



They See Me Rollin': Inherent Vulnerability of the Rolling Shutter in CMOS Image Sensors

Sebastian Köhler
University of Oxford
sebastian.kohler@cs.ox.ac.uk

Giulio Lovisotto
University of Oxford
giulio.lovisotto@cs.ox.ac.uk

Simon Birnbach
University of Oxford
simon.birnbach@cs.ox.ac.uk

Richard Baker
University of Oxford
richard.baker@cs.ox.ac.uk

Ivan Martinovic
University of Oxford
ivan.martinovic@cs.ox.ac.uk

ABSTRACT

In this paper, we describe how the electronic rolling shutter in CMOS image sensors can be exploited using a bright, modulated light source (e.g., an inexpensive, off-the-shelf laser), to inject fine-grained image disruptions. We demonstrate the attack on seven different CMOS cameras, ranging from cheap IoT to semi-professional surveillance cameras, to highlight the wide applicability of the rolling shutter attack. We model the fundamental factors affecting a rolling shutter attack in an uncontrolled setting. We then perform an exhaustive evaluation of the attack's effect on the task of object detection, investigating the effect of attack parameters. We validate our model against empirical data collected on two separate cameras, showing that by simply using information from the camera's datasheet the adversary can accurately predict the injected distortion size and optimize their attack accordingly. We find that an adversary can hide up to 75% of objects perceived by state-of-the-art detectors by selecting appropriate attack parameters. We also investigate the stealthiness of the attack in comparison to a naïve camera blinding attack, showing that common image distortion metrics can not detect the attack presence. Therefore, we present a new, accurate and lightweight enhancement to the backbone network of an object detector to recognize rolling shutter attacks. Overall, our results indicate that rolling shutter attacks can substantially reduce the performance and reliability of vision-based intelligent systems.

ACM Reference Format:

Sebastian Köhler, Giulio Lovisotto, Simon Birnbach, Richard Baker, and Ivan Martinovic. 2021. They See Me Rollin': Inherent Vulnerability of the Rolling Shutter in CMOS Image Sensors. In *Annual Computer Security Applications Conference (ACSAC '21)*, December 6–10, 2021, Virtual Event, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3485832.3488016>

1 INTRODUCTION

Digital cameras are undoubtedly ubiquitous today, with units now integrated into most mobile and IoT devices and already dominant in photography, monitoring, and surveillance. With the ongoing integration of cameras into the decision-making process of

intelligent systems, such as consumer home-security [18, 36], commercial security [3], blanket surveillance [31, 40], construction site safety [13], retail [43], and vehicle automation [33, 54], the integrity of the captured images is becoming a cornerstone of the correct behavior of the system.

Recent academic work has produced many examples of disrupting camera-based intelligent systems using projections on surfaces in the environment [25, 29, 58], or the placement of printed patches [14]. Attacks inhibiting the use of camera footage are well known and repeatedly shown in different settings [33, 54].

Naïve disruption of cameras using powerful light projected at the camera is straightforward and often effective, shown in academic work [33] and real-world settings [8]. However, this disruption has previously been coarse-grained, only blinding the camera entirely. While blinding attacks may cause significant disruption, the presence of blinding is easily detected as it generates substantial and sudden image changes. Additionally, blinding attacks trigger significant exposure adjustments, where the camera has to reduce the time the sensor is exposed to compensate for the increasing amount of light, and are therefore generally considered suspicious [38, 46].

Yet, we show that extremely widespread image sensors can be exploited to obtain small-scale control over the captured images. This is because Complementary Metal-Oxide Semiconductor (CMOS) image sensors typically implement an electronic rolling shutter mechanism that captures the image in sequential rows instead of all-at-once. As CMOS image sensors provide a convenient balance between production costs and image quality, they have become the major sensor technology used in modern cameras.

In this paper, we show how an adversary can exploit the electronic rolling shutter with a modulated light source and influence the images captured by the camera in a fine-grained way. The captured pixels' intensity can be varied arbitrarily in small bands of the image, down to a handful of pixel rows given sufficiently bright conditions. The attack is achievable with minimal preparation and no greater aiming requirements than a naïve blinding attack, yet it can be further optimized based on knowledge of the target camera to enhance its effectiveness. The attack is particularly effective when targeted at computer vision systems. As small-scale, high-frequency information plays a crucial role in these systems [50], e.g., for edge detection, the injection of fine-grained distortion has a more profound effect on performance than naïve blinding attacks. This is analogous to the injection of fine-grained, adversarial distortions into image data [17], or the physical world [42].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ACSAC '21, December 6–10, 2021, Virtual Event, USA

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8579-4/21/12...\$15.00

<https://doi.org/10.1145/3485832.3488016>

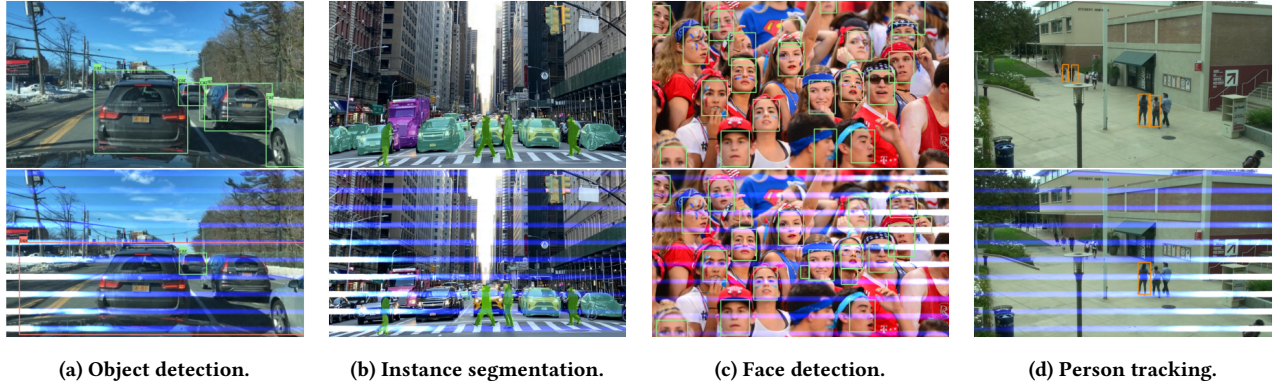


Figure 1: Computer vision use-cases and impact of rolling shutter attacks on predictions. The distortions introduced by the attack compromise the task’s accuracy. From left to right, models are FasterRCNN [35], MaskRCNN [19], RFCN [9] and SSD [24]. These distortions are more stealthy than simple blinding attacks, as demonstrated in Section 7.2.

In this work, we introduce rolling shutter attacks and study them in practice. We model the dynamics of the electronic rolling shutter, showing that all the parameters necessary to carry out an attack can typically be inferred from the camera’s technical specifications. We validate our modeling with empirical data collected from executing the rolling shutter attack on two cameras, showing how such data matches the expectations. We then investigate the stealthiness of the rolling shutter attack, showing that compared to a blinding attack, it causes only small changes in a set of commonly-used metrics for measuring image distortion. Nevertheless, we describe how rolling shutter artifacts have an outside effect on the performance of a range of computer vision systems (illustrated in Figure 1). We evaluate the attack for the object detection task: first, we physically reproduce the attack on seven different cameras, followed by a broader evaluation showing that the rolling shutter attack significantly degrades the performance of state-of-the-art object detection networks, as shown in Figure 1a. Finally, we discuss countermeasures for the attack detection, including introducing a new countermeasure that re-uses part of the deep network for detection, and outline the attack limitations. In summary, we make the following contributions:

- We model the mechanisms behind the rolling shutter, we show that an attacker can accurately control the attack-injected distortions simply by inferring parameters from the target camera’s technical specifications. We validate our modeling against empirical data by collecting a large dataset of attack executions with real hardware.
- We perform an extensive evaluation of the rolling shutter attack’s effect on the object detection task, by partially simulating it on the BDD100K and the VIRAT dataset, showing that the integrity of object detection can be compromised.
- We assess the stealthiness of the attack across several metrics on the BDD100K video dataset, our analysis highlights how the rolling shutter attack is harder to detect in comparison to a blinding attack.
- We discuss a range of countermeasures and propose an accurate and lightweight enhancement to the backbone network of an object detector to detect rolling shutter attacks.

2 RELATED WORK

Many cyber-physical systems rely on the physical measurements of a wide variety of sensors ranging from temperature sensors through cameras to Light Detection and Ranging (LiDAR) [7, 21]. Since the integrity of sensor measurements is crucial for systems to behave as intended, research has been conducted to evaluate the vulnerabilities of sensors against signal injections on the physical layer. Physical-layer attacks can be categorized by modality [15] and substantial attention has been paid to each: electromagnetic radiation [49, 54], acoustic waves [48, 53, 56] and optical emission [6, 32, 33, 41, 44]. In this work, we focus on the latter category.

