

# Event-based Star Tracking via Multiresolution Progressive Hough Transforms

Samya Bagchi

University of Adelaide

samya.bagchi@adelaide.edu.au

Tat-Jun Chin

University of Adelaide

tat-jun.chin@adelaide.edu.au

## Abstract

*Star trackers are state-of-the-art attitude estimation devices which function by recognising and tracking star patterns. Most commercial star trackers use conventional optical sensors. A recent alternative is to use event sensors, which could enable more energy efficient and faster star trackers. However, this demands new algorithms that can efficiently cope with high-speed asynchronous data, and are feasible on resource-constrained computing platforms. To this end, we propose an event-based processing approach for star tracking. Our technique operates on the event stream from a star field, by using multiresolution Hough Transforms to time-progressively integrate event data and produce accurate relative rotations. Optimisation via rotation averaging is then used to fuse the relative rotations and jointly refine the absolute orientations. Our technique is designed to be feasible for asynchronous operation on standard hardware. Moreover, compared to state-of-the-art event-based motion estimation schemes, our technique is much more efficient and accurate.*

## 1. Introduction

On many space missions, it is vital to estimate the attitude of the spacecraft [27], which is the 3DOF orientation (roll, pitch, yaw) of the body frame of the spacecraft w.r.t. an inertial frame, such as the celestial reference frame. The importance of attitude estimation derives from the need to control the bearing of the spacecraft or instruments on board, in order to achieve the mission objectives. Different types of sensors are available for calculating spacecraft attitude, such as sun sensors and magnetometers. It has been established, however, that star trackers are state-of-the-art in spacecraft attitude estimation [25], especially to support high precision orientation determination.

A star tracker is an optical device that estimates spacecraft attitude by recognising and tracking star patterns [27, Chap. 4]. Let  $I$  be an image of a star field captured by the camera of a star tracker. Let  $\mathbf{F}_{ref} \in O(3)$  and  $\mathbf{F} \in O(3)$  respectively define the inertial reference frame, and space-

craft body frame at image  $I$ . For simplicity, we assume calibrated cameras, thus  $\mathbf{F}$  is also the camera frame at  $I$ . The attitude at  $I$  is defined by the rotation matrix  $\mathbf{R}$ , where

$$\mathbf{F}_{ref} = \mathbf{R}\mathbf{F}. \quad (1)$$

In a typical star tracker, the process to determine  $\mathbf{R}$  begins by star identification [41]: matching the observed stars in  $I$  with known stars in a celestial catalogue expressed in  $\mathbf{F}_{ref}$ ; see Fig. 1. The matching can be done by comparing local descriptors [23], geometric voting [22], or sub-graph matching [18]. This establishes a set of 2D-3D correspondences, which are then used to compute  $\mathbf{R}$  via, e.g., SVD [27, Chap. 5] or more robust techniques [9]. In a practical system, a sequence of attitude estimates over time are then jointly optimised (e.g., using EKF with a kinematic model [27, Chap. 6]) to yield a set of refined attitudes.




Figure 1. In star identification, stars in an input image (top) are identified and matched with known stars in a star catalogue (bottom). This result was obtained using the technique of [23, 1].

### 1.1. Event cameras for star tracking

Recently, the feasibility of using event cameras for star tracking has been established [8]. Unlike a conventional optical sensor, an event sensor detects pixel-wise intensity changes asynchronously [24]. Potential benefits of using event cameras for star tracking are lower power consumption and higher speeds [8]. When observing a star field (a




Figure 2. (a) An event stream from a star field recorded using an event camera under ego-motion (event polarities not plotted; data provided by [8]). (b)(c) Motion compensated event images (7) before and after relative rotation estimation by the proposed HT method on the data in (a). Observe that (c) is a sharper star field image than (b), indicating that the proposed method has successfully estimated the ego-motion which generated the event stream. In the proposed relative motion model, the angular velocity in the time period  $\mathbb{T}$  is assumed constant.

scene with relatively few bright points in front of the space void), a vast majority of the pixel locations will not “fire”. This has been illustrated in fig. 2(a) which shows an event stream triggered by a star field. Hence, the event sensor may consume less power than a conventional sensor. An event sensor also has a high temporal resolution (e.g., inVation Davis 240C has  $\mu s$  resolution), which could enable higher-speed star tracking for ultra-fine attitude control.

A potential concern is that the sensor becomes “blind” if the event camera is static w.r.t. the star field. Apart from the orbital motion, a spacecraft (especially a nanosatellite) usually experiences wobbling, which inevitably leads to the generation of event streams.

**Challenges of event-based star tracking** To reap the potential benefits of event sensors for star tracking, several challenges must be met. First, the fundamentally different kind of data requires new attitude estimation algorithms that do not exist in the literature [27]. Second, the high data rate (e.g.,  $\mu s$  resolution) demands very efficient algorithms that can process the event stream, and that are also simple enough to be implemented on resource-constrained computing platforms or basic hardware such as FPGA or SoC.

## 1.2. Our contributions

We propose a novel event algorithm for star tracking. At the core our technique is the usage of a bank of Hough Transforms (HT) [28] at multiple time resolutions to incrementally process events for calculation of relative attitudes. The relative attitudes serve two purposes: track the camera motion, and integrate event streams for star identification. A rotation averaging formulation is then used to jointly optimise the attitude measurements to yield the final attitude estimates. As we will show in Sec. 5, our technique yields much more efficient and accurate star tracking than state-of-the-art event processing methods [16].

Crucially, the part of our algorithm (HT for camera tracking) that processes the high-speed event stream is designed

to be feasible for parallel computation on an FPGA [43]. While refinement by rotation averaging is still required in our pipeline, rotation averaging can be solved using cheap algorithms that conduct just a series of small matrix multiplications [15, Algorithm 1]. As we will show in Sec. 5, rotation averaging incurs only a small percentage of the total computational effort (i.e., 2s over a 45s event stream).

## 1.3. Previous work

Event sensing and processing are increasingly popular in computer and robotic vision [3]. Many important capabilities, such as optic flow computation, panomeric stitching, SfM, and SLAM have been shown to be feasible with event cameras. In fact, in settings where high-speed vision is necessary (e.g., UAV flight [30]), event cameras have proven to be superior than conventional cameras.

The application of event cameras to space engineering problems is nascent. In [10], the feasibility of using event cameras to detect objects in space was established. This was followed by [7], where a probabilistic multiple hypothesis tracker (PMHT) was used to track the objects through time. Our work is inspired by [8] who first proposed event cameras for star tracking. However, their algorithm is completely *synchronous*, in that event streams are converted into event images on which frame-based processing is conducted. Moreover, their method depends on solving relatively costly robust point cloud registration, whose event-driven version for resource-constrained platforms is unknown. In Sec. 5, we will compare the efficacy of our method against [8].

## 2. Event motion model

Over a time period  $\mathbb{T}$ , the output of an event camera is an *event stream*  $\mathcal{S}_{\mathbb{T}} = \{\mathbf{e}_i\}_{i=1}^N$ , where  $\mathbf{e}_i = (\mathbf{x}_i, t_i, p_i)$  is a single event,  $\mathbf{x}_i$  are the 2D pixel coordinates of  $\mathbf{e}_i$  on the image plane,  $t_i \in \mathbb{T}$  is the time when  $\mathbf{e}_i$  occurred, and  $p_i \in \{+, -\}$  is the polarity of  $\mathbf{e}_i$ , i.e., whether it was due to an

increase or decrease in intensity. Since stars are infinitely far away, the ego-motion that can be resolved from  $\mathcal{S}_{\mathbb{T}}$  is restricted to the rotation group, i.e., a continuous rotational motion [27]. For tractability, we first restrict  $\mathbb{T}$  to be small (e.g., 50–500 ms) relative to the angular rate of the camera, such that the motion can be modelled as a single rotation  $\mathbf{R}_{\mathbb{T}}$  called the *relative rotation*. Figs. 2(b) and 2(c) illustrate the feasibility of this model. In Secs. 2.1 and 2.2, we define  $\mathbf{R}_{\mathbb{T}}$  and survey existing methods to estimate it, before Secs. 3 and 4 describe our novel technique and overall star tracking architecture.

