

Test-Time Certifiable Self-Supervision to Bridge the Sim2Real Gap in Event-Based Satellite Pose Estimation

Mohsi Jawaid[†], Rajat Talak^{*}, Yasir Latif[†], Luca Carlone^{*} and Tat-Jun Chin[†]

Abstract—Deep learning plays a critical role in vision-based satellite pose estimation. However, the scarcity of real data from the space environment means that deep models need to be trained using synthetic data, which raises the Sim2Real domain gap problem. A major cause of the Sim2Real gap are novel lighting conditions encountered during test time. Event sensors have been shown to provide some robustness against lighting variations in vision-based pose estimation. However, challenging lighting conditions due to strong directional light can still cause undesirable effects in the output of commercial off-the-shelf event sensors, such as noisy/spurious events and inhomogeneous event densities on the object. Such effects are non-trivial to simulate in software, thus leading to Sim2Real gap in the event domain. To close the Sim2Real gap in event-based satellite pose estimation, the paper proposes a test-time self-supervision scheme with a certifier module. Self-supervision is enabled by an optimisation routine that aligns a dense point cloud of the predicted satellite pose with the event data to attempt to rectify the inaccurately estimated pose. The certifier attempts to verify the corrected pose, and only certified test-time inputs are backpropagated via implicit differentiation to refine the predicted landmarks, thus improving the pose estimates and closing the Sim2Real gap. Results show that our method outperforms established test-time adaptation schemes.

I. INTRODUCTION

On-orbit servicing (OOS) includes a range of space-based activities such as refuelling, maintenance, assembly and debris removal [1]. By promoting reusability, extendability and space environmental protection, OOS is crucial for sustainable space utilisation. However, the viability of OOS depends on conducting space operations safely and cost effectively, which argues for a high degree of autonomy.

Underpinning OOS are autonomous spacecraft rendezvous and docking capabilities, which require accurate and real-time estimation of the 6DoF pose of a target satellite relative to the servicer satellite in proximity range. To this end, vision-based satellite pose estimation has established its great potential [2], where the monocular case is a leading option due to its lower size, weight and power demands.

State-of-the-art vision-based satellite pose estimators employ deep learning [3], [4], [5], [6], either in an end-to-end manner or in combination with other computer vision techniques [7]. However, the scarcity of real data from the actual operating environment means that the deep models are often trained using synthetic data [8], [9]. This generates a simulation-to-real (Sim2Real) gap [10] between the training




Fig. 1. RGB frames (top) and corresponding event frames (bottom) of a textureless 3D printed satellite-like object (see Sec. V-A for justifications of using a textureless object) under different lighting conditions. The first column are synthetic, while the others are real. Observe that the real event frames under neutral and low are similar to the synthetic event frame, despite visible lighting variations in RGB. However, under harsh directional lighting, numerous events distributed non-uniformly on the object were generated by the event sensor. See also Sec. V-C on sensor tuning.

and testing domains, which raises doubts on whether the trained pose estimators can operate successfully in space.

A major cause of Sim2Real gap in vision-based satellite pose estimation are previously unseen before lighting conditions in space. It is well known that the major sources of illumination in orbit, *i.e.*, the Sun and the Earth’s albedo [11], can produce undesirable effects on the images of satellites in orbit, *e.g.*, high contrast, camera over-exposure, stray light and lens flare [10]. Such effects are also non-trivial to simulate in software, which further contributes to the Sim2Real gap. To bridge the gap, domain adaptation methods [12] have been applied in satellite pose estimation [9], [5], [6].

Another approach to mitigate the effects of novel lighting is by using optical sensors that are less susceptible to illumination variations, specifically, the event sensor [13]. Since event sensors asynchronously detect intensity changes, they have much higher dynamic ranges. Previous works showed that event sensors exhibited a high degree of sensing robustness across drastic lighting changes [14], [15]. The potential benefits of event sensing to space applications have been recognised [16], [17], [18]. In particular, Jawaid *et al.* [19] established the viability of training a 6DoF satellite pose estimator that operates on event frames, with initial results suggesting that the Sim2Real gap for event data is smaller; see Fig. 1 (cols. 2 and 3). We will provide empirical evidence to further support the claim in Sec. VI.

However, while event sensors offer a degree of resilience against lighting variability, harsh illumination conditions can still cause undesirable effects on commercial off-the-shelf

[†]Sentient Satellites Laboratory (SSL), Australian Institute for Machine Learning (AIML), The University of Adelaide, SA 5005, Australia.

^{*}Laboratory for Information and Decision Systems (LIDS), Massachusetts Institute of Technology, Cambridge, MA 02139, USA.




Fig. 2. (a) Real event frame captured under harsh lighting with the object detection bounding-box (blue) and projected wireframe (green) of the CAD model before test-time certifiable self-supervision. (b) Same event frame as (a) but with the projected wireframe after the test-time self-supervision.

(COTS) event sensors. In particular, non-uniform lighting due to strong directional light sources can lead to varying contrasts, which manifest as non-uniform event densities on an object; see Fig. 1 (col. 4). The pattern of real events generated are difficult to synthetically generate, which results in a domain gap for models trained on synthetic event data that did not include the effects of harsh lighting; see Fig. 2(a).

A. Contributions

To deal with the Sim2Real gap in the event domain due to harsh lighting on COTS event sensors, we propose a *certifiable self-supervised scheme* [20], [21] to adapt the event-based pose estimator at test time. The main idea is to employ a certifier that leverages prior knowledge of the target satellite to aim to verify the correctness of estimated poses in a self-supervised manner. Only certified testing inputs are backpropagated via implicit differentiation to bridge the Sim2Real gap at test time. Departing from [20], [21] who focussed on 3D point clouds, we developed a novel certifier that is optimised for the characteristics of 2D event frames (details in Sec. IV-B). See Fig. 2(b) for a preview.

Compared to previous unsupervised domain adaptation methods for satellite pose estimation [5], [6], our method is source-free (whereas [5] requires the original labelled training data in the adaptation) and refines all learnable weights (whereas [6] updates only the batch norm layers in the CNN [22]). Moreover, [5], [6] need to solve auxiliary tasks (e.g., semantic segmentation) to provide additional guidance, whereas our pipeline conducts only pose estimation.

To evaluate our ideas, we developed a novel dataset for event-based satellite pose estimation. Our new dataset is an improvement over the SEENIC dataset in [19], in that real event data from emulated satellite rendezvous trajectories were captured alongside real RGB data, with both channels accompanied with common calibrated ground truth poses.

