Reconstructing Spheres by Fitting Planes

Erol Ozgur[1]
erol.ozgur@sigma-clermont.fr
Mohammad Alkhatib[1]
mohammad.alkhatib@sigma-clermont.fr
Youcef Mezouar[1]
youcef.mezouar@sigma-clermont.fr
Adrien Bartoli[2]
adrien.bartoli@gmail.com

[1] Clermont Auvergne INP
Institut Pascal
Clermont-Ferrand, France
[2] University Hospital
Clermont-Ferrand, France

Abstract

We address the problem of reconstructing a sphere of a prescribed radius from a single calibrated view of its occluding contour. A sphere's occluding contour generally appears as an ellipse and existing methods use ellipse fitting. Most methods thus require ≥ 5 contour points though a few can also deal with the minimal case of 3 points. However they all share two shortcomings: *(i)* they fail for non-elliptic occluding contours, including parabola and hyperbola, and *(ii)* they use the point-to-ellipse distance, whose computation is not closed-form.

We make the observation that the spherically-normalised contour points form a circle in space, which we reconstruct by plane fitting. This handles minimal 3-point and redundant > 3 point fitting, copes with elliptic and non-elliptic contours, and benefits from the simple point-to-plane distance. The reconstructed circle then leads to a one-parameter sphere family from which the actual sphere of prescribed radius is uniquely retrieved. We robustify the method using random sampling at the plane fitting stage. We name our method `SpherO`, where letter 'O' depicts a circle. Experimental comparisons show that `SpherO` outperforms the current-best 3-point method.

1 Introduction

Sphere reconstruction forms a practical tool for applications requiring the pose of sphere-shaped objects, including visual servoing of robots. Sphere reconstruction, given its radius and its occluding contour in an image modelled by the calibrated perspective camera model, entails finding the sphere's 3D centre point. A sphere's occluding contour is observed as either an ellipse, a parabola or a hyperbola in the image. State of the art methods handle the elliptic case –the most common one– but fail for the other two cases. They approximate the point-to-ellipse distance, which is not closed-form, by its algebraic counterpart. We propose a novel simple and convex method resolving both limitations and handling minimal and redundant cases. It uses spherical normalisation, which places an image point on the camera-centred unit sphere, and an observation formalised in the following proposition and illustrated in figure 1.

Proposition 1. *A sphere's spherically-normalised occluding contour is a sphero-circle.*

A sphero-conic is a curve formed by intersecting a quadric cone with the unit sphere whose centre is the vertex of the cone. If the cone is circular then the sphero-conic is a circle [1] also known as 'sphero-circle'. Using this result with the camera centre taken as vertex and the sphere's projection lines taken as circular cone proves proposition 1.

The proposed reconstruction method has three steps. First, following proposition 1, we spherically normalise the sphere's occluding contour to form a sphero-circle. Second, we reconstruct the sphero-circle's support plane, centre and radius. This induces a one-parameter sphere family. Third, we retrieve the actual sphere of prescribed radius. We name our method `SpherO`, where 'O' depicts a circle. The use of spherical normalisation, following proposition 1, allows one to rid out the effect of the camera's image plane, only retaining the camera's projection centre. This gives `SpherO` the ability to use mere plane-fitting, making it the most compelling method against the state of the art. First, `SpherO` handles both minimal and redundant cases seamlessly. It can thus serve as minimal method in the inner loop of RANSAC-like robust methods and as final refinement method given the inlier set. Second, `SpherO` uses a geometric cost function, namely a point-to-plane distance. It entirely avoids the point-to-conic distance required by existing methods, providing a quick and sound way of determining the inlier set within RANSAC. Third, `SpherO` seamlessly deals with the settings where the sphere is seen as an ellipse, a parabola or a hyperbola. Fourth, `SpherO` provides a one-parameter family of spheres when the sphere's radius is unknown and uniquely reconstructs the sphere otherwise. Fifth, `SpherO` does not have artificial degeneracies. In other words, `SpherO` deals with all settings where the sphere's occluding contour is a proper conic. We review related work, formally describe `SpherO` and its robust version `Robust-SpherO`, present experimental results and conclude.

