WIP: Threat Modeling Laser-Induced Acoustic Interference in Computer Vision-Assisted Vehicles

Nina Shamsi*, Kaeshav Chandrasekar†, Yan Long†, Christopher Limbach†, Keith Rebello‡, Kevin Fu*
*Northeastern University, †University of Michigan, ‡The Boeing Company
{shamsi.n, k.fu}@northeastern.edu, {yanlong, kaeshavc, limbach}@umich.edu, keith.rebello@boeing.com

***Abstract*—Control or disablement of computer vision-assisted autonomous vehicles via acoustic interference is an open problem in vehicle cybersecurity research. This work explores a new threat model in this problem space: acoustic interference via high-speed, pulsed lasers to non-destructively affect drone sensors. Initial experiments verified the feasibility of laser-induced acoustic wave generation at resonant frequencies of MEMS gyroscope sensors. Acoustic waves generated by a lab-scale laser produced a 300-fold noise floor modification in commercial off-of-the-shelf (COTS) gyroscope sensor readings. Computer vision functionalities of drones often depend on such vulnerable sensors, and can be a target of this new threat model because of camera motion blurs caused by acoustic interference. The effect of laser-induced acoustics in object detection datasets was simulated by extracting blur kernels from drone images captured under different conditions of acoustic interference, including speaker-generated sound to emulate higher intensity lasers, and evaluated using state-of-the-art object detection models. The results show an average of 41.1% decrease in mean average precision for YOLOv8 across two datasets, and suggest an inverse relationship between an object detection model's mean average precision and acoustic intensity. Object detection models with at least 60M parameters appear more resilient against laser-induced acoustic interference. Initial characterizations of laser-induced acoustic interference reveal future potential threat models affecting sensors and downstream software systems of autonomous vehicles.**

I. Introduction

Advances in imaging hardware and computer vision (CV)-assisted controls have revolutionized capabilities for unmanned vehicles to navigate and monitor ground, underwater, and aerial environments. These advances, however, also raise concerns about lowering the cost of trespassing, malicious reconnaissance, and terrorism, and motivate the need for counter-unmanned vehicle technologies that can neutralize unwanted vehicles in protected spaces [23]. High-energy lasers can destroy targets from several kilometers away, and are thus far one of the most used unmanned aerial vehicle (UAV) mitigation techniques, but can neither be used for damage-free UAV disablement [40], nor for preventing lateral damage caused by crashing a UAV. Unmanned vehicle sensors are inherently vulnerable to modulated acoustic waves. Electronic speakers can generate non-destructive acoustic waves to modify outputs of micro-electromechanical system (MEMS) sensors, e.g., inertial measurement units (IMUs) including accelerometers, and gyroscopes, [36], and thus control or disable unmanned vehicles' functionalities. However, a major limitation of speaker-generated acoustic waves is their poor directionality and relatively short interference distance. The long-range capabilities of lasers have also been used to inject audio into voice-controllable systems due to a vulnerability of MEMS microphones [31], but the affect of laser-induced acoustic pressure on optical systems due to MEMS gyroscope destabilization is as yet unexplored. Given this, we explore a new threat model of acoustic interference attacks: acoustic wave generation using high-speed, pulsed lasers to non-destructively affect MEMS gyroscopes, and achieve potentially long-range disabling interference against computer vision-assisted drones.

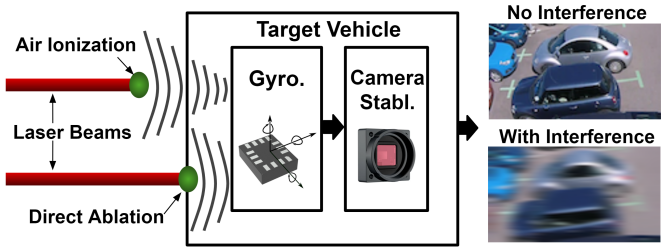

Fig. 1. Pulsed lasers generate acoustic waves that can change gyroscope sensor output to interfere with drone camera stabilization, resulting in motion blurs in images captured by the drone.

A key challenge for the laser-induced acoustic interference threat model is to characterize its impact on commercial devices with accessible and sustainable lab-scale setups. Initial experiments demonstrated 3 rad/s deviations with a 0.01 rad/s noise floor from acoustic wave generation in a COTS gyroscope under direct ablation and air ionization using a high-speed, pulsed laser. Moreover, evaluation of object detection models using acoustic blur degraded images suggests an inverse relationship between mean average precision (mAP) and acoustic interference intensity. For example, there was a 41.1% mean decrease in mAP [50-95] for YOLOv8 across two datasets. Experimental results and analysis of adversary capabilities enable identifying potential threat models. Future work to facilitate research in developing and assessing laser-induced acoustic interference against computer vision-assisted vehicles, and other sensor-based autonomous control systems

Symposium on Vehicles Security and Privacy (VehicleSec) 2024
26 February 2024, San Diego, CA, USA
ISBN 979-8-9894372-7-6
https://dx.doi.org/10.14722/vehiclesec.2024.23060
www.ndss-symposium.org