II. RELATED WORK

A. Satellite pose estimation

In uncooperative rendezvous and docking, the servicer spacecraft must estimate the 6DoF pose of the target space-

craft (e.g., a defunct satellite) without receiving communication or assistance from the target [23]. Sensing modalities that have been explored for the task include radar, LiDAR and optical [24]. Monocular vision is attractive due to its lower size, weight and power demands, especially on nanosatellites which are limited in those respects [2].

The importance of monocular satellite pose estimation as a research topic is underlined by several datasets [25], [10], [8] and challenges [26], [27], [9]. Early adopters of deep learning for the task include Sharma *et al.* [28], [3]. Pauly *et al.* [7] grouped existing deep learning approaches into two broad categories: hybrid modular approaches and direct end-to-end approaches. The former category of methods (e.g., [4], [29], [5]) employ classical vision techniques (e.g., perspective-n-point solver) alongside deep learning capabilities (e.g., object detection, landmark regression), while the latter category (e.g., [28], [3], [6]) directly regresses the pose using a CNN without intermediate non-learning stages.

B. Self-supervised object pose estimation

Self-supervision has been employed to bridge the domain gap between training and testing. Wang *et al.* [30], [31] trained an RGB-D-based pose estimator on synthetic data and refined it with self-supervised training on real unlabelled data, with the self-supervision signal created through differentiable rendering. An iterative student-teacher scheme was proposed by Chen *et al.* [32] to bridge the Sim2Real gap. Wang *et al.* [33] extracted pose-invariant features for category-level object pose estimation of unseen object instances. Differentiable rendering of signed distance fields was utilised by Zakharov *et al.* [34], together with normalised object coordinate spaces [33], to learn 9D cuboids in a self-supervised manner. Zhang *et al.* [35] proposed to jointly reconstruct the 3D shape of an object category and learn dense 2D-3D correspondences between the input image and the 3D shape, leveraging large-scale real-world object videos for training. Deng *et al.* [36] employed a manipulator to interact with the objects in the scene, which generates new data to fine-tune the object pose estimator.

Departing from previous approaches that institute multiple prediction heads that solve complementary tasks, Talak *et al.* [20], [21] advanced the certifiable self-supervision idea, where a geometric optimisation problem is solved to correct for prediction errors, while a certifier tries to verify the correction and only allows good testing samples to be backpropagated via implicit differentiation to adapt the network.

As mentioned in Sec. I-A, domain adaptation is crucial for satellite pose estimation [5], [6]. The advantages of our approach over [5], [6] have been outlined in Sec. I-A.

C. Event sensing for space applications

While event sensing has received a significant amount of attention in robotics and AI, the potential of event sensing for space applications is only starting to be recognised [16]. Many of the applications relate to guidance, navigation and control (GNC) systems of spacecraft and space robots, such

as star tracking and attitude estimation [37], [38], time-to-contact estimation for spacecraft landing [39], [18], and visual odometry for planetary robots [17]. Azzalini *et al.* [40] investigated the generation of synthetic events for optimal landing trajectories. To gain a better idea of the prospects of event sensing for space applications, see [16].

D. Event datasets for satellite pose estimation

As alluded to above, Jawaid *et al.* [19] first considered event sensing for satellite pose estimation, and a dataset was also contributed by them [41]. We address some shortcomings of this data, in particular we now provide equivalent RGB data as well as calibrated 6DoF object poses instead of just camera poses. Another dataset was recently proposed by Rathinam *et al.* [42], however at the time of submission, their dataset has not been released publicly, hence the precise characteristics were unknown, not to mention the inability to evaluate our method on their dataset.

III. EVENT-BASED POSE ESTIMATION

Here, we briefly describe the event-based pose estimation pipeline \mathbf{M}_{base} [19] that our work builds upon. Readers interested in more details of \mathbf{M}_{base} are referred to [19].

We assume that \mathbf{M}_{base} has been pre-trained to conduct 6DoF pose estimation based on event-frame observations of a target satellite-like object. During inference, an event stream \mathcal{E} is converted into event frames $\mathbf{I}^{ev} = \{I_n^{ev}\}_{n=1}^N$ via the E2F algorithm [19] with a fixed temporal batching window τ (see Sec. IV-D for details) where n is the timestamp index for an event frame I_n^{ev} . Each event frame I_n^{ev} is a normalised histogram of event positions where the brightest pixel indicates a pixel position with most contrast change events occurring in τ . Each I_n^{ev} is input to \mathbf{M}_{base} , which conducts object detection [43] that outputs a bounding box which is then passed to the landmark regression module [44]. The landmark regressor produces heatmaps $\hat{\Delta} = \{\hat{\delta}_z\}_{z=1}^Z$ of locations of Z pre-selected 3D landmark positions \mathcal{L} on the CAD model \mathbf{C} , which are converted into discrete 2D positions \mathcal{L}_{2D} . A robust object pose estimation step (PnP with RANSAC) then computes the object pose $(\hat{\mathbf{R}}, \hat{\mathbf{t}})$ given the 2D-3D correspondences between I_n^{ev} and \mathbf{C} .

In our work, \mathbf{M}_{base} was trained on synthetic event frames (details in Sec. V-B) with only the `RandomRotation` and `RandomTranslation` augmentations [19].

IV. TEST-TIME CERTIFIABLE SELF-SUPERVISION

Certifiable self-supervision for object pose estimation was first formalised in [20]. Here, we describe our extension to enable test-time certifiable self-supervision to bridge the Sim2Real gap in event-based satellite pose estimation.

A. Overall algorithm

Our test-time certifiable self-supervision method is summarised in Alg. 1, while the main processing pipeline \mathbf{M}_ζ is illustrated in Fig. 3. Note that \mathbf{M}_ζ subsumes \mathbf{M}_{base} . Our method devises a certifier that leverages prior knowledge of the target satellite in the form of \mathbf{C} to attempt to verify the

correctness of the estimated pose $(\hat{\mathbf{R}}, \hat{\mathbf{t}})$ for a testing event frame. Certified predictions are then used by the corrector to compute the correction signal, which is backpropagated through PnP to update all learnable weights of \mathbf{M}_{base} . We show in Sec. VI that after few epochs, the number of certified testing inputs converge to a high value, indicating that \mathbf{M}_{base} has been adapted to the testing (real) domain.

