

# Accuracy Characterization of the Vehicle State Estimation from Aerial Imagery

Eduardo Sánchez Morales<sup>1</sup>, Friedrich Kruber<sup>1</sup>, Michael Botsch<sup>1</sup>, Bertold Huber<sup>2</sup> and Andrés García Higuera<sup>3</sup>

**Abstract**—Due to their capability of acquiring aerial imagery, camera-equipped Unmanned Aerial Vehicles (UAVs) are very cost-effective tools for acquiring traffic information. However, not enough attention has been given to the validation of the accuracy of these systems. In this paper, an analysis of the most significant sources of error is done. This includes three key components. First, a vehicle state estimation by means of statistical filtering. Second, a quantification of the most significant sources of error. Third, a benchmark of the estimated state compared with state-of-the-art reference sensors. This work presents ways to minimize the errors of the most relevant sources. With these error reductions, camera-equipped UAVs are very attractive tools for traffic data acquisition. The test data and the source code are made publicly available.

## I. INTRODUCTION AND MOTIVATION

Given that cars can harm humans due to malfunction, they have to be tested extensively. Real-world data is a key element for the development, testing and validation of algorithms. However, the acquisition of real-world data is resource intensive, which can in turn limit the amount of available data for research.

Given their low cost, versatility and capability of acquiring aerial imagery, camera-equipped UAVs are popular data acquisition tools for automotive purposes. Having such video material is highly practical, considering that many vehicle motion models found in the literature describe the movement of cars precisely from this perspective [1][2][3]. Also, since the bird's eye-view is less prone to occlusion problems, it is a highly practical representation for traffic scenarios. Another advantage is the ability of recording various objects with a single device. Some relevant applications of aerial imagery for automotive data acquisition include the traffic flow analysis [4], the training of machine learning techniques [5] and the validation of autonomous vehicles [6].

On the other hand, the accuracy validation of camera-equipped UAVs has not received enough attention. When deploying these devices for automotive data acquisition, a series of sensor- and algorithm-related assumptions are made, which translate to inaccuracies. When considering the assumptions individually, the expected error might be negligible. But when considered together, lead to statistically significant errors.

The methodology used in the presented work evaluates the accuracy of camera-equipped UAVs when used for vehicle

state estimation. It also clarifies and quantifies the most significant sources of error. The results can then be used as a best practice guide to minimize errors in future works.

This work makes the following contributions:

It shows 1) a method for vehicle state estimation when using a camera-equipped UAV as sensor, 2) an analysis and quantification of the most significant sources of error, and 3) a benchmark of the estimated vehicle state using state-of-the-art reference sensors.

This paper is structured in the following manner: first, a brief review on vehicle state estimation by means of cameras is given (Section II). Then, the methodology for estimating the vehicle state using aerial imagery is detailed (Section III). Next, the most significant sources of error are analysed and quantified (Section IV). Later, the estimated vehicle state is benchmarked with state-of-the-art sensors (Section V). Finally, the conclusions are presented (Section VI).

Vectors are represented in boldface and matrices in boldface, capital letters. All the units are given in the International System of Units (SI), unless otherwise specified.

## II. RELATED WORK

The task of using cameras for estimating the vehicle state is not new, and the literature research reveals a wide variety of approaches. An example of a concept using a fixed camera is found on [7]. In that publication, the authors place a camera in the middle of a road and pointed it alongside the traffic flow. After detecting the surface plane, the distance from the vehicles to the camera is calculated. An average velocity is estimated from the calculated distance at each frame.

An example of inclusion of machine learning techniques for velocity estimation can be found in [8]. There, the authors document an algorithm that estimates the velocity and position of vehicles relative to a monocular camera mounted on a moving vehicle. Their method relies on extracting features from videos, and later regress the velocity and position of the vehicles by means of a Multi-Layer Perceptron.

Other works provide already processed video data. Examples of this are the HighD [9], the InD [10] and the INTERACTION [11] datasets. The first one provides traffic information for highway. The second and third provide traffic information for urban environments.

The publication of [12] approximates the most to the presented work. There, the authors equipped a test vehicle with a Satellite Navigation (SatNav) system, which allows to measure position and velocity. The main differences from [12] are that in the presented work: 1) More state variables are benchmarked. Not only the vehicle position and velocity, but the vehicle acceleration, orientation, course

<sup>1</sup>Eduardo Sánchez Morales, Friedrich Kruber and Michael Botsch are with the Department of Vehicle Safety, CARISSMA, Technische Hochschule Ingolstadt, Ingolstadt, Germany.

<sup>2</sup>Bertold Huber is with GeneSys Elektronik GmbH, Offenburg, Germany.

<sup>3</sup>Andrés García Higuera is with the Department of Electrical, Electronic and Automation Engineering, University of Castilla-La Mancha, Ciudad Real, Spain.

over ground and sideslip angle are compared too. 2) The relief displacement [13] is corrected. This greatly improves the accuracy of the estimated position. 3) The benchmark is done using more accurate sensors. The sensor used in [12] has an accuracy of 20 cm and 0.1 km/h in position and velocity respectively [14]; while the one used in this work has a position and velocity accuracy of 1 cm and 0.03 km/h accordingly. A better sensor accuracy allows a pixel-accurate comparison, which makes the benchmark more relevant. 4) Test data from a sample video and the source code are publicly available on [15]. This helps future works to easily generate a measurement vector from aerial imagery for the Kalman Filter (KF), and to fine tune the Kalman gains for specific applications.

### III. VEHICLE STATE ESTIMATION

The vehicle state is estimated by means of a KF because it allows to calculate state variables by system noise propagation. The detailed process is explained next.

#### A. Coordinate Systems

The coordinate systems used in this work are explained in what follows.

