Speed-Aware EBS-EKF for Event-Based Star Tracking

Gourishankar Mahadeo Bansode
School of Computing and Augmented Intelligence
Arizona State University
Tempe, Arizona, USA
gourishankar@asu.edu

***Abstract*—Event-based cameras are attractive for star tracking because they offer sparse asynchronous measurements, high temporal resolution, low latency, and potentially lower power consumption than conventional framing sensors. However, accurate event-based star tracking in the night sky remains difficult because low-light sensor behavior introduces systematic centroid bias, especially when stars differ in brightness and image-plane speed. Recent work on EBS-EKF addressed this problem with a magnitude-dependent offset correction and a 3D extended Kalman filter, but the supplementary discussion notes that the offset should theoretically vary with star speed even though a single calibration curve is used in practice. In this paper, we investigate Speed-Aware EBS-EKF, an extension that retains the strong 3D EKF backbone of EBS-EKF while replacing the fixed centroid correction with a speed-aware model and adaptive measurement uncertainty. Starting from a working event-tracking baseline that already performs event batching, centroid extraction, astrometric initialization, and batchwise weighted Wahba attitude estimation, we implement the proposed measurement model, validate it on synthetic star-motion tests, and evaluate it on public real-night-sky resources: the released EBS-EKF dataset, Earth-rotation-based event-camera sequences, and public variable-speed slew observations. Our goal is to improve robustness for fast-moving and bright stars without sacrificing the high update rates that make event-based star tracking promising for spacecraft attitude estimation.**

I. INTRODUCTION

Star trackers are a core component of spacecraft attitude determination because they estimate a camera's orientation relative to the celestial sphere using observed star positions. Conventional star trackers built on framing sensors can achieve high accuracy, but their exposure time, frame rate, and processing cadence limit how quickly they can respond to rapid motion. Event-based cameras offer a promising alternative because they output sparse asynchronous events at microsecond-scale temporal resolution, which is especially attractive for star fields where most of the image is dark and only a small number of pixels are active. Early event-based star tracking methods established feasibility by aggregating events into event images and then estimating relative or absolute rotation using geometric alignment, while later methods moved toward asynchronous filtering and higher-rate updates [2], [3], [9], [11].

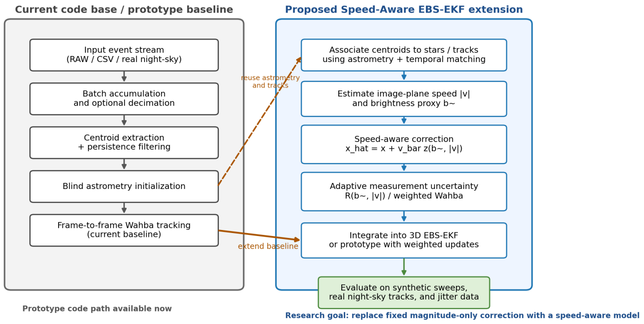

Fig. 1. System diagram for the proposed method. The left side summarizes the current prototype code base: batch-based event accumulation, centroid extraction, blind astrometry, and a batchwise weighted Wahba attitude update. The right side shows the speed-aware extension: estimate image-plane speed and a brightness proxy for each tracked star, apply a speed-aware centroid correction, assign adaptive measurement uncertainty, and integrate the corrected measurements into a weighted attitude update and ultimately a 3D EBS-EKF-style tracker.

At the same time, the main challenge in event-based star tracking is not only temporal estimation but also measurement quality. In low light, event timing and spatial distributions depend on sensor non-idealities, including bandwidth, bias settings, and dim-light triggering behavior [5]–[7]. EBS-EKF is a major step forward because it combines a 3D extended Kalman filter with a brightness-dependent centroid correction derived from low-light sensor modeling, and it validates the approach on real night-sky data rather than only on simulated monitor sequences [11]. However, the method still uses a fixed offset curve as a function of star magnitude, even though the authors note that the offset should theoretically vary with star speed. This leaves an important opening for improvement: if event centroid bias depends on both brightness and image-plane velocity, then fast-moving stars may still be modeled suboptimally under the current approximation.

In this paper, we build on that gap directly. Rather than replacing EBS-EKF with a completely different tracker, we keep its most valuable idea, a physically motivated event measurement model inside a 3D filtering framework, and focus on the part that still appears simplified. Concretely, we study replacing the fixed magnitude-only correction $z(m_s)$ with a speed-aware correction $z(m_s, \|v\|)$, where m_s is apparent magnitude and $\|v\|$ is image-plane speed. We also make the measurement uncertainty adaptive, since fast motion and brighter stars can change the reliability and spatial spread of positive-event centroids. This direction is well matched to the current state of the field: recent public datasets and evaluation protocols, including real night-sky benchmarking, variable-speed slew observations, and Earth-rotation-based accuracy checks, make it increasingly practical to test richer event measurement models under realistic conditions [8], [10], [11].

Our implementation starts from a working event-tracking baseline that already reads event streams, accumulates batches, extracts star centroids, plate-solves with astrometry, and performs batchwise attitude estimation via a weighted Wahba solve. We use that baseline as a debugging and visualization platform, then add a speed-aware correction module inspired by EBS-EKF's low-light event model. Synthetic star-motion experiments are used first to verify that the new correction behaves sensibly as speed and brightness vary, after which we evaluate the method on real night-sky data using across and about attitude error relative to a reference tracker. This paper's contributions include the following: (1) a speed-aware extension of the EBS-EKF centroid correction model, (2) an adaptive measurement-uncertainty formulation for event-based star tracking, and (3) an evaluation pipeline that compares the existing baseline, the original EBS-EKF model, and the proposed extension on synthetic and real event data.