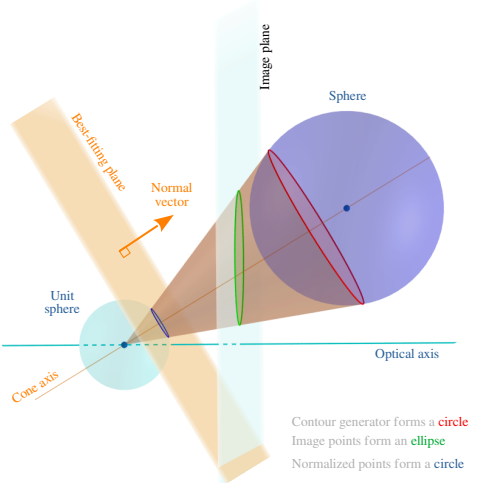

Figure 1: The geometric observation subtending the proposed `SpherO` method. Tilting the camera's image plane (in light blue) towards the cone axis leads to the sphere's occluding contour (in green) becoming a parabola and eventually a hyperbola. Importantly, the contour generator (in red) and normalised points (in blue) remain invariant, hence circles, showing that `SpherO` seamlessly handles all projection cases.

2 Related Work

There exist only three methods which use 3 points or more [6, 7, 8]. Methods [7] and [6] are not robust; they minimise the squared algebraic point-to-ellipse image distance. Methods in [8] minimise the geometric distance from a cone tangent to the sphere (method `Toth`); it provides a RANSAC-based robust solution specifically relying on the ellipse to form the consensus set (method `Robust-Toth`). It is shown in [8] that `Toth` outperforms the two methods [6, 7] and other existing methods requiring at least 5 points [3, 4, 5], and is thus the state of the art. Table 1 summarises the characteristics of the proposed `SpherO` and `Robust-SpherO` in comparison to `Toth` and `Robust-Toth`. We have that `SpherO` and `Toth` are both minimal and redundant, can handle any conical occluding contours and use a geometric cost. We will see that `SpherO` is simpler to derive and faster. We have that `Robust-SpherO` is the first robust method to handle any conical occluding contours, as `Robust-Toth` is restricted to ellipses.

Table 1: Comparison of $\geq$ 3-point sphere reconstruction methods.

	Handles all image conics	Fits	Geometric costs
`Toth` [8]	✓	cone	secant
`Robust-Toth` [8]	only ellipse	ellipse cone	point-to-ellipse distance* secant
`SpherO`	✓	plane	point-to-plane distance
`Robust-SpherO`	✓	plane	point-to-plane distance

**numerically approximated by algebraic distance*

3 Methodology

We give the three steps of `SpherO` and its robustification `Robust-SpherO`.

3.1 Notation and Spherical Normalisation

We express all 3-space coordinates in the standard pinhole camera coordinate frame. We denote the intrinsic matrix as $\mathbf{K} \in \Re^{3\times3}$, the optical centre as $\mathbf{O} \in \Re^3$ and the unit direction vector of the principal axis as $\underline{\mathbf{z}} \in \Re^{3\times1}$. We denote the sphere's centre as $\mathbf{C} \in \Re^3$ and its radius as $R \in \Re^+$. We denote the points on the sphere's contour generator as $\mathbf{P}_i \in \Re^3$ for $i = 1,\ldots,m \in \mathbb{N}$. We denote their corresponding points on the sphere's occluding contour in the image as $\mathbf{p}_i \in \Re^2$ in pixel units. They are related to each other proportionally as $\bar{\mathbf{p}}_i \propto \mathbf{K}\mathbf{P}_i$ where $\bar{\mathbf{p}}_i = \left[\mathbf{p}_i^\top, 1\right]^\top$ is the homogeneous coordinates of an occluding contour point. The backprojection of the occluding contour point forms a ray that starts from the optical centre $\mathbf{O}$ and passes through its corresponding point $\mathbf{P}_i$ on the sphere's contour generator.

Spherical normalisation of the occluding contour points is the first step of `SpherO`. It is equivalent to computing the ray's unit direction vector $\mathbf{q}_i \in S^2 \subset \Re^3$, $\|\mathbf{q}_i\| = 1$, as:

$$\mathbf{q}_i = \mathbf{K}^{-1}\bar{\mathbf{p}}_i / \|\mathbf{K}^{-1}\bar{\mathbf{p}}_i\|. \tag{1}$$

3.2 Sphero-circle Reconstruction by Plane Fitting

Sphero-circle reconstruction is the second step of `SpherO`. We give a solution using plane fitting. We first give the general case for $m \geq 3$ points. We then give an equivalent but lower computation cost version for $m = 3$, which is meant to be used in RANSAC's inner loop. We finally give closed-forms for the sphero-circle's centre and radius.