The vehicles move on the Local Tangent Plane (LTP). This plane can be defined in the East-North-Up (ENU) coordinate system, where  $x_{LTP}$  points east,  $y_{LTP}$  north and  $z_{LTP}$  upwards, with arbitrary origin  $o_{LTP}$  on the surface of the Earth. The Local Car Plane (LCP) is defined according to the ISO8855:2011 norm, where  $x_{LCP}$  points to the hood,  $y_{LCP}$  to the driver,  $z_{LCP}$  upwards, with origin  $o_{LCP}$  at the center of sprung mass of the car. For simplification purposes, it is assumed that 1) the  $x_{LCP}y_{LCP}$ -plane is parallel to the  $x_{LTP}y_{LTP}$ -plane, 2) the centre of mass of the car is the same as its geometrical centre, and 3) all sensors mounted on the vehicle measure in the LCP. The Pixel Coordinate Frame (PCF) is a bird's-eye image projection of the LTP. It is composed by the mutually perpendicular  $x_{PCF}$  and  $y_{PCF}$  axes, with origin  $o_{PCF}$  in one corner of the image. All quantities expressed in PCF are given in pixels (px).

#### B. PCF Mapping

The first step for estimating the vehicle state is the mapping of the PCF to the LTP. For this, Ground Control Points (GCPs) are placed on the  $x_{LTP}y_{LTP}$ -plane, in such a way that are visible on the PCF. The middle of the  $i$ -th GCP is defined in LTP as

$$\mathbf{g}_{i,LTP} = [x_{i,LTP} \quad y_{i,LTP}]^T, \quad (1)$$

and in PCF as

$$\mathbf{g}_{i,PCF} = [x_{i,PCF} \quad y_{i,PCF}]^T. \quad (2)$$

The PCF and the LTP have different scales. The spatial resolution  $\alpha$  that relates both scales is calculated from two GCPs by

$$\alpha = \frac{|\mathbf{g}_{i+1,LTP} - \mathbf{g}_{i,LTP}|}{|\mathbf{g}_{i+1,PCF} - \mathbf{g}_{i,PCF}|}. \quad (3)$$

The  $i$ -th GCP can then be expressed in PCF and in meters by

$$\mathbf{g}'_i = \mathbf{g}_{i,PCF} \cdot \alpha = [x'_i \quad y'_i]^T. \quad (4)$$

The orientation offset  $\xi_{\theta_{image}}$  from the LTP to the PCF is calculated as

$$\xi_{\theta_{image}} = \theta_i - \theta_{i,LTP}, \text{ with} \quad (5)$$

$$\theta_i = \text{atan2}(y'_{i+1} - y'_i, x'_{i+1} - x'_i), \text{ and} \quad (6)$$

$$\theta_{i,LTP} = \text{atan2}(y_{i+1,LTP} - y_{i,LTP}, x_{i+1,LTP} - x_{i,LTP}). \quad (7)$$

The GCP  $\mathbf{g}'_i$  is then rotated as follows

$$\mathbf{g}''_i = \mathbf{R}(\xi_{\theta_{image}})^T \mathbf{g}'_i, \quad (8)$$

where  $\mathbf{R}(\cdot)$  is a 2D rotation matrix. Finally, the linear offsets  $\xi_d$  from the LTP to the PCF are calculated by

$$\xi_d = \mathbf{g}''_i - \mathbf{g}_{i,LTP}. \quad (9)$$

So, a pixel  $\mathbf{p}_{PCF} = [x_{p,PCF} \quad y_{p,PCF}]^T$  on the PCF can be mapped to the LTP by

$$\mathbf{p}_{PCF}^{LTP} = \left( \mathbf{R}(\xi_{\theta_{image}})^T (\mathbf{p}_{PCF} \cdot \alpha) \right) - \xi_d. \quad (10)$$

#### C. Measurement Vector Generation

The next step is to generate a measurement vector from the videos recorded by the UAV and the parameters estimated in Section III-B. This requires to process the video sequence for detecting vehicles. The vehicle recognition can be achieved by a variety of machine learning methods that can be found in the literature, such as YOLO [16], DroNet [17] or  $R^3$  [18]. The first two output axes aligned bounding boxes; while the third one outputs rotated bounding boxes. The latter greatly improve the state estimation because they allow to calculate the vehicle orientation and a much more precise geometrical center.

The chosen detector for this work is Mask-RCNN [19][20]. The detailed methodology for processing video material to obtain an adequate bounding box, as well as open source code, can be found on [21].

The expected input from the UAV is a bounding box per frame of the video sequence that 1) has 4 sides, 2) covers all pixels corresponding to the detected shape of the vehicle, and 3) covers the least amount of pixels possible. The corners of the bounding box are given by the intersection of its sides. The  $i$ -th corner of the bounding box is defined in PCF as

$$\mathbf{b}_i = [x_{i,PCF} \quad y_{i,PCF}]^T, \quad (11)$$

and the bounding box is defined in PCF as

$$\mathbf{B}_{PCF} = [\mathbf{b}_1 \quad \mathbf{b}_2 \quad \mathbf{b}_3 \quad \mathbf{b}_4]. \quad (12)$$

The corners of the bounding box are mapped to the LTP as shown in the Equation (10) to obtain  $\mathbf{B}_{PCF}^{LTP}$ . Let  $\mathbf{b}_x$  and  $\mathbf{b}_y$  be horizontal vectors equal to the first and second rows of  $\mathbf{B}_{PCF}^{LTP}$  respectively. Then, the geometric centre of the vehicle

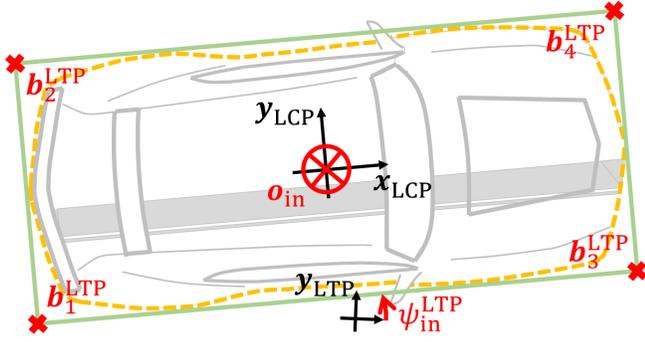


Fig. 1. Shown is: the shape of the car as detected by Mask-RCNN with an orange dashed line; the corresponding bounding box with a green solid line; the corners of the bounding box, the geometrical center and orientation of the car with red.

