

Aerial Robots, Visual-Inertial Odometry of

Davide Scaramuzza and Zichao Zhang

Synonyms

Visual-inertial state estimation, inertial-aided vision-based state estimation, Visual-Inertial Simultaneous Localization and Mapping (VISLAM), camera tracking, ego-motion estimation.

Definitions

Visual-Inertial odometry (VIO) is the process of estimating the state (pose and velocity) of an agent (e.g., an aerial robot) by using only the input of one or more cameras plus one or more Inertial Measurement Units (IMUs) attached to it. VIO is the only viable alternative to GPS and lidar-based odometry to achieve accurate state estimation. Since both cameras and IMUs are very cheap, these sensor types are ubiquitous in all today's aerial robots.

Overview

Cameras and IMUs are complementary sensor types. A camera accumulates the photons during the exposure time to get a 2D image. Therefore they are precise during slow motion and provide rich information, which is useful for other per-

Davide Scaramuzza and Zichao Zhang
Dep. of Neuroinformatics, ETH Zurich and University of Zurich
Dep. of Informatics, University of Zurich,
Zurich, Switzerland
e-mail: sdavide@ifi.uzh.ch
e-mail: zzhang@ifi.uzh.ch

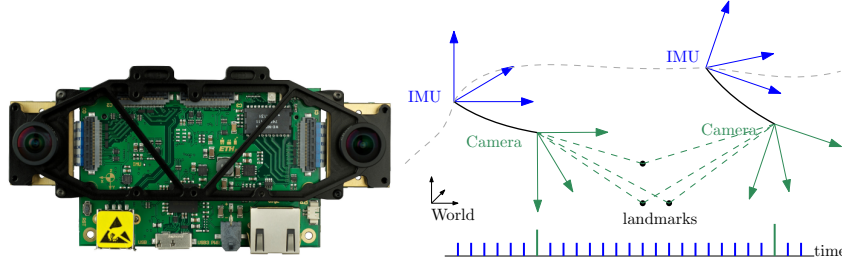


Fig. 1 Illustration of visual-inertial sensors and measurements. Left: Stereo VI-sensor from (Nikolic et al 2014). Right: Camera (green) and IMU (blue) measurements along a trajectory.

ception tasks, such as place recognition. However, they have limited output rate ($\sim 100\text{Hz}$), suffer from scale ambiguity in a monocular setup, and are not robust to scenes characterized by low texture, high speed motions (due to motion blur) and High Dynamic Range (HDR) (which may cause over- or under-exposure of the image). By contrast, an IMU is a proprioceptive sensor measuring the angular velocity and the external acceleration acting upon it. An IMU is scene-independent, which renders it unaffected by the aforementioned difficulties for cameras. Thus, it is the ideal complement to cameras to achieve robustness in low texture, high speed, and HDR scenarios. Additionally, an IMU has high output rate ($\sim 1,000\text{Hz}$). However, it suffers from poor signal-noise ratio at low accelerations and low angular velocities. Due to the presence of sensor biases, the motion estimated from an IMU alone tends to accumulate drift quickly. Therefore, a combination of both cameras and IMUs can provide accurate and robust state estimation in different situations.

The typical VIO configuration is illustrated in Fig. 1. The camera(s) and IMU(s) are rigidly attached, and the sensor suite outputs visual and inertial measurements at different rates. In VIO, the environment is represented as a set of 3D landmarks ${}^w\mathbf{p}$ that are projected by the camera to 2D image coordinates \mathbf{u} :

$$\mathbf{u} = \text{project}(\mathbf{T}_{CW} \cdot {}^w\mathbf{p}). \quad (1)$$

The IMU measures the angular velocity $\boldsymbol{\omega}$ and the external acceleration \mathbf{a} :

$$\boldsymbol{\omega} = {}_I\boldsymbol{\omega} + \mathbf{b}_g + \mathbf{n}_g, \quad \mathbf{a} = \mathbf{R}_{IW}({}^w\mathbf{a} - {}^w\mathbf{g}) + \mathbf{b}_a + \mathbf{n}_a, \quad (2)$$

where ${}_I\boldsymbol{\omega}$ is the angular velocity of the IMU expressed in the IMU frame, ${}^w\mathbf{a}$ the acceleration of the IMU in the world frame, and ${}^w\mathbf{g}$ the gravity in the world frame. \mathbf{b} and \mathbf{n} are the biases and additive noises respectively (see (Furgale et al 2013) for details).

