Event-based, 6-DOF Camera Tracking for High-Speed Applications

Guillermo Gallego, Jon E.A. Lund, Elias Mueggler, Henri Rebecq, Tobi Delbruck and Davide Scaramuzza

Abstract—In contrast to standard cameras, which produce frames at a fixed rate, event cameras respond asynchronously to pixel-level brightness changes, thus enabling the design of new algorithms for high-speed applications with latencies of microseconds. However, this advantage comes at a cost: because the output is composed by a sequence of events, traditional computer-vision algorithms are not applicable, so that a new paradigm shift is needed. We present an event-based approach for ego-motion estimation, which provides pose updates upon the arrival of each event, thus virtually eliminating latency. Our method is the first work addressing and demonstrating event-based pose tracking in six degrees-of-freedom (DOF) motions in realistic and natural scenes, and it is able to track high-speed motions. The method is successfully evaluated in both indoor and outdoor scenes.

Index Terms—Event-based vision, Pose tracking, Dynamic Vision Sensor, Bayes filter, Asynchronous processing, Conjugate priors.

◆

SUPPLEMENTARY MATERIAL

This paper is accompanied by a video showing the performance of our method in several indoor and outdoor sequences. A high-resolution version of the video is available at https://youtu.be/iZZ77F-hwzs.

1 INTRODUCTION

EVENT cameras [1, p.77], such as the Dynamic Vision Sensor (DVS [2]), are biologically inspired sensors that overcome many limitations of traditional cameras: they respond very fast (within microseconds) to brightness changes, have a very high dynamic range (120 dB vs 60 dB of standard cameras), and require low power and bandwidth (20 mW vs 1.5 W of standard cameras). Such advantages makes these sensors very attractive for low-powered and/or high-speed applications. However, because they convey the visual information in a radically different way than standard cameras (they do not provide grayscale values but only changes in intensity and the output is composed by a sequence of asynchronous events rather than frames), computer-vision algorithms that are conceived for conventional frame-based cameras do not work on event data. Therefore, new methods must be developed to leverage the advantages of event-driven vision [3].

Previous works on event cameras are still at an early stage of development since event cameras have become commercially available only since 2008 [2]. How to exploit the advantages of event cameras (i.e., high speed, low latency, and high dynamic range) is still an open research problem. The challenges we address in this paper are two: i) event-based 6-DOF pose tracking in natural scenes; ii) tracking the pose during very fast motions, e.g., where standard cameras suffer from motion blur, as shown in Fig. 1.

We present a novel probabilistic pose-tracking method for event-based vision sensors in a known environment. It is based on Bayesian filtering theory with three key contributions in the way that the events are processed: i) event-based pose update, meaning that the 6-DOF pose estimate is updated every time an event is generated, at *micro*second time resolution, ii) the design

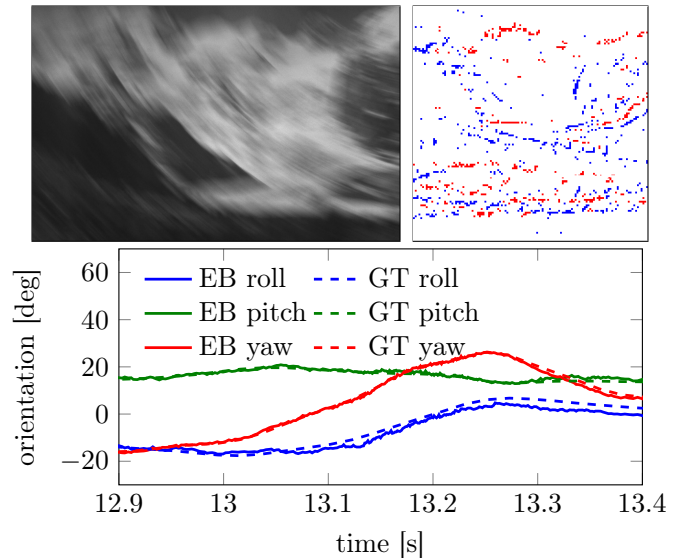• *All authors are with the University of Zurich, Zurich, Switzerland.*

Fig. 1: High-speed motion sequence. Top left: image from a standard camera, suffering from blur due to high-speed motion. Top right: set of asynchronous DVS events in an interval of 3 milliseconds, colored according to polarity. Bottom: estimated poses using our event-based (EB) approach, which provides low latency and high temporal resolution updates. Ground truth (GT) poses are also displayed.

of a sensor likelihood function using a mixture model that takes into account both the event generation process and the presence of noise and outliers (Section 4.3), and iii) the approximation of the posterior distribution of the system by a tractable distribution in the exponential family that is obtained by minimizing the Kullback-Leibler divergence (Section 4.4). The result is a filter adapted to the asynchronous nature of the DVS, which also incorporates an outlier detector that weighs measurements according to their confidence for improved robustness of the pose estimation. The approximation of the posterior distribution allows us to obtain a closed-form solution to the filter update equations

and has the benefit of being computationally efficient, unlike particle filtering. Localization of the DVS is achieved with respect to reference images (and their poses) of the scene. Our method can handle arbitrary, 6-DOF, high-speed motions of the DVS in natural scenes.

The paper is organized as follows: Section 2 reviews related literature on event-based ego-motion estimation methods. Section 3 describes the Dynamic Vision Sensor. The probabilistic approach developed for DVS 6-DOF pose tracking is described in Section 4, and it is empirically evaluated on natural scenes in Section 5. Conclusion and future work are highlighted in Section 6.

2 RELATED WORK ON EVENT-BASED EGO-MOTION ESTIMATION

The first work on pose tracking with a DVS was presented in [4]. The system design, however, was limited to slow planar motions (i.e., 3 DOF) and planar scenes parallel to the plane of motion consisting of artificial B&W line patterns. The method was extended to 3-D in [5] but relied on an external RGB-D sensor for depth estimation. However, a depth sensor introduces the same bottlenecks that exist in standard frame-based systems: depth measurements are outdated for very fast motions, and the depth sensor is still susceptible to motion blur.

In our previous work [6], a standard grayscale camera was attached to a DVS to estimate the small displacement between the current event and the previous frame of the standard camera. The system was developed for planar motion and artificial B&W striped background. This was due to the sensor likelihood being proportional to the magnitude of the image gradient, thus favoring scenes where large brightness gradients are the source of most of the event data. Because of the reliance on a standard camera, the system was again susceptible to motion blur and therefore limited to slow motions.