## 2.1. Problem formulation

Despite the fundamentally different sensing principle, the pinhole model can be used to describe the imaging process of an event camera [14]. A pixel location  $\mathbf{x}$  of an event  $\mathbf{e}$  can thus be backprojected to form a 3D ray

$$\vec{\mathbf{x}} = \frac{\mathbf{K}^{-1}\bar{\mathbf{x}}}{\|\mathbf{K}^{-1}\bar{\mathbf{x}}\|_2}, \quad (2)$$

where  $\bar{\mathbf{x}} = [\mathbf{x}^T, 1]^T$  is the augmented version of  $\mathbf{x}$ , and  $\mathbf{K}$  is the camera intrinsic matrix. Conceptually, the edge point in 3D space that generated  $\mathbf{e}$  lies along the ray. Existing calibration techniques for event cameras [3, Calibration] can be used to obtain  $\mathbf{K}$ , thus we assume that  $\mathbf{K}$  is known.

Let  $\mathbf{s}^{(t)}$  be the pixel coordinates of a star at time  $t \in \mathbb{T}$ . As mentioned above, by restricting  $\mathbb{T}$  to be small enough, the coordinates of the star at  $t = \alpha$  and  $\beta$  obey

$$\vec{\mathbf{s}}^{(\alpha)} = \mathbf{R}_{\mathbb{T}} \vec{\mathbf{s}}^{(\beta)}. \quad (3)$$

Re-expressing  $\mathbf{R}_{\mathbb{T}}$  using axis-angle representation yields

$$\vec{\mathbf{s}}^{(\alpha)} = \exp(\theta_{\mathbb{T}} \vec{\mathbf{a}}_{\mathbb{T}}) \vec{\mathbf{s}}^{(\beta)} \quad (4)$$

where  $\theta_{\mathbb{T}}$  and  $\vec{\mathbf{a}}_{\mathbb{T}}$  are respectively the angle and axis of  $\mathbf{R}_{\mathbb{T}}$ , and  $\exp$  denotes the exponential map.

Since the non-spurious events are generated by the edges of star blobs, it is reasonable [16] to expect that for each non-spurious event  $\mathbf{e}_i \in \mathcal{S}_{\mathbb{T}}$ , the following holds

$$\vec{\mathbf{x}}_i^{(\alpha)} = \exp(\theta_{\mathbb{T}} \vec{\mathbf{a}}_{\mathbb{T}}) \vec{\mathbf{x}}_i^{(\beta)}, \quad (5)$$

where  $\mathbf{x}_i^{(\alpha)}$  and  $\mathbf{x}_i^{(\beta)}$  are the pixel positions of the edge that generated  $\mathbf{e}_i$  if it was observed at  $t = \alpha$  and  $\beta$ . Further, by assuming that  $\mathbb{T}$  is small enough such that the angular velocity is constant in  $\mathbb{T}$  [16], we have that

$$\vec{\mathbf{x}}_i^{(\alpha)} = \exp\left(\frac{t_i - \alpha}{\beta - \alpha} \theta_{\mathbb{T}} \vec{\mathbf{a}}_{\mathbb{T}}\right) \vec{\mathbf{x}}_i. \quad (6)$$

Again, Fig. 2(c) illustrates the feasibility of assuming constant angular velocity when estimating  $\mathbf{R}_{\mathbb{T}}$  over small  $\mathbb{T}$ .

## 2.2. Previous event processing methods

**Using event images** Following [8], events near the start and end of  $\mathcal{S}_{\mathbb{T}}$ , say  $\{\mathbf{e}_i \in \mathcal{S}_{\mathbb{T}} \mid \alpha \leq t_i \leq \alpha + \Delta\}$  and  $\{\mathbf{e}_i \in \mathcal{S}_{\mathbb{T}} \mid \beta - \Delta \leq t_i \leq \beta\}$  are used to generate two event images  $I_{\alpha}$  and  $I_{\beta}$  [29]. Since the observed stars are points in the images, the rotation  $\mathbf{R}_{\mathbb{T}}$  can be estimated by robustly registering the point clouds [6] in the images.

A disadvantage of [8] is conceptual: converting event streams to images somewhat defeats the purpose of using event sensors. In Sec. 5, we will compare our method against [8].

**Contrast maximisation** The state-of-the-art technique for motion estimation from event streams is contrast maximisation (CM) [16], which readily applies to rotation computation. Unlike previous methods (e.g., [21, 12, 17, 35]), which rely on panoramic map and optic flow computation, CM is more elegant and shown to be superior.

CM finds the motion parameters that maximise the contrast of the “motion-compensated” event image

$$H(\mathbf{x} \mid \theta_{\mathbb{T}}, \vec{\mathbf{a}}_{\mathbb{T}}) = \sum_{i=1}^N p_i \kappa(\mathbf{x} - \mathbf{x}_i^{(\alpha)}) \quad (7)$$

where  $\kappa$  is a smoothing kernel (e.g., Gaussian), and

$$\mathbf{x}_i^{(\alpha)} = \frac{\mathbf{K}^{(1:2)} \vec{\mathbf{x}}_i^{(\alpha)}}{\mathbf{K}^{(3)} \vec{\mathbf{x}}_i^{(\alpha)}} \quad (8)$$

is the projection of  $\vec{\mathbf{x}}_i^{(\alpha)}$ , which is a function of  $\theta_{\mathbb{T}}$ ,  $\vec{\mathbf{a}}_{\mathbb{T}}$  (6), and  $\mathbf{K}^{(1:2)}$  and  $\mathbf{K}^{(3)}$  are the first-2 and 3rd rows of  $\mathbf{K}$ . The contrast of (7) is approximated by the variance  $\sigma^2$  of the image, which is a function of  $\theta_{\mathbb{T}}$ ,  $\vec{\mathbf{a}}_{\mathbb{T}}$ :

$$\sigma^2(\theta_{\mathbb{T}}, \vec{\mathbf{a}}_{\mathbb{T}}) = \frac{1}{D} \sum_{\mathbf{x}} [H(\mathbf{x} \mid \theta_{\mathbb{T}}, \vec{\mathbf{a}}_{\mathbb{T}}) - \mu_H]^2, \quad (9)$$

where  $D$  is the number of pixels in the image, and  $\mu_H$  is the mean  $\frac{1}{D} \sum_{\mathbf{x}} H(\mathbf{x} \mid \theta_{\mathbb{T}}, \vec{\mathbf{a}}_{\mathbb{T}})$ . To estimate  $\mathbf{R}_{\mathbb{T}}$ , we solve

$$\max_{\theta_{\mathbb{T}}, \vec{\mathbf{a}}_{\mathbb{T}}} \sigma^2(\theta_{\mathbb{T}}, \vec{\mathbf{a}}_{\mathbb{T}}) \quad (10)$$

under the constraints  $0 \leq \theta_{\mathbb{T}} \leq \pi$  and  $\|\vec{\mathbf{a}}_{\mathbb{T}}\| = 1$ . Problem (10) is usually solved gradient ascent methods [16].