II. RELATED WORK

Event-based star tracking. Conventional star trackers typically localize stars in frame-based images and then estimate camera attitude by matching observed stars to a catalog. Event-based star tracking departs from this setting by replacing synchronous frames with sparse, asynchronous brightness-change measurements. The first direct event-based star tracking pipeline by Chin et al. [3] aggregates events into short event images and estimates attitude through star identification, trimmed ICP, augmented rotation averaging, and rotation-only bundle adjustment. Bagchi and Chin [2] move closer to event-native processing by estimating relative rotations from the spatiotemporal event stream using multiresolution progressive Hough transforms, improving efficiency and making asynchronous operation more plausible. Ng et al. [9] extend this direction with an asynchronous Kalman filter that processes each event individually, but their formulation estimates only the 2D motion of the star field in the image plane and leaves full attitude estimation as future work. Reed et al. [11] introduce EBS-EKF, which combines a low-light event likelihood model, a magnitude-dependent centroid correction, and a 3D extended Kalman filter, and validates the method on synchronized real-night-sky data. Our work is closest to EBS-EKF. Rather than proposing an entirely new tracker, we build on its asynchronous 3D filtering backbone and investigate how its measurement model can be improved for greater robustness in low light and under challenging real-world motion.

Low-light DVS modeling and bias optimization. Several hardware-aware DVS papers are especially relevant because the remaining approximation gap in EBS-EKF lies largely in how individual events are modeled. Graça et al. [6] show that DVS behavior depends strongly on user-adjustable biases controlling bandwidth, sensitivity, refractory period, and noise, and argue that reliable event-camera performance requires hardware-aware modeling rather than purely black-box use. In a companion line of work, Graça et al. [5] analyze the physical limits of DVS noise under dim lighting and show that background activity depends jointly on illumination and bias settings, not illumination alone. Jiang and Zhou [7] further show that dim-light triggering exhibits a discontinuity in event timing caused by parasitic capacitance, and that this effect becomes more prominent when the rate of intensity change is slow. These papers do not solve star tracking directly, but they explain why fixed or overly simplified event likelihoods can become inaccurate in the low-light regime where star tracking operates. Our work uses these insights to motivate adaptive centroid correction and more principled bias and parameter tuning within an EBS-EKF-style tracker.

Calibration and event-based space imaging. Another important line of work studies how event cameras can be calibrated and interpreted for astronomical sensing. Ralph et al. [10] develop an automatic event-based star mapping and astrometric calibration pipeline using real observations, showing how motion compensation, clustering, and plate solving can convert raw event measurements into calibrated source measurements and relate spatiotemporal event-source structure to photometric source properties. This paper is important because it treats calibration and source characterization as first-class problems instead of assuming idealized measurements. Our work differs in scope: we focus on improving online star tracking rather than building a standalone calibration pipeline. However, this literature is still important because better source characterization can inform initialization, preprocessing, and the design of more accurate measurement models.

Datasets, jitter, and evaluation. A major weakness of earlier event-based star tracking papers is that most quantitative evaluations relied on simulated stars displayed on screens. Reed et al. [11] reduce this gap by releasing synchronized event streams and APS star-tracker estimates from real night-sky observations, enabling direct comparison on real data. Bagchi et al. [1] take a complementary step with e-STURT, a controlled dataset of real star observations collected with an event camera mounted on a piezoelectric stage, with hardware-recorded ground-truth jitter across multiple frequency bands and motion axes. Melamed et al. [8] provide another important evaluation protocol by using the Earth's rotation as a high-accuracy reference for assessing event-based star tracking on real starlight, reducing dependence on simulated scenes or on another star tracker as the only source of truth. Together, these papers shift the field from proof-of-concept demonstrations toward more realistic and reproducible evaluation.

III. METHOD

A. Overview

Our method is designed as an incremental extension of the current prototype pipeline rather than a complete replacement of it. The current code base already provides four useful building blocks: event accumulation, centroid extraction, blind

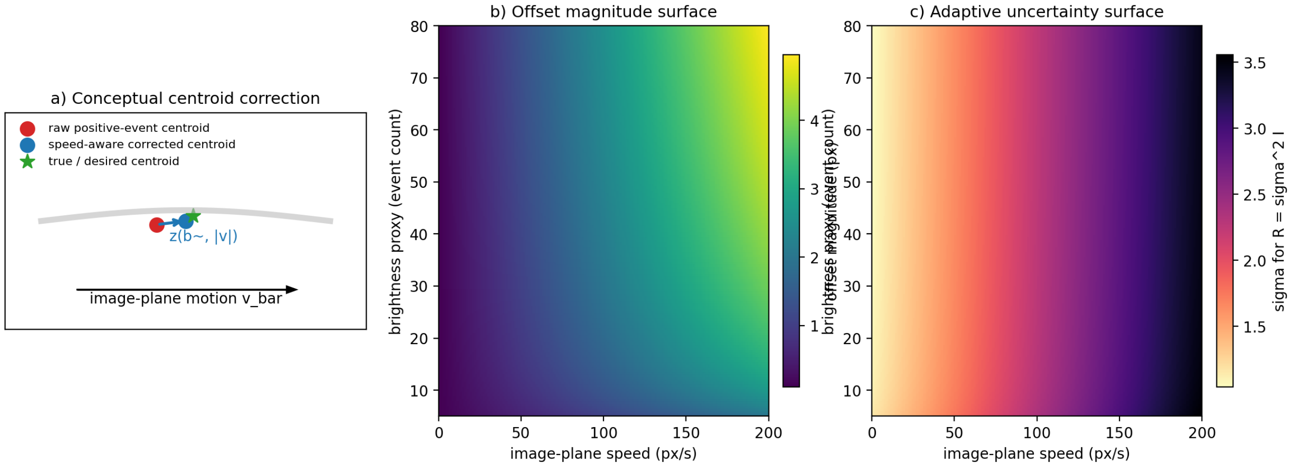

Fig. 2. Conceptual illustration of the proposed measurement model. Panel (a) shows the central idea: the observed positive-event centroid is shifted along the motion direction, and the proposed method applies a speed-aware correction toward the desired centroid. Panel (b) illustrates a plausible offset surface $z(\tilde{b}, \|v\|)$, where the required correction grows with image-plane speed and brightness proxy. Panel (c) illustrates a matching adaptive uncertainty surface, where measurement uncertainty increases for faster motion and weaker centroid support. These panels are conceptual and are included to illustrate the intended design of the calibrated measurement model.

