Efficient Real-Time Radial Distortion Correction for UAVs

Marcus Valtonen Örnhag[1], Patrik Persson[1], Mårten Wadenbäck[2], Kalle Åström[1], Anders Heyden[1]

[1]Centre for Mathematical Sciences
Lund University

[2]Department of Electrical Engineering
Linköping University

marcus.valtonen_ornhag@math.lth.se

Abstract

In this paper we present a novel algorithm for onboard radial distortion correction for unmanned aerial vehicles (UAVs) equipped with an inertial measurement unit (IMU), that runs in real-time. This approach makes calibration procedures redundant, thus allowing for exchange of optics extemporaneously. By utilizing the IMU data, the cameras can be aligned with the gravity direction. This allows us to work with fewer degrees of freedom, and opens up for further intrinsic calibration. We propose a fast and robust minimal solver for simultaneously estimating the focal length, radial distortion profile and motion parameters from homographies. The proposed solver is tested on both synthetic and real data, and perform better or on par with state-of-the-art methods relying on pre-calibration procedures. Code available at: `https://github.com/marcusvaltonen/HomLib`.[1]

1. Introduction

In epipolar geometry, the relative pose of two uncalibrated camera views is encoded algebraically as the fundamental matrix $\boldsymbol{F}$ concomitant with the two views. When trying to estimate $\boldsymbol{F}$ from point correspondences, it is well-known that the minimal case—*i.e.* the smallest number of point correspondences for which there exists at most finitely many solutions—uses seven point correspondences [19]. By using eight point correspondences instead of seven, the estimation problem results in a system of eight linear equations, which can be solved fast and in a numerically robust manner using the singular value decomposition (SVD) [18]. To solve the minimal case, *i.e.* using only seven point correspondences, one must conjoin the seven linear equations with one cubic equation emanating from the rank constraint $\det \boldsymbol{F} = 0$. In the case of calibrated cameras, the minimal case involves only five point correspondences; however, the corresponding system of polynomial equations now contains ten cubic equations, and the complexity of the solver increases further [40].

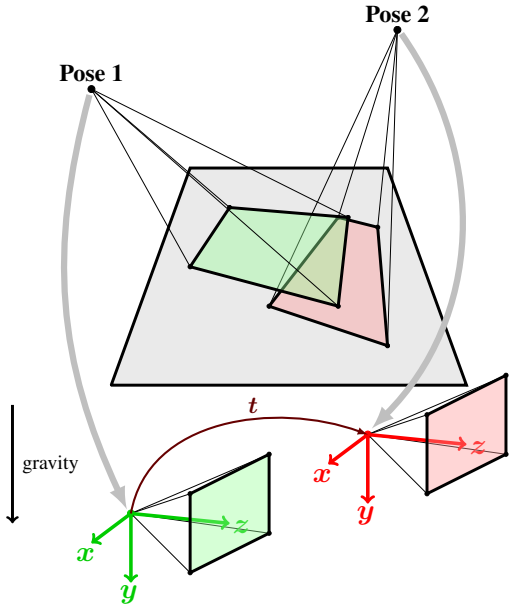

Figure 1: Utilizing the IMU data it is possible to align the camera views, only leaving an unknown translation $\boldsymbol{t}$. This assumes that the IMU drift is negligible, which is a realistic assumption if the measurements are not taken too far apart in time.

There are several benefits of reducing the number of point correspondences used to estimate the motion parameters. In most cases, this comes at the cost of increased complexity of the system to be solved. Solving systems of polynomial equations numerically, in a sufficiently fast and robust way, is a challenging task. One popular method, sometimes referred to as the *action matrix method*, works if there are finitely many solutions [10, 38]. The system of polynomial equations defines an ideal, for which a Gröbner basis can be computed, leading to an elimination template, where

[1] This work was supported by the Swedish Research Council (grant no. 2015-05639), the strategic research projects ELLIIT and eSSENCE, the Swedish Foundation for Strategic Research project, Semantic Mapping and Visual Navigation for Smart Robots (grant no. RIT15-0038), and Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by Knut and Alice Wallenberg Foundation.

the solutions to the original problem are obtained by solving an eigenvalue problem [23, 8, 6]. This process has been automated by several authors, with the automatic solver by Kukelova *et al*. [23] as one of the first. Recent advances use syzygies to make the elimination template smaller [30], as well as discarding spurious solutions by saturating the ideal [31]. Using Gröbner bases is not the only option; it has for example been shown that other bases can yield better performance [34], and that overcomplete spanning sets sometimes give better numerical stability [8]. Recently, alternative methods relying on resultants show promising results [2, 1].

Apart from adding extra constraints to the motion of the cameras, one may add scene requirements to reduce the number of necessary point correspondences. A classic example is when the point correspondences lie on a plane, in which case they are related through a homography. In applications where a planar environment is known to exist, such as indoor environments, the minimal number of point correspondences for the uncalibrated case is reduced from seven to four, and the corresponding system of equations—known as the Direct Linear Transform (DLT) equations—is linear in the entries of the homography, and can thus be solved using SVD.

If available, additional input data can be obtained from auxiliary sensors. In this paper we will consider UAVs equipped with an IMU, from which the gravity direction can be obtained. This, in turn, is assumed to be aligned with the ground plane normal.

Many commercially available UAVs are equipped with a camera that suffers from radial distortion to some degree. In order for the pinhole camera model to apply, such distortions must be compensated for, which is usually done in a pre-calibration process involving a calibration target. In contrast, we investigate a process for onboard radial distortion auto-calibration, *i.e*. a method capable of computing the radial distortion profile (and focal length) of the optics as well as the motion parameters, without a specific calibration target, thus eliminating the pre-calibration process. This enables the user of the UAV to exchange optics, without the need of intermediate calibration procedures, which may not be feasible without a calibration target. The main contributions of this paper are:

(i) a novel polynomial solver for simultaneous estimation of radial distortion profile, focal length and motion parameters, suitable for real-time applications,
(ii) new insights in how to handle IMU drift, and
(iii) extensive validation on synthetic and real data on a UAV system demonstrating the applicability of the proposed method.

While the algorithms proposed in this paper works for any application where a camera and IMU are available, we will use them exclusively for UAV positioning.

