



(51) International Patent Classification:

G06T 7/20 (2017.01) G06T 7/593 (2017.01)
G06T 7/246 (2017.01)

(21) International Application Number:

PCT/EP2018/073641

(22) International Filing Date:

03 September 2018 (03.09.2018)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

17189223.5 04 September 2017 (04.09.2017) EP

(71) Applicant: UNIVERSITÄT ZÜRICH [CH/CH]; Prorektorat VNW, Rämistr. 71, 8006 Zürich (CH).


(72) Inventors: **REBECQ, Henri**; Bahnhaldenstrasse 10, 8052 Zürich (CH). **SCARAMUZZA, Davide**; Jungstrasse 16, 8050 Zürich (CH).

(74) Agent: **DERTINGER, Thomas**; Schulz Junghans Patentanwälte PartGmbB, Großbeerenstraße 71, 10963 Berlin (DE).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA,

(54) Title: VISUAL-INERTIAL ODOMETRY WITH AN EVENT CAMERA

Fig. 2



(57) Abstract: The invention relates a method for generating a motion-corrected image for visual-inertial odometry comprising an event camera rigidly connected to an inertia measurement unit (IMU), wherein the event camera comprises pixels arranged in an image plane that are configured to output events in presence of brightness changes in a scene at the time they occur, wherein each event comprises the time at which it is recorded and a position of the respective pixel that detected the brightness change, the method comprising the steps of: Acquiring at least one set of events (S), wherein the at least one set (S) comprises a plurality of subsequent events (e); Acquiring IMU data (D) for the duration of the at least one set (S); Generating a motion-corrected image from the at least one set (S) of events (e), wherein the motion-corrected image is obtained by assigning the



SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- *as to the identity of the inventor (Rule 4.17(i))*
- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

Published:

- *with international search report (Art. 21(3))*

position (x_j) of each event (e_j) recorded at its corresponding event time (t_j) at an estimated event camera pose (T_{t_j}) to an adjusted event position (x'_j), wherein the adjusted event position (x'_j) is obtained by determining the position of the event (e_j) for an estimated reference camera pose (T_{t_k}) at a reference time (t_k), wherein the estimated camera pose (T_{t_j}) at the event time (t_j) and the reference camera pose (T_{t_k}) at the reference time (t_k) are estimated by means of the IMU data (D).

Visual-inertial odometry with an event camera

Specification

- 5 The invention relates to a method for generating a motion-corrected image from event cameras for visual-inertial odometry as well as to a method for visual-inertial odometry with an event camera.

Event cameras, such as the Dynamic Vision Sensor (DVS), work very differently from a traditional camera.

- 10 They have independent pixels that only send information referred to as events in presence of brightness changes in the scene at the time they occur.

- Thus, the output is not an intensity image but a stream of asynchronous events particularly at microsecond resolution, where each event comprises the time at which it is recorded and a position or an address of the respective pixel that detected the
15 brightness change as well as the sign of the brightness change (referred to as polarity).

Since events are caused by brightness changes over time, an event camera naturally responds to edges in the scene in presence of relative motion.

- Event cameras have numerous advantages over standard cameras: a latency in the
20 order of microseconds, low power consumption, and a very high dynamic range (130 dB compared to 60 dB of standard, frame-based cameras).

Most importantly, since all the pixels are independent, such sensors don't inherently suffer from motion blur.

- These properties render event cameras ideal sensors for applications where fast
25 response and high efficiency are crucial, and also for scenes with wide variations of illumination.

Additionally, since information is only sent in presence of brightness changes, the sensor removes inherent redundancy of standard cameras, thus requiring a very low

data rate (Kilobytes vs Megabytes). An example of such an event camera with an IMU is the DAVIS [3].

These properties make event cameras in applications where fast response and high efficiency are important and also in scenes with wide variations of illumination.

- 5 Additionally, since information is particularly only generated when the brightness in a pixel changes, the event camera removes inherent information redundancy of conventional frame-based cameras, thus requiring a lower data rate.

- 10 The output of the event camera is not an intensity image but a stream of particularly asynchronous events at microsecond resolution, where each event consists of its space-time coordinates and the sign of the brightness change (i.e. no intensity). Since events are caused by brightness changes over time, an event camera naturally responds to edges in the scene in presence of relative motion.

- Event cameras are used for visual odometry, i.e. the task for estimating a sensors position and orientation in space based on visual data. Visual odometry finds
15 applications particularly in simultaneous localization and mapping (SLAM) applications.

- In the following, orientation and position are summarized under the term "pose". The orientation and position each can comprise up to three degrees of freedom. The orientation can be described for example by the so-called Euler angles, wherein the
20 position can be specified by Cartesian coordinates in a specific frame of reference. Therefore, a pose comprises particularly six degrees of freedom.

In contrast to visual odometry or visual-based SLAM, one refers to visual-inertial odometry (VIO) or visual-inertial SLAM if not only visual data but also inertial data from an inertial measurement unit (IMU) are used for odometry or SLAM respectively.

- 25 State-of-the-art VIO and/or SLAM algorithms are capable of large-scale tracking, with an overall drift below 0.5% of the travelled distance [7]. However, VIO still fails in a number of situations, such as high-speed motions or high-dynamic range scenes.

- In the first case, large amounts of motion blur on the images spoil the visual information, forcing the method to rely on integration of the IMU data, resulting in
30 large amounts of accumulated drift.

In the second case, due to the limited dynamic range of frame-based cameras, large regions on the image are either over-, or under-exposed, which reduces drastically the amount of information exploitable.

5 SLAM and VIO methods can be categorized into methods comprising a non-linear optimization method, or a filtering-based method. Due to computational constraints, for a long time, real-time odometry or SLAM algorithms were only possible using a filtering approach, such as the Extended Kalman Filter (EKF) method [11].

However, non-linear optimization methods can provide better accuracy for a similar computational work, when compared to filtering methods.

10 A non-linear optimization or non-linear programming refers particularly to the use of a non-linear objective function with respect to the variable to be minimized or maximized. Furthermore, in the context of the current specification, non-linear optimization particularly refers to methods that optimize the variables for past data from the expedition, and thus iteratively or retrospectively adjusting for example the
15 generated three-dimensional map of the scene, wherein filtering methods are applied particularly only to a current measurement and the generated map is not adjusted retrospectively.

The use of event cameras for VIO and SLAM is known by the state of the art [11]. However, the disclosed method in [11] exhibits limited accuracy and processing
20 speed. In [11] a filtering approach is used for SLAM, rendering the method expectedly less accurate, as filtering methods are intrinsically less accurate than for example non-linear optimization approaches. Non-linear optimization approaches in turn, are known to be computational expensive and are not suited for the data structure provided by event cameras, as they require images rather than a stream of
25 events, limiting their use for real-time applications.

A computationally cost-efficient and accurate processing of the stream of events such that it becomes applicable for non-linear optimization VIO and SLAM algorithms while simultaneously maintaining the inherent higher degree of information is an unsolved problem.

30 Therefore, an object of the present invention is to provide a method and a computer program for generating a motion-corrected image from event cameras for visual-

inertial odometry and SLAM applications. The object is achieved by the method having the features of claim 1.

Advantageous embodiments are described in the subclaims.

According to claim 1, a method for generating a motion-corrected image for visual-inertial odometry (VIO) or visual-inertial SLAM comprises a particularly monocular event camera rigidly connected to an inertial measurement unit (IMU), wherein the event camera comprises pixels arranged in an image plane that are configured to particularly only output events in presence of brightness changes in a scene at the time they occur, wherein each event comprises the time at which it is recorded (also referred to as event time), a particularly two-dimensional position of the respective pixel (also referred to as an even position) that detected the brightness change as well as particularly a polarity value indicating the sign of the brightness change. The method according to the invention further comprises the steps of:

- Acquiring at least one set of events, wherein the at least one set comprises a plurality of subsequently or simultaneously recorded events from the event camera, particularly wherein the event time of the earliest event and the latest event delimit a time interval of the at least one set;
- Acquiring IMU data for the duration, particularly during the time interval of the at least one set;
- Generating a motion-corrected image from the at least one set of events, wherein the motion-corrected image is obtained by assigning the position of each event recorded at its corresponding event time t_j at an estimated event camera pose corresponding to the event time t_j to an adjusted event position, wherein the adjusted event position is obtained by determining the position of the event in a frame of reference of the event camera and an associated estimated reference camera pose at a reference time t_k^f , wherein the frame of reference of the event camera and the estimated camera pose at time t_j as well as the reference camera pose and the corresponding frame of reference of the event camera at time t_k^f are estimated by means of the IMU data.