$o_{in}$  is calculated by

$$o_{in} = \begin{bmatrix} \frac{\max(b_x) + \min(b_x)}{2} \\ \frac{\max(b_y) + \min(b_y)}{2} \end{bmatrix}, \quad (13)$$

where  $\max$  and  $\min$  are functions that choose respectively the maximum and minimum value of a vector.

The vehicle dimensions are calculated next. Let  $s$  be a vector containing the magnitude of the vectors that join the element 1 of  $B_{PCF}$  with all others, so that

$$s(1) < s(2) < s(3). \quad (14)$$

Then  $(s(1) \cdot \alpha)$  and  $(s(2) \cdot \alpha)$  are the estimated width and length of the car in meters. Knowing this, the orientation  $\psi_{in}^{LTP}$  of the vehicle is given by

$$\psi_{in}^{LTP} = \text{atan2}(y_{j,PCF}^{LTP} - y_{1,PCF}^{LTP}, x_{j,PCF}^{LTP} - x_{1,PCF}^{LTP}), \quad (15)$$

where  $j$  is the element of  $B_{PCF}$  associated with  $s(2)$ . The measurement vector for the KF is then defined as

$$z_{in} = [o_{in}^T \quad \psi_{in}^{LTP}]^T. \quad (16)$$

A graphical representation of the bounding box, vehicle geometric centre and orientation is shown on the Figure 1.

#### D. Kalman Filter

Having the measurement vector, the next step is to estimate the vehicle state. One of the most computationally efficient methods for estimating the *optimum* state of a mobile object, assuming Markovian-Gaussian random processes, is the KF [22]. It also allows to estimate state variables that are not part of the measurement vector by allowing the system noise to propagate. The specifics applicable to this work are described in the following.

The used state vector is defined as

$$\mathbf{x} = [x_{car}, y_{car}, v_{x,car}, v_{y,car}, a_{x,car}, a_{y,car}, \psi_{car}, \dot{\psi}_{car}]^T, \quad (17)$$

where  $x_{car}$  and  $y_{car}$  are the (x,y) coordinates of  $o_{LCP}$  in LTP,  $v_{x,car}$  and  $v_{y,car}$  are the velocities of  $o_{LCP}$  along the  $x_{LTP}$  and  $y_{LTP}$  axes,  $a_{x,car}$  and  $a_{y,car}$  are the accelerations of  $o_{LCP}$  along the  $x_{LTP}$  and  $y_{LTP}$  axes,  $\psi_{car}$  is the vehicle yaw in LTP, and  $\dot{\psi}_{car}$  is the yaw rate around  $z_{LCP}$ .



Fig. 2. Shown is a GCP as seen on aerial imagery. The UAV hovers at 50 m (left) and at 100 m (right). The blurring effect of the image prevents to unambiguously assign a pixel to a corner of the GCP.

Since the course over ground  $\theta_{cog}$  in LTP of the vehicle is defined as

$$\theta_{cog} = \text{atan2}(v_{y,car}, v_{x,car}), \quad (18)$$

the sideslip  $\beta_{car}$  of the car can then be calculated by

$$\beta_{car} = \theta_{cog} - \psi_{car}. \quad (19)$$

In this work, the sideslip angle is estimated by means of a Linear Kalman Filter (LKF) and Equation (19). This produces better results than using an Extended Kalman Filter (EKF). This is explained by the fact the sideslip angle is not part of the measurement vector. Using an EKF would imply the estimation of the sideslip angle by noise propagation, whereas Equation (19) allows a direct calculation.

#### IV. ANALYSIS OF SOURCES OF ERROR

Camera-equipped UAVs are practical, versatile and cost-effective for acquiring traffic information. But their sources of error have not been analysed deeply enough. The most significant sources of error are 1) the PCF to LTP Mapping, 2) the UAV stability, 3) the car labelling and detection, 4) the bounding box, and 5) the sensor synchronization. These sources of error are analysed and quantified in the following.

##### A. PCF to LTP Mapping

The literature review shows that the mapping of the GCP to the LTP is usually performed through the use of GCPs. Ideally, one point of the GCP can be associated with a specific pixel on the PCF. In reality, this rarely happens. An example of this can be seen on the Figure 2. The blurring effect that can be appreciated on the image can be caused by image compression, camera optics or light propagation. It is precisely this blurring effect that prevents to unambiguously associate a point of the GCP with a specific pixel. So, a pixel ambiguity  $\zeta = 1 \text{ px}$  in both, the  $x_{PCF}$  and  $y_{PCF}$  axes is not uncommon. So, the  $g_{i,PCF}$  is rewritten as

$$g_{i,PCF} = [x_{i,PCF} \quad y_{i,PCF}]^T \pm \zeta. \quad (20)$$

This association ambiguity has an effect on all three parameters that map the PCF to the LTP. The case of the spatial resolution is analysed first.

Considering a pixel ambiguity of  $\zeta = 1$  per axis and squared pixels, the distance between the true and associated positions of a GCP on the PCF can be of  $\sqrt{2}$  px. This mis-association causes an error on the spatial resolution. The similarity in percentage  $\eta_\alpha$  between the seen and the true values of the spatial resolution is calculated by

$$\eta_\alpha = \frac{|\mathbf{g}_{i+1,\text{PCF}} - \mathbf{g}_{i,\text{PCF}}|}{|\mathbf{g}_{i+1,\text{PCF}} - \mathbf{g}_{i,\text{PCF}}| + 2 \cdot \sqrt{2\zeta^2}}. \quad (21)$$

The Equation (4) is then rewritten as

$$\mathbf{g}'_i = \mathbf{g}_{i,\text{PCF}} \cdot \alpha \cdot \eta_\alpha = [x'_i \quad y'_i]^T. \quad (22)$$

From the Equations (21) and (22), it can be deduced that the effect of  $\eta_\alpha$  increases as  $\mathbf{g}_{i,\text{PCF}} \rightarrow \mathbf{g}_{i+1,\text{PCF}}$ . In a scenario, where the GCPs are only one pixel apart, the similarity between the seen and true spatial resolution can drop to 26%. In another scenario for a Full-HD image, where both GCPs are placed on opposite diagonal corners of the picture, the similarity can drop only to 99.87%. For example, if the true value of  $|\mathbf{g}_{i+1,\text{LTP}} - \mathbf{g}_{i,\text{LTP}}|$  is 100 m, a similarity of 99.87% will rescale it as 99.87 m, meaning a 13 cm difference.

