Automatic Analysis of Lens Distortions in Image Registration

Birgit Möller[1,2] and Stefan Posch[2]

[1] Underwater Vision Lab, Computer Vision & Robotics Group, University of Girona, 17003 Girona, Spain
[2] Pattern Recognition & Bioinformatics, Institute of Computer Science, Martin-Luther-University Halle-Wittenberg, 06120 Halle, Germany
{birgit.moeller,stefan.posch}@informatik.uni-halle.de

Abstract. Geometric image registration by estimating homographies is an important processing step in a wide variety of computer vision applications. The 2D registration of two images does not require an explicit reconstruction of intrinsic or extrinsic camera parameters. However, correcting images for non-linear lens distortions is highly recommended. Unfortunately, standard calibration techniques are sometimes difficult to apply and reliable estimations of lens distortions can only rarely be obtained. In this paper we present a new technique for automatically detecting and categorising lens distortions in pairs of images by analysing registration results. The approach is based on a new metric for registration quality assessment and facilitates a PCA-based statistical model for classifying distortion effects. In doing so the overall importance for lens calibration and image corrections can be checked, and a measure for the efficiency of accordant correction steps is given.

1 Introduction

A common task in modern computer vision and image sequence analysis is the *2D geometric registration* of image pairs. The principal goal of image registration is to estimate parameters of a suitable motion model that allows to relate two or more images to a common coordinate system. Applications include image mosaicing [1, 2], reconstruction of scene geometry [3], and navigation tasks [4].

Registration of two images relies on an appropriate, parameterised motion model to describe the camera motion. The choice of the model depends on the degrees of freedom of the camera and on the scene structure. Euclidean, affine or projective transformations are widely used [5]. For these motion models, 2D registration of two images does not require explicit reconstruction of either extrinsic or intrinsic camera parameters. However, the models usually assume real cameras as ideal pinhole projection devices, which is rarely true in real life. In particular, it has been shown that especially *non-linear lens distortions* significantly influence the quality of registration results (e.g., [6]). Accordingly, it is highly recommended to correct images for lens distortions prior to or during the registration process. Otherwise the registration results can seriously be degraded.

Proceedings of the ICVS Workshop on Camera Calibration Methods for Computer Vision Systems - CCMVS2007
Published in 2007 by Applied Computer Science Group, Bielefeld University, Germany

In literature a large variety of approaches to robustly calibrate lens distortion parameters are described. Basically, two different principal methodologies emerged over time. On the one hand, cameras and lenses are calibrated employing calibration patterns [7]. On the other hand specific structures in an image, such as straight lines, are used for parameter estimation [8]. Theses pattern- and structure-based calibration algorithms usually provide high accuracy and reliable correction of distortions, however can sometimes not be applied. Especially in domains where calibration patterns are difficult to use (e.g., underwater), or in case of changing internal parameters (e.g., zoom) camera calibration often has to rely on self-correction techniques (c.f. [9]). However, until now these have often shown a lack of robustness and stability. Accordingly, a reliable estimation of lens distortions in these fields is a challenging task, and errors frequently occur. In this situation it is advisable to check the quality of calibration and distortion correction results to facilitate appropriate measures in subsequent processing.

In this paper we present a new approach to facilitate such consistency checks. The overall goal of our work aims at an automatic assessment of the registration quality between two aligned images. This includes an identification of different underlying error sources that cause misalignments and visual artifacts (cf. [10]). Especially lens distortions have shown to have a significant influence on the quality of a registration result. Usually they cannot be handled adequately during registration. Accordingly, one fundamental building block in a general quality assessment procedure is the automatic detection and categorisation of lens distortions. The work discussed here presents a case study of a statistical approach towards this goal. Our proposed algorithm consists of two main phases. First, the overall quality of registration is assessed. For this step we propose a new quality metric with high local sensitivity [10]. The results of this first phase are so-called *quality maps* representing the local similarity between two registered images, and yielding the starting point for the second phase. Within this phase a thorough analysis of the error distributions in the maps is performed. Lens distortions turned out to cause striking error patterns in the quality maps. These patterns are used to classify the amount of lens distortions in the images. The various patterns are characterised in terms of a *PCA-based statistical model*, which allows to automatically assess lens distortions in registration results.