A downside of contrast maximisation is the relatively complex optimisation algorithm required to solve (10). In the following, we propose a much more efficient technique for star tracking, and compare against [16] in Sec. 5.

## 3. Event-based relative rotation estimation

By ignoring polarity, we convert  $\mathcal{S}_{\mathbb{T}}$  into the 3D point set

$$\mathcal{P}_{\mathbb{T}} = \{\mathbf{z}_i\}_{i=1}^N, \text{ where } \mathbf{z}_i = [\mathbf{x}_i^T, t_i]^T. \quad (11)$$

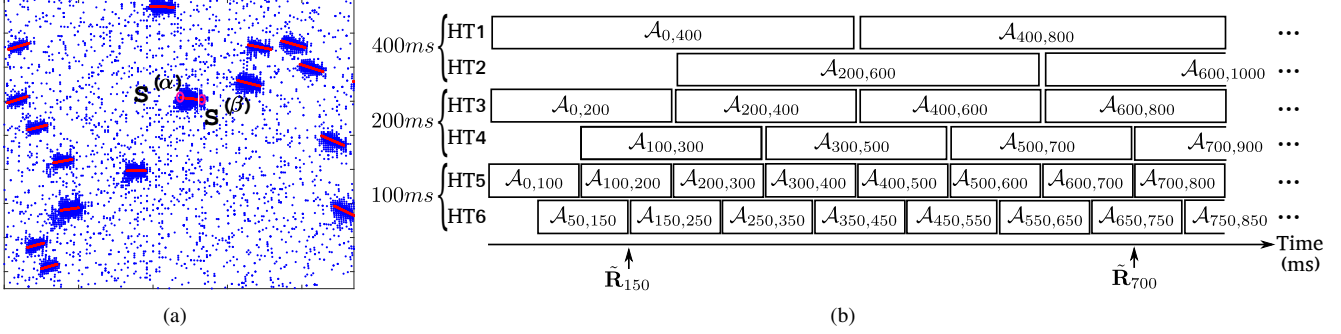


Figure 3. (a) Point correspondences extracted from the lines found in Fig. 4. (b) Overall architecture of our event-based attitude estimation method. Note that in practice the frequency of absolute attitude estimates  $\tilde{\mathbf{R}}_t$  is much lower than suggested by this diagram.




Figure 4. Lines in event stream found by Algorithm 1.

As shown in Fig. 4,  $\mathcal{P}_{\mathbb{T}}$  consists of points that form linear structures, as well as spurious points (gross outliers). Intuitively, in  $\mathcal{S}_{\mathbb{T}}$  from a short time span, the linear structures consist of events that correspond to stars. Our method for estimating  $\mathbf{R}_{\mathbb{T}}$  from  $\mathcal{S}_{\mathbb{T}}$  exploits this observation.

Our method does not use the polarity data, though it might be beneficial to do so—we leave this for future work.

### 3.1. Relative rotation from 3D lines

In the 3D spatio-temporal domain  $\mathbf{z} = [\mathbf{x}^T, t]^T$  obtained by ignoring polarity, a line  $\ell$  can be characterised by

$$\mathbf{z}_{\ell} + \lambda \bar{\mathbf{z}}_{\ell} \equiv \begin{bmatrix} \mathbf{x}_{\ell} \\ t_{\ell} \end{bmatrix} + \lambda \begin{bmatrix} \bar{\mathbf{x}}_{\ell} \\ \bar{t}_{\ell} \end{bmatrix}, \quad (12)$$

where  $\mathbf{z}_{\ell}$  is a point on  $\ell$ ,  $\bar{\mathbf{z}}_{\ell}$  is a unit vector that represents the direction of  $\ell$ , and  $\lambda$  is a scalar. Given noisy points  $\mathcal{Z} \subseteq \mathcal{P}_{\mathbb{T}}$  that belong to  $\ell$ , by orthogonal least squares [40],

$$\mathbf{z}_{\ell} = \frac{1}{|\mathcal{Z}|} \sum_{\mathbf{z}_i \in \mathcal{Z}} \mathbf{z}_i, \quad (13)$$

i.e., the sample mean of  $\mathcal{Z}$ . Let  $\mathbf{Z}$  be the matrix formed by horizontal concatenation of the vectors in  $\mathcal{Z}$ , and

$$\bar{\mathbf{Z}} = \mathbf{Z} - \mathbf{z}_{\ell} \otimes \mathbf{1}_{1 \times |\mathcal{Z}|} \quad (14)$$

be the mean-adjusted matrix. Then,  $\bar{\mathbf{z}}_{\ell}$  is the most significant left singular vector of  $\bar{\mathbf{Z}}$ . In other words, the procedure to estimate  $\ell$  is to perform PCA on  $\mathcal{Z}$  [40].

If  $\ell$  is due to a star, we seek the “end points” of  $\ell$  at time  $\alpha$  and  $\beta$ . Equating  $t_{\ell} + \lambda t'_{\ell}$  with  $\alpha$  and  $\beta$  respectively yields

$$\mathbf{s}^{(\alpha)} = \mathbf{x}_{\ell} + \frac{\alpha - t_{\ell}}{t'_{\ell}} \mathbf{x}'_{\ell}, \quad \mathbf{s}^{(\beta)} = \mathbf{x}_{\ell} + \frac{\beta - t_{\ell}}{t'_{\ell}} \mathbf{x}'_{\ell}. \quad (15)$$

These end points are exactly the image coordinates  $\mathbf{s}^{(\alpha)}$  and  $\mathbf{s}^{(\beta)}$  of the star at time  $\alpha$  and  $\beta$ ; see (3).

If we are able to extract  $K$  linear structures from  $\mathcal{P}_{\mathbb{T}}$ , by the above calculations we yield the point correspondences

$$\{(\mathbf{s}_k^{(\alpha)}, \mathbf{s}_k^{(\beta)})\}_{k=1}^K. \quad (16)$$

Fig. 3(a) illustrates correspondences from stars in  $\mathcal{S}_{\mathbb{T}}$ . The relative rotation  $\mathbf{R}_{\mathbb{T}}$  can then be estimated from such a correspondence set [4, 39]. Specifically, let

$$\mathbf{C} = [\vec{\mathbf{s}}_1^{(\alpha)}, \dots, \vec{\mathbf{s}}_K^{(\alpha)}] [\vec{\mathbf{s}}_1^{(\beta)}, \dots, \vec{\mathbf{s}}_K^{(\beta)}]^T \in \mathbb{R}^{3 \times 3}, \quad (17)$$

and  $\mathbf{C} = \mathbf{U}\mathbf{S}\mathbf{V}^T$  be the SVD of  $\mathbf{C}$ . Then,

$$\mathbf{R}_{\mathbb{T}} = \mathbf{V}\mathbf{U}^T \quad (18)$$

is the least squares estimate of the relative rotation.

### 3.2. Exact incremental HT for relative rotation

Our technique to estimate  $\mathbf{R}_{\mathbb{T}}$  finds the linear structures in  $\mathcal{P}_{\mathbb{T}}$  using HT, from which a set of star correspondences are extracted for estimation of  $\mathbf{R}_{\mathbb{T}}$ ; see Algorithm 1. More details are provided in this section.

Note that although HT has been used previously for event processing [11, 32, 31, 38, 42], we are the first to conduct event-based relative rotation estimation for star tracking. Second, our algorithm is designed to be feasible for asynchronous operation on standard hardware [43]. Specifically, each new event  $\mathbf{e}_i$  triggers an analytically exact update of the intermediate states, such that when all events are processed, matrix  $\mathbf{C} \in \mathbb{R}^{3 \times 3}$  as defined in (17) is complete. Moreover, each event is processed in constant time (per accumulator cell), which supports high-speed tracking.