Due to their wide range, the focused brightness, and the targeting accuracy, lasers are often the first choice for attacks against optical sensors. [32] showed that an infrared laser could be used to manipulate the measurements of a drop sensor in a medical infusion pump. In [44], an amplitude-modulated laser was leveraged to inject voice commands into the microphone of a voice-controlled assistant. Lasers have also been used to inject spoofed points into the sensor of a LiDAR system, leading to false-positive object detection [6, 33, 41, 45], and to naively blind cameras necessary for various Advanced Driver Assistance Systems (ADASs) [33, 54].

In addition to attacks that directly target sensors on the physical layer, intensive research has been conducted that highlights the vulnerability of environment perception systems of modern cars. [28] demonstrated that cameras recognize objects projected onto the road as genuine objects or obstacles leading to unintended maneuvers. Instead of projecting objects onto the road, [26] highlighted that it is possible to project objects directly onto the camera lens.

Recently, [23] showed that the rolling shutter can be abused to create a backdoor in the training dataset for face recognition systems. The closest work to our own is in [37], which demonstrated that the rolling shutter can be exploited to cause misclassifications in single-object scenes, in which the attacker controls the light source for an enclosed environment. We go beyond the findings in that work by relaxing some assumptions to a more generalized case, in which the attacker can not enjoy full control over the scene lighting and must contend with further uncontrolled environmental and camera specific parameters, such as the precise ambient light level,




Figure 2: Visualization of the frame acquisition in the time-domain for the two shutter mechanisms.

frame rate and exposure time, along with cluttered, dynamic environments. We analyze these factors in detail and develop a reliable model of required conditions for an effective real-world attack. We further validate our method on a range of cameras and show that it works consistently, beyond the single target evaluated in [37]. We study the problem of rolling shutter attacks in a general setting and describe their applicability and effectiveness in unconstrained conditions.

3 IMAGE SENSOR TECHNOLOGY

3.1 Photodiode Arrays

An image sensor is composed of a two-dimensional array of unit cells, also known as pixels. The central component of each pixel is a photodiode, which absorbs incident photons and converts them into a signal charge depending on the intensity of the light beam. An image is captured by exposing each pixel in the photosensitive array to incident light for a certain amount of time. This time is known as the integration period [12], or *exposure time*. The longer the integration period, the more photons are converted into signal charge and the brighter the resulting image. To capture an image and control its brightness, cameras implement an electronic shutter mechanism. With electronic shutters, the image sensor is constantly exposed to light and signal charge is continuously generated. However, the accumulated charge is flushed before reading, such that only the light gathered during the specified integration time, between the reset and readout scan, represents the final image [20]. The length of the integration time is typically managed by an automatic exposure control system, in which the measurements from an integrated light meter are used to optimize the exposure period.

3.2 Scanning Element

The two most common image sensor technologies, Charge-Coupled Device (CCD) and Complementary Metal-Oxide Semiconductor (CMOS), use different techniques to recover the accumulated signal charge from the image array. The main difference is the arrangement of scanning elements and measurement units in the chip. Measurement units quantify the signal charge and convert it into electrical signals, whereas scanning elements control the order in which the captured signal charges are read [22]. While CCD sensors have a single measurement unit, CMOS sensors have a measurement unit in each pixel [27]. As a result, different electronic shutter implementations exist — *global* and *rolling* electronic shutters.

Global Shutter. Image sensors with an electronic global shutter simultaneously expose the entire photodiode array to incident light and also read out the accumulated signal charge from it at once. In CCD image sensors the individual pixels are connected via vertical and horizontal shift registers that work as memory buffers. The signal charge is temporarily stored inside the memory buffers, while the new integration period has already begun. Figure 2a illustrates the readout of a CCD image sensor with an electronic global shutter in the time-domain. The use of such memory buffer weighs on production cost and chip size [51, 57], which lead to CCD image sensors being most common in industrial and professional use [12]. In comparison, in CMOS sensors, the individual pixels are connected via metallic wires that cannot store the signal charge [22]. Even though it is possible to implement an electronic global shutter in CMOS sensors by changing the circuit and adding additional memory buffer, it is often not economically viable and would defeat the main advantage — cheap manufacturing that enables widespread use in consumer and semi-professional applications.

Rolling Shutter. Image sensors with an electronic rolling shutter expose and read each row of the photodiode array one by one. As a result, the readout of the last pixel row is slightly delayed compared to the first one. An exemplary readout scan showing the time delay between the rows is depicted in Figure 2b. Since CMOS image sensors do not have a memory buffer, they can only implement an electronic rolling shutter and thus suffer from image distortions as a consequence of the delay between reading the first and last pixel row — banding and geometric distortion [22]. The most common phenomena are skewed or compressed objects, which can be observed particularly well when fast-moving objects are captured. Another frequently occurring distortion is horizontal banding, which is especially common in indoor environments where dimmable LED lights are used. Since LEDs are often dimmed by switching them on and off at a high frequency, bright horizontal bars might appear in the captured image [10]. In our proposed rolling shutter attack, we exploit this phenomenon to inject controlled distortions into the captured image.

4 THREAT MODEL

In this paper, we focus on defining a general basis for a rolling shutter attack on camera-based intelligent systems.

Goal. The adversary’s general goal is to reduce the reliability and performance of camera-based intelligent systems in a way that is *stealthier* than a blinding attack (denial-of-service). In the simplest case, where the camera is within the attacker’s reach, it is possible to use paint, bullets, or direct physical force to prevent a camera from working properly. However, these truly brute-force attacks are mitigated by keeping a camera out-of-reach or using tamper sensors in the housing. They are also unattractive as they risk the attacker being caught on camera while conducting the attack. While it is straightforward to blind a camera with a very bright light, this approach suffers from two key problems. First, it is easily detectable with even simple methods such as ‘scene too bright/dark’ thresholds or ‘image global change’ comparisons, or by monitoring automatic exposure levels [47]. Second, in bright environments, it often provides incomplete disruption. The image is




Figure 3: Illustration of the row-wise acquisition of a single frame. The arrow on the left indicates the row that is currently exposed and readout. While reading the frame, light is injected during the exposure of a row, leading to distortions in a small part of the image.

completely disrupted in one area, but the effect rapidly diminishes away from that point; leaving much of the image unobscured. The residual color tint has little effect on computer vision systems such as object detection or face recognition, that typically make greater use of high-frequency image components [50].

By contrast, ours is a *stealthy* attack that minimizes observable disruption to the system, while in fact impairing its function, as we demonstrate in more detail in Section 7.2.

Capabilities. We assume the attacker has access only to off-the-shelf equipment to drive and modulate the laser and to aim it appropriately. We consider adversaries that have no control or access over the target camera, although they *may* be able to obtain knowledge of the camera hardware specifications, either from a datasheet or from a replica device. However, knowledge of camera parameters is not a requirement, as we explain in Section 5 and demonstrate on a range of cameras in Section 6.1. Similar to previous work [6, 33, 45] and as shown in [5], we consider that the adversary has line-of-sight with the target camera and the capabilities of aiming the laser at it. In contrast to [23, 37], which assume a rolling shutter attack in a static scene (i.e., single object, fixed distance and static lighting), we focus on a more general threat model. Following the assumption that the adversary cannot access the video of the target camera, we consider the attack to be executed in a black-box setting without any feedback channel. As such, distortions can only be injected at random locations, making it infeasible to target specific objects or apply optimization techniques. However, our approach makes the attack applicable to dynamic multi-object scenes.

5 EXPLOITING THE ROLLING SHUTTER

Preliminaries. The rolling shutter attack is based on the consecutive exposure and readout of individual pixel rows. Each row is exposed for a short period t_{exp} , often in the range of microseconds, before the exposure of the next row starts. The time delay Δt_{rst} between the reset pulses of consecutive lines can be exploited to inject noise that causes partial distortion of the captured image. If a modulated bright light, e.g., a laser, is pointed at the target camera, additional light can be selectively injected into rows of the image. The proportion of the image that is disturbed is controlled by the time t_{on} , which determines how long the laser is switched on (i.e., is emitting light). The attack is illustrated in Figure 3, which depicts the row-wise acquisition of a single frame. In Figure 3a, the readout starts at the top-most row, the laser is inactive, and consequently, no image disturbances occur. At some point, the laser is turned on and

the row that is currently exposed, e.g., the 7th row in Figure 3b, is illuminated and the original image information is perturbed. Before the next line is exposed, the laser is switched off, and the rest of the frame is captured without any further interference, see Figure 3c.

5.1 Controlling the Distortion Size

Shining a modulated light onto the image sensor will inject light in sets of adjacent pixel rows; we refer to each injection as a *distortion*. Depending on the use-case, the attacker might want to introduce large or narrow distortions. The appearance of the distortions depends on several factors. Without loss of generality, we focus on rolling shutters that scroll top-to-bottom along the image sensor. The same modeling applies for left-to-right rolling shutters, although the distortion would appear vertically rather than horizontally (as in [23]).