Algorithm 1 Test-time certifiable self-supervision for event-based satellite pose estimation.

Require: Event-based pose estimation pipeline \mathbf{M}_{base} , N test-time event frames \mathbf{I}^{ev} , CAD model of satellite \mathbf{C} , 3D landmarks \mathcal{L} , self-training epochs E_{self} , self-training batch size b_{self} , learning rate α_{self} , event pixel filter γ , certifier threshold ε , threshold annealing rate β , Hausdorff distance percentile q .

- 1: **for** epoch = 1, 2, ..., E_{self} **do**
- 2: **for** iter = 1, 2, ..., $\lfloor N/b_{self} \rfloor$ **do**
- 3: $\mathbf{I}_b \leftarrow$ Sample b_{self} frames from \mathbf{I}^{ev} .
- 4: **for** I in \mathbf{I}_b **do**
- 5: $\hat{\Delta}, \hat{\mathbf{R}}, \hat{\mathbf{t}} \leftarrow \mathbf{M}_{base}(I)$.
- 6: $cert \leftarrow \text{certify}(I, \hat{\mathbf{R}}, \hat{\mathbf{t}} \mid \mathbf{C}, \gamma, \varepsilon, q)$.
- 7: **if** $cert$ is true **then**
- 8: $\tilde{\Delta} \leftarrow \text{corrector}(\mathbf{K}[\hat{\mathbf{R}} \hat{\mathbf{t}}], \mathcal{L})$.
- 9: $\text{backprop}(\mathbf{M}_{base}, \hat{\Delta}, \tilde{\Delta} \mid \alpha_{self})$.
- 10: **end if**
- 11: **end for**
- 12: **end for**
- 13: $\varepsilon \leftarrow \beta\varepsilon$.
- 14: **end for**
- 15: **return** $(\hat{\mathbf{R}}, \hat{\mathbf{t}})$ predicted by \mathbf{M}_{base} for all frames in \mathbf{I}^{ev} .

More details of the major components in Alg. 1 will be provided in the rest of this section.

B. Certification for event frames

The certifier module consists of two sub-blocks: the *event segmenter* and *alignment verifier*, which work together to verify whether the predicted pose $(\hat{\mathbf{R}}, \hat{\mathbf{t}})$ is consistent with the input event frame I . Details are as follows.

1) *Event segmenter*: The input event frame I is thresholded against value γ to produce a 2D point cloud

$$\mathbf{E}_{2D} = \{\mathbf{x} \in [1, W] \times [1, H] \mid I(\mathbf{x}) \geq \gamma\} := \{\mathbf{u}_p\}_{p=1}^P, \quad (1)$$

where W and H are respectively the width and height of image I . The assumption is that brighter pixels in I are more likely to correspond to valid foreground of the object.

2) *Alignment verifier*: The CAD model \mathbf{C} is posed in the estimated pose $(\hat{\mathbf{R}}, \hat{\mathbf{t}})$ via the camera projection

$$\mathbf{C}_{2D} = f(\mathbf{C} \mid \mathbf{K}, \hat{\mathbf{R}}, \hat{\mathbf{t}}) := \{\mathbf{v}_m\}_{m=1}^M, \quad (2)$$

where function f executes the pinhole projection. We then compute the Hausdorff Distance

$$\mathcal{H}(\mathbf{E}_{2D}, \mathbf{C}_{2D}) = \text{percentile}(\mathcal{D}(\mathbf{E}_{2D}, \mathbf{C}_{2D}), q), \quad (3)$$

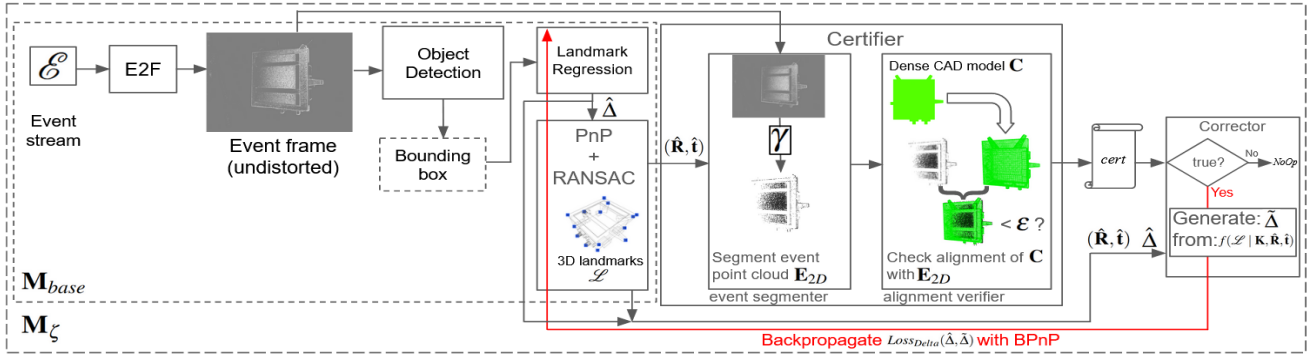


Fig. 3. Our test-time certifiable self-supervised for satellite pose estimation from event frames.

where $\text{percentile}(A, q)$ denotes the member value in the sorted set A below which $q\%$ of members of A lie, and

$$\mathcal{D}(\mathbf{E}_{2D}, \mathbf{C}_{2D}) = \left\{ \min_m \|\mathbf{u}_p - \mathbf{v}_m\|_2 \right\}_{p=1}^P \quad (4)$$

is the set of distances of the nearest neighbours of \mathbf{E}_{2D} in \mathbf{C}_{2D} . Using the q -th largest value of $\mathcal{D}(\mathbf{E}_{2D}, \mathbf{C}_{2D})$ enables a degree of robustness against spurious bright pixels in I .




Fig. 4. Alignment of \mathbf{E}_{2D} (black) and \mathbf{C}_{2D} (green) for a non-certified instance (left) and a certified instance (right). The hausdorff distance $\mathcal{H}(\mathbf{E}_{2D}, \mathbf{C}_{2D})$ is larger for the alignment on the left than the right.

Finally, we say that the estimated pose $(\hat{\mathbf{R}}, \hat{\mathbf{t}})$ for I is certified (*i.e.*, $(I, (\hat{\mathbf{R}}, \hat{\mathbf{t}}))$ is assigned a *true* certificate) if

$$\mathcal{H}(\mathbf{E}_{2D}, \mathbf{C}_{2D}) \leq \varepsilon. \quad (5)$$

Fig. 4 depicts results corresponding to a true and false certificate. Threshold ε is chosen to maximise the number of valid certified instances in the test-time self-supervision batch b_{self} . Moreover, ε is progressively reduced by a rate of β (Step 13 in Alg. 1) to encourage convergence.

