Demystifying Event-based Vision Sensor Biasing to Optimize Signal to Noise for Space Domain Awareness

Brian McReynolds, Rui Graca
Institute of Neuroinformatics, UZH/ETH Zürich
Rachel Oliver
Cornell University
Masashi Nishiguchi
Purdue University
Tobi Delbruck
Institute of Neuroinformatics, UZH/ETH Zürich

ABSTRACT

Event-based Vision Sensors (EVS) have recently garnered interest in the space sensing community due to performance advantages such as low-latency, wide-dynamic range, and minimal data requirements resulting from a sparse representation of dynamic visual information. Several studies to date have demonstrated their utility for SDA tasks, and two have even reported empirical measurements of limiting performance in terms of absolute sensitivity. In both studies, the EVS failed to achieve the same sensitivity compared to optimized frame-based, scientific CMOS cameras operating in the same environmental conditions and both reported a drop in sensitivity with increased target speeds, measured by scanning the sky at different rates. Notably, neither study employed a thorough exploration or description of the vast array of user defined sensor biases available in EVS. This paper applies recent advances in understanding EVS bias optimization and noise performance to explore the many degrees of freedom offered by tunable EVS biases. With a targeted exploration of the available parameter space, we attempt to push the performance limits of EVS in the challenging task of sub-pixel, dim target detection and identify general biasing techniques and principles that can be applied to any COTS EVS. A new simulation tool allows accurate simulation EVS responses to dim fast moving point sources. Using a DAVIS346 EVS and custom lab setup calibrated to resemble point source objects of varied brightness and speed, we present sensitivity results from eleven different hand-selected bias configurations. As a result, we approach nearly optimized EVS bias settings for SDA tasks, demonstrating as much as **1.6 m_V increase in sensitivity** ($\approx 4.3\times$ dimmer), and the ability to **detect objects moving up to 6.6 $\times$ faster** compared to default or naive bias configurations. Our results suggest that prior reports of EVS limiting magnitude and sensitivity while scanning could be significantly improved by employing more optimized bias configurations.

1. INTRODUCTION

Neuromorphic or Event-based Vision Sensors (EVS) are a relatively new class of camera gaining wide popularity in machine vision tasks. Inspired by the design of biological retinas, they achieve low latency, wide Dynamic Range (DR), high temporal resolution, minimal data requirements, and ultra-low power consumption [8]. They attain these performance benefits by a pixel circuit that independently reports changes in each pixel's log photocurrent. These detected changes are called 'events'. EVS performance advantages appear ideal for space-sensing tasks, and they have shown promise for astronomical applications including Space Domain Awareness (SDA) [4, 5, 2]. Despite many benefits of the technology and several studies demonstrating utility for space applications, initial estimates indicate available Commercial Off The Shelf (COTS) EVS are at a disadvantage to optimized, scientific CMOS Image Sensor

(CIS) by ≈ 1.5 of visual magnitudes (m_V) [16, 24]. Both studies also demonstrated a drop in sensitivity with increased telescope scan rates, which is consistent with reported pixel operation [14]. Optimistically, the temporal resolution of the EVS was potentially shown to be an advantage in the case of capturing a blinking Geosynchronous Orbit (GEO) object, but no thorough analysis was conducted on the performance limits [16].

While most of the ongoing research surrounding EVS applications to SDA involves exploring potential use cases and algorithm development for extracting useful information from the output, none thus far have conducted a detailed study of the EVS pixel architecture in order to optimize performance for SDA. The task of optimizing EVS sensitivity for SDA is complicated by two primary obstacles: elevated noise event rates at low illumination levels, and a complex multidimensional parameter space of user adjustable biases. EVS pixels sometimes record events in the absence of signal change due to thermal and shot noise [10], and in dim lighting conditions such as the night sky, the rate at which these noise events occur can increase dramatically. Distinguishing noise from signal is exceptionally challenging for SDA, as sparse, small objects like stars and satellites result in extremely low Signal to Noise Ratios (SNR). Noise rates can be controlled to some extent by tuning sensor settings or 'biases'. Most EVSs allow the user to manually adjust roughly six biases (each over a range typically spanning several orders of magnitude) that determine sensor performance such as sensitivity and response speed. The goal of this paper is to push the performance limits of the DAVIS346 EVS pixel architecture by exploring the biasing parameter space in order to find a more optimized configuration for SDA that balances sensitivity, speed and noise performance.

In Sec. 2, we present the basic operation of the EVS pixel, describe how the six main biases available in a typical EVS influence sensor performance, and discuss the influence of those biases on noise rates when the background is dark. Sec. 3 describes the methodology employed to test various bias configurations against a known repeatable stimulus. Sec. 3.1 details our custom lab set-up used to simulate calibrated point source targets of variable brightness and speeds, and describes the data collection process. Sec. 3.2 provides insight into selection criteria considered for eleven different bias configurations analyzed in this study, and Sec. 3.4 outlines the methods used to calibrate the point source targets to astronomical units (m_V) and Sec. 4.1 explains how we mapped target position ground truth to distinguish signal events from noise. Sec. 4 presents side-by-side comparison of the eleven bias combinations using both sensitivity and scan speed as quantitative detection metrics. To conclude, we prescribe which set of DAVIS346 biases analyzed is most optimal for detecting point source targets against dark backgrounds and, more importantly, identify general principles that can be applied to biasing any EVS model to optimize performance for SDA tasks.

2. EVS DESCRIPTION

The first practical EVS was reported in 2008 [14], and several prototypes have since been developed and released. Some have only been research grade prototypes [23, 20, 13], while others have been mass-produced and marketed as commercial products [3, 22, 7]. At the heart of all EVS cameras is an analog pixel design consisting of three stages: a logarithmic photoreceptor, a change amplifier, and a comparator stage. A simplified version of the basic circuit design is shown in Fig. 1.

A photodiode generates photocurrent (I_{ph}) in response to the pixel's incident light level, which is converted to an output voltage (V_p) proportional to the natural log of the photocurrent. This signal is buffered by a source-follower and amplified by a factor of $\frac{-C_1}{C_2}$, resulting in voltage signal (V_{diff}). In the comparator stage, V_{diff} is continuously monitored and compared to a reference level stored in analog memory which corresponds to the signal level of the last reported event. If the log intensity increases or decreases by a tunable global threshold value ($\theta_{ON/OFF}$), an event is recorded and read off chip by peripheral circuits. Increasing light levels produce 'ON' events, and decreasing light levels produce 'OFF' events. After reporting an event, the pixel ignores signal changes for a finite period of time known as 'refractory period', after which its memorized reference level is reset and the pixel again monitors the input signal to report subsequent changes relative to this new reference. The EVS output stream is a list of these individual events with each row containing four entries (x, y, t, p), where x and y indicate the 2D position of the pixel in the array, t is a (typically) μs resolution timestamp, and p is the polarity (ON or OFF) of signal change. This output format is called the Address Event Representation (AER) .

2.1 Bias Description

EVS performance can vary significantly depending on illumination level and selected bias configuration. Biases can be digitally programmed by the user in most EVS models to optimize sensor performance characteristics for a particular

task and environment [9]. The most critical biases for tuning pixel performance either directly control or influence sensitivity (thresholds), temporal response speed of the circuit, and refractory period.

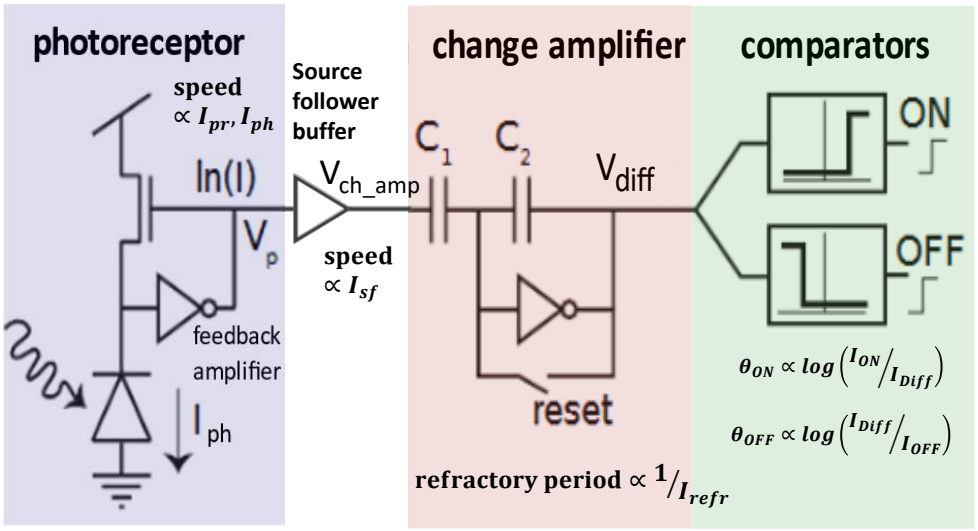

Fig. 1: Abstract EVS pixel architecture, adapted from [14].

