

Deep Visual Odometry with Events and Frames

Roberto Pellerito,¹ Marco Cannici,¹ Daniel Gehrig,¹ Joris Belhadj,² Olivier Dubois-Matra,² Massimo Casasco,² Davide Scaramuzza¹

¹ Robotics and Perception Group, University of Zurich, Switzerland
² European Space Agency

Abstract—Visual Odometry (VO) is crucial for autonomous robotic navigation, especially in GPS-denied environments like planetary terrains. To improve robustness, recent model-based VO systems have begun combining standard and event-based cameras. While event cameras excel in low-light and high-speed motion, standard cameras provide dense and easier-to-track features. However, the field of image- and event-based VO still predominantly relies on model-based methods and is yet to fully integrate recent image-only advancements leveraging end-to-end learning-based architectures. Seamlessly integrating the two modalities remains challenging due to their different nature, one asynchronous, the other not, limiting the potential for a more effective image- and event-based VO. We introduce RAMP-VO, the first end-to-end learned image- and event-based VO system. It leverages novel Recurrent, Asynchronous, and Massively Parallel (RAMP) encoders capable of fusing asynchronous events with image data, providing $8\times$ faster inference and 33% more accurate predictions than existing solutions. Despite being trained only in simulation, RAMP-VO outperforms previous methods on the newly introduced Apollo and Malapert datasets, and on existing benchmarks, where it improves image- and event-based methods by 58.8% and 30.6%, paving the way for robust and asynchronous VO in space.

Multimedial Material: For code and datasets visit <https://github.com/uzh-rpg/rampvo>.

I. INTRODUCTION

Visual Odometry (VO) is vital for robotic platforms but often fails in challenging scenarios involving low-light, high dynamic range, or high-speed motion. These shortcomings are often caused by the limitations of traditional frame-based camera, such as their susceptibility to motion blur, limited dynamic range, and unfavorable bandwidth-vs-latency tradeoff. While higher frame rates reduce latency, they come at the cost of higher bandwidth and increased processing power. Event-based cameras, which record per-pixel brightness changes asynchronously, address all these issues, offering high dynamic range (HDR), low latency, and low bandwidth and power usage, making them the ideal complement to regular cameras in VO systems. Their combination holds significant promise for critical VO applications, especially when sensors like GPS, LiDAR, and Inertial Measurement Units (IMUs) cannot be used or are ineffective due to radiation and temperature changes. These

This work was supported by the European Union’s Horizon Europe Research and Innovation Programme under grant agreement No. 101120732 (AUTOASSESS) and the European Research Council (ERC) under grant agreement No. 864042 (AGILEFLIGHT).

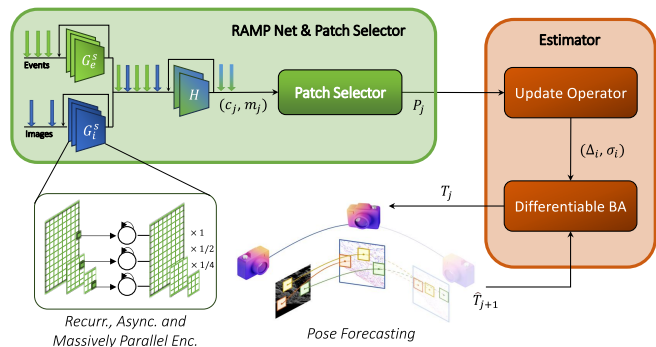


Fig. 1. Recurrent, Asynchronous, and Massively Parallel (RAMP) Encoders are used to process asynchronous events and images. Patches are extracted from the resulting encoding and used by the Estimator inspired by DPVO [1] to perform data-driven feature tracking and visual odometry. A simple pose forecasting module exploits previously extracted patches to initialize poses in the bundle adjustment, allowing for improved performance.

conditions are often encountered in planetary exploration and landing missions, where rapid motion and partial shadow are also common.

Recent model-based solutions [2] have shown potential in fusing images and events. However, the fields of image-only [1], [3] and event-only [4] VO have recently shown that learning-based pipelines, trained end-to-end, can surpass traditional model-based systems in accuracy and robustness. While the combination of data-driven approaches with systems that leverage images and events appears promising, effectively combining event data—with its distinctive asynchronicity and sparsity—with synchronous and dense frames is a non-trivial challenge in learning-based solutions. Traditional learning-based methods typically resort to artificially synchronizing events at image timestamps to facilitate data fusion, reducing the rate at which events are processed to that of the slower image modality. Nevertheless, in tasks such as VO, this simplification is not ideal and might limit the algorithm’s ability to exploit events received in between images, which is crucial for tracking features effectively.

To address these limitations, our work introduces an adaptive fusion approach that adjusts the frequency of event fusion based on the rate of incoming events, thus mirroring the pace of the scene dynamics. Our Recurrent, Asynchronous, and Massively Parallel (RAMP) encoders handle asynchronous events and images at varying rates and fuse them into a pyramidal memory that serves as a data-

agnostic feature space. We use these encoders in RAMP-VO, the first learning-based VO method that uses both events and frames. RAMP-VO leverages a motion-aware strategy based on event data to extract robust patch-based feature tracks. These features are then processed by a differentiable bundle adjustment module [1] which leverages a simple pose forecasting module for initialization.

We train RAMP-VO on an event-based version of TartanAir [5]. To address the lack of visual odometry datasets that feature image and event data in challenging space landing settings, we also introduce two novel datasets: the Malapert landing and the Apollo landing datasets, which feature challenging motion and lighting conditions due to stark shadows cast by the sun. The first dataset represents a realistic simulation of a spacecraft landing, covering several kilometers of descent near the Malapert crater in the south Moon pole. The second dataset, captured with real RGB and event cameras, features landings on a 3D scale model of the lunar surface and precise ground truth camera poses, making it a valuable resource for research and evaluation.

Despite being trained purely in simulation, RAMP-VO outperforms both image-based and event-based methods on traditional real-world benchmarks, as well as on the newly introduced Apollo and Malapert landing datasets. To summarize our contributions are:

- A novel massively parallel feature extractor, termed RAMP encoder, that asynchronously¹ fuses images and events, both spatially and temporally. Our encoder is 8 times faster and achieves a 33% higher performance than other state-of-the-art asynchronous solutions.
- RAMP-VO, the first learning-based VO using events and frames, which outperforms both image-based and event-based methods by 58.8% and 30.6%, respectively on traditional real-world benchmarks.
- Two novel datasets, Apollo and Malapert landing, targeting challenging planetary landing scenarios.