Next, the effect on the orientation offset is analysed. This is done in pixels to decouple errors caused by  $\eta_\alpha$ . To consider the pixel ambiguity, the Equation (6) is rewritten as

$$\xi_{\theta_{\text{image}}} = \text{atan2}(\Delta_y \pm 2\zeta, \Delta_x \pm 2\zeta), \quad (23)$$

where  $\zeta$  is multiplied by two because  $\xi_{\theta_{\text{image}}}$  is calculated using two GCPs. Similar as with  $\eta_\alpha$ , the effect of the pixel ambiguity increases as  $\mathbf{g}_{i,\text{PCF}} \rightarrow \mathbf{g}_{i+1,\text{PCF}}$ . In a scenario for a Full-HD image, where the GCPs are one pixel apart, the orientation error  $\eta_{\xi_{\theta_{\text{image}}}}$  could reach  $116^\circ$ . Else, if the GCPs are in opposing diagonal corners of the image, then  $\eta_{\xi_{\theta_{\text{image}}}}$  could reach only  $0.07^\circ$ .

The orientation error affects the rotation step of the PCF to LTP mapping. So, the effect of orientation error increases as the points to map are farther from the rotation axis. For the  $i$ -th corner of the bounding box, the mapping error  $\eta_{b_i}$  due to the orientation error is given by

$$\eta_{b_i} = \left| \mathbf{R} \left( \eta_{\xi_{\theta_{\text{image}}}} \right) \mathbf{b}_i - \mathbf{b}_i \right| \cdot \alpha. \quad (24)$$

For example, if the UAV records a Full-HD video while hovering at 50 m,  $\mathbf{b}_i = [1920 \quad 1079]^T$  and  $\eta_{\xi_{\theta_{\text{image}}}} = 116^\circ$ , then  $\eta_{b_i} \approx 124.8$  m. Else, if  $\eta_{\xi_{\theta_{\text{image}}}} = 0.07^\circ$ , then  $\eta_{b_i} \approx 0.08$  m.

The error propagation causes a deviation on the linear offsets as well. The linear offset error  $\eta_{\xi_d}$  due to the orientation and scaling errors is expressed by

$$\eta_{\xi_d} = \left( \left( \mathbf{R} \left( \eta_{\xi_{\theta_{\text{image}}}} \right) \right)^T \left( \mathbf{g}_{i,\text{PCF}} \cdot \eta_\alpha \right) - \mathbf{g}_{i,\text{PCF}} \right) \cdot \alpha. \quad (25)$$

In a scenario where the UAV records a Full-HD video while hovering at 50 m,  $\mathbf{g}_{i,\text{PCF}} = [1920 \quad 1079]^T$ , a similarity of 26% and orientation error of  $116^\circ$ , then  $\eta_{\xi_d} \approx [-2394 \text{ m} \quad -759 \text{ m}]^T$ . Otherwise, with a similarity of 99.87% and orientation error of  $0.07^\circ$ , then  $\eta_{\xi_d} \approx [-3.83 \text{ m} \quad -0.96 \text{ m}]^T$ .



Fig. 3. A vehicle with its bounding box.

From the previous, it is deduced that the best way to minimize errors caused by the PCF to LTP mapping, is to locate the GCPs as far from each other as possible. Also, since the direction of the pixel ambiguity is not stochastic, the errors can be compensated by using different combinations of various GCPs.

### B. UAV Stability

UAVs usually drift while hovering over the LTP. This drift introduces additional scaling, rotation and linear offsets. However, this can be corrected by means of image registration. The registration generally consists of three steps to find correspondences between two images: a feature detector, a descriptor and the matching. The goal of the detector is to find identical interest points under varying viewing conditions. The descriptor is a feature vector, which describes the local area around the point of interest. The SURF [23] detector and descriptor is used in this work. The distance between the feature vectors is computed to match the points between two images (matching). An interest point is found on two images when the distance fulfills a certain criterion, e. g. a nearest neighbor ratio matching strategy. The matches are then fed into the MLESAC algorithm [24] to eliminate outliers, i. e. incorrect pairs. From the remaining matching pairs, the scaling, rotation and linear offsets are calculated and applied to the video sequence to eliminate the effects of the drift of the UAV.

### C. Car Labelling and Detection

Another source of error that causes inaccuracies is the labelling and later detection of the vehicle. As stated above, Mask-RCNN is used for detecting vehicles on video sequences. This is done through a process in two steps: the labelling of vehicles for the training of the network, and the use of the trained network to detect vehicles on new video material. To make the detection more robust, neither tires nor mirrors are labeled as part of the vehicle; i. e., only the chassis is labeled as vehicle.

Similar as with the GCPs, often the pixels cannot be unambiguously associated with the car. So, a labelling/detection error of 1 px or more is common. The effect of this error is propagated to the bounding box. An example of this can be seen on Figure 3.

#### D. Bounding Box

As explained in Section III-C, the measurement vector for the KF is generated from the vehicle bounding box. Since the bounding box is created from the detected shape of the vehicle, this detection is the first source of inaccuracy. It is then realistic to have a pixel ambiguity of  $\zeta = 1$  px for each side of the bounding box.

The next bounding box-related source of error is the perspective. Figure 1 shows an ideal case for the creation of the bounding box, where the complete shape of the vehicle is orthonormal to the UAV, and so can be mapped to the LTP as shown in the Section III-B. However, this is not possible because of relief displacement. This optical effect occurs due to camera perspective and the height of the seen objects. It has two consequences for the filmed objects: First, their position shifts away from the principal point (image centre) as their height increases. Second, the dimensions of their bounding box grow if a vehicle side is occluded. A graphical representation of this is shown on Figure 4.