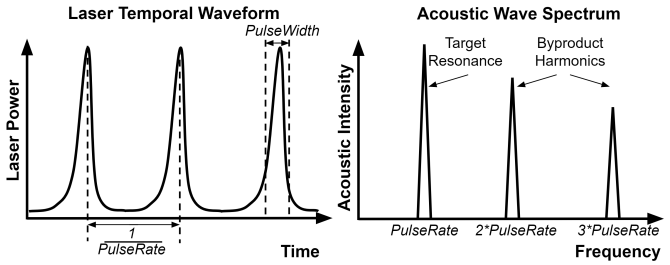

Fig. 2. A train of laser pulses with a pulse repetition rate of α Hz generates acoustic waves at α Hz and its harmonics. The adversary needs to set the pulse rate to the resonant frequency of the target sensors.

is discussed. Contributions of this work towards initial, proof-of-concept characterization of laser-induced acoustic interference are:

- Experimental analysis of acoustic interference generated through air ionization and direct ablation from high-speed, pulsed lasers to modify gyroscope sensor readings.
- A methodology to simulate and evaluate the impact of laser-induced acoustic interference on computer vision of unmanned vehicles.
- Exploration of potential laser-induced acoustic interference threat models and their limiting factors.

II. Background & Related Work

A. Laser-Induced Acoustic Waves

Lasers can generate acoustic waves via various photo-acoustic effects, which are the physical processes of optical energy being transformed into mechanical (acoustic) waves in gases, liquids, and solids. High-energy pulsed lasers can induce acoustic waves to affect MEMS sensors through both air ionization and direct material ablation.

Pulsed Lasers Different from continuous wave lasers that constantly output optical energy, pulsed lasers turn on and off with certain pulse repetition rates and pulse widths, as shown in Figure 2. A key parameter of acoustic interference in an inertial measurement unit (IMU) is the acoustic wave frequency which is resonant to a sensor's mechanical resonant frequency. For acoustic interference against MEMS IMU sensors, *the acoustic waves generated by a pulsed laser with a repetition rate of α Hz can be considered equivalent to the sinusoidal acoustic waves generated by electronic speakers at α Hz.*

Air Ionization A high-power laser beam focused onto a small area in the air can generate air-borne sound by imparting enough energy density to exceed the breakdown point of the air molecules [13]. An adversary could use air ionization to generate acoustic waves by setting the ionization breakdown point close to the surface of the target vehicle. The acoustic waves will propagate sequentially to the internal IMU sensor through airborne and structure-borne paths, and affect sensor readings, as shown in Figure 1.

Direct Ablation Laser ablation produces acoustic waves through a similar process, but with the plasma generated out of irradiated solid materials instead of air molecules. When generating acoustic interference through direct ablation, the adversary directly sets the ionization breakdown point on the surface of the target vehicle. The acoustic waves generate mechanical vibrations of the target vehicle's body and affect sensor readings mostly through structure-borne acoustics, as shown in Figure 1.

B. Acoustic Attacks on MEMS Sensors

MEMS IMUs are miniature electronic devices that can measure the movements of objects. Specifically, accelerometers and gyroscopes measure an object's acceleration and angular velocity respectively, and provide that measurement for downstream applications, including localization and navigation in semi or fully autonomous robots and vehicles. Gyroscope sensors are also widely used in modern camera systems for image stabilization (see Section II-C). Despite widespread use of MEMS IMU sensors, prior works have shown that such sensors are sensitive to acoustic interference at the sensors' mechanical resonant frequencies [12], [35], [36]. This is because the mechanical vibrations caused by sound waves are able to change the readings of the IMU sensors in a contactless way.

Conversely, Long et al. [18] showed that ambient acoustic waves near cameras can cause audio-specific image blurs, which allow adversaries to use computer vision techniques to eavesdrop on ambient audio by analyzing streams of camera images. The computer vision capabilities of a drone are integral for enabling many of its core functionalities, even for drones which utilize multi-modal sensing [28], but are not required for it to stay in flight, and as such are a reasonable target for damage-free drone disablement. While all previous works of acoustic interference attacks used electronic speakers to generate acoustic waves, this work explores first-of-its-kind methodology for laser-based acoustic interference attacks against MEMS gyroscopes and camera-based computer visions.

C. Image Stabilization in Drones

Drone computer vision capabilities may include object detection and tracking [2], [4], and obstacle avoidance [6], to name a few, and object detection enables many downstream computer vision capabilities [2], [15]. Drones use a combination of digital and mechanical/optical image stabilization to compensate for motion-induced image degradation [26]. Gyroscope-guided image stabilization is a well-established mechanical/optical technique which compensates for the shifts and rotations resultant from the movement of the optical system during image capture [17]. Object detection, used by either autonomous or operated drones, is a core capability for enabling drone function across different applications, including infrastructure inspection, crowd control, reconnaissance and surveillance, or counter-drone operations [2]. Many object detection techniques for drones rely on IMU sensor data, e.g., fast object detection [4], pose recognition [29], or gyro-informed tracking [15]. Object detection models rely on image quality [33], and are sensitive to sources of degradation