General plane fitting. We define the best-fitting plane as the minimizer of the sum of the squared orthogonal distances between the spherically-normalised points and the plane. Let $\mathbf{q}_0$ be the average of the points. We use a simple singular value decomposition of the $3 \times m$ matrix containing the centred points $\mathbf{q}_i - \mathbf{q}_0$, whose least singular vector gives the plane's unit normal vector $\underline{\mathbf{n}} \in S^2$ [2]. The plane's signed distance to origin is $d_\Pi = \mathbf{q}_0^\top \underline{\mathbf{n}}$ and the plane's coordinate vector is $\Pi = [\,\mathtt{sign}(d_\Pi)\underline{\mathbf{n}}^\top, -|d_\Pi|\,]^\top$, where $\mathtt{sign}(d_\Pi)$ ensures that the normal vector points toward the sphere.

Minimal 3-point plane recovery. Given 3 points $\mathbf{q}_1$, $\mathbf{q}_2$ and $\mathbf{q}_3$, the plane's normal is:

$$\underline{\mathbf{n}} = ((\mathbf{q}_3 - \mathbf{q}_1) \times (\mathbf{q}_2 - \mathbf{q}_1)) \,/\, \|(\mathbf{q}_3 - \mathbf{q}_1) \times (\mathbf{q}_2 - \mathbf{q}_1)\|. \tag{2}$$

The plane's signed distance to origin is then $d_\Pi = \mathbf{q}_i^\top \underline{\mathbf{n}}$ where $i \in \{1,2,3\}$ is any of the 3 points and the plane coordinates Π are formed as above.

Sphero-circle reconstruction. The best-fitting plane's normal $\underline{\mathbf{n}}$ yields the sphero-circle's support plane normal. We recall that the intersection of a plane orthogonal to the axis of a circular cone is a circle whose centre is the intersection point with the cone's axis. This implies that the sphero-circle's support plane is perpendicular to the cone's axis. It follows that the sphero-circle's centre $\mathbf{c} = d_\Pi \underline{\mathbf{n}}$ is the closest point of its support plane to the origin. We thus proceed to compute the sphero-circle's radius as follows. First, we form a right triangle $\measuredangle(\mathbf{O}\,\mathbf{c}\,\mathbf{q}_i)$ using the camera centre $\mathbf{O}$, the sphero-circle's centre $\mathbf{c}$ and any normalised point $\mathbf{q}_i$ of the sphero-circle. The length of hypotenuse $\|\,\mathbf{O} - \mathbf{q}_i\,\|$ is 1 since $\mathbf{q}_i$ is spherically normalised. Second, we apply Pythagoras' theorem to calculate the sphero-circle's radius $r = \sqrt{1 - d_\Pi^2}$. We finally write the sphero-circle's coordinate vector as $\odot = [\underline{\mathbf{n}}^\top, \mathbf{c}^\top, r]^\top$.

3.3 Sphere Family Extraction and Sphere Selection

We use the sphero-circle to form a one-parameter sphere family containing all spheres whose occluding contour is the observed one. We parameterise this family by the sphere radius ρ. If the reconstruction radius R is not prescribed, this family represents all potential reconstruction solutions; if it is prescribed, then `SpherO`'s third step uses it to select the uniquely reconstructed sphere from the family, by simply setting $\rho = R$.

We represent a sphere by a 4-vector in the affine space $\mathfrak{R}^4$ as $\mathbf{s}(\rho) = [\,\mathbf{C}(\rho)^\top, \rho\,]^\top$, with centre $\mathbf{C}(\rho) \in \mathfrak{R}^3$ and radius $\rho \in \mathfrak{R}^+$. The centre $\mathbf{C}(\rho)$ must be on the ray which starts from the camera centre $\mathbf{O}$. It thus passes through the sphero-circle's centre $\mathbf{c}$ and is also aligned with the sphero-circle's support plane normal, leading to:

$$\mathbf{C}(\rho) = d(\rho)\,\underline{\mathbf{n}}, \tag{3}$$

where $d(\rho)$ encodes the depth of the sphere's centre. This depth may be found by forming two similar right triangles. The first triangle $\measuredangle(\mathbf{O}\,\mathbf{c}\,\mathbf{q}_i)$ is the same one as explained in the

sphero-circle reconstruction. The second triangle is $\measuredangle(\mathbf{O}\,\mathbf{P}_i\,\mathbf{C})$. It follows that:

$$d(\rho) = \rho / r. \tag{4}$$

Finally, when the sphere radius is prescribed as R, the reconstructed sphere is $R[\underline{\mathbf{n}}^\top / r\,,\,1]^\top$.

