



Nonlinear ego-motion estimation from optical flow for online control of a quadrotor UAV

Volker Grabe^{1,2}, Heinrich H. Bühlhoff^{1,3}, Davide Scaramuzza² and Paolo Robuffo Giordano⁴

Abstract

For the control of unmanned aerial vehicles (UAVs) in GPS-denied environments, cameras have been widely exploited as the main sensory modality for addressing the UAV state estimation problem. However, the use of visual information for ego-motion estimation presents several theoretical and practical difficulties, such as data association, occlusions, and lack of direct metric information when exploiting monocular cameras. In this paper, we address these issues by considering a quadrotor UAV equipped with an onboard monocular camera and an inertial measurement unit (IMU). First, we propose a robust ego-motion estimation algorithm for recovering the UAV scaled linear velocity and angular velocity from optical flow by exploiting the so-called continuous homography constraint in the presence of planar scenes. Then, we address the problem of retrieving the (unknown) metric scale by fusing the visual information with measurements from the onboard IMU. To this end, two different estimation strategies are proposed and critically compared: a first exploiting the classical extended Kalman filter (EKF) formulation, and a second one based on a novel nonlinear estimation framework. The main advantage of the latter scheme lies in the possibility of imposing a desired transient response to the estimation error when the camera moves with a constant acceleration norm with respect to the observed plane. We indeed show that, when compared against the EKF on the same trajectory and sensory data, the nonlinear scheme yields considerably superior performance in terms of convergence rate and predictability of the estimation. The paper is then concluded by an extensive experimental validation, including an onboard closed-loop control of a real quadrotor UAV meant to demonstrate the robustness of our approach in real-world conditions.

Keywords

Sensor fusion, aerial robotics, visual-based control

1. Introduction

In recent years, inspection and search and rescue tasks have become one of the most important envisaged applications for unmanned aerial vehicles (UAVs). For instance, small-size UAVs, such as quadrotors, have been used to investigate disaster sites after earthquakes, such as Christchurch (New Zealand) and Emilia Romagna (Italy), and most prominently the Fukushima Daiichi power plant in Japan. Along similar lines, some large-scale research projects have been recently funded on these and related topics (see, e.g., AIRobots, 2010; ARCAS, 2011; Sherpa, 2013). Indeed, thanks to their high agility, pervasiveness and customizability, small UAVs represent an ideal robotic platform for navigating in harsh and cluttered environments, when operating in either full autonomy or under the (partial) remote control of skilled human operators.

In all cases, a widely recognized key component for the successful deployment of such systems is a reliable state estimation module able to deal with highly unstructured and/or GPS-denied indoor environments. This typically imposes a major requirement on the system: since the target environment for the considered applications cannot be

¹Max Planck Institute for Biological Cybernetics, Tübingen, Germany

²Robotics and Perception Group of the University of Zurich, Zurich, Switzerland

³Department of Brain and Cognitive Engineering, Korea University, Seoul, Korea

⁴CNRS at Iria and Inria Rennes Bretagne Atlantique, Rennes, France

Corresponding author:

Paolo Robuffo Giordano, CNRS at Iria and Inria Rennes Bretagne Atlantique, Campus de Beaulieu, 35042 Rennes Cedex, France.

Email: prg@irisa.fr

prepared before the deployment of the vehicle, the UAV is constrained to only rely on onboard sensing and processing capabilities. These constraints have motivated, over recent years, the extensive use of onboard cameras as the main sensory modality for state estimation purposes, see Weiss et al. (2013a) and Scaramuzza et al. (2014) for a recent overview. Vision indeed provides a rich sensory feedback which could, in principle, yield a full understanding of the surrounding environment. However, as is well known, an effective use of the visual information also presents many theoretical and practical difficulties. For instance, robust data extraction and association across multiple frames is often a major issue in real-world scenarios, especially when comparing images from distant points of view. Also, when relying on monocular cameras (as in most typical settings), any position/velocity information can only be retrieved up to an arbitrary scale factor, which must then be disambiguated by fusing the visual information with independent (metric) measurements from other onboard sensors. This has motivated many recent works on data fusion exploiting the concurrent (metric) measurements from onboard accelerometers embedded in the inertia measurement units (IMUs) present on most flying robotic systems (Martinelli, 2012; Li and Mourikis, 2013; Omari and Ducard, 2013).

Existing work on *metric* camera-based state estimation mostly relies on simultaneous localization and mapping (SLAM) techniques to build and maintain a map of visual features, with the position of the vehicle in the map being estimated in parallel. Acceleration measurements are then used to reconstruct the metric scale of the underlying map (Martinelli, 2012). This has been achieved by Nützi et al. (2011) and Faessler et al. (2015) using visual-SLAM algorithms, such as PTAM (Klein and Murray, 2007) and SVO (Forster et al., 2014), respectively, and, even more remarkably, through the observation of just a single feature over time by providing a closed-form solution for the computation of the metric scale factor (Martinelli, 2012). However, all of these approaches depend on the possibility to continuously track features for an extended period of time. Therefore, the level of robustness required for the UAV control can be hard to guarantee as the visual system can be affected by, e.g. unexpected occlusions or the need of heavy computations for a reliable feature matching. This is not the case, however, when relying on motion estimation from *optical flow*, as, in this case, data extraction and association is performed on consecutive, and, thus, spatially very close-by, acquired images. Motivated by these considerations, a first contribution of this paper is then the development of an approach based on optical flow decomposition for providing a reliable and robust ego-motion estimation module for safe UAV operation in unstructured environments.

Exploiting purely optical flow, a system capable of hovering and landing on a moving platform was presented by Herissé et al. (2012). However, this work did not consider the issue of determining the unknown scene scale factor since the proposed control approach was implemented by only relying on the non-metric linear velocity directly

obtained from a monocular camera. An approach combining optical flow decomposition and sensor fusion for metric scale estimation was instead presented by Honegger et al. (2013) by developing a small sensor for velocity estimation onboard UAVs. However, the proposed sensor relied on a ground facing sonar for metric scale estimation, thus limiting the vehicle to near ground operations within the range of the sonar. Finally, by exploiting an extended Kalman filter (EKF) to fuse optical flow measurements with IMU readings, Weiss et al. (2012) demonstrated the possibility of estimating the metric scale, sensor biases, and the IMU/camera relative pose. This system was later extended by Weiss et al. (2013b). However, the system was mainly designed for serving as an initialization of the PTAM framework in near hovering mode rather than for closed-loop UAV control over an extended period of time.

It can be noted that most of the presented camera-based state estimation methods rely on the classical EKF framework to fuse together the available sensor readings (e.g. from vision and IMU) for then extracting the metric scale. However, despite being widespread, the use of a EKF-based scheme presents two (often overlooked) drawbacks: (i) it necessarily involves a linearization of the (nonlinear) system dynamics (such as when dealing with visual-inertial estimation problems); and (ii) as a consequence, it does not allow for an explicit characterization of the estimation error behavior (e.g. its convergence rate). Therefore, as an additional contribution, in this paper we propose a novel visual-inertial estimation scheme exploiting optical flow and IMU measurements based on a recently developed nonlinear observation framework for *active* structure from motion (Spica and Robuffo Giordano, 2013; Spica et al., 2014). Compared with a classical EKF, the use of this nonlinear filter yields an estimation error dynamics with a *fully characterized convergence behavior* when the camera is moving with a constant acceleration norm with respect to the observed scene. It is then possible, for instance, to *actively impose* a desired error transient response by suitably acting on the estimation gains and on the UAV motion, as well as to *predict* the convergence time of the estimation error in terms of percentages of the initial error. Finally, the reported results will also extensively show that the use of this nonlinear filter yields a substantial faster (and more controlled) error convergence compared to a classical and ‘fully informed’ EKF, thus making it a viable and robust alternative to other consolidated schemes.

1.1. Summary of contributions

Part of the work presented in this paper is based on (Grabe et al., 2012a,b, 2013a). In this paper, we nevertheless provide a more comprehensive overview with additional details and several extensions. In particular, we introduce the following novel contributions: (i) a full characterization of the estimation error transient response for the proposed nonlinear estimation scheme used to retrieve the metric scale from optical flow and IMU measurements. As explained,

this analysis also determines the conditions on estimation gains and UAV motion which then allow to impose a *desired* convergence behavior to the scale estimation error; (ii) an extensive set of new simulative and experimental results investigating various aspects of the proposed ego-motion estimation scheme, including robustness against noise, trade-off between noise sensitivity and amount of ‘control effort’ (UAV linear acceleration), as well as an experimental validation of closed-loop velocity control with all computations run onboard and in real-time on a quadrotor UAV.

The rest of the paper is then structured as follows: in Section 2 we first review the proposed ego-motion estimation algorithm from optical flow which provides an estimation of the UAV (non-metric) linear and angular velocity. Subsequently, Section 3 introduces two estimation schemes meant to recover the unknown scale factor by fusing the visual information with the IMU readings: a filter based on the standard EKF machinery and a novel deterministic non-linear observer. The two filters are then compared by highlighting, for the latter, the possibility to characterize and actively shape its error transient response. Afterwards, Section 4 reports and discusses the results of several simulations and experiments aimed at validating and comparing the two ego-motion estimation approaches, and, finally, some experiments of closed-loop control on a quadrotor UAV are presented. Section 5 then concludes the paper and discusses some open points and future directions.

2. Ego-motion estimation from optical flow

The approach adopted in this work for ego-motion estimation from the perceived optical flow is based on the decomposition of the so-called *continuous homography matrix* (Ma et al., 2004), complemented with the typical measurements obtainable from an onboard IMU. A distinguishing feature of our method with respect to most of the previous literature is the use of a *continuous approach* for motion recovery. In fact, the typical incremental ego-motion estimation algorithms (e.g. the so-called visual odometry (Scaramuzza and Fraundorfer, 2011)) assume presence of small but *finite* camera displacements over frames, and are thus based on reconstruction methods involving the *discrete* epipolar/homography constraints. However, since most cameras acquire images at high rates, we judged more appropriate to adopt a *continuous* point-of-view in order to recover, at each step, the camera instantaneous linear/angular velocity rather than a finite displacement over time.

In the following we summarize, for the reader’s convenience, the main features of the employed ego-motion estimation method whose full details can be found in Grabe et al. (2012a,b).

2.1. Review of the continuous homography constraint

Seen from a moving camera, the apparent velocity of a point $X \in \mathbb{R}^3$ still in space as a result of the camera motion is

$$\dot{X} = [\omega]_{\times} X + v \quad (1)$$

where $v \in \mathbb{R}^3$, $\omega \in \mathbb{R}^3$ are the camera linear/angular velocity (both expressed in the camera frame), and $[\omega]_{\times} \in so(3)$ is the skew-symmetric matrix associated with vector $\omega \in \mathbb{R}^3$.

Consider a set of point features located on a common plane of equation $n^T X = d$ where $n \in \mathbb{S}^2$ is the unit normal vector to the plane, and $d \in \mathbb{R}$ the orthogonal distance of the plane to the camera frame. By rearranging the plane constraint as $\frac{1}{d} n^T X = 1$, equation (1) becomes

$$\dot{X} = [\omega]_{\times} X + v \frac{1}{d} n^T X = \left([\omega]_{\times} + \frac{1}{d} v n^T \right) X = H X \quad (2)$$

Matrix $H \in \mathbb{R}^{3 \times 3}$ is commonly referred to as the *continuous homography matrix*: it encodes the camera linear/angular velocity (v , ω), as well as the scene structure (n , d).

Defining $\lambda x = X$ for a scalar depth factor λ as the image of a point X , and exploiting the fact that $\dot{X} = \dot{\lambda} x + \lambda u$ and $\dot{x} = u$, where u is the *observed* velocity of the point x on the image plane, one obtains

$$u = Hx - \frac{\dot{\lambda}}{\lambda} x \quad (3)$$

The depth factor λ in (3) can be removed by pre-multiplication of $[x]_{\times}$. This results in the so-called *continuous homography constraint* (Ma et al., 2004)

$$[x]_{\times} Hx = [x]_{\times} u \quad (4)$$

which involves the measured (x, u) and the (unknown) continuous homography matrix H .

2.2. The 4-point algorithm (Algorithm V1)

Matrix H in (4) can be recovered from a set of N measured pairs (x_i, u_i) of detected features $x_i \in \mathbb{R}^3$ on the image plane and associated feature velocities $u_i \in \mathbb{R}^3$. This can be easily obtained as follows: by stacking the elements of H into the vector $H^S = [H_{11}, H_{21}, \dots, H_{33}] \in \mathbb{R}^9$, one can rewrite (4) as

$$a_i^T H^S = [x_i]_{\times} u_i \quad (5)$$

where $a_i = x_i \otimes [x_i]_{\times} \in \mathbb{R}^{9 \times 3}$, and \otimes stands for the Kronecker product. By then defining $A = [a_1, \dots, a_N]^T \in \mathbb{R}^{3N \times 9}$ and $B = [[x_1]_{\times} u_1, \dots, [x_N]_{\times} u_N]^T \in \mathbb{R}^{3N}$, one obtains the linear system

$$A H^S = B \quad (6)$$

Assuming presence of at least $N \geq 4$ measured pairs (x_i, u_i) , system (6) can be solved in a least-squares sense as $H^S = A^\dagger B$, with A^\dagger denoting the pseudo-inverse of matrix A .

After having recovered H , using standard techniques (Ma et al., 2004) it is further possible to algebraically decompose it into the scaled linear velocity v/d , the angular

velocity ω , and the plane normal \mathbf{n} . However, it can be shown that, in general, two physically equivalent solutions are compatible with a given homography matrix \mathbf{H} .

Thanks to this decomposition, one can then obtain the quantities $(\mathbf{v}/d, \omega, \mathbf{n})$ from two consecutive visual observations and without the use of additional sensor readings or previously acquired data structures such as, e.g., a map. This first version of the ego-motion algorithm will be referred to as *V1* in all the following developments.

2.3. Exploiting the known angular velocity (Algorithm V2)

Since any typical onboard IMU can directly measure the angular velocity ω_{IMU} , we can consider $\omega = \omega_{IMU}$ as *known* from external (i.e. not vision-based) sources. Knowledge of ω can then be used to derotate the perceived optical flow field as

$$\begin{bmatrix} u'_x \\ u'_y \end{bmatrix} = \begin{bmatrix} u_x \\ u_y \end{bmatrix} - \begin{bmatrix} -x_x x_y & 1 + x_x^2 & -x_y \\ -(1 + x_y)^2 & x_x x_y & x_x \end{bmatrix} \omega \quad (7)$$

where the interaction matrix relating \mathbf{u} to (\mathbf{v}, ω) was exploited (Chaumette and Hutchinson, 2006). This derotation step then reduces matrix \mathbf{H} to

$$\mathbf{H} = \frac{1}{d} \mathbf{v} \mathbf{n}^T \quad (8)$$

Since \mathbf{n} spans \mathbf{H}^T and $\|\mathbf{n}\| = 1$, we can obtain \mathbf{n} from the singular value decomposition $\mathbf{H} = \mathbf{U} \Sigma \mathbf{V}^T$ as the first column of matrix \mathbf{V} . The inherent sign ambiguity can be resolved by enforcing $n_z > 0$. Having retrieved \mathbf{n} , we then obtain $\mathbf{v}/d = \mathbf{H} \mathbf{n}$.

This algorithm, which will be referred to as *V2* in the following, requires the observation of at least three feature pairs $(\mathbf{x}_i, \mathbf{u}_i)$ and yields a *unique* solution for \mathbf{v}/d and \mathbf{n} .

2.4. Exploiting the known angular velocity and plane orientation (Algorithm V3)

In most indoor environments and when considering UAVs with downlooking cameras, the dominant plane can be safely taken as horizontal with, thus, its normal vector \mathbf{n} parallel to the gravity vector. Therefore, one can exploit the ability of onboard IMUs to estimate (via internal filtering) the local gravity vector, thus allowing to consider $\mathbf{n} \approx \mathbf{n}_{IMU}$ as measured independently from the visual input.