2. Related work

In most Simultaneous Localization and Mapping (SLAM) frameworks the distortion profile is pre-calibrated using a calibration target. This requires extra off-line processing, as well as scene requirements. For general scenes, there are a number of algorithms for simultaneously estimating the distortion profile and the motion parameters. Some authors propose methods based on large-scale optimization (bundle adjustment) [37], while others suggest using polynomial solvers [20, 7, 25, 29, 5, 24, 22, 26, 32, 41, 42, 33, 47]. A polynomial solver for the minimal case, *i.e*. the smallest number of point correspondences, is referred to as a minimal solver. There are several reasons to prefer minimal solvers, as they accurately encode intrinsic constraints, and transfer such properties to the final solution. Furthermore, they are suitable for robust estimation frameworks, *e.g*. RANSAC, as the number of necessary iterations (to obtain an inlier set with a pre-defined probability) is minimized.

There exists a number of different models for estimating the distortion profile. One classic approach, that is still frequently used in applications is the Brown–Conrady model [4]; however, although exceptions do exist [35], the division model by Fitzgibbon [12] is almost universally used in the construction of minimal solvers that deal with radial distortion. One reason for this is that the distortion profile can be accurately estimated using fewer parameters, which is consistent with the general theory behind minimal solvers. Other parametric models, recently *e.g*. [45], have been proposed, but are not suitable for minimal solvers for the same reason.

There are several methods that leverage the IMU data—or, simply, rely on the mechanical setup to be accurate enough—to assume a motion model with a known reference direction [13, 39, 21, 43, 46, 44, 15, 14, 11, 16, 48]. None of the mentioned papers, however, include simultaneous radial distortion correction, while only a handful consider the case of unknown focal length [14, 11, 16, 48]. To the best of our knowledge, we propose the first ever simultaneous distortion correction, focal length and motion estimation algorithm utilizing IMU data.

3. Embracing the IMU drift

Prevailing methods have been conceived under the assumption that only two angles can be compensated for using the IMU data, which is true under general conditions. The drift in the yaw angle, however, is often very small for consecutive frames. The idea is that we can disregard the error for the yaw angle initially, and instead correct for it later in the pipeline when enough time has passed for the drift to make a noticeable impact. This makes the equations significantly easier to handle, and allows for further intrin-

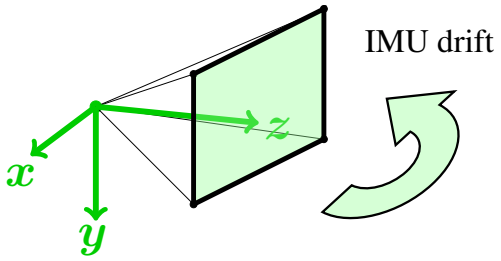

Figure 2: The pitch and roll angles can be accurately estimated using an IMU; however, the yaw angle (about the y-axis, also the gravity direction) often suffers from a drift that accumulates over time. By fusing angular velocity and accelerometer measurements, the drift is negligible for small time frames.

sic calibration, such as radial distortion correction.

Using [36], the orientation can be estimated both robustly and efficiently by fusing angular velocity and accelerometer measurements to form a single estimate of the orientation. The filter operates by integrating gyroscopic data and compensating for bias and integration errors by using the orientation that can be observed from the accelerometer. The rotation about the gravitational direction, however, is not observable and an inevitable drift will accumulate, see Figure 2. The drift is typically very small for short time frames since the sensor noise of the gyroscope is usually very low and the bias changes slowly.

3.1. New assumptions on the homography

Assume the reference direction is known, and aligned with the gravitational direction, chosen as the y-axis. Then, after a suitable change of coordinates, we may assume that

$$\boldsymbol{H}_y \sim \boldsymbol{I} + \frac{1}{d}\boldsymbol{t}\boldsymbol{n}^T, \tag{1}$$

where $\boldsymbol{I}$ is the identity matrix, $\boldsymbol{t}$ is the translation vector and $\boldsymbol{n}$ is the unit normal of the plane, see Figure 1. We will assume that the plane normal is aligned with the gravitational direction, which is a valid assumption when using the ground floor, thus $\boldsymbol{n} = [0,\ 1,\ 0]^T$. To ease notation, define

$$\boldsymbol{y}_i^{(j)} \coloneqq \boldsymbol{R}_j^T \boldsymbol{K}^{-1} \boldsymbol{x}_i, \tag{2}$$

where $\boldsymbol{R}_j$ is the rotation between the two coordinate systems (given by the IMU) and $\boldsymbol{K} = \mathrm{diag}(f,\ f,\ 1)$ is the calibration matrix, where f is the focal length, which is assumed to be constant. Then for two point correspondences $\boldsymbol{x}_1 \leftrightarrow \boldsymbol{x}_2$ the DLT equations can be written as

$$\boldsymbol{y}_2^{(2)} \times \boldsymbol{H}_y \boldsymbol{y}_1^{(1)} = \boldsymbol{0}\ . \tag{3}$$

The relation between the general (uncalibrated) homography $\boldsymbol{H}$ and $\boldsymbol{H}_y$ is thus given by

$$\boldsymbol{H}_y \sim \boldsymbol{R}_2^T \boldsymbol{K}^{-1} \boldsymbol{H} \boldsymbol{K} \boldsymbol{R}_1, \tag{4}$$

where $\boldsymbol{x}_2 \sim \boldsymbol{H}\boldsymbol{x}_1$. From this, the relative rotation $\boldsymbol{R}_{\text{rel}}$ and the direction of the relative translation $\boldsymbol{t}_{\text{rel}}$ can be extracted, and are given by

$$\boldsymbol{R}_{\text{rel}} = \boldsymbol{R}_2 \boldsymbol{R}_1^T \quad \text{and} \quad \boldsymbol{t}_{\text{rel}} \sim \boldsymbol{R}_2 \boldsymbol{t}\ . \tag{5}$$

Due to the global scale ambiguity, we may assume $d = 1$, and write

$$\boldsymbol{H}_y = \begin{bmatrix} 1 & h_1 & 0 \\ 0 & h_2 & 0 \\ 0 & h_3 & 1 \end{bmatrix}, \tag{6}$$

where $\boldsymbol{t}$ can be extracted directly through the entries h_i, given by

$$\boldsymbol{t} = \begin{bmatrix} h_1 \\ h_2 - 1 \\ h_3 \end{bmatrix}. \tag{7}$$