3.4 Robustification

We perform sampling-based robust sphere reconstruction using `SpherO`'s minimal plane-fitting solution. We give `Robust-SpherO` in algorithm 1. The inputs are the sphere's occluding contour pixels, the sphere's prescribed radius, the camera intrinsics and the RANSAC threshold. The output is the sphere of prescribed radius. Line 1 transforms the RANSAC threshold distance from pixel units to a distance in the normalised image plane as follows:

Algorithm 1 `Robust-SpherO`

Input: Sphere's occluding contour pixels $\{\,\mathbf{p}_i \mid \mathbf{p}_i \in \Re^2,\ i = \{1,\ldots,m\},\ m \geq 3\,\}$, sphere's radius $R \in \Re^+$ in metric units, and the camera intrinsics $\mathbf{K} \in \Re^{3\times3}$, RANSAC threshold $\tau_{pix} \in \Re^+$ in pixels (*e.g.*, 1 pixel)
Output: Sphere $\mathbf{s} \in \Re^4$

1: τ = `Threshold Normalisation`$(\tau_{pix}, \mathbf{K})$ // equation (5)
2: $\{\,\mathbf{q}_i\,\}$ = `Spherical Normalisation`$(\{\,\mathbf{p}_i\,\}, \mathbf{K})$ // equation (1)
3: Π = `RANSAC Plane Fit`$(\{\,\mathbf{q}_i\,\}, \tau)$ // section 3.2
4: $\odot$ = `Sphero-Circle Reconstruction`(Π) // section 3.2
5: $\mathbf{s}$ = `Sphere Recovery`$(\odot, R)$ // section 3.3

$$\tau = \tau_{pix} / \mathtt{max}(f_x, f_y) \tag{5}$$

where f_x and f_y are the focal lengths in pixels from the camera intrinsics. The RANSAC threshold τ_{pix} can be considered as the maximum reprojection error of a reconstructed sphere solution. Intuitively, one can choose it as equal to or greater than the camera's intrinsic calibration accuracy (*i.e.*, reprojection error in pixels). Line 2 spherically normalises the occluding contour. Line 3 uses RANSAC on the normalised occluding contour points. The algorithm iterates using the 3-point plane recovery formula (2) on random 3-point samples. Each plane uses a threshold scaled by its depth $\tau_\Pi = |d_\Pi|\,\tau$ to form its consensus set. The algorithm terminates with a best-fitting plane on the largest consensus set. Line 4 reconstructs the sphero-circle from the best-fitting plane. Line 5 retrieves the sphere of prescribed radius from the sphere family obtained through the sphero-circle.

4 Experiments and Results

We compare `Robust-SpherO` to the stat of the art `Robust-Toth`, *i.e.*, *3pFit + Direct3pFit* from [8], through synthetic and real data experiments.

4.1 Synthetic Data Experiments

We test the methods' robustness against elliptic and non-elliptic occluding contours. We used the intrinsics of a calibrated camera throughout the synthetic data experiments. The

focal length was $f_x = f_y = 1174$ pixels, the principal point was $x_0 = 1028.4$ and $y_0 = 673.4$ pixels and the skew was $s = 0$. For a fair comparison of the methods, they both used the same random samples from the sphere's occluding contour and the same RANSAC threshold τ_{pix}, chosen equal to the image noise level.

The graphs show the mean and standard deviation of the error, which is the Euclidean distance to the ground-truth sphere centres against the varying noise level, number of erroneous points, occlusion level and occluding contour type. We repeated the experiments 1000 times for each value of a varying parameter.

4.1.1 Robustness with Elliptic Occluding Contours

We formed the elliptic occluding contours ; the sphere radius was $R = 0.5$ metres and the sphere centre was generated randomly in each trial as $\mathbf{C} = [x \sim \mathcal{N}(\mu, \sigma^2), y \sim \mathcal{N}(\mu, \sigma^2), z \sim \mathcal{N}(\mu_z, \sigma_z^2)]^\top$ with $\mu = 0$ and $\sigma^2 = 2$ and $\mu_z = 5$ and $\sigma_z^2 = 1$ metres.

Accuracy versus image noise and number of correct points. We varied the white Gaussian noise level from 0 to 10 pixels with step size of 1 pixel. Experiments were performed with 100 correct points (figure 2 - left). We then varied the number of correct points from 10 points to 100 points with step size of 10 points. Experiments were performed with 2 pixels white Gaussian noise (figure 2 - right). `Robust-SpherO` outperforms `Robust-Toth` in both set of experiments, with errors consistently twice as low.

