

# Binocular Telescope for Neuromorphic Space Situational Awareness

**Alexandre Marcireau, Saeed Afshar, Nicholas Owen Ralph, Imogen Jones, and Gregory Cohen**

*International Centre for Neuromorphic Systems, The MARCS Institute for Brain, Behaviour and Development, Western Sydney University*

## ABSTRACT

Detecting the unexpected is essential in Space Situational Awareness (SSA) systems. A robust SSA system must be able to detect and correctly classify uncatalogued debris, spacecraft, and orbital anomalies in space. However, optical and atmospheric artifacts can readily lead to misclassifications. This is because, much like astronomy, SSA lacks strong priors on novel observations.

Neuromorphic event-based sensors excel at detecting fast-moving objects and unexpected changes in luminance, such as glints. Unlike conventional cameras, event-based sensors do not capture frames but detect luminance changes through independent pixels, allowing for an extremely high dynamic range and temporal resolution. These properties enable unique opportunities for space imaging, Resident Space Object (RSO) detection and tracking, high-speed adaptive optics, satellite identification, and real-time in-frame astrometry. With recent advances in sensor sensitivity, quality, and resolution, the neuromorphic vision sensor has become a serious contender in SSA applications.

While event-based space situational awareness provides significant advantages, the event-based sensing paradigm also introduces novel challenges that are not present in conventional frame-based SSA. Fast and faint point sources can be hard to identify amongst spurious change detections generated by other sources, particularly detections arising from insects, bats, and planes. The lack of absolute luminance information in a neuromorphic sensor makes it more challenging to differentiate between RSOs and atmospheric objects when their trajectories are similar from the observer's perspective. Spurious detections are not limited to atmospheric artifacts but may also be due to sensor noise. While recent neuromorphic sensors have dramatically improved noise characteristics compared to older models, it is still desirable to operate them as close to the noise floor as possible to detect ever-fainter objects.

To address these issues, a binocular telescope can be used in conjunction with neuromorphic sensors. By using two cameras and two telescopes attached to the same mount about 1 m apart we are able to capture "stereo event streams" allowing us to detect the presence of parallax and identify atmospheric objects. Additionally, combining neuromorphic cameras with traditional cameras can provide absolute luminance information, which can improve the accuracy of object detection and classification. Our current angular resolution, about one arc second per pixel, allows detection of parallax of objects within the atmosphere. Correlating the sensor's outputs will allow us to filter such objects as well as spurious detections while preserving isolated changes caused by faint satellites. We test several correlation algorithms which leverage the inherent high temporal resolution and the spatial sparsity of Neuromorphic sensors. The presented algorithmic and observational methodologies allow us to distinguish atmospheric objects and sensor noise from genuine detections of uncatalogued objects significantly improving on the current state of the art.

## 1. INTRODUCTION

Tracking and detecting objects that are known is a far easier task than detecting the unknown. The detection of unexpected events is an essential component of a robust space situational awareness system, and is a crucial source of information for characterisation and attribution in space domain awareness. The central question in detecting such anomalies is the question of validating that the discovery is signal rather than noise. This is particularly important in real-world problems, such as tracking, where speed and latency can be critical factors. This is made far easier if




Fig. 1: Two synchronized event-based cameras looking at the same star field using a binocular telescope generate highly correlated signal events but uncorrelated noise events. Leveraging the difference lets us generate a new denoised event stream.

you are able to image continuously with sub-millisecond temporal resolution, but high-speed cameras generate an un-manageable amount of data and can only operate for a few seconds.

Event-based cameras, which are a class of non-traditional optical sensor, provide a fundamentally different approach to this problem, introducing both novel capabilities and different challenges. One of the strengths of event-based cameras is their ability to always be operating at full-speed. They do not require cueing or triggering. Instead, they provide the ability for pervasive and continuous surveillance. They provide the ability to detect the unexpected glint from an otherwise stationary object, or the uncatalogued object passing through the field of view by happenstance.

Event-based cameras arose from the field of neuromorphic engineering, which seeks to build highly efficient and robust systems by leveraging the principles by which biology senses the world and how biological brains processes information. These sensors do not rely on uniform sampling through frames, but rather operate on asynchronous threshold detection. The output of an event-based camera is therefore a stream of contrast change events in which the data rate is wholly dependent on the photometric changes in the scene. This makes them highly efficient when applied to sparse scenes, and therefore extremely well suited to the task of optical space situational awareness.

However, changing the underlying imaging process complicates the analysis of the sensor and its characteristics. Conventional methods for detecting signal, rejecting noise, and characterising the performance of the sensors do not directly apply to event-based cameras. This paper explores the concepts of noise and noise rejection in event-based cameras when applied to the task of space imaging through a top-down approach that leverages ideas and techniques from conventional astronomy.

**Neuromorphic vision: the event-based camera.** The differences between a conventional optical camera and an event-based camera begin within the circuitry within each pixel. Both contain similar photodiodes, but the event-based camera has a complex analogue circuit within each pixel that responds to changes in luminance. There are far more transistors in each pixel of an event-based camera, which greatly limited their fill factor and sensitivity until the arrival of stack-chip solutions like the Prophesee Gen4 event-based camera [1].

The output of an event-based camera is a stream of change events, each representing either positive or negative changes in contrast at a particular pixel. Each event contains a timestamp with microsecond resolution that indicates when the contrast change occurred. These changes are generated asynchronously within in each pixel and transmitted out from the sensor as fast as possible, allowing for high-speed and low-latency systems. The lack of frames results in a data rate that is activity-driven, greatly reducing redundant information such as static and unchanging background. The reduction in data rate allows for far more efficient data processing and storage, allowing for significant power and computation efficiencies throughout a typical sensing pipeline.

Event-based cameras lend themselves particularly well to applications in space imaging [2]. The dark sky background creates extremely sparse scenes, allowing the camera to offer timing resolution comparable to cutting-edge commercial high-speed cameras whilst producing data rates that are orders of magnitude lower. Additionally, motion is ever present

in space imaging applications - either through motion of the telescope, or motion of the objects of interest in the field of view, thereby greatly simplifying the identification and detection of stars and resident space objects.

**Understanding noise in event-based cameras:** The sources of noise in conventional astronomy cameras is relatively well understood. However, understanding noise in event-based cameras represents a more complicated. This is evident by the number of neuromorphic publications on de-noising algorithms, and the relative dearth of publications exploring the underlying nature and characteristics of noise in these sensors. It is only recently that neuromorphic event-based camera simulators began to include models of noise and other non-ideal effects [3].