astrometry initialization, and batchwise attitude estimation from matched stars. We retain this structure and insert a measurement-model layer between star association and attitude update. The key idea is to treat the measured centroid of each star as a biased observation whose offset depends not only on brightness, as in EBS-EKF, but also on image-plane speed. In the full target method, this correction would be conditioned on catalog magnitude m_s and used inside an asynchronous 3D EBS-EKF update. In the current prototype, where per-star catalog brightness is not yet threaded through the whole pipeline, we use centroid support count as a brightness proxy $\tilde{b}$.

This design choice keeps the paper aligned with its central goal: improve the measurement model of EBS-EKF without discarding its core insight that physically meaningful event modeling should drive state estimation. It also keeps the implementation honest about what is currently available. Rather than claiming a finished asynchronous 3D tracker, the present system evaluates the proposed measurement model in a batchwise prototype that is simple enough to debug and visualize while still capturing the central geometric effect that motivates the method.

B. Speed-aware centroid correction

Let $x_i \in \mathbb{R}^2$ denote the measured centroid of star i in the current batch, and let $v_i \in \mathbb{R}^2$ denote its estimated image-plane velocity. In EBS-EKF, the corrected measurement is modeled as a magnitude-dependent shift along the direction of motion. We extend this idea with a speed-aware correction,

$$\hat{x}_i = x_i + \bar{v}_i \, z(\tilde{b}_i, \|v_i\|), \tag{1}$$

where $\bar{v}_i = v_i / \|v_i\|$ is the normalized image-plane velocity and $z(\cdot)$ is the proposed offset function. In the full method, z should be learned or calibrated from event data and parameterized by apparent magnitude m_s and speed. In the current prototype, we use a simple bilinear model,

$$z(\tilde{b}_i, \|v_i\|) = \beta_0 + \beta_1 \, \phi(\tilde{b}_i) \, \psi(\|v_i\|), \tag{2}$$

where $\phi(\cdot)$ maps centroid event count to a normalized brightness proxy and $\psi(\cdot)$ maps image-plane speed to a normalized motion scale.

The motivation for Eq. (1) is that EBS-EKF already shows that the positive-event centroid is systematically displaced from the true star location and that this displacement can be corrected by shifting the measurement along the motion direction. However, the EBS-EKF supplementary calibration notes that the offset should theoretically vary with star speed even though a single shared curve is used in practice. This formulation therefore keeps the same geometric interpretation as EBS-EKF but relaxes the assumption that the offset is a function of brightness alone [11]. This formulation preserves the central hypothesis of this paper: centroid bias should increase with motion and should not be treated as a fixed function of brightness alone.

C. Adaptive measurement uncertainty

A second component of the method is to make the reliability of each centroid measurement depend on the measurement conditions. Rather than using a fixed isotropic uncertainty for all stars, we define an adaptive covariance

$$R_i = \sigma_i^2 I_2, \qquad \sigma_i = \sigma_0 \left(1 + \alpha \, \psi(\|v_i\|) + \gamma \, \frac{1}{\sqrt{\tilde{b}_i + \epsilon}} \right), \tag{3}$$

where σ_0 is the base uncertainty, α controls how uncertainty grows with speed, and γ penalizes weak centroid support. Intuitively, fast-moving stars and dimmer or weaker clusters should contribute less strongly to the attitude update.

This adaptive form is motivated by the same low-light considerations that motivate the underlying EBS-EKF signal model. Prior hardware-aware studies show that dim-light DVS behavior depends on bandwidth, bias settings, and non-first-order timing effects, all of which can change the effective sharpness and trustworthiness of a centroid measurement [5]–[7]. In the current prototype code, Eq. (3) is converted into inverse-variance weights for a weighted Wahba update. In the final target system, the same covariance can be inserted directly into the measurement update of a 3D EKF following the EBS-EKF formulation [11].

D. Weighted Wahba attitude update in the prototype

The current prototype does not yet implement the full asynchronous 3D EBS-EKF update. Instead, it estimates attitude batch-by-batch using a weighted form of Wahba's problem. In classical star tracking, once observed stars are associated with known reference directions, the camera orientation is recovered by finding the rotation that best aligns the two sets of unit vectors. This is the role of Wahba's problem, and it is a standard geometric formulation in star tracking pipelines [3].

In our setting, each corrected centroid $\hat{x}_i$ is first backprojected through the camera intrinsics to form a unit observation ray

$$u_i = \frac{K^{-1}\bar{\hat{x}}_i}{\|K^{-1}\bar{\hat{x}}_i\|_2}, \tag{4}$$

where K is the intrinsic calibration matrix and $\bar{\hat{x}}_i = [\hat{x}_i^\top, 1]^\top$ is the homogeneous image coordinate. Let $s_i \in \mathbb{R}^3$ denote the corresponding reference unit direction, obtained either from the previous batch or from the star catalog after initialization. The prototype attitude update is then posed as the weighted least-squares problem

$$R^* = \arg\min_{R \in SO(3)} \sum_i w_i \, \|u_i - R s_i\|_2^2, \tag{5}$$

where $w_i \propto 1/\sigma_i^2$ are inverse-variance weights derived from Eq. (3). Intuitively, stars with higher predicted measurement uncertainty contribute less strongly to the attitude estimate.