The camera pose at time t_j and the reference camera pose at the time t_k^f each have a frame of reference of the event camera associated to it. The time t_k^f is also referred to as reference time.

Alternatively and equivalently, the motion-corrected image can be generated from the at least one set of events, by adjusting the event position recorded at its corresponding event time by transforming the event position in the frame of reference of the event camera at the corresponding event time t_j to the frame of reference of the event camera at the reference time t_k^f .

The reference time t_k^f can be the time of the earliest event comprised in the at least one set of events.

Pixel values of the motion-corrected image particularly correspond to the aggregated polarity of events from the at least one set at the adjusted event positions.

10 The motion-corrected image has a greatly reduced or no motion blur, such that any subsequent evaluation of the motion-corrected event image can yield more precise results. Particularly feature detection is greatly enhanced and can be performed with an increased reliability, as motion blur is absent in the motion-corrected images.

The processing time for generating the motion-corrected images is well within time scales that are considered as being real-time, i.e. particularly below 5 ms for 240 x 180 pixels, allowing the use of the method in time-sensitive applications.

The reference time for which the motion-corrected image is generated particularly corresponds to the earliest or latest event time of the set. The reference pose is particularly estimated from integrating the IMU data for a set of associated equations of motion. From an initial starting pose of the event camera it is particularly possible to estimate any later camera pose relative to the initial camera pose.

The estimated camera pose from the IMU data can be (re-)determined more precisely with subsequent image processing applications, such as VIO or SLAM. However, it is particularly not necessary to re-generate the motion-corrected images for a more precisely determined camera pose, even though it is possible. The event time is particularly provided with a microsecond resolution.

The events can be ordered within the at least one set by their corresponding event times.

According to an embodiment of the invention, from the IMU data the camera pose at the event time t_j is estimated particularly by means of a homogenous transformation T_{t_k, t_j}^f of the reference camera pose $T_{t_k}^f$ to the estimated camera pose T_{t_j} .

5 The transformation is particularly achieved by estimating a trajectory, comprising a pose of the event camera for a plurality of time points for the duration of the set, wherein the event camera poses can be interpolated for any specific time.

The transformation T_{t_k, t_j}^f is particularly a function that maps the frame of references of the event camera at the respective times to each other.

10 Unlike [11], a global motion correction, i.e. the generation of a motion correct image is performed globally (and not locally, independently for each feature tracked) for the pixels of the event camera. Instead of explicitly estimating the motion parameters ("optic flow") that allow warping of the events to a common frame (i.e. do motion correction), the method according to the invention particularly uses the following information to generate motion corrected image with a global transformation:

- 15
- Information regarding the IMU data, which particularly provides an estimate of the angular velocity.
 - Information from a previous odometry estimate, which particularly provides the global camera pose, and the camera linear velocity.
 - Information about the depth in the scene.

20 In contrast, [11] teaches a correction for the motion of the events by back-projecting said events on a current frame based on an estimated optical flow. The motion correction in [11] is part of a feature tracking method that estimates an optical flow, feature appearance and feature position by means of an optimization method. As such, the feature tracking method estimates the optical flow only locally, i.e. the
25 estimated optical flow is valid only for one feature and not all features in the image and thus the optical flow estimation has to be carried out once per feature, which is computationally expensive.

To simplify the problem and thus to reduce the computational load, [11] assumes that each feature is locally planar, i.e. that the 3D region corresponding to the feature
30 locally lies on a plane in 3D, which allows [11] to simplify the optimization. However,

even with this simplification, the optimization procedure must still be carried out for every feature tracked.

According to another embodiment of the invention, for the determination of the adjusted event position for the reference camera pose in the corresponding frame of reference of the event camera, a projection method is executed comprising the steps of:

- Estimating for the event time of the respective event a depth information for the event position of the event in a world reference frame for the estimated camera pose, wherein the depth information is estimated from a plurality of three-dimensional positions of landmarks in the world reference frame;
- Providing a camera projection model that is configured to describe the projection of a three-dimensional position onto a two-dimensional position in the image plane of the event camera;
- Using the camera projection model, the depth information and the estimated transformation to project the event position recorded at the estimated camera pose at the corresponding event time from the image plane onto a three-dimensional position in the world reference frame and re-projecting the three-dimensional position on which the event position has been projected from the world reference frame onto the image plane of the camera reference pose at the reference time t_k^f .

The projection method therefore comprises a transformation between the frame of reference of the event camera and the world reference frame. As the event camera potentially changes its pose over time, landmarks and features located in the scene which the camera records, change position and orientation in the frame of reference of the event camera.

The world reference frame is for example fixedly associated with the surrounding scene of the camera that is not attached to the camera.

The image plane of the camera reference pose is particularly the image plane of the event camera when it assumes the reference camera pose.

This embodiment particularly allows to make full use of the temporal information comprised in the events that is particularly absent in frame based cameras.

The projection model for the event camera can be for example determined by calibration measurements.

- 5 According to another embodiment of the invention, each pixel of the event camera has an associated depth information, such as a depth value, or is provided with an depth information. The depth information is particularly provided by landmarks located in the scene. In VIO and SLAM application these landmarks form the map within which the camera pose is estimated.
- 10 When no landmarks are yet identified, a constant depth information can be assumed for the events. In case the motion-corrected image is used in the context of a VIO or a SLAM algorithm, more precise information on the landmarks positions, and thus, more precise depth information is acquired as the algorithms process new motion-corrected images.
- 15 The depth information, particularly the depth information for each pixel, can be estimated for example by Delaunay triangulation, barycentric interpolation, or just the median depth of the landmarks and assume that the depth is constant over the scene.

- Particularly for the purpose of interpolation, the distribution of the features tracked in
- 20 3D is assumed to be planar.

- Such planar 3D feature distribution is different from [11] disclosing "locally planar feature" assumption. According to another embodiment of the invention, the transformation, particularly the homogenous transformation, for each event is determined from the depth information, the IMU data particularly providing the
- 25 camera angular velocity, and the state from the previous odometry estimate, particularly providing the camera linear velocity and camera pose.

Thus, an optical flow field is determined for the whole camera frame and not locally for each feature.

This allows a decoupling of the motion correction and the feature tracking.

This embodiment stands in strong contrast to [11], as [11] does not use depth information or previous odometry estimates to estimate a global optic flow field, valid on the entire camera frame. Rather, [11] estimates, for each feature, one

- 5 independent optic flow vector in each iteration of the method, in order to compensate for the motion of each feature independently. As a consequence, feature tracking and motion correction are heavily intertwined and cannot be decoupled.

According to the invention, it is assumed, particularly only assumed, that the distribution of landmarks in space is planar i.e. the 3D landmarks lie on a plane.

- 10 This approximation allows computing a global optical flow while not impacting the quality of estimated camera poses significantly.

An effect of estimating the depth information for each event position is that motion correction can be performed across all pixels of the camera particularly by means of the homogeneous transformation. I.e. the whole camera frame can be processed as

15 a whole. In contrast, [11] estimates a local optical flow for each landmark, i.e. for small windows having a typical size of a few pixels only, which, as pointed out above, is time and energy consuming.

- This embodiment allows for a decoupling of feature tracking and motion correction and allows for the processing of larger camera frames, while keeping the computational load comparably low. According to another embodiment of the
- 20 invention, the transformation is a global transformation such that a global motion correction for the events is achieved, particularly wherein the global transformation transforms the event positions for all pixels of the event camera by the same projection function, particularly wherein a depth information such as a depth value is
- 25 associated to each pixel of the event camera.