C. Corrector and backpropagation

The corrector takes the predicted pose $(\hat{\mathbf{R}}, \hat{\mathbf{t}})$ and projects the 3D landmarks onto I , *i.e.*,

$$\mathcal{L}_{2D} = f(\mathcal{L} | \mathbf{K}, \hat{\mathbf{R}}, \hat{\mathbf{t}}) \quad (6)$$

where $\mathcal{L}_{2D} := \{\mathbf{g}_z\}_{z=1}^Z$ are Z positions in I representing the projected 3D landmarks. We then create Z heatmaps $\tilde{\Delta} = \{\tilde{\delta}_z\}_{z=1}^Z$, where each $\tilde{\delta}_z$ is a matrix of size $W \times H$ with all 0's except at location \mathbf{g}_z where it has the value 1.

The differences between $\hat{\Delta}$ and $\tilde{\Delta}$ provide the correction signal to update the learnable weights in \mathbf{M}_{base} . To this end, we calculate the heatmap loss $Loss_{\Delta}$ between $\hat{\Delta}$ and $\tilde{\Delta}$

$$Loss_{\Delta}(\hat{\Delta}, \tilde{\Delta}) = \frac{1}{Z} \sum_{z=1}^Z MSE(\tilde{\delta}_z, \hat{\delta}_z), \quad (7)$$

which is the average mean-squared-error (*MSE*) between corresponding heatmaps $\tilde{\delta}_z$ and $\hat{\delta}_z$.

In order to update the landmark regressor to enable it to adapt to the testing domain, we employ BPNP [45] to backpropagate $Loss_{\Delta}$ from the corrector, through the PnP solver, to all the learnable weights in the landmark regressor. The learning rate of α_{self} is applied in the weight updates.

D. Hyperparameter settings

1) *Offline training*: The values of the hyperparameters when offline training on synthetic data were as follows:

- $Z = 14$ points on the surface of \mathbf{C} covering unique features *e.g.* corners, end points of the antennae, prominent edges.
- Training batch size $b_{offline} = 24$.
- Training epochs $E_{offline} = 40$.
- Learning-rate $\alpha_{offline} = 0.001$ in the first epoch then reduced by a factor of 0.1 on the 25th and 35th epoch.

2) *Certifiable self-training*: The values of the test-time self-training hyperparameters were as follows:

- Batch duration $\tau = 0.05$ s for the approach scenarios, 0.2 s for the orbit-slow scenarios and 0.01 s for the orbit-fast scenarios. These values were selected to optimise the signal-to-noise ratio in the event frames.
- Self-training batch size $b_{self} = 32$ was selected to maximise GPU usage.
- Self-training epochs $E_{self} = 20$.
- Learning-rate $\alpha_{self} = 0.001$ in the first epoch then reduced by a factor of 0.1 on the 15th epoch.
- Event pixel filter $\gamma = 0.9$.
- Epsilon $\varepsilon = 100$ in the first epoch of the self-training and was reduced by $\beta = 0.975$ at each subsequent epoch.
- Hausdorff distance percentile $q = 0.9997$.

V. EXPERIMENTAL SETUP

A. 3D model of target satellite

We designed a 3D model of a satellite-like object (see Fig. 1) that is geometrically similar to satellites employed in

previous satellite pose estimation works [10], [25].

In order to focus on the effects of lighting variations on the Sim2Real gap, the following design choices were made:

- The CAD model was textureless to preclude as much as possible the contributions of discrepancies in the digital and physical textures to the Sim2Real gap.
- The CAD model was specified to be of a homogeneous material which was hand tuned in Blender [46] to closely match the reflectance of 3D printed PLA surfaces.

B. Synthetic RGB and event datasets for training

The model in Sec. V-A was rendered using Blender. The model was moved smoothly towards the virtual camera while rotating the camera about the y and z axes simultaneously. Then, the camera was kept static and the model was rotated individually about the x, y and z axes and simultaneously about all pairs of the x, y and z axes. This yielded 2,520 sequential rendered RGB frames with ground-truth object pose for each frame. Successive poses differ by 1° on at most 2 axes. A lighting condition was created with 2 identical lights of 2 m radius positioned at distances of 2 m and -2 m in the world z axis, and 1 m in the world y axis. This created even diffused lighting around the object which casted minimal shadows; see Fig. 1 (top left). To generate the synthetic event frames, we employed V2E [47] with the sequential RGB frames as input, and set a matching framerate and exposure duration (e.g., 10 fps and 0.1 s respectively). This ensured that we obtain a one-to-one pairing of the synthetic RGB frames and event frames; see Fig. 1 (col 1).

C. Real RGB and event datasets for testing

1) *Model fabrication*: We 3D printed the model in Sec. V-A to the physical dims. of 26 cm \times 25 cm \times 10 cm.

2) *Sensors*: We employed a Prophesee EVK1 event camera with a 1280×720 pixel event sensor and a Basler Ace 2 conventional camera with a 1920×1200 pixel Sony IMX392 optical sensor. The same Arducam 4–12 mm Varifocal C-Mount lens was used for both cameras, with the focal length tuned separately to ensure similar pixel coverage of the target object given the object physical distance in our laboratory.

3) *Lighting conditions*: The printed 3D model was imaged under 3 lighting conditions:

- *neutral*: illuminated with evenly-lit ambient lighting.
 - *low*: illuminated with 1 weak light source.
 - *harsh*: illuminated with 1 ring light close to the object.
- See Fig. 1 (cols 2–4). *low* provided a dimly lit environment. *harsh* was challenging due to strong directional light, high sensor noise, occlusions due to shadows, and lens flares when the light source was within the field-of-view of the camera.

4) *Trajectories*: Following [19], two trajectories named *approach* and *orbit* were executed using a Universal Robots UR5 robot manipulator. *approach* simulated a docking manoeuvre where the camera linearly approached the satellite, which allowed to evaluate translation errors. The *orbit* trajectory simulated the chaser flying by the target, which allowed to evaluate rotation errors. See [19, Fig. 5] for an illustration of the trajectories. Both of the trajectories

were executed in *fast* and *slow* speeds, where *fast* was $5\times$ higher speed and acceleration than *slow*. The different speeds allowed to investigate the effects of motion blur.