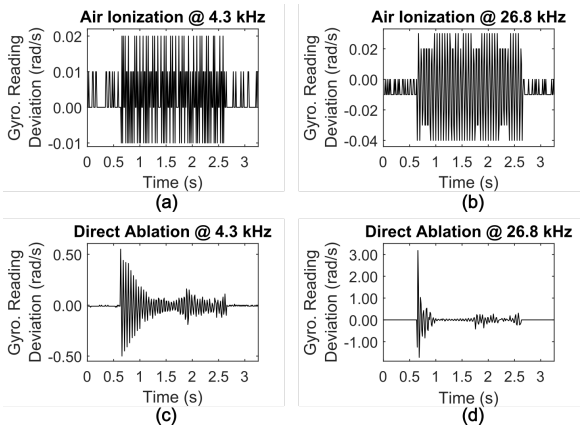

Fig. 3. Example of gyroscope reading deviations caused by laser air ionization and direct ablations at two acoustic resonant frequencies of the gyroscope.

affecting feature representations [16], [37]. The work herein focuses on inducing gyroscope measured deviation to cause camera destabilization, and determine at what acoustic intensity threshold sensor destabilization affects a drone's object detection capability.

III. Methodology

A. Pulsed Lasers for Acoustic Wave Generation

A Light Conversion Pharos DPSS femtosecond laser was used for all pulsed laser tests. This laser supports 1kHz to 200kHz pulse repetition rates, 3mJ/pulse max output energy, and 20W max output power. The laser wavelength is 1030nm. The optical system uses a focusing lens with a focal length of 7cm. By default, the highest output power of the laser was used. For attack targets, either a MPU6050 MEMS gyroscope sensor, or a DJI Mavic 2 Pro drone, was used near the laser focal point. A gyroscope test board was created to enable easier handling of the MPU6050 MEMS gyroscope during testing by attaching the sensor to a piece of plastic material commonly used to construct drones. The test board was used as the target for either air ionization and direct ablation-generated acoustic interference, and the gyroscope output was measured. When testing air ionization, the gyroscope sensor was placed near the focal point of the laser, but focal point was in the air. During air ionization, the generated acoustic waves propagate to the sensor through air. When testing direct ablation on the gyroscope test board, the laser focal point was on the test board plastic. In this case, the acoustic waves affect the sensor through direct mechanical vibrations. Different pulse rates were tested for the laser to generate acoustic interference at various frequencies.

B. Simulating Laser Acoustics Impact on Drone CV

As mentioned, object detection enables downstream computer vision capabilities, which are often required for enabling key functionality in a drone. Object detection can be enabled in current commercial drones [10], [21], but to decouple the effect of laser-induced acoustic interference on the drone's gyroscope from its effect on other components of the drone's hardware this work opted to capture acoustic-interfered data from the drone camera and evaluate the object detection models on a separate computer. However, object detection datasets with images captured under acoustic interference do not currently exist, so acoustic wave-induced motion blur was instead simulated in existing object detection datasets by extracting blur kernels from drone images captured under acoustic interference.

High Intensity Laser Simulation Electronic speakers were used to emulate higher intensity lasers for greater acoustic wave generation (i.e., >98dB). Acoustic intensity produced by air ionization has a theoretical limit of about 600Pa, equivalent to about 150dB SPL, which is significantly higher than that required by previous speaker-generated acoustic attacks [24]. Experimental results have also verified that lab-scale lasers can use air ionization to generate sound over 120dB SPL that can vibrate nearby objects in non-contact manners [11]. As such, we can gain insight into the potential effects of using higher powered lasers as emulated by electronic speakers. A drone was subjected to increasing acoustic intensity from 98dB to 110dB, with a step size of 3. A previously described setup was used for acoustic wave generation using an electronic speaker [12]. Videos recorded using the drone camera during acoustic interference were used to extract blur kernels. The most blur degraded frame from each video capture was identified using the BRISQUE metric [22] to simulate a blur kernel for each acoustic intensity level.

Acoustic Blur Kernel For this simulation method, the goal was to extract features from observed motion in a blurry image to inform a blur kernel, and use that blur kernel to degrade clean object detection dataset images in a localized manner. Feature descriptors in an image, such as localized intensity gradients, can be used to describe an image object's shape, and histogram of oriented gradients (HOG) is a commonly used technique for determining gradient orientations local to an object [7]. Motion blur can result in smoothing, effectively leading to a loss of gradients [3], and reduction in pixel intensities along edges [20]. To simulate acoustic blurred images, 3 second videos using the drone's camera were recorded while subjecting the drone to a direct laser ablation or electronic speaker generated acoustic waves, where acoustic wave generation was confirmed by using a previously determined resonant frequency. The drone remained in a stationary and stable position during laser ablation to preclude any other reasons for image motion blur. The most blurred frame from each video was identified using the BRISQUE metric [22], and used to determine blurred image intensity gradients using HOG. Finally, non-blurry and clean images were convolved using a 10×10 kernel on pixels mapped to corresponding blur image gradient locations with low-level intensities.

