Why Don't You Clean Your Glasses? Perception Attacks with Dynamic Optical Perturbations

Yi Han
Rutgers University
yh482@rutgers.edu

Matthew Chan
Rutgers University
matthew.chan@rutgers.edu

Eric Wengrowski
Steg AI
eric.wengrowski@gmail.com

Zhuohuan Li
Rutgers University
zl253@scarletmail.rutgers.edu

Nils Ole Tippenhauer
CISPA Helmholtz Center for Information Security
tippenhauer@cispa.de

Mani Srivastava
University of California, Los Angeles
mbs@ucla.edu

Saman Zonouz
Georgia Institute of Technology
saman.zonouz@gatech.edu

Luis Garcia
University of Utah
lgarcia@cs.utah.edu

Abstract—Camera-based autonomous systems that emulate human perception are increasingly being integrated into safety-critical platforms. Consequently, an established body of literature has emerged that explores adversarial attacks targeting the underlying machine learning models. Adapting adversarial attacks to the physical world is desirable for the attacker, as this removes the need to compromise digital systems. However, the real world poses challenges related to the "survivability" of adversarial manipulations given environmental noise in perception pipelines and the dynamicity of autonomous systems. In this paper, we take a sensor-first approach. We present EVILEYE, a man-in-the-middle perception attack that leverages transparent displays to generate dynamic physical adversarial examples. EVILEYE exploits the camera's optics to induce misclassifications under a variety of illumination conditions. To generate dynamic perturbations, we formalize the projection of a digital attack into the physical domain by modeling the transformation function of the captured image through the optical pipeline. Our extensive experiments show that EVILEYE's generated adversarial perturbations are much more robust across varying environmental light conditions relative to existing physical perturbation frameworks, achieving a high attack success rate (ASR) while bypassing state-of-the-art physical adversarial detection frameworks. We demonstrate that the dynamic nature of EVILEYE enables attackers to adapt adversarial examples across a variety of objects with a significantly higher ASR compared to state-of-the-art physical world attack frameworks. Finally, we discuss mitigation strategies against the EVILEYE attack.

I. INTRODUCTION

Safety-critical autonomous systems commonly rely on optical sensors coupled with artificial intelligence to mimic human perception across a variety of domains, including autonomous navigation [42], crowd surveillance and analysis [50], and facial recognition for authentication [33]. Although the associated deep learning algorithms enable promising performance across domains [12], they introduce vulnerabilities that are non-obvious and can be targeted by adversaries. There have already been real-world instances where perceptual models have failed, resulting in fatalities [4]. Researchers have shown various autonomous sensors to be exploitable [67], even production vehicles such as Tesla cars can be compromised [8].

Much of the initial research on adversarial machine learning [17] has focused on digital-domain software attacks, a holdover from more thoroughly-explored computer vision research. More recent work has explored the feasibility of adversarial machine learning attacks on sensors in the physical domain, where dynamic environmental conditions make precision attacks more difficult to execute [11, 6, 2, 53, 69, 64].

The development of robust, physical-domain attacks has trended from static, object-level attacks [11, 6, 2, 53, 69, 64, 21, 32, 57, 59] to dynamic, object-level attacks [15, 30, 39, 37] for misclassifying specific objects for non-static domains, e.g., autonomous vehicles detecting objects in various environments. However, object-level approaches hinge on augmented objects entering the target perception pipeline, e.g., an autonomous vehicle passing by a maliciously modified stop sign. Moreover, to attack multiple objects, each object needs to be modified individually. To overcome the aforementioned problems, recent attacks [20, 23, 32, 59, 26, 71] have started exploring sensor-modifying attacks. These attacks typically exploit vulnerable sensors (e.g. cameras, CMOS sensors or inertial sensors) in the perception pipeline of the victim system. Most relevant to this paper, in [26, 71] showed that stickers with static, adversarial perturbations could be placed on the lens of a target camera. However, such an approach cannot adapt to various scenes and environmental conditions due to the static nature of the stickers.

In this paper, we fill the gap by proposing an intelligent model, EVILEYE, to dynamically attack the sensor pipeline. EVILEYE utilizes a portable transparent display in front of an optical sensor to display a carefully crafted adversarial perturbation to alter the perception of the sensor. An auxiliary sensor is utilized to detect when a target object, e.g., a "Stop" sign, enters the victim camera's field of view, which dynamically informs the generation of an adversarial perturbation.

TABLE I: A comparison of the characteristics of closely related works. ● indicates taking this approach, ◐ indicates the work could plausibly take this approach, and ○ means that the work does not utilize the approach. OM: Object-Modifying; SM: Sensor-Modifying; SA: Static Attack; DA: Dynamic Attack; OA: Online Attack.

	OM	SM	SA	DA	OA
Li *et al.* [26] (2019)	○	●	●	○	○
Nassi *et al.* [37] (2020)	●	○	○	●	◐
Lovisotto *et al.* [30] (2020)	●	○	●	◐	○
Zolfi *et al.* [71] (2021)	○	●	●	○	○
Wang *et al.* [57] (2021)	●	○	●	○	○
Wang *et al.* [58] (2021)	●	○	●	○	○
Jia *et al.* [21] (2022)	●	○	●	◐	○
EVILEYE (this work)	○	●	●	●	●

To craft adversarial perturbations, EVILEYE first models the optical sensor pipeline, i.e., the optical path of a digitally generated adversarial perturbation from the transparent display to the optical sensor, using a feed-forward neural network. The sensor pipeline model enables EVILEYE to simulate physical adversarial attacks by implicitly accounting for the spectral sensitivity of the camera sensor, the spectral emittance of the transparent display, and the radiometric properties of light in free space [62, 61]. EVILEYE then utilizes a gradient-based iterative approach to generate adversarial perturbations. To make sure adversarial perturbations can survive dynamic conditions in a physical environment, EVILEYE simulates various environmental factors (e.g., perspective, distance, and illuminance) when searching for an adversarial perturbation.

EVILEYE has the following advantages over prior attack frameworks: 1) our attack does not require knowledge of when and where a target object will enter the periphery of a target sensor since perturbation decisions can be made at runtime, whereas prior works require a target sensor to follow a particular path; 2) our attack is more flexible than prior approaches because the adversarial perturbations can be adjusted dynamically (e.g., targeting different traffic signs for misclassification). According to our extensive empirical evaluations, this results in better attack performance; and 3) our attack is robust across various environmental lighting conditions. We empirically evaluate EVILEYE in the traffic sign recognition domain. We will make our code open source for reproduction and future research.

Contributions: The contributions set forth in the paper are summarized as follows:

- We introduce a new class of dynamic, sensor-first machine learning attacks in the physical domain.
- We formalize the contextualization of adversarial examples in the optical sensor pipeline.
- We develop and validate a novel, compact device designed to attack the optics of a camera system.
- We extensively evaluate the efficacy of our sensor-first attack on a traffic sign recognition task across various levels of illumination, significantly outperforming existing approaches.

II. PRELIMINARIES

In this section, we will describe common object perception pipelines in safety-critical, camera-based autonomous systems. We then provide an overview of the different classes of physical AEs presented in Table I. Finally, we provide the necessary preliminaries on the transparent display technology used to realize EVILEYE.

A. Camera-based Object Detection

Deep learning-enabled, camera-based safety-critical systems are typically compartmentalized into perception, planning, and control stages [42]. In this paper, we target the perception stage, where a camera captures the light reflecting off scene objects. The camera produces a single frame within a video sequence for a given time interval. Next, the object recognition module behind the optical sensors detects, regionally segments, and classifies the objects within each frame. Finally, the perceived objects are fed into the planning and control stages, which are domain-specific and tied to the underlying applications. Crowd surveillance and analysis [50] and autonomous navigation [3] are both examples of such systems.

B. Physical Adversarial Examples

A substantial body of research focuses on adversarial examples (AEs) in the digital domain [60]. Recent research has emerged where AEs are crafted in the physical domain to demonstrate the practicality of these attacks in the real world. In order to account for dynamic physical environmental conditions, many existing physical attacks employ Expectation over Transformation (EOT) [2] when generating adversarial perturbations. EOT is a data augmentation technique where data is synthesized using various transformations naturally found in a physical environment, such as varying lighting and camera angle. The adversarial perturbation pipeline highly depends on the placement of the perturbation relative to the victim sensor as well as the target perceived object.

As shown in Table I, we categorize physical domain attacks based on the location of the perturbation, i.e., object-level modification (OM) versus sensor-level modification (SM), as well as the dynamicity of the attack, i.e., static attacks (SA) calculated ahead of time versus dynamic attacks (DA) that can change at runtime. Additionally, we distinguish online attacks (OA) that generate dynamic perturbations based on the current context of the attack. Early static approaches targeted specific, individual objects [11, 6, 2, 53, 69, 64, 21, 58, 57]. These attacks change the appearance of a target object by attaching carefully crafted adversarial "patches" (e.g., stickers or posters) to the object. However, these static attacks have the disadvantage that the patches cannot be changed once applied.