This formulation is called a weighted Wahba solve. It estimates the single rotation R^* that best aligns the corrected measurements with their reference directions. In the prototype, tracking is performed by solving Eq. (5) repeatedly on consecutive event batches and then composing the resulting relative rotation with the previous attitude estimate. Therefore, the term "Wahba tracking" in this paper refers to repeated frame-to-frame attitude estimation via Wahba's problem, rather than to a dynamic state-space filter.

The weighted Wahba step is useful in the present work for two reasons. First, it gives a simple and interpretable way to test whether the proposed speed-aware centroid correction improves geometric alignment before introducing the additional complexity of a full nonlinear filter. Second, it provides a practical bridge from the existing batch-based code base to the long-term target system. Once the corrected centroids and uncertainty model are validated, the same measurement model can be inserted into a full asynchronous 3D EBS-EKF update, where rotation and angular velocity are propagated and corrected continuously rather than batchwise [11].

E. Synthetic data generation and scenario design

Before evaluating on real night-sky event streams, we generate controlled synthetic data to debug and calibrate the proposed speed-aware measurement model. This follows our evaluation strategy of using synthetic star-motion experiments as a short validation stage before the full real-data benchmark. Our synthetic pipeline is driven by a scenario table, `trade_space_scenarios_v2.csv`, and a catalog-query/simulation toolchain.

For each scenario, we first query a public star catalog over a rectangular sky region centered at a specified right ascension and declination. The scenario table specifies the field center, search width and height, magnitude limits, and the maximum number of catalog stars retained. We then simulate a moving event-camera view of that field using a pinhole/gnomonic projection onto the focal plane. Star brightness is converted into sensor electron rate using

$$R_e(m) = F_0 \, 10^{-0.4m} A \, \tau_{\text{opt}} \, \eta, \tag{6}$$

where m is star magnitude, F_0 is the zero-point photon flux, A is the aperture collecting area, τ_{opt} is optical throughput, and η is the quantum efficiency. Following standard star-imaging approximations, each star is modeled with an isotropic Gaussian point-spread function (PSF), and the synthetic per-pixel irradiance is computed by integrating the PSF over each pixel area using the error function rather than point sampling. This produces a smoother and more physically meaningful intensity sequence for high-speed event simulation, and it is consistent with the Gaussian star model used in both classical star centroiding and the EBS-EKF measurement model [4], [11], [12].

Given the intensity sequence $I_t(x, y)$, synthetic DVS events are emitted using a log-intensity contrast model. Positive and negative events are triggered when the accumulated log-intensity change crosses thresholds C^+ and C^-, respectively. The simulator writes `event_simulation.raw` / `event_simulation.csv`, frame-level attitude records, and `metadata.json` for each scenario. These outputs are directly useful for our prototype because they mirror the event-stream-plus-attitude structure needed for tracking, plotting, and ablation studies.

The scenario table is structured to support staged experiments rather than a single monolithic sweep. The V01–V24 scenarios emphasize validation sweeps at modest motion, varying star-field density, sky location, slew axis, focal length, detector thresholds, and optional roll angle. The R01–R24 scenarios emphasize speed stress-testing, varying spin rate from low-speed cases to aggressive high-speed motion while also including field-density and hardware-style variations. Future work will first run a reduced subset for debugging and calibration, then launch the full scenario sweep once the speed-aware centroid model is stable.

F. Prototype implementation and evaluation plan

The prototype implementation uses the existing `event_explorer` code base as a bridge between the current prototype and the final tracker. The newly added speed-aware module estimates image-plane speed from centroid displacement across consecutive batches, uses centroid event count as a brightness proxy, applies the speed-aware correction in Eq. (1), and converts Eq. (3) into weights for the weighted Wahba update. This provides an immediately testable version of the idea while preserving the current astrometry and visualization tools.

The experimental design has two stages. First, synthetic sweeps over controlled speed and centroid-support conditions verify that the correction behaves sensibly and identify regimes where the prototype begins to fail. Second, the method is compared against the unmodified baseline and the original EBS-EKF-style magnitude-only correction on real night-sky data, using across and about attitude error as the primary metrics, following the evaluation style used in the EBS-EKF literature [8], [11]. The full target system remains an asynchronous 3D EBS-EKF, but the present prototype is intentionally batchwise: it isolates the effect of the proposed measurement model and makes debugging simpler before introducing the full temporal-state machinery of an event-based EKF.

IV. EXPERIMENTAL SETUP

A. Prototype and Evaluation Protocol

All experiments in this report are conducted using the current batch-based prototype described in Section III. The prototype reads an event stream, accumulates short temporal batches, extracts candidate star centroids, initializes the field using blind astrometry, and then performs batchwise attitude estimation using the weighted Wahba update. The proposed speed-aware extension is inserted into this pipeline as a measurement-model change: after centroid extraction, the code estimates image-plane speed for each tracked star, applies the correction in Eq. (1), and assigns an adaptive measurement covariance through Eq. (3). This design allows the proposed measurement model to be studied without first reimplementing the full asynchronous 3D EBS-EKF, while still preserving the overall estimator structure motivated by EBS-EKF [11].

The initial evaluation protocol compares the existing baseline against the speed-aware prototype on synthetic and real event-based star-tracking data, using the same style of metrics emphasized by EBS-EKF, namely attitude behavior across time and eventual across/about error relative to a reference solution when available [11]. In the present study, the strongest available comparison is a paired baseline-versus-speed-aware run on a representative track, along with runtime statistics collected by the tracking and plotting pipeline. These results are therefore best understood as an initial ablation of the proposed measurement model, not yet as a final quantitative benchmark.