Mechanisms behind Distortion Size. The number of pixel rows that are illuminated by a distortion (*distortion size*) depends on four factors: (i) the absolute exposure time of individual rows t_{exp} , (ii) the time the laser is switched on t_{on} , (iii) the time between RST pulses Δt_{rst} , which is determined by the clock speed of the analog-to-digital converter, and (iv) the time delay between the last RST pulse and the moment the laser is switched on, o . Figure 4 shows how the four parameters interact to define the distortion size. Based on the four parameters, the adversary can precisely calculate the distortion size:

$$N_o = \left\lceil \frac{t_{exp} - o}{\Delta t_{rst}} \right\rceil + \left\lceil \frac{t_{on} + o}{\Delta t_{rst}} \right\rceil - 1, o \in [0, \Delta t_{rst}). \quad (1)$$

Since the attacked has no feedback channel, the laser cannot be synchronized with the rolling shutter of the target camera, so the offset o is unknown. Consequentially, the attacker can only learn an upper and lower bound for the number of consecutive rows that will be illuminated, that is:

$$N_{min} = \left\lceil \frac{t_{exp}}{\Delta t_{rst}} \right\rceil + \left\lceil \frac{t_{on}}{\Delta t_{rst}} \right\rceil - 2 \quad (2)$$

$$N_{max} = \left\lceil \frac{t_{exp}}{\Delta t_{rst}} \right\rceil + \left\lceil \frac{t_{on}}{\Delta t_{rst}} \right\rceil, \quad \text{with} \quad (3)$$

$$N_{min} \leq N_o \leq N_{max}.$$

In modern CMOS image sensors, the exposure time is generally much higher than the clock speed, $t_{exp} > \Delta t_{rst}$, meaning that the exposure times for consecutive rows overlap [34]. As a consequence, it is impossible to affect *individual* pixel rows with the laser, as the additional illumination will affect multiple rows at once, depending on the current t_{exp} . Figure 4 shows the frame acquisition for two different t_{exp} , a short and a long one, showing how larger t_{exp} leads to wider overlap across consecutive rows.

Distortion Repetition. The adversary can increase the number of injected distortions per-frame by increasing the laser frequency f . Intuitively, if the laser frequency f matches the frame rate F of the camera, each camera frame will present exactly one injected distortion. To cause interference at multiple parts in the frame at the




Figure 4: Overview of the parameters that influence the size of the injected distortion. The signals RST and OUT represent the clock cycles for the reset and readout scan, the PWM signal indicates when the laser is on.

same time, the attacker can increase f . The number of distortions per frame can be calculated as follows:

$$N_D = \frac{f}{F}. \quad (4)$$

When f is not an integer multiple of F , the location of the distortions will change over consecutive frames. For example, when $f > F$, the distortions will scroll downwards across consecutive frames. Depending on the attack scenario, the adversary can pick f accordingly to enhance, reduce or even prevent the movement of the distortions.

5.2 Estimating Parameters

While the mechanism to control the distortion size is consistent, in practice each camera has different parameters. This means that the adversary needs to estimate these parameters for a target camera to control the distortion size, using Equation 3 and Equation 4. While the attack can be carried out without accurately estimating these parameters, a better estimate will lead to the attack outcome being more predictable (i.e., more effective for the attacker). Across the introduced parameters, t_{on} and f are controlled directly by the adversary, Δt_{rst} and F are camera-dependent but static, t_{exp} is camera-dependent and dynamic. In the following, we describe how an adversary can infer the camera-dependent parameters using the information on the technical datasheet of a camera model.

Estimating Δt_{rst} and F . The operating frame rate F is generally available in the camera specifications (common examples are 25fps or 30fps). The clock speed Δt_{rst} is also sometimes available, nevertheless, when it is not, it can be estimated precisely using F and the full resolution of the image sensor:

$$\Delta t_{rst} = \frac{1}{F \cdot N_{rows}}, \quad (5)$$

where N_{rows} is the number of pixel rows of the image sensor. For some cameras, the size of the physical resolution of the photodiode array (total resolution) is larger than the maximum resolution used to record videos. For example, a 1920×1080 full-HD video may have been recorded on a camera with $N_{rows} > 1080$. The rolling shutter will still go through *all* N_{rows} pixel rows to capture additional image data, such as color information, but ignore a portion of them to read frames in a smaller resolution than the physical resolution; the ignored rows are known as *dead areas*. If the target image sensor has dead areas, Equation 4, used to calculate the number of injected

distortions, has to be extended by the ratio of visible rows $N_{visible}$ and total rows N_{rows} :

$$N_D = \frac{f}{F} \cdot \frac{N_{visible}}{N_{rows}}. \quad (6)$$

The sensor's total resolution N_{rows} and video resolution $N_{visible}$ can be found in the camera's technical specifications.

Estimating t_{exp} . During operation, t_{exp} is set automatically by an auto-exposure mechanism that balances the camera's t_{exp} with the current ambient light level: dimmer scenes will require larger t_{exp} to produce sufficient quality frames. Cameras have differing light sensitivities and lens apertures, and therefore use different t_{exp} for the same lighting conditions. To estimate t_{exp} , the adversary can use the *minimal luminous exposure* parameter H_v required by the camera to capture a usable frame and measure the current ambient light level E_v , with a light meter, at the time of the attack. Based on E_v and H_v , the exposure time t_{exp} can be calculated:

$$t_{exp} = \frac{H_v}{E_v}. \quad (7)$$

In Appendix A, we show examples of how we accurately estimated t_{exp} . In case H_v is not available, the adversary can measure the ambient light level E_v versus t_{exp} relationship by using a replica of the target camera and performing a set of measurements in controlled conditions. At the time of the attack, the adversary needs to measure E_v , lookup the respective t_{exp} that would be set by the target camera, and adjust the attack parameters accordingly. In Section 6.2, we demonstrate that imprecise estimates of t_{exp} do not affect the distortion size substantially.

6 ROLLING SHUTTER MODEL EVALUATION

This section describes the setup we used to execute the attack on a variety of cameras, a deeper analysis on the factors affecting the attack and the validation of our model presented in Section 5.

6.1 Preliminaries

Equipment. We only used off-the-shelf hardware that we bought online for less than \$90. An overview of the different components is depicted in Figure 5. We tested different lasers in terms of quality, wavelength (color), and power. The lasers ranged from cheap 650nm (red), 535nm (green) and 405nm (violet/blue) laser pointers with an output power of 5mW [2], to more expensive and powerful semi-professional 532nm (green) and 445nm (blue) continuous wave (CW) diodes with power ratings of up to 1.8W [1]. We noticed




Figure 5: The experimental setup we used for the collection of the different patterns and the evaluation of the attack on various cameras.

Table 1: Cameras we attempted to reproduce the attack on. The last column indicates whether we could reproduce the rolling shutter attack for the camera (✓) or not (X).

Camera	Resolution	F	Shutter	Success
Logitech C922	1920×1080	30	rolling	✓
Axis M3045-V	1920×1080	25	rolling	✓
YI Home Camera 1080p	1920×1080	24	rolling	✓
V380 Camera 720P	1280×720	25	rolling	✓
Google Nest	1920×1080	24	rolling	✓
Netgear Arlo	1280×720	24	rolling	✓
Imaging Source DFM 25G445	1280×960	25	global	X
Apple iPhone 7	3840×2160	30	rolling	✓

that an adversary might use green lasers to leverage the fact that most image sensors use a Bayer-Matrix color filter array and are therefore more sensitive to wavelengths at around 530nm [4].

Laser Modulation. The laser was modulated using a laser driver with Transistor-Transistor-Logic that was connected to the hardware Pulse Width Modulation (PWM) controller of a Raspberry Pi. A PWM modulated signal has two variables that can be changed — the frequency and the duty cycle. In our case, the frequency specifies how often the laser is switched on. The duty cycle, given in percent, determines the duration of a pulse, i.e., how long the laser is switched on each time the signal is set to high. In other words, the duty cycle specifies t_{on} . Since t_{on} depends on both frequency f and the duty cycle, we will use t_{on} in the rest of the paper rather than the duty cycle to make it easier to compare the injected distortions. Equation 8 shows how t_{on} is calculated, given the frequency f and the duty cycle D .

$$t_{on} = \frac{1}{f} D. \quad (8)$$

Applicability to Different Cameras. We tested whether we could reproduce the rolling shutter attack on seven different cameras with CMOS image sensors, ranging from cheap IoT cameras (YI Home Camera) over mid-range smart cameras (Google Nest) to semi-professional surveillance (Axis M3045-V) and smartphone cameras (Apple iPhone 7). For comparison, we also examined a camera with a CCD image sensor and an electronic global shutter mechanism. We printed an image showing a busy intersection with multiple objects, such as pedestrians and vehicles, placed it in front of each camera and captured a frame during normal operation as well as while pointing a modulated laser at it. The experimental

Table 2: Attack parameters we collected for the Logitech and Axis cameras. The last column reports whether the parameter can be controlled by the adversary at attack time. We used a blue-colored laser with 1.8W of power.