More recent efforts focus on enabling dynamic adversarial attacks to increase stealthiness or situational awareness. For example, an adversarial perturbation would only be displayed when a target vehicle is approaching a target street sign. Similarly, different perturbations could be generated for targeting different street signs. Lovisotto et al. [30] propose using a projector to cast adversarial perturbations onto the object. Such

an attack works well for dynamically generating perturbations for specific objects, like a stop sign at a particular intersection. Recent attacks also proposed sensor-first approaches. Instead of altering objects in a scene, they focus on what the optical sensor perceives. In the work of Li et al. [26] and Zolfi et al. [71], an adversary prints adversarial perturbations on a transparent paper and attaches the paper to a victim's camera. In this way, the adversarial perturbation follows the victim camera continuously. However, the sticker perturbations are static and do not have the advantages of stealthiness and situational awareness associated with dynamic attacks. In this paper, we aim to bridge the gap between both approaches by proposing a dynamic, sensor-first attack framework. We now describe the transparent mediums that enable such a framework.

C. Transparent Displays

A dynamic, sensor-first attack hinges on the availability of an appropriate attack medium. For camera-based sensing systems, our desired attack vector is light. We seek a medium that will project directly into a camera's optics. The ideal medium is a transparent display that 1) supports dynamic perturbations–as in the projector-based attacks [15, 30, 39, 37], 2) can generate perturbations based on the context of a target object [37], and 3) resides on or near the camera lens such that all light perceived by the camera is transmitted through the medium–as in the case of adversarial camera stickers [26]. A variety of transparent displays are available as consumer products, with a majority catered to augmented reality (AR) applications such as AR glasses or heads-up displays (HUDs). These products display visual objects and auxiliary information within the user's field of view for various applications, including gaming [52], driving assistance [1], surgical assistance [56], and construction safety [27].

At their core, existing AR solutions consist of two main components: a light source and a holographic combiner (i.e., a specialized lens) that renders images from the projected light into the user's field of view. Precise positioning of the combiner and projection source results in the user simultaneously viewing a scene along with the overlaid projection. Thus, AR technology provides an ideal transparent medium to project dynamic perturbations into the field-of-view of a target camera–serving as a premise for EVILEYE's design.

III. THREAT MODEL

This section describes the system and attacker models we consider for the EVILEYE framework, including a description of the use-cases we focus on and overall attack goals.

A. System Model

This work focuses on safety-critical autonomous systems that rely on camera-based perception. Examples include crowd surveillance where object recognition is used on camera feeds to detect human subjects of interest, or autonomous navigation, where a vehicle uses cameras to navigate through its environment. In general, these perception-to-acutation pipelines

are composed of several steps. Initially, the scene from the optical sensor is converted into a digital image. Next, a deep learning-based perception model is used to perform both object detection and classification on the image.

We assume several protection mechanisms on the autonomous system to preclude basic attacks. Autonomous vehicles are often equipped with obstruction detection mechanisms and will disable capabilities such as autonomous steering when an obstruction is detected [49]. Additionally, some applications have human-in-the-loop control [50] or may review footage in post-incident analyses, detecting attacks trying to directly fool the classifier (e.g., with overlaid images).

B. Attacker Model

We assume an adversary attaches a portable and low-profile gadget, consisting of an adversary-controlled transparent display and auxiliary sensor, in front of the victim's optical sensor or camera. We assume that an auxiliary sensor can sense when a target object enters the victim camera's field of view. The auxiliary sensor can take on many forms depending on the target application's real-time requirements and the attacker's resource constraints. For instance, the auxiliary sensor can be another camera attached alongside the victim camera with on-device object detection, or the sensor could be a GPS tracking device coupled with a map of traffic signs throughout the city. Alternatively, the "sensor" could be the attacker remotely observing the victim camera's movement while controlling the transparent gadget. The auxiliary sensor informs the transparent display gadget which adversarial perturbation to display. We assume that the adversary can prepare adversarial perturbations in advance for each target class or that the real-time requirements for the target application allow the controller to generate perturbations at runtime[1]. Previous physical attack frameworks require prior knowledge of a victim's intended route in order to physically modify objects (e.g. with stickers) before they are in view, limiting dynamicity. Using information from an auxiliary sensor (e.g. a GPS), we can situationally adapt perturbations to attack the current object in view.

Attacker goals. The attacker's primary objective is to fool the victim application's recognition module, leading to erroneous object predictions which compromise safety, all while evading detection. This means that 1) employing denial-of-service attacks (e.g. applying opaque obstructions) or directly presenting adversarial images (e.g. an incorrect traffic sign) will be ineffective; 2) perturbing system optics during other states of normal operation is also undesirable, as it increases risk of detection; and 3) the attack should be able function under varied environmental conditions, which especially important for light-based attacks. EVILEYE is designed to achieve this goal while navigating the aforementioned constraints, generating dynamic perturbations optimized to look like chromatic aberrations or distortions of the camera sensor.

---

[1]We note that current commercial off-the-shelf embedded technology may not be able to run real-time perturbation generation, so we assume that the adversary can prepare perturbations in advance as a practical workaround.
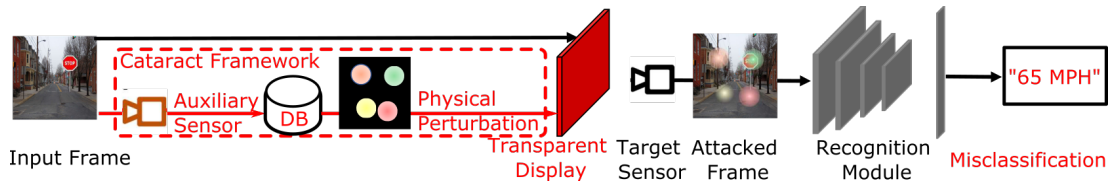
Fig. 1: Overview of EVILEYE framework. An attacker uses a transparent display and an auxiliary sensor to implement an optical man-in-the-middle attack against a target sensor. The attacker first uses the auxiliary sensor to determine the target object for an input frame. EVILEYE then looks up a pretrained digital perturbation to create an unsafe misclassification. In this example, a "Stop" sign is being perceived as a speed-limit sign by the victim system.
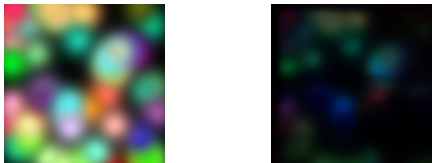
(a) Digital perturbation.   (b) Physical perturbation.

Fig. 2: A pair of digital and physical perturbations. The physical perturbation looks different from the digital perturbation due to the characteristics of the display and sensing pipeline.

IV. METHODOLOGY

In this section, we first formalize the research challenges for dynamically generating adversarial perturbations at the sensor-level. We then briefly provide an overview of how our design aims to address each research challenge. Finally, we describe the methodology of each design component for our proposed attack framework.

A. Problem Formalization

Research Challenge #1: Designing perturbations robust to optical transformations. Directly applying an adversarial perturbation computed in the digital domain fails because it does not account for the mutations introduced by the camera-display transfer function (CDTF) [62], i.e., the transformation of a pixel projected in the physical world to the pixel captured by the camera. In our case, an adversarial perturbation is displayed and perceived by the camera, going through a set of transformations as depicted in Figure 2. The color spectrum emitted by an electronic display may not match the color spectral sensitivity function of the camera [62], leading to less color-robust perturbations. Moreover, because the transparent display is in close proximity to the victim camera, the displayed perturbation is blurred in the perceived image due to the Bokeh effect [18]. The image formation pipeline processes within the camera converts the light received by the image sensor to a final digital image that is then affected by exposure, the sensor's sensitivity (ISO), and contrast transforms.

Research Challenge #2: Designing perturbations robust to environmental conditions. Multiple factors must be considered to make an adversarial perturbation robust to dynamically changing conditions in the physical environment. For instance, the camera's perspective, distance, and rotation relative to a target object will all impact the efficacy of a perturbation. The scene's environmental background and ambient illuminance also play a vital role in object recognition.

Research Challenge #3: Efficiently generating perturbations for real-time applications. In addition to optical and physical mutations, the transient nature of real-time applications introduces a significant challenge to AE generation. Environmental conditions can change significantly within a short period of time. Because generating a perturbation incurs latency in the attack pipeline, attack perturbations should be generated to perform well across this environmental variance, rather than attempt to optimize for a single frame.

Design Overview. Figure 3 depicts an overview of the EVIL-EYE framework. To address the digital-to-physical mapping of **Research Challenge #1** and **Research Challenge #2**, we employ an end-to-end deep learning approach for modeling the optical transformations (we provide details in Section IV-B). We incorporate environmental background noise and illuminance variations into an automated data collection framework for perturbation training and generation (details in Section IV-C). Finally, to address the real-time dynamics discussed in **Research Challenge #3**, we adopt a universal adversarial perturbation (UAP) model [35] that is robust against variance across multiple time frames for a given scene. Moreover, we leverage semantic information to efficiently generate this UAP (detailed in Section IV-C).

B. Modeling Digital-to-Physical Perturbations

Under the additive adversarial perturbation model, the effect of a physical adversarial perturbation can be abstracted as

$$y' = F(x + I_p), \quad y \neq y' \tag{1}$$

where x is a frame of the video stream perceived by the victim camera. I_p is a physical adversarial perturbation. $F(\cdot)$ is the recognition model. y is the true label of the perceived object. y' is an intentionally misclassified label. I_p can be produced in different ways such as printed on a sticker and attached to the object [11], printed on a transparent paper and attached to the camera sensor [26] or projected onto the object using a projector [30]. In our attack, I_p is presented on a transparent display and attached to the camera sensor. The camera sensor can see through the display.

