Rolling Colors: Adversarial Laser Exploits against Traffic Light Recognition

Chen Yan[1], Zhijian Xu[1,2], Zhanyuan Yin[3], Xiaoyu Ji[1], Wenyuan Xu[1]

[1]Zhejiang University, [2]The Chinese University of Hong Kong,

[3]The University of Chicago

Traffic Light Recognition

- Enables vehicles to detect and recognize traffic light signals
- Essential for full autonomous driving in urban areas

How does traffic light recognition work?

What if traffic light recognition goes wrong?

Green Light

Run a red light

Intersection Accident!

Red Light

Sudden stop

Rear-end Crash!

Spoof traffic light recognition?

- Use fake traffic lights? … Probably not the best idea.

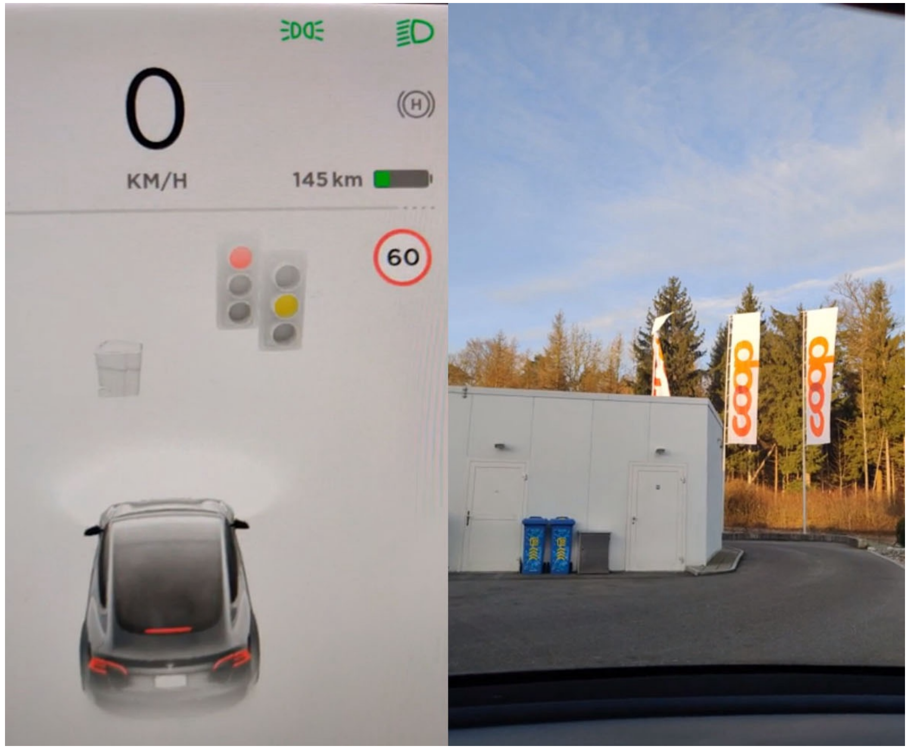

Let’s use laser!

- Narrow beam of radiation → *Travel a long distance, hard to detect*
- Previous studies have shown laser’s capability on interfering cameras

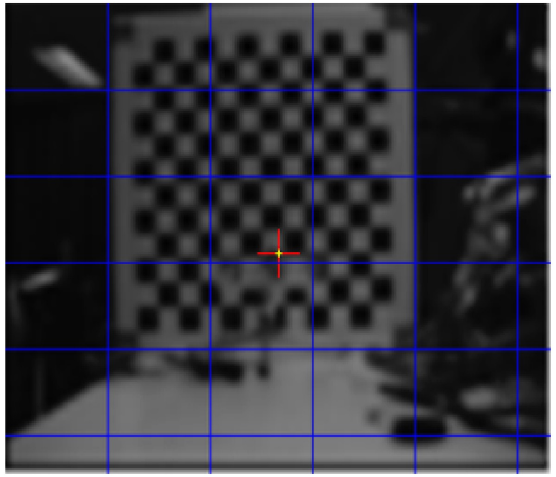

Laser off

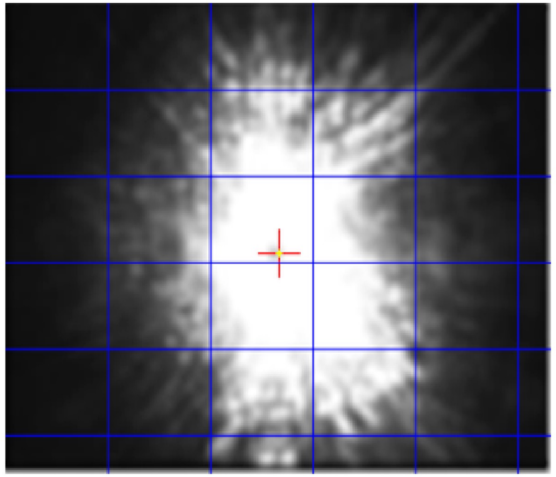

Laser on

Laser off

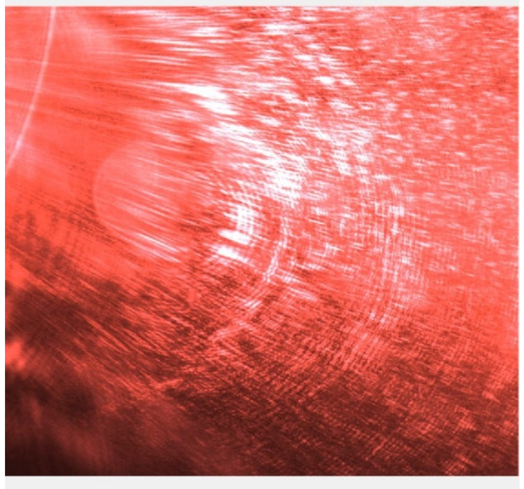

Laser on

Petit et al., Remote Attacks on Automated Vehicles Sensors: Experiments on Camera and LiDAR (Blackhat 2015)
Yan et al., Can You Trust Autonomous Vehicles: Contactless Attacks against Sensors of Self-Driving Vehicles (DEFCON 2016)