In order to apply the pinhole camera model, radially distorted feature points must be rectified. Assuming the distortion can be modeled by the division model [12], using only a single distortion parameter λ, the distorted (measured) image point $\boldsymbol{x}_i$ in camera i obeys the relationship

$$\boldsymbol{x}_i^u = \phi(\boldsymbol{x}_i, \lambda) = \begin{bmatrix} x_i \\ y_i \\ 1 + \lambda(x_i^2 + y_i^2) \end{bmatrix}, \tag{8}$$

where $\boldsymbol{x}_i = [x_i,\ y_i,\ 1]^T$, and $\boldsymbol{x}_i^u$ are the undistorted image points compatible with the pinhole camera model. Here we implicitly assume that the distortion center is at the center of the image. The modified DLT equations, can therefore be written as

$$\phi(\boldsymbol{x}_i, \lambda) \times \boldsymbol{H}\phi(\boldsymbol{x}_j, \lambda) = \boldsymbol{0}, \tag{9}$$

for two point correspondences $\boldsymbol{x}_i \leftrightarrow \boldsymbol{x}_j$.

3.2. Benefits of this approach

The homography described in Section 3.1 is greatly simplified compared to a general homography and has fewer parameters that need to be determined. In the case of unknown radial distortion profile, the competing methods [26, 12] return a general homography, *i.e.* with eight degrees of freedom. Unless one makes assumptions about the motion of the cameras—for example that it consists only of pure rotations—it is not possible to extract the motion parameters, even in the partially calibrated case. To see this, note that a Euclidean homography

$$\boldsymbol{H}_{\text{euc}} \sim \boldsymbol{R} + \boldsymbol{t}\boldsymbol{n}^T, \tag{10}$$

has eight degrees of freedom—three in $\boldsymbol{R}$, three in $\boldsymbol{t}$ and two in $\boldsymbol{n}$ (since the length of $\boldsymbol{n}$ is arbitrary). This is to be compared to a general homography that also has eight degrees of freedom. We conclude that a partially calibrated

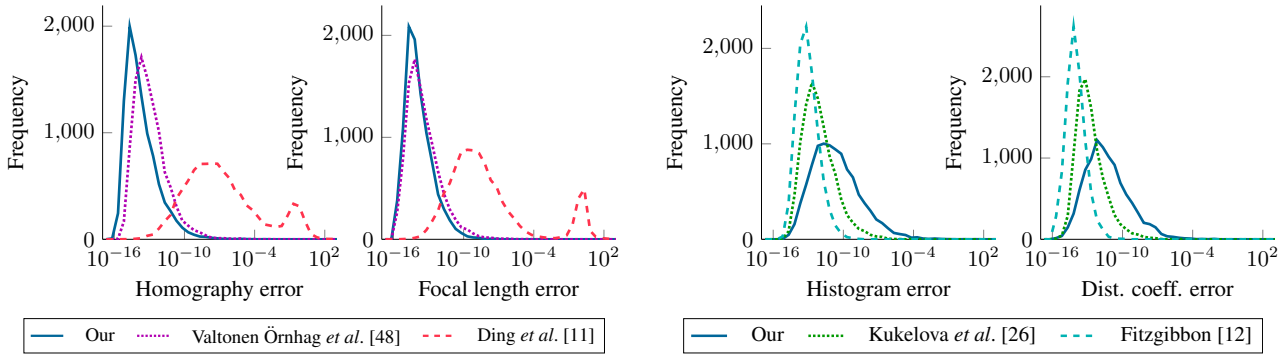

Figure 3: Error histogram for 10,000 randomly generated problem instances for the proposed solvers: (left) fHf and (right) $frHfr$. The solver [26] estimates two focal lengths, and we calculate the error for both and report the geometric mean. Most solvers have an acceptable error distribution, since in practice it rarely has an impact if the error is of the magnitude 10^{-14} or 10^{-10}.

homography on the form $\boldsymbol{K}\boldsymbol{H}_{\text{euc}}\boldsymbol{K}^{-1}$ must have nine degrees of freedom (the focal length f in $\boldsymbol{K}$ and the eight from $\boldsymbol{H}_{\text{euc}}$), hence is over-parametrized, *i.e*. there exists a one-dimensional family of possible decompositions. For this reason, we cannot extract the pose of the methods [26, 12], unless we assume that we know the focal length *a priori*, or constrain the motion. This, in itself, makes the methods infeasible to include in a SLAM framework, where we want to estimate the camera positions.

4. Polynomial solvers

In this section we present two-sided solvers, *i.e*. when the same intrinsic parameters (focal length and/or radial distortion parameter) are assumed for both cameras.

4.1. Calibrated case (1.5 point)

This case does not model an unknown focal length or distortion parameter, and is essentially the same approach as in [15], but is given here for completeness. Given 1.5 point correspondences it is possible to form the linear system $\boldsymbol{A}\boldsymbol{h} = \boldsymbol{b}$, where $\boldsymbol{A}$ is a 3×3 matrix and $\boldsymbol{h}$ contains the h_i from (6). For non-degenerate configurations, the matrix $\boldsymbol{A}$ has full rank, and the solution can be obtained immediately as $\boldsymbol{h} = \boldsymbol{A}^{-1}\boldsymbol{b}$. This is a very fast solver, since it is linear and can be solved without SVD.

4.2. Equal and unknown focal length (fHf, 2 point)

Parameterize the inverse of the unknown calibration matrix as $\boldsymbol{K}^{-1} = \text{diag}(1, 1, w)$, and consider the rectified points (2), which now depend linearly on the unknown parameter w. Parameterizing $\boldsymbol{H}_y$ as in (6), it is clear that the equations obtained from (3) are linear in h_1, h_2 and h_3 and quadratic in w. This system of equations has infinitely many solutions, if we allow $w = 0$. Such solutions, however, do not yield geometrically meaningful reconstructions, and should therefore be excluded. This can be achieved using saturation, through the method suggested in [31].