B. Data Sources

Three sources of data are used in this work. First, the real-data path uses the released EBS-EKF real-night-sky benchmark, which provides synchronized event streams and APS reference solutions for direct comparison on real star fields [11]. Second, we use Earth-rotation-based evaluation sequences to support stronger real-sky validation under a cleaner reference-motion model [8]. Third, we use public variable-speed slew observations and synthetic trade-space scenarios for speed-stress testing and calibration of the proposed measurement model [10]. In addition to these public resources, we construct a synthetic trade-space for controlled debugging and calibration. Synthetic scenarios are defined in a scenario table and instantiated by querying a star catalog over a specified sky region, projecting the stars through a gnomonic camera model, converting apparent magnitude to photon/electron rate, modeling each star as a Gaussian point-spread function, and then emitting ON/OFF events through a log-intensity threshold model. This synthetic path is used to explore how the proposed correction behaves as a function of image-plane speed, brightness proxy, field density, and optical configuration before moving to larger real-data sweeps.

C. Metrics

The main quantities examined in these preliminary experiments are: (1) the estimated attitude trajectory in right ascension, declination, and roll; (2) the mean and maximum processing latency per batch; (3) how often the processing time exceeds the nominal per-batch budget; and (4) the number of tracked centroids per batch. These are natural metrics for the current prototype because they directly reveal whether the proposed measurement model changes computational burden or the qualitative stability of the estimated attitude. Future work will pair these plots with across/about attitude error and aggregate ablations across multiple tracks, matching the reporting style of EBS-EKF [11].

V. RESULTS AND DISCUSSION

A. Baseline versus Speed-Aware Prototype

Figures 3 and 4 compare the current magnitude-only baseline against the proposed speed-aware prototype on one representative track. Both runs process the same 3.0 s sequence, consisting of 372 batches and approximately 8.7 million events, with the same average centroid count of 38.3 centroids per batch. This is important because it indicates that the comparison is not being driven by a different number of detected stars; instead, the difference is isolated to the measurement model inserted between centroid extraction and attitude update.

From the runtime plots, the speed-aware prototype does not introduce a measurable computational penalty on this track. In fact, the mean processing latency decreases slightly from 8.47 ms to 8.17 ms, the maximum observed latency decreases from 124.49 ms to 118.39 ms, and the number of budget overruns decreases from 50 to 49. Both runs remain above the real-time budget and are explicitly flagged as *not real-time*,

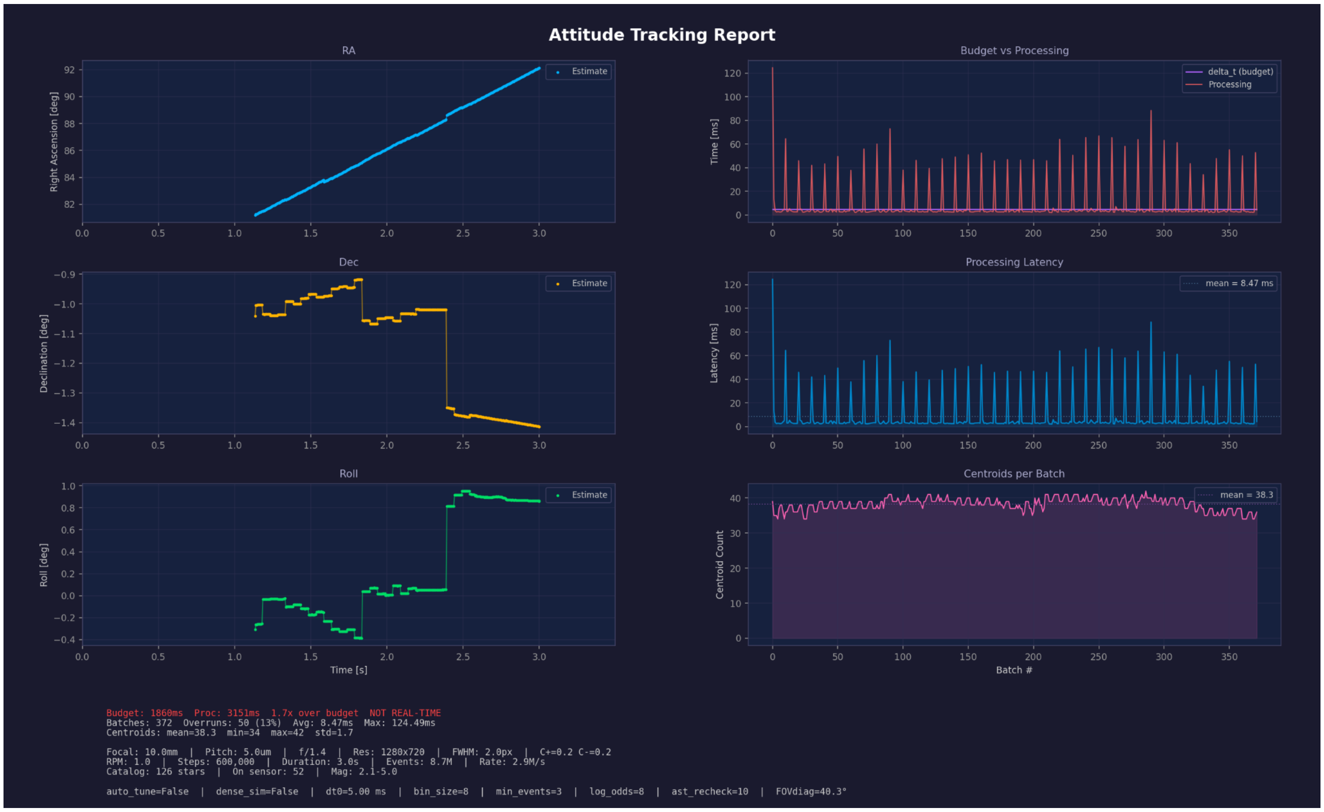

```
Budget: 1860ms  Proc: 3151ms  1.7x over budget  NOT REAL-TIME
Batches: 372  Overruns: 50 (13%)  Avg: 8.47ms  Max: 124.49ms
Centroids: mean=38.3  min=34  max=42  std=1.7

Focal: 10.0mm  |  Pitch: 5.0um  |  f/1.4  |  Res: 1280x720  |  FWHM: 2.0px  |  C+=0.2 C-=0.2
RPM: 1.0  |  Steps: 600,000  |  Duration: 3.0s  |  Events: 8.7M  |  Rate: 2.9M/s
Catalog: 126 stars  |  On sensor: 52  |  Mag: 2.1-5.0

auto_tune=False  |  dense_sim=False  |  dt0=5.00 ms  |  bin_size=8  |  min_events=3  |  log_odds=8  |  ast_recheck=10  |  FOVdiag=40.3°
```