II. RELATED WORK

Learning-based Visual Odometry. Recent advancements in VO have witnessed a paradigm shift towards learning-based approaches [6], surpassing traditional methods in accuracy and robustness [7], [8]. Unsupervised methods [9], [10], [11] exploit additional depth and optical flow predictors while recent methods employ neural radiance fields [12], [13], [14] to further improve performance. Supervised methods, on the other hand, either rely on end-to-end camera-motion regressors [15], [16], [17], [18] or exploit hybrid solutions that combine geometric models with deep neural networks [19], [20], [21]. Among these works, DROID-SLAM [3] recently proposed to combine an iterative learning-based optimization inspired by RAFT [22] with a

¹Notice that in this work, the term “asynchronous” refers to our network’s ability to handle data streams at different and varying rates. Our network processes images or packets of events (as frame-like event representations) as soon as they are available, without stream synchronization. This differs from networks operating on event-by-event processing, where “asynchronous” refers to the network layers’ functioning.

differentiable bundle adjustment layer. The follow-up work, DPVO [1], further improves its efficiency by replacing dense feature tracking with a sparse patch-based variant. Our work builds upon DPVO, but significantly improves its robustness by effectively fusing images and events together.

Event-based Motion Estimation. Although full 6DOF pose estimation using only events has been successfully demonstrated in the literature [23], [24], most event-based VO rely on additional sensors. While some systems incorporate depth estimates from stereo or depth cameras [25], [26], [27], others integrate IMU measurements to improve robustness and scale recovery [28], [29], [30], [31], [32], [33], [29]. Standard image frames have also been incorporated to extract features and then track them with events [34] or to optimize additional residual errors [28], [2], often exploiting DAVIS cameras [35] or beamsplitter setups [2].

Despite these advancements, challenges persist in low-texture environments, directing toward exploration into deep learning approaches. While early attempts focused on unsupervised techniques [36], [37], DEVO [4] has recently shown promise in transferring to out-of-distribution scenarios by training on large simulated datasets. Notably, however, DEVO only utilizes events and still depends on encoders primarily optimized for images. In contrast, our work employs a recurrent and pyramidal feature extractor that effectively fuses images with events, preserving their incremental nature and exploiting the best of both modalities.

Fusing events and frames. While a great variety of approaches leverage, optimize, or fuse images and events for different downstream tasks [38], [39], [40], [41], the topic of effectively fusing the two data modalities while considering their different nature has thus far been underexplored. Some methods [38], [42], [39] synchronize and concatenate both modalities and process them together with a shared encoder. Others [40], [43], instead, use specialized feature extractors, but still resort to data synchronization for processing.

To date, only one prior study, RAM Net [44], has proposed a specialized and asynchronous way of fusing events and frames. However, RAM Net relies on a sequential hierarchical feature extraction process and utilizes slow Conv-GRU modules. Our asynchronous encoders build upon RAM Net but exploit pixel-wise operations and parallel extraction of multi-scale features, demonstrating both higher performance as well as improved efficiency.

III. METHODOLOGY

Our end-to-end event- and frame-based visual odometry algorithm, RAMP-VO, builds on deep patch visual odometry (DPVO) [1] and takes inspiration from recurrent asynchronous multimodal (RAM) networks [44]. The main innovations of this work reside in how events and frames are fused together. An ideal vision encoder should fuse sensor measurements at an adaptable rate, processing more information during fast motion for better tracking, and ensuring reliability even when a sensor experiences temporary outages. Our RAMP encoders achieve this by creating a sensor-agnostic representation that is recurrently

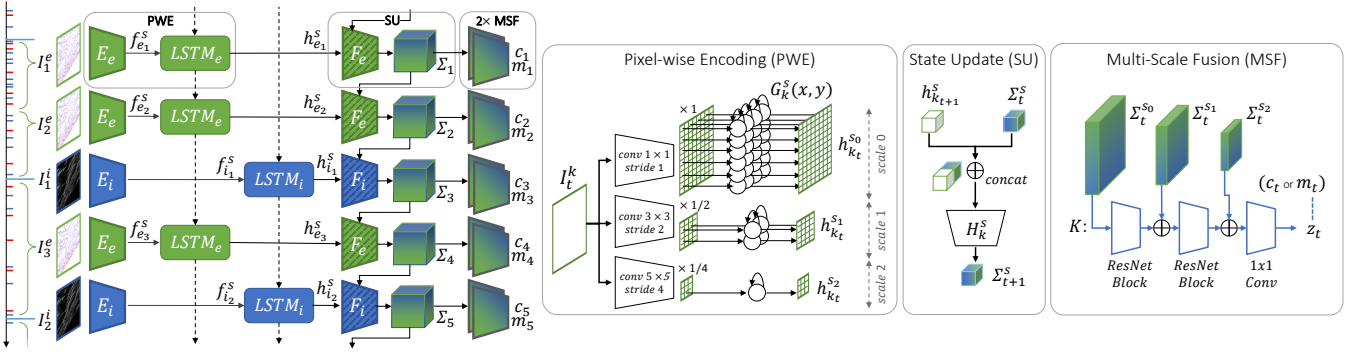


Fig. 2. An overview of the proposed RAMP Net encoder. Events and images are first asynchronously processed by two parallel pixel-wise, multi-scale, encoding branches (PWE) made of a set of convolutional layers followed by pixel-wise LSTMs $G_k^s(x, y)$. A shared state Σ_t^s is then updated (SU) with features coming from different data modalities k by employing sensor-specific encoders at each scale. The multi-scale features are then finally combined through two separate fusion modules (MSF) to produce the matching and context features, m_t and c_t

updated whenever new data becomes available, either images or events. We present the main components of RAMP-VO in Section III-A, and the RAMP encoder in Section III-B.