Applying adversarial perturbation computed in the digital domain fails due to the mutations/transformations introduced
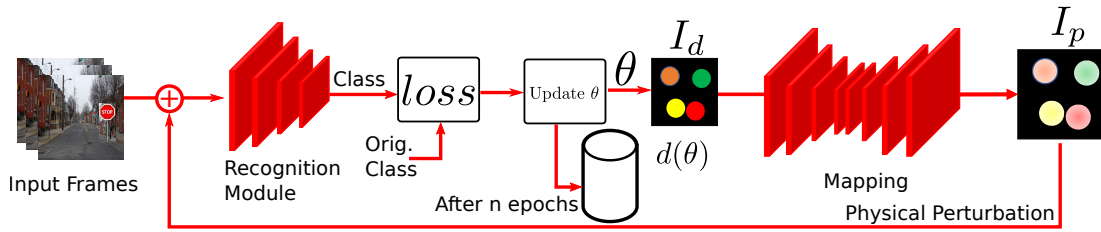
Fig. 3: Overview of training process for physical perturbations in the EVILEYE framework. EVILEYE generates a digital perturbation to create an unsafe misclassification. EVILEYE's digital-to-physical mapping then transforms the digital perturbation into a physical perturbation, i.e., the projection of the digital perturbation onto the transparent display.

by the CDTF [62]. Such mutations are caused by the characteristics of the displaying and camera sensing pipeline. On the transparent display side, the optics of the holographic combiner and projection source such as color spectrum compression and shifting, distortion will already produce an image different from the source image. On the sensor side, the optics of the camera, such as focal length, lens distortion, exposure, and white balance also affect the final perceived image. Therefore, in order to produce effective adversarial perturbations in the physical domain, it is necessary to understand and consider these mutations in the design loop. There is prior effort attempting to model these mutations. Lovisotto et al. [30] used a multi-layer perceptron (MLP) to model the color spectrum mapping between intended and perceived colors given the color of the projection surface. while this mapping model works fine for their attack purpose, it requires discretizing the color space to reduce complexity, which could lead to imprecision. Besides, to achieve a more precise digital-to-physical mapping and thus a more powerful attack, other mutating factors mentioned above also need to be considered. In this paper, we propose establishing a digital-to-physical mapping that takes all these mutations into consideration.

Digital-to-physical mapping. Given a pair of digital and physical perturbation image I_d and I_p, we wish to find a mapping $T(\cdot)$ such that

$$I_p = T_{\Theta, D, C, E}(I_d) \qquad (2)$$

where Θ represents the domain-constraint parameters of the perturbation generation (e.g., the size, shape, and color constraints of a perturbation), D represents the transparent display, C denotes the camera, and E denotes environmental parameters (e.g., ambient light). It would be too complicated to model each component, i.e., aforementioned mutations, individually. Therefore, we take a data-driven approach. We employ a convolutional neural network to approximate $T(\cdot)$.

Fitting the model. $T(\cdot)$ is fitted with pairs of digital and physical perturbations. We generate a dataset of I_d and I_p pairs as the training set. For each pair, we first randomly generate and display I_d on the transparent display. We then take two images, one with I_d turned on (mutated image) and one with I_d turned off (the background). We derive I_p by subtracting the background from the mutated image. When we collect data for fitting $T(\cdot)$, no object is present in the scene.

As for generating I_d, we use highly structured patterns consisting of dots with different colors. Extending the idea of using small dots to create adversarial perturbations in [26], we define a dot-based adversarial perturbation generation function $I_{d,\theta}$. The parameters contained in θ describe the dots, and are defined as 1) $c(i^{(k)}, j^{(k)})$ - center coordinates of the kth dot; 2) r_k - radius of the kth dot; 3) γ_k - RGB color of the kth dot; 4) $\alpha_{\max}$ - maximum alpha blending value; 5) β - exponential dropoff of alpha value; 6) n - maximum number of dots. I_d can be computed as

$$I_d(i,j) = \sum_{k=1}^{n} \alpha_{\max} e^{-\beta \cdot d_k} \cdot \gamma_k \qquad (3)$$

where $d_k(i,j)$ is the distance of the pixel (i,j) with respect to the center of the kth dot

$$d_k(i,j) = \frac{(i - i^{(k)})^2 + (j - j^{(k)})^2}{r_k^2}. \qquad (4)$$

I_d superimposes n dots together to create a pattern with various colors and shapes. Figure 2a shows the example of a digital perturbation. The number of dots in an adversarial perturbation affects its capacity; an adversarial perturbation with fewer dots is less powerful than one with more dots. A perturbation with more dots can create more complex patterns and, thus, can alter the frames and the prediction of the target recognition model more easily. The advantages of using such highly structured patterns are 1) small modulations to pixel values, as in pixel level adversarial perturbations, do not survive when displayed and perceived on an object; 2) It would require significantly more data in the training set in order to ensure generalizability, i.e., to enable $T(\cdot)$ to output physical perturbations based on unseen digital perturbations.

To fit $T(\cdot)$, we want the output of $T(\cdot)$ to be as similar as possible to the ground truth physical perturbation I_p. To achieve this goal, we combine both mean square error (MSE) and Learned Perceptual Image Patch Similarity (LPIPS)–a metric optimized for perceptual loss [68], to form the loss function l

$$l = (1 - a) * \text{MSE}(T(I_d), I_p) + a * \text{LPIPS}(T(I_d), I_p) \qquad (5)$$

where a controls the ratio of the two terms. The MSE term encourages the colors to be as close as possible, the LPIPS loss term emphasizes maintaining the structural patterns.

5

C. Crafting Adversarial Perturbations

In this section, we describe how we design a digital and physical co-optimization framework to find the adversarial perturbation that can evade the target recognition model.

Attack Objective. As is mentioned in research challenge #3, UAP fits better in our attack scenario: an adversary prepares adversarial perturbations for the classes of images they intend to attack beforehand. Each one of these UAPs can cause an incorrect prediction for all the instances of the object in the corresponding class. The adversary then attaches the attack gadget to the victim camera, and remotely controls which UAP to display on the gadget. Thus, we define our attack objective as

$$\arg\min_{\theta} \quad \mathbf{E}_{x \sim D(x|y)} l[(F(x + I_p))] \tag{6}$$

where θ represents the free parameters of the perturbation generation function (e.g., c and γ). l is the loss function of the attack. A sample x is drawn from a distribution $D(x|y)$ (e.g., the class of all stop signs). Depending on the type of attacks the adversary wants to perform and the target recognition model, l can take various forms. For instance, for an untargeted attack on a traffic sign classifier (i.e., to classify a frame into any class other than the true class), l is the reciprocal of the cross entropy loss with respect to the true class. For an untargeted attack on a traffic sign detector, l is the average class score of all the predicted bounding boxes with respect to the class under attack (see Section V-B for details).

Optimizing a Successful Attack. As each module in our attack pipeline is differentiable, it is theoretically possible to approach a solution of Equation 6 using gradient descent. However in practice, we found the effectiveness of an optimized perturbation to be highly sensitive to the initialization of free parameters, namely, c and γ. Also, as shown in [26], the gradients with respect to the free parameters present a highly non-convex loss surface. Therefore, we first find a good initialization using a coarse grained greedy block coordinate descent search. We then apply fine-grained gradient descent using this initialization. Specifically, we split a perturbation into blocks of the same size. The center of these blocks are the candidate locations of the dots c. We also discretize the RGB color space to obtain a fixed set of candidate colors. We then optimize for one dot at a time. For each dot, we try all the candidate locations and colors and pick the one that gives the maximum loss (to have a higher chance of misleading the recognition model). We repeat this process until convergence. Next, using the dots computed above as initialization, we iteratively compute the gradients of the loss with respect to the free parameters and extract the sign of the gradient. We then add a small step in the direction of the sign. Detailed steps of our optimization are shown in Algorithm 1. Note during the optimization of the attack the weights of $T(\cdot)$ remain fixed.

Serving the Attack. Attack optimization is performed in advance for the objects under attack. The generated perturbations are stored in a database. At runtime, the attacker first uses the auxiliary sensor to determine the context of the runtime environment (e.g., a target object entering the scene).

EVILEYE then displays a pre-generated perturbation fetched from the database to create an unsafe classification.

Algorithm 1: Physical AE Generation

Input: Images under attack $X = \{x1, x2, ...x_n\}$, attack objective l, maximum number of iterations *maxiter*, perturbation generation function θ

Output: Digital perturbation I_d

1 Initialize θ with coarse grained greedy coordinate descent search

2 **while** NOT *converge & iter $\leq$ maxiter* **do**

3   1. Compute digital perturbation I_d

4   2. Compute physical perturbation $I_p = T(I_d)$

5   3. Apply I_p to X

6   4. Query the victim model and compute the attack objective l

7   5. Perform a step of gradient descent to update θ

Robust Physical Adversarial Examples To make an adversarial perturbation robust against dynamic physical environmental conditions, we consider the following factors when preparing images for the adversarial perturbation generation process:

- **Background.** The context of the background plays a role in recognizing an object. We use images with various backgrounds within the same context.
- **Perspective and rotation.** The object might come in a different perspective and rotated in front of the camera. We prepare images with various perspectives and rotations. This is application-specific, e.g., when attacking a traffic sign recognition model, we consider perspectives between $-30°$ and $30°$, with rotations between $-5°$ and $5°$.
- **Distance.** The object might also be located at various distances from the camera. To account for this, we consider various sizes of the object.
- **Illuminance.** To account for different light conditions, we vary the ambient light of the environment when collecting images. Specifically, we consider an illuminance ranging from 30 lux to 3000 lux.