The remainder of this paper is organised as follows. In the subsequent section a brief overview over state-of-the-art research in registration quality assessment and lens distortion treatment is given. Section 3 introduces the basic principals of our new registration quality metric, and in Section 4 our statistical approach for lens distortion analysis is discussed. The outcomes of several experimental studies are presented in Section 5, and the paper is closed with a conclusion and an outlook on ongoing and future work in Section 6.

2 Related Work

Non-linear lens distortions and their impact in geometric scene reconstruction and camera motion recovery are well-studied topics in the computer vision com-

Proceedings of the ICVS Workshop on Camera Calibration Methods for Computer Vision Systems - CCMVS2007
Published in 2007 by Applied Computer Science Group, Bielefeld University, Germany
This document and other contributions archived and available at: http://biecoll.ub.uni-bielefeld.de

munity (e.g., [6, 11]). Over the years countless approaches for lens calibration and distortion compensation in images emerged [7], which in many situations allow for robust and reliable camera calibration and image correction. Accordingly, Zhang [12] states that the best way of dealing with lens distortions in image processing is to perform an explicit calibration of the camera *before* acquiring the data to be processed. However, if such techniques cannot be applied, the estimation of calibration parameters becomes a significantly more challenging task. In addition, if parameters can be reconstructed using self-calibration, usually no assessment of their reliability is supplied. Especially in case of an integrated recovery of camera motion and lens distortions, as e.g. proposed in [13], a unique parameter reconstruction is often impossible, since linear and projective motion models can also compensate for lens distortions to a certain degree [14].

Consequently, in areas of applications where state-of-the-art techniques may fail to provide reliable results, complementing approaches are of great interest. They allow to assess the amount of lens distortions present in given images after correction. To this end we propose to perform a detailed analysis of registration results, and to exploit cues within image differences for assessment. Unfortunately, objectively assessing the quality of an image registration is still an unsolved problem. Common registration techniques [15] and related optimisation criteria often do not yield reliable quality measures for an overall objective assessment of the result. Neither feature-based criteria like the *geometric reprojection error* [5] nor featureless measures like the *Mean Squared Error (MSE)* [16] provide an *objective* metric for the final registration quality. Therefore the final quality assessment of registration results is often left to the user for visual inspection which is obviously not feasible for automatic procedures. Only few work has been published so far towards this direction [17].

We propose a new metric for objective registration quality assessment inspired by structural quality metrics from the field of *image* quality assessment [18]. The overall goal in this field is to quantify differences between images as they result, e.g., from image compressions. Related quality metrics often exploit structural image properties like local gradients [19], image entropy or mutual information [20]. In [21] a new *Measure of Structural Similarity* based on local image intensity statistics is proposed. With regard to registration quality assessment, however, all these measures have shown a lack of local sensitivity, since resulting error maps are usually interpreted by performing global error averaging. Our metric builds on [19], however, we follow a pattern recognition approach for analysis of the quality maps and avoid unspecific averaging schemes. The maps are interpreted as images with characteristic spatial intensity distributions. In analogy to statistical approaches using PCA to describe spatial patterns [22] *eigen error patterns* are calculated and serve as basis for an automatic analysis.

3 Objective Registration Quality Assessment

The first phase of our approach for automatic analysis of radial lens distortions in image registration consists of an objective assessment of the registration quality between two images. We propose a new quality metric that aims to provide high local sensitivity. This is achieved laying strong emphasis on pixel-wise calculations. At the same time the metric allows for meaningful global assessments. The overall registration quality between two images is characterised in terms of a *quality map* that points out local structural differences between the images. The global energy of this map correlates well with the overall quality of the registration result. Accordingly, this can be used directly as a quality index for an objective assessment [10]. In addition and with regard to lens distortions, however, also the spatial error distribution contains valuable information.

