Demo: Physically Hijacking Object Trackers

Raymond Muller*, Yanmao Man†, Z. Berkay Celik*, Ming Li†, and Ryan Gerdes‡
*Purdue University †University of Arizona ‡Virginia Tech

Videos: https://youtube.com/playlist?list=PL1wf-CLdUk8KFhgFAHHfaUaku-8IL3z_h

Introduction. This demo is based on AttrackZone [1], a method to physically launch tracker hijacking attacks against Siamese object trackers [3]. Object trackers are used in domains including autonomous driving, pedestrian detection, and mobile robot navigation to improve the accuracy and robustness of object detectors. Previous attacks against object tracking either lacked real-world applicability or did not work against Siamese trackers, which are gaining prevalence due to their high accuracy in real-world applications.

Launching physical attacks against Siamese trackers is challenging: *(i)* attack perturbations must respect physical constraints on where they can be applied, *(ii)* perturbations must be crafted to be as imperceptible as possible to the casual observer while achieving the desired behavior, *(iii)* attacks must influence a group of frames over a given period, necessitating rapid adjustments to the attack to match changes in the surrounding environment, and *(iv)* attacks must evade existing methods (e.g., Kalman filtering) implemented to correct object tracking errors.

Attacking Siamese Trackers. We address these challenges with AttrackZone. As input, it takes the attacker's desired tracker behavior along with camera and 3D point cloud data of the target's environment, taken from a surrogate source, such as a drone flying near the victim in real-time or publicly available mapping data. Figure 1 illustrates the attack zone generation process. AttrackZone uses a worklist-based algorithm to distill the 3D point cloud data into points that are within the camera's field of view. The 3D points are then projected onto the 2D camera image via the camera's intrinsic matrix. The 2D points are grown horizontally and vertically by a scalar defined by the attacker in order to produce contiguous regions throughout the image. These regions encode the projector-perturbable areas and are passed into an adversarial optimization function to yield perturbations that affect the desired behavior.

We evaluated AttrackZone on three different Siamese trackers, DaSiamRPN, SiamRPN, and DaSiamRPN+, against autonomous driving and video surveillance. Using these three models, we conduct (1) emulated attacks against test samples of trackers, (2) simulated attacks against a Drivetruth [2] generated dataset, and (3) real-world attacks using a real camera, projector, and vehicles in a controlled environment. On average, AttrackZone achieves the attacker-desired behavior 92% of the time, requiring between 0.3-3 seconds of continuous perturbation to be successful depending on the domain (e.g., digital or physical).

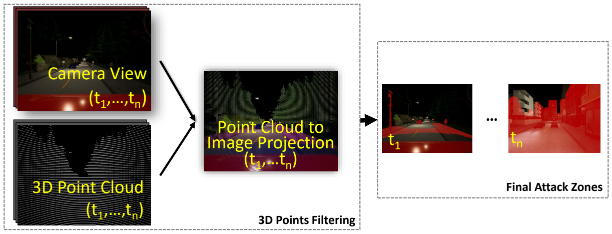

Fig. 1: AttrackZone's zone generation process.

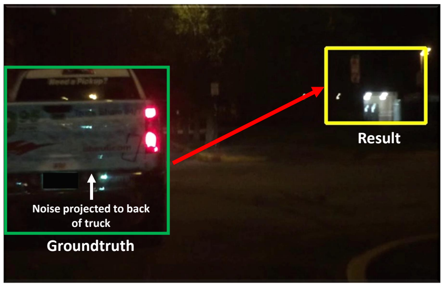

Fig. 2: A tracker hijacking attack moving a truck's tracker off the road, causing an autonomous vehicle to believe the way is clear. As the vehicle accelerates, it will collide with the truck.

Figure 2 demonstrates a successful attack against autonomous driving. An attacker aims to cause an autonomous vehicle to collide with another vehicle. The attacker projects adversarial perturbations onto the back of a truck, causing the truck's object tracker to move off the road. Thinking that the road is now clear, the autonomous vehicle will accelerate, rear-ending the truck.

Demonstration Plan. We provide a video playlist demonstrating two attacks against separate datasets, as well as one real-world attack each for autonomous driving and video surveillance. Attacks against autonomous driving work to either stop an autonomous vehicle or cause it to collide, while the attack against video surveillance work to disguise an attacker's entry into a restricted area. Our project website breaks down the code used to conduct the attack and evaluate AttrackZone: **https://github.com/purseclab/AttrackZone**.

Acknowledgment

This work has been partially supported by the National Science Foundation (NSF) under grants CNS-2144645 and CNS-1801611, the Army Research Office (ARO) under grant W911NF2110320, and through a gift by Qualcomm. The views expressed are those of the authors only.

References

[1] R. Muller, Y. Man, Z. B. Celik, R. Gerdes, and M. Li, "Physical Hijacking Attacks against Object Trackers," in *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, 2022.

[2] R. Muller, Y. Man, Z. B. Celik, M. Li, and R. Gerdes, "DriveTruth: Automated autonomous driving dataset generation for security applications," in *International Workshop on Automotive and Autonomous Vehicle Security (AutoSec), collocated with NDSS*, 2022.

[3] M. Ondrasovic and P. Tarabek, "Siamese Visual Object Tracking: A Survey," *IEEE Access*, 2021.

Symposium on Vehicles Security and Privacy (VehicleSec) 2023
27 February 2023, San Diego, CA, USA
ISBN 1-891562-88-6
https://dx.doi.org/10.14722/vehiclesec.2023.23049
www.ndss-symposium.org