Plugging (8) in (4) yields

$$[\mathbf{x}]_{\times} \frac{1}{d} \mathbf{v} = \frac{[\mathbf{x}]_{\times} \mathbf{u}}{\mathbf{n}^T \mathbf{x}} \quad (9)$$

where now $\mathbf{n}^T \mathbf{x}$ is a known quantity. Letting $\delta = ([\mathbf{x}]_{\times} \mathbf{u}) / (\mathbf{n}^T \mathbf{x}) \in \mathbb{R}^3$, one then obtains the following equation linear in \mathbf{v}/d (the only unknown left)

$$[\mathbf{x}_i]_{\times} \frac{\mathbf{v}}{d} = \delta_i, \quad i = 1 \dots N \quad (10)$$

A least-squares approximation of \mathbf{v}/d over all N tracked features can be obtained by stacking all $[\mathbf{x}_i]_{\times}$ into the matrix $\mathbf{\Gamma} = [[\mathbf{x}_1]_{\times}, \dots, [\mathbf{x}_N]_{\times}]^T \in \mathbb{R}^{3N \times 3}$ and all δ_i into the vector $\Delta = [\delta_1, \dots, \delta_N]^T \in \mathbb{R}^{3N}$ resulting in the linear system

$$\mathbf{\Gamma} \frac{\mathbf{v}}{d} = \Delta \quad (11)$$

which can be solved as $\mathbf{v}/d = \mathbf{\Gamma}^\dagger \Delta$.

Note that any two distinct feature point vectors $\mathbf{x}_i, \mathbf{x}_j$ will never be parallel due to the perspective projection of the camera. Thus, in principle, only two flow vectors are required to obtain a solution for \mathbf{v}/d . However, a more robust estimation is of course obtained by incorporating all observed flow vectors.

This third algorithm will be referred to as *V3* in the following.

2.5. Segmentation of features

In order to further improve the robustness of the velocity estimation system, we developed a thresholding method to test whether a set of features belongs to a common plane, and whether a new feature is a member of an existing plane (Grabe et al., 2012b).

This method is based on the observation that the ‘reprojection’ vector

$$\mathbf{E} = \mathbf{B} - \mathbf{A} \mathbf{H}^S = \mathbf{B} - \mathbf{A} \mathbf{A}^\dagger \mathbf{B} \in \mathbb{R}^{3N}$$

must vanish when all observed features belong to a common plane, since (6) and its variants admit an exact solution only in this case. Therefore, the quantity $\|\mathbf{E}\|/N$ (where N is the number of tracked features) can be taken as a ‘planarity measure’ to decide whether the observed features belong to a common plane. This approach can then be generalized to test newly tracked features against a previously known plane.

Use of this thresholding method can clearly improve the robustness of the velocity estimation algorithms presented in the previous section, since outliers (of the planarity assumption) can be more easily rejected from the maintained set of visual features. Therefore, we proposed in Grabe et al. (2012b) a random sample consensus (RANSAC)-based approach meant to find an initial set of features belonging to a predominant plane which is then extended as the UAV explores the scene. In parallel, the system constantly monitors the set of rejected features to detect new or more consistent planes in the environment. To maintain the quality of the feature set, the agreement of each feature with the main segmented plane is also constantly monitored, with outliers being removed as the camera traverses the environment.

Using this filtering technique, we showed in Grabe et al. (2012b) that the estimation of the scaled linear velocity \mathbf{v}/d could be improved by a factor of 25% in the presence of a dominant ground plane partially covered by additional

objects compared to an unfiltered version of the algorithm. This then validated the use of the proposed ego-motion estimation algorithms (which are based on the continuous *homography* constraint) also in non-perfectly planar scenes as those one can typically find in indoor environments.

2.6. Final discussion

We conclude with some final considerations about the choice of basing the proposed ego-motion estimation algorithm upon the *homography* constraint instead of the other well-known *epipolar* constraint often exploited in computer and robotic vision. In general, the applicability of the homography versus epipolar constraint depends on the structure of the environment: methods relying on the *epipolar constraint* are more appropriate for highly unstructured scenes where a collection of points in *general position* can always be detected.¹ On the other hand, solutions based on the *homography constraint* should be favored when dealing with approximately planar scenes. Indeed, in the ideal case of features extracted from a perfect planar scene, the epipolar constraint ‘loses rank’ and one cannot obtain any longer a unique solution (up to a scalar factor) for the ego-motion recovery (Ma et al., 2004). The homography constraint, on the other hand, allows us to correctly address the case of planar scenes (which is the case considered in this paper).

Furthermore, compared with epipolar-based algorithms, the homography constraint remains better conditioned also in case of stationary flight, e.g. with a small amount of translational motion (indeed, the epipolar constraint vanishes for a zero translation/linear velocity). Therefore, while the epipolar constraint can be more suited for persistently translating maneuvers or vehicles (e.g. fixed-wing), the homography constraint ought to be more robust for typical indoor (close to stationary) flight regimes of vertical take-off and landing vehicles such as quadrotor UAVs.

These considerations then motivated our choice of the *homography* constraint for dealing with the ego-motion recovery. Indeed, as explained, in many scenarios, such as indoor hallways or areal coverage, one can safely assume the presence of a dominant ground plane spanning most of the observed features. Coupled with the RANSAC-based outlier rejection described in the previous section, our proposed solution then proved to yield accurate results as it will be discussed in Section 4.2.

3. Online scale estimation

The ego-motion algorithm (in its three variants) presented in the previous section allows us to obtain a reliable estimation of the scaled camera linear velocity \mathbf{v}/d . In this section, we discuss two estimation schemes meant to recover the plane distance d by fusing the optical flow decomposition with the onboard accelerometer measurements.




Fig. 1. Locations of the IMU (\mathcal{I}), camera (\mathcal{C}), body (\mathcal{B}) and world frames (\mathcal{W}) relative to each other. Frames \mathcal{I} , \mathcal{B} and \mathcal{C} are assumed to be rigidly linked to each other. The world frame \mathcal{W} is oriented horizontally with its z -axis pointing down, following the NED convention.

3.1. Equations of motion

We start by deriving the equations of motion relevant to the case under consideration. In the following, we will denote with $\mathcal{B}, \mathcal{C}, \mathcal{I}$ and \mathcal{W} the body, camera, IMU and inertial world frames, respectively. The origin of frame \mathcal{B} is assumed to be located at the quadrotor barycenter, while frames \mathcal{C} and \mathcal{I} are supposed to be rigidly attached to \mathcal{B} , see Figure 1. Throughout the text, left superscripts will be exploited to indicate the frames where quantities are expressed in. The symbol ${}^{\mathcal{X}}\mathbf{R}_{\mathcal{Y}} \in SO(3)$ will be used to denote the rotation matrix from frame \mathcal{X} to frame \mathcal{Y} , and ${}^{\mathcal{Z}}\mathbf{p}_{\mathcal{X}\mathcal{Y}} \in \mathbb{R}^3$ to represent the vector from the origin of frame \mathcal{X} to the origin of frame \mathcal{Y} , expressed in frame \mathcal{Z} . We also introduce the following quantities instrumental for the next developments: $\mathbf{g} \in \mathbb{R}^3$ as the gravity vector, and ${}^{\mathcal{I}}\mathbf{f} \in \mathbb{R}^3$, ${}^{\mathcal{I}}\boldsymbol{\omega} \in \mathbb{R}^3$ as the specific acceleration and angular velocity at the origin of \mathcal{I} .

Define ${}^{\mathcal{C}}\mathbf{v} = {}^{\mathcal{C}}\mathbf{R}_{\mathcal{W}} {}^{\mathcal{W}}\dot{\mathbf{p}}_{\mathcal{WC}}$ as the camera linear velocity in camera frame, and ${}^{\mathcal{B}}\mathbf{v} = {}^{\mathcal{B}}\mathbf{R}_{\mathcal{W}} {}^{\mathcal{W}}\dot{\mathbf{p}}_{\mathcal{WB}}$ as the body linear velocity in body frame. Since ${}^{\mathcal{B}}\mathbf{R}_{\mathcal{C}}$, ${}^{\mathcal{B}}\mathbf{R}_{\mathcal{I}}$, ${}^{\mathcal{B}}\mathbf{p}_{\mathcal{BC}}$ and ${}^{\mathcal{B}}\mathbf{p}_{\mathcal{BT}}$ are assumed constant, from standard kinematics the following relationships hold:

$$\begin{aligned} {}^{\mathcal{B}}\mathbf{v} &= {}^{\mathcal{B}}\mathbf{R}_{\mathcal{C}}({}^{\mathcal{C}}\mathbf{v} + [{}^{\mathcal{C}}\boldsymbol{\omega}]_{\times} {}^{\mathcal{C}}\mathbf{p}_{\mathcal{CB}}) = {}^{\mathcal{B}}\mathbf{R}_{\mathcal{C}} {}^{\mathcal{C}}\mathbf{v} + [{}^{\mathcal{B}}\boldsymbol{\omega}]_{\times} {}^{\mathcal{B}}\mathbf{p}_{\mathcal{CB}}, \\ {}^{\mathcal{C}}\dot{\mathbf{v}} &= {}^{\mathcal{C}}\mathbf{R}_{\mathcal{I}}({}^{\mathcal{I}}\mathbf{a} + [{}^{\mathcal{I}}\boldsymbol{\omega}]_{\times} {}^{\mathcal{I}}\mathbf{p}_{\mathcal{IC}} + [{}^{\mathcal{I}}\boldsymbol{\omega}]_{\times}^2 {}^{\mathcal{I}}\mathbf{p}_{\mathcal{IC}}) - [{}^{\mathcal{C}}\boldsymbol{\omega}]_{\times} {}^{\mathcal{C}}\mathbf{v} \\ &= {}^{\mathcal{C}}\mathbf{R}_{\mathcal{I}} {}^{\mathcal{I}}\mathbf{a} + [{}^{\mathcal{C}}\dot{\boldsymbol{\omega}}]_{\times} {}^{\mathcal{C}}\mathbf{p}_{\mathcal{IC}} + [{}^{\mathcal{C}}\boldsymbol{\omega}]_{\times}^2 {}^{\mathcal{C}}\mathbf{p}_{\mathcal{IC}} - [{}^{\mathcal{C}}\boldsymbol{\omega}]_{\times} {}^{\mathcal{C}}\mathbf{v} \end{aligned} \quad (12)$$

$$= {}^{\mathcal{C}}\boldsymbol{\kappa} - [{}^{\mathcal{C}}\boldsymbol{\omega}]_{\times} {}^{\mathcal{C}}\mathbf{v} \quad (13)$$

$${}^{\mathcal{C}}\boldsymbol{\omega} = {}^{\mathcal{C}}\mathbf{R}_{\mathcal{I}} {}^{\mathcal{I}}\boldsymbol{\omega} \quad (14)$$

$${}^{\mathcal{C}}\dot{\boldsymbol{\omega}} = {}^{\mathcal{C}}\mathbf{R}_{\mathcal{I}} {}^{\mathcal{I}}\dot{\boldsymbol{\omega}} \quad (15)$$

where ${}^{\mathcal{I}}\mathbf{a} = {}^{\mathcal{I}}\mathbf{R}_{\mathcal{W}} {}^{\mathcal{W}}\ddot{\mathbf{p}}_{\mathcal{WI}}$ is the linear acceleration experienced by the IMU and

$${}^{\mathcal{C}}\boldsymbol{\kappa} = {}^{\mathcal{C}}\mathbf{R}_{\mathcal{I}} {}^{\mathcal{I}}\mathbf{a} + [{}^{\mathcal{C}}\dot{\boldsymbol{\omega}}]_{\times} {}^{\mathcal{C}}\mathbf{p}_{\mathcal{IC}} + [{}^{\mathcal{C}}\boldsymbol{\omega}]_{\times}^2 {}^{\mathcal{C}}\mathbf{p}_{\mathcal{IC}} \quad (16)$$

is the camera linear acceleration with respect to \mathcal{W} . We note that ${}^{\mathcal{I}}\mathbf{a} = {}^{\mathcal{I}}\mathbf{f} + {}^{\mathcal{I}}\mathbf{g}$ and ${}^{\mathcal{I}}\mathbf{g} = {}^{\mathcal{I}}\mathbf{R}_{\mathcal{W}}[0, 0, g]^T$ in the case of a horizontal orientation of the world frame, see Figure 1.

In the presence of a planar scene ${}^{\mathcal{C}}\mathbf{n}^T \mathbf{c} \mathbf{P} + d = 0$ one also has (see, e.g., Robuffo Giordano et al., 2008)

$${}^{\mathcal{C}}\dot{\mathbf{n}} = -[{}^{\mathcal{C}}\boldsymbol{\omega}]_{\times} {}^{\mathcal{C}}\mathbf{n} \quad (17)$$

$$\dot{d} = {}^{\mathcal{C}}\mathbf{v}^T {}^{\mathcal{C}}\mathbf{n} \quad (18)$$

Finally, according to this notation, the decomposition of the optical flow summarized in the previous section allows us to directly measure the scaled linear velocity ${}^{\mathcal{C}}\tilde{\mathbf{v}} = {}^{\mathcal{C}}\mathbf{v}/d$. The estimation schemes presented in the following are then meant to recover the (unmeasurable) value of the plane distance d and the metric linear velocity vector ${}^{\mathcal{C}}\mathbf{v}$.

3.2. Scale ambiguity

As in any monocular structure from motion (SfM) problem, the inherent *scale ambiguity* of the observed scene can be resolved only if the camera motion is sufficiently exciting with respect to the estimation task. For instance, the depth of a point feature can be recovered only if the camera travels with a non-zero linear velocity not aligned with the projection ray of the observed point, and other constraints exist for different geometrical primitives (Spica et al., 2014).

In the case under consideration, a camera traveling with zero linear acceleration (${}^{\mathcal{C}}\boldsymbol{\kappa} = \mathbf{0}$) does not allow us to infer the plane distance d and the linear velocity ${}^{\mathcal{C}}\mathbf{v}$ from the measured ${}^{\mathcal{C}}\mathbf{v}/d$: an accelerated motion with respect to the observed plane is necessary in order to recover the correct pair $({}^{\mathcal{C}}\mathbf{v}, d)$. This (known) requirement can be briefly justified as follows: by setting ${}^{\mathcal{C}}\boldsymbol{\kappa} = \mathbf{0}$, the dynamics (13)–(18) reduce to

$$\begin{cases} {}^{\mathcal{C}}\dot{\mathbf{v}} = -[{}^{\mathcal{C}}\boldsymbol{\omega}]_{\times} {}^{\mathcal{C}}\mathbf{v} \\ \dot{d} = {}^{\mathcal{C}}\mathbf{v}^T {}^{\mathcal{C}}\mathbf{n} \end{cases} \quad (19)$$

that is, an expression homogeneous in the pair $({}^{\mathcal{C}}\mathbf{v}, d)$. Therefore, given a solution $({}^{\mathcal{C}}\mathbf{v}^*(t), d^*(t))$ of system (19), any other pair $(\lambda {}^{\mathcal{C}}\mathbf{v}^*(t), \lambda d^*(t))$ is also a valid solution, with $\lambda \in \mathbb{R}$ being an arbitrary scale factor. The measured output ${}^{\mathcal{C}}\mathbf{v}/d$, however, does *not* depend on the scale factor λ : hence, whatever the adopted filter, a non-accelerating camera cannot disambiguate the correct pair $({}^{\mathcal{C}}\mathbf{v}, d)$ from the perceived optical flow as any other pair $(\lambda {}^{\mathcal{C}}\mathbf{v}(t), \lambda d(t))$ would be equally consistent with the dynamics (19) and the measured ${}^{\mathcal{C}}\mathbf{v}/d$.