---

**Algorithm 1** Event-based HT to estimate relative rotation.

---

**Require:** Time period  $\mathbb{T} = [\alpha, \beta]$ , camera intrinsic matrix

$\mathbf{K}$ , Hough cells  $\mathcal{D} \times \mathcal{U} \times \mathcal{V}$ , peak threshold  $\delta$ .

```

1: for each Hough cell  $\tau \in \mathcal{D} \times \mathcal{U} \times \mathcal{V}$  do
2:    $votes^{[\tau]} \leftarrow 0$ ,  $\mathbf{z}_\ell^{[\tau]} \leftarrow \mathbf{0}$ ,  $\bar{\mathbf{z}}_\ell^{[\tau]} \leftarrow \mathbf{0}$ .
3: end for
4:  $\mathbf{C} \leftarrow \mathbf{0}_{3 \times 3}$ .
5: while current time is within  $\mathbb{T}$  do
6:   for each newly arrived event  $\mathbf{e}_i = [\mathbf{x}_i^T, t_i, p_i]^T$  do
7:      $\mathbf{z}_i \leftarrow [\mathbf{x}_i^T, t_i]^T$ .
8:     for each  $\bar{\mathbf{z}} \in \mathcal{D}$  do
9:        $(u, v) \leftarrow$  Project  $\mathbf{z}_i$  onto plane along  $\bar{\mathbf{z}}$  (19).
10:       $[u, v] \leftarrow$  Cell in  $\mathcal{U} \times \mathcal{V}$  that contains  $(u, v)$ .
11:       $\tau \leftarrow [\bar{\mathbf{z}}, u, v]$ .
12:       $votes^{[\tau]} \leftarrow votes^{[\tau]} + 1$ .
13:      if  $votes^{[\tau]} < \delta$  then
14:        UpdatePCA( $\mathbf{z}_\ell^{[\tau]}, \bar{\mathbf{z}}_\ell^{[\tau]}, \mathbf{z}_i$ ) (20)–(25).
15:      else if  $votes^{[\tau]} = \delta$  then
16:        UpdatePCA( $\mathbf{z}_\ell^{[\tau]}, \bar{\mathbf{z}}_\ell^{[\tau]}, \mathbf{z}_i$ ) (20)–(25).
17:         $\mathbf{s}^{(\alpha)}, \mathbf{s}^{(\beta)} \leftarrow$  End points of  $\mathbf{z}_\ell^{[\tau]} + \lambda \bar{\mathbf{z}}_\ell^{[\tau]}$  (15).
18:         $\mathbf{C} \leftarrow \mathbf{C} + \bar{\mathbf{s}}^{(\alpha)}(\bar{\mathbf{s}}^{(\beta)})^T$ .
19:      else
20:         $\mathbf{s}^{(\alpha)}, \mathbf{s}^{(\beta)} \leftarrow$  End points of  $\mathbf{z}_\ell^{[\tau]} + \lambda \bar{\mathbf{z}}_\ell^{[\tau]}$  (15).
21:         $\mathbf{C} \leftarrow \mathbf{C} - \bar{\mathbf{s}}^{(\alpha)}(\bar{\mathbf{s}}^{(\beta)})^T$ .
22:        UpdatePCA( $\mathbf{z}_\ell^{[\tau]}, \bar{\mathbf{z}}_\ell^{[\tau]}, \mathbf{z}_i$ ) (20)–(25).
23:         $\mathbf{s}^{(\alpha)}, \mathbf{s}^{(\beta)} \leftarrow$  End points of  $\mathbf{z}_\ell^{[\tau]} + \lambda \bar{\mathbf{z}}_\ell^{[\tau]}$  (15).
24:         $\mathbf{C} \leftarrow \mathbf{C} + \bar{\mathbf{s}}^{(\alpha)}(\bar{\mathbf{s}}^{(\beta)})^T$ .
25:      end if
26:    end for
27:  end while
28:  $(\mathbf{U}, \mathbf{S}, \mathbf{V}) \leftarrow$  SVD of  $\mathbf{C} \in \mathbb{R}^{3 \times 3}$ .
29: return  $\mathbf{R}_\mathbb{T} \leftarrow \mathbf{V}\mathbf{U}^T$ .

```

---

**Hough domain parametrisation** We follow the Hough parametrisation of [20, 13] for the line  $\mathbf{z} + \lambda \bar{\mathbf{z}}$  in 3D. The line direction  $\bar{\mathbf{z}} = [\bar{z}_1, \bar{z}_2, \bar{z}_3]^T$ , with  $\|\bar{\mathbf{z}}\|_2 = 1$ , is discretised as a set  $\mathcal{D}$  of 1281 vertices of an icosahedron after 4 recursive steps of polygonal subdivision of each triangular mesh (see [13, Fig. 2]), which is sufficient for our problem.

Instead of discretising  $\mathbb{R}^3$  for the point  $\mathbf{z} = [z_1, z_2, z_3]^T$  on the line (which leads to a non-minimal parametrisation), we project  $\mathbf{z}$  along  $\bar{\mathbf{z}}$  onto the plane that passes through the origin that is orthogonal to the line, yielding the 2D point

$$\begin{aligned} u &= \left(1 - \frac{\bar{z}_1^2}{1 + \bar{z}_3}\right) z_1 - \left(\frac{\bar{z}_1 \bar{z}_2}{1 + \bar{z}_3}\right) z_2 - \bar{z}_1 z_3, \\ v &= \left(\frac{\bar{z}_1 \bar{z}_2}{1 + \bar{z}_3}\right) z_1 + \left(1 - \frac{\bar{z}_2^2}{1 + \bar{z}_3}\right) z_2 - \bar{z}_2 z_3; \end{aligned} \quad (19)$$

see [36] for details. The 2D space  $(u, v)$  is then discretised as  $\mathcal{U} \times \mathcal{V}$ . By keeping the duration of  $\mathcal{S}_\mathbb{T}$  constant (e.g., to

100ms) and re-centring  $\mathcal{P}_\mathbb{T}$  such that the centroid is at the origin,  $\mathcal{U} \times \mathcal{V}$  is kept within a fixed bounded region.

**Exact updating** Each new event  $\mathbf{e}_i$  within  $\mathbb{T}$  votes for the discrete set of line parameters in  $\mathcal{D} \times \mathcal{U} \times \mathcal{V}$ . Apart from the usual vote accumulator, in each cell  $\tau \in \mathcal{D} \times \mathcal{U} \times \mathcal{V}$ , we also maintain the *least squares-refined* line parameters  $\mathbf{z}_\ell^{[\tau]}$  and  $\bar{\mathbf{z}}_\ell^{[\tau]}$  that fit the points that voted for  $\tau$ .

The key to exact event-triggered updating is to estimate the refined line parameters incrementally. This can be achieved using incremental PCA [37]. For each  $\mathbf{e}_i$  that voted for  $\tau$ , updating  $\mathbf{z}_\ell^{[\tau]}$  is straightforward:

$$\mathbf{z}_\ell^{[\tau]+} = \frac{votes^{[\tau]} - 1}{votes^{[\tau]}} \mathbf{z}_\ell^{[\tau]-} + \frac{1}{votes^{[\tau]}} \mathbf{z}_i, \quad (20)$$

where  $votes^{[\tau]}$  is the number of votes to  $\tau$  inclusive of  $\mathbf{e}_i$ , and  $\mathbf{z}_\ell^{[\tau]-}$  and  $\mathbf{z}_\ell^{[\tau]+}$  are vector  $\mathbf{z}_\ell^{[\tau]}$  before and after updating.