A. Overview

RAMP-VO processes a temporally ordered, asynchronous stream of data made of ‘frames’. These are either standard images or frames made of events with $H \times W$ resolution. Similar to DPVO [1], we extract a set of patches and predict their optical flow across frames by computing correlation volumes over *matching features* \mathbf{m}_j and additional *context features* \mathbf{c}_j . We iteratively refine the predicted patch flow and finally use a differentiable bundle adjustment layer to update the camera poses. For each new frame with index j we first compute \mathbf{m}_j and \mathbf{c}_j through the RAMP encoder described in Section III-B, and then extract N patches with dimension $p \times p$ from these feature maps. While DPVO [1] and DEVO [4] select these patches randomly or through a learning-based strategy, we opt for a simpler yet effective alternative. In particular, we rely on events’ activity to extract patches from textured areas, as these regions usually generate a higher number of events compared to uniform ones. We first compute an $H \times W$ density map by counting the number of events per pixel. We then apply non-maximum suppression to remove candidates that are not the highest in their 11×11 neighborhood and finally obtain the patch centers by selecting the N locations with the highest counts.

We denote the l -th patch \mathbf{P}_j^l extracted from frame j as

$$\mathbf{P}_j^l = [\mathbf{x} \ \mathbf{y} \ \mathbf{1} \ \mathbf{d}]^T, \quad \mathbf{x}, \mathbf{y}, \mathbf{d} \in \mathbb{R}^{1 \times p^2}. \quad (1)$$

As in DPVO, we model patches as collections of contiguous pixels, with \mathbf{x}, \mathbf{y} the coordinates of the pixels in the patch and \mathbf{d} their depth (constant within the patch).

Following [1], we build a bipartite *patch graph* \mathcal{E} , depicted in Figure 3, by connecting a patch \mathbf{P}_j^l extracted from frame j to every frame i within a distance r from j , each containing the projection \mathbf{P}_{ji}^l of the original patch onto the frame:

$$\mathbf{P}_{ji}^l \sim \bar{\mathbf{K}} \mathbf{T}_i \mathbf{T}_j^{-1} \bar{\mathbf{K}}^{-1} \mathbf{P}_j^l, \quad \bar{\mathbf{K}} = \begin{pmatrix} \mathbf{K} & \mathbf{0} \\ \mathbf{0}^T & 1 \end{pmatrix} \quad (2)$$

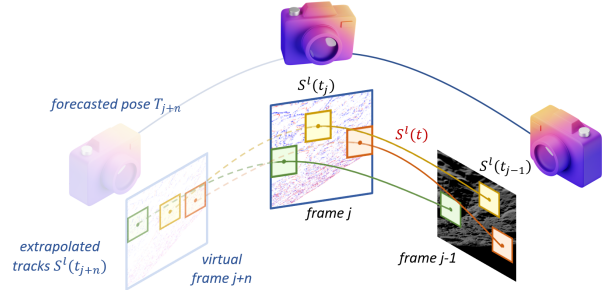


Fig. 3. Illustration of pose initialization. Through patch extraction and projection into future frames we construct feature tracks for frames $j, j-1, \dots$ which we use to construct the splines $S^l(t_j)$. To perform pose initialization, we extrapolate the feature tracks to time t_{j+n} , and apply bundle adjustment to solve for the forecasted pose T_{j+n} .

where $\bar{\mathbf{K}}$ is a 4×4 matrix built from the 3×3 camera matrix \mathbf{K} and $\mathbf{T}_i, \mathbf{T}_j$ are poses of frames i, j . We summarize this operation as $\mathbf{P}_{ji}^l = \omega(\mathbf{T}_i, \mathbf{T}_j, \mathbf{P}_j^l)$.

Next, RAMP-VO computes camera motion by estimating 2D corrections $\Delta_{li} \in \mathbb{R}^2$ for each projected patch \mathbf{P}_{ji}^l , as well as importance values $\sigma_{li} \in \mathbb{R}^{2 \times 2}$ through a series of blocks involving a correlation lookup, 1D Convolution, Soft Aggregation, Transition Block and Factor Head. These steps extract the features of each patch in frame j from \mathbf{m}_j and compare them with those obtained by cropping \mathbf{m}_i with the reprojection of the patch in frame i , together with additional context features \mathbf{c}_i . Since these operations, described in [1], are out of the scope of the current work, we summarize them with the *update operator* F in Fig. 1:

$$\Delta_{li}, \sigma_{li} = F(\mathbf{P}_{ji}^l, \mathbf{c}_i, \mathbf{m}_i, \mathbf{m}_j) \quad (3)$$

Finally, given the corrected positions $\mathbf{P}_{ji}^l + \Delta_{li}$ and their weights σ_{li} , RAMP-VO performs a differentiable bundle adjustment (BA) step, which minimizes the projection error:

$$\sum_{(l,i) \in \mathcal{E}} \left\| \left[\hat{\mathbf{P}}_{ji}^l + \Delta_{li} \right] - \hat{\omega}(\mathbf{T}_i, \mathbf{T}_j, \mathbf{P}_j^l) \right\|_{\sigma_{li}}^2. \quad (4)$$

Here the left-hand side is kept fixed, while the optimization solves for the camera poses $\mathbf{T}_i, \mathbf{T}_j$ and depths \mathbf{d}_j^l , refining the

camera trajectory to match the predicted patch trajectories. The $\hat{\cdot}$ operator selects the central pixel of the patch, and \mathcal{E} is the path graph. This operation is fully differentiable, and thus used during training to backpropagate errors.

Contrary to [1] which uses a simple linear initialization, we bootstrap initial poses in our RAMP-VO by fitting two cubic univariate splines modeling the 2D motion of the patch center along the last 11 frames, such that $(\mathbf{x}_j^l, \mathbf{y}_j^l) \sim \mathbf{S}^l(t_j) = (S_x^l(t_j), S_y^l(t_j))$. We then extrapolate the location of the camera at time t_{j+1} , by evaluating $\mathbf{S}^l(t_{j+1})$, assuming the first and second derivatives to be zero at the splines' boundaries and constant depth. We obtain the 6 DOF pose at time t_{j+1} through BA, optimizing Equation 4 in which we use $\mathbf{S}^l(t_{j+1})$ in place of the corrected patch position.

B. Asynchronous and Massively Parallel Encoders

Denote the stream of data $\{\mathbf{I}_{k_j}(t_j)\}_{j=1}^T$ captured at timestamps t_j . Here \mathbf{I} (henceforth denoted as "frame") denotes either a $5 \times H \times W$ sized event stack [45], in the case of events, or a $C \times H \times W$ sized image². The variable $k_j \in \{e, i\}$ denotes the sensor at timestamp t_j . We encode these data structures using a Recurrent, Asynchronous and Massively Parallel (RAMP) encoder. An overview of the architecture is provided in Figure 2. We employ two different Multi-Scale Fusion modules, with identical structures, to generate either the context or the matching feature maps.