On the other hand, the presence of a non-zero (and known) linear acceleration ${}^{\mathcal{C}}\boldsymbol{\kappa} \neq \mathbf{0}$ breaks the homogeneity of (19) with respect to $({}^{\mathcal{C}}\mathbf{v}, d)$ and makes it possible to recover the correct scale factor λ during motion. Section 3.4 will revisit these considerations in the context of persistency of excitation of the camera motion during the estimation task.

3.3. Scale estimation based on the EKF

As a first approach to estimate the distance to the planar scene d , we develop a classical EKF upon the equations of motion of the system. In particular, we adopt the discrete version of the EKF, and let index $k \in \mathbb{N}$ denote the k th iteration step. For clarity, we will append a right subscript m to identify all of those quantities that are directly available through one of the onboard sensors, e.g. specific force ${}^{\mathcal{I}}\mathbf{f}_m$ and angular velocity ${}^{\mathcal{I}}\boldsymbol{\omega}_m$ from the IMU, and the scaled linear velocity ${}^{\mathcal{C}}\tilde{\mathbf{v}}_m = ({}^{\mathcal{C}}\mathbf{v}/d)_m$ from the camera.

We define the EKF state vector \mathbf{x} to consist of the metric camera linear velocity in camera frame ${}^{\mathcal{C}}\mathbf{v}$ and the *inverse* $\rho = 1/d$ of the camera distance to the planar scene:

$$\mathbf{x} = \begin{bmatrix} {}^{\mathcal{C}}\mathbf{v} \\ \rho \end{bmatrix}, \quad {}^{\mathcal{C}}\mathbf{v} \in \mathbb{R}^3, \rho \in \mathbb{R} \quad (20)$$

The choice of taking an *inverse parameterization* $\rho = 1/d$ (in place of just d) for the scene scale is indeed known to yield better results in the context of structure estimation from motion, see, e.g., De Luca et al. (2008) and Civera et al. (2008). This inverse parameterization will also be instrumental in the design of the nonlinear observer of Section 3.4.

Since the onboard IMU directly outputs the specific acceleration ${}^{\mathcal{I}}\mathbf{f}_m$ and angular velocity ${}^{\mathcal{I}}\boldsymbol{\omega}_m$ in its own frame,² we rewrite (13) as

$$\begin{aligned} {}^{\mathcal{C}}\dot{\mathbf{v}} &= {}^{\mathcal{C}}\mathbf{R}_{\mathcal{I}}({}^{\mathcal{I}}\mathbf{a} + [{}^{\mathcal{I}}\dot{\boldsymbol{\omega}}]_{\times} {}^{\mathcal{I}}\mathbf{p}_{IC} + [{}^{\mathcal{I}}\boldsymbol{\omega}_m]_{\times}^2 {}^{\mathcal{I}}\mathbf{p}_{IC}) - [{}^{\mathcal{C}}\boldsymbol{\omega}_m]_{\times} {}^{\mathcal{C}}\mathbf{v} \\ &\approx {}^{\mathcal{C}}\mathbf{R}_{\mathcal{I}}({}^{\mathcal{I}}\mathbf{f}_m + {}^{\mathcal{I}}\mathbf{g} + [{}^{\mathcal{I}}\boldsymbol{\omega}_m]_{\times}^2 {}^{\mathcal{I}}\mathbf{p}_{IC}) - [{}^{\mathcal{C}}\boldsymbol{\omega}_m]_{\times} {}^{\mathcal{C}}\mathbf{v} \end{aligned} \quad (21)$$

As no direct measurement of ${}^{\mathcal{I}}\dot{\boldsymbol{\omega}}$ is possible on most quadrotor setups, and as ${}^{\mathcal{I}}\boldsymbol{\omega}$ is usually a noisy signal, we approximate ${}^{\mathcal{I}}\dot{\boldsymbol{\omega}} \approx \mathbf{0}$ in (21) rather than attempting to recover ${}^{\mathcal{I}}\dot{\boldsymbol{\omega}}$ via a numerical differentiation. Exploiting (18), the dynamics of ρ is instead given by

$$\dot{\rho} = -\frac{\dot{d}}{d^2} = -\rho^2 {}^{\mathcal{C}}\mathbf{v}^T {}^{\mathcal{C}}\mathbf{n} \quad (22)$$

Consequently, the following equations govern the predicted state $\bar{\mathbf{x}}[k]$ in terms of the previous estimated state $\hat{\mathbf{x}}[k-1]$:

$${}^{\mathcal{C}}\bar{\mathbf{v}}[k] = {}^{\mathcal{C}}\hat{\mathbf{v}}[k-1] + T {}^{\mathcal{C}}\dot{\mathbf{v}}[k-1] \quad (23)$$

$$\bar{\rho}[k] = \hat{\rho}[k-1] - T \hat{\rho}^2[k-1] {}^{\mathcal{C}}\hat{\mathbf{v}}[k-1]^T {}^{\mathcal{C}}\mathbf{n}[k-1] \quad (24)$$

where T denotes the sampling time of the filter.

Although most quantities derived in the following steps are time varying, from now on, for the sake of exposition clarity, we will omit the time dependency $[k]$ wherever possible.

To compute the predicted covariance matrix of the system uncertainty $\bar{\Sigma}[k] \in \mathbb{R}^{4 \times 4}$, we first derive the Jacobian matrix $\mathbf{G}[k] \in \mathbb{R}^{4 \times 4}$

$$\mathbf{G} = \begin{bmatrix} \frac{\partial^c \bar{\mathbf{v}}[k]}{\partial^c \hat{\mathbf{v}}[k-1]} & \frac{\partial^c \bar{\mathbf{v}}[k]}{\partial \hat{\rho}[k-1]} \\ \frac{\partial \bar{\rho}[k]}{\partial^c \hat{\mathbf{v}}[k-1]} & \frac{\partial \bar{\rho}[k]}{\partial \hat{\rho}[k-1]} \end{bmatrix} \quad (25)$$

$$= \begin{bmatrix} \mathbf{I}_3 - T[\mathcal{I}\boldsymbol{\omega}_m]_{\times} & \mathbf{0}_{3 \times 1} \\ -T\hat{\rho}^2 \mathbf{c}_n^T & 1 - 2T\hat{\rho}^c \hat{\mathbf{v}}^T \mathbf{c}_n \end{bmatrix}$$

Matrix $\hat{\Sigma}[k-1]$ from the previous step is then propagated as

$$\bar{\Sigma}[k] = \mathbf{G}\hat{\Sigma}[k-1]\mathbf{G}^T + \mathbf{R}. \quad (26)$$

Here, matrix $\mathbf{R} \in \mathbb{R}^{4 \times 4}$ is obtained from

$$\mathbf{R} = \mathbf{V} \begin{bmatrix} \text{cov}(\mathcal{I}\mathbf{f}_m) & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \text{cov}(\mathcal{I}\boldsymbol{\omega}_m) \end{bmatrix} \mathbf{V}^T \quad (27)$$

where

$$\mathbf{V} = \begin{bmatrix} \frac{\partial^c \bar{\mathbf{v}}[k]}{\partial^{\mathcal{I}} \mathbf{f}_m} & \frac{\partial^c \bar{\mathbf{v}}[k]}{\partial^{\mathcal{I}} \boldsymbol{\omega}_m} \\ \frac{\partial \bar{\rho}[k]}{\partial^{\mathcal{I}} \mathbf{f}_m} & \frac{\partial \bar{\rho}[k]}{\partial^{\mathcal{I}} \boldsymbol{\omega}_m} \end{bmatrix} \quad (28)$$

$$= \begin{bmatrix} T^c \mathbf{R}_I & T(\mathbf{R}_I \mathbf{M} + [\hat{\mathbf{v}}]_{\times}^c \mathbf{R}_I) \\ \mathbf{0}_{1 \times 3} & \mathbf{0}_{1 \times 3} \end{bmatrix} \in \mathbb{R}^{4 \times 6}$$

$$\mathbf{M} = (\mathcal{I}\boldsymbol{\omega}_m^T \mathbf{p}_{IC}) \mathbf{I}_3 + \mathcal{I}\boldsymbol{\omega}_m \mathbf{p}_{IC}^T - 2\mathcal{I}\mathbf{p}_{IC} \mathcal{I}\boldsymbol{\omega}_m^T \quad (29)$$

and $\text{cov}(\mathcal{I}\mathbf{f}_m) \in \mathbb{R}^{3 \times 3}$, $\text{cov}(\mathcal{I}\boldsymbol{\omega}_m) \in \mathbb{R}^{3 \times 3}$ are the covariance matrices of the accelerometers/gyroscopes sensors in the IMU, which can be experimentally determined.

The predicted state $\bar{\mathbf{x}}$ is then updated whenever a new scaled visual velocity estimate $\mathbf{z}_m = \tilde{\mathbf{v}}_m = (\mathbf{c}_v/d)_m$ becomes available from the optical flow decomposition. Let $\bar{\mathbf{z}}$ be the predicted scaled visual velocity estimation based on the predicted state $\bar{\mathbf{x}}$

$$\bar{\mathbf{z}} = \mathbf{c}_v \bar{\rho} \quad (30)$$

The Kalman gain $\mathbf{K} \in \mathbb{R}^{4 \times 3}$ is obtained as

$$\mathbf{K} = \bar{\Sigma} \mathbf{J}^T (\mathbf{J} \bar{\Sigma} \mathbf{J}^T + \text{cov}(\mathbf{z}_m))^{-1} \quad (31)$$

where $\text{cov}(\mathbf{z}_m) \in \mathbb{R}^{3 \times 3}$ is the covariance matrix of the scaled visual velocity measurement, and the Jacobian $\mathbf{J} \in \mathbb{R}^{3 \times 4}$ is given by

$$\mathbf{J} = \begin{bmatrix} \frac{\partial \bar{\mathbf{z}}[k]}{\partial^c \bar{\mathbf{v}}[k]} & \frac{\partial \bar{\mathbf{z}}[k]}{\partial \bar{\rho}[k]} \\ \mathbf{I}_3 \bar{\rho} & \mathbf{c}_v \end{bmatrix} \quad (32)$$

Finally, the predicted state $\bar{\mathbf{x}}[k]$ is updated to the estimated state $\hat{\mathbf{x}}[k]$, together with uncertainty matrix $\hat{\Sigma}[k]$, as

$$\hat{\mathbf{x}} = \begin{bmatrix} \mathbf{c}_v \\ \hat{\rho} \end{bmatrix} = \begin{bmatrix} \mathbf{c}_v \\ \bar{\rho} \end{bmatrix} + \mathbf{K}(\mathbf{z}_m - \bar{\mathbf{z}}) \quad (33)$$

$$\hat{\Sigma} = (\mathbf{I}_4 - \mathbf{K}\mathbf{J})\bar{\Sigma} \quad (34)$$

3.3.1. Discussion. Several quantities are needed for the implementation of the proposed EKF. Apart from the estimated state $\hat{\mathbf{x}}[k]$, one needs knowledge of:

1. the constant IMU/camera rotation matrix ${}^{\mathcal{I}}\mathbf{R}_c$ and displacement vector ${}^{\mathcal{I}}\mathbf{p}_{IC}$;
2. the IMU angular velocity ${}^{\mathcal{I}}\boldsymbol{\omega}_m$;
3. the scaled camera linear velocity $\mathbf{c}_v \tilde{\mathbf{v}}_m = (\mathbf{c}_v/d)_m$;
4. the IMU linear acceleration ${}^{\mathcal{I}}\mathbf{a} = {}^{\mathcal{I}}\mathbf{f}_m + {}^{\mathcal{I}}\mathbf{g}$;
5. the plane normal \mathbf{c}_n .

The quantities in item 1 are assumed to be known from a preliminary IMU/camera calibration phase, see, e.g., Lobo and Dias (2007), while vector ${}^{\mathcal{I}}\boldsymbol{\omega}_m$ in item 2 is available directly from the IMU gyroscope readings. Similarly, vector $\mathbf{c}_v \tilde{\mathbf{v}}_m$ in item 3 is retrieved from the optical flow decomposition described in Section 2.

Measurement of the linear acceleration ${}^{\mathcal{I}}\mathbf{a}$ in item 3 requires the specific acceleration ${}^{\mathcal{I}}\mathbf{f}_m$ (directly available through the IMU accelerometer readings) and knowledge of the gravity vector ${}^{\mathcal{I}}\mathbf{g}$ in IMU frame. An estimation of this latter quantity is also provided by standard IMUs in *near-hovering conditions*, a fact largely exploited when recovering the UAV attitude from onboard sensing, see, e.g., Mahony et al. (2008). In our case, we found the employed IMU able to provide a reliable estimation of ${}^{\mathcal{I}}\mathbf{g}$ even when undergoing moderate accelerations thanks to its internal filtering scheme. An experimental validation of these observations is provided in Section 4.4.

Finally, \mathbf{c}_n can be directly recovered from the optical flow decomposition as discussed in Sections 2.2–2.3 (this step can also be complemented by the use of filtering techniques such as those discussed in Eudes et al. (2013)), or, in the case of a horizontal planar scene as often found in indoor environments, by identifying its direction with that of ${}^{\mathcal{I}}\mathbf{g}$. This latter possibility was exploited in all of the experiments reported in Section 4, which indeed involved a plane normal \mathbf{c}_n parallel to \mathbf{g} (horizontal ground plane).

3.4. Scale estimation based on a nonlinear estimation scheme

In this section, we summarize the derivations of the alternative *nonlinear observer* for retrieving the plane distance d based on a framework originally proposed in De Luca et al. (2008) and Robuffo Giordano et al. (2008) for dealing with SfM problems, and recently revisited in Spica and Robuffo Giordano (2013) and Spica et al. (2014) in the context of *nonlinear active estimation*. Compared with the previously discussed EKF, this estimation scheme does not require any linearization/approximation step of the system

dynamics. This results in an overall cleaner design which, following the analysis reported in Spica and Robuffo Giordano (2013), also allows for a *full characterization* of the estimation error transient response in the case of a camera traveling with a constant linear acceleration norm ($\|{}^c\boldsymbol{\kappa}(t)\| = \text{const}$). In particular, we will discuss how one can impose to the estimation error a transient response equivalent to that of a *reference linear second-order system with desired poles* by suitably acting on the estimation gains and on the UAV acceleration. This possibility enables the designer, for instance, to choose the needed combination of estimation gains/UAV motion yielding a desired estimation performance, as well as to predict in advance the convergence time of the filter in terms of percentages of the initial error.³ However, we also note that, as opposed to the EKF, the filter design assumes deterministic dynamics. Therefore, it does not explicitly account for the noise present in the system concerning state prediction and sensor measurements. This point will be thoroughly addressed in Section 4.

For the sake of exposition, the following developments are formulated here in continuous time. We first recall a classical result of the adaptive control literature, also known as the Persistency of Excitation Lemma (Marino and Tomei, 1995), and upon which the next developments are based.