The main purpose for understanding noise is to capture more signal, which requires a thorough understanding of the characteristics of the sources of noise in the sensing system. Noise in conventional integrating cameras is primarily additive, combining with the true signal to form the output of the sensor. The output of a conventional optical sensor can therefore be split into measurement error and signal. This is neatly characterised by metrics such as signal-to-noise ratio, and can be easily characterised and mitigated. Unfortunately, this is not the case for neuromorphic event-based cameras.

As event-based cameras do not produce absolute measures of illumination, the noise manifests differently. It appears as spurious measurements, in the form of false change events in the data output. These false detections arise from several sources, such as unavoidable transistor mismatch in each pixel, ill-understood temperature effects on the analog feedback loops within each pixel, and a multitude of different leakage currents. Each pixel in the camera operates independently and therefore has its own performance characteristics.

**Binocular telescope setups:** Using two bore-sighted telescopes and two synchronised event-based cameras can help to address the discovery problem. Parallax can be used to calculate the distance of an object and estimate whether it is atmospheric or orbital, while correlations across sensors can provide valuable information to remove noise (figure 1) quickly confirm new detections. The benefits of binocular setups have been shown on multiple occasions with conventional sensors, but have never been used with event-based cameras. The Dragonfly Telescope Array comprises eight frame-based cameras with refracting lenses [4]. The Huntsman Telescope Array is based on the Dragonfly design and principles, and comprises ten cameras with refracting lenses[5]. The Condor Array comprises six cameras, also with refracting lenses [6]. These instruments were built to better resolve low-surface-brightness galaxies, with absolute magnitudes at least one magnitude lower than the ambient night sky [4–6]. The same concept can be applied to improve our ability to detect low-luminance RSOs.

In this work, we describe the development of a binocular telescope mount setup equipped with two event-based camera whose timestamps are synchronized in hardware. We show how this system can be used to calculate the distance of atmospheric objects and evaluate its noise characteristics by comparing star fields recordings with the Gaia DR3 catalogue.

## 2. EXPERIMENTAL SETUP

The data presented in this paper was collected with a specialised observatory built specifically for event-based sensors, called the Astrosite. We used two identical telescopes attached to the same mount, and two identical event-based cameras.

The Astrosite is a mobile observatory that consists of a 2ft standard shipping container with a sliding roof, a scissor lift that supports a motorised mount, and three bore sighted telescopes (figure 2). It is primarily designed for Space Situational Awareness applications, which do not require extremely large mirrors. We used two Officina Stellare RH200 Veloce telescopes mounted on both sides of the center telescope for our experiments. The distance between the scopes' optical axes is 0.685 m and the telescopes have an aperture of 200 mm and a focal length of 600 mm ( $f/3$ ).

The event-based cameras are Prophesee Event Camera Evaluation Kits 4 (EVK4), which use the IMX636 sensor designed by Prophesee and Sony. The sensor array contains  $1280 \times 720$  pixels in a  $1/2.5''$  format, with a pixel size of  $4.86\mu\text{m} \times 4.86\mu\text{m}$ . A plate solve gives a field of view of  $35.6' \times 20.0'$  with a pixel scale of  $1.669 \text{ arcsec/px}$ . The cameras are independently connected over USB to a computer that records their events. They are also connected together with a synchronization cable<sup>1</sup> that ensures that they use the same timestamp reference and tick at the same

<sup>1</sup><https://docs.prophesee.ai/stable/hw/manuals/synchronization.html>




Fig. 2: The Astrosite is a movable observatory that fits in a standard 20ft shipping container. It is specifically designed to operate event-based sensors. The data presented in this paper was collected in Sydney with Astrosite 1. Astrosite 2, also shown on this figure, is currently deployed near Adelaide.



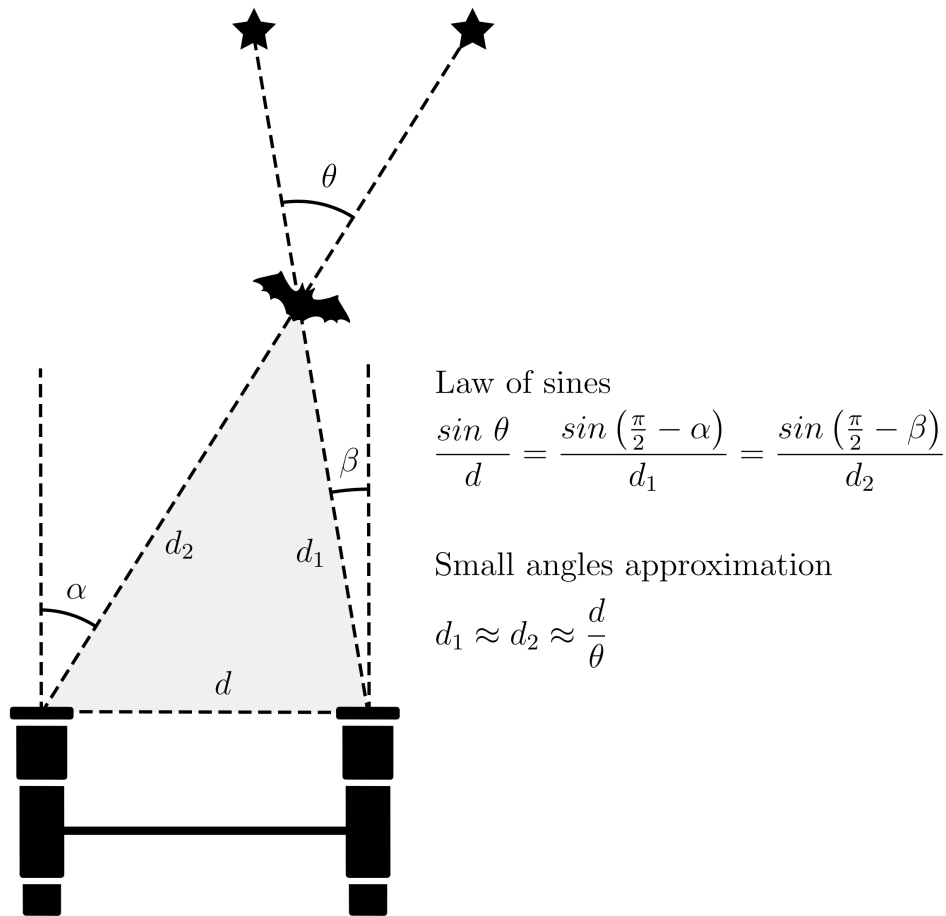


Fig. 3: Simplified model of parallax applied to atmospheric objects. In general, the atmospheric object detected by the sensors is not guaranteed to be in the same plane as the telescopes' optical axes. However, the more general case yields the same equations under the small angles approximation.

rate. Events from the two cameras thus have matching timestamps down to the microsecond and only need to be aligned spatially.

