

Gyrento: Event-based Omnidirectional Visual Gyroscope in a Manhattan World - Extended Version

Daniel Rodrigues Da Costa, Pascal Vasseur, Fabio Morbidi

Abstract— In this paper, we study the problem of estimating the orientation of an event omnidirectional camera mounted on a robot and observing 3D parallel lines in a man-made environment (Manhattan world). We present *Gyrento*, the first event-based omnidirectional visual gyroscope. *Gyrento* does not require any initialization, provides certifiably globally optimal solutions, and is scalable, since the size of the nonlinear least-squares cost function is independent of the number of lines. Thanks to the Cayley-Gibbs-Rodrigues parameterization of a 3D rotation, this cost function is a degree-four rational function in three variables, which can be efficiently minimized via off-the-shelf polynomial optimization software. Numerical simulations and real-world experiments with a robot manipulator, as well as an IMU, show the effectiveness of our visual gyroscope and elucidate the impact of camera velocity on the attitude estimation error.

I. INTRODUCTION

A. Motivation and related work

In the last decade, *event cameras* have been increasingly used in computer vision and robotics [1], and they have made serious inroads into real-world applications [2]. However, their potential with agile robots, has yet to be fully exploited. Event cameras asynchronously measure per-pixel brightness changes, and output a stream of events that encode the time, location and sign of these variations. They have several attractive properties, compared to the standard frame-based cameras: high temporal resolution (in the order of microseconds), very high dynamic range (up to 140 dB), low power consumption, and high pixel bandwidth (in the order of kHz), which results in reduced motion blur. Finally, the spatial resolution of consumer-grade event cameras is comparable to that of traditional cameras (1280×720 pixels in Prophesee EVK4 - HD).

An area of active research in event-based vision, is *visual odometry*. This is a challenging task, since event cameras do not output images, and traditional computer vision algorithms cannot be applied directly. Numerous event-based visual odometry algorithms have appeared in the recent literature [3]–[10]. However, with the exception of [3], [4], [7], [9], which only take a stream of events from a monocular camera as input, the other methods rely on an Inertial Measurement Unit (IMU) [5], [6], a second extrinsically-calibrated

The authors are with the MIS laboratory, University of Picardie Jules Verne, 33 rue Saint-Leu, 80039 Amiens Cedex, France. Corresponding author: D. Rodrigues Da Costa, Email: daniel.rodrigues.da.costa@u-picardie.fr

This work was supported by AID (Agence de l’Innovation de Défense), through the research project EVENTO, “*Omnidirectional Event Cameras for High-Speed Robots*” (2021–2024) and by the ANR-FWF EVELOC “*Event-based Visual Localization*” project (ANR-23-CE33-0011).




Figure 1: Three parallel and mutually orthogonal groups of 3D lines $\{L_1, L_2\}$ (red), $\{L_3, L_4, L_5\}$ (green), $\{L_6, L_7, L_8\}$ (blue), are observed by an omnidirectional event camera. *Gyrento* estimates the orientation of the camera frame $\{C\}$ with respect to the Manhattan (or world) frame $\{M\}$.

depth camera [10], or combine events and greyscale frames from a DAVIS sensor [8].

Line features have a simple mathematical representation and they are ubiquitous in urban environments (e.g., doors, windows, sidewalks, lamp posts, etc.). Moving event cameras can easily detect edges in a 3D scene and several methods have been recently proposed for event-based line detection and tracking. In [11], a variant of Line Segment Detection (LSD) [12] algorithm tailored to event data, called ELiSeD, is introduced, while in [13], the authors presented a method based on the detection of event planes in the (x, y, t) space and tracking of these planes over time. On the other hand, in [14], an algorithm for line and segment detection which relies on iterative event-based weighted least-squares fitting, is proposed. In [15], a geometric model of line-generated manifolds is introduced and the linear velocity of an event camera is estimated from given angular rates measured by an IMU. This work has been recently improved in [16] (see also [17], where a stereo event camera and an IMU are used). In [18], the authors described a line-based PTAM (Parallel Tracking and Mapping) algorithm that estimates structure and motion, and produces accurate camera pose estimates in real-time, by using an error-state Kalman filter. The line extraction procedure is based on the Hough transform. Similarly in [19], a trifocal tensor based incidence relation that can be solved in closed-form, is used for direct estimation of camera linear velocity.

Unfortunately, event cameras, as conventional vision sensors, suffer from a narrow field of view, which is especially critical when the camera is mounted on fast-moving robots. Relatively few works in the literature, have addressed this

problem and dealt with *omnidirectional event cameras*. In [20], the authors developed a stereo system consisting of two event cameras mounted on a rotating head, which provides real-time 360° panoramic vision. In [21], a panoramic tracking and mapping algorithm that only relies on the pixel positions of the events in the continuous stream is proposed. In [22] a multi-view stereo method, called EOMVS, is designed to estimate a depth map and reconstruct a 3D scene using an omnidirectional (fisheye) event camera. Finally, fisheye event cameras observing synthetic road scenes from a moving car, are considered in the SynWoodScape dataset [23] as well.

This paper presents *Gyrevento*, a new event-based omnidirectional visual gyroscope¹ (see Fig. 1). The combination of a 360° view and an event-based mechanism offers two distinctive advantages which have not been jointly explored so far, in the literature : the primitives (lines, in our case) always lie within the FoV of the camera (persistence of measurement) and the gyroscope is robust to strong accelerations (no motion blur). Unlike conventional IMUs, our visual gyroscope is drift-free, thus providing more reliable long-term attitude estimates. Our work builds upon a large body of research on vanishing point extraction and attitude estimation in a Manhattan world, for perspective and omnidirectional color cameras [24]–[27]. By using the unified central projection model [28], these results have been adapted to meet the special requirements of an omnidirectional event camera. Gyrevento processes an asynchronous stream of events, and detects and tracks great circles on the unit sphere (the projection of groups of 3D parallel lines in the Manhattan world) to estimate the orientation of the camera frame. A RANSAC-based classifier generates hypotheses for orthogonal vanishing points. Our nonlinear least-squares formulation relies on the Cayley-Gibbs-Rodrigues (CGR) parameterization [29], [30] of a 3D rotation matrix and on the minimization of a multivariate rational function, whose degree is independent of the number of observed 3D lines. Compared to the solution based on the eigendecomposition of a large multiplication matrix in [24], [31], our approach is simpler yet provably optimal (in fact, the minima are certified to be global [32]). In summary, the original contributions of this paper can be listed as follows:

- 1) We propose the first visual gyroscope for event-based omnidirectional cameras only based on the detection and tracking of line features. Gyrevento is initialization-free, provides globally optimal certified solutions and is scalable (it is independent of the number of 3D lines observed by the camera).
- 2) We take advantage of the CGR parameterization of a 3D rotation to define a nonlinear least-squares cost function. This multivariate rational cost function is efficiently minimized with GloptiPoly [33], for global attitude estimation.
- 3) In our real-world experiments with a manipulator, for a given trajectory, we study the effect of camera velocity on the attitude estimation error. The corresponding omnidirectional event datasets have been publicly released.

¹By *visual gyroscope*, we mean an algorithm which exploits the visual information to estimate the 3D orientation of a camera (but *not* its velocity).

We also describe a generic calibration procedure for event central panoramic cameras, which we believe is of independent interest for the robotic and computer vision communities.

The remainder of this paper is organized as follows. In Sect. II, we set the mathematical foundations and briefly introduce some tools for later reference. In Sect. III, the problem studied in the article is formulated, and the different functional blocks of the proposed event-based omnidirectional visual gyroscope are described. The theory is validated via extensive numerical and hardware experiments in Sect. IV. Finally, in Sect. V, the main contributions of the paper are summarized and several avenues for future research are outlined.

Notation: Throughout this paper, \mathbf{I}_n denotes the $n \times n$ identity matrix, $\|\mathbf{x}\|$ the Euclidean norm of vector $\mathbf{x} \in \mathbb{R}^n$ and $\{e_1, e_2, e_3\}$ the canonical basis for \mathbb{R}^3 . The cross product of two vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^3$ is denoted by $\mathbf{u} \times \mathbf{v}$ and $\text{SO}(n) \triangleq \{\mathbf{R} \in \mathbb{R}^{n \times n} : \mathbf{R}^T \mathbf{R} = \mathbf{I}_n, \det(\mathbf{R}) = 1\}$ is the special orthogonal group in dimension n . Finally, $g(\mathbf{x}) \in \mathbb{R}[\mathbf{x}]$ indicates a multivariate polynomial in $\mathbf{x} \in \mathbb{R}^n$.

II. PRELIMINARIES

In this section, we review some useful mathematical tools, that we will use in the rest of this paper.

A. CGR parameterization

Using the Cayley-Gibbs-Rodrigues (CGR) parameterization, a rotation matrix $\mathbf{R} \in \text{SO}(3)$ is expressed as

$$\mathbf{R}(\mathbf{s}) = \frac{\bar{\mathbf{R}}(\mathbf{s})}{1 + \mathbf{s}^T \mathbf{s}}, \quad \bar{\mathbf{R}}(\mathbf{s}) \triangleq (1 - \mathbf{s}^T \mathbf{s})\mathbf{I}_3 + 2[\mathbf{s}]_{\times} + 2\mathbf{s}\mathbf{s}^T, \quad (1)$$

where $\mathbf{s} = [s_1, s_2, s_3]^T$ is the vector of CGR parameters [34, p. 469], and

$$[\mathbf{s}]_{\times} = \begin{bmatrix} 0 & s_3 & -s_2 \\ -s_3 & 0 & s_1 \\ s_2 & -s_1 & 0 \end{bmatrix}, \quad (2)$$

is the corresponding skew-symmetric matrix². The rotation matrix \mathbf{R} can be equivalently expressed as

$$\mathbf{R}(\mathbf{s}) = (\mathbf{I}_3 + [\mathbf{s}]_{\times})(\mathbf{I}_3 - [\mathbf{s}]_{\times})^{-1},$$

which is generally known as Cayley’s formula. The CGR parameterization is closely related to the other well-known attitude representations: unit quaternion and axis-angle [30]. In fact, if $\mathbf{v} = \sin(\vartheta/2)\mathbf{r}$ is the vector (or imaginary) part of the unit quaternion, and $\eta = \cos(\vartheta/2)$ is the scalar (or real) part, with (\mathbf{r}, ϑ) the unit vector and angle of the axis-angle representation, respectively, the vector of CGR parameters is defined as $\mathbf{s} \triangleq \mathbf{v}/\eta = \tan(\vartheta/2)\mathbf{r}$. By consequence,

$$\begin{bmatrix} \mathbf{v} \\ \eta \end{bmatrix} = \frac{1}{\sqrt{1 + \mathbf{s}^T \mathbf{s}}} \begin{bmatrix} \mathbf{s} \\ 1 \end{bmatrix}.$$

If it is desired to solve the *inverse problem* of computing the vector of CGR parameters \mathbf{s} corresponding to a given rotation

²From (2), it follows that $\mathbf{u} \times \mathbf{v} = -[\mathbf{u}]_{\times} \mathbf{v}$ for $\mathbf{u}, \mathbf{v} \in \mathbb{R}^3$. Note that it is more common in the literature to define $[\mathbf{u}]_{\times}$ in order that the sign in the previous equation be reversed, i.e. such that $\mathbf{u} \times \mathbf{v} = [\mathbf{u}]_{\times} \mathbf{v}$.