Thresholds (θ_{ON} and θ_{OFF}) are directly controlled by three biases: I_{diff}, I_{ON}, and I_{OFF}. Nozaki and Delbruck [21] derived the theoretical ON and OFF threshold levels as a function of bias currents for the typical EVS pixel as

$$\theta_{\mathrm{ON}} = \frac{\kappa_n}{\kappa_p^2}\frac{C_2}{C_1}\log\left(\frac{I_{\mathrm{ON}}}{I_{\mathrm{diff}}}\right) \approx \frac{1}{14}\log\left(\frac{I_{\mathrm{ON}}}{I_{\mathrm{diff}}}\right) \tag{1}$$

and

$$\theta_{\mathrm{OFF}} = \frac{\kappa_n}{\kappa_p^2}\frac{C_2}{C_1}\log\left(\frac{I_{\mathrm{diff}}}{I_{\mathrm{OFF}}}\right) \approx \frac{1}{14}\log\left(\frac{I_{\mathrm{diff}}}{I_{\mathrm{OFF}}}\right). \tag{2}$$

In the DAVIS346 camera used for this study, the dimensionless term to the left of the log is approximately 1/14, but this value likely differs slightly in other camera models. To first order, when thresholds are low, this value corresponds to the fractional signal change required to generate an ON or OFF event. Based on these equations, it is tempting to assume that thresholds can be set arbitrarily low; however, in reality, mismatch in transistor thresholds and noise set a practical limit on the lowest usable threshold levels. This value is typically between 0.11 and 0.15 fractional change in good lightning conditions, but can be much higher in the dark, especially if an application requires low noise rates.

Photoreceptor response speed is also a critical performance limitation for EVS response characteristics and typically depends on one of three variables: photocurrent, I_{pr} bias, and I_{sf} bias. This response speed is often referred to as 'pixel bandwidth' in EVS literature. Under low illumination levels, pixel bandwidth depends primarily on photocurrent and increases monotonically with increasing light levels [12]. To first order, when the background photocurrent is low, increasing I_{pr} and I_{sf} should have only a minor influence on extending pixel bandwidth. With a more detailed analysis including parasitic Miller capacitance in the feedback transistor, there is a weak dependency on photoreceptor bias I_{pr}, and increasing this bias by roughly 2 orders of magnitude was shown to increase the pixel bandwidth by a factor of roughly 3 $\times$, even when the photocurrent is the dominant factor in determining response speed [17]. At higher illumination levels, either I_{pr} or I_{sf} will determine the pixel bandwidth, depending on how these two bias currents are balanced.

Graca *et. al.* recently showed that, for the DAVIS346 photoreceptor, optimal noise performance is achieved when I_{pr} is set high and I_{sf} is used to limit the pixel bandwidth to achieve the desired noise rate [11]. Pixel bandwidth is

difficult to measure, but has been estimated to be as low as several Hz for the DAVIS346 EVS at very low illumination levels [17]. This order of magnitude is consistent with estimates for newer generation EVS as well, with manufacturer specifications reporting a maximum response latency of 18 ms near the sensor's reported low-light limit of 40 mlux on chip [7]. If response latency is assumed to be $\approx \frac{1}{2\pi f_{3\text{dB}}}$, where $f_{3\text{dB}}$ represents the pixel bandwidth, this corresponds to an estimate of ≈ 1 Hz.

Refractory period is uniquely defined by I_{refr}. This bias current determines the rate at which the voltage node that controls the reset switch shown in Fig. 1 is charged after each event. When this node voltage is low, the reset switch is connected, and the output of the change amplifier is shorted directly to its input. Once the node is sufficiently charged, the reset switch disconnects so the pixel can again report changes. A higher I_{refr} bias, therefore, results in a shorter refractory period. Although this parameter is referenced throughout EVS literature, precise values are not reported in manufacturer specifications. In [17], the value was measured to range from ≈ 330 μs to 10 ms for an earlier DAVIS variant, and [15] used a similar measurement technique to report values from tens of μs to several ms for newer generation sensors. In astronomical scenarios, a longer refractory period may reduce the amount of information recorded about each detected object, but could benefit noise performance.

2.2 Noise Performance

Because EVS typically exhibit elevated noise rates under low illumination, the evolution of noise rates with varied biases is a crucial consideration for exploring their performance limits in night-time SDA. Fine tuning biases for SDA is a multivariate process, and requires consideration of sensing goals as well as the second-order effects of noise. Finding optimal biases for detecting dim, sub-pixel, point source targets against a dark background thus requires a thorough exploration of both the intended first-order effects of tuning each bias, as well as the unavoidable influence on noise rates when the background is dark. Table 1 outlines the typical first-order impacts of each parameter adjustment on both detection performance and noise.

Parameter	Bias(es)	Adjustment	Main Effect	Noise Rate Effect
Threshold(s)	$I_{\text{diff}}, I_{\text{ON}}, I_{\text{OFF}}$	Increase*	Reduced sensitivity	Decrease
		Decrease*	Improved sensitivity	Increase
Pixel Bandwidth	$I_{\text{pr}}, I_{\text{sf}}$	Increase	Faster photoreceptor response	Increase
		Decrease	Slower photoreceptor response	Decrease
Refractory Period	I_{refr}	Increase	Shorter interval between consecutive events	Increase
		Decrease	Longer interval between consecutive events	Decrease

Table 1: Effects of bias adjustments, adapted from [17]. The intended influence of tuning each controllable EVS performance parameter is listed above, as well as the unavoidable impact on noise rates. Optimization for SDA is a delicate balance of biasing for sensitivity, speed and noise management.
*ON and OFF thresholds are tuned independently and scale with $\log\left(\frac{I_{\text{ON}}}{I_{\text{diff}}}\right)$ and $\log\left(\frac{I_{\text{diff}}}{I_{\text{OFF}}}\right)$ respectively.

The general principles and relationships outlined in Table 1 are useful for a coarse bias adjustment, but true optimization requires a deeper understanding of the pixel operation and architecture, including non-ideal behaviors. Two recent studies hint at more optimal biasing techniques for noise management. McReynolds *et. al.* [18] explains the phenomenon of pixel-level noise events tending to occur in pairs of opposite polarity and demonstrates two novel techniques for reducing noise rates: 1) increasing the refractory period, and 2) intentionally applying imbalanced ON and OFF thresholds. The first decouples noise pairs by allowing the noisy background signal to settle nearer to its mean value before the pixel resets its new reference level. The second reduces the probability that an opposite polarity noise event occurs after an event corresponding to the lower of the two thresholds, by making the threshold much larger. This technique was shown to reduce noise rates by as much as 80% when the ON threshold is $\approx 3\times$ more sensitive than the OFF.

Graca *et. al.* [11] examine node voltage recordings at various pixel stages in response to noisy input signals to prescribe

more optimal biasing of the photoreceptor circuit. The authors show that setting a high I_{pr} bias shifts uninformative circuit noise components added by the photoreceptor transistor into higher frequency ranges. For a given light level, total integrated noise power remains constant regardless of I_{pr}. The noise power contained in lower frequency ranges is therefore reduced in when I_{pr} is set very high. The source follower buffer defines an independent low-pass corner frequency, and by setting I_{sf} to a low value, the noise power entering the pixel's change amplifier can be reduced significantly, approaching a theoretical limit of $\approx 2\times$ the photon shot noise. The techniques demonstrated in each of these studies show promise for pushing the limits of EVS in dim-lighting conditions. While both reports demonstrate improved noise rates, neither includes measurements to quantify performance improvements in terms of sensitivity or response speed.

2.3 Simulator for Biasing Influence on SDA Task Performance

Biasing to manage noise is an important goal, but only useful for SDA tasks if it also improves the sensor's capability to detect dimmer or faster moving objects. Using advanced biasing techniques to manage noise rates should allow the sensor to be optimized for more sensitive thresholds or faster bandwidth settings. To build intuition by visualizing the theoretical performance in SDA tasks as a function of these biasing parameters, we developed a MATLAB GUI tool for analyzing the expected pixel response to a point source object against a background illumination at or below the dark current limit of the photodiode[1]. We incorporated estimates of dark current, quantum efficiency, minimum allowable thresholds, pixel bandwidth, and refractory period to illustrate the theoretical response of the DAVIS346. With minor changes to these parameters, the simulations could be adapted to other EVS models. A screenshot of the GUI is shown in Fig. 2. The figure provides a visual example of how threshold, bandwidth, and refractory period influence the expected response of an individual EVS pixel as a point source object with finite Point Spread Function (PSF) passes through its Field of View (FOV).