According to another embodiment of the invention, the motion-corrected image is given by

$$I(\mathbf{x}) = \sum_{e_j \in S} \delta(\mathbf{x} - \mathbf{x}'_j) \quad \text{Eq. (1)}$$

wherein $I(\mathbf{x})$ is the value of the motion-corrected image at the position $\mathbf{x}$, e_j is the j^{th} event in the set S of events, $\mathbf{x}'_j$ is the adjusted event position and δ stands for the spatial delta function, which equals one when $\mathbf{x}_j = \mathbf{x}'_j$ and which is zero otherwise.

According to another embodiment of the invention, the projection method is adjusting
5 the event position according to

$$\mathbf{x}'_j = \pi \left(T_{t_k^f, t_j} \left(Z(\mathbf{x}_j) \pi^{-1}(\mathbf{x}_j) \right) \right) \quad \text{Eq. (2)}$$

wherein $\mathbf{x}_j$ is the position of the j^{th} event from the set with an associated event time t_j and an estimated associated camera pose T_{t_j} , wherein $T_{t_k^f, t_j}^f$ is the transformation from the estimated camera pose at the time t_j to the reference pose of the event
10 camera $T_{t_k^f}$ at the time t_k^f and wherein π is the camera projection model and Z is the depth information for the respective event position $\mathbf{x}_j$.

Accordingly, π^{-1} refers to the inverted camera projection model.

This embodiment details the global approach of motion correction that allows for a decoupling of the motion correction and the feature tracking.

15 It can be seen that Eq. (2) is an equation that does not involve locally varying optical flows or the like, but particularly a homogenous global transformation is applied to each event.

According to another embodiment of the invention, a plurality of sets is generated, particularly wherein each set of the plurality of sets comprises events recorded
20 temporally subsequently or simultaneously, and particularly wherein the resulting associated time intervals are delimited by the respective earliest and latest events comprised in the respective set, particularly wherein the plurality of sets is temporally ordered such that two subsequent sets are adjacent i.e. they do not comprise identical events, or the sets are overlapping, i.e. they comprise the same events by a
25 predefined number of events, wherein for each set the corresponding motion-corrected image is generated from the respective events comprised in the set, such that a temporal series of the motion-corrected images is obtained.

The sets can comprise the same number of events per set.

The duration of the sets can vary depending on the rate the events are recorded by the event camera. Therefore, the motion-corrected images are particularly also generated at a variable rate. In contrast, the IMU data is particularly acquired at a fixed frequency. Thus, the acquisition of IMU data and the generation of sets particularly happen asynchronous.

This embodiment allows for subsequent VIO and SLAM algorithms to be performed on the series of motion-corrected images.

According to another embodiment of the invention, on a series of keyframe images that consists of a plurality of particularly selected motion-corrected images, a simultaneous localization and mapping (SLAM) method or a visual-inertial odometry (VIO) method is performed.

The concept of particularly selecting a subset of motion-corrected images as keyframe images has the advantage that the computational load can be controlled and is kept comparably low.

The term "keyframe image" is used herein as commonly used in the art. The selection of keyframe images is well-known to a person skilled in the art.

In the example section it is elaborated how a selection of motion-corrected images can be made.

According to another embodiment of the invention, SLAM method is a visual-inertial SLAM method, wherein the VIO method and/or the visual-inertial SLAM method employs the keyframe images and the IMU data to obtain a three-dimensional map comprising the landmark positions and an estimated keyframe image camera pose for each keyframe image, particularly wherein the landmark positions and the estimated keyframe image camera pose are estimated with respect to the world reference frame.

The keyframe image camera pose is particularly the reference camera pose of the respective motion-corrected image that corresponds to the keyframe image.

In the context of the application the terms “SLAM”, “SLAM method” and “SLAM algorithm” are used synonymously.

According to another embodiment of the invention, the visual-inertial SLAM method and/or the VIO method comprises a non-linear optimization method, wherein the non-linear optimization method minimizes a joint cost function JC for all keyframe images particularly by varying the estimated keyframe image camera pose T^m for each keyframe image and the three-dimensional landmark positions l^n . Here, m is an index for the M keyframe images and n is an index for the N landmarks.

This global optimization with the joint cost function (in the literature also referred to with the term “objective function”) that particularly involves also the estimation of previously estimated camera poses and landmark positions (i.e. a re-evaluation of past results) allows for more precise SLAM and/or VIO results, i.e. more precise maps and camera pose estimations, which in turn allows for longer expeditions, while the mapping and localization error remains comparably small.

The combination of the generation of the motion-corrected images, and the non-linear optimization in a SLAM and/or VIO algorithm allows for the fast and precise execution of the method according to the invention. This embodiment and the dependent embodiments allow a more precise and a faster execution of SLAM and/or VIO on event based cameras particularly in fast moving environments, where the event camera is moving rapidly relative to the surrounding scene and the lighting conditions might be highly dynamic within the field of view of the event camera.

According to another embodiment of the invention, the joint cost function comprises a sum of an inertial error term for each keyframe image and a re-projection error term for each landmark and for each keyframe image, wherein the re-projection error term and the inertial error term are minimized, wherein the cost function is particularly given by:

$$JC := \sum_{m=1}^M \sum_{n \in J(m)} \epsilon_r^{n,mT} W_r^{n,m} \epsilon_r^{n,m} + \sum_{m=1}^{M-1} \epsilon_s^{mT} W_s^m \epsilon_s^m \quad \text{Eq. (3)}$$

wherein JC is the joint cost function, M is the number of keyframe images, $J(m)$ are the landmarks comprised in the respective m^{th} keyframe image, $\varepsilon_r^{n,m}$ is a re-projection error of the n^{th} landmark in the m^{th} keyframe, $W_r^{n,m}$ is an associated weighting factor for the re-projection error, ε_s^m is an inertial error of the m^{th} keyframe image and W_s^m is the associated weighting factor of the inertial error, wherein the re-projection error is given by $\varepsilon_r^{n,m} = z^{n,m} - \pi(T_{CW}^m l^n)$ wherein $z^{n,m}$ is an estimated two-dimensional landmark position of the n^{th} landmark in the m^{th} keyframe image, π is the camera projection model, and T_{CW}^m is particularly a homogenous transformation of a camera frame of reference to the world reference frame for the m^{th} keyframe image for the varied camera pose, l^n is the estimated three-dimensional landmark position, wherein T_{CW}^m (and thus the camera pose T^m for the m^{th} keyframe image) and l^n are varied, wherein the inertial error ε_s^m is estimated from the difference between the estimated camera pose from the IMU data for the m^{th} keyframe image and the varied camera pose T^m .

- 15 According to another embodiment of the invention, the method is executed and executable more than 100 times per second, particularly more than 200 times per second.

Such execution speed is also referred to with the term “real-time”, as the time lag between camera acquisition and the generation of the motion-corrected image as well as the determination of the camera pose and the updated map (i.e. landmark positions) is not or almost unnoticeable by a person.

According to another embodiment of the invention, the number of events per set is more than 10^3 and particularly less than 10^6 particularly for an event camera comprising 240 x 180 pixels.

- 25 With this number of events a reliable processing for subsequent SLAM algorithms is granted.

According to another embodiment of the invention, the sets are overlapping by more than 10^2 events or wherein the sets not overlapping, i.e. particularly adjacent to each other, particularly for an event camera comprising 240 x 180 pixels.