An event-based algorithm to track the 6-DOF pose of a DVS alone and during very high-speed motion was presented in [7]. However, the method was developed specifically for artificial, B&W line-based maps. Indeed, the system worked by minimizing the point-to-line reprojection error.

Estimation of the 3-D orientation of a DVS to generate high-resolution panoramas of natural scenes was presented in [8]. However, the system was restricted to rotational motions, and, thus, did not account for translation and depth.

Contrarily to all previous works, the approach we present in this paper tackles full 6-DOF motions, does not rely on external sensors, can handle arbitrary fast motions, and is not restricted to specific texture or artificial scenes.

3 EVENT-BASED CAMERAS. THE DYNAMIC VISION SENSOR (DVS)

Event-based vision constitutes a paradigm shift from conventional (e.g., frame-based) vision. In standard cameras, pixels are acquired and transmitted simultaneously at fixed rates; this is the case of both global-shutter or rolling-shutter sensors. Such sensors provide little information about the scene in the "blind time" between consecutive images. Instead, event-based cameras such as the DVS [2] (Fig. 2a) have independent pixels that respond asynchronously to relative contrast changes. If $I(\mathbf{u}, t)$ is the intensity sensed at a pixel $\mathbf{u} = (x, y)^\top$ of the DVS, an event

(a) The Dynamic Vision Sensor (DVS).
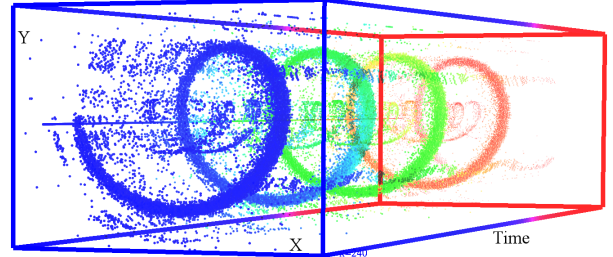
(b) Visualization of the output of a DVS (event stream) while viewing a rotating scene, which generates a spiral-like structure in space-time. Events are represented by colored dots, from red (far in time) to blue (close in time). Event polarity is not displayed. Noise is visible by isolated points.

Fig. 2: Event-based camera.

is generated if the temporal visual contrast (in log scale) exceeds a nominal threshold C_{th}:

$$\Delta \ln I := \ln I(\mathbf{u}, t) - \ln I(\mathbf{u}, t - \Delta t) \gtrless C_{\text{th}}, \qquad (1)$$

where Δt is the time since the last event was generated at the same pixel. Different thresholds may be specified for the cases of contrast increase (C_{th}^+) or decrease (C_{th}^-). An event $e = (x, y, t, p)$ conveys the spatio-temporal coordinates and sign (i.e., polarity) of the brightness change, with $p = +1$ (ON-event: $\Delta \ln I > C_{\text{th}}^+$) or $p = -1$ (OFF-event: $\Delta \ln I < C_{\text{th}}^-$). Events are time-stamped with microsecond resolution and transmitted asynchronously when they occur, with very low latency. A sample output of the DVS is shown in Fig. 2b. Another advantage of the DVS is its very high dynamic range (120 dB), which notably exceeds the 60 dB of high-quality, conventional frame-based cameras. This is a consequence of events triggering on log-intensity changes (1) instead of absolute intensity. The spatial resolution of the DVS is 128×128 pixels, but newer sensors, such as the Dynamic and Active-pixel VIsion Sensor (DAVIS) [9], and color DAVIS (C-DAVIS) [10] will have higher resolution (640×480 pixels), thus overcoming current limitations.

4 PROBABILISTIC APPROACH

Consider a DVS moving in a known static scene. The map of the scene is described by a sparse set of reference images $\{I_l^r\}_{l=1}^{N_r}$, poses $\{\boldsymbol{\xi}_l\}_{l=1}^{N_r}$, and depth map(s). Suppose that an initial guess of the location of the DVS in the scene is also known. The problem we face is that of exploiting the information conveyed by the event stream to track the pose of the DVS in the scene. Our goal is

to handle arbitrary 6-DOF, high-speed motions of the DVS in realistic (i.e., natural) scenes.

We design a robust filter combining the principles of Bayesian estimation, posterior approximation, and exponential family distributions with a sensor model that accounts for outlier observations. In addition to tracking the kinematic state of the DVS, the filter also estimates some sensor parameters automatically (e.g., event triggering threshold C_{th}) that would otherwise be difficult to tune manually. [1]

The outline of this section is as follows. First, the problem is formulated as a marginalized posterior estimation problem in a Bayesian framework. Then, the motion model and the measurement model (a robust likelihood function that can handle both good events and outliers) are presented. Finally, the filter equations that update the parameters of an approximate distribution to the posterior probability distribution are derived.

4.1 Bayesian Filtering

We model the problem as a time-evolving system whose state s consists of the kinematic description of the DVS as well as sensor and inlier/outlier parameters. More specifically,

$$s = (\boldsymbol{\xi}_c, \boldsymbol{\xi}_i, \boldsymbol{\xi}_j, C_{\mathrm{th}}, \pi_m, \sigma_m^2)^\top, \qquad (2)$$

where $\boldsymbol{\xi}_c$ is the current pose of the DVS (at the time of the event, t in (1)), $\boldsymbol{\xi}_i$ and $\boldsymbol{\xi}_j$ are two poses along the DVS trajectory that are used to interpolate the pose of the last event at the same pixel (time $t - \Delta t$ in (1)), C_{th} is the contrast threshold, and π_m and σ_m^2 are the inlier parameters of the sensor model, which is explained in Section 4.3.2.

Let the state of the system at time t_k be s_k, and let the sequence of all past observations (up to time t_k) be $o_{1:k}$, where o_k is the current observation (i.e., the latest event).