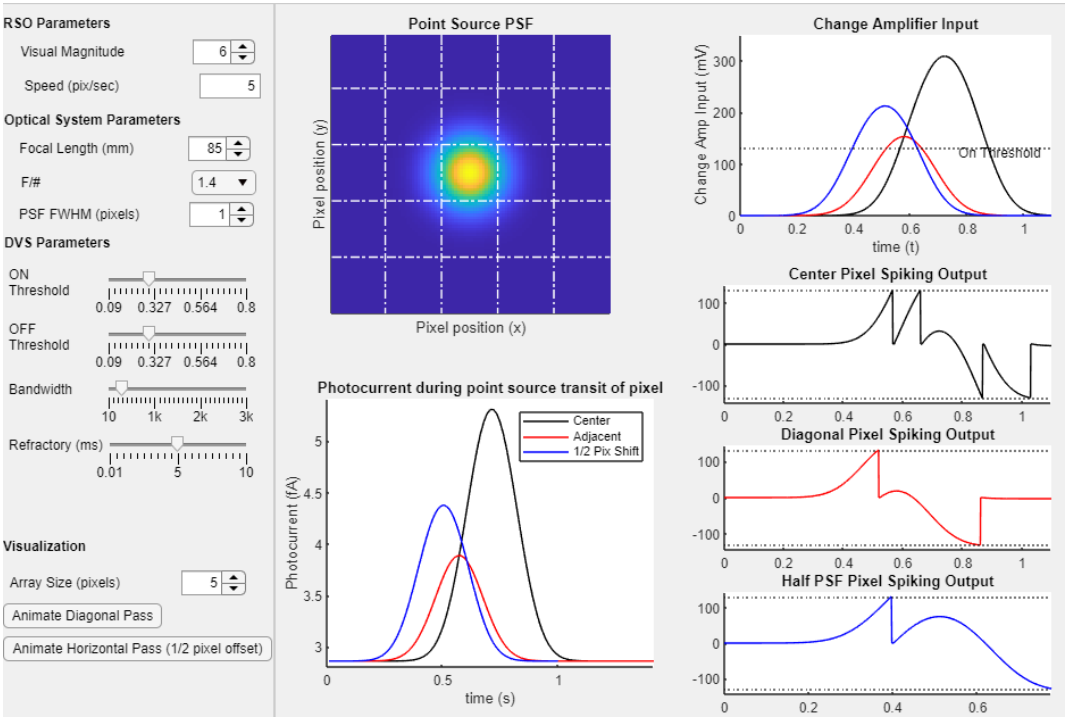

Fig. 2: Our custom MATLAB GUI models various bias controlled parameters (thresholds, bandwidth, and refractory period), and optical system parameters (focal length, F/#, and PSF) to predict EVS performance limits for dim/fast sub-pixel object detection. The visualization includes a model of photocurrent, change amplifier input voltage V_{sf}, and spiking output of the change amplifier V_{diff} for three pixel transit scenarios: 1- object passing through center of pixel (black), 2 - object passing along edge of pixel (blue), 3 - object passing diagonally touching only corner of pixel (red).

[1]Available at: https://github.com/SensorsINI/sda-dvs-sim

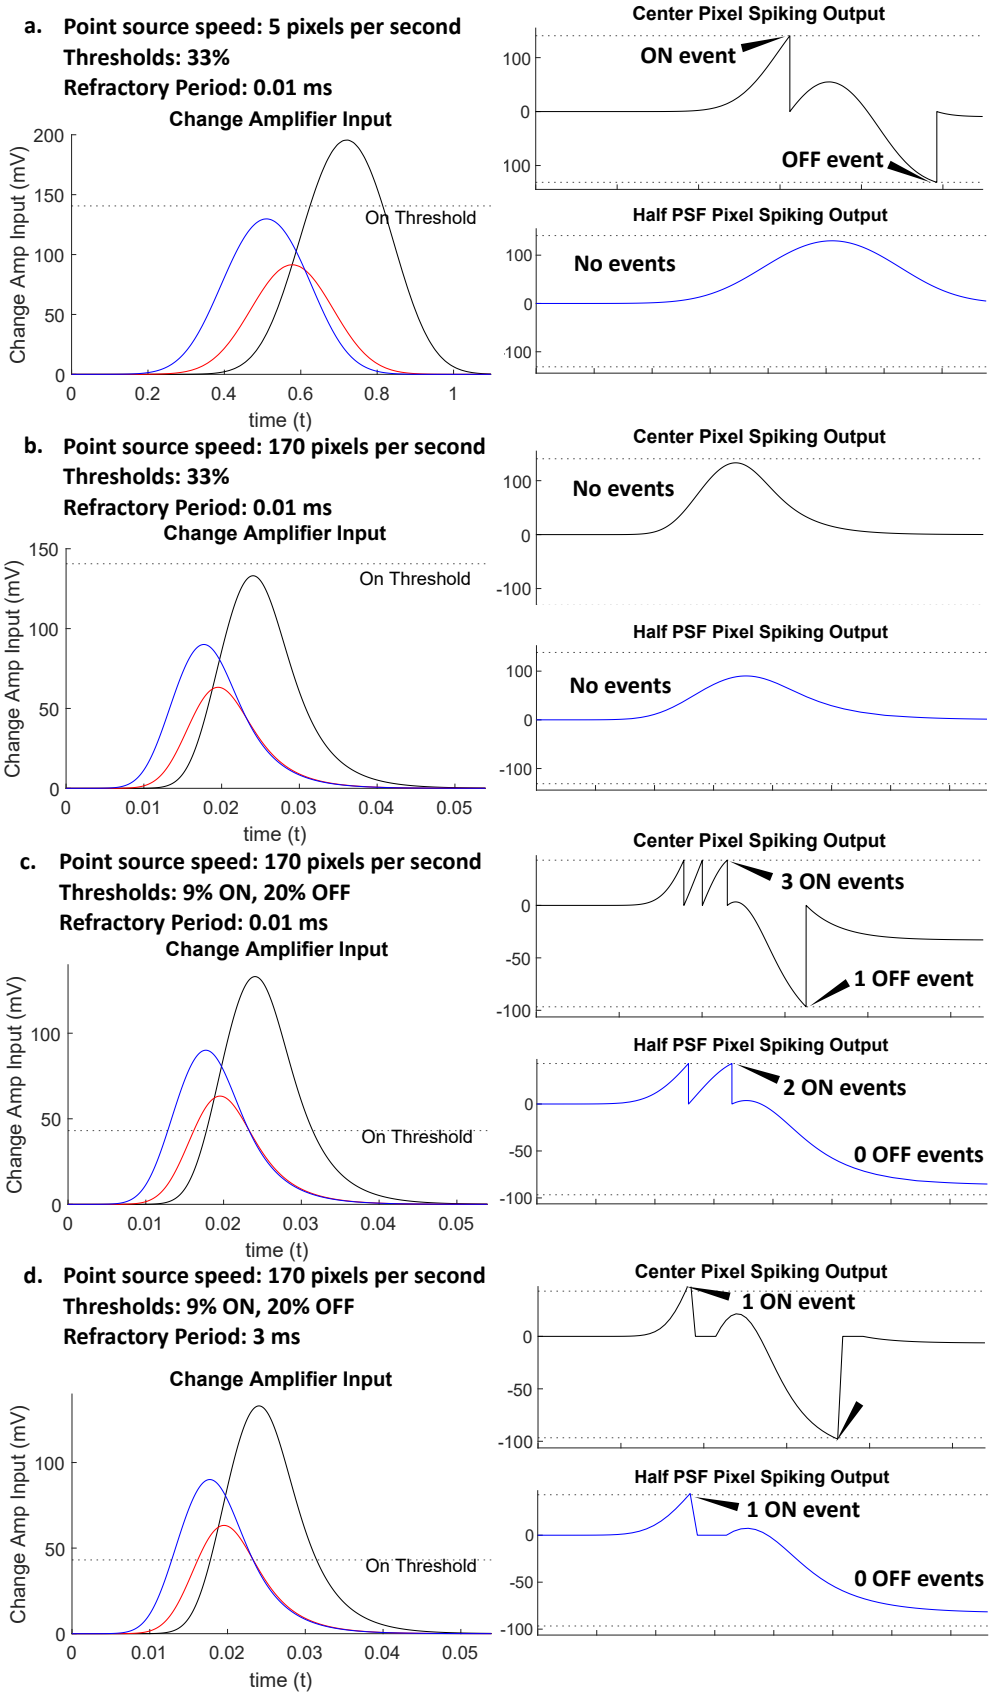

Fig. 3: Pixel Simulation Results. Pixel response to a 4.5 m_V, 1 pixel FWHM PSF point source was simulated for two transit paths: through pixel center (black) and along pixel edge (blue). Different speeds and sensor settings are shown to demonstrate how settings can be tuned to influence sensor response in SDA tasks. For a moving, sub-pixel object to generate a continuous streak in space, both black and blue curves should generate at least one event.

Fig. 3 provides a visual example of how sensor settings influence the expected response of an individual EVS pixel as point source objects of varied speed pass through its FOV along different trajectories. Using our custom MATLAB simulator, we simulated a point source object of 4.5 m_V passing through the center (black) and edge (blue) of a DAVIS346 pixel, using a focal length (41 mm) and f/# (1.8) matching our lab set-up and setting the PSF Full Width Half Max (FWHM) to 1 pixel. Fig. 3**a** demonstrates a condition where the thresholds are sensitive enough to detect the object both entering and exiting the pixel's FOV when it passes through the center, but insufficient for objects passing along the pixel's edge. Part **b** demonstrates how increasing objects speed attenuates the input signal due to the low-pass characteristic of the photoreceptor so the pixel is not able to respond, and **c** shows that decreasing the thresholds allows the pixel to respond to both transit paths, even at faster speeds. Part **d** then uses the same parameters used in **c** but increases the refractory period to show how information may be suppressed when adjusting this parameter. If an application does not require an estimation of the exact brightness change, this may be acceptable, and setting a long refractory period is a technique that has shown potential utility for noise rate management [6].

3. METHODOLOGY

To empirically test bias optimization for SDA, we developed a custom lab set-up for simulating point-source targets of variable speed and brightness. For direct comparison to other studies, the set-up was carefully calibrated to photometric (illuminance, [lux]) and astronomical (apparent visual magnitude, [m_V]) units. We then used the techniques described in Sec. 2.2 to develop eleven custom bias configurations to compare performance in detecting both the dimmest and fastest moving point source objects. This set-up allows for rapid, automated data collection and accurate mapping of ground truth (object position) to separate signal from noise.