Trajectory	Scene	No.events	No.poses	Duration (s)
approach	harsh-fast	5842584	41	1.6
	harsh-slow	6463394	90	3.56
	lowlight-fast	1399075	41	1.6
	lowlight-slow	1578355	90	3.56
	neutral-fast	1475790	41	1.6
orbit	neutral-slow	1814205	91	3.6
	harsh-fast	13465922	68	2.68
	harsh-slow	15056827	165	6.56
	lowlight-fast	2126339	68	2.68
	lowlight-slow	2119757	165	6.56
	neutral-fast	2171772	68	2.68
	neutral-slow	2568447	165	6.56

TABLE I

SUMMARY OF THE REAL EVENT DATA.

Trajectory/ RGB-setting	Scene	No. frames	No. poses	Duration (s)
approach/ exposure priority	harsh-fast	34	34	1.651
	harsh-slow	72	72	3.552
	lowlight-fast	33	33	1.601
	lowlight-slow	73	73	3.601
	neutral-fast	33	33	1.601
	neutral-slow	73	73	3.601
orbit/ exposure priority	harsh-fast	55	55	2.752
	harsh-slow	138	138	6.854
	lowlight-fast	57	57	2.802
	lowlight-slow	138	138	6.855
	neutral-fast	57	57	2.802
	neutral-slow	139	139	6.904
approach/ motion priority	harsh-fast	97	97	1.603
	harsh-slow	217	217	3.607
	lowlight-fast	96	96	1.653
	lowlight-slow	218	218	3.623
	neutral-fast	97	97	1.603
	neutral-slow	217	217	3.606
orbit/ motion priority	harsh-fast	169	169	2.805
	harsh-slow	414	414	6.900
	lowlight-fast	169	169	2.805
	lowlight-slow	414	414	6.899
	neutral-fast	170	170	2.821
	neutral-slow	408	408	6.882

TABLE II

SUMMARY OF THE REAL RGB DATA.

5) *Data capture*: With the event camera we captured an event stream \mathcal{E} for each motion and light setting, where an event is a tuple with image coordinates, polarity and timestamp indicating intensity change. Note that tuning the event sensor bias settings did not resolve the undesirable outputs under the *harsh* condition, while tuning for the *low* condition would produce many spurious events making the data unusable. Thus, we utilised only the default event sensor settings. Adaptive tuning of event sensors is an open problem [48] and is beyond the scope of this paper.

With the conventional camera we captured J RGB frames $\{I_j^{rgb}\}_{j=1}^J$ where j is the timestamp index for a captured RGB frame I_j^{rgb} . To examine the behaviour of the optical sensor under lighting variations, two exposure settings were used: *motionpriority* which had a low shutter speed resulting in 60 fps, and *exposurepriority* which optimised the

exposure to the low lighting condition resulting in 20 fps. Fig. 5 shows images under the two settings.

The robot arm was polled at 60 Hz to obtain the robot base to gripper pose $(\mathbf{R}_k^{b2g}, \mathbf{t}_k^{b2g})$, where k is the robot timestamp.

6) *Intrinsic calibration*: The camera intrinsics and lens radial distortion parameters were estimated with the aid of a checkerboard pattern displayed on a screen (using respectively the Metavision SDK [49] for the event camera and OpenCV [50] for the frame-based camera).




Fig. 5. Zoomed in view to demonstrate the difference between RGB frames captured with motionpriority and exposurepriority settings.




Fig. 6. (a) Rendered frame of a satellite under normal lighting, with projected wireframe model. (b) Real image of a 3D model of the satellite under neutral lighting; to the human eye (a) and (b) look similar. (c) Event frame corresponding to (a), generated using V2E [51]. (d) Real event frame corresponding to (b). We demonstrate that the domain gap between (c) and (d) is lower than that between (a) and (b).

7) *Ground-truth from calibration*: We obtain the gripper to camera constant rigid transform $(\mathbf{R}^{g2c}, \mathbf{t}^{g2c})$ and the base to checkerboard constant transform $(\mathbf{R}^{b2w}, \mathbf{t}^{b2w})$ by passing in base to gripper poses $(\mathbf{R}_k^{b2g}, \mathbf{t}_k^{b2g})$ polled from the robot arm to OpenCV's `calibrateRobotWorldHandEye` with the estimated transforms from the checkerboard to the camera $(\mathbf{R}_k^{w2c}, \mathbf{t}_k^{w2c})$ from `solvePnP`. A satellite to camera transformation $(\mathbf{R}_k^{s2c}, \mathbf{t}_k^{s2c})$ is computed by manually annotating the landmarks in one frame under neutral lighting. Finally we obtain the satellite to camera poses $\{(\mathbf{R}_k, \mathbf{t}_k)\}_{k=1}^K$ as

$$T(\mathbf{R}_k, \mathbf{t}_k) = T(\mathbf{R}_k^{b2s}, \mathbf{t}_k^{b2s})^{-1} T(\mathbf{R}_k^{b2g}, \mathbf{t}_k^{b2g}) T(\mathbf{R}^{g2c}, \mathbf{t}^{g2c}),$$

where $T(\mathbf{R}, \mathbf{t})$ indicates the 4×4 matrix form of the rigid transformation (\mathbf{R}, \mathbf{t}) , and

$$T(\mathbf{R}_k^{b2s}, \mathbf{t}_k^{b2s}) = T(\mathbf{R}_k^{b2g}, \mathbf{t}_k^{b2g}) T(\mathbf{R}^{g2c}, \mathbf{t}^{g2c}) T(\mathbf{R}_k^{s2c}, \mathbf{t}_k^{s2c})^{-1}.$$

8) *Dataset statistics*: Tabs. I and II summarise the real event and RGB datasets produced¹.

VI. RESULTS

We evaluated our approach on the datasets described in Sec. V-A. As alluded to in Sec. II-D, SEENIC [41] was unsuitable since it did not have simultaneous event and RGB data, while SPADES [42] was not publicly available yet.