Figure 2: Great circle G (green) and its unit normal vector \mathbf{n} (red). A generic point on the unit sphere is denoted by \mathbf{p} . For the sake of clarity, only the upper hemisphere is shown in the figure.

matrix $\mathbf{R} = [R_{ij}]$, we can take advantage of the axis-angle representation [35, Sect. 2.5]. In fact, for $\sin \vartheta \neq 0$, we have that $\vartheta = \arccos((\text{trace}(\mathbf{R}) - 1)/2)$ and $\mathbf{r} = \frac{1}{2 \sin \vartheta} [R_{32} - R_{23}, R_{13} - R_{31}, R_{21} - R_{12}]^T$, from which $\mathbf{s} = \tan(\vartheta/2)\mathbf{r}$ is obtained.

Remark 1: (Singularities of CGR representation): The vector of CGR parameters is a minimal representation of orientation, but it has the disadvantage that $\|\mathbf{s}\| \rightarrow \infty$ as $\vartheta \rightarrow \pm\pi$ (non-global representation). Thus, rotations through $\pm\pi$ cannot be represented (except formally as $\pm\infty\mathbf{r}$ [34, p. 470]). \diamond

B. Fitting great circles on the unit sphere

A great circle of a unit sphere (the intersection of the sphere and a plane that passes through the center point of the sphere), can be specified by two parameters α and β . These parameters correspond to the directional angles of the normal vector \mathbf{n} to the plane containing the great circle in the Cartesian coordinate system whose origin is located at the center of the sphere [36, Sect. 4.1], see Fig. 2. A generic point \mathbf{p} on the unit sphere can be expressed as

$$\mathbf{p} = \begin{bmatrix} \sin \theta \cos \varphi \\ \sin \theta \sin \varphi \\ \cos \theta \end{bmatrix},$$

where $\theta \in (0, \pi)$ is the polar angle and $\varphi \in [0, 2\pi)$ is the azimuthal angle. Given a point \mathbf{p} and a great circle G with unit normal vector,

$$\mathbf{n} = \begin{bmatrix} \sin \alpha \cos \beta \\ \sin \alpha \sin \beta \\ \cos \alpha \end{bmatrix},$$

the distance from \mathbf{p} to the plane containing G is given by $\delta = |\mathbf{p}^T \mathbf{n}|$.

Let us now consider N points \mathbf{p}_k , $k \in \{1, \dots, N\}$, on the unit sphere. The great circle fitting problem can be stated as

the problem of finding a plane in order to minimize the sum of squares of distances between the N points and the plane. The objective function to minimize is then

$$V(\mathbf{n}) = \sum_{k=1}^N (\mathbf{p}_k^T \mathbf{n})^2,$$

where \mathbf{n} is the normal vector to the plane containing the best-fit great circle. In other words, the optimal normal \mathbf{n} (in the least-squares sense) is the one which minimizes the sum of squared residuals $\delta_k^2 = (\mathbf{p}_k^T \mathbf{n})^2$. As observed in [37], this problem can be converted into a standard eigenvalue problem. In fact, by introducing the matrix $\mathbf{A} = [\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_N]^T \in \mathbb{R}^{N \times 3}$, we can equivalently rewrite the objective function as

$$V(\mathbf{n}) = \mathbf{n}^T \mathbf{A}^T \mathbf{A} \mathbf{n}. \quad (3)$$

The solution for \mathbf{n} is thus the eigenvector of matrix $\mathbf{B} \triangleq \mathbf{A}^T \mathbf{A} \in \mathbb{R}^{3 \times 3}$ corresponding to its smallest eigenvalue. Having found the unit-norm vector \mathbf{n} , one may readily describe the great circle by its two angles α and β .

Remark 2: (Properties of matrix B): It is easy to verify that matrix \mathbf{B} is symmetric and positive semidefinite, hence the eigenvalues of \mathbf{B} are non-negative real numbers. Moreover, matrix \mathbf{B} can be expressed as the sum of N rank-one matrices, i.e.

$$\mathbf{B} = \sum_{k=1}^N \mathbf{p}_k \mathbf{p}_k^T.$$

This shows that for \mathbf{B} to be full-rank, i.e. of rank 3, at least three different non-antipodal points on the unit sphere, are necessary. \diamond

Remark 3: (Weighted least squares): Additional knowledge about the points on the unit sphere can be incorporated into the least-squares problem by introducing the objective function,

$$V_w(\mathbf{n}) = \sum_{k=1}^N w_k (\mathbf{p}_k^T \mathbf{n})^2 = \mathbf{n}^T \mathbf{A}^T \mathbf{W} \mathbf{A} \mathbf{n},$$

where $\mathbf{W} = \text{diag}(w_1, w_2, \dots, w_N)$ and w_k is a positive weight associated with point \mathbf{p}_k . Similarly to matrix \mathbf{B} , we can now define the matrix $\mathbf{B}_w = \mathbf{A}^T \mathbf{W} \mathbf{A} \in \mathbb{R}^{3 \times 3}$. A natural choice for the weights is $w_k = 1/\varrho_k^2$, $k \in \{1, \dots, N\}$, where ϱ_k^2 is the variance associated with the observation of point \mathbf{p}_k (see Sect. III-B for further discussion on the choice of ϱ_k). \diamond

III. DESIGN OF THE VISUAL GYROSCOPE

In this section, we first formulate the problem and then present the overall architecture of our event-based omnidirectional visual gyroscope (see Fig. 3). The five constituent blocks of our algorithm are described in more detail below.

A. Problem formulation

Consider an event-based omnidirectional camera rigidly attached to a mobile robotic platform (for the sake simplicity, we will assume that the origin of the camera frame coincides with the origin of the robot frame). The robot moves in a man-made environment (Manhattan world), and the event camera observes line features in the surrounding 3D space. Our goal

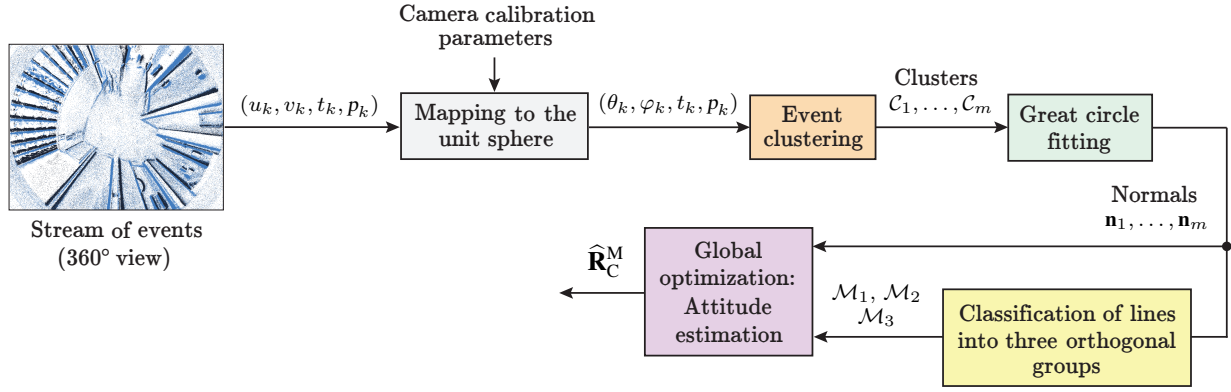


Figure 3: Flowchart of Gyrento.

is to design a visual gyroscope which takes the asynchronous stream of events from the camera as input, and outputs an estimate of the orientation of camera frame $\{C\}$ with respect to the world or Manhattan frame $\{M\}$ (cf. Fig. 1).

A stream of asynchronous events generated by a neuro-morphic camera can be mathematically represented as $\mathbf{e}_k = [u_k, v_k, t_k, p_k]^T$ [1], [14]. This 4-tuple describes an event occurring at time t_k (a.k.a. the *timestamp* of the event) at pixel $[u_k, v_k]^T$. The polarity $p_k \in \{-1, +1\}$ is the sign of brightness change detected by the camera since the last event at pixel $[u_k, v_k]^T$.

B. Mapping to the unit sphere

a) Camera model : Thanks to their 360° FoV on the horizontal plane, omnidirectional cameras are ideal for fast-moving robots since the observed features are always visible. To describe these sensors, we use the unified central projection model of [28], which encompasses a breadth of cameras (perspective, catadioptric, and fisheye [38]). Event $\mathbf{e}_k = [u_k, v_k, t_k, p_k]^T$ is thus mapped into $\mathbf{e}_k^\ominus = [\theta_k, \varphi_k, t_k, p_k]^T$ on the surface of the unit sphere of the unified model (see Fig. 4). The design of our visual gyroscope greatly benefits from the spherical representation. In fact:

- 3D lines in the Manhattan world, become great circles on the unit sphere,
- All parallel lines in the 3D space intersect at two points on the sphere, called *vanishing points* (green squares in Fig. 4). These points can be obtained by computing the cross product of the normal vectors to two great circles.

With reference to Fig. 4, let \mathbf{d}_i be the 3D direction of the i -th 3D line L_i expressed in the Manhattan frame $\{M\}$, and \mathbf{n}_i the normal vector to the corresponding great circle G_i on the unit sphere, expressed in the camera frame $\{C\}$ (\mathbf{n}_i is also referred to as “normalized moment vector” in [24]). Both \mathbf{d}_i and \mathbf{n}_i are unit vectors.

