

Towards Dynamic Camera Calibration for Constrained Flexible Mirror Imaging

Aubrey K. Dunne, John Mallon, Paul F. Whelan

► **To cite this version:**

Aubrey K. Dunne, John Mallon, Paul F. Whelan. Towards Dynamic Camera Calibration for Constrained Flexible Mirror Imaging. The 8th Workshop on Omnidirectional Vision, Camera Networks and Non-classical Cameras - OMNIVIS, Oct 2008, Marseille, France. 2008. <inria-00325391>

HAL Id: inria-00325391

<https://hal.inria.fr/inria-00325391>

Submitted on 29 Sep 2008

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Towards Dynamic Camera Calibration for Constrained Flexible Mirror Imaging

Aubrey K. Dunne, John Mallon, and Paul F. Whelan

Vision Systems Group
Dublin City University
aubrey.dunne@eeng.dcu.ie

Abstract. Flexible mirror imaging systems consisting of a perspective camera viewing a scene reflected in a flexible mirror can provide direct control over image field-of-view and resolution. However, calibration of such systems is difficult due to the vast range of possible mirror shapes and the flexible nature of the system. This paper proposes the fundamentals of a dynamic calibration approach for flexible mirror imaging systems by examining the constrained case of single dimensional flexing. The calibration process consists of an initial primary calibration stage followed by in-service dynamic calibration. Dynamic calibration uses a linear approximation to initialise a non-linear minimisation step, the result of which is the estimate of the mirror surface shape. The method is easier to implement than existing calibration methods for flexible mirror imagers, requiring only two images of a calibration grid for each dynamic calibration update. Experimental results with both simulated and real data are presented that demonstrate the capabilities of the proposed approach.

1 Introduction and Background

Flexible mirror imaging systems, recently introduced in [1], are non-conventional catadioptric cameras typically consisting of a perspective camera viewing a scene reflected in a bendable mirror. The characteristics of the camera image can be continuously and non-linearly altered by directly flexing the mirror. This enables control over both the field-of-view of the image and the image resolution. By suitable flexing of the mirror, certain portions of the scene being imaged can be attributed more sensor resources than other portions, and the field-of-view can be easily altered as required. Flexible mirror imagers may thus be beneficial in active vision monitoring and security applications, where higher resolution could be obtained in image areas containing objects of interest without sacrificing field-of-view (as would normally be the case in pan-tilt-zoom camera systems).

The introduction of flexible mirror imaging systems naturally leads to the requirement for methods for their calibration, so that tasks such as tracking in surveillance can be accomplished. Calibration of flexible mirror imaging systems is difficult firstly because the mirror deflection is generally unconstrained, and secondly because the camera configuration, and thus calibration, alters each time

that the mirror is flexed. The calibration is intrinsically linked to the estimation of the flexible mirror surface itself, since knowledge of local mirror surface shape allows surface normals and thus reflected rays to be determined for given incident rays. However, since the mirror shape can change, catadioptric calibration methods that assume some prior knowledge of mirror shape, such as [2][3][4], are not applicable, and rather methods that can estimate the mirror shape are required. Much work has been presented in the literature on the recovery of surface shape from images of a diffuse surface, using either structured light, or epipolar geometry. The recovery of specularly reflective surface shape has received relatively less attention from the vision community. This situation differs from the diffuse case since 'features' seen in the image are virtual features caused by the reflection, which do not obey the epipolar constraint. Halstead et al. [5] presented one of the first techniques for determining specular surface shape from surface reflection in a single image. They used a bespoke conical calibration object and concentric camera to determine corneal surface estimates. Savarese et al. [6] describe a method for specular surface recovery from a single image of a planar calibration target when at least two local orientations are available at each target point. Several approaches to estimating mirror shape from motion have also been presented. Swaminathan et al. [7] examine the dependence on surface geometry of specularities in static scenes with constant velocity camera motion, whilst Roth et al. [8] go further by estimating the specular surface geometry from specular flow and using the result to improve surface estimation from diffuse flow. Oren et al. [9] recover information on surface profile from specular reflections whose paths overlap under constrained camera motions. However, all of these methods assume a fixed mirror shape with relative mirror scene motion, and so are not suitable for flexible imaging systems calibration.

