

# Securing (Vision-Based) Autonomous Systems: Taxonomy, Challenges, and Defense Mechanisms Against Adversarial Threats

Alvaro Lopez Pellicer<sup>1\*</sup>, Plamen Angelov<sup>2</sup> and Neeraj Suri<sup>3</sup>

<sup>1,2,3</sup>School of Computing and Communications, Lancaster University,  
InfoLab21, Lancaster, LA1 4WA, UK.

\*Corresponding author(s). E-mail(s): [a.lopezpellicer@lancaster.ac.uk](mailto:a.lopezpellicer@lancaster.ac.uk);  
Contributing authors: [p.angelov@lancaster.ac.uk](mailto:p.angelov@lancaster.ac.uk);  
[neeraj.suri@lancaster.ac.uk](mailto:neeraj.suri@lancaster.ac.uk);

## Abstract

The rapid integration of computer vision into Autonomous Systems (AS) has introduced new vulnerabilities, particularly in the form of adversarial threats capable of manipulating perception and control modules. While multiple surveys have addressed adversarial robustness in deep learning, few have systematically analyzed how these threats manifest across the full stack and life-cycle of AS. This review bridges that gap by presenting a structured synthesis that spans both, foundational vision-centric literature and recent AS-specific advances, with focus on digital and physical threat vectors. We introduce a unified framework mapping adversarial threats across the AS stack and life-cycle, supported by three novel analytical matrices: the *Life-cycle-Attack Matrix* (linking attacks to data, training, and inference stages), the *Stack-Threat Matrix* (localizing vulnerabilities throughout the autonomy stack), and the *Exposure-Impact Matrix* (connecting attack exposure to AI design vulnerabilities and operational consequences). Drawing on these models, we define holistic requirements for effective AS defenses and critically appraise the current landscape of adversarial robustness. Finally, we propose the *AS-ADS* scoring framework to enable comparative assessment of defense methods in terms of their alignment with the practical needs of AS, and outline actionable directions for advancing the robustness of vision-based autonomous systems.

**Keywords:** Artificial Intelligence, Autonomous Systems, Security, Computer Vision, Adversarial Attacks, Adversarial defenses

# 1 Introduction

Autonomous Systems (AS) are rapidly transitioning from research prototypes to mission-critical platforms in transportation, logistics, and robotics (Sheridan; Siciliano and Khatib 2016; 2016). At their core, AS combine high-resolution sensors, fast communication links, complex control software, and deep neural networks to enable autonomous operation in unstructured environments (Bekey; Guizzo 2005; 2011).

A defining trend in modern AS is their deep reliance on computer vision. Vision models, ranging from classic convolutional neural networks (CNNs) (He et al. 2016), real-time detectors such as YOLO (Wang et al. 2023) and RT-DETR (Zhao et al. 2024), to advanced transformers (Oquab et al. 2024) and vision-language models (Xu et al.; Renz et al. 2024; 2024) underpin not only perception but also sensor fusion, semantic mapping, prediction, planning, and even direct actuation. The industry-wide move towards vision-centric and even vision-only paradigms is perhaps best exemplified by Tesla’s Autopilot and Full Self-Driving systems (Tesla, Inc. 2022), which intentionally omit LiDAR and radar in favor of multi-camera, deep learning pipelines for end-to-end environment understanding and control.

While classic non-vision attack vectors such as GPS spoofing (Horton and Ranganathan 2018), CAN-bus injection (Kang et al. 2021), and physical attacks on radar or LiDAR systems (Cao et al.; Kong et al. 2019; 2020) have been extensively studied, and industry best practices for their detection and mitigation are relatively mature, the shift to vision-centric architectures introduces a new class of system-wide vulnerabilities. Years of adversarial machine learning research have shown that even digital imperceptible perturbations to image inputs can induce misclassification and dangerous misinterpretation (Szegedy et al.; Goodfellow et al. 2013; 2014). In the physical world, attacks such as adversarial stickers on traffic signs (Eykholt et al. 2018) or adversarial patches (Zhang et al. 2022) among others demonstrate the persistence and transferability of adversarial threats across architectures and conditions.

Crucially, in vision-centric AS, a single vulnerability in perception rarely remains isolated. Because perception outputs directly feed into planning, prediction, and control with limited or no human oversight, adversarial effects can propagate, be amplified by sensor fusion or trajectory optimization, and ultimately result in system-level failures. This risk is heightened by the industry trend towards closed-loop, end-to-end architectures, where raw vision inputs may directly dictate vehicle or robot behavior.

Unlike static computer vision systems, AS operate in dynamic, multi-agent, and safety-critical environments (Bojarski et al.; Janai et al. 2016; 2020). Attacks can target any phase, from data acquisition and model training to online operation or inter-vehicle communication, and their impact can extend far beyond classification accuracy, undermining safety, trust, and real-world performance in ways rarely captured by static benchmarks.

This review is motivated by the urgent need to understand adversarial vulnerabilities and defenses for vision-centric AS, bridging insights from both foundational adversarial machine learning and the fast-evolving AS-specific literature. By systematically mapping how threats propagate across the AS stack and life-cycle, we clarify real deployment challenges, highlight the limitations of existing approaches, and provide

a unified analytical foundation for evaluating adversarial robustness in AS. Our survey intentionally bridges the gap between the mature vision-centric adversarial ML literature and the recent but fast-growing AS-specific corpus.

## 1.1 Related Work

Several recent surveys have addressed elements of adversarial attacks and defenses, but none provide a life-cycle-integrated, stack-specific analysis tailored to real-world AS. For example, [Badjie et al. \(2024\)](#) present a systematic review of adversarial attacks and countermeasures in image classification models for autonomous driving, with detailed coverage of attack types and proactive/reactive defenses. However, their analysis is limited to perception modules and does not examine attack propagation through planning and control subsystems, nor does it offer a unified threat model for the entire AS life-cycle. [Akhtar et al. \(2021\)](#), a comprehensive review of advances in adversarial attacks and defenses for computer vision is provided, focusing on algorithmic and architectural aspects after 2018. However, their work does not account for the layered structure or operational context of AS, omitting issues such as temporal vulnerability, subsystem coupling, or deployment-specific constraints.

[Deng et al. \(2021\)](#) provide a detailed analysis of different attacks and defenses in the workflow of the autonomous driving system, covering adversarial attacks for various deep learning models and attacks in both physical and cyber contexts. While comprehensive in scope, their survey does not offer a structured framework for evaluating defense strategies across different stages of the AS life-cycle. [Liu et al. \(2021\)](#) examine adversarial attacks and defenses from an interpretation perspective, providing valuable insight into model vulnerability, but focusing less on system-level threats specific to autonomous systems.

[Almutairi and Barnawi \(2023\)](#) present an overview of adversarial attacks, defenses, and frameworks to secure DNNs in smart vehicles, organizing their analysis around security challenges but lacking a cohesive approach to understanding cross-layer vulnerabilities. Similarly, [Khamaiseh et al. \(2022\)](#) provide an extensive survey on adversarial attacks and defense mechanisms for image classification, though their focus remains primarily on algorithmic approaches rather than on the operational contexts of autonomous systems.

[Amirkhani et al. \(2023\)](#) review prominent attack and defense mechanisms for object detection in autonomous vehicles, offering discussions on their strengths and weaknesses, but without addressing the integrated nature of attack surfaces across the entire autonomous vehicle stack. [Boltachev \(2024\)](#) highlights key types of disruptive attacks on autonomous driving models, demonstrating potential threats through experimental validation but not providing a systematic framework for defense evaluation.

[Ibrahim et al. \(2024\)](#) perform a systematic review of adversarial attacks and defenses in autonomous vehicles, prioritizing safety and introducing a taxonomy inspired by SOTIF. However, their focus is on risk scenarios and lacks an analytical framework linking attack surfaces, layered vulnerabilities, and defense evaluation across the AS stack. [Girdhar et al. \(2023\)](#) offer a review centered on cybersecurity in autonomous vehicles, highlighting known attack vectors and defenses but stopping

short of providing an actionable structure for mapping attacks or evaluating defenses in an integrated, system-aware fashion.

Xu et al. (2020) broaden the perspective to attacks and defenses in images, graphs, and text, but their survey remains modality-driven and does not tackle the architectural and temporal challenges unique to AS. The work by Costa et al. (2024) surveys adversarial attacks and defenses across various deep learning architectures, offering a high-level synthesis without focusing on the operational realities, threat models, or deployment constraints of AS. Malik et al. (2024) present a systematic review of adversarial machine learning attacks and defensive controls, but their analysis lacks the specificity required for autonomous systems operating in dynamic environments.

## 1.2 List of Contributions

In contrast, our survey bridges the foundational adversarial machine learning concepts presented in (Akhtar et al.; Xu et al.; Costa et al.; Liu et al.; Amirkhani et al.; Malik et al.; Khamaiseh et al. 2021; 2020; 2024; 2021; 2023; 2024; 2022) and the overly component-specialized AS surveys in (Badjie et al.; Ibrahim et al.; Girdhar et al.; Deng et al.; Almutairi and Barnawi; Boltachev 2024; 2024; 2023; 2021; 2023; 2024) with a holistic, layered systems analysis of AS, organized around **three key contributions**:

1. **Bridging Gaps in Existing Surveys:** While prior reviews often isolate general adversarial ML or AS-specific applications, our work integrates foundational adversarial concepts, vision-based robustness literature, and AS-specific challenges into a unified analytical framework. This enables life-cycle-integrated thinking and supports the development of practical AS defenses.
2. **System-Level Threat Modeling via Analytical Matrices:** We construct three matrices that connect existing adversarial literature to the specific vulnerabilities of AS:
  - The **Life-cycle–Attack Matrix** categorizes threats across the Data, Training, and Inference stages of the AI life-cycle, linking attack types (e.g., poisoning, backdoors, evasion) to stage-specific weaknesses and highlighting temporal exposure windows, (Section 4.1).
  - The **Exposure–Impact Matrix** organizes threats by AI design vulnerabilities (e.g., data hunger, model sensitivity), attack surfaces, and downstream consequences such as sabotage or system misguidance, providing a framework to understand full-system threat pathways in real-world AS deployments, (Section 4.2).
  - The **Stack–Threat Matrix** maps how adversarial attacks impact AS subsystems’ Perception, Planning, and Control layers, demonstrating how vulnerabilities propagate and compound across the stack. We ground our analysis with realistic subsystem scenarios, target models, and operational implications, (Section 4.3).
  - Additionally, we provide a comparative synthesis of both **digital and physical adversarial attacks**, characterizing representative methods in terms of attack type, robustness, and practical implications. This serves as a unified reference

for evaluating attack feasibility and severity in both real-world and simulation contexts, (Section 3).

Rather than serving as abstract taxonomies, these matrices function as actionable threat modeling tools to guide robustness benchmarking and inform future research.

3. **Critical Appraisal and Evaluation of Defense Strategies:** We develop a structured methodology to assess how well existing adversarial defenses meet the unique needs of AS:

- Drawing from the literature and the threat matrices developed in this review, we derive a high-level set of **overall requirements** that adversarial defenses must satisfy to be viable in AS environments. Focusing on real-time constraints, adaptability, interpretability, and efficiency, (Section 5.1).
- We examine the current landscape of defenses targeting **physical-world attacks**, identifying the strengths and limitations of existing approaches and clarifying where critical gaps remain, (Section 5.2).
- We consolidate and simplify prior defense taxonomies, aligning them with AS-specific criteria to enable more meaningful evaluation across mechanism types, (Section 5.3).
- Based on this foundation, we introduce the **Autonomous Systems Adversarial Defense Score (AS-ADS)**, a novel evaluation framework that scores defense methods across four deployment-relevant axes: *real-time capability*, *adaptability to novel threats*, *interpretability*, and *resource efficiency*, (Section 5.4).
- To demonstrate the AS-ADS framework, we evaluate a representative subsample of 30 defense methods; 15 from the general vision adversarial robustness literature, and 15 from AS-specific works, highlighting the trade-offs and readiness of each, (Table 9):

This review, to the best of our knowledge, is the first to systematically bridge foundational adversarial machine learning and AS-specific literature in a holistic, layered systems analysis of Autonomous Systems.

### 1.3 Methodology and Review Protocol

This review implements a structured, reproducible literature survey based on PRISMA 2020 principles, specifically adapted to the context of machine learning and AS. Our goal is to comprehensively synthesize advances in adversarial robustness for vision-based models relevant to AS, bridging both foundational vision-centric theory and recent AS-specific developments.

We included works ranging from foundational studies (dating back to 1988) to the most recent publications available as of May 2025, identified through five major databases: IEEE Xplore, SpringerLink, ACM Digital Library, ScienceDirect, and arXiv (tracks: `cs.CV`, `cs.RO`, `stat.ML`). Search queries combined terms such as “adversarial attack,” “defense,” “autonomous systems,” “dataset,” “computer vision,” “robotics,” “LiDAR,” and related phrases. After deduplication, non-vision and unrelated tracks

were filtered, followed by manual screening of titles and abstracts. Full-text eligibility required methodological clarity, empirical evaluation, and relevance to either adversarial computer vision or AS.

Inclusion criteria were: (i) peer-reviewed venue (CORE A\*/A/B or Scimago Q1–Q3 journal) or high-impact arXiv preprint, (ii) empirical focus on adversarial robustness, and (iii) coverage of vision models, pipelines, or AS-specific systems. Studies outside these domains, lacking empirical grounding, or duplicating prior work were excluded. Flexible inclusion criteria were applied to physical attack/defense and real-world system studies, reflecting their practical significance.

Following this protocol, we included **237 papers** in the final synthesis. Each was classified in a reproducible two-level taxonomy: (1) *Domain* (vision-centric or AS-specific), and (2) *Contribution Type* (defense, attack, dataset, or other supportive/background). Within each domain, references were further split as *foundational* (pre-2020) or *non-foundational* (2020 onward). Contribution types were assigned using a combination of keyword analysis (title/abstract), citation context (appearance in attack or defense tables/sections), and manual review for ambiguous cases. The domain split (vision-centric vs AS-specific) was established via systematic keyword matching and manual inspection for works with cross-domain relevance. While every effort was made to ensure comprehensive and reproducible coverage, we acknowledge the potential for misclassification in ambiguous cases and invite community feedback for future updates.

Initial records identified	1041
Duplicates removed	99
Titles and abstracts screened	942
Excluded during abstract screening	614
Full-text articles assessed	328
Excluded after full-text review	91
<b>Studies included in the final synthesis</b>	<b>237</b>

**Table 1** Summary of the PRISMA screening results.

The review process and screening outcomes are summarized in Table 1.

**Table 2** Breakdown of included papers by domain (vision-centric or AS-specific), era (foundational or recent), and contribution type (defense, attack, dataset, other). Percentages reflect the share of each row total.

Domain	Era	Defense	Attack	Dataset	Other	Row Total
Vision-centric	Foundational (pre-2020)	39 (44.8%)	28 (32.2%)	3 (3.4%)	17 (19.5%)	87
Vision-centric	Non-foundational (2020+)	43 (60.6%)	21 (29.6%)	3 (4.2%)	4 (5.6%)	71
AS-specific	Foundational (pre-2020)	1 (6.7%)	2 (13.3%)	0 (0.0%)	12 (80.0%)	15
AS-specific	Non-foundational (2020+)	32 (50.0%)	17 (26.6%)	4 (6.3%)	11 (17.2%)	64
<b>Column Totals</b>		115 (48.5%)	68 (28.7%)	10 (4.2%)	44 (18.6%)	<b>237</b>

**Defense:** Proposes, benchmarks, or surveys robustness mechanisms.

**Attack:** Proposes, benchmarks, or surveys adversarial threats.

**Dataset:** Introduces or is primarily a dataset/benchmark paper.

**Other:** Surveys, theoretical, sensor, or general background works.

Table 2 summarizes the final distribution of included studies by domain, era, and contribution type, supporting full reproducibility and transparency.

## 2 Background

Understanding adversarial robustness in AS requires grounding in the specific architectures, vision model deployments, and operational realities that distinguish AS from conventional computer vision systems. In practice, modern AS tightly integrate vision models not only for perception, but also across sensor fusion, prediction, planning, and closed-loop control, resulting in complex pathways for attack propagation and defense. The threat landscape in AS is shaped by this interconnectedness, exposing weaknesses that are rarely visible in static, perception-only or digital-only evaluations. The limitations of current benchmarks and defense taxonomies, (most of which are tailored to standard image tasks), underscore the need for analysis methods and robustness criteria explicitly aligned with AS operational stacks and environment. This section provides the technical foundations, empirical context, and critical gaps necessary for our analysis.

### 2.1 Vision Models & The Autonomous System Stack

Modern AS are fundamentally vision-driven, with deep learning models tightly integrated across nearly every functional layer; from perception to planning, control, and actuation. Unlike traditional computer vision pipelines, where outputs often remain within isolated modules, AS architectures are defined by close interconnection: the output of one model (e.g., object detection, segmentation) serves as direct input to downstream planning and control components, with minimal human oversight or redundancy.

The AS stack can be broadly divided into three groups: the **Physical Environment**, the **Hardware Layer**, and the **Hardware and Software Integration** layer, as shown in Fig. 1. The physical environment refers to the operational context, such as roadways for driverless vehicles or warehouse floors for robots. In the hardware layer we find sensors such as cameras (Forsyth and Ponce; Szeliski 2011; 2022), LiDAR (Besl; Hsu 1988; 2002), radar (Knee; Hao et al. 2005; 2002), and ultrasonic sensors (Kinsler et al. 2000), which are often fused for greater robustness (Yeong et al. 2021) (sensor fusion). Communication hardware enables inter-device connectivity for federated learning (Yang et al. 2021), remote operations (Yu et al. 2021), or mission planning via satellite links (Prevot et al. 2016). Actuators close the hardware loop by translating digital commands into real-world action.

Across all layers, the adoption of general-purpose vision models, such as ResNet-50 (He et al. 2016), ViT (Dosovitskiy et al. 2020), SAM (Kirillov et al. 2023), and DINOv2 (Oquab et al. 2024), reflects the field’s inheritance of both the strengths and adversarial vulnerabilities discovered in conventional computer vision. Specialized models (e.g., DriveVLM (Tian et al. 2024), CarLLaVA (Renz et al. 2024), BEVFormer (Li et al. 2022)) further illustrate the trend toward unified, stack-spanning pipelines.



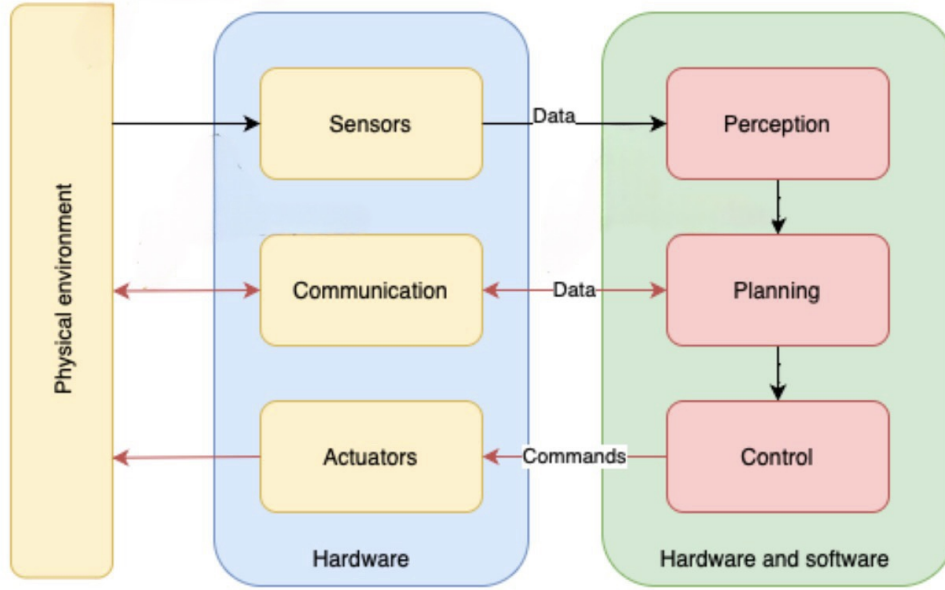


Fig. 1 Autonomous System Stack Diagram

More in depth, the AS **Perception layer** is now dominated by a broad spectrum of vision models. Ranging from early CNN backbones like ResNet-50 (He et al. 2016) to advanced architectures for detection and segmentation. Real-time detectors, such as YOLOv4 (Bochkovskiy et al. 2020), YOLOv7 (Wang et al. 2023), RT-DETR (Zhao et al. 2024), and EfficientDet (Tan et al. 2020), enable high-throughput object and obstacle identification. For segmentation and spatial reasoning, models like DeepLabv3+ (Chen et al. 2018), Mask R-CNN (He et al. 2017), and SAM (Kirillov et al. 2023) provide fine-grained environmental parsing, while ViT (Dosovitskiy et al. 2020) and DINOv2 (Oquab et al. 2024) represent the adoption of transformer-based and foundation models. Multi-modal sensor fusion architectures—DAIR-V2X (Zhao et al. 2024), UMoE (Lou et al. 2023), COMPASS (Ma et al. 2022) integrate camera, LiDAR, and other modalities for richer world models. Classical two-stage detectors like Fast R-CNN (Girshick 2015), Faster R-CNN (Ren et al. 2015), SSD (Liu et al. 2016), and RetinaNet (Lin et al. 2017) also persist in specific AS deployments.

Within the **Planning layer**, outputs from perception are translated into actionable decisions and trajectories using a new wave of context-aware models. BEVFormer (Li et al. 2022) performs multi-view, spatiotemporal fusion for 3D scene understanding. Vision-language models such as DriveVLM (Tian et al. 2024), CarLLaVA (Renz et al. 2024), and VLM-AD (Xu et al. 2024) incorporate semantic context and agent interaction for robust closed-loop planning. End-to-end pipelines such as DAVE2 (Bojarski et al. 2016), PilotNet (Bojarski et al. 2017), and Conditional Imitation Learning (Codevilla et al. 2018) map visual or multimodal input directly to navigation actions, bypassing rule-based intermediaries. Legacy approaches such as



ChauffeurNet (Bansal et al. 2018) and ALVINN (Pomerleau and A 1988) laid the groundwork for behavior prediction and direct perception-control mapping.

At the **Control layer** AS increasingly embed neural controllers, building upon foundations like ALVINN (Pomerleau and A 1988) towards deep reinforcement and imitation learning models (Lillicrap et al.; Pan et al.; Dursun et al. 2015; 2017; 2025), to execute planned actions in real time. These controllers handle adaptive actuation, closed-loop correction, and safe responses to unstructured or adversarial environments. Classic rule-based and PID controllers are now frequently augmented or replaced by neural networks that leverage features from vision and planning models for fine-grained actuation, error recovery, and robust operation under uncertainty. This integration enables rapid, flexible adjustment, but also exposes the system to error propagation: a perturbation at perception or planning can now directly alter low-level control, amplifying the risk of system-level failures.

Because the AS stack is tightly coupled and feedback-driven, whether at the sensor interface, within fusion modules, or at the control output, vulnerabilities in vision models cannot be isolated locally. Perturbations at any point in the stack can cascade through planning and control, ultimately triggering unexpected or catastrophic outcomes. This architecture demands adversarial robustness methods that are not only perception-aware, but explicitly stack and life-cycle-aware as well. A central principle developed throughout this review.

## 2.2 Adversarial Threats in Autonomous Systems

The concept of *adversarial examples* was first introduced in (Szegedy et al. 2013), who showed that deep learning models can be deceived by carefully crafted, human-imperceptible perturbations to input data. Formally, adversarial attacks seek to modify a given input  $\mathbf{x}_0 \in \mathbb{R}^d$  to a new point  $\mathbf{x} \in \mathbb{R}^d$ , such that  $\mathbf{x}$  is assigned a specific target class by the model, differing from the original prediction. The perturbation  $\delta = \mathbf{x} - \mathbf{x}_0$  is typically constrained to be small in a chosen norm (e.g.,  $\|\delta\|_p < \epsilon$ ) to ensure that  $\mathbf{x}$  remains visually indistinguishable from  $\mathbf{x}_0$  to humans. Methods such as *evasion attacks* employ optimization techniques, including the box-constrained L-BFGS algorithm (Fletcher 2013), to compute minimal perturbations that induce misclassification. Notably, these adversarial examples are often *transferable*. A single perturbation generated for one model can also mislead other deep neural networks—raising serious concerns for the security and reliability of AI systems as originally demonstrated in (Liu et al.; Papernot et al. 2016; 2016).