Lemma 1 (Persistency of Excitation). *Consider the system*

$$\begin{cases} \dot{\boldsymbol{\xi}} = -\mathbf{D}\boldsymbol{\xi} + \boldsymbol{\Omega}^T(t)\mathbf{z} \\ \dot{\mathbf{z}} = -\boldsymbol{\Lambda}\boldsymbol{\Omega}(t)\mathbf{P}\boldsymbol{\xi} \end{cases} \quad (35)$$

where $\boldsymbol{\xi} \in \mathbb{R}^m$, $\mathbf{z} \in \mathbb{R}^p$, $\mathbf{D} > 0$, $\mathbf{P} = \mathbf{P}^T > 0$ such that

$$\mathbf{D}^T\mathbf{P} + \mathbf{P}\mathbf{D} = \mathbf{Q}, \quad \text{with } \mathbf{Q} > 0,$$

and $\boldsymbol{\Lambda} = \boldsymbol{\Lambda}^T > 0$. If $\|\boldsymbol{\Omega}(t)\|$, $\|\dot{\boldsymbol{\Omega}}(t)\|$ are uniformly bounded and the persistency of excitation condition is satisfied, i.e. there exist a $T > 0$ and $\gamma > 0$ such that

$$\int_t^{t+T} \boldsymbol{\Omega}(\tau)\boldsymbol{\Omega}^T(\tau) d\tau \geq \gamma\mathbf{I} > 0, \quad \text{for all } t \geq t_0 \quad (36)$$

then $(\boldsymbol{\xi}, \mathbf{z}) = (\mathbf{0}, \mathbf{0})$ is a globally exponentially stable equilibrium point.

This result can be exploited as follows: assume a vector $\mathbf{x} = [\mathbf{x}_1^T \mathbf{x}_2^T]^T \in \mathbb{R}^{m+p}$ can be split into a *measurable* component \mathbf{x}_1 and an *unmeasurable* component \mathbf{x}_2 . Defining an estimation vector $\hat{\mathbf{x}} = [\hat{\mathbf{x}}_1^T \hat{\mathbf{x}}_2^T]^T \in \mathbb{R}^{m+p}$, and the corresponding estimation error

$$\mathbf{e} = \begin{bmatrix} \boldsymbol{\xi} \\ \mathbf{z} \end{bmatrix} = \begin{bmatrix} \mathbf{x}_1 - \hat{\mathbf{x}}_1 \\ \mathbf{x}_2 - \hat{\mathbf{x}}_2 \end{bmatrix}$$

the goal is to design an update rule for $\hat{\mathbf{x}}$ such that the closed-loop error dynamics matches formulation (35). When this manipulation is possible, Lemma 1 ensures global exponential convergence of the estimation error

$\mathbf{e} = [\boldsymbol{\xi}^T \mathbf{z}^T]^T$ to $\mathbf{0}$, thus allowing us to infer the unmeasurable value of \mathbf{x}_2 from knowledge of \mathbf{x}_1 . The PE condition (36) plays the role of an observability constraint: estimation of \mathbf{x}_2 is possible iff matrix $\boldsymbol{\Omega}(t) \in \mathbb{R}^{p \times m}$ is *sufficiently exciting* over time in the sense of (36). We finally note that, as $\boldsymbol{\Omega}(t)$ is a generic (but known) *time-varying* quantity, formulation (35) is not restricted to only span the class of linear systems, but it can easily accommodate nonlinear terms as long as they are embedded in matrix $\boldsymbol{\Omega}(t)$.

We now detail how to tailor (35) to the case under consideration. We start by defining (similarly to what done in Section 3.3)

$$x_2 = \rho = \frac{1}{d} \quad (37)$$

and

$$\mathbf{x}_1 = {}^c\tilde{\mathbf{v}} = \frac{{}^c\mathbf{v}}{d} = {}^c\mathbf{v}x_2 \quad (38)$$

with, therefore, $m = 3$ and $p = 1$. Exploiting (22), the dynamics of x_2 is given by

$$\dot{x}_2 = -x_2^2 {}^c\mathbf{v}^T \mathbf{c} \mathbf{n} = -x_2 \mathbf{x}_1^T \mathbf{c} \mathbf{n} \quad (39)$$

As for the dynamics of \mathbf{x}_1 , using (13) we have

$$\begin{aligned} \dot{\mathbf{x}}_1 &= {}^c\dot{\mathbf{v}}x_2 + {}^c\mathbf{v}\dot{x}_2 = {}^c\dot{\mathbf{v}}x_2 - {}^c\mathbf{v}x_2\mathbf{x}_1^T \mathbf{c} \mathbf{n} \\ &= {}^c\dot{\mathbf{v}}x_2 - \mathbf{x}_1\mathbf{x}_1^T \mathbf{c} \mathbf{n} = {}^c\boldsymbol{\kappa}x_2 - [{}^c\boldsymbol{\omega}]_{\times} \mathbf{x}_1 - \mathbf{x}_1\mathbf{x}_1^T \mathbf{c} \mathbf{n} \end{aligned} \quad (40)$$

We then define

$$\begin{aligned} \boldsymbol{\Omega}^T(t) &= {}^c\boldsymbol{\kappa} = {}^c\mathbf{R}_I^T \mathbf{a} + [{}^c\boldsymbol{\omega}]_{\times} {}^c\mathbf{p}_{IC} + [{}^c\dot{\boldsymbol{\omega}}]_{\times} {}^c\mathbf{p}_{IC} \\ &\approx {}^c\mathbf{R}_I^T \mathbf{a} + [{}^c\boldsymbol{\omega}]_{\times} {}^c\mathbf{p}_{IC} \end{aligned} \quad (41)$$

where, analogously to the EKF case, ${}^I\dot{\boldsymbol{\omega}}$ is neglected since no direct measurement of the UAV angular acceleration is available onboard. With these settings, the update rule for the estimated state $\hat{\mathbf{x}}$ can be designed as

$$\begin{cases} \dot{\hat{\mathbf{x}}}_1 = \boldsymbol{\Omega}^T(t)\hat{\mathbf{x}}_2 - [{}^c\boldsymbol{\omega}]_{\times} \mathbf{x}_1 - \mathbf{x}_1\mathbf{x}_1^T \mathbf{c} \mathbf{n} + \mathbf{D}\boldsymbol{\xi} \\ \dot{\hat{\mathbf{x}}}_2 = -\hat{\mathbf{x}}_2\mathbf{x}_1^T \mathbf{c} \mathbf{n} + \boldsymbol{\Lambda}\boldsymbol{\Omega}(t)\boldsymbol{\xi} \end{cases} \quad (42)$$

with $\mathbf{D} > 0$ and $\boldsymbol{\Lambda} > 0$ being symmetric and positive-definite gain matrices. Note that (42) involves only measured and estimated quantities, including a feedback action on $\boldsymbol{\xi} = \mathbf{x}_1 - \hat{\mathbf{x}}_1$, i.e. the measurable component of the error vector. With this choice, the dynamics of the estimation error $\mathbf{e} = [\boldsymbol{\xi}^T \mathbf{z}^T]^T$ then becomes

$$\begin{cases} \dot{\boldsymbol{\xi}} = -\mathbf{D}\boldsymbol{\xi} + \boldsymbol{\Omega}^T(t)\mathbf{z} \\ \dot{\mathbf{z}} = -\boldsymbol{\Lambda}\boldsymbol{\Omega}(t)\boldsymbol{\xi} - \mathbf{z}\mathbf{x}_1^T \mathbf{c} \mathbf{n} \end{cases} \quad (43)$$

It is easy to verify that, by letting $\mathbf{P} = \mathbf{I}_3$, formulation (35) is almost fully recovered apart from the spurious scalar term

$$g(\mathbf{e}, t) = -z\mathbf{x}_1^T \mathbf{c}\mathbf{n} \quad (44)$$

in (43). Nevertheless, exponential convergence of the estimation error $\mathbf{e}(t)$ to $\mathbf{0}$ can still be proven by resorting to Lyapunov theory and by noting that the spurious term $g(\mathbf{e}, t)$ is a *vanishing perturbation* of an otherwise globally exponentially stable nominal system, i.e. $g(\mathbf{0}, t) = \mathbf{0}$, for all t . We refer the reader to De Luca et al. (2008), Robuffo Giordano et al. (2008), and Spica and Robuffo Giordano (2013) for additional discussion and proofs of these facts.

We note that the design of observer (42) did not require any linearization step as for the previous EKF thanks to the more general class of (nonlinear) systems spanned by formulation (35). It is also worth analyzing, in our specific case, the meaning of the PE condition (36) necessary for obtaining a converging estimation. Being $\mathbf{\Omega}^T(t) \in \mathbb{R}^3$ a vector, condition (36) requires that $\|\mathbf{\Omega}(t)\|^2$ (i.e. at least one component) does not ultimately vanish over time. On the other hand, vector $\mathbf{\Omega}^T(t)$ represents the camera linear acceleration through space with respect to the inertial world frame \mathcal{W} , see (41). Therefore, we recover the requirement of Section 3.2 stating that estimation of d is possible if and only if the camera undergoes a physical acceleration with respect to the observed plane, and, consequently, moving at constant velocity with respect to \mathcal{W} cannot allow the estimation to converge.

Finally, as done in the previous Section 3.3.1, we list the quantities necessary for implementing the proposed observer (42). In addition to the estimated state $\hat{\mathbf{x}}$, these are:

1. the constant IMU/camera rotation matrix ${}^I\mathbf{R}_C$ and displacement vector ${}^I\mathbf{p}_{IC}$;
2. the IMU angular velocity ${}^I\boldsymbol{\omega}_m$;
3. the scaled camera linear velocity $\mathbf{x}_1 = {}^c\tilde{\mathbf{v}}_m$;
4. the IMU linear acceleration ${}^I\mathbf{a} = {}^I\mathbf{f}_m + {}^I\mathbf{g}$;
5. the plane normal $\mathbf{c}\mathbf{n}$.

Thus, the same quantities discussed in Section 3.3.1 are also required for the nonlinear estimation scheme (42).

3.4.1. Shaping the estimation transient response. We now apply the theoretical analysis developed in Spica and Robuffo Giordano (2013) to the case at hand, and aimed at characterizing the transient response of the error system (43) in the unperturbed case, i.e. with $g(\mathbf{e}, t) = \mathbf{0}$.

Let $\mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^T = \mathbf{\Omega}$ be the singular value decomposition of matrix $\mathbf{\Omega}$ where $\boldsymbol{\Sigma} = [\mathbf{S} \ \mathbf{0}] \in \mathbb{R}^{p \times m}$, $\mathbf{S} = \text{diag}(\sigma_i) \in \mathbb{R}^{p \times p}$, and $0 \leq \sigma_1 \leq \dots \leq \sigma_p$ are the singular values of $\mathbf{\Omega}$. In the case under consideration ($p = 1$, $m = 3$), it is $\mathbf{U} = 1$ and $\sigma_1 = \|\mathbf{\Omega}\|$. Let also $\boldsymbol{\Lambda} = \alpha\mathbf{I}$, $\alpha > 0$ (scalar gain $\boldsymbol{\Lambda}$). By then designing the gain $\mathbf{D} \in \mathbb{R}^{3 \times 3}$ as

$$\mathbf{D} = \mathbf{V} \begin{bmatrix} D_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{D}_2 \end{bmatrix} \mathbf{V}^T \quad (45)$$

with $D_1 \in \mathbb{R} > 0$ and $\mathbf{D}_2 \in \mathbb{R}^{2 \times 2} > 0$, it is possible to show that, under the change of coordinates

$$\boldsymbol{\eta} = \frac{1}{\sqrt{\alpha}\sigma_1} z \quad (46)$$

and in the approximation $\sigma_1(t) = \|\mathbf{\Omega}(t)\| \approx \text{const}$, the estimation error $z(t)$ (when expressed in the coordinates $\boldsymbol{\eta}(t)$) obeys the following second-order linear dynamics

$$\ddot{\boldsymbol{\eta}} = -D_1 \dot{\boldsymbol{\eta}} - \alpha\sigma_1^2 \boldsymbol{\eta} \quad (47)$$

that is, a (unit-)mass-spring-damper system with stiffness $\alpha\sigma_1^2$ and damping D_1 .

The transient response of the estimation error $z(t) = x_2(t) - \hat{x}_2(t) = 1/d(t) - 1/\hat{d}(t)$ can then be imposed by properly ‘placing the poles’ of system (47). This can be achieved by the following.

1. Taking $D_1 = 2\sqrt{\alpha}\sigma_1$ so as to obtain a critically damped state evolution for (47) (real and coincident poles). Note that this choice univocally determines D_1 as a function of the system state (the current value of $\sigma_1 = \|\mathbf{\Omega}\|$).
2. *Actively* enforcing $\alpha\sigma_1^2(t) = \alpha\|\mathbf{\Omega}(t)\|^2 = \sigma_d^2 = \text{const}$ over time for some desired $\sigma_d^2 \neq 0$, that is, by flying with a given non-zero and constant norm of the camera linear acceleration $\mathbf{\Omega}$ scaled by gain α (a free parameter whose role will be better detailed in the following).

From standard linear system theory, these choices then result in the following behavior for the estimation error $z(t) = 1/d(t) - 1/\hat{d}(t)$

$$z(t) = (1 + \sigma_d(t - t_0))e^{-\sigma_d(t-t_0)}z(t_0) \quad \text{for all } t \geq t_0 \quad (48)$$

We conclude by noting that, in general, equation (48) represents a *reference evolution* for $z(t)$ since (i) in real-world conditions it may be hard to maintain *exactly* a $\|\mathbf{\Omega}(t)\| = \text{const}$ over time (condition also needed to render exactly valid the change of coordinates (46)), and (ii) this analysis omits the effect of the (vanishing) disturbance $g(\mathbf{e}, t)$ in (44). Nevertheless, Section 4 will show an excellent match between the reported results and the reference behavior (48).

Finally, we comment about the role of gain α (the only free parameter of the nonlinear observer): being $\sigma_d = \sqrt{\alpha}\|\mathbf{\Omega}\|$, the same convergence rate for $z(t)$ (dictated by the value of σ_d , see (48)) can be equivalently obtained by either accelerating faster or by increasing gain α . While increasing gain α may always appear more convenient in terms of reduced control effort, practical issues such as noise, discretization or quantization errors may impose an upper limit on the possible value of α , thus necessarily requiring a larger $\|\mathbf{\Omega}\|$ (larger acceleration norm) for

obtaining the desired convergence speed. More details about this point will be presented in the next Section 4.

4. Experimental evaluation

This section reports the results of several experiments meant to illustrate and validate the proposed framework for ego-motion estimation on simulated and real scenarios. After the description of the experimental setup in Section 4.1, we replicate the structure of the previous sections by first comparing in Section 4.2 the three different variants discussed in Section 2 for retrieving the scaled linear and angular velocity from the optical flow decomposition. Then, we proceed to present in Section 4.3 a series of experiments involving the two scale estimation techniques illustrated in Section 3. An assessment of the ability of the employed IMU in estimating the gravity direction while the quadrotor undergoes the accelerations needed to recover the unknown scale is also presented in Section 4.4. Further considerations about the possibility of predicting (and imposing) an estimation transient behavior for the non-linear observer of Section 3.4, and about its robustness against noise, are then discussed in Sections 4.5 and 4.6, respectively. The experimental evaluation is then finally concluded in Section 4.7 by reporting the results of a closed-loop velocity control on a real quadrotor UAV with all of the computations performed onboard.

4.1. Experimental setup

For our experiments, we made use of a quadrotor UAV from MikroKopter as experimental platform, and of the free *TeleKyb* software framework (Grabe et al., 2013b) as middleware and control software. An external motion tracking system (Vicon, see <http://www.vicon.com>) was employed for obtaining the ground truth data and for allowing closed-loop control of the UAV in all of the following experiments apart from the results reported in Section 4.7 where autonomous flight was instead achieved without any external aid.