In general, registration quality is deeply linked to pixel-wise differences between two registered images. Accordingly, detecting such differences yields cues for possible misalignments and registration failures. However, as outlined in more detail in [10], not all pixel-wise differences are directly related to registration failures, i.e., misalignments of corresponding pixels. This is due to image differences caused by other effects which cannot be compensated by improving the geometric registration of the images. Examples for these effects are vignetting or global illumination changes in the images. Hence, a quality metric has to distinguish between these two kinds of image differences, i.e., between *registration* errors that are related to structural image differences originating from the registration process, and *visual* errors that are due to differences between correctly aligned pixels. Visual errors may for example be due to changes in lighting conditions or moving objects. For lens distortion analysis mainly registration errors are of interest.

3.1 Local Quality Criteria

To locally assess registration quality and visual appearance of two aligned images I_1 and I_2, three different pixel-wise local quality criteria are initially calculated:

(i) Absolute Intensity Difference D. Differences within the pixel-wise intensity values of two images always provide cues for possible misalignment. The local intensity difference for each pixel position (x, y) is defined as follows:

$$D(x, y) = |I_1(x, y) - I_2(x, y)| \tag{1}$$

(ii) Edge Preservation Map E. Besides difference of intensities also the orientation of the gradient yields valuable cues for image comparisons. In [19] a metric for image fusion performance was proposed that exploits gradient magnitude and orientation for a perceptually motivated assessment of how well edge information is preserved during image fusion. In our experiments, especially the analysis of gradient orientation has turned out to provide important information for assessment of registration quality. We use the *edge preservation map* E proposed in [19], but only exploit the gradient orientation of two images:

$$E(x, y) = \frac{\Gamma_\alpha}{1 + e^{k_\alpha (A(x,y) - \sigma_\alpha)}} \quad , \quad A(x, y) = 1 - \frac{|\alpha_1(x, y) - \alpha_2(x, y)|}{\pi/2} \quad , \tag{2}$$

where $A(x, y)$ is a measure of difference in local orientation, $\alpha_k = \tan^{-1}\left(\frac{s_k^y(x,y)}{s_k^x(x,y)}\right)$, and $s_k^y(x,y)$ and $s_k^x(x,y)$ are the results of the Sobel operator. The constants in the formulas are chosen according to the default values suggested by Xydeas et al. with $\Gamma_\alpha = 0.9879$, $k_\alpha = -22$ and $\sigma_\alpha = 0.8$, leading to values of E within the range of $[0,1]$ where 1 indicates identity in gradient orientation.

(iii) Structural Risk Map R. One criterion for assessing the image structure at a given pixel is the magnitude of the gradient. We use this value to assess the extent to which single pixels may give reliable cues for registration quality. At positions where the gradient magnitude is quite small in both images, structure is only weakly distinctive and its analysis may lead to wrong conclusions. Hence, these positions are excluded from structural analysis. The binary *risk map* R for marking non-relevant pixels in structural analysis is calculated as follows:

$$R(x,y) = \begin{cases} 1, & \text{if } G_1(x,y) \leq \theta_G \wedge G_2(x,y) \leq \theta_G \\ 0, & \text{otherwise} \end{cases} \tag{3}$$

where G_1 and G_2 are the local gradient magnitudes in both images and θ_G a suitable threshold. Morphological dilation with a 3×3 squared mask is applied to the risk map to also exclude pixels close to a homogeneous neighbourhood.

3.2 Block-wise Error Pooling and Quality Map Generation

Given the local quality criteria, for each pixel position within the aligned images two binary decisions are carried out to detect registration and/or visual errors:

(a) Registration Error Detection. We assume a pixel (x,y) in a structured image section ($R(x,y) = 0$, Eq. 3) to show a registration error if its edge preservation value $E(x,y)$ is smaller than a threshold θ_E. This is chosen to indicate a significant mismatch in the local gradient orientation between both images.