3.1 Lab Set-up

The physical lab set-up shown in Fig. 4 is centered around a precision-fabricated, 17 cm radius, anodized aluminum disk containing small holes of 100, 120, 150, 200, and 250 μm in diameter. Five holes of each diameter are located at varied radial distances from the center mounting point: 7, 8.75, 10.5, 12.25, and 14 cm. The hole sizes and distances were carefully chosen to compare objects of differing brightness and speed, resulting in brightness increments of approximately 0.5 m_V and speed increments of 1.25×, 1.5×, 1.75×, and 2× for increasing size and radial distances. The disk is mounted to a digitally programmable motor by Henschel Robotics which, in stepper mode, is capable of precise speed control down to a few hundredths of a degree per second. The upper speed limit was not tested, but was sufficient to observe the limits in photoreceptor response speed for the dimmest (smallest) point source targets with all eleven bias configurations considered.

The motor/disk combination is placed in front of an 'integrating' light box, with a variable input illumination level provided by a high-power, white LED (Cree Inc. MHDGWT-0000-000N0HM230GCT) and fed into the box by a fiber optic cable. The power circuit for the LED is described in [19], and the opening of the light box is $\approx 7.5 \times 10$ cm to match the format of the DAVIS346 Focal Plane Array (FPA). A Fujinon 7-70mm, f/1.8 lens is used to image the portion of the disk corresponding to the opening of the light box, creating the illusion of moving stars or Resident Space Objects (RSO) as the holes on the disk rotate through the FOV. The viewing distance (63 cm) and focal length (41 mm) were selected to ensure that even the largest point source objects (holes) are still sub-pixel compared to the DAVIS' 18.5 μm pixel pitch when the image is in focus (i.e. $\frac{41mm}{63cm} \times 250\mu m = 15.8\mu m$).

3.2 Bias Selection

In this study, we propose that truly pushing the performance limits of an EVS when the background is dark requires accepting some degree of elevated noise rates. The eleven hand-selected bias combinations we crafted for are shown in Table 2, and labeled according to the main performance parameter(s) adjusted to limit background noise rates. For fair comparison across configurations, we tuned the indicated biases while continuously monitoring noise rates while viewing a very dark, uniform background illumination level ($\approx$ 5 mlux) until the average noise rate across the array settled to 0.75 Hz/pixel (noise events per second per pixel). This value was selected both to limit the size of recordings and to facilitate differentiating signal from noise for later analysis.

In row 1 of Table 2 (Slow BW), we start with the simplest methods of noise management to illustrate baseline performance, first limiting pixel bandwidth by reducing I_{pr} and I_{sf} proportionally from their default settings. Row 6 (High TH) represents another simple noise management technique of increasing both ON and OFF thresholds while keeping them approximately equal in magnitude (as evidenced by balanced ON and OFF noise event rates). Configurations

in rows 2-4 are inspired by the observation that noise rates can be reduced with imbalanced ON and OFF thresholds [18]. Row 3 tests a slightly less imbalanced configuration than row 2, and row 4 combines imbalanced thresholds with faster pixel bandwidth settings. Row 5 tests the method of increasing refractory period by significantly reducing I_{refr} to decouple opposite polarity event pairs [18].

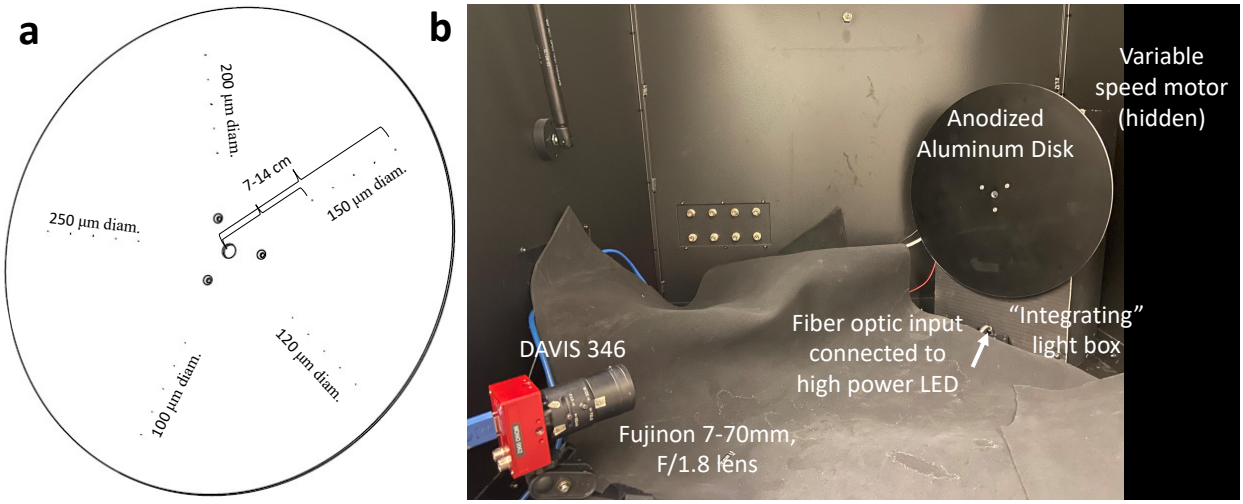

Fig. 4: Disk and lab set-up. **(a)** Diagram of our custom-fabricated, anodized aluminum disk with precision drilled holes of 100-250 μm diameter located at radial distances between 7-14 cm. **(b)** Moving point source targets are simulated by rotating the disk in front of an 'integrating' light box with variable input intensity and viewed by a DAVIS346 EVS through a 7-70mm, F/1.8 lens. The entire assembly was placed into a dark box to approximate night sky illumination, which for typical viewing conditions is below the dark current limit of the photodiode at $\approx$ 1-10 mlux.

Rows 7-10 draw inspiration from [11] by testing various combinations with extremely high values of I_{pr}. Row 7 uses high I_{sf} (Fast SF), and row 8 uses a low (Slow SF) setting. Rows 9 and 10 test the extreme (Fast) I_{pr} setting with both balanced and imbalanced threshold levels. A significant drawback of these four combinations is a drastic increase in static power consumption of $15\mu\text{A} \times 1.8\text{V} \times 90\text{k pixels} \approx 2.4\text{W}$. Therefore, we tested one additional bias configuration (row 11) with a less extreme I_{pr} value that still leverages the combination of high photoreceptor bias and low source follower bias. To refine noise rates, the selected I_{pr}-I_{sf} combination used in configuration 11 was also paired with imbalanced ON and OFF thresholds.

3.3 Data Collection

Data collection was automated by a Python script run in a Jupyter notebook capable of interfacing directly with the motor through its web interface and the DAVIS camera through jAER. With the motor running at a constant speed, the angular position is continuously monitored, and when the disk reaches a pre-determined angular position corresponding to the closest object entering into the camera's FOV, a command is sent to jAER to initiate a recording. At the same time, the time and angular position of the motor are recorded and saved for post-processing. Once the streaks have passed through the FOV, the recording is saved, and again, angular position and time are recorded from the motor.

To ensure the capability of distinguishing signal from noise in our analysis, we selected a minimum object motion rate of $\approx$ 5 pixels per second (center streak), to guarantee the ability to visually and programmatically distinguish signal events from noise. At this speed, we should be able to look for signal events in an individual pixel within a 0.2 s window, and during that same window, the expected number of noise events is only $0.2\text{s} \times 0.75$ events/sec ≈ 0.15 events, compared to at least 1 expected signal event if the object is bright enough to trigger an event. When averaging pixel response over the five streaks, this balance allows for a reasonable method to separate signal from noise and assign a quantitative metric corresponding to pixel response for differing brightness levels.

Bias Configuration	I_{pr}	I_{sf}	I_{refr}	I_{ON}	I_{OFF}	I_{diff}
1 Slow BW	21.1 pA	834.7 fA	4.7 nA	1.7 μA	225.6 pA	19.6 nA
2 Imbalanced TH	638.1 pA	24.5 pA	4.7 nA	449.8 nA	11.9 pA	100.3 nA
3 Less Imbalanced TH	638.1 pA	24.5 pA	4.7 nA	1.9 μA	68.3 pA	100.3 nA
4 Imbalanced TH, Fast BW	2.3 nA	90.5 pA	4.7 nA	1.0 μA	11.9 pA	100.3 nA
5 Longer RP	682.6 pA	26.0 pA	187.0 pA	4.7 μA	225.6 pA	19.6 nA
6 High TH	638.1 pA	24.5 pA	4.7 pA	5.6 μA	67.9 pA	20.5 nA
7 Fast PR and Fast SF	16.8 μA	9.1 nA	4.9 nA	1.7 μA	225.6 pA	19.6 nA
8 Fast PR and Slow SF	1.3 μA	324.6 fA	4.0 nA	1.7 μA	225.6 pA	19.6 nA
9 Fast PR, Slow SF, Balanced TH	15.0 μA	834.7 fA	4.7 nA	186.9 nA	2.0 nA	8.2 nA
10 Fast PR, Slow SF, Imbalanced TH	15.0 μA	834.7 fA	4.7 nA	205.1 nA	1.9 nA	19.6 nA
11 Optimal PR, Imbalanced TH	40.1 nA	11.9 pA	4.7 nA	328.2 nA	5.9 pA	19.6 nA

Table 2: Bias configurations evaluated. Eleven custom bias configurations were crafted to examine the performance limits of the DAVIS346. A brief description of each configuration is shown along with estimated current values indicated by jAER [1] for each tunable bias.

Each set of streaks was recorded for a range of 16 LED voltage (brightness) levels for each hole size, resulting in 80 distinct recordings for each bias. An additional two voltage levels were included for mapping target position ground truth (extremely bright) and average noise rates (0V input) for the pixels along the path of each point source. Additionally, to examine the drop in sensitivity with speed, the recordings were repeated over a range of 12 different speed settings with a constant, mid-low brightness input of 15 mV. After an initial visual assessment of performance, the 5 best performing bias configurations were further analyzed at 5 mV for an additional 6 speed settings. This resulted in 60 and 30 variable speed recordings respectively for the two voltage levels assessed.