The trick to update  $\bar{\mathbf{z}}_\ell^{[\tau]}$  is to also maintain the left singular vectors  $\mathbf{P}^{[\tau]} \in \mathbb{R}^{3 \times 3}$  and non-zero singular values  $\Sigma^{[\tau]} \in \mathbb{R}^{3 \times 3}$  of the mean-adjusted matrix (14) of the points that voted for  $\tau$  (recall that  $\bar{\mathbf{z}}_\ell^{[\tau]}$  is the left-most vector of  $\mathbf{P}^{[\tau]}$ ). Given the new  $\mathbf{z}_i$ , compute “difference” matrix

$$\hat{\mathbf{B}} = \left[ \mathbf{0}_{3 \times 1} \quad \sqrt{\frac{votes^{[\tau]} - 1}{votes^{[\tau]}}} (\mathbf{z}_i - \mathbf{z}_\ell^{[\tau]-}) \right] \in \mathbb{R}^{3 \times 2} \quad (21)$$

and its orthogonal projection onto  $\text{span}(\mathbf{P}^{[\tau]})$

$$\tilde{\mathbf{B}} = \mathbf{Q}\mathbf{R} \left( \hat{\mathbf{B}} - \mathbf{P}^{[\tau]}(\mathbf{P}^{[\tau]})^T \hat{\mathbf{B}} \right) \in \mathbb{R}^{3 \times 2}. \quad (22)$$

Then, compose the mean-adjusted singular value matrix

$$\mathbf{E} = \begin{bmatrix} \Sigma^{[\tau]} & (\mathbf{P}^{[\tau]})^T \tilde{\mathbf{B}} \\ \mathbf{0}_{2 \times 3} & \tilde{\mathbf{B}}^T (\hat{\mathbf{B}} - \mathbf{P}^{[\tau]}(\mathbf{P}^{[\tau]})^T \hat{\mathbf{B}}) \end{bmatrix} \in \mathbb{R}^{5 \times 5} \quad (23)$$

and compute its SVD

$$\mathbf{E} = \tilde{\mathbf{P}} \tilde{\Sigma} \tilde{\mathbf{Q}}^T. \quad (24)$$

Then,  $\mathbf{P}^{[\tau]}$  and  $\Sigma^{[\tau]}$  are updated as

$$\mathbf{P}^{[\tau]} = [\mathbf{P}^{[\tau]} \quad \tilde{\mathbf{B}}] \tilde{\mathbf{P}} \quad \text{and} \quad \Sigma^{[\tau]} = \tilde{\Sigma}, \quad (25)$$

and the revised  $\bar{\mathbf{z}}_\ell^{[\tau]}$  is taken as the left-most column of  $\mathbf{P}^{[\tau]}$ . Note that the mean-adjusted matrix  $\tilde{\mathbf{Z}}$  (14) (which grows with the number of events) need not be maintained. For more information of the above procedure, see [37].

In Algorithm 1, if the number of votes in a cell exceeds the pre-determined threshold  $\delta$ , the least squares-fitted line of the cell is used to extract a star correspondence, following (15). However, instead of simply collecting the star correspondences, to facilitate asynchronous operation, the “covariance” matrix (17) is directly updated. Note that

$$\mathbf{C} = \sum_{k=1}^K \bar{\mathbf{s}}_k^{(\alpha)} (\bar{\mathbf{s}}_k^{(\beta)})^T, \quad (26)$$

hence, the contribution of a new correspondence  $(\mathbf{s}^{(\alpha)}, \mathbf{s}^{(\beta)})$  to  $\mathbf{C}$  can simply be introduced by adding the outer product  $\vec{\mathbf{s}}^{(\alpha)}(\vec{\mathbf{s}}^{(\beta)})^T$  to  $\mathbf{C}$ . If the correspondence from the pre-update line of the cell has contributed to  $\mathbf{C}$ , the outer product term due to the old correspondence is first subtracted.

As can be seen in Algorithm 1, as soon as all events within  $\mathbb{T}$  are processed, matrix  $\mathbf{C}$  is complete and the rotation  $\mathbf{R}_{\mathbb{T}}$  can directly be extracted from  $\mathbf{C}$  via SVD.

**Complexity analysis** For each incoming event  $\mathbf{e}_i$ , at most  $|\mathcal{D}|$  Hough cells ( $|\mathcal{D}| = 1281$  in our experiments) are visited (see loop starting in Step 8 in Algorithm 1). As is standard in hardware implementation of HT [43], Hough cell voting and its associated operations can be parallelised. Also note that the SVDs can be computed efficiently on FPGA [26].

For the per-event computations to be constant time, it remains to show that the least squares line and covariance matrix updating are constant time. The main cost in the former requires the QR decomposition of a  $3 \times 2$  matrix  $\tilde{\mathbf{B}}$  (22) and SVD of a  $5 \times 5$  matrix  $\mathbf{E}$  (24). Note that although (25) seems to increase the sizes of  $\mathbf{P}^{[\tau]}$  and  $\Sigma^{[\tau]}$ , only the top-left  $3 \times 3$  submatrix of the updated results are meaningful for data in 3D space, hence, the sizes of  $\mathbf{P}^{[\tau]}$  and  $\Sigma^{[\tau]}$  can be kept constant at  $3 \times 3$ . Lastly, updating the  $3 \times 3$  covariance matrix  $\mathbf{C}$  by adding/subtracting the outer product of a single operation is clearly constant time. Note that there are also no “racing” issues with updating  $\mathbf{C}$  in parallel, since each cell contributes to  $\mathbf{C}$  independently.

#### 4. Overall architecture for attitude estimation

Fig. 3(b) shows the overall architecture of our event-based attitude estimation technique. Symbol  $\mathcal{A}_{\alpha,\beta}$  indicates the instance of the proposed HT (Algorithm 1) for the time period  $[\alpha, \beta]$ . We use a bank of HTs at multiple time resolutions (specifically,  $400ms$ ,  $200ms$ ,  $100ms$ ; these can be changed to accommodate data with different angular rates) to track rotational motions over longer time periods.

For each time resolution, two HT instances separated by a stride half of the resolution (e.g., for  $400ms$ , the two HTs differ by a stride of  $200ms$ ) are employed, leading to a total of six HTs. Each HT processes incoming events asynchronously (Sec. 3). When the period of a HT is finished, a relative rotation over that period is returned (e.g.,  $\mathcal{A}_{100,300}$  outputs the relative rotation  $\mathbf{R}_{100,300}$ ), and the HT is immediately restarted for the next period (e.g.,  $\mathcal{A}_{300,500}$ ).

A relative rotation describes the relative attitude change across a time period. In star tracking, however, the quantity of interest is the “absolute” attitude [27] (see Sec. 1). Moreover, simply chaining the relative attitudes will lead to drift errors. Thus, in our pipeline, motion-compensated event images (7) are generated using the HT results (we use the  $100ms$  HTs for this purpose) and subject to star identification and absolute attitude estimation using [23, 1]. The

frequency of absolute attitude estimation is much lower, due to the higher cost of star identification, e.g., minutes.