If the height of the object is known, the relief displacement can be corrected. However, because of the combination of camera perspectives, vehicle shapes and poses; the sides of the bounding box could be projections of points with different heights. The corners would then be intersections of projections at different heights. So, the height in LTP of the corners of the bounding box is not known. Even when manually analysing the video material, it is common that pixels cannot be conclusively associated to specific vehicle parts. An example of this is shown in Figure 3.

A statement about the height can be done only for the corner that is closest to the vertical of the drone. Given that only the body is labelled as a vehicle (Section IV-C), it can be deduced that the two sides of the bounding box that are closest to the middle of the picture are tangent to the bottom of the car. So, the intersection of these lines is at this height as well. Knowing the height of this corner, its relief displacement can be corrected as follows.

Defining the horizontal and vertical resolution of the image as  $r_x$  and  $r_y$  respectively, the coordinates in PCF of  $\mathbf{b}_i$  with respect to the middle of the picture are given by

$$\begin{bmatrix} x_{i,\text{seen}} \\ y_{i,\text{seen}} \end{bmatrix} = \begin{bmatrix} x_{i,\text{PCF}} - \frac{r_x}{2} \\ y_{i,\text{PCF}} - \frac{r_y}{2} \end{bmatrix}. \quad (26)$$

The shift correction  $\Delta_{x,\text{PCF}}$  along the  $x_{\text{PCF}}$  axis is calculated on the PCF as follows

$$\Delta_{x,\text{PCF}} = \frac{x_{i,\text{seen}} \cdot h_{i,\text{LTP}}}{h_{\text{UAV}}}, \quad (27)$$

where  $h_{i,\text{LTP}}$  is the height of the  $i$ -th corner on LTP and  $h_{\text{UAV}}$  is the hovering altitude of the UAV in LTP. The correction  $\Delta_{y,\text{PCF}}$  along the  $y_{\text{PCF}}$  axis is calculated by analogy. The corrected coordinates  $\mathbf{b}_{i,c}$  of  $\mathbf{b}_i$  are then given by

$$\mathbf{b}_{i,c} = \mathbf{b}_i - [\Delta_{x,\text{PCF}} \quad \Delta_{y,\text{PCF}}]^T. \quad (28)$$

Since the detection error of  $\pm 1$  px per axis is transferred to the bounding box, then  $\mathbf{b}_i$  could have an error of  $\sqrt{2}$  px. If this error occur towards the principal point, then  $\Delta_{x,\text{PCF}}$  and

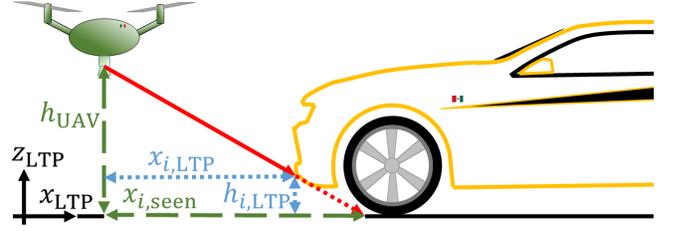


Fig. 4. Shown is: the line of sight of the UAV with red solid line, the seen vehicle position with green dashed line and the true vehicle position with blue dotted line.

$\Delta_{y,\text{PCF}}$  add an error towards the principal point as well. The error  $\eta_{b,i}$  of  $\mathbf{b}_i$  due to the shape detection and correction of the relief displacement is given by

$$\eta_{b,i} = \left( \sqrt{2} + \frac{h_{i,\text{LTP}}}{h_{\text{UAV}}} \right) \cdot \alpha. \quad (29)$$

As an example for a sedan with  $\pm 15$  cm ground clearance and a UAV hovering at 100 m, then  $\eta_{b,i} \approx 0.04$  m.

The corrected bounding box is reconstructed to find the centre of the vehicle. Let the width  $w_{\text{car}}$  and length  $l_{\text{car}}$  of the vehicle be known, the Equation (14) hold, the spatial resolution have no errors,  $\mathbf{b}_1$  be the corner closest to the image centre and its relief displacement be corrected. Then,  $\mathbf{b}_1$  is used as base for scaling  $\mathbf{b}_j$  and  $\mathbf{b}_k$  as follows

$$\mathbf{b}_{j,\text{scaled}} = \left( \frac{l_{\text{car}}}{s(2)} \cdot \alpha \cdot (\mathbf{b}_j - \mathbf{b}_1) \right) + \mathbf{b}_1, \quad \text{and} \quad (30)$$

$$\mathbf{b}_{k,\text{scaled}} = \left( \frac{w_{\text{car}}}{s(1)} \cdot \alpha \cdot (\mathbf{b}_k - \mathbf{b}_1) \right) + \mathbf{b}_1, \quad (31)$$

where  $k$  is the element of  $\mathbf{B}_{\text{PCF}}$  associated with  $s(1)$ . The corrected centre of the vehicle is then calculated by

$$\mathbf{o}_{\text{in}} = \frac{\mathbf{b}_{j,\text{scaled}} + \mathbf{b}_{k,\text{scaled}}}{2}. \quad (32)$$

As stated before, the height cannot be easily determined for the two sides of the bounding box farthest from the vertical of the drone. Considering passenger vehicles, this height can range from  $h_{i,\text{min}} = 0.11$  m (ground clearance of sport vehicles [25][26]) and up to  $h_{i,\text{max}} = 1.85$  m (roof of multi-purpose vehicles [27][28]). If the relief displacement is not corrected, the maximum positioning error  $\eta_{\text{scale}}$  due to wrong vehicle dimensions is given by

$$\eta_{\text{scale}} = \frac{\sqrt{x_{i,\text{seen}}^2 + y_{i,\text{seen}}^2} \cdot \alpha \cdot (h_{i,\text{max}} - h_{i,\text{min}})}{2 \cdot h_{\text{UAV}}}. \quad (33)$$

For example, if the UAV hovers at 50 m and  $\mathbf{b}_i$  is on one corner of the image, then  $\eta_{\text{scale}} \approx 0.63$  m.