Figures 1 and 2 depict the location of all relevant sensors and frames. The quadrotor was equipped with an additional 3DM-GX3-25 IMU from MicroStrain (frame \mathcal{I}) to provide the measurements of the specific acceleration ${}^{\mathcal{I}}\mathbf{f}_m$, of the angular velocity ${}^{\mathcal{I}}\boldsymbol{\omega}_m$, and of the gravity vector ${}^{\mathcal{I}}\mathbf{g}_m$ at 200 Hz. The gravitational vector was estimated internally by the IMU via fusion of the measurements from the accelerometer, rate gyroscopes, temperature sensor and a 3D magnetometer. The visual input was provided by a pre-calibrated MatrixVision mvBlueFox camera (frame \mathcal{C}) mounted downfacing on the UAV, and delivering images of size 752×480 pixel. The onboard processing was carried out on a small PC with an Intel Atom processing unit, while velocity command inputs were transmitted to the vehicle using a wireless serial interface. When processing recorded data offline on a desktop computer, the optical flow could be decomposed at 50 Hz. When, instead, relying on the onboard processing unit (as in the experiments




Fig. 2. Experimental setup with the highlighted location of IMU and camera. The x -axis of the body frame is oriented along the red metal beam.

of Section 4.7), the acquired images were processed online at a frame rate of 17 Hz. In this case, we also found that, in average, the images were delayed by 100 ms with respect to the IMU measurements. In order to compensate for this effect, we delayed the processing of the IMU readings accordingly.

In order to allow for an efficient feature tracking, the point features were initially detected using a FAST tracker (Rosten and Drummond, 2005). An efficient pyramidal implementation of the Lucas–Kanade algorithm (Lucas and Kanade, 1981) was then used to track the features over all consecutive images. For further guaranteeing a minimum number of observed features while avoiding unnecessary and costly feature re-sampling, the point features were only sampled once the set of currently tracked features did drop below a given threshold. New FAST features were then detected and added to the set until a desired maximum set size was reached.

4.2. Comparison of algorithms $V1$ – $V3$ for ego-motion estimation from optical flow

We start by presenting a comparison of the three possibilities discussed in Sections 2.2–2.4 for recovering \mathbf{v}/d and $\boldsymbol{\omega}$ during flight. For the reader's convenience, we recall that algorithm $V1$ (Section 2.2) is only based on the perceived optical flow for obtaining $(\mathbf{v}/d, \boldsymbol{\omega})$, algorithm $V2$ (Section 2.3) exploits the additional (independent) measurement of $\boldsymbol{\omega}$ from the onboard gyros, and algorithm $V3$ (Section 2.4) further exploits the independent knowledge of the scene normal \mathbf{n} taken coincident with the gravity vector.

To allow for a controlled and direct comparison of the three possibilities $V1$ – $V3$, we recorded data while flying along a circular trajectory of 2 m in diameter as shown in Figure 3. The trajectory was chosen in order to have the UAV accelerating sinusoidally along the three Cartesian directions with a period of 10 s, and a maximum speed and acceleration of 0.6 m/s and 0.296 m/s², respectively. The height varied from 0.5 to 1.5 m along the trajectory.




Fig. 3. Circular trajectory used in the experiments for comparing the three algorithms *V1*, *V2*, and *V3* of Section 2. The trajectory has a diameter of 2 m and lies on a plane at 26.5° with respect to the ground. The UAV completes one revolution every 10 s. The height varies between 0.5 and 1.5 m. A similar trajectory, but lying on a horizontal plane, has also been used for the experiments of Section 4.3

In addition, the vehicle was periodically rotated around its body z -axis in the range of $[-70^\circ \dots 70^\circ]$. The quadrotor relied on an external motion tracking system to track this precomputed trajectory. Onboard hardware was used to record vision and IMU data during flight. Afterwards, all three algorithms of Section 2.2–2.4 were run on this common dataset to allow for a direct comparison.

Figure 4 summarizes the results of this first experiment. The plots in Figure 4a–(c) show a superimposition of the three estimations of \mathbf{v} against the ground truth provided by the external tracking system. In all three cases, the linear velocity \mathbf{v} was recovered from the estimated \mathbf{v}/d scaled by the ground truth measurement of d . Figure 4d then reports the corresponding error norms.

From these results we can conclude that all three algorithms *V1*–*V3* yield a satisfactory performance. However, one can also note how the integration of the IMU measurements (i.e. algorithms *V2*–*V3*) reduces the mean error of the linear velocity estimation compared with algorithm *V1*, see Table 1. A comparison of the angular velocity $\boldsymbol{\omega}$ as retrieved from the visual system (algorithm *V1*) against the IMU readings and the ground truth of the external tracking system is given in Figures 4e–(f), with Figure 4g reporting the corresponding error norm with respect to the ground truth. We can note how the IMU provides a less noisy estimate of $\boldsymbol{\omega}$ almost free of spikes, while the camera-based estimation is less accurate during low-altitude flights.

Figure 4h indicates the origin of the irregularities in the recovered $\boldsymbol{\omega}$. Since the vehicle is moving with a constant speed in the world frame, \mathbf{v}/d increases with lower altitudes. Therefore, in this case the point features tend to vanish faster from the field of view and can only be tracked over a shorter amount of time. In particular, at 0.5 m altitude and under the given experimental conditions, motion blur starts to affect the tracking and feature re-sampling considerably. Therefore, the system can only maintain a set of 30 to

65 features instead of the, in this case, intended 50 to 100 with a negative effect on the retrieved $\boldsymbol{\omega}$. On the other hand, the estimated \mathbf{v}/d results almost insensitive (in the algorithms *V2* and *V3*) with respect to the number of tracked features as can be seen in Figure 4a–(c): here, the estimated \mathbf{v}/d in the *V2* and *V3* cases tracks very well the ground truth despite the various drops in the number of tracked features (Figure 4h). Table 1 finally provides a quantitative summary of the reported results. One can again notice the improvements obtained when fusing the decomposition of the perceived optical flow with measurements from the IMU with, in particular, algorithm *V3* performing slightly better than *V2*.

In view of these considerations, all of the following results were obtained by relying on algorithm *V2* for retrieving \mathbf{v}/d from the decomposition of the optical flow, with $\boldsymbol{\omega}$ instead provided by the onboard IMU. This choice is motivated by the fact that algorithm *V2* grants a performance comparable with algorithm *V3* but under milder assumptions on the surrounding environment (no need of a horizontal dominant plane and of an accurate sensing of the gravity direction).

4.3. Comparison of the two scale estimation schemes

In the following, the two scale estimation schemes of Sections 3.3 and 3.4 are first evaluated in simulation to compare their performance in ideal conditions. Subsequently, the results obtained on recorded sensor data from real flights are presented. We also compare our solution against a recent related work by qualitatively reproducing the experiments presented in Weiss et al. (2012).

Before starting the discussion, we note that each filter allows for two distinct possibilities for retrieving the metric linear velocity \mathbf{v} . Indeed, a first possibility is to just set $\hat{\mathbf{v}} = (\mathbf{v}/d)_m \hat{d}$ exploiting the estimated \hat{d} and the measured $(\mathbf{v}/d)_m$ from the optical flow decomposition. However, a second possibility is to exploit the *internal* estimation of \mathbf{v} maintained by both filters, that is $\hat{\mathbf{v}}$ from (33) in the EKF case, and $\hat{\mathbf{v}} = \hat{\mathbf{x}}_1/\hat{\mathbf{x}}_2$ for the nonlinear observer case (see Section 3.4). Both possibilities have their pros/cons: in general the second possibility may result in a less noisy but also potentially more delayed estimation of \mathbf{v} (because of the ‘filtering’ action of both estimation schemes). Nevertheless, for the sake of obtaining a less noisy estimate, all of the following simulation and experiment results rely on this second possibility.

4.3.1. Simulated data. In order to illustrate the convergence of the proposed estimation schemes, we generated a synthetic acceleration profile together with the corresponding (simulated) sensor readings. This resulted in a camera motion similar to the horizontal circular trajectory used for the experimental results of Section 4.3.2, with an associated constant acceleration norm $\|\boldsymbol{\Omega}\| \approx 0.296 \text{ m/s}^2$.

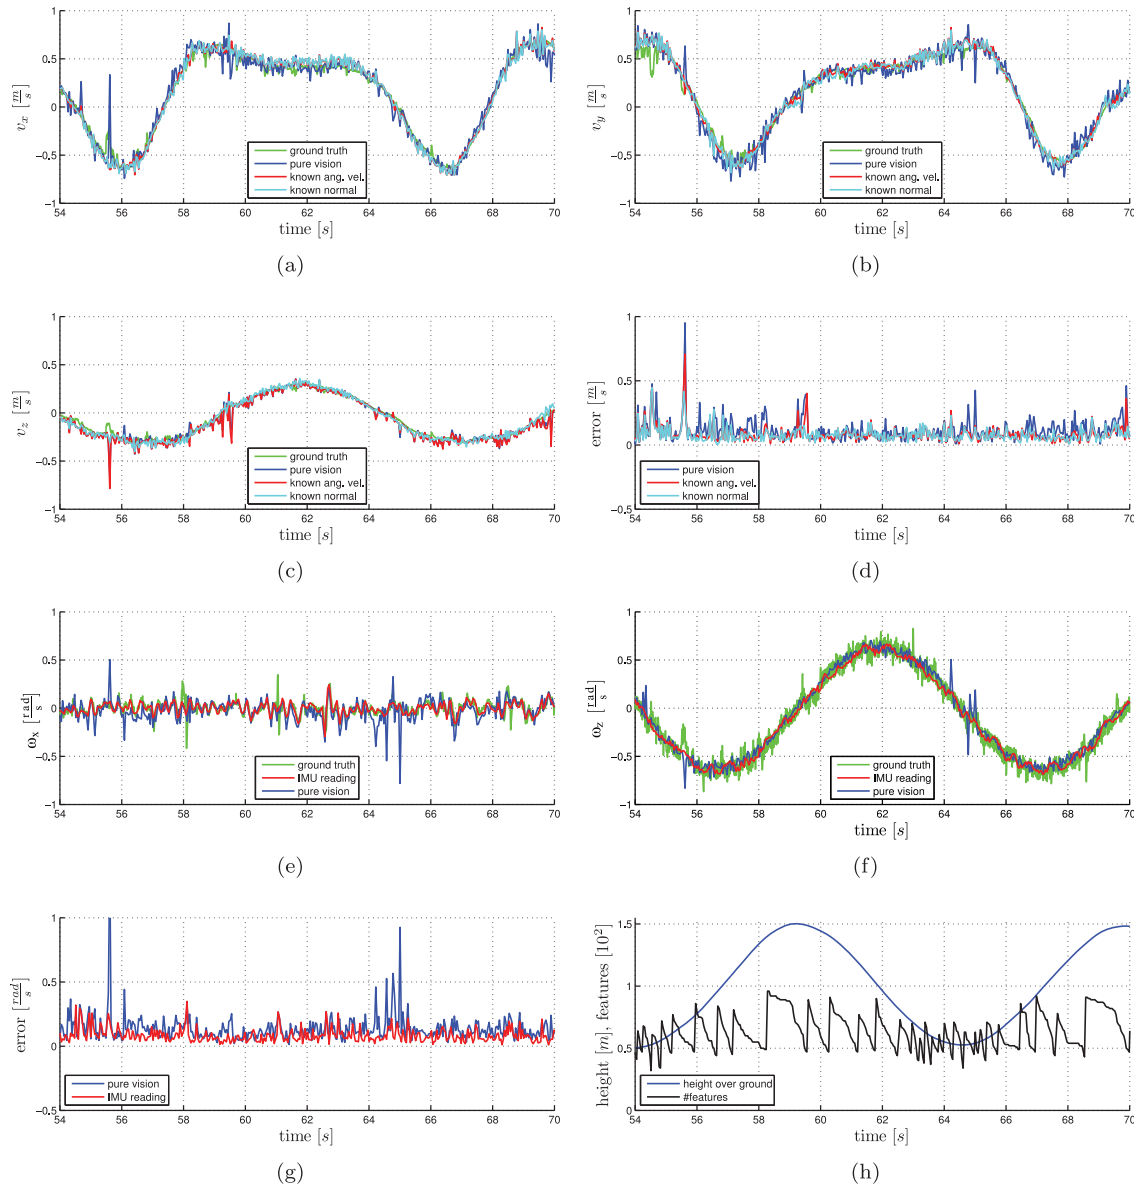


Fig. 4. Experimental results: comparison against a ground truth of the algorithms $V1$ ('pure vision'), $V2$ ('known ω '), and $V3$ ('known ω and n ') for retrieving the scaled linear velocity and angular velocity (Section 2). For the purpose of visualization only, the estimated non-metric linear velocity has been scaled using the real distance to the plane d obtained from an external optical tracking system. (a)–(c) Estimated linear velocity along the x , y and z axes. (d) Norm of the error between ground truth and the estimated linear velocity for the three cases. (e)–(f) Estimated angular velocity along the x and z axes. (g) Norm of the error between the ground truth and the estimated angular velocities. (h) Altitude of the vehicle and number of tracked features.

Table 1. Comparison of the errors in retrieving the scaled linear velocity and angular velocity for the three algorithms of Sections 2.2–2.4: the pure visual estimation ($V1$), the inclusion of the readings from the rate gyros ($V2$), and the additional integration of the measured normal vector taken coincident with the gravity direction ($V3$). In the case of algorithms $V2$ and $V3$, the angular velocity is measured from the IMU.

Employed algorithm:	$V1$ (2.2)	$V2$ (2.3)	$V3$ (2.4)
Linear velocity m/s			
Mean error	0.134	0.117	0.113
Standard deviation	0.094	0.093	0.088
Angular velocity rad/s			
Mean error	0.151	IMU	IMU
Standard deviation	0.110	IMU	IMU

All generated sensor readings were perturbed with an additive zero-mean Gaussian noise with covariance matrices taken from the real sensor characteristics: $0.00003 \mathbf{I}_3 \text{ m/s}^2$, $0.00002 \mathbf{I}_3 \text{ rad/s}$ and $0.00005 \mathbf{I}_3 \text{ 1/s}$ for $\mathcal{I}f_m$, $\mathcal{I}\omega_m$, and for the non-metric linear velocity from optical flow $(\mathcal{C}v/d)_m$, respectively. The same covariance matrices were employed in the steps (27)–(31) of the EKF. As for the nonlinear observer, its only free parameter (gain α) was taken as $\alpha = 12$: this value was chosen so as to obtain, for both filters, the same noise level (that is, variance of the estimation error) in the estimated $\hat{d}(t)$ after convergence. This was meant to ensure a ‘fair’ comparison among the EKF and the nonlinear observer in terms of estimation accuracy. The choice of $\alpha = 12$ then resulted in a value of $\sigma_d = \sqrt{\alpha} \|\Omega\| \approx 1.025$ for the ideal estimation error convergence in (48).

In order to test the robustness against an unknown scene depth, the initial estimation of the plane distance was chosen as $\hat{d}(t_0) = 5 \text{ m}$ for both filters against a real distance of $d(t_0) = 1 \text{ m}$. Furthermore, in the EKF case the initial $\mathcal{C}\hat{v}(t_0)$ was taken as $\mathcal{C}\hat{v}(t_0) = \hat{v}_m(t_0)\hat{d}(t_0)$ (i.e. exploiting the first measurement of $\hat{v}_m = (\mathcal{C}v/d)_m$ scaled by the initial estimation $\hat{d}(t_0)$). In the nonlinear observer case, we instead directly set $\hat{x}_1(t_0) = \hat{v}_m(t_0)$ (see (38)). Finally, the EKF was initialized with $\hat{\Sigma}(t_0) = \mathbf{I}_4$ for representing the initial level of uncertainty, although we will later show that the EKF performance is basically insensitive to the initial choice of $\hat{\Sigma}(t_0)$.

Figure 5a reports a comparison of the estimation performance of both the EKF and the nonlinear observer run in parallel on the same UAV trajectory and sensor readings, together with the ground truth for the plane distance d . One can then appreciate the faster convergence of the estimation error for the nonlinear observer case with respect to the ‘fully informed’ EKF (in the sense of being aware of the exact noise characteristics). This better performance can be ascribed to the lack of linearization of the system dynamics in the nonlinear observer case.

