

Identifying Lens Distortions in Image Registration by Learning from Examples

Birgit Möller and Stefan Posch

Institute of Computer Science, Faculty of Natural Sciences III
Martin-Luther-University Halle-Wittenberg, 06099 Halle/Saale, Germany
{birgit.moeller, stefan.posch}@informatik.uni-halle.de

Abstract

Automatic quality assessment of image registration results is an important issue in image processing. Many applications strongly depend on accurate registration results, sometimes even requiring automatic self-recovery from registration failures. In doing so it is not sufficient just to detect registration errors. Also a distinction between different *error sources* is necessary, as each source has an individual impact on the final registration result and requires specific compensation strategies. We present a new approach to automatically identify *lens distortions* in image pairs, known to have a significant impact on registration. The key idea is to analyse registration residuals and to learn a model of spatial residual distributions typical for distorted images. Our approach relies on a new metric for registration quality assessment and implements a regression scheme based on SVMs for predicting distortions in unknown data. The potential of the approach is demonstrated by experimental results on synthetic as well as real image data.

1 Introduction

The analysis of similarities and differences between pairs of overlapping images, and the explicit modelling of image changes due to camera motion by (rigid) image transformations is a well-studied topic in computer vision. This so called *2D geometric image registration* yields the basis for a wide variety of applications, ranging from image sequence compression and mosaicing [9], up to the estimation of camera motion in space and geometric reconstruction of 3D scene structure [18]. The aim of registration is to determine an alignment of corresponding parts in partially overlapping images by estimating parameters of an appropriate model for the camera motion during image acquisition.

Within the last two decades large amounts of registration algorithms emerged [23]. On the one hand featureless approaches were proposed for parameter estimation (e.g., [13, 9]) that aim to minimise colour or intensity differences between images, e.g., using the well-known Mean Squared Error (MSE) as optimisation criterion. On the other hand, feature-based approaches are very common that build on an explicit detection of corresponding image points and, e.g., aim to minimise the reprojection error [7]. In general both classes of techniques allow for an accurate estimation of registration parameters. Accordingly, nowadays the basic task of 2D image registration is a quite mature field. But, nevertheless

there is still no guarantee for correct estimation results in any case, and registration errors as well as complete failures may still occur, and require appropriate treatment.

Despite promising advances in registration algorithms, an *automatic* and *objective assessment* of the final quality of a registration result is still an open issue. Although optimisation criteria used within the registration process itself in principal also provide a measure for the overall quality of the registration result, they often show a lack of local sensitivity (cf. [21, 14]). Also, the visual appearance of registered images as, e.g., observed by humans, does often not correlate very well to the registration quality as imposed by the optimisation function. Consequently, assessing the quality of registration results and especially determining the underlying error source in case of failure of the registration process is still a task usually left to the human user. This renders a fully automatic online image registration with automatic recovery from registration failures infeasible.

Geometric registration of two images may fail due to a wide variety of error sources, e.g., a lack of distinctive structures in the images, the choice of an inappropriate motion model or even simply image noise. One of the most important reasons, however, is given by *non-linear lens distortions* that are known to have a serious impact on the final result of a registration [8]. Accordingly, correcting images for these distortions prior to a geometric registration step is highly recommended. In literature several calibration algorithms are proposed that allow for robust camera calibration and lens distortion estimation. Among these approaches offline techniques based on calibration patterns [19, 1] turn out to form a quasi-standard with regard to accuracy. However, there are acquisition situations where these standard approaches are – if at all – difficult to apply, e.g., due to challenging environmental conditions (like underwater), or where a large flexibility is required as in case of continuously changing internal camera parameters due to zoom.

Image registration often forms a constitutional building block in image processing hierarchies, i.e., higher-level analysis modules strongly rely on accurate results of the registration stage. Due to this, in situations where a reliable calibration of the camera and an explicit compensation for non-linear lens distortions is not possible, it is at least advisable to thoroughly check the quality of any registration result. In particular, assessing the overall amount of distortions present in given images yields a valuable and often indispensable reliability measure for the outcomes of subsequent processing stages.

In this paper we present a new approach towards a fully automatic detection and quantification of lens distortions in registered images. The proposed method is part of an integrated approach for an automatic assessment of registration results [14]. As the final registration quality may be deteriorated by a wide range of potential error sources, any objective quality metric requires not only to detect remaining misalignments between two images, but also needs to distinguish between various underlying reasons. As these have individual impacts on the result they also request for individual compensation strategies.