We designed the RAMP encoders to be easily parallelizable and to effectively fuse asynchronous data streams. To do so, we encode each pixel independently, in contrast with [46] which requires sequential processing and expensive two-stage ConvGRU operations, and recurrently fuse information into a shared state. First, RAMP encoders process the data stream with pixel-wise, sensor-specific networks at multiple scales s , creating a sequence of features $\{\mathbf{f}_{k_j}^s\}_{j=1}^T$:

$$\mathbf{f}_{k_j}^s(t_j) = E_{k_j}^s(\mathbf{I}_{k_j}(t_j)). \quad (5)$$

These encoders have a kernel size 1×1 , 3×3 , and 5×5 respectively, and a stride of 2^s with $s = 0, 1, 2$.

This stream of features is then further encoded into a sensor-agnostic state through two recurrent stages: (1) intra-sensor update, and (2) inter-sensor update. In the intra-sensor update step, we process features originating from a single sensor and scale using an LSTM operating on individual pixels, inspired by [47]

$$\mathbf{h}_{k_j}^s(t_j) = G_{k_j}^s(\mathbf{f}_{k_j}^s(t_j)), \quad (6)$$

where we have omitted the cell state for brevity.

The inter-sensor update step, finally combines together the hidden states from separate sensors asynchronously using only a single depth-wise convolution in the following way:

$$\Sigma_j^s = H_{k_j} \left(\|\mathbf{h}_{k_j}^s(t_j)\| \|\Sigma_{j-1}\| \right) \quad (7)$$

where $\|$ denotes concatenation. A different depth-wise convolution H_{k_j} is used depending on which sensor the data

being encoded is generated from. Notably, updates to the Σ_j^s state can occur at any frequency, enabling our encoder to seamlessly handle asynchronous data streams.

After generating the sensor-agnostic states Σ_j at each time t_j , two hierarchical encoders generate either the context or the matching feature maps, each at $1/4$ the original resolution

$$\mathbf{z}_j = K(\{\Sigma_j^s\}_{s=0}^2). \quad (8)$$

Here \mathbf{z}_j can be either \mathbf{c}_j or \mathbf{m}_j , *i.e.*, the context and matching features required by the DPVO backend.

Multi-Scale Fusion. The Multi-scale Fusion (MSF) module follows the single-scale feature extractors in [1] and consists of a 7×7 convolution with stride 2, four residual blocks (two at $1/2$ and two at $1/4$ of the initial resolution), and a final 1×1 convolution. We follow the configuration in [1], using full-scale features Σ_j^0 as input to the encoder, but additionally injecting Σ_j^1 and Σ_j^2 after the second and fourth residual blocks respectively. We do so by concatenating these features to the residual block's outputs and adjusting the channels of the next operation to accommodate the additional features.

IV. EXPERIMENTS

Training. We train RAMP-VO on the synthetic TartanAir [5] dataset, with same train and test-split in [1], which we augment with events using VID2E [48]. During training, we follow the DPVO [1] training scheme to select and filter frames for training. Additionally, we also enforce a minimum of 1.2M events between every pair of frames and remove additional frames if this condition is not satisfied. Finally, we feed RAMP-VO by interleaving two event stacks for every pair of selected frames. We create the first event stack by stacking the 600,000 events received just before the mid-timestamp between the pair of images, and the second event stack by aggregating the 600,000 events preceding the second frame. Since ground truth poses are only given for frames, we do not compute a loss for the mid-frame events. At inference time, we feed the model with asynchronous images and events, by creating a new event stack each time M events are received. As the number of triggered events depends on the scene dynamics, with more events triggered with faster motion, the number of event stacks collected varies adaptively. For comparisons, we also test our model using the data feeding strategy used for training, which we call *synchronized* as events are collected at regular rates.

We train RAMP-VO for 350,000 steps with sequences of 15 images and 30 event stacks on a Quadro RTX 8000 GPU. The remaining hyperparameters are set as in DPVO [1]. Full training takes around 8 days on our hardware.

Datasets. We use the Stereo DAVIS [49] and EDS [2] datasets, and the ECCV 2020 SLAM competition TartanAir [5] test-split to benchmark our method. Stereo DAVIS and EDS are real-camera datasets, allowing us to test RAMP-VO's robustness to domain shift and noise originating from the sensor or potential miscalibration. To ablate the crucial components of RAMP-VO architecture, we use instead the TartanAir test-split used in [1]. Additionally, we introduce

²For color images $C = 3$ and for gray-scale images $C = 1$.




Fig. 4. Comparisons of the ablated models (a) and of RAMP with asynchronous and synchronized data (b) on the full TartanAir test set. We show the importance of using a RAMP encoder as in RAMP-VO over a sequential, single-scale encoder and feed-forward encoder. We also show the benefits of using the full event information in RAMP-VO with finer discretizations. The RAMP encoder is better at maintaining memory than the RAM-Net-like encoder, as highlighted in low framerate VO experiments (c) on the *carwelding* sequences of TartanAir.




Fig. 5. Preview, and qualitative trajectory comparison on Malapert dataset (a), and Apollo dataset (b). Note that while the Malapert sequence is measured in kilometers, the Apollo sequence, recorded at a miniature scale of the Moon’s surface, is in centimeters.

two new benchmarks that replicate lunar surface landings and feature challenging lighting conditions: Apollo landing and Malapert landing. These benchmarks feature rapid motion, high dynamic range, and textureless terrains.

Malapert landing. It features a total of 20 minutes of simulated recordings from the Malapert south Moon region, divided into 2 sequences. We use the planets and satellites simulator PANGU³ to generate realistic descent trajectories, each 250 km long in altitude and 40 km in translation, featuring partial or complete darkness. Ground truth poses of the spacecraft’s center of mass are provided at 5Hz, together with 640×480 synchronized RGB images. We generate synthetic events using Vid2E [48] with default settings.

Apollo landing. A real dataset consisting of a total of 5 minutes of recording, split into 6 trajectories, featuring both vertical and lateral descent trajectories on a 260×260 cm scale replica of the Apollo 17 landing site. Frames and events are recorded with a beam-splitter similar to the one used in [2]. Poses are recorded with an OptiTrack motion capture system. We downsample frames and events to a resolution of 640×480 before processing.