The mapping to the unit sphere requires the camera to be intrinsically calibrated, i.e. the parameters $\{f_u, f_v, u_0, v_0, \xi\}$ need to be known, where (f_u, f_v) are the sizes of unit length in horizontal and vertical pixels, respectively, (u_0, v_0) are the coordinates of the projection of the optical center of the camera onto the image plane in pixels, and ξ is the distance between the unit sphere’s first projection center and the perspective second projection center of the camera [39].

b) Propagation of uncertainty : To determine the variance ϱ_k^2 associated with the spatial location of k -th event \mathbf{e}_k^\ominus on the unit sphere (cf. Remark 3), an uncertainty propagation analysis can be performed. Following [40], [41], let \mathbf{x}'_k be the uncertain observation in the image plane (metric units) of a 3D world point, with associated covariance matrix $\Sigma_{\mathbf{x}'_k \mathbf{x}'_k} \in \mathbb{R}^{2 \times 2}$. The application of the inverse projection map of the camera $\pi^{-1} : \mathbb{R}^2 \rightarrow \mathbb{R}^3$ (cf. [39]), yields

$$\mathbf{x}_k = \pi^{-1}(\mathbf{x}'_k), \quad \Sigma_{\mathbf{x}_k \mathbf{x}_k} \simeq \mathbf{J}_{\pi^{-1}} \Sigma_{\mathbf{x}'_k \mathbf{x}'_k} \mathbf{J}_{\pi^{-1}}^T,$$

where $\Sigma_{\mathbf{x}_k \mathbf{x}_k} \in \mathbb{R}^{3 \times 3}$ is the covariance matrix associated with homogeneous vector \mathbf{x}_k (a rank-two matrix), and $\mathbf{J}_{\pi^{-1}} = \partial \pi^{-1} / \partial \mathbf{x}'_k$ is the Jacobian matrix of the inverse projection map. The subsequent normalization of \mathbf{x}_k leads to

$$\mathbf{p}_k = \frac{\mathbf{x}_k}{\|\mathbf{x}_k\|}, \quad \Sigma_{\mathbf{p}_k \mathbf{p}_k} \simeq \mathbf{J}_k \Sigma_{\mathbf{x}_k \mathbf{x}_k} \mathbf{J}_k^T,$$

where $\Sigma_{\mathbf{p}_k \mathbf{p}_k} \in \mathbb{R}^{3 \times 3}$ is the covariance matrix associated with point \mathbf{p}_k on the unit sphere and $\mathbf{J}_k = \frac{1}{\|\mathbf{x}_k\|} (\mathbf{I}_3 - \mathbf{p}_k \mathbf{p}_k^T)$, [40]. Unfortunately, $\Sigma_{\mathbf{p}_k \mathbf{p}_k}$ remains a singular matrix. There exist then two options for computing the variance ϱ_k^2 .

The *first option* is to set $\varrho_k^2 = \text{trace}(\Sigma_{\mathbf{p}_k \mathbf{p}_k})$ or $\varrho_k^2 = \lambda_{\max}(\Sigma_{\mathbf{p}_k \mathbf{p}_k})$, where $\lambda_{\max}(\Sigma_{\mathbf{p}_k \mathbf{p}_k})$ denotes the maximum eigenvalue of $\Sigma_{\mathbf{p}_k \mathbf{p}_k}$. The *second option* consists in representing the uncertainty by a 2×2 covariance matrix in the tangent plane to the unit sphere at point \mathbf{p}_k (in this way, a minimal representation of uncertainty is obtained). To this end, let the columns of $\mathbf{J}_{\tilde{\mathbf{p}}_k}(\mathbf{p}_k) \in \mathbb{R}^{3 \times 2}$ form an orthonormal basis for the kernel of \mathbf{p}_k^T , i.e. $\mathbf{J}_{\tilde{\mathbf{p}}_k}^T(\mathbf{p}_k) \mathbf{J}_{\tilde{\mathbf{p}}_k}(\mathbf{p}_k) = \mathbf{I}_2$. Note that $\mathbf{J}_{\tilde{\mathbf{p}}_k}$ also represents the Jacobian of the transformation from the tangent plane to the original vector \mathbf{p}_k . Hence, the transpose of $\mathbf{J}_{\tilde{\mathbf{p}}_k}$ transforms the vector \mathbf{p}_k into its reduced counterpart $\tilde{\mathbf{p}}_k$, i.e.

$$\tilde{\mathbf{p}}_k = \mathbf{J}_{\tilde{\mathbf{p}}_k}^T(\mathbf{p}_k) \mathbf{p}_k = \mathbf{0},$$

with associated covariance matrix

$$\Sigma_{\tilde{\mathbf{p}}_k \tilde{\mathbf{p}}_k} = \mathbf{J}_{\tilde{\mathbf{p}}_k}^T(\mathbf{p}_k) \Sigma_{\mathbf{p}_k \mathbf{p}_k} \mathbf{J}_{\tilde{\mathbf{p}}_k}(\mathbf{p}_k).$$

Differently from $\Sigma_{\mathbf{p}_k \mathbf{p}_k} \in \mathbb{R}^{3 \times 3}$, the *reduced* covariance matrix $\Sigma_{\tilde{\mathbf{p}}_k \tilde{\mathbf{p}}_k} \in \mathbb{R}^{2 \times 2}$ is nonsingular. Therefore, one can set $\varrho_k^2 = \det(\Sigma_{\tilde{\mathbf{p}}_k \tilde{\mathbf{p}}_k})^{1/2}$, by taking the area of the uncertainty ellipse associated with $\Sigma_{\tilde{\mathbf{p}}_k \tilde{\mathbf{p}}_k}$, into account.

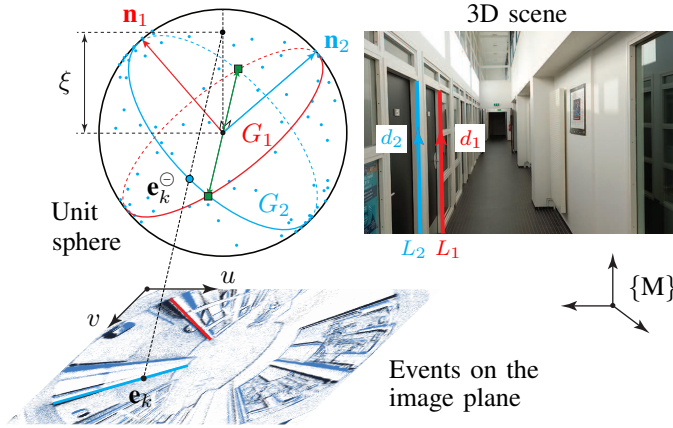


Figure 4: [Top right] Parallel lines L_1 (red) and L_2 (cyan) with directions \mathbf{d}_1 and \mathbf{d}_2 in the Manhattan frame $\{M\}$. [Left] Great circles G_1 and G_2 corresponding to L_1 and L_2 , with normal vectors \mathbf{n}_1 and \mathbf{n}_2 in the camera frame $\{C\}$. The two vanishing points are marked as green squares on the unit sphere.

C. Event clustering and great circle fitting

Our visual gyroscope relies on 3D lines (which map into great circles on the unit sphere), as primitives. To detect the circles from the continuous stream of data, we take all the events with the same polarity over a small time window T . In other words, we slice the event cloud into batches, whose size depends on the camera’s event rate. We then apply DBSCAN (“Density-Based Spatial Clustering of Applications with Noise”), a well-known data clustering algorithm [42], to find m clusters of events, $\mathcal{C}_1, \dots, \mathcal{C}_m$, which represent candidate great circles on the unit sphere. DBSCAN is easy-to-use and it only requires three parameters:

- `minPts`: the minimum number of nearest neighbors required to form a dense region,
- `ρ` : the radius of a neighborhood with respect to a point,
- `distFcn`: the distance function between two points.

If the ρ -neighborhood of a point contains at least `minPts` neighbors, then DBSCAN identifies the point as the *center point*. Note that the choice of distance function is tightly coupled to the choice of ρ , and the most common metric is the Euclidean distance. However, since the events \mathbf{e}_k are mapped into the unit sphere (cf. Sect. III-B), the *great-circle distance* turns out to be the most appropriate choice in our case. More precisely, we set `distFcn` = $1 - \cos\gamma$, where γ amounts to the angle subtended by two points along a great arc of the sphere.

Once the m clusters have been estimated by DBSCAN, the procedure described in Sect. II-B is used to fit m great circles to them. Let $\mathbf{p}_k = [\sin\theta_k \cos\varphi_k, \sin\theta_k \sin\varphi_k, \cos\theta_k]^T$ be the 3D coordinates of event \mathbf{e}_k^\ominus , and let $\mathcal{C}_i = \{\mathbf{p}_{i_1}, \mathbf{p}_{i_2}, \dots, \mathbf{p}_{N_i}\}$, $i \in \{1, \dots, m\}$, be the i -th cluster. Then, the normal vector \mathbf{n}_i to G_i is computed via the eigendecomposition of matrix $\mathbf{B}_i = \mathbf{A}_i^T \mathbf{A}_i$ where $\mathbf{A}_i = [\mathbf{p}_{i_1}, \mathbf{p}_{i_2}, \dots, \mathbf{p}_{N_i}]^T \in \mathbb{R}^{N_i \times 3}$. Note that if additional information about the events on the unit sphere is available (cf. Sect. III-B), one can alternatively use the weighted version of matrix \mathbf{B}_i presented in Remark 3.

D. Classification of 3D lines

To estimate the orientation of the camera frame with respect to the Manhattan frame, the knowledge of the normal vectors to the great circles $\mathbf{n}_1, \dots, \mathbf{n}_m$ alone, is not sufficient. In fact, a crucial step consists in partitioning the 3D lines in the Manhattan world, into three parallel and mutually-orthogonal groups $\mathcal{M}_1, \mathcal{M}_2$ and \mathcal{M}_3 . As in prior research [24], [25], to classify the lines, we opted for a variant of RANSAC (generically known as adaptive RANSAC, such as 3-line RANSAC), which guarantees more accurate results.

Our variant of RANSAC serves the purpose of identifying the dominant directions. A dominant direction is characterized by the vanishing point which best describes one of the directions in the Manhattan frame. The first step consists in selecting all possible pairs of distinct great circles. For each pair, we find one of the vanishing points, by calculating the cross product of the normals to the two great circles (see the green squares in Fig. 4). Once all the vanishing points have been obtained, those that are sufficiently close to the directions estimated in the previous step, are grouped into three categories. The centroid of each group of vanishing points is computed to make the algorithm less sensitive to outliers. Finally, to determine the three sets of normals associated with the three newly-established dominant directions, we choose those which are close enough to the great circles of these dominant directions (see Sect. IV for more details).

E. Global optimization: Attitude estimation

In the previous sections, we have seen how to estimate the normal vectors to the great circles and how to partition the 3D parallel lines into different orthogonal groups. The last module of our visual gyroscope takes these data as input, and provides an estimate $\hat{\mathbf{R}}_C^M$ of $\mathbf{R}_C^M \in \text{SO}(3)$, the rotation matrix representing the orientation of the camera in the Manhattan frame $\{M\}$, see Fig. 3. Following [24, Sect. 3], we can formulate this problem, in general, as

$$\hat{\mathbf{R}}_C^M = \arg \min_{\mathbf{R}} \frac{1}{2} \sum_{i=1}^m \frac{1}{\sigma_i^2} (\mathbf{d}_i^T \mathbf{R} \mathbf{n}_i)^2, \quad (4)$$

s.t. $\mathbf{R}^T \mathbf{R} = \mathbf{I}_3, \det(\mathbf{R}) = 1,$

where σ_i is a positive weight which reflects the uncertainty in each line-normal observation, \mathbf{d}_i is the 3D direction of line L_i in $\{M\}$, and \mathbf{n}_i is the noisy estimate of the normal to the great circle G_i in $\{C\}$ (cf. Sects. III-B and III-D). The nonlinear weighted constrained least-squares problem (4) (for $m \geq 3$) is known in the literature as the *orientation-from-line-correspondences problem*. In what follows, we will make the simplifying assumption that $\mathbf{d}_i \in \{e_1, e_2, e_3\}$ (i.e., we restrict ourselves to the three cardinal directions).