Attack Scenario and Requirements

Requirement ①

The laser interference affects the image as little as possible

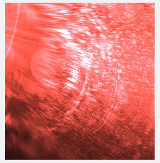

Requirement ②

The traffic light can be detected under laser interference

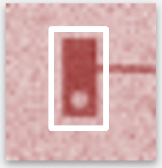

Requirement ③

The traffic light is recognized as being in a targeted color

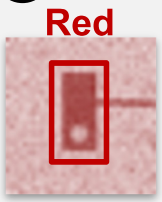

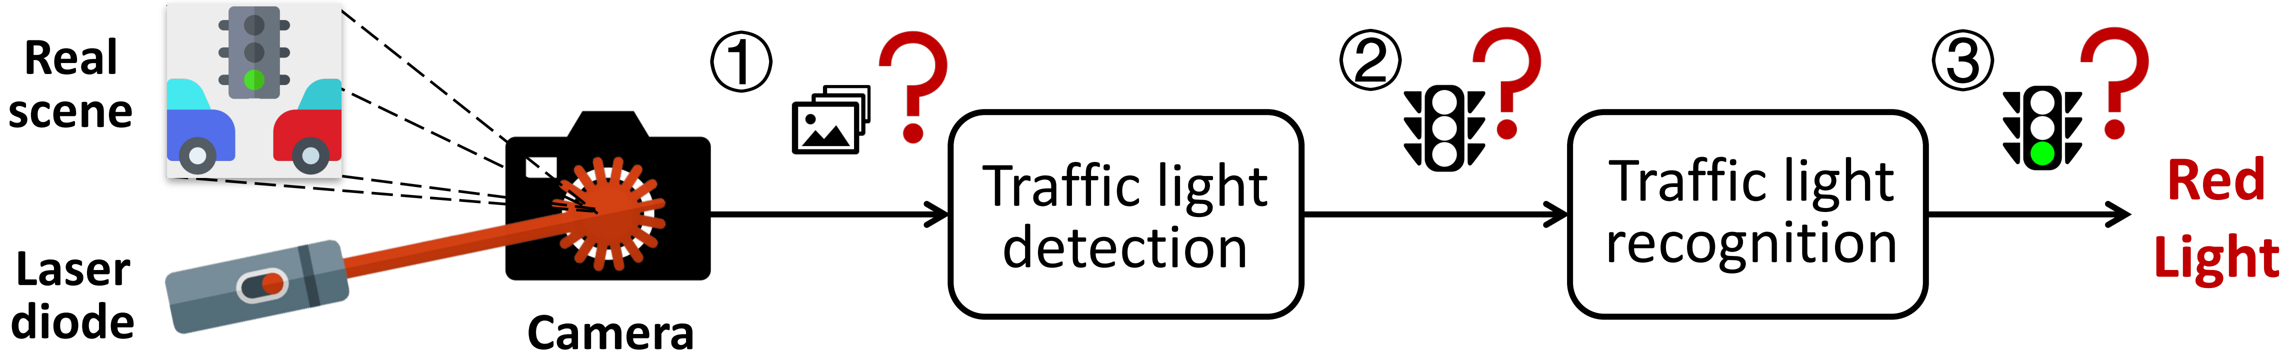

R1: Laser Interference Study

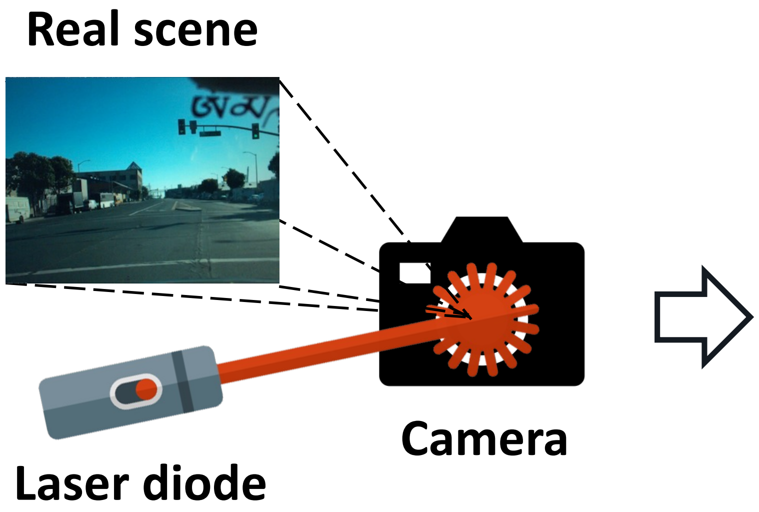

Overexposure

Color shift

White unbalance

Affect the entire image → Can be easily detected

Color stripe

Only affect the traffic light region → Stealthier attack

R1: Exploiting the Rolling Shutter

- A rolling shutter is a type of image capture in cameras that **records the frame line by line on an image sensor** instead of capturing the entire frame all at once.

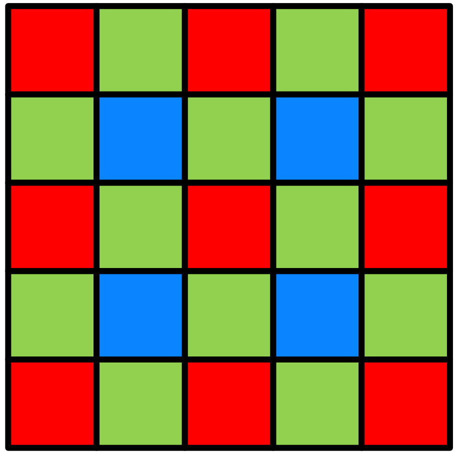

Image sensor

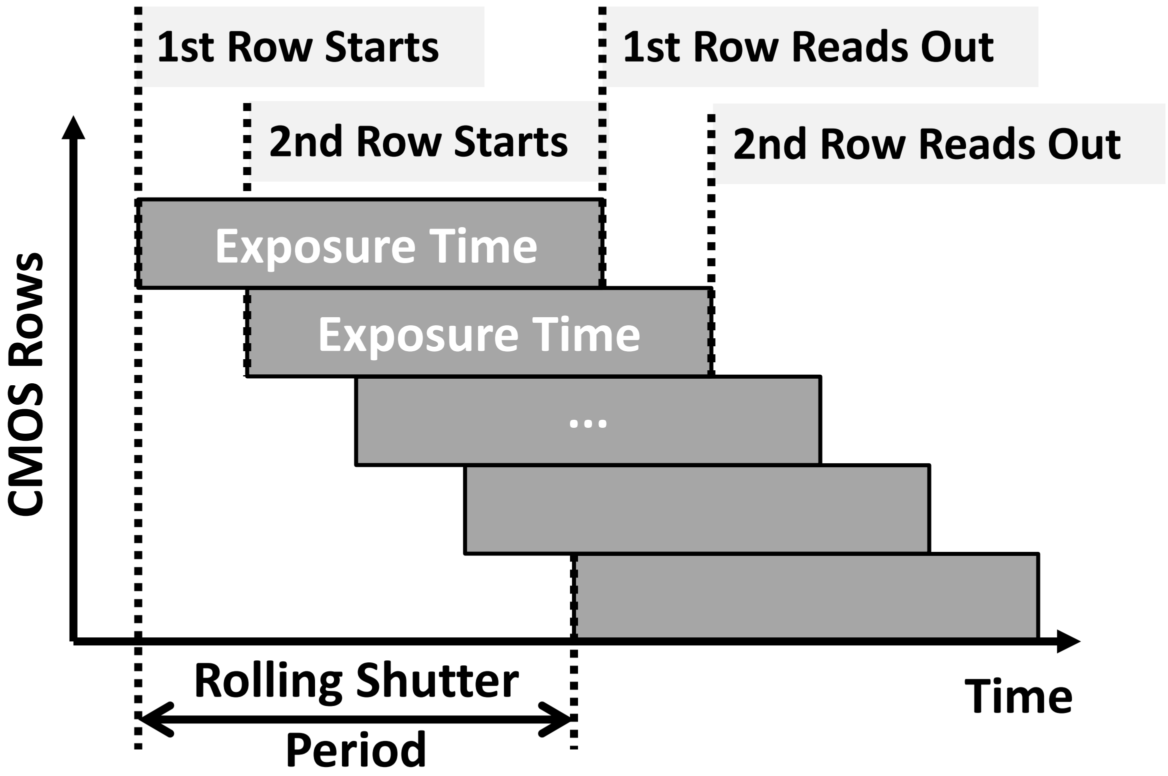

R1: Exploiting the Rolling Shutter

- A rolling shutter is a type of image capture in cameras that **records the frame line by line on an image sensor** instead of capturing the entire frame all at once.

Video source: SmarterEveryDay https://www.youtube.com/watch?v=dNVtMmLlnoE&t=348s

R1: Exploiting the Rolling Shutter

- Inject a color stripe into the captured image with **a laser pulse signal**
- **Full control over the stripe's number, width, and position**
- Synchronize the laser pulse with the rolling shutter period (frame rate)