(b) Visual Error Detection. A pixel is marked to show a visual error if the local image intensity difference exceeds a threshold θ_D, provided that the pixel is allowed to vote for visual errors according to the risk map R ($R(x,y) = 1$, cf. Eq. 3), i.e., lying in a predominantly homogeneous neighbourhood.

The outcome of these two thresholding processes are two binary images indicating image pixels where an error of the given type exists between the registered images (for more details see [10]). In the context of lens distortion analysis only the registration error image is considered in subsequent steps. To preserve the local sensitivity provided by the pixel-wise calculations, a block-based voting scheme is applied for error pooling. This is in contrast to common image quality assessment techniques [18] where mainly a global assessment by averaging is carried out. In our case the binary registration error image is divided into blocks of size 8×8 pixels, and for each block the relative amount of pixels with registration errors is determined. The result is then given by a final *quality map* where each pixel represents a single block, with its intensity value proportional to the ratio of pixels with registration errors in the related image section (see Fig. 1 for examples). Blocks with an error ratio below 10 percent are assumed to be free of errors and set to an intensity value of zero.

4 Statistical Analysis of Lens Distortions

The formerly generated quality maps clearly indicate image sections where structural image differences occur, that are most likely related to misalignments of the images. With regard to registration quality evaluation and lens distortion detection especially the spatial distribution of these errors yields valuable cues.

If lens distortions are present in a pair of images, usually not all image regions are registered equally well. Linear transformations and homographies can only partially compensate for lens distortions. Accordingly, characteristic spatial distributions of the blocks with errors, i.e., specific patterns, can be observed within the quality maps (Fig. 1). Depending on the actual amount of distortions present in both images these patterns are more or less pronounced and show a radial symmetric shape. Thus, a more detailed analysis of these patterns allows to categorise the amount of radial distortion present in the images and to assess their influence on the final registration result.

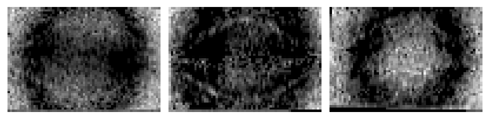

Fig. 1. Example error maps (normalised) as resulting from a registration of distorted images.

We apply a statistical approach to implicitly extract meaningful and discriminative characteristics of the patterns, based on a *principal component analysis*. The 2D quality maps are transformed into 1D feature vectors concatenating their rows. Given the vectors of a suitable set of training samples, the eigen-spectrum of the vector space is then calculated and a set of representative eigenvectors chosen [22], yielding the base for classifying unknown data (Subsec. 4.2).

4.1 Training and Test Set Generation

To enable the statistical classification of unknown data a representative training set of image pairs with known distortions and related registration quality maps is required. Lens distortions in image analysis are usually modelled using a radial-symmetric non-linear *polynomial mapping* between ideal points $\mathbf{p}$ in undistorted images and related distorted points $\mathbf{p}_d$ as observable in distorted images [5]:

$$\mathbf{p} = \mathbf{p}_0 + (1 + k_1 \cdot r + k_2 \cdot r^2 + \ldots)(\mathbf{p}_d - \mathbf{p}_0) \quad (4)$$

The k_i are called *distortion coefficients*, $\mathbf{p}_0 = (x_0, y_0)$ denotes the center of radial distortion (which is often assumed to be identical to the principal point of the camera and the image center), and $r = \parallel \mathbf{p}_d - \mathbf{p}_0 \parallel$ gives the Euclidean distance of any point $\mathbf{p}_d$ to the center of distortion. In practice it is very common to consider only the terms with even exponentials.

The automatic quantification of lens distortions in image data requires an appropriate metric. Basically, the coefficients k_i of the distortion model (Eq. 4) yield such a measure and could be used directly. However, as the direct reconstruction of several distortion coefficients is sometimes ambiguous, and since our primary goal here is to categorise the amounts of lens distortions in registration rather than to reconstruct the coefficients of a specific distortion model, we apply two alternative metrics. On the one hand the distortion in the image pairs is characterised in terms of the maximum pixel offset caused by the applied distortion (Δ_{max}), and on the other hand we use the average pixel offset of a selection of pixels resulting from a certain distortion (Δ_{avg}) as metric in our experiments.

