wireless sensing theory, further demonstrating the explainability of *Loki*.
- Comprehensive evaluations of various localization models across various real-world scenarios highlight the *Loki*'s effectiveness and adaptability.

The rest of this paper is organized as follows. Section II introduces the background and motivation of *Loki*. Section III explains the threat model. Section IV details the attack design of *Loki*. Section V provides *Loki*'s implementation. Section VI reports the evaluation results, followed by security analysis and possible defense methods in Section VII. Section VIII and Section IX presents related works and discussion, respectively. Finally, we conclude the paper in Section X.

II. Background and Motivation

In this section, we introduce the background of wireless indoor localization and the motivation of *Loki*'s design.

A. Wireless Indoor Localization Preliminaries

Wireless indoor localization can be broadly categorized into two categories: device-based [23]–[25] and device-free [26], [27], depending on whether the target carries a device that interacts with the wireless indoor localization system. This study focuses exclusively on device-based localization, the most widely adopted and commercialized approach. Potential attacks on device-free localization will also be discussed in Section VIII. In wireless localization systems, signals propagate from the transmitter (Tx) to the receiver (Rx), interacting with the environment. At the Rx, we measure how the received signal differs from the transmitted signal, thus obtaining the CSI, characterizing the environment, and performing localization of the Tx or Rx. Specifically, the Orthogonal Frequency Division Multiplexing (OFDM) signal exhibits fluctuations in the frequency spectrum, while Multiple Input Multiple Output (MIMO) antenna arrays result in variations in the spatial spectrum.

By applying the Fast Fourier Transform (FFT) to the frequency and spatial dimensions of the CSI, we can derive two key localization metrics, i.e., Time-of-Flight (ToF) and Angle-of-Arrival (AoA) [28]. ToF specifies the propagation distance of the signal, while AoA defines the angle of the received signal. Together, these metrics enable point localization on a 2-D plane using a model-based approach [29], [30]. However, model-based methods are effective only under ideal conditions and face challenges such as multipath interference and background clutter: they are often limited by narrow bandwidth and insufficient antennas and struggle to distinguish among these paths, often becoming impractical [31]. To address these

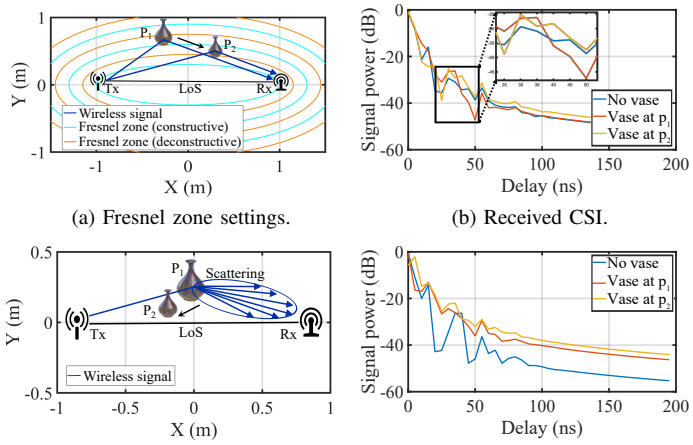

(a) Fresnel zone settings. (b) Received CSI.

(c) Forward scattering region settings. (d) Received CSI.

Fig. 2. The impact of object placement in Fresnel zone and forward scattering region on CSI.

issues, deep learning (DL)-based wireless indoor localization has been proposed [24], [25]. Leveraging the exceptional feature extraction and data processing capabilities of DL, these systems can effectively learn from massive amounts of CSI samples and ground truth location pairs. This approach incorporates environment-relevant prior information into the localization system, reducing bias caused by multipath effects and achieving more precise localization.

B. Vulnerabilities of Indoor Localization

Recent studies have claimed improved precision in wireless indoor localization systems [24]. However, these systems may still be susceptible to certain environment factors. In this section, we analyze the potential vulnerabilities of such localization systems, focusing on the impact of environment changes on CSI. During wireless signal propagation, it interacts with surroundings in complex ways, including reflection, scattering, and diffraction. Any changes in the environment directly affect CSI measurements, potentially introducing variations in the input to DL-based localization systems. Wireless sensing theory identifies two critical "weak spots" that can significantly affect CSI: the Fresnel zone and the forward scattering region. The Fresnel zone consists of elliptical areas around the Line of Sight (LoS) path between Tx and Rx, where obstructions can cause notable signal changes. The forward scattering region is an area where electromagnetic waves are scattered towards the Rx after encountering an obstacle, potentially altering signal propagation even when direct LoS is partially blocked. We illustrate the impact of object placement in these two regions in Fig. 2.

To verify the impact of these "weak spots" on CSI, we conduct experiments on one Tx and Rx. We set the length of the LoS path to 2 m and placed a common household vase with a diameter of 0.5 m at two different points (p_1 and p_2) within 2 m and 0.5 m of the Rx, corresponding to the Fresnel zone and forward scattering region, respectively. Fig. 2 shows the collected CSI samples plotted as power delay profiles. In Fig. 2b, one may observe that placing objects at adjacent ellipses of Fresnel zones can cause significant CSI change, due to constructive/destructive interference caused by phase reversals at different ellipses of the Fresnel zone. Fig. 2d illustrates that when objects are positioned near the forward scattering region, the CSI strength is greatly influenced by the blocking/unblocking of the LoS. These results demonstrate that the signal space, which serves as input to DL-based localization systems, can be significantly altered by strategically placing an everyday object such as a vase, and potentially be used for attacking DL-based localization systems.

III. Threat Model

We present the threat model of *Loki*, including the adversary goal, requirements, capabilities, and constraints in this section.

A. Adversary Goal

In this paper, we target device-based wireless indoor localization systems. A device-based wireless indoor localization model $\mathcal{F}_\theta(H) \rightarrow \hat{\Lambda}$ with parameters θ, which takes the CSI H as an input and outputs the target device's location $\hat{\Lambda} = (x, y)$. The primary objective of wireless localization models is to minimize the expected localization error across diverse positions within an indoor environment as follows:

$$\min_{\theta} \mathbb{E}_{(H,\Lambda)} \mathcal{L}(\mathcal{F}_\theta(H), \Lambda),$$

where $\mathcal{L}(\cdot)$ is the loss function (e.g., root mean square error (RMSE)) that measures the difference between the output of $\mathcal{F}_\theta(\cdot)$ and the ground truth Λ.

The adversary's goal is to fool the wireless indoor localization models to generate erroneous locations of the target device by moving objects in the scene. The object movement will impact the wireless channel and compromise the measurements of CSI. The Rx then feeds the compromised CSI into $\mathcal{F}_\theta$, negatively affecting its prediction. Formally, the adversary goal of *Loki* can be expressed as follows. Assuming the sensing scene is composed of objects numbered from 1 to N, and $P_i = (x_i, y_i)$ is the 2-D position of the i-th object, and $P = \{P_1, P_2, \ldots, P_N\}$ is a set representing all objects' positions in the target scene. Let $P_\text{s} = (x_\text{s},\ y_\text{s})$ be the position of the selected object. $\mathcal{R}$ is the process of CSI generation in the scene and $H = \mathcal{R}(P)$. *Loki* aims to craft an adversarial position P_s^adv for the selected object, resulting in perturbed CSI, maximizing the difference between the faulty prediction and the ground truth of the target device's location. Formally, the adversary's goal can be formulated as:

$$\arg\max_{P'_\text{s}} \mathcal{L}(\mathcal{F}_\theta(\mathcal{R}(P'), \Lambda)),$$

where $P' = P \setminus P_\text{s} \cup P'_\text{s}$ constitutes all the objects in the scene after moving the selected object.

B. Attack Requirements

As mentioned in Section I, conventional attacks in the signal space face challenges such as inaccessibility, conspicuousness, and implementation difficulties. In contrast, we envision *Loki* as a novel physical-world attack with the following properties, making it a genuine threat to wireless indoor localization systems:

- *Effectiveness.* The attack should successfully manipulate the target device's position estimation, significantly impacting the quality of location-based services.

- *Accessibility.* The attack surface should be readily accessible to the adversary, requiring no complex manipulation of wireless devices.
- *Inconspicuity.* The manipulated object should be pre-existing in the scene, with its subtle repositioning limited to a few centimeters, making the movement imperceptible to observers.
- *Ease of operation.* The object's position should be easily alterable by any ordinary individual without external assistance.

To illustrate our approach in line with these requirements, we only require repositioning objects that are already present in the scene and intend to utilize ordinary medium-sized objects (e.g., a household vase). The displacement applied to these objects should remain within a reasonable range, typically in the range of 10cm. Moving already-present objects within a small range is usually feasible and unnoticeable, especially in a public area. It is worth noting that while our design allows for object rotation in conjunction with position shifting, we deliberately choose not to explore this aspect in the current paper, as it could significantly compromise the attack's stealthiness.

C. Adversary Capabilities

In the following discussion, we explain the adversary capabilities in the *Loki* attack. We assume that the adversary can access the target indoor environment, obtain its spatial layout, collect CSI data using their own transceivers or existing infrastructure, and manipulate movable indoor objects. The spatial layout can be acquired through various ways, such as floor plans, photographs, or LiDAR scanning, which facilitate the construction of a digital replica of the indoor scene. The collected CSI data are then used to calibrate material properties, as detailed in Section IV-A2. These assumptions are realistic because indoor spaces requiring location-based services are generally public areas, freely accessible to individuals. Furthermore, it is assumed that the adversary has access to the locations and orientations of wireless transceivers, information that is also readily accessible.