Object Detection Datasets and Models Images degraded by acoustic blur kernels were used to evaluate performance for different object detection models. Only the validation set images were degraded to emulate an object detection model

trained for normal applications. Object detection baseline model performance was established using Detectron2's RCNN X101-FPN and Faster RCNN X101-FPN pretrained model checkpoints on COCO2017 [38]. Datasets with objects for common drone applications, e.g. infrastructure surveillance or detecting other drones, were also used for evaluating the effect of acoustic wave generation on an object detection model [9], [27] using pretrained model checkpoints available for each dataset [14]. The metrics used for object detection include mean average recall (mAR), and mean average precision at different intersection over union (IoU) thresholds, 0.5, 0.6, and the interval of [.5,0.5-0.9] [25]. The metrics AP_M and AP_L are average precision for ground truth objects with different pixel areas [25].

IV. Experimental Results

TABLE I
YOLOv8 Object Detection Evaluation

Dataset	Experiment	mAR	mAP50	mAP60	mAP 50-95
ParkingLot [9]	Baseline	0.993	0.994	0.993	0.918
	Laser	0.834	0.895	0.910	0.562
	98 dB	0.712	0.831	0.964	0.605
	101 dB	0.706	0.823	0.895	0.594
	104 dB	0.694	0.806	0.889	0.576
	107 dB	**0.673**	0.794	0.903	0.563
	110 dB	0.871	**0.779**	**0.874**	**0.535**
DroneSegment [27]	Baseline	0.999	0.995	0.999	0.999
	Laser	**0.834**	0.895	0.910	0.562
	98 dB	0.843	0.912	0.898	0.572
	101 dB	0.858	0.921	0.916	0.567
	104 dB	0.883	0.908	**0.891**	0.572
	107 dB	0.898	0.826	0.902	0.563
	110 dB	0.927	**0.813**	0.900	**0.555**

A. Laser-Induced Acoustic Wave Generation

Acoustic wave generation from high-speed, pulsed laser was experimentally verified to modify gyroscope readings through both air ionization and direct ablation. For example, Figure 3 shows gyroscope reading deviations when the laser produces acoustic waves at two acoustic resonant frequencies (audible at 4.3kHz, and inaudible at 26.8kHz) of the gyroscope sensors compared to when there is no laser-generated acoustic interference. In each case, the laser is turned on for 2 seconds.

Out-of-focus Ablation When the target object is in the optical path of the laser, ablation can happen not only when the target is right at the focal point of the laser beam but also when the target deviates from the focal point along the path, i.e., being out-of-focus. Out-of-focus ablation can happen if the adversary fails to aim the laser at the target, or purposely adjusts the focal point away from the focal point. To investigate the impact of out-of-focus ablation, the gyroscope test board was moved away from the focal point, which is 7cm away from the focusing lens. Figure 6 shows how changing the test board's position relative to the focus point changes the sensor readings, where a location deviation of 0cm represents being at the focal point. When the ablation target is 21cm away from the focal point, the maximum sensor reading deviation reduces from about 3rad/s (Figure 3 (d)) to about 0.08rad/s (Figure 6 (d)). This is about 2.7% of the original deviation amplitude, but is still higher than the 0.01 rad/s noise floor.

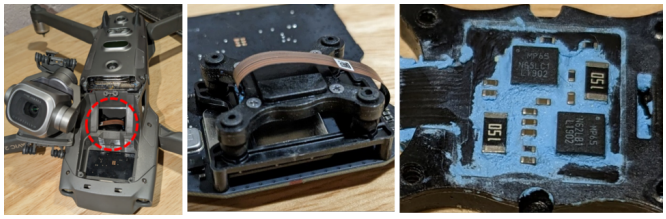

Fig. 4. A tear-down view of the DJI Mavic 2 Pro drone's gyroscope sensor.

While interference strength, as measured by gyroscope reading deviation from a 0.01 rad/s noise floor, is lower when the target is out-of-focus, the stability of interference improves. In contrast to in-focus ablation, which may quickly burn a hole through the target material, out-of-focus ablation only causes slow melting of the ablated material, which is less destructive due to lower energy-per-area (Section II-A). Figure 6 shows a comparison between the burn marks generated by in-focus and out-of-focus ablation, and Figure 6 (d) shows that the acoustic interference could last for over 5 seconds with out-of-focus ablation. The results suggest that *advanced adversaries could deliberately create out-of-focus conditions to achieve more stable and persistent controls over the induced acoustic interference at the cost of lower interference strength.*

B. Impact on Object Detection

After verifying the feasibility of using lasers to generate acoustic interference to modify gyroscope sensor measurements, the effect of acoustic interference on drone camera hardware was investigated.

TABLE II
COCO2017 Object Detection Evaluation

Model	Experiment	mAP50-95	mAP50	AP_M	AP_L
RCNN	Baseline	65.076	86.222	60.236	73.577
	Laser	**55.897**	**77.658**	**49.687**	**65.676**
	98 dB	57.650	79.581	51.542	67.604
	101 dB	57.633	79.578	51.503	67.586
	104 dB	57.631	79.549	51.462	67.556
	107 dB	57.623	79.572	51.388	67.554
	110 dB	57.624	79.597	51.452	67.566
Faster RCNN	Baseline	37.455	52.935	40.59	50.573
	Laser	**30.375**	**44.036**	**33.248**	**41.039**
	98 dB	31.038	44.914	34.023	42.291
	101 dB	30.988	44.877	33.967	42.219
	104 dB	30.988	44.828	33.941	42.223
	107 dB	30.994	44.824	34.000	42.267
	110 dB	30.998	44.853	34.002	42.240