By using the CGR parameterization (cf. Sect. II-A), problem (4) can be equivalently rewritten as

$$\hat{\mathbf{s}} = \arg \min_{\mathbf{s}} J, \quad J = \frac{1}{2} \sum_{i=1}^m \frac{1}{\sigma_i^2} (\mathbf{d}_i^T \mathbf{R}(\mathbf{s}) \mathbf{n}_i)^2. \quad (5)$$

Note that compared to (4), the optimization problem (5) is *unconstrained* (the constraint is dropped since the CGR parameterization ensures that $\mathbf{R}(\mathbf{s})$ is a rotation matrix).

To algebraically find the global minimum of problem (5), the authors in [24] first determined all the critical points of J , by solving the optimality conditions, and then chose the one(s) than minimize (5). The optimality conditions are three quintic polynomials in the variable \mathbf{s} , whose solutions comprise the critical points of (5). Directly solving these polynomial equations is challenging, and in [24], the authors employed a method developed by [43] that computes the *multiplication matrix* (a generalization of the companion matrix to multivariate polynomial systems), whose eigenvalues are the roots of the associated (saturated) polynomial system. While this elegant method has the advantage of being not iterative and of not requiring any initialization, the construction of the multiplication matrix (a 40×40 matrix, in the problem under investigation) and the computation of its eigenvalues, are far from trivial. As pointed out in [24, Sect. 3.4], see also [31], symbolic reduction methods based on the computation of the system's Gröbner basis, could be applied as well, in principle, but they are susceptible to numerical instability, in practice.

To circumvent this difficulty, in this paper we propose to numerically solve problem (5) with *GloptiPoly*, a publicly-available optimization toolbox for Matlab [33]. *GloptiPoly* solves, or at least tries to approximate, the Generalized Problem of Moments (GPM), an infinite-dimensional optimization problem which can be viewed as an extension of the classical problem of moments [44], [45]. It allows to build up a hierarchy of semidefinite programming (SDP) or linear matrix inequality (LMI) relaxations of the GPM, whose associated monotone sequence of optimal values converges to the global optimum. This is not the only option: in fact, other specialized toolboxes for global polynomial optimization, such as *GpoSolver* [46] or *REPOP* [47], exist in the literature. Note that the cost function in (5) can be expressed as

$$J = \frac{g(\mathbf{s})}{h(\mathbf{s})} = \frac{\sum_{i=1}^m \frac{1}{\sigma_i^2} (\mathbf{d}_i^T \bar{\mathbf{R}}(\mathbf{s}) \mathbf{n}_i)^2}{2(1 + \mathbf{s}^T \mathbf{s})^2},$$

where $g(\mathbf{s})$ and $h(\mathbf{s})$ are polynomials of degree 4 in \mathbf{s} (no matter the number m of lines). The minimization of a *rational function*³ as J , can also be formulated as a linear moment problem [48], and *GloptiPoly* can easily handle this class of optimization problems. In fact, the toolbox applies the LMI relaxation technique described in [49], to write the corresponding moment problem. *GloptiPoly* is able to extract a certifiably globally optimal solution (which corresponds to the best attitude estimation of the camera) from the solution of SDP relaxations. Once the optimal vector of CGR parameters $\hat{\mathbf{s}}$ has been computed, the corresponding rotation matrix $\hat{\mathbf{R}}_C^M$ can be determined via equation (1).

In contrast to [24], an explicit expression for the weights $\sigma_1, \dots, \sigma_m$, appearing in problem (5), is given in [31]. However, this leads to a computationally intractable optimization problem, since the degree of the rational cost function rapidly

grows with each new summand. To get around this problem, the authors relaxed the formulation, by assuming that the variances $\sigma_1^2, \dots, \sigma_m^2$ are approximately the same for all measurements.

It is also worth pointing out that the rational cost function J in (5) can be transformed into an equivalent polynomial cost function, by introducing the auxiliary variable $s_4 = [2(1 + \mathbf{s}^T \mathbf{s})^2]^{-1}$. A polynomial of degree 5 in four unknowns, s_1, s_2, s_3 and s_4 , has then to be minimized. We will explore this alternative formulation in future works.

To conclude this section, three remarks are in order.

Remark 4: (Characterization of the solutions): In [24, Sect. 3.5], the authors have studied the existence of solutions to problem (5), and determined the number of global minima. They have shown that regardless of the number of measurements, there exist at least four solutions for a camera's orientation from observation of lines with known directions in the Manhattan frame. Specifically, assuming noisy-free measurements in the camera's frame, if $\hat{\mathbf{R}}_C^{M^*}$ is one solution, then $\hat{\mathbf{R}}_C^M = \Pi \hat{\mathbf{R}}_C^{M^*}$ with

$$\begin{aligned} \Pi \in & \left\{ \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \begin{bmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & -1 \end{bmatrix}, \begin{bmatrix} -1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -1 \end{bmatrix}, \begin{bmatrix} -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \right\} \\ & = \{ \mathbf{I}_3, \mathbf{R}_x(\pi), \mathbf{R}_y(\pi), \mathbf{R}_z(\pi) \}, \end{aligned}$$

is a solution too, where $\mathbf{R}_x(\pi)$ denotes the 3×3 elementary rotation of an angle π about the x -axis. In other words, each solution $\hat{\mathbf{R}}_C^{M^*}$ corresponds to an *equivalence class* of solutions in the form of

$$\hat{\mathbf{R}}_C^{M^*} = \{ \Pi \hat{\mathbf{R}}_C^{M^*}, \Pi \in \{ \mathbf{I}_3, \mathbf{R}_x(\pi), \mathbf{R}_y(\pi), \mathbf{R}_z(\pi) \} \}. \quad (6)$$

The equivalence relation (6) shows that (5) is an optimization problem defined over a quotient manifold $\mathcal{Q} = \bar{\mathcal{Q}}/\sim$, where $\bar{\mathcal{Q}} = \text{SO}(3)$ is the total space and \sim denotes the equivalence relation underlying (6) [50, Sect. 9.2].

The *minimum requirements* for estimating the orientation of a camera from line observations were studied in [51]. The author showed that it is necessary to observe three or more lines, and at least one of the 3D lines must be non-parallel to the others, to have a finite number of solutions for the camera's orientation. In a Manhattan world, this corresponds to observing three lines that pass through at least two (out of three) vanishing points. \diamond

Remark 5: (Singular configurations): In Remark 1, we have seen that when the rotation angle corresponding to $\mathbf{R} \in \text{SO}(3)$ is equal to $\pm\pi$, the CGR parameterization is degenerate. In practice, even in the proximity of $\pm\pi$, we can have extremely large values for the CGR parameters. To circumvent this problem, in [31], the authors simply proposed to rotate the measurements or 3D lines to an arbitrary (randomly generated) reference frame, find the global minimum (or minima), and then rotate the solution(s) back to the original frame. \diamond

Remark 6: (Global minimization of rational functions): It is well known that many geometric computer vision problems can be formulated as a minimization problem, where the objective function is a rational function in the unknown variables [32], [48]. These rational functions arise due to the perspective projection of a pinhole camera.

³Let $g(\mathbf{x})$ and $h(\mathbf{x}) \neq 0$ be two multivariate polynomials. Their quotient $f(\mathbf{x}) = g(\mathbf{x})/h(\mathbf{x})$ is called a *multivariate rational function*. The *degree* of a multivariate rational function is the maximum of the degrees of $g(\mathbf{x})$ and $h(\mathbf{x})$, when the fraction is reduced to lowest terms.

In particular, in the last decade, the CGR parameterization has been widely used to recast several problems of interest in computer vision and robotics, as (typically non-convex) polynomial optimization or fractional programming problems. By resting upon a least-squares formulation, two notable examples are the PnP (Perspective- n -Point) [29], [52] and the PnL (Perspective- n -Line) problems [31], [53], [54].

Generally speaking, any robustified constrained nonlinear least-squares problem of the form:

$$\begin{aligned} \min_{\mathbf{R}, \mathbf{q} \in \mathbb{R}^q} \quad & \frac{1}{2} \sum_{i=1}^m \mathcal{L}_i(\|\mathbf{f}_i(\mathbf{R}, \mathbf{q})\|^2), \\ \text{s.t.} \quad & \mathbf{R}^T \mathbf{R} = \mathbf{I}_3, \det(\mathbf{R}) = 1, \end{aligned}$$

where $\mathcal{L}_i(\cdot)$ is a loss function (a scalar-valued function used to reduce the influence of outliers on the solution), can be equivalently rewritten as the unconstrained least-squares problem,

$$\min_{\mathbf{s} \in \mathbb{R}^3, \mathbf{q} \in \mathbb{R}^q} \frac{1}{2} \sum_{i=1}^m \mathcal{L}_i(\|\mathbf{f}_i(\mathbf{R}(\mathbf{s}), \mathbf{q})\|^2), \quad (7)$$

using the CGR parameterization. To take advantage of the global optimization tools mentioned above, special care should be taken when choosing the cost functions \mathbf{f}_i and the loss function \mathcal{L}_i (note that if \mathcal{L}_i is the identity function, the usual nonlinear least-squares problem is obtained). In fact, one should guarantee that under the CGR parameterization, the objective function of problem (7) is a sum of m rational functions in \mathbf{s} and \mathbf{q} . \diamond

IV. EXPERIMENTAL VALIDATION

In this section, our omnidirectional visual gyroscope is validated via numerical (Sects. IV-A and IV-B) and real-world experiments (Sect. IV-C) with an event camera mounted on different robotic platforms.

A. Numerical simulations: Global optimization

In our first test, we evaluated the accuracy and robustness of the global optimization module described in Sect. III-E (magenta block in Fig. 3), in a synthetic environment. To have a realistic setup, we manually extracted 30 mutually orthogonal lines (from windows, door frames, fixtures, etc.) in a $7 \times 2 \times 2.5 \text{ m}^3$ hallway located in our lab building (5 lines are parallel to the x axis, 10 are parallel to the y axis, and 15 are parallel to the z axis). Two trajectories of the omnidirectional event camera have been considered in this environment. The first trajectory is a *figure-of-eight* (see Fig. 5(a)), whose parametric representation is,

$$\begin{cases} x = 60 \sin(2aj), \\ y = 120 \cos(aj) + 300, \\ z = 60, \\ \phi_1 = \phi_2 = 0, \\ \phi_3 = -\arctan\left(\frac{\cos(2aj)}{\sin(aj)}\right) - \frac{\pi}{2}, \end{cases} \quad j \in [-180^\circ, 180^\circ], \quad (8)$$

where (x, y, z) are the coordinates of the optical center of the camera in centimeters, (ϕ_1, ϕ_2, ϕ_3) are the roll, pitch and yaw angles of the camera in radians, and $a = \pi/180^\circ$. The second trajectory is a *helix* (see Fig. 5(b)), expressed by,

$$\begin{cases} x = 30 (\sin(aj) - 1), \\ y = 30 \cos(aj), \\ z = j/6, \\ \phi_1 = 50 a(1 - e^{-j/180}), \\ \phi_2 = -3 a(1 - e^{-j/180}), \\ \phi_3 = aj. \end{cases} \quad j \in [0, 3 \times 360^\circ], \quad (9)$$

In Figs. 5(a),(b), a set of 15 lines, marked in blue, is shown. We used equation (8) to describe the planar motion of a camera mounted on a wheeled robot, and equation (9), the ascending motion of a quadrotor.