Param	Range [min, max]	Symbol	Adversary Controlled
Camera Exposure		t_{exp}	X
- Logitech	[100μs, 2,500μs]		
- Axis	[32μs, 1,000μs]		
Laser on-time		t_{on}	✓
- Logitech	[320μs, 16,000μs]		
- Axis	[50μs, 400μs]		
Laser Frequency		f	✓
- Logitech	[30Hz, 900Hz]		
- Axis	[25Hz, 750Hz]		
Camera fps		F	X
- Logitech	30fps		
- Axis	25fps		

setup is depicted in Figure 5 and examples of resulting attack frames are shown in Figure 15. Table 1 provides an overview of all tested cameras and whether we were able to physically inject rolling shutter patterns.

In-depth Rolling Shutter Pattern Collection. To further analyze the factors affecting the attack, we picked two different cameras which allowed us to control the camera exposure time t_{exp} : the Logitech C922, a common webcam, and the Axis M3045-V, a dome surveillance camera. We systematically reproduced the attack in controlled conditions by changing t_{exp} and t_{on} , which affect the resulting distortions. We placed the camera and the laser in a closed environment with little or no light, and collected a set of videos by varying the parameters of the attack, which are reported in Table 2. We set the laser frequency f to be slightly offset from the camera frame rate F so that over consecutive video frames, the location of the pattern slowly iterated over the image rows. This way, the set of collected frames covered all possible locations of the distortion in the image. Using a mask that filtered out pixels with a value per color channel ≤ 10 (range [0, 255]), we extracted the generated *pattern* of distortions from the captured video for each parameter configuration.

6.2 Accuracy of the Shutter Model

Here, we show how the equations introduced in Section 5.1 accurately model the number of rows affected by the distortion.

Estimated vs Empirical N_o . To estimate N_o with Equation 1 the adversary only needs Δt_{rst} ; as t_{on} is set by them and t_{exp} is known in this experiment (in the next paragraph we analyze the estimation of t_{exp}). We used Equation 5 to estimate Δt_{rst} . For the Logitech camera, this led to $\Delta t_{rst} = 46.3\mu s$. In contrast, the Axis camera has dead areas and requires N_{rows} to be determined before Δt_{rst} can be calculated. For most cameras and image sensors, manufacturers provide detailed information about the number of total, effective and recording pixels. However, for the Axis camera used in our experiment this information was not available. By calculating the ratio between visible and invisible distortions, we estimated N_{rows} to be 2,160 (1,080 visible and 1,080 invisible rows). We therefore used $N_{rows} = 2160$ for the Axis camera, resulting in $\Delta t_{rst} = 18.5\mu s$.

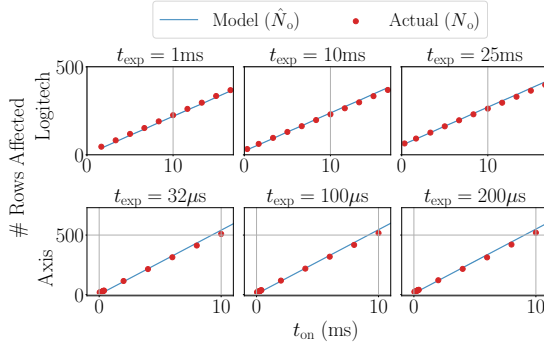


Figure 6: Comparison between the model-predicted number of affected rows and the empirical value found by measuring the number in practice on the two cameras.

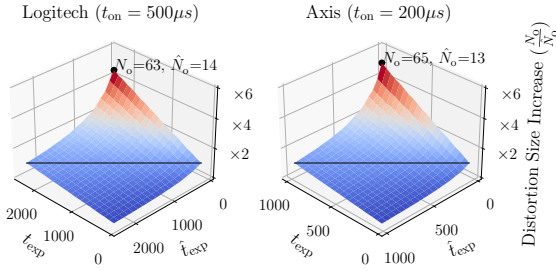


Figure 7: Increment in the expected distortion size \hat{N}_o and actual size N_o as the adversary's exposure time estimate t_{exp} diverges from the true exposure time value \hat{t}_{exp} .

The results in Figure 6, show how the prediction of our model closely matches the empirical outcome of the attack, for various t_{exp} , t_{on} , and the two considered cameras. We find this model to be precise regardless of the part of the image that is being hit. Figure 6 also shows the confidence intervals for the number of rows affected by the attack, computed over every frame of the collected videos.

Incorrect Estimation of t_{exp} . As mentioned in Section 5.2, the estimation of t_{exp} can be imprecise if information from the datasheet is missing, leading to changes in the resulting distortion size N_o . Using the rolling shutter model of the previous section, we computed the effect of an incorrect estimate of t_{exp} on the resulting N_o . We used N_{max} (see Equation 3) as the effect of o is negligible. Given an estimate of t_{exp} as \hat{t}_{exp} , and the resulting actual and estimated number of affected rows as N_o and \hat{N}_o , we report the relationship between these four variables in Figure 7. In the worst case, where the adversary underestimates \hat{t}_{exp} and t_{exp} is longest, this can lead to a distortion that is over 4× larger than expected. The relationship depicted in Figure 7 also depends on t_{on} . Additional plots for different t_{on} are shown in Appendix B. Nevertheless, we find that roughly 70% of the difference in distortion size is within a factor of two of the expected value (i.e., either twice as small or twice as large). This shows that large deviations occur only when the adversary's estimate is far off the real exposure value, showing that the attack's outcome is quite predictable under an educated guess.



Figure 8: Results for SSD under normal operation (left) and with an attack pattern overlaid (right). The different colors indicate how an object is affected by the attack: red=hidden, orange=misplaced, green=unaltered.

7 ATTACK EVALUATION

This section presents the impact of the rolling shutter attack on the *object detection* task, i.e., locating and identifying semantic objects. In addition, it compares the amount of interference introduced by our attack with a blinding attack. As described in our threat model in Section 4, fine-grained optimizations are infeasible, thus the adversary chooses attack parameters by quantifying the attack success offline.

7.1 Targeting Object Detection

Defining Attack Success. The goal of object detectors is to identify and locate objects in the input image with 2D boxes, which we refer to as $\{b\}_i^{(o)}$ when they belong to a legitimate image and $\{b\}_i^{(c)}$ when they belong to the attack-corrupted image. Figure 8 shows an example of the attack effect on the detected objects: in the rolling shutter attack-corrupted image, many objects are misdetected compared to a clean input. In practice, to compute the attack effect, for each box, we categorized the effect of the attack on each object box based on the Intersection-over-Union (IoU) between original and corrupted boxes. IoU is a measure of the overlapping area of two boxes, as a proportion of the total, combined area; it is commonly used to evaluate the accuracy of a box prediction compared to a ground truth box. We measured the attack outcome with the effect that the distortions have on the detected boxes with the following:

- **Hidden:** for a box $b_i^{(o)}$ its object has been *hidden* if there is no box in the corrupted image with $\text{IoU}(b_i^{(o)}, b_j^{(c)}) > 0.5$.
- **Misplaced:** for a box $b_i^{(o)}$ its object has been *misplaced* if there is a box in the corrupted image with IoU smaller than 0.95 and same output class.
- **Appeared:** for a box (in the corrupted image) $b_j^{(c)}$, its object has *appeared* if there is no box in the original image with $\text{IoU}(b_j^{(c)}, b_i^{(o)}) > 0.5$.

7.1.1 Method.

Attack Pattern Collection. We simulated the attack to find the best configuration of adversary-controlled parameters (i.e., t_{on} and f): the one that leads to the highest number of hidden objects. We collected the patterns as described in Section 6.1. Given that the Axis camera's frame rate is 25 fps ($F=25$), we collected patterns for four different laser frequencies, $f = 25, 250, 500$, and 750Hz. Since

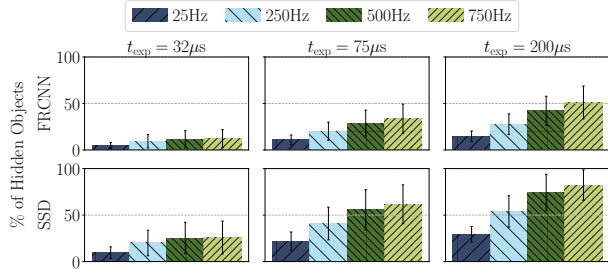


Figure 9: Percentage of hidden objects in the two video datasets for $f = 25, 250, 500$, and 750Hz , various t_{exp} and the two object detection models. Error bars show the standard deviation over the results per-video.

the image sensor of the Axis camera has dead areas, the number of distortions can be calculated using Equation 6. We collected patterns with duty cycles from 0.1% to 40% (from $1.3\mu s$ to $533.3\mu s$), because a change in f or duty cycle leads to different t_{on} (see Equation 8). It should be noted that t_{on} refers to the time the laser is on per-distortion rather than per-frame. Since t_{exp} affects the attack but is not controlled by the attacker — as it is set by the auto-exposure mechanism — we tested empirically which t_{exp} the camera would set on a sunny or cloudy day outdoors, finding $32\mu s$ and $200\mu s$, respectively. We used this range of t_{exp} to collect the patterns.