VIO is the process of estimating the state of the sensor suite using the camera and IMU measurements (1) and (2). Typically, the quantities to estimate are N states at different times $\{t_i\}_{i=1}^N$

$$\mathbf{X}_i = [\mathbf{T}_{WI}^i, \mathbf{v}_{WI}^i, \mathbf{b}_a^i, \mathbf{b}_g^i], \quad i = 1, 2, 3, \dots, N \quad (3)$$

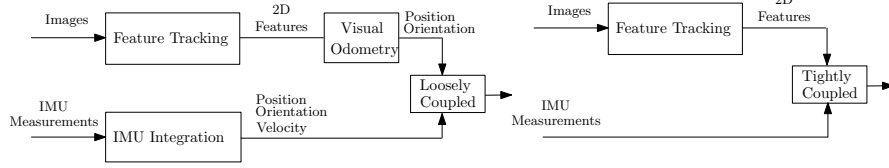


Fig. 2 Comparison of loosely (left) and tightly coupled (right) paradigms for VIO.

where \mathbf{T}_{WI}^i is the 6-DoF pose of the IMU, \mathbf{v}_{WI}^i is the velocity of the IMU, \mathbf{b}_a^i and \mathbf{b}_g^i are the biases of the accelerometer and gyroscope respectively. In contrast to visual-only odometry, the velocity and the biases are essential to utilize the IMU measurements and have to be estimated in addition to the 6-DoF pose. Specifically, biases are necessary for computing the actual sensor angular velocity and acceleration from the raw measurements (2), and velocity is needed for integrating acceleration to get position.

VIO can utilize multiple cameras and IMUs; however, the minimum number of cameras and IMUs that is sufficient to perform VIO is one. Indeed, a single moving camera allows us to measure the geometry of the 3D scene and the camera motion up to an unknown metric scale: the projection function in (1) satisfies $project(\mathbf{p}) = project(s \cdot \mathbf{p})$ for an arbitrary scalar s and an arbitrary point \mathbf{p} ; a single IMU, instead, renders metric scale and gravity observable (due to the presence of gravity in (2)) (Martinelli 2013).

Depending on the specific information used to fuse visual and inertial measurements, VIO approaches can be categorized into two paradigms: *loosely coupled* and *tightly coupled* (Corke et al 2007). Conceptually, loosely coupled methods process visual and inertial measurements separately by computing two independent motion estimates that are then fused to get the final output. By contrast, tightly coupled methods compute the final output directly from the raw camera and IMU measurements, e.g., the tracked 2D features, angular velocities, and linear accelerations (2). The difference between these two approaches is conceptually illustrated in Fig. 2.

Tightly coupled approaches are more accurate than the loosely coupled ones. First, using IMU integration to predict the 2D feature locations in the next frame can be used to facilitate feature tracking. Second, loosely-coupled approaches do not consider the visual and inertial information coupling, making them incapable of correcting drift in the vision-only estimator.

Key Research Findings

The Three Major VIO Paradigms

Existing VIO approaches can be categorized by the number of camera-poses involved in the estimation, which is highly correlated with the computational demand

and accuracy. Full smoothers (or batch nonlinear least-squares algorithms) estimate the complete history of states, fixed-lag smoothers (or sliding window estimators) consider a window of the latest states, and filtering methods only estimate the latest state. Full smoothers, by keeping the whole pose history, allow re-linearization when the estimate is updated. By contrast, sliding window estimators and filters marginalize older states (which locks the linearization error permanently) and are therefore less accurate but more efficient. Early research on VIO focuses on filters because of their efficiency. Recently, the research focus has shifted to fixed-lag/full smoothers due to their superior accuracy and the availability of more powerful computers.

Note that there are different criteria that can be used to characterize VIO algorithms. VIO approaches can use different representations of the uncertainty for the measurements and the Gaussian priors: the Extended Kalman Filter (EKF) represents the uncertainty using a covariance matrix; instead, information filters and smoothers resort to the information matrix (the inverse of the covariance) or the square-root of the information matrix (Kaess et al 2012; Wu et al 2015). The number of times in which the measurement model is linearized is also an important criterion. While a standard EKF (in contrast to the iterated EKF) processes a measurement only once, a smoothing approach allows linearizing multiple times.

While the terminology is vast, the underlying algorithms are tightly related. For instance, it can be shown that the iterated Extended Kalman filter equations are equivalent to the Gauss-Newton algorithm, commonly used for smoothing (Bell and Cathey 1993).

Filtering

Filtering algorithms enable efficient estimation by restricting the inference process to the latest state of the system. Classic approaches estimate both the poses and landmarks, and the complexity of the filter (e.g., the Extended Kalman Filter) grows quadratically in the number of estimated landmarks. Therefore, a small number of landmarks are typically tracked to allow real-time operation (Davison et al 2007; Bloesch et al 2015; Jones and Soatto 2011).