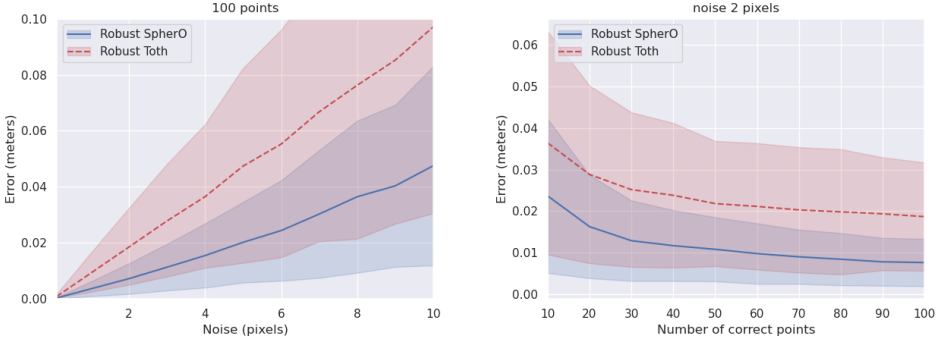

Figure 2: Mean and standard deviation of errors versus image noise levels (left) and number of correct points (right). Correct points belong to the sphere's occluding contour.

Accuracy versus number of erroneous points. We varied the erroneous point rate from 5% to 75% with step size of 5% and using 100 points in each trial. Figure 3 presents the results with the following fixed parameters: 1 pixel white Gaussian noise (left) and 2 pixels white Gaussian noise (right). `Robust-SpherO` outperforms `Robust-Toth`, though a breakpoint appears at 50% and 35% of erroneous point rates, respectively.

Accuracy versus occlusion. We varied the occlusion level from 10% to 70% with step size of 10% over 100 contour points. Figure 4 presents the results with the following fixed parameters: 1 pixel white Gaussian noise and 10% erroneous points (left); 2 pixels white Gaussian

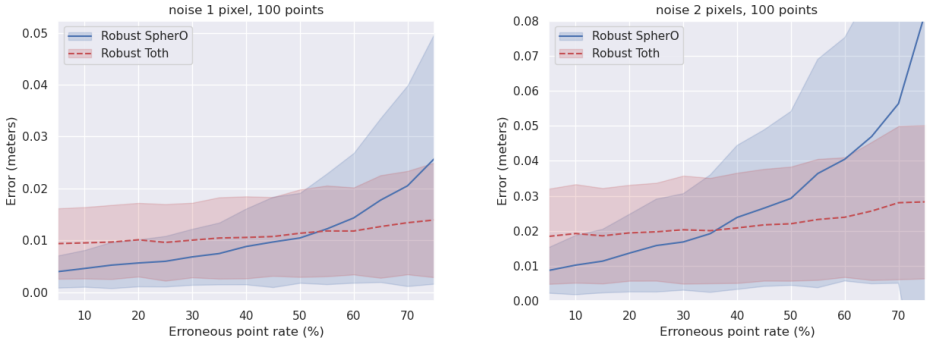

Figure 3: Mean and standard deviation of errors versus erroneous point rates with 1 pixel noise (left) and with 2 pixels noise (right). Erroneous points do not belong to the sphere's occluding contour.

noise and 20% erroneous points (right). `Robust-SpherO` outperforms `Robust-Toth`, though a breakpoint appears at 50% and 40% of occlusions, respectively.

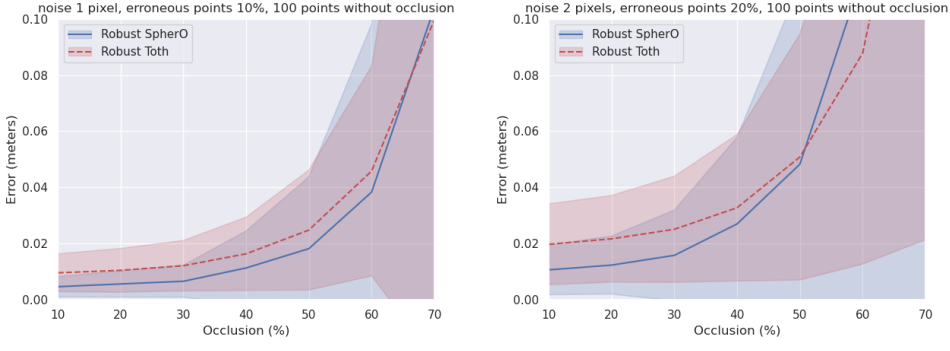

Figure 4: Mean and standard deviation of errors versus occlusion with 1 pixel noise and 10% erroneous points (left) and with 2 pixels noise and 20% erroneous points (right).