A. Details of evaluation metrics

Following Sec. V-C (Data capture), we converted each real event stream \mathcal{E} to event frames $\{I_n^{ev}\}_{n=1}^N$. Each ground truth object pose $(\mathbf{R}_k, \mathbf{t}_k)$ was associated with an event frame I_n^{ev} nearest in timestamp to obtain a sequence $\mathbf{I}_K^{ev} = \{I_k^{ev}\}_{k=1}^K \subset \mathbf{I}^{ev}$ synchronised with the ground truth poses. In a similar way, a sequence of RGB frames $\mathbf{I}_K^{rgb} = \{I_k^{rgb}\}_{k=1}^K \subset \mathbf{I}^{rgb}$ with ground truth was obtained. Given poses $\{(\mathbf{R}_k, \mathbf{t}_k)\}_{k=1}^K$ estimated from \mathbf{I}_K^{ev} or \mathbf{I}_K^{rgb} , following [25], [10] we compare them to the ground truth poses using

$$\phi_k = \frac{\|\mathbf{t}_k - \hat{\mathbf{t}}_k\|_2}{\|\mathbf{t}_k\|_2}, \quad \psi_k = 2 \arccos(|\langle \mathbf{q}_k, \hat{\mathbf{q}}_k \rangle|),$$

which resp. measure translation error and rotation error, and \mathbf{q}_k and $\hat{\mathbf{q}}_k$ are the quaternion form of \mathbf{R}_k and $\hat{\mathbf{R}}_k$. The overall error for an event stream \mathcal{E} or RGB sequence \mathbf{I}^{rgb} are

$$\Phi = \text{mean}(\{\phi_k\}_{k=1}^K), \quad \Psi = \text{mean}(\{\psi_k\}_{k=1}^K).$$

B. Is the domain gap smaller in event data?

In order to test the lighting Sim2Real domain gap in event data, we compared it against conventional RGB data. We trained the pose estimation pipeline \mathbf{M}_{base} separately on synthetic RGB and event data described in Sec. V-B and then run inference of \mathbf{M}_{base} trained on synthetic RGB on real RGB data and \mathbf{M}_{base} trained on synthetic event on real event data. Φ and Ψ are computed for each real event stream \mathcal{E} and real RGB sequence \mathbf{I}^{rgb} which are used as metrics to quantify the domain gap between the synthetic training data and real testing data. Table III shows the values of Φ and Ψ for each permutation of trajectory, lighting and speed setting. From this table we can clearly see that the \mathbf{M}_{base} trained on synthetic event data performs well when tested on real event data whereas, the \mathbf{M}_{base} trained on synthetic RGB data has a notably wider domain gap given the much larger error values with the exception of 3 scenes where the mean value is skewed by few outlying frames. From this we can conclude that indeed the Sim2Real gap is smaller in event data as compared to RGB data. A qualitative result of the same experiment is shown in Fig. 6.

While the Sim2Real gap is smaller in the event domain, there was still appreciable error particularly on sequences with harsh lighting. The next experiment will investigate

¹Our datasets will be released at: <https://github.com/mohsij/space-opera>.

the effectiveness of the proposed certifiable self-supervision method in bridging the remaining Sim2Real gap.

Trajectory/ RGB-setting	Scene	Φ		Ψ (rad)	
		RGB	Event	RGB	Event
approach/ exposure priority	harsh-fast	6.544	1.580	2.900	0.182
	harsh-slow	2.545	0.040	2.977	0.054
	low-fast	28.446	0.035	2.846	0.062
	low-slow	62.276	0.035	2.862	0.061
	neutral-fast	1.703	0.041	3.084	0.060
orbit/ exposure priority	neutral-slow	1.267	0.036	3.084	0.053
	harsh-fast	1.635	10.074	2.509	1.004
	harsh-slow	1.600	42.112	2.827	0.907
	low-fast	125.378	39.917	2.661	0.727
	low-slow	29.354	8.909	2.713	0.814
approach/ motion priority	neutral-fast	1.405	0.120	2.843	0.253
	neutral-slow	1.715	1.936	2.870	0.200
	harsh-fast	41.749	1.580	2.758	0.182
	harsh-slow	58.674	0.040	2.727	0.054
	low-fast	18.292	0.035	2.934	0.062
orbit/ motion priority	low-slow	10.344	0.035	3.056	0.061
	neutral-fast	1.764	0.041	3.060	0.060
	neutral-slow	56.705	0.036	3.078	0.053
	harsh-fast	537.502	10.074	2.627	1.004
	harsh-slow	471.970	42.112	2.611	0.907
orbit/ motion priority	low-fast	62.245	39.917	2.570	0.727
	low-slow	37.387	8.909	2.612	0.814
	neutral-fast	2.116	0.120	2.839	0.253
	neutral-slow	43.310	1.936	2.868	0.200

TABLE III

COMPARISON OF THE ERRORS IN THE RGB AND EVENT DOMAINS.

C. Is test-time certifiable self-supervision effective?

Trajectory	Scene	Φ		Ψ (rad)	
		M_{base}	M_{ζ}	M_{base}	M_{ζ}
approach	harsh-fast	1.580	21.529	0.059	0.098
	harsh-slow	0.040	0.021	0.054	0.085
	low-fast	0.035	0.056	0.062	0.058
	low-slow	0.035	0.024	0.061	0.106
	neutral-fast	0.041	0.049	0.060	0.053
orbit	neutral-slow	0.036	0.012	0.053	0.110
	harsh-fast	10.074	0.021	0.210	0.055
	harsh-slow	42.112	79.796	0.264	0.143
	low-fast	39.917	0.030	0.085	0.045
	low-slow	8.909	0.028	0.152	0.061
orbit	neutral-fast	0.120	0.019	0.069	0.039
	neutral-slow	1.936	0.031	0.081	0.224

TABLE IV

ABLATION STUDY FOR CERTIFIABLE SELF-TRAINING.

We evaluate the difference between pose estimation pipelines M_{base} and M_{ζ} to determine whether the test-time certifiable self-supervision is an effective method to close the remaining domain gap in event data. We train M_{base} and M_{ζ} on synthetic event data and then conduct inference on both pipelines to obtain the overall translation and rotation errors Φ and Ψ respectively. These errors shown in Table IV indicate that the certifiable self-supervision routine definitely improves performance to estimate the translations and although the rotation scores Ψ seem to suggest that M_{base} has equivalent performance in determining the orientation of the satellite, we remind readers as mentioned in Sec. V-C (Trajectories) that the `orbit` scenes are intended to effectively evaluate the pipeline’s ability to estimate the rotation as consecutive poses in the `orbit` scenes display significant

rotation changes. This trajectory provides a clearer idea of how well the pipeline performs in estimating rotations. From these results we can see that our test-time certifiable self-supervision for event data is effective. Fig. 7 is also provided to demonstrate that over the self-supervision routine, the number of certified instances (for initial $\epsilon = 100$) increases as the landmark regressor weights are updated. This indicates that the Sim2Real gap has been alleviated.