Fig. 3. Attitude tracking report for the magnitude-only baseline on a representative track. The figure shows the estimated right ascension (RA), declination (Dec), and roll over time, along with processing time relative to the batch budget, per-batch latency, and the number of centroids extracted in each batch. This run serves as the reference behavior before introducing the proposed speed-aware centroid correction and adaptive measurement covariance.

but the important result for this phase is that the proposed correction does not make the current prototype slower. This is encouraging because one of the motivations for event-based star tracking is high update rate; a richer measurement model is only attractive if it preserves that advantage.

The attitude traces show that the speed-aware model changes the estimated motion in all three channels. The RA curve exhibits stronger correction and a visibly different late-track trend under the speed-aware model, while the declination and roll plots show larger deviations after approximately 2.4 s. At the present stage, these changes should be interpreted cautiously. They demonstrate that the proposed measurement layer is active and materially affects the final attitude estimate, but they do not yet prove an accuracy improvement by themselves. A stronger claim will require either across/about comparison against a trusted reference or consistent gains across several tracks. This is exactly why future work focuses on paired multi-track evaluation rather than on a single qualitative example.

B. What These Results Mean

Even though the current comparison is preliminary, it already supports two useful conclusions. First, the speed-aware module is technically integrated into the end-to-end pipeline and can alter the final attitude estimate in a meaningful way. Second, this change comes with essentially no additional runtime burden in the present prototype. Together, these two observations justify moving on to a larger ablation study. The most natural next step is to repeat the same baseline-versus-speed-aware comparison on a small set of tracks spanning different motion regimes, then summarize the results with a compact table of latency, budget overrun count, and attitude error where ground truth is available.

At the same time, the paired figures also highlight the main limitation of the present study: we are still evaluating a batch-wise weighted-Wahba prototype, not the full asynchronous 3D EBS-EKF target system. Thus, the present results should be read as evidence that the proposed measurement model is viable and worth pursuing, rather than as final evidence that the full speed-aware EBS-EKF has already surpassed existing approaches. This distinction is important and should be made explicit in the final paper.

C. Broader Design Implications

The broader design motivation remains the same throughout this paper: low-SWaP, high-speed star tracking is fundamentally constrained by photon budget, dwell time, and compute. Small apertures reduce collected signal, high angular rates reduce effective observation time, and both factors make centroiding more fragile. In that setting, improvements to the measurement model are especially valuable because they target one of the few algorithmic levers available without changing

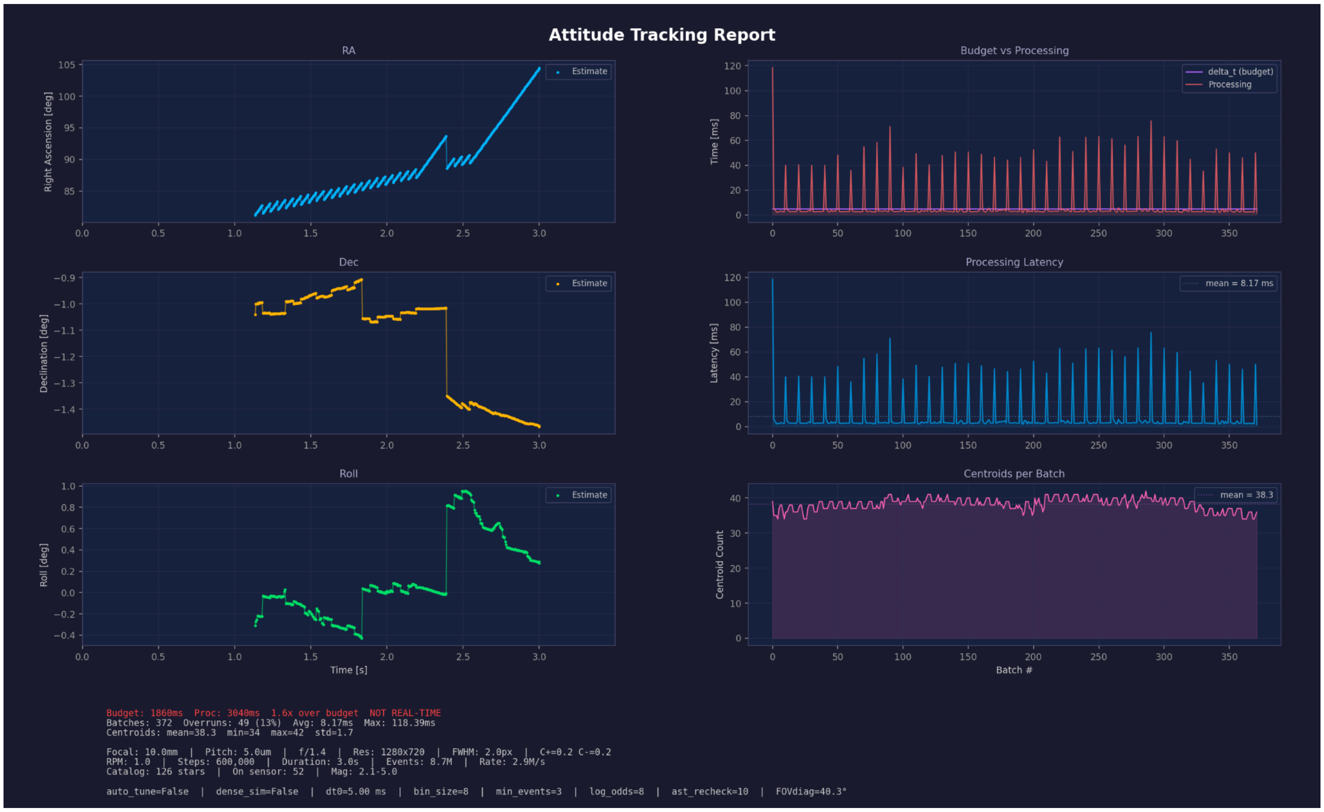

```
Budget: 1860ms  Proc: 3040ms  1.6x over budget  NOT REAL-TIME
Batches: 372  Overruns: 49 (13%)  Avg: 8.17ms  Max: 118.39ms
Centroids: mean=38.3  min=34  max=42  std=1.7

Focal: 10.0mm  |  Pitch: 5.0um  |  f/1.4  |  Res: 1280x720  |  FWHM: 2.0px  |  C+=0.2 C-=0.2
RPM: 1.0  |  Steps: 600,000  |  Duration: 3.0s  |  Events: 8.7M  |  Rate: 2.9M/s
Catalog: 126 stars  |  On sensor: 52  |  Mag: 2.1-5.0

auto_tune=False  |  dense_sim=False  |  dt0=5.00 ms  |  bin_size=8  |  min_events=3  |  log_odds=8  |  ast_recheck=10  |  FOVdiag=40.3°
```