Regarding the wireless localization model, we focus on two primary scenarios: the white-box and black-box settings. In the white-box setting, the attacker possesses full knowledge of a victim's localization model, including its architecture and weights. This enables the attacker to perfectly replicate the user's model, thus facilitating a more effective attack. This assumption aligns with real-world scenarios, as supported in [17]. In the white-box scenario, attackers can employ gradient-based algorithms to execute precise attacks on the model. In the black-box setting, however, the adversary can only interact with the model by providing inputs and observing outputs, without any details of the model's internal mechanisms. This scenario poses a more challenging yet realistic task. Black-box attacks leverage model transferability; adversarial attacks are generated on a pre-trained surrogate model using its specific architecture and CSI dataset. These attacks are then transferred to compromise the target's previously unseen wireless indoor localization model.

D. Adversary Constraints

The adversary of the *Loki* attack faces the following constraints: first, it does not possess the ability to embed adversarial signals within the system, as doing so would require hacking into wireless devices, which is both impractical and highly conspicuous. Second, the adversary is restricted from introducing new objects into the target environment; they are limited to manipulating existing objects only. Any introduction of unfamiliar objects in the scene would be noticeable and could raise suspicion. Instead, the adversary can only work with the movement of already-present objects. Furthermore, any manipulation of these objects must adhere to physical laws. For example, optimized placement of objects cannot result in overlap with other physical entities, nor should objects float in the air. Finally, the complexity of the optimization process for object movement is constrained. A brute-force traversal of the entire sensing environment for each possible position is infeasible due to the high computational cost, both in real-world measurements and in simulations.

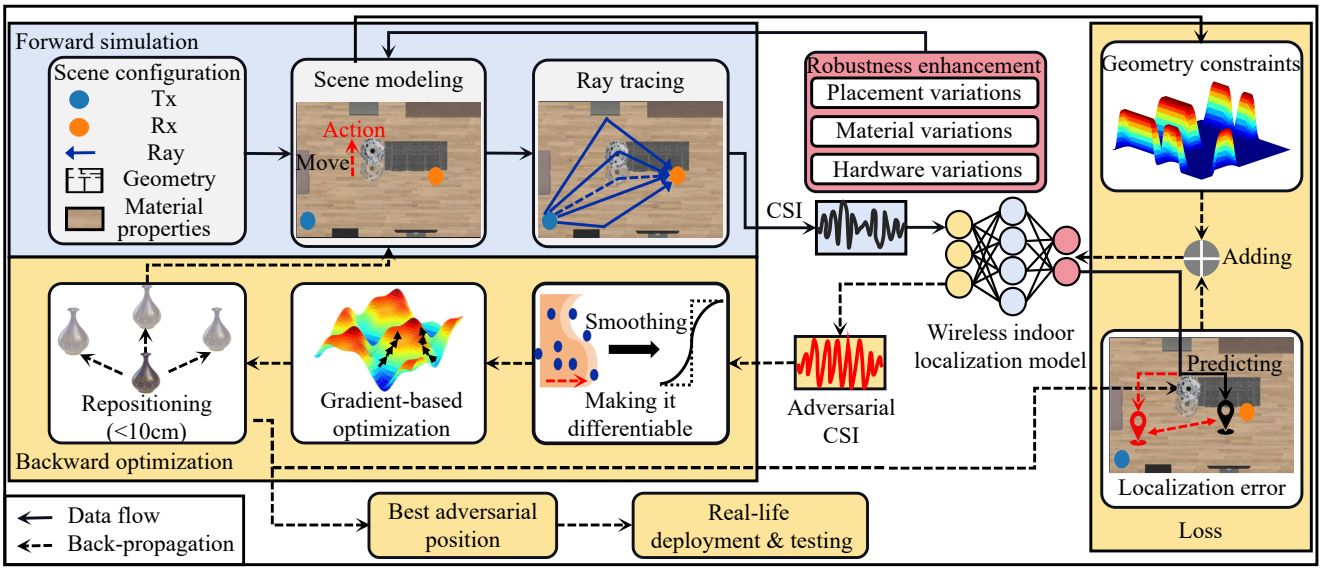

Fig. 3. The workflow of *Loki*'s attack strategy.

IV. Attack Design

In this section, we introduce the attack design of *Loki*. As shown in Fig. 3, the workflow consists of four steps: scene configuration and modeling, radio propagation modeling, backward optimization, and robustness enhancement. The key symbols and their meanings are summarized in Table I.

TABLE I
Summary of Key Symbols

Symbol	Meaning	Symbol	Meaning
$\mathcal{R}(\cdot)$	RT model	$\mathcal{E}$	Digital twin
$P^{\mathrm{R,T}}$	Rx/Tx position	$C^{\mathrm{R,T}}$	Rx/Tx antenna pattern
O	Scene objects	$t_{\max}$	Max ray bounces
H	CSI	A^{r}	Reflection spread factor
$T(\cdot)$	Transfer function	R_{s}	Scattering matrix
E^{in}	Incident field	Q	Diffraction matrix
E^{out}	Outgoing field	D	Basis transformation matrix
Λ	True location	$P_{\mathrm{s}}^{\mathrm{adv}}$	Adversarial position
$\mathcal{P}$	Feasible region	ϵ	Max movable distance

A. Scene Configuration and Modeling

1) Scene Configuration: We consider three properties of the real-world environment, including the transceivers, scene geometry, and the characteristics of the materials involved, which are key factors influencing CSI. The geometric properties can be obtained either by manual measurement or by using LiDAR to replicate a real-world environment in the digital world. Besides, the electromagnetic (EM) parameters of the corresponding material are sourced from the international telecommunication union recommendation [32]. As for Tx and Rx, we manually record their positions and publicly obtain relevant antenna settings from online technical specifications.

2) Scene Modeling: After getting scene configurations, we generate a digital scene model suitable for RT. We formulate the modeling process of a digital twin as $\mathcal{E} = \mathcal{G}(P^{\mathrm{R}}, P^{\mathrm{T}}, C^{\mathrm{R}}, C^{\mathrm{T}}, O, t_{\max})$. Here, $t_{\max}$ is the maximum number of allowed interactions between the environment and ray, discussed in Section IV-B1. P^{R} and C^{R} are the Rx's positions and antenna patterns, while P^{T} and C^{T} are the Tx's positions and antenna patterns. O represents all the N objects or obstacles in the scene. $O_i = \{x_i, y_i, z_i, \zeta\}$ is the i-th objects located at (x_i, y_i, z_i) with materials ζ. High-precision modeling hinges on the precise measurement of these key features, particularly material characteristics. We adopt the calibration method to enable physical properties to converge to real-world values, eliminating "mismatches" with minimal effort [33], [34]. However, even with advanced calibration techniques, some biases inevitably remain. Therefore, to ensure effective real-world scalability, we address strategies for enhancing the robustness of *Loki* in Section IV-D.

B. Radio Propagation Modeling

1) Shooting-and-Bouncing Ray Method: To acquire the estimation of wireless signals, the RT approach is widely adopted to accurately predict the propagation of high-frequency EM waves [35]. In this paper, our simulation is conducted using the shooting-and-bouncing ray (SBR) forward RT algorithm [36]. The basic principle of SBR is to track individual rays emitted from the source, tracing both direct and reflected paths by applying Snell's law and geometrical optics. According to the RT theory, CSI H can be seen as an integral over the rays from the Tx, which can be mathematically represented as: $H = \int_{\Omega} h(\omega, O, P^{\mathrm{R}}, P^{\mathrm{T}}, C^{\mathrm{R}}, C^{\mathrm{T}}) d\omega$, where Ω is the unit sphere, ω denotes a ray path shooting from Ω. $h(\cdot)$ is the function quantifying how the single ray propagates in the indoor scene. However, there is no closed-form solution to the equation. Therefore, Monte Carlo sampling [37] is used to approximate the integral by shooting M rays from the Tx, and bouncing rays on intersected objects until the maximum depth $t_{\max}$ is reached or the ray is captured by the Rx. The resulting Monte Carlo estimation is:

$$H = \frac{1}{M} \sum_{i=1}^{M} h(\omega_i, O, P^{\mathrm{R}}, P^{\mathrm{T}}, C^{\mathrm{R}}, C^{\mathrm{T}}). \tag{1}$$

2) CSI Generation: The CSI H between a Tx and Rx at frequency f can be formulated as the sum of M propagation paths (rays) [38]:

$$H = \frac{c_0}{4\pi f} \sum_{i=1}^{M} C^{\mathrm{R}}(\phi_i^{\mathrm{R}}) T_i (C^{\mathrm{T}}(\phi_i^{\mathrm{T}})) e^{-j2\pi f \tau_i}, \tag{2}$$

where c_0 is the vacuum speed of light, ϕ_i^{R} and ϕ_i^{T} are variables related with the intersection between the transceiver and i-th path with time delay of arrival τ_i. T_i is the transfer function of i-th path. A propagation path i undergoes t_{max} intersections with real objects in the scene, experiencing different propagation phenomena, including reflection, diffraction, and diffuse scattering. Each intersection is represented as a hit point $P_j = (x_j, y_j, z_j)$, for $j = 0, \ldots, t_{max} + 1$, and the P_0 and $P_{t_{max}+1}$ are P^T and P^R, respectively. For each hit point, we need to establish a relationship between the incident field at current hitting point P_j, denoted as $E_j^{\mathrm{in}}(P_j)$, and the created field $E_{j+1}^{\mathrm{in}}(P_{j+1})$ at the next hit point P_{j+1}, involving the relation of $E_j^{\mathrm{in}}(P_j)$ to $E_j^{\mathrm{out}}(P_j)$, the outgoing field at P_j, and the relation of $E_j^{\mathrm{out}}(P_{j+1})$ to $E_{j+1}^{\mathrm{in}}(P_{j+1})$. The incoming and outgoing fields at the hit point P_j are related as follows:

$$E_j^{\mathrm{out}}(P_{j+1}) = F_j(E_j^{\mathrm{in}}(P_j)), \tag{3}$$

where $F_j(\cdot)$ describes different types of wireless signal interaction (e.g., reflection, scattering, and diffraction).