Baselines. We evaluate the proposed RAMP-VO architecture against several VO state-of-the-art methods making use of images only (I), events only (E) as well as based on both images and events (I+E). We follow the evaluation in [2], [1], [4] and select ORB-SLAM2 [50], ORB-SLAM3 [51], COLMAP [52], [53], DROID [3] and DPVO [1] as image-only baselines, while we use EDS [2] for comparison

against methods fusing images and events, being the only VO system of this kind in the literature. We implement an event-only DPVO baseline, EDPVO, that directly processes event stacks, as well as one that processes images and events concatenated together, DPVO+events. Finally, we also compare against DEVO [4] which makes use of a similar DPVO-based architecture but only relies on events, contrary to our method that also uses images.

A. Effects of RAMP blocks

We start by analyzing the individual contribution of RAMP-VO modules and input modalities on the TartanAir [5] test set. In this section, we adopt the evaluation protocol in [1], which analyzes the percentage of sequences from the TartanAir test set below a given absolute trajectory error threshold, producing plots like the one in Fig. 4. This is done to discount the effect of individual diverging sequences, which report abnormally high trajectory errors, and would skew the average. We use 5000 different thresholds equally spaced between an ATE[m] of 0 to 1. We summarize the results by computing the the area under the curve (AUC), and use it to compare between ablations.

Results: From Figure 4 (a), DPVO+events has the lowest performance among the methods we tested (AUC of 0.56), highlighting that a trivial event-image fusion is not sufficient to exploit events. Moreover, our RAMP-VO encoder in synchronous mode outperforms the baseline that uses RAM-Net both when we use multiple scales, as in RAM-Net, but also when just one scale is used. The single-scale RAMP-VO (AUC of 0.71) achieves a 11% increase over

³<https://pangu.software/>

TABLE I
AVERAGE ABSOLUTE TRAJECTORY ERROR ON THE MALAPERT AND APOLLO DATASETS

	Input	Malapert [km]		Apollo [cm]		
		cam-1	cam-2	rec-1	rec-3	rec-4
DPVO [1]	I	73.1	48.2	0.9	0.3	1.3
DPVO [1]	I+E	40.1	34.0	0.9	0.2	1.1
RAMP-VO SS (Ours)	I+E	1.1	9.4	0.7	0.2	1.0
RAMP-VO (Ours)	I+E	0.6	4.3	0.8	0.2	0.9

the RAM-Net encoder (AUC 0.64), while the multi-scale version (AUC 0.81) improves by 27% over RAM-Net. When events are processed asynchronously, using $M = 250000$, the performance increases even more, reaching a 0.85 AUC and a 33% improvement.

In Figure 4 (b), we further analyze the performance of RAMP-VO using asynchronous images and events while varying the number of events M from 100000 to 750000. Although RAMP-VO is trained with synchronized events stacks and frames, it demonstrates consistent performance with asynchronous data. Notably, as the event sequence’s discretization becomes finer (i.e., smaller M values), performance improves, peaking at $M = 250,000$, where it achieves a 5% improvement compared to synchronous event feeding.

Low Framerate VO: In this section, we demonstrate that our method is not only able of asynchronous processing, but it can also cope with frames provided at very low rates. Since our RAMP encoders build sensor-agnostic features, we can exploit the faster data stream, i.e., event voxels, to keep the feature embedding updated, and generate a consistent pose.

To test this capability, we design an experiment on the *carwelding hard p003* and *carwelding easy p007* sequences of the TartanAir, where we artificially reduce the framerate of the images, subsampling them, but keeping event stacks fixed at 20 Hz. We then evaluate the trajectory error against ground truth poses at 20 Hz for DPVO+events, RAMP-VO, and RAMP-VO with RAM-Net-like encoder, and report the results in Fig. 4 (c). Our RAMP encoder consistently outperforms the baselines, including the RAM-Net-like encoder, indicating its superior performance on asynchronous data.

Timing Results: To further motivate the use of the proposed RAMP encoders, as opposed to RAM Net [44], we time the two encoders on the TartanAir test set. We measure the average time required to extract features Σ_j from a single frame (or event) on a Quadro RTX 8000 GPU. RAM Net takes 370 ms on average, while the proposed RAMP Net only takes 47 ms, resulting in a $8\times$ speedup. Leveraging pixel-wise feature processing, our encoder exploits significantly higher parallelism than RAM Net.

B. Results on Space Data

Next, we validate RAMP-VO on the low light and low frame rates Malapert and Apollo landing datasets, where report the average absolute trajectory error over 5 runs.

Results: In Table I, we report performance on the Malapert landing dataset, where RAMP-VO recovers accurate poses deviating only 0.2% to 1.7% from the ground truth over 250km trajectories. DPVO, instead, can not recover a valid

TABLE II
AVERAGE ABSOLUTE TRAJECTORY ERROR (M) ON ECCV 2020 SLAM COMPETITION MONOCULAR TEST-SPLIT. METHODS MARKED WITH (✓) USE GLOBAL OPTIMIZATION / LOOP CLOSURE. TOP PERFORMING (NON-GLOBAL) METHOD IN BOLD, SECOND BEST UNDERLINED.

	ORB-SLAM3 [51]	COLMAP [52]	DROID-SLAM [3]	DSO [54]	DROID-VO [3]	DPVO [1]	RAMP-VO (Ours)
Input	I	I	I	I	I	I	I+E
Global	✓	✓	✓	✗	✗	✗	✗
ME00	13.61	15.20	0.17	9.65	0.22	0.16	<u>0.20</u>
ME01	16.86	5.58	0.06	3.84	0.15	<u>0.11</u>	0.04
ME02	20.57	10.86	0.36	12.20	0.24	<u>0.11</u>	0.10
ME03	16.00	3.93	0.87	8.17	1.27	<u>0.66</u>	0.46
ME04	22.27	2.62	1.14	9.27	1.04	<u>0.31</u>	0.16
ME05	9.28	14.78	0.13	2.94	0.14	<u>0.14</u>	0.13
ME06	21.61	7.00	1.13	8.15	1.32	<u>0.30</u>	0.12
ME07	7.74	18.47	0.06	5.43	0.77	<u>0.13</u>	0.12
MH00	15.44	12.26	0.08	9.92	<u>0.32</u>	0.21	0.36
MH01	2.92	13.45	0.05	0.35	0.13	0.04	<u>0.06</u>
MH02	13.51	13.45	0.04	7.96	<u>0.08</u>	0.04	0.04
MH03	8.18	20.95	0.02	3.46	0.09	<u>0.08</u>	0.04
MH04	2.59	24.97	0.01	-	1.52	<u>0.58</u>	0.41
MH05	21.91	16.79	0.68	12.58	0.69	0.17	<u>0.25</u>
MH06	11.70	7.01	0.30	8.42	<u>0.39</u>	0.11	0.11
MH07	25.88	7.97	0.07	7.50	0.97	<u>0.15</u>	0.07
Average	14.38	12.50	0.33	7.32	0.58	<u>0.21</u>	0.17