The telescopes' optical axes are aligned within 50arcsec. While mechanical fine-tuning could reduce this error, mechanical flexing in the mounting plate induces alignment errors that depend on the mount's absolute position. It is thus extremely difficult to perfectly align the telescopes in all orientations. We instead use objects at infinity, namely stars or the Moon, to visually align the cameras' fields of view after collecting data. We ignore optical distortions and assume that, given our small field of view, angles are proportional to pixel distances. This allows us to align the images with a rigid transformation. The translation component of that transformation corrects the telescopes' alignment, the rotation component corrects the sensor arrays' rotation, and the scale component accounts for slightly different focus settings. We calculate a different transformation for every field captured with the binocular setup.

The Astrosite was located in Werrington South campus, Western Sydney University when data was collected for the experiments described in this paper. Star recordings were collected on 2023-08-21 and Moon recordings were collected on 2023-08-25.

### 3. ATMOSPHERIC OBJECT IDENTIFICATION FROM PARALLAX

Event-based cameras do not provide a measure of absolute luminance since their pixels only report relative changes. This intentional design choice directly contributes to the high dynamic range of the cameras and their low data rate, while providing enough information for most detection and tracking applications. However, the lack of absolute luminance can complicate the identification of unexpected objects coming across the field of view. Satellites, flying objects such as planes and birds, and optical artifacts can have a very similar appearance when recorded with an event-based camera.

Despite the small distance between the two telescopes, the high angular resolution of our setup yields significant parallax for objects in the atmosphere. A simple geometric model (figure 3) provides a formula to calculate an object's distance from the angular difference between the sight-lines to the object. A difference of one pixel, or 1.669 arcsec, corresponds to a distance of 84.6 km. Hence, only objects that are less than 84.6 km away from the telescopes may have visible parallax. This is enough to calculate the altitude of any atmospheric object in the zenith direction. Commercial planes and clouds have a maximum altitude in the order of 10 km and must thus have an elevation of at least  $6.79^\circ$  above the horizon to have visible parallax (ignoring the Earth's curvature and atmospheric effects).

The minimum measurable angular difference significantly impacts the maximum measurable distance. If noise, atmospheric perturbations, or poor alignment with background stars increase the smallest measurable difference to five pixels instead of one, the maximum altitude becomes 16.9 km and the minimum elevation for planes at 10.0 km becomes  $36.3^\circ$  above the horizon.

We detected an object with visible parallax while staring at the Moon at 8:56 pm local time with the binocular setup (figure 4). The object stayed in the field of view of each camera for about 120 ms but was not visible in both fields of view at the same time, which complicates parallax calculations. Fortunately, the cameras' high temporal resolution allows us to estimate the object's speed and predict its position on the first camera at different instants, and in particular at a time when the object is visible on the second camera. We align the sensors' fields of view by accumulating the events over 4 ms and manually estimating the rigid transformation that best aligns the Moon images.

After temporal propagation and spatial alignment, the pixel distance between the two images of the object is 159 px, which corresponds to an angular difference of 265 arcsec. This places the object 533 m away from the telescopes with a speed of 15.7 km/h. These values seem to indicate a flying animal, such as a bird or a bat. We believe that we detected the latter, given the presence of large bat colonies on Werrington campus and the time of the experiment (three hours and a half after sunset).

This simple example demonstrates that a binocular event-based camera setup can be used to identify atmospheric objects by their parallax. Conversely, the absence of parallax provides a lower bound on the object's distance and can be used to confirm that the object is not in the atmosphere. This is particularly important to confirm the detection of objects that are not listed in an existing catalogue. Parallax detection also illustrates the importance of accurate time synchronization to precisely align event streams.

From a statistical perspective, we are able to detect objects with parallax because the signals generated by objects at infinity are strongly correlated in space and time across cameras (after spatial alignment), but the signals generated by atmospheric objects are not. Interestingly, this applies to camera noise as well.

### 4. BINOCULAR STREAMS MERGING AND NOISE REDUCTION

#### 4.1 Methodology

##### 4.1.1 Event-based camera noise

Event-based cameras report changes in luminance in the form of events. However, different noise sources in the pixel circuit [7] result in spurious events that are not caused by a change in scene luminance. The number and temporal distribution of noise events can be hard to predict as it depends on the intrinsic properties of each pixel, on the absolute scene luminance, and on the bias settings of the camera. Remarkably, increased scene activity causes the pixel to reset more often and reduces the relative number of events caused by pixel memory leaks, thus making noise partially anti-correlated with signal changes. While event-based pixel models exist [3], they predict complex noise distributions