Interfered Video Collection. The laser setup used for this research only enables in-focus direct ablation for generating acoustic interference which can produce motion blurs in a drone camera video. Figure 5 (b) and (c) demonstrate an example of the video blurs when the laser generates acoustic waves at the gyroscope's resonant frequency via direct ablation. The insufficient interference strength is the main reason that air ionization and out-of-focus ablation could not generate observable changes in the camera videos. It should be noted

that the drone's gyroscope is installed in the center of the drone body and covered by thick plastic and metal materials (Figure 4), making it challenging for adversaries to aim the laser's focal point closely at the drone's gyroscope, as well as for air-borne acoustic waves to reach the sensor. Potential threat model modifications which could enable air ionization and out-of-focus ablation to affect drone cameras are discussed in Section VI.

Object Detection Evaluation Direct ablation of the drone gyroscope with the laser causes the drone's camera to vibrate, which induces motion blurs in the drone camera stream, as shown in Figure 5 (c). The most likely cause of the camera vibration is mechanical vibration of the optical system due to destabilization caused in the gyroscope via the resonant frequency of the laser. The DJI Mavic 2 Pro drone used in this research uses a triaxial gimbal [1], and a system of sensors dubbed *FlightAutonomoy* [6] for camera stabilization. Thus, it is likely that mechanical vibrations caused by the laser-induced acoustic interference against the drone's gyroscope overpowered additional onboard mechanisms for camera stabilization, but this requires more experimentation to confirm. The BRISQUE [22] score of the images produced using the blur kernels matched that of the most blurry frame identified from the drone videos, but a more precise method to confirm accurate blur kernel extraction is required. Current methods utilize neural nets for both point spread function estimation and image deblurring [5], [39], and a deblurring neural net may potentially be used for extracting a more precise blur kernel. The object detection pretrained model checkpoints were trained on relatively low noise images without motion blurs, and using blurred images for object detection dropped the performance of the object detection model as expected, as seen in Table I. However, the blur kernels likely contain information that is unique to the mechanical vibrations experienced by the drone's optical system, which is why our experiments utilize data captured from the drone for object detection model evaluation.

The lowest values for Table I and II are highlighted. The YOLOv8 [34] model used in this research has 3.2M parameters, while the the Faster RCNN and RCNN [38] have greater than 60M parameters, which is why YOLOv8 experiences relatively more degradation as acoustic intensity is increased Table II. Furthermore, it's likely that the pretrained checkpoints for Faster RCNN and RCNN cannot degrade further due to inherent learned feature representations in the model weights even with increased acoustic intensity. In Table I, the lowest mean average precision values are found for higher acoustic intensities, implying poorer identification of relevant objects. But high mean average recall (mAR) for higher intensities implies that models were still able to find relevant bounding boxes despite the motion blur produced at these intensities. We find that the 98dB acoustic intensity of the laser (non-simulated) was sufficient to produce a 38.8% and 43.7% percent decrease in mAP [50-95] across two datasets using YOLOv8. During direct ablation, the laser focal point was on the gyroscope test board, which can result in destruction of the drone body, making it unlikely for the destroyed plastic to propagate generated acoustic waves. As such, we believe that the image blurs produced during direct ablation are most likely by vibrations from laser-induced sounds which do not propagate through the plastic. However, we aim to confirm this via stable interference from less damaging methods.

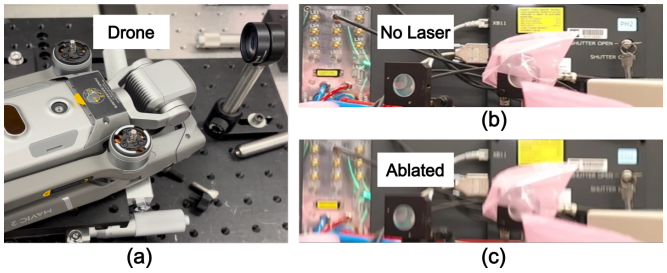

Fig. 5. Ablation test setup with a DJI Mavic 2 Pro drone, and the drone camera images with and without laser-induced acoustic interference. The laser ablation causes motion blurs in the camera output.

V. Discussion

A. Air Ionization and Direct Ablation

Interference Strength At the same acoustic frequency and laser-target distance, direct ablation causes larger changes to the sensor readings than air ionization, often by over one order of magnitude. Direct ablation possibly has a more efficient coupling path, or is amplified by the inherent particle compositions of plastics and air molecules. For air ionization, closer distances between the focal point and the sensor almost always lead to larger deviations in sensor readings, corresponding with prior research on acoustic attacks using electronic speakers (Section II-B). However sensor reading deviations during direct ablations are more unpredictable, possibly due to the non-monotonic acoustic energy responses in complex mechanical structures. Such an effect was also observed in previous works investigating sensor-related attacks which exploit structure-borne acoustic wave propagation [18]. In both air ionization and direct ablation, the strength of interference decreases proportionally with the laser output power.