TABLE III
AVERAGE ABSOLUTE TRAJECTORY ERROR (CM) ON STEREO DAVIS

	Input	Bin	Boxes	Desk	Monitor
ORB-SLAM2 [50]	I	2.5	7.0	9.3	10.3
DPVO [1]	I	3.6	6.7	7.0	11.5
DPVO [1]	E	5.8	7.2	7.8	14.4
EDS [2]	I+E	2.6	5.8	5.0	8.0
DPVO [1]	I+E	4.2	5.3	4.9	6.2
RAMP-VO (Ours)	I+E	3.1	5.1	3.1	4.2

trajectory leading to ATE errors of several kilometers, equivalent to 20% to 30%. When events are added, DPVO decreases the error by 45.14% to 29.46%, highlighting the importance of events in dark regions. We report a qualitative comparison on a Malapert sample in Figure 5.

On the Apollo landing dataset, both RAMP-VO versions, multi- and single-scale (SS), outperform image and image+event DPVO baselines by up to 30.77%, reaching an error from 2% to 6% of the ground truth. Single-scale and multi-scale RAMP-VO achieve similar results on Apollo, contrary to Malapert’s low-light environments, where single-scale errors are twice as high, indicating the need to focus on both global and fine-grained details for improved robustness.

C. Comparison with State of the Art

We conclude by comparing our RAMP-VO with state-of-the-art methods on the Stereo DAVIS [49] dataset, the EDS [2] benchmark, and the Tartan Air test-split from the ECCV 2020 SLAM competition. We generate events for the ECCV2020 competition using VID2E [48] with default settings. We run RAMP net asynchronously for all tests, using $M = 20000$ for Stereo DAVIS [49] and $M = 250000$

TABLE IV

AVG. ABSOLUTE TRAJECTORY ERROR (CM) AND R_{rmse} (DEG) ON EDS. (*) INDICATES METHODS TRAINED ON TARTANAIR TRAIN+TEST DATA. TOP PERFORMING (NON-GLOBAL) VO METHOD IN BOLD, SECOND BEST UNDERLINED.

Input Global	ORB-SLAM3 [51]	DPVO [1]	DEVO* [4]	RAMP-VO (Ours)
	I ✓	I ✗	E ✗	I+E ✗
	ATE / R_{rmse}	ATE / R_{rmse}	ATE / R_{rmse}	ATE / R_{rmse}
pean. dark	6.15 / 11.40	<u>1.26</u> / <u>1.83</u>	4.78 / 2.49	1.20 / 0.64
pean. light	27.26 / 6.88	<u>12.99</u> / 2.66	21.07 / 3.84	9.03 / 6.40
pean. run	16.83 / 5.78	<u>25.48</u> / 11.19	38.10 / 18.28	13.19 / <u>11.43</u>
rocket dark	10.12 / 9.75	27.41 / 5.23	8.78 / 4.16	7.20 / 2.63
rocket light	32.53 / 11.39	63.11 / 10.44	<u>59.83</u> / <u>9.28</u>	17.53 / 4.04
ziggy	26.92 / 4.42	<u>14.86</u> / <u>3.45</u>	11.84 / 2.32	19.05 / 7.66
ziggy hdr	81.98 / 17.67	66.17 / 10.32	22.82 / <u>9.07</u>	<u>28.78</u> / 5.13
ziggy fly.	20.57 / 8.02	<u>10.85</u> / <u>3.66</u>	10.92 / 3.39	6.35 / 5.07
all chars	21.37 / 9.02	95.87 / 29.00	10.76 / 3.62	<u>28.61</u> / <u>9.89</u>
Average	27.08 / 9.37	35.33 / 8.64	<u>21.00</u> / <u>6.27</u>	14.57 / 5.88

events with the ECCV2020 competition test-split and EDS. Poses are collected after processing a frame.

Results: Results for ECCV 2020 competition are available in Table II, where we follow the evaluation in [1], reporting the ATE[m] of the median of 5 runs with scale alignment. RAMP-VO is able to recover a better pose compared to all other image-based state-of-art methods [51], [54], [52], [3], [1] in most cases, outperforming by 19% and 48% DPVO [1] and DROID-SLAM using loop closure [3], respectively.

Results on Stereo DAVIS are reported in Table III. Given our emphasis on space applications, where space-graded cameras such as the AURICAMTM 4 often operate at low-FPS, we evaluate performance using the 5 FPS benchmark introduced in the supplementary analysis of [2]. Except for *Bin*, the proposed RAMP-VO outperforms all other baselines that use either both or just one of the two modalities on Stereo DAVIS, surpassing the top-performing method, EDS, by 27.6%. On the EDS benchmark reported in Table IV, our method consistently improves over DPVO and achieves an average 30% improvement over DEVO which, contrary to our method, is trained on additional data from the TartanAir test set. These benchmarks represent completely new scenarios compared to the TartanAir training setting. Stereo DAVIS features gray-scale frames and a lower 180×240 resolution, while EDS has a higher resolution from Prophesee Gen3.1 and FLIR cameras, and an increased variety of test cases, including light and dark scenes and wider motions. Similar to the Apollo landing dataset, RAMP-VO is thus required to generalize from simulated to real sensor data.

It is worth noting how effective processing and fusion of event data is particularly important to achieve high performance in these benchmarks. Indeed, naive adaptations of DPVO for event processing fall short, while specialized event fusion techniques, like ours, can achieve better generalization

⁴AURICAM’s datasheet available on www.sodern.com.

and transfer to real-world data.