Furthermore, Figure 5b shows the behavior of the estimation error $z(t) = 1/d(t) - 1/\hat{d}(t)$ for both filters, with superimposed the ‘ideal’ transient response (48). We can then note the very good match of this latter ideal response with the behavior of the estimation error $z(t)$ for the nonlinear observer case, thus confirming the theoretical analysis of Section 3.4. The plot also shows a horizontal band at 10% of the initial error: according to the ideal response (48), the estimation error of the nonlinear observer should have dropped below 10% of the initial error $z(t_0)$ after 3.79 s. The result in Figure 5b is slightly off this prediction because of the presence of the perturbation term $g(z, t)$ in (44) (also shown in the plot). Indeed, presence of this term (neglected in the analysis) initially slows down the convergence rate of $\hat{d}(t)$ with respect to its predicted behavior. Note that, ideally, one should have had $g(z, t) = 0$ in this case, since v was always perpendicular to the plane normal n (see again (44)). However, the simulated noisy measurements resulted in the presence of a small v_z which gave rise

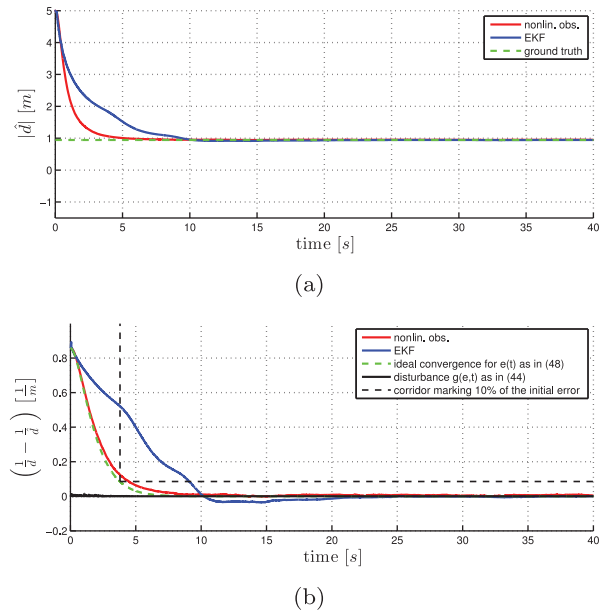


Fig. 5. Simulated results: estimated plane distance d using a simulated dataset. (a) Estimated distance $d(t)$ and (b) estimation error $z(t) = (1/d(t) - 1/\hat{d}(t))$ for the nonlinear observer and the EKF filter. In (b) the ideal convergence of the nonlinear observer as per (48) is superimposed together with the influence of the neglected disturbance $g(e, t)$ given in (44). The vertical dashed line in (b) denotes the predicted time at which the estimation error of the nonlinear observer should have dropped below the threshold of 10% of the initial error (highlighted with a horizontal dashed line).

to the initial non-null value of $g(z, t)$. Additional details about the possibility of predicting the convergence time of the nonlinear observer are also given in Section 4.5.

The linear velocities estimated by both observers are shown in comparison in Figure 6. As with the estimation of the scale factor d , the estimation for the metric velocity v converges faster in the nonlinear observer case. However, both scale estimation schemes provide very reliable velocity measurements after 20 s as shown in Figure 7 where the norm of the velocity estimation errors are reported. Finally, Figure 8 shows the behavior of the distance $\hat{d}(t)$ estimated by the EKF when initializing the state covariance $\hat{\Sigma}(t_0)$ with values between $\hat{\Sigma}(t_0) = 10^{-6} \mathbf{I}_4$ and $\hat{\Sigma}(t_0) = 10^9 \mathbf{I}_4$ against the estimation provided by the nonlinear observer. One can then verify how the EKF convergence is always slower than in the nonlinear observer case, and how any ‘reasonable’ initial $\hat{\Sigma}(t_0)$ has only a small effect on the EKF performance.

Table 2 finally summarizes the quantitative results for the comparison of both scale estimation approaches. The reported root mean square (RMS) error for both systems was computed after 20 s of flight, i.e. after both estimates have converged.

4.3.2. Recorded data. The two scale estimation approaches were also compared by processing a dataset collected

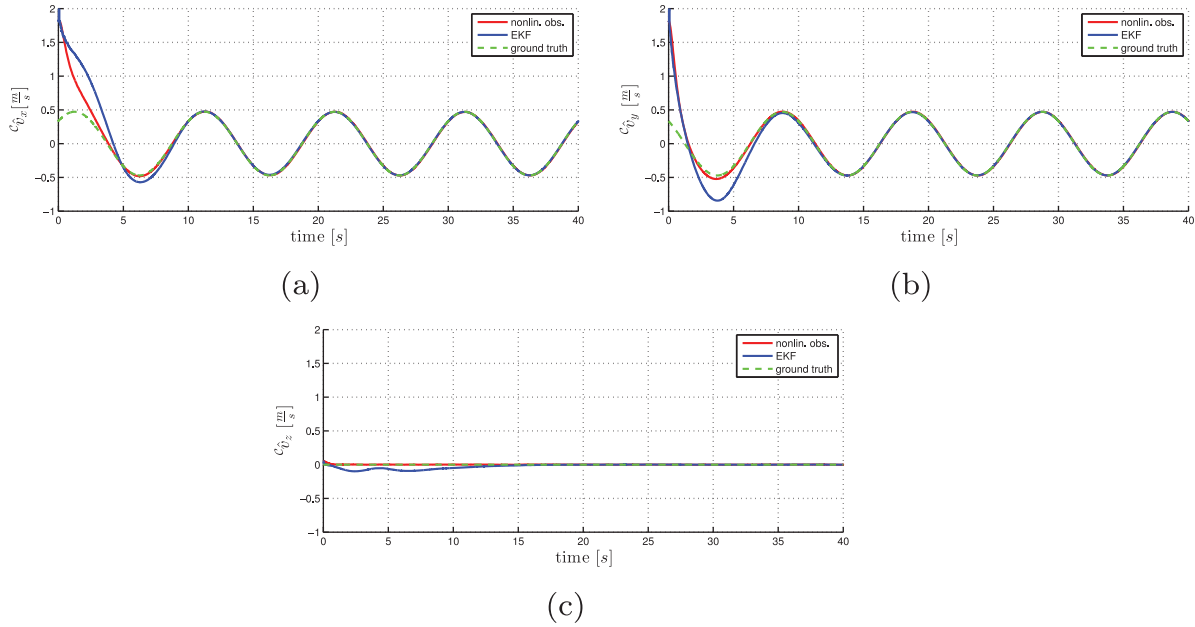


Fig. 6. Simulated results: estimated linear velocity from simulated data in the (a) x , (b) y and (c) z components. The estimation provided by the nonlinear observer converges faster, but both state estimators provide reliable results after convergence is reached.




Fig. 7. Simulated results: norm of the error for the linear velocity estimation for the nonlinear observer and EKF with respect to the ground truth.

during a real quadrotor flight. The chosen trajectory was similar to the circular one used in Section 4.2 but rotated on a horizontal plane in order to yield a constant height during motion. The chosen trajectory was again characterized by an acceleration of $\|\mathbf{\Omega}\| \approx 0.296 \text{ m/s}^2$, thus resulting in a $\sigma_d = \sqrt{\alpha}\|\mathbf{\Omega}\| \approx 1.025 \text{ m/s}^2$ for the ideal system (48). However, on the real quadrotor the actual accelerations were found slightly higher than those expected from the ideal circular trajectory due to additional constant regulations necessary to overcome small real-world disturbances. All sensor offsets were calibrated before the recording of the dataset, and the covariance matrices of ${}^I\mathbf{f}_m$, ${}^I\mathbf{\omega}_m$ and $({}^C\mathbf{v}/d)_m$ were estimated over a period of 60 s. Both filters were initialized as described in the previous Section 4.3.1, thus with an initial $\hat{d}(t_0) = 5 \text{ m}$ against a real height of $d(t_0) \approx 1 \text{ m}$. The gain for the nonlinear observer was again selected as $\alpha = 12$ to ensure a similar noise level for the scale estimation from the EKF and the nonlinear observer after convergence.




Fig. 8. Simulated results: influence of the initial choice of $\hat{\Sigma}(t_0)$ on the convergence rate of the plane distance $\hat{d}(t)$ estimated by the EKF. Note how the EKF convergence is always slower than in the nonlinear observer case (dashed line) despite the wide range of initial $\hat{\Sigma}(t_0)$. Furthermore, the EKF convergence results are almost insensitive to any ‘reasonable’ initial choice of $\hat{\Sigma}(t_0)$ (this plot should also be compared against Figure 5a).

Figure 9a shows the behavior of the estimated plane distance $\hat{d}(t)$ for both filters, with superimposed the ground truth $d(t)$. The behavior of the estimation error $z(t) = 1/d(t) - 1/\hat{d}(t)$ is shown in Figure 9b, and the behavior of $d(t) - \hat{d}(t)$ is also reported in Figure 9c for a better appreciation of the convergence behavior. Again, from these plots one can note the good match among the estimation error of the nonlinear observer against the ideal

Table 2. Comparison of the two scale estimation approaches presented in this work. The results are discussed in Section 4.3.

	EKF	Nonlinear observer
Simulated data		
RMS error after convergence (scale)	0.0075 m	0.0078 m
RMS error after convergence (velocity)	0.0071 m/s	0.0111 m/s
Recorded data		
RMS error after convergence (scale)	0.0923 m	0.0988 m
RMS error after convergence (velocity)	0.1160 m/s	0.0945 m/s
Comparison against Weiss et al. (2012)		
RMS error x -axis (velocity)	0.0101 m/s	0.0074 m/s
RMS error y -axis (velocity)	0.0141 m/s	0.0095 m/s
RMS error z -axis (velocity)	0.0107 m/s	0.0114 m/s
Improvement over Weiss et al. (2012) (average)	250%	323%

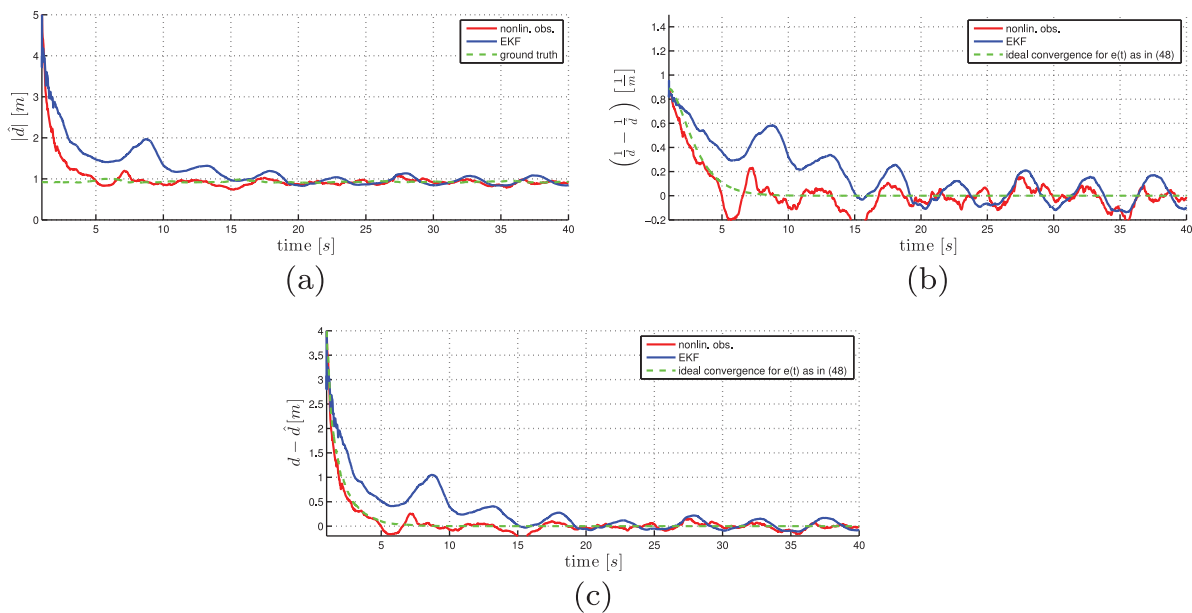


Fig. 9. Experimental results: (a) estimation of $d(t)$, (b) estimation error $1/d(t) - 1/\hat{d}(t)$ and (c) estimation error $d(t) - \hat{d}(t)$ against the ground truth exploiting a recorded dataset. The initial $\hat{d}(t_0)$ is set to 5 m against a real distance of approximately $d(t_0) = 1$ m. The ideal transient response of the nonlinear observer as per (48) is also shown in (b) and (c) with a dashed line. Note, again, the very good match between the predicted and actual behavior of the estimation error for the nonlinear observer case, as well as the faster overall convergence with respect to the EKF scheme.

response (48). On this dataset from real sensor measurements, the mismatch between the predicted and actual convergence rate of the nonlinear estimator was smaller than in the previous case of simulated data. This can be explained by the fact that the real acceleration of the vehicle was slightly higher than the acceleration expected from the commanded circular trajectory, and this in turn compensated for the initial ‘disturbing’ effects of the perturbation g (thanks to the higher level of excitation of the actual trajectory).

The estimated metric linear velocities are shown in Figure 10. The error against the ground truth (depicted in Figure 11) shows a fast convergence rate for the nonlinear observer. Nevertheless, the EKF yielded a reliable output after approximately 20 s as well. All quantitative results can be compared in Table 2.

As an additional validation of our approach, we also tested both algorithms on an inclined circular trajectory similar to that depicted in Figure 3 in order to deal with the case of a *time-varying* plane distance $d(t)$. The plots shown in Figures 12–13 again fully confirm that both the scene depth and the vertical component of the linear velocity can be reliably estimated despite the more challenging UAV motion (time-varying $d(t)$).

Motivated by these good results in comparison with the EKF, all of the following experiments only involve the nonlinear observer scheme.

4.3.3. Comparison to previous work. For a direct comparison with a state-of-the-art scale estimation approach, we applied the nonlinear observer scheme under experimental conditions similar to those used by Weiss et al. (2012).

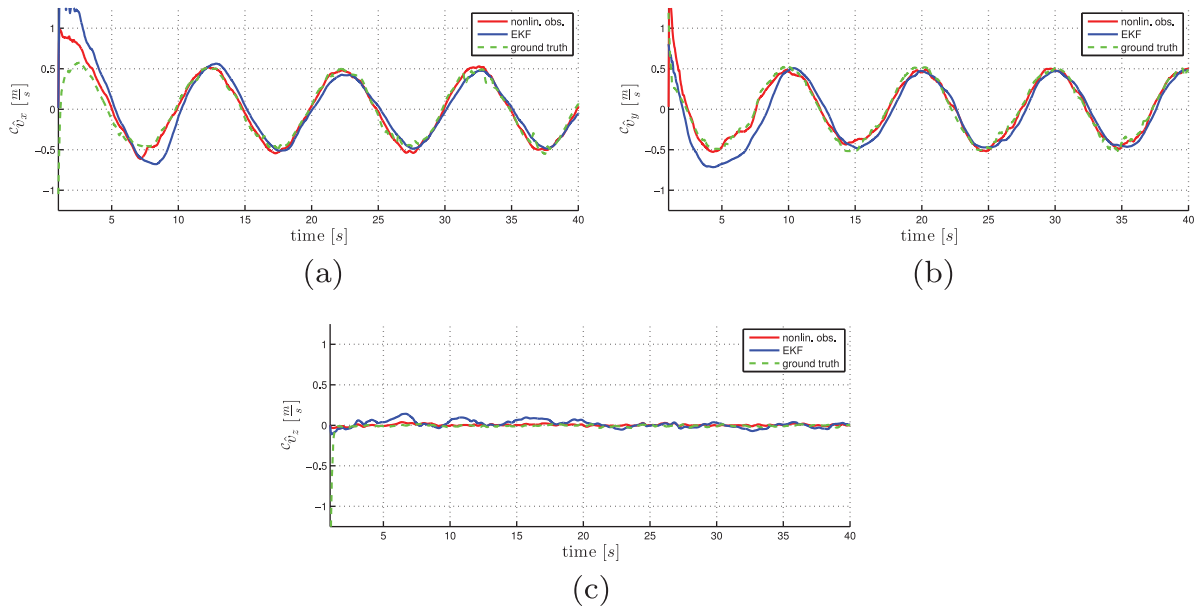


Fig. 10. Experimental results: estimated linear velocity from a recorded dataset in the (a) x , (b) y , and (c) z components. The ground truth was obtained by numerically differentiating the position information from the motion capture system.