Our knowledge of the system state is contained in the posterior probability distribution $p(s_k|o_{1:k})$, also known as *belief* [11, p.27], which is the marginalized distribution of the smoothing problem $p(s_{1:k}|o_{1:k})$. The Bayes filter recursively estimates the system state from the observations in two steps: prediction and correction. The correction step updates the posterior by:

$$p(s_k|o_{1:k}) \propto p(o_k|s_k)p(s_k|o_{1:k-1}), \qquad (3)$$

where $p(o_k|s_k)$ is the likelihood function (sensor model) and we used independence of the events given the state. The prediction step, defined by

$$p(s_k|o_{1:k-1}) = \int p(s_k|s_{k-1})p(s_{k-1}|o_{1:k-1})ds_{k-1}, \qquad (4)$$

incorporates the motion model $p(s_k|s_{k-1})$ from t_{k-1} to t_k.

We incorporate in our state vector not only the current DVS $\boldsymbol{\xi}_c^k$ pose but also the other relevant poses for contrast calculation (poses $\boldsymbol{\xi}_i^k, \boldsymbol{\xi}_j^k$ in (2)), so that we may use the filter to partially correct errors of already estimated poses. Past events that are affected by the previous pose are not re-evaluated, but future events that reference back to such time will have better previous-pose estimates.

To have a computationally feasible filter, we approximate the posterior (3) by a tractable distribution with parameters η_{k-1} that condense the history of events $o_{1:k-1}$,

$$p(s_k|o_{1:k}) \approx q(s_k; \eta_k). \qquad (5)$$

---

1. Today's event-based cameras, such as the DVS [2] or the DAVIS [9], have almost a dozen tuning parameters that are neither independent nor linear.

Assuming a motion model with slowly varying zero-mean random diffusion, so that most updates of the state are due to the events, the recursion on the approximate posterior becomes, combining (3)-(5),

$$q(s_k; \eta_k) \approx C\, p(o_k|s_k)q(s_k; \eta_{k-1}) \qquad (6)$$

for some normalizing constant C. The approximate posterior q is computed by minimization of the Kullback-Leibler (KL) divergence between both sides of (6). As tractable distribution we choose one in the exponential family because they are very flexible and have nice properties for sequential Bayes estimation. The KL minimization gives the update equations for the parameters of the approximate posterior.

4.2 Motion model

The diffusion process leaves the state mean unchanged and propagates the covariance. How much process noise is added to the evolving state is determined by the trace of the covariance matrix (sum of the eigenvalues): each incoming event adds white noise to the covariance diagonal, thus increasing its trace, up to some allowed maximum. This works gracefully across many motion speeds. More specifically, we used a maximum standard deviation of 0.03 for poses parametrized in normalized twist coordinates (with translation in units relative to the mean scene depth), to factor out the metric scale in the diffusion process.

4.3 Measurement Model

Here we elaborate on the choice of likelihood function $p(o_k|s_k)$ in (6) that is used to model the DVS events. Our contributions are, starting from an ideal sensor model, *i*) to define a dimensionless implicit function based on the contrast residual to measure how well the DVS pose and the a priori information (e.g., a map of the scene) explain an event (Section 4.3.1), and *ii*) to build upon such measurement function taking into account noise and outliers, yielding a mixture model for the likelihood function (Section 4.3.2).

4.3.1 Ideal Sensor Model

In a noise-free scenario, an event is triggered as soon as the temporal contrast reaches the threshold (1). Such a measurement would satisfy $\Delta \ln I - C_{\mathrm{th}} = 0$. For simplicity, let us assume that the polarity has already been taken into account to select the appropriate threshold $C_{\mathrm{th}}^+, C_{\mathrm{th}}^-$. Defining the measurement function

$$M := \frac{\Delta \ln I}{C_{\mathrm{th}}} - 1, \qquad (7)$$

the event-generation condition becomes $M = 0$ in a dimensionless formulation. Assuming a prediction of the temporal contrast is generated using the system state, $\Delta \ln I(s_k)$, then (7) depends on both the system state and the observation, $M(o_k, s_k)$. More precisely, denoting by

$$\tilde{s} = (\boldsymbol{\xi}_c, \boldsymbol{\xi}_i, \boldsymbol{\xi}_j, C_{\mathrm{th}})^\top, \qquad (8)$$

the part of the state (2) needed to compute (7), we have $M(o_k, \tilde{s}_k)$. The likelihood function that characterizes such an ideal sensor model is

$$p(o_k|s_k) = \delta(M(o_k, \tilde{s}_k)), \qquad (9)$$

where δ is the Dirac delta distribution.

All deviations from ideal conditions can be collectively modeled by a noise term in the likelihood function. Hence, a more realistic yet simple choice than (9) that is also supported by the bell-shaped form of the threshold variations observed in the DVS [2] is a Gaussian distribution,

$$p(o_k|s_k) = \mathcal{N}(M(o_k, \tilde{s}_k); 0, \sigma_m^2). \tag{10}$$

Most previous works in the literature do not consider an implicit measurement function (7) or Gaussian model (10) based on the contrast residual. Instead, they use explicit measurement functions that evaluate the goodness of fit of the event either in the spatial domain (reprojection error) [4], [7] or in the temporal domain (event-rate error), e.g., image reconstruction thread of [8], assuming Gaussian errors. Our measurement function (7) is based on the event-generation process and combines in a scalar quantity all the information contained in an event (space-time and polarity) to provide a measure of its fit to a given state and a priori information. However, models based on a single Gaussian distribution (10) are very susceptible to outliers. Therefore, we opt for a mixture model to explicitly account for them, as explained next.

4.3.2 Resilient Sensor Model. Likelihood Function

Based on the empirical observation that there is a significant amount of outliers in the event stream, we propose a likelihood function consisting of a normal-uniform mixture model. This model is typical of robust sensor fusion problems [12], where the output of the sensor is modeled as a distribution that mixes a good measurement (normal) with a bad one (uniform):

$$p(o_k|s_k) = \pi_m \mathcal{N}(M(o_k, \tilde{s}_k); 0, \sigma_m^2) \tag{11}$$
$$+ (1 - \pi_m)\mathcal{U}(M(o_k, \tilde{s}_k); M_{\min}, M_{\max}),$$

where π_m is the inlier probability (and $(1 - \pi_m)$ is the outlier probability). Inliers are normally distributed around 0 with variance σ_m^2. Outliers are uniformly distributed over a known interval $[M_{\min}, M_{\max}]$. The measurement parameters σ_m^2 and π_m are considered unknown and are collected in the state vector s_k to be estimated.