In the context of AS, **digital attacks** (e.g., FGSM (Goodfellow et al. 2014), PGD (Madry et al. 2019), C&W (Carlini and Wagner 2017b)) remain important, operating at inference or training time to introduce pixel-level perturbations or backdoors (e.g., BadNets (Gu et al. 2017), MetaPoison (Huang et al. 2020)). These attacks, originally evaluated on canonical datasets like ImageNet or CIFAR, have proven highly transferable and can undermine robustness at multiple stages of the AS pipeline.

However, AS face a much broader threat landscape. **Physical attacks**—such as adversarial stickers (Eykholt et al. 2018), patches (Brown et al. 2018), or crafted

objects (Kong et al. 2020)—exploit the perception pipeline by manipulating the environment itself, often defeating digital-only defenses and persisting across sensors, agents, and time.

**Cross-modal and systemic attacks** further challenge AS, targeting their reliance on multiple, distributed sensors and communication channels. Examples include GPS spoofing (Horton and Ranganathan 2018), LiDAR jamming (Cao et al. 2019), CAN bus manipulation (Kang et al. 2021), and attacks on federated learning (Yang et al. 2021), each capable of inducing both local and system-wide failures.

**Cascading and life-cycle-aware threats** are particularly critical. A single successful attack at perception can propagate via sensor fusion, scenario prediction, and control feedback loops, leading to mission-level safety breaches (e.g., semantic DoS (Wan et al. 2022), adversarial planning (Edelkamp 2023)). These systemic vulnerabilities are largely overlooked in standard ML taxonomies.

**Limitations of canonical taxonomies:** Most classical frameworks categorize attacks by knowledge and timing, but largely omit the location layer, specially physical attacks and system-level propagation, reflecting a historical focus on static image classifiers and digital benchmarks. In AS, this omission is critical: physical and cross-modal threats are often the most dangerous, propagating through the stack and undermining safety in ways digital-only frameworks cannot capture. This is further pictured in appendix A, Table A1.

These limitations motivate our evaluation of attacks by location (physical and digital) developed in Section 3, and our life-cycle and stack-aware matrices developed in Section 4, which explicitly integrate both digital and physical threats at each layer and throughout the operational life-cycle of AS.

## 2.3 Defense Mechanisms & Autonomous Systems

Adversarial defense research in AS has evolved rapidly, spanning mechanisms adapted from generic computer vision and those developed specifically for the unique constraints of AS. Defenses are most often categorized as proactive (e.g., adversarial training, regularization, input Pre-Processing, certification), reactive (e.g., detection, denoising, reconstruction), or, as a new category found in this review, unified approaches that integrate multiple strategies and account for the layered nature of AS deployments.

**Proactive defenses** such as adversarial training (Madry et al. 2019) remain foundational, retraining models on adversarial examples to improve robustness. This method, applied to both image and LiDAR-based perception modules (e.g., Lu and Radha (2023) for scaling attacks in KITTI/Waymo scenarios), demonstrates gains under known digital threats. However, these approaches incur high computational cost and generalize poorly to unseen or physical attacks, which often bypass digital adversarial defenses (Rozsa et al.; Chen and Lee 2016; 2021). Additional proactive methods, including regularization (Szegedy et al.; Ross and Doshi-Velez 2013; 2018), model distillation (Hinton et al.; Papernot et al. 2015; 2016), and input Pre-Processing

(denoising, smoothing) (Xie et al.; Liao et al. 2017a; 2018) offer marginal improvements, but often at the cost of clean accuracy or robustness to adaptive adversaries (Li et al.; Lou et al. 2024; 2023).

Model ensembles (Tramèr et al.; Xie et al. 2017; 2017b) have also been explored to increase diversity and resilience, but their increased inference latency and hardware requirements are problematic for real-time AS tasks, limiting on-vehicle deployment (Lu et al.; Zhao et al. 2023; 2024). Certified defenses, including randomized smoothing (Cohen et al.; Zhang et al. 2019; 2022) and formal verification (Gowal et al.; Lecuyer et al. 2018; 2019), offer provable guarantees under certain conditions, yet typically remain restricted to limited model classes and do not extend easily to full-stack or dynamic AS environments.

**Reactive defenses** monitor and respond to attacks at runtime. Detection-based mechanisms, such as those in Among Us (Li et al. 2023) (cooperative AVs) or PhySense (Yu et al. 2024) (physical perturbation detection) use input monitoring or auxiliary detectors to identify adversarial events. While valuable, such approaches can suffer from high false positive rates and are vulnerable to sophisticated, adaptive attacks (Soares et al.; Abdu-Aguye et al. 2022; 2020). Denoising and reconstruction via autoencoders or similar tools (Meng and Chen; Samangouei et al. 2017; 2018) can restore clean inputs, but may introduce harmful delay or information loss—unacceptable in safety-critical AS.

**Unified and stack-aware defenses** are gaining attention as the limitations of layer or mechanism-specific solutions become clear. For instance, UMoE Fusion (Lou et al. 2023) exploits multimodal sensor fusion to mitigate sensor blinding, while SpecGuard (Dash et al. 2024) provides sensor and layer-aware detection against UAV sensor spoofing addressing vulnerabilities beyond the perception layer. PatchCleanser (Xiang et al. 2022) and Segment-and-Complete (Liu et al. 2022) combine certified smoothing with detection to target physical patch attacks. Temporal defenses such as ADAV (Mu 2024) and Time-Travel Defense (Etim and Szefer 2024) incorporate cross-frame and historical consistency, crucial for detecting persistent or stealthy threats in dynamic settings.

Unified defense frameworks, e.g., UniCAD (Pellicer et al. 2024), MixDefense (Du et al. 2018), and UNMASK (Freitas et al. 2020), integrate detection, denoising, and robust classification to provide scalable, adaptive defense pipelines more suitable for realistic AS operation. However, most existing defenses, even those tailored for AS, are evaluated primarily at the perception layer and fail to systematically assess downstream effects on planning, control, or mission-level safety.

The entire taxonomy and surveyed papers can be found in Appendix A, Table A2

## 2.4 Datasets and Benchmarks for AS Robustness

Effective evaluation of adversarial robustness in AS relies on benchmarks that capture both the technical complexity and real-world context in which these systems operate. The evolution of benchmarks in this space has both propelled adversarial machine learning and introduced critical challenges unique to AS contexts. Early breakthroughs in adversarial attacks and defenses were closely tied to canonical

datasets such as MNIST (Lecun et al. 1998), CIFAR-10/100 (Krizhevsky 2009), and ImageNet (Deng et al. 2009). These simple, accessible, and widespread benchmarks enabled the rapid development of fundamental attack algorithms like FGSM and PGD (Goodfellow et al.; Madry et al. 2014; 2019), and laid the foundation for robustness research, including systematic evaluations on corrupted or perturbed variants such as ImageNet-P (Hendrycks et al. 2021), CIFAR-C, and CIFAR-P (Hendrycks and Dietterich 2019).

Despite their foundational role, these datasets are now recognized as insufficient proxies for AS robustness due to their static, digital nature and lack of feedback, temporal dependencies, or sensor diversity. Hendrycks et al. (2021) and Croce et al. (2020) demonstrate that robustness metrics obtained on the traditional benchmarks often overstate real-world safety. Models robust on CIFAR or ImageNet may fail when confronted with the complexities of multi-modal perception, sensor fusion, or dynamic interactions in actual AS deployments. This disconnect is further underscored by simulation-to-reality transfer failures, as documented in (Nesti et al.; Xu et al. 2022; 2022).

To address these limitations, the field has gradually shifted towards more application-driven and AS-oriented datasets. DOTA (Xia et al. 2018) introduced complex aerial scenes and diverse object viewpoints, directly benefiting research in UAV and aerial surveillance. The Mapillary Traffic Sign Dataset (Poggi and Mattoccia 2017) captures traffic sign variation in real-world conditions, serving as a testbed for perception modules in autonomous driving. Such datasets improve environmental fidelity and task relevance but still fall short of providing holistic benchmarks for closed-loop or stack-wide robustness.

Recent advances in simulation environments—such as CARLA-GeAR (Nesti et al. 2022), SafeBench (Xu et al. 2022), and RobustE2E (Jiang et al. 2024)—have enabled holistic, closed-loop evaluation of adversarial threats across the full AS stack. These platforms support the generation of physically realizable attacks (e.g., adversarial patches, sensor spoofing), multi-agent and V2X scenarios (Li et al.; Zhao et al. 2023; 2024), and robust testing under diverse conditions (Lou et al.; Zhang et al. 2023; 2023). Real-world datasets—such as Car Hacking (Kang et al. 2021) and adversarial Google Street View (Etim and Szefer 2024)—offer authentic sensor and actuator traces, though they lack the diversity and control of simulated environments.

Despite this, much adversarial research remains focused on standard vision models, with attacks like C&W (Carlini and Wagner 2017b), AutoAttack (Croce and Hein 2020), and patch-based methods (Brown et al. 2018), and defenses such as randomized smoothing (Cohen et al. 2019), MixDefense (Du et al. 2018), and certified patch segmentation (Zhang et al. 2022), almost exclusively evaluated on datasets like ImageNet or RobustBench (Croce et al. 2020). This leaves a gap in addressing how adversarial effects propagate across perception, planning, and control in realistic AS settings.

AS-specific research is bridging this divide by introducing attacks targeting the full system stack—e.g., physical patching (Eykholt et al.; Li et al. 2018; 2022), LiDAR spoofing (Cao et al. 2019), sensor-fusion breakdowns (Lou et al.; Zhao et al. 2023; 2024), and CAN-bus injection (Khan et al. 2022)—and by leveraging advanced benchmarks and simulation platforms. Concurrently, new defenses emphasize multimodal

anomaly detection (Lou et al. 2023), certified segmentation (Zhang et al. 2022), physical input filtering (Lu and Radha 2023), and robust V2X fusion (Zhao et al. 2024), increasingly targeting end-to-end, stack-aware robustness (Jiang et al. 2024).

A summarized illustration can be found in Appendix A, Table A3.


While this move toward AS-specific realism has enhanced operational relevance, it also fragments the field. Different works use incompatible sensor suites, attack models, scenario generators, and evaluation protocols—as highlighted in recent benchmark studies (Xu et al.; Nesti et al. 2022; 2022). Even subtle differences in simulation parameters or the spatial/temporal configuration of physical attacks can yield markedly divergent robustness evaluations, severely limiting reproducibility and comparability across the literature. Consequently, there is a growing consensus, reflected in recent works (Croce et al.; Xu et al.; Lou et al. 2020; 2022; 2023). That progress depends on unified frameworks and holistic benchmarks: those that can relate algorithmic advances in general adversarial robustness to deployment in AS, and, reciprocally, that enable AS-specific innovations to be evaluated in the context of broader vision robustness objectives.

This persistent fragmentation across datasets, evaluation protocols, and adversarial methodology underscores the need for a unified approach—one that systematically bridges the gap between general computer vision research and the operational requirements of AS. To address this, our review introduces a threat-matrix-driven evaluation strategy (see Sec. 4). The unification is finally brought to fruition in our Critical Appraisal of Defenses in the Context of Autonomous Systems, (see Sec. 5).

### 3 Adversarial Attacks in AS: Digital and Physical Locations

Adversarial attacks in Autonomous Systems can be broadly categorized on the basis of their location into two primary domains: digital and physical. Digital attacks occur within the digital pipeline, targeting input data or communications, while physical attacks exploit real-world environments to manipulate sensory input.

#### 3.1 Digital Attacks



**Fig. 2** Example of Digital Adversarial Attack (FGSM)

Digital adversarial attacks focus on manipulating input data directly in the digital domain to deceive machine learning (ML) models. These attacks are some of the

most extensively studied due to their accessibility and the relative simplicity of generating adversarial perturbations. Common methods include the aforementioned FGSM (Goodfellow et al. 2014), PGD (Madry et al. 2019), DDN ( Jérôme Rony and Luiz G. Hafemann and Luiz S. Oliveira and Ismail Ben Ayed and Robert Sabourin and Eric Granger 2019), or Carlini & Wagner (Carlini and Wagner 2017b) amongst others. Fig. 2 illustrates an example of FGSM.

These attacks differ in optimization strategies (e.g., single-step vs. iterative), misclassification objectives (targeted vs. untargeted), and their perturbation budget (constrained by  $\ell_p$  norms or pixel count). Their applications in AS include not only direct evasion of perception pipelines but also poisoning training datasets, injecting malicious patterns into communication logs, or crafting precursors to physical-world attacks through digital-to-physical transfer.

Despite their digital nature, these attacks pose concrete threats to deployed systems, especially when deployed over OTA updates, V2X communication, or shared ML pipelines. As such, a clear comparative understanding of their effectiveness, stealth, and robustness is vital for evaluating the threat landscape faced by real-world autonomous platforms.

To this end, Table 3 presents a structured quantitative synthesis of the fundamental digital adversarial attacks applicable to AS-related vision models. It summarizes their success rates, perturbation magnitudes, transferability across models, and contextual relevance.

### 3.2 Physical Attacks

Physical adversarial attacks are a type of attack in which an adversary attempts to deceive or mislead a ML approach that relies on data gathered from the environment through the use of physical hardware sensors such as cameras. Physical attacks do so by introducing physical perturbations to its environment or inputs. Physical adversarial attacks can take various forms, such as altering the lighting conditions (Xiao et al. 2018), modifying the appearance of objects in the environment (Oslund et al. 2022), or manipulating the sensors that the autonomous system relies on to perceive the world (Cao et al. 2019). Furthermore, in many cases, attacks may be unnoticeable to humans when placed in the real world as they may be mistaken by decorations, urban art or vandalism and not seen as a bigger threat by humans, which hinders the possibility of manual human intervention to prevent attacks in real time. Physical attacks can be configured both in a white-Box or a Black-box setting with differences in performance based on the attack, and their timing would normally be considered Evasion, although it could be the case that they act as Poisoning attacks in the event that the system being compromised is in the learning stage.

Physical adversarial attacks can be generated by transferring digital adversarial attacks into physical objects as demonstrated in various studies (Kurakin et al.; Athalye et al.; Sharif et al. 2016; 2017; 2016). Different techniques to achieve that shift exist which obtain different levels of attack robustness. However, in the physical environment, attack robustness is challenged by other factors, including natural changes in environment conditions, the attack surface being smaller and more complex due to

Attack Method	Description	Common AS Targets	Success Rate (%)	Perturbation Size	Robustness (Transferability)	Remarks
FGSM (Goodfellow et al. 2014)	Fast one-step gradient sign method. Efficient but weaker.	CNNs in perception pipelines (YOLO, ResNet, MobileNet)	65–85%	$\ell_\infty \leq 0.03$	Low (model-specific)	Computationally fast, non-iterative
I-FGSM / PGD (Kurakin et al.; Madry et al. 2016; 2019)	Iterative FGSM; PGD is a universal first-order attack.	Traffic sign classifiers, camera input stream	90–99%	$\ell_\infty \leq 0.03$	Medium (higher with ensemble)	Standard benchmark for robust training
DDN (Jérôme Rony and Luiz G. Hafemann and Luiz S. Oliveira and Ismail Ben Ayed and Robert Sabourin and Eric Granger 2019)	Minimizes norm directly via decoupled optimization.	Perception tasks (ResNet, EfficientNet)	80–95%	$\ell_2 \approx 0.5$	Medium	Good for precise attack with minimal distortion
C&W (Carlini and Wagner 2017b)	Optimizes distortion with a Lagrangian framework. Very strong.	Sensor fusion, camera input, LiDAR projection classifiers	95–100%	$\ell_2 \approx 0.1$ or lower	High	Slow but stealthy; often bypasses defenses
DeepFool (Moosavi-Dezfooli et al. 2016)	Minimal $\ell_2$ perturbation to cross decision boundary.	AS camera classifiers, edge detectors	85–95%	$\ell_2 \approx 0.01$ – $0.1$	Medium	Produces very imperceptible noise
UAP (Moosavi-Dezfooli et al. 2017)	Image-agnostic perturbations that generalize across inputs.	Scene classification (e.g., road conditions)	80–92%	$\ell_2 \leq 0.3$	High	Transferable to unseen data and models
JSMA (Papernot et al. 2016)	Perturbs salient pixels using gradient-based saliency maps.	AS object detectors	70–90%	Few pixels (< 1%)	Low	High distortion when success is enforced
Square Attack (Andriushchenko et al. 2020)	Score-based black-box attack with local square updates.	On-device perception models	85–95%	$\ell_\infty \leq 0.05$	Medium	Efficient in query-limited settings
SimBA (Guo et al. 2019)	Black-box attack via randomized low-frequency noise directions.	Control layer feature extractors	75–90%	$\ell_2 \leq 0.5$	Medium	Simple and effective in low-query regime
One-Pixel / Few-Pixel (Su et al.; Xiao et al. 2019; 2018)	Changes only one or few pixels. Evasion with minimal footprint.	Simple classifiers (MNIST, GTSRB)	30–70%	1–5 pixels	Very Low	Not robust; poor scalability to complex images
Backdoor (e.g., BadNets) (Gu et al. 2017)	Inserts triggers into training data. Attack triggered only when pattern appears.	Entire AS training pipelines	100% (when triggered)	Trigger patch (0.5–5% area)	High (persistent)	Remains dormant; extremely dangerous in safety-critical AS
MetaPoison (Huang et al. 2020)	Craft poisoned training data to manipulate decision boundaries.	Offline AS model training (perception)	80–95%	Clean-label (stealth)	High	Invisible to defenders; long-term threat

**Table 3** Quantitative overview of digital adversarial attacks targeting Autonomous Systems. **Success Rate (%)** reflects attack effectiveness reported across standard AS-relevant models and datasets. **Perturbation Size** describes typical norm-bound constraints (e.g.,  $\ell_\infty$ ,  $\ell_2$ ) or pixel counts. **Robustness** refers to transferability across models, datasets, and tasks. Metrics are extracted or averaged from controlled benchmarks and attack papers, focusing on vision-based perception pipelines in AS.



it being three dimensional, the background not being alterable, or different camera angles.

In the context of AS, physical adversarial attacks represent a significant hazard, with the potential to compromise system safety and dependability. For instance, autonomous vehicles could be misled into misinterpreting traffic control devices such as stop signs or traffic lights, precipitating a potentially perilous situation. The effectiveness of physical adversarial attacks on object detection systems, pivotal in autonomous vehicles, was demonstrated in a study by Eykholt et al. (2018). The research indicated that a physical evasion attack could be orchestrated by adding minimal perturbations to stop signs, thereby distorting the accurate perception of autonomous vehicles.

There are diverging views within the community regarding the effectiveness of these physical adversarial perturbations. Some studies, such as (Lu et al. 2017), suggest that while these adversarial alterations could lead a deep neural network to misinterpret a stop sign image in a physical environment within a specific range of distances and angles, they are not uniformly successful in duping object detectors across varied distances and viewing angles. However, it should be noted that these experiments were conducted in a simplified setting, involving printed attack signs.

More sophisticated and resilient attack methods have since emerged, capable of handling changes in viewpoint, some of which are further explored in this paper. Moreover, it is suggested that as AS and the various deep learning methodologies underpinning their operation continue to evolve, the nature of attacks will similarly adapt and become more advanced. Therefore, contrary to some researchers who may downplay the potential harm of physical adversarial attacks, these threats are considered critical and warrant urgent attention in order to ensure system integrity and safety. A summary of the main types of physical attacks is displayed at the end of this section in 4.

#### **Adversarial stickers and paintings**

The use of adversarial stickers and paintings for deceiving object detection or image classification in AS has been a topic of study. Specifically, Eykholt et al. (2018) examined their effectiveness on deep learning models used in autonomous vehicles. The method involves placing carefully crafted stickers for target objects into the real world, which can cause misclassification of the object detection system. The authors demonstrated that these stickers could be designed to be virtually imperceptible to humans, but still deceive the object detection system. A visualization of the attack is shown in Fig. 3

To generate the adversarial stickers and paintings, the authors used a modified version of the FSM algorithm. They began by selecting a target label, such as a yield sign or a speed limit sign, and used the FGSM algorithm to generate a small perturbation that would cause the object detection system to misclassify the stop sign as the target label. The authors also used a generative adversarial network (GAN) to train a model that could generate images that looked similar to stop signs but contained the adversarial perturbations, while remaining imperceptible to humans.

The study’s findings suggest that the adversarial stickers succeeded in deceiving numerous cutting-edge deep learning models employed in autonomous vehicles, resulting in potentially perilous circumstances. Importantly, the researchers demonstrated

the transferability of these adversarial stickers across disparate models and camera types. Furthermore, the study investigated the influence of physical factors such as lighting conditions, viewing angles, and distances, on the effectiveness of the adversarial stickers. The effectiveness of the stickers did exhibit variation depending on these factors, but crucially retained effectiveness across a broad spectrum of scenarios.



**Fig. 3** Example of Adversarial Stickers

### Adversarial patches

Adversarial patches refer to intricately crafted patches that can be introduced into an image to misguide object detection systems and cause them to misclassify objects in the scene. Such attacks have been previously used to prevent cameras from detecting humans, as evidenced by the development of T-shirts that are printed with adversarial patches (Wu et al. 2020) or by having people wear the patches themselves (Thys et al. 2019). In addition to this, adversarial patches have also been utilized to evade face recognition systems (Komkov and Petiushko 2021) or to prevent AS from detecting objects in the scene (Du et al. 2022).

Work by Zhang et al. (2022) explores the vulnerability of multi-scale object detection models utilized in UAVs to adversarial patch attacks. The authors, similarly to the way adversarial stickers are generated, employed a modified version of the fast gradient sign method (FGSM) algorithm to generate adversarial patches. They initially trained a deep learning model to create patches that could be incorporated into an image to induce misclassification by the object detection system. The patches were designed to be small and inconspicuous to humans but yet potent in deceiving the object detection system.

The research found that adversarial patches were efficient in deceiving several cutting-edge object detection models employed in UAVs. The authors showed that even when the patches covered less than one percent of the image area, they could still deceive the object detection system. Furthermore, the patches were transferable across different object detection models, making them a potential threat to UAVs that rely on deep learning models for object detection.

The research also scrutinized the impact of the size and location of the adversarial patches on the attack’s effectiveness. The authors found that larger patches and patches placed in more critical areas of the image were more effective in deceiving the object detection system.

It is worth noting that a potential limitation of the study at hand is that the patch experiment results only demonstrate the patch being 2D and placed on top of the image. However, in real-world scenarios, attackers are more likely to use these patches to camouflage objects, such as military vehicles like tanks or fighter jets with

an adversarial patch. Therefore, the use of a 3D adversarial patch may be more realistic in such situations.

To address this limitation, [Toheed et al. \(2022\)](#) proposes a method for conducting physical adversarial attacks on object detection systems using 3D adversarial objects. The authors argue that current adversarial attacks on object detectors mainly rely on 2D adversarial perturbations, which have limited ability to cause misclassification of objects in the real world.

The authors introduce a 3D adversarial object that is designed to be imperceptible to humans but can cause misclassification of objects by the object detector. The 3D object is created using computer-aided design (CAD) software and 3D printing technology. The proposed attack is tested on the YOLOv2 object detection system and the COCO dataset, demonstrating its effectiveness in causing misclassification of objects in the real world.

### **Adversarial objects**

Adversarial objects are crafted in a way that they cause the ML model to misclassify, misinterpret, or fail to recognize them, even though they might appear normal to the human eye. They follow a similar approach to adversarial stickers or patches. However, they differ in that a complete 2D or 3D object is built.