For an incoming EM wave with direction $\hat{k}_i$ incident on a surface, the field transformation for reflection is:

$$F_j(E_j^{\mathrm{in}}(P_j)) = R(\hat{k}_i, \zeta) E_j^{\mathrm{in}}(P_j) A^{\mathrm{r}}(P_{j+1}, P_j) e^{-j2\frac{\pi}{\lambda} d_{j+1}}, \tag{4}$$

where $A^{\mathrm{r}}(P_{j+1}, P_j)$ is the spreading factor related with the shape of the wavefront [39] and $d_j = \|P_j - P_{j-1}\|$ represents the distance between hit points P_j and P_{j-1}. $R(\cdot)$ is a differentiable function dependent on direction of incidence $\hat{k}_i$ and material properties ζ. In the scattering process, the $F_j(\cdot)$ is given as:

$$F_j(E_j^{\mathrm{in}}(P_j)) = \left\| R_s E_j^{\mathrm{in}}(P_j) \right\| \frac{Z(\hat{k}_i, \hat{k}_s, dA)}{d_{j+1}} e^{-j2\frac{\pi}{\lambda} d_{j+1}}, \tag{5}$$

where $\hat{k}_s$ is the direction of the scattered ray, dA is the size of a small area element wrapped around the hit point P_j,

and $Z(\hat{k}_i, \hat{k}_s, dA)$ is a function of dA, the incidence direction $\hat{k}_i$, the scattering direction $\hat{k}_s$, and the scattering pattern [40]. R_s is the Fresnel coefficients matrix [32]. According to [41], diffraction can be calculated as:

$$F_j(E_j^{\text{in}}(P_j)) = QE_j^{\text{in}}(P_j)A^{\text{d}}(P_{j+1}, P_j)e^{-j2\frac{\pi}{\lambda}d_{j+1}}, \quad (6)$$

where Q is the diffraction matrix and $A^{\text{d}}(P_{j+1}, P_j)$ is the spreading factor for diffraction. In particular, when $j = 0$ in $F_j(\cdot)$, we have:

$$E_0^{\text{in}}(P_0) = C^{\text{T}}(\phi_i^{\text{T}}), \quad F_0(E_0^{\text{in}}(P_0)) = \frac{e^{-j2\frac{\pi}{\lambda}d_1}}{d_1}E_0^{\text{in}}(P_0). \quad (7)$$

The incoming field at the hit point P_{j+1} can be calculated as:

$$E_{j+1}^{\text{in}}(P_{j+1}) = D_j E_j^{\text{out}}(P_{j+1}), \quad (8)$$

where D_j is the basis transformation matrix detailed in [21].

From Eqn. (3)–(8), we can derive the relationship between the transmitted and received field:

$$\begin{aligned} E_{t_{max}+1}^{\text{in}}(P_{t_{max}+1}) &= D_{t_{max}+1}F_{t_{max}}(\dots D_1 F_0(E_0^{\text{in}}(P_0)) \\ &= T(C_T(\phi^{\text{T}}))e^{-j2\pi f\tau}, \end{aligned} \quad (9)$$

where the $\tau = (c_0)^{-1}\sum_j d_j$ represents the total propagation delay. $T(\cdot)$ is the transfer function in Eqn. (2) for the given propagation path. $C^{\text{T}}(\cdot)$, as well as C^{R}, can be obtained as in [42]. We can successfully model the generation of CSI based on the aforementioned procedures.

C. Backward Optimization

In this section, we detail the gradient-based optimization for object placement, in which it is crucial to obtain the CSI's gradient w.r.t. the object position. Furthermore, the geometry constraint must also be considered for moving the object to prevent its placement from violating the laws of physics. While existing simulators like Sionna RT represent advancements in differentiable ray tracing for radio propagation modeling [33], they lack the full differentiability and support for object position optimization required for our wireless indoor localization application. This limitation necessitates the development of our own fully differentiable wireless ray tracing simulator.

1) Making it Differentiable: To achieve back-propagation and acquire the partial derivative of the scene parameters O, especially the object positions, we need to cast our attention back to Eqn. (1). We can find that the differentiability of CSI to scene parameters O depends on whether $h(\cdot)$ is differentiable. As indicated by Eqn. (2) and Eqn. (9), it is clear that CSI is differentiable w.r.t. the hit points P, so the partial derivatives, $\partial_P H$, can be obtained leveraging the automatic differentiation (AD) computation capabilities of machine learning frameworks, like TensorFlow. In most cases, the above strategy gives correct gradients if the visibility of objects in the scene does not change, i.e., the M sampled paths would not vary. However, evaluating the impact of strategically placed objects requires shifting their positions, inevitably altering visibility and necessitating path resampling. This introduces a significant challenge: geometric edges and occlusions cause discontinuous changes in the RT paths used for wireless signal propagation modeling. These discontinuities disrupt the optimization process by hindering gradient-based methods. While similar challenges have been addressed in computer graphics, particularly in scene reconstruction [43], [44], these techniques are not readily adaptable to the intricacies of RT in wireless propagation. A promising solution involves smoothing these discontinuities. By replacing the discontinuous changes in RT with smoothly parameterized functions [45], we can obtain meaningful gradients that guide the identification of optimal attack positions. For instance, the binary validity check of a path, typically represented by an indicator function $\mathbb{1}(\cdot)$, exhibits abrupt changes with object movement. Approximating this with a smooth function, such as $\frac{1}{1+e^{\alpha x}}$ (where α controls the smoothing), eliminates the discontinuity and enables effective gradient-based optimization.

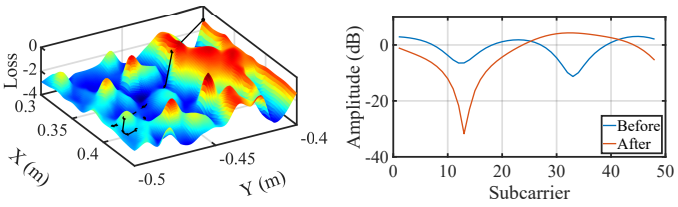

(a) Process of gradient descent. (b) Altered CSI subcarriers.

Fig. 4. Illustration of CSI modification through strategic manipulation of vase position with differentiability.

To validate the effectiveness of differentiable placement of the object, we move the vase to control the CSI, discussed in Section II-B, through the gradients provided by the module. Previously, the random placement of the vase caused significant disturbances to CSI, but these disturbances are random. Here, we attempt to leverage the differentiable placement of the object with the gradient-based method to more intelligently control the position of the vase, thereby generating interference with CSI that aligns with our expectations. Specifically, we take the amplitude of CSI as the optimization objective, with the aim of shifting the vase using gradient descent to modify CSI. The optimization process is illustrated in Fig. 4. With differentiability, one can observe that it is indeed possible to generate adversarial object positions that result in significant CSI variations.

2) Geometry Constraints: To move the selected object properly, we have to adhere to two rules: i) the placement has to respect physical laws and fall into the feasible region $\mathcal{P}$, and ii) the new position must be close to the original location to avoid being conspicuous. We will now elaborate on these constraints in detail.

a) Physical feasibility constraint: The feasible region should be sufficient to hold the object and should not result in a physical overlap of objects. The valid region can be obtained either by a human annotator, given that we have established a digital twin of the target scene. The adversary's objective can be reformulated with these inherent constraints: $P_{\text{s}}^{\text{adv}} = \arg\max_{P_{\text{s}}' \in \mathcal{P}} \mathcal{L}(\mathcal{F}_\theta(\mathcal{R}(P \setminus P_{\text{s}} \cup P_{\text{s}}')), \Lambda)$. Generally, rather than solving the constrained optimization, the Lagrangian-relaxed form of the problem is preferred to use, as in the work [46]:

$$P_{\text{s}}^{\text{adv}} = \arg\max \mathcal{L}(\mathcal{F}_\theta(\mathcal{R}(P \setminus P_{\text{s}} \cup P_{\text{s}}')), \Lambda) - \lambda\mathcal{P}.$$

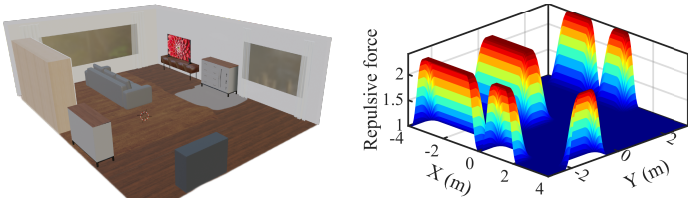

(a) Room geometry. (b) Potential function.

Fig. 5. APF of geometry constraints.

