

ACCURATE OPTICAL TARGET POSE DETERMINATION FOR APPLICATIONS IN AERIAL PHOTOGRAMMETRY

D. A. Cucci

École Polytechnique Fédérale de Lausanne (EPFL), Switzerland - (davide.cucci)@epfl.ch

Commission III, WG III/3

KEY WORDS: Optical Targets, Targets Orientation, Image Processing Algorithms, Navigation, Geometry

ABSTRACT:

We propose a new design for an optical coded target based on concentric circles and a position and orientation determination algorithm optimized for high distances compared to the target size. If two ellipses are fitted on the edge pixels corresponding to the outer and inner circles, quasi-analytical methods are known to obtain the coordinates of the projection of the circles center. We show the limits of these methods for quasi-frontal target orientations and in presence of noise and we propose an iterative refinement algorithm based on a geometric invariant. Next, we introduce a closed form, computationally inexpensive, solution to obtain the target position and orientation given the projected circle center and the parameters of the outer circle projection. The viability of the approach is demonstrated based on aerial pictures taken by an UAV from elevations between 10 to 100 m. We obtain a distance RMS below 0.25 % under 50 m and below 1 % under 100 m with a target size of 90 cm, part of which is a deterministic bias introduced by image exposure.

1 INTRODUCTION

Optical coded targets consist in distinctive visual patterns designed to be automatically detected and measured in camera images. A number of different designs have been proposed in the literature, all of them having the following properties: i) low or no similarity with common environmental features, ii) robust disambiguation between multiple instances, usually by means of signals encoded in the marker, iii) cheap and robust detection algorithm. Some examples are shown in Figure 1. Particular attention is currently being given to square coded markers due to the availability of powerful and ready to use augmented reality libraries, such as ARToolKit (Wagner et al., 2008), AprilTags (Olson, 2011), and arUco (Garrido-Jurado et al., 2014). Even tough visual targets are well studied in the literature, this field is still fertile for research in specific applications such as camera calibration, photogrammetry, robotics, augmented reality and in machine vision techniques in general.

In this work we propose a design for an optical target and a novel detection and measuring algorithm for which the requirements are driven by a specific application within the mapKITE project¹. In this project, a new mapping paradigm is introduced in which a terrestrial-aerial tandem, composed by a human driven terrain vehicle (TV), and an autonomous micro aerial vehicle (UAV), performs the surveys, augmenting terrestrial laser range-finder data with a aerial photography (Molina et al., 2015).

The UAV is required to autonomously follow the terrain vehicle during mapping. Normally, the TV transmits its position to the UAV via a radio link. However, the connection between the two vehicles may be lost or the position of the TV might become temporarily unavailable, e.g., because of GPS outages. Thus, optical guidance is also considered. In this context, the tracking algorithm must primarily exhibit a very low percentage of false positive tracks, which could catastrophically deceive guidance strategies, and secondarily, a low percentage of false negative tracks, to allow for smooth following and reduced probability of losing the TV. The optical target is also beneficial in post-processing:

¹See acknowledgments.




Figure 1: Examples of fiducial targets commonly employed in robotics and photogrammetry applications.

if the accuracy of the target center and distance measurements is high enough, the TV position can be employed to build *kine-matic* ground control point (KGCP) to be used in assisted aerial triangulation (AAT) (Molina et al., 2016).

Our solution, depicted in Figure 2, is a multi-scale design that makes use of a code for robust detection and multiple concentric circles for the determination of the target center projection and of the target orientation with respect to the camera. We prefer to use large circles instead of relying on corner features, as it is common in square coded markers, e.g. the one in Figure 1(c), since circle projections, ellipses on the image plane, are cheaply and robustly detected and their defining parameters can be accurately determined by means of closed form least-squares curve fitting. Furthermore, while the information on the position of a corner tend to concentrate in one single pixel as the distance increases, the ellipse center and dimensions are embodied in the coordinates of tents to hundreds of edge pixels, even from high distance.

This work is organized as follows: in Section 2 we introduce our optical target design and we discuss the detection algorithm. In Section 3 we present how the position and orientation of the target can be recovered from the outer circle projection and from the projected circle center, which in turn is located by means of an algorithm discussed in Section 4. We conclude the paper presenting an experimental evaluation based on simulated and real-world data, and some directions for further accuracy improvement.