The small set of absolute attitudes are used to “ground” the relative rotations via a rotation averaging process, which also denoises the measured quantities and produce the final set of refined (absolute) attitude estimates. Let

$$\{\tilde{\mathbf{R}}_{\alpha,\beta}\}_{\langle\alpha,\beta\rangle \in \mathcal{N}} \quad (27)$$

be the set of measured relative attitudes, where  $\mathcal{N}$  encodes the periods  $\mathbb{T} = [\alpha, \beta]$  processed by a HT instance, and let

$$\{\tilde{\mathbf{R}}_{\gamma}\}_{\gamma \in \mathcal{M}} \quad (28)$$

be the set of absolute attitude measurements. Note that  $\mathcal{N}$  must be a connected graph, in that it is always possible to find a path between any two “time stamps”  $\alpha$  and  $\beta$ . Define

$$\mathcal{T} = \{0, \Delta t, 2\Delta t, 3\Delta t, \dots\} \quad (29)$$

be the time stamps of the set of absolute attitudes that will be optimised by rotation averaging, where the value of  $\Delta t$  is a common denominator of the time resolutions of the HTs employed. For the architecture in Fig. 3(b), we set  $\Delta t = 50ms$ , which also ensures  $\mathcal{M} \subset \mathcal{T}$ .

Following [8], we formulate rotation averaging as

$$\begin{aligned} \min_{\mathbf{R}_{\beta} \mathbf{R}_{\mathcal{G}}} \quad & \sum_{\langle\alpha,\beta\rangle \in \mathcal{N}} \left\| \mathbf{R}_{\alpha} - \tilde{\mathbf{R}}_{\alpha,\beta} \mathbf{R}_{\beta} \right\|_F \\ & + \alpha \sum_{\mathbf{R}_{\gamma} \in \mathcal{M}} \left\| \mathbf{R}_{\gamma} - \tilde{\mathbf{R}}_{\gamma} \mathbf{R}_{\mathcal{G}} \right\|_F \end{aligned} \quad (30)$$

subject to  $\mathbf{R}_{\mathcal{G}} = \mathbf{I}$ ,

where  $\mathbf{R}_{\mathcal{G}}$  is a “dummy” attitude variable,  $\mathbf{I}$  is the identity matrix, and  $\alpha$  is a positive constant that defines the relative importance of the relative and absolute rotations. Intuitively, adding error terms of the form

$$\left\| \mathbf{R}_{\gamma} - \tilde{\mathbf{R}}_{\gamma} \mathbf{R}_{\mathcal{G}} \right\|_F = \left\| \mathbf{R}_{\gamma} - \tilde{\mathbf{R}}_{\gamma} \right\|_F, \quad \gamma \in \mathcal{M} \quad (31)$$

encourage consistency between some of the attitude estimates and the measured absolute rotations, which is then propagated to the rest of the sequence.

To solve (30), we temporarily ignore the constraint  $\mathbf{R}_{\mathcal{G}} = \mathbf{I}$  optimise the attitudes using an existing rotation averaging algorithm (we used [5] in our work). Then, right multiply each of the optimised absolute attitudes  $\hat{\mathbf{R}}_t$  with  $\hat{\mathbf{R}}_{\gamma}^{-1} = (\hat{\mathbf{R}}_{\gamma})^T$  to re-orient the system. It has been shown that rotation averaging is quite insensitive to initialisations [33, 15], thus, when solving (30) we simply initialise all rotation variables as the identity matrix.

In a practical system, instead of solving ever-growing instances of (30), rotation averaging can be executed over a fixed-sized temporal sliding window. However, in this case, we solved the rotation averaging for the entire sequence of 45 seconds.



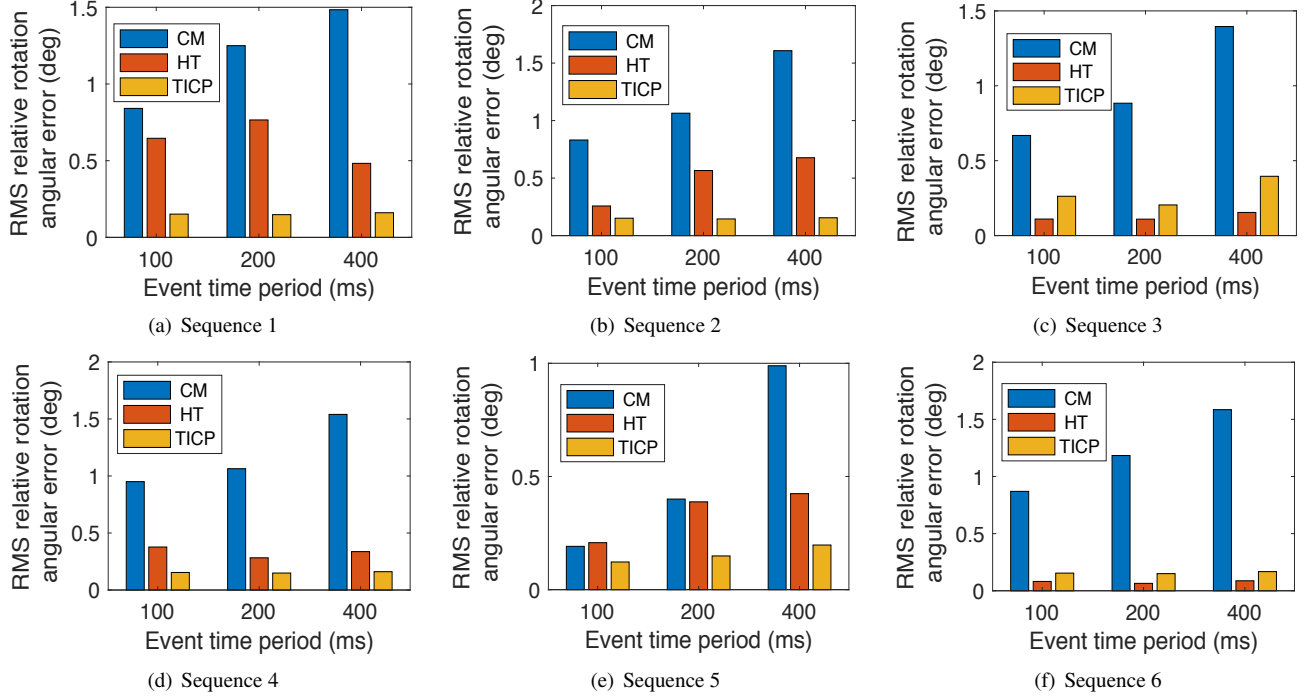


Figure 5. Comparison of relative rotation estimation accuracy using three methods (HT, TICP and CM) on 6 event streams/sequences.

## 5. Results

To evaluate the proposed event-based star tracking technique, we used the event data from [8, 2], where there are 11 event streams from observing star fields using an iniVation Davis 240C event camera (see details in [8] on data generation and event camera calibration). Each event stream contains ground truth absolute attitudes across the full stream.

Our evaluation focussed on the accuracy and runtime of relative attitude estimation and absolute attitude estimation. For both experiments, we compared against:

- The baseline method of [8] conducts robust registration using a variant of ICP called trimmed ICP [6] (henceforth referred as “TICP”) on event images to estimate relative rotations.
- The state-of-the-art event-based motion estimation algorithm of contrast maximisation [16] (henceforth, “CM”). See Sec. 2.2 for a summary of these methods. Due to space constraints, we were able to fit the results of only 6 event streams; see supplementary material for more results.

To compare estimated  $\tilde{\mathbf{R}}$  and ground truth  $\mathbf{R}^*$  rotations in our evaluation, we used the angular distance [19]

$$\angle(\tilde{\mathbf{R}}, \mathbf{R}^*) = 2 \arcsin \left( \frac{2}{\sqrt{2}} \|\tilde{\mathbf{R}} - \mathbf{R}^*\|_F \right). \quad (32)$$

### 5.1. Accuracy of relative attitude estimation

For each event stream, as per the multi-resolution architecture in Fig. 3(b), we estimated relative rotations from

event “chunks” using HT, TICP and CM. Since each stream is about 45s long, there are 224 chunks of 400ms long, 449 chunks of 200ms long, and 899 chunks of 100ms long.