The problem according to the invention is furthermore solved by a method for visual-inertial odometry with an event camera, wherein the event camera is rigidly connected to an inertial measurement unit (IMU), wherein the event camera comprises pixels that are configured to output events in presence of brightness changes in a scene at the time they occur, wherein each event comprises the time at which it is recorded (also referred to as event time) and a position (also referred to as event position) of the respective pixel that detected the brightness change, the method comprising the steps of:

- Acquiring IMU data with the IMU unit;
- Acquiring a plurality of sets each comprising a plurality of events;
- For each set generating an image from the events comprised in the respective set, wherein each image is obtained by assigning the event position to a pixel coordinate of the image, such that a series of images is obtained;
- Perform on a series of keyframe images that consists of a plurality of images from the series of images, a visual inertial simultaneous localization and mapping (SLAM) method and/or a visual inertial odometry (VIO) method, wherein the VIO method and/or the visual-inertial SLAM method employ the keyframe images and the IMU data to obtain a three-dimensional map comprising the landmark positions and an estimated keyframe image camera pose for each keyframe image wherein, the VIO method and/or the visual-inertial SLAM method comprises a non-linear optimization method, wherein the non-linear optimization method minimizes a joint cost function for all keyframe images, particularly by varying the estimated keyframe image camera pose for each keyframe image and the three-dimensional landmark positions.

The method allows for a rapid and accurate execution of SLAM and/or VIO algorithms such that particularly in real time or with a time lag of less than 10 ms per processed image (with an camera size of 240 x 180 pixels), a map of the surrounding scene of the event camera and its pose in the scene can be obtained.

The obtained accuracy and processing speed is exceeding state of the art methods for event camera based SLAM or VIO algorithms. This is particularly due to the use of the non-linear optimization method.

5 The concept of particularly selecting a subset of motion-corrected images as keyframe images has the advantage that the computational load can be controlled and is kept lower.

The keyframe image camera pose is particularly the reference camera pose of the respective image that corresponds to the keyframe image.

10 In the example section it is elaborated how a selection of motion-corrected images can be made.

According to another embodiment of the invention, the joint cost function comprises a sum of an inertial error term for each keyframe image and a re-projection error term for each keyframe image and for each landmark comprised in the respective keyframe image, wherein the re-projection error term and the inertial error term are
15 minimized.

According to another embodiment of the invention, the cost function is particularly given by:

$$JC := \sum_{m=1}^M \sum_{n \in \mathcal{I}(m)} \boldsymbol{\varepsilon}_r^{n,mT} \mathbf{W}_r^{n,m} \boldsymbol{\varepsilon}_r^{n,m} + \sum_{m=1}^{M-1} \boldsymbol{\varepsilon}_S^mT \mathbf{W}_S^m \boldsymbol{\varepsilon}_S^m$$

wherein JC is the joint cost function, M is the number of keyframe images, $\mathcal{I}(m)$ are the landmarks comprised in the respective m^{th} keyframe image, $\boldsymbol{\varepsilon}_r^{n,m}$ is a re-
20 projection error of the n^{th} landmark in the m^{th} keyframe, $\mathbf{W}_r^{n,m}$ is an associated weighting factor for the re-projection error, $\boldsymbol{\varepsilon}_S^m$ is an inertial error of the m^{th} keyframe image and $\mathbf{W}_S^m$ is the associated weighting factor of the inertial error.

According to another embodiment of the invention, the re-projection error is given by $\boldsymbol{\varepsilon}_r^{n,m} = \mathbf{z}^{n,m} - \pi(T_{CW}^m \mathbf{l}^n)$ wherein $\mathbf{z}^{n,m}$ is an estimated two-dimensional landmark
25 position of the n^{th} landmark in the m^{th} keyframe image, π a camera projection model, and T_{CW}^m is particularly a homogenous transformation of a camera frame of reference

to a world reference frame for the m^{th} keyframe image for the varied camera pose, l^n is the estimated three-dimensional landmark position, wherein T_{CW}^m and l^n are varied, wherein the inertial error ϵ_S^m is estimated from the difference between the estimated camera pose from the IMU data for the m^{th} keyframe image and the varied camera
5 pose T^m .

According to another embodiment of the invention, the method is executed and executable more than 100 times per second, particularly more than 200 times per second.

Such execution speed is also referred to with the term "real-time", as the time lag
10 between camera acquisition and the generation of the motion-corrected image as well as the determination of the camera pose and the updated map (i.e. landmark positions) is not or almost unnoticeable by a person.

According to another embodiment of the invention, the number of events per set is more than 10^3 and particularly less than 10^6 particularly for an event camera
15 comprising 240 x 180 pixels.

With this number of events a reliable processing for subsequent SLAM algorithms is granted.

According to another embodiment of the invention, the sets are overlapping by more than 10^2 events or wherein the sets not overlapping, i.e. particularly adjacent to each
20 other, particularly for an event camera comprising 240 x 180 pixels.

The problem according to the invention is also solved by a computer program comprising computer program code, wherein the computer program code is configured to execute the method according to the invention, when the computer program is executed on a computer.

25 According to another embodiment of the invention, the computer on which the computer program is executed is comprised or connected to the IMU and/or the event camera.

The connection of the computer to the IMU can be wireless, particularly via electromagnetic wave, such as radiofrequency or optic transmission, or via electric conducting means, such as cables or a printed circuit board.

5 According to another embodiment of the invention, the computer program is configured to transmit data from or to the IMU and/or the event camera to the computer, particularly wirelessly.

The terms 'processor' or 'computer', or system thereof, are used herein as ordinary context of the art, such as a general purpose processor or a micro-processor, RISC processor, or DSP, possibly comprising additional elements such as memory or communication ports. Optionally or additionally, the terms 'processor' or 'computer' or derivatives thereof denote an apparatus that is capable of carrying out a provided or an incorporated program and/or is capable of controlling and/or accessing data storage apparatus and/or other apparatus such as input and output ports. The terms 'processor' or 'computer' denote also a plurality of processors or computers connected, and/or linked and/or otherwise communicating, possibly sharing one or more other resources such as a memory.

20 The terms 'software', 'program', 'software procedure' or 'procedure' or 'software code' or 'code' or 'application' or 'app' may be used interchangeably according to the context thereof, and denote one or more instructions or directives or circuitry for performing a sequence of operations that generally represent an algorithm and/or other process or method. The program is stored in or on a medium such as RAM, ROM, or disk, or embedded in a circuitry accessible and executable by an apparatus such as a processor or other circuitry.

25 The processor and program may constitute the same apparatus, at least partially, such as an array of electronic gates, such as FPGA or ASIC, designed to perform a programmed sequence of operations, optionally comprising or linked with a processor or other circuitry.

30 As used herein, without limiting, a module represents a part of a system, such as a part of a program operating or interacting with one or more other parts on the same unit or on a different unit, or an electronic component or assembly for interacting with one or more other components.

As used herein, without limiting, a process represents a collection of operations for achieving a certain objective or an outcome.

The term 'configuring' and/or 'adapting' for an objective, or a variation thereof, implies using at least a software and/or electronic circuit and/or auxiliary apparatus designed
5 and/or implemented and/or operable or operative to achieve the objective.

A device storing and/or comprising a program and/or data constitutes an article of manufacture. Unless otherwise specified, the program and/or data are particularly stored in or on a non-transitory medium.

In the following Figure and Example section, the invention is explained in detail with
10 reference to exemplary embodiments shown in the figures. It is noted that the drawings are not to scale. Features and embodiments disclosed in the example section even though disclosed in combination with other exemplary features can be separately combined with other features disclosed in another embodiment. It is noted
15 that the processing pipeline for the methods according to the invention laid out in the following is particularly also applicable for the event images that have not been motion-corrected. For this purpose the steps and the prerequisites for generating the motion-corrected images (particularly laid out in claim 1 and its dependent subclaims) can be left out or replaced by a step that solely generates an image $I_{conv}(x)$ (as set
20 out below and in claims 11 to claim 14) from the acquired events. Such an imager can be processed with the non-linear optimization method as laid below and throughout the description. This is particularly suitable in case rapid scene changes are not to be expected and/or when the illumination conditions are favourable.

Preliminaries. In this section, a notation is introduced that is used throughout the rest of the description. An IMU model is disclosed that is used for the calculations
25 comprising IMU data. Furthermore, formulas for discrete integration of the equations of motion are provided for the IMU data.