Calibration methods for the general camera model such as [10] and [11] can be used to determine reflected ray directions and thus calibrate any catadioptric camera, but these methods are time consuming and when applied to a flexible mirror imager would require repeating the entire calibration process each time the mirror is flexed. Gonçalves and Araújo [12] also present a method for calibrating a catadioptric system consisting of a quadric mirror surface from a single image of a calibration target. A completely non-parametric method for specular surface estimation based on voxel carving was presented in [13], in which normal vectors are accumulated for each scene voxel, and the voxels with the normals in best agreement are considered to be on the specular surface. The method requires images of the reflective surface from many different viewpoints to achieve good results. All the above methods assume a fixed specular surface geometry. If they were to be applied to calibrate a flexible imaging system, the system would have to be completely recalibrated after each mirror flexion, which is not practical in the envisaged system.

Kuthirummal and Nayar [1], who introduced the concept of flexible mirror imaging with a nominally planar mirror surface, presented the only calibration method specifically for such systems of which the authors are aware. Their calibration is performed offline by acquiring an image of the mirror boundary, and

an accompanying mirror surface measurement, for each mirror deformation that is likely to be imposed. They present results for a calibration with greater than 30,000 image-estimate pairs. The one-to-one mapping between the mirror surface shape and a descriptor of its outline is then stored in a look-up table, so that once the mirror outline is visible in the image the calibration is available.

This paper focuses on the constrained case of a flexible mirror fixed along one edge that allows flexing in the plane perpendicular to this fixed edge. In practice, this system is easy to implement as it requires only a single point of deflection, yet it allows variation in field-of-view and image resolution across one dimension of the image. Such a setup has potential application in surveillance systems, where variations in the horizontal field-of-view are important but a fixed vertical field of view is typically adequate. The major contribution of this paper is the introduction of a novel method of calibration for this type of constrained flexible mirror imaging system. The principals of an alternative calibration method to [1] are established, where the aim is to dynamically update the calibration rather than attempt to completely recalibrate after each mirror deflection.

The proposed calibration method takes its inspiration from the field of mirror design for catadioptric systems. Swaminathan et al. [14] developed a linear method for determining the catadioptric mirror surface shape necessary to implement a desired scene to image map. The method uses constraints on the incident and reflected ray directions to determine the parameters of a B-spline surface model of the mirror. Similarly to [5], the surface is determined directly from the set of surface normals. This method is adapted for the purpose of calibration in this paper, so that in conjunction with scene information obtained by imaging a calibration grid both before and after the mirror deflection, the new mirror surface can be estimated and thus the calibration determined. In comparison to the calibration method of [1], the proposed method has the significant advantage of a much simpler calibration setup. It does not require a stereo rig for surface shape estimation, nor is there a large and laborious data collection procedure required. The proposed method is adaptive where [1] is exhaustive. The mirror shapes that it can calibrate are not limited to the set of pre-calibrated deformations nor is there a requirement for the entire mirror boundary to be visible in the image, although unlike [1] the presented method can not perform calibration in real time.

Section 2 presents the camera model and the method for mirror shape estimation. In Section 3 the calibration process, consisting of an offline primary stage and a dynamic stage that operates to update the calibration as the mirror flexes, are presented. Results for both simulated and real data demonstrating the performance of the calibration method are detailed in Section 4. Finally, directions for future work and conclusions are presented.

2 Camera Model

This paper deals with a constrained type of flexible mirror imaging system, in which a perspective camera images a scene reflected in a thin, flexible, nominally

planar mirror surface, as used in [1]. Furthermore, one edge of the flexible mirror is assumed to be fixed in position and in first derivative relative to the perspective camera, i.e. the mirror is clamped at one edge. It will be seen that this constraint removes the ambiguity in the location of the mirror in the world coordinate system. Flexing can occur in a direction perpendicular to the edge, and so the mirror can be modelled by a 2D curve in the plane. For simplicity it is assumed that the fixed edge of the mirror is vertical with respect to the perspective camera viewing it, and is perpendicular to the camera’s principal axis.