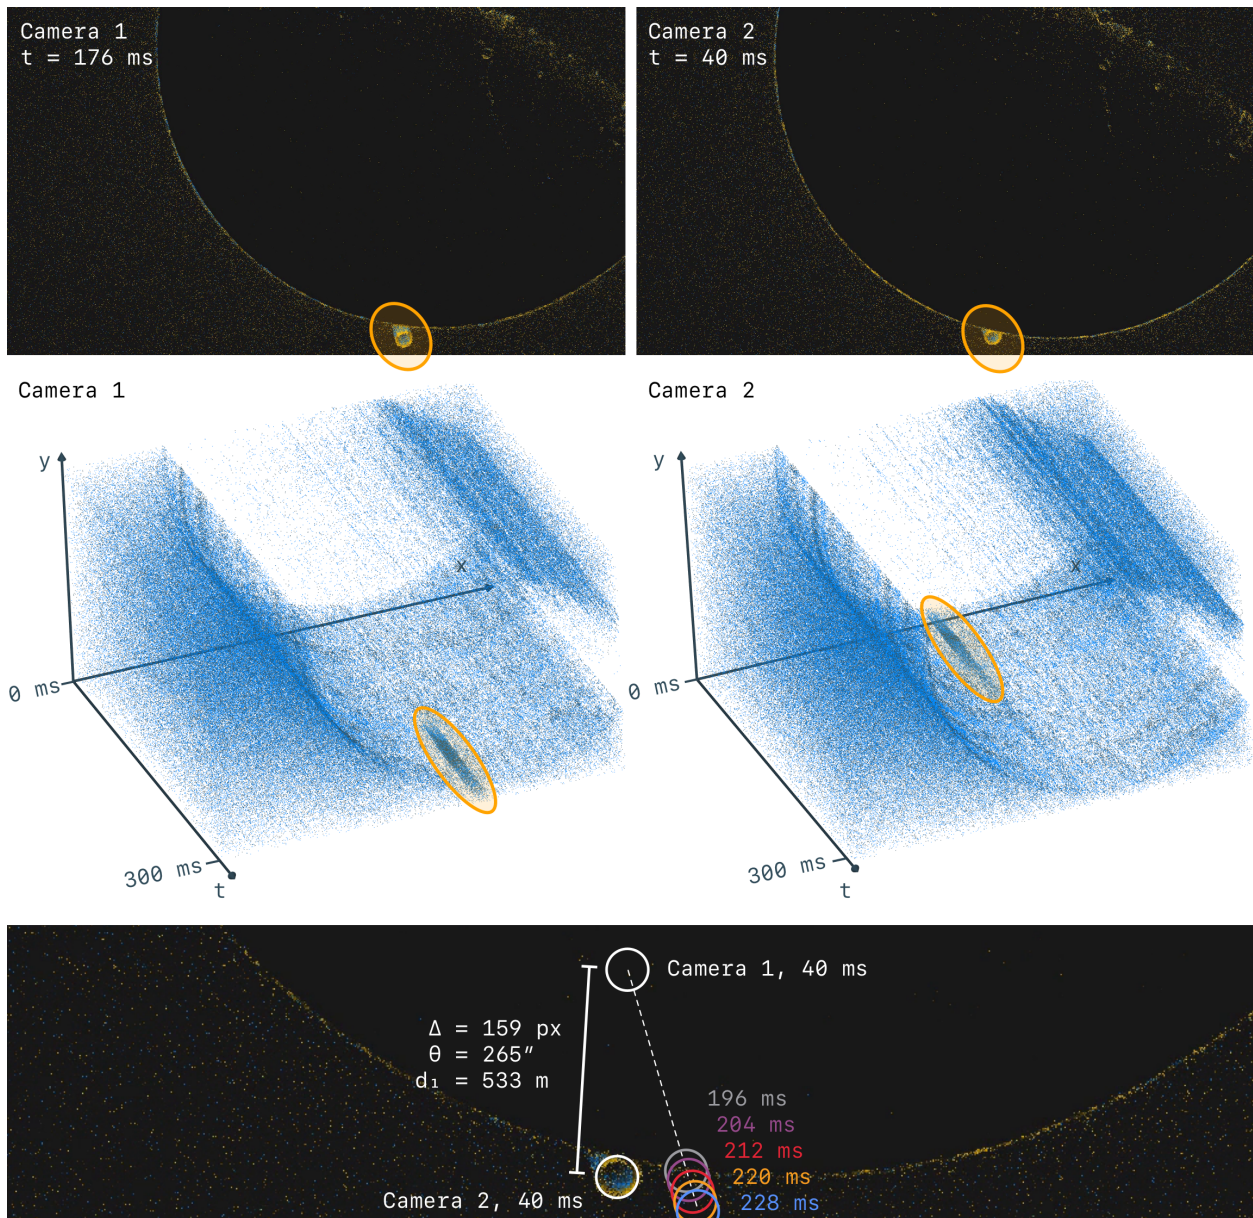


Fig. 4: The top row shows a slice generated from event-based data. Black pixels indicate no recent activity, yellow pixels indicate ON events (increase in luminance) and blue pixels indicate OFF events (decrease in luminance). The object has high parallax but was not visible in both sensors at the same time, as shown by the middle row (event point cloud). The bottom row shows how the high temporal resolution can be used to estimate the object's speed and calculate its position in both fields of view at the same time. The parallax equation places the objects about 500 metres from the telescope, which is compatible with a bird or a bat.

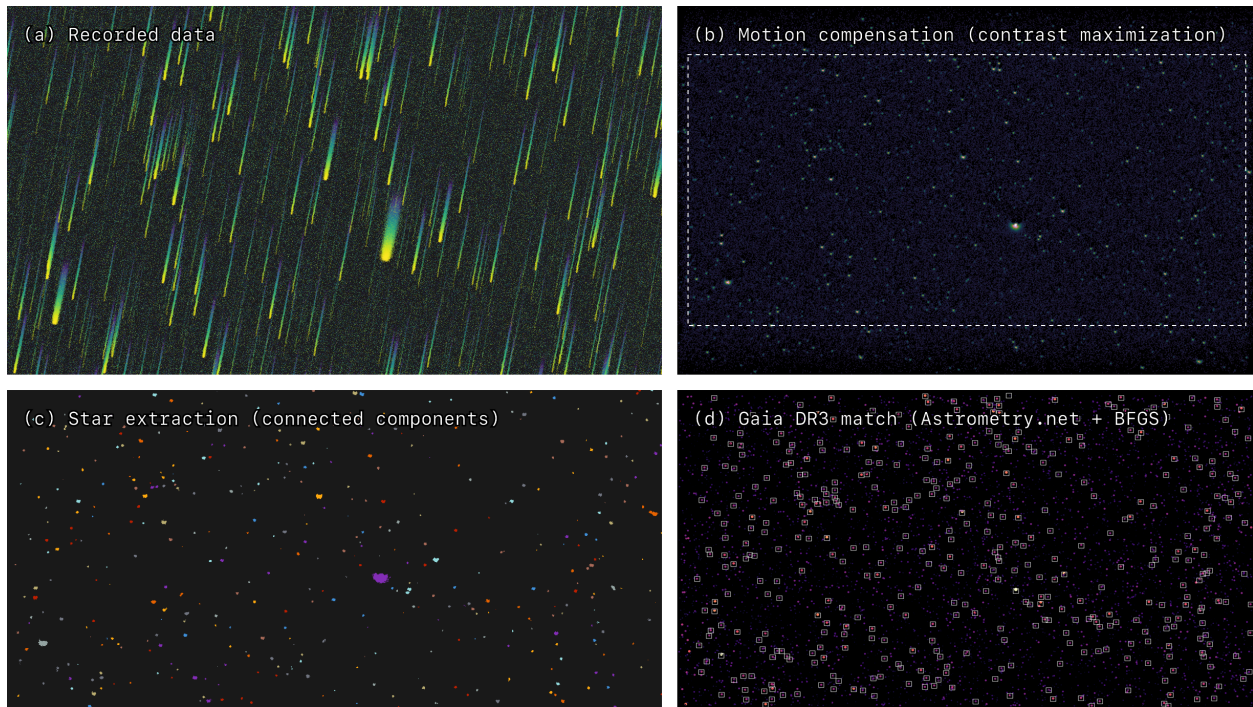


Fig. 5: These four panels illustrate the step used by our pipeline to calculate the difference between extracted stars and catalogue stars. (a) shows ten seconds of recorded data (non-moving mount). The colour gradient represents time. (b) is the accumulated frame generated by compensating global motion with the Contrast Maximization algorithm. (c) uses colours to represent the clusters calculated by the connected components algorithm. (d) represents stars from the Gaia DR3 catalogue. The white squares are the extracted clusters centers.

that are not easy to remove without losing events triggered by faint, but real scene changes in the process.