Figure 2: The optical target used in this work. The target code is obtained sampling on the red dashed line.

2 OPTICAL TARGET DESIGN AND DETECTION

The optical target employed in this work is depicted in Figure 2. Its main features are: i) two large circles replace corners or point features for relative position and orientation determination, ii) the position of the white dots in the black ring encode a signal that is employed for robust detection and 2D orientation determination, iii) multiple (one, in this case) replica can be superimposed with different scales and distinguished thanks to a different target code. This fractal design is necessary to allow at least one target to be fully visible from low distances and it has been considered for automated landing based on optical guidance.

In the following we summarize the main image processing steps employed in the target detection.

2.1 Contours Detection

The first step involves isolating edges, i.e., pixels with high intensity gradient. We employ the Canny detector and then we group positive pixels to form closed contours as in (Suzuki et al., 1985).

2.2 Circles Detection

In the general case, the circular target will appear as an ellipse once projected on the image plane. However, in our application the nominal distance from the target and the limitations on the allowed UAV orientations limit the perspective deformation, so that we can select candidate targets searching for approximately circular contours. To do so we use a computationally inexpensive, yet very effective, test: for each contour, we compute its center D_i averaging the coordinates of the edge pixels, then we compute the mean μ_i and the standard deviation σ_i of the distance between each edge pixel and D_i . We accept a contour if: i) $\mu_i > \mu_{MIN}$ and ii) $\sigma_i < k\mu_i$, where μ_{MIN} and k are two parameters of the algorithm. Note that if the i -th contour is exactly a circle, σ_i is zero, while moderate values of k allow ellipsoidal contours to be also accepted. Finally, we form clusters of concentric circles with similar center coordinates and for each cluster we mark the circle with largest radius as the cluster representative.

2.3 Code Test and Target Heading Determination

We test the inside of each cluster representative for the presence of the target code. We obtain a one dimensional signal sampling the image along the red line in Figure 2, which is defined according to the measured center and radius of the cluster representative. This line intersect the white dots in the black outer ring at specific




Figure 3: Target signal test and orientation determination.

angles and the black/white alternation forms a periodic signal, as a function of the angle. Different codes can be obtained changing the position of the dots, allowing for multiple target tracking and for target scale resolution in case multiple replicas are detected. The signal obtained sampling the image is matched with the reference signal maximizing their normalized cross-correlation (see Figure 3). The cluster representative for which the normalized cross correlation is maximum, and exceeding a minimum threshold, is selected as the final target candidate.

The target 2D orientation on the image plane is also measured: it corresponds to the lag between reference and image signals for which the normalized cross correlation is maximum. While encoded signals are widely employed in optical target detection to reject false matches and to allow for multiple target tracking, the authors are not aware of other target designs that employ signals to measure 2D orientation.

The target 2D orientation is a desirable output at early stages of the tracking pipeline as a naive optical-based TV following could be implemented employing only the information extracted so far, e.g. target position and 2D orientation on the image plane, if an external elevation control strategy is available.

2.4 Sub-pixel Edge Detection and Ellipse Fitting

Once the correct outer circle has been identified, we search in the corresponding circles cluster for the inner one. Next, we fit two preliminary ellipses on the contours corresponding to the two concentric circles, marked in green in Figure 2. This involves finding the parameters Q' of the ellipse equation $x^T Q' x = 0$, with Q' being a symmetric 3×3 matrix with eigenvalues $\lambda_1, \lambda_2 > 0$ and $\lambda_3 < 0$, such that it best fits the coordinates of the edge pixels. This problem can be solved in least-square sense and a closed form solution exists, e.g., see (Fitzgibbon et al., 1996).

The determined ellipses are then refined with sub-pixel edge detection. We determine $\lceil 4\sqrt{a^2 + b^2} \rceil$ points on the ellipse, where a and b are the ellipse major and minor axis. The number of points is a lower bound for the ellipse perimeter. For each of those points, we sample the original image at four locations on a line orthogonal to the ellipse contour (see Figure 4) and we fit a 3-rd order polynomial on the obtained intensities. The inflection point of this polynomial gives the position of the sub-pixel edge point. Finally, we correct the sub-pixel edge points coordinates for lens radial and tangential distortion and we repeat the ellipse curve fitting steps employing the set of refined edge points.