To objectively test the efficacy of the algorithms, we executed them on a common platform: Matlab on an Intel i7 2.7 GHz machine, and using `fmincon` to solve CM (10). Since our focus here is relative rotation accuracy, it was not necessary to execute HT in a streaming fashion.

To assess the accuracy of a relative rotation estimate  $\tilde{\mathbf{R}}_{\alpha,\beta}$  from a time period  $\mathbb{T} = [\alpha, \beta]$ , we computed the angular distance of  $\tilde{\mathbf{R}}_{\alpha,\beta}$  to the ground truth relative rotation

$$\mathbf{R}_{\alpha,\beta}^* = \mathbf{R}_\alpha^* (\mathbf{R}_\beta^*)^T, \quad (33)$$

where  $\mathbf{R}_\alpha^*$  and  $\mathbf{R}_\beta^*$  ground truth absolute attitudes. Fig. 5 shows the root mean square (RMS) angular distance

$$\sqrt{\sum_{\langle \alpha, \beta \rangle} \angle(\tilde{\mathbf{R}}_{\alpha,\beta}, \mathbf{R}_{\alpha,\beta}^*)^2} \quad (34)$$

separated according to the length of the time periods  $\mathbb{T}$  (200, 300 and 400 ms). The RMS error of HT is conclusively lower than CM on most of the input cases, followed by TICP which has lower error than HT in 4 of the 6 sequences.

It should be reminded, however, that TICP does not intrinsically support asynchronous processing of event streams; moreover, it depends on much more complex routines (sorting, nearest neighbour search, alternating optimisation, etc.). In contrast, HT is designed to be simple

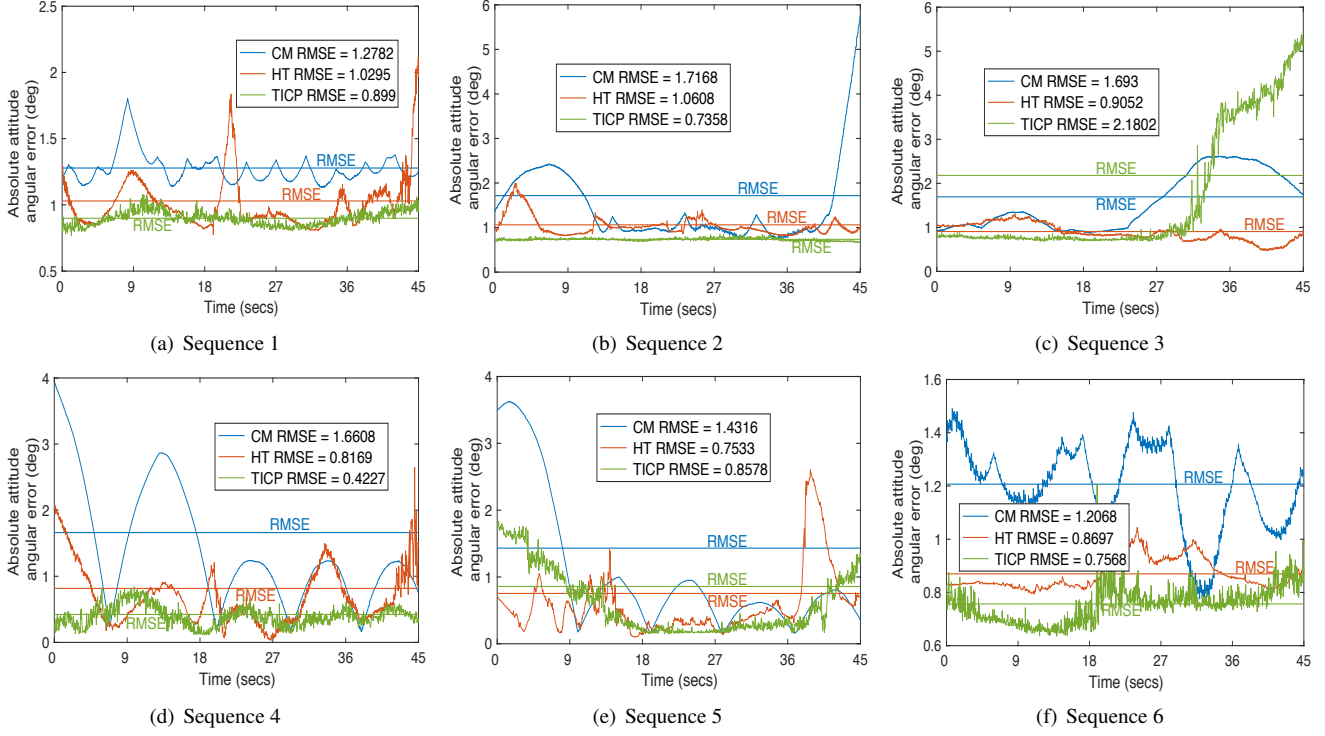


Figure 6. Comparison of absolute attitude estimation accuracy using three methods (HT, TICP and CM) on 6 event streams/sequences.

$ T $ (ms)	Me- thod	1	2	3	4	5	6
100	HT	1.23	1.85	1.49	2.11	1.31	1.67
	CM	16.46	28.31	21.89	33.15	19.43	24.36
200	HT	1.88	2.39	2.01	2.14	1.93	2.09
	CM	45.68	61.45	53.46	65.78	48.91	61.24
400	HT	2.51	2.98	2.67	3.11	2.63	2.83
	CM	82.45	114.78	103.56	118.32	89.12	109.25

Table 1. Runtime (in seconds) of event-based processing methods (HT, CM) for relative rotation estimation on six event streams.

enough to be parallelised on standard hardware (e.g., FPGA and SoC). Thus, it is not surprising that TICP can be more accurate than HT. In any case, as we show later, the accuracy of the final attitude estimates of HT is very competitive.

Table 1 presents the average runtime of the event-processing methods to compute relative rotations. Note that as the time period increases, the runtime of CM increases much more rapidly than HT. While a better solver than *fmincon* can potentially speed up CM, even with an order of magnitude speedup, HT is still faster than CM. Note that the runtime of TICP does not change with the size of the event chunk, since it always aligns two event images; due to this different computational paradigm, the runtime of TICP is not shown. The runtime reported for HT can be further optimised with an implementation that exploits the inherent parallelism of the algorithm. The cost of rotation

averaging is around 2 seconds per sequence.

## 5.2. Accuracy of absolute attitude estimation

Fig. 6 plots the “raw” and RMS angular distances between ground truth and estimated absolute attitudes, based on performing the proposed rotation averaging technique (Sec. 4) on the relative rotations (computed by HT, TICP and CM) and a small set of “grounding” absolute attitudes (28) (5 per event stream, distributed uniformly across 45 seconds, and computed using [23, 1] on the motion compensated event image by HT on 100 ms chunks). As can be seen, the accuracy of the optimised absolute attitudes using HT satisfies the requirements of a commercial star tracker [34] ( $1^\circ$  RMS angular error).

Towards the end of Sequence 3, the error of TICP increases drastically unlike the proposed HT method. This was because there were much fewer stars in the FOV in that period.

## 6. Conclusion

In this paper, we proposed an event-based star tracking algorithm using a novel multiresolution asynchronous HT technique to accurately and efficiently estimate relative rotations on event streams. Our results show that our technique is superior to existing event-based processing schemes for motion estimation with its accuracy sufficiently fulfilling the requirements of a commercial star tracker.