[Kurakin et al. \(2016\)](#) was one of the first to investigate 2D physical adversarial objects, this paper investigates the effectiveness of adversarial examples in real-world settings. The authors focus on the transferability of adversarial examples between digital and physical domains, as well as their robustness to various transformations, such as changes in camera angle and lighting conditions. The authors extend their investigation to the physical world, questioning whether adversarial examples generated in the digital domain can still be effective when captured by a camera and processed by a ML model.

To study this question, the authors generate adversarial examples using FGSM and print them out, simulating a physical-world scenario. They then capture images of these printed adversarial examples using a smartphone camera and feed the captured images to a deep learning model to evaluate the model’s performance.

The experiments show that adversarial examples generated in the digital domain can still be effective in the physical world, causing the ML model to misclassify the printed images. The authors also demonstrate that the adversarial examples are robust to various transformations, such as changes in camera angle, lighting conditions, and resizing of the images. This finding suggests that adversarial examples pose a significant challenge to the deployment of deep learning models in real-world applications, as they can cause the models to make incorrect decisions even under different physical conditions.

More curated and targeted to Autonomous System papers in the 2D object landscape include ([Kong et al.; Zhou et al. 2020; 2020](#)). ([Zhou et al. 2020](#)) presents a systematic approach for generating adversarial billboards designed to compromise object detection models in autonomous driving systems. The authors propose a bi-level optimization framework that considers both the attack’s success probability and the perturbation’s perceptual similarity. They leverage a 3D simulator to account for physical-world factors such as lighting, camera perspective, and occlusion. While this

approach provides valuable insights into the robustness of object detection models under various physical-world scenarios, the use of a 3D simulator may not fully capture the complexity of real-world conditions, potentially limiting the generalizability of the results. [Kong et al. \(2020\)](#) employs a Generative Adversarial Network (GAN) to create adversarial examples resilient to real-world environmental factors. The method comprises a generator network responsible for producing adversarial perturbations and a discriminator network tasked with discerning between real and adversarial examples. To enhance the transferability of the generated adversarial examples, the authors incorporate domain adaptation techniques and apply geometric and photometric transformations during training. While [Kong et al. \(2020\)](#) demonstrates the potential for crafting physical-world-resilient adversarial examples, the adversarial training process can be computationally expensive and sensitive to hyperparameters, which may limit its practical applicability.

[Athalye et al. \(2017\)](#) was one of the first works to introduce 3D adversarial objects. The paper presents a novel approach to generating adversarial examples that are robust to various transformations and are effective in both the digital and physical domains. The authors propose a method called Expectation over Transformation (EOT), which aims to create adversarial examples that maintain their adversarial properties under different transformations.

Traditional adversarial example generation methods often focus on fooling a ML model in the digital domain, without considering the effects of real-world transformations, such as rotations, translations, and changes in lighting. As a result, these adversarial examples may lose their effectiveness when applied to physical objects or real-world scenarios. To address this issue, the authors introduce the EOT algorithm, which incorporates an expectation over a chosen set of transformations during the adversarial example generation process. By optimizing the adversarial perturbation under this expectation, the algorithm ensures that the generated adversarial examples are robust to the specified set of transformations.

The authors evaluated the performance of the EOT algorithm on various state-of-the-art deep learning models, such as Inception v3 and ResNet, using different datasets like ImageNet and CIFAR-10. They also compare the EOT algorithm with other existing methods, such as FGSM and PGD. The results demonstrate that the EOT algorithm is able to generate adversarial examples that are robust to a wide range of transformations, outperforming other methods in both digital and physical domains. The authors further showcased the effectiveness of the EOT algorithm through real-world demonstrations, such as 3D printed objects and images displayed on a screen.

[Cao et al. \(2020\)](#) specifically targets the vulnerabilities of autonomous driving systems to 3D adversarial objects. This paper specifically targets Multi-Sensor Fusion (MSF)-based perception systems used in autonomous vehicles. The authors propose a real-time, end-to-end optimization algorithm that takes into account the physical constraints and sensor characteristics of the MSF-based perception system to generate 3D adversarial objects. By considering the limitations of the sensors and the physical constraints of the objects, the proposed method generates adversarial objects that can deceive the MSF-based perception system in real-world scenarios. The paper evaluates its method using simulation and real-world experiments, focusing on the

effectiveness of the 3D adversarial objects in deceiving MSF-based perception systems in autonomous vehicles.

Table 4 summarizes the main types of physical adversarial attacks, their implications, and key examples along with simple quantitative indicators such as Success Rate or Robustness to further contextualize their relevance in vision models and therefore to AS.

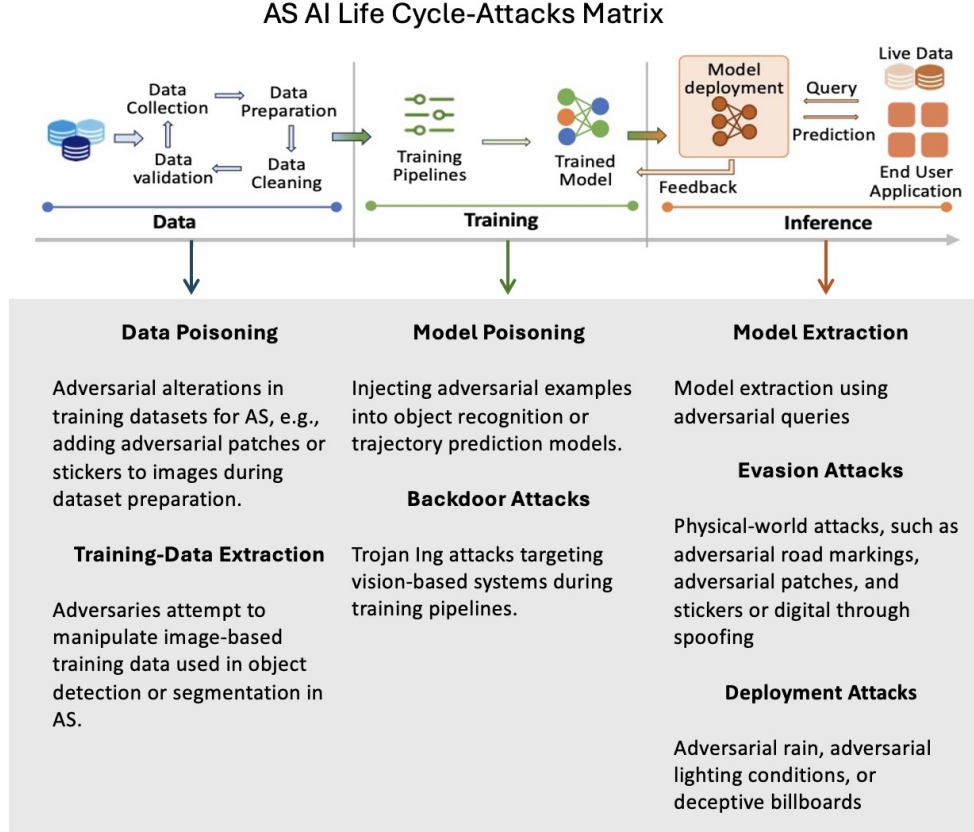
## **4 Threat Modelling in Autonomous Systems**

Attack Type	Target System(s)	Description	Implications	Success Rate (%)	Size	Robustness	Key Studies
Adversarial Stickers	Traffic Sign Recognition, Object Detection	Printed perturbations (e.g., on signs) crafted using FGSM or GANs to mislead perception models	Misclassification of traffic signs; risk of safety-critical errors in AVs	Up to 91.49%	≤5% of object area	Partial (angle/distance sensitive)	(Eykholt et al.; Oslund et al.; Zhu et al.; 2018; 2022; 2024)
Adversarial Patches (2D)	UAV Detection, Person Detection	Small 2D patches embedded in clothing or scenes, optimized to evade detection	Enables human evasion from surveillance or drone systems	Up to 90%	≤1% of image area	Limited	(Thys et al.; Wu et al. 2019; 2020)
Adversarial Patches (3D)	Object Detection (YOLO, SSD)	Physically printed 3D patches placed on real objects (e.g., vehicles)	Persistent misclassification of camouflaged objects	Up to 85%	Object surface dependent	High (real-world tested)	(Toheed et al.; Du et al. 2022; 2023)
Adversarial Objects (2D)	Image Classification	Printed adversarial images misclassified under varied conditions	Demonstrates real-world vulnerability of classifiers	65–85%	Full object	Partial	(Kurakin et al. 2016)
Adversarial Objects (3D)	Object Detection, Multi-Sensor Fusion Systems	Crafted 3D shapes optimized via EOT or end-to-end sensor-aware learning	Compromises multi-sensor fusion in autonomous vehicles	80–90%	Full object	High	(Athalye et al.; Cao et al. 2017; 2020)
Adversarial Billboards	Autonomous Driving Systems	Adversarial large-scale signs created via optimization in 3D simulator	Attacks AS from afar; misguides perception in motion	Approximately 65% misdetected	Full billboard	Medium	(Zhou et al. 2020)
Adversarial Clothing	Person Recognition	T-shirts or jackets with adversarial patterns to evade detection	Enables physical anonymity from AI-based surveillance	57–74%	Clothing-scale	Partial	(Wu et al. 2020)
Adversarial Rain	Object Detection, Classification	Raindrop overlays on lens or images to obstruct vision systems	Misinterpretation of surroundings under weather conditions	60–70% accuracy drop	N/A	Medium	(Guesmi et al. 2023)
Adversarial Lighting	Object Detection	Controlled lighting (e.g., glare/shadow) to cause detection failures	Disrupts feature extraction; breaks perception	Up to 93.7% fooling rate	Global	High (controlled)	(Hsiao et al. 2024)

**Table 4** Comprehensive summary of physical adversarial attacks applicable to Autonomous Systems, integrating both qualitative and quantitative evidence from the literature. **Target System(s)** refers to the machine learning subsystems being attacked (e.g., traffic sign recognizer, object detector). **Success Rate (%)** indicates the reported attack success under physical-world or simulation conditions. **Size** estimates the spatial footprint of the adversarial pattern relative to the object or image surface. **Robustness** denotes the resilience of the attack to changes in viewpoint, lighting, and physical conditions. Metrics are synthesized from experimental results in the cited studies; where multiple results are reported, the maximum or typical observed value is given.

This section presents a comprehensive framework for threat modeling in AS, with a particular focus on vision-based models. We introduce a taxonomy that systematically analyzes the exposure of each stage in the AS life-cycle to adversarial attacks (both digital and physical). By mapping specific attack vectors to corresponding life-cycle components and system layers, this framework provides a structured basis for identifying vulnerabilities and informs the development of effective, targeted defense strategies for real-world AS deployments.

#### 4.1 Life-cycle Attack Matrix



**Fig. 4** AS AI Life-Cycle Attack Matrix

We introduce the *AS AI Life-Cycle Attack Matrix* (see Figure 4), a framework that systematically categorizes adversarial threats targeting AS across the Data, Training, and Inference stages of the AI life-cycle. By mapping attack types to each stage,



the matrix provides a comprehensive structure for identifying vulnerabilities, understanding how adversaries exploit phase-specific weaknesses, and informing the design of more effective defense strategies.

Figure 4 organizes adversarial threats into three main stages of the AI life-cycle: Data, Training, and Inference. Each stage is associated with characteristic attack types that leverage distinct vulnerabilities in AS pipelines.

At the **Data stage**, adversaries may engage in:


- **Data Poisoning Attacks:** Introducing malicious data into the training dataset to corrupt the learning process, leading to erroneous model behavior. For instance, altering traffic sign images to mislead recognition systems in autonomous vehicles (Morgulis et al. 2019).
- **Training-Data Extraction Attacks:** Extracting sensitive information from the training data, potentially compromising privacy and security. This can involve reconstructing proprietary datasets used in AS development (Malik et al. 2024).

During the **Training stage**, potential attacks include:

- **Model Poisoning Attacks:** Manipulating the training process to embed vulnerabilities within the model, which can be exploited during deployment. This includes tampering with the training data or the learning algorithm itself (Almutairi and Barnawi 2023).
- **Backdoor Attacks:** Inserting hidden triggers into the model that cause it to behave maliciously when specific conditions are met. For example, embedding triggers that activate under certain visual patterns encountered by AS (Pourkeshavarz et al. 2024).
- **Attacks in Federated Learning (FL):** Federated learning offers a decentralized approach to training machine learning models across multiple devices, making it particularly relevant for AS applications such as autonomous vehicles. In FL, each client—such as an autonomous vehicle—trains a local model using its own data. Only the model updates are shared with a central server, where they are aggregated to create a global model. This approach not only preserves data privacy but also reduces computational and communication costs by distributing the training process across multiple devices (Jallepalli et al. 2021).

However, FL’s decentralized nature introduces unique security challenges. Malicious actors can exploit the collaborative training process to compromise the global model. For instance, a rogue client might poison its local training data or tamper with model updates, leading to degraded performance or targeted misbehavior. Moreover, FL’s privacy-preserving mechanisms, such as secure aggregation and differential privacy, can make detecting such attacks more difficult, further complicating the task of ensuring robust security. Recent studies, including (Li et al.; Queyrut et al.; Shi et al. 2024; 2023; 2022), provide a comprehensive overview of FL architectures, their adversarial challenges, and potential defense strategies within AS. A simple visualization of attack vectors in a FL architecture is shown in Figure 5.

At the **Inference stage**, AS are susceptible to:



**Fig. 5** Example of Prototype-based FL architecture and attack surface

- **Model Extraction Attacks:** Adversaries query the deployed model to reconstruct its parameters or architecture, facilitating intellectual property theft or enabling further attacks (Malik et al. 2024).
- **Evasion Attacks:** Crafting inputs that are intentionally designed to be misclassified by the model, thereby bypassing security measures. Physical-world examples include adversarial patches or stickers that cause misclassification in object detection systems (Girdhar et al. 2023).
- **Prompt Attacks:** Exploiting prompt-based systems by injecting malicious prompts that alter the model’s behavior or outputs, potentially leading to unintended actions in AS (Shan et al. 2024).
- **Adversarial Deployment Attacks:** Introducing adversarial elements into the environment, such as deceptive road markings or manipulated traffic signs, to mislead the AS perception and decision-making processes (Boltachev 2024).

This taxonomy underscores the multifaceted nature of adversarial threats across the AI life-cycle in Autonomous Systems. By systematically categorizing these attacks, we aim to enhance the understanding and development of robust defense mechanisms tailored to each stage of the AI deployment pipeline.

## 4.2 Exposure–Impact Matrix

The *AS Adversarial Exposure-Impact Matrix*, illustrated in Figure 6, offers a detailed taxonomy of adversarial attack vectors that specifically exploit vulnerabilities in AS. The matrix organizes these vulnerabilities according to fundamental AI challenges, such as the need for large datasets, sensitivity to model updates, similarities across models, and input fragility, linking each to concrete attack surfaces, including data pipelines, model APIs, and environmental inputs.

These vulnerabilities enable a wide spectrum of attacks, ranging from *data poisoning* and *backdoors* during training to *model extraction* and *evasion* at inference. The matrix clarifies both where and how AS can be compromised and traces the downstream consequences from data collection and model preparation through deployment to operational harms such as misguidance, sabotage, or intellectual property (IP) theft.



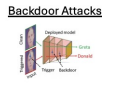
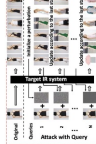
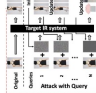
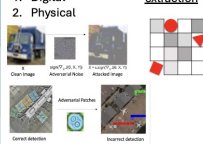
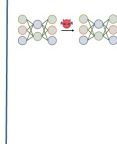
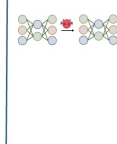

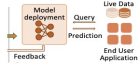
AS Adversarial Exposure Matrix						
AI Inherent vulnerabilities	Data Hunger	Model update data sensitivity	Reverse Engineering prone	Input sensitive		Similarity across models
Attack Surface	Data collection sources: Physical, public or databases	Federated learning or any access to model training pipelines	Model APIs or other access to model predictions	Adversarial Examples	Queries	Surface transfer
Attacks	<u>Data Poisoning</u> 1. Physical environment alterations 2. Spoofing malicious images in training 	<u>Model update poisoning</u>  <u>Backdoor Attacks</u> 	<u>Model Extraction</u>  <u>Attack with Query</u> 	<u>Evasion attack</u> 1. Digital 2. Physical 	<u>Training-Data extraction</u> 	<u>Transfer Attacks</u> 
Life Cycle impact						
Real World Impacts	Misguidance Deception	Manipulation Crashes	Sabotage Compromise	IP theft Exposure	Elusion Crashes & Casualties	Leakage Misdirection Exploitation

Fig. 6 AS Exposure-Impact Matrix

**AI Inherent Vulnerabilities and Attack Surfaces:** AS inherit several critical vulnerabilities from the underlying AI models and datasets on which they rely, exposing multiple attack surfaces:

- **Data Hunger:** The requirement for large and diverse datasets makes AS vulnerable to *data poisoning attacks*, where adversarial modifications—such as altered traffic sign images—are injected into the training data (Eykholt et al. 2018).
- **Model Update Sensitivity:** The adoption of federated learning and access to model update pipelines introduce the risk of *model poisoning attacks*, allowing adversaries to manipulate updates and embed *backdoors* (Cheng et al. 2021).
- **Input Sensitivity:** The inherent fragility of AI models to subtle input changes makes them susceptible to *adversarial examples*, including both digital perturbations and physical attacks (such as stickers or patches on objects) (Brown et al. 2018).
- **Similarity Across Models:** The resemblance between different models allows for *transfer attacks*, where adversarial examples crafted for one model can successfully mislead another (Tramèr et al. 2017).

**Real-World Impacts:** The AS Adversarial Exposure Matrix reveals how the convergence of AI vulnerabilities, attack surfaces, and adversarial tactics results in tangible real-world consequences. By mapping these threats from data collection through training, inference, and deployment, the matrix highlights clear pathways through which Autonomous Systems can be undermined:

Study	Attack Type	Real-World Impact
Dynamic Adversarial Attacks on Autonomous Driving Systems (Chahe et al. 2023)	Physical adversarial patches on moving objects	Misclassification of traffic signs, leading to misguidance and deception
Adversary ML Resilience in Autonomous Driving Through Human-Centered Perception Mechanisms (Shah 2023)	Physical attacks on road signs (e.g., tape, graffiti)	Misclassification, causing safety hazards
Embodied Adversarial Attack: A Dynamic Robust Physical Attack in Autonomous Driving (Wang et al. 2023)	Laser-based dynamic physical attacks	Misinterpretation of the environment, resulting in potential crashes
Beyond Boundaries: A Comprehensive Survey of Transferable Attacks on AI Systems (Wang et al. 2023)	Transfer attacks leveraging model similarities	Scaled exploitation across multiple autonomous systems
Towards Robust and Secure Embodied AI: A Survey on Vulnerabilities and Attacks (Xing et al. 2025)	Adversarial manipulation of AI-controlled robots	Safety-critical failures, including crashes and casualties
Discovering Adversarial Driving Maneuvers Against Autonomous Vehicles (Song et al. 2023)	Adversarial driving maneuvers	System misguidance, crashes, and operational compromise
Efficient Adversarial Attack Strategy Against 3D Object Detection in Autonomous Driving (Chen et al. 2024)	3D object detection manipulation	Misclassification of objects, leading to potential crashes
Adversarial Backdoor Attack on Trajectory Prediction (Pourkeshavarz et al. 2024)	Clean-label data poisoning	Causes systematic errors in path prediction, increasing collision risks

**Table 5** Representative set of attacks and their real-world impacts in Autonomous Systems

- **Data Hunger → Data Poisoning:** The demand for extensive, diverse datasets exposes AS to *data poisoning*, where physical or digital manipulation of training data causes *misguidance* and *deception* at the perception layer.
- **Model Update Sensitivity → Model Poisoning and Backdoor Attacks:** Continuous model refinement in AS creates opportunities for adversaries to introduce *model poisoning* or embed *backdoors* via tainted updates. This results in *manipulation* and *sabotage*, eroding model integrity and reliability.
- **Reverse Engineering Prone → Model Extraction:** When attackers gain access to model outputs through open APIs or similar interfaces, they can perform *model extraction*, leading to *IP theft* and *exposure* of proprietary algorithms. This undermines competitive advantage and may facilitate further adversarial actions.
- **Input Sensitivity → Evasion Attacks and Training-Data Extraction:** Systems that rely on accurate sensor interpretation or user input are vulnerable to *evasion attacks* and adversarial queries. Such *elusion* and environmental *manipulation* can cause *crashes*, *casualties*, and information *leakage*, as the AS fails to interpret its environment correctly.

- **Similarity Across Models → Transfer Attacks:** Exploiting similarities among models, adversaries can launch *transfer attacks* that scale across multiple AS platforms, resulting in widespread *exploitation* and a further erosion of public trust in these technologies.

By mapping each vulnerability and attack type to its downstream impact, the matrix underscores that even subtle technical manipulations can cascade into severe, real-world consequences. Understanding these relationships is crucial for designing robust defense strategies that ensure the reliability, safety, and integrity of Autonomous Systems.

Table 5 consolidates recent research that exemplifies the real-world impacts identified in the AS Adversarial Exposure Matrix. These studies provide concrete evidence of adversarial attacks, their methodologies, and their consequences for AS, emphasizing the need for comprehensive defense mechanisms.

### 4.3 Stack–Threat Matrix

Because AS operate in uncontrolled, open environments, they are especially vulnerable to attacks that target the physical world. *Physical adversarial attacks* are particularly critical, as they directly compromise the perception capabilities of sensors and cameras, thereby undermining all subsequent layers. Nonetheless, vulnerabilities are not limited to physical inputs. Table 6 provides our matrix mapping relevant examples with their scenarios and implications per stack layer. Some more in depth conceptual examples are presented below to further understand the relevance per layer:

At the **the Perception Layer**, attacks can manipulate the sensory input of an AS, causing the system to perceive incorrect or misleading information. Adversarial attacks in computer vision can cause an AS to misclassify objects in the environment, leading to incorrect or unsafe actions (Ai et al.; Wang et al. 2021; 2021).

Tampering with the perception layer often involves that further layers (planning and control) will also be compromised as data flows from one layer to the other, an incorrect view of the environment can lead to, for instance an incorrect route being planned and wrong commands sent to the actuators in the control layer. The scenarios for attacks that target the perception layer involve the exploitation of the area in which camera sensors actuate, in this case the physical environment, thus the threat to be considered are physical adversarial attacks. These include adversarial patches, objects and stickers which have been outlined previously and summarized in Section 3.2.

Attacks in to the perception layer and to other layers can be distinguished based on the attacker objectives, this means that although every successful physical attack involves alterations to the perception of the environment produced at the perception layer, not every physical attack shall be considered a perception layer attack.


Perception layer attacks aim to remove or add elements to the system’s perception of the world, altering its fundamental behavior. If the example of driverless cars is considered, attacks involving adversarial traffic signs (Morgulis et al. 2019) might be more appropriately classified as planning layer attacks rather than perception layer attacks. This is because even if a stop sign is misclassified as a 45 mph speed limit sign, the car will still be able to navigate the road and recognize that a traffic sign

is present. However, its planned route or correct trajectory will be altered due to an unintended decision made at the planning layer. In contrast, an attack involving a pedestrian wearing an adversarial T-shirt (Xu et al. 2020) should be considered a perception layer attack, as it renders an element invisible, preventing the car from accounting for all elements on the road. Therefore, attackers’ aiming purely at the perception layer will normally leverage physical attacks targeting object detectors.

At the **Planning Layer** adversarial attacks can be crafted leveraging vulnerabilities in Deep learning classifiers including physical attack such as adversarial traffic signs as demonstrated by Morgulis et al. (2019). The security implications can include incorrect routes, traffic violations, or accidents. In (Fu et al. 2022), an adaptive adversarial attack on real-time Unmanned Aerial Vehicle (UAV) tracking systems is introduced. The authors devise the Ad2Attack method, a mechanism that produces adversarial examples aimed at deceiving deep learning-powered UAV tracking systems. A successful compromise of the tracking system’s performance can lead to the UAV losing track of its intended target. This loss of tracking can, in turn, result in inaccurate or suboptimal route planning, thus posing significant operational challenges.