Coordinate Frame Notation. The notation of [5] is adopted. A point P represented in a frame of reference A is written as a position vector ${}_A r_P$. A transformation between frames of reference such as the world reference frame W and the frame of reference
30 of the event camera C is represented by a homogeneous matrix T_{WC} that transforms points from the world reference frame W to reference frame C . The rotational part of

T_{WC} is expressed as a rotation matrix $R_{WC} \in SO(3)$. In experiments a DAVIS sensor [3], comprising the event camera and the IMU, is used. As the event camera and the IMU are rigidly coupled to each other, the camera frame of reference and a frame of reference of the IMU can be readily transformed to each other. The transformation

5 can be obtained by an extrinsic calibration of the event camera and the IMU, using for example the Kalibr toolbox [6]. For reasons of simplicity and without loss of generality it is assumed that the frame of references of the IMU and the event camera are identical.

IMU Model and Motion Integration. The IMU includes a 3-axis accelerometer and a 3-axis gyroscope, and allows measuring the rotational rate and the acceleration of the

10 IMU and thus the event camera with respect to an inertial frame such as the world reference frame W . These measurements are referred to as IMU data.

The IMU data for acceleration $\mathbf{a}$ and the gyroscopic data $\boldsymbol{\omega}$, are affected by additive white noise $\boldsymbol{\eta}$ and slowly varying sensor biases $\mathbf{b}$

15
$${}_c\boldsymbol{\omega}(t) = {}_w\boldsymbol{\omega}(t) + \mathbf{b}_g(t) + \boldsymbol{\eta}_g(t) \text{ and } \mathbf{a}(t) = \mathbf{R}_{WC}^t(t) ({}_w\mathbf{a}(t) - {}_w\mathbf{g}) + \mathbf{b}_a(t) + \boldsymbol{\eta}_a(t)$$

Wherein ${}_w\mathbf{g}$ is the gravity vector in the world reference frame. Denoting the position vector and the velocity as ${}_w\mathbf{r}(t)$ and ${}_w\mathbf{v}(t)$, respectively, the equations of motion can be numerically integrated as shown in [4].

Eq. 4:

$$\mathbf{R}_{WB}(t + \Delta t) = \mathbf{R}_{WB}(t) \exp({}_c\boldsymbol{\omega}(t) - \mathbf{b}_g(t) - \boldsymbol{\eta}_{gd}(t)\Delta t)$$

20
$${}_w\mathbf{v}(t + \Delta t) = {}_w\mathbf{v}(t) + {}_w\mathbf{g}\Delta t + \mathbf{R}_{WB}(t) ({}_c\mathbf{a}(t) - \mathbf{b}_a(t) - \boldsymbol{\eta}_{ad}(t))\Delta t$$

$${}_w\mathbf{r}(t + \Delta t) = {}_w\mathbf{r}(t) + {}_w\mathbf{v}(t)\Delta t + 0.5 {}_w\mathbf{g}\Delta t^2 + 0.5 \mathbf{R}_{WB}(t) ({}_c\mathbf{a}(t) - \mathbf{b}_a(t) - \boldsymbol{\eta}_{ad}(t))\Delta t^2$$

where $\boldsymbol{\eta}_{ad}(t)$ and $\boldsymbol{\eta}_{gd}(t)$ are noise variables.

This integration scheme performs very well in practice, provided the IMU sampling rate is comparably high, e.g. above 100 Hz.

Visual-inertial odometry (VIO) is typically performed in two parallel threads:

- A front-end that takes the events from the event camera as an input. The front-end generates the motion-corrected images and detects and tracks features and landmarks in the motion-corrected images.
- 5 - A back-end fuses the feature tracks, the landmarks and the IMU data to continuously update the current and past camera poses.

Generation of motion-corrected images. A stream of events is provided by the event camera. The events are split in sets S_k to generate motion-corrected images. The k^{th} set S_k comprises a predefined number of events. Optionally the sets can overlap by a
 10 predefined number of events that are shared by two adjacent sets. However, it is also possible to generate non-overlapping sets that are adjacent to each other, i.e. they do not share events.

A start time t_k^f and a duration Δt_k^f of each set are controlled by the events, which preserves the data driven nature of the method. The k^{th} set S_k spans the time interval
 15 $[t_k^f, t_k^f + \Delta t_k^f]$. The duration is determined by the first and last event of the set.

In order to arrive at an image from the set, conventionally one can accumulate the events as follows: $I_{conv}(\mathbf{x}) = \sum_{e_j \in S} \delta(\mathbf{x} - \mathbf{x}_j)$, i.e. the intensity of the image $I_{conv}(\mathbf{x})$ is simply the sum of the events that comprise an event position $\mathbf{x}_j = \mathbf{x}$. However, this
 20 yields an event image that is not usable for reliable feature detection or tracking, as illustrated in Fig. 3. Small set sizes (Fig. 3A) do not convey enough information, while large set sizes (Fig. 3B) induce motion blur.

In contrast, the k^{th} motion-corrected image is generated from the k^{th} set S_k of events as follows (see e.g. Fig. 2 and Fig. 3C).

$$I_k(\mathbf{x}) = \sum_{e_j \in S_k} \delta(\mathbf{x} - \mathbf{x}_j') \quad \text{Eq. (5)}$$

25 where $\mathbf{x}_j'$ is the adjusted event position, obtained by transferring the j^{th} event recorded at time t_j with an event position $\mathbf{x}_j$ in a frame of reference C_{t_j} of the event

camera at time t_j to a frame of reference $C_{t_k^f}$ of the event camera associated to a reference camera pose $T_{t_k^f}$ at the time t_k^f . This transfer is given by:

$$\mathbf{x}'_j = \pi \left(T_{t_k^f, t_j} \left(Z(\mathbf{x}_j) \pi^{-1}(\mathbf{x}_j) \right) \right), \quad \text{Eq. (6)}$$

where $\mathbf{x}_j$ is the position of the j^{th} event from the set with an associated event time t_j and an estimated associated camera pose T_{t_j} , wherein $T_{t_k^f, t_j}$ is the transformation between the frames of reference C_{t_j} , $C_{t_k^f}$ of the event camera for the respective time points from the estimated camera pose T_{t_j} at the time t_j to the reference pose of the event camera $T_{t_k^f}$ at the time t_k^f and wherein π is the camera projection model and Z is the depth information for the respective event position $\mathbf{x}_j$.

- 10 The incremental transformation $T_{t_k^f, t_k^f + \Delta t_k^f}$ is obtained through integration of the equations of motion (see. Eq. 4) for the IMU data acquired during the time interval $[t_k^f, t_k^f + \Delta t_k^f]$. The reference pose $T_{t_k^f}$, and the IMU biases $\mathbf{b}$ are known from the state estimation thread. The remaining quantities required to evaluate Eq. (6) are the transformation $T_{t_k^f, t_j}$, which can be linearly interpolated from $T_{t_k^f}$ and the camera
- 15 pose $T_{t_k^f + \Delta t_k^f}$ at the time $t_k^f + \Delta t_k^f$, in the space of rigid-body motion transformations. Furthermore, $Z(\mathbf{x}_j)$ is estimated using two-dimensional linear interpolation (on the image plane) of the landmarks, re-projected on the frame of reference C_{t_j} of the camera pose T_{t_j} . Alternatively, the median can be used instead of interpolating the depth information $Z(\mathbf{x}_j)$. Another option is to use three-dimensional Delaunay
- 20 triangulation to interpolate the sparse depth information $Z(\mathbf{x}_j)$.

Feature detection. New features are detected whenever the number of feature tracks falls below a certain threshold, or if the current motion-corrected image is selected as a keyframe image (see below). A so-called FAST corner detector [9] on the motion-corrected image can be used for this purpose. A bucketing grid is used to ensure that

25 the features are evenly distributed over the motion-corrected image.

Feature Tracking and Landmark Triangulation. Two types of landmarks are maintained: candidate landmarks, and persistent landmarks, whose three-dimensional position in space has been successfully triangulated. Newly extracted features are initialized as candidate landmarks, and are tracked across the motion-corrected images. As soon as a candidate landmark can be reliably triangulated (i.e. with bearing rays enclosing an angle greater than e.g. 5°), it is converted to a persistent landmark, and added to the map.