However, the geometry constraints of scenes are discrete. Furthermore, there is no gradient information available for the coordinates of the selected objects, given these constraints, which prevents us from optimizing. To tackle this problem, we smooth the discrete constraints into a continuous function and establish a differentiable relationship between geometry constraints and the location of the selected object. Inspired by [47], we use the artificial potential field (APF) approach to introduce gradients between the location and the cost function. The basic idea of APF is that the destination of the moving object exerts an attractive force U^{a} on it while the obstacles apply a repulsive force U^{r}, and the resultant force navigates the object to the destination. The U^{a} and U^{r} are all represented by differentiable functions. We focus on the repulsive force U^{r} imposed by obstacles, as in our context, the mobile object does not have a definite destination. Thus, we can achieve automatic obstacle avoidance while optimizing the position of objects with the following objective:

$$P_{\mathrm{s}}^{\mathrm{adv}} = \arg\max \mathcal{L}(\mathcal{F}_\theta(\mathcal{R}(P \setminus P_{\mathrm{s}} \cup P_{\mathrm{s}}')), \Lambda) - \lambda U^{\mathrm{r}},$$

where U^{r} can be expressed as follows:

$$U^{\mathrm{r}}(P) = \begin{cases} \frac{1}{2} g \left(\frac{1}{\eta(P)} - \frac{1}{\eta_0} \right) & \eta(p) \leq \eta_0 \\ 0 & \text{otherwise}, \end{cases}$$

where $\eta(P) = \min_{P \in P_i^{\mathrm{e}}} \|P - P_s\|$ is the smallest distance between the mobile object and the i-th obstacle edge P_i^{e}, and η_0 and g are appropriately chosen constants. When the object is not near the obstacle edge, the repulsive force is 0, which does not have any impact on *Loki*'s optimization decision. The repulsive force only pushes the object away when it moves extremely close to the obstacles, as illustrated in Fig. 5. If necessary, U^{r} can be replaced with any other proper APF functions that do not negatively impact *Loki*'s performance.

b) Inconspicuity constraint: The second geometric constraint ensures that the repositioned spot remains close to its original location, making it visually inconspicuous to human observers. To quantify this, we introduce ϵ as the maximum allowable distance between the repositioned position P_s' and P_s, as will be defined in Section V-A. With this constraint, the optimization problem can be formulated as follows:

$$\begin{aligned} P_{\mathrm{s}}^{\mathrm{adv}} = \arg\max_{P_{\mathrm{s}}'} & \Big\{ \mathcal{L}\big(\mathcal{F}_\theta(\mathcal{R}(P \setminus \{P_{\mathrm{s}}\} \cup P_{\mathrm{s}}')), \Lambda\big) - \lambda U^{\mathrm{r}} \Big\} \\ \text{subject to} \quad & \|P_{\mathrm{s}}' - P_{\mathrm{s}}\|_2 \leq \epsilon. \end{aligned} \tag{10}$$

3) Gradient-based Optimization: The subsequent step of the *Loki* attack involves determining the adversarial position of the object to fool the wireless indoor localization model. This begins by feeding generated CSI into the model to make a precise prediction of the location. To enlarge the localization error, we perform a gradient ascent approach. Specifically, we calculate its gradients w.r.t. the object position P_s', and then update P_s' using these gradients. Note that the loss function takes both the localization error and the geometry constraints into consideration. This iterative procedure is conducted until the maximum number of optimization steps, T_p, is reached, at which point the most effective adversarial position is recorded. Finally, we test the adversarial attacks in real-world scenarios.

D. Robustness Enhancement

While optimizing object placement within a perfectly replicated digital scene offers a valuable starting point, real-world deployments introduce complexities that can undermine the effectiveness of adversarial attacks on wireless indoor localization models. Specifically, CSI is highly susceptible to variations in placement, materials, and hardware [48]. Therefore, to enhance the robustness of P^{adv}, we propose a universal adversarial attack designed to withstand these real-world variations. Inspired by [49], we leverage the concept of Expectation over Signal Variation (EoSV). This involves introducing a set of transformations T to the signal and optimizing the expectation of the objective function:

$$\arg\max \mathbb{E}_{T \in \mathcal{T}} \mathcal{L}(\mathcal{F}_\theta(\mathcal{R}(P \setminus P_{\mathrm{s}} \cup P_{\mathrm{s}}')), \Lambda) - \lambda U^{\mathrm{r}},$$

where $\mathcal{T}$ represents the set of transformations T affecting the wireless signal. To implement EoSV, we introduce random discrepancies in the virtual object's position, on the scale of a fraction of the carrier wavelength. This compensates for minor deviations between the ground truth geometry and its digital representation. A similar approach addresses variations in material characteristics. Due to the computational cost of exploring all possible discrepancies, we approximate the expectation in Eqn.(11) by averaging over five randomly sampled variations. These variations encompass object positions, materials, and signal variations introduced by hardware. Furthermore, mirroring the technique in [17], we incorporate real-world CSI noise collected in the target environment at different times of day into the adversarial attack optimization process. This strengthens the attack's resilience against random signal fluctuations encountered in practical deployments.

V. Implementation, Setup, and Metrics

In this section, we provide details of *Loki*'s implementation, experiment setup, and metrics used in our study.

A. System Implementation

We construct a high-fidelity digital twin of a real-world indoor environment using Blender 3.6 and import it into Mitsuba 3 [50] for wireless RT. All deep learning components, including camera calibration, wireless indoor localization models, and our novel differentiable object placement module, are implemented using TensorFlow 2.13 and Python 3.8. The largest repositioning distance ϵ is set to 10 cm to achieve

inconsipuity to human eyes. For training the differentiable object placement module, we utilize the Root Mean Square Propagation (RMSprop) optimizer with a learning rate of 0.01 and a maximum of 20 iterations. This configuration, found to be effective through empirical testing, is maintained for all subsequent experiments. The geometry constraint weighting factor λ is set to 0.1 based on hyperparameter tuning. Parameters g and η_0 within U^{r} are set to 4 and 1, respectively. Finally, for EoSV, we introduce random discrepancies in geometry within two wavelengths and variations in material characteristics within 5%.

TABLE II
EXPERIMENTAL SETUP SUMMARY

Scene	Meeting Room $3.53\times2.85\ \mathrm{m}^2$	Corridor $31.14\times2.65\ \mathrm{m}^2$	Classroom $9.46\times7.8\ \mathrm{m}^2$	Laboratory $10.79\times5.59\ \mathrm{m}^2$
Model	CiFi	AAR	SIABR	DLoc
Tx count	1	3	4	4
Antenna count	3	12	16	16

B. Experiment Setup

In the experiment setup, to illustrate the effectiveness of *Loki* across various indoor environments, we select 4 scenes from the real world for our analysis. The real-world scenes include a meeting room (3.53 m × 2.85 m), a corridor (31.14 m × 2.65 m), a classroom (9.46 m × 7.8 m), and a laboratory (10.79 m × 5.59 m). Fig. 6 shows the layout of the four scenarios. In addition to this, we employ the ASUS RT-AX86U wireless APs in our setting and obtain CSI with the AX-CSI tool [51]. We employ the models CiFi [52], AAR [53], SIABR [25], and DLoc [24] for our wireless indoor localization security research due to their practicality and representative nature in the field. CiFi formulates wireless indoor localization as a classification problem, utilizing DCNN as the localization model. The system first builds a CSI database of various locations, then generates localization results by weighted averaging of the DCNN's highest-probability classification outcomes. In contrast, AAR and DLoc tackle indoor localization as a regression task, enhancing the CNN model complexity through ResNet. Lastly, SIABR is the first to leverage broadband wireless signals for indoor localization. These models demand diverse experiment setups regarding the number of Tx devices and antennas: DLoc requires 4 Tx devices with a total of 16 antennas, SIABR needs 3 Tx devices, while CiFi and AAR require only 1 Tx with 3 antennas. We implement these models and collect CSI data by adhering closely to the methodologies outlined in the respective papers. The experimental parameters used in our settings are summarized in Table II.

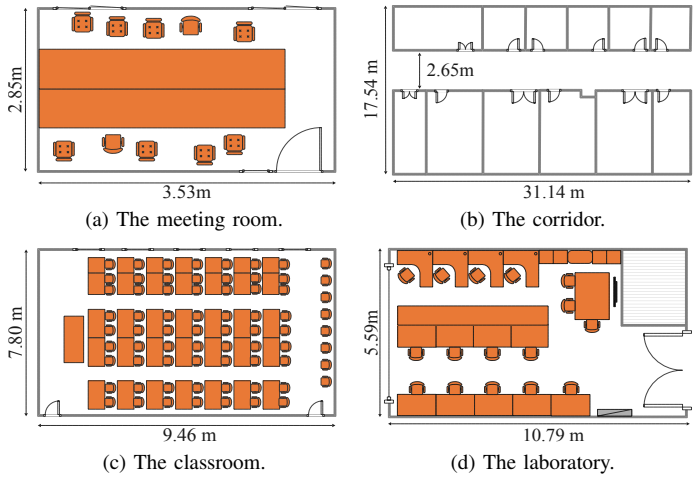

(a) The meeting room. (b) The corridor.

(c) The classroom. (d) The laboratory.

Fig. 6. The real-world experiment scene layouts.

C. Metrics

To evaluate the performance of our proposed *Loki*, we use various metrics, drawing extensively from prior studies [16], [19], [24], [25], [52], [53]. Specifically, the selected metrics include:

Localization error. For a specific device in the scene, the localization error is defined as the discrepancy between the original estimated location $\hat{\Lambda}$ and the estimated location $\hat{\Lambda}^{\mathrm{adv}}$ after *Loki*'s attack, i.e., $\|\hat{\Lambda} - \hat{\Lambda}^{\mathrm{adv}}\|_2$.

Attack Success Rate (ASR). The ASR is calculated as $\frac{N_{\mathrm{success}}}{N_{\mathrm{total}}}$, where N_{success} indicates the number of successful attacking, and N_{total} indicates the total number of the attacking. We identify a successful attack as the alteration in estimated location caused by *Loki* is greater than a threshold δ. Specifically, once an attack causes a localization error greater than δ, it is recognized as a successful attack and could be counted. The metric is first applied to wireless localization in [19].

VI. ATTACK EVALUATION

In this section, we conduct extensive experiments to show the superiority and robustness of *Loki*. After that, we study how *Loki* extends to multiple objects, targeted attacks, and indoor tracking, and investigate the effectiveness of *Loki* under black-box settings. Finally, we study the impact of the EoSV module and the maximum repositioning distance.