We exploit the linear relation of h_1, h_2, h_3, making it possible to write the equations as

$$\boldsymbol{M}\begin{bmatrix}\boldsymbol{h}\\ 1\end{bmatrix} = \boldsymbol{0}, \tag{11}$$

where $\boldsymbol{M}$ is a 4×4 matrix depending on w, and $\boldsymbol{h}$ is the vector containing the elements h_i. Thus, we may consider finding a non-trivial nullspace of $\boldsymbol{M}$, which exists if and only if $\det \boldsymbol{M} = 0$. This equation reduces to a sextic polynomial in the unknown w, thus has six solutions, which can be found using a simple root finding algorithm (action matrix method is not necessary). Since we know from before that the original problem has four solutions, we conclude that two spurious solutions have been added; in fact, these can easily be disregarded as a pre-processing step, as they correspond to nullspace basis vectors with last element equal to zero. Numerical tests confirm that this is the case.

When the (up to) four possible real solutions of w have been obtained, the unknowns h_1, h_2 and h_3 can be obtained using SVD.

4.3. Equal and unknown focal length and radial distortion coefficient($frHfr$, 2.5-point)

Let us now consider the case with equal and unknown focal length and radial distortion coefficient. We use the division model introduced in [12], with a single distortion parameter λ.

Given two point correspondences $\boldsymbol{x}_1 \leftrightarrow \boldsymbol{x}_2$, the modified DLT equations (9) hold true. Building an elimination template from these equations yields a large and numerically unstable solver, and therefore, we reparameterize the problem. Applying $\boldsymbol{H}_y = \boldsymbol{I} + \boldsymbol{t}\boldsymbol{n}^T$ to (4), we get

$$\boldsymbol{K}^{-1}\boldsymbol{H}\boldsymbol{K} \sim \boldsymbol{R}_2\boldsymbol{H}_y\boldsymbol{R}_1^T = \boldsymbol{R}_2\boldsymbol{R}_1^T + \boldsymbol{R}_2\boldsymbol{t}\boldsymbol{n}^T\boldsymbol{R}_1^T\ . \tag{12}$$

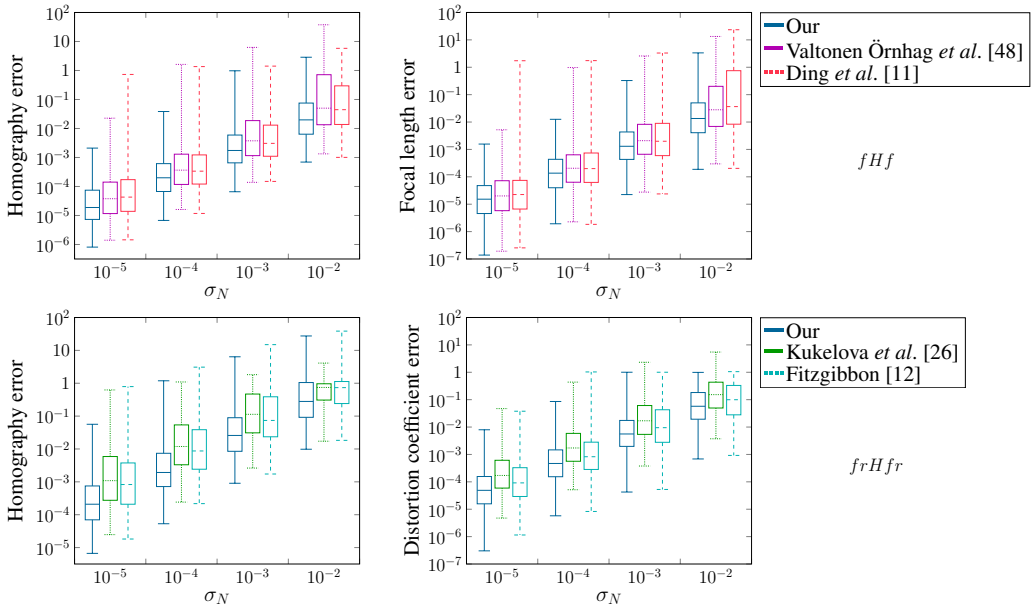

Figure 4: Noise sensitivity comparison for Gaussian noise with standard deviation σ_N. For each noise level 1,000 random problem instances were generated. The geometric mean error is shown for [26].

Denote the elements of $\boldsymbol{R}_1$ by r_{ij}. Introducing $\hat{\boldsymbol{R}} = \boldsymbol{R}_2\boldsymbol{R}_1^T$, $\hat{\boldsymbol{t}} = \boldsymbol{R}_2\boldsymbol{t}$ and $\hat{\boldsymbol{n}} = \boldsymbol{R}_1\boldsymbol{n} = [r_{12}, r_{22}, r_{32}]^T$, the general homography can now be written as

$$\boldsymbol{H} \sim \boldsymbol{K}\left(\hat{\boldsymbol{R}} + \hat{\boldsymbol{t}}\hat{\boldsymbol{n}}^T\right)\boldsymbol{K}^{-1}\,. \tag{13}$$

This accomplishes two things: (1) we have replaced several multiplications, (2) we have reduced the number of input data necessary. Analyzing the quotient ring of the corresponding ideal, we conclude that there are three possible solutions.

We parameterize the calibration matrix as $\boldsymbol{K} = \mathrm{diag}(f,\, f,\, 1)$ and its inverse $\boldsymbol{K}^{-1} = \mathrm{diag}(f^{-1},\, f^{-1},\, 1)$, respectively. From here on it would be possible to construct an elimination template; however, we may eliminate one variable in order to get a reduced system. Using only the third row of (9), from three point correspondences, one obtains a system on the form $\boldsymbol{M}\boldsymbol{v} = \boldsymbol{0}$, where $\boldsymbol{M}$ is a 3×9 coefficient matrix, and $\boldsymbol{v}$ is the vector of monomials, more precisely[2]

$$\boldsymbol{v} = \begin{bmatrix}\hat{t}_1 f\lambda & \hat{t}_1 f & \hat{t}_1 & \hat{t}_2 f\lambda & \hat{t}_2 f & \hat{t}_2 & f\lambda & f & 1\end{bmatrix}^T\,. \tag{14}$$