We set the number of observed lines either to 15 or to 30. In the first case, to guarantee that the Manhattan world assumption is satisfied, 3 lines are imposed (one per coordinate axis), while the other 12 are drawn at random out of the remaining 27 lines. In the second case, all 30 lines are taken. The measurements of the normal vectors $\mathbf{n}_1, \dots, \mathbf{n}_m$, were corrupted with additive zero-mean white Gaussian noise of increasing magnitude. The noise acting on the three components of the normal vectors is independent and the standard deviation $\zeta^* = \kappa \zeta$, $\kappa = \sin(3^\circ)$, $\zeta \in \{0, 0.5, 1\}$. Figs. 5(c),(d) report the mean and standard deviation of the angular estimation error of the camera over 50 trials, for the two trajectories and for different values of m and ζ^* . More specifically, given $\widehat{\mathbf{R}}_C^M = \mathbf{R}_z(\widehat{\phi}_3) \mathbf{R}_y(\widehat{\phi}_2) \mathbf{R}_x(\widehat{\phi}_1)$, we computed the statistics of the errors $\epsilon_j = |\phi_j - \widehat{\phi}_j|$, $j \in \{1, 2, 3\}$, in degrees. Problem (5), with weights $\sigma_1 = \dots = \sigma_m = 1$, was numerically solved with GloptiPoly, using the conic programming solver SeDuMi [55]. To single out the “best solution” from the four possible minima of the cost function J in (5) during the camera motion (cf. Remark 4), we used the *chordal distance* $d_{\text{chor}}[\cdot]$ in $\text{SO}(3)$ between the current and previous rotation estimate⁴. In fact, at time $\ell \in \{1, 2, \dots\}$, we compute,

$$\begin{aligned} \arg \min_{\Pi} \quad & d_{\text{chor}}[\Pi \widehat{\mathbf{R}}_C^{M^*}(\ell), \widehat{\mathbf{R}}_C^{M^*}(\ell-1)] = \\ \arg \min_{\Pi} \quad & \|\Pi \widehat{\mathbf{R}}_C^{M^*}(\ell) - \widehat{\mathbf{R}}_C^{M^*}(\ell-1)\|_F, \end{aligned}$$

where $\|\cdot\|_F$ denotes the Frobenius norm of a matrix, i.e. for $\mathbf{A} \in \mathbb{R}^{n \times n}$, $\|\mathbf{A}\|_F = \text{trace}(\mathbf{A}^T \mathbf{A})^{1/2}$. While this method is easy to implement, it is rather sensitive to the accuracy of the initial camera orientation estimate.

From an inspection of Figs. 5(c),(d), we can notice that the worst mean estimation error is 5.04° (4.45°) for the figure-of-eight (helical) trajectory. In both cases, it corresponds to the yaw angle: this is due to the unequal repartition of 3D parallel lines in the specific simulation environment considered. Moreover, for the same magnitude of noise, the

⁴Other distance measures on $\text{SO}(3)$ are possible as well [56, Sect. 4]. We opted for the chordal distance because of its simplicity.

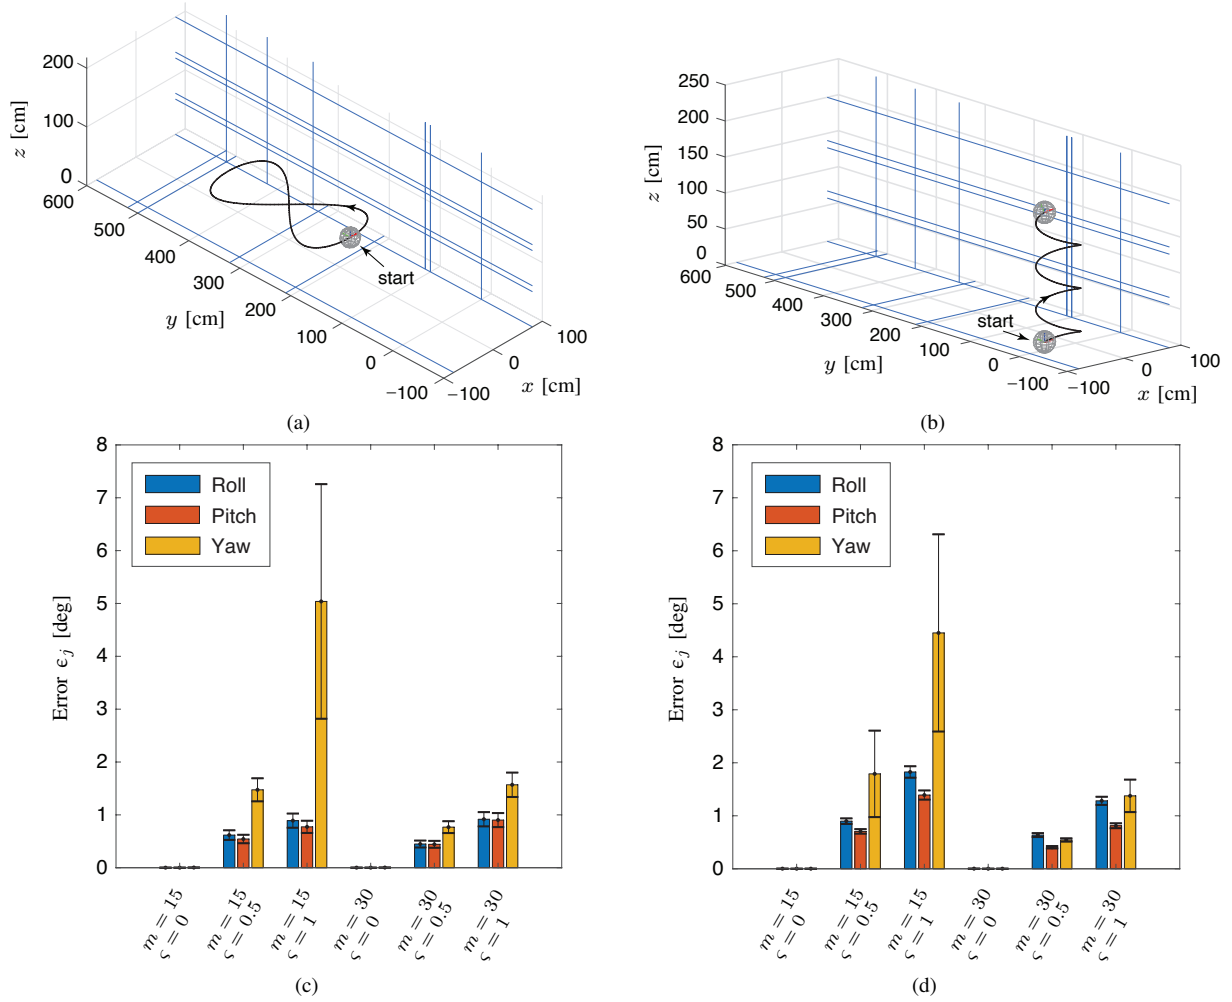


Figure 5: *Numerical simulations*: (a) Figure-of-eight trajectory; (b) Helicoidal trajectory; (c),(d) Mean and standard deviation of the estimation error $\epsilon_j = |\phi_j - \hat{\phi}_j|$, $j \in \{1, 2, 3\}$, over 50 trials, for the figure-of-eight and helicoidal trajectory, respectively. The errors for different numbers of lines, $m \in \{15, 30\}$, and different levels of noise, $\zeta^* = \kappa\zeta$, $\kappa = \sin(3^\circ)$, $\zeta \in \{0, 0.5, 1\}$, are compared.

larger the number of observed lines, the more accurate the visual gyroscope, as expected.

B. Simulated event camera in CARLA

To evaluate the effectiveness of our visual gyroscope in large-scale environments under realistic conditions, an omnidirectional event camera was created in the autonomous driving simulator, CARLA (“Car Learning to Act”, [57]). The camera was mounted on the roof of a car moving along a closed trajectory of 780 m, in a synthetic urban environment (see Figs. 6(a) and 6(c)). CARLA’s default map 10 was modified to comply with the Manhattan world assumption (in particular, the layout of the buildings was simplified). The maximal linear and angular velocities of the car are 57.53 km/h and 24.44 deg/s, respectively, and the total runtime is 66.75 s. As in [23], we adopted a three-step procedure to generate the stream of events fed into the visual gyroscope. First, we extract the events by computing the difference between consecutive RGB frames from five orthogonal viewpoints. These correspond to five perspective cameras centered on

the car, which are available by default in CARLA. Cubemap projection is then followed by ray tracing, which we use to build a lookup table that stores the corresponding pixels in the cubemap and in the omnidirectional image. A circular mask was applied to the omnidirectional image, in order to extract only the informative pixels of the scene. The car was also equipped with a simulated compass and GNSS, which we used to obtain an accurate ground-truth pose of the vehicle.

Fig. 6(b) shows an example of catadioptric event frame. The calibration parameters used to generate it, are: $f_u = 309.845$ pixels, $f_v = 310.012$ pixels, $u_0 = 600.935$ pixels, $v_0 = 372.102$ pixels and $\xi = 0.9521$.