A. Overall Performance

We first evaluate *Loki*'s performance by analyzing the localization error and ASR across various models and real-world scenes. In our study, we randomly select 100 reference points within these scenes, target them using *Loki*, and present the results in Fig. 7. As illustrated in Fig. 7a, the median localization errors for the meeting room, the corridor, the classroom, and the laboratory consistently exceed 1.95 m, 2.15 m, 3.13 m, and 3.62 m across all models, respectively. These errors are substantial when considered in relation to the dimensions of the respective environments. We attribute these substantial errors to several factors. Firstly, CSI, which reflects communication channels, is highly susceptible to environmental changes caused by object movement. Additionally, while machine learning models generally perform well across various tasks, they are prone to vulnerabilities from non-random disturbances [54]. Last but not least, our proposed *Loki* attack can effectively identify the adversarial object positions that can attack the localization model via optimization. One may see that SIABR and DLoc almost achieve the lowest localization errors across the four scenarios. This can be attributed to more Tx devices and antennas, which constitute

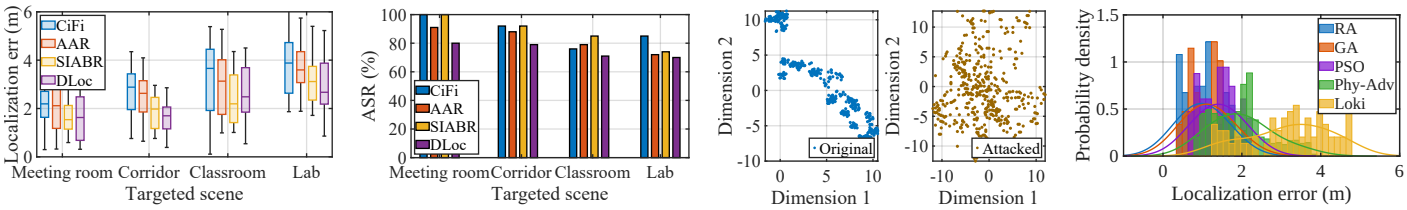

(a) Localization errors across scenes. (b) ASR across scenes. (c) t-SNE of sample distribution. (d) Baseline comparisons.

Fig. 7. Overall performance of *Loki*.

a potential defense method. We postpone the discussion on defense strategies to Section VII.

We also measure the ASR of all these attacks in each scene. We set the threshold for a successful attack in these environments to 1.5 m, 2 m, 2.5 m, and 2.5 m according to the relative sizes of each environment. The results are reported in Fig. 7b. We can observe that the *Loki*'s attack achieves an average success rate of 84.5% across all four environments and throughout the four models. Notably, the ASR in all the scenes exceeds 60%, underscoring the resilience of our attack across diverse environments and various models. To dive deeper into *Loki*'s capabilities, we introduce the t-SNE method to visualize the feature embeddings of the original features and the features after the attack. In detail, we use CiFi as the target model and sample various reference points to attack with *Loki*. We project the output of the last layer of the CNN of CiFi into two dimensions using t-SNE to visualize the distribution of these samples in both their original state and under attack. The results are depicted in Fig. 7c. It can be observed that the t-SNE of the samples after *Loki*'s attack appears more scattered compared with the original samples due to increased scene diversity caused by object movement and confusion of the wireless indoor localization model.

To demonstrate *Loki*'s effectiveness, we compare it against three baseline object placement methods for implementing the attack: i) random attack (RA), which randomly places the object within a 20 cm range; 2) genetic algorithm (GA) [55]; and 3) particle swarm optimization (PSO) [56]. Both GA and PSO are established gradient-free optimization algorithms that maximize their respective fitness functions. We utilize the objective function defined in Section IV-C2 as the fitness function for both methods. We also use Phy-Adv [19] as another baseline method to show the superiority of Loki. As shown in Fig. 7d, *Loki* achieves a localization error of up to 4.86 m, significantly outperforming all these baselines.

B. Robustness Analysis

In this section, we analyze *Loki*'s robustness to different real-world factors. These factors include the size of the object, the distance between the selected object and the device to be localized, as well as the various types of objects. Here, we use a range of objects (e.g., a vase, chair, cabinet, bookshelf, TV set, and sofa) as attack vectors across all four scenarios, unless otherwise stated.

1) Effect of Object Size: We select objects with different volumes (length $\times$ width $\times$ height) to investigate the impact of their size on localization estimation errors. The results are shown in Fig. 8. Overall, the object size is strongly correlated with the localization error. As the volume continues to increase, *Loki*'s attack effectiveness improves significantly. This is reasonable, as the larger the size of the object, the more EM waves it interacts with. Therefore, a larger object can induce more pronounced interference with CSI when it moves, thereby enabling a broader scope for descent during gradient search. Meanwhile, we are surprised to find when the volume of the object is only 0.2 m^3, the localization errors can also be within a range close to 3 m, which demonstrates the effectiveness of *Loki*.

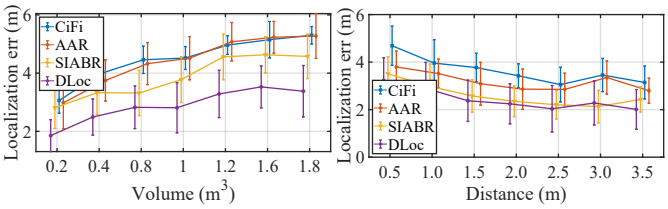

Fig. 8. Varying size. Fig. 9. Varying distance.

2) Effects of Distance: We investigate the impact of the distance from the moved object to the target device. The results are presented in Fig. 9. It becomes evident that when the distance to the target device falls within the range of 2.5 m, the performance of the *Loki* attack is inversely proportional to the distance, resulting in averages for mean localization estimation errors above 3.09 m. Moreover, as the distance continues to increase, the effect of the attack remains almost unchanged. The phenomenon can be explained by the fact that as the distance increases, the variation in wireless signals induced by object movement progressively diminishes, resulting in a reduced impact on the received CSI, thereby leading to getting stuck in local optima during optimization. However, even as the distance increases, the localization error is still above 2.1 m, underscoring the robustness of *Loki* under diverse distances.

3) Effects of Object Types: We further study the effect of different objects, whose different sizes and material properties influence *Loki*'s attack performance. We evaluate *Loki*'s performance with six everyday objects: a vase (VA), a cabinet (CB), a bookshelf (BS), a display cabinet (DC), a TV set (TS), and a sofa (SF), utilizing them as attack vectors to determine *Loki*'s effectiveness across various conditions. The results are detailed in Fig. 10. Our findings reveal that the type of object significantly impacts *Loki* 's performance. Notably, using a table as the attack vector yields optimal results due to its larger size and consequent amplification of wireless signal disruption. Interestingly, the vase, despite its smaller size, exhibits a similar attack performance to the display cabinet.

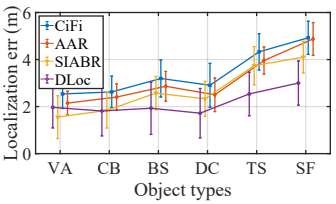

Fig. 10. Varying types.

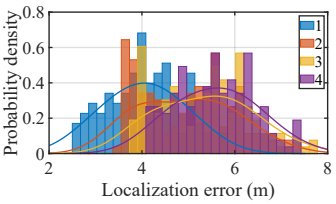

Fig. 11. Multi-objects.

This is attributed to the vase's smooth surface, which enhances reflectivity and magnifies its impact on wireless signals.

C. Extension to Multiple Objects

In this section, we explore the enhancement of *Loki* by simultaneously moving multiple objects. Our experiment involves controlling the movement of 2, 3, and 4 objects at once, with results depicted in Fig. 11. The data reveal a clear trend: as the number of objects increases, the effectiveness of the attack improves. Remarkably, in extreme cases, *Loki* induces localization errors exceeding 6 m. This heightened impact arises because moving multiple objects at once affects the CSI more significantly than a single object's movement. However, the relationship between the number of objects and localization errors is not linearly proportional. The gain from moving two objects is greater than that from moving four, likely due to the increased complexity of the optimization problem, which becomes more challenging and non-convex with more objects involved. Nonetheless, achieving localization errors beyond 5 m clearly demonstrates the effectiveness of *Loki*'s attack strategy.

D. Extension to Targeted Attack

While our initial system design focused on untargeted attacks, we extended *Loki* to explore targeted attacks against localization methods that treat the problem as a classification task (e.g., CiFi). For targeted attacks, we reformulate the objective as: $P_s^{adv} = \arg\min_{P'_s \in \mathcal{P}} \mathcal{L}(\mathcal{F}_\theta(\mathcal{R}(P \backslash P_s \cup P'_s)), \Lambda_{anchor})$ where Λ_{anchor} represents the targeted anchor point's location. As shown in Fig. 13, the localization errors in Fig. 13a are slightly smaller compared to Fig. 7a, likely due to the more restricted parameter space in targeted attacks. Nevertheless, *Loki* still achieves an impressive 75.2% success rate on average in these targeted attacks.

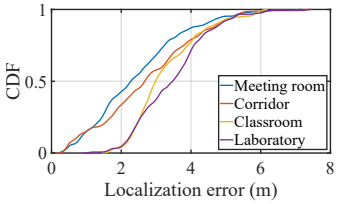

(a) Localization errors.

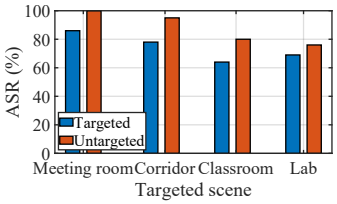

(b) Probability of success.

Fig. 12. Extending *Loki* to targeted attack.

E. Attack on Indoor Tracking

We extend our analysis of *Loki* to evaluate its effectiveness against the indoor tracking system DLoc, the sole baseline supporting indoor tracking, moving beyond static targets. Unlike the single optimization problem described in Eqn. (11), tracking moving targets requires *Loki* to simultaneously attack all potential locations within the target's feasible region and compute the expectation to determine the optimal adversarial positioning. As illustrated in Fig. 13a, initially, repositioning a single object induces a median tracking error of 1.2 m. Notably, the attack's effectiveness scales dramatically with the number of repositioned objects. When increased to 4 objects, *Loki* achieves tracking errors escalating to 4.3 m. To further illustrate the attack's potency, Fig. 13b presents the trajectories before and after the attack. The visualization clearly reveals how *Loki* substantially deviates from the trajectory estimated by DLoc, underscoring its remarkable capability to disrupt indoor tracking systems.