Fig. 4. Attitude tracking report for the proposed speed-aware prototype on the same representative track. Compared with the magnitude-only baseline in Fig. 3, this run applies the proposed speed-aware centroid correction and adaptive measurement covariance before the attitude update. The figure is intended as a qualitative comparison of trajectory behavior, latency, and centroid statistics on the same data sequence.

the hardware. The synthetic trade-space study is therefore not an auxiliary exercise; it is a necessary tool for understanding how the proposed correction behaves as speed, brightness, field density, and optical parameters change.

VI. CONCLUSION

This paper investigates whether the centroid correction in EBS-EKF can be improved by making it speed-aware rather than brightness-only. Starting from a working batch-based event-tracking prototype, we introduced a speed-aware centroid correction and an adaptive measurement-uncertainty model, and we integrated these into a weighted Wahba-based attitude update. The resulting prototype is able to process a representative track without increasing runtime cost, while producing meaningfully different attitude trajectories from the unmodified baseline. These results suggest that the proposed measurement model is practical to evaluate in the current code base and motivates broader experiments on synthetic sweeps and multiple real night-sky tracks.

The main limitation of the present study is that the evaluation is still preliminary. The current results are sufficient to show that the speed-aware module is operational and computationally feasible, but not yet sufficient to claim a consistent accuracy improvement. Future work will therefore focus on larger paired evaluations, beginning with synthetic trade-space sweeps and then moving to public real-night-sky resources and stronger validation protocols such as Earth-rotation-based comparison. In the longer term, the same corrected measurements and adaptive uncertainty model can be integrated into a full asynchronous 3D EBS-EKF, bringing the implementation closer to the original design goal.

REFERENCES

[1] Samya Bagchi, Peter Anastasiou, Matthew Tetlow, Tat-Jun Chin, and Yasir Latif. Event-based star tracking under spacecraft jitter: the e-STURT dataset. *IEEE Transactions on Aerospace and Electronic Systems*, pages 1–18, 2026. Early-access journal version; preprint available as arXiv:2505.12588.

[2] Samya Bagchi and Tat-Jun Chin. Event-based star tracking via multiresolution progressive hough transforms. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 2132–2141, 2020.

[3] Tat-Jun Chin, Samya Bagchi, Anders Eriksson, and André van Schaik. Star tracking using an event camera. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1646–1655, 2019.

[4] Tjorven Delabie, Joris De Schutter, and Bart Vandenbussche. An accurate and efficient gaussian fit centroiding algorithm for star trackers. *The Journal of the Astronautical Sciences*, 61:60–84, 2014.

[5] Rui Graça, Brian McReynolds, and Tobi Delbruck. Optimal biasing and physical limits of DVS event noise. *arXiv preprint arXiv:2304.04019*, 2023. Poster presentation at the 37th International Image Sensor Workshop (IISW 2023).

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

[7] Xiao Jiang and Fei Zhou. Characterization of dim light response in DVS pixel: Discontinuity of event triggering time. *arXiv preprint arXiv:2404.17771*, 2024.

[8] Dennis Melamed, Connor Hashemi, and Scott McCloskey. Quantifying accuracy of an event-based star tracker via earth's rotation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, pages 4696–4703, 2025.

[9] Yonhon Ng, Yasir Latif, Tat-Jun Chin, and Robert Mahony. Asynchronous kalman filter for event-based star tracking. In *Computer Vision – ECCV 2022 Workshops*, volume 13801 of *Lecture Notes in Computer Science*, pages 66–79. Springer, Cham, 2023.

[10] Nicholas Owen Ralph, Alexandre Marcireau, Saeed Afshar, Nicholas Tothill, André van Schaik, and Gregory Cohen. Astrometric calibration and source characterisation of the latest generation neuromorphic event-based cameras for space imaging. *Astrodynamics*, 7(4):415–443, 2023.

[11] Albert W. Reed, Connor Hashemi, Dennis Melamed, Nitesh Menon, Keigo Hirakawa, and Scott McCloskey. EBS-EKF: Accurate and high frequency event-based star tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6510–6519, 2025.

[12] Xiaowei Wan, Gangyi Wang, Xinguo Wei, Jian Li, and Guangjun Zhang. Star centroiding based on fast gaussian fitting for star sensors. *Sensors*, 18(9):2836, 2018.