An alternative is to adopt a structureless approach where landmark positions are marginalized out of the state vector (see, for instance, the Multi-State Constraint Kalman filter (MSCKF) (Mourikis and Roumeliotis 2007)). A drawback of the structureless filter is that the processing of landmark measurements needs to be delayed until all measurements of a landmark are obtained (Mourikis and Roumeliotis 2007). This hinders accuracy as the filter cannot use all current visual information.

There are two major error sources for filtering approaches. First, filters absorb the information of the older states into the estimation of the latest state and drop the older states permanently. Therefore, linearization error and erroneous outlier measurements (see (Tsotsos et al 2015)) are locked in the filter state. Second, linearization error renders filters inconsistent. Generally, the VIO problem has four unobservable directions: the global position and the orientation around the gravity direction (yaw) (Martinelli 2013; Kottas et al 2012). In (Kottas et al 2012) it is shown

that linearization at the wrong estimate adds spurious information in unobservable directions. To address this problem, the first-estimates jacobian approach (Huang et al 2008) is often adopted to ensure that a state is not updated with different linearization points, which is a source of inconsistency.

Fixed-lag Smoothing

Fixed-lag smoothers estimate the states that fall within a given time window, while marginalizing out older states (Mourikis and Roumeliotis 2008; Sibley et al 2010; Dong-Si and Mourikis 2011; Leutenegger et al 2015). For VIO, which is highly non-linear, fixed-lag smoothing approaches are generally more accurate than filtering, since they relinearize part of the past measurements as the estimate is updated Maybeck (1979). Moreover, these approaches are more robust to outliers by explicit outlier rejection after the optimization or using robust cost functions (Hartley and Zisserman 2003). However, since fixed-lag smoothers still resort to marginalization, they, similar to filters, suffer from inconsistency and linearization errors (Hesch et al 2014; Dong-Si and Mourikis 2011; Huang et al 2011).

Fixed-lag smoothers are more computationally expensive than filters, since multiple states instead of the latest one are estimated. In addition, the marginalization of the states outside the estimation window can lead to dense Gaussian priors, which hinders efficient matrix operations. For this reason, it has been proposed to drop certain measurements instead of marginalizing them to maintain the sparsity of the problem (Leutenegger et al 2015).

Full Smoothing

Full smoothing methods estimate the entire history of the states by solving a large nonlinear optimization problem (Jung and Taylor 2001; Sterlow and Singh 2004; Bryson et al 2009; Indelman et al 2013; Patron-Perez et al 2015). Full smoothing guarantees the highest accuracy, since it can update the linearization point of the complete state history as the estimate evolves. However, because the optimization problem is approximately cubic with respect to the dimension of the states, real-time operation quickly becomes infeasible as the trajectory and the map grow over time. Common practice (also widely used in fixed-lag smoothers) is to only keep selected keyframes (Leutenegger et al 2015; Qin et al 2017; Strasdat et al 2010; Nerurkar et al 2014) and/or run the optimization in a parallel tracking and mapping architecture (Mourikis and Roumeliotis 2008; Klein and Murray 2009). A breakthrough has been the development of incremental smoothing techniques (iSAM (Kaess et al 2008), iSAM2 (Kaess et al 2012)). They leverage the expressiveness of factor graphs to maintain sparsity and to identify and update only the typically small subset of variables affected by a new measurement. VIO using the incremental smoothing framework has been demonstrated in (Forster et al 2017a).

The different update rates of cameras and IMUs bring additional difficulties for both full-smoothing and fixed-lag smoothing approaches. Filtering methods usually use IMUs for the process model and cameras for the measurement model, and, thus, they handle the different rates of IMUs and cameras naturally. In smoothing approaches, however, it is infeasible for real-time applications to add a state at every IMU measurement, since the problem complexity grows with the dimension of the states. Therefore, the IMU measurements are typically integrated between frames to form relative motion constraints. This requires the integration to be repeated when the state estimate changes (i.e., after each optimization iteration). (Lupton and Sukkarieh 2012) show that this repeated integration can be avoided by a reparametrization of the relative motion constraints. Such reparametrization is called IMU preintegration. (Forster et al 2017a) build upon [2] and bring the theory of IMU preintegration to maturity by properly addressing the manifold structure of the rotation group $SO(3)$.

Camera-IMU Calibration

The knowledge about the spatial transformations and temporal offsets between camera(s) and IMU(s) is crucial to obtain good performance from VIO. Off-line spatial calibration of cameras and IMUs is a well-studied problem and can be solved using both filters (Kelly and Sukhatme 2011) and batch optimization (Furgale et al 2013). For on-line self calibration, state-of-the-art VIO algorithms often include the unknown spatial transformation (between cameras and IMUs) in the states and estimate it together with the motion of the sensor suite ((Leutenegger et al 2015; Li and Mourikis 2013). If the visual-inertial sensor suite is not hardware synchronized as in (Nikolic et al 2014), the temporal offset of the cameras and IMUs also needs to be estimated. While different approaches have been proposed for off-line calibration ((Kelly and Sukhatme 2014; Furgale et al 2013)), very few works have been done for on-line processing ((Li and Mourikis 2013)).