Since $\hat{t}_1$ and $\hat{t}_2$ are present in only three monomials, either of the two can be eliminated; we will proceed by eliminating the latter, as it yields a smaller elimination template. After Gauss–Jordan elimination, the coefficient matrix is given by

$$\hat{\boldsymbol{M}} = \begin{array}{c} \begin{matrix}\hat{t}_2 f\lambda & \hat{t}_2 f & \hat{t}_2 & \hat{t}_1 f\lambda & \hat{t}_1 f & \hat{t}_1 & f\lambda & f & 1\end{matrix} \\ \begin{bmatrix} 1 & & & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet \\ & 1 & & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet \\ & & 1 & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet \end{bmatrix} \end{array}, \tag{15}$$

from which we establish the following relations

$$\begin{aligned} \hat{t}_2 f\lambda + g_1(\hat{t}_1, f, \lambda) &= 0, \\ \hat{t}_2 f + g_2(\hat{t}_1, f, \lambda) &= 0, \\ \hat{t}_2 + g_3(\hat{t}_1, f, \lambda) &= 0, \end{aligned} \tag{16}$$

where g_i are polynomials of three variables. Furthermore, the following constraints must be fulfilled

$$\begin{aligned} g_1(\hat{t}_2, f, \lambda) - \lambda g_2(\hat{t}_2, f, \lambda) &= 0, \\ g_2(\hat{t}_2, f, \lambda) - f g_3(\hat{t}_2, f, \lambda) &= 0, \\ g_1(\hat{t}_2, f, \lambda) - \lambda f g_3(\hat{t}_2, f, \lambda) &= 0\,. \end{aligned} \tag{17}$$

We can now use the first row of (9), from which we get two equations (which must be multiplied by f to make it polynomial). Together with (17) we have five equations in four unknowns.

To build a solver we saturate f, to remove spurious solutions corresponding to zero focal length. Analyzing the quotient ring we have again three solutions and the basis heuristic [34] yields a template size of 26×29. Using the hidden variable trick, as in Section 4.2, we were able to construct a solver with a template size of 17×20; however, this solver was not as numerically stable as the one proposed, nor faster.

[2] It turns out that the third row does not contain any reciprocal f.

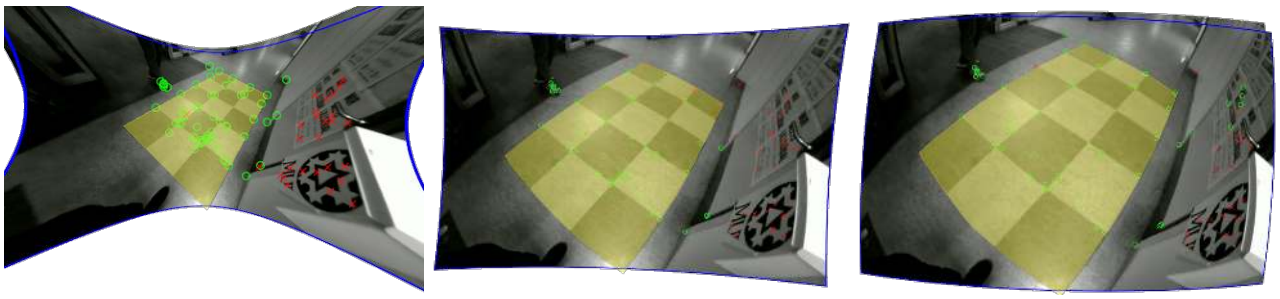

(a) Our, 70 % inliers. (b) Kukelova *et al*. [26], 62 % inliers. (c) Fitzgibbon [12], 67 % inliers.

Figure 5: Selection of panoramas created with the competing methods. The blue frame is added for visualization, as well as inliers (green circles) and outliers (red crosses). Note that none of the methods require a checkerboard to be visible in the scene, but is simply chosen to ease the ocular inspection of the stitching. A correct rectification will map physically straight lines to straight lines, *i.e*. the yellow area should be a quadrilateral. Only our method is capable of producing this result.

5. Experiments

5.1. Numerical stability and noise sensitivity

We compare the proposed methods with other state-of-the-art methods on synthetic data, to evaluate the numerical stability. For the case of unknown focal length we compare to the 2.5 point method [48] and the 3.5 pt method [11], and in the case of unknown radial distortion we compare to the 5 point methods [26, 12]. We generate noise free problem instances, by generating homographies and rotation matrices, and project a random set of points to establish point correspondences. In the case of radial distortion, these points are distorted using the division model. The error histograms are shown in see Figure 3. In the case of unknown focal length our method is superior to the others; however, with unknown radial distortion, we are not as stable as others. The accuracy, however, is in the order of 10^{-10}. This is sufficient for most applications. We will see in future experiments, that this does not cause any practical issues. The homography error is measured as the difference between the estimated homography and the ground truth in the Frobenius norm, normalized with the Frobenius norm of the ground truth homography, where the homographies are chosen such that $h_{33} = 1$. The errors of the focal length and radial distortion coefficient are measured as the absolute difference divided by the ground truth value.

Lastly, Gaussian noise is added to the image correspondences, in order to compare the noise sensitivity of the methods. The standard deviation σ_N is varied for a number of different noise levels. For all noise levels, our solvers perform superior to the other methods, for both the case with and without radial distortion, see Figure 4.

5.2. Speed evaluation

Next we compare the execution time for the considered methods. We compare the mean execution time given a minimal set of point correspondences until the putative homographies, and other parameters are obtained, *i.e*. including all pre-processing and post-processing steps. Furthermore, for the 2.5 and 3.5 point methods, we discard false solutions using the previously unused DLT equation.

As we are interested in performing the computations onboard the UAV, we evaluate the performance on a Raspberry Pi 4, and the mean execution times are listed in Table 1. All solvers are implemented in C++ using Eigen [17] and compiled in `g++` with the `-O2` optimization flag. Lastly, we list the maximal number of iterations possible on a 30 fps system, which we will use in Section 5.3 to compare real-time performance.

Table 1: Mean execution time on a Raspberry Pi 4 for 100,000 randomly generated problem instances in C++. The last column is the maximal number of iterations possible when running 30 fps.