At this stage we have fitted on the image two ellipses, corresponding to the projection of the two concentric circle of the target, and we have obtained the related parameters Q'_o and Q'_i . From now on, we will not perform any further measurement on the image.

3 TARGET POSE DETERMINATION

In this section we discuss how the relative position and orientation of the target with respect to the camera can be recovered.




Figure 4: Sub-pixel edge detection and ellipse fitting. The outer ellipse major and minor axis are respectively 13.320 and 12.876 pixel long, while the contour is composed of 150 edge pixels. The image correspond to real-world data and it is taken from a distance of 98.6 m.

To this end, let us assume for the moment that we know very precisely the position of the projection of the concentric circles center. We will discuss how this can be obtained in Section 4.

In Figure 5 we have sketched the target pose determination problem: O is the optical center of the camera, C is the center of the target, segment \overline{AB} is the target diameter that is coplanar with the y axis of the camera frame. The sketch is a projection of the scene on the yz plane. As we have assumed, we know very precisely the position of the projected circle center in the image plane, C' . Moreover, we have estimated the parameters of the ellipse Q'_o , being the projection on the image plane of the target outer circle. We determine analytically the intersection of Q'_o with a line trough C' and parallel to the y axis of the camera frame, obtaining the points A' and B' in the image plane. This points correspond to the projection of A and B .

Three viewing rays, directions in camera frame, correspond to the three points A' , B' , and C' and they can be determined by pre-multiplying image point coordinates with the inverse of the intrinsic camera calibration matrix, and then normalizing, e.g.:

$$\vec{v}_A = \frac{K^{-1}A'}{\|K^{-1}A'\|}. \quad (1)$$

Note that in general \vec{v}_A , \vec{v}_B and \vec{v}_C do not belong to the yz plane, as it might appear from the drawing in Figure 5, where we have drawn their projection on that plane.

Next, the angles θ_1 and θ_2 can be determined from the *dot* product of the viewing rays:

$$\theta_1 = \cos^{-1}(\vec{v}_A \cdot \vec{v}_C), \quad (2)$$

$$\theta_2 = \cos^{-1}(\vec{v}_B \cdot \vec{v}_C). \quad (3)$$

We are now ready to determine the length of \overline{OC} , which is the target distance. This problem can be seen as a two-dimensional instance of the well known Perspective-3-Points (P3P) problem (Fischler and Bolles, 1981, Wu and Hu, 2006). A closed form solution can be obtained by applying the cosine theorem to the trian-




Figure 5: Target pose determination problem, projection on the yz plane of the camera frame. O is the camera optical center, C is the target center and the segment \overline{AB} is the target diameter which is coplanar to the y axis. Clearly $\overline{AC} = \overline{BC} = r$.

gles AOC , COB and AOB , i.e.,

$$\begin{cases} r^2 = \overline{OA}^2 + \overline{OC}^2 - 2\overline{OA} \overline{OC} \cos(\theta_1) \\ r^2 = \overline{OB}^2 + \overline{OC}^2 - 2\overline{OB} \overline{OC} \cos(\theta_2) \\ (2r)^2 = \overline{OA}^2 + \overline{OB}^2 - 2\overline{OA} \overline{OB} \cos(\theta_1 + \theta_2) \end{cases}, \quad (4)$$

where r is the target radius and the length of the segments \overline{AC} and \overline{BC} . Solving in \overline{OC} yields:

$$\overline{OC} = \frac{\sqrt{2}r \sin(\theta_1 + \theta_2)}{\sqrt{3 - 2 \cos(2\theta_1) - 2 \cos(2\theta_2) + \cos(2(\theta_1 + \theta_2))}}, \quad (5)$$

We can also determine the angle ϕ , which is the misalignment of the considered target diameter \overline{AB} with respect to the y axis of the camera frame: first we determine the angles γ_1 and β_2 , applying the sine theorem to the triangles AOC and COB :

$$\gamma_1 = \sin^{-1}\left(\frac{\overline{OC}}{r} \sin(\theta_1)\right), \quad (6)$$

$$\beta_2 = \sin^{-1}\left(\frac{\overline{OC}}{r} \sin(\theta_2)\right). \quad (7)$$