Both types of landmarks are tracked in subsequent motion-corrected images $I_k(\mathbf{x})$ and $I_{k+1}(\mathbf{x})$ using pyramidal Lukas-Kanade tracking [2]. The incremental transformation $T_{t_k^f, t_{k+1}^f}^f$ (integrated from the IMU data) for the time interval between $[t_k^f, t_{k+1}^f]$ is used to predict the feature position in $I_{k+1}(\mathbf{x})$. The patches around each feature are warped through an affine warp, computed using $T_{t_k^f, t_{k+1}^f}^f$, prior to pyramidal alignment. Landmarks that are not successfully tracked in the current motion-corrected event image are discarded immediately. The depth associated to persistent landmarks is used in that step. If it is not known, then only the incremental rotation is used.

Outlier Filtering. Two-point RANSAC [10] is used (using the relative orientation between the current motion-corrected image and the last keyframe image) to further filter out outlier feature tracks. This is done between the current motion-corrected image and the last keyframe image.

Keyframe Image Selection. A new keyframe image is selected either when the number of tracked features falls below a predefined threshold, or when the distance to the last keyframe image (scaled by the median scene depth) reaches a minimum threshold.

Initialization. To initialize the SLAM method, the first motion-corrected images are added to the back-end without initializing any feature track. The back-end in turn estimates the initial attitude of the sensor by observing the gravity direction. The displacement between the subsequent motion-corrected images is then estimated by integrating IMU data.

Back-End. In this section, it is exemplary described how feature tracks from the event stream obtained by the feature tracking and triangulation is fused to update the full sensor state over time.

As opposed to the Extended Kalman Filter (EKF)-based filtering employed in [11], a full smoothing approach based on non-linear optimization on selected keyframe images is used.

This is computationally tractable by use of pre-integration theory [8, 5], that consists of combining many inertial measurements (IMU data) between two keyframe images into a single relative motion constraint, thus avoiding to re-integrate IMU data in each step of the optimization. This approach considerably accelerates the optimization procedure. The back-end implementation can be based on OKVIS [7].

Formulation of Visual-Inertial Odometry. The visual-inertial localization and mapping problem is formulated as a joint optimization of a cost function JC that contains weighted re-projection errors and inertial error terms:

$$JC := \sum_{m=1}^M \sum_{n \in J(m)} \boldsymbol{\varepsilon}_r^{n,mT} \mathbf{W}_r^{n,m} \boldsymbol{\varepsilon}_r^{n,m} + \sum_{m=1}^{M-1} \boldsymbol{\varepsilon}_s^{mT} \mathbf{W}_s^m \boldsymbol{\varepsilon}_s^m \quad \text{Eq. (7)}$$

where JC is the joint cost function, M is the number of keyframe images, $J(m)$ are the landmarks comprised in the respective m^{th} keyframe image, $\boldsymbol{\varepsilon}_r^{n,m}$ is a re-projection error of the n^{th} landmark in the m^{th} keyframe, $\mathbf{W}_r^{n,m}$ is the information matrix of the landmark measurement for the re-projection error, $\boldsymbol{\varepsilon}_s^m$ is an inertial error of the m^{th} keyframe image and $\mathbf{W}_s^m$ is the information matrix of the inertial error.

The optimization is carried out, not on all the keyframe images, but on a bounded set of keyframe image composed of M keyframe images, and a sliding window containing the last H frames. In between keyframe images, the prediction of the event camera state and pose is propagated using the IMU data that fall in between the keyframe images. Google Ceres [1] optimizer is used to carry out the optimization.

Re-projection Error. The re-projection error is given by $\boldsymbol{\varepsilon}_r^{n,m} = \mathbf{z}^{n,m} - \pi(T_{CW}^m \mathbf{l}^n)$ wherein $\mathbf{z}^{n,m}$ is an estimated two-dimensional landmark position of the n^{th} landmark in the m^{th} keyframe image, π is the camera projection model, and T_{CW}^m is particularly a

homogenous transformation of a camera frame of reference to the world reference frame for the m^{th} keyframe image for the varied camera pose, $\mathbf{l}^n$ is the estimated three-dimensional landmark position, wherein T_{CW}^k (and thus the camera pose T^m for the m^{th} keyframe image) and $\mathbf{l}^n$ are varied.

- 5 *Inertial Error.* The inertial error ε_S^m is estimated from the difference between the estimated camera pose from the IMU data for the k^{th} keyframe image and the varied camera pose T^m . The IMU equations of motion (Eq. 4) are used to predict the current state based on the previous state. Then, the inertial errors are simply computed as the difference between the predictions based on the previous state and the actual state. For orientation, a simple multiplicative minimal error is used.

Keyframe Image Marginalization. Keeping all keyframe images in the Gauss-Newton system state quickly becomes intractable. However, simply discarding measurements from past keyframe images neglects valuable information. To overcome this problem, old keyframe images are partially marginalized out using the Schur complement on the corresponding observations. This turns old measurements into a prior for the system, represented as summands in the construction of the Schur complement. Details are disclosed in [7].

Parameters. The sets size can be selected for all datasets, particularly in the range of 10^3 to 10^5 events. This translates to a set duration of about 5 to 10 milliseconds. The patch size used for feature tracking can be 32 x 32 pixels, with 2 pyramid levels. The event camera has 240 x 180 pixels with a 1 kHz IMU (DAVIS sensor [3]).

These are exemplary parameters that can be adapted depending on the event camera and IMU used, the environment, the desired trade-off between accuracy and performance, etc..

25 Fig. 1 shows a schematic time line representation (indicated by the parallel horizontal arrows) for events recorded at different times. The events are grouped in a plurality of overlapping sets (indicated by the horizontal lines above the time line arrow), wherein each set consists of the same number of events (4 events in the schematic case) and the overlap of the sets is constant (namely two events on both ends of each set). The events are depicted as dots on the upper time line arrow, wherein the earliest and latest event times of each set are written below the dots and connected

with a dashed line. Note that the lower index of the earliest and latest event times refers to the corresponding set number. The duration of each set depends on the event time of the first (earliest) event of each set and the last (latest) event of each set, i.e. the event rate.

- 5 Fig. 2 schematically shows how the event position is adjusted for the motion of the event camera during the set acquisition of the k^{th} set. Landmarks are indicated as l_1 to l_4 . The trajectory of the event camera and its associated camera pose is depicted as a solid, curved arrow line, wherein the camera pose $T_{t_k^f}$ at the time t_k^f is indicated by a triangle. Accordingly, the camera pose $T_{t_k^f + \Delta t_k^f}$ at the latest event time $t_k^f + \Delta t_k^f$ of the set is shown. On the trajectory dots represent the recording time of an event (e), and squares indicate the acquisition of IMU data (D).

The camera poses $T_{t_k^f}$ and $T_{t_k^f + \Delta t_k^f}$ are estimated by integration (see Eq. 4) of the acquired IMU data. The intermediate camera pose T_{t_j} at the event time t_j can be estimated by interpolation.

- 15 In order to adjust the position of the event recorded at time t_j for a motion-corrected image, the event position is projected on a three-dimensional world reference frame, by means of a depth information Z obtained from the landmark positions, and a projection model of the event camera. From the world reference frame the three-dimensional position of the projected event is re-projected and transformed in a frame of reference of the camera at the reference time t_k^f that corresponds to a camera pose at the reference time.

By adjusting all event positions to the frame of reference of the event camera for the reference time and the associated reference camera pose, the motion-corrected image is obtained.

- 25 Fig. 3 shows various examples of image formation from events. Each event is depicted as a black dot at its recorded position. Fig. 3A and Fig. 3B are images generated without adjusting the event positions. Fig. 3A shows an example where only 3.000 events have been used to generate an image. While a motion blur is not readily detectable (because the set time is so short, that the camera hardly moves during the time interval) the image is corrupted by noise, owed to the small number of events.

If a greater number of events, here 30.000, is collected per set, motion blur becomes visible and prominent (see Fig. 3B). On such images VIO or SLAM methods will not be able to reliably identify features and feature positions, which render such images problematic for VIO or SLAM.