Author	Time (μs)	No. iter.
Our fHf	215	155
Our $frHfr$	149	223
Valtonen Örnhag *et al*. [48]	80	416
Ding *et al*. [11]	3301	10
Kukelova *et al*. [27]	371	89
Fitzgibbon [12]	428	77
Kukelova *et al*. [26]	226	147

5.3. Real data

In this section, we compare the proposed methods on real data. We use the datasets from [48], captured using a UAV with a monochrome global shutter camera (OV9281) with resolution 640×480. The UAV is equipped with an inertial measurement unit (MPU-9250). In the experiments with only unknown focal length, the extracted features where undistorted using a pre-calibrated distortion profile (using the OpenCV [3] camera calibration procedure); for the case

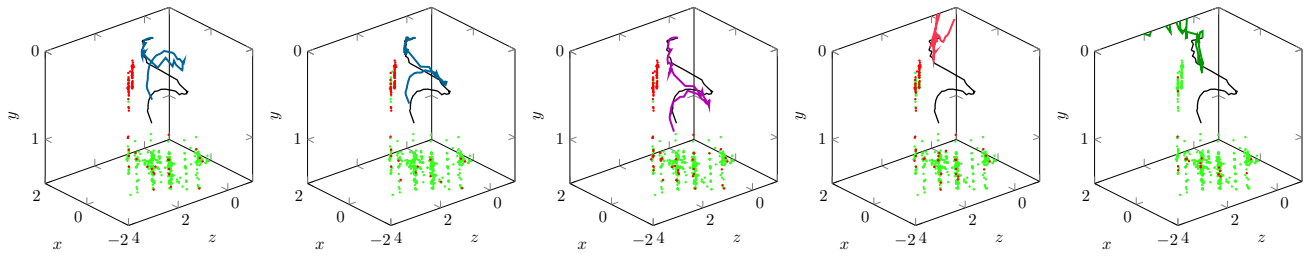

Figure 6: Estimated trajectories from the *Indoor* dataset. From left to right: Our fHf, our $frHfr$, Valtonen Örnhag *et al*. [48], Ding *et al*. [11] and Kukelova *et al*. [27]. The green dots indicate points that have been selected as inliers at least once, and the red points those which have been consistently rejected. Rectified images were used as input to all solvers except for $frHfr$, which received the raw input images.

with unknown radial distortion profile, the raw unprocessed coordinates were used as input.

The ground truth was obtained using a complete SLAM system where the reprojection error and IMU error were minimized. No scene requirements are enforced by the system, hence feature points from non-planar structures will be present—such feature points should be discarded by a robust framework as outliers.

The dataset consists of both indoor and outdoor sequences containing planar surfaces, and includes varying motions and length of sequences. Example images from the sequences are shown in Figure 7.

We exclusively use consecutive keyframes for extracting and matching features. The following keyframe insertion heuristic is used: (i) If the frame has moved more than 2.5% of its median depth from the previous keyframe, then proceed to next step, (ii) If we are tracking 10 % fewer 3D points in the current frame compared to the closest keyframe, or if the uncertainty of the projection of a 3D point is higher than a predefined threshold, then increase new information counter, otherwise, reset counter. (iii) If the new information counter is above 3 and the current frame tracks at least 10 3D points, add frame as keyframe.

We use the IMU filter technique [36], described in Section 3, to obtain the estimates. Since these measurements are noisy, we propose to use the novel solvers in a LO-RANSAC framework [9]. As a first step the solvers are used to discard outliers and in the inner LO loop we propose to optimize over the space of Euclidean homographies with unknown focal length. This refinement step allows for correction of the errors initially caused by the IMU filter. Empirically, we have seen that this improves the accuracy.

Figure 7: Example images from the dataset [48].

For a fair comparison, we simulate a scenario where the UAV uses a frame rate of 30 fps, and limit the number of RANSAC cycles to fit this time frame. For simplicity, we use the values from Table 1 alone, acknowledging there are other parts in the pipeline—such as image capturing, feature extraction and matching, LO cycles, etc.—that would affect a complete system; however, we argue that this overhead time is roughly independent of the solver used.

When working with radial distortion correction, one may choose to minimize the reprojection error in the undistorted image space, or in the distorted image space. In [28], it was shown that it is beneficial to perform triangulation in the distorted image space. Therefore, we chose to measure the reprojection in the distorted space.

5.3.1 Image stitching

As argued in Section 3.2 we cannot decompose the homographies obtained from [26, 12], into a relative pose; however, we can still test the ability of the methods to return an accurate distortion profile.

We simulate a scenario where a UAV is navigating in 30 fps, and limit the number of iterations for each method according to Table 1. We use the same pixel threshold for all methods, for two consecutive keyframes of the *Indoor* sequence. We chose this sequence, because it naturally contains a checkerboard pattern, which facilitates in making an accurate ocular evaluation of the quality of the estimated distortion profile. Physically straight lines should be mapped to straight lines if the rectification is successful.

In Figure 5 we show the results of the estimated distortion profile. We notice that the distortion profile is correct for the proposed method as the yellow area is a quadrilat-

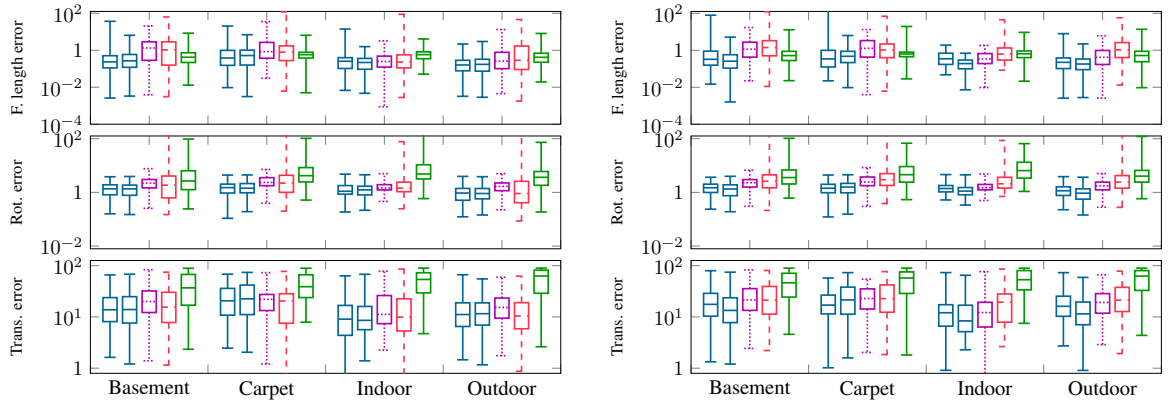

Figure 8: Errors for the different methods—from left to right: fHf, $frHfr$, Valtonen Örnhag *et al*. [48], Ding *et al*. [11], Kukelova *et al*. [27]—using the metrics (18). (Left) rectified input images were used for all but the $frHfr$. (Right) unrectified images were used for all methods.

eral, whereas this is not the case for other methods. Furthermore, we note that the method by Kukelova *et al*. [26] does not contain all inliers of the ground plane, and that the method by Fitzgibbon [12] pick incorrect matches of the wall.