To evaluate $M(o_k, \tilde{s}_k)$, we need to compute the contrast $\Delta \ln I(\tilde{s}_k)$ in (7). We do so based on a known reference image I^r (and pose) and both relevant DVS poses for contrast calculation. Assuming the depth of the scene is known, the point $\mathbf{u}'$ in the reference image corresponding to the event location $\mathbf{u}$ in the DVS can be computed (with calibrated cameras). Then, we transfer intensities $I(\mathbf{u}, t) \approx I^r(\mathbf{u}'(t))$ to approximate $\Delta \ln I \approx \ln I^r(\mathbf{u}'(t_k)) - \ln I^r(\mathbf{u}'(t_k - \Delta t))$, where t_k is the time of the current event and Δt is the time since the last event at the same pixel. This approach is more accurate than linearizing $\Delta \ln I$. We assume that for a small pose change there is a relatively large number of events from different pixels. In this case the information contribution of a new event to an old pose will be negligible, and the new event will mostly contribute to the most recent pose.

Next, we linearize the measurement function in (11) around the expected state $\bar{s}_k = E_{p(s_k|o_{1:k-1})}[s_k]$, prior to incorporating the measurement correction:

$$M(o_k, \tilde{s}_k) \approx M(o_k, \bar{\bar{s}}_k) + \nabla_{\tilde{s}} M(o_k, \bar{\bar{s}}_k) \cdot (\tilde{s}_k - \bar{\bar{s}}_k)$$
$$= \bar{M}_k + J_k \cdot \Delta \tilde{s}_k, \tag{12}$$

where $\bar{M}_k$ and J_k are the predicted measurement and Jacobian at $\bar{\bar{s}}_k$, respectively. Substituting in (11) we get:

$$p(o_k|s_k) = \pi_m \mathcal{N}(\bar{M}_k + J_k \cdot \Delta \tilde{s}_k; 0, \sigma_m^2) + (1 - \pi_m)\mathcal{U}. \tag{13}$$

We assume that the linearization is a good approximation to the original measurement function.

Finally, we may re-write the likelihood (13) in a more general and convenient form for deriving the filter equations, as a sum of exponential families for the state parameters s_k (see the Appendix):

$$p(o_k|s_k) = \sum_j h(s_k) \exp(\eta_{o,j} \cdot T(s_k) - A_{o,j}). \tag{14}$$

4.4 Posterior Approximation and Filter Equations

Our third contribution pertains to the approximation of the posterior distribution using a tractable distribution. For this, we consider variational inference theory [13], and choose a distribution in the exponential family as well as conjugate priors, minimizing the relative entropy error in representing the true posterior distribution with our approximate distribution, as we explain next.

Exponential families of distributions are useful in Bayesian estimation because they have *conjugate priors* [13]: if a given distribution is multiplied by a suitable prior, the resulting posterior has the same form as the prior. Such a prior is called a conjugate prior for the given distribution. The prior of a distribution in the exponential family is also in the exponential family, which clearly simplifies recursion. A mixture distribution like (14) does not, however, have a conjugate prior: the product of the likelihood and a prior from the exponential family is not in the family. Instead, the number of terms of the posterior doubles for each new measurement, making it unmanageable. Nevertheless, for tractability and flexibility, we choose as conjugate prior a distribution in the exponential family and approximate the product, in the sense of the Kullback-Leibler (KL) divergence [14], by a distribution of the same form, as expressed by (6). This choice of prior is optimal if either the uniform or the normal terms of the likelihood dominates the mixture; we expect that small deviations from this still gives good approximations.

Letting the KL divergence (or relative entropy) from a distribution f to a distribution g be

$$D_{\mathrm{KL}}(f\|g) = \int f(x) \ln \frac{f(x)}{g(x)} dx, \tag{15}$$

which measures the information loss in representing distribution f by means of g, the posterior parameters η_k are calculated by minimization of the KL divergence from the distribution on the right hand side of (6) to the approximating posterior (left hand side of (6)):

$$\eta_k = \arg \min_\eta D_{\mathrm{KL}}\big(C\, p(o_k|s_k)q(s_k; \eta_{k-1})\|q(s_k; \eta)\big).$$

It can be shown [13, p.505] that for g in the exponential family, the necessary optimality condition $\nabla_\eta D_{\mathrm{KL}}(f\|g) = 0$ gives the system of equations (in η)

$$E_{f(s)}[T(s)] = E_{g(s)}[T(s)], \tag{16}$$

i.e., the expected sufficient statistics must match. Additionally, the right hand side of (16) is $\nabla A \equiv \nabla_\eta A = E_{g(s)}[T(s)]$ since g is in the exponential family. In our case, $g \equiv q(s_k; \eta)$, $f \propto p(o_k|s_k)q(s_k; \eta_{k-1})$ and (16) can also be written in terms of the

Algorithm 1 Event-based pose tracking

Initialize state variables (DVS pose, contrast threshold, inlier ratio). Then, for each incoming event:
- propagate state covariance (zero-mean random diffusion)
- transfer the event to the map, compute the depth and evaluate the measurement function M function (10).
- compute K_k in (16), the inlier probability π_m, the weight w_k in (17), and the gain $w_k K_k$.
- update filter variables and covariance (e.g., (18)-(19)).

parameters of (14) [(3)-(6) in the Appendix], the log-normalizer A and its gradient:

$$0 = \sum_j \exp\big(A(\eta_{o,j} + \eta_{k-1}) - A(\eta_{k-1}) - A_{o,j})\big)$$
$$\times \big(\nabla A(\eta_{o,j} + \eta_{k-1}) - \nabla A(\eta)\big). \qquad (17)$$

Equation (17) describes a system of equations that can be solved for η, yielding the update formula for η_k in terms of η_{k-1} and the current event o_k. For a multivariate Gaussian distribution over the DVS poses, explicit calculation of all update rules has the simple form of an Extended Kalman Filter (EKF) [11], [15] weighted by the inlier probability of that event:

$$K_k = P_k J_k^\top \big(J_k P_k J_k^\top + \sigma_m^2\big)^{-1} \qquad (18)$$

$$w_k = \frac{\pi_m \mathcal{N}(\bar{M}_k; 0, \sigma_m^2)}{\pi_m \mathcal{N}(\bar{M}_k; 0, \sigma_m^2) + (1 - \pi_m)\mathcal{U}} \qquad (19)$$

$$\boldsymbol{\xi}_{k+1} = \boldsymbol{\xi}_k + w_k K_k \bar{M}_k \qquad (20)$$

$$P_{k+1} = (\mathbb{1} - w_k K_k J_k)P_k, \qquad (21)$$

where $\mathbb{1}$ is the identity, $\bar{M}_k$ and J_k are given in (12), $\boldsymbol{\xi}$ are the 6-DOF twist coordinates of the DVS pose, P is the pose covariance matrix, and $w_k K_k$ acts as the Kalman gain. A pseudocode of the approach is outlined in Algorithm 1.