Dynamic Adversarial Attack. One major advantage of our proposed attack compared to prior static attacks is that in our attack approach, the adversary can change the perturbation based upon different scenarios. For example, in a traffic sign recognition use case, if a car is following a particular route, we can assume there are several different safety-critical traffic signs along the route. To maximize the damage of the attack, The adversary wants to successfully attack as many critical traffic signs as possible. In the static attack setting, the perturbation cannot be changed, and the adversary can only compute a single UAP to accommodate all different kinds of traffic signs. In our dynamic attack setting, we can change the perturbation based on external information such as GPS of the car and a traffic sign map [9, 55] to maximize the attack

success rate. Moreover, an attacker can enable perturbations only when desired to minimize alerting the vehicle to a camera obstruction or unwanted behavior during normal operation.

V. IMPLEMENTATION AND EVALUATION

In this section, we provide details on the implementation and evaluation of the EVILEYE attack. First, we describe the design of our neural-based approach to digital-to-physical mapping $T(\cdot)$, which we call *TNet*. We then describe our prototype and experimental setup. Finally, we evaluate each component of EVILEYE, including how the attack fares against existing defenses.

A. Digital-to-Physical Mapping via TNet

To approximate the non-linear mapping between digital-to-physical perturbations, we design and train a feed-forward neural network model called TNet. Beyond boasting good performance accuracy, feed-forward neural networks are also differentiable. So a pre-trained TNet can be included as a component of the training procedure of EVILEYE, where differentiation of the mapping function $T(\cdot)$ is essential to the back-propagation step (see Algorithm 1).

The inputs to TNet are RGB images representing digital perturbations, and the outputs are RGB images that illustrate how these digital perturbations will be perceived by the victim camera in the physical world. Two architectural backbones were initially selected for TNet: a fully-convolutional auto-encoder, and the UNet architecture [46]. Empirically, we found that the UNet architecture's multiple skip connections enabled the model to converge faster and with better accuracy across a range of tested batch sizes.

Model Training. To train the mapping, we randomly vary the center and the RGB vector of each dot and generate $10,000$ digital perturbations. We then collect their corresponding physical perturbations using the aforementioned hardware setup. During our initial investigation, we allowed other parameters of the alpha blending model to vary (e.g., r, α_{max}, and γ). With too many free parameters and insufficient training data to cover the combinatorially growing space, TNet cannot learn a robust mapping. Additionally, if the training dataset is not large enough, TNet overfits the training data, and during the attack phase, resulting in a less-precise computed gradient. The optimization will deviate significantly from the correct direction, leading to an ineffective perturbation. Given the data collection bottleneck with a physical setup, we wanted to minimize the reliance on a large dataset. Thus, we fixed all the parameters except the center c and the RGB vector γ of each dot. We found that a dataset with $10,000$ pairs was sufficient to train TNet.

We used 80% of the image pairs for training and 20% for validation. As mentioned in Section IV, the number of dots affects the capacity of our attack: too few dots cannot perturb an image effectively, while effective perturbations with too many dots are difficult to train without a sufficiently large dataset. To show the effect of different number of dots, we consider 10, 30 and 50 dots when training TNet. We set α in

TABLE II: Accuracy of TNet models. The total loss of TNet-UNet on the validation set as well as four other image similarity metrics is significantly better than the other two baseline models. There are not fixed ranges for the metrics (of course they all need to be non-negative numbers) except that SSIM is between 0 and 1. However, a smaller MSE/LPIPS or a larger PSNR/SSIM value indicates better similarity between the predicted physical perturbation and the ground truth.

	Validation loss	MSE	LPIPS	PSNR	SSIM
TNet-UNet	**2.26e-4**	**1.95e-4**	**0.31**	**37.17**	**0.92**
TNet-CNN	2.39e-3	1.42e-3	9.75	28.46	0.83
TNet-MLP	0.01	8.74e-3	53.76	20.58	0.58

Equation 5 to be 0.0004. We trained TNet for $1,000$ epochs with a batch size of 32 and a learning rate of 0.003. We used Adam [22] as the optimizer.

Evaluating TNet. Besides MSE and LPIPS [68]–which we used to train TNet–we also consider Peak Signal-to-Noise Ratio (PSNR) and the Structural Similarity Index (SSIM) [16] to measure the accuracy of TNet. PSNR is derived from MSE, and takes into account the peak intensity of the input images to describe the similarity between two images. SSIM is designed based on luminance, contrast, and structure, so a higher SSIM indicates higher similarity in these factors.

We compared our UNet-based TNet (TNet-UNet) with two baseline models: a CNN Autoencoder-based model (TNet-CNN) and a multi-layer perceptron (MLP)-based model (TNet-MLP), since prior work [30] used an MLP-based approach. For both baseline models, we configure the architecture such that the overall number of weights is similar to TNet-UNet.

Table II shows various metrics of the trained models on the test dataset. TNet-UNet achieves significantly better results in terms of the validation loss and the four metrics compared to the two baseline models, implying that the UNet architecture is the correct choice to approximate the digital-to-physical mapping with high accuracy. Figure 5 depicts an example of different TNet models' output and the corresponding ground truth. A more precise physical mapping guarantees more precise gradient values of the pattern generating function in the attack stage, increasing the probability of a successful adversarial perturbation. We select the TNet-UNet architecture for the rest of the paper and refer to it as TNet.

B. Prototype and Experimental Setup

Implementation. The entire adversarial perturbation generation pipeline is implemented in PyTorch [41]. We utilize an open-source UNet implementation[2] for the TNet architecture. All training and adversarial perturbation generation jobs are run on an Ubuntu 18.10 computer with 4 NVIDIA RTX A5000 graphic cards and 132GB memory. Note in the attack phase the adversary only needs to serve pre-generated perturbations.

Transparent Display Gadget. Borrowing the mechanism behind AR glasses and HUDs, we attach a $50:50$ ratio beam

---

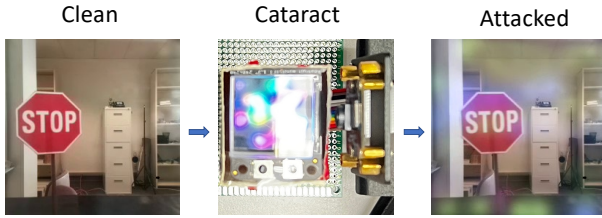[2]https://github.com/zhixuhao/unet

Fig. 4: Our physical evaluation setup, consisting of a Raspberry Pi victim camera and a low-cost transparent display analogue, consisting of an organic light-emitting diode (OLED) display and beam splitter. Adversarial perturbations are displayed on the screen and reflected by the beam splitter, overlaying them onto the natural scene as observed by the victim camera.

splitter to an Adafruit Mini PiTFT 1.3" LED display (4). We control the display using the Python ST7789 library[3]. The beam splitter partially reflects light emitted from the display into the victim camera lens. Light from the perceived object or scene can simultaneously reach the victim camera lens. The gadget is already very compact - slightly larger than an apple watch. Commercial transparent displays for AR can be even smaller (as small as a piece of len in a pair of glasses[4]).

Victim Camera. A Raspberry Pi Camera v2 was selected as the victim camera for all experiments. The camera is controlled by the PiCamera module[5]. The camera's parameters (e.g., auto exposure, white balance) are fixed for all captured images.

Controlling the Light Conditions. Indoor experiments were performed in a controlled laboratory environment. We measured the ambient light with a Urceri MT-912 lux meter. By adjusting the scene illumination, we were able to simulate different ambient environments, e.g., sunny versus cloudy conditions. The default environmental luminous flux was approximately 50 lux. A 65 Watt floodlight placed at varying distances was used to achieve different levels of illumination, with a maximum illuminance of approximately 3000 lux.

C. Traffic Sign Recognition Results

In this experiment, the goal of EVILEYE is to cause misclassifications of traffic sign recognition modules.

Recognition Modules. We consider 3 recognition modules: one classifier and two object detectors. For the classifier, we finetune a Resnet-50 [14] using the combined dataset of the LISA traffic sign dataset [34] and 3400 traffic sign images that we collected. The LISA traffic sign dataset is a well-known benchmark for traffic sign recognition applications. Similar to the LisaCNN [11] model, we use 17 classes from the dataset. The total number of images in all 17 classes is 6966. The fine-tuned classifier achieves 99% accuracy, which matches the results of LisaCNN. For object detectors, we consider Faster R-CNN [45] and YOLO v3 [44]. For Faster R-CNN, we use ResNet-50 and a feature pyramid network [28]
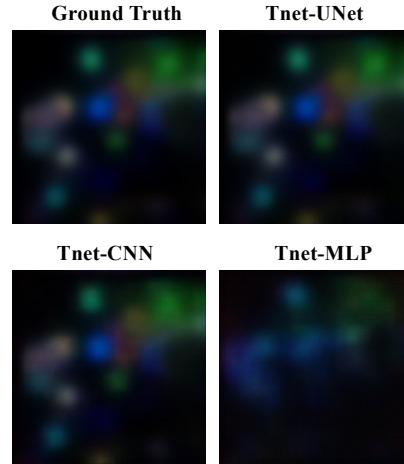
Fig. 5: Visual comparison between the output of all TNet models and the ground truth of the 50-dot case. TNet-UNet output approximates the ground truth perturbation best.

as the backbone. For YOLO v3 we use Darknet-53 as the backbone [44]. Similar to [30], we set the detection threshold of bounding boxes for Faster R-CNN and YOLO v3 to be 0.6 and 0.4 respectively. The input image size is set to be 240x240 for Faster R-CNN and ResNet-50, and 224x224 for YOLO v3. Both of these two detectors are pretrained on the MS COCO dataset [29]. For Faster R-CNN we obtained the pretrained checkpoint from PyTorch [41]. For Yolov3, we download the checkpoint from its official GitHub page[6]. Since the MS COCO dataset only contains stop signs but excludes all of the other traffic signs we otherwise consider, for the detectors, we present attack results only for stop signs.