A set of 12.400 image pairs was generated from a real image sequence, formerly corrected for radial lens distortions using a calibration pattern (Fig. 2). The images of each pair were related to each other by projective transformations, dominated by vertical translations with moderate offsets between 0 and 40 pixels, which usually enables a robust registration. The pairs were then artificially distorted with varying amounts of radial distortion, assuming the distortion center in the image center. k_2 was chosen from the interval $\left]0, 2.5 \cdot 10^{-6}\right]$, $k_4 \in \left]0, 2.5 \cdot 10^{-11}\right]$, and all other coefficients were set to zero. Accordingly, $\Delta_{max} \in]0, 80]$ and $\Delta_{avg} \in]0, 30]$. Finally, the distorted pairs were registered by estimating homographies [1], and the registration quality was assessed applying our metric. The resulting maps resembled well patterns observed in registration of real images acquired in an underwater environment.

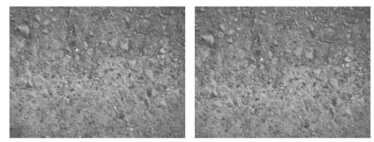

Fig. 2. An exemplary image pair as used in training and test, showing a planar, stony yard.

4.2 PCA-based Classification

As described, the primary goal of our approach is to categorise the amounts of lens distortions in pairs of registered images. To this end we define four classes of different distortion levels (from small (index 1) to very strong distortion (index 4)) by splitting the intervals of Δ_{max} and Δ_{avg}, respectively, in four equidistant sections each. For learning a subset of 12.000 pairs was chosen from the sample data and partitioned into these classes according to the amount of distortion. For each class a principal component analysis was performed. According to the eigen-spectrum of each class a number of eigenvectors between 5 and 10 appeared reasonable to form the corresponding eigenspaces, covering approximately between 60% and 80% of the overall variance within the data. For our experiments we use 5 or 10 eigenvectors for each class, respectively (Fig. 3). For classifying unknown quality maps we employ the original eigen-space approach as proposed in [22]. An unknown quality map is projected into each of the four eigen-spaces. Subsequently, the Euclidean distances between the original map and the reprojected maps are calculated. The minimum distance of all classes defines the final classification result. In addition, for the corresponding distorted image pair also the value of Δ_{avg} is predicted. For this,

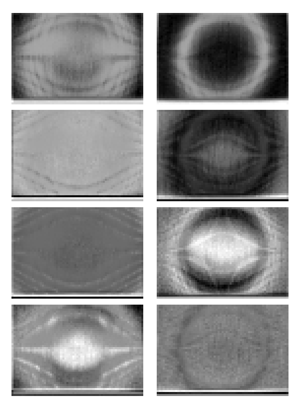

Fig. 3. The first 4 eigenvectors of the 10-D eigenspaces for the classes with least (left) and most (right) distortion, using Δ_{avg} as distortion metric.

the four nearest neighbours within the winning class are determined. Subsequently, Δ_{avg} is calculated by weighted interpolation from the distortion values of these neighbours.

5 Results

The overall performance of our approach was tested on the remaining 400 images of the sample set of 12.400 images, not being part of the training set. In addition, we wanted to rule out that the approach just classifies the images according to the overall amount of errors present, but not due to its local distribution. For this, we used another 100 undistorted image pairs with low registration quality not caused by lens distortion. Low registration quality was simulated by adding error terms to individual parameters of the estimated homographies, with a magnitude of up to 1%. During classification image pairs with a minimal distance above a threshold were rejected and assumed to not belong to any of the four classes, hence, being free of distortions (class R). Figure 1 summarises the outcomes of the classification experiments in terms of the confusion matrices.