Accuracy versus depth. We varied the sphere's depth along the camera's optical axis from 1 to 10 metres with a step size of 1 metre and using 100 points in each trial with 1 pixel and 2 pixels white Gaussian noise, respectively. The error substantially increases beyond 4 metres for both methods. `Robust-SpherO` outperforms `Robust-Toth` with errors consistently at least twice as low.

Synthesis. `Robust-SpherO` is almost always substantially more accurate than its contender `Robust-Toth`. Nonetheless, there exist extreme cases, where the number of erroneous points or the level of occlusion are very large, for which the performance order reverts. `Robust-SpherO` uses 3D points to form a plane's consensus set while `Robust-Toth` uses 2D points to form an ellipse's consensus set. Both methods use the same RANSAC threshold, which is an in-plane distance. `Robust-SpherO` thus includes more points in its

consensus set than `Robust-Toth` because some out-of-plane 3D points remain within the range of the RANSAC threshold although they would be out of range if projected onto the consensus plane. The higher the percentage of erroneous points, the higher the probability of `Robust-SpherO` to include a higher percentage of erroneous points in its consensus set. Yet, `Robust-SpherO` (*i.e.*, plane fitting) reconstructs stably up to a large rate of erroneous points. These extreme cases are however unlikely to happen in real scenarios.

4.1.2 Robustness with Non-elliptic Occluding Contours

`Robust-Toth` is not designed to handle the parabola (*i.e.*, $\mathbf{C}^\top \underline{\mathbf{z}} = R$) nor the hyperbola (*i.e.*, $\mathbf{C}^\top \underline{\mathbf{z}} < R$), only the ellipse. Subsequently, a parabola yields an `Inf` through division by zero in equation (21) in [8] by a projection parallel to the normalised image plane for a point which otherwise would yield an endpoint of the major axis of an ellipse. However, `Robust-Toth` rarely breaks down in practice for a parabola. This is because a mild noise on a sample of the contour points can easily shift the fit to an ellipse well aligned with the contour points. On the other hand, in a hyperbola case, an endpoint of the major axis of an ellipse is projected to the opposite side of the image from the backside of the camera in equation (21) in [8]. This yields a flipped and scaled ellipse away from the contour points. Subsequently, `Robust-Toth` produces wrong results with a hyperbola. We formed a parabolic occluding contour by placing the sphere centre at $\mathbf{C} = [1.2, 0, 1]^\top$ with radius $R = 1$ in metres. We formed a hyperbolic occluding contour by placing the sphere centre at $\mathbf{C} = [0, -1.2, 0.8]^\top$ again with radius $R = 1$ in metres. We made 10 trials on each occluding contour with 1 pixel image noise and 5% erroneous points. In each trial, the erroneous points were randomly chosen over 100 points which were also randomly chosen from the occluding contour. Figure 5 presents the results. `Robust-SpherO` substantially outperforms `Robust-Toth`.

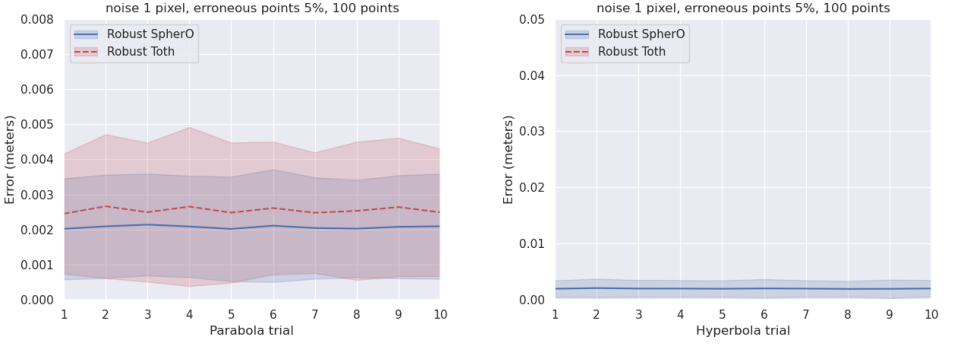

Figure 5: Mean and standard deviation of errors for the parabola (left) and hyperbola (right) occluding contours with 1 pixel noise and 5% erroneous points. For the hyperbola case (right), `Robust-Toth` produced out-of-bound results with a mean of about 0.42 metres.