Each equation gives two solutions, since $\sin(\alpha) = \sin(\pi - \alpha)$, yet only one couple (γ_1, β_2) satisfies

$$\theta_1 + \gamma_1 + (\pi - \gamma_2) = \theta_2 + \gamma_2 + \beta_2 = \pi. \quad (8)$$

We can now determine \overline{OB} , the vector \vec{v}_{OB} and ϕ :

$$\overline{OB} = r \frac{\sin(\pi - \beta_2 - \theta_2)}{\sin \theta_2}, \quad (9)$$

$$\vec{v}_{OB} = \overline{OB} \vec{v}_B - \overline{OC} \vec{v}_C, \quad (10)$$

$$\phi = \tan^{-1}(\vec{v}_{OB,y}, \vec{v}_{OB,z}). \quad (11)$$

The angle ϕ is the roll angle of the relative orientation of the target with respect to the camera. Indeed this angle is formed by the z axis of the camera frame and the target diameter \overline{AB} , that has been chosen to lie in the yz plane. To determine the pitch angle θ , the procedure just discussed is repeated, this time intersecting the ellipse Q'_o with a line trough C' and parallel to the x axis, thus obtaining the projection of a diameter lying in the xz plane.

4 PROJECTION OF THE CIRCLE CENTER

It remains to discuss how the coordinates of the projection of C can be determined. Note that this point is *not* the center of either Q'_o or Q'_i , as it can be seen remembering that perspective transformations do not preserve ratios between distances. We chose




Figure 6: $d(\tilde{C}', \gamma, Q'_o)$ as a function of γ for five random points \tilde{C}' in the neighborhood of C' . The ellipse Q'_o is the outer circle projection of a target placed a true distance of 15.56 m.

not to employ corner features to locate C' since these would not be sharp or visible from high distance, as it is clear from Figure 4.

One known solution to determine the image coordinates of C' is to employ two concentric circles, such as Q_o and Q_i . Once the parameters of two ellipses have been fitted to the pixels composing the corresponding contours, a quasi-analytical solution exists for C' : let λ be the repeated eigenvalue of $Q'_i Q'^{-1}_o$, any non-zero row of $Q'^{-1}_o - \lambda Q'^{-1}_i$ is the homogeneous coordinate of the projected circle center. For the proof see (Kim et al., 2005).

The afore-discussed result holds only if the *estimated* ellipses Q'_o and Q'_i are projections of concentric circles. In other words, in order for Q'_o and Q'_i to be the projections of two concentric circles, certain constraints must be satisfied by their parameters, see again (Kim et al., 2005). So constrained estimation should be employed to jointly fit Q'_o and Q'_i on the edge pixels. If this is not the case, as in (Yang et al., 2014) and in this work, i.e., the two parameter sets are estimated independently, geometric constraints between the two ellipses may not hold exactly. In practical applications this causes $Q'_i Q'^{-1}_o$ not to have any repeated eigenvalue, but two very similar ones, and the non-zero rows is $Q'^{-1}_o - \lambda Q'^{-1}_i$ being slightly different, whereas they should be equal. This gives up to six candidates for C' . As we will show in the following, in general these solutions are not accurate enough for being employed in the method discussed in Section 3.

We discuss here a technique to improve the estimate of C' obtained with the Kim's method. We first define a function d that, for a given ellipse Q' , estimates \overline{OC} as a function of the coordinates of the projected circle center, as in Section 3, yet, to obtain A' and B' , we intersect Q'_o with an arbitrary line through C' , and parametrized with an angle γ . Thus we have:

$$\overline{OC} = d(C', \gamma, Q'), \quad (12)$$

where γ is the angle formed by the line used to intersect Q' and the x axis of the image plane. Note that in Section 3 we used a vertical line, i.e., $\gamma = \pi/2$ to determine ϕ , and an horizontal one, i.e., $\gamma = 0$, to determine θ . We would like to stress the fact that the function $d(\cdot)$ is analytical.