From the discussion above it follows that the best way to minimize positioning errors due to relief displacement, is to correct it only for the corner nearest to the centre of the image, and to re-scale the bounding box. If the vehicle measures are not known, generic values can be taken from the literature [21].

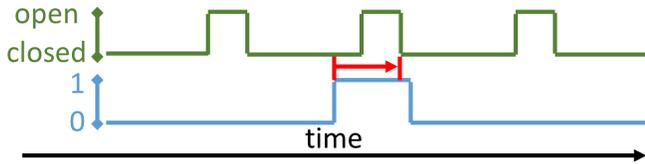


Fig. 5. Shown is: the PPS pulse from a SatNav receiver with blue, the shutter exposure time with green, and the synchronization error with red.

### E. Sensor Synchronization

Another relevant source of error is the synchronization of the sensors. A best practice approach is to use of the Pulse Per Second (PPS) from SatNav receivers and thus to employ the Universal Time Coordinated (UTC-time) as a common timeline. Atomic clocks aboard satellites [29] and negligible delays between receivers make this the best option. Unfortunately, most consumer-grade UAVs have no option for triggering the camera with a PPS pulse, nor do they provide a frame-accurate UTC timestamp.

In this work, the synchronization process is performed using as reference the light of the in-built LED of a Neo-6M [30] SatNav receiver as follows. The rising edge of the PPS pulse indicates the start of every second. This rising edge is used as a trigger for lighting up a LED that stays on for a determined period of time so that the light can be seen on the image to be processed. Since this LED is part of the SatNav module, the latency between the PPS reception and LED lighting up is negligible. So, the start of each second can be known with frame-accuracy by recording this LED. A graphical representation of this is shown in Figure 5. The limitation of this technique is that videos are not a continuous image, but a series of static pictures. So it is not known if the LED lights up when the shutter is closed, creating a synchronization error. The maximum synchronization error  $\eta_\tau$  is given by

$$\eta_\tau = \frac{1}{\tau_{FR}}, \quad (34)$$

where  $\tau_{FR}$  is the camera frame rate in frames per second (fps). In this work, 50 fps are used. So  $\eta_\tau \leq .02$  s. The positioning error  $\eta_{pos}$  caused by  $\eta_\tau$  is given by

$$\eta_{pos} = \sqrt{v_{x,car}^2 + v_{y,car}^2} \cdot \eta_\tau, \quad (35)$$

and the velocity error  $\eta_{vel}$  caused by  $\eta_\tau$  is given by

$$\eta_{vel} = \sqrt{a_{x,car}^2 + a_{y,car}^2} \cdot \eta_\tau. \quad (36)$$

As an example with a vehicle moving with 50 km/h, braking with 5 m/s<sup>2</sup> and a UAV hovering at 100 m recording a video with 50 fps, then  $\eta_{pos} \leq 0.25$  m and  $\eta_{vel} \leq .09$  m/s.

It can be deduced from what is discussed above that synchronization errors, even in the millisecond order, have a statistically significant influence. Also, if no PPS trigger is available, the best way to minimize errors due to synchronization is to film the on-board LED of SatNav receivers.

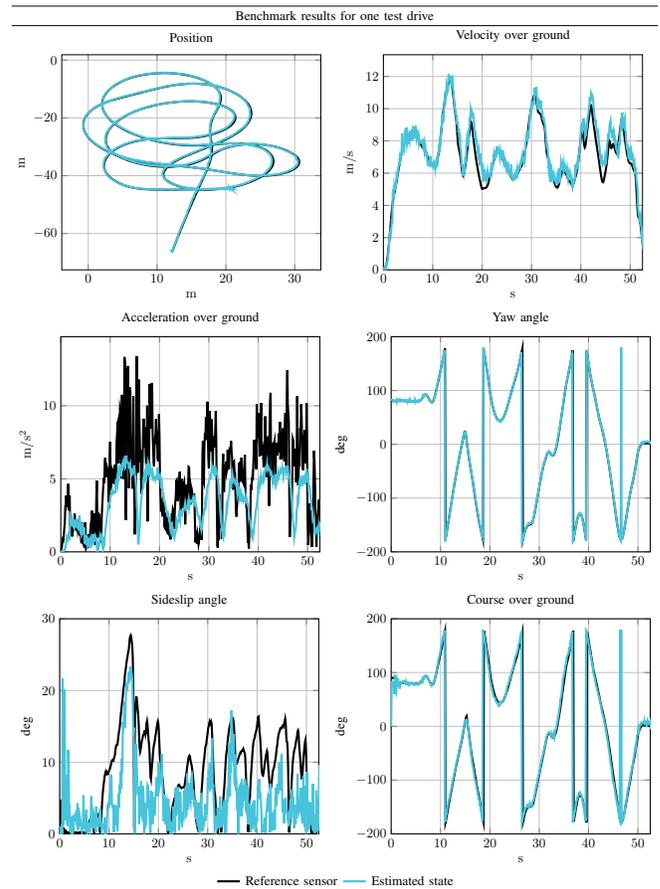


Fig. 6. Shown are the benchmark results for one test drive with the UAV hovering at 75 m.

## V. BENCHMARK OF THE ESTIMATED VEHICLE STATE

The benchmark process is detailed in the following. The estimated state vector is compared to the outputs of an Automotive Dynamic Motion Analyzer (ADMA) G-Pro+ [31]. This Inertial Navigation System (INS) is equipped with servo-accelerometers, optical gyroscopes and receives Real-Time Kinematic (RTK) correction data. The estimated sideslip angle is compared to the output of a Correvit S-Motion [32]. This sensor estimates velocities in the longitudinal and lateral axes by means of an optical grid [33]. The reference sideslip is computed analog to the Equation (18). Finally, the UAV used is the DJI Phantom 4 Pro V2 [34].