Reflection of light rays from specular surfaces is governed by the law of reflection, which states that the angle of incidence is equal to the angle of reflection. Knowledge of the surface normal at any point on the surface allows the reflected ray to be calculated for any incident ray. Referring to Fig. 1, reflected light rays




Fig. 1. Flexible mirror imaging system.

entering the perspective camera are termed viewing rays and are represented by unit vectors V_i . The corresponding light rays that fall on the mirror from the scene are called scene rays and are represented by unit vectors V_r . For any given mirror surface shape, the flexible imager can be modelled by the general camera model, which maps image pixels (u, v) to scene ray directions V_r via a look up table. This model is non-parametric and so can effectively cope with the non-linearity and non-centrality of the flexible mirror imager.

2.1 Mirror Shape Model

The mirror surface is modelled by a B-spline curve with quadratic basis functions and an open uniform knot vector. The mirror spline model is

$$D(u) = \sum_{i=1}^{K_f} c_i f_i(u) \quad (1)$$

$$= C^T F(u), \quad (2)$$

where u is the horizontal image pixel coordinate, $f_i(u)$ are the spline basis functions, and K_f is the number of basis functions. $K_f = 5$ in this paper. As in [14, 1], the complete mirror surface is described parametrically by the distances of the mirror points from the perspective camera centre, given by the spline model, as measured along the viewing rays V_l of the primary optics. The perspective camera centre is assumed without loss of generality to be at the world coordinate origin. The mirror surface, $S(u)$, is therefore given by

$$S(u) = D(u)V_l(u) . \quad (3)$$

Note that for the experiments with real data in Section 4.2, an additional B-spline that is a function of the vertical image pixel coordinate, and that has a single quadratic basis function, is incorporated into the mirror surface model. This is used to remove the vertical curvature that would otherwise result from Eqn. 3. The modification also accounts for any misalignment that exists in the experimental setup between the fixed mirror edge and the perspective camera, where the fixed mirror edge may not be perpendicular with respect to the camera principal axis.

2.2 Mirror Shape Estimation

The mirror surface is estimated from a set of linear equations as in [14]. For known V_l s and V_r s, the surface normals, $N(u)$, are given by

$$N(u) = \frac{V_l(u) - V_r(u)}{|V_l(u) - V_r(u)|} . \quad (4)$$

The tangents to the surface are given by the first derivative of the surface. Enforcing orthogonality between normals and the tangents leads to the following

$$\frac{dS(u)}{du} \cdot N(u) = 0 . \quad (5)$$

Combining Eqns. 3 and 5 gives

$$\left(D \frac{\partial V_l}{\partial u} + V_l \frac{dD}{du}\right) \cdot N(u) = 0 . \quad (6)$$

Expanding out with the B-spline basis functions and coefficients gives

$$\frac{dV_l}{du} \cdot N(u)(C^T F(u)) + V_l \cdot N(u)(C^T \frac{dF(u)}{du}) = 0 , \quad (7)$$

which is conveniently rewritten as

$$\left[\left(\frac{dV_l}{du} \cdot N(u) \right) F(u)^T + (V_l \cdot N(u)) \frac{dF(u)^T}{du} \right] C = 0 . \quad (8)$$

From the pinhole model for the perspective camera V_l is given by $\frac{[u \ f]^T}{|[u \ f]|}$, where f is the camera focal length, and thus

$$\frac{dV_l}{du} = \frac{1}{(u^2 + f^2)^{\frac{3}{2}}} \begin{bmatrix} f^2 \\ -uf \end{bmatrix}. \quad (9)$$

Each image point and corresponding mirror surface normal allow an equation of the form of Eqn. 8 to be formed. By stacking at least K_f such equations into a matrix a homogeneous equation of the form $Ax = 0$ is obtained. This can then be solved using the SVD with equilibration to determine the B-spline coefficients. Note that A must contain at least one row for an image point in each knot interval.

2.3 Mirror Edge Constraint

The above solution for the B-spline coefficients is up to scale. For complete calibration it is necessary to remove the scale ambiguity. The fixed mirror edge constraint imposed in Section 2 provides the required constraint and so must be incorporated into the mirror shape estimation equations.