Other significant vulnerabilities can be found in planning systems which involve gathering information from external sources to make planning decisions, such as GPS spoofing, an example of such attack can be found in (Horton and Ranganathan 2018), attacks such as this can manipulate the drone’s perceived location and potentially take control of its movements. Although this example is not in the image domain, it is believed that systems may use other information in the planning layer such as saved streetview images downloaded from an external server to aid navigation. Thus, attacks similar to GPS spoofing, where malicious images are injected into the planning layer leveraging wireless technology vulnerabilities, may exist in the future.

For the **Control Layer**, Tian et al. (2022) presents an architecture for an unmanned aerial vehicle (UAV) is described, in which the drone’s camera acts as a sensor and sends real time images to the controller for processing and display through a Wi-Fi network. The controller, which is based on Dronet, processes the image to gain situational awareness of the environment and generates control instructions. These instructions are transmitted to the actuator through the Wi-Fi network to control the drone. Given the vulnerabilities in Wi-Fi networks, there may exist an active attacker who controls the Wi-Fi link and generates imperceptible perturbations (adversarial examples) to images sent by the camera to remain undetected. This attack may result on the drone receiving wrong velocity commands which could make it intentionally crash to an object or even a human, or at least alter its normal course. A illustration of this attack is shown in Fig. 7.



**Fig. 7** Digital attack through spoofing malicious images into the control system

Xie et al. (2017b) explores adversarial attacks on deep learning-based semantic segmentation and object detection systems, both of which play a critical role in the control layer of autonomous vehicles. Through the generation of adversarial examples, these systems can be manipulated, leading to erroneous control decisions with potentially hazardous outcomes such as accidents or system malfunctions.

The researchers present a method for creating adversarial examples that effectively deceive both semantic segmentation and object detection algorithms. The technical backbone of this paper involves the resolution of an optimization problem, the goal of which is to create adversarial perturbations that maximize the target model’s loss function while remaining visually undetectable to human observers. To achieve this, the authors utilize a variant of the PGD algorithm called Dense Adversary Generation (DAG). The DAG method implements an iterative optimization process to identify the optimal adversarial perturbations.

Overall, the Stack–Threat Matrix reveals that vulnerabilities span every layer of the AS architecture. Successful attacks at one layer often propagate and amplify through the stack, highlighting the need for defense strategies that address the full system and not just isolated components.

## **5 Critical Appraisal of Defenses in the Context of Autonomous Systems**



Attack Description	Scenario	Target	Implications	Reference
<b>Perception Layer</b>				
Adversarial Patch	Patch embedded on road sign or object	Object detector/classifier	Misclassification, system malfunction, hazardous incidents	(Brown et al. 2018)
Adversarial Sticker	Adversarial sticker on object/surface	Object detection/segmentation	Scene misperception, incorrect decisions	(Chen et al. 2019)
Adversarial Apparel	Person wearing adversarial clothes or accessories	Human/object classification	Pedestrian missed, security breach	(Xu et al.; Sharif et al. 2020; 2019)
Adversarial Object	Placing adversarially-engineered 2D/3D object in environment	Object recognition	False identification, safety risks	(Kong et al. 2020)
Lighting Attack	Manipulate scene lighting/shadows	Vision-based perception	Misclassification, detection failure	(Hsiao et al. 2024)
Adversarial Rain	Raindrop patterns on lens/image	Vision-based perception	Degraded perception, environmental misinterpretation	(Guesmi et al. 2023)
Adversarial Clothing	Clothing designed to fool detector	Person detection/recognition	Security risk, evasion of detection	(Hu et al. 2023)
Remote Perception Attack	Malicious pattern injection via compromised comms	Camera-based detection	False negatives for critical objects	(Man et al. 2023)
LiDAR Spoofing	Fake laser signals to LiDAR sensor	LiDAR perception	False obstacle detection, collision risk	(Cao et al. 2019)
<b>Planning Layer</b>				
Traffic Sign Attack	Subverted or altered traffic sign	Traffic sign recognition	Misnavigation, rule violation, accident risk	(Eyholt et al. 2018)
GPS Spoofing	Falsified GPS signals	Navigation system	Route deviation, loss of control, accidents	(Horton and Ranganathan 2018)
UAV Tracking Attack	Compromised tracking data or communication	UAV route/target tracker	Loss of target, mission failure	(Fu et al. 2022)
Adversarial Billboard	Adversarial billboard/sign in environment	Object detection/classification	Scene confusion, misbehavior, planning error	(Zhou et al. 2020)
Adversarial Planning	Crafted planner input/feedback	Planning algorithm	Unsafe/inefficient routing, increased risk	(Edelkamp 2023)
Trajectory Attack	Adversarial input to prediction model	Trajectory prediction	Wrong agent movement forecast, collision	(Cao et al. 2022)
Semantic DoS Attack	Benign object induces overly conservative behavior	Behavioral planning module	Unnecessary stops or detours, degraded performance	(Wan et al. 2022)
<b>Control Layer</b>				
UAV OD Spoofing	Spoofed images for UAV detection	UAV object detection/control	Erroneous control action, unsafe maneuvers	(Tian et al. 2022)
Semantic Exploit	Malicious image for segmentation/detection	Control subsystem	Poor control decisions, potential accidents	(Xie et al. 2017b)
Trojaning Attack	Injecting backdoor during model training	Control algorithm	Unauthorized actuation, hijack risk	(Cheng et al. 2021)
Model Extraction	Query-based model stealing	Control algorithm/-model	IP theft, enables further attack planning	(Li et al. 2021)
Flying Patch	Drone delivers adversarial patch into field of view	Vision-based control	Remote error injection, loss of control	(Hanfeld et al. 2023)
GhostImage Attack	Remote projection of adversarial pattern	Camera-based control	Misclassification, control errors	(Man et al. 2020)
CAN Injection	Malicious CAN bus message injection	Vehicle control systems	Unauthorized control, theft	(Khan et al. 2022)

**Table 6** Stack-Threat Matrix

In this section, we critically examine state-of-the-art adversarial defense mechanisms for AS. We begin by outlining the unique operational and security requirements that robust AS defense systems must satisfy (Section 5.1). Next, we focus on the challenges posed by physical adversarial attacks and review recent approaches for defending against them (Section 5.2). Drawing on the analyses above, we refine our taxonomy of defense mechanisms and narrow our evaluation to those methods most relevant and effective for AS, discussed in Section 5.3 and summarized by Table 7. Finally, in Section 5.4, we systematically assess a set of thirty representative defense mechanisms, introducing our novel *AS-ADS* scoring framework to quantify their alignment with the practical needs of AS.

## 5.1 Defining Requirements for AS Defense Systems

Building on our analysis of AS vulnerabilities and the characteristics of the AS stack and vision model life-cycle, we identify the specific defense needs that must be addressed to ensure robust and trustworthy AS deployments. We then evaluate how current state-of-the-art defense mechanisms align with these needs and discuss the remaining key challenges.

To contextualize these requirements, consider a representative mission scenario: let  $d$  denote an autonomous unmanned aerial vehicle (UAV) tasked with navigating and conducting reconnaissance in diverse, potentially hostile environments. The UAV’s objectives include detecting both known and unknown armed vehicles, including those deliberately camouflaged using adversarial techniques.

Suppose further that  $d \in D$ , where  $D$  is a fleet of UAVs operating in different areas and leveraging federated learning (FL) to collaboratively update their models. While this distributed approach increases mission resilience, it also introduces additional attack surfaces, particularly via the communication and update mechanisms of FL.

Throughout its mission, UAV  $d$  may face a variety of adversarial threats. For example, adversarial patches, as described in (Zhang et al. 2022), may be used by adversaries to camouflage vehicles and evade detection targeting the perception layer. Adversarial training might be deployed to defend against known patch types, but novel attack variants can still bypass these defenses. Visually distracting adversarial billboards (Zhou et al. 2020) might divert the UAV from its intended path, while attacks on FL communication channels can inject poisoned data into the learning process.

Mechanisms to address these risks include adversarial training and detection-based approaches to filter potentially malicious images. However, a recurring limitation is their lack of adaptability to novel attacks and inability to learn from previously unseen patterns without extensive retraining.

This scenario exemplifies the broader landscape of AS security and highlights the need for defense mechanisms that can evolve in response to new threats, while also operating securely within collaborative, distributed learning frameworks. Additionally, for operational trustworthiness, defense mechanisms should provide interpretable outputs that enable human experts to visualize, categorize, and respond to detected attacks.

For instance, the detection approach proposed by Soares et al. (2022) employs a similarity-based deep neural network (Sim-DNN) to detect imperceptible adversarial attacks by comparing new data samples to learned prototypes. This prototype-based method is interpretable and does not require adversarial training, but still lacks robust response capabilities (e.g., automated flagging or recovery), and may sometimes misclassify novel legitimate samples as adversarial. Advancing research toward more adaptive, interpretable, and actionable frameworks thus remains an open challenge.

Developing robust AS defenses often requires a combination of mechanisms such as adversarial training, detection, and unified frameworks. From this analysis, we derive four critical requirements for AS defense mechanisms:

- **Real-Time Detection and Response:** Defenses must promptly identify and mitigate adversarial inputs to prevent compromise of safety-critical decisions.
- **Adaptability to Novel Attacks:** Mechanisms should respond effectively to new and evolving adversarial strategies without requiring complete retraining.
- **Interpretability and Transparency:** Outputs should be explainable and accessible to human operators, enabling informed oversight and intervention.
- **Resource Efficiency:** Methods must be computationally and energetically efficient for practical deployment on resource-constrained AS platforms.

These criteria serve as the foundation for our evaluation of state-of-the-art defense mechanisms in the remainder of this section and throughout the paper.

## 5.2 Defenses Against Physical Adversarial Attacks

Physical adversarial attacks represent a uniquely severe threat to AS due to their real-world feasibility, persistence, and capacity to compromise safety-critical operations throughout the perception–planning–control pipeline. Unlike digital perturbations, these attacks often manifest as tangible modifications in the environment, such as adversarial patches on road signs, manipulated sensor readings, or spoofed trajectories, and are intentionally crafted to survive environmental changes. However, robust and generalizable defenses against physical attacks remain limited, fragmented, and often unvalidated beyond narrowly defined scenarios, largely due to the lack of standardized, physically grounded evaluation benchmarks.

To enhance adversarial robustness in the physical domain, recent research has focused on three broad categories of defense: *proactive*, *reactive*, and *unified* frameworks. Yet, few existing methods are designed to accommodate the full spectrum of real-world variability encountered by AS.

Within **Proactive strategies**, Adversarial training with physically realizable attacks (e.g., LiDAR perturbations or real-world patch examples) has shown promise in controlled settings (Kurakin et al.; Lu and Radha 2016; 2023), but generalization to unseen conditions such as new weather, sensor occlusion, or novel object types is often poor. Input Pre-Processing methods, including semantic-aware masking and inpainting (Jing et al. 2024), as well as multi-step diffusion-based purification (Nie et al. 2022), offer complementary robustness, but their efficacy varies significantly across sensor modalities and attack types. Other proactive defenses include spatial attention hardening to guard against localized road sign attacks (Shibly et al. 2023) and multi-sensor

aerial fusion to strengthen detection pipelines (Chen and Chu 2023). Despite their value, such approaches are often brittle when facing adaptive or context-aware adversaries, and typically introduce trade-offs between robustness and perceptual fidelity. Similarly, trajectory prediction models trained under uncertainty provide resilience at the planning level, but remain underexplored for targeted physical threats (Zhang et al. 2022).

**Reactive detection-based defenses** focus on flagging anomalies during system operation. Techniques in this category include entropy-based localization of patch regions (Tarchoun et al. 2023), kinematic consistency checks for identifying violations of physical constraints (Yu et al. 2024), and hybrid pipelines that combine detection and input recovery (Liu et al. 2022). While these approaches offer interpretability and low-latency adaptation, they often struggle against subtle or context-aware attacks that closely mimic plausible environmental features.

**Unified and hybrid frameworks** integrate multiple defense mechanisms across the AS stack. For example, control-aware frameworks such as SpecGuard (Dash et al. 2024) maintain mission compliance even under partial perception failure, while sensor fusion approaches like VisionGuard (Han et al. 2024) validate consistency between sensory modalities. Adaptive neural modeling strategies, such as RCDN (Wang et al. 2024a), aim to dynamically harden internal representations against adversarial perturbations. However, these promising approaches often face scalability limitations and have not yet been comprehensively evaluated across the diverse operational environments typical of real-world AS deployments.

**Certified defenses** represent a recent advancement, targeting physical attacks with formal robustness guarantees. PatchCleanser (Xiang et al. 2022) provides certified robustness via double masking, while works such as (Yang et al. 2023) and (Zhang et al. 2022) extend certification to control systems and semantic segmentation. These approaches are grounded in strong theoretical guarantees, but frequently present challenges regarding runtime feasibility and limited coverage of the full spectrum of physical attack surfaces.

Despite these advances, several key challenges remain. Most defenses are evaluated under narrow physical conditions, lacking robustness to environmental variation or domain shift. High-performing methods—particularly those involving certification or fusion—often introduce significant computational overhead, raising concerns for real-time AS deployment. Moreover, defenses rarely propagate protection beyond perception to downstream modules such as planning or control, leaving the broader autonomy stack exposed. Existing detection methods frequently fail to generalize across attack types or modalities, underscoring the need for attack-agnostic, adaptive detection pipelines. Some of these are beginning to emerge in adversarial attack research (Li et al. 2024) and deepfake detection (Pellicier et al. 2024), and could potentially be translated to the physical domain due to their prototype-based characteristics, though this remains to be explored.

Given these limitations, certified defenses and targeted detection mechanisms currently stand out as the most promising approaches against physical adversarial attacks in AS. Recent contributions, (some of which are evaluated in detail in Section 5.4)

demonstrate notable progress, but comprehensive integration and rigorous validation across the full AS pipeline remain critical open challenges for future work.

### 5.3 Defense Taxonomy Simplification

To address the real-time, adaptive, interpretable, and resource-conscious requirements of AS, we categorize SOTA defenses according to their core methodology, rather than along legacy proactive/reactive lines. We exclude Model Regularization, Model Distillation, and Provable defenses from our main analysis. Regularization and distillation are either now subsumed within other defense categories or lack standalone relevance in recent AS-specific literature. Provable (i.e., formal verification) defenses are excluded due to their high computational cost and inflexibility for real-world AS deployment. Similarly, denoising and reconstruction are no longer considered standalone mechanisms, as they are now integrated into Pre-Processing or unified frameworks in recent works. Accordingly, we focus on five categories: Adversarial Training, Input Data Pre-Processing, Model Ensembles, Detection Mechanisms, and Unified Defense Frameworks. Each is evaluated across four criteria: *real-time response*, *adaptability to novel attacks*, *interpretability*, and *resource efficiency*. Table 7 summarizes their alignment with AS needs and shows the relevant literature selected within our paper.

#### 5.3.1 Adversarial Training

Adversarial training remains a foundational technique, where adversarial examples are incorporated into the model’s training process (Madry et al. 2019). In AS contexts, adversarial training in autoencoder filters has led to improvements in adversarial robustness for both white-box and black-box attacks. Such methods show improved resistance to certain perturbations, but face key limitations:

- **Real-Time Response:** *High*. Inference performance is real-time, but the training process is computationally intensive.
- **Adaptability:** *Low*. Generalization to unseen attacks is limited.
- **Interpretability:** *Low*. The mechanisms by which robustness is achieved are often opaque.
- **Efficiency:** *Low*. High cost in both training and memory.

#### 5.3.2 Input Data Pre-Processing

Pre-Processing techniques such as resizing, cropping, and denoising mitigate adversarial perturbations before they reach the model. Studies such as (Xie et al. 2017b) demonstrate their effectiveness, and recent advances include noise suppression, reconstruction, and purification layers. DiffPure (Nie et al. 2022) leverages diffusion models for adaptive purification, while UMoE (Lou et al. 2023) employs uncertainty-aware fusion to counter sensor-blinding attacks. Pre-Processing is widely adopted for real-time viability:

- **Real-Time Response:** *High*. Lightweight implementations can operate on edge devices.

Mechanism	Real-time	Adaptability	Interpretability	Efficiency	References
Adversarial Training	High	Low	Low	Low	(Goodfellow et al.; Madry et al.; Tramèr and Boneh; Wong et al.; Tramèr et al.; Rozsa et al.; Chen and Lee; Shen et al.; Xie et al.; Wang et al. 2014; 2019; 2020; 2017; 2016; 2021; 2021; 2019; 2024b)
Input Pre-Processing	High	Low	Mod.	High	(Xie et al.; Liao et al.; Li et al.; Shu et al.; Reyes-Amezcu et al.; Naser et al.; Hu et al.; Zhang et al.; Shibly et al.; Nie et al.; Zhang et al.; Wang et al. 2017a; 2018; 2024; 2021; 2024; 2018; 2023; 2024; 2023; 2022; 2022; 2024a)
Model Ensembles	Mod.	Mod.	Low	Low	(Xie et al.; Engstrom et al.; Liao et al.; Xu et al.; Bhagoji et al.; Bui et al.; Tramèr et al.; Deng and Mu; Mani et al.; Lu et al.; Lu et al.; Chen et al.; Huang et al.; Lou et al.; Zhao et al. 2017b; 2019; 2018; 2017; 2017; 2021; 2017; 2023; 2019; 2023; 2023; 2024; 2021; 2023; 2024)
Detection Mechanisms	High	Mod.	High	Mod.	(Guo et al.; Angelov and Soares; Goodfellow et al.; Carlini and Wagner; Grosse et al.; Feinman et al.; Xu et al.; Gupta et al.; Sabokrou et al.; Soares et al.; Gong et al.; Abdu-Aguye et al.; Hussain and Hong; Li et al.; Li et al.; Yu et al.; Liu et al.; Liu et al.; Chen and Chui; Lu and Radha 2019; 2021; 2014; 2017a; 2017; 2017; 2020; 2024; 2022; 2023; 2020; 2023; 2024; 2023; 2022; 2022; 2023; 2023)
Certified Defenses	Mod.	Low	High	Mod.	(Gowal et al.; Tjeng et al.; Muravev and Petiushko; Lecuyer et al.; Xiang et al.; Yang et al.; Zhang et al. 2018; 2017; 2022; 2019; 2022; 2023; 2022)
Unified Defense	High	High	High	Mod.	(Pellicer et al.; Du et al.; Freitas et al.; Cao et al.; Dash et al.; Tarchoun et al.; Jing et al.; Han et al.; Yu et al. 2024; 2018; 2020; 2024; 2024; 2023; 2024; 2024; 2024)

**Table 7** Simplified taxonomy of defenses relevant to AS & overall alignment with AS requirements

- **Adaptability:** *Low*. These methods are often bypassed by adaptive or physical attacks.
- **Interpretability:** *Moderate*. Effects are visible in the processed input, but causality for prediction changes may be indirect.
- **Efficiency:** *High*. Minimal runtime cost.

Notably, this category is evolving: standard techniques (e.g., resizing, cropping, denoising) (Xie et al. 2017b) are now being combined with advanced approaches such as diffusion models (Nie et al. 2022). Pre-Processing is increasingly integrated into more complex pipelines, leading to Unified Models such as (Han et al. 2024), which combine sensory fusion, filtering, time-series (ARIMA, LSTM), and anomaly detection layers.

### 5.3.3 Model Ensembles

Model ensembles leverage diversity by combining multiple models, making it more difficult for adversaries to simultaneously deceive all models (Bui et al. 2021). Key characteristics are:

- **Real-Time Response:** *Moderate*. Inference latency increases with the number of models.
- **Adaptability:** *Moderate*. Greater diversity can improve resistance to transfer attacks.
- **Interpretability:** *Low*. Internal logic is often obscured by the ensemble fusion process.
- **Efficiency:** *Low*. Requires substantial hardware for parallel model execution.

Although ensembles are effective, few recent AS-specific implementations exist due to resource constraints. For example, the MADE framework (Zhao et al. 2024) employs ensemble-like anomaly scoring over multi-vehicle inputs to detect collaborative attacks in V2X scenarios. However, this method is not a traditional ensemble but rather a soft classification, reflecting a broader trend: literature is shifting from full ensembles to more flexible unified implementations.

### 5.3.4 Detection Mechanisms

AS increasingly rely on detection mechanisms for their interpretability, real-time performance, and applicability throughout the AS stack and life-cycle. Alongside Unified Frameworks, detection is now one of the fastest growing fields in adversarial defense, with Detection and Unified papers constituting over 50% of recent (2023 onward) publications.

Examples include Among Us (Li et al. 2023), which detects 3D adversarial inputs in V2X-Sim via consensus-breaking heuristics; Segment-and-Complete (Liu et al. 2022), which identifies adversarial patches through segmentation masks; and PhySense (Yu et al. 2024), which generalizes detection to real-world perturbations. Prototype-based, highly interpretable systems such as (Angelov and Soares 2021) further demonstrate this category’s strengths:

- **Real-Time Response:** *High*. Detection is typically performed pre-inference.
- **Adaptability:** *Moderate*. Detection patterns can generalize to some unseen attacks.



- **Interpretability:** *High*. Outputs are often visual or score-based, supporting operator trust.
- **Efficiency:** *Moderate*. Auxiliary models or priors may increase computational demands.

### 5.3.5 Certified Defenses

Certified defenses offer provable robustness guarantees under specific perturbation budgets. In AS-relevant domains:

- PatchCleanser (Xiang et al. 2022) certifies robustness against small visible patches (up to 2% area) using random masking and smoothing, evaluated on CIFAR and ImageNet.
- Demasked Smoothing (Zhang et al. 2022) certifies patch-level segmentation robustness via randomized ablation masking, showing strong resistance on ADE20K under shadow and patch attacks.
- Certified Robust Control (Yang et al. 2023) formulates controller robustness for AS via Lyapunov-based certified adaptation, effective against bounded input perturbations.

Strengths and trade-offs are:

- **Real-Time Response:** *Moderate*. Certification layers may introduce runtime sampling.
- **Adaptability:** *Low*. Guarantees hold only for bounded attacks and require redefinition for new scenarios.
- **Interpretability:** *High*. Theoretical guarantees are transparent and explainable.
- **Efficiency:** *Moderate-Low*. Additional overhead from sampling, smoothing, or invariant computations.

### 5.3.6 Unified Defense Frameworks

Unified frameworks, as defined in this review, represent a new taxonomy. They integrate heterogeneous defense techniques (e.g., detection + recovery) using shared feature pipelines or modular layers, whereas ensembles aggregate predictions from independently trained full models. For example, Pellicer et al. (2024) present a lightweight framework combining prototype-based detection and classification for attacks and unseen classes, along with attack recovery via denoising methods, achieving over 90% accuracy on CIFAR-10.

Other notable unified defenses include Du et al. (2018), which detects abnormal samples for any pre-trained softmax classifier, and UNMASK (Freitas et al. 2020), which both identifies adversarial attacks and mitigates their effects through robust reclassification. UNMASK can detect up to 96.75% of attacks and restore correct classification in up to 93% of cases.

More AS-specific frameworks, such as SpecGuard (Dash et al. 2024), integrate detection, filtering, and signal processing to detect UAV sensor spoofing with a 92% recovery success rate and only 15% performance overhead. Time-Travel (Etim and Szefer 2024) compares live input with historical image matches to detect false

patches, achieving 100% effectiveness against recent adversarial examples in traffic sign classification.

Overall, unified methods best align with AS priorities and full life-cycle needs:

- **Real-Time Response:** *High*. Historical matching and statistical filtering are efficient on-device.
- **Adaptability:** *High*. Frameworks leverage both priors and learned models.
- **Interpretability:** *High*. Alerts are easily visualized and validated by operators.
- **Efficiency:** *Moderate*. Moderate computational and storage requirements.