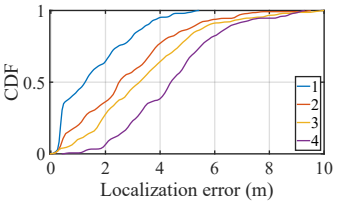

(a) Effects of different number of repositioned objects.

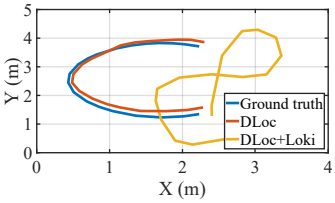

(b) Tracking trajectories.

Fig. 13. Attack on indoor tracking.

F. Black-box Attack

In this section, we evaluate *Loki*'s black-box performance when the attacker is unaware of the specific architecture of the target model. We first train a surrogate model and utilize it to generate an adversarial object position. By physically moving the object to the position in the real environment, we obtain perturbed CSI measurements and then feed it into the target black-box model for location prediction. Two surrogate models are designed based on the representative structures commonly adopted in wireless indoor localization systems [24], [25], including DCNN and LSTM. Specifically, we use a DCNN surrogate by default for all targets. For the SIABR model, we additionally test an LSTM surrogate (denoted as SIABR-L in evaluation) to exploit its structural similarity. These models learn a mapping from CSI to spatial coordinates. Initially, we deploy a CNN model featuring three ResNet blocks to train and then launch attacks on models such as CiFi, AAR, and DLoc, all based on the CNN architecture. Furthermore, we apply the same adversarial strategy to SIABR, which integrates LSTM networks with ResNet, to evaluate the impact of varying model architectures under black-box scenarios.

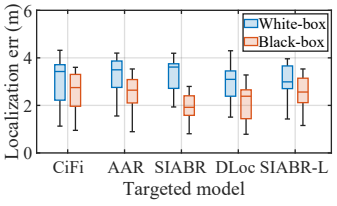

(a) Localization error.

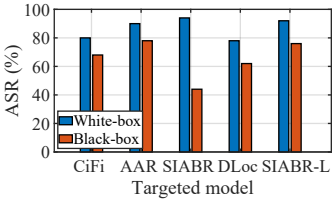

(b) ASR.

Fig. 14. Black-box attack.

The results of these experiments are presented in Fig. 14. One may observe that black-box attacks using CNN with 3 ResNet blocks against CiFi, AAR, and DLoc achieve an average localization error of 2.5 m, 2.51 m and 2.13 m,

respectively; all the ASRs exceed 60%, even under black-box attack. In the case of attacking different model architectures, the average localization error and ASR of SIABR are 1.868 m and 44%. This is because SIABR has a more complex network, including ResNet, LSTM, and an intra-transmitter attention mechanism. Without knowing the specific model architecture, the attack scenario is harder only using the CNN with 3 ResNet blocks. Nonetheless, when we employ the black-box attack using an LSTM combined with ResNet against SIABR (SIABR-L), it achieves a higher localization error of more than 2.5 m. This result shows that while access to the model helps, we can still get good performance with a surrogate.

G. Ablation Study

We further conduct an ablation study in this section to understand the contribution of EoSV. We compare the original *Loki* attack with the following cases: without the expectation of object placement variation (w/o EoPV), without the expectation of material variation (w/o EoMV), and without the expectation induced by hardware variation (w/o EoHV). The comparison results are shown in Table III, illustrating that EoSV significantly impacts the effectiveness of *Loki*'s attack. Specifically, neglecting the influence of materials yields the greatest impact, with the localization errors decreasing by approximately 2.00 m. The cause of this phenomenon is largely due to the material properties significantly affecting wireless propagation, with mechanisms of reflection, scattering, and diffraction. Overall, the median localization error (over 3.89 m) of *Loki* is effectively enhanced with EoSV, showing a high practicality of our attack.

TABLE III
ABLATION STUDY

Component	Loki	w/o EoPV	w/o EoMV	w/o EoHV
Loc. err.	3.89 m	2.21 m	1.89 m	2.25 m

H. Sensitivity Analysis

In this section, we analyze the impact of maximum repositioning distance and environment knowledge on *Loki*'s attack.

1) Analysis of maximum repositioning distance: We conduct experiments to investigate how the maximum repositioning distance affects the attack performance. We vary the maximum repositioning distance to 5 cm, 10 cm, 15 cm, and 20 cm and optimize the object placement within each range. As shown in Fig. 15, a substantial performance gain occurs when the movable distance increases up to 10 cm, increasing localization error by up to 1.47 m. Although larger repositioning distances can further improve the attack performance, the gains are only 0.55 m and 0.72 m with 15 cm and 20 cm maximum repositioning distance, respectively. This is largely due to the periodicity of wireless signal propagation. A displacement of 10 cm is comparable to the signal wavelength (approx. 6-12 cm), which is sufficient to induce maximal phase shifts (e.g., shifting from constructive to destructive interference) within the sensitive Fresnel zones. Extending the movement range further yields diminishing marginal perturbations to the multipath structure.

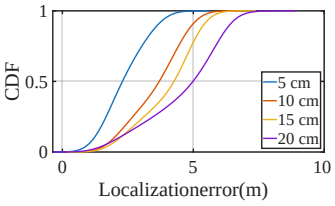

Fig. 15. Analysis of maximum repositioning distance.

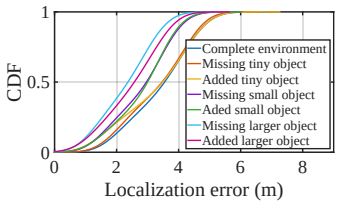

Fig. 16. Analysis of environment knowledge.

2) Analysis of Environment Knowledge: We also conduct experiments to explore the effectiveness of Loki under incomplete environment knowledge. Specifically, we investigate Loki's performance in seven scenarios: 1) with complete environment knowledge; 2) with missing tiny objects (e.g., books and backpacks); 3) with added tiny objects; 4) with missing specific small objects (e.g., desktop computers, chairs); 5) with added small objects; 6) with missing larger objects (e.g., sofa and bookshelf); 7) with added larger objects. The results are reported in Fig. 16. We observe that discrepancies involving tiny objects exhibit negligible impact on attack performance. Variations in small objects lead to a slight degradation, with a reduction of roughly 0.66 m in median localization error, indicating Loki's robustness to such variations. In contrast, discrepancies involving larger objects result in a more noticeable effect, with the median localization error decreasing by approximately 1.16 m. Nevertheless, even in these challenging scenarios, the attack retains sufficient efficacy to compromise indoor localization models, demonstrating that Loki can still perform effectively under incomplete environment knowledge.

VII. Security Analysis and Defense Discussion

In this section, we conduct a comprehensive security analysis of *Loki*, with the objective of uncovering the fundamental reasons behind its effectiveness and resilience. Our findings aim to further inform the development of defensive strategies for safeguarding wireless indoor localization models against such attacks. To filter out irrelevant factors and streamline the experiment process, our tests are carried out in a vacant classroom. In this controlled setting, the Tx and Rx devices are positioned at coordinates (0 m, -2.5 m) and (0 m, 1.5 m), respectively, with a strategically placed vase serving as the attack vector. Note that the single Tx device is equipped with three antennas. The vase is positioned at varying distances of 0.5 m and 1 m from the Rx and is moved to deceive the localization model. We plot the CSI variations as a contour map, and provide a visualization of *Loki*'s optimization trajectories on top of the CSI contour maps in Fig. 17, to examine how the

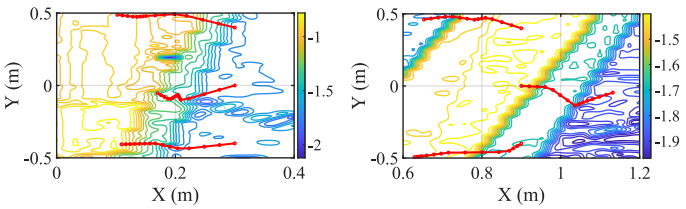

(a) Forward scattering region. (b) Fresnel zone.

Fig. 17. Security analysis.

adverse locations identified by *Loki* correspond to the "weak spots" indicated by wireless sensing theory.

One may readily observe in Fig. 17 that moving an object within the forward scattering region and the Fresnel zone results in significant variations in CSI strength. The optimization trajectories of *Loki* align with the normal direction on the contour map, where CSI strength undergoes the fastest changes. Nevertheless, some discrepancies between the normal direction and these trajectories suggest the need for *Loki*'s ray-tracing approach instead of relying solely on conventional wireless sensing models. These findings reveal that *Loki*'s optimization process consistently identifies areas that can disrupt wireless indoor localization, regardless of an object's initial position. These "weak spots" cause pronounced changes in CSI, posing a considerable risk for errors in localization models. Moreover, these weak spots are inherent to the physical characteristics of wireless signals, rendering them unavoidable. This highlights the effectiveness and robustness of *Loki*. In the following section, we further explore strategies to enhance wireless indoor localization based on this analysis.

One effective strategy to defend against *Loki* is to enhance localization models by incorporating more input data, and a practical approach to achieve this is by increasing the number of Tx devices and antennas. This creates a richer and more resilient dataset for the localization models. To evaluate the efficacy of this defense method, we introduce one or two additional Tx devices into the scene and customize the models to accommodate the increased CSI input. As detailed in Table IV, our findings show that the localization error decreases as the number of Tx devices increases: from an average error of 3.57 m to 2.96 m and further down to 2.52 m. Nevertheless, adding more transceivers is not always feasible in the real world, thus limiting the practicality of such defense.