A significant contribution to the community is the open source calibration toolbox *Kalibr* (Furgale et al 2013). It uses a continuous representation of the trajectory instead of discrete states, which has, therefore, the ability to model the temporal offset between the cameras and IMUs. *Kalibr* is widely used for both spatial and temporal calibration of camera-IMU systems.

Examples of Applications

Pioneering work on VIO-based autonomous navigation of aerial robots was done in the context of the European project sFly (2009-2012) (Bloesch et al 2010; Weiss et al 2013; Lynen et al 2013; Forster et al 2013; Meier et al 2012; Scaramuzza et al 2014).

Nowadays, there are several open-source VIO software packages available that have explicitly been developed for and successfully deployed on aerial robots. Most of these pipelines are monocular, since a single camera and an IMU are the minimal sensor suite necessary for reliable state estimation; the monocular setup is a convenient choice for flying robots due to its low weight and power consumption, with respect to other sensor configurations, such as stereo or multi-camera systems. The open source monocular VIO software packages currently available are:

- MSCKF (Mourikis and Roumeliotis 2007) - The Multi-State Constraint Kalman Filter (MSCKF) forms the basis of many modern, proprietary VIO systems (such as Google ARCore, the former Google Tango), but until recently no sufficient, publicly available implementation existed. The original MSCKF algorithm in (Mourikis and Roumeliotis 2007) proposed a measurement model that expressed the geometric constraints between all of the camera poses that observed a particular image feature, without the need to maintain an estimate of the 3D feature position in the state. The extended Kalman filter backend in (Zhu et al 2017) implements this formulation of the MSCKF for event-based camera inputs, but has been adapted to feature tracks from standard cameras. The code is publicly available at: https://github.com/daniilidis-group/msckf_mono.
- OKVIS (Leutenegger et al 2013, 2015) - Open Keyframe-based Visual-Inertial SLAM (OKVIS) utilizes non-linear optimization on a sliding window of keyframe poses. The cost function is formulated with a combination of weighted reprojection errors for visual landmarks and weighted inertial error terms. The frontend uses a multi-scale Harris corner detector (Harris and Stephens 1988) to find features, and then computes BRISK descriptors (Leutenegger et al 2011) on them in order to perform data association between frames. Keyframes older than the sliding window are marginalized out of the states being estimated. OKVIS uses Googles Ceres solver (Agarwal et al 2010) to perform non-linear optimization. It should be noted that OKVIS is not optimized for monocular VIO, and in (Leutenegger et al 2015) it shows superior performance using a stereo configuration. The software is available in a ROS-compatible package at: https://github.com/ethz-asl/okvis_ros.
- ROVIO (Bloesch et al 2015) - Robust Visual Inertial Odometry (ROVIO) is a visual-inertial state estimator based on an extended Kalman Filter (EKF), which proposed several novelties. In addition to FAST corner features (Rosten et al 2010), whose 3D positions are parameterized with robot-centric bearing vectors and distances, multi-level patches are extracted from the image stream around these features. The patch features are tracked, warped based on IMU-predicted motion, and the photometric errors are used in the update step as innovation terms. Unlike OKVIS, ROVIO was developed as a monocular VIO pipeline. The pipeline is available as an opensource software package at: <https://github.com/ethz-asl/rovio>.
- VINS-Mono - VINS-Mono (Qin et al 2017) is a non-linear optimization-based sliding window estimator, tracking robust corner features (Shi and Tomasi 1994), similar to OKVIS. However, VINS-Mono introduces several new features to this

class of estimation framework. The authors propose a loosely-coupled sensor fusion initialization procedure to bootstrap the estimator from arbitrary initial states. IMU measurements are pre-integrated before being used in the optimization, and a tightly-coupled procedure for relocalization is proposed. VINS-Mono additionally features modules to perform 4DoF pose graph optimization and loop closure. The software is available in both a ROS compatible PC version and an iOS implementation for state estimation on mobile devices. VINS-Mono is available at: <https://github.com/HKUST-Aerial-Robotics/VINS-Mono>.