Fig. 6(d) reports the orientation of the car estimated by the visual gyroscope along the trajectory (solid), and the corresponding ground truth (dashed). The time evolution of the magnitude of the estimation error (solid) and the mean estimation error (dashed) are shown in Fig. 6(e). In Fig. 6(d) and 6(e), the following parameters were used: time window $T = 13$ ms, $\rho = 0.75^\circ$, $\minPts = 3$, $r_g = 5^\circ$, $\mathcal{L}_{arc} = 7^\circ$, $\mathcal{L}_d = 1^\circ$. The last three detection parameters, r_g , \mathcal{L}_{arc} and \mathcal{L}_d ,

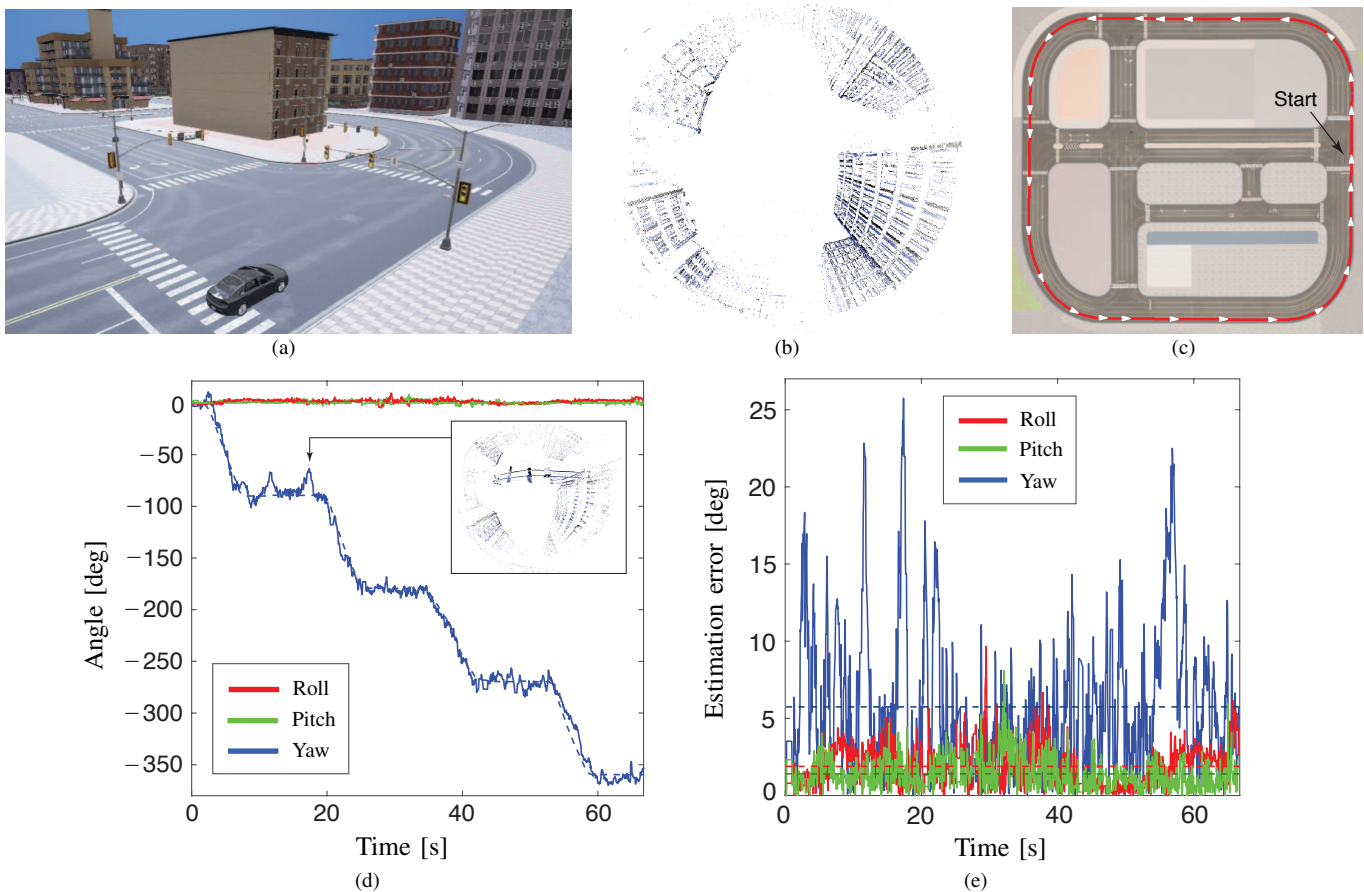


Figure 6: *CARLA simulator*: (a) Top view of the car in the synthetic urban environment; (b) Example of catadioptric event frame; (c) Closed trajectory (red) of the car, marked as a white triangle every 2.5 s; (d) Actual orientation of the car (dashed) and orientation estimated by Gyrevento (solid); (e) Time evolution of the magnitude of the estimation error (solid), and mean estimation error (dashed).

all expressed in degrees, have been introduced for enhanced flexibility. r_g is the angle of the spherical cone with apex at the center of the sphere, containing the vanishing points used to identify the new dominant directions in RANSAC (cf. Sect III-D), \mathcal{L}_{arc} denotes the minimum length of an arc on the unit sphere (corresponding to a straight line in the Manhattan frame), and \mathcal{L}_d is the maximum thickness of an arc.

From Fig. 6(e), we can see that the mean estimation error is 1.889° (roll), 1.402° (pitch) and 5.743° (yaw). The spikes on the estimation error occurring at 11.55 s, 17.4 s and 56.85 s, can be ascribed to the non-perfect conformity of our simulation environment with the Manhattan world assumption. In fact, street furniture (traffic lights, street lamps, sunshades, etc.) perturbs parallel line detection and negatively impacts the accuracy of the visual gyroscope.

C. Real-world experiments: Event camera on a robot arm

a) Camera specifications : In our real-world experiments, we used a Prophesee Evaluation Kit 3 - HD (EVK3). This event camera, released in 2021, has a resolution of 1280×720 pixels, a latency of $220 \mu\text{s}$, a dynamic range exceeding 110 dB, a pixel size of $4.86 \times 4.86 \mu\text{m}$, and a maximum

bandwidth of 1.6 Gbps. An event $\mathbf{e}_k = [u_k, v_k, t_k, p_k]^T$ is encoded on 64 bits: 28 bits for u_k and v_k , 32 bits for t_k , and 4 bits for p_k . A catadioptric event camera was obtained by screwing a VStone VS-C450MRTK objective with hyperbolic mirror, on top of EVK3.

b) Camera calibration : The calibration of an event camera is more challenging than that of a traditional camera, and it remains a subject of ongoing research. Some authors have simply proposed to use blinking patterns or active illumination with static checkerboards, for the calibration of perspective event cameras. Recently, a dynamic and a learning-based approach have been described in [58] and [59], respectively, to simplify and make the calibration procedure more robust. The *E-Calib* toolbox [60], has also been designed with the same goal in mind: fast, robust and accurate calibration of event cameras. However, with the exception of [22] where a fisheye event camera is considered, we are not aware of any other generic calibration algorithm for *omnidirectional event cameras*. To calibrate different typologies of panoramic event sensors (catadioptric, fisheye, etc.), we developed a generic protocol which consists in filming a flashing pattern in front of the static camera. A grayscale image is generated




Figure 7: (a), (b) Two views of the experimental setup: The catadioptric event camera is mounted on the end-effector of a Staubli TX-60 robot and observes the parallel lines available in the environment. (c) Event frame captured by the catadioptric event camera, 2.5 seconds after the beginning of Sequence 1.a (a pixel is light blue if the polarity $p_k = +1$, and black if $p_k = -1$).

by accumulating the events over a fixed time window (that the user can select), which, in turn, is exploited to estimate the intrinsic parameters $\{f_u, f_v, u_0, v_0, \xi\}$ of the camera via conventional calibration algorithms, available in Matlab or OpenCV (in our experiments, we used HySCaS [61]).

c) Experimental setup : To validate the overall pipeline reported in Fig. 3, experiments have been carried out in a laboratory environment. Similarly to the *RobotEvt* dataset [62], in which an iniVation DAVIS 240C event camera was installed on a UR-5 arm, our catadioptric event camera was mounted on the end-effector of a 6-axis Staubli TX-60 robot (see Fig. 7). The rotary encoders on the joints of the robot provide accurate measurements of angular positions and velocities at 250 Hz. Forward kinematics of the robot was used to compute the position and orientation of the end-effector (our ground truth), from specified values of joint parameters. The stream of events and the measurements from the encoders were synchronized offline, and the estimated orientation of the camera was updated at a fixed (arbitrary) frequency of 25 Hz. As in Sect. IV-A, problem (5), with weights $\sigma_1 = \dots = \sigma_m = 1$, was numerically solved with GloptiPoly, using SeDuMi.

We evaluated the accuracy of Gyrevento, by considering two sequences of growing complexity. For a given trajectory, to study the effect of speed, three sets of camera velocities were tested. In what follows, we use $\mathbf{v}_C^{\max} = [v_x^{\max}, v_y^{\max}, v_z^{\max}, \omega_x^{\max}, \omega_y^{\max}, \omega_z^{\max}]^T$ the maximum velocity of the camera (expressed in the camera frame), where the linear velocities are in mm/s and the angular velocities in deg/s.

- **Sequence 1:** 1-DoF motion (pure rotation about the optical axis of the camera):

- $\mathbf{v}_C^{\max} = [0, 0, 0, 0, 0, 47.8]^T$, duration 18.20 s, $T = 10$ ms, $r_g = 30^\circ$, (see Fig. 7(c)),
- $\mathbf{v}_C^{\max} = [0, 0, 0, 0, 0, 91.8]^T$, duration 10.24 s, $T = 10$ ms, $r_g = 30^\circ$,
- $\mathbf{v}_C^{\max} = [0, 0, 0, 0, 0, 136.6]^T$, duration 7.56 s, $T = 5$ ms, $r_g = 45^\circ$.

- **Sequence 2:** 4-DoF motion (rotation about the optical axis of the camera and 3D translation in the $[-20, 290] \times [0, 230] \times [380, 670]$ mm³ range):

- $\mathbf{v}_C^{\max} = [46.6, 55.6, 36.3, 0, 0, 19.2]^T$, duration 27.52 s, $T = 22.5$ ms, $r_g = 18^\circ$,
- $\mathbf{v}_C^{\max} = [69.4, 80.8, 56.0, 0, 0, 28.8]^T$, duration 19.08 s, $T = 20$ ms, $r_g = 20^\circ$,
- $\mathbf{v}_C^{\max} = [92.1, 108.0, 72.0, 0, 0, 38]^T$, duration 14.80 s, $T = 15$ ms, $r_g = 22.5^\circ$.

The parameters \minPts , ρ , \mathcal{L}_{arc} , \mathcal{L}_d have been chosen as in Sect. IV-B, while $f_u = 310.2723$ pixels, $f_v = 308.8265$ pixels, $u_0 = 601.7725$ pixels and $v_0 = 372.3330$ pixels and $\xi = 1.1099$. As in Sect. IV-B, we took advantage of a circular mask to discard the non-informative pixels in the outer rim of the omnidirectional image.

d) Discussion of the results : Figs 8(a) and 8(b) show the time evolution of the actual orientation of the camera (dashed) and the orientation estimated by Gyrevento (solid) in Sequence 1a and Sequence 2c, respectively. The impact of camera velocity on the attitude estimation error is evaluated in Fig. 8(c) for Sequence 1 and in Fig. 8(d) for Sequence 2. For the sake of clarity, the outliers are not shown in the box charts in Figs. 8(c) and 8(d). They are 6.7° , 9.2° , 12° and 11.5° , 8.2° , 4.0° , respectively, and the inertia of the camera during sharp accelerations/decelerations of the manipulator, is the main responsible for them. Figs. 8(c) and 8(d) indicate that Gyrevento works best for angular velocities in the $[20, 50]$ deg/s range (the maximum mean error is less than 2.5°). However, it continues to deliver a satisfactory performance for speeds up to 136 deg/s. Our visual gyroscope can thus be used onboard agile robotic platforms (e.g. quadrotors). The mean error in Sequence 2 is about twice as small as that in Sequence 1, since the 4-DoF camera motion offers richer visual information (and, ultimately, more exploitable parallel lines). Finally, we can notice that the time-evolution of the estimation error in Figs. 8(c) and 8(d), is smoother than that in Figs. 6(d) and 6(e). In fact, the scene in Sequences 1 and 2 contains less artifacts, and the Manhattan world assumption is more likely to hold than in CARLA's simulated urban environment.