V. CONCLUSION

In this work, we introduce RAMP-VO, an end-to-end VO system tailored for challenging environments such as those encountered during lunar descents. RAMP-VO uses RAMP encoders to fuse event data into frames, achieving $8 \times$ speedup and 33% improvement over state-of-the-art asynchronous encoders. Moreover, by incorporating events, RAMP-VO reduces the trajectory error of existing image-only deep-learning-based solutions by 58.8%, as well as event-only methods by 30.6% on average. Experiments show that RAMP-VO can transfer zero-shot to real data despite being trained only on synthetic one, while still outperforming the other baselines. We designed RAMP-VO for space applications focusing on improving accuracy and sensor fusion. Future efforts should prioritize optimizing the method for deployment on resource-constrained, space-graded hardware leveraging techniques such as model quantization, network compression, and efficient hardware implementations [55]. Despite these limitations, we view this work as a milestone in event data fusion for VO, and we believe it can spark new interest in the use of event cameras and learning-based approaches for robust navigation.

REFERENCES

- [1] Z. Teed, L. Lipson, and J. Deng, “Deep patch visual odometry,” in *Annual Conference on Neural Information Processing Systems*, 2023.
- [2] J. Hidalgo-Carrió, G. Gallego, and D. Scaramuzza, “Event-aided direct sparse odometry,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2022, pp. 5781–5790.
- [3] Z. Teed and J. Deng, “Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 16 558–16 569, 2021.
- [4] S. Klenk, M. Motzet, L. Koestler, and D. Cremers, “Deep event visual odometry,” in *IEEE International Conference on 3D Vision*. IEEE, 2023.
- [5] W. Wang, D. Zhu, X. Wang, Y. Hu, Y. Qiu, C. Wang, Y. Hu, A. Kapoor, and S. Scherer, “Tartanair: A dataset to push the limits of visual slam,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2020, pp. 4909–4916.
- [6] C. Chen, B. Wang, C. X. Lu, N. Trigoni, and A. Markham, “Deep learning for visual localization and mapping: A survey,” *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- [7] R. Song, R. Zhu, Z. Xiao, and B. Yan, “Contextavo: Local context guided and refining poses for deep visual odometry,” *Neurocomputing*, vol. 533, pp. 86–103, 2023.
- [8] I. A. Kazerouni, L. Fitzgerald, G. Dooly, and D. Toal, “A survey of state-of-the-art on visual slam,” *Expert Systems with Applications*, vol. 205, p. 117734, 2022.
- [9] A. Ranjan, V. Jampani, L. Balles, K. Kim, D. Sun, J. Wulff, and M. J. Black, “Competitive collaboration: Joint unsupervised learning of depth, camera motion, optical flow and motion segmentation,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 12 240–12 249.
- [10] G. Lu, “Deep unsupervised visual odometry via bundle adjusted pose graph optimization,” in *IEEE International Conference on Robotics and Automation*. IEEE, 2023, pp. 6131–6137.
- [11] W. Zhao, Y. Wang, Z. Wang, R. Li, P. Xiao, J. Wang, and R. Guo, “Self-supervised deep monocular visual odometry and depth estimation with observation variation,” *Displays*, vol. 80, p. 102553, 2023.
- [12] Z. Zhu, S. Peng, V. Larsson, W. Xu, H. Bao, Z. Cui, M. R. Oswald, and M. Pollefeys, “Nice-slam: Neural implicit scalable encoding for slam,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2022, pp. 12 786–12 796.
- [13] E. Sucar, S. Liu, J. Ortiz, and A. J. Davison, “imap: Implicit mapping and positioning in real-time,” in *IEEE International Conference on Computer Vision*, 2021, pp. 6229–6238.