The posterior approximation described in this section allows us to fuse the measurements and update the state-vector efficiently, without keeping multiple hypothesis in the style of particle filters, which would quickly become intractable due to the dimension of the state-vector.

5 EXPERIMENTAL RESULTS

For our pose estimation algorithm to work, it requires an existing photometric map of the scene. As mentioned at the beginning of Section 4, without loss of generality we describe the map in terms of depth maps with associated reference frames. These can be obtained from a previous mapping stage by means of classical dense reconstruction approaches using standard cameras (e.g., DTAM [16] or REMODE [17]) or even using a DVS (future research). For simplicity, in this work we use reference images from a standard camera and assume nearly planar scenes; however, this is not a constraint of the method.

We evaluated the performance of our algorithm on several indoor and outdoor datasets. The datasets also contain fast motion with excitations in all six degrees of freedom (DOF).

First, we assessed the accuracy of our method against ground truth obtained by a motion-capture system. We placed the DVS in front of a scene consisting of rocks (Fig. 3) at a mean scene depth of 60 cm and recorded eight datasets. Fig. 3 shows the position and orientation errors (i.e., difference between the estimated ones
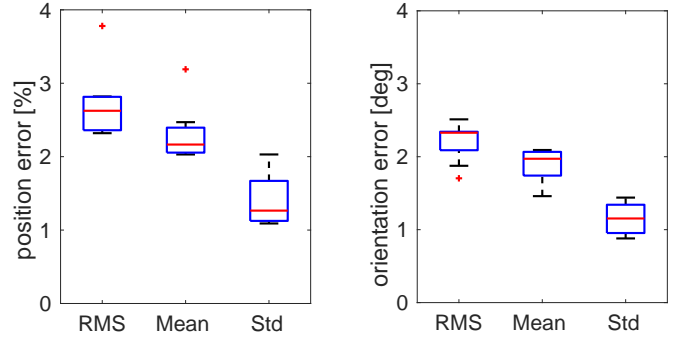

Fig. 4: Error in position (relative to a mean scene depth of 60 cm) and orientation (in degrees) of the trajectories recovered by our method for *all* "rocks" sequences (ground truth is given by a motion capture system). We provide box plots of the root-mean-square (RMS) errors, the mean errors and the standard deviation (Std) of the errors.
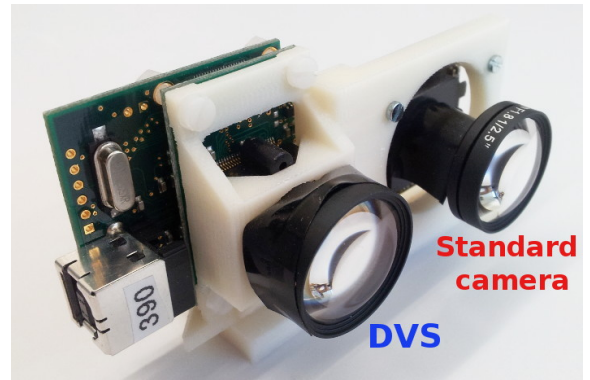

Fig. 5: A DVS and a standard camera mounted on a rig. The standard camera was only used for comparison.

and ground truth)[2] for one of the datasets, while Fig. 8 shows the actual values of the estimated trajectory and ground truth over time. Fig. 4 summarizes the errors of the estimated trajectories for all sequences. The mean RMS errors in position and orientation are 1.63 cm and 2.21°, respectively, while the mean and standard deviations of the position and orientation errors are $\mu = 1.38$ cm, $\sigma = 0.84$ cm, and $\mu = 1.89°$, $\sigma = 1.15°$, respectively. Notice that the RMS position error corresponds to 2.71 % of the average scene depth, which is very good despite the poor spatial resolution of a DVS.

Next, we show the results of our algorithm in three outdoor datasets. In this case, ground truth is obtained via pose tracking with a standard camera running the SVO visual-odometry algorithm, which is available open source [19].[3]

To acquire accurate data for the evaluation, we rigidly mounted the DVS and the standard camera on a rig (see Fig. 5), and the same lens model was mounted on both sensors. The DVS has a spatial resolution of 128×128 pixels and operates asynchronously, in the microsecond scale. The standard camera is a global shutter MatrixVision Bluefox camera with a resolution of 752×480 pixels and a frame rate of up to 90 Hz. Both camera and DVS were