Evaluation Protocol. We created physical models of all 17 classes of traffic signs that we used to train the recognition module for evaluating our attack. Each model traffic sign is approximately 3cm x 3cm. Since an actual traffic sign is much larger (an actual stop sign is about 76cm x 76cm), we scaled down the distance between the victim camera and the traffic sign accordingly. For instance, 47cm in our setup corresponds to 12m for an actual traffic sign. To collect data for evaluation, we placed a model traffic sign in front of the victim camera as shown in Section 4. We recorded videos while moving the traffic sign to different distances and angles. Similar to [30], we consider a maximum distance of 47cm (which is equivalent to 12m for a full-size traffic sign). For varying the angles, we keep the traffic sign within the field of view of the camera. Each video contains on average 500 frames.

Generating Perturbations. When generating a perturbation using the dot-based perturbation generating function (Section IV-B), we set the number of dots n to be 100, the radius of the k^{th} dot r_k to be $1/10$ of the perturbation's height, the alpha blending value α_{max} to be 1, and the alpha dropoff β to be 1. We define the attack objective for the traffic sign classifier to minimize the reciprocal of the cross-entropy loss with respect to the true label. For detectors, the attack objective
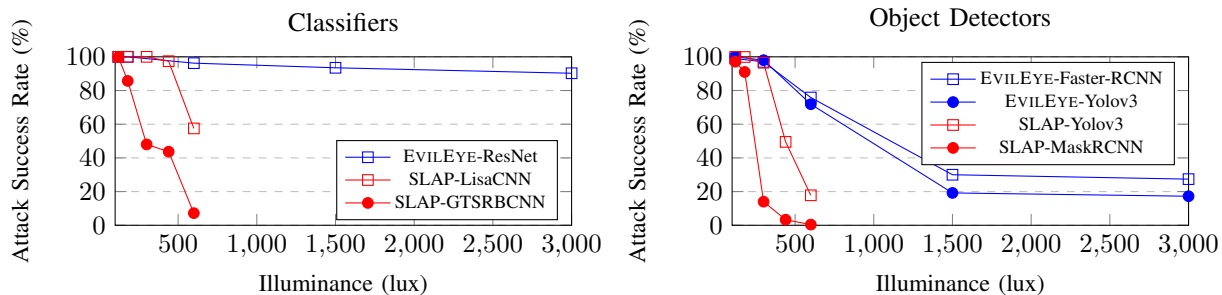
Fig. 6: EVILEYE indoor attack results of classifiers and object detectors. EVILEYE works under various light conditions and can survive significantly stronger ambient light comparing to prior work.

is to minimize the average class score of all the bounding boxes with respect to the class under attack. Our training dataset for generating the perturbations contains $1,000$ images for each type of traffic sign. We use a batch size of 16. The initial learning rate for the dot centers c and color γ are 1 and 0.1 respectively. Each batch is divided by 10 every 200 epochs. We train each perturbation for 500 epochs but we observe the optimization usually converges much earlier.

Indoor Results. For indoor experiments, we evaluated the adversarial perturbations under various illuminance values (produced by moving the floodlight) ranging from 120 lux to 3000 lux. We define the attack success rate (ASR) to be the rate of misclassified frames across all recorded video frames with the target object. For the traffic sign classifier, this corresponds to the rate of misclassification into any other class. For the traffic sign detector, it corresponds to any missed detections. Figure 6 shows the results of our indoor experiments. For a given model, we report the average ASR over all the traffic signs across different illuminance values. As can be seen from the figures, our attack is effective in a wide range of environmental lighting conditions, including much stronger ambient light when compared to prior work. We attribute this success to our hardware and software co-designed attack approach, which enables a strong perturbation to be computed and precisely imposed on the victim camera. Because our attack is sensor-based, i.e., the adversarial perturbation is placed in close proximity to the victim camera, the ambient light has less of an impact on the perturbation performance. Although we observe that the ASR for object detectors degrades in strong ambient light, we note that fooling an object detector is more difficult than fooling a classifier since the adversarial perturbation needs to account for both the bounding box and the class of the object in the bounding box. Specifically, only the pixel mutations inside and around the bounding box of the object can effectively alter the final prediction of the model.

Outdoor Results. Outdoor experiments allow us to test our attack in a more realistic environment as well as in stronger natural light. To perform outdoor experiments, we transported the entire experimental setup to an outdoor road. We measured the illuminance using the same lux meter. Figure 7 shows the outdoor attack results. For the classifier, EVILEYE is able to maintain a high ASR across all tested illuminance levels. For

detectors, EVILEYE has nearly perfect attack performance up to 600 lux. This is equivalent to an overcast afternoon with some overhead cover. As illuminance increases, EVILEYE is still effective up to 3,000 lux, which is equivalent to noon on an overcast day. In general, EVILEYE can survive much stronger ambient light compared to prior light-based attacks [30], which were tested to be effective only up to 120 lux. To push the limit of EVILEYE in terms of illuminance, we exposed the setup to direct sunlight. In the most extreme case, we pointed the victim camera towards the sun. The measured illuminance was 60,000 lux. To reduce the amount of light traveling into the camera lens, we utilized a neutral-density (ND) filter in front of the transparent display. The attack results are shown in Figure 8. From the results with the ND filter, we found that EVILEYE can work under direct sunlight.
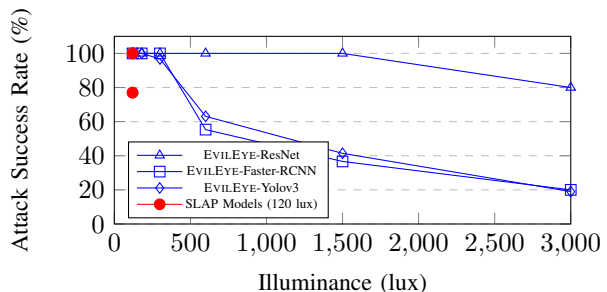

Fig. 7: EVILEYE outdoor attack results. The object classification ASR remains high, while the object detection ASR drops off at higher illuminance values.

Comparison to Static Approach. As mentioned in Section IV, one major advantage of our attack approach compared to prior work is that our attack is dynamic. The adversary can change the displayed perturbation based on different situations. For example, the adversary can change the perturbations for different traffic signs based on the location of the victim vehicle to maximize the attack damage. For static approaches, once a perturbation is attached to the victim camera, the attack cannot be changed or removed. To evaluate the efficacy of our approach, we compared the following two cases: 1) dynamic attack: an adversarial perturbation is crafted for each traffic sign and can be changed dynamically during the attack phase; and 2) static attack: an adversarial perturbation is crafted to maximize the misclassification rate of all types of traffic signs. To prepare such a perturbation, we perform the same
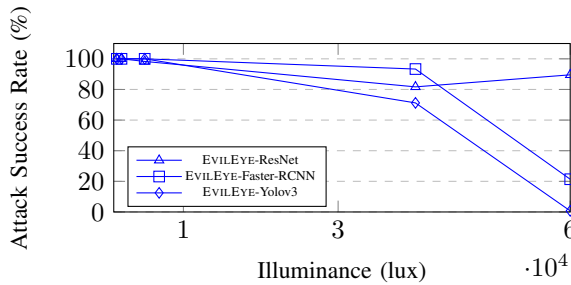
Fig. 8: Improving the EVILEYE outdoor results with the neutral density (ND) filter that attenuates the impact of illuminance. EVILEYE can work in direct sunlight (40k to 60k lux).

optimization process with the same attack objective: maximize the loss with respect to the true label. However, instead of using images from a single class (e.g., the stop sign), we use images from all the classes. We randomly sample 1,000 routes from a traffic sign map[7]. Each route contains approximately 100 traffic signs. We then test how effective the perturbations in each case are on these traffic signs. For the dynamic attack, the adversary is able to locate the vehicle through GPS and has access to a traffic sign map. For each encountered traffic sign, the perturbation specifically designed for it can be displayed accordingly. For the static attack, since there is only one single perturbation, it will be used for all the traffic signs. We report the average ASR over all the traffic signs in a route and all the routes for both cases in Figure 9. Attempting to use one single perturbation for all the classes results in degraded attack performance because designing a single perturbation to accommodate too many cases is more challenging. EVILEYE has the ability to change the perturbation dynamically so each perturbation can focus on its own target class.

Transferability. We also investigate the transferability of EVILEYE across different models–in addition to the 3 models we already evaluated. We also consider Google Vision[8], a commercial pre-trained Image content recognizer. We queried the Google Vision API with our perturbed images. We define the ASR in this case as the percentage of queried images that do not return "Stop Sign" in their top predictions. We report the cross-model ASR in Figure III. According to the results, perturbations are more transferable between Faster R-CNN and YOLO v3, while not as transferable with ResNet-50. This is because Faster R-CNN and YOLO v3 are both object detectors and are both trained on the same dataset. This makes it easier for AEs to transfer between them. On the contrary, ResNet-50 is performing a different task (i.e., only classification) and trained on a different dataset. Finally, Google Vision is completely unable to defend against EVILEYE. Results are shown in Table III.