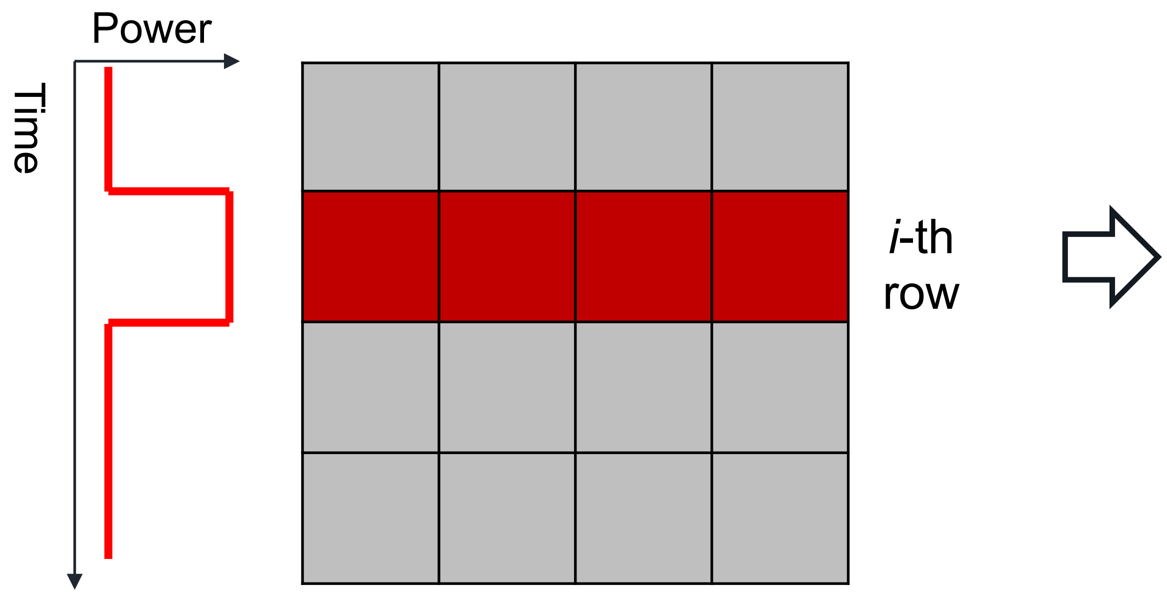

Laser pulse **CMOS grid**

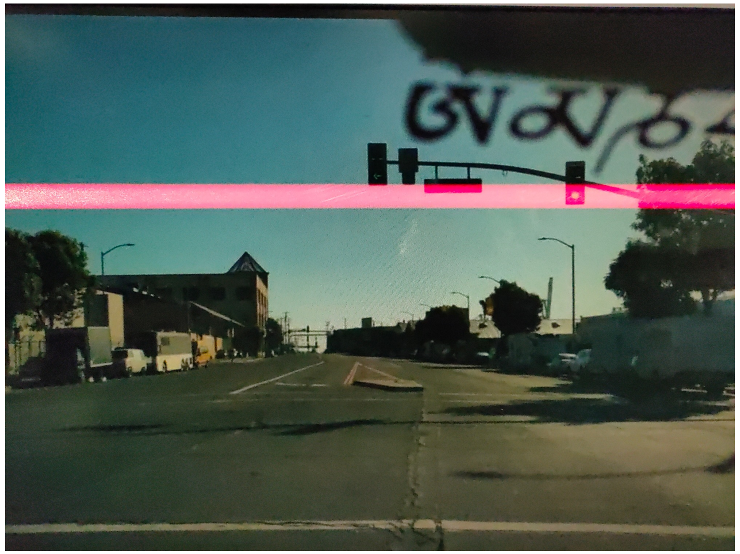

Color stripe injected

R2: Passing Traffic Light Detection

The injected color stripe must **NOT** affect traffic light detection

- **Requisite 1:** proper control of laser intensity
- **Requisite 2:** proper design of the stripe's width and position

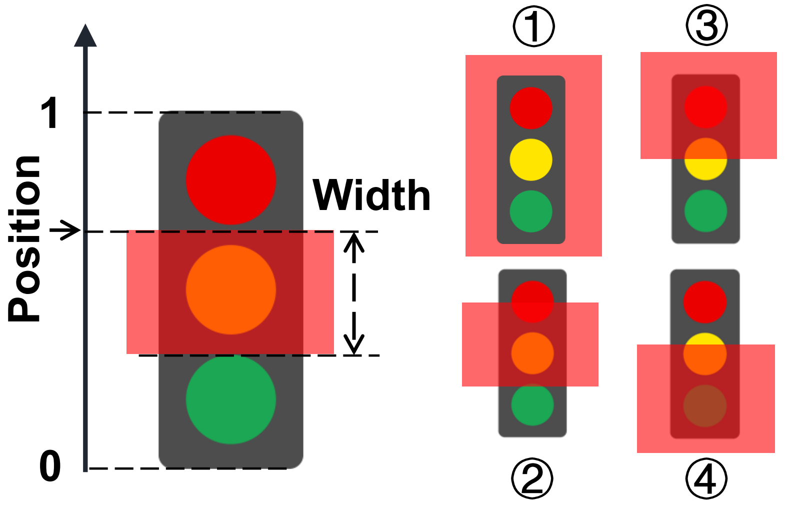

4 color stripe strategies

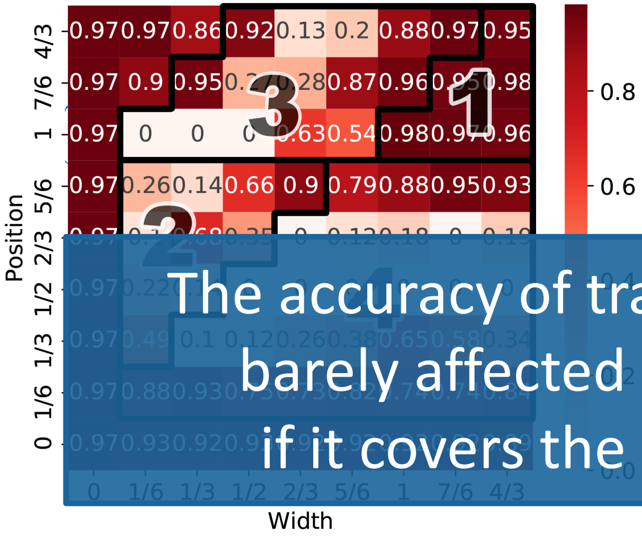

Red laser (650nm)

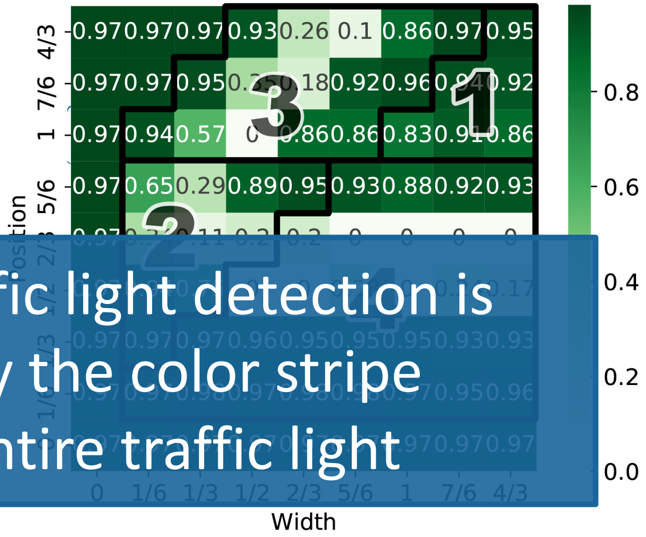

Green laser (520nm)

The accuracy of traffic light detection is barely affected by the color stripe if it covers the entire traffic light

R3: Spoofing Traffic Light Recognition

Spoof traffic light detection to a targeted color (Red → Green, Green → Red)

- **Requisite:** fine-tune the laser parameters according to specific traffic lights and attack scenarios