Assume that a solution for the mirror surface and scale is available for some position and deflection of the mirror (see Section 3). In this case, for an image pixel, u_e , viewing the fixed edge of the mirror the spline value $D(u_e)$ is known. This value is constant for all mirror deflections so that after a deflection of the mirror the spline equation associated with u_e is

$$D(u_e) = F'(u_e)^T C', \quad (10)$$

where the prime indicates new spline coefficients and basis functions. A least squares solution for C' can thus be extracted in a subspace spanned by $K_f - 1$ basis vectors as

$$C' = C_P + C_B \Phi, \quad (11)$$

where C_P is the particular solution, C_B is a matrix of basis vectors, and Φ is the new vector of unknowns. Finally, C' is incorporated into Eqn. 8, and after some rearrangement the new equation to be solved that incorporates the edge constraint is

$$GC_B \Phi = -GC_P, \quad (12)$$

where

$$G = \left(\frac{\partial V_l}{\partial u} \cdot N(u) \right) F'(u)^T + (V_l \cdot N(u)) \frac{dF'(u)^T}{du}. \quad (13)$$

A solution to Eqn. 12 can be determined using standard least squares techniques.

3 Calibration

The goal of calibration is to determine the mapping between viewing rays V_l and scene rays V_r . This mapping is given by the mirror surface, so that calibration reduces to the estimation of the mirror surface. The calibration consists of

two stages: (1) calibration of primary optics and determination of initial mirror surface, and (2) dynamic updating of calibration after the mirror is flexed. The primary calibration stage is only executed once, whereas the dynamic stage is applied after each mirror flexion.

3.1 Primary Calibration

Primary calibration aims is to determine a complete calibration of the imager for some initial mirror position and shape. It employs two basic calibration techniques. Firstly the perspective camera is calibrated using a standard method such as [15]. This information is used to both remove radial distortion from all subsequent perspective camera images, and to enable the V_l s, which remain fixed, to be calculated as described in Section 2.2. The complete flexible mirror imager is then calibrated as a general camera for the initial mirror shape and position using general calibration [10], which gives the mapping between u and V_r for that configuration. Although general calibration is a time consuming process, it only needs to be carried out once. Interestingly, the primary calibration stage could be performed in a single step, by calibration of the perspective camera, if the initial mirror shape is known to be planar. Once the V_l s and V_r s are known, their intersection points on the mirror surface are estimated so that the mirror position can be determined. The position estimate is found by selecting the points that minimise the distance to V_l and V_r for each corresponding pair of rays.

3.2 Dynamic Calibration

Any deflection of the mirror requires the calibration to be updated, which is equivalent to estimating the updated mirror shape, $S(u)'$. In Section 2.2 it is shown that knowledge of the scene and viewing rays is sufficient for reconstructing the mirror surface. The viewing rays V_l' for $S(u)'$ can be directly determined from the location of the features in the second image as in Section 2.2, but the V_r' s are unknown. To determine them some information about the scene is required, and this is acquired by imaging a calibration grid, placed in the imager field-of-view, both before and after mirror deflection. The pose of this grid is estimated from the first image, which has already been calibrated (either in primary calibration stage or in the previous application of dynamic calibration). A non-central pose estimation technique, such as [11], is necessary to estimate the pose. The estimate of the calibration grid pose allows the 3-space locations of the grid corner scene points, $P(u)$, to be determined.

Referring to Fig. 1, it is clear that if the scene points $P(u)$ are at infinity, then V_r' is equal to V_r for the corresponding image points. This infinite scene assumption allows us to approximate V_r' as V_r for distant scene points, and enables a linear estimation of the mirror surface using Eqns. 3, 5 and 12. The linearly estimated result, $\hat{S}(u)'$, is used to initialise a non-linear minimisation to determine the final mirror surface estimate, $S(u)'$. For the minimisation V_r' is

represented as

$$V_r'(u) = P(u) - \hat{S}(u)' \quad (14)$$

$$= P(u) - \hat{D}(u)' V_i(u)' \quad (15)$$

so that the estimated surface normal, $\hat{N}(u)$, is

$$\hat{N}(u) = V_i(u)' - [P(u) - \hat{D}(u)' V_i(u)'] . \quad (16)$$

Non-linear minimisation is then applied to minimise the following cost function

$$\Gamma(D(u)') = \frac{d\hat{S}(u)'}{du} \cdot \frac{\hat{N}(u)}{|\hat{N}(u)|} \quad (17)$$