Fig. 7. Plot for two representative scenes showing percentage of instances where the certification is true across the self-supervision epochs.

D. Does the certifiable pipeline surpass other TTA methods?

Trajectory	Scene	Φ		Ψ (rad)	
		SPNv2	M_{ζ}	SPNv2	M_{ζ}
approach	harsh-fast	0.285	21.529	1.322	0.098
	harsh-slow	0.229	0.021	1.193	0.085
	low-fast	0.242	0.056	1.159	0.058
	low-slow	0.186	0.024	1.231	0.106
	neutral-fast	0.210	0.049	1.285	0.053
orbit	neutral-slow	0.228	0.012	1.307	0.110
	harsh-fast	0.243	0.021	1.107	0.055
	harsh-slow	0.280	79.796	1.098	0.143
	low-fast	0.253	0.030	1.063	0.045
	low-slow	0.271	0.028	1.048	0.061
orbit	neutral-fast	0.249	0.019	0.996	0.039
	neutral-slow	0.224	0.031	1.008	0.224

TABLE V

COMPARISON OF OUR PIPELINE M_{ζ} AND SPNv2.

We compared the pose estimation performance of our pipeline against SPNv2 proposed by Park *et al.* [6] which amalgamates additional concepts like multi-task learning [5] and test time adaptation (TTA) [22]. We trained both SPNv2 and our pipeline M_{ζ} on synthetic event data and performed TTA as well as inference on real event data. Φ and Ψ are computed for all the real event scenes for SPNv2 and M_{ζ} and summarised in Tab. V. Although, the mean was skewed due to outlying event frames in 2 scenes however in general these results demonstrated the superior performance of M_{ζ} for permutations of laboratory settings and SPNv2 which was designed for RGB data does not cope well with real event data.

VII. CONCLUSIONS

We captured a novel synthetic dataset consisting of RGB and event frames as well as accompanying ground-truth poses extensively covering the expected orientations of the satellite model for training. A new event and RGB dataset was also captured in real-life lab settings with several new

conditions to extend our previous work in [19]. These datasets were evaluated on the same base pipeline to conclusively demonstrate the lower Sim2Real domain gap of event sensors as compared to RGB sensors for the task of satellite pose estimation. Moreover, we proposed a test-time certifiable self-supervision scheme to further close the domain gap for event frames. A comparison against established TTA method indicated superior performance of our pipeline.