Our proposed method for lens distortion detection builds on a new registration quality metric with high local sensitivity [14]. The key idea of the approach is to analyse spatial error distributions in so called *quality maps* that result from applying quality metrics to registered images. Lens distortions have shown to cause striking spatial patterns in these maps. Thus the quality maps serve as input for a machine learning approach that aims to predict the amount of lens distortions in image pairs by learning related effects from ground-truth examples. In this paper we extend an earlier case study dedicated to an initial categorisation of distortion artefacts into discrete classes [15] towards a continuous regression scheme. It allows for a direct and more distinctive prediction of distortions

using PCA-based dimension reduction and *support vector machines*.

The remainder of this paper is organised as follows. In Section 2 a brief overview of related work is provided. Section 3 summarises the basic principals of our metric for objective quality assessment. We predict distortions present in image pairs by performing a regression based on SVMs as outlined in Sec. 4. In Section 5 results of our approach are discussed. The paper finishes with a conclusion and an outlook on future work (Sec. 6).

2 Related Work

Although image registration algorithms play an important role in a wide variety of computer vision applications, the development of procedures for a fully automatic assessment of the quality of registration results has not yet been in the focus of the research community. Only few steps towards this direction were proposed [10], and work often focused on theoretically optimal boundaries of registration algorithms instead [3]. Consequently, objective metrics that would also allow for comparisons between different algorithms without requiring ground-truth data (which is often difficult to acquire) are not available, and quality assessment is usually done manually by humans.

In contrast to this lack of objective quality metrics with regard to image registration, within the field of *image quality assessment* over the years various metrics for measuring differences between images emerged [21]. These metrics usually aim at quantifying differences between images that show the same contents, however, may have undergone certain transformations or filtering operations, e.g., being compressed for data transmission purposes or degraded by various types of noise. Also, in virtual reality applications these metrics are common to compare artificially rendered scenes with real images.

Metrics in the field of image quality assessment usually exploit either visual image properties, emulating capabilities of the human eye (e.g., [12]), or structural properties of the images, e.g., local gradients [22], mutual information [16] or local image statistics of gray values [21]. Especially the second class of metrics is closely linked to an assessment of registration quality since image registration mainly aims at a structural alignment of images, i.e., minimising structural image differences. Accordingly, our metric builds on these approaches. However, as they have shown a lack of local sensitivity with regard to assessing registration results, we follow a pattern based approach for global analysis of image differences rather than applying common averaging schemes. This yields an increased local sensitivity indispensable in registration quality assessment.

Due to the lack of objective quality metrics in registration, also the undoubtedly serious impact of lens distortions in image registration has up to now mostly been evaluated qualitatively. E.g., in [8] the influence of lens distortions on accurately closing 360° panoramas is analysed, however, the results are presented in terms of visual maps to be interpreted manually by the human observer. This coincides with the fact, that even an objective assessment of the outcomes of stand-alone algorithms for recovering lens distortion coefficients is still an open question, if no ground-truth data is available.

As the effects of lens distortions and camera motion in image pairs are often tightly coupled, it is quite difficult to separate them uniquely from each other [11], as, e.g., aimed at by algorithms for an integrated recovery of camera motion and (radial) lens distortions [17, 6]. Hence, techniques to assess the amounts of distortions in image pairs as discussed in this paper are of significant interest, especially with regard to the goal of achieving high-quality image analysis results that rely on an accurate image registration.

3 A Metric for Registration Quality Assessment

The overall quality of a registration result is usually associated with the amount of pixel-wise differences between two registered images, remaining after alignment. Accordingly, analysing residuals between registered images yields a suitable starting point for any objective assessment. However, in doing so it has to be taken into account that image residuals not only emerge from registration failures, but may also be caused by other error sources, not directly related to the registration process itself.

We distinguish between *visual errors* and *structural errors* (Fig. 1). Visual errors are exclusively due to visual image differences like illumination changes or vignetting. Structural errors are caused by geometric misalignment of the images, i.e., due to a true registration failure.

Our objective metric accounts for both classes of errors by thoroughly analysing local image structure and intensity differences. In Fig. 2 an overview of the two-step algorithm is provided. In the first stage three different pixel-wise quality criteria are calculated. Subsequently these are used as input in two independent analysis processes for detecting structural and visual differences between the images. Below the two stages of the algorithm are outlined. As for an analysis of lens distortions only structural differences are of interest, details about the visual error detection are omitted here, but can be found in [14].