4.2 Real Data Experiments

4.2.1 Relative Accuracy Evaluation

We evaluate the relative accuracy of the compared methods by reconstructing multiple spheres whose centre distances are known. We place a target sphere with known radius R in con-

tact with at least three auxiliary spheres whose radii R_i are also known. This yields the ground-truth distances between the target sphere's centre $\mathbf{C}$ and the auxiliary spheres' centres $\mathbf{C}_i$ as $\|\mathbf{C}-\mathbf{C}_i\| = R+R_i$. Three auxiliary spheres fully constrain the target sphere's centre $\mathbf{C}$, although two solutions exist. We use the compared methods to reconstruct the respective spheres' centres $\widehat{\mathbf{C}}$ and $\widehat{\mathbf{C}}_i$ from which a pairwise error is computed as $e_i = |\,\|\widehat{\mathbf{C}}-\widehat{\mathbf{C}}_i\| - (R+R_i)\,|$ for each method. Figure 6 reports the results with 3 auxiliary spheres. The relative error is computed as $\|\mathbf{e}\| = \|[e_1,e_2,e_3]^\top\|$. The target sphere's radius was 8.5 *cm* (football) and the auxiliary spheres' radii were 3.25 *cm* (tennis ball), 3.5 *cm* (green ball) and 4.75 *cm* (rainbow ball), respectively. The target sphere's centre was about 40 *cm* away from the camera. We set the RANSAC threshold $\tau_{pix} = 0.5$ pixels which was about twice of the intrinsic calibration accuracy of the camera. We repeated the experiment 1000 times with both methods. `Robust-SpherO` yielded 4.76 *mm* mean error with 0.01 *mm* std. `Robust-Toth` yielded 6.77 *mm* mean error with 1.9 *mm* std. We also evaluate the relative

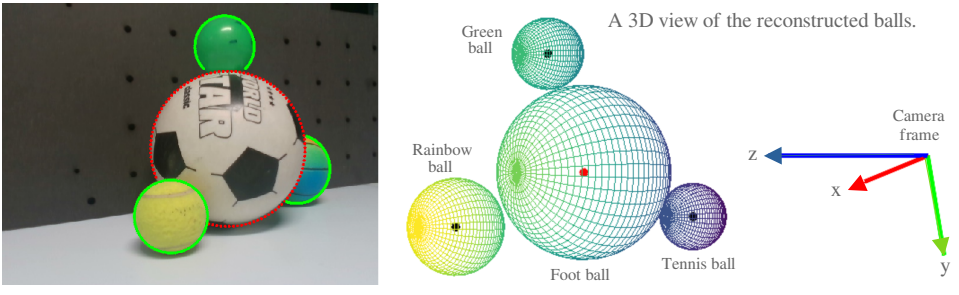

Figure 6: The input image with the spheres' occluding contours (left). A 3D view from behind the balls reconstructed by `Robust-SpherO` (right). The rainbow ball was mostly occluded by the football. We used 100 contour points for each ball.

accuracy of the compared methods at different depths. Figure 7 shows the relative reconstruction errors in millimetres of the football in contact with a tennis ball, a small basket ball and a yellow ball. The balls are placed at 1 metre, 2 metres and 3 metres away from the camera, respectively. `Robust-SpherO` outperformed `Robust-Toth` with 5.9 *mm*, 84.6 *mm* and 231.1 *mm* mean errors versus 6.9 *mm*, 91 *mm* and 232.2 *mm* at 1 metre, 2 metres and 3 metres, respectively. We observe that the reconstruction accuracy mostly depends on the quality of the segmented occluding contours. The yellow ball is mostly occluded and segmenting its occluding contours at 2 and 3 metres was significantly difficult. This yielded the worst relative errors for both methods.

	Robust-SpherO	Robust-Toth
$e_{tennisball}$	2.0 ± 1.4	2.0 ± 1.4
$e_{basketball}$	5.6 ± 0.0	6.6 ± 1.0
$e_{yellowball}$	0.3 ± 0.0	0.3 ± 0.0

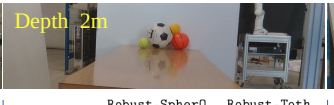

	Robust-SpherO	Robust-Toth
$e_{tennisball}$	3.5 ± 1.5	3.5 ± 1.5
$e_{basketball}$	0.7 ± 0.0	2.9 ± 2.6
$e_{yellowball}$	84.6 ± 0.0	90.9 ± 12

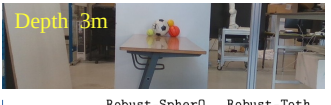

	Robust-SpherO	Robust-Toth
$e_{tennisball}$	5.2 ± 0.2	5.2 ± 0.2
$e_{basketball}$	4.5 ± 0.0	23 ± 15
$e_{yellowball}$	231 ± 0.0	231 ± 0.4

Figure 7: Mean and standard deviation of relative reconstruction errors (*mm*) of the football with respect to auxiliary balls. The experiments are repeated 1000 times.