[7]https://help.mapillary.com/hc/en-us/articles/360003021432-Exploring-traffic-signs

[8]https://cloud.google.com/vision

TABLE III: Transferability Results. We test generated perturbations across different model architectures. We additionally test transferability on the Google Vision Classifier API, achieving 100% ASR in all cases. (*) indicates the usage of a neutral-density filter (approx. 90% light reduction).

| lux	Source Model	Target Model (ASR)			
		ResNet-50	YOLO v3	Faster R-CNN	Google Vision
120	ResNet-50	-	68.01%	98.32%	100%
	YOLO v3	0%	-	100%	100%
	Faster R-CNN	0%	95.2%	-	100%
600	ResNet-50	-	13.77%	33.33%	100%
	YOLO v3	0%	-	86.61%	100%
	Faster R-CNN	0%	32.23%	-	100%
1400*	ResNet-50	-	100%	100%	100%
	YOLO v3	0%	-	100%	100%
	Faster R-CNN	0%	92.22%	-	100%
5000*	ResNet-50	-	100%	100%	100%
	YOLO v3	0%	-	100%	100%
	Faster R-CNN	0%	100%	-	100%
40000*	ResNet-50	-	26.77%	45.07%	100%
	YOLO v3	0%	-	63.31	100%
	Faster R-CNN	3.33%	82.23%	-	100%
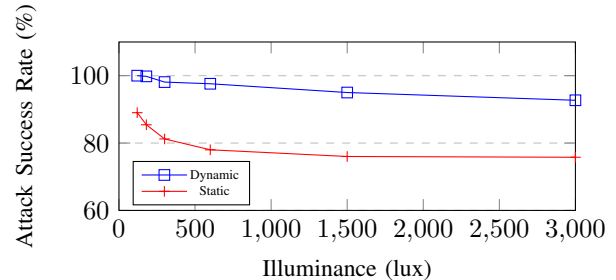
Fig. 9: Comparing EVILEYE to a static approach. EVILEYE's dynamic approach switch perturbations for different type of traffic signs while the static approach crafts one single perturbation to fool all the traffic signs. EVILEYE's dynamic approach gives significantly better ASR.

D. Evaluation Against Existing Adversarial ML Defenses

In this section, we evaluate our physical adversarial perturbations against existing ML defenses. We consider two main criteria in choosing defenses to evaluate. The first criteria is non-specificity, as defenses targeting specific attack signatures or features are less likely to be successful against attacks they weren't designed for. The second defense criteria is low computational overhead, allowing it to be used in real-time applications. As such, we evaluate these defenses on traffic sign identification. Specifically, we test our adversarial perturbations against SentiNet [7] – a defense designed to detect physical AEs (e.g., physically-placed patches or stickers) – as well as feature-squeezing [66] and input randomization [65], two adversarial defenses which transform classifier inputs and meet the criteria of generality and low overhead. Addition-

TABLE IV: ASR of our AEs evaluated against existing defense solutions under different light levels. We present averaged runs for Input Randomization, and only the best defensive result for Feature Squeezing (full results can be seen in Figure 11).

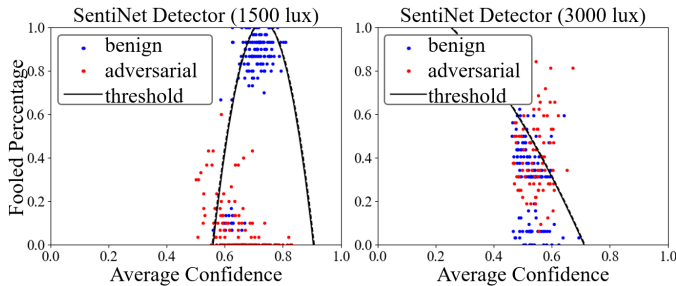| Model	Ambient Light (lux)	Baseline ASR	Input Randomization	SentiNet	Feature Squeezing	Physical Adv. Training
Resnet-50	120	100%	98.5%	100%	94.1%	100%
	180	100%	98.3%	100%	59.2%	100%
	300	100%	90.5%	100%	84.1%	100%
	600	100%	95.8%	100%	89.9%	75.3%
	1500	92.7%	92.2%	87.0%	28.3%	26.1%
	3000	88.7%	91.4%	67.5%	86.8%	34.3%

Fig. 10: A visualization of SentiNet attack detection at different lux levels. Adversarial samples fall within the distribution of benign samples in terms of Average Classifier Confidence and Fool Percentage when salient areas are masked by SentiNet, and are therefore difficult to detect.
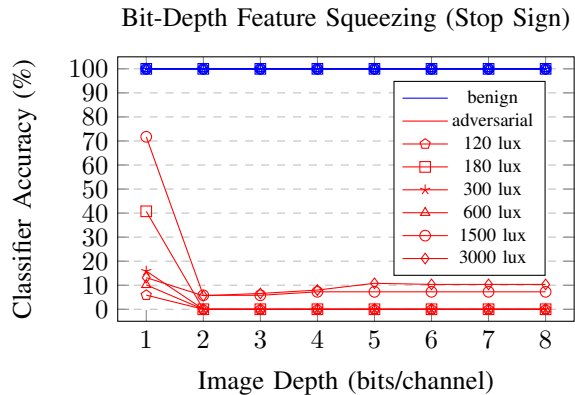

Fig. 11: A comparison of classifier accuracy for benign and adversarial samples when varying image bit-depth of the RGB channels, ranging from native 8-bit color resolution (256 possible values/channel) to 1-bit (2 possible values/channel). Color represents benign (blue) and adversarial (red) performance. Markers indicate different lux levels.

ally, we evaluate adversarial training [13] as a general and preemptive defense simulating a knowledgeable or adaptive defender. A summary of these results can be found in Table IV. In general, the trend of increasing brightness suppresses the effectiveness of EVILEYE, resulting in decreased ASR as lux levels increase.

SentiNet. SentiNet [30] detects adversarial attacks using a comparative saliency mechanism. Using Classifier Confidence and classifier Fool Rate (%) to differentiate between benign and adversarial salient regions over several datasets, SentiNet can detect adversarial patches and perturbations. Results in Figure 10 show a low detection rate, with the adversarial samples' distribution heavily overlapping and falling below the attack detection threshold polynomial fit by the behavior of the benign samples. We suspect that our adversarial perturbations are more difficult to detect using SentiNet because they are less centralized than the original patches studied. This is suggested by the low Fool Percentage of our AEs – individually, a single dot or localized dot cluster is not salient enough to produce adversarial behavior.

Feature-Squeezing. Feature squeezing reduces the dimensionality of classifier inputs, restricting the space for AEs. This technique offers a simple and general strategy to attenuate both physical-domain and digital-domain attacks. We evaluate the effectiveness of bit-depth reduction on our adversarial perturbations. It can be seen that feature squeezing works only at low bit-depth.

Input Randomization. Input randomization modifies the classifier input during inference with a random re-scaling and padding in order to disrupt possible attacks which may rely on precise positioning. We rescale input images randomly between 224-240 pixels from their original size of 224, then randomly pad each edge to reach the final size of 240. We present averaged results over several runs with different random initializations, however, we note that the variance was less than 10% among each group. We find that input randomization leads to only minor changes in the effectiveness of our adversarial perturbations.

Adversarial Training. Wu *et al.* [63] propose an occlusion-based adversarial training process to increase classifier robustness against physically-realizable attacks. This method of *physical* adversarial training was designed to combat patch-based and sticker-based attacks [48, 11] with the idea that the characteristics of digital and physical attacks differ greatly, and therefore conventional adversarial training is not well-suited to physical attacks. We find physical adversarial training to be a promising method for partially mitigating our attack, especially at higher lux levels. We hypothesize that the higher effectiveness is hindered by the slight mismatch in threat models, as physical adversarial training focuses on "occlusive" perturbations like stickers and patches, whereas light-based perturbations are additive and do not occlude.

VI. RELATED WORK

Since the concept of AEs was initially proposed [51], many techniques have been put forward to generate AEs against state-of-the-art machine learning classifiers in the digital space [60], where attackers have precise control over adversarial modifications. In this paper, we focus on AEs in the physical domain.

Physical Adversarial Examples. Kurakin et al. [24] study how digital-domain attacks perform in the physical world using photographs of printed-out AEs, finding that they can survive certain transformations like sensor noise, rescaling, and brightness changes. However, Lu et al. [31] showed that the robustness of adversarial examples generated using this approach is decreased over multiple viewings from different angles and distances. In order to generate more robust AEs in the real world, Athalye et al. [2] propose Expectation over Transformation, optimizing AEs over a set of transformations to account for varying physical environmental conditions by projecting 2-dimensional perturbations onto 3D-printed objects. [21] also, consider a set of environmental transformations when fabricating malicious traffic signs. In all

cases, these approaches focused on static, object-modifying attacks.

Many works create physical AEs by applying adversarial patterns directly to objects themselves. For example, Sharif et al. [48] print adversarial patterns on eyeglass frames to fool facial recognition, and Eykholt et al. [11] place stickers on street signs mimicking graffiti to cause misclassification. Another set of approaches [53, 5, 69] generate adversarial patches. Li et al. [26] demonstrate an adversarial sticker not limited to a single object by applying a translucent sticker in front of the camera itself. Similarly, a key constraint of all these attacks is that they are object-modifying, static attacks, and cannot adapt or be modified once applied.