Simulation Setup. We simulated the attack effect on a subset of two different video datasets, namely BDD100K [55], a large and diverse driving video dataset, and VIRAT [30], a dataset of surveillance footage. We randomly picked 50 videos from the BDD100K and, due to the longer duration, 25 videos from the VIRAT dataset, for each video, we extracted every 10^{th} frame. For each parameter configuration (f , t_{exp} , t_{on}), we randomly selected ten patterns. Next, we generated approximately 7 million corrupted frames by overlaying the extracted patterns over the legitimate video frames. Finally, we assessed the simulated attack outcome by comparing legitimate and corrupted frames, as showcased in Figure 8. We evaluated two well-known state-of-the-art object detection models from the Tensorflow Model Zoo, namely Single Shot Detector (SSD) [24] and Faster RCNN [35] (FRCNN), both using Inception v2 backbone network. For performance reasons, we resized all frames to 640×360 pixels, before they were passed into the object detectors.¹

7.1.2 Results.

Choosing Attack Parameters. The percentage of objects hidden in the input frames for the various frequencies and exposure times is presented in Figure 9. It should be noted that for all four frequency settings, the image area covered by the laser-injected light is the same. As the frequency increases, t_{on} decreases, which results in smaller, but more frequent distortions (see Equation 8). However, the sum of t_{on} per frame stays the same. Figure 9 shows a clear increment in the ratio of hidden objects when using larger f . For $f = 750\text{Hz}$, we obtained the largest number of hidden objects. At $t_{exp} = 200\mu s$, 51% and 82% of objects were hidden for FRCNN and




Figure 10: Percentage of hidden, misplaced, and appeared objects in the evaluation of the BDD100K and the VIRAT dataset, for various t_{on} , exposure times t_{exp} , and the two models used in the evaluation (SSD and FRCNN). Shaded areas show 99% confidence intervals.

SSD, respectively. We noted that the narrower distortions generated by higher frequencies strongly affect the object detection task, in particular, for small objects. Given these results, an adversary would select 750Hz as the modulation frequency for the attack, as this selection better generalizes to different input scenes. In line with this, we limit the presentation of the remaining results to this 750Hz configuration.

Effect of Exposure. As the exposure time is not controlled by the adversary we investigated the effect of t_{exp} . Figure 10 shows the ratio of hidden, misplaced, or appeared objects for various exposure times t_{exp} and laser on-times t_{on} . With increasing t_{exp} and thus greater distortion, more objects are misplaced or hidden. As expected, an increasing t_{on} also leads to more objects being hidden. However, the number of misplaced or appeared objects does not change significantly. In contrast, for short t_{exp} less of the injected light is absorbed, leading to less intensive distortions and fewer hidden objects. While the weak illumination could be compensated with a more powerful laser, this highlights how in dimmer ambient light settings, the attack requires fewer resources on the adversary's side. For the largest evaluated exposure ($t_{exp} = 200\mu s$) and $t_{on} = 533\mu s$, our attack can hide up to 90% of objects in an image for SSD and up to 65% for FRCNN. This is in line with the fact that SSD is designed with more focus on speed of execution rather than prediction accuracy: SSD can execute one inference every 30ms, while FRCNN needs 89ms.

The impact that the rolling shutter attack can have on an autonomous driving platform as a whole, where object detection is used together with other components, is illustrated in Appendix C.

7.2 Comparison with Blinding Attack

7.2.1 Method. To measure the perturbation caused by our attack, we repeated parts of the simulation on the BDD100K dataset as described in Section 7.1. We randomly selected 100 videos from the validation set and extracted 10 evenly spaced frame pairs, whereas each pair consisted of two consecutive frames. For clarity,

¹Implementation and evaluation code is available at <https://github.com/ssloxford/they-see-me-rollin>

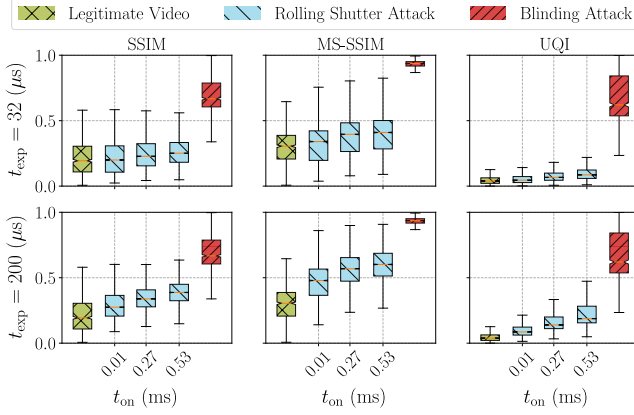


Figure 11: Image perturbation metrics between consecutive frames for $f = 750\text{Hz}$ during normal operation, a rolling shutter attack and a blinding attack. High values correspond to large perturbations.

we refer to those frames as *previous* and *current*. We then calculated the amount of perturbation between those consecutive frames to evaluate how the *current* frame has changed compared to the *previous* frame during (i) normal operation, (ii) a rolling shutter attack, and (iii) a blinding attack. For (ii) and (iii), we generated corrupted frames by overlaying the malicious patterns, as described in Section 7.1, onto the *current* frame of each pair. We measured the image perturbation with several metrics from adversarial machine learning literature [39]. For each scenario and pair of frames, we calculated three different metrics: the Structural Similarity (SSIM), Multi-Scale SSIM (MS-SSIM) and Universal Image Quality Index (UQI). The values for SSIM, MS-SSIM and UQI are bounded in $[0, 1]$ and were inverted in our case to represent *dissimilarity* rather than similarity. In addition, we removed perturbations with a value of 0, as these correspond to stationary frames. To measure the detection performance for pairs of anomalous frames, we set up a simple binary-threshold approach, i.e., if the dissimilarity metric value is higher than a certain threshold, we label the frame pair as anomalous. We report detection results with the area under the receiver operating characteristic curve (ROC-AUC), where a score of 1.0 corresponds to perfect detection while a score of 0.5 corresponds to random guess.

7.2.2 Results. Figure 11 shows the frame-to-frame perturbation introduced by the blinding and the rolling shutter attack for $t_{exp} = 32\mu\text{s}$ and $t_{exp} = 200\mu\text{s}$. The results indicate that for short t_{exp} our attack is causing interference that is similar across all metrics to the expected level of perturbation seen in consecutive legitimate video frames. In contrast, the blinding attack generates much larger perturbations compared to our attack, which are easily detected. We found that SSIM, MS-SSIM and UQI could detect almost 100% of blinding attacks with the binary-threshold approach, i.e., ROC-AUC of 0.99, 1.0 and 1.0, respectively. While higher t_{exp} causes more perturbation, the level of interference is still closer to that for legitimate video frames than for blinding. In comparison, for most

cases our attack was indistinguishable from legitimate frame-to-frame changes when using the same threshold approach. The ROC-AUC scores for the detection of the rolling shutter attack with $t_{exp} = 32\mu\text{s}$ and $t_{on} = 0.53\text{ms}$ are 0.61, 0.68 and 0.80, for SSIM, MS-SSIM and UQI, respectively. Similar results (0.67, 0.79, 0.81) were observed for $t_{exp} = 200\mu\text{s}$ and $t_{on} = 0.01\text{ms}$. However, as expected, the attack is more easily detected with increasing t_{on} . Nevertheless, choosing the worst performing scenario for the rolling shutter attack at the 0% false positive rate threshold ($t_{on} = 0.53\text{ms}$, $t_{exp} = 32\mu\text{s}$ and using MS-SSIM), we obtain that 100% of blinding attacks are detected while only 6.55% of rolling shutter attacks are.

8 COUNTERMEASURES

The vulnerability that enables the rolling shutter attack is inherent to CMOS image sensors, meaning that attacks cannot be prevented entirely without changing the hardware design. However, detection in software can mitigate the attacks' impact, while other countermeasures are applicable in some special cases.

8.1 Hardware Countermeasures

A clear design option is to use a camera with a global shutter mechanism. This does not stop the injected light but prevents rolling shutter attacks as the image rows are exposed at the same time, leading to a simple blinding attack. However, this requires a design-time choice and also carries additional component costs for incorporating frame memory or using CCD image sensors. These costs increase dramatically for existing systems, as the replacements must be retrofitted as well.