Interference Stability and Destruction Weaker sensor reading deviations due to air ionization are often more stable than those caused by direct ablation. For example, Figure 3 (bottom) shows that the gyroscope deviations under direct ablation start with a large value, and then quickly reduce to a small value. This is because there is a limited amount of plastic material that can be ablated. While air ionization consumes air molecules that can be replenished by ambient airflow, laser ablation consumes the solid plastic material directly, causing permanent destruction of burn marks (see Figure 6 for example). Once laser ablation has consumed all the plastic material around the laser's focal point area, there will be a hole in the plastic, and further laser energy will start consuming the air in the hole, essentially transforming the process into air ionization. As a result, *the strength of interference caused by direct ablation gradually decreases to a constant lower*

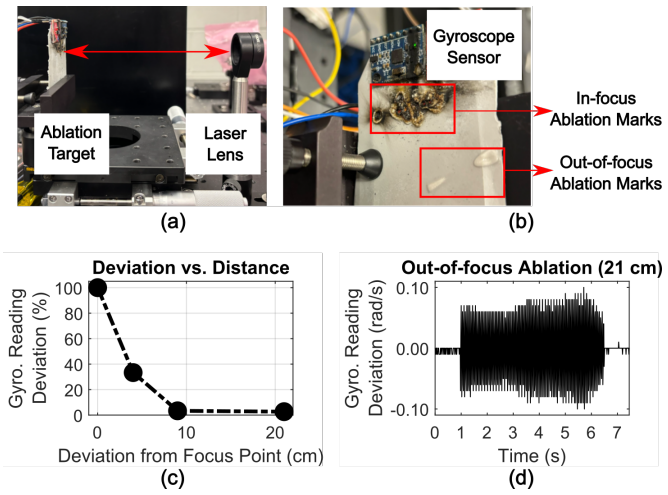

Fig. 6. Gyroscope reading deviations caused by in-focus and out-of-focus ablations. Out-of-focus ablation causes smaller degrees of deviation, but produces more stable interference with less damaging effects.

bound value that is decided by the interference generated by air ionization.

B. Potential Threat Models

Attack Surface and Goal Three levels of target complexity should be considered for identifying attack goals of laser-induced acoustic interference: 1) sensor type, 2) vehicle subassemblies, and 3) level of vehicle autonomy (Figure 7). The adversary may interfere with high-level functionalities by changing the output of sensors, e.g., IMUs, microphones, pressure sensors etc., that are sensitive to acoustic interference. This work provides an illustrative example of interference against CV-assisted aerial vehicles by affecting the gyroscope sensors used by drones and cause image blurs in drone captured images. Theoretically, adversaries can also affect the autonomous navigation control systems that directly depend on IMU sensors, as has been demonstrated in previous works (Section II-B). By modulating the output power of laser pulses, adversaries should be able to generate diverse acoustic waveforms that can inject more fine-grained false information into downstream software systems.

Adversary Characteristics Different attack goals require specific requirements for designing and manufacturing laser systems. Key factors affecting a laser system used for acoustic interference include the laser's maximum attack distance, portability, power level, acoustic generation mechanisms, etc. Lightweight, high-energy lasers which an individual can carry and use to produce sound have already been demonstrated, suggesting the possibility of more diverse and low-profile application scenarios of this attack [8], [19]. The required output power of lasers depends on the desired attack scenarios. Short-distance attacks can use lower-power lasers because of less energy dissipation of the laser in air. Additionally, direct ablation allows adversaries to generate the same intensity of sensor reading variations with lower laser power, but at the cost of gradually damaging the ablated target. In contrast, using

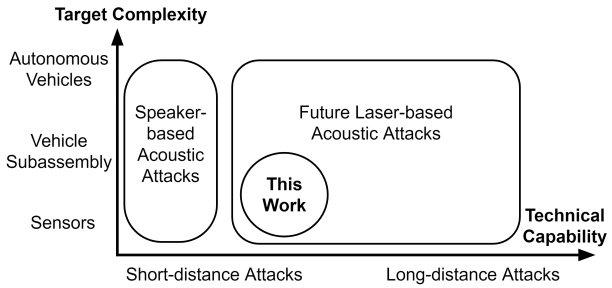

Fig. 7. This work explores the threat model of using lasers to perform acoustic attacks against sensors and the subassembly of computer vision-assisted autonomous vehicles. In contrast to speaker-based attacks, laser-based attacks could allow adversaries to achieve long-distance interference.

air ionization keeps the target device intact but requires the adversary to use a higher-power laser to compensate for the loss of acoustic energy when the acoustic waves propagate from air to the target. The adversary thus needs to carefully align the laser system design with the desired attack effects.

VI. Future Work

Laser Signal Modulation The laser setup used for this research produces an acoustic signal limited to a near-sinusoidal waveform at a single frequency. More complex laser-based acoustic attacks, for precise control over target systems, such as steering an unmanned vehicle [36] or making CV algorithms misclassify detected objects to certain classes [12], require adversaries to produce acoustic waves with varying designated amplitudes and frequencies. The use of additional optical modulation devices to control the intensity of individual pulses of a laser pulse train may achieve acoustic waves with varying designated amplitudes and frequencies. The feasibility of this is demonstrated by using lasers to produce low-intensity music and speech audio [13], [32].