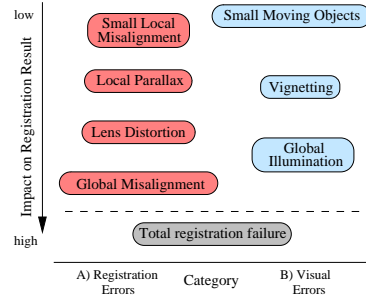


Figure 1: A taxonomy of image differences in registration (cf. [14]).




Figure 2: Overview of the metric for assessing the quality of a registration result.

3.1 Pixel-wise Quality Criteria

In the first stage three different quality criteria are calculated within the overlapping area of two registered images. These are the pixel-wise intensity difference, an *edge preservation map* E (cf. [22]), and finally a *structural risk map* R (cf. [5]) to separate structural and visual errors. Regarding structural error detection, only E and R are relevant.

Edge Preservation Map E

Image misalignment coincides with structural image differences. The edge preservation map exploits variations in the local gradient orientation between two images for detecting such defects, applying a perceptually motivated distance function:

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

$s_k^x(x,y)$ and $s_k^y(x,y)$ are the results of the Sobel operator, and $\Gamma_\alpha = 0.9879$, $k_\alpha = -22$ and $\sigma_\alpha = 0.8$ are constants set according to the defaults suggested in [22].

Structural Risk Map R

At positions where the gradient magnitude is small in both images, structure is only weakly distinctive and the analysis of gradient orientations may lead to wrong conclusions. Accordingly, these positions are marked in the binary risk map to be excluded from structural analysis, and exclusively considered in visual error analysis:

$$R(x,y) = \begin{cases} 1, & \text{if } G_1(x,y) \leq \theta_G \wedge G_2(x,y) \leq \theta_G \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

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

3.2 Global Assessment for Lens Distortion Analysis

For assessing the overall registration quality of an image pair, as outlined in our initial approach [14], a block-wise error pooling of edge preservation maps is performed, followed by a global voting procedure that yields four different error ratios for global error assessment. Regarding an identification and quantification of lens distortions as considered in this paper, however, the algorithm is slightly modified to better preserve local information.

Basically, the edge preservation map E is downsampled by a factor of eight, applying Gaussian masks for low pass filtering. In doing so, only pixels with distinctive gradient magnitudes in both images as indicated by the risk map R are considered. Figure 3 depicts some representative example *quality maps*, resulting from registering images degraded by varying amounts of non-linear distortions.




Figure 3: Example maps (contrast enhanced for visualisation), covering small (left) to big (right) distortion.

4 Predicting Lens Distortions with SVMs

The quality maps as calculated in the previous section reflect the overall registration quality in terms of their energy. However, with regard to lens distortion analysis, which is the main goal of this paper, not only the total amount of structural errors, but also their spatial distribution plays an important role. In particular, as can be seen in the example maps in Fig. 3, lens distortions cause striking spatial error patterns in the quality maps. The larger the degree of distortion within the registered images, the more pronounced are the radial symmetric patterns observed in related maps. Consequently, learning the relationship between the various characteristic error patterns and related amounts of radial distortion from ground-truth data should allow to quantify lens distortions in pairs of images.

To automatically quantify distortions according to spatial error distributions, an adequate measure for the degree of distortion in image pairs is required. Also, a suitable mapping function is necessary that relates the quality maps or accordant feature representations of these maps to the distortion measure. In our approach, we use a model-independent distortion measure (Sec. 4.1), and apply support vector machines to learn the transformation from the quality maps to this measure as a regression function (Sec. 4.2).

4.1 Characterising Lens Distortions

Radial 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 from distorted images [7]:

$$\mathbf{p} = \mathbf{p}_0 + (1 + k_2 \cdot r^2 + k_4 \cdot r^4 + \dots)(\mathbf{p}_d - \mathbf{p}_0), \quad (3)$$

where the k_i are called *distortion coefficients*, \mathbf{p}_0 denotes the center of radial distortion (which is often assumed to be identical to the image center), and $r = \|\mathbf{p}_d - \mathbf{p}_0\|$ gives the Euclidean distance of point \mathbf{p}_d to this center of distortion. Basically, the coefficients k_i of this model would yield a straight-forward measure for the distortion present in given images. Estimating a mapping function between spatial error distributions is in this case directly related to the reconstruction of parameters for a specific model of distortion.