- SVO+MSF - Multi-Sensor Fusion (MSF) (Lynen et al 2013) is a general EKF framework for fusing data from different sensors in a state estimate. Semi-Direct Visual Odometry (SVO) (Forster et al 2014, 2017b) is a computationally lightweight visual odometry algorithm that aligns images by tracking FAST corner features and edgelets and minimizing the photometric error of patches around them. This sparse alignment is then jointly optimized with the scene structure by minimizing the reprojection error of the features in a nonlinear least-squares optimization. The pose estimated from the vision-only SVO is provided to MSF as the output of a generic pose sensor, where it is then fused in a loosely-coupled manner with the IMU data, as proposed in (Faessler et al 2016). Both MSF and SVO are publicly available, and communicate through a ROS interface. SVO 2.0 is available at: https://github.com/uzh-rpg/rpg_svo_example. MSF is available at: https://github.com/ethz-asl/ethzasl_msf.
- SVO+GTSAM (Forster et al 2017a) - The same visual odometry frontend as in the SVO+MSF system has also been paired with a full-smoothing backend performing online factor graph optimization using iSAM2 (Kaess et al 2012). In (Forster et al 2017a), the authors present results using this integrated system and propose the use of pre-integrated IMU factors in the pose graph optimization. Both components of this approach, SVO and the GTSAM 4.0 optimization toolbox (Dellaert 2012), are publicly available. SVO 2.0 is available at: https://github.com/uzh-rpg/rpg_svo_example. GTSAM is available at: <https://bitbucket.org/gtborg/gtsam/>

A benchmark comparison of all the aforementioned open-source monocular VIO pipelines on common flying robot hardware (Odroid, Up Board, and Intel NUC) has been recently published (Delmerico and Scaramuzza 2018). The evaluation considers the pose estimation accuracy, per-frame processing time, and CPU and memory load while processing the EuRoC Micro Aerial Vehicle datasets (Burri et al 2016), which contain several 6DoF trajectories typical of flying robots. Note that quantitatively evaluating the accuracy is a non-trivial task due to the unobservable DoFs in VIO systems, for which a tutorial can be found in Zhang and Scaramuzza (2018).

Very recently, a VIO pipeline combining an event-camera, a standard camera, and an IMU (called UltimateSLAM) has been published (Rosinol Vidal et al 2018). It is shown on public datasets that this hybrid pipeline leads to an accuracy improvement of 85% over standard, frame-based VIO systems. Furthermore, it is shown that it can be used for autonomous quadrotor flight in scenarios inaccessible with traditional VIO, such as high-speed motion, low-light environments and high dynamic range scenes.

Future Directions for Research

New research directions are the integration of complementary sensors, such as event cameras, and new algorithmic tools, such as deep learning.

Contrarily to standard cameras, which send entire images at fixed frame rates, event cameras, such as the dynamic vision sensor (DVS) (Lichtsteiner et al 2008) or the asynchronous time-based image sensor (ATIS) (Posch et al 2011), only send the local pixel-level brightness changes caused by movement in a scene at the time they occur. Event cameras have four key advantages compared to standard cameras: a very low temporal latency (microseconds), a very high output rate (up to 1 MHz vs 100Hz of standard cameras), a very high dynamic range (up to 140 dB vs 60 dB of standard cameras), and a very low power consumption (10mW vs 1W of standard cameras). These properties enable the design of a new class of VIO and VISLAM algorithms that can operate in scenes characterized by high-speed motion (Gallego et al 2017; Rebecq et al 2017a; Rosinol Vidal et al 2018) and high-dynamic range (Kim et al 2014, 2016; Rebecq et al 2017b,a; Rosinol Vidal et al 2018), where standard cameras fail. However, since the output is composed of a sequence of asynchronous events, traditional frame-based computer-vision algorithms are not directly applicable, so that novel algorithms must be developed to deal with these cameras.

A robust VIO architecture should not solely exploit geometry and sensor measurement models but should also be able to exploit semantic/contextual information about the environment and application-specific priors about the motion dynamics. In this respect, the recent development of deep visual odometry (Costante et al 2016; Wang et al 2017; Zhou et al 2017) has shown promising initial results, especially in addressing open challenges with standard cameras, such as dealing with the aperture problem, motion blur, defocus, and low visibility scenarios.

Cross-References

Visual Simultaneous Localization and Mapping; Visual Odometry.