---

2. The rotation error is measured using the angle of their relative rotation (i.e., geodesic distance in $SO(3)$ [18]).

3. SVO reports relative errors of 0.1%; hence it is justified to use its pose estimates as ground truth.
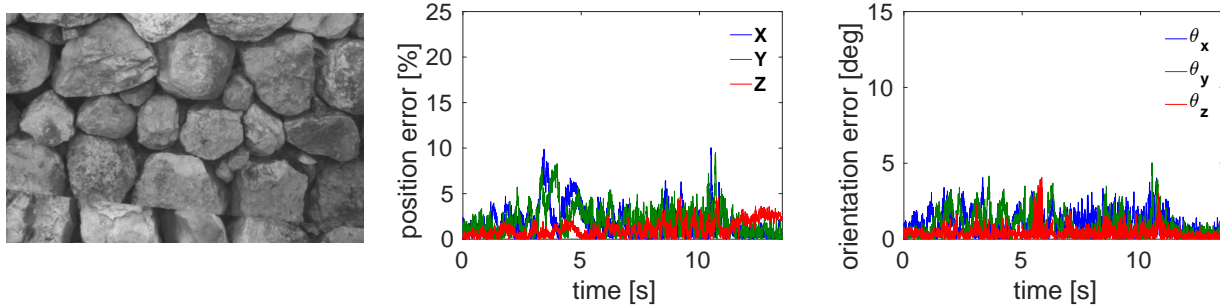
Fig. 3: Error plots in position (relative to a mean scene depth of $60\,\text{cm}$) and in orientation (in degrees) for one of the test sequences with ground truth provided by a motion capture system with sub-millimeter accuracy.
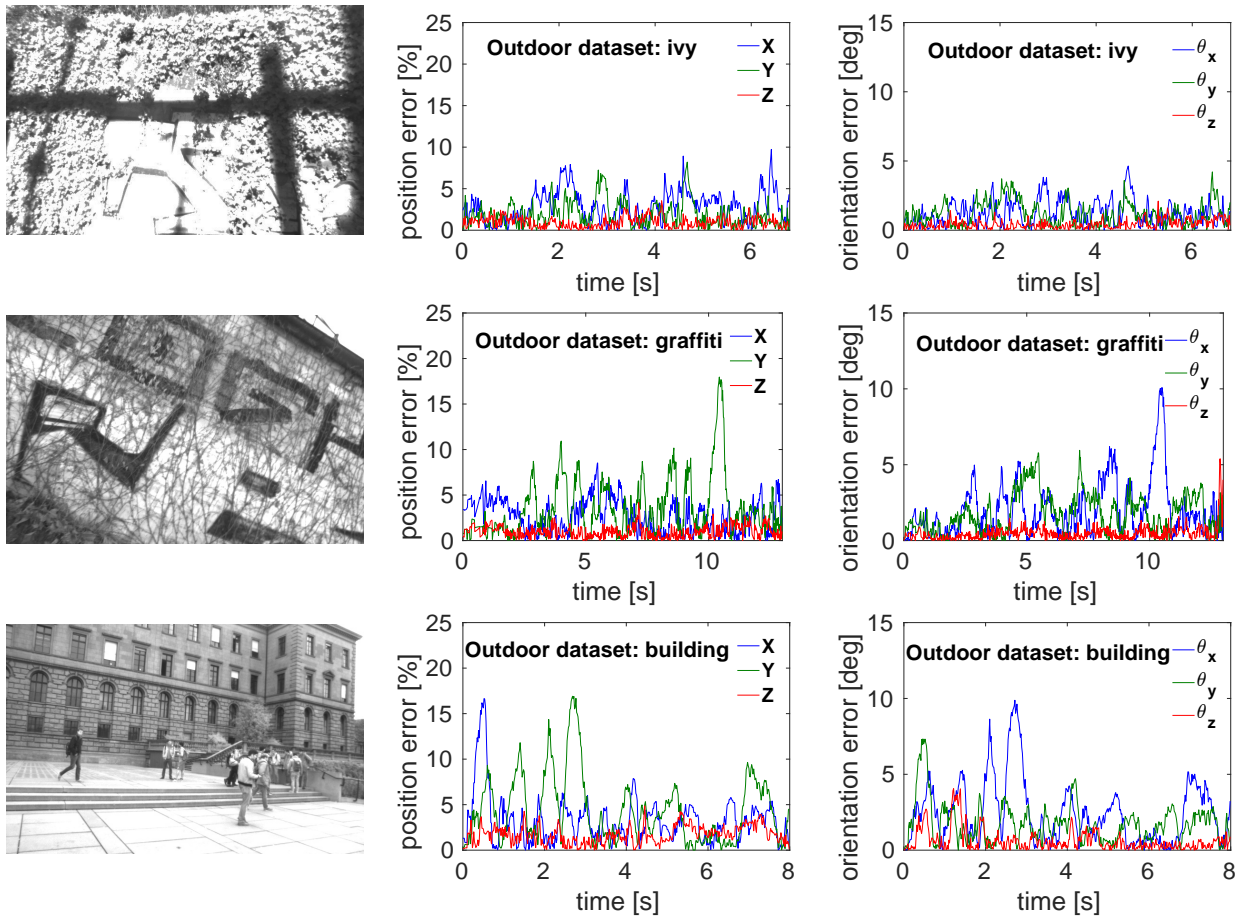

Fig. 6: Error plots in position (2nd column, relative to the mean scene depth) and in orientation (3rd column, in degrees) for three outdoor test sequences (1st column): "ivy", "graffiti", and "building". The mean scene depths are $2.5\,\text{m}$, $3\,\text{m}$, and $30\,\text{m}$, respectively.

calibrated intrinsically and extrinsically.

The three outdoor datasets ("ivy", "graffiti", and "building") were recorded with the DVS-plus-camera rig viewing an ivy, a graffiti covered by some plants, and a building with people moving in front of it, respectively (see Fig. 6, 1st column and accompanying video submission). The rig was moved by hand with increasing speed. All sequences exhibit significant translational and rotational motion. The error plots in position and orientation of all 6-DOFs are given in Fig. 6. The reported error peaks in the "graffiti" and "building" datasets are due to a decrease of overlap between the DVS frustum and the reference map, thus making pose estimation ambiguous for some motions (e.g., Y-translation vs. X-rotation).

Table 1 summarizes the statistics of the pose tracking error for the three outdoor sequences. For the "ivy" dataset, the mean and standard deviation of the position error are $9.93\,\text{cm}$ and $4.60\,\text{cm}$, which correspond to $3.97\,\%$ and $1.84\,\%$ of the average scene depth ($2.5\,\text{m}$), respectively. The mean and standard deviation of

TABLE 1: Error measurements of three outdoor sequences. Translation errors are relative (i.e., scaled by the mean scene depth).