4.2.2 Absolute Accuracy Evaluation

We used an Intel RealSense 3D camera to measure 3D surface points of the spheres from which the spheres' centres $\mathbf{C}_i$ were computed. We used the methods to reconstruct the sphere centres $\widehat{\mathbf{C}}_i$. We then computed an absolute error $e_i = \|\widehat{\mathbf{C}}_i - \mathbf{C}_i\|$ per sphere. We used the image shown in figure 6 for the experiment, which was taken by the depth camera. We repeated the experiment 1000 times with both methods. Table 2 lists the mean and std of the errors in millimetres with respect to the 3D camera measured centres. We observe that the

Table 2: Reconstruction error of the four balls (mean $\pm$ std in *mm*).

	Football	Tennis ball	Green ball	Rainbow ball
`Robust-Toth` [8]	3.3 ± 0.0	4.1 ± 1.9	8.2 ± 1.6	15.2 ± 2.5
`Robust-SpherO`	3.3 ± 0.0	1.9 ± 0.0	4.2 ± 0.0	7.4 ± 0.0

relative accuracy, as explained in section 4.2.1, computed from the spheres' centres measured by the 3D camera was $11.4\,mm$. However, we also observe that the reconstructed spheres are close to their absolute true locations. `Robust-SpherO` outperforms `Robust-Toth` by a large margin, having a lower absolute error, both on average and in standard deviation.

In some of the 1000 repetition experiments, the standard deviation values computed from the reconstruction errors of `Robust-SpherO` remain 0 because in every trial the consensus set included all the occluding contour points of a ball for the given RANSAC threshold.

4.3 Runtimes

The mean and standard deviation of the runtimes combining the synthetic and real data experiments for `Robust-SpherO` and `Robust-Toth` were $7.5 \pm 0.7\,ms$ and $13 \pm 0.5\,ms$, respectively. `Robust-SpherO` thus brings an overall speed up factor of approximately 1.5 compared to `Robust-Toth`.

5 Conclusion

We have proposed `SpherO`, a method to reconstruct a sphere of a prescribed radius from a single calibrated view of its occluding contour. It is accurate, simple and fast. Its robustified version `Robust-SpherO` outperforms the state of the art. In addition, `SpherO` provides a closed-form, convex solution for the image conic from the sphere's occluding contour points, minimising a geometric error criterion. It may also be possible to find a theoretical error bound for `SpherO`, which is however a non-trivial problem. As future work, we will improve `SpherO` and `Robust-SpherO` by replacing the plane-fitting step by an optimal space-circle-fitting.

Acknowledgement

This work is funded by the projects ANR JCJC - IMMORTALLS and Interreg Sudoe Programme (European Regional Development Fund) - REMAIN (S1/1.1/EO111).

References

[1] M. Chasles and C. Graves. *Two geometrical memoirs on the general properties of cones of the second degree, and on the spherical conics*. Dublin : For Grant and Bolton, 1841.

[2] O. Edlund. *Some Notes on Least Squares, QR-factorization, SVD and Fitting*. Lecture notes on Numerical Analysis, Luleå University of Technology, 2013.

[3] R. Halir and J. Flusser. Numerically stable direct least squares fitting of ellipses. In *WSCG98*, 1998.

[4] C. Lu, S. Xia, M. Shao, and Y. Fu. Arc-support line segments revisited: An efficient high-quality ellipse detection. *IEEE Transactions on Image Processing*, 29:768–781, 2020.

[5] C. Meng, Z. Li, X. Bai, and F. Zhou. Arc adjacency matrix-based fast ellipse detection. *IEEE Transactions on Image Processing*, 29:4406–4420, 2020.

[6] K. Shi, X. Li, and H. et al. Xu. Sphere localization from a minimal number of points in a single image. In *Proceedings of the international conference on advances in computer technology, information science and communications–Volume 1: CTISC, INSTICC. SciTePress*, pages 65–70, 2019.

[7] J. Sun, H. He, and D. Zeng. Global calibration of multiple cameras based on sphere targets. *Sensors*, 16, 2016.

[8] T. Toth and L. Hajder. A minimal solution for image-based sphere estimation. *International Journal of Computer Vision*, 131:1428–1447, 2023.