3.4 Photometric and Astronomical Calibration

To accurately estimate the true apparent magnitude of the holes, we used the following procedure.

In addition to event-based output, the DAVIS346 includes an Active Pixel Sensor (APS) option to record standard exposure measurement frames by integrating each pixel's photocurrent for over a fixed time window. This exposure measurement feature was used to verify that the light box created a flat-field illumination level across its opening. Then, illuminance values at the output of the integrating light box were calibrated against the available range of LED power source voltages using a Tektronix J17 photometer with the J1811 illuminance head. Rather than assuming the total light intensity incident on the back of the disk is collected over the surface area of each hole, we performed an additional step to approximate the amount of light lost to back reflection or absorption. Again, using the DAVIS exposure measurement, we estimated that approximately 90% of the luminous flux incident on each hole was projected outwards, with only $\approx 10\%$ loss. Next, the luminous flux [lumens] of each point source (hole size) was calibrated by multiplying the illuminance measured by the photometer [lumens/m^2] by the area of each hole.

In the final step of photometric calibration, we assume each hole appears to the observer (or sensor) as a uniform flat disk. As a result, the luminous flux escaping the hole emanates into 2π steradians with a $\cos\theta$ dependency, where θ is the angle between the object and image planes. This assumption results in a simple relationship between the luminous flux (ϕ_v) of each point source and the illuminance (E_v) at the observer's location,

$$E_v = \frac{\phi_v}{\pi d^2}, \tag{3}$$

where d is the distance from the light source to the entrance pupil of the sensor. The conversion between illuminance at the sensor's location and apparent magnitude is then straightforward. For astronomical objects, apparent magnitude

in the visible band is calculated as a function of the illuminance from that object at the observer's position (typically the earth's surface) by the relationship,

$$m_V = -14.18 - 2.5\log_{10}(E_v). \tag{4}$$

A simple validation was conducted by observing the dimmest visible hole-size/voltage combination after ≈ 30 min of dark adaptation. The limit was found to be 6.25 m_V, which is squarely in the expected range of limiting magnitude for unaided human vision (6-6.5). Precise calibration is not vital for the bias configuration comparisons at the heart of this study; however, confirming the accuracy of the calibration will allow us to make approximate comparisons to previous EVS limiting magnitude estimates and better parameterize physics-based sensor models using the limits observed in this study. Once confirming the calibration produced viable results, Eqs. (3) and (4) were used to convert photometer-measured illuminance values and generate the calibration curves shown in Fig. 5.

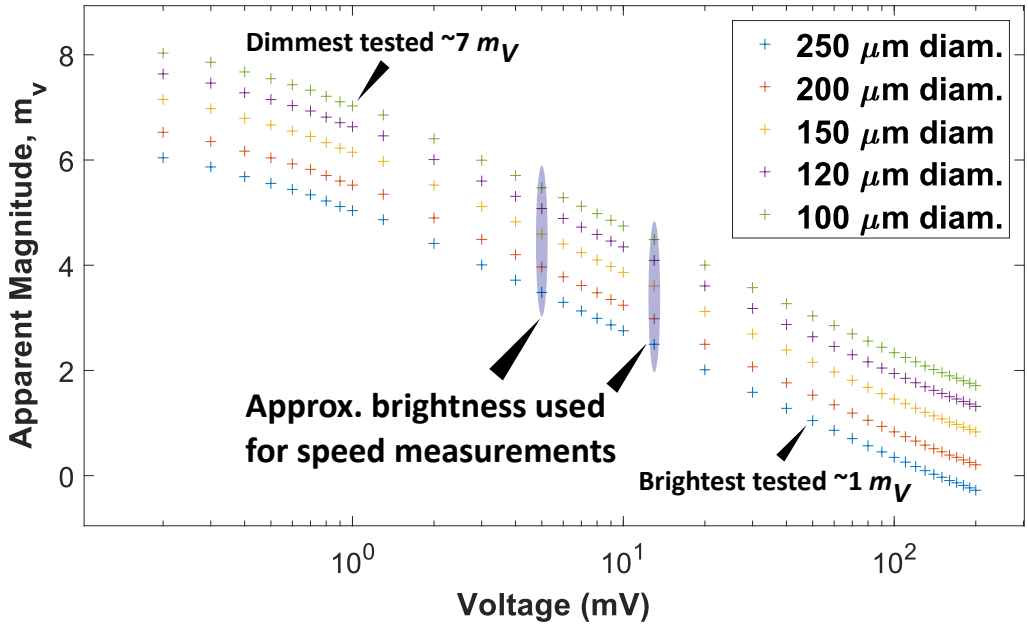

Fig. 5: Point source calibration. Each point source was calibrated to apparent magnitude across a range of LED driver input voltages. EVS output was collected across 16 voltage levels between 50mV and 1mV to examine limiting magnitude, resulting in apparent visual magnitude estimates from 1 to 7 m_V. Voltage and brightness levels selected for subsequent measurements to characterize drop in sensitivity with increased target speed are highlighted.

4. ANALYSIS

Initial analysis was conducted by viewing the recorded data in jAER to determine if the streak(s) were visually detected for each calibrated m_V level and recording the dimmest observed object for each bias configuration. The same validation was done for the variable speed recordings. Beyond this initial assessment, the collection methodology produced data points to map object position as a function of time. Post-processing the data for further analysis consisted of four key steps: 1 - mapping the pixel addresses along each streak as a function of motor angular position, 2 - examining the noise event profile of the mapped pixels, 3 - extracting matched events occurring within a small temporal window of the expected time that the object passes through each pixel, and 4 - subtracting the expected number of noise events during that same temporal window. This mapping produced a quantitative result for comparing the limiting magnitude of various biases; however, for the faster speed settings, variable trigger delay times led to too much uncertainty in the motor's angular position readout to apply the same quantitative analysis. Therefore, we only report the maximum speed detected for bias/brightness combination based on our own visual analysis of the jAER recordings.

4.1 Mapping Target Position Ground Truth to Determine Signal Events Per Pixel

Pixel addresses along the streak were mapped by using the 150 mV (brightest LED) recording to select the pixels with the highest response. Events were binned into 400 ms time-slices (frames) by summing the total number of events per-pixel in each time-slice and placing the result into a 346×260 array. These time slices were stacked to create a 3D array of 'pseudo-frames'. The array was then compressed into a single image by extracting each pixel's maximum value (event count) along the time axis, and several iterations of spatial correlation filters were applied to remove noisy pixels and identify only pixels belonging to the point source streaks. This method produced a list of $\approx$ 1.3k pixel addresses for each set of holes. Using the recorded start and stop angles and times of the motor, the discretized temporal information in the pseudo-frame stack was used to map each pixel's maximum value along the time axis to the corresponding angular position of the motor. To validate correct spatio-temporal mapping, the vertical (y) position of each pixel was plotted against the time sample corresponding to when the maximum value was observed. Because the targets move roughly vertically through the FPA in time, this produced five distinct lines, one for each streak. An example of the spatial map is shown in Fig. 6, along with the plot of y position as a function of time for all mapped pixels in a single recording. After verifying the accuracy of the pixel maps, the x, y pixel address and corresponding angular position were determined by linearly interpolating the motor angle vs. time read directly from the Henschel motor. The pixel addresses and angles were saved and used as ground truth for subsequent analysis of the sensor's response to varied illumination levels.

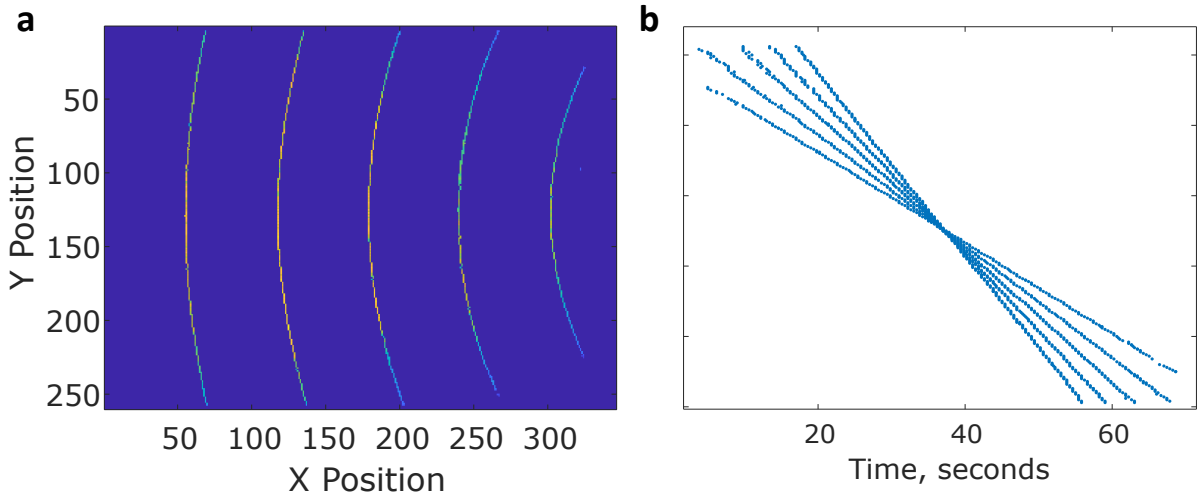

Fig. 6: Example Pixel Map. **a.** Point source position ground truth was constructed by creating pixel maps of maximum event counts recorded in response to the highest illumination level for each set of streaks. **b.** The vertical (y) position of each pixel's maximum value was then plotted against the time sample during which it was observed to verify correlation with a moving object. The temporal information was then combined with angular and temporal information extracted from the motor to extract the motor angular position corresponding to the target's transit of each pixel.