Other works use light-based perturbations projected onto surfaces in facial recognition attacks [38, 70] and object detection and classification [30, 39]. Projecting AEs onto objects has the benefits of being more dynamic and not leaving behind physical artifacts. However, they suffer from increased sensitivity to environmental lighting conditions and are still difficult to scale to multiple objects. The idea of using light projection to attack has also been applied to the camera sensors. Exploiting ghost effect and auto-exposure control in optical imaging systems, Man et al. [32] use a projector to inject arbitrary patterns (e.g. a stop sign) into the victim camera's field of view. This attack suffers from increased sensitivity to environmental lighting conditions as well. Also, it is not easy to apply the attack in a real-world setting (e.g. a moving camera on an autonomous vehicle). Wang et al. [59] utilized infrared (IR) lights to attack autonomous vehicles. However, IR lights can be strongly interfered by solar radiation. This significantly limits the capability of the attack. Along the direction of attacking the vulnerabilities in sensors, Kohler et al. [23, 47] exploited the rolling shutter in CMOS image sensors using a bright, modulated light source to cause image disruptions. Ji et al. [20] exploited the inertial sensors meant for image stabilization. In their attack, an adversary controls the output of an inertial sensor by emitting deliberately designed acoustic signals. This produces blurred images that will be misclassified in the decision-making pipeline.

VII. LIMITATIONS AND FUTURE WORK

In this section, we discuss the limitations of EVILEYE and enumerate future research directions toward bridging the gap between digital domain- and physical domain-based adversarial machine learning.

Practicality of transparent display implementation. The prototype implementation of EVILEYE's transparent provided a low-cost and portable implementation to enable proof-of-concept research for sensor-level, dynamic adversarial attacks. Emerging technologies such as the Microsoft HoloLens [40] or the Google Glass [36] have shown the industry trend to develop compact and computationally efficient implementations of transparent displays. Future work could also incorporate advancements in attenuating the impact of ambient light on perception, especially in the context of outdoor displays [25, 10]. Thus, the emergence of such technologies would enable the adversaries to implement EVILEYE in real-world settings at scale. Also, Itoh et. al [19] provide a comprehensive overview of optical see-through head-mounted displays, comparing designs that tradeoff between various factors such as form factor, size, and light efficiency. Future work can explore a more optimal design across these tradeoffs.

Attack imperceptibility for autonomous systems. A common notion in adversarial examples research is to ensure perturbations are *imperceptible* to humans in the perception loop. For instance, adversarial examples for spam detectors aimed to bypass machine learning classifiers while not looking suspicious to the target human [54]. For autonomous systems, humans are not expected to be in the loop, i.e., they are not expected to monitor the video camera feed that is being fed into the perception pipeline. Even if a human was monitoring the video feed, the perturbations and objects would most likely be fleeting. Thus, imperceptibility in the context of cyber-physical autonomous systems may target humans who are performing post-incident analyses to investigate likely causes for a malfunction. However, certain autonomous systems domains and applications still require imperceptibility of perturbations, e.g., automatic speech recognition attacks on smart home assistants should be imperceptible to humans who are in the same room [43]. Future work can investigate the need for perceptibility across autonomous systems to understand to *whom* and *when* perturbations need to be imperceptible.

Real-world considerations. Several challenges arise when applying EVILEYE to the real-world. First, the impact of the additional display on lighting, real-world impact, such as additional reflections need to be considered. EVILEYE employs the digital-to-physical mapping to deal with these factors. Moreover, the data used for generating the perturbations and the actual video feed frames at runtime can be different, known as the distributional drift problem. One solution to this can be obtaining more comprehensive training dataset. Also, in some scenarios accessing the victim camera can be challenging for EVILEYE . But it remains a big threat to those systems where camera sensors are accessible (e.g. surveillance, authentication, etc.). Also setting changes on the victim system, such as hardware revisions, and software updates can degrade the accuracy of EVILEYE . In this case, iterations of the attack will be needed. Currently, EVILEYE generates perturbations that can only work on a specific camera. it can be beneficial to make the attack transferable across different cameras in future work. Finally, more complex perception systems employ more complicated sensor arrays (e.g., multi-camera, multi-modal) as the perceptual module. Future work can explore camera pipeline attacks on such systems.

VIII. CONCLUSION

This paper introduced EVILEYE, a man-in-the-middle perception attack on safety-critical cyber-physical systems. EVILEYE is the first sensor-first, dynamic adversarial machine learning framework for physical-domain attacks. EVILEYE leverages transparent displays to generate dynamic physical adversarial examples. The digital-to-physical perturbation

pipeline is enabled by modeling the environmental noise due to optical transformations and environmental factors. We show the efficacy of EVILEYE on the real-world use case of traffic sign recognition, demonstrating that EVILEYE can significantly outperform existing attack frameworks across varying levels of ambient illumination, including over 80% ASR at 60,000 lux–whereas prior works can only achieve a similar ASR at less than 120 lux. EVILEYE provides a practical approach to dynamically modeling optical transformations in the context of adversarial machine learning attacks in the real world.

IX. ACKNOWLEDGEMENT

REFERENCES

[1] L. Abdi, F. B. Abdallah, and A. Meddeb, "In-vehicle augmented reality traffic information system: a new type of communication between driver and vehicle," *Procedia Computer Science*, vol. 73, pp. 242–249, 2015.

[2] A. Athalye, L. Engstrom, A. Ilyas, and K. Kwok, "Synthesizing robust adversarial examples," in *International conference on machine learning*. PMLR, 2018, pp. 284–293.

[3] S. A. Bagloee, M. Tavana, M. Asadi, and T. Oliver, "Autonomous vehicles: challenges, opportunities, and future implications for transportation policies," *Journal of modern transportation*, vol. 24, no. 4, pp. 284–303, 2016.

[4] V. A. Banks, K. L. Plant, and N. A. Stanton, "Driver error or designer error: Using the perceptual cycle model to explore the circumstances surrounding the fatal tesla crash on 7th may 2016," *Safety science*, vol. 108, pp. 278–285, 2018.

[5] T. B. Brown, D. Mané, A. Roy, M. Abadi, and J. Gilmer, "Adversarial patch," *arXiv preprint arXiv:1712.09665*, 2017.

[6] S.-T. Chen, C. Cornelius, J. Martin, and D. H. P. Chau, "Shapeshifter: Robust physical adversarial attack on faster r-cnn object detector," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2018, pp. 52–68.

[7] E. Chou, F. Tramer, and G. Pellegrino, "Sentinet: Detecting localized universal attacks against deep learning systems," in *2020 IEEE Security and Privacy Workshops (SPW)*. IEEE, 2020, pp. 48–54.

[8] Y. Deng, X. Zheng, T. Zhang, C. Chen, G. Lou, and M. Kim, "An analysis of adversarial attacks and defenses on autonomous driving models," in *2020 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, 2020, pp. 1–10.

[9] C. Ertler, J. Mislej, T. Ollmann, L. Porzi, G. Neuhold, and Y. Kuang, "The mapillary traffic sign dataset for detection and classification on a global scale," in *European Conference on Computer Vision*. Springer, 2020, pp. 68–84.

[10] L. H. Eun and K. J. Hong, "How bright of luminance is needed for outdoor commercial display?" in *2016 IEEE 6th International Conference on Consumer Electronics-Berlin (ICCE-Berlin)*. IEEE, 2016, pp. 141–144.

[11] K. Eykholt, I. Evtimov, E. Fernandes, B. Li, A. Rahmati, C. Xiao, A. Prakash, T. Kohno, and D. Song, "Robust physical-world attacks on deep learning visual classification," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 1625–1634.

[12] I. Goodfellow, Y. Bengio, and A. Courville, *Deep learning*. MIT press, 2016.

[13] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," *arXiv preprint arXiv:1412.6572*, 2014.

[14] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.

[15] S. Hoory, T. Shapira, A. Shabtai, and Y. Elovici, "Dynamic adversarial patch for evading object detection models," *arXiv preprint arXiv:2010.13070*, 2020.

[16] A. Hore and D. Ziou, "Image quality metrics: Psnr vs. ssim," in *2010 20th international conference on pattern recognition*. IEEE, 2010, pp. 2366–2369.

[17] L. Huang, A. D. Joseph, B. Nelson, B. I. Rubinstein, and J. D. Tygar, "Adversarial machine learning," in *Proceedings of the 4th ACM workshop on Security and artificial intelligence*, 2011, pp. 43–58.

[18] A. Ignatov, J. Patel, and R. Timofte, "Rendering natural camera bokeh effect with deep learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020, pp. 418–419.

[19] Y. Itoh, T. Langlotz, J. Sutton, and A. Plopski, "Towards indistinguishable augmented reality: A survey on optical see-through head-mounted displays," *ACM Computing Surveys (CSUR)*, vol. 54, no. 6, pp. 1–36, 2021.

[20] X. Ji, Y. Cheng, Y. Zhang, K. Wang, C. Yan, W. Xu, and K. Fu, "Poltergeist: Acoustic adversarial machine learning against cameras and computer vision," in *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2021, pp. 160–175.

[21] W. Jia, Z. Lu, H. Zhang, Z. Liu, J. Wang, and G. Qu, "Fooling the eyes of autonomous vehicles: Robust phys-

ical adversarial examples against traffic sign recognition systems," *arXiv preprint arXiv:2201.06192*, 2022.

[22] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.

[23] S. Köhler, G. Lovisotto, S. Birnbach, R. Baker, and I. Martinovic, "They see me rollin': Inherent vulnerability of the rolling shutter in cmos image sensors," in *Annual Computer Security Applications Conference*, 2021, pp. 399–413.