(a)	1	2	3	4	R
1	17	7	0	0	0
2	6	52	11	4	0
3	0	25	79	26	0
4	1	2	30	140	0
R	36	20	16	2	26

(a) #EV= 5, Δ_{max}

(b)	1	2	3	4	R
1	18	5	0	1	0
2	6	49	13	5	0
3	0	33	70	27	0
4	0	4	22	147	0
R	29	1	20	2	48

(b) #EV= 10, Δ_{max}

(c)	1	2	3	4	R
1	23	7	1	0	0
2	14	56	12	5	0
3	1	23	115	21	0
4	1	2	39	80	0
R	45	20	3	0	32

(c) #EV= 5, Δ_{avg}

(d)	1	2	3	4	R
1	26	4	1	0	0
2	8	58	18	2	1
3	0	23	102	32	3
4	0	2	32	88	0
R	28	3	2	2	65

(d) #EV= 10, Δ_{avg}

Table 1. Results of the PCA-based quality map classification in terms of confusion matrices (rows, groundtruth, and columns, classification results), using either 5 or 10 eigenvectors (EV) for modelling the class-specific eigen-spaces, and Δ_{max} or Δ_{avg}.

The best classification results were obtained using 10 eigenvectors and Δ_{avg} as metric, leading to a classification rate of almost 70%. A number of 10 eigenvectors has shown to be preferable compared to using only 5 vectors for each class for both metrics. From the confusion matrices it is obvious that miss classification occurs predominantly between neighbouring classes. This can be explained at least in part by the discretization of the amount of distortion into distinct classes. Furthermore, in case that quality maps from distortion-free image pairs are miss classified, they are mainly classified into class 1. This is consistent as this class represents small distortions, hence, hard to distinguish from nearly error-free non distorted samples.

In addition to these categorisation tests, we performed an extended experiment to recover directly the values of the distortion metric Δ_{avg} for distorted test samples. Δ_{avg} was interpolated from the training vectors that were most similar to the distorted test pattern. The roots of the mean squared errors for comparing recovered values for Δ_{avg} for the 400 distorted test patterns with groundtruth are shown in Table 2, separately for each class.

#Eigenvectors per Class	Class 1 (small dist.)	Class 2 (moderate dist.)	Class 3 (strong dist.)	Class 4 (very str. dist.)
5	3.48	4.73	5.13	4.85
10	3.56	4.79	5.19	4.93

Table 2. Results of interpolating exact distortion values from similar training patterns. The values comprise the roots of the MSE between groundtruth and predicted values, given a dynamic range of Δ_{avg} in our experiments equal to an interval of $]0, 30]$.

These results provide a deeper insight into the distinctiveness of the approach. The prediction results support the potential of our approach to adequately capture the amount of distortions in image pairs. In the majority of test cases the tendency of the algorithm to assess the distortion amount is quite satisfying, and a useful accuracy in reconstruction of distortion measurements can be achieved.

6 Conclusion

Lens distortions are known to have a serious impact on image registration results. In addition to common approaches for determining adequate correction coefficients for distortion compensation from calibration patterns or scene structures, we propose a new paradigm for detecting and categorising non-linear distortion effects in registered images. The main idea of our approach is to extract information about distortions present in aligned images from an analysis of registration quality. Based on a local quality assessment metric and principal component based modelling of spatial error distributions, our approach allows to robustly categorise the amount of lens distortions present in images. Also first attempts for a direct prediction of distortion values have underlined the potential of the approach. Ongoing work aims at extended tests on real data, and a regression-based prediction of distortion values from the data, presumably without explicitly defining different distortion categories in advance. Furthermore, additional error types, for example due to dynamic scene parts, will be considered and added to the overall objective quality assessment procedure.

Acknowledgements

This work was supported by a fellowship within the Postdoc-Programme of the German Academic Exchange Service (DAAD). The authors also would like to thank Rafael Garcia and Armagan Elibol from the Underwater Vision Lab Girona for providing test image sequences and related camera calibration data.