This gives the solution, $S(u)'$, for which the viewing ray for the image of each scene point $P(u)$ maps to the scene ray passing through that point. Once the calibration has been updated ($S(u)'$ determined), the mirror can be flexed again and dynamic calibration reapplied using the V_r 's estimated in the last iteration. Note that although there is an inherent bias due to the infinity assumption, it is found that for small mirror deflections the linear estimate is sufficient for the convergence of the non-linear estimation. Large mirror deflections should be carried out in stages, with dynamic calibration applied at each stage.

Fig. 2 a flow diagram outlining the proposed calibration process.




Fig. 2. Proposed flexible mirror imaging calibration method.

4 Experimental Results

Experiments were carried out using simulated data, to evaluate the effect on performance of configuration variations, and using real data, so as to characterise the overall performance of the calibration method.

4.1 Simulated Experiments

Experiments with simulated data were conducted for variations in image noise, scene depth (distance between mirror and calibration grid) and mirror curvature. The simulation results are presented in Fig. 3, where the datapoints in each plot are the average of 100 random trials. Row 1 shows the error plots for increasing




Fig. 3. Error in linear and final mirror surface estimates for varying parameters. The simulated camera has focal length 26mm and image size 516×656 pixels. The flexible mirror is 420mm in length and is oriented at 67.5° to the camera. Unless otherwise stated, the mirror deflection is 16.7mm , the scene depth is 1650mm and noise with 1 pixel standard deviation is added.

additive Gaussian image noise. Additive image noise has a greater impact on the final surface estimate than on the linear estimate, since the error due to the infinite scene assumption far outweighs the error introduced by the noise.

The final surface estimate error is linearly proportional to the noise. The error plots for increasing scene depth are shown in row 2. As expected, the error decreases as the calibration grid moves further from the mirror, in line with the infinite scene assumption. Note that for large scene depths ($> 3000mm$) the final error becomes independent of the scene depth (infinity assumption is effective approximation). The error plots for increasing mirror curvature, κ , are shown in row 3. It can be seen that the error in the linear surface estimate increases with increasing curvature due to increasing weakness of the infinite scene assumption. The error in the final mirror surface estimate is larger than the maximum error for either of the other two simulated experiments principally due to the poor linear estimate.

4.2 Real Experiments

The calibration method was evaluated for a real flexible imaging system set up as described in Section 2. A flexible plastic mirror¹ was attached to a sheet aluminium substrate in order to improve rigidity. Deflections were applied to the back of the substrate using optical translation stages with $0.01mm$ resolution.

In the first set of experiments, high-accuracy comparative reference measurements of the mirror surface were obtained using a 3D laser scanner². Fig. 4 shows the estimated and reference mirror surface before and after a mirror deflection of approximately $7.98mm$. It can be seen that the primary calibration error is of similar magnitude to the error for the final mirror surface estimate after flexing, indicating that the accuracy of the dynamic calibration is very sensitive to the primary calibration. The primary calibration error can be attributed to the higher order components of the mirror surface, caused by mirror defects, that are not modelled by Eq. 2. The distance from the mirror surface to the calibration grid was approximately $280mm$, and as a result the linear estimate of the flexed mirror surface in Fig. 4 is relatively inaccurate, although the final estimate is still good.

A second set of experiments were conducted with the same flexible mirror imager configuration as above in order to separately assess the primary and dynamic calibration stages. The overall accuracy of the primary calibration and the effect of the surface estimation using B-splines were evaluated by comparing directly measured scene rays with scene rays estimated from the model. Approximately 160 pixels with corresponding V_I s were chosen uniformly from the image, and the errors in the angles between the measured V_r s and the V_r s determined by reflecting these V_I s in the estimated mirror surface were measured. The mean and standard deviation of these errors were 0.5552° and 0.3285° , respectively. Note that all angular errors are measured in the horizontal plane (u direction), in accordance with Eqn. 3.