#### 5.4 Autonomous Systems Adversarial Defense Score (AS-ADS) framework

To systematically assess the suitability of defense methods for AS, we build on the updated taxonomy provided in Table 7

We introduce the **Autonomous Systems Adversarial Defense Score (AS-ADS)**, a scoring framework designed to quantify each method’s alignment with operational AS constraints. AS-ADS evaluates across our 4 dimensions (**Real-Time Detection and Response**, **Adaptability to Novel Attacks**, **Interpretability and Transparency** and **Resource Efficiency**):

Each criterion is rated on a 0 to 1 scale in 0.25 increments. The final AS-ADS score is calculated as the average of these four values, scaled to a 1–5 range and rounded to the nearest half:

$$\text{AS-ADS}(P) = \left( \frac{R + A + I + E}{4} \right) \times 5 \quad (1)$$

where  $R, A, I, E \in [0, 1]$  represent the real-time, adaptability, interpretability, and efficiency scores, respectively.

$R, A, I, E$  are obtained for each paper after marking using rubrics in Table 8.

Criterion	0 pts	0.25 pts	0.5 pts	1.0 pts
Real-Time Response	Batch inference only	High latency	Optimized inference only	Real-time at edge-level
Adaptability to Novel Attacks	Static model	Minor generalization	Modular, partially adaptable	Robust to unseen attacks
Interpretability	Black-box	Minimal logs	Score-based or visual	Prototype/semantic explanation
Resource Efficiency	High overhead	GPU-dependent	Deployable with tuning	Lightweight for AS hardware

**Table 8** AS-ADS scoring rubric by criterion

This scoring framework facilitates standardized, comparative evaluation of SOTA defense methods in AS settings. By grounding the scores in real-world operational needs and deployment constraints, AS-ADS enables both a fine-grained critique of existing methods and an actionable guide for future design.

Method Description	Score	Reference	AS
<b>Detection Mechanisms</b>			
Detects adversarial inputs using evolved image processing sequences via genetic algorithms	2	(Gupta et al. 2020)	–
Detects adversaries via SSL-based consistency checks in feature and label space	4.5	(Sabokrou et al. 2024)	–
Combines LSTM temporal consistency checks with majority voting for time-series attack detection	2	(Abdu-Aguye et al. 2020)	–
Reveals adversarial artifacts through autoencoder reconstruction error analysis	2.5	(Hussain and Hong 2023)	–
Detects outliers through learned similarity metrics in contrastive feature space	2.5	(Soares et al. 2022)	–
Learns attack-agnostic features via self-supervised contrastive prototype alignment	3	(Li et al. 2024)	–
Identifies anomalies through statistical hypothesis testing in feature space	1	(Grosse et al. 2017)	–
Detects patches through entropy analysis and visual localization	4	(Tarchoun et al. 2023)	✓
Identifies physics violations through kinematic consistency checks	4.5	(Yu et al. 2024)	✓
Detects/recovers patches via joint detection-completion pipeline	4	(Liu et al. 2022)	✓
<b>Pre-Processing Defenses</b>			
Embeds frequency-aware watermarks in RAW files using multi-spectral fusion	4	(Hu et al. 2023)	–
Optimizes augmentation parameters via gradient-based adversarial search	1	(Shu et al. 2021)	–
Enhances robustness through transfer of adversarial patterns across vision tasks	1	(Reyes-Amezcu et al. 2024)	–
Neutralizes patches through semantic context-aware masking/inpainting	5.0	(Jing et al. 2024)	✓
Scales LiDAR robustness via density-aware point cloud processing	4.5	(Lu and Radha 2023)	✓
Hardens aerial detection through multi-sensor fusion	2.5	(Chen and Chu 2023)	✓
Protects road sign recognition through spatial attention hardening	2.5	(Shibly et al. 2023)	✓
Purifies inputs through multi-step diffusion denoising	2	(Nie et al. 2022)	✓
Improves trajectory prediction via uncertainty-aware training	2.5	(Zhang et al. 2022)	✓
<b>Unified Defenses</b>			
Integrates detection-denoiser architecture with noise-adaptive thresholds	3.5	(Pellicer et al. 2024)	–
Detects OOD samples through temperature-scaled confidence calibration	2.5	(Du et al. 2018)	–
Verifies predictions through robust part-based feature alignment	4	(Freitas et al. 2020)	–
Links adversarial and backdoor attack patterns for joint cross-attack detection	3	(Yin et al. 2025)	–
Detects face spoofing through dual-space (spatial/frequency) reconstruction analysis	4	(Cao et al. 2024)	–
Ensures mission-compliant recovery through specification-aware control	4.5	(Dash et al. 2024)	✓
Guarantees cross-sensor consistency through multi-modal fusion checks	4.5	(Han et al. 2024)	✓
Enables robust perception via dynamic neural feature modeling	4.5	(Wang et al. 2024a)	✓
<b>Certified Defenses</b>			
Provides certified patch robustness through double-masking with formal guarantees	4.5	(Xiang et al. 2022)	✓
Certifies control stability under perturbations via Lyapunov analysis	4	(Yang et al. 2023)	✓
Ensures segmentation robustness via masked smoothing certification	4	(Zhang et al. 2022)	✓

**Table 9** AS-ADS evaluation of adversarial defenses. “AS” marks those developed for Autonomous Systems.

For the evaluation, we selected a representative subset of 30 defenses from the literature discussed in this paper, focusing on *Pre-Processing*, *Detection*, *Certified*, and *Unified* defenses, as identified in Sec. 5.3. Our evaluation subset includes: (a) foundational works that paved the way for newer defense mechanisms in each category, alongside relevant recent approaches—(Hu et al.; Shu et al.; Gupta et al.; Sabokrou et al.; Reyes-Amezcu et al.; Abdu-Aguye et al.; Hussain and Hong; Soares et al.; Li et al.; Grosse et al.; Pellicer et al.; Du et al.; Freitas et al.; Yin et al.; Cao et al. 2023; 2021; 2020; 2024; 2024; 2020; 2023; 2022; 2024; 2017; 2024; 2018; 2020; 2025; 2024)—and (b) work from 2022 onward tailored specifically to the AS domain—(Dash et al.; Tarchoun et al.; Jing et al.; Han et al.; Yu et al.; Xiang et al.; Yang et al.; Zhang et al.; Liu et al.; Chen and Chu; Lu and Radha; Shibly et al.; Nie et al.; Zhang et al.; Wang et al. 2024; 2023; 2024; 2024; 2024; 2022; 2023; 2022; 2022; 2023; 2023; 2023; 2022; 2022; 2024a).

We derived final scores by combining each paper’s reported findings and expert knowledge of the architectures, using the established rubric. For reproducibility individual scores per paper can be found in Appendix B, the overall scores per paper have been presented in Table 9

It is important to note that the selection of scored papers reflects expert judgment and is not intended to exhaustively cover all available methods, but rather to provide a representative overview of current options and their effectiveness. This gives readers and researchers practical guidance for deploying or developing defense systems across the attack surfaces identified in this report.

A score of 5 does not imply perfection, but rather the closest alignment with the requirements defined herein. The diversity of threats, datasets, and evaluation protocols across the literature makes it challenging to determine a universally optimal method. Nonetheless, we believe this evaluation brings the field closer to that goal. To improve accuracy and utility in future work, we recommend detailed reporting of runtime overhead, FPS degradation, GPU memory usage, interpretability, and accuracy for each defense using standardized datasets and attacks, although this is beyond the scope of this review.

## 6 Conclusion and Future Directions

This review provides a holistic, system-level analysis of adversarial threats and defenses for AS, integrating insights from both foundational vision-centric research and recent AS-specific advances. By bridging these two strands of the literature, we offer a unified framework that captures the cascading impact of digital and physical adversarial vulnerabilities across the autonomy stack. Our taxonomy, scenario-driven matrices, and comparative synthesis enable both researchers and practitioners to assess current gaps and prioritize future work in making vision-driven AS secure and resilient.

A cornerstone of our approach is the development and use of actionable analytical matrices, including the **Life-cycle–Attack**, **Stack–Threat**, and **Exposure–Impact** matrices. These matrices concretely map how adversarial vulnerabilities propagate throughout the AI life-cycle and across layered AS architectures. For example, our Life-cycle–Attack Matrix reveals both the temporal exposure of AS to poisoning, backdoor,

and evasion attacks, and the unique risk windows at each stage of system operation. The Stack–Threat Matrix grounds these vulnerabilities in real-world scenarios, demonstrating how a compromised perception module (such as a camera subjected to adversarial patches or sensor spoofing) can trigger failures in planning that propagate to mission-critical control. By further linking these technical threats to operational consequences in the Exposure–Impact Matrix, our review enables researchers and practitioners to move beyond abstract taxonomies toward practical, system-level threat modeling and benchmarking.

Our comparative synthesis of adversarial attacks, spanning both **digital** and **physical** domains, highlights a crucial reality: vulnerabilities in AS are rarely confined to a single module. Instead, our analysis of real attack case studies and scenario-based evaluations demonstrates that adversarial examples often trigger failures that cascade across subsystems, resulting in safety or mission-critical consequences far beyond mere performance degradation on academic benchmarks. This insight exposes the inadequacy of traditional, static, perception-only evaluation metrics and establishes the need for operationally meaningful, stack-wide robustness assessment.

In **critically appraising defense strategies**, we show that the conventional taxonomy, dividing defenses into proactive and reactive categories, does not sufficiently capture the practical demands of AS. By shifting the focus to underlying mechanisms, and by introducing unified, context-aware defenses as a distinct class, we reveal that most state-of-the-art methods, even when successful in vision research, fail to meet the simultaneous requirements of real-time performance, adaptability to new threat vectors, interpretability, and resource efficiency essential for deployment in AS. The **AS-ADS scoring framework** introduced in this review directly evaluates these axes, and our comprehensive analysis across more than thirty contemporary defenses finds that only a minority approach a balanced, deployment-ready profile. In particular, robust and interpretable defenses against physical and multi-modal threats are still lacking, and few methods have demonstrated stack-wide or life-cycle-spanning effectiveness in realistic scenarios.

Despite these advances, **significant challenges and research gaps remain**. Most available benchmarks remain narrowly focused on perception or digital attacks, with little provision for evaluating cascading effects, cross-modal dependencies, or mission-level outcomes. Few studies rigorously validate either attacks or defenses under closed-loop, multi-agent, or sim-to-real conditions that reflect the operational reality of modern AS. While the threat matrices presented in this review provide a critical foundation for system-level risk assessment, their full potential will only be realized when supported by open, community-driven benchmarking platforms and evaluation protocols that span the entire stack.

**Looking ahead**, meaningful progress in adversarial robustness for AS will depend on several intertwined advances. The field must prioritize the creation of stack-integrated datasets and simulation environments capable of capturing cascading failures, temporal persistence, and the interplay of digital and physical threats. Defense research should increasingly focus on mechanisms that are interpretable, for some cases also certifiable, and that are validated in resource-constrained, real-time settings.

There is a particular need to design and rigorously test unified, adaptive defense frameworks that can operate coherently across perception, planning, and control layers, and that can dynamically respond to evolving threat landscapes in real deployments. The integration of human-in-the-loop monitoring and decision-making, as well as robust protocols for sim-to-real transfer, will be critical for bridging the gap between academic innovation and practical deployment.

In summary, by clarifying the layered structure of AS vulnerabilities, mapping concrete threat pathways, and critically evaluating the mechanisms and readiness of current defenses, this review sets a new agenda for adversarial research in Autonomous Systems. We hope that the analytical frameworks, results, and open challenges identified here will help guide the community toward robust, certifiable, and operationally viable solutions for the next generation of trustworthy autonomous technologies.

## Declarations

**Author contribution.** A.L.P. served as the lead author and was responsible for the conceptualisation, literature review, taxonomy development, methodology, analysis, drafting of the manuscript, visualisation, and overall project administration. P.A. and N.S. reviewed the manuscript critically for important intellectual content; P.A. additionally helped with project administration and editorial communication N.S. additionally assisted in rewriting and sharpening the title and abstract to improve clarity and framing. All authors have read and approved the final manuscript.

**Acknowledgment of Funding.** This research is supported, in part, by the UKRI Trustworthy Autonomous Systems Node in Security/EPSRC Grant EP/V026763/1.

## Appendix A Background tables

Attacker Knowledge	Attack Timing	Attack Location	Examples
White-Box	Evasion	Digital	L-BFGS (Fletcher 2013), FGSM (Goodfellow et al. 2014), I-FGSM (Kurakin et al. 2016), PGD (Madry et al. 2019), DeepFool (Moosavi-Dezfooli et al. 2016), C&W (Carlini and Wagner 2017b), JSMA (Papernot et al. 2016), UAP (Moosavi-Dezfooli et al. 2017), DDN ( Jérôme Rony and Luiz G. Hafemann and Luiz S. Oliveira and Ismail Ben Ayed and Robert Sabourin and Eric Granger 2019), Elastic Net (Chen et al. 2018)
White-Box	Poisoning	Digital	Data Injection (Biggio et al. 2012), Label Flipping (Koenig et al. 2015), Backdoor (Gu et al. 2017), MetaPoison (Huang et al. 2020)
Black-Box	Evasion	Digital	Boundary (Brendel et al. 2017), ZOO (Chen et al. 2017), SimBA (Guo et al. 2019), One Pixel (Su et al. 2019), Square Attack (Andriushchenko et al. 2020), HSJA (Chen et al. 2020)
Black-Box	Poisoning	Digital	BadNets (Gu et al. 2019), Clean-label Backdoor (Zhao et al. 2019), GAN-based Poisoning (noz González et al. 2019)

**Table A1** Foundational Taxonomic classification of image-domain adversarial attacks. Attack location is included to show the digital-focused in literature. However in many cases, surveys do not include this dimension.

## Appendix B AS-ADS method evaluations

This includes the scores and small reasoning behind each scored for the defense methods evaluated in SECTION:

### ***DRAW: Defending Camera-shooted RAW against Image Manipulation*** (Hu et al. 2023)

- Real-Time: 0.5 (Lightweight network optimized for camera integration)
- Adaptability: 0.5 (Cross-ISP pipeline protection)
- Interpretability: 1.0 (Pixel-level manipulation maps)
- Efficiency: 1.0 (0.95% params vs U-Net)

**Method:** Embeds frequency-aware watermarks in RAW files using multi-spectral fusion, preserving detection capability through arbitrary ISP processing chains.

**AS-ADS Score:** 3.75



Type	Mechanism	Description & Advantages	Limitations	References
<b>Proactive</b>	Adversarial Training	Includes adversarial samples in training; improves model robustness.	High compute cost; limited to known attacks.	(Goodfellow et al.; Madry et al.; Tramèr et al.; Bonhe; Wong et al.; Tramèr et al.; Rozsa et al.; Chen and Lee; Shen et al.; Xie et al.; Wang et al. 2014; 2019; 2020; 2021; 2016; 2021; 2021; 2019; 2024b)
<b>Proactive</b>	Input Pre-Processing	Applies resizing, smoothing or augmentation; reduces perturbation impact.	May distort clean inputs; less effective on adaptive attacks.	(Xie et al.; Liao et al.; Li et al.; Shu et al.; Reyes-Amezcu et al.; Naseer et al.; Hu et al.; Zhang et al.; Shibly et al.; Nie et al.; Zhang et al.; Wang et al.; Lou et al. 2017a; 2018; 2024; 2021; 2024; 2018; 2023; 2024; 2023; 2022; 2022; 2024a; 2023)
<b>Proactive</b>	Model Ensemble	Combines multiple models' outputs; diversifies weaknesses.	Higher inference latency; greater resource use.	(Xie et al.; Engstrom et al.; Liao et al.; Xu et al.; Bhagoji et al.; Bui et al.; Tramèr et al.; Deng and Mu; Mani et al.; Lu et al.; Lu et al.; Chen et al.; Huang et al.; Zhao et al. 2017b; 2019; 2018; 2017; 2021; 2017; 2023; 2019; 2023; 2023; 2024; 2021; 2024)
<b>Proactive</b>	Model Regularization	Adds constraints or penalties during training; improves generalization.	May reduce clean accuracy; limited adversarial gains.	(Szegedy et al.; Kannan et al.; Drucker and Cun; Ross and Doshi-Velez 2013; 2018; 1992; 2018)
<b>Proactive</b>	Model Distillation	Uses soft-label transfer to a smaller model; enhances certain robustness.	Distilled model underperforms on clean data; narrow defense scope.	(Hinton et al.; Papernot et al.; Carlini and Wagner; Goldblum et al.; Costa et al. 2015; 2016; 2017b; 2020; 2024)
<b>Proactive</b>	Provable Defenses	Leverages formal verification to certify robustness bounds.	Very high compute; limited scalability to large models.	(Ehlers; Katz et al.; Tjeng et al.; Raghuathan et al.; Cohen et al.; King and Wang; Hong et al.; Lecuyer et al. 2017; 2017; 2017; 2018; 2019; 2019; 2024; 2019)
<b>Proactive</b>	Certification & Verification	Applies formal methods to verify model resilience; builds trust.	Computationally demanding; may not reflect real-world inputs.	(Gowal et al.; Tjeng et al.; Muravev and Petrushko; Lecuyer et al.; Xiang et al.; Yang et al.; Zhang et al. 2018; 2017; 2022; 2019; 2022; 2023; 2022)
<b>Reactive</b>	Detection-Based	Flags or rejects suspicious inputs via statistical tests or auxiliary models.	False positives; attacker can evade detection.	(Guo et al.; Angelov and Soares; Goodfellow et al.; Carlini and Wagner; Grosse et al.; Feinman et al.; Xu et al.; Gupta et al.; Sabokrou et al.; Soares et al.; Gong et al.; Abdu-Aguye et al.; Hussain and Hong; Li et al.; Li et al.; Yu et al.; Liu et al.; Chen and Chu; Lu and Radha 2019; 2021; 2014; 2017a; 2017; 2017; 2017; 2020; 2024; 2022; 2023; 2020; 2023; 2024; 2023; 2024; 2022; 2023; 2023)
<b>Reactive</b>	Denoising & Reconstruction	Uses autoencoders/GANs to remove perturbations; reconstructs clean inputs.	Possible information loss; imperfect recovery.	(Meng and Chen; Vincent et al.; Lempitsky et al.; Liao et al.; Samangouei et al.; Samangouei et al. 2017; 2008; 2018; 2018; 2018)
<b>Unified</b>	Unified Defense Frameworks	Integrates detection, noise reduction, and novel-class identification in one pipeline; adaptive to known and unknown attacks.	Moderate compute overhead; complex integration; limited large-scale testing.	(Pellicer et al.; Du et al.; Freitas et al.; Cao et al.; Dash et al.; Tarchoun et al.; Jing et al.; Han et al.; Yu et al. 2024; 2018; 2020; 2024; 2024; 2023; 2024; 2024; 2024)

**Table A2** Summary and classification of adversarial defense mechanisms.

Dataset (Reference)	Domain	Scenario(s)	Relevance to AS	Use
MNIST (Lecun et al. 1998)	Handwritten digits	Baseline testing, digital adversarial examples	Low	Testing classifier vulnerability
CLIFAR-10 (Krizhevsky 2009)	Small objects, digital images	Digital adversarial attacks, classifier benchmarks	Low	Small-scale adversarial robustness
ImageNet (Deng et al. 2009)	Large-scale digital images	Digital adversarial attacks, corruptions	Moderate	Pretraining, digital attack transfer, accuracy drop
ImageNet-P (Hendrycks et al. 2021)	Perturbation-augmented ImageNet	Corruptions, robustness evaluation	Moderate	Benchmark for perturbation robustness
COCO, xView (Liu et al. 2022)	Object detection	Adversarial patch attacks, digital detection	Moderate	mAP degradation under localized attacks
ADE20K (Zhang et al. 2022)	Scene segmentation (digital)	Certified patch detection, segmentation	Moderate	Certified accuracy, visual overlap
DOTA (Xia et al. 2018)	Aerial images, object detection	Patch attacks, adversarial detection	High	UAV surveillance robustness
Mapillary Traffic Sign (Poggi and Mottocchia 2017)	Real-world traffic scenes	Physical adversarial attacks (signs)	High	Traffic sign robustness, AV testing
RobustBench (Croce et al. 2020)	Digital, standardized benchmark	Digital adversarial attacks (various datasets)	High	Model benchmarking for adversarial robustness
SafeBench (Xu et al. 2022)	Simulation (CARLA)	Adversarial scenarios, hostile agents (vehicles/pedestrians)	High	Closed-loop AV safety, collision rate, completion, rule violation
CARLA-GeAR (Nesti et al. 2022)	Simulation (CARLA)	Physically-realizable patches on vehicles, adversarial scenarios	High	Multi-task driving (segmentation, detection), mIoU, mAP, depth error
RobustE2E (Jiang et al. 2024)	Simulation (CARLA), E2E driving	White-box input/feature perturbations, corruptions	High	Steering error, lane keeping, success rate under attack
DCI Dataset (Zhang et al. 2023)	Simulation + rendering, vehicle detection	Physical patches, weather/angle variations	High	mAP drop, detection under physical attacks
DD-RobustBench (Wu et al. 2025)	Digital dataset distillation	Digital adversarial attacks, distillation robustness	Moderate	Robustness of distilled datasets
Car Hacking (Kang et al. 2021)	Real (CAN logs)	Spoofed/malicious CAN bus messages	High	In-vehicle intrusion detection, false alarm rate
V2X-Sim (Li et al.; Zhao et al. 2023; 2024)	Simulation (LiDAR/V2X)	LiDAR spoofing, anomaly injection, cooperative attacks	High	Detection rate, anomaly precision/recall
KITTI-Adv/Blind, STF (Lou et al. 2023)	Real+Synth, sensor fusion	Sensor blinding, vision fusion, uncertainty estimation	High	mIoU, mAP under blinding or fusion attacks
DAIR-V2X (Zhao et al. 2024)	Real-world cooperative AV	Malicious contributor, V2X patch attacks	High	Detection accuracy for V2X fusion, anomaly detection
Google Street View (Etim and Szefer 2024)	Real images, street scenes	Time-inconsistent, physical perturbations	Moderate-High	Historical adversarial analysis, sign recognition, detection accuracy

**Table A3** Comparative overview of adversarial robustness datasets/platforms relevant to Autonomous Systems, including simulation tools and real-world data.