For the same pair of images we measure the inliers as a function of time, see Figure 9. The only method converging to the correct number of inliers in the allotted time is our method, which it does by a large margin.

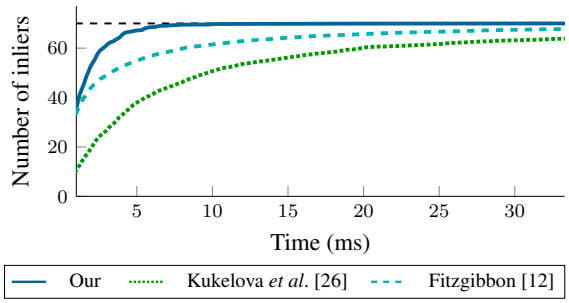

Figure 9: Number of inliers over time (average of 100 tests). The proposed method is the only method to converge to the correct number of inliers (black dashed line) within the 33.3 ms allotted per frame when running at 30 fps.

5.3.2 Pose estimation

In this final section, we compare both proposed methods to [27, 11, 48]. Note that only one of these methods (our $frHfr$) estimates the radial distortion profile. As argued in Section 3.2, the methods [26, 12] cannot estimate the motion parameters without additional requirements that are not applicable for UAVs, hence cannot be compared in this section.

The error metrics are defined as in [44, 11, 48], and are given by

$$\begin{aligned} e_{\boldsymbol{R}} &= \arccos\left(\frac{\mathrm{tr}(\boldsymbol{R}_{\mathrm{GT}}\boldsymbol{R}_{\mathrm{est}}^T)-1}{2}\right), \\ e_{\boldsymbol{t}} &= \arccos\left(\frac{\boldsymbol{t}_{\mathrm{GT}}^T\boldsymbol{t}_{\mathrm{est}}}{\|\boldsymbol{t}_{\mathrm{GT}}\|\|\boldsymbol{t}_{\mathrm{est}}\|}\right), \\ e_f &= \frac{|f_{\mathrm{GT}}-f_{\mathrm{est}}|}{f_{\mathrm{GT}}}. \end{aligned} \tag{18}$$

In Figure 6 we compare the estimated trajectories for all methods. It can be seen that there are only small differences between the methods using pre-calibrated radial distortion profile, and the proposed $frHfr$ solver. Furthermore, we measure the errors, according to (18) for all four sequences. In the left part of Figure 8 we use rectified images for all except the $frHfr$ method, which still performs best or on par with the other methods in terms of all errors. In the right part of the figure, we run the same experiment, but all methods are given the raw (unrectified) images as input—here it is clear that out method achieves superior results.

6. Conclusions

We have presented the first ever method capable of simultaneously estimating the distortion profile, focal length and motion parameters from a pair of homographies, while incorporating the IMU data. The method relies on a novel assumption that the IMU data is accurate enough, to disregard the IMU drift for small time frames, allowing for simpler equations and faster solvers. We have shown that this assumption is true on both synthetic and real data, and that the proposed methods are robust. The method has been shown to give accurate reconstructions, and performs on par or better than state-of-the-art methods relying on pre-calibration procedures, while being fast enough for real-time applications.

References

[1] Snehal Bhayani, Zuzana Kukelova, and Janne Heikkilä. Computing stable resultant-based minimal solvers by hiding a variable, 2020. arXiv preprint arXiv:2007.10100.

[2] Snehal Bhayani, Zuzana Kukelova, and Janne Heikkilä. A sparse resultant based method for efficient minimal solvers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.

[3] Gary Bradski. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*, 2000.

[4] Duance C. Brown. Decentering distortion of lenses. *Photogrammetric Engineering*, 32:444–462, 03 1966.

[5] Martin Bujnak, Zuzana Kukelova, and Tomas Pajdla. New efficient solution to the absolute pose problem for camera with unknown focal length and radial distortion. In *Asian Conference on Computer Vision (ACCV)*, 11 2011.

[6] Martin Bujnak, Zuzana Kukelova, and Tomas Pajdla. Making minimal solvers fast. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 06 2012.

[7] Martin Byröd, Matthew Brown, and alle. Åstrom. Minimal Solutions for Panoramic Stitching with Radial Distortion. In *Proceedings of the British Machine Vision Conference (BMVC)*, London, 2009.

[8] Martin Byröd, Klas Josephson, and Kalle Åström. Fast and Stable Polynomial Equation Solving and Its Application to Computer Vision. *International Journal of Computer Vision*, 84(3):237–256, 9 2009.

[9] Ondřej Chum, Jiří Matas, and Josef Kittler. Locally optimized ransac. In *Pattern Recognition*, pages 236–243. Springer Berlin Heidelberg, 2003.

[10] David A. Cox, John Little, and Donald O'Shea. *Using Algebraic Geometry*. Graduate Texts in Mathematics. Springer New York, 2005.

[11] Yaqing Ding, Jian Yang, Jean Ponce, and Hui Kong. An efficient solution to the homography-based relative pose problem with a common reference direction. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.

[12] Andrew W. Fitzgibbon. Simultaneous linear estimation of multiple view geometry and lens distortion. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, Dec 2001.

[13] Friedrich Fraundorfer, Petri Tanskanen, and Marc Pollefeys. A minimal case solution to the calibrated relative pose problem for the case of two known orientation angles. In *European Conference on Computer Vision (ECCV)*, pages 269–282, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.

[14] Banglei Guan, Yu Qifeng, and Friedrich Fraundorfer. Minimal solutions for the rotational alignment of imu-camera systems using homography constraints. *Computer vision and image understanding*, pages 79 – 91, 2018.

[15] Banglei Guan, Pascal Vasseur, Cédric Demonceaux, and Friedrich Fraundorfer. Visual odometry using a homography formulation with decoupled rotation and translation estimation using minimal solutions. In *International Conference on Robotics and Automation (ICRA)*, pages 2320–2327, 2018.

[16] Banglei Guan, Ji Zhao, Zhang Li, Fang Sun, and Friedrich Fraundorfer. Minimal solutions for relative pose with a single affine correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.

[17] Gaël Guennebaud, Benoît Jacob, et al. Eigen v3. http://eigen.tuxfamily.org, 2010.