Suppose now that we have applied the Kim's method based on the repeated eigenvalue of $Q'_i Q'^{-1}_o$ and we have obtained an initial guess for the projection of the circle center, \tilde{C}' . As anticipated, this estimate is problematic in practical cases and it is not accurate enough to be employed for computing \overline{OC} . To understand why, consider the following reasoning: if \tilde{C}' is the true projection of the circle center, the function $d(\tilde{C}', \gamma, Q')$ must return the same value for every choice of the angle γ . In fact, if \tilde{C}' is the projected circle center, whatever line passing through \tilde{C}' will intersect the




Figure 7: Landscape of $D(\tilde{C}')$ for an ellipse with major and minor axis 91.02 and 76.46 px, angle -113.81° and center (379.43, 120.55). Marked points are the two candidates for the projected circle center C' and the solution obtained employing the method in (Kim et al., 2005).

ellipse Q' in two points being the projection of one diameter of the circle and thus enabling the reasoning discussed in Section 3. Instead, if \tilde{C}' is not the projected circle center, the segment \overline{AB} is not a diameter and in general $\overline{AC} \neq \overline{BC} \neq r$. Thus, even for subtle deviations of \tilde{C}' from its true coordinates, $d(\tilde{C}', \gamma, Q')$ returns different distances for different values of γ . As an example, we have chosen an ellipse Q'_o corresponding to a target placed at 15.56 m and we have plotted $d(\tilde{C}', \gamma, Q'_o)$ for five random points \tilde{C}' sampled in the close neighborhood of the true projected circle center ($\sigma = 0.1$ pixels). See Figure 6. Variations in $d(\tilde{C}', \gamma, Q'_o)$ of the order of 0.5 m are obtained.

The aforesaid reasoning gives us a very sensitive geometric invariant that we can employ to improve our initial guess for C' :

$$D(\tilde{C}) := \text{std}_\gamma [d(\tilde{C}, \gamma, Q'_o)], \quad (13)$$

$$C' = \underset{\tilde{C}}{\text{argmin}} \{D(\tilde{C})\}. \quad (14)$$

Namely, for a given ellipse Q'_o we search for the point \tilde{C}' for which the dispersion of the values of $d(\tilde{C}', \gamma, Q'_o)$ is minimum. As we know, there exist at least one point C' for which $D(\tilde{C}')$ is identically zero, being the true projection of the circle center.

The landscape of the function $D(\tilde{C}')$ is depicted in Figure 7 for the ellipse Q'_o employed in the previous example. It is immediately evident that in this case there are two global minima for $D(\tilde{C}')$. To see why, consider again Figure 5 and take the symmetry of A and B around the bisector of the angle $\angle AOB$: the same projections A' and B' are obtained but the orientation of the target and the projection of the circle center C' are different. It is not possible to discriminate between the two orientations, or decide for one of the two attractors in Figure 7 employing only the projection of one circle contour. However, for what concerns target distance estimation, the existence of two projected circle center candidates is not an issue since the same distance is associated to both points (swap θ_1 and θ_2 in Equation 5). A very related problem exists in every planar target design and it is often referred as *pose ambiguity*: there exist multiple target orientations that correspond to almost indistinguishable projections in the image plane, see (Schweighofer and Pinz, 2006). Solutions for this problem involve moving to 3D targets, or in general employing techniques other than image processing to resolve the ambiguity, e.g., see (Wu et al., 2014, Tanaka et al., 2014).




Figure 8: Estimate of the detection probability as a function of the true distance from the camera.

To solve the minimization problem in Equation 14 we employ a gradient descent search strategy with adaptive learning rate and we initialize the search with the solution obtained with Kim’s method. The function to be optimized is reasonably regular and the initial guess is already close to one global optimum. Note that the initial guess is computed employing both the inner and the outer ellipses, in which case the solution is unique. Thus we conjecture that in most of the cases the initial guess lies in the attractor basin of the true projected circle center.

In this section we have presented very sensitive method to determine the coordinates of the projected circle center, which enables the pose determination method in Section 3. We would like to stress the fact that in the presented algorithm the only quantities that are measured on the image are the parameters of the ellipses Q'_o and Q'_i . All the other quantities are analytical or quasi-analytical functions of these parameters. Indeed, from Q'_o and Q'_i we derive an initial guess for C' with the Kim’s method. Then, only Q'_o is employed in the numerical minimization of $D(\tilde{C}')$, yielding C' . The coordinates of C' are finally plugged in the expressions in Section 3 to determine OC , ϕ and θ . So the accuracy and precision of the pose estimation greatly depend on the procedure employed to measure Q'_o and Q'_i in the image, from the edge detection to the curve fitting.