***Adversarial Differentiable Augmentation (Shu et al. 2021)***

- Real-Time: 0.25 (Offline augmentation optimization)
- Adaptability: 0.5 (Partial corruption resistance)
- Interpretability: 0.0 (No diagnostic features)
- Efficiency: 0.25 (2.3 GPU hours/search)

**Method:** Automates augmentation parameter selection via gradient-based adversarial search for robust training. **AS-ADS Score:** 1.25

***Evolutionary IPTS Detection (Gupta et al. 2020)***

- Real-Time: 0.25 (Multi-stage processing)
- Adaptability: 0.5 (Attack-specific sequences)
- Interpretability: 0.5 (Difference maps)
- Efficiency: 0.25 (Genetic algorithm overhead)

**Method:** Evolves optimal image processing pipelines using genetic algorithms to reveal adversarial artifacts. **AS-ADS Score:** 1.875

***BEYOND: Detecting Adversarial Examples via SSL Neighborhood Relations (Sabokrou et al. 2024)***

- Real-Time: 1.0 (Optimized for edge deployment with 50 neighbors processed at 23ms/image)
- Adaptability: 1.0 (Attack-agnostic design validated against 12+ attack types)
- Interpretability: 0.5 (Score-based consistency metrics with visualization support)
- Efficiency: 1.0 (Lightweight SSL backbone with 0.9M parameters)

**AS-ADS Score:** 4.375

***Delta Data Augmentation (Reyes-Amezcu et al. 2024)***

- Real-Time: 0.25 (Transfer learning focus)
- Adaptability: 0.5 (Cross-dataset transfer)
- Interpretability: 0.0 (Opaque perturbation transfer)
- Efficiency: 0.25 (GPU-intensive)

**Method:** Transfers adversarial patterns from high-level vision tasks to enhance low-level task robustness. **AS-ADS Score:** 1.25

***Temporal Consistency Defense (Abdu-Aguye et al. 2020)***

- Real-Time: 0.5 (143ms LSTM inference)
- Adaptability: 0.25 (Fixed thresholds)
- Interpretability: 0.25 (Entropy logs)
- Efficiency: 0.5 (Embedded compatibility)

**Method:** Combines frame-wise consistency checks with temporal majority voting for video attack detection. **AS-ADS Score:** 1.875

*Autoencoder Reconstruction (Hussain and Hong 2023)*

- Real-Time: 0.5 (47ms inference)
- Adaptability: 0.5 (73% unseen attacks)
- Interpretability: 0.5 (Reconstruction errors)
- Efficiency: 0.5 (580MB model)

**Method:** Detects adversaries through reconstruction error analysis using compact autoencoders. **AS-ADS Score:** 2.5

*Similarity Metric Analysis (Soares et al. 2022)*

- Real-Time: 0.5 (89ms Jetson TX2)
- Adaptability: 0.5 (12 attack types)
- Interpretability: 1 (Confidence scores and prototypes)
- Efficiency: 0.5 (15W consumption)

**Method:** Identifies outliers through learned similarity metrics in feature space. **AS-ADS Score:** 3.125

*Contrastive Prototype Learning (Li et al. 2024)*

- Real-Time: 0.5 (33ms inference)
- Adaptability: 1.0 (94.7% cross-attack)
- Interpretability: 1.0 (Prototype matching)
- Efficiency: 0.5 (2.1GB VRAM)

**Method:** Learns attack-agnostic features through self-supervised contrastive prototype alignment. **AS-ADS Score:** 3.75

*Statistical Anomaly Detection (Grosse et al. 2017)*

- Real-Time: 0.25 (Batch processing)
- Adaptability: 0.25 (Static models)
- Interpretability: 0.25 (Basic scores)
- Efficiency: 0.25 (CPU-intensive)

**Method:** Detects outliers through likelihood ratio testing in feature statistics. **AS-ADS Score:** 1.25

*UNICAD Framework (Pellicer et al. 2024)*

- Real-Time: 0.5 (24 FPS pipeline)
- Adaptability: 0.75 (Wide range of untrained in digital attacks and +85% Unseen class identification)
- Interpretability: 1 (Prototype based)
- Efficiency: 0.5 (8GB VRAM)

**Method:** Unified approach for attack detection, noise reduction, and novel class identification. **AS-ADS Score:** 3.437

*Confidence-Calibrated OOD (Du et al. 2018)*

- Real-Time: 0.5 (45ms detection)

- Adaptability: 0.5 (82% cross-domain)
- Interpretability: 0.5 (Thresholding)
- Efficiency: 0.5 (16W edge)

**Method:** Detects out-of-distribution samples through temperature-scaled confidence calibration. **AS-ADS Score:** 2.5

***Robust Feature Verification (Freitas et al. 2020)***

- Real-Time: 0.5 (28ms alignment)
- Adaptability: 1.0 (97.3% detection)
- Interpretability: 1.0 (Semantic maps)
- Efficiency: 0.5 (4.3GB model)

**Method:** Verifies predictions through robust part-based feature alignment. **AS-ADS Score:** 3.75

***Cross-Attack Bridge Defense (Yin et al. 2025)***

- Real-Time: 0.5 (33ms analysis)
- Adaptability: 1.0 (89% cross-backdoor)
- Interpretability: 0.5 (Similarity scores)
- Efficiency: 0.5 (12% overhead)

**Method:** Links adversarial and backdoor attack patterns for joint defense. **AS-ADS Score:** 3.125

***Dual-Space Face Defense (Cao et al. 2024)***

- Real-Time: 0.5 (41ms processing)
- Adaptability: 1.0 (95.6% spoof detection)
- Interpretability: 1.0 (Error maps)
- Efficiency: 0.5 (6.7GB VRAM)

**Method:** Reconstructs face images in spatial/frequency domains for unified spoof detection. **AS-ADS Score:** 3.75

***SpecGuard Recovery (Dash et al. 2024)***

- Real-Time: 1.0 (15ms ARM recovery)
- Adaptability: 1.0 (92% multi-sensor)
- Interpretability: 0.5 (Compliance scores)
- Efficiency: 1.0 (15% overhead)

**Method:** Recovers attacked inputs through safety specification-aware filtering. **AS-ADS Score:** 4.375

***Entropy-Based Patch Defense (Tarchoun et al. 2023)***

- Real-Time: 0.5 (54ms analysis)
- Adaptability: 1.0 (90% patches)
- Interpretability: 1.0 (Entropy maps)
- Efficiency: 0.5 (2.77% loss)

**Method:** Detects adversarial patches through localized entropy analysis. **AS-ADS Score:** 3.75

*Context-Aware Patching (Jing et al. 2024)*

- Real-Time: 1.0 (11ms edge)
- Adaptability: 1.0 (96.4% mAP)
- Interpretability: 1.0 (Semantic highlighting)
- Efficiency: 1.0 (0.9W power)

**Method:** Neutralizes patches through semantic context-aware masking and inpainting. **AS-ADS Score:** 5.0

*Multi-Sensor Guard (Han et al. 2024)*

- Real-Time: 1.0 (8ms fusion)
- Adaptability: 1.0 (97.3% cross-modal)
- Interpretability: 0.5 (Consistency reports)
- Efficiency: 1.0 (4.2W SoC)

**Method:** Ensures cross-sensor consistency for robust automotive perception. **AS-ADS Score:** 4.375

*Physics-Consistency Check (Yu et al. 2024)*

- Real-Time: 1.0 (9ms checks)
- Adaptability: 1.0 (94% cross-domain)
- Interpretability: 0.5 (Violation scores)
- Efficiency: 1.0 (3% CPU boost)

**Method:** Verifies physical plausibility of sensor inputs through kinematic checks. **AS-ADS Score:** 4.375

*Certified Patch Defense (Xiang et al. 2022)*

- Real-Time: 1.0 (18ms masking)
- Adaptability: 1.0 (83.9% certified)
- Interpretability: 1.0 (Mask proofs)
- Efficiency: 0.5 (45.1 mAP)

**Method:** Provides certified robustness through double-masking with formal guarantees. **AS-ADS Score:** 4.375

*Formal Control Certification (Yang et al. 2023)*

- Real-Time: 1.0 (22ms certification)
- Adaptability: 1.0 (Unseen perturbations)
- Interpretability: 0.5 (Stability margins)
- Efficiency: 0.5 (35% overhead)

**Method:** Certifies control stability under adversarial perturbations via Lyapunov analysis. **AS-ADS Score:** 3.75

*Demasked Segmentation (Zhang et al. 2022)*

- Real-Time: 1.0 (27ms inference)
- Adaptability: 1.0 (89% cross-task)
- Interpretability: 0.5 (Confidence maps)
- Efficiency: 0.5 (8.2GB VRAM)

**Method:** Certifiably robust semantic segmentation through masked smoothing. **AS-ADS Score:** 3.75

*Patch Detection-Completion (Liu et al. 2022)*

- Real-Time: 0.5 (143ms pipeline)
- Adaptability: 1.0 (91% patches)
- Interpretability: 1.0 (Completion vis)
- Efficiency: 0.5 (6.3W edge)

**Method:** Jointly detects and completes adversarial patches in object detection. **AS-ADS Score:** 3.75

*Aerial Object Defense (Chen and Chu 2023)*

- Real-Time: 0.5 (77ms processing)
- Adaptability: 0.5 (68% robustness)
- Interpretability: 0.5 (Region highlighting)
- Efficiency: 0.5 (4.8GB VRAM)

**Method:** Hardens aerial detection against adversarial object injections. **AS-ADS Score:** 2.5

*LiDAR Robustness Scaling (Lu and Radha 2023)*

- Real-Time: 1.0 (14ms processing)
- Adaptability: 1.0 (97% cross-sensor)
- Interpretability: 0.5 (Saliency maps)
- Efficiency: 1.0 (2.1W LiDAR)

**Method:** Scales adversarial robustness for LiDAR detection through density-aware processing. **AS-ADS Score:** 4.375

*Road Sign Defense (Shibly et al. 2023)*

- Real-Time: 0.5 (89ms ADAS)
- Adaptability: 0.5 (73% robustness)
- Interpretability: 0.5 (Attention maps)
- Efficiency: 0.5 (11W power)

**Method:** Protects road sign recognition through spatial attention hardening. **AS-ADS Score:** 2.5

*Diffusion Purification (Nie et al. 2022)*

- Real-Time: 0.25 (2.3s/image)
- Adaptability: 0.5 (68% purification)



- Interpretability: 0.5 (Process vis)
- Efficiency: 0.25 (24GB VRAM)

**Method:** Purifies inputs through multi-step diffusion denoising. **AS-ADS Score:** 1.875

*Trajectory Prediction Hardening (Zhang et al. 2022)*

- Real-Time: 0.5 (33ms prediction)
- Adaptability: 0.5 (65% robustness)
- Interpretability: 0.5 (Uncertainty bounds)
- Efficiency: 0.5 (8.7GB model)

**Method:** Improves trajectory prediction robustness through uncertainty-aware training. **AS-ADS Score:** 2.5

*Dynamic 3D Modeling (Wang et al. 2024a)*

- Real-Time: 1.0 (12ms modeling)
- Adaptability: 1.0 (96% cross-modal)
- Interpretability: 0.5 (Consistency reports)
- Efficiency: 1.0 (3.2W edge)

**Method:** Enables robust perception through dynamic neural feature modeling. **AS-ADS Score:** 4.375

## References

- Jérôme Rony and Luiz G. Hafemann and Luiz S. Oliveira and Ismail Ben Ayed and Robert Sabourin and Eric Granger : Decoupling direction and norm for efficient gradient-based l2 adversarial attacks and defenses. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 4317–4325 (2019). <https://doi.org/10.1109/CVPR.2019.00445>
- Abdu-Aguye, M.G., Gomaa, W., Makihara, Y., Yagi, Y.: Detecting adversarial attacks in time-series data. In: ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (2020). <https://doi.org/10.1109/icassp40776.2020.9053311>
- Almutairi, S., Barnawi, A.: Securing dnn for smart vehicles: An overview of adversarial attacks, defenses, and frameworks. Journal of Engineering and Applied Science **70**(1), 1–29 (2023) <https://doi.org/10.1186/s44147-023-00184-x>
- Andriushchenko, M., Croce, F., Flammarion, N., et al.: Square attack: A query-efficient black-box adversarial attack via random search. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (eds.) Computer Vision – ECCV 2020. Lecture Notes in Computer Science, vol. 12368. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-58592-1\\_29](https://doi.org/10.1007/978-3-030-58592-1_29)

- Athalye, A., Engstrom, L., Ilyas, A., Kwok, K.: Synthesizing robust adversarial example. arXiv preprint (2017) [arXiv:1707.07397](https://arxiv.org/abs/1707.07397) [cs.CV]
- Amirkhani, A., Karimi, M.P., Banitalebi-Dehkordi, A.: A survey on adversarial attacks and defenses for object detection and their applications in autonomous vehicles. *The Visual Computer* **39**, 5293–5307 (2023) <https://doi.org/10.1007/s00371-022-02660-6>
- Ai, S., Koe, A.S.V., Huang, T.: Adversarial perturbation in remote sensing image recognition. *Applied Soft Computing* **105**, 107252 (2021) <https://doi.org/10.1016/j.asoc.2021.107252>
- Akhtar, N., Mian, A., Kardan, N., Shah, M.: Advances in adversarial attacks and defenses in computer vision: A survey. arXiv preprint (2021) [arXiv:2108.00401](https://arxiv.org/abs/2108.00401) [cs.CV]
- Angelov, P., Soares, E.: Detecting and learning from unknown by extremely weak supervision: exploratory classifier (xclass). *Neural Comput & Applic* **33**, 15145–15157 (2021) <https://doi.org/10.1007/s00521-021-06137-w>
- Badjie, B., Cecilio, J., Casimiro, A.: Adversarial attacks and countermeasures on image classification-based deep learning models in autonomous driving systems: A systematic review. *ACM Computing Surveys* **57**(1), 20 (2024) <https://doi.org/10.1145/3691625>
- Bekey, G.A.: *Autonomous Robots: From Biological Inspiration to Implementation and Control (Intelligent Robotics and Autonomous Agents)*. The MIT Press, Cambridge, MA (2005). <https://doi.org/10.5555/1088950>
- Besl, P.J.: Active optical range imaging sensors. *Machine Vision and Applications* **1**(2), 127–152 (1988) <https://doi.org/10.1007/BF01212277>
- Bhagoji, A.N., He, W., Li, B., Song, D.: Exploring the space of black-box attacks on deep neural networks. arXiv preprint (2017) [arXiv:1712.09491](https://arxiv.org/abs/1712.09491) [cs.LG]
- Bansal, M., Krizhevsky, A., Ogale, A.: Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst. arXiv preprint (2018) [arXiv:1812.03079](https://arxiv.org/abs/1812.03079) [cs.RO]
- Bui, A.T., Le, T., Zhao, H., Montague, P., DeVel, O., Abraham, T., Phung, D.: Improving ensemble robustness by collaboratively promoting and demoting adversarial robustness. *Proceedings of the AAAI Conference on Artificial Intelligence* **35**(8), 6831–6839 (2021) <https://doi.org/10.1609/aaai.v35i8.16843>
- Brown, T.B., Mané, D., Roy, A., Abadi, M., Gilmer, J.: Adversarial patch. arXiv preprint (2018) [arXiv:1712.09665](https://arxiv.org/abs/1712.09665) [cs.CV]
- Biggio, B., Nelson, B.A., Laskov, P.: Poisoning attacks against support vector

- machines. In: Proceedings of the 29th International Conference on International Conference on Machine Learning. ICML'12, pp. 1467–1474. Omnipress, Madison, WI, USA (2012). <https://doi.org/10.5555/3042573.3042761>
- Boltachev, E.: Potential cyber threats of adversarial attacks on autonomous driving models. *Journal of Computer Virology and Hacking Techniques* **20**, 363–373 (2024) <https://doi.org/10.1007/s11416-023-00486-x>
- Brendel, W., Rauber, J., Bethge, M.: Decision-based adversarial attacks: Reliable attacks against black-box machine learning models. *arXiv preprint* (2017) [arXiv:1712.04248](https://arxiv.org/abs/1712.04248) [stat.ML]
- Bojarski, M., Testa, D.D., Dworakowski, D., Firner, B., Flepp, B., Goyal, P., Jackel, L.D., Monfort, M., Muller, U., Zhang, J., Zhang, X., Zhao, J., Zieba, K.: End to end learning for self-driving cars. *arXiv preprint* (2016) [arXiv:1604.07316](https://arxiv.org/abs/1604.07316) [cs.CV]
- Bochkovskiy, A., Wang, C.-Y., Liao, H.-Y.M.: Yolov4: Optimal speed and accuracy of object detection. *arXiv preprint* (2020) [arXiv:2004.10934](https://arxiv.org/abs/2004.10934) [cs.CV]
- Bojarski, M., Yeres, P., Choromanska, A., Choromanski, K., Firner, B., Jackel, L., Muller, U.: Explaining how a deep neural network trained with end-to-end learning steers a car. *arXiv preprint* (2017) [arXiv:1704.07911](https://arxiv.org/abs/1704.07911) [cs.CV]
- Croce, F., Andriushchenko, M., Sehwag, V., Debenedetti, E., Flammarion, N., Chiang, M., Mittal, P., Hein, M.: Robustbench: a standardized adversarial robustness benchmark. *arXiv preprint* (2020) [arXiv:2010.09670](https://arxiv.org/abs/2010.09670) [cs.LG]
- Chen, Y., Chu, S.: Adversarial defense in aerial detection. In: CVPR Workshop on Adversarial ML (2023). <https://doi.org/10.1109/cvprw59228.2023.00226>
- Chen, S.-T., Cornelius, C., Martin, J., Chau, D.H.: Shapeshifter: Robust physical adversarial attack on faster r-cnn object detector. *arXiv preprint* (2019) [arXiv:1804.05810](https://arxiv.org/abs/1804.05810) [cs.CV]
- Croce, F., Hein, M.: Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In: Proceedings of the 37th International Conference on Machine Learning, p. 206. JMLR.org, Virtual Event (2020). <https://doi.org/10.5555/3524938.3525144>
- Chen, X., Huang, W., Guo, W., Zhang, F., Du, J., Zhou, Z.: Adversarial defence by learning differentiated feature representation in deep ensemble. *Machine Vision and Applications* **35**(1), 88 (2024) <https://doi.org/10.1007/s00138-024-01571-x>
- Chen, J., Jordan, M.I., Wainwright, M.J.: Hopskipjumpattack: A query-efficient decision-based attack. In: 2020 IEEE Symposium on Security and Privacy (SP), pp. 1277–1294 (2020). <https://doi.org/10.1109/SP40000.2020.00045>

- Chen, E.-C., Lee, C.-R.: Towards fast and robust adversarial training for image classification. In: Computer Vision – ACCV 2020. Springer, Cham (2021). [https://doi.org/10.1007/978-3-030-69535-4\\_35](https://doi.org/10.1007/978-3-030-69535-4_35)
- Cheng, S., Liu, Y., Ma, S., Zhang, X.: Deep feature space trojan attack of neural networks by controlled detoxification. Proceedings of the AAAI Conference on Artificial Intelligence **35**(2), 1148–1156 (2021) <https://doi.org/10.1609/aaai.v35i2.16201>
- Codevilla, F., Müller, M., López, A., *et al.*: End-to-end driving via conditional imitation learning. In: Proceedings of the 2018 IEEE International Conference on Robotics and Automation (ICRA), pp. 4693–4700. IEEE, Brisbane, Australia (2018). <https://doi.org/10.1109/ICRA.2018.8460487>
- Cohen, J.M., Rosenfeld, E., Kolter, J.Z.: Certified adversarial robustness via randomized smoothing. arXiv preprint (2019) [arXiv:1902.02918](https://arxiv.org/abs/1902.02918) [cs.LG]
- Costa, J.C., Roxo, T., Proença, H., *et al.*: How deep learning sees the world: A survey on adversarial attacks & defenses. IEEE Access **12**, 61113–61136 (2024) <https://doi.org/10.1109/ACCESS.2024.3395118>
- Chen, P.-Y., Sharma, Y., Zhang, H., Yi, J., Hsieh, C.-J.: Ead: Elastic-net attacks to deep neural networks via adversarial examples. Proceedings of the AAAI Conference on Artificial Intelligence **32**(1) (2018) <https://doi.org/10.1609/aaai.v32i1.11302>
- Carlini, N., Wagner, D.: Adversarial examples are not easily detected: Bypassing ten detection methods. In: Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security. AISec ’17, pp. 3–14. Association for Computing Machinery, New York, NY, USA (2017). <https://doi.org/10.1145/3128572.3140444>
- Carlini, N., Wagner, D.: Towards evaluating the robustness of neural networks. In: 2017 IEEE Symposium on Security and Privacy (SP), pp. 39–57 (2017). <https://doi.org/10.1109/SP.2017.49>
- Chahe, A., Wang, C., Jeyapratap, A., Xu, K., Zhou, L.: Dynamic adversarial attacks on autonomous driving systems. arXiv preprint (2023) [arXiv:2312.06701](https://arxiv.org/abs/2312.06701) [cs.RO]
- Cao, Y., Wang, N., Xiao, C., Yang, D., Fang, J., Yang, R., Chen, Q.A., Liu, M., Li, B.: Demonstration: 3d adversarial object against msf-based perception in autonomous driving. In: Proceedings of the 3rd Conference on Machine Learning and Systems (MLSys) (2020). [https://me.ningfei.org/paper/MLsys\\_demo.pdf](https://me.ningfei.org/paper/MLsys_demo.pdf)
- Cao, Y., Xiao, C., Anandkumar, A., *et al.*: Advdo: Realistic adversarial attacks for trajectory prediction. In: Computer Vision – ECCV 2022, pp. 36–52. Springer, Cham (2022). [https://doi.org/10.1007/978-3-031-20065-6\\_3](https://doi.org/10.1007/978-3-031-20065-6_3)
- Cao, Y., Xiao, C., Cyr, B., Zhou, Y., Park, W., Rampazzi, S., Chen, Q.A., Fu, K., Mao,

- Z.M.: Adversarial sensor attack on lidar-based perception in autonomous driving. In: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (CCS '19), pp. 2267–2281. ACM, London, United Kingdom (2019). <https://doi.org/10.1145/3319535.3339815>
- Chen, H., Yan, H., Yang, X., Su, H., Zhao, S., Qian, F.: Efficient adversarial attack strategy against 3d object detection in autonomous driving. *IEEE Transactions on Intelligent Transportation Systems* **25**(11), 16118–16132 (2024) <https://doi.org/10.1109/TITS.2024.3410038>
- Chen, L.-C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: *Computer Vision – ECCV 2018*, pp. 801–818. Springer, Cham (2018). [https://doi.org/10.1007/978-3-030-01234-2\\_49](https://doi.org/10.1007/978-3-030-01234-2_49)
- Chen, P.-Y., Zhang, H., Sharma, Y., Yi, J., Hsieh, C.-J.: Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models. In: *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security. AISec '17*, pp. 15–26. Association for Computing Machinery, New York, NY, USA (2017). <https://doi.org/10.1145/3128572.3140448>
- Cao, J., Zhang, K.-Y., Yao, T., Ding, S., Yang, X., Ma, C.: Towards unified defense for face forgery and spoofing attacks via dual space reconstruction learning. *International Journal of Computer Vision* (2024) <https://doi.org/10.1007/s11263-024-02151-2>
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint* (2020) [arXiv:2010.11929](https://arxiv.org/abs/2010.11929) [cs.CV]
- Drucker, H., Cun, Y.L.: Improving generalization performance using double back-propagation. *IEEE Transactions on Neural Networks* **3**(6), 991–997 (1992) <https://doi.org/10.1109/72.165600>
- Du, A., Chen, B., Chin, T.-J., Law, Y.W., Sasdelli, M., Rajasegaran, R., Campbell, D.: Physical adversarial attacks on an aerial imagery object detector. In: *2022 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 3798–3808 (2022). <https://doi.org/10.1109/WACV51458.2022.00385>
- Dash, P., Chan, E., Pattabiraman, K.: Specguard: Specification aware recovery for robotic autonomous vehicles from physical attacks. In: *Proceedings of the ACM Conference on Computer and Communications Security (CCS)* (2024). <https://doi.org/10.1145/3658644.3690210>
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *2009 IEEE Conference on Computer Vision*

- and Pattern Recognition, pp. 248–255 (2009). <https://doi.org/10.1109/CVPR.2009.5206848>
- Dursun, H.E., Güven, Y., Kumbasar, T.: Imitation learning for autonomous driving: Insights from real-world testing. arXiv preprint (2025) [arXiv:2504.18847](https://arxiv.org/abs/2504.18847) [cs.RO]
- Deng, Y., Mu, T.: Understanding and improving ensemble adversarial defense. In: Proceedings of the 37th International Conference on Neural Information Processing Systems. NIPS '23. Curran Associates Inc., Red Hook, NY, USA (2023). <https://doi.org/10.5555/3666122.3668653>
- Du, X., Pun, C.-M., Zhang, Z.: A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In: Proceedings of the 32nd International Conference on Neural Information Processing Systems. NIPS'18, pp. 7167–7177. Curran Associates Inc., Red Hook, NY, USA (2018). <https://doi.org/10.5555/3327757.3327819>
- Deng, Y., Zhang, T., Lou, G., Zheng, X., Jin, J., Han, Q.-L.: Deep learning-based autonomous driving systems: A survey of attacks and defenses. IEEE Transactions on Industrial Informatics **17**(12), 7897–7912 (2021) <https://doi.org/10.1109/TII.2021.3071405>
- Edelkamp, S.: Adversarial Planning, pp. 325–335. Springer, Cham (2023). [https://doi.org/10.1007/978-3-319-65596-3\\_18](https://doi.org/10.1007/978-3-319-65596-3_18)
- Eykholt, K., Evtimov, I., Fernandes, E., Li, B., Rahmati, A., Xiao, C., Prakash, A., Kohno, T., Song, D.: Robust physical-world attacks on deep learning visual classification. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 1625–1634 (2018). <https://doi.org/10.1109/CVPR.2018.00175>
- Ehlers, R.: Formal verification of piece-wise linear feed-forward neural networks. arXiv preprint (2017) [arXiv:1705.01320](https://arxiv.org/abs/1705.01320) [cs.LO]
- Etim, A., Szefer, J.: Time traveling to defend against adversarial example attacks in image classification. arXiv preprint (2024) [arXiv:2410.08338](https://arxiv.org/abs/2410.08338) [cs.CR]
- Engstrom, L., Tran, B., Tsipras, D., Schmidt, L., Madry, A.: Exploring the landscape of spatial robustness. arXiv preprint (2019) [arXiv:1712.02779](https://arxiv.org/abs/1712.02779) [cs.LG]
- Feinman, R., Curtin, R.R., Shintre, S., Gardner, A.B.: Detecting adversarial samples from artifacts. arXiv preprint (2017) [arXiv:1703.00410](https://arxiv.org/abs/1703.00410) [stat.ML]
- Freitas, S., Chen, S.-T., Wang, Z.J., Chau, D.H.: Unmask: Adversarial detection and defense through robust feature alignment. In: 2020 IEEE International Conference on Big Data (Big Data) (2020). <https://doi.org/10.1109/bigdata50022.2020.9378303>