After primary calibration, the mirror was deflected in three stages, with a $5mm$ deflection at each stage, and dynamic calibration was applied at each stage.

¹ Plastic wing mirror replacement mirror from www.carpointeurope.com

² The non-specular back surface of the mirror was scanned instead of its specular surface, since a mirrored surface cannot be accurately measured by a laser scanner.




Fig. 4. Mirror surface estimates for real data before and after flexing.

Assessment of the accuracy of each dynamic calibration was made by again comparing directly measured V_r s with V_r s estimated from the calibrated model. An image of an active grid (spatio-temporally varying calibration patterns displayed on a flat-screen monitor [16]) was captured by the imager. The grid was then translated 100mm away from the mirror and imaged again. At each position, the world pose of the active grid was determined by reflection in a planar mirror [17]. Knowing the two poses of the active grid, and decoding the active grid data, the V_r s for any image pixel can be determined. For 168 uniformly selected points in the image, their associated directly determined V_r s were compared to the V_r s estimated from the calibration data. The magnitudes of the angular errors are shown in Fig. 5, where the maximum directly determined ray deflections for stages 1, 2 and 3 were 4.1366° , 8.4113° , and 12.9171° respectively. The mean and standard deviation of the angular errors, and the estimated deflection of the mirror determined from the calibration data, are shown in Table 1 for each stage of deflection. The estimated deflection values show good agreement with the actual deflections, with the maximum error being 7.76%. The reprojection errors at each stage for 160 corner points on a $252 \times 261\text{mm}$ chessboard grid are also shown in Table 1. Note that the magnitude of the reprojection error is dependent on the distance of the calibration grid from the mirror (scene depth). Nevertheless, there is no apparent trend in any of the results in Table 1 to indicate that there is significant cumulative error in the dynamic calibration.

Finally, a distortion correction experiment was performed for the calibrated camera. Exactly perspectively correcting an image of a non-planar scene taken

Table 1. Measured and estimated mirror deflections, and u direction angular errors between estimated and directly determined V_r s, for each deflection stage (SD = standard deviation).

Stage No	Deflection magnitude (mm)		Angular error ($^\circ$)	Reprojection error (mm)
	Measured	Estimate	Mean (SD)	Mean (SD)
1	5.0	4.7104	0.3578 (0.1751)	0.5122 (0.3027)
2	5.0	4.6121	0.1088 (0.0833)	0.4446 (0.2486)
3	5.0	4.7414	0.2318 (0.1477)	0.3992 (0.2582)




Fig. 5. Angular errors in the calibrated V_r s across the image at each stage of deflection. Quadrilateral on ground plane indicates mirror boundary in image. Fixed mirror edge is rightmost side of quadrilateral.

by a flexible mirror imager is not possible in the general case since the image is multi-perspective. However, by selecting a suitable synthetic plane a near-perspective image can be formed. Fig. 6 shows a face image from the flexible mirror imager both before and after such approximate perspective correction, in which the non-linear distortion of the face is seen to have been removed.

5 Future Directions

The calibration method presented in this paper applies to a constrained flexible mirror imager in which the mirror is deflected in only one dimension. Current work is focussing on generalising the framework to allow the imager to be calibrated for unconstrained deflections. Further work will examine auto-calibration for flexible mirror imagers, whereby existing scene points rather than calibration grid points are used in the dynamic calibration stage. This is a difficult problem as the unknown depth of scene points introduces a further degree of freedom into the calibration task.




Fig. 6. a) Face image from flexible mirror imager after 12mm deflection; b) Near-perspective face image with non-linear distortion removed using calibration data. Note the large distortion at top right of the image that was not captured by the calibration

6 Conclusion

Flexible mirror imaging systems provide unique benefits over existing imagers because they allow the field-of-view and the image resolution to vary within the image. However, the only existing method for specifically calibrating such systems operates offline by directly acquiring 3D mirror surface shape measurements for a very large set of possible mirror deformations. This paper investigates the fundamentals of a novel alternative method for the calibration of flexible imaging systems that does not require a significant offline stage. The proposed method facilitates the dynamic update of the current calibration after each mirror deflection. The flexible mirror, fixed in position along one edge and flexed in a single dimension, is modelled by a B-spline curve. A primary calibration stage determines the optical configuration for an initial mirror position. After the mirror flexes, dynamic calibration, requiring only a single image of a planar grid in each deflected mirror position, is applied to update the calibration. Simulated results demonstrate the performance of the method under variations in image noise, scene depth and mirror curvature, and validate the use of the linear estimate for initialising the non-linear minimisation process. Experiments with a real flexible imaging system demonstrates good results when compared to ground truth data. The presented principles of dynamic flexible mirror calibration and the proposed calibration method are important steps towards complete unconstrained calibration of flexible imaging systems.