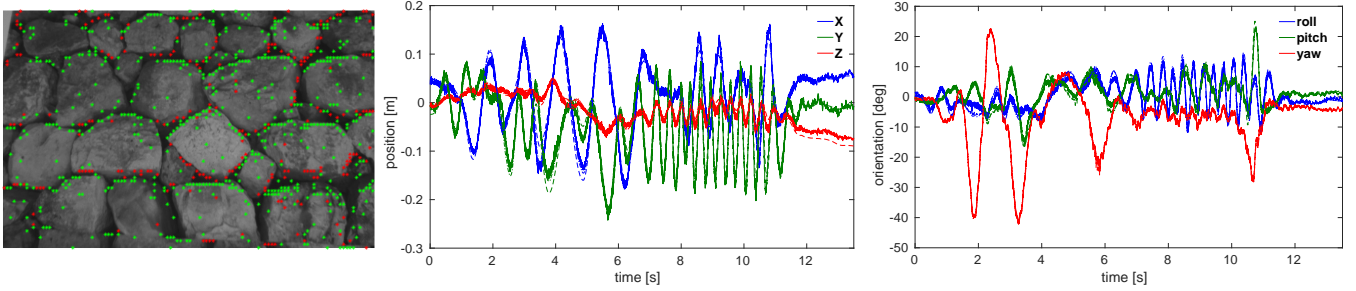|	Position error [%]			Orientation error [°]		
	RMS	μ	σ	RMS	μ	σ
"ivy"	4.37	3.97	1.84	2.21	2.00	0.94
"graffiti"	5.88	5.23	2.70	3.58	3.09	1.80
"building"	7.40	6.47	3.60	3.99	3.43	2.05

Fig. 8: Indoor experiment with 6-DOF motion. Left: Image of the standard camera overlaid with events (during mild motion). Events are displayed in red and green, according to polarity. Estimated position (center) and orientation (right) from our event-based algorithm (solid line), a frame-based method (dash-dot line) and ground truth (dashed line) from a motion capture system.
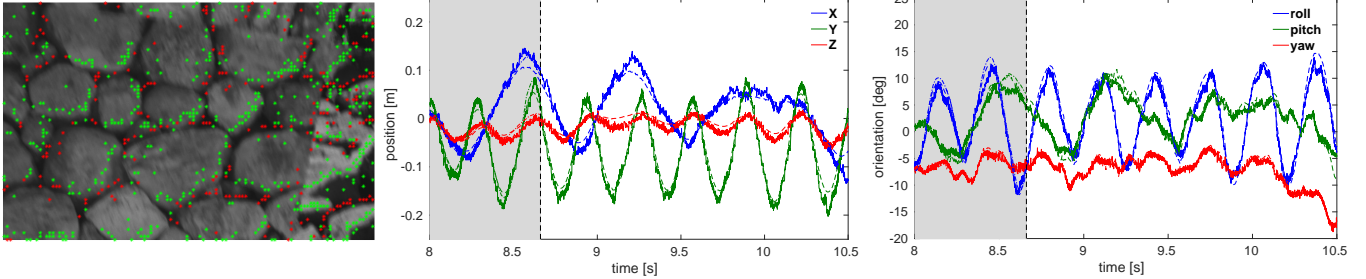

Fig. 9: Zoom of Fig. 8. Left: Image of the standard camera overlaid with events (red and green points, according to polarity) during high-speed motion. Center and right: estimated trajectories. Due to the very high temporal resolution, our algorithm can still track the motion even when the images of the standard camera are sufficiently blurred so that the frame-based method (FB) failed. The event-based method (EB) provides pose updates even in high-speed motions, whereas the frame-based method loses track (it only provides pose updates in the region marked with the shaded area, then it fails).
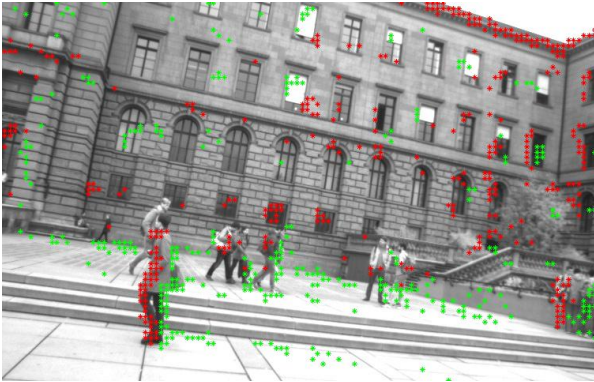

Fig. 7: The algorithm is able to track the DVS pose in spite of the considerable amount of events generated by moving objects (e.g., people) in the scene.

the orientation error are $2.0°$ and $0.94°$, respectively. For the "building" dataset, which presents the largest errors, the mean and standard deviation of the orientation error are $3.43°$ and $2.05°$, respectively, while, in position error, the corresponding figures are $1.94\,\mathrm{m}$ and $1.08\,\mathrm{m}$, that correspond to $6.47\,\%$ and $3.60\,\%$ of the average scene depth ($30\,\mathrm{m}$), respectively.

As reported by the small errors in Table 1, overall our event-based algorithm is able to accurately track the pose of the DVS also outdoors. This shows that, in spite of the limited resolution of the DVS (128×128 pixels), the accuracy of the results provided by our event-based algorithm is comparable to that obtained by a standard camera processing $20\times$ higher resolution images

(752×480 pixels). This is made possible by the DVS temporal resolution being ten thousand times larger than the standard camera. We expect that the results provided by our approach would be even more accurate with the next generation of event-based sensors currently being developed [9], [10], which will have higher spatial resolution (640×480 pixels). Finally, observe that in the "building" sequence (Fig. 6, bottom row), our method gracefully tracks the pose in spite of the considerable amount of events generated by moving objects (e.g., people) in the scene (see Fig. 7).

5.1 Tracking during high-speed motions

In addition to the error plots in Fig. 3, we show in Fig. 8 the actual values of the trajectories (position and orientation) acquired by the motion capture system (dashed line) and estimated by the event-based method (solid line) and frame-based method (dash-dot). Notice that they are all are almost indistinguishable relative to the amplitude of the motion excitation, which gives a better appreciation of the small errors reported in Figs. 3 and 4.

Figure 9 shows a magnified version of the estimated trajectories during high-speed motions (occurring at $t \geq 7\,\mathrm{s}$ in Fig. 8). The frame-based method is able to track in the shaded region, up to $t \approx 8.66\,\mathrm{s}$ (indicated by a vertical dashed line), at which point it loses tracking due to motion blur, while our event-based method continues to accurately estimate the pose.

6 CONCLUSION

We have presented a novel, event-based probabilistic approach to track the pose of an arbitrarily moving event camera in 6-DOF in

natural scenes. Our approach follows a Bayesian filtering methodology: the sensor model is given by a mixture-model likelihood that takes into account both the event-generation process and the presence of noise and outliers; the posterior distribution of the system state is approximated according to the relative-entropy criterion using distributions in the exponential family and conjugate priors. This yields a robust EKF-like filter that provides pose updates for every incoming event, at microsecond time resolution.