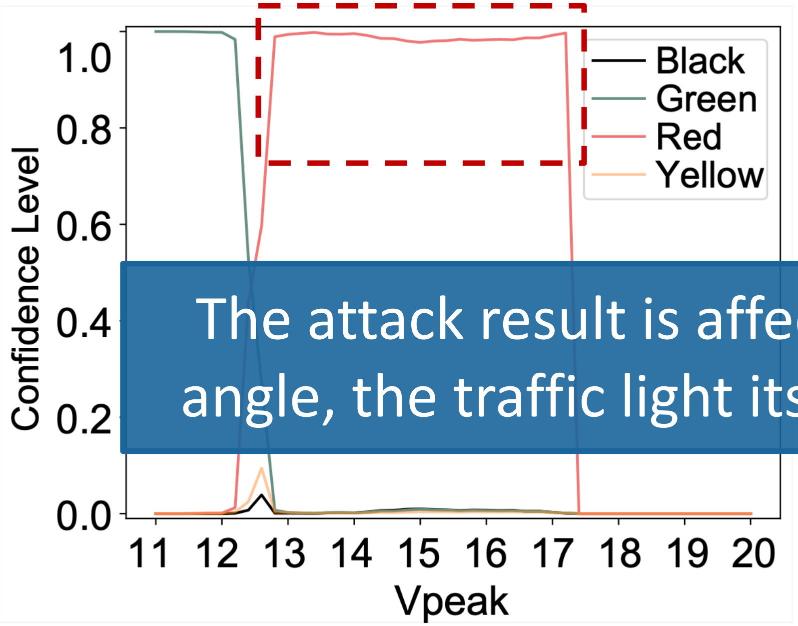

Green → Red

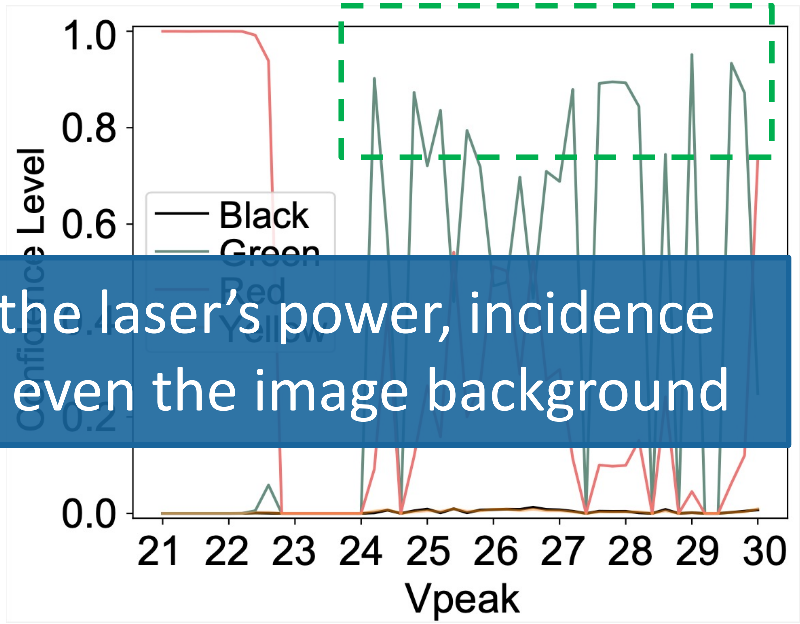

Red → Green

The attack result is affected by the laser's power, incidence angle, the traffic light itself, and even the image background