Since noise sources do not correlate with scene changes, we expect noisy events to be mostly uncorrelated across the two sensors of the binocular setup. Conversely, we expect events caused by scene changes to be generated at similar times and locations. It should thus be possible to filter noise using a simple algorithm that merges two aligned event streams by only selecting events from the first stream if they match events in the second. This empirical approach does not require an advanced pixel noise model and solely relies on signal correlations.

#### 4.1.2 Star matching pipeline

Comparing the noise properties of single camera streams and merged camera streams is challenging. Usual metrics, such as signal-to-noise ratio, are ill-defined for event-based cameras. In the context of Space Situational Awareness, limiting magnitude is easier to define and more directly applicable to practical scenarios [8]. This metric can be estimated empirically by comparing detected star with an existing catalogue. This section describes the pipeline used to process star field recordings. It can be applied to normal event-based camera streams and merged camera streams, and aims to rely as much as possible on simple event-based computer vision and star extraction techniques (figure 5).

Since event-based cameras detect changes, we consider non-sidereally-tracked star fields as a means to generate contrast. This has the added benefit of removing pointing model artifacts and mechanical vibrations from the data. We use Contrast Maximization [9] to estimate global scene motion and warp the events to turn star streaks into points. This operation is equivalent to a shear transformation of the three dimensional  $(x, y, t)$  event point cloud. It converts pixel coordinates into "scene pixels" coordinates. Counting the number of events generated by each "scene pixel" yields an accumulated frame that share many properties with conventional astronomy images. Since the camera is moving relative to the stars, objects near the edge of the accumulated frame spend less time in the field of view than objects near the center. Hence, they generate fewer events and may bias statistical measurements. As we consider relatively



slow and short scenes, we simply crop the edges of the accumulated frame to remove this bias and ensure that all the stars are visible from start to end. Analytical techniques that correct the imbalance near the edges rather than cropping could be considered for longer scenes [10]. We mask the events generated by the most active pixels, also known as *hot* pixels, before using Contrast Maximization. The events reported by such pixel are responsible for a local contrast maximum at a speed of zero that hinders the convergence of the optimizer. We also remove OFF events from the stream since pixel latency is higher for such events. We use the BFGS optimization algorithm [11] to find the speed  $(v_x, v_y)$  that maximizes contrast, where  $v_x$  and  $v_y$  are expressed in pixels per second.

We extract stars from the accumulated image with a technique inspired by robust star extraction for conventional images [12]. We apply a median filter to remove isolated noise events and calculate the  $k$ -th percentile of the pixels' intensities (that is, the event counts of "scene pixels"), for pixels with a non-zero number of events. We use the  $k$ -th percentile to threshold the filtered image and use a connected components algorithm to identify pixel clusters. We define  $\rho = 1 - k/100$  to describe the sensitivity of the star extraction algorithm.  $\rho = 0$  does not select any pixel for clustering while  $\rho = 1$  selects all the pixels with a non-zero event count.

We transfer the pixel labels generated by connected components to the events warped by contrast maximization. This allows us to account for the event count distribution within each cluster and better estimate their center, which we calculate as the center of mass of the events. We also calculate the radius  $r$  of each cluster by approximating the cluster area with the cluster's event count  $c$ , hence  $r = \sqrt{c/\pi}$ .

The list of extracted centers can be directly given to the Astrometry.net algorithm [12] by skipping the star extraction step<sup>2</sup>. We configure the algorithm to estimate a WCS transformation [13] without SIP correction, since our sensor is small and uses only the center of the telescope's light cone. While the algorithm yields a good match between extracted stars and catalogue stars, it uses an index built from Gaia DR2 (epoch 2015.5), it does not account for proper motion, and it only returns index stars rather than all the stars in the field of view.

We use the center of the field of view estimated by Astrometry.net to download stars from the Gaia DR3 catalogue [14, 15]. Previous estimations [16] place the limiting magnitude of Prophesee EVK4 cameras around 14 (with a similar telescope). We download all sources with magnitude 18 and below to make sure that any detectable source is listed in the catalogue. We calculate their position at the time of the recording by using the proper motion provided by Gaia. To minimize the position error between extracted stars and catalogue stars, we optimize the WCS transformation with the following heuristic.

1. Given a WCS transformation, calculate the pixel coordinates of catalogue stars.
2. Initialize an empty list of matches.
3. Calculate pair-wise distances between the  $n$  catalogue stars and the  $m$  extracted stars, where *distance* is defined as the Euclidean distance between the positions in sensor coordinates minus the radius of the extracted star. Store the distances in an  $n \times m$  table.
4. Find the smallest distance in the table, add the corresponding pair to the list of matches, and remove the corresponding row and column from the table.
5. Repeat step 4 until the table is empty.
6. Calculate and return the median distance among matches.

We use BFGS to optimize the six parameters of the WCS transform (sky coordinates reference point, reference pixel, scale, and rotation). We start this optimization process with the WCS transform provided by Astrometry.net, which makes the problem significantly simpler than blind astronomy calibration.

After optimizing the WCS transformation, we match catalogue stars and extracted stars in order of increasing catalogue magnitude. We test whether the position of each star in the catalogue falls within the radius of an extracted star. If it does, we remove the extracted star from the list of unmatched sources and mark the catalogue star as successfully detected (true positive). Otherwise, we count the catalogue star as a missed detection (false negative).

---

<sup>2</sup><https://github.com/neuromorphicsystems/astrometry>



### 4.1.3 Binocular event streams merging

The streams generated by cameras of the binocular setup are readily aligned within 50 arcsec. Like the Moon recording in section 3, we calculate a rigid transformation to fine-tune the streams' spatial alignment. Instead of manually estimating the transformation, we use extracted star positions and an optimization algorithm similar to the WCS optimization process described in the previous section.