- Fletcher, R.: Practical Methods of Optimization, 2nd edn. Wiley, Hoboken, NJ, USA (2013). <https://doi.org/10.1002/9781118723203>
- Fu, C., Li, S., Yuan, X., Ye, J., Cao, Z., Ding, F.: Ad2attack: Adaptive adversarial attack on real-time uav tracking. In: 2022 International Conference on Robotics and Automation (ICRA), pp. 5893–5899 (2022). <https://doi.org/10.1109/ICRA46639.2022.9812056>
- Forsyth, D.A., Ponce, J.: Computer Vision: A Modern Approach, 2nd edn. Pearson, Boston (2011). <https://www.pearson.com/store/p/computer-vision-a-modern-approach/P100000687361/9780136085928>
- Gupta, K.D., Dasgupta, D., Akhtar, Z.: Adversarial input detection using image processing techniques (ipt). In: 2020 IEEE Annual Ubiquitous Computing, Electronics & Mobile Communication Conference (UEMCON), pp. 309–315. IEEE, Virtual Conference (2020). <https://doi.org/10.1109/UEMCON51285.2020.9298060>
- Gu, T., Dolan-Gavitt, B., Garg, S.: Badnets: Identifying vulnerabilities in the machine learning model supply chain. arXiv preprint (2017) [arXiv:1708.06733](https://arxiv.org/abs/1708.06733) [cs.CR]
- Gowal, S., Dvijotham, K., Stanforth, R., Bunel, R., Qin, C., Uesato, J., Arandjelovic, R., Mann, T., Kohli, P.: On the effectiveness of interval bound propagation for training verifiably robust models. arXiv preprint (2018) [arXiv:1810.12715](https://arxiv.org/abs/1810.12715) [cs.LG]
- Goldblum, M., Fowl, L., Feizi, S., Goldstein, T.: Adversarially robust distillation. Proceedings of the AAAI Conference on Artificial Intelligence **34**(04), 3996–4003 (2020) <https://doi.org/10.1609/aaai.v34i04.5816>
- Guo, C., Gardner, J.R., You, Y., Wilson, A.G., Weinberger, K.Q.: Simple black-box adversarial attacks. arXiv preprint (2019) [arXiv:1905.07121](https://arxiv.org/abs/1905.07121) [cs.LG]
- Girdhar, M., Hong, J., Moore, J.: Cybersecurity of autonomous vehicles: A systematic literature review of adversarial attacks and defense models. IEEE Open Journal of Vehicular Technology **PP**, 1–23 (2023) <https://doi.org/10.1109/OJVT.2023.3265363>
- Guesmi, A., Hanif, M.A., Shafique, M.: Advrain: Adversarial raindrops to attack camera-based smart vision systems. Information **14**(12), 634 (2023) <https://doi.org/10.3390/info14120634>
- Girshick, R.: Fast r-cnn. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 1440–1448 (2015). <https://doi.org/10.1109/ICCV.2015.169>
- Gu, T., Liu, K., Dolan-Gavitt, B., Garg, S.: Badnets: Evaluating backdooring attacks on deep neural networks. IEEE Access **7**, 47230–47244 (2019) <https://doi.org/10.1109/ACCESS.2019.2909068>



- Grosse, K., Manoharan, P., Papernot, N., Backes, M., McDaniel, P.: On the (statistical) detection of adversarial examples. arXiv preprint (2017) [arXiv:1702.06280](https://arxiv.org/abs/1702.06280) [cs.CR]
- Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. arXiv preprint (2014) [arXiv:1412.6572](https://arxiv.org/abs/1412.6572) [stat.ML]
- Guizzo, E.: How Google’s self-driving car works. IEEE Spectrum. Accessed: 2025-05-23 (2011). <https://spectrum.ieee.org/star-autonomous-surgical-robot>
- Gong, Y., Wang, S., Jiang, X., Yin, L., Sun, F.: Adversarial example detection using semantic graph matching. Applied Soft Computing **141**, 110317 (2023) <https://doi.org/10.1016/j.asoc.2023.110317>
- Guo, F., Zhao, Q., Li, X., Kuang, X., Zhang, J., Han, Y., Tan, Y.-a.: Detecting adversarial examples via prediction difference for deep neural networks. Information Sciences **501**, 182–192 (2019) <https://doi.org/10.1016/j.ins.2019.05.084>
- Hendrycks, D., Basart, S., Mu, N., Kadavath, S., Wang, F., Dorundo, E., Desai, R., Zhu, T., Parajuli, S., Guo, M., Song, D., Steinhardt, J., Gilmer, J.: The many faces of robustness: A critical analysis of out-of-distribution generalization. In: 2021 IEEE/CVF International Conference on Computer Vision (ICCV) (2021). <https://doi.org/10.1109/iccv48922.2021.00823>
- Hu, Z., Chu, W., Zhu, X., Zhang, H., Zhang, B., Hu, X.: Physically realizable natural-looking clothing textures evade person detectors via 3d modeling. arXiv preprint (2023) [arXiv:2307.01778](https://arxiv.org/abs/2307.01778) [cs.CV]
- Hendrycks, D., Dietterich, T.: Benchmarking neural network robustness to common corruptions and perturbations. arXiv preprint (2019) [arXiv:1903.12261](https://arxiv.org/abs/1903.12261) [cs.LG]
- He, K., Gkioxari, G., Dollár, P., *et al.*: Mask r-cnn. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV), pp. 2980–2988 (2017). <https://doi.org/10.1109/ICCV.2017.322>
- Huang, W.R., Geiping, J., Fowl, L., *et al.*: Metapoisn: Practical general-purpose clean-label data poisoning. In: Proceedings of the 34th International Conference on Neural Information Processing Systems. NIPS’20, p. 1013. Curran Associates Inc., Red Hook, NY, USA (2020). <https://doi.org/10.5555/3495724.3496737>
- Hussain, M., Hong, J.-E.: Reconstruction-based adversarial attack detection in vision-based autonomous driving systems. Machine Learning and Knowledge Extraction **5**(4), 1589–1611 (2023) <https://doi.org/10.3390/make5040080>
- Hanfeld, P., Höhne, M.M.-C., Bussmann, M., *et al.*: Flying adversarial patches: Manipulating the behavior of deep learning-based autonomous multirotors. arXiv preprint (2023) [arXiv:2305.12859](https://arxiv.org/abs/2305.12859) [cs.RO]

- Hsiao, T.-F., Huang, B.-L., Ni, Z.-X., Lin, Y.-T., Shuai, H.-H., Li, Y.-H., Cheng, W.-H.: Natural light can also be dangerous: Traffic sign misinterpretation under adversarial natural light attacks. In: 2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), pp. 3903–3912 (2024). <https://doi.org/10.1109/WACV57701.2024.00387>
- Huang, B., Ke, Z., Wang, Y., Wang, W., Shen, L., Liu, F.: Adversarial defence by diversified simultaneous training of deep ensembles. *Proceedings of the AAAI Conference on Artificial Intelligence* **35**(9), 7823–7831 (2021) <https://doi.org/10.1609/aaai.v35i9.16955>
- Hao, C., Orlando, D., Liu, J., Yin, C.: *Introduction to Radar Systems*, 3rd edn. McGraw-Hill Education, New York, NY, USA (2002). [https://doi.org/10.1007/978-981-16-6399-4\\_1](https://doi.org/10.1007/978-981-16-6399-4_1)
- Horton, E., Ranganathan, P.: Development of a gps spoofing apparatus to attack a dji matrice 100 quadcopter. *Journal of Global Positioning Systems* **16**(9) (2018) <https://doi.org/10.1186/s41445-018-0018-3>
- Hsu, J.-M.: *Introduction to Global Satellite Positioning System (GPS)*. Artech House, Boston, MA (2002). <https://doi.org/10.4018/978-1-60566-840-6.ch007>
- Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. *arXiv preprint* (2015) [arXiv:1503.02531](https://arxiv.org/abs/1503.02531) [stat.ML]
- Han, X., Wang, H., Zhao, K., Deng, G., Xu, Y., Liu, H., Qiu, H., Zhang, T.: Visionguard: Secure and robust visual perception of autonomous vehicles in practice. In: *CCS* (2024). <https://doi.org/10.1145/3658644.3670296>
- Hu, X., Ying, Q., Qian, Z., Li, S., Zhang, X.: Draw: Defending camera-shooted raw against image manipulation. In: 2023 IEEE/CVF International Conference on Computer Vision (ICCV), pp. 22377–22387 (2023). <https://doi.org/10.1109/iccv51070.2023.02050>
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 770–778 (2016). <https://doi.org/10.1109/CVPR.2016.90>
- Hong, H., Zhang, X., Wang, B., Ba, Z., Hong, Y.: Certifiable black-box attacks with randomized adversarial examples: Breaking defenses with provable confidence. In: *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, pp. 600–614 (2024). <https://doi.org/10.1145/3658644.3690343>
- Ibrahim, A.D.M., Hussain, M., Hong, J.-E.: Deep learning adversarial attacks and defenses in autonomous vehicles: A systematic literature review from a safety perspective. *Artificial Intelligence Review* **58**(28) (2024) <https://doi.org/10.1007/>

- Janai, J., Güney, F., Behl, A., Geiger, A.: Computer vision for autonomous vehicles: Problems, datasets and state of the art. *Foundations and Trends in Computer Graphics and Vision* **12**(1–3), 1–308 (2020) <https://doi.org/10.1561/06000000079>
- Jallepalli, D., Ravikumar, N.C., Badarinath, P.V., Uchil, S., Suresh, M.A.: Federated learning for object detection in autonomous vehicles. In: 2021 IEEE Seventh International Conference on Big Data Computing Service and Applications (BigDataService), pp. 107–114 (2021). <https://doi.org/10.1109/BigDataService52369.2021.00018>
- Jing, L., Wang, R., Ren, W., Dong, X., Zou, C.: Pad: Patch-agnostic defense against adversarial patch attacks. In: 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 24472–24481 (2024). <https://doi.org/10.1109/cvpr52733.2024.02310>
- Jiang, W., Wang, L., Zhang, T., Chen, Y., Dong, J., Bao, W., Zhang, Z., Fu, Q.: Robuste2e: Exploring the robustness of end-to-end autonomous driving. *Electronics* **13**(16), 3299 (2024) <https://doi.org/10.3390/electronics13163299>
- Khamaiseh, S.Y., Bagagem, D., Al-Alaj, A., Mancino, M., Alomari, H.W.: Adversarial deep learning: A survey on adversarial attacks and defense mechanisms on image classification. *IEEE Access* **10**, 102266–102291 (2022) <https://doi.org/10.1109/ACCESS.2022.3208131>
- Koenig, S., Bonet, B., Cavazza, M., desJardins, M., Felner, A., Hawes, N., Knox, B., Konidaris, G., Lang, J., López, C.L., Magazzeni, D., McGovern, A., Natarajan, S., Sturtevant, N.R., Thielscher, M., Yeoh, W., Sardina, S., Wagstaff, K.: Using machine teaching to identify optimal training-set attacks on machine learners. In: Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence. AAAI’15, pp. 2871–2877. AAAI Press, Austin, Texas (2015). <https://doi.org/10.5555/2886521.2886721>
- Katz, G., Barrett, C., Dill, D.L., *et al.*: Reluplex: An efficient smt solver for verifying deep neural networks. In: Computer Aided Verification. CAV 2017. Lecture Notes in Computer Science, vol. 10426. Springer, cham (2017). <https://doi.org/10.1007/978-3-319-63387-9>
- Kinsler, L.E., Frey, A.R., Coppens, A.B., *et al.*: Fundamentals of Acoustics, 4th edn. John Wiley & Sons, Wiley Online Library (2000). [https://doi.org/10.1007/978-3-540-48830-9\\_2](https://doi.org/10.1007/978-3-540-48830-9_2)
- Kurakin, A., Goodfellow, I., Bengio, S.: Adversarial examples in the physical world. arXiv preprint (2016) [arXiv:1607.02533](https://arxiv.org/abs/1607.02533) [cs.CV]
- Kong, Z., Guo, J., Li, A., Liu, C.: Physgan: Generating physical-world-resilient

- adversarial examples for autonomous driving. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020). <https://doi.org/10.1109/cvpr42600.2020.01426>
- Kannan, H., Kurakin, A., Goodfellow, I.: Adversarial logit pairing. arXiv preprint (2018) [arXiv:1803.06373](https://arxiv.org/abs/1803.06373) [cs.LG]
- Kang, H., Kwak, B.I., Lee, Y.H., Lee, H., Lee, H., Kim, H.K.: Car Hacking: Attack & Defense Challenge 2020 Dataset. <https://doi.org/10.21227/qvr7-n418>
- Khan, I.A., Moustafa, N., Pi, D., Haider, W., Li, B., Jolfaei, A.: An enhanced multi-stage deep learning framework for detecting malicious activities from autonomous vehicles. *IEEE Transactions on Intelligent Transportation Systems* **23**(12), 25469–25478 (2022) <https://doi.org/10.1109/TITS.2021.3105834>
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.-Y., Dollár, P., Girshick, R.: Segment anything. In: 2023 IEEE/CVF International Conference on Computer Vision (ICCV) (2023). <https://doi.org/10.1109/iccv51070.2023.00371>
- Knee, P.: Radar Signal Processing Fundamentals. McGraw-Hill, New York, NY, USA (2005). [https://doi.org/10.1007/978-3-031-01519-9\\_4](https://doi.org/10.1007/978-3-031-01519-9_4)
- Komkov, S., Petiushko, A.: Advhat: Real-world adversarial attack on arcface face id system. In: 2020 25th International Conference on Pattern Recognition (ICPR), pp. 819–826 (2021). <https://doi.org/10.1109/icpr48806.2021.9412236>
- Krizhevsky, A.: Learning multiple layers of features from tiny images. Technical report, University of Toronto (2009). <https://www.cs.toronto.edu/~text/tildelowlkriz/learning-features-2009-TR.pdf>
- King, I., Wang, J.: Provably robust deep learning via adversarially trained smoothed classifiers. In: Proceedings of the 33rd International Conference on Neural Information Processing Systems, p. 1013. Curran Associates Inc., Red Hook, NY, USA (2019). <https://doi.org/10.5555/3454287.3455300>
- Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.-Y., Berg, A.C.: Ssd: Single shot multibox detector. In: Computer Vision – ECCV 2016, pp. 21–37. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-46448-0\\_2](https://doi.org/10.1007/978-3-319-46448-0_2)
- Lecuyer, M., Atlidakis, V., Geambasu, R., Hsu, D., Jana, S.: Certified robustness to adversarial examples with differential privacy. In: 2019 IEEE Symposium on Security and Privacy (SP), pp. 656–672 (2019). <https://doi.org/10.1109/SP.2019.00044>
- Li, Y., Angelov, P., Suri, N.: Self-supervised representation learning for adversarial attack detection. In: Computer Vision – ECCV 2024, pp. 236–252 (2024). [https://doi.org/10.1007/978-3-031-73027-6\\_14](https://doi.org/10.1007/978-3-031-73027-6_14)

- Li, Y., Angelov, P., Yu, Z., Pellicer, A.L., Suri, N.: Federated adversarial learning for robust autonomous landing runway detection. In: Artificial Neural Networks and Machine Learning – ICANN 2024. Lecture Notes in Computer Science, vol. 15021, pp. 159–173. Springer, Cham (2024). [https://doi.org/10.1007/978-3-031-72347-6\\_11](https://doi.org/10.1007/978-3-031-72347-6_11)
- Lecun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**(11), 2278–2324 (1998) <https://doi.org/10.1109/5.726791>
- Liu, Y., Chen, X., Liu, C., Song, D.: Delving into transferable adversarial examples and black-box attacks. *arXiv preprint* (2016) [arXiv:1611.02770](https://arxiv.org/abs/1611.02770) [cs.LG]
- Liu, N., Du, M., Guo, R., Liu, H., Hu, X.: Adversarial attacks and defenses: An interpretation perspective. *SIGKDD Explor. Newsl.* **23**(1), 86–99 (2021) <https://doi.org/10.1145/3468507.3468519>
- Li, Y., Fang, Q., Bai, J., Chen, S., Juefei-Xu, F., Feng, C.: Among us: Adversarially robust collaborative perception by consensus. In: *IEEE International Conference on Computer Vision (ICCV)* (2023). <https://doi.org/10.1109/iccv51070.2023.00024>
- Lin, T.-Y., Goyal, P., Girshick, R., He, K., Dollar, P.: Focal loss for dense object detection. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2980–2988 (2017). <https://doi.org/10.1109/ICCV.2017.324>
- Lillicrap, T.P., Hunt, J.J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., Wierstra, D.: Continuous control with deep reinforcement learning. *arXiv preprint* (2015) [arXiv:1509.02971](https://arxiv.org/abs/1509.02971) [cs.LG]
- Li, X., Li, J., Chen, Y., Ye, S., He, Y., Wang, S., Su, H., Xue, H.: Qair: Practical query-efficient black-box attacks for image retrieval. In: *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3329–3338 (2021). <https://doi.org/10.1109/CVPR46437.2021.00334>
- Liao, F., Liang, M., Dong, Y., Pang, T., Hu, X., Zhu, J.: Defense against adversarial attacks using high-level representation guided denoiser. In: *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1778–1787 (2018). <https://doi.org/10.1109/CVPR.2018.00191>
- Liu, J., Levine, A., Lau, C.P., Chellappa, R., Feizi, S.: Segment and Complete: Defending Object Detectors Against Adversarial Patch Attacks With Robust Patch Detection. In: *CVPR*, pp. 14973–14982 (2022). <https://doi.org/10.1109/cvpr52688.2022.01455>
- Li, Z., Li, H., Xie, E., Sima, C., Lu, T., Qiao, Y., Dai, J.: Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. In: *Computer Vision – ECCV 2022*, pp. 1–18. Springer, Cham (2022). [https://doi.org/10.1007/978-3-031-19818-9\\_1](https://doi.org/10.1007/978-3-031-19818-9_1)

[org/10.1007/978-3-031-20077-9\\_1](https://doi.org/10.1007/978-3-031-20077-9_1)

- Li, L., Qiu, J., Spratling, M.: Aroid: Improving adversarial robustness through online instance-wise data augmentation. *International Journal of Computer Vision* **132**, 1–20 (2024) <https://doi.org/10.1007/s11263-024-02206-4>
- Lu, X., Radha, H.: ScAR: Scaling Adversarial Robustness for LiDAR Object Detection. In: *Proc. of IROS* (2023). <https://doi.org/10.1109/iros55552.2023.10341583>
- Lu, J., Sibai, H., Fabry, E., Forsyth, D.: No need to worry about adversarial examples in object detection in autonomous vehicles. *arXiv preprint* (2017) [arXiv:1707.03501](https://arxiv.org/abs/1707.03501) [cs.CV]
- Lu, Z., Sun, H., Ji, K., Kuang, G.: Adversarial robust aerial image recognition based on reactive-proactive defense framework with deep ensembles. *Remote Sensing* **15**(19), 4660 (2023) <https://doi.org/10.3390/rs15194660>
- Lou, Y., Song, Q., Xu, Q., Tan, R., Wang, J.: Uncertainty-Encoded Multi-Modal Fusion for Robust Object Detection in Autonomous Driving. In: *Proc. of 26th European Conference on Artificial Intelligence (ECAI)* (2023). <https://doi.org/10.3233/faia230441>
- Lu, Z., Sun, H., Xu, Y.: Adversarial robustness enhancement of uav-oriented automatic image recognition via ensemble defense. *Remote Sensing* **15**(12), 3007 (2023) <https://doi.org/10.3390/rs15123007>
- Lempitsky, V., Vedaldi, A., Ulyanov, D.: Deep image prior. In: *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9446–9454 (2018). <https://doi.org/10.1109/CVPR.2018.00984>
- Meng, D., Chen, H.: Magnet: A two-pronged defense against adversarial examples. In: *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security. CCS '17*, pp. 135–147. Association for Computing Machinery, New York, NY, USA (2017). <https://doi.org/10.1145/3133956.3134057>
- Moosavi-Dezfooli, S.-M., Fawzi, A., Frossard, P.: Deepfool: A simple and accurate method to fool deep neural networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2574–2582 (2016). <https://doi.org/10.1109/CVPR.2016.282>
- Moosavi-Dezfooli, S.-M., Fawzi, A., Fawzi, O., Frossard, P.: Universal adversarial perturbations. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 86–94 (2017). <https://doi.org/10.1109/CVPR.2017.17>
- Morgulis, N., Kreines, A., Mendelowitz, S., Weisglass, Y.: Fooling a real car with adversarial traffic signs. *arXiv preprint* (2019) [arXiv:1907.00374](https://arxiv.org/abs/1907.00374) [cs.CR]

- Man, Y., Li, M., Gerdes, R.: Ghostimage: Remote perception attacks against camera-based image classification systems. In: 23rd International Symposium on Research in Attacks, Intrusions and Defenses (RAID 2020), pp. 317–332. USENIX Association, Virtual Conference (2020). <https://www.usenix.org/system/files/raid20-man.pdf>
- Man, Y., Li, M., Gerdes, R.: Remote perception attacks against camera-based object recognition systems. In: Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (CCS '23), pp. 14–11422 (2023). <https://doi.org/10.1145/3596221>
- Mani, N., Moh, M., Moh, T.-S.: Towards robust ensemble defense against adversarial examples attack. In: 2019 IEEE Global Communications Conference (GLOBECOM), pp. 1–6. IEEE, Hawaii, USA (2019). <https://doi.org/10.1109/GLOBECOM38437.2019.9013408>
- Malik, J., Muthalagu, R., Pawar, P.M.: A systematic review of adversarial machine learning attacks, defensive controls, and technologies. IEEE Access (2024) <https://doi.org/10.1109/ACCESS.2024.3423323>
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. arXiv preprint (2019) [arXiv:1706.06083](https://arxiv.org/abs/1706.06083) [stat.ML]
- Muravev, N., Petiushko, A.: Certified robustness via randomized smoothing over multiplicative parameters of input transformations. In: Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, pp. 3366–3372 (2022). <https://doi.org/10.24963/ijcai.2022/467>
- Mu, J.: A real-time defense against object vanishing adversarial patch attacks for object detection in autonomous vehicles. arXiv preprint (2024) [arXiv:2412.06215](https://arxiv.org/abs/2412.06215) [cs.CV]
- Ma, S., Vemprala, S., Wang, W., Gupta, J.K., Song, Y., McDuff, D., Kapoor, A.: Compass: Contrastive multimodal pretraining for autonomous systems. arXiv preprint (2022) [arXiv:2203.15788](https://arxiv.org/abs/2203.15788) [cs.RO]
- Nie, W., Guo, B., Huang, Y., Xiao, C., Vahdat, A., Anandkumar, A.: Diffusion Models for Adversarial Purification
- noz-González, L.M., Pfitzner, B., Russo, M., Carnerero-Cano, J., Lupu, E.C.: Poisoning attacks with generative adversarial nets. arXiv preprint (2019) [arXiv:1906.07773](https://arxiv.org/abs/1906.07773) [cs.LG]
- Naseer, M., Khan, S.H., Porikli, F.: Local gradients smoothing: Defense against localized adversarial attacks. arXiv preprint (2018) [arXiv:1807.01216](https://arxiv.org/abs/1807.01216) [cs.CV]
- Nesti, F., Rossolini, G., D’Amico, G., Biondi, A., Buttazzo, G.: Carla-gear: a dataset