To generate the test data, a test vehicle is equipped with the reference sensors and is driven by a human in a random manner on a paved test track. No driving robot is used and no specific maneuver is driven in order to avoid tuning KF parameters to fit a specific trajectory, maneuver or driving style. The test drives include standstill, walking velocity, high acceleration, hard braking, tractive and non-tractive driving (drifting) [35]. The velocity is limited to 50 km h<sup>-1</sup>. No markers are placed on the vehicle to approach real-testing conditions on open roads. To give variety to the test dataset, the videos are taken at three hovering altitudes: 50 m, 75 m and 100 m. The videos are recorded with 50 fps and 4K resolution (3840 px x 2160 px), which are later compressed

to FullHD (1920 px x 1080 px). Four videos per height are recorded.

Figure 6 shows the benchmark results for one of the test drives with the biggest errors. For this trial run, the test vehicle is driven in such a manner to make it drift. This includes full-throttle acceleration, hard braking and sudden steering.

The estimated position, course over ground and yaw for this trial have a mean error of 19.16 cm, 4.71° and 1.02° respectively. This precision is equivalent to that of consumer-grade INSs.

Figure 6 shows that the estimated velocity is affected by a dampening effect and by a time delay. Both are caused by the Kalman gains. A test-specific tuning of the gains could help to reduce the velocity error for this trial, but would increase the error for other tests with lower vehicle dynamics.

The deviation of the estimated sideslip that is shown on Figure 6 is caused mainly by the velocity error. This is because the sideslip angle is estimated using the course over ground, which is calculated from the velocity (Section III-D). During this test, a sideslip angle of 27.69° is reached, which clearly indicates that the vehicle is drifting.

The error of the acceleration is explained by two facts: First, the estimated acceleration is calculated by system noise propagation (Section III-D), so it is low-pass filtered. Second, the reference acceleration that is measured by the INS includes vibrations from the drivetrain, the tyres and the suspension, as well as from the pitch and roll of the vehicle.

Figure 7 shows the cumulative frequency diagrams for all performed tests. The benchmark results show that, once steps are taken to minimize errors, the precision of the estimated vehicle state is comparable to the precision of consumer-grade sensors, such as silicon-based INSs or SatNav receivers with no correction data. This with the advantage of being able to record information for various traffic participants with a single UAV. Also, the precision of the estimated sideslip angle allows to differentiate between tractive and non-tractive driving. This is a relevant state variable to determine the vehicle stability.

## VI. CONCLUSIONS

In this work, an evaluation of the accuracy of camera-equipped UAVs as tools for traffic data acquisition is presented. This includes a detailed description of the vehicle state estimation from aerial imagery, an analysis and quantification of the most relevant sources of error, and a benchmark of the estimated vehicle state with state-of-the-art reference sensors. The findings can then serve as a best-practice guide to minimize errors in future works. The source code is publicly available to aid researchers adapt KF gains for specific tasks.

The results of the benchmark show that, once steps are taken to minimize errors, camera-equipped UAVs are very attractive tools for automotive data acquisition. This because relatively low effort is required, while still obtaining a vehicle state with a precision comparable to that of consumer-grade INSs. This with the added benefit of recording various vehicles with a single device.

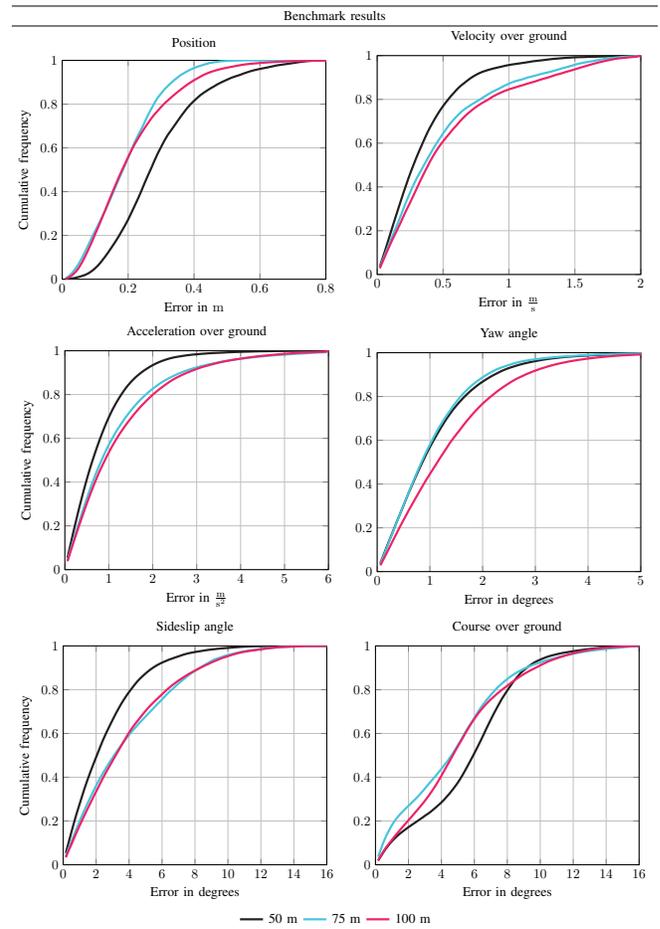


Fig. 7. Shown are the cumulative frequency diagrams of the benchmark results for three hovering altitudes.

Being the correction of the relief displacement the source of error with the most unknowns, it has the biggest improving potential. If the height of the vehicle pixels is known in a deterministic manner, the positioning error can be greatly reduced.

## VII. ACKNOWLEDGEMENT

The authors acknowledge the financial support by the Federal Ministry of Education and Research of Germany (BMBF) in the framework of FH-Impuls (project number 03FH7102IA).

## REFERENCES

- [1] D. Schramm, M. Hiller, R. Bardini *et al.*, *Modellbildung und Simulation der Dynamik von Kraftfahrzeugen*. Springer, 2010, vol. 124.
- [2] M. Abe, *Vehicle handling dynamics: theory and application*. Butterworth-Heinemann, 2015.
- [3] D. Bastow, G. Howard, and J. P. Whitehead, *Car suspension and handling*. SAE international Warrendale, 2004.
- [4] Data from sky. RCE systems s.r.o. [Online]. Available: <https://datafromsky.com/>
- [5] O. Gallitz, O. D. Candido, M. Botsch, and W. Utschick, "Interpretable feature generation using deep neural networks and its application to lane change detection," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, Oct 2019, pp. 3405–3411.
- [6] F. Kruber, J. Wurst, and M. Botsch, "An unsupervised random forest clustering technique for automatic traffic scenario categorization," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, Nov 2018, pp. 2811–2818.