REFERENCES

- [1] A. J. Cavaciuti, J. H. Heying, and J. Davis, "In-space servicing, assembly, and manufacturing for the new space economy," Center for Space Policy and Strategy, Aerospace Corporation, Tech. Rep., 2022.
- [2] L. Pasqualetto Cassinis, R. Fonod, and E. Gill, "Review of the robustness and applicability of monocular pose estimation systems for relative navigation with an uncooperative spacecraft," *Progress in Aerospace Sciences*, vol. 110, p. 100548, 2019.
- [3] S. Sharma and S. D'Amico, "Pose estimation for non-cooperative rendezvous using neural networks," *arXiv:1906.09868*, 2019.
- [4] B. Chen, J. Cao, A. Parra, and T.-J. Chin, "Satellite pose estimation with deep landmark regression and nonlinear pose refinement," in *ICCV Workshops*, 2019.
- [5] Z. Wang, M. Chen, Y. Guo, Z. Li, and Q. Yu, "Bridging the domain gap in satellite pose estimation: a self-training approach based on geometrical constraints," *CoRR*, vol. abs/2212.12103, 2022.
- [6] T. H. Park and S. D'Amico, "Robust multi-task learning and online refinement for spacecraft pose estimation across domain gap," *Advances in Space Research*, 2023.
- [7] L. Pauly, W. Rharbaoui, C. Shneider, A. Rathinam, V. Gaudillière, and D. Aouada, "A survey on deep learning-based monocular spacecraft pose estimation: Current state, limitations and prospects," *Acta Astronautica*, vol. 212, pp. 339–360, 2023.
- [8] P. F. Proença and Y. Gao, "Deep learning for spacecraft pose estimation from photorealistic rendering," in *ICRA*. IEEE, 2020.
- [9] "Spark challenge: Spacecraft recognition leveraging knowledge of space environment," <https://cvi2.uni.lu/spark-2021/>.
- [10] T. H. Park, M. Märtens, G. Lecuyer, D. Izzo, and S. D'Amico, "SPEED+: Next-generation dataset for spacecraft pose estimation across domain gap," *arXiv preprint arXiv:2110.03101*, 2021.
- [11] C. Beierle and S. D'Amico, "Variable-magnification optical stimulator for training and validation of spaceborne vision-based navigation," *Journal of Spacecraft and Rockets*, vol. 56, pp. 1060–1072, 2019.
- [12] M. Wang and W. Deng, "Deep visual domain adaptation: A survey," *Neurocomputing*, vol. 312, pp. 135–153, 2018.
- [13] C. Posch, D. Matolin, and R. Wohlgenannt, "A qvga 143 db dynamic range frame-free pwm image sensor with lossless pixel-level video compression and time-domain cds," *IEEE Journal of Solid-State Circuits*, vol. 46, no. 1, pp. 259–275, 2010.
- [14] A. R. Vidal, H. Rebecq, T. Horstschaefer, and D. Scaramuzza, "Ultimate slam? combining events, images, and imu for robust visual slam in hdr and high-speed scenarios," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 994–1001, 2018.
- [15] H. Rebecq, R. Ranftl, V. Koltun, and D. Scaramuzza, "High speed and high dynamic range video with an event camera," *IEEE TPAMI*, vol. 43, no. 6, pp. 1964–1980, 2019.
- [16] D. Izzo, A. Hadjiivanov, D. Dold, G. Meoni, and E. Blazquez, "Neuromorphic computing and sensing in space," *arXiv preprint arXiv:2212.05236*, 2022.
- [17] F. Mählknecht, D. Gehrig, J. Nash, F. M. Rockenbauer, B. Morrell, J. Delaune, and D. Scaramuzza, "Exploring event camera-based odometry for planetary robots," *arXiv preprint arXiv:2204.05880*, 2022.
- [18] S. McLeod, G. Meoni, D. Izzo, A. Mergy, D. Liu, Y. Latif, I. Reid, and T.-J. Chin, "Globally optimal event-based divergence estimation for ventral landing," in *ECCV Workshops*, 2022.
- [19] M. Jawaid, E. Elms, Y. Latif, and T.-J. Chin, "Towards bridging the space domain gap for satellite pose estimation using event sensing," in *ICRA*, 2023.
- [20] R. Talak, L. R. Peng, and L. Carlone, "Certifiable object pose estimation: Foundations, learning models, and self-training," *IEEE Trans. Robotics*, vol. 39, no. 4, pp. 2805–2824, 2023.
- [21] J. Shi, R. Talak, D. Maggio, and L. Carlone, "A correct-and-certify approach to self-supervise object pose estimators via ensemble self-training," in *Robotics: Science and Systems*, 2023.
- [22] D. Wang, E. Shelhamer, S. Liu, B. A. Olshausen, and T. Darrell, "Tent: Fully test-time adaptation by entropy minimization," in *ICLR*, 2021.
- [23] W. Fehse, "Rendezvous with and capture / removal of non-cooperative bodies in orbit: The technical challenges," *Journal of Space Safety Engineering*, vol. 1, no. 1, pp. 17–27, 2014.
- [24] NASA, "On-orbit satellite servicing study," Goddard Space Flight Center, NASA, Tech. Rep., October 2010.
- [25] M. Kisantal, S. Sharma, T. H. Park, D. Izzo, M. Märtens, and S. D'Amico, "Satellite pose estimation challenge: Dataset, competition design, and results," *IEEE TAES*, vol. 56, no. 5, 2020.
- [26] "Kelvins satellite pose estimation competition 2019," <https://kelvins.esa.int/satellite-pose-estimation-challenge/challenge/>.
- [27] "Kelvins satellite pose estimation competition 2021." [Online]. Available: <https://kelvins.esa.int/pose-estimation-2021/challenge/>
- [28] S. Sharma, C. Beierle, and S. D'Amico, "Pose estimation for non-cooperative spacecraft rendezvous using convolutional neural networks," in *2018 IEEE Aerospace Conference*. IEEE, 2018, pp. 1–12.
- [29] Y. Huo, Z. Li, and F. Zhang, "Fast and accurate spacecraft pose estimation from single shot space imagery using box reliability and keypoints existence judgments," *IEEE Access*, vol. 8, 2020.
- [30] G. Wang, F. Manhardt, J. Shao, X. Ji, N. Navab, and F. Tombari, "Self6d: Self-supervised monocular 6d object pose estimation," in *ECCV*, vol. 12346, 2020, pp. 108–125.
- [31] G. Wang, F. Manhardt, X. Liu, X. Ji, and F. Tombari, "Occlusion-aware self-supervised monocular 6d object pose estimation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [32] K. Chen, R. Cao, S. James, Y. Li, Y. Liu, P. Abbeel, and Q. Dou, "Sim-to-real 6d object pose estimation via iterative self-training for robotic bin picking," in *ECCV 2022*, vol. 13699, 2022, pp. 533–550.
- [33] H. Wang, S. Sridhar, J. Huang, J. Valentin, S. Song, and L. J. Guibas, "Normalized object coordinate space for category-level 6d object pose and size estimation," in *CVPR 2019*, 2019.
- [34] S. Zakharov, W. Kehl, A. Bhargava, and A. Gaidon, "Autolabeling 3d objects with differentiable rendering of SDF shape priors," in *CVPR*, 2020.
- [35] K. Zhang, Y. Fu, S. Borse, H. Cai, F. Porikli, and X. Wang, "Self-supervised geometric correspondence for category-level 6d object pose estimation in the wild," in *ICLR*, 2023.
- [36] X. Deng, Y. Xiang, A. Mousavian, C. Eppner, T. Bretl, and D. Fox, "Self-supervised 6d object pose estimation for robot manipulation," in *ICRA 2020*. IEEE, 2020, pp. 3665–3671.
- [37] T.-J. Chin, S. Bagchi, A. Eriksson, and A. Van Schaik, "Star tracking using an event camera," in *CVPR Workshops*, 2019.
- [38] Y. Ng, Y. Latif, T. Chin, and R. E. Mahony, "Asynchronous kalman filter for event-based star tracking," in *ECCV Workshops*, 2022.
- [39] O. Sikorski, D. Izzo, and G. Meoni, "Event-based spacecraft landing using time-to-contact," in *CVPR 2021*, 2021, pp. 1941–1950.
- [40] L. J. Azzalini, E. Blazquez, A. Hadjiivanov, G. Meoni, and D. Izzo, "On the generation of a synthetic event-based vision dataset for navigation and landing," *CoRR*, vol. abs/2308.00394, 2023.
- [41] E. Elms, M. Jawaid, Y. Latif, and T.-J. Chin, "SEENIC: dataset for Spacecraft pose Estimation with Neuromorphic vision," Oct. 2022. [Online]. Available: <https://doi.org/10.5281/zenodo.7214231>
- [42] A. Rathinam, H. Qadadri, and D. Aouada, "Spades: A realistic spacecraft pose estimation dataset using event sensing," *arXiv preprint arXiv:2311.05310*, 2023.
- [43] Y. Wu, A. Kirillov, F. Massa, W.-Y. Lo, and R. Girshick, "Detectron2," <https://github.com/facebookresearch/detectron2>, 2019.
- [44] K. Sun, B. Xiao, D. Liu, and J. Wang, "Deep high-resolution representation learning for human pose estimation," in *CVPR*, 2019.
- [45] B. Chen, A. Parra, J. Cao, N. Li, and T.-J. Chin, "End-to-end learnable geometric vision by backpropagating pnp optimization," in *CVPR*, 2020.
- [46] <https://www.blender.org/>.
- [47] M. L. Katz, K. Nikolic, and T. Delbruck, "Live demonstration: Behavioural emulation of event-based vision sensors," in *ISCAS*, 2012.
- [48] R. Graça, B. McReynolds, and T. Delbruck, "Shining light on the DVS pixel: A tutorial and discussion about biasing and optimization," in *CVPR*, 2023.
- [49] "Metavision sdk docs," <https://docs.prophesee.ai/stable/index.html>.
- [50] G. Bradski, "The OpenCV Library," 2000.
- [51] Y. Hu, S.-C. Liu, and T. Delbruck, "v2e: From video frames to realistic dvs events," in *CVPR*, 2021, pp. 1312–1321.