Additionally, camera redundancy is a straightforward protection approach, but also quickly increases manufacturing costs for little improvement in security. In some scenarios, the existing use of multiple sensors allows the system to continue operating effectively despite the performance degradation caused by the rolling shutter attack (e.g., in autonomous driving, with multiple cameras or LiDAR). However, it is not trivial to reliably determine which sensor input is malicious or untrustworthy in order to stop relying on it. Furthermore, while a system may be tolerant to the degradation of its sensor inputs, there is still some finite budget for degradation that it can withstand before it becomes unsafe or ineffective.

8.2 Deep Learning Based Detection

Here, we explore the re-use of the deep network used for the computer vision task to detect the presence of the attack. We re-use the backbone network using the insight that appearance of the distortions can be captured by early convolutional filters, designed to capture low-level input features.

Setup and Training. Given the backbone network, we added a network head as follows: on top of the first convolutional layer (after pooling, when present) we added one additional convolution layer with 32 filters, then we used global maximum pooling and added a binary classification layer which predicts whether the input is corrupted by a rolling shutter attack. We also used dropout and selected leaky ReLu as the activation function, which has been found to lead to faster convergence and improved performance [52]. We created a dataset to train our detection by splitting both collected patterns and video frames into training, validation and testing

Table 3: Rolling shutter attack detection performance for the backbone networks. False Negatives (FN) and False Positives (FP) are calculated over 3,560 test frames. The column *Overhead* represents the computational overhead for the detection mechanism in percentages with the standard deviation in parentheses.

Model	Backbone	Accur.	FN	FP	Overhead in % ($\pm\sigma$)
SSD	MobileNet v2	0.98820	32	10	+0.57 (± 4.059)
SSD	Inception v2	0.99494	18	0	+0.13 (± 4.242)
FRCNN	Inception v2	0.99551	7	9	-0.11 (± 5.045)
FRCNN	Resnet-50	0.98989	25	11	+1.38 (± 5.119)

subsets, with ratios 60%, 20% and 20%, respectively. As in Section 7.1, corrupted frames were obtained by overlaying the pattern onto the legitimate frame. We used all patterns collected for the Axis camera ($\sim 6,500$, see Section 6.1) and 180 videos randomly selected from BDD100K, which we sampled every 10^{th} frame (20,880 frames). We used cross entropy loss and trained the added parameters with Adam, with learning rate 10^{-4} and for 10 epochs.

Results. The test set accuracies for the two object detectors with different backbones are presented in Table 3. The detection is efficient, all networks obtain over 98.8% accuracy on the task, although MobileNet v2 performs slightly worse due to design compromises for faster execution speed. Table 3 also reports the number of False Negatives (FN) and False Positives (FP). This shows that setting the detection threshold to 0.5 would generally tend to favor low false positives, but a different trade-off can be sought by changing the threshold for detection.

To analyze the performance impact of the proposed detection mechanism, we measured the required inference time for all images in the test set (batch size = 1) with and without the added layers on an NVIDIA TITAN RTX. The last column of Table 3 shows the computational overhead of the detection, indicating that it is negligible. As such, this detection measure could be applied as a software patch to protect existing systems. Additionally, the method represents a general approach that can be applied to protect any computer vision network from rolling shutter attacks.

9 LIMITATIONS

In this section, we provide an overview of the limitations of our evaluation and the rolling shutter attack itself.

9.1 Evaluation Limitations

The evaluation presented in this paper is a first step in raising awareness about the rolling shutter attack and its impact. Nevertheless, our attack evaluation has some limitations that an attacker might need to consider when performing the attack. Although the patterns we used have been collected from real cameras under close to real-world settings, our simulations do not perfectly reflect the dynamic behavior of the target camera under attack. Depending on the camera, numerous factors and mechanisms may influence the appearance of the injected distortion. We mimicked the behavior by exhaustively simulating the attack with different patterns that correspond to different exposure times. Nevertheless, depending

on the scenario, an attacker might not always freely choose parameters f and t_{on} while maintaining the desired brightness of the injected distortions. For instance, in bright settings, short t_{on} and non-powerful lasers might not generate enough light to be captured by the camera. Furthermore, we have not explicitly considered the effect of auto-exposure and auto-focus mechanisms. While we found that these are generally not triggered for most of our attack configurations, we cannot exclude that these mechanisms might come into play as the attack is executed (this will depend on camera-specific factors). This will have an effect on the appearance of the distortions and may cause warnings in the vision-based system.

9.2 Attack Limitations

Due to the characteristics of the rolling shutter attack, there are some limitations that cannot be easily solved.

Limitations in Dim Conditions. Contrary to what humans perceive, the amount of ambient light does not scale linearly between what is perceived as bright versus what is perceived as dark. For example, on a clear day, the amount of ambient light at midday is $>30\text{k lx}$, at sunset, it is $\sim 400\text{ lx}$ while at night it is close to 0 lx . Cameras are no exception: in poor lighting conditions, cameras have to increase the exposure time *dramatically* to capture enough light. This leads to a large overlap of consecutive pixel rows resulting in fine-grained distortions not being achievable. Therefore, in dim conditions, even a very short t_{on} will impact a significant proportion of image rows, limiting the amount of high-frequency distortions from which the attack benefits.

Timing & Synchronization. While it is, in principle, possible to target specific rows during the attack, tight synchronization between the modulation signal and the row reset signal of the image sensor is required. Altering the content of the i -th row to hide a specific object, requires exact knowledge about the exposure timing. In case the attacker has access to the video feed, such synchronization could be obtained through an iterative process. However, as in our threat model, it is generally reasonable to assume that the attacker does not have this information and therefore needs to design their attack accordingly.

Frame Rate Accuracy. To control the movement of distortions between consecutive frames, the laser modulation must vary in relation to the camera frame rate, or a multiple of it. Accurate knowledge of the frame rate allows the attacker to keep the progression of the pattern in line with their intentions. The frame rate F of a camera is typically specified in the datasheet, although minor timing inaccuracies might cause the frame rate to deviate slightly (e.g., 30.0025 fps instead of 30 fps). However, this only causes the pattern to drift slowly if the attack is carried out over an extended period of time and is negligible for attacks of short duration.

10 CONCLUSION

In this paper, we have highlighted an inherent vulnerability of the electronic rolling shutter mechanism implemented in most modern CMOS image sensors, which can be exploited to inject fine-grained distortions into the captured frames. We showed how an adversary can construct an attack in practice by introducing a

model to accurately predict the size of the injected distortions. We describe how this model can be used to account for uncontrolled camera-specific and environmental variations, in order to make the attack more reliable. Using only off-the-shelf equipment, we reproduced the attack steps in practice on two separate cameras. We then investigated the attack within the context of an object detection task, where the camera-captured images are used by computer vision models to locate and classify objects. Based on two well-known video datasets, we evaluated the performance of two state-of-the-art object detection models under the rolling shutter attack, showing that for SSD, our attack can hide up to 75% of objects present in the input frames. At the same time, we showed that our attack is stealthier than a global blinding attack across several anomaly metrics, making it harder to detect. Our findings outline a weakness in using cameras with CMOS images sensors in settings with potentially-adversarial camera inputs, and particularly for computer vision applications in safety-critical systems.

ACKNOWLEDGMENTS

This work was supported by a grant from Mastercard and by the Engineering and Physical Sciences Research Council (EPSRC).