Threat Model & Attack Workflow

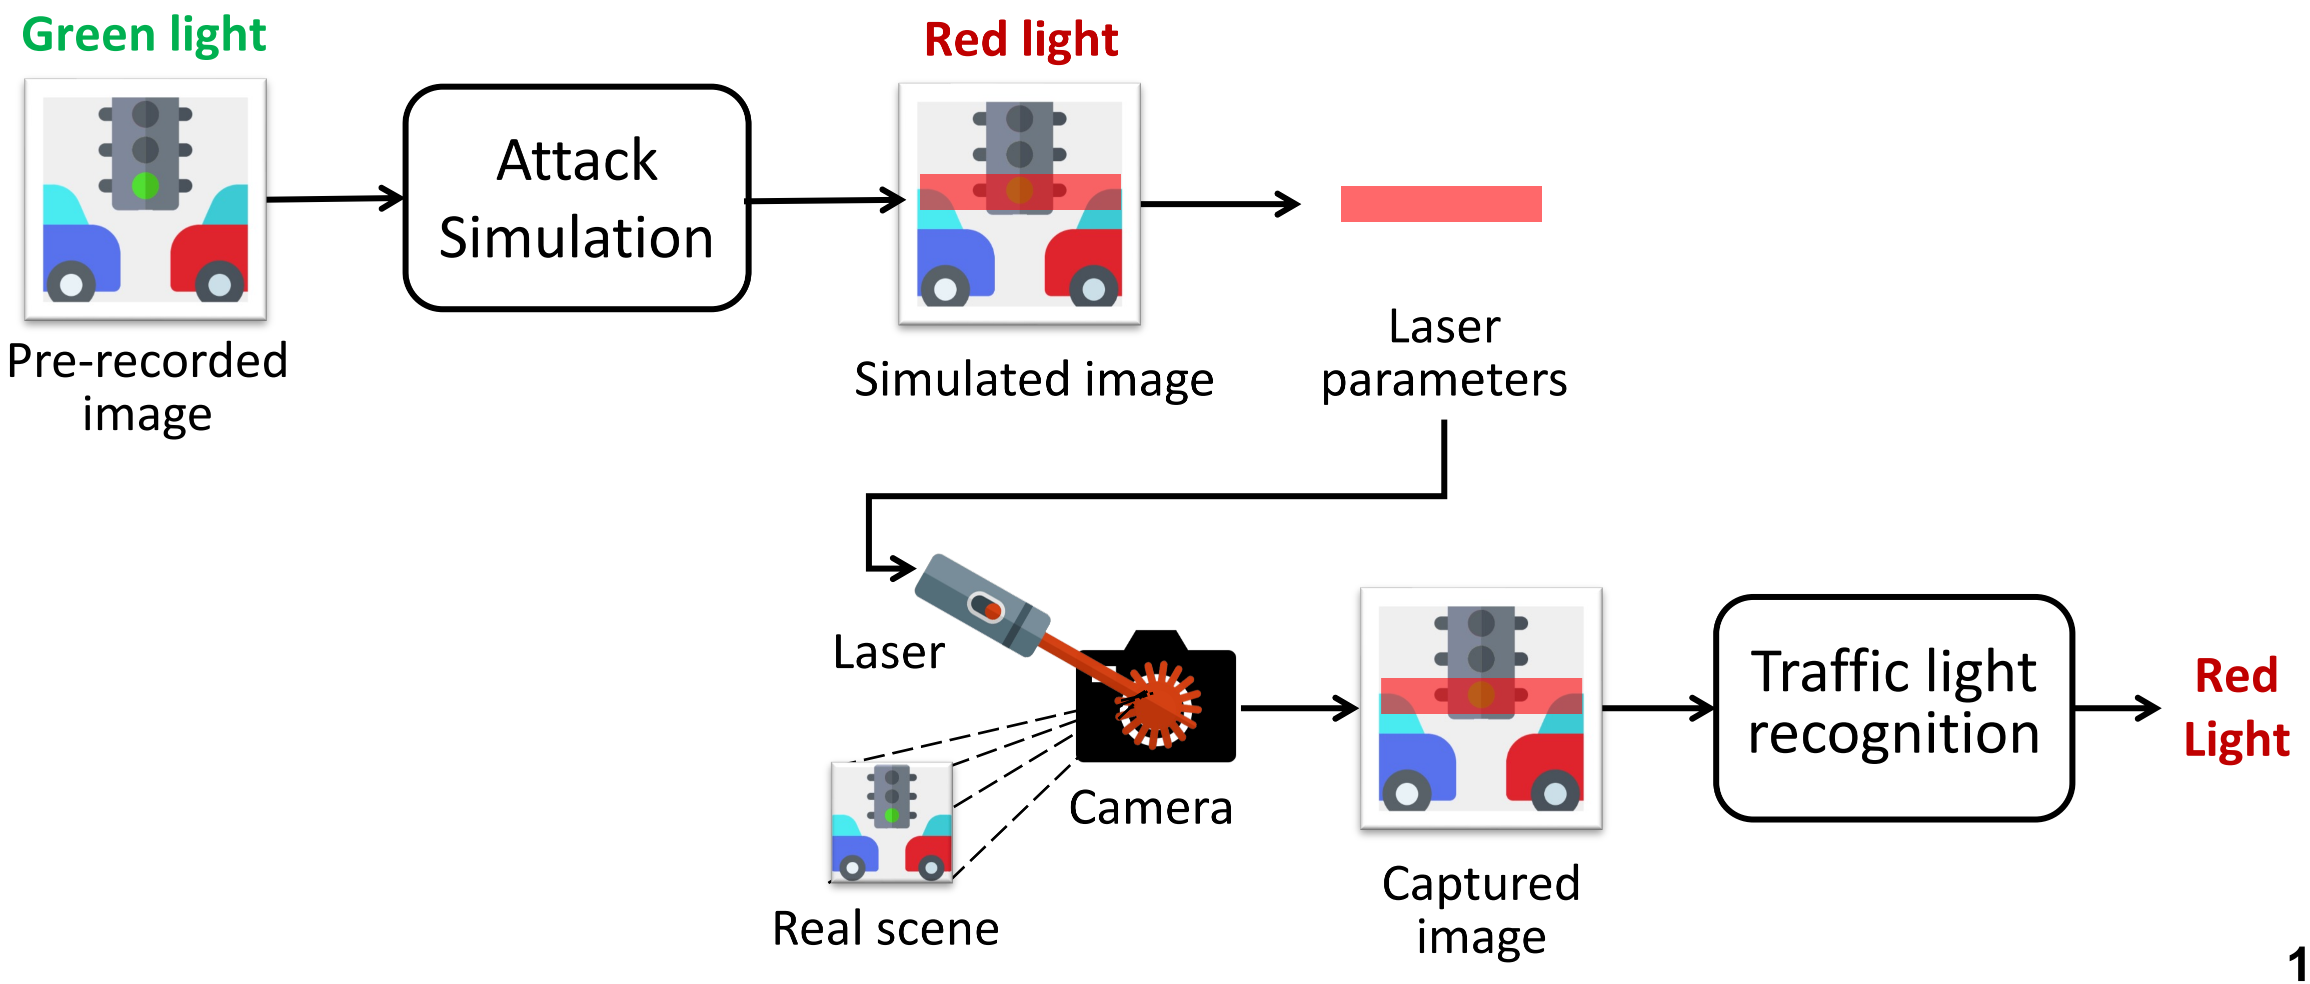

Attack Design

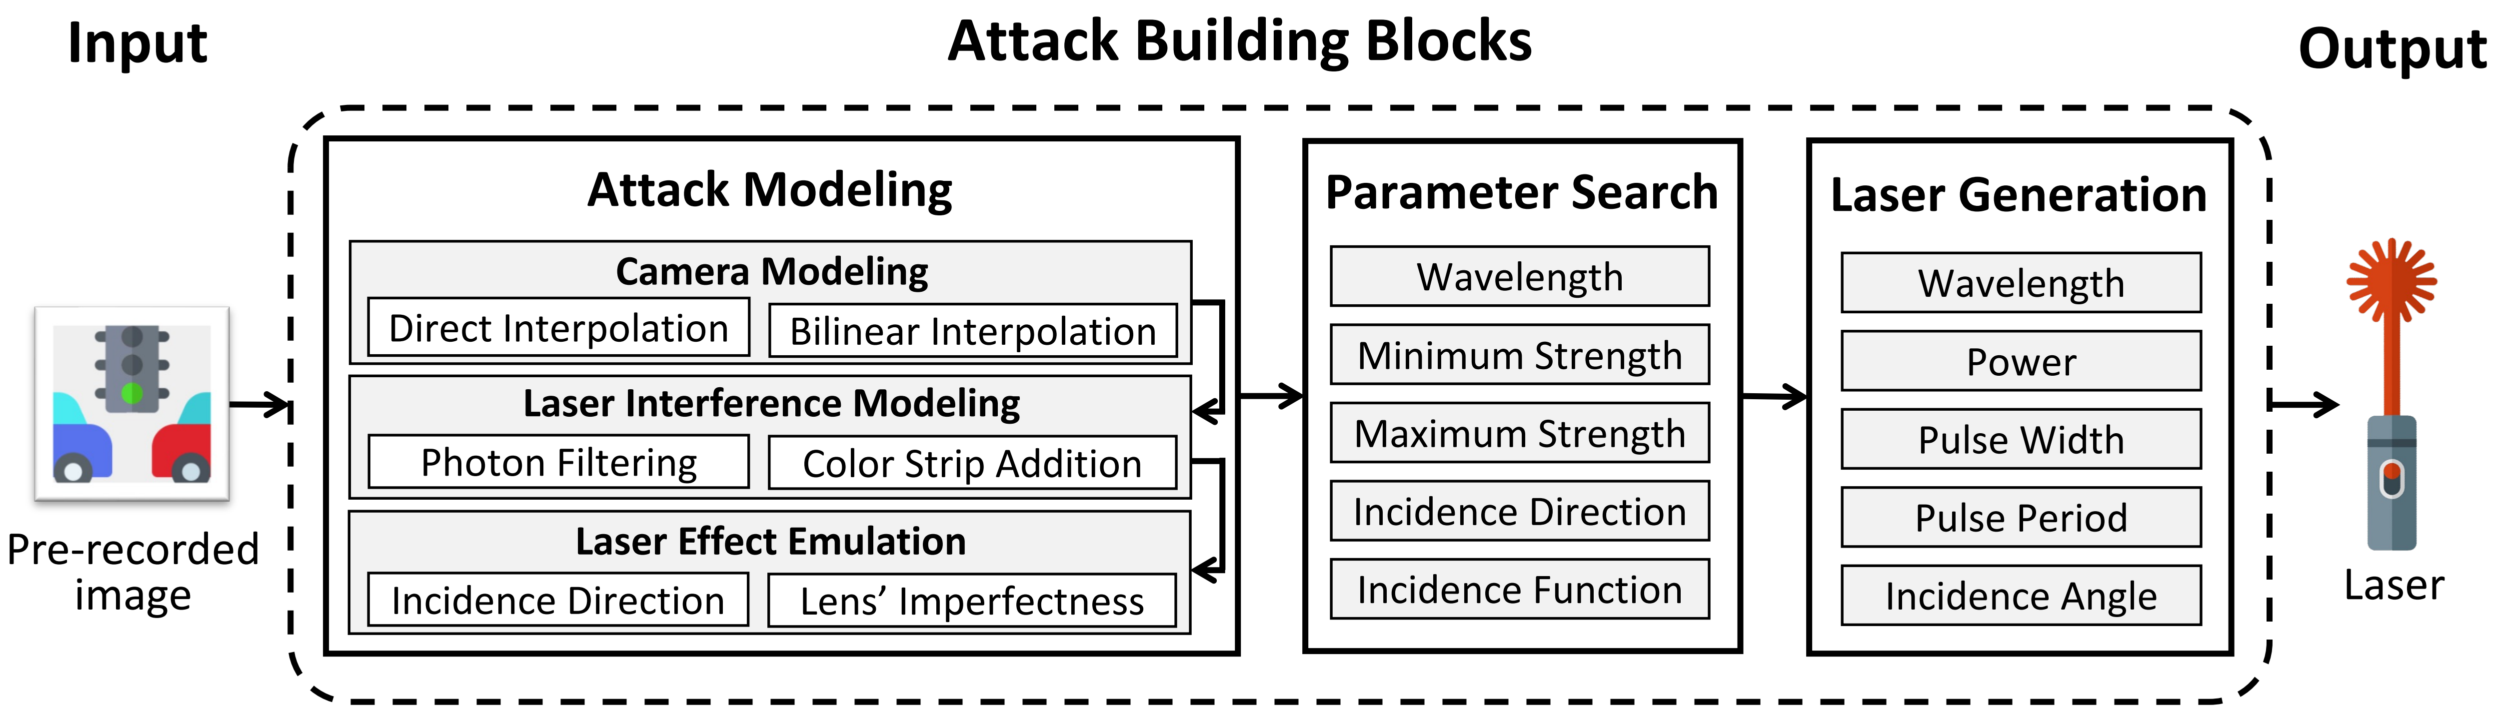

Simulate laser interference

Find effective parameters

Map parameters to laser signals

Evaluation

- Emulated Attacks
- Real-World Attacks in Stationary Setups
- Real-World Attacks in Motion

5 Cameras

Used on Tesla vehicles

2 Models

apollo nexar

Real-World Attacks in Stationary Setups

- **Experiment Setup**
- **Overall Performance**
  - Red → Green
  - Green → Red
- **Impact of the Traffic Light**
  - Distances: 5m - 25m
  - 5 Positions: [L, ML, M, MR, R]
  - 3 Directions: [L, M, R]