V. CONCLUSION AND FUTURE WORK

In this paper, we have presented Gyrevento, the first omnidirectional visual gyroscope for event cameras only based

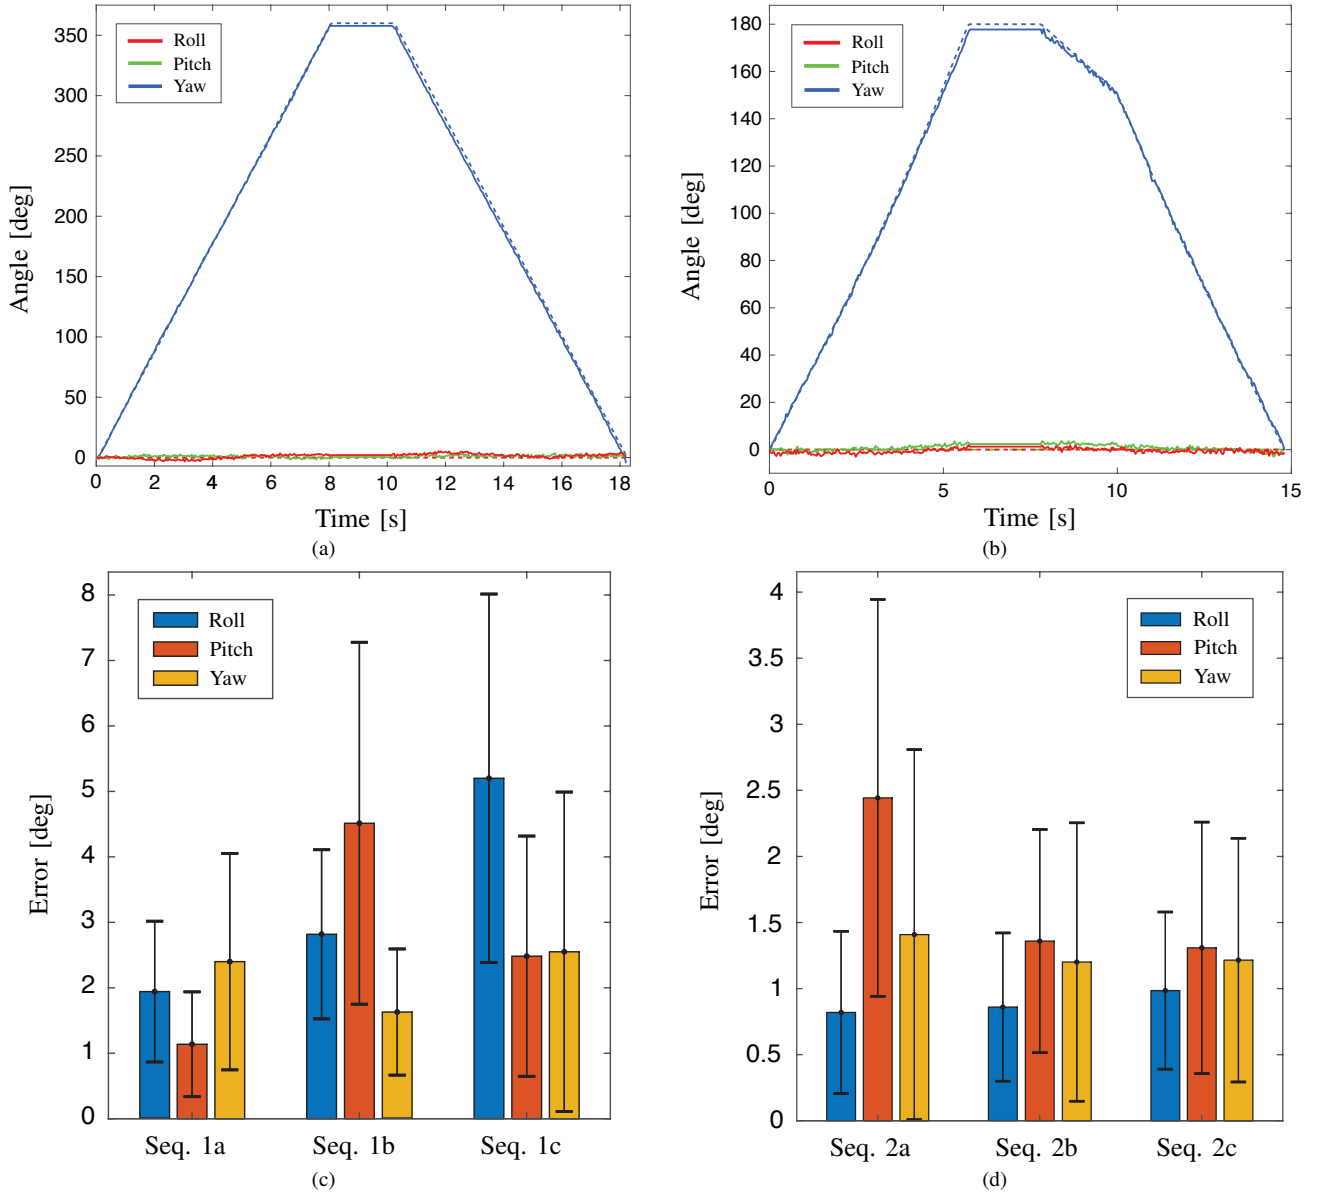


Figure 8: *Experimental results*: Actual orientation of the camera in degrees (dashed) and orientation estimated by Gyrevento (solid): (a) Sequence 1a, (b) Sequence 2c. Mean and standard deviation of the angular estimation error in degrees: (c) Sequences 1a-1c, (d) Sequences 2a-2c.

on the observation of 3D lines in a structured environment. The gyroscope has been successfully validated via simulation and real-world experiments with a robot manipulator. The experimental results have also shed light on the impact of camera velocity on the attitude estimation error.

In future research, we will work towards the design of an event-based omnidirectional visual odometer and we plan to relax the Manhattan world assumption about the existence of three mutually orthogonal directions (cf. Atlanta world assumption [27]). It would also be interesting to automatically incorporate the Manhattan constraints into the event-clustering step, which is, at present, the computational bottleneck of Gyrevento (its complexity is $O(n_e \log(n_e))$, where n_e is the number of detected events). On the other hand, the com-

plexity of the great circle fitting and classification steps (cf. Fig. 3) is $o(n_e)$ and $O(m^2)$, respectively, with $m \ll n_e$. To robustify our multi-line tracker, we could exploit the polarity of the events and the spatio-temporal variant of DBSCAN, ST-DBSCAN [63]. Finally, in future research, we will weigh the pros and cons of other vectorial parameterizations of a 3D rotation [30], besides CGR.

Supplemental material

The event datasets, calibration parameters, and videos relative to the simulation and real-world experiments in Sect. IV, and to tests performed with CARLA simulator and an IMU rigidly attached to the catadioptric event camera (Sequence 3), are available at the following link:

<http://home.mis.u-picardie.fr/~fabio/Gyrevento.html>

In the video which compares the three speeds of the Stäubli TX-60, Sequence 1a, 1b, 1c and Sequence 2a, 2b, 2c are called Sequence 1, Speed 25%, 50%, 75% and Sequence 2, Speed 50%, 75%, 100%, respectively. These percentages refer to the maximum achievable speed of the robot.