REFERENCES

- [1] Amazon.com. [n. d.]. 1.8 W A M140 445 Nm blue Diode. <https://www.amazon.com/dp/B00UECEMEE>.
- [2] Amazon.com. [n. d.]. JING TAI Powerful Pointer. <https://amazon.com/dp/B0932TQMNZ>.
- [3] Axis Communications. 2020. Network video solutions for all industries. <https://www.axis.com/en-us/solutions-by-industry>.
- [4] BE Bayer. 1976. Colour filter array. *United States of America patent* 3971065 (1976).
- [5] Yulong Cao, Jiaxiang Ma, Kevin Fu, Rampazzi Sara, and Morley Mao. [n. d.]. Automated Tracking System For LiDAR Spoofing Attacks On Moving Targets. ([n. d.]).
- [6] Yulong Cao, Yimeng Zhou, Qi Alfred Chen, Chaowei Xiao, Won Park, Kevin Fu, Benjamin Cyr, Sara Rampazzi, and Z. Morley Mao. 2019. Adversarial sensor attack on LiDAR-based perception in autonomous driving. *Proceedings of the ACM Conference on Computer and Communications Security* (2019), 2267–2281. <https://doi.org/10.1145/3319535.3339815> arXiv:1907.06826
- [7] Yaobin Chen and Lingxi Li. 2013. *Advances in Intelligent Vehicles*. <https://doi.org/10.1016/C2011-0-07400-8>
- [8] Anthony Cuthbertson. 2019. Hong Kong protesters use lasers to avoid facial recognition cameras and blind police. <https://www.independent.co.uk/news/world/asia/hong-kong-protests-lasers-facial-recognition-ai-china-police-a9033046.html>.
- [9] Jifeng Dai, Yi Li, Kaiming He, and Jian Sun. 2016. R-fcn: Object detection via region-based fully convolutional networks. In *Advances in neural information processing systems*. 379–387.
- [10] Henry Dietz and Paul Eberhart. 2019. Shuttering methods and the artifacts they produce. *Electronic Imaging* 2019, 4 (2019), 590–591.
- [11] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. 2017. CARLA: An Open Urban Driving Simulator. In *Proceedings of the 1st Annual Conference on Robot Learning*. 1–16.
- [12] Daniel Durini. 2019. *High performance silicon imaging: Fundamentals and applications of CMOS and CCD sensors*. <https://doi.org/10.1016/C2017-0-01564-1>
- [13] EVERGUARD.ai. 2020. sentri360. <https://everguard.ai/sentri360.php>.
- [14] Kevin Eykholt, Ivan Evtimov, Earlene Fernandes, Bo Li, Amir Rahmati, Florian Tramèr, Atul Prakash, Tadayoshi Kohno, and Dawn Song. 2018. Physical adversarial examples for object detectors. In *12th USENIX Workshop on Offensive Technologies (WOOT 18)*.
- [15] Ilias Giechaskiel and Kasper Rasmussen. 2019. Taxonomy and challenges of out-of-band signal injection attacks and defenses. *IEEE Communications Surveys & Tutorials* 22, 1 (2019), 645–670.
- [16] Ionel Gog, Sukrit Kalra, Peter Schaffhalter, Matthew A Wright, Joseph E Gonzalez, and Ion Stoica. 2021. Pylot: A Modular Platform for Exploring Latency-Accuracy Tradeoffs in Autonomous Vehicles. *arXiv preprint arXiv:2104.07830* (2021).
- [17] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572* (2014).
- [18] Google Nest. 2020. Nest Cam. https://store.google.com/us/product/nest_cam_outdoor?hl=en-US.
- [19] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. 2017. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*. 2961–2969.
- [20] T Kinugasa, M Noda, T Imaide, I Aizawa, Y Todaka, and M Ozawa. 1987. An Electronic Variable-Shutter System in Video Camera Use. *IEEE Transactions on Consumer Electronics* CE-33, 3 (1987), 249–255.
- [21] Eric D Knapp and Joel Thomas Langill. 2014. *Industrial Network Security: Securing critical infrastructure networks for smart grid, SCADA, and other Industrial Control Systems*. Syngress.
- [22] Takao Kuroda. 2017. *Essential principles of image sensors*. <https://doi.org/10.1201/b17411>
- [23] Haoliang Li, Yufei Wang, Xiaofei Xie, Yang Liu, Shiqi Wang, Renjie Wan, Lap-Pui Chau, and Alex C Kot. 2020. Light Can Hack Your Face! Black-box Backdoor Attack on Face Recognition Systems. *arXiv preprint arXiv:2009.06996* (2020).
- [24] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. 2016. Ssd: Single shot multibox detector. In *European conference on computer vision*. Springer, 21–37.
- [25] Giulio Lovisotto, Henry Turner, Ivo Slugaonovic, Martin Strohmeier, and Ivan Martinovic. 2021. SLAP: Improving Physical Adversarial Examples with Short-Lived Adversarial Perturbations. In *30th USENIX Security Symposium (USENIX Security 21)*.
- [26] Yanmao Man, Ming Li, and Ryan Gerdes. 2020. GhostImage: Remote Perception Attacks against Camera-based Image Classification Systems. In *23rd International Symposium on Research in Attacks, Intrusions and Defenses (RAID 2020)*. USENIX Association, San Sebastian, 317–332. <https://www.usenix.org/conference/raid2020/presentation/man>
- [27] Junichi Nakamura. 2006. *Image sensors and signal processing for digital still cameras*. 1–336 pages. <https://doi.org/10.1201/9781420026856>
- [28] Ben Nassi, Dudi Nassi, Raz Ben-Netanel, Yisroel Mirsky, Oleg Drokun, and Yuval Elovici. 2020. Phantom of the ADAS: Phantom Attacks on Driver-Assistance Systems.
- [29] Dudi Nassi, Raz Ben-Netanel, Yuval Elovici, and Ben Nassi. 2019. MobilBye: Attacking ADAS with Camera Spoofing. *arXiv preprint arXiv:1906.09765* (2019).
- [30] Sangmin Oh, Anthony Hoogs, Amitha Perera, Naresh Cuntoor, Chia-Chih Chen, Jong Taek Lee, Saurajit Mukherjee, JK Aggarwal, Hyungtae Lee, Larry Davis, et al. 2011. A large-scale benchmark dataset for event recognition in surveillance video. In *CVPR 2011*. IEEE, 3153–3160.
- [31] BBC Online. 2020. Met Police to deploy facial recognition cameras. <https://www.bbc.co.uk/news/uk-51237665>.
- [32] Youngseok Park, Yunmok Son, Hocheol Shin, Dohyun Kim, and Yongdae Kim. 2016. This ain't your dose: Sensor spoofing attack on medical infusion pump. In *10th USENIX Workshop on Offensive Technologies (WOOT 16)*.
- [33] Jonathan Petit, Bas Stottelaar, Michael Feiri, and Frank Kargl. 2015. Remote attacks on automated vehicles sensors: Experiments on camera and lidar. *Black Hat Europe* 11 (2015), 2015.
- [34] Qlming. 2014. Rolling Shutter vs. Global Shutter. <http://www.tecnicaenlaboratorios.com/Qlming/brochures/RollingvsGlobalShutter.pdf>.
- [35] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*. 91–99.
- [36] Ring. 2020. Ring Camera. <https://ring.com>.
- [37] Athena Sayles, Ashish Hooda, Mohit Gupta, Rahul Chatterjee, and Earlene Fernandes. 2021. Invisible Perturbations: Physical Adversarial Examples Exploiting the Rolling Shutter Effect. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 14666–14675.
- [38] Blue Iris Security. 2020. Video Security. <https://blueirissoftware.com/>.
- [39] Mahmood Sharif, Lujo Bauer, and Michael K Reiter. 2018. On the suitability of lp-norms for creating and preventing adversarial examples. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. 1605–1613.
- [40] Sam Shead. 2019. Chinese residents worry about rise of facial recognition. <https://www.bbc.co.uk/news/technology-50674909>.
- [41] Hocheol Shin, Dohyun Kim, Yujin Kwon, and Yongdae Kim. 2017. Illusion and dazzle: Adversarial optical channel exploits against lidars for automotive applications. In *International Conference on Cryptographic Hardware and Embedded Systems*. Springer, 445–467.
- [42] Dawn Song, Kevin Eykholt, Ivan Evtimov, Earlene Fernandes, Bo Li, Amir Rahmati, Florian Tramèr, Atul Prakash, and Tadayoshi Kohno. 2018. Physical Adversarial Examples for Object Detectors. In *12th USENIX Workshop on Offensive Technologies (WOOT 18)*. USENIX Association, Baltimore, MD. <https://www.usenix.org/conference/woot18/presentation/eykholt>
- [43] Nick Statt. 2020. Amazon is expanding its cashierless Go model into a full-blown grocery store. <https://www.theverge.com/2020/2/25/21151021/amazon-go-grocery-store-expansion-open-seattle-cashier-less>.
- [44] Takeshi Sugawara, Benjamin Cyr, Sara Rampazzi, Daniel Genkin, and Kevin Fu. 2020. Light commands: laser-based audio injection attacks on voice-controllable systems. In *29th USENIX Security Symposium (USENIX Security 20)*. 2631–2648.