After spatially aligning the second stream with the first, we iterate over the events of the first stream. For each event  $(x_1, y_1, t_1)$ , we count the events  $(x_2, y_2, t_2)$  of the second stream such that  $x_1 - w \leq x_2 \leq x_1 + w$ ,  $y_1 - w \leq y_2 \leq y_1 + w$ , and  $t_1 - \tau \leq t_2 \leq t_1$ .  $w$  is the spatial window of the algorithm and  $\tau$  is its temporal window. We only keep events from the first stream if at least one event in the second stream verifies these inequalities.

## 4.2 Results

We collected data from two star fields centered on RA 14h 27' 22" Dec -65° 52.0' 51" (field A) and RA 16h 56' 0" Dec -63° 19.1' 0" (field B). The brightest star in A has magnitude 5.34 and the brightest in B has magnitude 5.98.

Figure 6 (top left) shows the number of matching sources as a function of the magnitude for star field A. We calculate each curve by defining an array of ones and zeroes, with one entry per star in the catalogue. An entry is one if the catalogue star matches an extracted star, and zero otherwise. The plotted curve is a smoothed view of that array, obtained by convolving it with a Hamming window of size 41. The window size was determined empirically and it mostly impacts the overall smoothness of the curve. The number of successful matches can be interpreted as recall, or the probability to generate a detection given the magnitude of a star. For small values of the sensitivity parameter  $\rho$ , the curve is close to a step function with a sharp drop. We thus define limiting magnitude as the point where the recall curve drops below 0.5.

Figure 6 (bottom) shows the limiting magnitude for both star fields as a function of the sensitivity parameter  $\rho$ . The limiting magnitude increases with the sensitivity for small values of  $\rho$ , but it starts decreasing past a certain point. We believe that the decrease is caused by clusters mergers. This causes a decrease in the number of extracted stars and shifts the center of clusters to the wrong location.

We merge binocular streams from field A and field B with the algorithm described section 4.1.3, with  $w = 1$  px and  $\tau = 0.2$  s. Figure 7 shows two small regions of field A recorded by sensor 1, before and after filtering using data from sensor 2. The circles are not extracted stars but sources from the Gaia catalogue (magnitudes up to 15.5). The filtering algorithm removes almost all the background noise while preserving faint features up to magnitude 14.5. The recall curves in figure 6 (top right) are compatible with these observations and become sharper after filtering. However, the bottom figure seems to indicate reduced performance in terms of peak limiting magnitude, although performance depends less on the sensitivity parameter  $\rho$  for merged streams. Even though  $\rho$  controls the sensitivity, it is not normalised across recordings. For this reason, we use a different range of values on single-camera streams and merged streams.

Increasing the sensitivity increases the number of successful matches beyond the limiting magnitude. For instance, a recall value of 0.1 indicates that 10% of stars with a given magnitude were detected. We believe such matches to be false positives caused by the large density of stars in Gaia DR3. Plotting the number of false positives as a function of the limiting magnitude (figure 8) illustrates the cost of increasing the sensitivity. While merged streams slightly outperform single streams on this metric (they yield fewer false positives for a given limiting magnitude), the results are not as telling as the noise visualisations in figure 7.

## 5. CONCLUSION AND DISCUSSION

We presented a binocular telescope setup with two event-based cameras whose timestamps are synchronized in hardware. We showed how this setup can be used to calculate parallax for atmospheric objects. Parallax is critical in telling whether new and unexpected detections can be attributed to an orbital object or may be a bird crossing the field of view. We demonstrated the principle by manually aligning two recordings of the Moon. This allowed us to identifying an atmospheric object which we believe to be a bat. We plan to develop algorithms to automate the parallax detection process for event-based cameras in the future.

Combining binocular streams can also be used to distinguish spurious events (noise) from events caused by actual scene changes (signal), since only the latter are correlated across sensors. While the algorithm introduced to combine