After mapping signal ground truth for each pixel as a function of motor angular position, correlating expected true signal events for the remaining recordings was straightforward. For each recording, we spatio-temporally matched signal to ground truth by summing events reported by mapped pixels within a temporal window of $\pm$ 0.5 seconds of the target's expected pixel transit time using the recorded start and stop angular positions from the motor. Because the bias configurations were selected to push performance limits and thus allowed elevated noise rates, it was necessary to differentiate signal from noise. Noise rates can vary dramatically from pixel to pixel; however, an individual pixel's noise rate for a given bias configuration is relatively stable in time. Considering this, we used the pure background noise recordings (0 mV LED setting) to find each mapped pixel's average noise rate over $\approx$ 360 seconds. For the group of mapped pixels, we then multiplied the average noise rate by the assessed temporal window of 1s to find the expected number of erroneously matched noise events in the absence of any signal. Finally, we subtracted the number of expected noise events from the number of spatio-temporally matched events and divided the result by the total number of matched pixels to obtain the number of matched signal events per pixel for each recording.

4.2 Sensitivity

Using jAER to play back each recording and manually adjusting the frame-rate, rendering-rate, and rendering scale (number of events required for maximum visual response), we analyzed all 80 recordings for each bias. Table 3 reports the highest apparent magnitude observed for each of the eleven bias configurations. In some cases, the streak was only visible for a short segment due to the limited depth of field of the optical set-up, causing point source objects to go slightly out of focus near the edges of the FPA. Still, even in these cases, the target generated enough sequential events in adjacent pixels to be visually detected. Based on this initial rough estimate, it is clear that some bias configurations clearly outperform others on the basis of limiting visual magnitude, with a range spanning $3.77m_V$ to $7.03m_V$, a factor of about $20\times$ in intensity. Biases 1 (Slow bandwidth), 5 (Long refractory period) and 6 (High thresholds) are the simplest configurations selected, and with the exception of 1, generally perform much worse the more complex bias combinations.

The technique of increasing refractory period (config. 5) is the worst performing of all the biases selected. Although shown to be an effective method for controlling noise rates, recent work suggests the refractory period must be $\approx 4\times$ the characteristic time constant of the photoreceptor's temporal response in order to decouple noise event pairs [15]. Because the photoreceptor is quite slow in the dark, this requires a refractory period near the practical maximum limit of the circuit, and it's likely a large fraction of the pixels are no longer operating properly with this configuration. Despite much higher power consumption, configurations 7-10 only perform roughly on par with config. 1 (slow bandwidth), and config. 3 (slightly imbalanced ON and OFF thresholds) does not appear to have a significant advantage. On the other hand, the three best performing biases for the task of dim target detection all employ some degree of imbalanced thresholds.

Bias Configuration	1	2	3	4	5	6	7	8	9	10	11
Dimmest Observed (m_V)	5.7	**7.03**	5.47	6.4	3.77	4.98	5.52	5.47	5.99	5.99	6.4

Table 3: Dimmest visible object for each bias configuration determined by visualizing recordings in jAER.

From the initial visual analysis, **setting imbalanced threshold levels in favor of more ON events is by far the most effective technique for improving dim object detection for SDA**. Configuration 2 applies thresholds with the highest degree of imbalance, and achieves the best performance observed at 7.02 m_V. Configurations 4 and 11 each also have a high degree of imbalance (ON more sensitive than OFF), and achieve the same limiting performance based on visual inspection. Compared to the two most common noise management methods (1-Slow Bandwidth, and 6-High Thresholds), by initial estimate, config. 2 attains a performance advantage of 1.33 m_V ($3.4\times$ dimmer) and 2.05 m_V ($6.6\times$ dimmer) respectively!

Because the visual detection method is only qualitative and lacks resolution near the limiting magnitude, we used the methodology described in Sec. 4.1 to generate plots of mapped events per pixel vs. calibrated apparent magnitude. A sample of the resulting plots is shown in Fig. 7. Different hole sizes are represented by distinct symbols. While the general shape is extremely similar, the varied hole sizes do generate slightly different curves. We suspect the main reason to be difficulty in obtaining perfect focus of the lens, resulting in some minor image blur. As a result, the image of smallest holes remains sub-pixel in size, but some fraction of the photon flux from the larger holes begins to bleed into adjacent pixels, reducing the likelihood of response for the main (center) pixel. As a result, we used the data points from the 100 μm holes to quantitatively compare the pixel response.

We first generated a polynomial fit for pixel response over the range of brightness values measured (see Fig. 7). A third order polynomial was found to match closely to the measured data points over the range of interest – where the curve drops below 1 matched signal event per pixel – for all bias configurations. We interpolated these curves to find the brightness level in m_V corresponding to a matched event rate of 1 event per pixel. This value was chosen as a conservative estimate of limiting magnitude, but it's likely that a robust detection algorithm would be capable of picking out a target, even if some fraction of pixels fails to respond. The results from the two best performing biases (configs 2 and 11) are shown in Fig. 8, along with the two simplistic bias tweaks commonly used to manage noise rates when the background is dark. Applying this quantitative analysis interpolates between the gaps from the data recordings to provide a higher resolution comparison, but still points to a considerable advantage for the more advanced biasing methods. Configuration 2 outperforms the high threshold (6) and slow bandwidth (1) settings by **1.6 and 1.04** m_V respectively. Configuration 11 achieves slightly lower performance, but still shows marked improvement

over the simple biasing techniques with advantages of 1.29 and 0.73 m_V.

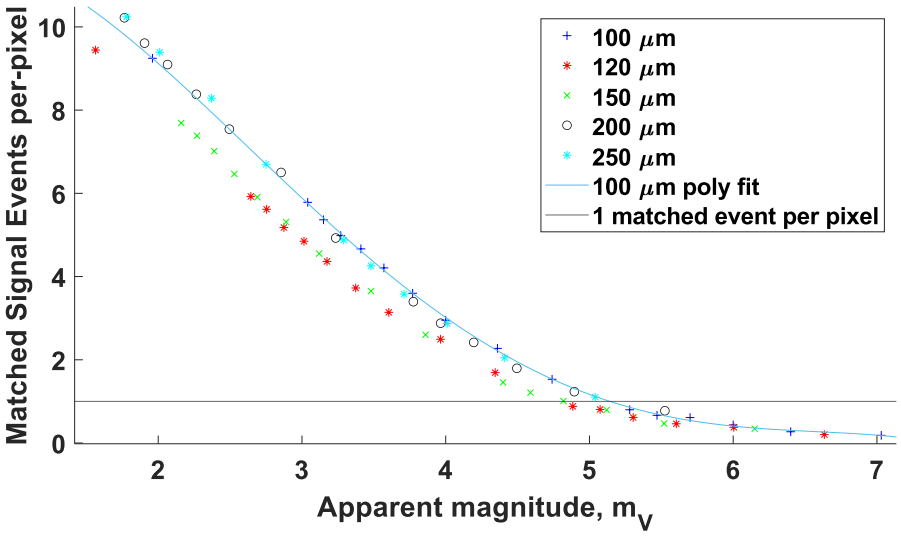

Fig. 7: Example plot of matched signal events vs. calibrated apparent magnitude. Different hole sizes generated some slight variation in the exact position of the curve, but the shape of the response curve was closely matched across all recordings. The smallest hole size was selected for further analysis, and a 3rd order polynomial was fit to the raw data points. By inspection, the polynomial fit models the pixel response well in the region corresponding to 1 event per pixel, so can we use this curve to interpolate the pixel response and compare approximate limiting magnitude for different bias configurations.

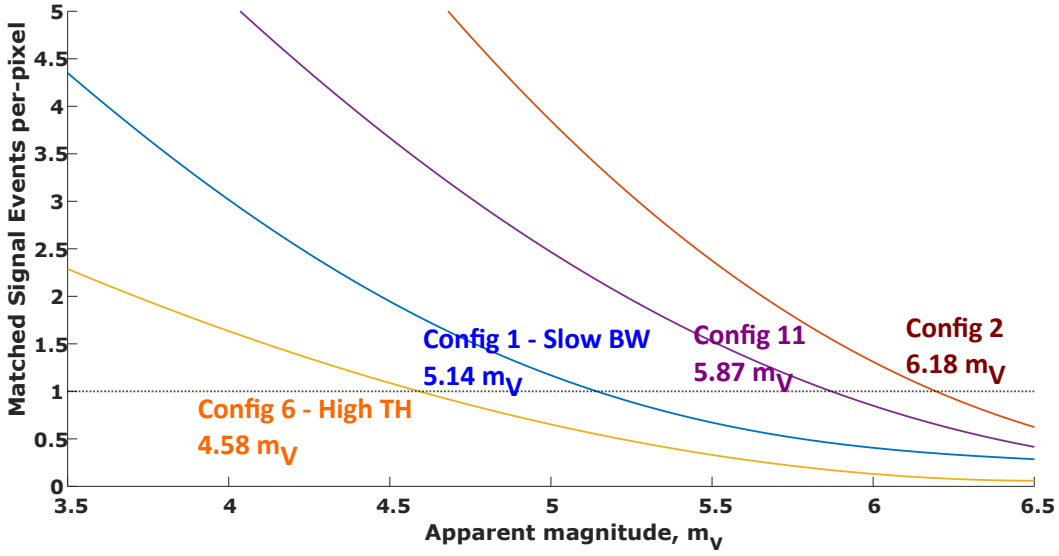

Fig. 8: A third order polynomial was fit to the matched pixel response across the range of evaluated tested brightness. The curves were interpolated to find the brightness level corresponding to 1 matched event per pixel to compare sensitivity across bias configurations. The two best (2 and 11) are plotted alongside two simplistic, baseline configurations.