- [14] E. Sandström, K. Ta, L. Van Gool, and M. R. Oswald, "Uncle-slam: Uncertainty learning for dense neural slam," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2023, pp. 4537–4548.
- [15] S. Wang, R. Clark, H. Wen, and N. Trigoni, "Deepvo: Towards end-to-end visual odometry with deep recurrent convolutional neural networks," in *IEEE International Conference on Robotics and Automation*. IEEE, 2017, pp. 2043–2050.
- [16] F. Xue, X. Wang, S. Li, Q. Wang, J. Wang, and H. Zha, "Beyond tracking: Selecting memory and refining poses for deep visual odometry," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 8575–8583.
- [17] M. R. U. Saputra, P. P. De Gusmao, Y. Almalioğlu, A. Markham, and N. Trigoni, "Distilling knowledge from a deep pose regressor network," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 263–272.
- [18] X.-Y. Kuo, C. Liu, K.-C. Lin, and C.-Y. Lee, "Dynamic attention-based visual odometry," in *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2020, pp. 36–37.
- [19] L. Sun, W. Yin, E. Xie, Z. Li, C. Sun, and C. Shen, "Improving monocular visual odometry using learned depth," *IEEE Transactions on Robotics*, vol. 38, no. 5, pp. 3173–3186, 2022.
- [20] H. Zhan, C. S. Weerasekera, J.-W. Bian, and I. Reid, "Visual odometry revisited: What should be learnt?" in *IEEE International Conference on Robotics and Automation*. IEEE, 2020, pp. 4203–4210.
- [21] N. Yang, L. v. Stumberg, R. Wang, and D. Cremers, "D3vo: Deep depth, deep pose and deep uncertainty for monocular visual odometry," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020, pp. 1281–1292.
- [22] Z. Teed and J. Deng, "Raft: Recurrent all-pairs field transforms for optical flow," in *European Conference on Computer Vision*. Springer, 2020, pp. 402–419.
- [23] H. Kim, S. Leutenegger, and A. J. Davison, "Real-time 3d reconstruction and 6-dof tracking with an event camera," in *European Conference on Computer Vision*. Springer, 2016, pp. 349–364.
- [24] H. Rebecq, T. Horstschaefer, G. Gallego, and D. Scaramuzza, "Evo: A geometric approach to event-based 6-dof parallel tracking and mapping in real time," *IEEE Robotics and Automation Letters*, vol. 2, pp. 593–600, 2017.
- [25] Y. Zhou, G. Gallego, and S. Shen, "Event-based stereo visual odometry," *IEEE Transactions on Robotics*, vol. 37, no. 5, pp. 1433–1450, 2021.
- [26] Y.-F. Zuo, J. Yang, J. Chen, X. Wang, Y. Wang, and L. Kneip, "Devo: Depth-event camera visual odometry in challenging conditions," in *IEEE International Conference on Robotics and Automation*. IEEE, 2022, pp. 2179–2185.
- [27] D. Weikersdorfer, D. B. Adrian, D. Cremers, and J. Conrath, "Event-based 3d slam with a depth-augmented dynamic vision sensor," in *IEEE International Conference on Robotics and Automation*. IEEE, 2014, pp. 359–364.
- [28] 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.
- [29] A. Zihao Zhu, N. Atanasov, and K. Daniilidis, "Event-based visual inertial odometry," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 5391–5399.
- [30] H. Rebecq, T. Horstschaefer, and D. Scaramuzza, "Real-time visual-inertial odometry for event cameras using keyframe-based nonlinear optimization," in *British Machine Vision Conference*, 2017.
- [31] F. Mählknecht, D. Gehrig, J. Nash, F. M. Rockenbauer, B. Morrell, J. Delaune, and D. Scaramuzza, "Exploring event camera-based odometry for planetary robots," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 8651–8658, 2022.
- [32] T. Qin, J. Pan, S. Cao, and S. Shen, "A general optimization-based framework for local odometry estimation with multiple sensors," *arXiv preprint arXiv:1901.03638*, 2019.
- [33] P. Chen, W. Guan, and P. Lu, "Esvio: Event-based stereo visual inertial odometry," *IEEE International Conference on Robotics and Automation*, 2023.
- [34] B. Kueng, E. Mueggler, G. Gallego, and D. Scaramuzza, "Low-latency visual odometry using event-based feature tracks," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2016, pp. 16–23.
- [35] C. Brandli, R. Berner, M. Yang, S.-C. Liu, and T. Delbruck, "A 240x180 130dB 3 μ s latency global shutter spatiotemporal vision sensor," *IEEE Journal of Solid-State Circuits*, vol. 49, no. 10, pp. 2333–2341, 2014.
- [36] A. Z. Zhu, L. Yuan, K. Chaney, and K. Daniilidis, "Unsupervised event-based learning of optical flow, depth, and egomotion," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 989–997.
- [37] C. Ye, A. Mitrokhin, C. Fermüller, J. A. Yorke, and Y. Aloimonos, "Unsupervised learning of dense optical flow, depth and egomotion with event-based sensors," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2020, pp. 5831–5838.
- [38] I. Alonso and A. C. Murillo, "EV-SegNet: Semantic segmentation for event-based cameras," in *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2019.
- [39] S. Tulyakov, D. Gehrig, S. Georgoulis, J. Erbach, M. Gehrig, Y. Li, and D. Scaramuzza, "Time lens: Event-based video frame interpolation," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2021, pp. 16 155–16 164.
- [40] S. Tulyakov, A. Bochicchio, D. Gehrig, S. Georgoulis, Y. Li, and D. Scaramuzza, "Time lens++: Event-based frame interpolation with parametric non-linear flow and multi-scale fusion," *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [41] L. Sun, C. Sakaridis, J. Liang, Q. Jiang, K. Yang, P. Sun, Y. Ye, K. Wang, and L. V. Gool, "Event-based fusion for motion deblurring with cross-modal attention," in *European Conference on Computer Vision*. Springer, 2022, pp. 412–428.
- [42] Y. Hu, J. Binas, D. Neil, S.-C. Liu, and T. Delbruck, "Ddd20 end-to-end event camera driving dataset: Fusing frames and events with deep learning for improved steering prediction," in *IEEE Intelligent Transportation Systems Conference*. IEEE, 2020, pp. 1–6.
- [43] N. Messikommer, C. Fang, M. Gehrig, and D. Scaramuzza, "Data-driven feature tracking for event cameras," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2023, pp. 5642–5651.
- [44] D. Gehrig, M. Rügge, M. Gehrig, J. Hidalgo-Carrió, and D. Scaramuzza, "Combining events and frames using recurrent asynchronous multimodal networks for monocular depth prediction," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2822–2829, 2021.
- [45] S. Mostafavi I., L. Wang, Y.-S. Ho, and K.-J. Y. Yoon, "Event-based high dynamic range image and very high frame rate video generation using conditional generative adversarial networks," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [46] M. Gehrig, W. Aarents, D. Gehrig, and D. Scaramuzza, "Dsec: A stereo event camera dataset for driving scenarios," in *IEEE Robotics and Automation Letters*.
- [47] M. Cannici, M. Ciccone, A. Romanoni, and M. Matteucci, "A differentiable recurrent surface for asynchronous event-based data," in *European Conference on Computer Vision*. Springer, 2020, pp. 136–152.
- [48] D. Gehrig, M. Gehrig, J. Hidalgo-Carrió, and D. Scaramuzza, "Video to Events: Recycling video datasets for event cameras," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [49] Y. Zhou, G. Gallego, H. Rebecq, L. Kneip, H. Li, and D. Scaramuzza, "Semi-dense 3D reconstruction with a stereo event camera," in *European Conference on Computer Vision*, 2018, pp. 242–258.
- [50] R. Mur-Artal and J. D. Tardós, "ORB-SLAM2: An open-source SLAM system for monocular, stereo, and RGB-D cameras," *IEEE Transactions on Robotics*, vol. 33, no. 5, pp. 1255–1262, Oct. 2017.
- [51] C. Campos, R. Elvira, J. J. G. Rodríguez, J. M. Montiel, and J. D. Tardós, "Orb-slam3: An accurate open-source library for visual, visual-inertial, and multimap slam," *IEEE Transactions on Robotics*, vol. 37, no. 6, pp. 1874–1890, 2021.
- [52] J. L. Schönberger and J.-M. Frahm, "Structure-from-motion revisited," in *IEEE Conference on Computer Vision and Pattern Recognition*, June 2016.
- [53] J. L. Schönberger, E. Zheng, M. Pollefeys, and J.-M. Frahm, "Pixel-wise view selection for unstructured multi-view stereo," in *European Conference on Computer Vision*, 2016.
- [54] J. Engel, V. Koltun, and D. Cremers, "Direct sparse odometry," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 3, pp. 611–625, 2017.
- [55] G. Furano, G. Meoni, A. Dunne, D. Moloney, V. Ferlet-Cavrois, A. Tavoularis, J. Byrne, L. Buckley, M. Psarakis, K.-O. Voss *et al.*, "Towards the use of artificial intelligence on the edge in space systems: Challenges and opportunities," *IEEE Aerospace and Electronic Systems Magazine*, vol. 35, no. 12, pp. 44–56, 2020.