References

- Agarwal A, Mierle K, Others (2010) Ceres solver. <http://ceres-solver.org>
- Bell BM, Cathey FW (1993) The iterated kalman filter update as a gauss-newton method. *IEEE Trans Autom Control* 38(2):294–297, DOI 10.1109/9.250476
- Bloesch M, Weiss S, Scaramuzza D, Siegwart R (2010) Vision based MAV navigation in unknown and unstructured environments. In: *IEEE Int. Conf. Robot. Autom. (ICRA)*
- Bloesch M, Omari S, Hutter M, Siegwart R (2015) Robust visual inertial odometry using a direct EKF-based approach. In: *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*

- Bryson M, Johnson-Roberson M, Sukkarieh S (2009) Airborne smoothing and mapping using vision and inertial sensors. In: IEEE Int. Conf. Robot. Autom. (ICRA), pp 3143–3148
- Burri M, Nikolic J, Gohl P, Schneider T, Rehder J, Omari S, Achtelik MW, Siegwart R (2016) The EuRoC micro aerial vehicle datasets. *Int J Robot Research* 35:1157–1163, DOI 10.1177/0278364915620033
- Corke P, Lobo J, Dias J (2007) An introduction to inertial and visual sensing. *Int J Robot Research* 26(6):519–535, DOI 10.1177/0278364907079279
- Costante G, Mancini M, Valigi P, Ciarfuglia T (2016) Exploring representation learning with cnns for frame-to-frame ego-motion estimation. In: IEEE Int. Conf. Robot. Autom. (ICRA)
- Davison AJ, Reid ID, Molton ND, Stasse O (2007) MonoSLAM: Real-time single camera SLAM. *IEEE Trans Pattern Anal Machine Intell* 29(6):1052–1067
- Dellaert F (2012) Factor graphs and GTSAM: A hands-on introduction. Tech. Rep. GT-RIM-CP&R-2012-002, Georgia Institute of Technology
- Delmerico J, Scaramuzza D (2018) A benchmark comparison of monocular visual-inertial odometry algorithms for flying robots. In: IEEE Int. Conf. Robot. Autom. (ICRA)
- Dong-Si TC, Mourikis A (2011) Motion tracking with fixed-lag smoothing: Algorithm consistency and analysis. In: IEEE Int. Conf. Robot. Autom. (ICRA)
- Faessler M, Fontana F, Forster C, Mueggler E, Pizzoli M, Scaramuzza D (2016) Autonomous, vision-based flight and live dense 3D mapping with a quadrotor MAV. *J Field Robot* 33(4):431–450, DOI 10.1002/rob.21581
- Forster C, Lynen S, Kneip L, Scaramuzza D (2013) Collaborative monocular SLAM with multiple micro aerial vehicles. In: IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS), pp 3962–3970, DOI 10.1109/IROS.2013.6696923
- Forster C, Pizzoli M, Scaramuzza D (2014) SVO: Fast semi-direct monocular visual odometry. In: IEEE Int. Conf. Robot. Autom. (ICRA), pp 15–22, DOI 10.1109/ICRA.2014.6906584
- Forster C, Carlone L, Dellaert F, Scaramuzza D (2017a) On-manifold preintegration for real-time visual-inertial odometry. *IEEE Trans Robot* 33(1):1–21, DOI 10.1109/TRO.2016.2597321
- Forster C, Zhang Z, Gassner M, Werlberger M, Scaramuzza D (2017b) SVO: Semidirect visual odometry for monocular and multicamera systems. *IEEE Trans Robot* 33(2):249–265, DOI 10.1109/TRO.2016.2623335
- Furgale P, Rehder J, Siegwart R (2013) Unified temporal and spatial calibration for multi-sensor systems. In: IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)
- Gallego G, Lund JEA, Mueggler E, Rebecq H, Delbruck T, Scaramuzza D (2017) Event-based, 6-DOF camera tracking from photometric depth maps. *IEEE Trans Pattern Anal Machine Intell* DOI 10.1109/TPAMI.2017.2658577
- Harris C, Stephens M (1988) A combined corner and edge detector. In: Proc. Fourth Alvey Vision Conf., Manchester, UK, vol 15, pp 147–151
- Hartley R, Zisserman A (2003) Multiple View Geometry in Computer Vision. Cambridge University Press, second Edition
- Hesch JA, Kottas DG, Bowman SL, Roumeliotis SI (2014) Camera-IMU-based localization: Observability analysis and consistency improvement. *Int J Robot Research* 33(1):182–201
- Huang GP, Mourikis AI, Roumeliotis SI (2008) A first-estimates jacobian EKF for improving SLAM consistency. In: Int. Symp. Experimental Robotics (ISER)
- Huang GP, Mourikis AI, Roumeliotis SI (2011) An observability-constrained sliding window filter for SLAM. In: IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS), pp 65–72
- Indelman V, Williams S, Kaess M, Dellaert F (2013) Information fusion in navigation systems via factor graph based incremental smoothing. *J Robot and Auton Syst* 61(8):721–738
- Jones ES, Soatto S (2011) Visual-inertial navigation, mapping and localization: A scalable real-time causal approach. *Int J Robot Research* 30(4)
- Jung SH, Taylor C (2001) Camera trajectory estimation using inertial sensor measurements and structure from motion results. In: IEEE Int. Conf. Comput. Vis. Pattern Recog. (CVPR)
- Kaess M, Ranganathan A, Dellaert F (2008) iSAM: Incremental smoothing and mapping. *IEEE Trans Robot* 24(6):1365–1378