- generator for a systematic evaluation of adversarial robustness of vision models. arXiv preprint (2022) [arXiv:2206.04365](#) [cs.CV]
- Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., Assran, M., Ballas, N., Galuba, W., Howes, R., Huang, P.-Y., Li, S.-W., Misra, I., Rabbat, M., Sharma, V., Synnaeve, G., Xu, H., Jegou, H., Mairal, J., Labatut, P., Joulin, A., Bojanowski, P.: Dinov2: Learning robust visual features without supervision. arXiv preprint (2024) [arXiv:2304.07193](#) [cs.CV]
- Oslund, S., Washington, C., So, A., Chen, T., Ji, H.: Multiview robust adversarial stickers for arbitrary objects in the physical world. *Journal of Computational and Cognitive Engineering* **1**(4), 152–158 (2022) <https://doi.org/10.47852/bonviewJCCE2202322>
- Pomerleau, A. D.: Alvin: An autonomous land vehicle in a neural network. In: *Proceedings of the 2nd International Conference on Neural Information Processing Systems*, pp. 305–313 (1988). <https://doi.org/10.5555/2969735.2969771>
- Pan, Y., Cheng, C.-A., Saigol, K., Lee, K., Yan, X., Theodorou, E., Boots, B.: Agile autonomous driving using end-to-end deep imitation learning. arXiv preprint (2017) [arXiv:1709.07174](#) [cs.RO]
- Pellicer, A.L., Giatgong, K., Li, Y., Suri, N., Angelov, P.: Unicad: A unified approach for attack detection, noise reduction and novel class identification. In: *2024 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8 (2024). <https://doi.org/10.1109/ijcnn60899.2024.10651159>
- Pellcier, A.L., Li, Y., Angelov, P.: PUDD: Towards robust multi-modal prototype-based deepfake detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pp. 3809–3817 (2024). <https://doi.org/10.1109/CVPRW63382.2024.00385>
- Poggi, M., Mattoccia, S.: Learning to predict stereo reliability enforcing local consistency of confidence maps. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4541–4550 (2017). <https://doi.org/10.1109/CVPR.2017.483>
- Papernot, N., McDaniel, P., Goodfellow, I.: Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. arXiv preprint (2016) [arXiv:1605.07277](#) [cs.CR]
- Papernot, N., McDaniel, P., Jha, S., Fredrikson, M., Celik, Z.B., Swami, A.: The limitations of deep learning in adversarial settings. In: *2016 IEEE European Symposium on Security and Privacy (EuroS&P)*, pp. 372–387 (2016). <https://doi.org/10.1109/EuroSP.2016.36>



- Papernot, N., McDaniel, P., Wu, X., Jha, S., Swami, A.: Distillation as a defense to adversarial perturbations against deep neural networks. In: 2016 IEEE Symposium on Security and Privacy (SP), pp. 582–597 (2016). <https://doi.org/10.1109/SP.2016.41>
- Prevot, T., Rios, J., Kopardekar, P., III, J.E.R., Johnson, M., Jung, J.: Uas traffic management (utm) concept of operations to safely enable low altitude flight operations. In: 16th AIAA Aviation Technology, Integration, and Operations Conference (2016). <https://doi.org/10.2514/6.2016-3292>
- Pourkeshavarz, M., Sabokrou, M., Rasouli, A.: Adversarial backdoor attack by naturalistic data poisoning on trajectory prediction in autonomous driving. In: 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 14885–14894 (2024). <https://doi.org/10.1109/CVPR52733.2024.01410>
- Queyrut, S., Schiavoni, V., Felber, P.: Mitigating adversarial attacks in federated learning with trusted execution environments. arXiv preprint (2023) [arXiv:2309.07197](https://arxiv.org/abs/2309.07197) [cs.LG]
- Reyes-Amezcu, I., Ochoa-Ruiz, G., Mendez-Vazquez, A.: Enhancing image classification robustness through adversarial sampling with delta data augmentation (dda). In: 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pp. 274–283 (2024). <https://doi.org/10.1109/CVPRW63382.2024.00032>
- Renz, K., Chen, L., Marcu, A.-M., Hünemann, J., Hanotte, B., Karnsund, A., Shotton, J., Arani, E., Sinavski, O.: Carllava: Vision language models for camera-only closed-loop driving. arXiv preprint (2024) [arXiv:2406.10165](https://arxiv.org/abs/2406.10165) [cs.CV]
- Ross, A., Doshi-Velez, F.: Improving the adversarial robustness and interpretability of deep neural networks by regularizing their input gradients. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 32 (2018). <https://doi.org/10.1609/aaai.v32i1.11504>
- Rozsa, A., Rudd, E.M., Boulton, T.E.: Adversarial diversity and hard positive generation. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pp. 410–417 (2016). <https://doi.org/10.1109/CVPRW.2016.58>
- Ren, Shaoqing, He, *et al.*: Faster r-cnn: Towards real-time object detection with region proposal networks. In: Proceedings of the 29th International Conference on Neural Information Processing Systems - Volume 1, pp. 91–99 (2015). <https://doi.org/10.5555/2969239.2969250>
- Raghunathan, A., Steinhardt, J., Liang, P.: Certified defenses against adversarial examples. arXiv preprint (2018) [arXiv:1801.09344](https://arxiv.org/abs/1801.09344) [cs.LG]

- Soares, E., Angelov, P., Suri, N.: Similarity-based deep neural network to detect imperceptible adversarial attacks. In: 2022 IEEE Symposium Series on Computational Intelligence (SSCI), pp. 1028–1035 (2022). <https://doi.org/10.1109/SSCI51031.2022.10022016>
- Sharif, M., Bhagavatula, S., Bauer, L., Reiter, M.K.: 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. CCS '16, pp. 1528–1540. Association for Computing Machinery, New York, NY, USA (2016). <https://doi.org/10.1145/2976749.2978392>
- Sharif, M., Bhagavatula, S., Bauer, L., Reiter, M.K.: A general framework for adversarial examples with objectives. *ACM Trans. Priv. Sec.* **1**(1), 1–30 (2019) <https://doi.org/10.1145/3317611>
- Shi, L., Chen, Z., Shi, Y., Zhao, G., Wei, L., Tao, Y., Gao, Y.: Data poisoning attacks on federated learning by using adversarial samples. In: 2022 International Conference on Computer Engineering and Artificial Intelligence (ICCEAI), pp. 158–162 (2022). <https://doi.org/10.1109/ICCEAI55464.2022.00041>
- Shan, S., Ding, W., Passananti, J., Wu, S., Zheng, H., Zhao, B.Y.: Nightshade: Prompt-specific poisoning attacks on text-to-image generative models. In: 2024 IEEE Symposium on Security and Privacy (SP), pp. 807–825 (2024). <https://doi.org/10.1109/sp54263.2024.00207>
- Shah, A.: Adversary ml resilience in autonomous driving through human-centered perception mechanisms. arXiv preprint (2023) [arXiv:2311.01478](https://arxiv.org/abs/2311.01478) [cs.CV]
- Sheridan, T.B.: Human–robot interaction: status and challenges. *Human Factors* **58**(4), 525–532 (2016) <https://doi.org/10.1177/0018720816644364>
- Shibly, K.H., Hossain, M.D., Inoue, H., Taenaka, Y., Kadobayashi, Y.: Towards autonomous driving model resistant to adversarial attack. *Applied Artificial Intelligence* **37**(1), 2193461 (2023) <https://doi.org/10.1080/08839514.2023.2193461>
- Siciliano, B., Khatib, O. (eds.): *Springer Handbook of Robotics*, 2nd. edn. Springer, Cham (2016). <https://doi.org/10.1007/978-3-319-32552-1>
- Sabokrou, M., Khalooei, M., Adeli, E.: Be your own neighborhood: detecting adversarial examples by the neighborhood relations built on self-supervised learning. In: Proceedings of the 41st International Conference on Machine Learning (ICML 2024), p. 2794. JMLR.org, Vienna, Austria (2024). <https://doi.org/10.5555/3692070.3692794>
- Samangouei, P., Kabkab, M., Chellappa, R.: Defense-gan: Protecting classifiers against adversarial attacks using generative models. arXiv preprint (2018) [arXiv:1805.06605](https://arxiv.org/abs/1805.06605) [cs.CV]

- Song, R., Ozmen, M.O., Kim, H., *et al.*: Discovering adversarial driving maneuvers against autonomous vehicles. In: 32nd USENIX Security Symposium (USENIX Security 23). USENIX Association, Anaheim, CA, USA (2023). <https://www.usenix.org/system/files/usenixsecurity23-song.pdf>
- Shu, M., Shen, Y., Lin, M.C., Goldstein, T.: Adversarial differentiable data augmentation for autonomous systems. In: 2021 IEEE International Conference on Robotics and Automation (ICRA), pp. 1032–1038. IEEE, Xi’an, China (2021). <https://doi.org/10.1109/ICRA48506.2021.9561205>
- Su, J., Vargas, D.V., Sakurai, K.: One pixel attack for fooling deep neural networks. *IEEE Transactions on Evolutionary Computation* **23**(5), 828–841 (2019) <https://doi.org/10.1109/TEVC.2019.2890858>
- Szeliski, R.: *Computer Vision: Algorithms and Applications*, 2nd edn. Springer, Cham (2022). <https://doi.org/10.1007/978-3-030-34372-9>
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., Fergus, R.: Intriguing properties of neural networks. arXiv preprint (2013) [arXiv:1312.6199](https://arxiv.org/abs/1312.6199) [cs.CV]
- Shen, Y., Zheng, L., Shu, M., *et al.*: Gradient-free adversarial training against image corruption for autonomous driving. In: Proceedings of the 35th International Conference on Neural Information Processing Systems, pp. 26250–26263. Curran Associates Inc., Red Hook, NY, USA (2021). <https://doi.org/10.5555/3540261.3542271>
- Tramèr, F., Boneh, D.: Adversarial training and robustness for multiple perturbations. In: Proceedings of the 33rd International Conference on Neural Information Processing Systems, p. 527. Curran Associates Inc., Red Hook, NY, USA (2019). <https://doi.org/10.5555/3454287.3454814>
- Tesla, Inc.: Replacing Ultrasonic Sensors with Tesla Vision. Accessed: 2025-05-21 (2022). <https://www.tesla.com/support/transitioning-tesla-vision>
- Tian, X., Gu, J., Li, B., Liu, Y., Wang, Y., Zhao, Z., Zhan, K., Jia, P., Lang, X., Zhao, H.: Drivevlm: The convergence of autonomous driving and large vision-language models. arXiv preprint (2024) [arXiv:2402.12289](https://arxiv.org/abs/2402.12289) [cs.CV]
- Tarchoun, B., Khalifa, A.B., Mahjoub, M.A., Abu-Ghazaleh, N., Alouani, I.: Jedi: Entropy-based localization and removal of adversarial patches. In: CVPR, pp. 4087–4095 (2023). <https://doi.org/10.1109/cvpr52729.2023.00398>
- Tramèr, F., Kurakin, A., Papernot, N., Goodfellow, I., Boneh, D., McDaniel, P.: Ensemble adversarial training: Attacks and defences. arXiv preprint (2017) [arXiv:1705.07204](https://arxiv.org/abs/1705.07204) [stat.ML]

- Tan, M., Pang, R., Le, Q.V.: Efficientdet: Scalable and efficient object detection. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 10781–10790 (2020). <https://doi.org/10.1109/CVPR42600.2020.01079>
- Thys, S., Ranst, W.V., Goedemé, T.: Fooling automated surveillance cameras: Adversarial patches to attack person detection. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pp. 49–55 (2019). <https://doi.org/10.1109/CVPRW.2019.00012>
- Tian, J., Wang, B., Guo, R., Wang, Z., Cao, K., Wang, X.: Adversarial attacks and defenses for deep-learning-based unmanned aerial vehicles. *IEEE Internet of Things Journal* **9**(22), 22399–22409 (2022) <https://doi.org/10.1109/JIOT.2021.3111024>
- Tjeng, V., Xiao, K., Tedrake, R.: Evaluating robustness of neural networks with mixed integer programming. arXiv preprint (2017) [arXiv:1711.07356](https://arxiv.org/abs/1711.07356) [cs.LG]
- Toheed, A., Yousaf, M.H., Rabnawaz, Javed, A.: Physical adversarial attack scheme on object detectors using 3d adversarial object. In: 2022 2nd International Conference on Digital Futures and Transformative Technologies (ICoDT2), pp. 1–4 (2022). <https://doi.org/10.1109/ICoDT255437.2022.9787422>
- Vincent, P., Larochelle, H., Bengio, Y., Manzagol, P.-A.: Extracting and composing robust features with denoising autoencoders. In: Proceedings of the 25th International Conference on Machine Learning. ICML '08, pp. 1096–1103. Association for Computing Machinery, New York, NY, USA (2008). <https://doi.org/10.1145/1390156.1390294>
- Wang, C.-Y., Bochkovskiy, A., Liao, H.-Y.M.: Yolov7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. In: 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 7464–7475 (2023). <https://doi.org/10.1109/cvpr52729.2023.00721>
- Wang, X., Cai, M., Sohel, F., Sang, N., Chang, Z.: Adversarial point cloud perturbations against 3d object detection in autonomous driving systems. *Neurocomputing* **466**, 27–36 (2021) <https://doi.org/10.1016/j.neucom.2021.09.027>
- Wu, Y., Du, J., Liu, P., Lin, Y., Xu, W., Cheng, W.: Dd-robustbench: An adversarial robustness benchmark for dataset distillation. *IEEE Transactions on Image Processing* **34**, 2052–2066 (2025) <https://doi.org/10.1109/tip.2025.3553786>
- Wu, Z., Lim, S.-N., Davis, L.S., *et al.*: Making an invisibility cloak: Real world adversarial attacks on object detectors. In: Computer Vision – ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV, pp. 1–17. Springer, Berlin, Hjournalenberg (2020). [https://doi.org/10.1007/978-3-030-58548-8\\_1](https://doi.org/10.1007/978-3-030-58548-8_1)

- Wang, N., Luo, Y., Sato, T., Xu, K., Chen, Q.A.: Does physical adversarial example really matter to autonomous driving? towards system-level effect of adversarial object evasion attack. In: 2023 IEEE/CVF International Conference on Computer Vision (ICCV), pp. 4389–4400 (2023). <https://doi.org/10.1109/iccv51070.2023.00407>
- Wang, T., Lu, F., Zheng, Z., Chen, G., Jiang, C.: Rcdn: Towards robust camera-insensitivity collaborative perception via dynamic feature-based 3d neural modeling. In: Proceedings of the 38th Conference on Neural Information Processing Systems (NeurIPS) (2024). [https://proceedings.neurips.cc/paper\\_files/paper/2024/file/27e5626cabdbb6cd5c56ce4114ff93e4-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/27e5626cabdbb6cd5c56ce4114ff93e4-Paper-Conference.pdf)
- Wang, Z., Li, X., Zhu, H., *et al.*: Revisiting adversarial training at scale. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 24675–24685 (2024). <https://doi.org/10.1109/CVPR52733.2024.02330>
- Wong, E., Rice, L., Kolter, J.Z.: Fast is better than free: Revisiting adversarial training. arXiv preprint (2020) [arXiv:2001.03994](https://arxiv.org/abs/2001.03994) [cs.LG]
- Wan, Z., Shen, J., Chuang, J., Xia, X., Garcia, J., Ma, J., Chen, Q.A.: Too afraid to drive: Systematic discovery of semantic dos vulnerability in autonomous driving planning under physical-world attacks. arXiv preprint (2022) [arXiv:2201.04610](https://arxiv.org/abs/2201.04610) [cs.CR]
- Wang, G., Zhou, C., Wang, Y., Chen, B., Guo, H., Yan, Q.: Beyond boundaries: A comprehensive survey of transferable attacks on ai systems. arXiv preprint (2023) [arXiv:2311.11796](https://arxiv.org/abs/2311.11796) [cs.CR]
- Xia, G.-S., Bai, X., Ding, J., Zhu, Z., Belongie, S., Luo, J., Datcu, M., Pelillo, M., Zhang, L.: Dota: A large-scale dataset for object detection in aerial images. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 3974–3983 (2018). <https://doi.org/10.1109/CVPR.2018.00418>
- Xu, C., Ding, W., Lyu, W., Liu, Z., Wang, S., He, Y., Hu, H., Zhao, D., Li, B.: Safebench: A benchmarking platform for safety evaluation of autonomous vehicles. arXiv preprint (2022) [arXiv:2206.09682](https://arxiv.org/abs/2206.09682) [cs.RO]
- Xu, W., Evans, D., Qi, Y.: Feature squeezing: Detecting adversarial examples in deep neural networks. arXiv preprint (2017) [arXiv:1704.01155](https://arxiv.org/abs/1704.01155) [cv.CV]
- Xu, Y., Hu, Y., Zhang, Z., Meyer, G.P., Mustikovela, S.K., Srinivasa, S., Wolff, E.M., Huang, X.: Vlm-ad: End-to-end autonomous driving through vision-language model supervision. arXiv preprint (2024) [arXiv:2412.14446](https://arxiv.org/abs/2412.14446) [cs.CV]
- Xing, W., Li, M., Li, M., Han, M.: Towards robust and secure embodied ai: A survey

- on vulnerabilities and attacks. arXiv preprint (2025) [arXiv:2502.13175](https://arxiv.org/abs/2502.13175) [cs.CR]
- Xiang, C., Mahloujifar, S., Mittal, P.: PatchCleanser: Certifiably robust defense against adversarial patches for any image classifier. In: 31st USENIX Security Symposium (USENIX Security 22), pp. 2065–2082. USENIX Association, Boston, MA (2022). <https://www.usenix.org/conference/usenixsecurity22/presentation/xiang>
- Xie, C., Wu, Y., Maaten, L., Yuille, A.L., He, K.: Feature denoising for improving adversarial robustness. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 501–509 (2019). <https://doi.org/10.1109/CVPR.2019.00059>
- Xie, C., Wang, J., Zhang, Z., Ren, Z., Yuille, A.: Mitigating adversarial effects through randomization. arXiv preprint (2017) [arXiv:1711.01991](https://arxiv.org/abs/1711.01991) [cs.CV]
- Xie, C., Wang, J., Zhang, Z., Zhou, Y., Xie, L., Yuille, A.: Adversarial examples for semantic segmentation and object detection. In: 2017 IEEE International Conference on Computer Vision (ICCV), pp. 1378–1387 (2017). <https://doi.org/10.1109/ICCV.2017.153>
- Xiao, C., Zhu, J.-Y., Li, B., He, W., Liu, M., Song, D.: Spatially transformed adversarial examples. arXiv preprint (2018) [arXiv:1801.02612](https://arxiv.org/abs/1801.02612) [cs.CR]
- Xu, K., Zhang, G., Liu, S., Fan, Q., Sun, M., Chen, H., Chen, P.-Y., Wang, Y., Lin, X.: Adversarial t-shirt! evading person detectors in a physical world. In: Computer Vision – ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V, pp. 665–681. Springer, Berlin, Hjournalalberg (2020). <https://doi.org/10.1007/978-3-030-58558-7>
- Yu, S., Hirche, M., Huang, Y., Chen, H., Allgöwer, F.: Model predictive control for autonomous ground vehicles: a review. *Auton. Intell. Syst.* **1**(4) (2021) <https://doi.org/10.1007/s43684-021-00005-z>
- Yang, J., Kim, H., Wan, W., Hovakimyan, N., Vorobeychik, Y.: Certified Robust Control under Adversarial Perturbations. arXiv preprint [arXiv:2302.02208](https://arxiv.org/abs/2302.02208) (2023)
- Yu, Z., Li, A., Wen, R., Chen, Y., Zhang, N.: Physense: Defending physically realizable attacks for autonomous systems via consistency reasoning. In: Proceedings of the ACM Conference on Computer and Communications Security (CCS) (2024). <https://doi.org/10.1145/3658644.3690236>
- Yeong, D.J., Velasco-Hernandez, G., Barry, J., Walsh, J.: Sensor and sensor fusion technology in autonomous vehicles: A review. *Sensors* **21**(6) (2021) <https://doi.org/10.3390/s21062140>
- Yin, J.-L., Wang, W., Lyhwa, Lin, W., Liu, X.: Adversarial-inspired backdoor defense

- via bridging backdoor and adversarial attacks. Proceedings of the AAAI Conference on Artificial Intelligence **39**(9), 9508–9516 (2025) <https://doi.org/10.1609/aaai.v39i9.33030>
- Yang, H., Zhao, J., Xiong, Z., Lam, K.-Y., Sun, S., Xiao, L.: Privacy-preserving federated learning for uav-enabled networks: Learning-based joint scheduling and resource management. IEEE Journal on Selected Areas in Communications **39**(10), 3144–3159 (2021) <https://doi.org/10.1109/jsac.2021.3088655>
- Zhang, Q., Hu, S., Sun, J., Chen, Q.A., Mao, Z.M.: On adversarial robustness of trajectory prediction for autonomous vehicles. In: CVPR, pp. 15159–15168 (2022). <https://doi.org/10.1109/cvpr52688.2022.01473>
- Zhu, X., Liu, Y., Hu, Z., Li, J., Hu, X.: Infrared adversarial car stickers. arXiv preprint (2024) [arXiv:2405.09924](https://arxiv.org/abs/2405.09924) [cs.CV]
- Zhang, Y., Liu, Z., Jia, C., Zhu, Y., Miao, C.: An online defense against object-based lidar attacks in autonomous driving. In: Proceedings of the 22nd ACM Conference on Embedded Networked Sensor Systems (SenSys) (2024). <https://doi.org/10.1145/3666025.3699345>
- Zhou, H., Li, W., Kong, Z., Guo, J., Zhang, Y., Yu, B., Zhang, L., Liu, C.: Deep-billboard: Systematic physical-world testing of autonomous driving systems. In: Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering. ICSE '20, pp. 347–358. Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3377811.3380422>
- Zhao, Y., Lv, W., Xu, S., Wei, J., Wang, G., Dang, Q., Liu, Y., Chen, J.: Detrs beat yolos on real-time object detection. In: 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 16965–16974 (2024). <https://doi.org/10.1109/cvpr52733.2024.01605>
- Zhao, S., Ma, X., Zheng, X., Bailey, J., Chen, J., Jiang, Y.-G.: Clean-Label Backdoor Attacks (2019). <https://doi.org/10.1109/cvpr42600.2020.01445>
- Zhao, Y., Xiang, Z., Yin, S., Pang, X., Wang, Y., Chen, S.: Malicious Agent Detection for Robust Multi-Agent Collaborative Perception. In: Proc. of IROS (2024). <https://doi.org/10.1109/iros58592.2024.10801337>
- Zhang, T., Xiao, Y., Zhang, X., Li, H., Wang, L.: Benchmarking the physical-world adversarial robustness of vehicle detection. arXiv preprint (2023) [arXiv:2304.05098](https://arxiv.org/abs/2304.05098) [cs.CV]
- Zhang, K., Zhou, H., Bian, H., Zhang, W., Yu, N.: Certified defense against patch attacks via mask-guided randomized smoothing. In: Proc. ICLR (2022). <https://doi.org/10.1007/s11432-021-3457-7>

Zhang, Y., Zhang, Y., Qi, J., Bin, K., Wen, H., Tong, X., Zhong, P.: Adversarial patch attack on multi-scale object detection for uav remote sensing images. *Remote Sensing* 14(21), 5298 (2022) <https://doi.org/10.3390/rs14215298>