Acknowledgements

This research is funded by the Irish Research Council for Science, Engineering and Technology: funded by the National Development Plan.

References

1. Kuthirummal, S., Nayar, S.K.: Flexible mirror imaging. In: OMNIVIS07, Rio de Janeiro, Brazil. (2007) 1–8
2. Tardif, J., Sturm, P., Roy, S.: Self-calibration of a general radially symmetric distortion model. In: Proceedings of the 9th European Conference on Computer Vision, Graz, Austria. Volume 4 of Lecture Notes in Computer Science., Springer (2006) 186–199
3. Thirithala, S., Pollefeys, M.: Multi-view geometry of 1d radial cameras and its application to omnidirectional camera calibration. In: ICCV '05: Proceedings of the Tenth IEEE International Conference on Computer Vision, IEEE Computer Society (2005) 1539–1546
4. Scaramuzza, D., Martinelli, A., Siegwart, R.: A Toolbox for Easily Calibrating Omnidirectional Cameras. In: Proc. 2007 IEEE International Conference on Intelligent Robots and Systems. (2006)
5. Halstead, M.A., Barsky, B.A., Klein, S.A., Mandell, R.B.: Reconstructing curved surfaces from specular reflection patterns using spline surface fitting of normals. In: Proc. 23rd Conf. Computer Graphics and Interactive Techniques, New York, USA (1996) 335–342
6. Savarese, S., Chen, M., Perona, P.: Local shape from mirror reflections. *IJCV* **64** (2005) 31–67
7. Swaminathan, R., Kang, S.B., Szeliski, R., Criminisi, A., Nayar, S.K.: On the motion and appearance of specularities in image sequences. In: ECCV (1). (2002) 508–523
8. Roth, S., Black, M.J.: Specular flow and the recovery of surface structure. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition, Washington, DC, USA, IEEE Computer Society (2006) 1869–1876
9. Oren, M., Nayar, S.K.: A theory of specular surface geometry. *IJCV* **24** (1997) 105–124
10. Grossberg, M.D., Nayar, S.K.: A general imaging model and a method for finding its parameters. In: Proc. 8th ICCV. Volume 2. (2001) 108–115
11. Sturm, P., Ramalingam, S.: A generic concept for camera calibration. In: Proc. ECCV, Prague, Czech Republic. Volume 2., Springer (2004) 1–13
12. Gonçalves, N., Araújo, H.: Low-cost method for the estimation of the shape of quadric mirrors and calibration of catadioptric cameras. *Optical Engineering* **46** (2007) 1 – 12
13. Bonfort, T., Sturm, P.: Voxel carving for specular surfaces. In: ICCV '03: Proceedings of the Ninth IEEE International Conference on Computer Vision, Washington, DC, USA, IEEE Computer Society (2003) 591
14. Swaminathan, R., Nayar, S.K., Grossberg, M.D.: Framework for designing catadioptric projection and imaging systems. In: In Proceedings IEEE Conf. on Computer Vision-PROCAMS, Nice, France. (2003)
15. Zhang, Z.: A flexible new technique for camera calibration. *IEEE Trans. PAMI* **22** (2000) 1330 – 1334
16. Dunne, A.K., Mallon, J., Whelan, P.F.: Efficient generic calibration method for general cameras with single centre of projection. In: Proceedings of the IEEE International Conference on Computer Vision, Rio de Janeiro, Brazil. (2007)
17. Bonfort, T., Sturm, P., Gargallo, P.: General specular surface triangulation. In: Proceedings of the Asian Conference on Computer Vision, Hyderabad, India. Volume II. (2006) 872–881