- Kaess M, Johannsson H, Roberts R, Ila V, Leonard J, Dellaert F (2012) iSAM2: Incremental smoothing and mapping using the Bayes tree. *Int J Robot Research* 31:217–236
- Kelly J, Sukhatme GS (2011) Visual-inertial sensor fusion: Localization, mapping and sensor-to-sensor self-calibration. *Int J Robot Research* 30(1):56–79, DOI 10.1177/0278364910382802
- Kelly J, Sukhatme GS (2014) A General Framework for Temporal Calibration of Multiple Proprioceptive and Exteroceptive Sensors, Springer Berlin Heidelberg, Berlin, Heidelberg, pp 195–209. DOI 10.1007/978-3-642-28572-1_14, URL https://doi.org/10.1007/978-3-642-28572-1_14
- Kim H, Handa A, Benosman R, Ieng SH, Davison AJ (2014) Simultaneous mosaicing and tracking with an event camera. In: *British Machine Vis. Conf. (BMVC)*, DOI 10.5244/C.28.26
- Kim H, Leutenegger S, Davison A (2016) Real-time 3d reconstruction and 6-dof tracking with an event camera. In: *Eur. Conf. Comput. Vis. (ECCV)*
- Klein G, Murray D (2009) Parallel tracking and mapping on a camera phone. In: *IEEE ACM Int. Sym. Mixed and Augmented Reality (ISMAR)*
- Kottas DG, Hesch JA, Bowman SL, Roumeliotis SI (2012) On the consistency of vision-aided inertial navigation. In: *Int. Symp. Experimental Robotics (ISER)*
- Leutenegger S, Chli M, Siegwart R (2011) BRISK: Binary robust invariant scalable keypoints. In: *Int. Conf. Comput. Vis. (ICCV)*, pp 2548–2555, DOI 10.1109/ICCV.2011.6126542
- Leutenegger S, Furgale P, Rabaud V, Chli M, Konolige K, Siegwart R (2013) Keyframe-based visual-inertial SLAM using nonlinear optimization. In: *Robotics: Science and Systems (RSS)*
- Leutenegger S, Lynen S, Bosse M, Siegwart R, Furgale P (2015) Keyframe-based visual-inertial SLAM using nonlinear optimization. *Int J Robot Research*
- Li M, Mourikis AI (2013) 3-d motion estimation and online temporal calibration for camera-imu systems. In: *2013 IEEE International Conference on Robotics and Automation*, pp 5709–5716, DOI 10.1109/ICRA.2013.6631398
- Lichtsteiner P, Posch C, Delbruck T (2008) A 128×128 120 dB 15 μ s latency asynchronous temporal contrast vision sensor. *IEEE J Solid-State Circuits* 43(2):566–576, DOI 10.1109/JSSC.2007.914337
- Lupton T, Sukkariyeh S (2012) Visual-inertial-aided navigation for high-dynamic motion in built environments without initial conditions. *IEEE Trans Robot* 28(1):61–76
- Lynen S, Achtelik M, Weiss S, Chli M, Siegwart R (2013) A robust and modular multi-sensor fusion approach applied to MAV navigation. In: *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*
- Martinelli A (2013) Observability properties and deterministic algorithms in visual-inertial structure from motion. *Foundations and Trends in Robotics* pp 1–75
- Maybeck P (1979) *Stochastic Models, Estimation and Control*, vol 1. Academic Press, New York
- Meier L, Tanskanen P, Heng L, Lee GH, Fraundorfer F, Pollefeys M (2012) PIXHAWK: A micro aerial vehicle design for autonomous flight using onboard computer vision. *Auton Robots* 33(1–2):21–39
- Mourikis AI, Roumeliotis SI (2007) A multi-state constraint Kalman filter for vision-aided inertial navigation. In: *IEEE Int. Conf. Robot. Autom. (ICRA)*, pp 3565–3572
- Mourikis AI, Roumeliotis SI (2008) A dual-layer estimator architecture for long-term localization. In: *Proc. of the Workshop on Visual Localization for Mobile Platforms at CVPR*, Anchorage, Alaska
- Nerurkar E, Wu K, Roumeliotis S (2014) C-KLAM: Constrained keyframe-based localization and mapping. In: *IEEE Int. Conf. Robot. Autom. (ICRA)*
- Nikolic J, Rehder J, Burri M, Gohl P, Leutenegger S, Furgale P, Siegwart R (2014) A synchronized visual-inertial sensor system with FPGA pre-processing for accurate real-time SLAM. In: *IEEE Int. Conf. Robot. Autom. (ICRA)*
- Patron-Perez A, Lovegrove S, Sibley G (2015) A spline-based trajectory representation for sensor fusion and rolling shutter cameras. *Int J Comput Vis* 113(3):208–219, DOI 10.1007/s11263-015-0811-3
- Posch C, Matolin D, Wohlgenannt R (2011) A QVGA 143 dB dynamic range frame-free PWM image sensor with lossless pixel-level video compression and time-domain CDS. *IEEE J Solid-State Circuits* 46(1):259–275, DOI 10.1109/JSSC.2010.2085952