References

1. Mann, S., Picard, R.W.: Video orbits of the projective group: A new perspective on image mosaicing. Technical Report 338, MIT Media Laboratory Perceptual Computing Section (1996)

2. Brown, M., Lowe, D.G.: Recognising panoramas. In: Proc. of the Ninth IEEE International Conference on Computer Vision. (2003) (2):1218–1225
3. Torr, P.H.S., Fitzgibbon, A.W., Zisserman, A.: The Problem of Degeneracy in Structure and Motion Recovery from Uncalibrated Image Sequences. International Journal of Computer Vision **32**(1) (1999) 27–44
4. Gracias, N., Santos-Victor, J.: Underwater video mosaics as visual navigation maps. Computer Vision and Image Understanding **79**(1) (2000) 66–91
5. Hartley, R., Zisserman, A.: Multiple View Geometry in Computer Vision. Cambridge University Press (2004)
6. Hsu, S.C., Sawhney, H.S.: Influence of global constraints and lens distortion on pose and appearance recovery from a purely rotating camera. In: 4th IEEE Workshop on Applications of Computer Vision (WACV), Princeton, NJ, USA (1998) 154 ff.
7. Tsai, R.: A versatile camera calibration technique for high-accuracy 3d machine vision metrology using off-the-shelf tv cameras and lenses. IEEE Journal of Robotics and Automation **3**(4) (1987) 323–344
8. Thormählen, T., Broszio, H., Wassermann, I.: Robust line-based calibration of lens distortion from a single view. In: Proc. of Int. Conf. Mirage, INRIA Rocquencourt, France (2003) 105–112
9. Stein, G.P.: Lens distortion calibration using point correspondences. In: Int. Conf. on Comp. Vis. and Patt. Recogn. (1997) 602–608
10. Möller, B., Garcia, R., Posch, S.: Towards objective quality assessment of image registration results. In: Proc. of Int. Conf. on Computer Vision Theory and Applications (VISAPP '07), Barcelona, Spain (2007) 233–240
11. Tordoff, B., Murray, D.: Violating rotating camera geometry: The effect of radial distortion on Self-Calibration. In: Proc. of ICPR. (2000) 423–427
12. Zhang, Z.: On the epipolar geometry between two images with lens distortion. In: Proc. of the Int. Conf. on Pattern Recognition. Volume I. (1996) 407–411
13. Fitzgibbon, A.W.: Simultaneous linear estimation of multiple view geometry and lens distortion. In: CVPR. (2001) (1):125–132
14. Li, H., Hartley, R.: Self-calibration of radial lens distortion using two-view invariants. In: Proc. of ICCV, Omnivis Workshop, Beijing, China (2005)
15. Zitová, B., Flusser, J.: Image registration methods: a survey. Image and Vision Computing **21**(11) (2003) 977–1000
16. Irani, M., Anandan, P., Hsu, S.: Mosaic based representations of video sequences and their applications. In: Proc. Int. Conf. on Computer Vision. (1995) 605–611
17. Kim, H.S., Kim, H.C., Lee, W.K., Kim, C.H.: Stitching reliability for estimating camera focal length in panoramic image mosaicing. In: Int. Conf. on Pattern Recognition, Barcelona, Spain (2000) 1:596–599
18. Wang, Z., Bovik, A.: Modern Image Quality Assessment. Synth. Lectures on Image, Video & Multimedia Processing. Morgan & Claypool Publishers (2006)
19. Xydeas, C., Petrović, V.: Objective image fusion performance measure. Electr. Letters **36**(4) (2000) 308–309
20. Qu, G., Zhang, D., Yan, P.: Information measure for performance of image fusion. Electronics Letters **38**(7) (2002) 313–315
21. Wang, Z., Bovik, A., Sheikh, H., Simoncelli, E.: Image quality assessment: From error visibility to structural similarity. IEEE Trans. on Image Processing **13**(4) (2004) 600–612
22. Turk, M., Pentland, A.: Eigenfaces for recognition. Journal of Cognitive Neuroscience **3**(1) (1991) 71–86