Fig. 11. Experimental results: norm of the velocity estimation error against the ground truth for the nonlinear filter and EKF scheme. Note, again, the faster convergence rate of the nonlinear observer with respect to the EKF.

Therefore, the camera–IMU system was moved along a trajectory consisting of small sinusoidal hand-held motions. This resulted in a vehicle trajectory with an amplitude of about 0.1 m/s and a frequency of approximately 1 Hz at a height of 0.5 m. All three directions were tested individually while motion along the other axes was avoided.

Figure 14 shows the estimated metric velocity against ground truth for the nonlinear observer. On this trajectory, we found an RMS value of [0.0074, 0.0095, 0.0114] m/s for the three Cartesian directions, respectively. This allows us to compare our results to the figures of [0.028, 0.035, 0.025] m/s reported in Weiss et al. (2012). Therefore, using the nonlinear observer, we could obtain an average improvement of 320% compared with the results reported in Weiss et al. (2012). All results are summarized in Table 2. Although the experimental conditions are obviously different since the original dataset is not available, we believe these results do indicate the effectiveness of the proposed nonlinear observer in dealing with visual–inertial

scale and ego-motion estimation versus other consolidated approaches.

4.4. Estimation of the gravity vector

Both scale estimation schemes require knowledge of the direction of the gravity vector \mathbf{g} , which, as explained in Section 3.3.1, is assumed provided by the IMU itself. In order to verify the accuracy of this gravity estimation, we compared the IMU estimation of \mathbf{g} against the ground truth obtained from the motion tracking system during the circular flights conducted for the experiments of Section 4.3.2.

The comparison of the x and y components of the normalized gravity vector $\mathbf{g}/\|\mathbf{g}\|$ are shown in Figure 15. On average, the error of the gravity direction estimated by the IMU versus the ground truth is smaller than 1° . This error is smaller than the accuracy with which we were able to define the object frame of the quadrotor in the motion tracking system. Therefore, we can consider the estimated \mathbf{g} from the onboard IMU suitable for our needs.

4.5. Predicting the error convergence time with the nonlinear observer

As shown in Figure 5b–(c) and discussed in Sections 4.2–4.3, keeping a constant acceleration norm $\|\mathbf{a}(t)\|$ allows to impose to the estimation error $\mathbf{z}(t) = 1/d(t) - 1/\hat{d}(t)$ of the nonlinear observer a behavior equivalent to the ideal response (48). This possibility allows for a more general observation. Given a desired time for reaching some percentage of the initial error $\mathbf{z}(t_0)$, and choosing a gain α , one can exploit (48) to determine the needed *acceleration norm* to achieve this goal. Similarly,

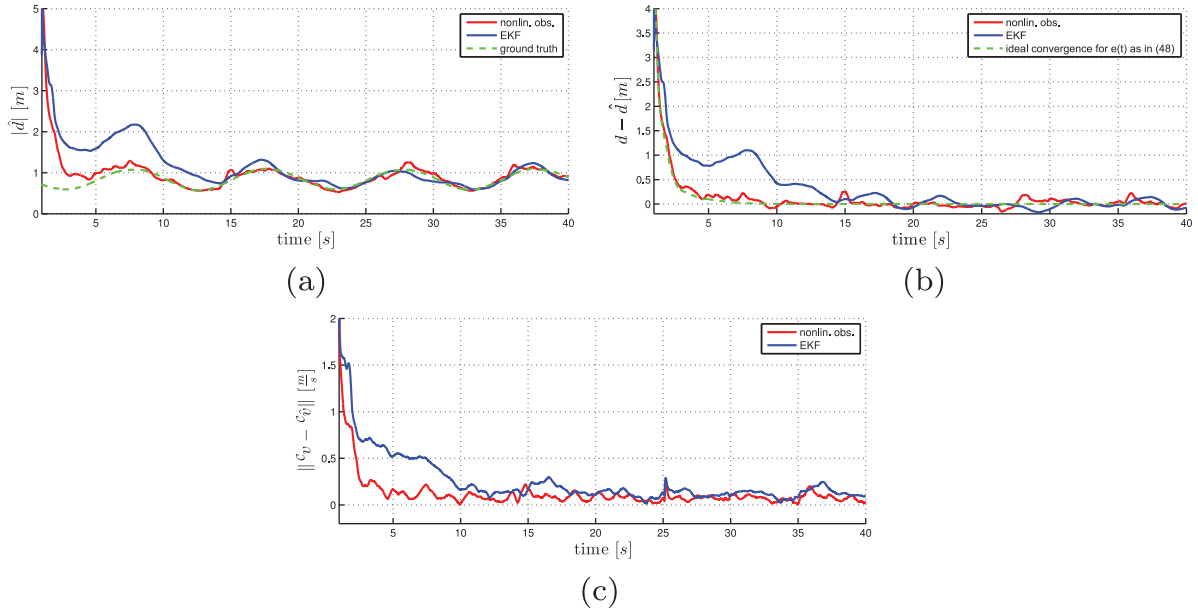


Fig. 12. Experimental results: results from flying along an inclined circular trajectory, thus involving a time-varying $d(t)$. (a) Estimated $\hat{d}(t)$ for the nonlinear observer and EKF, (b) Estimation error $d(t) - \hat{d}(t)$ and ideal convergence behavior (48) for the nonlinear observer, and (c) norm of the velocity estimation error against the ground truth. Note, again, the very good performance of the nonlinear observer in recovering $d(t)$ and $\mathbf{v}(t)$ despite the more challenging UAV motion (as well as, again, the almost perfect match between the actual and ideal transient response of the estimation error).

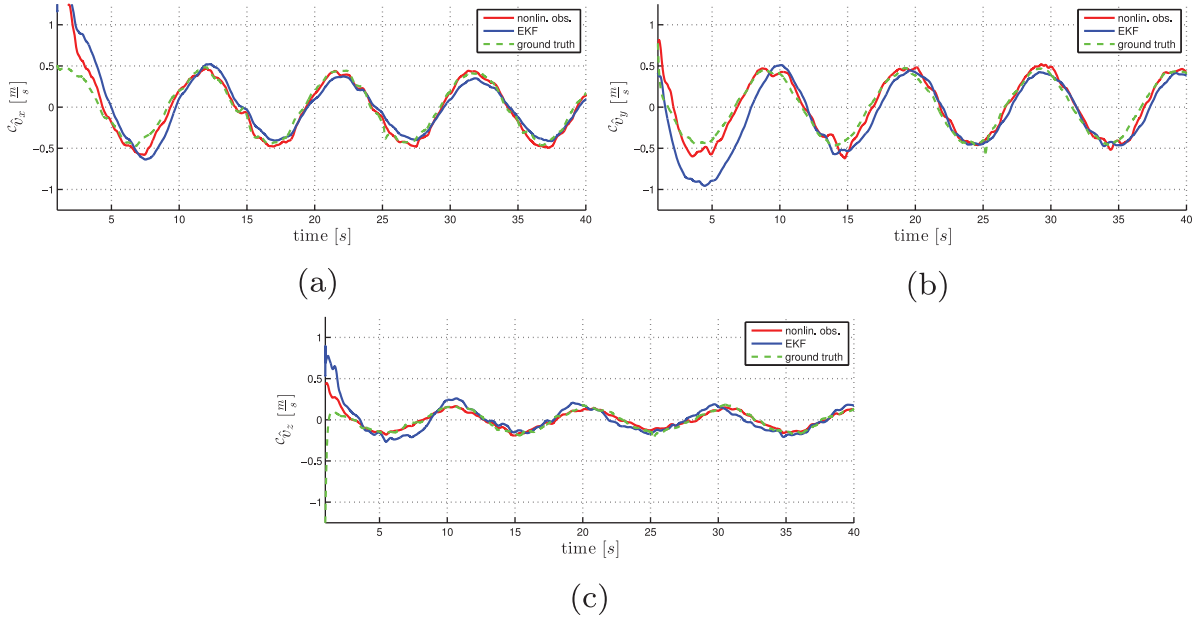


Fig. 13. Experimental results: results from flying along an inclined circular trajectory, thus involving a time-varying $d(t)$. Estimated linear velocities in the (a) x , (b) y , and (c) z components.

for a given trajectory with a known acceleration norm, and a given gain α , one can determine the time needed to reach a desired percentage of the initial estimation error. Such a relation between acceleration norm and convergence time is plotted in Figure 16 for the three cases of reaching 10%, 1%, and 0.1% of the initial error under the (arbitrary)

choice $\alpha = 12$. Thus, for a bounded initial estimation error, i.e. with the maximum distance to the planar scene known before launch, one can univocally predict how long to fly (with constant acceleration norm) for having a guaranteed accuracy in the estimated distance. We note that this analysis can, of course, be done for any choice of gain α .

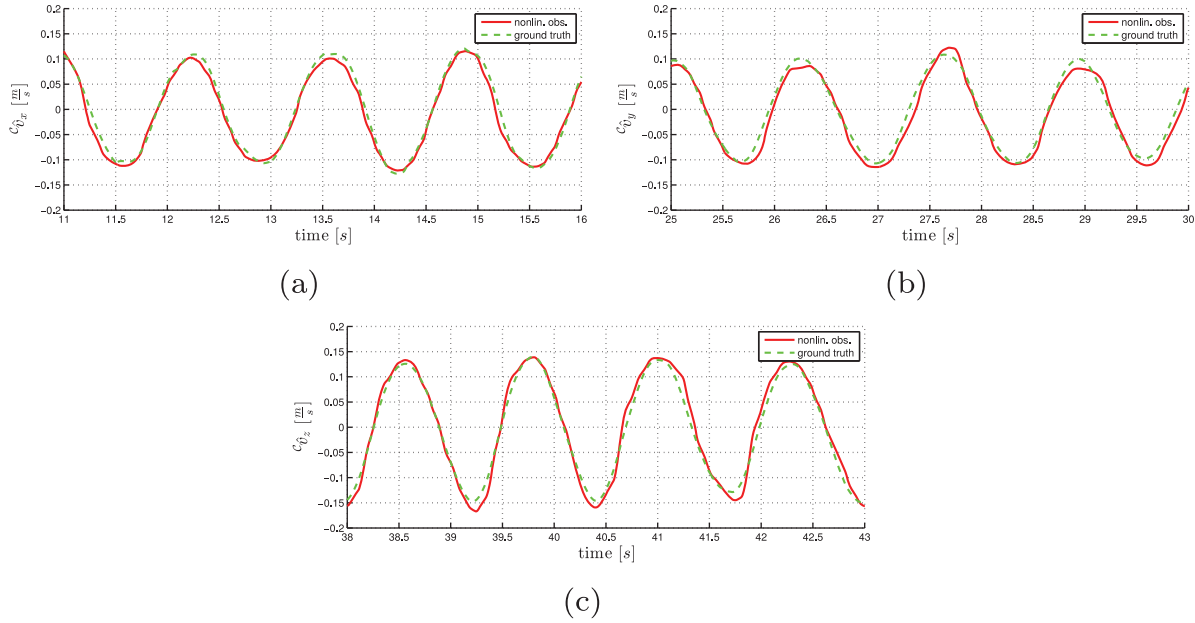


Fig. 14. Experimental results: estimated linear velocity applying the nonlinear observer on the experimental conditions similar to those used in Weiss et al. (2012). The results for the (a) x , (b) y and (c) z directions are plotted against the ground truth obtained from the tracking system.




Fig. 15. Experimental results: x and y components of the normalized gravity vector $\mathbf{g}/\|\mathbf{g}\|$ estimated by the internal sensor fusion algorithm of the employed IMU against the ground truth obtained from the motion capture system. On average, the error between the IMU estimation and the ground truth is in the range of 1° despite the accelerated motion undergone by the UAV.

4.6. Noise robustness of the nonlinear observer

As opposed to the fully informed EKF, the nonlinear observer handles presence of noise in the sensor readings in an indirect way. In particular, the single gain α regulates the ‘aggressiveness’ of the filter, influencing both the convergence speed and its sensitivity to noise. As explained, in all of the previous experiments we tuned α so as to obtain the same noise level (variance of the estimation error after convergence) in the nonlinear observer and EKF cases.

To test the robustness of the nonlinear observer to increased levels of sensor noise for the same choice of α , we increased the simulated noise by a factor of 50 for both the IMU readings and the scaled velocity \mathbf{v}/d from the optical flow decomposition. Figure 17a shows the influence of different sensor noise levels on the estimation error




Fig. 16. Relation between the required ‘excitement’ $\|\Omega\|\sqrt{\alpha}$ and convergence time t needed to reach the same estimation accuracy. The plot was generated for a gain $\alpha = 12$. The three curves show the relation between the time needed to reach a percentage $\epsilon = [10\%, 1\%, 0.1\%]$ of the initial error vs $\|\Omega\|\sqrt{\alpha}$. The plot can be read as follows: on a trajectory with an acceleration norm of 0.296 m/s^2 , the error will drop below 10% of its initial value within 3.79 s and below 1% within 6.47 s.

$d(t) - \hat{d}(t)$. Despite the high noise level, the filter demonstrates a good level of robustness in generating a consistent state estimation.

Furthermore, from the theoretical analysis of Section 3.4, the convergence rate of the nonlinear observer in (48) is determined by the quantity $\sigma_d = \|\Omega\|\sqrt{\alpha}$. Thus, one can always trade smaller accelerations with a higher gain α for obtaining the same (ideal) convergence rate. In this sense, we compared a circular flight with an acceleration norm of $\|\Omega\| \approx 0.296 \text{ m/s}^2$ and a choice of $\alpha = 12$ to two flights with acceleration norms 0.148 and 0.037 m/s^2 , and the gain α chosen so as to yield the same σ_d (i.e. same error convergence rate). The results are reported in Figure 17b where we can note, as expected, that by keeping the same σ_d one

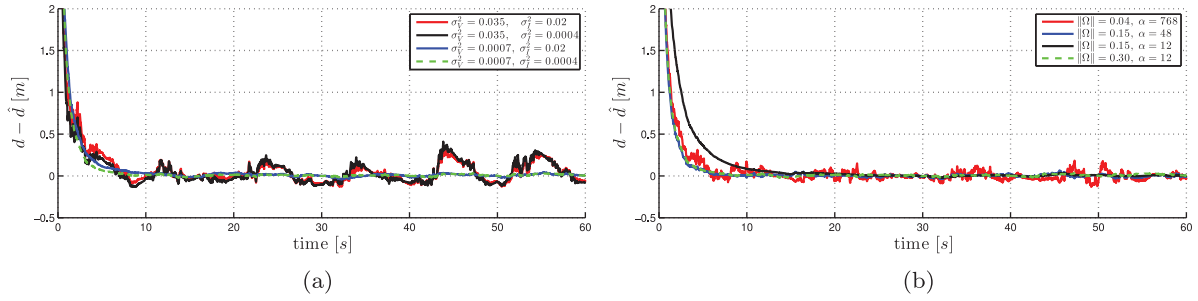


Fig. 17. Experimental results: influence of sensor noise on the nonlinear observer scheme. (a) Robustness of the nonlinear observer to different noise levels. The variance values of $\sigma_v^2 = 0.0007$ and $\sigma_l^2 = 0.0004$ for the visual and inertial sensor data, respectively, correspond to the actual noise found on the real sensors. The noise was then increased 50 times for a comparison. (b) Plot showing the trade-off between 'control effort' (norm of the acceleration $\|\Omega\|$) and gain α . In the case of lower accelerations, the same convergence rate can be re-established using a higher gain α but at the expense of an increased noise level as clear from the plot. Gain α can then be used to tune the 'aggressiveness' of the filter versus a given amount of control effort $\|\Omega\|$.