Another promising defense against the *Loki* attack is adversarial training [57], [58]. The technique involves integrating a subset of adversarial examples into the training data, thereby enhancing model robustness against adversarial attacks during inference. Notably, adversarial training serves as an equivalent mechanism to retraining with post-deployment data. In our approach, we gather CSI that includes adversarial samples and use the data to train wireless indoor localization models. By following the process outlined in Section VI, we generate adversarial CSI specifically targeting the *Loki* attack. We expect this training method to reduce the models' sensitivity to such attacks. Upon applying adversarial training and assessing the models' performance against *Loki* attacks, the results in Table IV reveal a reduction in average localization errors from 3.57 m to 2.85 m. We attribute the inferior performance of adversarial training, compared with adding transceivers, to the insufficient amount of adversarial samples. Nevertheless, these findings underscore the potential of adversarial training to create more resilient wireless indoor localization models.

TABLE IV
EFFECTS OF POTENTIAL DEFENSE METHODS

Localization error (m)	Adding 1 Tx	Adding 2 Tx	Adversarial training
CiFi	3.37 (↓0.58)	2.91(↓1.04)	3.19 (↓0.76)
AAR	3.04 (↓0.61)	2.70 (↓0.95)	2.98 (↓0.67)
SIABR	2.79 (↓0.56)	2.28 (↓1.07)	2.51 (↓0.84
DLoc	2.65 (↓0.70)	2.21 (↓1.14)	2.73 (↓0.62)

VIII. Related Works

Wireless indoor localization has become increasingly susceptible to various types of attacks. One such attack is spoofing [59], [60], where a malicious party introduces one or more transmitters into the target area to inject fake signals. This deceit results in the localization model producing incorrect position estimations. Additionally, attackers may deploy signal jammers that intentionally emit distorting signals omnidirectionally, significantly compromising localization accuracy [11]. The threat intensifies with collaborative jammer attacks, where multiple jammers are used, thereby amplifying the range and intensity of the jamming signals [13]. However, spoofing and jamming are often impractical in real-world scenarios due to two key limitations: i) they require expensive equipment, hindering widespread use, and ii) attackers need specialized expertise to execute complex operations like signal synchronization, fabrication, and replay.

The successful application of DL in wireless indoor localization has brought new security challenges. An adversary can deliberately perturb input data to deceive an underlying DL model [54]. Typically, the perturbation is directly added into the input of these wireless indoor localization models [14], which is usually deemed unrealistic. The work [61] develops a reinforcement learning-based attack, though its effectiveness is limited by challenges in predicting device movements. A more practical approach is to manipulate interference signals through a third-party device to produce adversarial samples [15]–[17], [19]. Nowadays, physical adversarial attacks have garnered significant attention, especially in the realm of wireless indoor localization. These attacks are not only more realistic but also pose greater challenges, as they demand robust perturbations capable of withstanding various environment factors. The Phy-Adv [19] attack targets the susceptibility of wireless indoor localization to environment conditions by ingeniously crafting a shield using common materials to create perturbations. However, this approach is merely a preliminary step and suffers from limitations such as inaccessibility and conspicuousness. Consequently, there remains substantial potential for further exploration in the domain of physical adversarial attacks on wireless indoor localization.

IX. Discussion

While *Loki* demonstrates superior and robust attacks on wireless indoor localization, it still bears certain limitations. First, RT technique is computationally expensive. Nevertheless, our differentiable ray tracer remains tractable in our experimental setup. Specifically, the time and memory cost to simulate the CSI for a single transmitter-receiver pair takes approximately 374 ms on average and 6.8 GB GPU memory on an NVIDIA RTX 4090 GPU, respectively. Second, manual construction of indoor environments remains labor-intensive. When feasible, LiDAR can be employed to quickly capture the geometry of the environment. In practice, a complete LiDAR

scan can achieve a position accuracy of under 2 cm. Subsequently, we apply data-driven calibration to automatically refine the environment model's accuracy. Finally, as wireless indoor localization systems become increasingly widespread, the societal implications of our work become more significant. The ability to manipulate these systems presents security risks across various applications. To mitigate these risks, we recommend deploying additional wireless transceivers along with periodic retraining or adversarial training. It is also important to emphasize that the objective of our work is to expose vulnerabilities present in these systems and raise awareness, thereby stimulating the development of more robust security measures.

X. Conclusion

In this paper, we propose *Loki* as the first physical-world adversarial attack on wireless indoor localization utilizing differentiable object placement. Exploiting existing objects in the scene as attack vectors, and only reposition them by a few centimeters, *Loki* offers superior effectiveness, accessibility, inconspicuity, and ease of operation. Our experiments validate the large localization error and high ASR achieved by *Loki*, demonstrating its successful application in both ray-tracing simulation and real-world scenes. Our results expose critical vulnerabilities in widely used wireless indoor localization models and hence underscore an urgent need for enhanced security measures against such attacks.

Acknowledgments

The study is supported by Shenzhen Science and Technology Program (No. 20231120215201001) and National Natural Science Foundation of China (No. 62502191).

References

[1] F. Zafari, A. Gkelias, and K. K. Leung, "A Survey of Indoor Localization Systems and Technologies," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 3, pp. 2568–2599, 2019.

[2] J. Hu, Z. Chen, T. Zheng, R. Schober, and J. Luo, "HoloFed: Environment-adaptive Positioning via Multi-band Reconfigurable Holographic Surfaces and Federated Learning," *IEEE Journal on Selected Areas in Communications*, vol. 41, no. 12, pp. 3736–3751, 2023.

[3] T. Ehrens, "iBeacon," https://www.techtarget.com/iotagenda/definition/Apple-iBeacon, 2019.

[4] Google, "Geolocation API Overview," https://developers.google.com/maps/documentation/geolocation/overview?hl=zh-cn, 2021.

[5] Cisco, "Cisco Spaces Data Sheet," https://www.cisco.com/c/en/us/products/collateral/wireless/dna-spaces/datasheet-c78-741786.html, 2022.

[6] L. Guo, Z. Lu, S. Zhou, X. Wen, and Z. He, "When Healthcare Meets Off-the-shelf WiFi: A Non-wearable and Low-costs Approach for In-home Monitoring," *arXiv preprint arXiv:2009.09715*, 2020.

[7] T. Kulshrestha, D. Saxena, R. Niyogi, and J. Cao, "Real-time Crowd Monitoring using Seamless Indoor-outdoor Localization," *IEEE Transactions on Mobile Computing*, vol. 19, no. 3, pp. 664–679, 2019.

[8] N. Ghourchian, M. Allegue-Martinez, and D. Precup, "Real-time Indoor Localization in Smart Homes using Semi-supervised Learning," in *Proc. of AAAI*, vol. 31, no. 2, 2017, pp. 4670–4677.

[9] D. Balakrishnan and A. Nayak, "An Efficient Approach for Mobile Asset Tracking using Contexts," *IEEE Transactions on Parallel and Distributed Systems*, vol. 23, no. 2, pp. 211–218, 2011.

[10] H. Cao, W. Huang, G. Xu, X. Chen, Z. He, J. Hu, H. Jiang, and Y. Fang, "Security Analysis of Wifi-based Sensing Systems: Threats from Perturbation Attacks," *arXiv preprint arXiv:2404.15587*, 2024.

[11] S. Gezici, M. R. Gholami, S. Bayram, and M. Jansson, "Jamming of Wireless Localization Systems," *IEEE Transactions on Communications*, vol. 64, no. 6, pp. 2660–2676, 2016.

[12] D. Moser, P. Leu, V. Lenders, A. Ranganathan, F. Ricciato, and S. Capkun, "Investigation of Multi-device Location Spoofing Attacks on Air Traffic Control and Possible Countermeasures," in *Proc. of the 22nd ACM MobiCom*, ser. MobiCom '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 375–386. [Online]. Available: https://doi.org/10.1145/2973750.2973763

[13] C. Goztepe, S. Büyükçorak, G. K. Kurt, and H. Yanikomeroglu, "Localization Threats in Next-Generation Wireless Networks," *IEEE Communications Magazine*, vol. 59, no. 9, pp. 51–57, 2021.

[14] X. Wang, X. Wang, S. Mao, J. Zhang, S. C. Periaswamy, and J. Patton, "Adversarial Deep Learning for Indoor Localization with Channel State Information Tensors," *IEEE Internet of Things Journal*, vol. 9, no. 19, pp. 18 182–18 194, 2022.

[15] P. Huang, E. Gönültaş, M. Arnold, K. P. Srinath, J. Hoydis, and C. Studer, "Attacking and Defending Deep-Learning-Based Off-Device Wireless Positioning Systems," *IEEE Transactions on Wireless Communications*, 2024.

[16] Z. Liu, C. Xu, Y. Xie, E. Sie, F. Yang, K. Karwaski, G. Singh, Z. L. Li, Y. Zhou, D. Vasisht *et al.*, "Exploring Practical Vulnerabilities of Machine Learning-based Wireless Systems," in *Proc. of the 20th USENIX NSDI*, 2023, pp. 1801–1817.

[17] C. Li, M. Xu, Y. Du, L. Liu, C. Shi, Y. Wang, H. Liu, and Y. Chen, "Practical Adversarial Attack on WiFi Sensing Through Unnoticeable Communication Packet Perturbation," in *Proc. of the 30th ACM MobiCom*, 2024, pp. 373–387.

[18] T. Zheng, J. Hu, R. Tan, Y. Zhang, Y. He, and J. Luo, "{π-Jack}:{Physical-World} Adversarial Attack on Monocular Depth Estimation with Perspective Hijacking," in *Proc. of the 33rd USENIX Security*, 2024, pp. 7321–7338.

[19] J. Wang, Y. Tao, Y. Zhang, W. Liu, Y. Kong, S. Tan, R. Yan, and X. Liu, "Adversarial Examples against WiFi Fingerprint-based Localization in the Physical World," *IEEE Transactions on Information Forensics and Security*, 2024.

[20] K. Wu, J. Xiao, Y. Yi, D. Chen, X. Luo, and L. M. Ni, "CSI-Based Indoor Localization," *IEEE Transactions on Parallel and Distributed Systems*, vol. 24, no. 7, pp. 1300–1309, 2013.