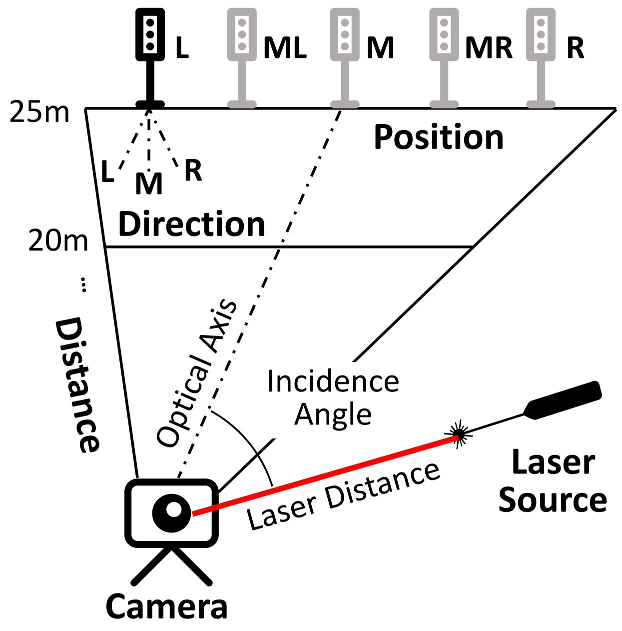

(a) Illustration of setup

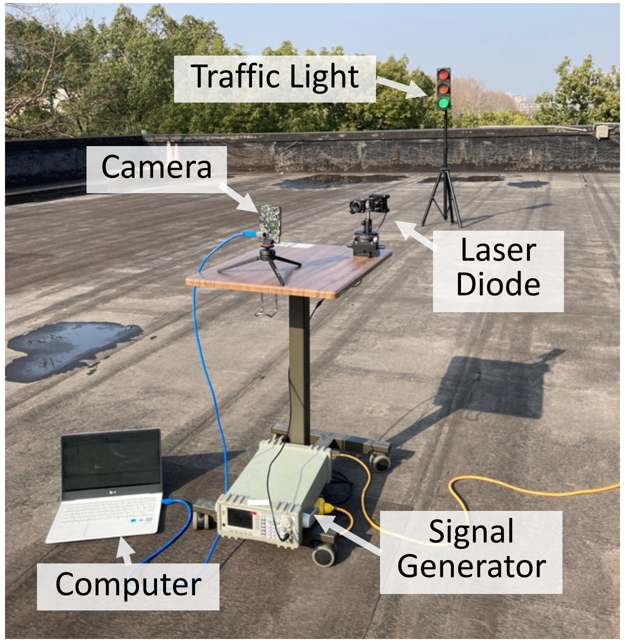

(b) Real setup

Real-World Attacks in Stationary Setups

- **Overall Performance**

Table 2: Success rates of attacking 2 systems and 5 cameras.

Sys.	Attack Scenario	Target Camera					Avg.
		Tesla	Xiaomi	Hikv	OPPO	OpMV	
Apollo	R→G	[illegible]	[illegible]	[illegible]	[illegible]	7.39%	20.27%
	R→DoS	[illegible]	[illegible]	[illegible]	[illegible]	3.04%	32.74%
	G→R	[illegible]	[illegible]	[illegible]	[illegible]	7.37%	41.86%
	G→DoS	[illegible]	[illegible]	[illegible]	[illegible]	2.63%	36.46%
Nexar	R→G	[illegible]	[illegible]	[illegible]	[illegible]	0%	20.78%
	R→DoS	[illegible]	[illegible]	[illegible]	[illegible]	100%	48.04%
	G→R	[illegible]	[illegible]	[illegible]	[illegible]	100%	65.94%
	G→DoS	[illegible]	[illegible]	[illegible]	[illegible]	0%	32.72%