[24] A. Kurakin, I. J. Goodfellow, and S. Bengio, "Adversarial examples in the physical world," in *Artificial intelligence safety and security*. Chapman and Hall/CRC, 2018, pp. 99–112.

[25] C. Lanca, A. Teo, A. Vivagandan, H. M. Htoon, R. P. Najjar, D. P. Spiegel, S.-H. Pu, and S.-M. Saw, "The effects of different outdoor environments, sunglasses and hats on light levels: Implications for myopia prevention," *Translational vision science & technology*, vol. 8, no. 4, pp. 7–7, 2019.

[26] J. Li, F. Schmidt, and Z. Kolter, "Adversarial camera stickers: A physical camera-based attack on deep learning systems," in *International Conference on Machine Learning*, 2019, pp. 3896–3904.

[27] X. Li, W. Yi, H.-L. Chi, X. Wang, and A. P. Chan, "A critical review of virtual and augmented reality (vr/ar) applications in construction safety," *Automation in Construction*, vol. 86, pp. 150–162, 2018.

[28] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2117–2125.

[29] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *European conference on computer vision*. Springer, 2014, pp. 740–755.

[30] G. Lovisotto, H. Turner, I. Sluganovic, M. Strohmeier, and I. Martinovic, "Slap: Improving physical adversarial examples with short-lived adversarial perturbations," *arXiv preprint arXiv:2007.04137*, 2020.

[31] J. Lu, H. Sibai, E. Fabry, and D. Forsyth, "No need to worry about adversarial examples in object detection in autonomous vehicles," *arXiv preprint arXiv:1707.03501*, 2017.

[32] Y. Man, M. Li, and R. Gerdes, "{GhostImage}: Remote perception attacks against camera-based image classification systems," in *23rd International Symposium on Research in Attacks, Intrusions and Defenses (RAID 2020)*, 2020, pp. 317–332.

[33] M. Mehdipour Ghazi and H. Kemal Ekenel, "A comprehensive analysis of deep learning based representation for face recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2016, pp. 34–41.

[34] A. Mogelmose, M. M. Trivedi, and T. B. Moeslund, "Vision-based traffic sign detection and analysis for intelligent driver assistance systems: Perspectives and survey," *IEEE Transactions on Intelligent Transportation Systems*, vol. 13, no. 4, pp. 1484–1497, 2012.

[35] S.-M. Moosavi-Dezfooli, A. Fawzi, O. Fawzi, and P. Frossard, "Universal adversarial perturbations," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1765–1773.

[36] O. J. Muensterer, M. Lacher, C. Zoeller, M. Bronstein, and J. Kübler, "Google glass in pediatric surgery: an exploratory study," *International journal of surgery*, vol. 12, no. 4, pp. 281–289, 2014.

[37] B. Nassi, D. Nassi, R. Ben-Netanel, Y. Mirsky, O. Drokin, and Y. Elovici, "Phantom of the adas: Phantom attacks on driver-assistance systems." *IACR Cryptol. ePrint Arch.*, vol. 2020, p. 85, 2020.

[38] D.-L. Nguyen, S. S. Arora, Y. Wu, and H. Yang, "Adversarial light projection attacks on face recognition systems: A feasibility study," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020, pp. 814–815.

[39] N. Nichols and R. Jasper, "Projecting trouble: Light based adversarial attacks on deep learning classifiers," *arXiv preprint arXiv:1810.10337*, 2018.

[40] A. K. Noor, "The hololens revolution," *Mechanical Engineering*, vol. 138, no. 10, pp. 30–35, 2016.

[41] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds. Curran Associates, Inc., 2019, pp. 8024–8035. [Online]. Available: http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf

[42] S. D. Pendleton, H. Andersen, X. Du, X. Shen, M. Meghjani, Y. H. Eng, D. Rus, and M. H. Ang, "Perception, planning, control, and coordination for autonomous vehicles," *Machines*, vol. 5, no. 1, p. 6, 2017.

[43] Y. Qin, N. Carlini, G. Cottrell, I. Goodfellow, and C. Raffel, "Imperceptible, robust, and targeted adversarial examples for automatic speech recognition," in *International conference on machine learning*. PMLR, 2019, pp. 5231–5240.

[44] J. Redmon and A. Farhadi, "Yolov3: An incremental improvement," *arXiv preprint arXiv:1804.02767*, 2018.

[45] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," *Advances in neural information processing systems*, vol. 28, pp. 91–99, 2015.

[46] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp.

234–241.

[47] A. Sayles, A. Hooda, M. Gupta, R. Chatterjee, and E. Fernandes, "Invisible perturbations: Physical adversarial examples exploiting the rolling shutter effect," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 14 666–14 675.

[48] M. Sharif, S. Bhagavatula, L. Bauer, and M. K. Reiter, "Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition," in *Proceedings of the 2016 acm sigsac conference on computer and communications security*, 2016, pp. 1528–1540.

[49] J. Souman, M. van Weperen, J. Hogema, M. Hoedemaeker, F. Westerhuis, A. Stuiver, and D. de Waard, "Human factors guidelines report 2: Driver support systems overview," 2021.

[50] G. Sreenu and M. S. Durai, "Intelligent video surveillance: a review through deep learning techniques for crowd analysis," *Journal of Big Data*, vol. 6, no. 1, pp. 1–27, 2019.

[51] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," *arXiv preprint arXiv:1312.6199*, 2013.

[52] B. H. Thomas, "A survey of visual, mixed, and augmented reality gaming," *Computers in Entertainment (CIE)*, vol. 10, no. 1, pp. 1–33, 2012.

[53] S. Thys, W. Van Ranst, and T. Goedemé, "Fooling automated surveillance cameras: adversarial patches to attack person detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2019, pp. 0–0.

[54] J. Tygar, "Adversarial machine learning," *IEEE Internet Computing*, vol. 15, no. 5, pp. 4–6, 2011.

[55] J. E. Vargas-Munoz, S. Srivastava, D. Tuia, and A. X. Falcao, "Openstreetmap: Challenges and opportunities in machine learning and remote sensing," *IEEE Geoscience and Remote Sensing Magazine*, vol. 9, no. 1, pp. 184–199, 2020.

[56] P. Vávra, J. Roman, P. Zonča, P. Ihnát, M. Němec, J. Kumar, N. Habib, and A. El-Gendi, "Recent development of augmented reality in surgery: a review," *Journal of healthcare engineering*, vol. 2017, 2017.

[57] D. Wang, C. Li, S. Wen, Q.-L. Han, S. Nepal, X. Zhang, and Y. Xiang, "Daedalus: Breaking nonmaximum suppression in object detection via adversarial examples," *IEEE Transactions on Cybernetics*, 2021.

[58] J. Wang, A. Liu, Z. Yin, S. Liu, S. Tang, and X. Liu, "Dual attention suppression attack: Generate adversarial camouflage in physical world," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 8565–8574.

[59] W. Wang, Y. Yao, X. Liu, X. Li, P. Hao, and T. Zhu, "I can see the light: Attacks on autonomous vehicles using invisible lights," in *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, 2021, pp. 1930–1944.

[60] X. Wang, J. Li, X. Kuang, Y.-a. Tan, and J. Li, "The security of machine learning in an adversarial setting: A survey," *Journal of Parallel and Distributed Computing*, vol. 130, pp. 12–23, 2019.

[61] E. Wengrowski, K. J. Dana, M. Gruteser, and N. Mandayam, "Reading between the pixels: Photographic steganography for camera display messaging," in *2017 IEEE International Conference on Computational Photography (ICCP)*. IEEE, 2017, pp. 1–11.

[62] E. Wengrowski, W. Yuan, K. J. Dana, A. Ashok, M. Gruteser, and N. Mandayam, "Optimal radiometric calibration for camera-display communication," in *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2016, pp. 1–10.

[63] T. Wu, L. Tong, and Y. Vorobeychik, "Defending against physically realizable attacks on image classification," in *International Conference on Learning Representations*, 2020. [Online]. Available: https://openreview.net/forum?id=H1xscnEKDr

[64] Z. Wu, S.-N. Lim, L. S. Davis, and T. Goldstein, "Making an invisibility cloak: Real world adversarial attacks on object detectors," in *European Conference on Computer Vision*. Springer, 2020, pp. 1–17.

[65] C. Xie, J. Wang, Z. Zhang, Z. Ren, and A. Yuille, "Mitigating adversarial effects through randomization," *arXiv preprint arXiv:1711.01991*, 2017.

[66] W. Xu, D. Evans, and Y. Qi, "Feature squeezing: Detecting adversarial examples in deep neural networks," *arXiv preprint arXiv:1704.01155*, 2017.

[67] C. Yan, W. Xu, and J. Liu, "Can you trust autonomous vehicles: Contactless attacks against sensors of self-driving vehicle," *DEF CON*, vol. 24, no. 8, p. 109, 2016.

[68] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *CVPR*, 2018.

[69] Y. Zhao, H. Zhu, R. Liang, Q. Shen, S. Zhang, and K. Chen, "Seeing isn't believing: Towards more robust adversarial attack against real world object detectors," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2019, pp. 1989–2004.

[70] Z. Zhou, D. Tang, X. Wang, W. Han, X. Liu, and K. Zhang, "Invisible mask: Practical attacks on face recognition with infrared," *arXiv preprint arXiv:1803.04683*, 2018.

[71] A. Zolfi, M. Kravchik, Y. Elovici, and A. Shabtai, "The translucent patch: A physical and universal attack on object detectors," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 15 232–15 241.