Diverse Targets and Scenarios While this work provides a preliminary evaluation of how laser-induced acoustic interference can affect the output of individual sensor components in a controlled lab setup, future work may include diverse targets and scenarios, e.g., vehicle subassemblies. Adversaries wishing to use direct ablation to generate acoustic interference face trade-offs between damage from ablation to target devices vs. laser-induced interference strength and duration. Such trade-offs are dominated by practical factors such as the ablated material (e.g., plastic vs. metal), the mechanical construct of target vehicles, the wavelengths and pulse widths of lasers, etc. Adversaries may develop simulation models similar to ones currently used for destructive high-energy lasers [30], [40]. Overall, sophisticated software as well as physical testbeds need to be developed to support future evaluations.

VII. Acknowledgments

We appreciate the valuable suggestions and remarks from our reviewers. This research was supported by grant 1916762 from the NSF Center for Hardware and Embedded Systems Security and Trust (CHEST).

REFERENCES

[1] *Mavic 2 Pro/Zoom User Manual*, https://dl.djicdn.com/downloads/Mavic_2/Mavic_2_Pro_Zoom_User_Manual_v2.2_en.pdf, 2020, accessed: 2023-12-18.

[2] M. W. Ashraf, W. Sultani, and M. Shah, "Dogfight: Detecting drones from drones videos," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 7067–7076.

[3] P. Balasubramanian, S. Pathak, and A. Mittal, "Improving gradient histogram based descriptors for pedestrian detection in datasets with large variations," in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2016, pp. 104–113.

[4] W. Budiharto, A. A. Gunawan, J. S. Suroso, A. Chowanda, A. Patrik, and G. Utama, "Fast object detection for quadcopter drone using deep learning," in *2018 3rd international conference on computer and communication systems (ICCCS)*. IEEE, 2018, pp. 192–195.

[5] L. Chen, X. Chu, X. Zhang, and J. Sun, "Simple baselines for image restoration," *arXiv preprint arXiv:2204.04676*, 2022.

[6] F. Corrigan, "Mavic pro - highlights you need to know," https://store.dji.com/guides/mavic-pro-highlights-need-know/, 2017, accessed: 2023-12-18.

[7] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05)*, vol. 1. Ieee, 2005, pp. 886–893.

[8] Dark Footage. (2023) The Non-Lethal Weapon that Uses Phantasmagoric Sound Effects to Terrify the Enemy. https://www.youtube.com/watch?v=4cI235yAmPA. [Online; accessed 09-Dec-2023].

[9] P. R. de Almeida, L. S. Oliveira, A. S. Britto, E. J. Silva, and A. L. Koerich, "Pklot – a robust dataset for parking lot classification," *Expert Systems with Applications*, vol. 42, no. 11, p. 4937–4949, Jul. 2015. [Online]. Available: http://dx.doi.org/10.1016/j.eswa.2015.02.009

[10] DJI, "Advanced sensing - object detection sample," https://developer.dji.com/onboard-sdk/documentation/sample-doc/advanced-sensing-object-detection.html, 2017, accessed: 2023-12-18.

[11] N. Hosoya, M. Nagata, I. Kajiwara, and R. Umino, "Nano-second laser-induced plasma shock wave in air for non-contact vibration tests," *Experimental Mechanics*, vol. 56, pp. 1305–1311, 2016.

[12] X. Ji, Y. Cheng, Y. Zhang, K. Wang, C. Yan, W. Xu, and K. Fu, "Poltergeist: Acoustic adversarial machine learning against cameras and computer vision," in *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2021, pp. 160–175.

[13] K. Kaleris, B. Stelzner, P. Hatziantoniou, D. Trimis, and J. Mourjopoulos, "Laser-sound: optoacoustic transduction from digital audio streams," *Scientific reports*, vol. 11, no. 1, p. 476, 2021.

[14] A. Kanametov, "yolov8-face," https://github.com/akanametov/yolov8-face/tree/dev, 2022.

[15] K. Lee and C. Kim, "A gyro-based tracking assistant for drones with uncooled infrared camera," *IEEE Sensors Journal*, pp. 1–1, 2023.

[16] H. Liu, F. Jin, H. Zeng, H. Pu, and B. Fan, "Image enhancement guided object detection in visually degraded scenes," *IEEE Transactions on Neural Networks and Learning Systems*, p. 1–14, 2023. [Online]. Available: http://dx.doi.org/10.1109/TNNLS.2023.3274926

[17] S. Liu, H. Li, Z. Wang, J. Wang, S. Zhu, and B. Zeng, "Deepois: Gyroscope-guided deep optical image stabilizer compensation," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 32, no. 5, pp. 2856–2867, 2021.

[18] Y. Long, P. Naghavi, B. Kojusner, K. Butler, S. Rampazzi, and K. Fu, "Side eye: Characterizing the limits of pov acoustic eavesdropping from smartphone cameras with rolling shutters and movable lenses," in *2023 IEEE Symposium on Security and Privacy (SP)*. Los Alamitos, CA, USA: IEEE Computer Society, may 2023, pp. 1857–1874. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/SP46215.2023.10179313

[19] Mary-Ann Russon. (2023) US military test a non-lethal gun that can terrify an enemy with sound. https://www.ibtimes.co.uk/us-military-test-non-lethal-gun-that-can-scare-enemy-sound-1513539. [Online; accessed 09-Dec-2023].