Red→Green: 20.53%

Green→Red: 44.95%

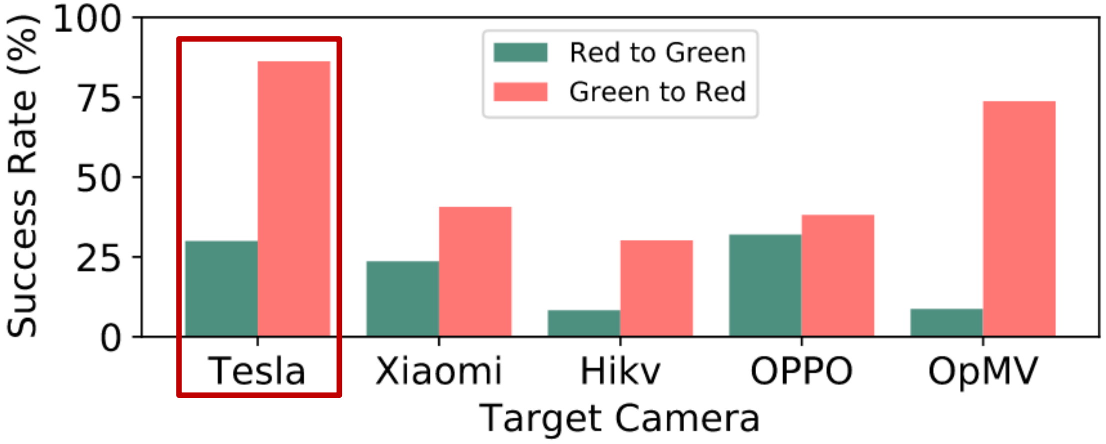

Green→Red is easier than Red→Green

Tesla camera is the most vulnerable

Real-World Attacks in Stationary Setups

- **Impacts of the traffic light's distance, position and direction.**

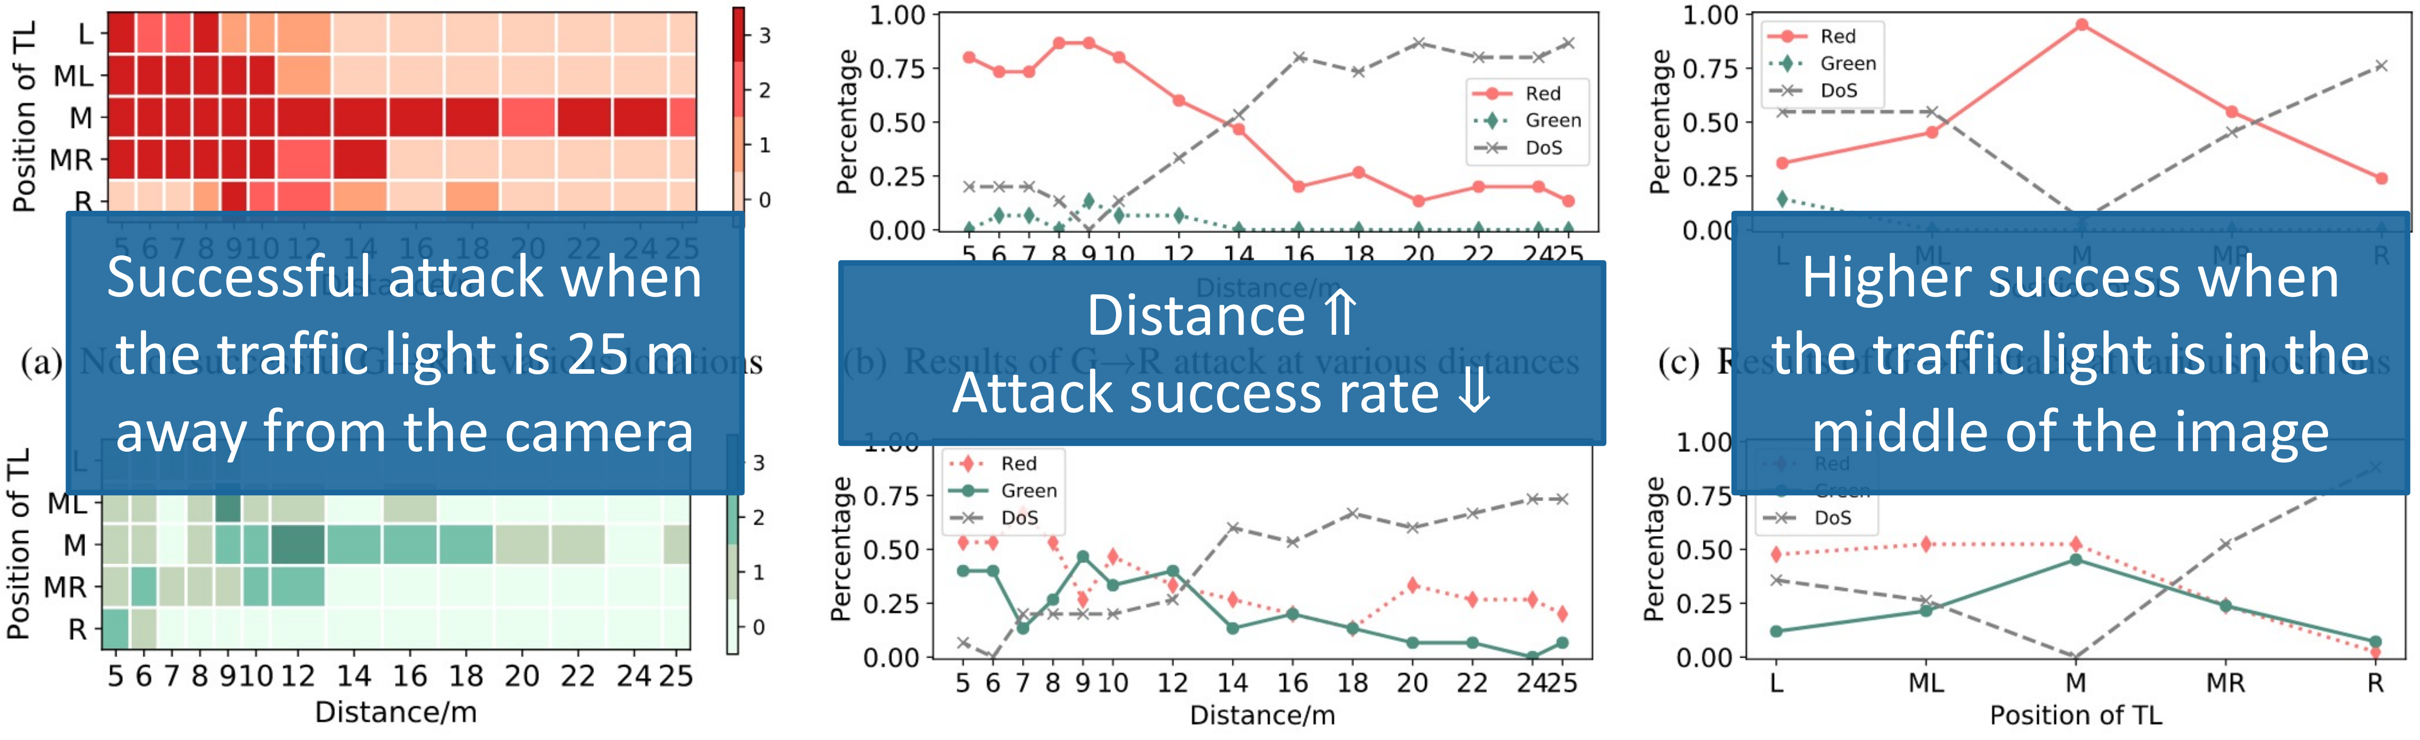

(d) No. of successful R→G at various locations (e) Results of R→G attack at various distances (f) Results of R→G attack at various positions

Successful attack when the traffic light is 25 m away from the camera

Distance ⇑
Attack success rate ⇓

Higher success when the traffic light is in the middle of the image

Real-World Attacks in Motion

- Effectiveness across Continuous Video Frames
- Feasibility of Tracking and Laser Aiming
- End-to-End Impact on Driving

Effectiveness across Continuous Video Frames

Experiment setup

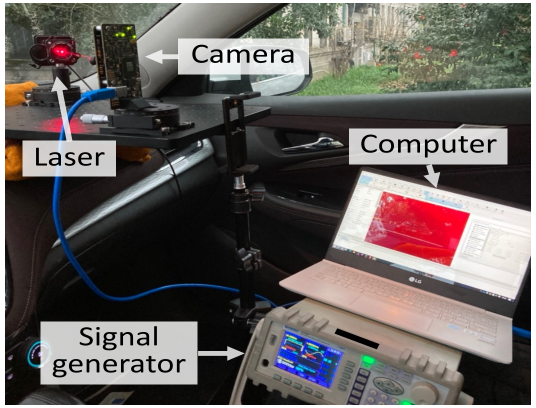

Attack videos

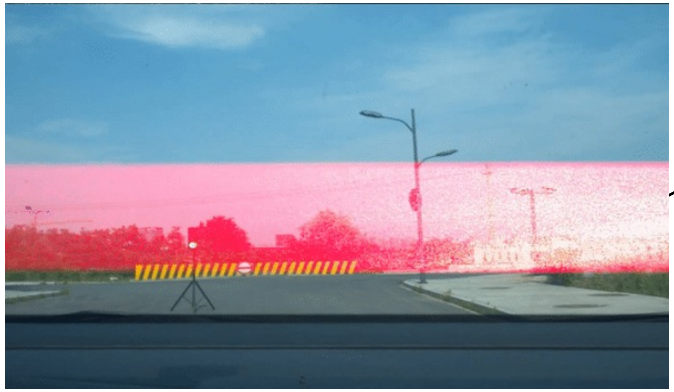

Attack results across continuous frames

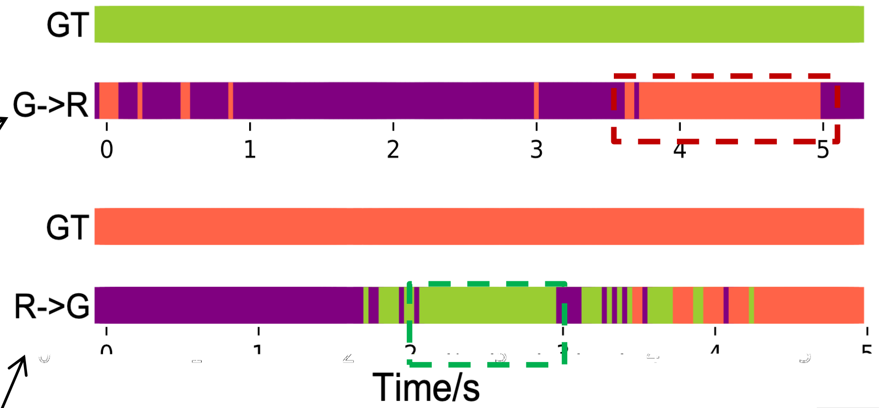

The attack can continuously spoof traffic light recognition for more than 1 second with a success rate of 85.2%

Feasibility of Tracking and Laser Aiming

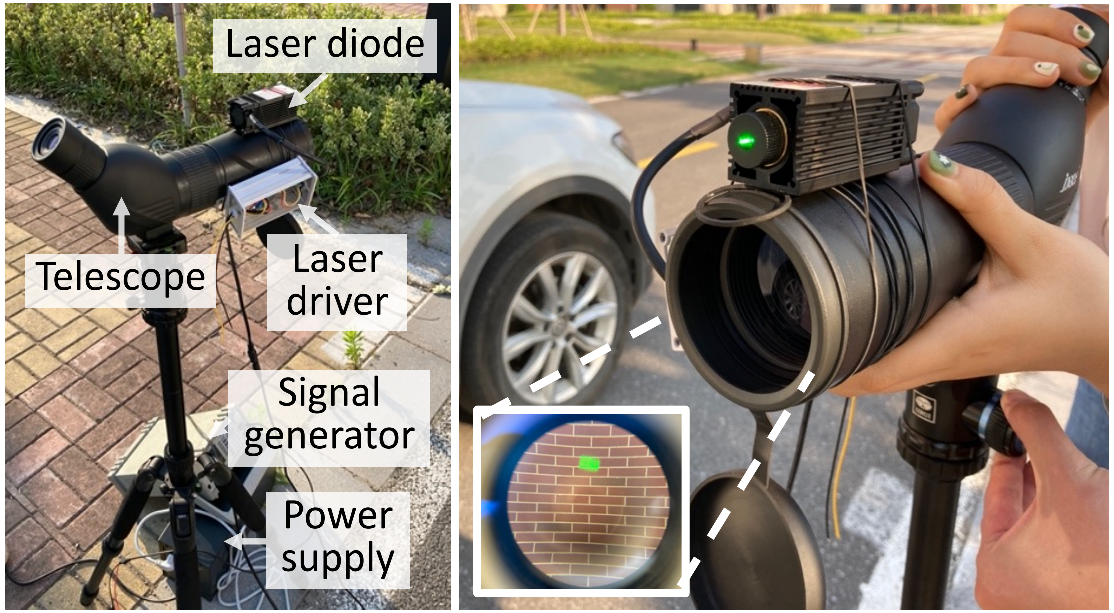

Manual tracking and aiming equipment

Setup for long-range laser aiming experiment
(the attacker was on the roadside and 40-80 m away from the vehicle)