## References

- [1] <http://nova.astrometry.net/>.
- [2] <https://cs.adelaide.edu.au/~tjchin/startracking/>.
- [3] Event-based vision resources. [https://github.com/uzh-rpg/event-based\\_vision\\_resources](https://github.com/uzh-rpg/event-based_vision_resources).
- [4] K. S. Arun, T. S. Huang, and S. D. Blostein. Least-squares fitting of two 3-d point sets. *IEEE TPAMI*, 9(5):698–700, 1987.
- [5] A. Chatterjee and V. Madhav Govindu. Efficient and robust large-scale rotation averaging. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 521–528, 2013.
- [6] D. Chetverikov, D. Svirko, D. Stepanov, and P. Krsek. The trimmed iterative closest point algorithm. In *International Conference on Pattern Recognition (ICPR)*, 2002.
- [7] B. Cheung, M. Rutten, S. Davey, G. Cohen, and N. Gordon. Probabilistic multi hypothesis tracker for an event based sensor. In *International Conference on Information Fusion (FUSION)*, 2018.
- [8] T.-J. Chin, S. Bagchi, A. Eriksson, and A. van Schaik. Star tracking using an event camera. *arXiv:1812.02895*, 2018.
- [9] T.-J. Chin and M. Tetlow. Robust attitude estimation to support space monitoring using nano-satellites. In *AIAA Space Forum*, 2014.
- [10] G. Cohen, S. Afshar, A. van Schaik, A. Wabnitz, T. Bessell, M. Rutten, and B. Morreale. Event-based sensing for space situational awareness. In *Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, 2017.
- [11] J. Conradt, M. Cook, R. Berner, P. Lichtsteiner, R. J. Douglas, and T. Delbruck. A pencil balancing robot using a pair of AER dynamic vision sensors. In *IEEE International Symposium on Circuits and Systems*, 2009.
- [12] M. Cook, L. Gugelmann, F. Jug, C. Krautz, and A. Steger. Interacting maps for fast visual interpretation. In *International Joint Conference on Neural Networks (IJCNN)*, 2011.
- [13] C. Dalitz, T. Schramke, and M. Jeltsch. Iterative Hough transform for line detection in 3D point clouds. *Image Processing On Line*, 7:184–196, 2017.
- [14] T. Delbrück, B. Linares-Barranco, E. Culurciello, and C. Posch. Activity-driven, event-based vision sensors. In *Proceedings of 2010 IEEE International Symposium on Circuits and Systems*, pages 2426–2429, 2010.
- [15] A. Eriksson, C. Olsson, F. Kahl, and T.-J. Chin. Rotation averaging and strong duality. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [16] G. Gallego, H. Rebecq, and D. Scaramuzza. A unifying contrast maximization framework for event cameras, with applications to motion, depth, and optical flow estimation. In *Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [17] G. Gallego and D. Scaramuzza. Accurate angular velocity estimation with an event camera. *IEEE Robot. Autom. Lett.*, 2:632–639, 2017.
- [18] G. Galvizo and L. Lim. Empirical analysis of common subgraph isomorphism approaches to the lost-in-space star identification problem. *arXiv:1808.08686*, 2018.
- [19] R. Hartley, J. Trumpf, Y. Dai, and H. Li. Rotation averaging. *International Journal of Computer Vision (IJCV)*, 103(3):267–305, 2013.
- [20] M. Jeltsch, C. Dalitz, and R. Pohle-Frölich. Hough parameter space regularisation for line detection in 3d. In *VISAPP*, 2016.
- [21] H. Kim, A. Handa, R. Benosman, S.-H. Ieng, and A. Davison. Simultaneous mosaicing and tracking with an event camera. In *British Machine Vision Conference (BMVC)*, 2014.
- [22] M. Kolomenkin, S. Pollak, I. Shimshoni, and M. Lindenbaum. Geometric voting algorithm for star trackers. *IEEE Transactions on Aerospace and Electronic Systems*, 44(2):441–456, 2008.
- [23] D. Lang, D. W. Hogg, K. Mierle, M. Blanton, and S. Roweis. Astrometry.net: blind astrometric calibration of arbitrary astronomical images. *The Astronomical Journal*, 139:1782–1800, 2010.
- [24] P. Lichtsteiner, C. Posch, and T. Delbruck. A 128×128 120 dB 15μs latency asynchronous temporal contrast vision sensor. In *IEEE Journal of Solid-State Circuits*, volume 43, pages 566–576, 2008.
- [25] C. C. Liebe. Accuracy performance of star trackers - a tutorial. *IEEE Trans. Aero. Electron. Syst.*, 38(2):587–599, 2002.
- [26] W. Ma, M. Kaye, D. Luke, and R. Doraiswami. An fpga-based singular value decomposition processor. In *2006 Canadian Conference on Electrical and Computer Engineering*, pages 1047–1050. IEEE, 2006.
- [27] F. L. Markley and J. L. Crassidis. *Fundamentals of Spacecraft Attitude Determination and Control*, volume 33 of *Space Technology Library*. Springer-Verlag New York, 2014.
- [28] J. Matas, C. Galambos, and J. Kittler. Progressive probabilistic Hough transform. In *BMVC*, 1998.
- [29] E. Mueggler, C. Forster, N. Baumli, G. Gallego, and D. Scaramuzza. Lifetime estimation of events from dynamic vision sensors. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2015.
- [30] E. Mueggler, B. Huber, and D. Scaramuzza. Event-based, 6-DOF pose tracking for high-speed maneuvers. In *IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, 2014.
- [31] Z. Ni. *Asynchronous event based vision: algorithms and applications to microrobotics*. PhD thesis, Université Pierre et Marie Curie, 2013.
- [32] Z. Ni, C. Pacoret, R. Benosman, S. Ieng, and S. Régnier. Asynchronous eventbased high speed vision for microparticle tracking. *Journal of Microscopy*, 245(3):236–244, 2012.
- [33] C. Olsson and O. Enqvist. Stable structure from motion for unordered image collections. In *Scandinavian Conference on Image Analysis (SCIA)*, 2011.
- [34] R. Oromolla, G. Fasano, G. Rufino, M. Grassi, C. Pernechele, and C. Dionisio. A new star tracker concept for satellite attitude determination based on a multi-purpose panoramic camera. *Acta Astronautica*, 140:166–175, 2017.
- [35] C. Reinbacher, G. Munda, and T. Pock. Real-time panoramic tracking for event cameras. In *ICCP*, 2017.
- [36] K. S. Roberts. A new representation for a line. In *CVPR*, 1988.

- [37] D. A. Ross, J. Lim, R.-S. Lin, and M.-H. Yang. Incremental learning for robust visual tracking. *IJCV*, 77(1–3):125–141, 2008.
- [38] S. Seifozzakerini, W.-Y. Yau, B. Zhao, and K. Mao. Event-based Hough transform in a spiking neural network for multiple line detection and tracking using a dynamic vision sensor. In *BMVC*, 2016.
- [39] O. Sorkine-Hornung and M. Rabinovich. Least-squares rigid motion using SVD. Technical report, ETH Zurich, 2017.
- [40] H. Späth. Orthogonal least squares fitting with linear manifolds. *Numerische Mathematik*, 48:441–445, 1986.
- [41] B. B. Spratling and D. Mortari. A survey on star identification algorithm. *Algorithms*, 2(1):93–107, 2009.
- [42] D. R. Valeiras, X. Clady, S.-H. Ieng, and R. Benosman. Event-based line fitting and segment detection using a neuromorphic visual sensor. *IEEE Transactions on Neural Networks and Learning Systems*, 2018.
- [43] X. Zhou, N. Tomagou, Y. Ito, and K. Nakano. Implementations of the Hough Transform on embedded multicore processors. *International Journal of Networking and Computing*, 4(1):174–188, 2014.