However, since the direct reconstruction of coefficients k_i is sometimes ambiguous and we do not want to link our approach to a specific distortion model, we use a model-independent distortion metric Δ_{avg} . The distortion within an image is characterised by the average pixel offset caused by the applied distortion. Δ_{avg} is calculated from a subset of pixels located at equidistant positions on a ray through the center of distortion, which is sufficient to adequately characterise radial symmetric distortions as assumed here.

4.2 Learning Distortion Effects from Examples

The prediction of distortion from quality maps may be cast in the framework of function approximation. We use support vector machines (SVMs) in ε -regression mode [4] to achieve good generalisation abilities. The width of the insensitivity tube of the ε -insensitive loss function $|\xi|_\varepsilon$,

$$|\xi|_\varepsilon := \begin{cases} 0 & \text{if } |\xi| \leq \varepsilon \\ \varepsilon - |\xi| & \text{otherwise} \end{cases}, \quad (4)$$

is adjusted with the parameter ε . As usual, the parameter C denotes the tradeoff between the accuracy of the approximation on the training set and the capacity of the regression system. As kernels we use the a linear kernel, a polynomial kernel with degree 2 and a Radial Basis Function with γ controlling the width of the kernel.

In principle, the quality maps could directly be used as input for the SVM. However, as the maps usually have a size of about 60 to 80 pixels in width and 30 to 50 in height, this would result in a few thousands of input variables. This usually poses problems for generalisation, even for SVMs, and also for the computational load. Thus, we use PCA for feature extraction to project the quality maps into lower dimensional subspace. Subsequently each component of the resulting input vector is scaled to the interval $[0, 1]$.


5 Experiments & Results

Given our objective metric for registration quality assessment and the model-independent distortion measure Δ_{avg} to quantify the amount of distortions present in images, the overall goal of our approach is to automatically identify lens distortions in pairs of registered images, and to assess their influence on the quality of the registration result from spatial residual distributions. To this end the SVMs are trained on quality maps generated by assessing registration results of image pairs with known distortion (Subsec. 5.1 and 5.2), and then applied to maps of registered synthetic as well as real test image pairs (Subsec. 5.3).

5.1 Sample Image Data

Learning a regression function requires a representative set of sample image pairs and related quality maps with known radial distortion. Unfortunately, real image data suffering from a large variety of different kinds and amounts of radial distortions, for which also a ground-truth calibration is available, is rare. Accordingly, we created ground-truth data from real image sequences with sufficient structure to allow for representative spatial error patterns. First, the raw images were corrected for radial lens distortions using a calibration pattern and the Matlab toolbox [1]. Subsequently pairs of images to be registered were selected and subjected to synthetic radial lens distortions of varying amount.

In detail, a set of 6.400 image pairs, 520×360 pixels in size, was selected from a real sequence showing a stony yard scanned by a camera moving in vertical direction in parallel to the ground. In Figure 4 an example image pair is depicted. The two images of each pair were randomly chosen from the sequence to ensure a reasonable variety of different offsets and image transformations in the data set.



However, the offset in each pair was restricted to a maximum of 40 pixels to enable robust registration. Accordingly, given a moderate camera motion, the images of each pair were separated by up to 7 frames.

Figure 4: An exemplary image pair as used in training and test, showing a planar, stony yard.

To distort the images we applied varying radial distortion as defined in Eq. (3), assuming the distortion center in the image center. k_2 was sampled uniformly from the range $]0, 2.5 \cdot 10^{-6}]$ and k_4 from $]0, 2.5 \cdot 10^{-11}]$, giving values for Δ_{avg} between 0 and ≈ 30 .

5.2 Experimental Setup and Training Phase

Once the sample set of distorted image pairs is given, each pair is registered by estimating a homography using the method proposed in [13]. Subsequently, the registration quality is assessed applying our metric, resulting in a quality map for each registered image pair of the sample set. They yield the base for training and testing the SVMs. 6000 quality maps were used in the training, while the remaining 400 formed the synthetic test set.

The first training stage consisted of determining a suitable feature subspace for reducing the dimensionality of the input data by PCA. To reduce the overall computational load only a subset of 2760 training samples was used in this step. In doing so the maximum number of relevant eigenvectors is bounded by the number of input images, restricting the size of the covariance matrix and simplifying the eigen analysis (details can be found in [20]). According to the eigenspectrum of the input data set we used between 10 and 60 eigenvectors (EVs) in our experiments to represent the feature space (see also Table 1).