- Qin T, Li P, Shen S (2017) VINS-Mono: A robust and versatile monocular visual-inertial state estimator. arXiv e-prints URL <https://arxiv.org/abs/1708.03852>
- Rebecq H, Gallego G, Mueggler E, Scaramuzza D (2017a) EMVS: Event-based multi-view stereo—3D reconstruction with an event camera in real-time. *Int J Comput Vis* pp 1–21, DOI 10.1007/s11263-017-1050-6
- Rebecq H, Horstschäfer T, Gallego G, Scaramuzza D (2017b) EVO: A geometric approach to event-based 6-DOF parallel tracking and mapping in real-time. *IEEE Robot Autom Lett* 2:593–600, DOI 10.1109/LRA.2016.2645143
- Rosinol Vidal T, Rebecq H, Horstschaefer T, Scaramuzza D (2018) Ultimate slam? combining events, images, and imu for robust visual slam in hdr and high speed scenarios. *IEEE Robotics and Automation Letters (RA-L)*, 2018 DOI 10.1109/lra.2018.2793357
- Rosten E, Porter R, Drummond T (2010) Faster and better: A machine learning approach to corner detection. *IEEE Trans Pattern Anal Machine Intell* 32(1):105–119, DOI 10.1109/TPAMI.2008.275
- Scaramuzza D, Achtelik M, Doitsidis L, Fraundorfer F, Kosmatopoulos EB, Martinelli A, Achtelik MW, Chli M, Chatzichristofis S, Kneip L, Gurdan D, Heng L, Lee G, Lynen S, Meier L, Pollefeys M, Renzaglia A, Siegwart R, Stumpf JC, Tanskanen P, Troiani C, Weiss S (2014) Vision-controlled micro flying robots: from system design to autonomous navigation and mapping in GPS-denied environments. *IEEE Robot Autom Mag*
- Shi J, Tomasi C (1994) Good features to track. In: *IEEE Int. Conf. Comput. Vis. Pattern Recog. (CVPR)*, pp 593–600, DOI 10.1109/CVPR.1994.323794
- Sibley G, Matthies L, Sukhatme G (2010) Sliding window filter with application to planetary landing. *J Field Robot* 27(5):587–608
- Sterlow D, Singh S (2004) Motion estimation from image and inertial measurements. *Int J Robot Research*
- Strasdat H, Montiel J, Davison A (2010) Real-time monocular SLAM: Why filter? In: *IEEE Int. Conf. Robot. Autom. (ICRA)*
- Tsotsos K, Chiuso A, Soatto S (2015) Robust inference for visual-inertial sensor fusion. In: *IEEE Int. Conf. Robot. Autom. (ICRA)*
- Wang S, Clark R, Wen H, Trigoni N (2017) Deepvo: Towards end-to-end visual odometry with deep recurrent convolutional neural networks. In: *IEEE Int. Conf. Robot. Autom. (ICRA)*
- Weiss S, Achtelik MW, Lynen S, Achtelik MC, Kneip L, Chli M, Siegwart R (2013) Monocular vision for long-term micro aerial vehicle state estimation: A compendium. *J Field Robot* 30(5):803–831, DOI 10.1002/rob.21466
- Wu KJ, Ahmed AM, Georgiou GA, Roumeliotis SI (2015) A square root inverse filter for efficient vision-aided inertial navigation on mobile devices. In: *Robotics: Science and Systems (RSS)*
- Zhang Z, Scaramuzza D (2018) A tutorial on quantitative trajectory evaluation for visual(-inertial) odometry. *IEEE/RSJ Int Conf Intell Robot Syst (IROS)*
- Zhou T, Brown M, Snavely N, Lowe DG (2017) Unsupervised learning of depth and ego-motion from video. In: *IEEE Int. Conf. Comput. Vis. Pattern Recog. (CVPR)*, pp 6612–6619, DOI 10.1109/CVPR.2017.700
- Zhu AZ, Atanasov N, Daniilidis K (2017) Event-based visual inertial odometry. In: *IEEE Int. Conf. Comput. Vis. Pattern Recog. (CVPR)*, pp 5816–5824