APPENDIX: AI DETECTION REPORT

Figure 5 shows the Grammarly report for this submission. At the time of submission, Grammarly indicated 2% text matching external sources and 1% text containing patterns that resemble AI-generated text. The flagged overlap appears to come primarily from bibliography entries, citations, paper titles, author names, and standard technical phrasing.

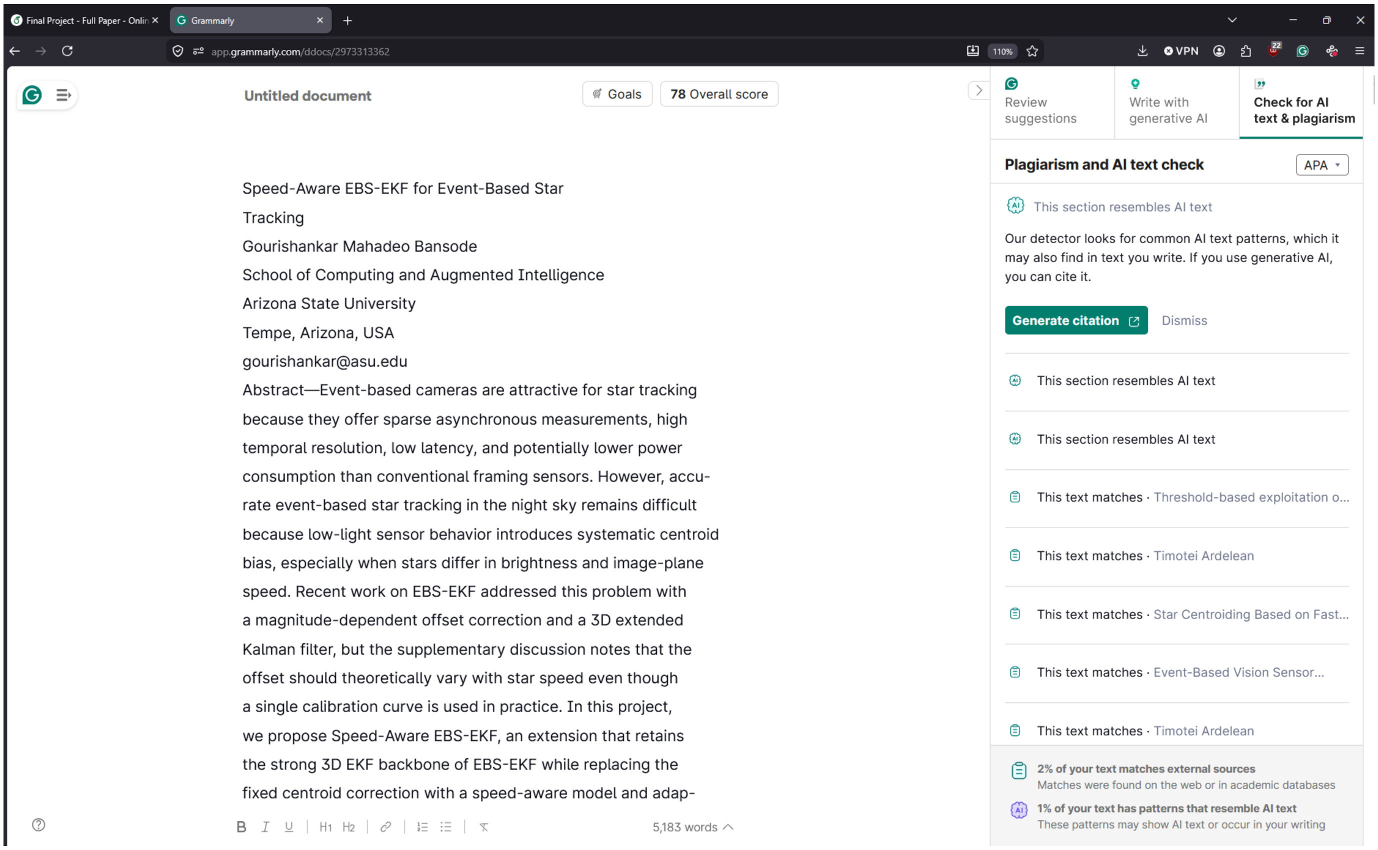

Speed-Aware EBS-EKF for Event-Based Star Tracking
Gourishankar Mahadeo Bansode
School of Computing and Augmented Intelligence
Arizona State University
Tempe, Arizona, USA
gourishankar@asu.edu

Abstract—Event-based cameras are attractive for star tracking because they offer sparse asynchronous measurements, high temporal resolution, low latency, and potentially lower power consumption than conventional framing sensors. However, accurate event-based star tracking in the night sky remains difficult because low-light sensor behavior introduces systematic centroid bias, especially when stars differ in brightness and image-plane speed. Recent work on EBS-EKF addressed this problem with a magnitude-dependent offset correction and a 3D extended Kalman filter, but the supplementary discussion notes that the offset should theoretically vary with star speed even though a single calibration curve is used in practice. In this project, we propose Speed-Aware EBS-EKF, an extension that retains the strong 3D EKF backbone of EBS-EKF while replacing the fixed centroid correction with a speed-aware model and adap-

Plagiarism and AI text check

This section resembles AI text

Our detector looks for common AI text patterns, which it may also find in text you write. If you use generative AI, you can cite it.

This section resembles AI text

This section resembles AI text

This text matches · Threshold-based exploitation o...

This text matches · Timotei Ardelean

This text matches · Star Centroiding Based on Fast...

This text matches · Event-Based Vision Sensor...

This text matches · Timotei Ardelean

2% of your text matches external sources
Matches were found on the web or in academic databases

1% of your text has patterns that resemble AI text
These patterns may show AI text or occur in your writing

Fig. 5. Grammarly report for the final submission, indicating 2% external-source matches and 1% AI-like text patterns.