We have compared our method against ground truth provided by a motion capture system or a state-of-the-art frame-based pose-tracking pipeline. The experiments revealed that the proposed method accurately tracks the pose of the event-based camera. In future, we plan to extend the proposed framework to a full event-based SLAM (Simultaneous Localization and Mapping) in 6-DOF.

APPENDIX A
REWRITING THE LIKELIHOOD FUNCTION

A distribution in the exponential family can be written as

$$p(x; \eta) = h(x) \exp\left(\eta \cdot T(x) - A(\eta)\right), \qquad (22)$$

where η are the natural parameters, $T(x)$ are the sufficient statistics of x, $A(\eta)$ is the log-normalizer, and $h(x)$ is the base measure.

The likelihood (13) can be rewritten as:

$$p(o_k | s_k) = \frac{1}{\sqrt{2\pi}} \exp(\ln(\pi_m) - \ln(\sigma_m)) \qquad (23)$$
$$- \frac{1}{2}\left[J_k^i J_k^j \frac{\tilde{s}_k^i \tilde{s}_k^j}{\sigma_m^2} + 2\bar{M}_k J_k^i \frac{\tilde{s}_k^i}{\sigma_m^2} + \frac{\bar{M}_k^2}{\sigma_m^2} \right]$$
$$+ \exp(\ln((1 - \pi_m)/(M_{\max} - M_{\min}))),$$

where we use the Einstein summation convention for the indices of $J_k = (J_k^i)$ and $\tilde{s}_k = (\tilde{s}_k^i)$. Collecting the sufficient statistics into

$$T(s_k) = \left[\frac{\tilde{s}_k^i \tilde{s}_k^j}{\sigma_m^2}, \frac{\tilde{s}_k^i}{\sigma_m^2}, \frac{1}{\sigma_m^2}, \ln(\sigma_m), \ln(\pi_m), \ln(1 - \pi_m) \right],$$

the likelihood can be conveniently rewritten as a sum of two exponential families (14), $j = 1, 2$, with $h(s) = 1$,

$$\eta_{o,1} = \left[-\frac{1}{2} J_k^i J_k^j, -\bar{M}_k J_k^i, -\frac{1}{2}\bar{M}_k^2, -1, 1, 0 \right] \qquad (24)$$

$$\eta_{o,2} = [0_{ij}, 0_i, 0, 0, 1] \qquad (25)$$

$$A_{o,1} = \ln\sqrt{2\pi} \qquad (26)$$

$$A_{o,2} = -\ln(M_{\max} - M_{\min}). \qquad (27)$$

REFERENCES

[1] A. Belbachir, *Smart Cameras*. Springer US, 2009. 1
[2] P. Lichtsteiner, C. Posch, and T. Delbruck, "A 128x128 120 dB 15 μs latency asynchronous temporal contrast vision sensor," *IEEE J. of Solid-State Circuits*, vol. 43, no. 2, pp. 566–576, 2008. 1, 2, 3, 4
[3] S. Schraml, A. Belbachir, and H. Bischof, "Event-driven stereo matching for real-time 3D panoramic vision," in *Proc. IEEE Int. Conf. Computer Vision and Pattern Recognition*, June 2015, pp. 466–474. 1
[4] D. Weikersdorfer and J. Conradt, "Event-based Particle Filtering for Robot Self-Localization," in *IEEE Intl. Conf. on Robotics and Biomimetics (ROBIO)*, 2012. 2, 4
[5] D. Weikersdorfer, D. B. Adrian, D. Cremers, and J. Conradt, "Event-based 3D SLAM with a depth-augmented dynamic vision sensor," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, Jun. 2014, pp. 359–364. 2
[6] A. Censi and D. Scaramuzza, "Low-Latency Event-Based Visual Odometry," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2014. 2
[7] E. Mueggler, B. Huber, and D. Scaramuzza, "Event-based, 6-DOF Pose Tracking for High-Speed Maneuvers," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2014. 2, 4
[8] H. Kim, A. Handa, R. Benosman, S.-H. Ieng, and A. J. Davison, "Simultaneous Mosaicing and Tracking with an Event Camera," in *British Machine Vision Conf. (BMVC)*, 2014. 2, 4
[9] C. Brandli, R. Berner, M. Yang, S.-C. Liu, and T. Delbruck, "A 240x180 130dB 3us Latency Global Shutter Spatiotemporal Vision Sensor," *IEEE J. of Solid-State Circuits*, 2014. 2, 3, 7
[10] C. Li, C. Brandli, R. Berner, H. Liu, M. Yang, S. Liu, and T. Delbruck, "An RGBW Color VGA Rolling and Global Shutter Dynamic and Active-Pixel Vision Sensor," in *International Image Sensor Workshop (IISW)*, Vaals, Netherlands, June 2015. 2, 7
[11] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. The MIT Press, Cambridge, MA, 2005. 3, 5
[12] G. Vogiatzis and C. Hernández, "Video-based, Real-Time Multi View Stereo," *Image and Vision Computing*, vol. 29, no. 7, 2011. 4
[13] C. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006. 4
[14] S. Kullback and R. A. Leibler, "On Information and Sufficiency," *Ann. Math. Statist.*, vol. 22, no. 1, pp. 79–86, 03 1951. 4
[15] R. Kalman, "A New Approach to Linear Filtering and Prediction Problems," *ASME Journal of Basic Engineering*, 1960. 5
[16] R. Newcombe, S. Lovegrove, and A. Davison, "DTAM: Dense tracking and mapping in real-time," in *Intl. Conf. on Computer Vision (ICCV)*, Barcelona, Spain, Nov. 2011, pp. 2320–2327. 5
[17] M. Pizzoli, C. Forster, and D. Scaramuzza, "REMODE: Probabilistic, monocular dense reconstruction in real time," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2014. 5
[18] D. Q. Huynh, "Metrics for 3D Rotations: Comparison and Analysis," *Journal of Mathematical Imaging and Vision*, vol. 35, no. 2, pp. 155–164, 2009. 5
[19] C. Forster, M. Pizzoli, and D. Scaramuzza, "SVO: Fast Semi-Direct Monocular Visual Odometry," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2014. 5