- [7] H.-J. Song, Y. Chen, and Y.-Y. Gao, "Velocity calculation by automatic camera calibration based on homogenous fog weather condition," *International Journal of Automation and Computing*, vol. 10, 04 2013.
- [8] M. Kampelmühler, M. G. Müller, and C. Feichtenhofer, "Camera-based vehicle velocity estimation from monocular video," *CoRR*, vol. abs/1802.07094, 2018. [Online]. Available: <http://arxiv.org/abs/1802.07094>
- [9] R. Krajewski, J. Bock, L. Kloeker, and L. Eckstein, "The highd dataset: A drone dataset of naturalistic vehicle trajectories on german highways for validation of highly automated driving systems," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, 2018, pp. 2118–2125.
- [10] J. Bock, R. Krajewski, T. Moers, L. Vater, S. Runde, and L. Eckstein, "The ind dataset: A drone dataset of naturalistic vehicle trajectories at german intersections."
- [11] W. Zhan, L. Sun, D. Wang, H. Shi, A. Clausse, M. Naumann, J. Kümmerle, H. Königshof, C. Stiller, A. de La Fortelle, and M. Tomizuka, "INTERACTION Dataset: An INTERnational, Adversarial and Cooperative moTION Dataset in Interactive Driving Scenarios with Semantic Maps," *arXiv:1910.03088*, 2019.
- [12] G. Guido, V. Gallelli, D. Rogano, and A. Vitale, "Evaluating the accuracy of vehicle tracking data obtained from unmanned aerial vehicles," *International Journal of Transportation Science and Technology*, vol. 5, no. 3, pp. 136 – 151, 2016, unmanned Aerial Vehicles and Remote Sensing. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S2046043016300272>
- [13] V. Sunita Devi, "Measurement of relief displacement from vertical photograph." *International Journal of Science*.
- [14] Video vbox pro | 10 or 20hz. RACELOGIC. [Online]. Available: <https://www.vboxautomotive.co.uk/index.php/en/products/video-data-loggers/video-vbox-pro>
- [15] F. Kruber and E. Sánchez Morales, "Vehicle Detection and State Estimation with Aerial Imagery." [Online]. Available: <https://github.com/fkthi>
- [16] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [17] C. Kyrkou, G. Plastiras, T. Theodoridis, S. I. Venieris, and C. Bouganis, "DroNet: Efficient convolutional neural network detector for real-time UAV applications," in *2018 Design, Automation Test in Europe Conference Exhibition (DATE)*, 2018.
- [18] Q. Li, L. Mou, Q. Xu, Y. Zhang, and X. X. Zhu, "R<sup>3</sup>-net: A deep network for multi-oriented vehicle detection in aerial images and videos," *CoRR*, 2018. [Online]. Available: <http://arxiv.org/abs/1808.05560>
- [19] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," in *Proceedings of the International Conference on Computer Vision (ICCV)*, 2017.
- [20] Waleed Abdulla, "Mask R-CNN for object detection and instance segmentation on Keras and TensorFlow," 2017. [Online]. Available: [https://github.com/matterport/Mask\\_RCNN](https://github.com/matterport/Mask_RCNN)
- [21] F. Kruber, E. Sánchez Morales, S. Chakraborty, and M. Botsch, "Vehicle position estimation with imagery from unmanned aerial vehicles," in *2020 IEEE Intelligent Vehicles Symposium (IV)*, 2020.
- [22] R. Kalman, "A new approach to linear filtering and prediction problems," no. 1, pp. 1–12, 1960.
- [23] H. Bay, T. Tuytelaars, and L. Van Gool, "Surf: Speeded up robust features," in *Computer Vision – ECCV 2006*, A. Leonardis, H. Bischof, and A. Pinz, Eds.
- [24] P. Torr and A. Zisserman, "MLESAC: A New Robust Estimator with Application to Estimating Image Geometry," *Computer Vision and Image Understanding*, 2000.
- [25] V. der TÜV e.V., "Merkblatt 751," 2008.
- [26] J. M. BARROSO, *COMMISSION REGULATION (EU) No 678/2011*. Office for Official Publications of the European Communities, 14 July 2011.
- [27] K. Van Miert, *REGULATION (EEC) No 4064/89 MERGER PROCEDURE*. Office for Official Publications of the European Communities, 1999.
- [28] G. Heydinger, R. Bixel, W. R. Garrott, M. Pyne, J. G. Howe, and D. A. Guenther, *Measured Vehicle Inertial Parameters-NHTSAs Data Through November 1998*. Society of Automotive Engineers, 1999.
- [29] J. Nelson. What is an atomic clock? National Air Space Agency. [Online]. Available: <https://www.nasa.gov/feature/jpl/what-is-an-atomic-clock>
- [30] NEO-6, u-blox AG, 12 2011, rev. E.
- [31] Adma family gps/inertial system automotive/railway. GeneSys Elektronik GmbH. [Online]. Available: <https://www.genesys-offenburg.de/en/products/adma-family-gpsinertial-system-automotiverailway/>
- [32] Correvit s-motion datenblatt. Kistler. [Online]. Available: <https://www.kistler.com/de/produkt/type-2053--2055/>
- [33] J. Haus and N. Lauinger, "Optische gitter: Die abbildung der realität – 75 jahre berührungslose dynamische meßtechnik auf der basis optischer gitter," *Laser Technik Journal*, vol. 4, pp. 43–47, 04 2007.
- [34] Techn. daten. SZ DJI Technology Co., Ltd. [Online]. Available: <https://www.dji.com/de/phantom-4-pro-v2/specs>
- [35] M. Abdulrahim, "On the dynamics of automobile drifting," *SAE Mobilus*, 04 2006.