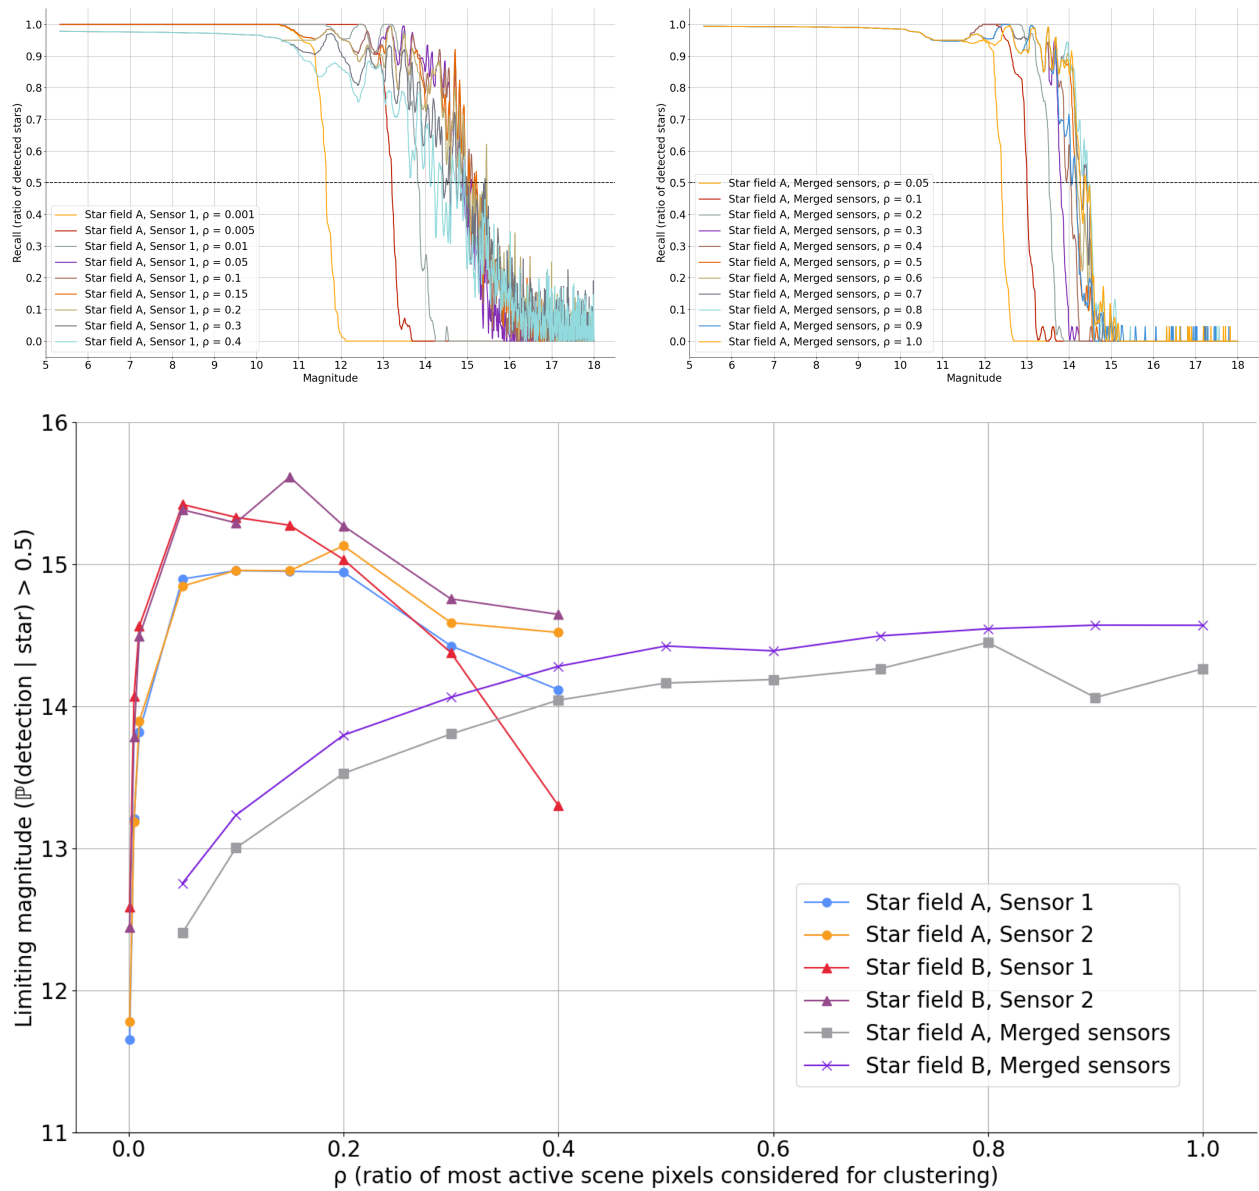


Fig. 6: The top row shows recall curves for different values of the sensitivity parameter  $\rho$ . The top left plot corresponds to a single camera stream and the top right plot corresponds to merged binocular streams. We define the limiting magnitude as the smallest magnitude such that the recall curve is smaller than 0.5. The bottom plot shows the limiting magnitude as a function of the sensitivity parameter for both star fields.




Fig. 7: The top row shows two small regions of the first star field after motion compensation (single camera stream). The circles indicate stars from the Gaia DR3 catalogue (ground truth). The bottom row shows the same regions after combining the binocular streams.

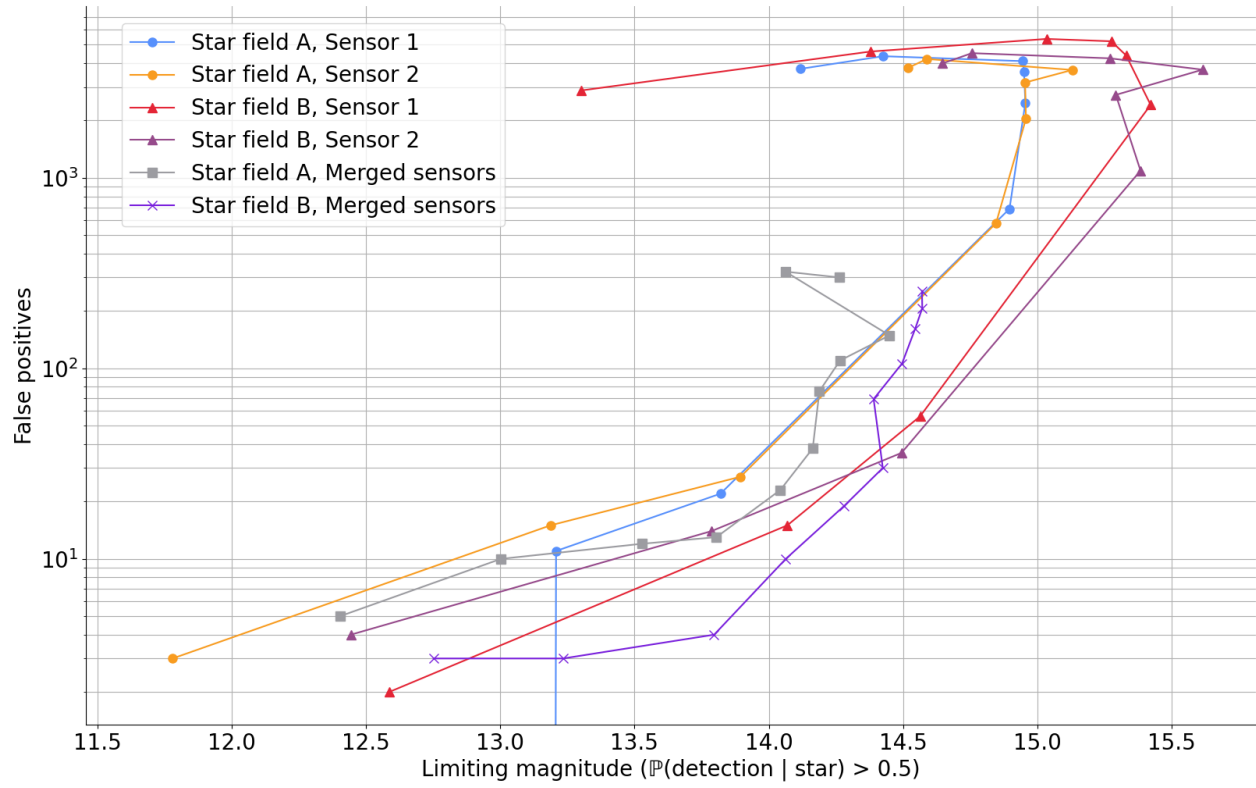


Fig. 8: False positives are extracted stars that do not have a match in Gaia. Increasing the sensitivity increases the limiting magnitude but also yields a larger number of spurious star detections. Merged streams slightly outperform single camera streams on this metric.

binocular event streams is extremely simple, assessing its performance is challenging. Noise models and metrics for event-based cameras often struggle to predict the performance of algorithms on real data. Hence, we used a statistical metric to estimate the limiting magnitude of stars recorded by event-based cameras. This allows us to assess the benefits of binocular systems compared to single sensors.