REFERENCES

- [1] G. Gallego, T. Delbrück, G. Orchard, C. Bartolozzi, B. Taba, A. Censi, S. Leutenegger, A.J. Davison, J. Conratt, K. Daniilidis, and D. Scaramuzza. Event-Based Vision: A Survey. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(1):154–180, 2022.
- [2] D. Gehrig and D. Scaramuzza. Low-latency automotive vision with event cameras. *Nature*, 629(8014):1034–1040, 2024.
- [3] H. Kim, S. Leutenegger, and A.J. Davison. Real-Time 3D Reconstruction and 6-DoF Tracking with an Event Camera. In *Proc. Europ. Conf. Comput. Vis.*, pages 349–364, 2016.
- [4] 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 Rob. Autom. Lett.*, 2(2):593–600, 2017.
- [5] A.Z. Zhu, N. Atanasov, and K. Daniilidis. Event-based Visual Inertial Odometry. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, pages 5391–5399, 2017.
- [6] E. Mueggler, G. Gallego, H. Rebecq, and D. Scaramuzza. Continuous-Time Visual-Inertial Odometry for Event Cameras. *IEEE Trans. Robot.*, 34(6):1425–1440, 2018.
- [7] H. Kim and H.J. Kim. Real-Time Rotational Motion Estimation With Contrast Maximization Over Globally Aligned Events. *IEEE Rob. Autom. Lett.*, 6(3):6016–6023, 2021.
- [8] J. Hidalgo-Carrió, G. Gallego, and D. Scaramuzza. Event-aided Direct Sparse Odometry. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, pages 5781–5790, 2022.
- [9] D. Liu, A. Parra, Y. Latif, B. Chen, T.-J. Chin, and I. Reid. Asynchronous Optimisation for Event-based Visual Odometry. In *Proc. IEEE Int. Conf. Robot. Automat.*, pages 9432–9438, 2022.
- [10] Y.F. Zuo, J. Yang, J. Chen, X. Wang, Y. Wang, and L. Kneip. DEVO: Depth-Event Camera Visual Odometry in Challenging Conditions. In *Proc. IEEE Int. Conf. Robot. Automat.*, pages 2179–2185, 2022.
- [11] C. Brändli, J. Strubel, S. Keller, D. Scaramuzza, and T. Delbruck. ELiSeD – An Event-Based Line Segment Detector. In *Proc. 2nd Int. Conf. Event-based Contr. Comm. and Sign. Proc.*, 2016.
- [12] R.G. Von Gioi, J. Jakubowicz, J.-M. Morel, and G. Randall. LSD: A Fast Line Segment Detector with a False Detection Control. *IEEE Trans. Pattern Anal. Mach. Intell.*, 32(4):722–732, 2008.
- [13] L. Everding and J. Conratt. Low-Latency Line Tracking Using Event-Based Dynamic Vision Sensors. *Front. Neurobot.*, 12:4, 2018.
- [14] D.R. Valeiras, X. Clady, S.-H. Ieng, and R. Benosman. Event-Based Line Fitting and Segment Detection Using a Neuromorphic Visual Sensor. *IEEE Trans. Neural Networks Learn. Syst.*, 30(4):1218–1230, 2019.
- [15] L. Gao, H. Su, D. Gehrig, M. Cannici, D. Scaramuzza, and L. Kneip. A 5-Point Minimal Solver for Event Camera Relative Motion Estimation. In *Proc. IEEE Int. Conf. Comput. Vis.*, pages 8049–8059, 2023.
- [16] L. Gao, D. Gehrig, H. Su, D. Scaramuzza, and L. Kneip. An N-Point Linear Solver for Line and Motion Estimation with Event Cameras. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, pages 14596–14605, 2024.
- [17] X. Lu, Y. Zhou, J. Niu, S. Zhong, and S. Shen. Event-based Visual Inertial Velometer. In *Proc. Robotics: Science and Systems*, 2024. Paper n. 88.
- [18] W. Chamorro, J. Solà, and J. Andrade-Cetto. Event-based Line SLAM in Real-time. *IEEE Rob. Autom. Lett.*, 7(3):8146–8153, 2022.
- [19] W. Xu, X. Peng, and L. Kneip. Tight Fusion of Events and Inertial Measurements for Direct Velocity Estimation. *IEEE Trans. Robot.*, 40:240–256, 2024.
- [20] S. Schraml, A.N. Belbachir, and H. Bischof. An Event-Driven Stereo System for Real-Time 3-D 360° Panoramic Vision. *IEEE Trans. Ind. Electron.*, 63(1):418–428, 2016.
- [21] C. Reinbacher, G. Munda, and T. Pock. Real-Time Panoramic Tracking for Event Cameras. In *Proc. IEEE Int. Conf. Comput. Photog.*, 2017.
- [22] H. Cho, J. Jeong, and K.-J. Yoon. EOMVS: Event-Based Omnidirectional Multi-View Stereo. *IEEE Rob. Autom. Lett.*, 6:6709–6716, 2021.
- [23] A.R. Sekkat, Y. Dupuis, V.R. Kumar, H. Rashed, S. Yogamani, P. Vasseur, and P. Honeine. SynWoodScape: Synthetic Surround-View Fisheye Camera Dataset for Autonomous Driving. *IEEE Rob. Autom. Lett.*, 7(3):8502–8509, 2022.
- [24] F.M. Mirzaei and S.I. Roumeliotis. Optimal Estimation of Vanishing Points in a Manhattan World. In *Proc. IEEE Int. Conf. Comput. Vis.*, pages 2454–2461, 2011.
- [25] J.-C. Bazin, C. Demonceaux, P. Vasseur, and I. Kweon. Rotation estimation and vanishing point extraction by omnidirectional vision in urban environment. *Int. J. Robot. Res.*, 31(1):63–81, 2012.
- [26] P. Kim, H. Li, and K. Joo. Quasi-Globally Optimal and Real-Time Visual Compass in Manhattan Structured Environments. *IEEE Rob. Autom. Lett.*, 7(2):2613–2620, 2022.
- [27] D. Yan, H. Jiang, T. Li, and C. Shi. Efficient Vanishing Point Estimation for Accurate Camera Rotation Estimation in Indoor Environments. *IEEE Rob. Autom. Lett.*, 8(11):6899–6906, 2023.
- [28] C. Geyer and K. Daniilidis. A Unifying Theory for Central Panoramic Systems and Practical Implications. In *Proc. Europ. Conf. Comput. Vis.*, pages 445–461, 2000.
- [29] J.A. Hesch and S.I. Roumeliotis. A Direct Least-Squares (DLS) Method for PnP. In *Proc. IEEE Int. Conf. Robot. Automat.*, pages 383–390, 2011.
- [30] O.A. Bauchau and L. Trainelli. The Vectorial Parameterization of Rotation. *Nonlinear Dyn.*, 32(1):71–92, 2003.
- [31] F.M. Mirzaei and S.I. Roumeliotis. Globally Optimal Pose Estimation from Line Correspondences. In *Proc. IEEE Int. Conf. Robot. Automat.*, pages 5581–5588, 2011.
- [32] F. Bugarin, D. Henrion, and J.B. Lasserre. Minimizing the sum of many rational functions. *Math. Prog. Comp.*, 8(1):83–111, 2016.
- [33] D. Henrion, J.-B. Lasserre, and J. Löfberg. GloptiPoly 3: moments, optimization and semidefinite programming. *Optim. Methods Software*, 24(4-5):761–779, 2009.
- [34] M.D. Shuster. A Survey of Attitude Representations. *J. Astronaut. Sci.*, 41(4):439–517, 1993.
- [35] B. Siciliano, L. Sciacivco, L. Villani, and G. Oriolo. *Robotics: Modelling, Planning and Control*. Advanced Textbooks in Control and Signal Processing. Springer, 2009.
- [36] X. Ying, Z. Hu, and H. Zha. Fisheye Lenses Calibration Using Straight-Line Spherical Perspective Projection Constraint. In P.J. Narayanan, S.K. Nayar, and H.-Y. Shum, editors, *Proc. Asian Conf. Comput. Vis.*, pages 61–70. Springer, 2006.
- [37] C.F. Marcus. A Note on Fitting Great Circles by Least Squares. *Commun. ACM*, 4(8):353, 1961.
- [38] J. Courbon, Y. Mezouar, and P. Martinet. Evaluation of the Unified Model of the Sphere for Fisheye Cameras in Robotic Applications. *Adv. Robotics*, 26(8-9):947–967, 2012.
- [39] J. Barreto. A unifying geometric representation for central projection systems. *Comput. Vis. Image Underst.*, 103(3):208–217, 2006.
- [40] W. Förstner. Minimal Representations for Uncertainty and Estimation in Projective Spaces. In R. Kimmel, R. Klette, and A. Sugimoto, editors, *Proc. 10th Asian Conf. Comput. Vis.*, pages 619–632. Springer, 2010.
- [41] S. Urban, J. Leitloff, and S. Hinz. MLPnP - A Real-Time Maximum Likelihood Solution to the Perspective-n-Point Problem. *ISPRS Ann. Photogramm. Remote Sens. Spatial Inf. Sci.*, 3(3):131–138, 2016.
- [42] E. Schubert, J. Sander, M. Ester, H.P. Kriegel, and X. Xu. DBSCAN Revisited, Revisited: Why and How You Should (Still) Use DBSCAN. *ACM Trans. Database Syst.*, 42(3), 2017. Article n. 19.
- [43] W. Auzinger and H.J. Stetter. An Elimination Algorithm for the Computation of All Zeros of a System of Multivariate Polynomial Equations. In R.P. Agarwal, Y.M. Chow, and S.J. Wilson, editors, *Proc. Int. Conf. Numerical Mathematics*, volume 86, pages 11–30. 1988.
- [44] J.B. Lasserre. *Moments, Positive Polynomials and Their Applications*, volume 1 of *Optimization Series*. Imperial College Press, 2010.
- [45] D. Henrion, M. Korda, and J.B. Lasserre. *The Moment-SOS Hierarchy: Lectures In Probability, Statistics, Computational Geometry, Control And Nonlinear PDEs*, volume 4 of *Series on Optimization and its Applications*. World Scientific, 2020.
- [46] J. Heller and T. Pajdla. GpoSolver: A MATLAB/C++ Toolbox for Global Polynomial Optimization. *Optim. Methods Software*, 31:405–434, 2016.
- [47] O. Karaca, G. Darivianakis, P. Beuchat, A. Georghiou, and J. Lygeros. The REPOP Toolbox: Tackling Polynomial Optimization Using Relative Entropy Relaxations. *Proc. IFAC World Cong.*, 50(1):11652–11657, 2017.
- [48] F. Kahl and D. Henrion. Globally Optimal Estimates for Geometric Reconstruction Problems. *Int. J. Comput. Vision*, 74(1):3–15, 2007.
- [49] D. Jibeteau and E. de Klerk. Global optimization of rational functions: a semidefinite programming approach. *Math. Program.*, 106(1):93–109, 2006.

- [50] N. Boumal. *An Introduction to Optimization on Smooth Manifolds*. Cambridge University Press, 2023.
- [51] H.H. Chen. Pose Determination from Line-to-Plane Correspondences: Existence Condition and Closed-Form Solutions. *IEEE Trans. Pattern Anal. Mach. Intell.*, 13(6):530–541, 1991.
- [52] C. Sweeney, V. Fragoso, T. Höllerer, and M. Turk. gDLS: A Scalable Solution to the Generalized Pose and Scale Problem. In *Proc. Europ. Conf. Comput. Vis.*, pages 16–31, 2014.
- [53] Q. Yu, G. Xu, and Y. Cheng. An efficient and globally optimal method for camera pose estimation using line features. *Mach. Vision Appl.*, 31(6):1–11, 2020.
- [54] L. Zhou, D. Koppel, and M. Kaess. A Complete, Accurate and Efficient Solution for the Perspective-N-Line Problem. *IEEE Rob. Autom. Lett.*, 6(2):699–706, 2021.
- [55] J.F. Sturm. Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optim. Methods Software*, 11:625–653, 1999.
- [56] R. Hartley, J. Trumpf, Y. Dai, and H. Li. Rotation Averaging. *Int. J. Comput. Vision*, 103(3):267–305, 2013.
- [57] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun. CARLA: An Open Urban Driving Simulator. In *Proc. 1st Annual Conf. Robot Learn.*, pages 1–16, 2017. [Online] <https://carla.org>.
- [58] K. Huang, Y. Wang, and L. Kneip. Dynamic Event Camera Calibration. In *Proc. IEEE/RSJ Int. Conf. Intel. Robots Syst.*, pages 7021–7028, 2021.
- [59] M. Muglikar, M. Gehrig, D. Gehrig, and D. Scaramuzza. How to Calibrate Your Event Camera. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshop*, pages 1403–1409, 2021.
- [60] M. Salah, A. Ayyad, M. Humais, D. Gehrig, A. Abusafieh, L. Seneviratne, D. Scaramuzza, and Y. Zweiri. E-Calib: A Fast, Robust and Accurate Calibration Toolbox for Event Cameras. *IEEE Trans. Image Process.*, 33:3977–3990, 2024.
- [61] G. Caron and D. Eynard. HySCaS: Hybrid Stereoscopic Calibration Software. [Online] <https://mis.u-picardie.fr/~g-caron/software>, 2011.
- [62] D. Liu, Á. Parra, and T.-J. Chin. Spatiotemporal Registration for Event-based Visual Odometry. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, pages 4937–4946, 2021.
- [63] D. Birant and A. Kut. ST-DBSCAN: An algorithm for clustering spatial-temporal data. *Data Knowl. Eng.*, 60(1):208–221, 2007.