[20] G. Mather, "The use of image blur as a depth cue," *Perception*, vol. 26, no. 9, pp. 1147–1158, 1997.

[21] MathWorks, "Fly a parrot minidrone and detect objects - matlab & simulink example," https://in.mathworks.com/help/supportpkg/parrot/ref/color-detection-and-landing-parrot-example.html, 2017, accessed: 2023-12-18.

[22] A. Mittal, A. K. Moorthy, and A. C. Bovik, "No-reference image quality assessment in the spatial domain," *IEEE Transactions on image processing*, vol. 21, no. 12, pp. 4695–4708, 2012.

[23] B. Nassi, R. Bitton, R. Masuoka, A. Shabtai, and Y. Elovici, "Sok: Security and privacy in the age of commercial drones," in *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2021, pp. 1434–1451.

[24] M. Oksanen and J. Hietanen, "Photoacoustic breakdown sound source in air," *Ultrasonics*, vol. 32, no. 5, pp. 327–331, 1994.

[25] R. Padilla, W. L. Passos, T. L. B. Dias, S. L. Netto, and E. A. B. da Silva, "A comparative analysis of object detection metrics with a companion open-source toolkit," *Electronics*, vol. 10, no. 3, p. 279, Jan. 2021. [Online]. Available: http://dx.doi.org/10.3390/electronics10030279

[26] S. Pant, P. Nooralishahi, N. P. Avdelidis, C. Ibarra-Castanedo, M. Genest, S. Deane, J. J. Valdes, A. Zolotas, and X. P. V. Maldague, "Evaluation and selection of video stabilization techniques for uav-based active infrared thermography application," *Sensors*, vol. 21, no. 5, p. 1604, Feb. 2021. [Online]. Available: http://dx.doi.org/10.3390/s21051604

[27] Projects, "Dronesegment dataset," https://universe.roboflow.com/projects-s5hzp/dronesegment, jan 2023, visited on 2023-12-24. [Online]. Available: https://universe.roboflow.com/projects-s5hzp/dronesegment

[28] G. Rizzoli, F. Barbato, M. Caligiuri, and P. Zanuttigh, "Syndrone-multi-modal uav dataset for urban scenarios," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 2210–2220.

[29] P. Shan, R. Yang, H. Xiao, L. Zhang, Y. Liu, Q. Fu, and Y. Zhao, "Uavpnet: A balanced and enhanced uav object detection and pose recognition network," *Measurement*, vol. 222, p. 113654, Nov. 2023. [Online]. Available: http://dx.doi.org/10.1016/j.measurement.2023.113654

[30] J. Stupl and G. Neuneck, "Assessment of long range laser weapon engagements: The case of the airborne laser," *Science & Global Security*, vol. 18, no. 1, pp. 1–60, 2010.

[31] T. Sugawara, B. Cyr, S. Rampazzi, D. Genkin, and K. Fu, "Light commands:{Laser-Based} audio injection attacks on {Voice-Controllable} systems," in *29th USENIX Security Symposium (USENIX Security 20)*, 2020, pp. 2631–2648.

[32] R. M. Sullenberger, S. Kaushik, and C. M. Wynn, "Photoacoustic communications: delivering audible signals via absorption of light by atmospheric h2o," *Optics Letters*, vol. 44, no. 3, pp. 622–625, 2019.

[33] S. Sun, W. Ren, T. Wang, and X. Cao, "Rethinking image restoration for object detection," in *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, Eds., vol. 35. Curran Associates, Inc., 2022, pp. 4461–4474. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2022/file/1cac8326ce3fbe79171db9754211530c-Paper-Conference.pdf

[34] J. Terven, D.-M. Córdova-Esparza, and J.-A. Romero-González, "A comprehensive review of yolo architectures in computer vision: From yolov1 to yolov8 and yolo-nas," *Machine Learning and Knowledge Extraction*, vol. 5, no. 4, pp. 1680–1716, 2023.

[35] T. Trippel, O. Weisse, W. Xu, P. Honeyman, and K. Fu, "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, 2017, pp. 3–18.

[36] Y. Tu, S. Rampazzi, and X. Hei, "Towards adversarial process control on inertial sensor systems with physical feedback side channels," in *Proceedings of the 5th Workshop on CPS&IoT Security and Privacy*, 2023, pp. 39–51.

[37] Y. Wang, X. Yan, K. Zhang, L. Gong, H. Xie, F. L. Wang, and M. Wei, "Togethernet: Bridging image restoration and object detection together via dynamic enhancement learning," in *Computer Graphics Forum*, vol. 41, no. 7. Wiley Online Library, 2022, pp. 465–476.

[38] Y. Wu, A. Kirillov, F. Massa, D. W.-Y. Lo, and R. Girshick, "Detectron2," https://github.com/facebookresearch/detectron2, 2019.

[39] X. Yang, L. Huang, Y. Luo, Y. Wu, H. Wang, Y. Rivenson, and A. Ozcan, "Deep-learning-based virtual refocusing of images using an engineered point-spread function," *ACS Photonics*, vol. 8, no. 7, pp. 2174–2182, 2021.

[40] B. Zohuri and B. Zohuri, *Directed energy weapons*. Springer, 2016.