All experiments were based on the *libSVM* [2] that supports easy training and testing as well as parameter optimisation. For each SVM initially appropriate parameters C , γ and ϵ (Subsec. 4.2) were determined by performing a grid search on the parameter space. A discrete sampling scheme with equidistant spacing in the logarithmic scale was applied to each of the parameters, and the resulting SVM-configurations were evaluated on the given training data subset using 5-fold cross-validation. Once the optimal SVM parameter set according to a minimal MSE on the training data was found for each SVM, they were trained on the full set of 6000 samples, and tested as outlined below.

#Eigenvectors (EVs)	SVM _{linear}	SVM _{poly2}	SVM _{RBF}
10	5.92	5.46	4.57
20	5.54	4.89	4.33
30	5.35	4.66	4.04
50	5.19	n.a.	3.82
60	5.13	n.a.	4.08

Table 1: RMSE of the SVM-regression approach on the set of 400 synthetic test images. "n.a." denotes experiments that were not carried out due to very high training times.

5.3 Results and Discussion

Different test runs were carried out using different dimensionalities for the feature vectors, and either linear kernels (SVM_{linear}), polynomial kernels with degree 2 (SVM_{poly2}), or kernels based on radial basis functions (SVM_{RBF}).

Since the times for grid search and training of polynomial kernels turned out to be significantly higher than for the others, i.e., sometimes taking several days to terminate, for high dimensional data only linear and RBF kernels were applied. For testing two different data sets were used: i) the remaining 400 artificially distorted image pairs not being part of the training set, and ii) ten real image pairs acquired in a different domain, i.e., underwater. The latter ones were mainly tested to assess the generalisation capabilities of our approach.

In Table 1 the results of the test runs on the 400 synthetic images are summarised in terms of the roots of the Mean Squared Error (RMSE) for each run. In general RBF kernels yield the best results, and a selection of 50 EVs appears to be favourable. Polynomial kernels clearly outperform linear ones, but extraordinary high training times on high-dimensional data impair their practical relevance.

The SVM that performed best on the synthetic test data, i.e., using the RBF kernel with 50D vectors, was in the second test phase used to predict distortions $\hat{\Delta}_{avg}$ in unknown image pairs from a different domain, and also acquired with another autofocus camera (Fig. 5). Note that the focus was nearly constant during image acquisition according to camera motion and scene structure. As no ground-truth was available for this data set, the sequence was manually calibrated for lens distortions, based on a calibration pattern visible in some of the images. The estimated radial distortion relates to $k_2 \approx 1.5 \cdot 10^{-6}$, so that $\Delta_{avg} \approx 19$. In Table 2 the prediction results $\hat{\Delta}_{avg}$ for all ten real image pairs are summarised. In general the predicted values for the image pairs appear quite reasonable and meet the manual calibration result quite well, given a slight tendency to underestimate the amount of distortions. Although local image structure sometimes varies significantly within these images, the approach is able to interpret the overall residual distribution in a satisfying way, and to relate it to the known training samples for distortion prediction. This underlines the overall potential of our new technique, and particularly emphasises its generalisation capabilities with regard to unknown domains.




Figure 5: Example real test image acquired underwater.

Pair	$\hat{\Delta}_{avg}$
0	16.1
1	14.0
2	13.7
3	20.0
4	19.9
5	18.2
6	17.0
7	16.4
8	15.8
9	17.6

Table 2: $\hat{\Delta}_{avg}$ for the set of 10 real image pairs.

6 Conclusions

Objective registration quality assessment and an identification of underlying error sources is a key aspect in the development of fully automatic and robust image registration procedures. In this paper we have presented a new approach for automatically identifying lens distortions in image registration by analysing registration residuals and learning links between spatial error patterns and amounts of distortion from examples. Based on a new metric for objective quality assessment and a regression scheme using SVMs, our approach allows to predict the amounts of lens distortions in pairs of images, and to assess their impact on image registration. Tests on synthetic as well as real image data have shown the high potential of this technique, and also its generalisation capabilities with regard to various domains and acquisition devices. Future work will focus on a bigger variation of image transformations and offsets, including extended tests on real data. Finally, the approach will be integrated into the overall framework for automatic registration error analysis [14], yielding a fully automatic system for objective quality assessment.