4.3 Speed Limitations

In addition to finding the limiting visual magnitude for each bias configuration, we also examined the maximum detected speed for two LED power settings, 15 mV and 5 mV. Results based on visual analysis are shown Table 4. Again, bias configurations 2 and 11 vastly outperform the rest of the field. Given the intuition gained by simulating different parameter combinations in our MATLAB GUI, this result can be rationalized. Even if the object's speed exceeds the photoreceptor's low-pass corner frequency, a more sensitive (lower threshold) pixel will respond deeper into the high frequency range before the amplitude of signal pulse caused by the object's motion drops below a single threshold level. Comparing the results to the simple, single parameter bias tweaks, the optimal settings **outperform configs 1 and 6 by 6.6× and 2.5× respectively**.

Although the same streaks were visible for identical voltage settings between the configurations 2 and 11, the streaks from config. 11 were subjectively slightly easier to see. Given the limiting visual magnitude estimates, the opposite should be expected. Even more notably, config. 6 (high threshold) outperforms config. 1 (slow bandwidth) by a significant margin (82 to 31 pixels per second), despite a much lower sensitivity observed in Table 3 (4.98 to 5.7 m_V). These results imply that the photoreceptor and source follower biases indeed have a significant effect on response speed, and the optimization technique implemented in configuration 11 has at least a minor influence on speeding up the pixel response compared to baseline biases.

Bias Configuration	1	2	3	4	5	6	7	8	9	10	11
Fastest Observed											
@ 3.98 m_V (pixels/sec)	36	**243**	36	206	N/A	97	145	67	97	145	**243**
@ 5.47 m_V (pixels/sec)	N/A	**67**	N/A	36	N/A	N/A	N/A	N/A	12	19	**67**

Table 4: Fastest visible object for each bias configuration for objects of 3.98 and 5.47 m_V determined by viewing recordings in jAER.

Quantitative analysis was not possible for all data sets due to hardware and software limitations causing uncertainty in the timing of the higher speed recordings. However, a few sets produced clean enough results for a more thorough analysis on the influence of biases on the sensor's temporal response. The same method described in Sec. 4.1 was applied to the varied speed data recordings for configurations 1, 6 and 11 at 15 mV LED power supply voltage. The brightest (250 μm) targets produced the best results for analysis, which correspond to 2.39 m_V point sources. To examine the drop in sensitivity vs. speed across the three selected bias configurations, the responses were first normalized by dividing the number of matched events per pixel by the maximum value observed (i.e. at low speed) for the same bias and illumination level. The normalized pixel response (fraction of peak response) is plotted against log speed in Fig. 9. For each data-set, a linear fit is applied to the region between the corner frequency (f_{3dB}) and the point where the curve begins to level out near zero response.

The slope of -0.36/e-fold of speed observed in all cases has a simple yet enlightening explanation. We assume event count is a discrete approximation of the magnitude of the current pulse generated as a point source object transits a pixel's FOV. At transit speeds above the photoreceptor corner frequency, the signal amplitude downstream of the low-pass filter is attenuated as a function of speed, and thus generates fewer events per pass. For a first-order system, the attenuation rate is -10dB per decade, and our experimental result shows that for an increase in speed of 1 e-fold ($\approx$ 2.71), the number of detected events is reduced by 0.36, or $\approx 1/e = 0.37$. This result is consistent with the theoretical drop in signal amplitude as a function of speed for a first order system – an e-fold increase in speed results in a $1/e$ decrease in the amount of time the object's photon flux resides on the pixel, thus a $1/e$ drop in signal amplitude.

The experimental results help explain the qualitative visual analysis contained in Table 4 and highlight the fact that the ability to detect faster objects is a combination of both sensitivity and temporal response. Notably, config 1 was previously found to be more sensitive than config 6 by a factor of 0.46 m_V; however, a higher pixel bandwidth (corner frequency) extends 6's response into higher speed ranges for a stimulus that is visible for both configurations (82 vs. 31 pix/sec). On the other hand, configs 11 and 6 have very similar temporal response profiles in terms of low-pass corner frequency, f_{3dB}, but the higher sensitivity of 11 results in the same object being visible at speeds roughly 2.5 times faster (206 vs 82 pix/sec). The approximate f_{3dB} location inferred from the plotted data is indicated for the three bias configurations shown in Fig. 9.

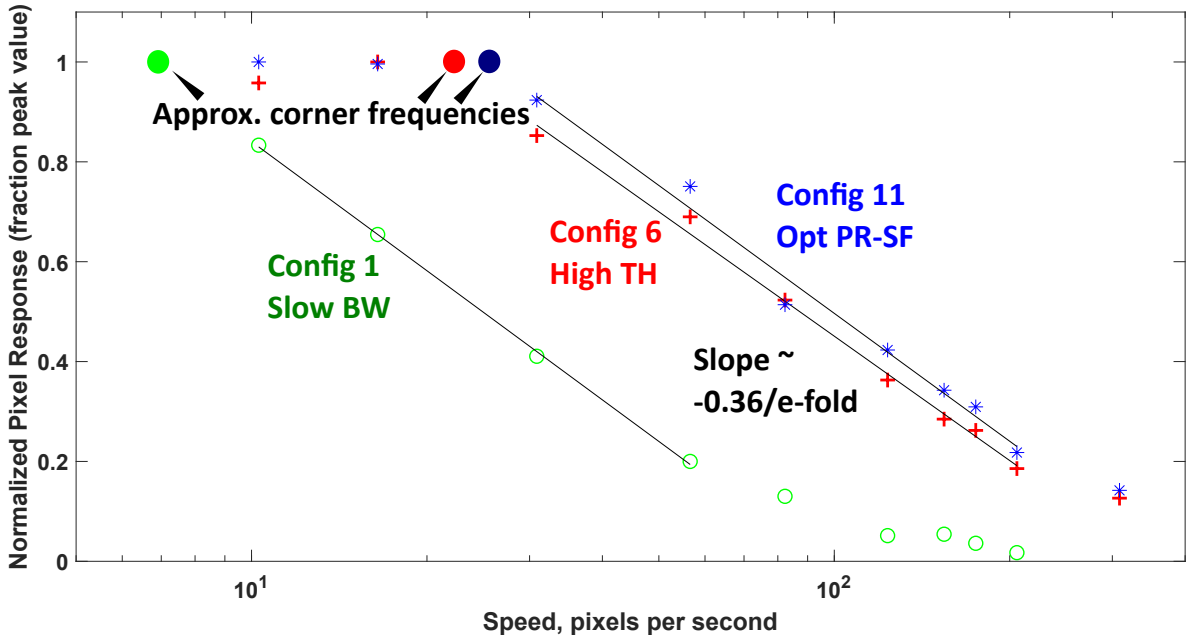

Fig. 9: Drop in sensitivity with speed. For each configuration, the normalized pixel response (number of detected signal events divided by peak pixel response) drops predictably as a function of speed above a particular value. This value is assumed to be corner frequency of the photoreceptor circuit, f_{3dB} (Hz), and its value depends on sensor biases.

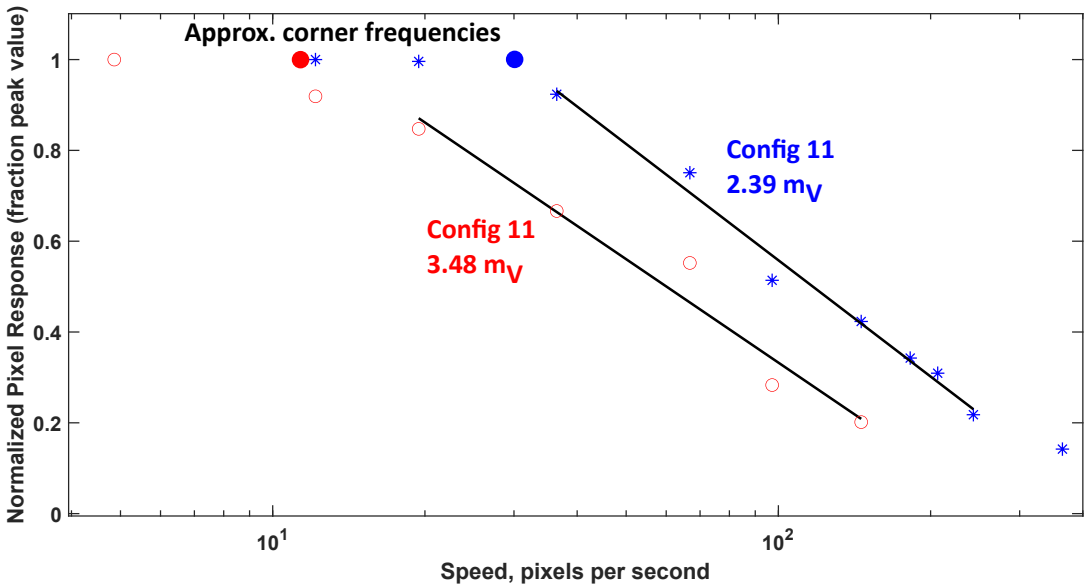

Fig. 10: Objects of different brightness levels result in different corner frequency estimates, even with the same bias configuration. This implies the brightness of the point source itself determines the pixel's temporal response as it transits the FOV.