- [45] Jiachen Sun, Yulong Cao, Qi Alfred Chen, and Z Morley Mao. 2020. Towards robust lidar-based perception in autonomous driving: General black-box adversarial sensor attack and countermeasures. In *29th USENIX Security Symposium (USENIX Security 20)*. 877–894.
- [46] Synology. 2020. Synology Surveillance Station. <https://synology.com/en-us/surveillance>.
- [47] Bosch Security Systems. 2016. FW 6.30 Tamper Detection. https://resources-boschsecurity-cdn.azureedge.net/public/documents/TN_VCA_tamper_detect_WhitePaper_enUS_22996235531.pdf.
- [48] Timothy Trippel, Ofir Weisse, Wenyuan Xu, Peter Honeyman, and Kevin Fu. 2017. WALNUT: Waging doubt on the integrity of MEMS accelerometers with acoustic injection attacks. In *2017 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 3–18.
- [49] Yazhou Tu, Sara Rampazzi, Bin Hao, Angel Rodriguez, Kevin Fu, and Xiali Hei. 2019. Trick or heat? Manipulating critical temperature-based control systems using rectification attacks. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. 2301–2315.
- [50] Haohan Wang, Xindi Wu, Zeyi Huang, and Eric P Xing. 2020. High-frequency Component Helps Explain the Generalization of Convolutional Neural Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 8684–8694.
- [51] Martin Wany and Georg P Israel. 2003. CMOS image sensor with NMOS-only global shutter and enhanced responsivity. *IEEE Transactions on electron devices* 50, 1 (2003), 57–62.
- [52] Bing Xu, Naiyan Wang, Tianqi Chen, and Mu Li. 2015. Empirical evaluation of rectified activations in convolutional network. *arXiv preprint arXiv:1505.00853* (2015).
- [53] Wenyuan Xu, Chen Yan, Weibin Jia, Xiaoyu Ji, and Jianhao Liu. 2018. Analyzing and enhancing the security of ultrasonic sensors for autonomous vehicles. *IEEE Internet of Things Journal* 5, 6 (2018), 5015–5029.
- [54] Chen Yan, Wenyuan Xu, and Jianhao Liu. 2016. Can you trust autonomous vehicles: Contactless attacks against sensors of self-driving vehicle. *DEF CON* 24, 8 (2016), 109.
- [55] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. 2020. BDD100K: A diverse driving dataset for heterogeneous multitask learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2636–2645.
- [56] Guoming Zhang, Chen Yan, Xiaoyu Ji, Tianchen Zhang, Taimin Zhang, and Wenyuan Xu. 2017. Dolphinattack: Inaudible voice commands. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. 103–117.
- [57] Zhimin Zhou, Bedabrata Pain, and Eric R. Fossum. 1997. Frame-transfer CMOS active pixel sensor with pixel binning. *IEEE Transactions on Electron Devices* (1997). <https://doi.org/10.1109/16.628834>
- [58] Zhe Zhou, Di Tang, Xiaofeng Wang, Weili Han, Xiangyu Liu, and Kehuan Zhang. 2018. Invisible mask: Practical attacks on face recognition with infrared. *arXiv preprint arXiv:1803.04683* (2018).

Table 4: Results of the cross-correlation between the color channel histograms of the frames captured with auto-exposure, manual exposure setting and manual exposure setting with an offset of 60 μ s.

Color Channel	Manual	Manual + Offset
Red	0.9961	0.9858
Green	0.9986	0.9865
Blue	0.9939	0.9545

A ESTIMATING t_{exp}

On the basis of the Axis M3045-V, we demonstrate how the adversary can estimate the exposure time t_{exp} . Moreover, we show that the method leads to very accurate results, by comparing frames captured with auto-exposure and manually set exposure.

Method. In the first step, we extracted $H_v = 0.25lx$ from the datasheet. We then collected multiple frames with the auto-exposure mechanism. Subsequently, we deactivated the auto-exposure mechanism, measured the current ambient light level E_v with a lux meter and used Equation 7 to estimate t_{exp} . We repeated the frame collection for different lighting conditions, such as at sunset, on a cloudy day and a sunny day.

Accuracy of Estimation. On a first glance, there was no optical difference between the frames collected with auto-exposure and the ones collected with manual settings. To analyze the accuracy, we calculated histograms for each color channel (R, G and B) of the collected frames. We then calculated the cross-correlation between the color histograms of the frames captured with auto-exposure and the ones with manually set exposure. The results in Table 4 show that there is only a small difference between the frames captured with auto-exposure and the ones with manual exposure. To make the results more comparable, we also collected frames with an exposure of $t_{exp} + 60\mu s$. We picked an offset of 60 μs , because this was the first value where we noticed considerable changes in brightness. As can be seen in Table 4, an offset of 60 also led to a noticeable difference in the cross-correlation of all three color channels. This suggests that the estimation of t_{exp} is accurate.

B INCORRECT t_{exp} - ADDITIONAL RESULTS

As reported in Section 6.2, incorrect estimates of the exposure time can lead to distortion sizes that diverge from the expected size. Here we report the best and worst divergence observed for the parameters used in our data collection, which depends on t_{on} . In fact, larger t_{on} will lead to smaller variation, and vice-versa. The shortest t_{on} is reported in Figure 12, the longest t_{on} is reported in Figure 13. The figures show that for larger t_{on} (which also generally produce better attack success in our evaluation, Section 7), incorrect estimates tend to lead to diminishing divergence in expected versus actual distortion size. For Logitech we get a worst-case distortion size increase of 1.6 (down from 6 in the shortest t_{on} setting), while the same increase for Axis goes from a factor of 11 (in the shortest t_{on} setting) down to a factor of 3.

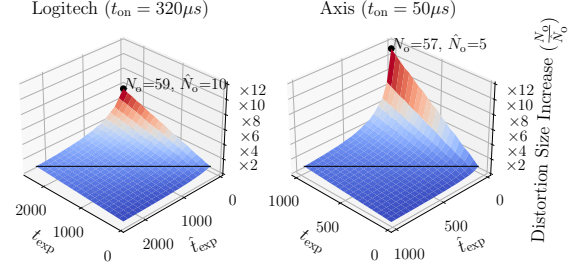


Figure 12: As in Figure 7, increment in the expected distortion size \hat{N}_o and actual size N_o as the adversary's exposure time estimate t_{exp} diverges from the true exposure time value \hat{t}_{exp} . Shortest t_{on} used in the data collection of Section 6.

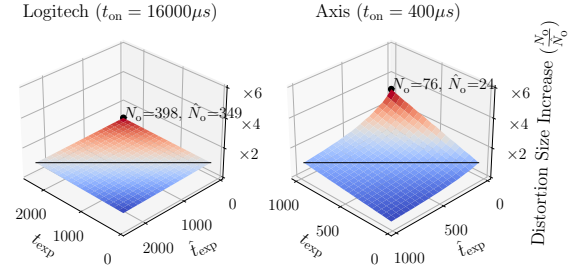


Figure 13: As in Figure 7, increment in the expected distortion size \hat{N}_o and actual size N_o as the adversary's exposure time estimate t_{exp} diverges from the true exposure time value \hat{t}_{exp} . Longest t_{on} used in the data collection of Section 6.

C EFFECT ON AUTONOMOUS DRIVING

We used the open-source simulator CARLA [11], in combination with the self-driving car platform Pyrot [16] to illustrate the effects of the rolling shutter attack on an autonomous system. By default, the Pyrot agent is equipped with four sensors: (i) a wide-angle RGB camera for lane detection and object detection/tracking, (ii) a telephoto camera for traffic light detection, (iii) a LiDAR sensor for localization, and (iv) GPS for route planning. We uniquely targeted the wide-angle RGB camera with our attack, to measure how the system as a whole is affected.

CARLA Setup. We utilized two predefined scenarios from the CARLA scenario runner, where an entity is crossing the road in front of the car: *Scenario 1* involves a pedestrian and *Scenario 2* involves a cyclist. We collected the baseline performance of Pyrot by simulating each scenario 50 times. We then repeated the simulation under the presence of the rolling shutter attack. We intercepted the frames captured by the wide-angle camera as they were forwarded to Pyrot and overlaid the rolling shutter pattern, as described in Section 7.1.1. The attack parameters were as follows: $f = 750Hz$, $t_{exp} = 200\mu s$ and $t_{on} = 0.53ms$. In line with the previous evaluation, we again used two object detectors, FRCNN and SSD, which are among the pre-implemented options in Pyrot.




Figure 14: Velocity curve of the agent during normal operation (baseline) and during the rolling shutter attack. The grey area shows the time when the pedestrian/cyclist crosses the road. The dashed line marks the time when the agent collides with the non-player character.

Results. Most of the key findings from Section 7 can be transferred to the CARLA simulations. In particular, the performance difference between SSD and FRCNN stands out. While for *Scenario 1* during normal operation for both object detectors the probability of a collision was equal (5%), under the rolling shutter attack, SSD showed a 67% probability of a collision compared to 47% for FRCNN. Similar behavior can be observed for the baseline of *Scenario 2*. FRCNN showed a 0% probability of a collision, while using SSD led to an increase to 15%. However, interestingly, under attack, FRCNN experienced a strong performance drop and fell behind SSD. The probability of a safety infraction increased to 97% compared to 65% for SSD.

In Figure 14, the velocity curve of the autonomous vehicle for the two scenarios and object detectors is presented. The figure shows that in the baseline, the car immediately reduces speed as soon as the non-player character starts moving. In contrast, as the braking gap shows, during the attack, a noticeable delay in braking can be observed.



Figure 15: Physical evaluation of the attack. These images were collected during a physical experiment where we printed images, placed them in front of the camera, and carried out the attack with the laser (see Figure 5). The first column reports the object detection results on FRCNN, the second column on SSD.