Feasibility of Tracking and Laser Aiming

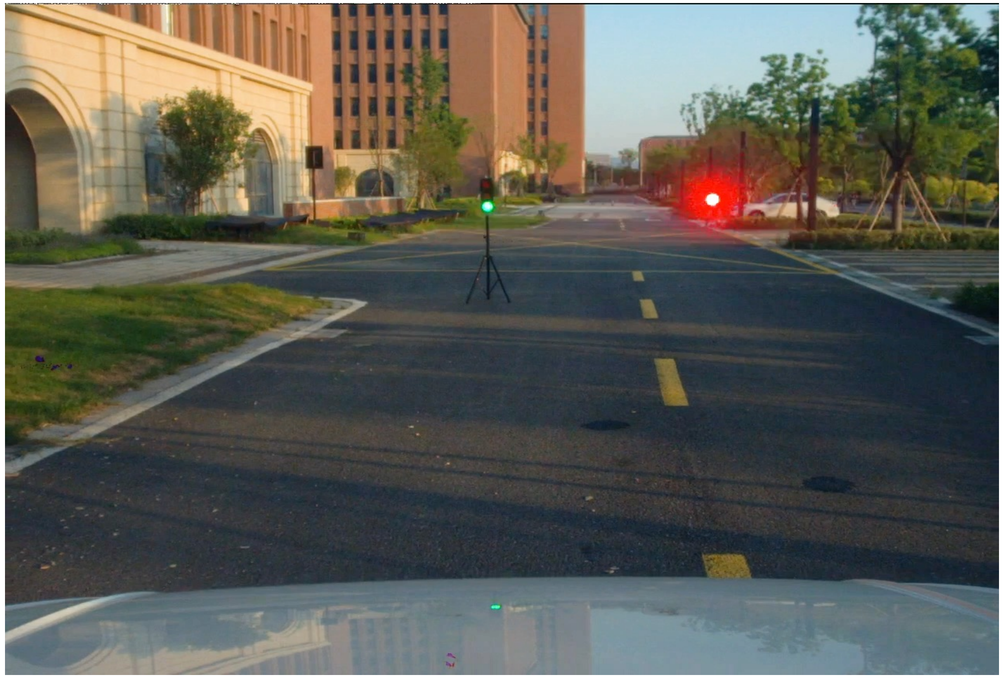

1. Attacker can track the target camera and aim the laser at the same time even when the vehicle is moving at 20 km/h.
2. The average attack success rate of spoofing traffic light recognition is 28.4%.

End-to-End Impact on Driving

Attack Scenario 1:
Running a red light

Attack Scenario 2:
Emergency stop

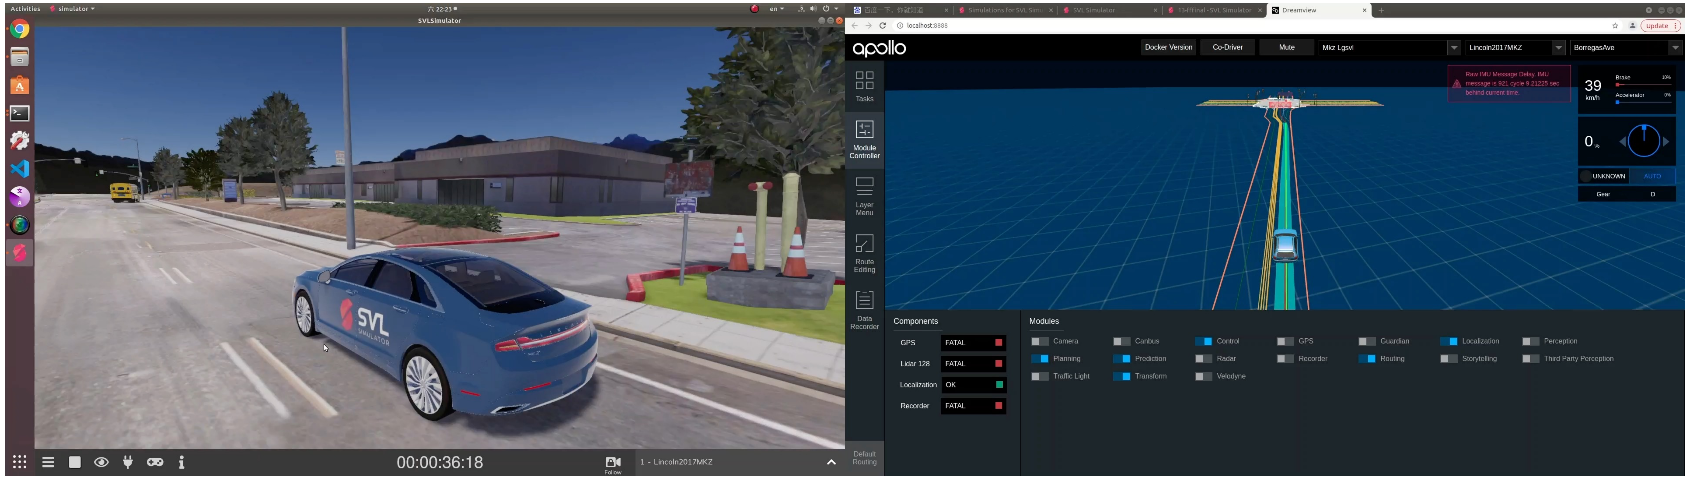

Countermeasures

- Use global shutters instead of rolling shutters
- **Rolling shutter improvement**: expose the CMOS rows in a *random sequence*

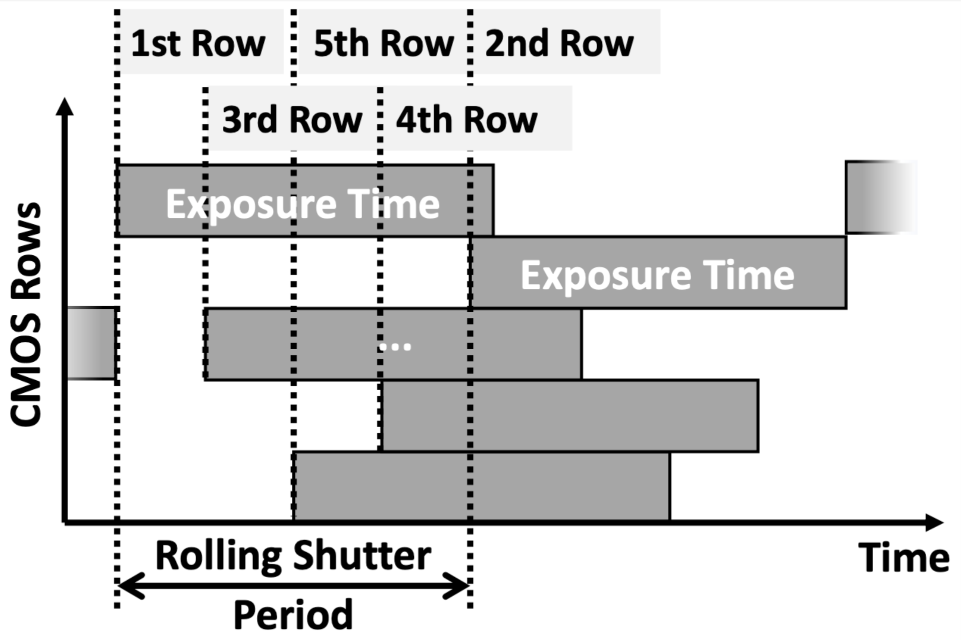

Sequential rolling (before defense)

Random rolling (after defense)

Summary

- A new approach to injecting adversarial images by exploiting an **inherent vulnerability of the rolling shutters** in CMOS cameras
- Experimentally validated the feasibility of **fooling traffic light recognition using laser**
- Evaluated the attack in real-world setups on 2 traffic light recognition systems, 5 cameras, and a moving vehicle

Questions?

Attack demos: https://sites.google.com/view/rollingcolors

USSLAB homepage: http://usslab.org