Short Description of the Proposed Demo

License plate detection (**LPD**) underpins many intelligent transportation and public-safety applications, yet its resilience to **physically realizable adversarial perturbations** remains underexplored. This study introduces an **attack methodology** for the **physical domain**, enforcing that perturbations remain effective under real-world materialization, including printing, plate mounting, and camera recapture.

Methodologically, we optimize perturbations in the CIE–*Lab* color space under two realism constraints. First, modifications are confined to the **Luminance (*L*) channel** to preserve chromatic fidelity; second, the license plate's numeric content is masked to maintain human readability. By **digital-to-physical transfer**, we mean that adversarial patterns computed in the digital domain remain effective once materialized in the real world (printed, mounted, and re-acquired). To promote this transfer, the optimization uses **Expectation Over Transformation (EOT)** attack and integrates a differentiable **Print-and-Scan (P&S) simulator** (CycleGAN based) directly in the loop, modeling key reproduction artifacts during printing: halftoning, dot gain/ink spread; and during camera capture: optical blur, illumination non-uniformity. Evaluation targets an **SSD-300**–based LPD model on a dataset of real Italian license plates, complemented by **on-vehicle physical trials**.

Empirically, both RGB and *L*-channel variants achieve 100% digital Attack Success Rate (ASR), confirming the vulnerability of standard LPD pipelines in the digital/computational setting (i.e., evaluation performed entirely within the model's inference pipeline, without physical rendering, printing, or recapture). In physical tests, the combined **EOT + P&S** strategy attains **63.15% ASR (RGB)** and **65% ASR (*L*-channel)**. The **luminance-only constraint** also improves perceptual quality (**PSNR ≈ 22 dB vs ≈ 15.6 dB** for RGB) and reduces visible chromatic artifacts, yielding a favorable **stealth and effectiveness** trade-off under realistic conditions.

Overall, this study demonstrates a practical, physically realizable attack methodology against LPD, quantifies its impact through digital evaluations and **vehicle-mounted field tests**, and motivates **luminance-constrained perturbations** with **in-loop P&S modeling** as a strong baseline for future robustness assessments in safety-critical AI.

Demo Setup and What Will Be Shown

In the demo, we will showcase physical adversarial license plates attacking an SSD-300 based LPD system. We will bring:

- A laptop running the LPD software and visualization interface (with the possibility of connecting to an external display, if available),
- A set of printed physical adversarial license plates,
- A physical car bumper mock-up on which the plates are mounted,
- A tripod and a camera device.

The car bumper with the mounted adversarial plate will be placed at a fixed location. Using the tripod-mounted camera, we will capture images of the plate from three different viewpoints, varying both distance and angle to emulate realistic surveillance. The acquired images will be processed live by the LPD system on the laptop, and the output (detections or missed detections) will be shown to attendees. The demo will illustrate how the proposed physically realizable perturbations cause the LPD model to fail, despite the plate remaining human-readable and visually natural, thereby highlighting concrete robustness challenges for safety-critical AI systems deployed in real-world scenarios.