Although data was insufficient to generate plots for config 2, the photoreceptor and source follower biases are identical for configurations 2 and 6. The slightly faster response of configuration 11 evidenced by Fig. 9 helps explain how it achieves similar speed performance to 6 (see Table 4) despite slightly lower sensitivity (Table 3). This result hints that

combining the highly imbalanced thresholds of config 2 with the more optimized PR-SF balance of 11 could further extend the performance limits observed in this study to approach a truly optimized bias configuration for SDA.

Examining drop in sensitivity with speed for two different brightness levels, we observe another phenomenon that is predicted by the physics of the photoreceptor circuit, but often overlooked. Fig. 10 shows that even for the exact same bias configuration and background illumination, the approximate low-pass corner frequency varies by a factor of $\approx$ three – the ratio of the two brightness levels tested. This observation is consistent with the prediction that the corner frequency varies monotonically with increasing photocurrent for low illumination; however, the result in this case contains an interesting nuance. Typically the corner frequency is considered a function of DC background illumination, but our measurements show that the transient impulse in photocurrent caused by the passing object itself changes the temporal response of the photoreceptor. Therefore, **when the background is dark, point source objects of varying brightness levels themselves determine $f_{\mathbf{3dB}}$ as they transit through a pixel's FOV!**

5. CONCLUSION

Our experimental results clearly demonstrate that advanced biasing techniques can significantly extend the utility of the DAVIS346 camera for detecting both dimmer and faster moving point source objects in the SDA environment. We leveraged a deeper understanding of EVS pixel architecture and operation to explore the large bias parameter space, selecting 11 custom bias configurations intended to simultaneously increase detection performance and manage background noise rates. To compare performance, we allowed a consistent baseline background noise rate across all configurations. We found the method of deliberately applying imbalanced ON and OFF thresholds to be most effective for increasing sensitivity, demonstrating an increase in sensitivity of up to 1.6 m_{V} compared to simply increasing thresholds to limit noise rates, and an increase of 1.04 m_{V} compared to the technique of restricting pixel bandwidth. Imbalanced thresholds in combination with techniques to optimize photoreceptor biases demonstrated the ability to detect objects moving 6.6$\times$ faster than the restricted bandwidth configuration, and 2.5$\times$ faster than the high threshold configuration.

This work also opens several future directions. First, the biasing techniques demonstrated should be tested on newer EVS variants to see if similar performance improvements are achievable. Although newer generation models feature higher fill factor and quantum efficiency, utilize more advanced fabrication technology, and incorporate more advanced readout schemes, the general pixel operating principle remains the same. Second, the data collected contains mapped ground truth for sub-pixel, point source objects of varied brightness, suitable for developing, testing, and training EVS de-noising, detection, tracking, and characterization algorithms in this challenging sensing environment where varied conditions can drastically change the profiles of signals of interest. Additionally, this study's data should inform improvements to current low illumination EVS modeling by providing both another avenue of model evaluation and increasing our understanding of the factors affecting the corner frequency of the photoreceptor circuit. Finally, knowledge about the coupled effect of motion and visual magnitude should be incorporated into operational SDA concepts involving EVS to include event readout limitations and additional resulting system requirements.

6. REFERENCES

[1] jaer: Java tools for Address-Event representation (AER) neuromorphic vision and audio sensor processing.

[2] Saeed Afshar, Andrew P Nicholson, Andre van Schaik, and Gregory Cohen. Event-based object detection and tracking for space situational awareness. November 2019.

[3] Christian Brandli, Raphael Berner, Minhao Yang, Shih-Chii Liu, and Tobi Delbruck. A 240$\times$ 180 130 db 3 μs latency global shutter spatiotemporal vision sensor. *IEEE J. Solid-State Circuits*, 49(10):2333–2341, 2014.

[4] Gregory Cohen, Saeed Afshar, Brittany Morreale, Travis Bessell, Andrew Wabnitz, Mark Rutten, and André van Schaik. Event-based sensing for space situational awareness. *The Journal of the Astronautical Sciences*, 66(2):125–141, June 2019.

[5] Gregory Cohen, Saeed Afshar, and André van Schaik. Approaches for astrometry using event-based sensors. In *Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, page 25, 2018.

[6] Tobi Delbruck, Rui Graca, and Marcin Paluch. Feedback control of event cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1324–1332, 2021.

[7] Thomas Finateu, Atsumi Niwa, Daniel Matolin, Koya Tsuchimoto, Andrea Mascheroni, Etienne Reynaud, Pooria Mostafalu, Frederick Brady, Ludovic Chotard, Florian LeGoff, Hirotsugu Takahashi, Hayato Wakabayashi, Yusuke Oike, and Christoph Posch. A 1280×720 Back-Illuminated stacked temporal contrast Event-Based vision sensor with 4.86µm pixels, 1.066GEPS readout, programmable Event-Rate controller and compressive Data-Formatting pipeline, 2020.

[8] Guillermo Gallego, Tobi Delbruck, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew J Davison, Jorg Conradt, Kostas Daniilidis, and Davide Scaramuzza. Event-based vision: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(1):154–180, January 2022.

[9] R Graça, B McReynolds, and T Delbruck. Shining light on the DVS pixel: A tutorial and discussion about biasing and optimization. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 4045–4053, 2023.

[10] Rui Graca and Tobi Delbruck. Unraveling the paradox of intensity-dependent DVS pixel noise. In *2021 International Image Sensor Workshop (IISW)*, 2021.

[11] Rui Graca, Brian McReynolds, and Tobi Delbruck. Optimal biasing and physical limits of DVS event noise. In *2023 International Image Sensor Workshop (IISW)*, 2023.

[12] Yuhuang Hu, Shih-Chii Liu, and Tobi Delbruck. v2e: From video frames to realistic DVS events. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1312–1321, 2021.

[13] Juan Antonio Leñero-Bardallo, Teresa Serrano-Gotarredona, and Bernabé Linares-Barranco. A 3.6 μ s latency asynchronous Frame-Free Event-Driven Dynamic-Vision-Sensor. *IEEE J. Solid-State Circuits*, 46(6):1443–1455, June 2011.

[14] Patrick Lichtsteiner, Christoph Posch, and Tobi Delbruck. A 128 × 128 120 db 15μs latency asynchronous temporal contrast vision sensor. *IEEE J. Solid-State Circuits*, 43(2):566–576, 2008.

[15] Peter N. McMahon-Crabtree, Lucas Kulesza, Brian J. McReynolds, Daniel S. O'Keefe, Anirvin Puttur, Diana Maestas, Christian P. Morath, and Matthew G. McHarg. Event-based camera refractory period characterization and initial clock drift evaluation. In *Proc. SPIE 12693, Unconventional Imaging, Sensing, and Adaptive Optics*, volume 12693. SPIE, To be published Sept. 2023.

[16] Peter N McMahon-Crabtree and David G Monet. Commercial-off-the-shelf event-based cameras for space surveillance applications. *Appl. Opt., AO*, 60(25):G144–G153, September 2021.

[17] B McReynolds, R Graca, and T Delbruck. Experimental methods to predict dynamic vision sensor event camera performance. *Opt. Eng.*, 2022.

[18] Brian McReynolds, Rui Graca, and Tobi Delbruck. Exploiting alternating DVS shot noise event pair statistics to reduce background activity. In *2023 International Image Sensor Workshop (IISW)*, 2023.

[19] Diederik Paul Moeys. *Analog and digital implementations of retinal processing for robot navigation systems*. PhD thesis, 2016.

[20] Diederik Paul Moeys, Federico Corradi, Chenghan Li, Simeon A Bamford, Luca Longinotti, Fabian F Voigt, Stewart Berry, Gemma Taverni, Fritjof Helmchen, and Tobi Delbruck. A sensitive dynamic and active pixel vision sensor for color or neural imaging applications. *IEEE Trans. Biomed. Circuits Syst.*, 12(1):123–136, February 2018.

[21] Yuji Nozaki and Tobi Delbruck. Temperature and parasitic photocurrent effects in dynamic vision sensors. *IEEE Trans. Electron Devices*, 64(8):3239–3245, August 2017.

[22] Yunjae Suh, Seungnam Choi, Masamichi Ito, Jeongseok Kim, Youngho Lee, Jongseok Seo, Heejae Jung, Dong-Hee Yeo, Seol Namgung, Jongwoo Bong, and Others. A 1280× 960 dynamic vision sensor with a 4.95-μm pixel pitch and motion artifact minimization. In *2020 IEEE international symposium on circuits and systems (ISCAS)*, pages 1–5, 2020.

[23] Minhao Yang, Shih-Chii Liu, and Tobi Delbruck. A dynamic vision sensor with 1% temporal contrast sensitivity and In-Pixel asynchronous delta modulator for event encoding. *IEEE J. Solid-State Circuits*, 50(9):2149–2160, September 2015.

[24] Michał Żołnowski, Rafal Reszelewski, Diederik Paul Moeys, T Delbrück, and Krzysztof Kamiński. Observational evaluation of event cameras performance in optical space surveillance. In *NEO and Debris Detection Conference, Darmstadt, Germany*, 2019.