[18] Richard I. Hartley. In Defense of the Eight-Point Algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(6):580–593, 6 1997.

[19] Richard I. Hartley and Andrew Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, Cambridge, England, UK, second edition, 2004.

[20] Klas Josephson and Martin Byröd. Pose Estimation with Radial Distortion and Unknown Focal Length. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2419–2426, Miami, FL, USA, 6 2009. IEEE Computer Society.

[21] Mahzad Kalantari, Amir Hashemi, Franck Jung, and Jean-Pierre Guedon. A New Solution to the Relative Orientation Problem Using Only 3 Points and the Vertical Direction. *Journal of Mathematical Imaging and Vision*, 39:259–268, 3 2011.

[22] Yubin Kuang, Jan Erik Solem, Fredrik Kahl, and Karl Åström. Minimal solvers for relative pose with a single unknown radial distortion. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 33–40. IEEE - Institute of Electrical and Electronics Engineers Inc., 2014.

[23] Zuzana Kukelova, Martin Bujnak, and Tomas Pajdla. Automatic generator of minimal problem solvers. *European Conference on Computer Vision (ECCV)*, pages 302–315, October 2008.

[24] Zuzana Kukelova, Martin Bujnak, and Tomas Pajdla. Real-time solution to the absolute pose problem with unknown radial distortion and focal length. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, December 2013.

[25] Zuzana Kukelova, Martin Byröd, Klas Josephson, Tomas Pajdla, and Karl Åström. Fast and robust numerical solutions to minimal problems for cameras with radial distortion. *Computer Vision and Image Understanding*, 114(2):234–244, 2010.

[26] Zuzana Kukelova, Jan Heller, Martin Bujnak, and Tomas Pajdla. Radial distortion homography. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 639–647, June 2015.

[27] Zuzana Kukelova, Joe Kileel, Bernd Sturmfels, and Tomás Pajdla. A clever elimination strategy for efficient minimal solvers. *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3605–3614, 2017.

[28] Zuzana Kukelova and Viktor Larsson. Radial distortion triangulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

[29] Zuzana Kukelova and Tomas Pajdla. A minimal solution to radial distortion autocalibration. *IEEE Transactions on Pat-*

tern Analysis and Machine Intelligence, 33(12):2410–2422, 2011.

[30] Viktor Larsson and Kalle Åström. Uncovering symmetries in polynomial systems. *European Conference on Computer Vision (ECCV)*, pages 252–267, October 2016.

[31] Viktor Larsson, Kalle Åström, and Magnus Oskarsson. Polynomial solvers for saturated ideals. *International Conference on Computer Vision (ICCV)*, pages 2307–2316, October 2017.

[32] Viktor Larsson, Zuzana Kukelova, and Yinqiang Zheng. Making minimal solvers for ab-solute pose estimation compact and robust. *International Conference on Computer Vision (ICCV)*, pages 2335–2343, October 2017.

[33] Viktor Larsson, Zuzana Kukelova, and Yinqiang Zheng. Camera pose estimation with unknown principal point. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.

[34] Viktor Larsson, Magnus Oskarsson, Kalle Åström, Alge Wallis, Zuzana Kukelova, and Tomas Pajdla. Beyond gröbner bases: Basis selection for minimal solvers. *Computer Vision and Pattern Recognition (CVPR)*, pages 3945–3954, 2018.

[35] Viktor Larsson, Torsten Sattler, Zuzana Kukelova, and Marc Pollefeys. Revisiting radial distortion absolute pose. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.

[36] Sebastian O. H. Madgwick, Andrew J. L. Harrison, and Ravi Vaidyanathan. Estimation of imu and marg orientation using a gradient descent algorithm. In *2011 IEEE International Conference on Rehabilitation Robotics*, pages 1–7, 2011.

[37] Holger Meuel, Stephan Ferenz, Marco Munderloh, Hanno Ackermann, and Jörn Ostermann. In-loop radial distortion compensation for long-term mosaicking of aerial videos. *Proc. of the 23rd IEEE International Conference on Image Processing (ICIP)*, pages 2961–2965, Sept. 2016.

[38] H. Michael Möller and Hans J. Stetter. Multivariate polynomial equations with multiple zeros solved by matrix eigenproblems. *Numerische Mathematik*, 70(3):311–329, 1995.

[39] Oleg Naroditsky, Xun Zhou, Jean Gallier, Stergios Roumeliotis, and Kostas Daniilidis. Two efficient solutions for visual odometry using directional correspondence. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34:818–24, 11 2011.

[40] David Nistér. An Efficient Solution to the Five-Point Relative Pose Problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(6):756–770, 6 2004.

[41] James Pritts, Zuzana Kukelova, Viktor Larsson, and Ondrej Chum. Radially-distorted conjugate translations. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[42] James Pritts, Zuzana Kukelova, Viktor Larsson, and Ondřej Chum. Rectification from radially-distorted scales. In *Asian Conference of Computer Vision (ACCV)*, pages 36–52, 2018.

[43] Olivier Saurer, Friedrich Fraundorfer, and Marc Pollefeys. Homography based visual odometry with known vertical direction and weak Manhattan world assumption. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Vilamoura, Portugal, 2012.

[44] Olivier Saurer, Pascal Vasseur, Rémi Boutteau, Cédric Demonceaux, Marc Pollefeys, and Friedrich Fraundorfer. Homography based egomotion estimation with a common direction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(2):327–341, Feb 2017.

[45] Thomas Schops, Viktor Larsson, Marc Pollefeys, and Torsten Sattler. Why having 10,000 parameters in your camera model is better than twelve. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.

[46] Chris Sweeney, John Flynn, and Matthew Turk. Solving for relative pose with a partially known rotation is a quadratic eigenvalue problem. In *2014 2nd International Conference on 3D Vision*, volume 1, pages 483–490, 2014.

[47] Marcus Valtonen Örnhag. Radially distorted planar motion compatible homographies. In *International Conference on Pattern Recognition Applications and Methods (ICPRAM)*, pages 280–288, Valletta, Malta, February 2020.

[48] Marcus Valtonen Örnhag, Patrik Persson, Mårten Wadenbäck, Kalle Åström, and Anders Heyden. Minimal solvers for indoor uav positioning, 2020. arXiv preprint arXiv:2003.07111.