Combining sensors significantly reduces visual noise but only provides mild statistical benefits. This may indicate that the merging algorithm discards very faint sources, like most single-sensor de-noising algorithms. It is also possible that our statistical metric over-estimates the limiting magnitude of noisy event streams. In particular, the metric assumes that all the pixels have the same limiting magnitude, which is unlikely to be true due to transistor mismatch, and it uses relatively large radii to match extracted stars and catalogue stars. Improving alignment and accounting for pixel variation may be required to build a stronger and more conclusive measure of the limiting magnitude. Alternatively, it may be possible to study these properties by developing more predictive models of event-based pixel noise, combined with physical models of diffraction and atmospheric perturbations.

Importantly, the visual results presented in this paper should not be used to conclude that merging streams is the best way to de-noise event-based data. Indeed, we did not include single stream de-noising techniques in our study. These techniques usually rely on correlations between pixels and need to accumulate evidence before labelling an event as signal or noise. Hence, properly comparing them with a binocular approach requires careful measurements of accuracy and latency for transient and non-transient signals, and we deemed the latter out of the scope of this study. These measurements are particularly important to evaluate the sensitivity of event-based cameras to faint glinting satellites.

Besides distinguishing atmospheric from non-atmospheric objects, binocular event-based setups show promise to calculate the altitude of fast-moving atmospheric objects. This includes planes, drones, flying animals, and atmospheric re-entries. Extending the altitude estimation to orbital objects is possible but it would require a much longer baseline between cameras, in the order of a kilometre rather than a metre. Unfortunately, synchronization cables can only provide reliable signals over a few metres. A potential alternative is the use of GPS receivers, which provide a pulse-per-second clock with sub-microsecond precision anywhere on the planet.

## 6. ACKNOWLEDGEMENTS

This study is based on data obtained at Western Sydney University's Werrington campus. We acknowledge the traditional owners of the land on which the campus stands, the Dharug people, and pay our respects to elders past, present, and emerging.

This work has made use of data from the European Space Agency (ESA) mission *Gaia* (<https://www.cosmos.esa.int/gaia>), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the *Gaia* Multilateral Agreement.

## REFERENCES

- [1] Guillermo Gallego, Tobi Delbruck, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew Davison, Joerg Conradt, Kostas Daniilidis, and Davide Scaramuzza. Event-based Vision: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(1):154–180, January 2022. ISSN 0162-8828, 2160-9292, 1939-3539. doi: 10.1109/TPAMI.2020.3008413. URL <http://arxiv.org/abs/1904.08405>. arXiv:1904.08405 [cs].
- [2] 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: 125–141, 2019.
- [3] Damien Joubert, Alexandre Marcireau, Nic Ralph, Andrew Jolley, André van Schaik, and Gregory Cohen. Event camera simulator improvements via characterized parameters. *Frontiers in Neuroscience*, 15, July 2021. doi: 10.3389/fnins.2021.702765. URL <https://doi.org/10.3389/fnins.2021.702765>.



- [4] Roberto G Abraham and Pieter G van Dokkum. Ultra-low surface brightness imaging with the dragonfly telephoto array. *Publications of the Astronomical Society of the Pacific*, 126(935):55, 2014.
- [5] Lee R. Spitler, Fergus D. Longbottom, Jaime A. Alvarado-Montes, Amir E. Bazkiaei, Sarah E. Caddy, Wilfred T. Gee, Anthony Horton, Steven Lee, and Daniel J. Prole. The huntsman telescope, 2019.
- [6] Kenneth M Lanzetta, Stefan Gromoll, Michael M Shara, Stephen Berg, David Valls-Gabaud, Frederick M Walter, and John K Webb. Introducing the condor array telescope. i. motivation, configuration, and performance. *Publications of the Astronomical Society of the Pacific*, 135(1043):015002, 2023.
- [7] R. Graca, 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, Los Alamitos, CA, USA, jun 2023. IEEE Computer Society. doi: 10.1109/CVPRW59228.2023.00423. URL <https://doi.ieeecomputersociety.org/10.1109/CVPRW59228.2023.00423>.
- [8] Peter N. McMahon-Crabtree and David G. Monet. Commercial-off-the-shelf event-based cameras for space surveillance applications. *Applied Optics*, 60(25):G144, July 2021. doi: 10.1364/ao.425623. URL <https://doi.org/10.1364/ao.425623>.
- [9] Guillermo Gallego, Henri Rebecq, and Davide Scaramuzza. A unifying contrast maximization framework for event cameras, with applications to motion, depth, and optical flow estimation. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, June 2018. doi: 10.1109/cvpr.2018.00407. URL <https://doi.org/10.1109/cvpr.2018.00407>.
- [10] Sami Arja, Alexandre Marcireau, Richard L. Balthazor, Matthew G. McHarg, Saeed Afshar, and Gregory Cohen. Density invariant contrast maximization for neuromorphic earth observations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 3983–3993, June 2023.
- [11] Yu-Hong Dai. Convergence properties of the bfgs algorithm. *SIAM Journal on Optimization*, 13(3):693–701, 2002.
- [12] Dustin Lang, David W. Hogg, Keir Mierle, Michael Blanton, and Sam Roweis. Astrometry.net: Blind astrometric calibration of arbitrary astronomical images. *The Astronomical Journal*, 139(5):1782–1800, Mar 2010. ISSN 1538-3881. doi: 10.1088/0004-6256/139/5/1782. URL <http://dx.doi.org/10.1088/0004-6256/139/5/1782>.
- [13] E. W. Greisen and M. R. Calabretta. Representations of world coordinates in fits. *Astronomy & Astrophysics*, 395(3):1061–1075, Nov 2002. ISSN 1432-0746. doi: 10.1051/0004-6361:20021326. URL <http://dx.doi.org/10.1051/0004-6361:20021326>.
- [14] Gaia Collaboration, T. Prusti, and J.H.J. de Bruijne et al. The gaia mission. *Astronomy & Astrophysics*, 595:A1, nov 2016. doi: 10.1051/0004-6361/201629272. URL <https://doi.org/10.1051/0004-6361/201629272>.
- [15] R. Andrae, M. Fouesneau, and R. Sordo et al. Gaia data release 3. *Astronomy & Astrophysics*, 674:A27, jun 2023. doi: 10.1051/0004-6361/202243462. URL <https://doi.org/10.1051/0004-6361/202243462>.
- [16] Nicholas Owen Ralph, Alexandre Marcireau, Saeed Afshar, Nicholas Tohill, André van Schaik, and Gregory Cohen. Astrometric calibration and source characterisation of the latest generation neuromorphic event-based cameras for space imaging. *Astrodynamics*, Aug 2023. ISSN 2522-0098. doi: 10.1007/s42064-023-0168-2. URL <http://dx.doi.org/10.1007/s42064-023-0168-2>.