- 5 In Fig. 3C, a motion-corrected image according to the invention is shown. The event number is 30.000 per set. The motion-corrected image exhibits neither motion blur nor excessive noise and is therefore ideally suited for VIO or SLAM.

References

- [1] A. Agarwal, K. Mierle, and Others. Ceres solver. <http://ceres-solver.org>.
- 10 [2] S. Baker and I. Matthews. Lucas-kanade 20 years on: A unifying framework. *Int. J. Comput. Vis.*, 56(3):221–255, 2004.
- [3] Christian Brandli, Raphael Berner, Minhao Yang, Shih-Chii Liu, and Tobi Delbruck. A 240x180 130dB 3us latency global shutter spatiotemporal vision sensor. *IEEE J. Solid-State Circuits*, 49(10):2333–2341, 2014. ISSN 0018-9200. doi: 10.1109/JSSC.2014.2342715.
- 15 [4] Christian Forster, Luca Carlone, Frank Dellaert, and Davide Scaramuzza. IMU preintegration on manifold for efficient visual-inertial maximum-a-posteriori estimation. In *Robotics: Science and Systems (RSS)*, 2015. doi: 10.15607/RSS.2015.XI.006.
- 20 [5] Christian Forster, Luca Carlone, Frank Dellaert, and Davide Scaramuzza. On-manifold preintegration for real-time visual-inertial odometry. *IEEE Trans. Robot.*, 33(1):1–21, 2017. doi: 10.1109/TRO.2016.2597321.
- [6] P. Furgale, J. Rehder, and R. Siegwart. Unified temporal and spatial calibration for multi-sensor systems. In *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2013.
- 25 [7] S. Leutenegger, S. Lynen, M. Bosse, R. Siegwart, and P. Furgale. Keyframe-based visual-inertial odometry using nonlinear optimization. *Int. J. Robot. Research*, 2015.

- [8] T. Lupton and S. Sukkarieh. Visual-inertial-aided navigation for high-dynamic motion in built environments without initial conditions. *IEEE Trans. Robot.*, 28(1):61–76, February 2012.
- 5 [9] E. Rosten and T. Drummond. Machine learning for high-speed corner detection. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 430–443, 2006. doi: 10.1007/11744023_34.
- [10] C. Troiani, A. Martinelli, C. Laugier, and D. Scaramuzza. 2-point-based outlier rejection for camera-imu systems with applications to micro aerial vehicles. In *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2014.
- 10 [11] A. Zhu, N. Atanasov, and K. Daniilidis. Event-based visual inertial odometry. In *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recog.*, 2017.

Claims

1. A method for generating a motion-corrected image for visual-inertial odometry using an event camera rigidly connected to an inertial measurement unit (IMU), wherein the event camera comprises pixels arranged in an image plane that are configured to output events (e) in presence of brightness changes in a scene at the time they occur, wherein each event (e) comprises the time at which it is recorded and a position of the respective pixel that detected the brightness change, the method comprising the steps of:
 - Acquiring at least one set of events (S), wherein the at least one set (S) comprises a plurality of subsequent events (e);
 - Acquiring IMU data (D) for the duration of the at least one set (S);
 - Generating a motion-corrected image from the at least one set (S) of events (e), wherein the motion-corrected image is obtained by assigning the position (x_j) of each event (e_j) recorded at its corresponding event time (t_j) at an estimated camera pose (T_{t_j}) to an adjusted event position (x'_j), wherein the adjusted event position (x'_j) is obtained by determining the position of the event (e_j) for an estimated reference camera pose ($T_{t_k^f}$) at a reference time (t_k^f), wherein the estimated camera pose (T_{t_j}) at the event time (t_j) and the estimated reference camera pose ($T_{t_k^f}$) at the reference time (t_k^f) are estimated by means of the IMU data (D).
2. Method according to claim 1, wherein from the IMU data the estimated camera pose (T_{t_j}) at the event time (t_j) is estimated particularly by means of a homogenous transformation ($T_{t_k^f, t_j}^f$) of the estimated reference camera pose ($T_{t_k^f}$) at the reference time (t_k^f) to the estimated camera pose (T_{t_j}).
3. Method according to claim 2, wherein for the determination of the adjusted event position (x'_j) for the estimated reference camera pose ($T_{t_k^f}$), a projection method is executed comprising the steps of:

- Estimating a depth information (Z) for each event position in a world reference frame for the estimated camera pose (T_{t_j}), wherein the depth information (Z) is estimated from a plurality of three-dimensional positions of landmarks;
- 5 - Providing a camera projection model (π) that is configured to project a three-dimensional position onto a two-dimensional position in the image plane of the event camera;
- Using the camera projection model (π), the depth information (Z) and the transformation ($T_{t_k^f, t_j}$) to project each event position recorded at the
- 10 estimated camera pose at its corresponding event time from the image plane onto a three-dimensional position in the world reference frame and re-projecting the event position from the world reference frame onto the image plane of the estimated reference camera pose ($T_{t_k^f}$) at the reference time (t_k^f).

- 15 4. Method according to one of the claims 1 to 3, wherein the motion-corrected image is given by

$$I(\mathbf{x}) = \sum_{e_j \in S} \delta(\mathbf{x} - \mathbf{x}_j')$$

wherein $I(\mathbf{x})$ is the value of the motion-corrected image at the position $\mathbf{x}$, e_j is the j^{th} event in the set S of events, and $\mathbf{x}_j'$ is the adjusted event position, wherein δ is the delta-function.

- 20 5. Method according to claim 3, wherein the projection method adjusts each event position according to

$$\mathbf{x}_j' = \pi \left(T_{t_k^f, t_j} \left(Z(\mathbf{x}_j) \pi^{-1}(\mathbf{x}_j) \right) \right)$$

wherein $\mathbf{x}_j$ is the position of the j^{th} event from the set with an corresponding event time t_j and an estimated camera pose (T_{t_j}), wherein $T_{t_k^f, t_j}$ is the

transformation ($T_{t_k^f t_j}$) from the estimated camera pose (T_{t_j}) at the event time (t_j) to the estimated reference camera pose ($T_{t_k^f}$) at the reference time (t_k^f), wherein π^{-1} refers to the inverted camera projection model (π), and Z is the depth information (Z), for the respective event position $\mathbf{x}_j$.

- 5 6. Method according to one of the preceding claims, wherein a plurality of sets (S) is acquired, wherein each set particularly comprises the same predefined number of events (e), particularly wherein resulting associated time intervals are delimited by the respective earliest and latest events comprised in the respective set, particularly wherein the plurality of sets (S) is temporally ordered
10 such that two subsequent sets are temporally adjacent or overlapping by a predefined number of events, wherein for each set the corresponding motion-corrected image is generated, such that a temporal series of the motion-corrected images is obtained.
- 15 7. Method according to claim 6, wherein on a series of keyframe images that consists of a plurality of motion-corrected images from the series of motion-corrected images, a simultaneous localization and mapping (SLAM) method and/or a visual inertial odometry method is performed.
- 20 8. Method according to claim 7, wherein the SLAM method is a visual-inertial SLAM method, wherein the visual inertial odometry method and/or the visual-inertial SLAM method employ the keyframe images and the IMU data to obtain a three-dimensional map comprising landmark positions and an estimated keyframe image camera pose for each keyframe image, particularly wherein the landmark positions and the estimated keyframe image camera pose are estimated with respect to the world reference frame.
- 25 9. Method according to claim 8, wherein the visual-inertial SLAM method and/or the VIO method comprises a non-linear optimization method, wherein the non-linear optimization method minimizes a joint cost function for all keyframe images particularly by varying the estimated keyframe image camera pose for each keyframe image and the three-dimensional landmark positions.