5 EXPERIMENTAL EVALUATION

In this section we evaluate the target detection algorithm based on synthetic and real-world images. To this end, we flew with a custom octo-copter over the optical target and we acquired pictures at relative elevations between 10 and 100 m.

We employed a Ximea MQ042MG 4 Mpx camera equipped with 16mm lenses, pointing down, which gives us a GSD of 3.5 cm and a field of view of 70 m at 100 m elevation. The radius of the target outer circle is 0.9 m. We determined the intrinsic camera calibration matrix and lens distortion coefficients during a separate flight over a dedicated calibration field. To obtain ground truth positions for the camera optical center and for the marker center we oriented the collected frames by means of AAT.

The detection algorithm proved to be very robust: the target was successfully detected in 1807 frames out of 2134 (84.7 %). An estimate of the probability that the target is detected as a function of the true distance is shown in Figure 8. Based on these empirical results we argue that the detection is reliable if the target diameter size in the image is bigger than 40 pixels, which translates, given the considered choice of camera and lenses, to a flight elevation lower than 80 m. Moreover, a deceiving target, i.e., a target with the same geometry but different code, was placed nearby. The deceiving target was selected in only 5 images, confirming the viability of the approach in multi-target applications.




Figure 9: Error in distance measurement as a function of the true distance. We have also plotted a moving average and standard deviations computed with window size 10 m.




Figure 10: Effect of different exposures on the estimated distance for two frames at nominal distance of 38.41 m and 97.07 m (darker on the left, brighter on the right, original in the center).

In Figure 9 we have plotted the difference between the reference distance, i.e., OC in Figure 5, and the one computed with the described tracking algorithm. The red curves are windowed average and standard deviations. It is possible to see that the algorithm yields unbiased and very precise distance measurements up to 50 m while at higher elevation the distance measurement appears to be biased by some deterministic effect, even though mean error remains below 1 % of the distance. We don’t experience a comparable effect in synthetic data.

One possible explanation is the following: sharp edges become gradients in camera images due to lens point spread function and blurring. The sub-pixel edge detection presented in Section 2.4 aims at recovering the precise edge location from these gradients. However, different camera exposures yield different gradients, making the ellipses appear slightly bigger or smaller and thus giving different distance estimation. An example is given in Figure 10, where we show that if we artificially change the exposure we obtain different distance estimates, depending on the elevation. This can motivate the deterministic bias visible in Figure 9. In order to increase the distance accuracy we have to either precisely control the exposure of each frame or have a sub-pixel edge detector that is invariant with respect to exposure.

Other possible reasons are: i) the geometry of the photogrammetric problem is rather weak as the flight involved a climb and a descent always pointing the same area. It is possible that the obtained reference positions from AAT are not sufficiently accurate. Unfortunately, camera positions uncertainties are not provided by the commercial software used for image orientation; ii) the intrinsic camera calibration is not constant and can change because of temperature, vibrations and other effects. It is possible that some distortion between reference positions and estimated target distance is due to slight variations in the camera calibration.

In Figure 11 we have plotted the ϕ angle estimation error for synthetic images obtained rendering the target as it would be seen




Figure 11: Error in ϕ measurement on synthetic images.

from random camera positions and orientations. It is possible to see that while the estimates are unbiased, the error standard deviation is quite high even at low distances. This is expected as the marker distance is 10 to 50 times bigger than the marker radius.

6 CONCLUSIONS AND FUTURE WORK

In this work we have introduced a novel optical target design developed with the goal of determining with high accuracy its center and its distance from UAV-based mapping elevations. An experimental evaluation based on real-world data has shown remarkable accuracy and has shown that further improvements can be obtained correcting for the exposure induced bias.