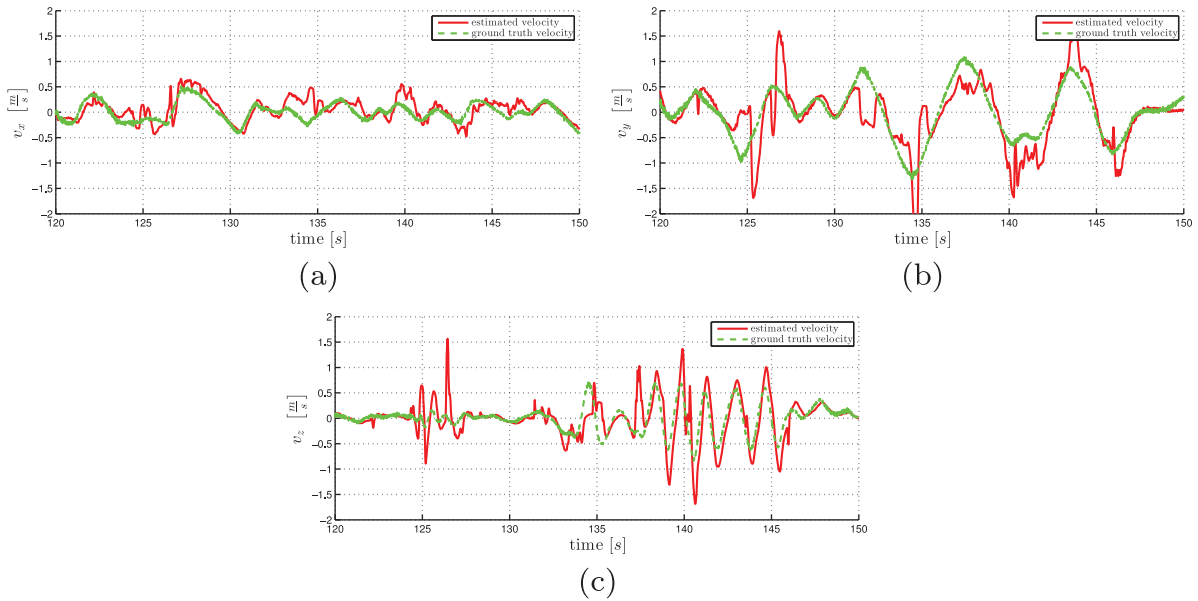


Fig. 18. Experimental results: estimated and ground truth linear velocity during a closed loop flight under velocity control for the (a) x, (b) y, and (c) z axes. An object was moved through the field of view at $t = 125$ s inducing a transient disturbance on the estimated velocity.

obtains the same transient behavior for the estimation error $z(t)$, although a higher noise level is clearly induced by the larger employed α .

4.7. Closed-loop control of a real quadrotor UAV

As a final validation of the overall framework, we made use of the estimated metric velocity to 'close the loop' and control the UAV motion in real time. For this experiment, we combined algorithm V2 for the optical flow decomposition with the scale estimation obtained from the nonlinear observer. The vehicle was commanded using a gamepad to send velocity commands via a wireless link.

Figure 18 shows a section of a longer flight relying purely on onboard sensors. To test the robustness of the

system, we moved an object through the field of view of the down-facing camera at time $t = 125$ s, causing some disturbances in the estimated velocity. Starting from time $t = 134$ s, the vehicle was additionally commanded to move in the vertical direction, causing the height, and therefore the distance to the plane, to change. The plots show some temporary over- and under-estimations of the linear velocity due to the abrupt commanded motions, but otherwise the proposed approach is always able to adapt to the changing height and recover a reliable estimation.

A video from one of these experiments (Multimedia Extension 1) is also attached to the paper. Both external views of the vehicle and the image stream from the down-facing camera are shown while the vehicle navigates purely based on onboard sensors and processing capabilities.

We then believe these experiments contribute to demonstrate that the presented velocity-based state estimation pipeline can be used for the closed-loop control of a real quadrotor UAV using solely onboard sensors.

5. Conclusions and future work

In this paper, we have addressed the need of a reliable onboard ego-motion estimation for quadrotor UAVs to overcome the boundaries of controlled indoor environments (such as when relying on external motion tracking systems for recovering the UAV state). To this end, we first discussed three variants of an algorithm based on the continuous homography constraint to obtain an estimation of the UAV scaled linear velocity and angular velocity from the decomposition of the perceived optical flow. This step, indeed, allows to retrieve ego-motion information independently of maps, known landmarks, and the need for tracking features over an extended period of time. Subsequently, we extensively discussed the issue of estimating the (unknown) metric scale factor by fusing the scaled velocity retrieved from optical flow with the high-frequency accelerometer readings of an onboard IMU. Scale estimation was achieved by proposing two estimation schemes: a first one based on a classical EKF and a second one on a novel nonlinear observation framework. Simulated and real experimental data were presented to assess and compare the performance of both filters in ideal and real conditions. When compared with the EKF, the nonlinear observer demonstrated a consistent better performance in terms of convergence rate of the scale estimation error. Furthermore, the proposed theoretical analysis showed the possibility to actively impose (and thus predict) the error transient response of the nonlinear observer by suitably acting on the estimation gains and UAV motion (the norm of its acceleration). This analysis was, again, fully confirmed by the reported simulative/experimental results in several conditions, also involving different levels of sensor noise for testing the robustness of the approach.

With the advantage of a fast and predictable convergence for recovering the metric UAV linear velocity, the nonlinear observer proved to be a suitable choice for then implementing a fully onboard control of flying vehicles. In particular, we successfully demonstrated the reliability of the proposed framework in closed-loop experiments on a quadrotor UAV.

5.1. Future work

Despite the convincing results, we are considering extending our work in several ways. A first possibility is to further exploit the awareness of the error convergence behavior when using the nonlinear filter. Inspired by what was presented in Achtelik et al. (2013) in the context of *offline* path planning, one could devise an *online* strategy meant to adjust the input UAV trajectory (e.g. the velocity commands) in order to maintain a desired level of *excitement*, in our case represented by the quantity $\sigma_d = \sqrt{\alpha} \|\Omega\|$. This

could result in, e.g. execution of small circular trajectories when the vehicle is commanded to hover in place with, instead, a more precise tracking of the desired/commanded velocities when undergoing more aggressive manoeuvres.

On a similar note, we are currently investigating the possibility to dynamically adapt gain α for the nonlinear observer. Indeed, in the presence of noisy sensor readings, it might be desirable to start with a high value of α for imposing an initial quick error convergence, and to then reduce it when a sufficient convergence level is reached so as to obtain a smoother state estimation (see also Figure 17b). Finally, it would also be worth investigating the possibility to extend the structure of the nonlinear observer in order to account for possible *biases* in the employed measurements: this is, for instance, the case of the squared term $x_1 x_1^T$ in (42) which, being $x_1 = v/d$ a noisy measurement, will introduce a non-compensated bias in the observer dynamics. While presence of this bias had, in practice, a negligible effect in the reported results, an extension able to compensate for it online would clearly provide an interesting theoretical and practical improvement (especially in situations involving higher level of noise in the measurements).

Acknowledgements

The authors would like to thank Dr Antonio Franchi and Martin Riedel for their valuable suggestions and contribution on the development of the underlying software framework.

Funding

Part of Heinrich Bülthoff's research was supported by the Brain Korea 21 PLUS Program through the National Research Foundation of Korea funded by the Ministry of Education. Part of Davide Scaramuzza's research was supported by the Swiss National Centre of Competence in Research (NCCR) Robotics.

Notes

1. Specifically, the points must not lie on some particular quadratic surfaces which, however, include planes as a special case, see Ma et al. (2004).
2. We assume the IMU is able to self-calibrate by an initial estimation of the measurement biases (indeed, this is often the case of many commercial IMUs, including the one employed in the experiments reported in Section 4). The interested reader is referred to, e.g., Martinelli (2012), Weiss et al. (2012), and Weiss et al. (2013a) for other estimation approaches addressing the concurrent IMU bias and structure estimation during flight.
3. Therefore, given an upper bounded initial error $d(t_0) - \hat{d}(t_0)$, one can, for instance, plan in advance the duration of the UAV motion so as to yield a guaranteed accuracy in the estimated plane distance.

References

- Achtelik MW, Weiss S, Chli M and Siegwart R (2013) Path planning for motion dependent state estimation on micro aerial vehicles. In: *Proceedings of the international conference on robotics and automation*, pp. 3926–3932.

- AIRobots (2010) EU Collaborative Project ICT-248669. <http://www.airobots.eu>.
- ARCAS (2011) EU Collaborative Project ICT-287617. <http://www.arcas-project.eu>.
- Chaumette F and Hutchinson S (2006) Visual servo control, part I. Basic approaches. *Robotics and Automation Magazine* 13(4): 82–90.
- Civera J, Davison AJ and Montiel J (2008) Inverse depth parametrization for monocular SLAM. *Transactions on Robotics* 24(5): 932–945.
- De Luca A, Oriolo G and Robuffo Giordano P (2008) Feature depth observation for image-based visual servoing: theory and experiments. *The International Journal of Robotics Research* 27(10): 1093–1116.
- Eudes A, Morin P, Mahony R and Hamel T (2013) Visuo-inertial fusion for homography-based filtering and estimation. In: *Proceedings of the international conference on intelligent robots and systems*, Tokyo, Japan, pp. 5186–5192.
- Faessler M, Fontana F, Forster C and Scaramuzza D (2015) Automatic re-initialization and failure recovery for aggressive flight with a monocular vision-based quadrotor. In: *Proceedings of the international conference on robotics and automation*, Seattle, WA, USA.
- Forster C, Pizzoli M and Scaramuzza D (2014) SVO: fast semi-direct monocular visual odometry. In: *Proceedings of the international conference on robotics and automation*, Hong Kong, China.
- Grabe V, Bühlhoff HH and Robuffo Giordano P (2012a) On-board velocity estimation and closed-loop control of a quadrotor uav based on optical flow. In: *Proceedings of the international conference on robotics and automation*, St Paul, MN, USA, pp. 491–497.
- Grabe V, Bühlhoff HH and Robuffo Giordano P (2012b) Robust optical-flow based self-motion estimation for a quadrotor UAV. In: *Proceedings of the international conference on intelligent robots and systems*, Vilamoura, Portugal, pp. 2153–2159.
- Grabe V, Bühlhoff HH and Robuffo Giordano P (2013a) A comparison of scale estimation schemes for a quadrotor UAV based on optical flow and IMU measurements. In: *Proceedings of the international conference on intelligent robots and systems*, Tokyo, Japan, pp. 5193–5200.
- Grabe V, Riedel M, Bühlhoff HH, Robuffo Giordano P and Franchi A (2013b) The TeleKyb framework for a modular and extendible ROS-based quadrotor control. In: *Proceedings of the European conference on mobile robots*, Barcelona, Spain, pp. 19–25.
- Herissé B, Hamel T, Mahony R and Russotto FX (2012) Landing a VTOL unmanned aerial vehicle on a moving platform using optical flow. *Transactions on Robotics* 28(1): 77–89.
- Honegger D, Meier L, Tanskanen P and Pollefeys M (2013) An open source and open hardware embedded metric optical flow CMOS camera for indoor and outdoor applications. In: *Proceedings of the international conference on robotics and automation*, Karlsruhe, Germany, pp. 1736–1741.
- Klein G and Murray D (2007) Parallel tracking and mapping for small AR workspaces. In: *Proceedings of the international symposium on mixed and augmented reality*, Nara, Japan, pp. 225–234.
- Li M and Mourikis AI (2013) High-precision, consistent EKF-based visual-inertial odometry. *The International Journal of Robotics Research* 32(6): 690–711.
- Lobo J and Dias J (2007) Relative pose calibration between visual and inertial sensors. *The International Journal of Robotics Research* 26(6): 561–575.
- Lucas BD and Kanade T (1981) An iterative image registration technique with an application to stereo vision. In: *Proceedings of the international joint conference on artificial intelligence*, Vancouver, BC, Canada, pp. 674–679.
- Ma Y, Soatto S, Kosecka J and Sastry SS (2004) *An Invitation to 3-D Vision*. New York: Springer.
- Mahony R, Hamel T and Pflimlin JM (2008) Nonlinear complementary filters on the special orthogonal group. *Transactions on Automatic Control* 53(5): 1203–1218.
- Marino R and Tomei P (1995) *Nonlinear Control Design: Geometric, Adaptive and Robust*. Englewood Cliffs: Prentice-Hall.
- Martinelli A (2012) Vision and IMU data fusion: closed-form solutions for attitude, speed, absolute scale, and bias determination. *Transactions on Robotics* 28(1): 44–60.
- Nützi G, Weiss S, Scaramuzza D and Siegwart R (2011) Fusion of IMU and vision for absolute scale estimation in monocular SLAM. *Journal of Intelligent and Robotic Systems* 61(1-4): 287–299.
- Omari S and Ducard G (2013) Metric visual-inertial navigation system using single optical flow feature. In: *Proceedings of the European control conference*, Zurich, Switzerland, pp. 1310–1316.
- Robuffo Giordano P, De Luca A and Oriolo G (2008) 3D structure identification from image moments. In: *Proceedings of the international conference on robotics and automation*, pp. 93–100.
- Rosten E and Drummond T (2005) Fusing points and lines for high performance tracking. In: *Proceedings of the international conference on computer vision*, Beijing, China, pp. 1508–1515.
- Scaramuzza D, Achtelik MC, Doitsidis L, et al. (2014) Vision-controlled micro flying robots: from system design to autonomous navigation and mapping in GPS-denied environments. *IEEE Robotics and Automation Magazine* 21(3): 26–40.
- Scaramuzza D and Fraundorfer F (2011) Visual odometry: part I - the first 30 years and fundamentals. *IEEE Robotics and Automation Magazine* 18(4): 80–92.
- Sherpa (2013) EU Collaborative Project ICT-600958. <http://www.sherpa-project.eu/sherpa>.
- Spica R and Robuffo Giordano P (2013) A framework for active estimation: application to structure from motion. In: *Proceedings of the conference on decision and control*, pp. 7647–7653.
- Spica R, Robuffo Giordano P and Chaumette F (2014) Active structure from motion: application to point, sphere, and cylinder. *IEEE Transactions on Robotics* 30(6): 1499–1513.
- Weiss S, Achtelik MW, Lynen S, et al. (2013a) Monocular vision for long-term micro aerial vehicle state estimation: a compendium. *Journal of Field Robotics* 30(5): 803–831.
- Weiss S, Achtelik MW, Lynen S, Chli M and Siegwart R (2012) Real-time onboard visual-inertial state estimation and self-calibration of MAVs in unknown environments In: *Proceedings of the International Conference on Robotics and Automation*, St Paul, MN, USA, pp. 957–964.
- Weiss S, Brockers R and Matthies L (2013b) 4DoF drift free navigation using inertial cues and optical flow. In: *Proceedings of the international conference on intelligent robots and systems*, Tokyo, Japan, pp. 4180–4186.

Appendix: Index to Multimedia Extension

Archives of IJRR multimedia extensions published prior to 2014 can be found at <http://www.ijrr.org>, after 2014 all videos are available on the IJRR YouTube channel at <http://www.youtube.com/user/ijrrmultimedia>

Table of Multimedia Extension

Extension	Media type	Description
1	Video	Experiments of UAV closed-loop control with only onboard sensing and processing power