[21] Z. Yun and M. F. Iskander, "Ray Tracing for Radio Propagation Modeling: Principles and Applications," *IEEE Access*, vol. 3, pp. 1089–1100, 2015.

[22] X. Han, T. Zheng, T. X. Han, and J. Luo, "RayLoc: Wireless Indoor Localization via Fully Differentiable Ray-tracing," *arXiv preprint arXiv:2501.17881*, 2025.

[23] M. Kotaru, K. Joshi, D. Bharadia, and S. Katti, "SpotFi: Decimeter Level Localization using WiFi," in *Proc. of ACM SIGCOMM*, 2015, pp. 269–282.

[24] R. Ayyalasomayajula, A. Arun, C. Wu, S. Sharma, A. R. Sethi, D. Vasisht, and D. Bharadia, "Deep Learning based Wireless Localization for Indoor Navigation," in *Proc. of the 26th ACM MobiCom*, 2020, pp. 1–14.

[25] S. Fan, Y. Wu, C. Han, and X. Wang, "SIABR: A Structured Intra-Attention Bidirectional Recurrent Deep Learning Method for Ultra-Accurate Terahertz Indoor Localization," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 7, pp. 2226–2240, 2021.

[26] J. Huang, J.-X. Bai, X. Zhang, Z. Liu, Y. Feng, J. Liu, X. Sun, M. Dong, and M. Li, "KeystrokeSniffer: An Off-the-Shelf Smartphone Can Eavesdrop on Your Privacy from Anywhere," *IEEE Transactions on Information Forensics and Security*, 2024.

[27] T. Zheng, Z. Chen, S. Zhang, and J. Luo, "Catch Your Breath: Simultaneous RF Tracking and Respiration Monitoring with Radar Pairs," *IEEE Transactions on Mobile Computing*, vol. 22, no. 11, pp. 6283–6296, 2022.

[28] Z. Chen, T. Zheng, and J. Luo, "Octopus: A Practical and Versatile Wideband MIMO Sensing Platform," in *Proc. of the 27th ACM MobiCom*, 2021, pp. 601–614.

[29] J. Xiong and K. Jamieson, "ArrayTrack: A Fine-Grained Indoor Location System," in *Proc. of the 10th USENIX NSDI*, 2013, pp. 71–84.

[30] T. Zheng, Z. Chen, S. Ding, and J. Luo, "Enhancing RF Sensing with Deep Learning: A Layered Approach," *IEEE Communications Magazine*, vol. 59, no. 2, pp. 70–76, 2021.

[31] A. Yassin, Y. Nasser, M. Awad, A. Al-Dubai, R. Liu, C. Yuen, R. Raulefs, and E. Aboutanios, "Recent Advances in Indoor Localization: A Survey on Theoretical Approaches and Applications," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 2, pp. 1327–1346, 2017.

[32] P. Series, "Effects of Building Materials and Structures on Radiowave Propagation Above About 100 MHz," *Recommendation ITU-R*, pp. 2040–1, 2015.
[33] J. Hoydis, F. A. Aoudia, S. Cammerer, F. Euchner, M. Nimier-David, S. t. Brink, and A. Keller, "Learning Radio Environments by Differentiable Ray Tracing," *arXiv preprint arXiv:2311.18558*, 2023.
[34] C. Ruah, O. Simeone, J. Hoydis, and B. Al-Hashimi, "Calibrating Wireless Ray Tracing for Digital Twinning Using Local Phase Error Estimates," *IEEE Transactions on Machine Learning in Communications and Networking*, vol. 2, pp. 1193–1215, 2024.
[35] D. He, B. Ai, K. Guan, L. Wang, Z. Zhong, and T. Kürner, "The Design and Applications of High-Performance Ray-Tracing Simulation Platform for 5G and Beyond Wireless Communications: A Tutorial," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 1, pp. 10–27, 2019.
[36] H. Ling, R.-C. Chou, and S.-W. Lee, "Shooting and Bouncing Rays: Calculating the RCS of An Arbitrarily Shaped Cavity," *IEEE Transactions on Antennas and Propagation*, vol. 37, no. 2, pp. 194–205, 1989.
[37] W. K. Hastings, "Monte Carlo Sampling Methods Using Markov Chains and Their Applications," *Biometrika*, vol. 57, no. 1, pp. 97–109, 04 1970.
[38] T. Fugen, J. Maurer, T. Kayser, and W. Wiesbeck, "Capability of 3-D Ray Tracing for Defining Parameter Sets for the Specification of Future Mobile Communications Systems," *IEEE Transactions on Antennas and Propagation*, vol. 54, no. 11, pp. 3125–3137, 2006.
[39] G. Deschamps, "Ray Techniques in Electromagnetics," *Proc. of the IEEE*, vol. 60, no. 9, pp. 1022–1035, 1972.
[40] V. Degli-Esposti, F. Fuschini, E. M. Vitucci, and G. Falciasecca, "Measurement and Modelling of Scattering From Buildings," *IEEE Transactions on Antennas and Propagation*, vol. 55, no. 1, pp. 143–153, 2007.
[41] R. Kouyoumjian and P. Pathak, "A Uniform Geometrical Theory of Diffraction for An Edge in A Perfectly Conducting Surface," *Proc. of the IEEE*, vol. 62, no. 11, pp. 1448–1461, 1974.
[42] C. A. Balanis, *Antenna Theory: Analysis and Design*. John Wiley & Sons, 2016.
[43] G. Loubet, N. Holzschuch, and W. Jakob, "Reparameterizing Discontinuous Integrands for Differentiable Rendering," *ACM Transactions on Graphics (TOG)*, vol. 38, no. 6, pp. 1–14, 2019.
[44] Z. Zhang, N. Roussel, and W. Jakob, "Projective Sampling for Differentiable Rendering of Geometry," *ACM Transactions on Graphics (TOG)*, vol. 42, no. 6, pp. 1–14, 2023.
[45] J. Eertmans, L. Jacques, and C. Oestges, "Fully Differentiable Ray Tracing via Discontinuity Smoothing for Radio Network Optimization," in *Proc. of the 18th EuCAP*, 2024, pp. 1–5.
[46] N. Carlini and D. Wagner, "Towards Evaluating the Robustness of Neural Networks," in *Proc. of IEEE S&P*, 2017, pp. 39–57.
[47] O. Khatib, "Real-time Obstacle Avoidance for Manipulators and Mobile Robots," in *Proc. of IEEE ICRA*, vol. 2, 1985, pp. 500–505.
[48] S. Zhou and G. Giannakis, "Adaptive Modulation for Multiantenna Transmissions with Channel Mean Feedback," *IEEE Transactions on Wireless Communications*, vol. 3, no. 5, pp. 1626–1636, 2004.
[49] A. Athalye, L. Engstrom, A. Ilyas, and K. Kwok, "Synthesizing Robust Adversarial Examples," in *Proc. the 35th ICML*, ser. Proceedings of Machine Learning Research, J. Dy and A. Krause, Eds., vol. 80. PMLR, 10–15 Jul 2018, pp. 284–293.
[50] W. Jakob, S. Speierer, N. Roussel, M. Nimier-David, D. Vicini, T. Zeltner, B. Nicolet, M. Crespo, V. Leroy, and Z. Zhang, "Mitsuba 3 Renderer," 2022, https://mitsuba-renderer.org.
[51] F. Gringoli, M. Cominelli, A. Blanco, and J. Widmer, "AX-CSI: Enabling CSI Extraction on Commercial 802.11 Ax Wi-Fi Platforms," in *Proc. of the 15th ACM Workshop on WiNTECH*, 2022, pp. 46–53.
[52] X. Wang, X. Wang, and S. Mao, "Deep Convolutional Neural Networks for Indoor Localization with CSI Images," *IEEE Transactions on Network Science and Engineering*, vol. 7, no. 1, pp. 316–327, 2020.
[53] B. Zhang, H. Sifaou, and G. Y. Li, "CSI-Fingerprinting Indoor Localization via Attention-Augmented Residual Convolutional Neural Network," *IEEE Transactions on Wireless Communications*, vol. 22, no. 8, pp. 5583–5597, 2023.
[54] C. Szegedy, "Intriguing Properties of Neural Networks," *arXiv preprint arXiv:1312.6199*, 2013.
[55] J. H. Holland, "Genetic Algorithms," *Scientific American*, vol. 267, no. 1, pp. 66–73, 1992.
[56] J. Kennedy and R. Eberhart, "Particle Swarm Optimization," in *Proc. of ICNN*, vol. 4. IEEE, 1995, pp. 1942–1948.
[57] A. Shafahi, M. Najibi, M. A. Ghiasi, Z. Xu, J. Dickerson, C. Studer, L. S. Davis, G. Taylor, and T. Goldstein, "Adversarial Training for Free!" in *Proc. of NeurIPS*, vol. 32. Curran Associates, Inc., 2019.
[58] T. Zheng, Z. Chen, S. Ding, C. Cai, and J. Luo, "Adv-4-Adv: Thwarting Changing Adversarial Perturbations via Adversarial Domain Adaptation," *NeuroComputing*, vol. 569, p. 127114, 2024.
[59] J. Yang, Y. Chen, W. Trappe, and J. Cheng, "Detection and Localization of Multiple Spoofing Attackers in Wireless Networks," *IEEE Transactions on Parallel and Distributed Systems*, vol. 24, no. 1, pp. 44–58, 2013.
[60] Z. Sun, S. Balakrishnan, L. Su, A. Bhuyan, P. Wang, and C. Qiao, "Who is in Control? Practical Physical Layer Attack and Defense for Mmwave-based Sensing in Autonomous Vehicles," *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 3199–3214, 2021.
[61] R. Li, H. Hu, and Q. Ye, "RFTrack: Stealthy Location Inference and Tracking Attack on Wi-Fi Devices," *IEEE Transactions on Information Forensics and Security*, 2024.