The presented position and orientation determination algorithm works on the estimated parameters of the two ellipses corresponding to the outer and inner circles in the marker. These parameters are robustly estimated in least-squares sense from a set of sub-pixel edge points. This step is crucial for the accuracy of the distance and orientation measurements, as no other quantities are measured on the image. One direct way to improve this step is the following: at the present stage, the two ellipses are estimated independently. This does not exploit the fact that they are the projection of two concentric circles. The curve fitting problem can be reformulated so that the constraints existing between the ellipse parameters are kept into account during estimation. Moreover, the effects of varying exposure on the edge positions on the image has to be investigated as biases might exist in the employed sub-pixel detection method.

After the two ellipses are determined, the projection of the circle center is estimated numerically minimizing an analytical non linear function of the outer ellipse parameters. Two attractor basins, and thus two candidates for the circle center projection, exist for this function. These two candidates correspond to two possible orientations for the target. The orientation ambiguity can be solved employing the second ellipse. We argue that it should be possible to obtain the coordinates of the two candidates in closed form as a function of the outer ellipse equation only. This would remove the need for the numerical minimization.

After the projected circle center coordinates have been determined, the distance and the orientation of the target are determined analytically. A error propagation analysis has to be performed to provide confidence interval for these quantities as a function of the uncertainty on ellipse parameters, which comes from least-squares curve fitting, and as a function of the intrinsic camera calibration parameters.

ACKNOWLEDGEMENTS

This work was partially funded by the European Community Horizon 2020 Programme under grant agreement no. 641518 (project

mapKITE, www.mapkite.com) managed by the European GNSS Agency (GSA).

REFERENCES

- Fischler, M. A. and Bolles, R. C., 1981. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* 24(6), pp. 381–395.
- Fitzgibbon, A. W., Fisher, R. B. et al., 1996. A buyer's guide to conic fitting. *DAI Research paper*.
- Garrido-Jurado, S., Muñoz-Salinas, R., Madrid-Cuevas, F. and Marn-Jimnez, M., 2014. Automatic generation and detection of highly reliable fiducial markers under occlusion. *Pattern Recognition* 47(6), pp. 2280 – 2292.
- Kim, J.-S., Gurdjos, P. and Kweon, I.-S., 2005. Geometric and algebraic constraints of projected concentric circles and their applications to camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 27(4), pp. 637–642.
- Molina, P., Blázquez, M., Sastre, J. and Colomina, I., 2015. A method for simultaneous aerial and terrestrial geodata acquisition for corridor mapping. *The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences XL-1/W4*, pp. 227–232.
- Molina, P., Blázquez, M., Sastre, J. and Colomina, I., 2016. Precision analysis of point-and-scale photogrammetric measurements for corridor mapping: Preliminary results. *The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences XL-3/W4*, pp. 85–90.
- Olson, E., 2011. AprilTag: A robust and flexible visual fiducial system. In: *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, pp. 3400–3407.
- Schweighofer, G. and Pinz, A., 2006. Robust pose estimation from a planar target. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 28(12), pp. 2024–2030.
- Suzuki, S. et al., 1985. Topological structural analysis of digitized binary images by border following. *Computer Vision, Graphics, and Image Processing* 30(1), pp. 32–46.
- Tanaka, H., Sumi, Y. and Matsumoto, Y., 2014. A solution to pose ambiguity of visual markers using moiré patterns. In: *Intelligent Robots and Systems (IROS 2014)*, 2014 IEEE/RSJ International Conference on, IEEE, pp. 3129–3134.
- Wagner, D., Reitmayr, G., Mulloni, A., Drummond, T. and Schmalstieg, D., 2008. Pose tracking from natural features on mobile phones. In: *Proceedings of the 7th IEEE/ACM International Symposium on Mixed and Augmented Reality*, IEEE Computer Society, pp. 125–134.
- Wu, P.-C., Tsai, Y.-H. and Chien, S.-Y., 2014. Stable pose tracking from a planar target with an analytical motion model in real-time applications. In: *Multimedia Signal Processing (MMSp)*, 2014 IEEE 16th International Workshop on, IEEE, pp. 1–6.
- Wu, Y. and Hu, Z., 2006. Pnp problem revisited. *Journal of Mathematical Imaging and Vision* 24(1), pp. 131–141.
- Yang, X., Fang, S., Kong, B. and Li, Y., 2014. Design of a color coded target for vision measurements. *Optik-International Journal for Light and Electron Optics* 125(14), pp. 3727–3732.