10. Method according to claim 9, wherein the joint cost function comprises a sum of an inertial error term for each keyframe image and a re-projection error term for each landmark and for each keyframe image, wherein the re-projection error term and the inertial error term are minimized, wherein the cost function is particularly given by:

$$JC := \sum_{m=1}^M \sum_{n \in l(m)} \boldsymbol{\varepsilon}_r^{n,mT} \boldsymbol{W}_r^{n,m} \boldsymbol{\varepsilon}_r^{n,m} + \sum_{m=1}^{M-1} \boldsymbol{\varepsilon}_s^m \boldsymbol{W}_s^m \boldsymbol{\varepsilon}_s^m$$

wherein JC is the joint cost function, M is the number of keyframe images, $l(m)$ are the landmarks comprised in the respective m^{th} keyframe image, $\boldsymbol{\varepsilon}_r^{n,m}$ is a re-projection error of the n^{th} landmark in the m^{th} keyframe, $\boldsymbol{W}_r^{n,m}$ is an associated weighting factor for the re-projection error, $\boldsymbol{\varepsilon}_s^m$ is an inertial error of the m^{th} keyframe image and $\boldsymbol{W}_s^m$ is the associated weighting factor of the inertial error, particularly wherein the re-projection error is given by $\boldsymbol{\varepsilon}_r^{n,m} = \boldsymbol{z}^{n,m} - \pi(T_{CW}^m \boldsymbol{l}^n)$ wherein $\boldsymbol{z}^{n,m}$ is an estimated two-dimensional landmark position of the n^{th} landmark in the m^{th} keyframe image, π is the camera projection model, and T_{CW}^m is particularly a homogenous transformation of a camera frame of reference to the world reference frame for the m^{th} keyframe image for the varied camera pose, $\boldsymbol{l}^n$ is the estimated three-dimensional landmark position, wherein T_{CW}^m and $\boldsymbol{l}^n$ are varied, wherein the inertial error $\boldsymbol{\varepsilon}_s^m$ is estimated from the difference between the estimated camera pose from the IMU data for the m^{th} keyframe image and the varied camera pose T^m .

11. A method for visual-inertial odometry with an event camera, wherein the event camera is rigidly connected to an inertial measurement unit (IMU), wherein the event camera comprises pixels that are configured to output events in presence of brightness changes in a scene at the time they occur, wherein each event comprises the time at which it is recorded and a position of the respective pixel that detected the brightness change, the method comprising the steps of:

- Acquiring IMU data (D) with the IMU unit;
- Acquiring a plurality of sets (S) each comprising a plurality of events (e);

- For each set (S) generating an image from the events comprised in the respective set (S), wherein each image is obtained by assigning the event position (x_j) to a pixel coordinate of the image, such that a series of images is obtained;
- 5 - Perform on a series of keyframe images that consists of a plurality of images from the series of images, a visual inertial simultaneous localization and mapping (SLAM) method and/or a visual inertial odometry (VIO) method, wherein the VIO method and/or the visual-inertial SLAM method employ the keyframe images and the IMU data to
10 obtain a three-dimensional map comprising the landmark positions and an estimated keyframe image camera pose for each keyframe image

characterized in that,

the VIO method and/or the visual-inertial SLAM method comprises a non-linear optimization method, wherein the non-linear optimization method minimizes a
15 joint cost function for all keyframe images, particularly by varying the estimated keyframe image camera pose for each keyframe image and the three-dimensional landmark positions.

12. Method according to claim 11, wherein the joint cost function comprises a sum of an inertial error term for each keyframe image and a re-projection error term
20 for each keyframe image and for each landmark comprised in the respective keyframe image, wherein the re-projection error term and the inertial error term are minimized.

13. Method according to claim 12, wherein the cost function is particularly given by:

$$JC := \sum_{m=1}^M \sum_{n \in l(m)} \epsilon_r^{n,mT} W_r^{n,m} \epsilon_r^{n,m} + \sum_{m=1}^{M-1} \epsilon_s^{mT} W_s^m \epsilon_s^m$$

25 wherein JC is the joint cost function, M is the number of keyframe images, $l(m)$ are the landmarks comprised in the respective m^{th} keyframe image, $\epsilon_r^{n,m}$ is a re-projection error of the n^{th} landmark in the m^{th} keyframe, $W_r^{n,m}$ is an associated weighting factor for the re-projection error, ϵ_s^m is an inertial error of

the m^{th} keyframe image and W_S^m is the associated weighting factor of the inertial error.

14. Method according to claim 12 or 13, wherein the re-projection error is given by $\epsilon_r^{n,m} = \mathbf{z}^{n,m} - \pi(T_{CW}^m \mathbf{l}^n)$ wherein $\mathbf{z}^{n,m}$ is an estimated two-dimensional landmark position of the n^{th} landmark in the m^{th} keyframe image, π a camera projection model, and T_{CW}^m is particularly a homogenous transformation of a camera frame of reference to a world reference frame for the m^{th} keyframe image for the varied camera pose, $\mathbf{l}^n$ is the estimated three-dimensional landmark position, wherein T_{CW}^m and $\mathbf{l}^n$ are varied, wherein the inertial error ϵ_S^m is estimated from the difference between the estimated camera pose from the IMU data for the m^{th} keyframe image and the varied camera pose T^m .
15. Computer program comprising computer program code, wherein the computer program code is configured to execute the method according to one of the claims 1 to 10 and/or to one of the claims 11 to 14, when the computer program is executed on a computer.

Fig. 1




Fig. 2




Fig. 3


A)



B)



C)



INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2018/073641

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06T7/20 G06T7/246 G06T7/593
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06T

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	Alex Zihao Zhu ET AL: "Event-based Visual Inertial Odometry", CVPR 2017 Open Access, 26 July 2017 (2017-07-26), XP055417091, Retrieved from the Internet: URL: http://openaccess.thecvf.com/content_cvpr_2017/papers/Zhu_Event-Based_Visual_Inertial_CVPR_2017_paper.pdf [retrieved on 2017-10-19]	1-8,15
Y	cited in the application abstract Sections 1-6 ----- -/-	9-15



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

12 November 2018

Date of mailing of the international search report

27/11/2018

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Tessens, Linda

INTERNATIONAL SEARCH REPORT

International application No

PCT/EP2018/073641

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	STEFAN LEUTENEGGER ET AL: "Keyframe-based visual-inertial odometry using nonlinear optimization", INTERNATIONAL JOURNAL OF ROBOTICS RESEARCH., vol. 34, no. 3, 15 December 2014 (2014-12-15), pages 314-334, XP055417249, US ISSN: 0278-3649, DOI: 10.1177/0278364914554813 cited in the application	9-15
A	abstract Sections 1-7.1.1	1-8
A	----- ELIAS MUEGGLER ET AL: "Continuous-Time Visual-Inertial Trajectory Estimation with Event Cameras", ARXIV.ORG, CORNELL UNIVERSITY LIBRARY, 201 OLIN LIBRARY CORNELL UNIVERSITY ITHACA, NY 14853, 23 February 2017 (2017-02-23), XP080748368, the whole document	1-15
A	----- US 2016/209217 A1 (BABU BENZUN WISELY [US] ET AL) 21 July 2016 (2016-07-21) the whole document	1-15
A	----- HENRI REBECQ ET AL: "EVO: A Geometric Approach to Event-Based 6-DOF Parallel Tracking and Mapping in Real Time", IEEE ROBOTICS AND AUTOMATION LETTERS, vol. 2, no. 2, 25 January 2017 (2017-01-25), pages 593-600, XP055522513, DOI: 10.1109/LRA.2016.2645143 abstract Section III -----	1-15

INTERNATIONAL SEARCH REPORT

International application No.
PCT/EP2018/073641

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

see additional sheet

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☒ As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☐ No protest accompanied the payment of additional search fees.

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

This International Searching Authority found multiple (groups of) inventions in this international application, as follows:

1. claims: 1-15

Method of generating a motion-corrected image for visual-inertial odometry (VIO) from events captured using an event camera.

Method of performing accurate visual-inertial odometry with an event camera.

1.1. claims: 1-10(completely); 15(partially)

Method of generating a motion-corrected image for visual-inertial odometry (VIO) from events captured using an event camera.

1.2. claims: 11-14(completely); 15(partially)

Method of performing accurate visual-inertial odometry with an event camera.

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2018/073641

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2016209217 A1	21-07-2016	US 2016209217 A1	21-07-2016
		WO 2016077703 A1	19-05-2016