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 Dr. Rafael Garcia and the team of the Underwater Vision Lab at the University of Girona, Spain, for providing test images and related camera calibration data.

References

- [1] J.-Y. Bouguet. Camera calibration toolbox for matlab, 2006. http://www.vision.caltech.edu/boujuetj/calib_doc/index.html.
- [2] C.-C. Chang and C.-J. Lin. Libsvm: a library for support vector machines, 2001. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- [3] A. Criminisi, I. Reid, and A. Zisserman. A plane measuring device. *Image and Vision Computing*, 17(8):625–634, 1999.
- [4] N. Cristianini and J. Shawe-Taylor. *An introduction to Support Vector Machines and other kernel-based learning methods*. Cambridge Univ. Press, NY, USA, 2000.
- [5] D. Farin and P.H.N. de With. Misregistration errors in change detection algorithms and how to avoid them. In *Int. Conf. on Image Proc.*, pages II:438–441, 2005.
- [6] A. Fitzgibbon. Simultaneous linear estimation of multiple view geometry and lens distortion. In *Int. Conf. on Comp. Vis. and Patt. Recog.*, pages (1):125–132, 2001.
- [7] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2004.
- [8] S. Hsu and H. Sawhney. Influence of global constraints and lens distortion on pose and appearance recovery from a purely rotating camera. In *4th IEEE Works. on Appl. of Comp. Vision (WACV)*, page 154 ff., Princeton, NJ, USA, October 1998.

- [9] M. Irani, P. Anandan, J. Bergen, R. Kumar, and S. Hsu. Efficient representations of video sequences and their applications. *Sign. Proc:Image Comm.*, 8:327–351, 1996.
- [10] H.-S. Kim, H.-C. Kim, W.-K. Lee, and C.-H. Kim. Stitching reliability for estimating camera focal length in panoramic image mosaicing. In *Int. Conf. on Pattern Recognition*, pages 1:596–599, Barcelona, Spain, 2000.
- [11] H. Li and R. Hartley. Self-calibration of radial lens distortion using two-view invariants. In *Proc. of ICCV, Omnivis Workshop*, Beijing, China, October 2005.
- [12] J. Lubin. A visual discrimination model for imaging system design and evaluation. In E. Peli, editor, *Visual Models for Target Detection and Recognition*, pages 245–283. World Scientific, Singapore, 1995.
- [13] S. Mann and R.W. Picard. Video orbits of the projective group: A new perspective on image mosaicing. Technical Report 338, MIT Media Laboratory Perceptual Computing Section, Boston, USA, 1996.
- [14] B. Möller, R. Garcia, and S. Posch. Towards objective quality assessment of image registration results. In *Proc. of Int. Conf. on Computer Vision Theory and Applications (VISAPP '07)*, pages 233–240, Barcelona, Spain, 2007.
- [15] B. Möller and S. Posch. Automatic analysis of lens distortions in image registration. In *Proc. of Workshop on Camera Calibration Methods for Computer Vision Systems (CCMVS '07), ICVS '07*, Bielefeld, Germany, March 2007.
- [16] G. Qu, D. Zhang, and P. Yan. Information measure for performance of image fusion. *Electronics Letters*, 38(7):313–315, March 2002.
- [17] H. Sawhney and R. Kumar. True multi-image alignment and its application to mosaicing and lens distortion correction. 21(3):235–243, 1999.
- [18] P. H. S. Torr, A. W. Fitzgibbon, and A. Zisserman. The Problem of Degeneracy in Structure and Motion Recovery from Uncalibrated Image Sequences. *International Journal of Computer Vision*, 32(1):27–44, August 1999.
- [19] R.Y. Tsai. 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):323–344, 1987.
- [20] M. Turk and A. Pentland. Eigenfaces for recognition. *Journal of Cognitive Neuroscience*, 3(1):71–86, 1991.
- [21] Z. Wang and A.C. Bovik. *Modern Image Quality Assessment*. Synth. Lectures on Image, Video & Multimedia Processing. Morgan & Claypool Publishers, 2006.
- [22] C.S. Xydeas and V. Petrović. Objective image fusion performance measure. *Electr. Letters*, 36(4):308–309, Februar 2000.
- [23] B. Zitová and J. Flusser. Image registration methods: a survey. *Image and Vision Computing*, 21(11):977–1000, 2003.