DOI: 10.2352/ISSN.2470-1173.2016.13.IQSP-009

Adaptive Geometric Calibration Correction for Camera Array

Florian Ciurea, Dan Lelescu, Priyam Chatterjee, Kartik Venkataraman; Pelican Imaging; Santa Clara, CA/USA

Abstract

Multi-camera systems are increasingly gaining popularity for various applications and their correct functionality depends on precise registration. The complexity of registering the various images to each other is reduced significantly by rectifying the images. This usually relies on an offline calibration process. In reality, components of the camera module respond differently to various factors such as temperature variations, field conditions, etc. Therefore, changes in geometric camera calibration, unless accounted for, can affect the proper registration, which in turn leads to severe degradation of the imaging system or can lead to artifacts. We present a method that can assess the geometric calibration of an array camera and perform an adaptive adjustment of geometric calibration by robust feature matching in any imaged scene. Assuming a gradual degradation of geometric calibration from their previously calibrated values, we exploit the redundancy of a camera array system to recover from the variation of calibrated parameters. Compared to other online calibration methods mostly used for stereo systems, our proposed method is efficient and robust, and derives a solution for multi-camera systems. We illustrate the usefulness of our geometric calibration compensation approach through a super-resolution application where we recover significant image details that are lost due to errors in calibration.

Introduction

In a camera array system such as the monolithic camera array system presented in [1], as well as in a general camera array architecture such as [10], accurate registration is essential to the functionality of the imaging system. We use a 4×4 monolithic camera array of low resolution (1000×750 pixels each) with monochromatic Red/Green/Blue filters for each individual camera in the array.

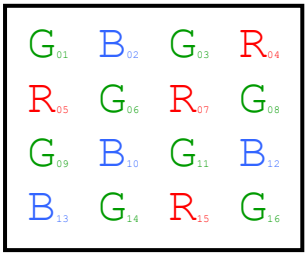

1.41	1	1.41	2.24
1	0	1	2
1.41	1	1.41	2.24
2.24	2	2.24	2.83

Figure 1: PiCam 4×4 camera array (viewed towards scene) and baseline multipliers

As detailed in [1], we compute both a high resolution super-resolved image and a depth map from the reference viewpoint. Sub-pixel registration is achieved through two main stages: 1) Normalization, which includes photometric and geometric correction of cameras using calibration data (captured offline) [5][6] and 2) Parallax estimation. The overall accuracy of the system further into the pipeline is predicated by precise sub-pixel parallax correction. As is the case with general stereo setups, we perform parallax search on epipolar lines [2][3][4] by using rectified geometry employed at the geometric calibration stage. The rectification process achieves two main objectives by reducing the search space from 2D (entire image) to 1D (epipolar lines) for each non-reference camera. First, it dramatically reduces complexity – this is compounded in multi-camera setups because of the number of cameras in the array. Second, it severely constraints the matches to be on corresponding epipolar lines, therefore decreasing the likelihood of false matches. Additionally, in multi-baseline stereo, searches are performed simultaneously in various non-reference cameras at disparities proportional to the camera baselines, improving accuracy [9].

The accuracy of our geometric calibration process is a critical step in a camera array and it represents the foundation for correct registration. In contrast to other methods for auto-calibration in stereo setups [4][7][8], we use a nonparametric approach to geometric calibration: our geometric calibration consists of a dense vector field for every non-reference image which denotes the corresponding pixel in the reference image. This smooth, dense vector field represents the perturbation that needs to be applied to pixels in a non-reference image such that they are aligned onto epipolar lines with respect to the reference camera pixels.

Overview of the method

Our method is based on exploiting the calibrated-stereo constraint: in rectified multi-camera stereo setups, world points are imaged in various cameras on epipolar lines, with the amount of disparity encountered on epipolar lines being determined by corresponding baselines, focal length, and depth of world points.

If calibration information was provided and accurate, the orthogonal projections $\{d_{toEPI}\}_k^i$ of these imaged points i onto epipolar lines should be close to zero for each non-reference camera k in the array (subject to accuracy of the geometric calibration).

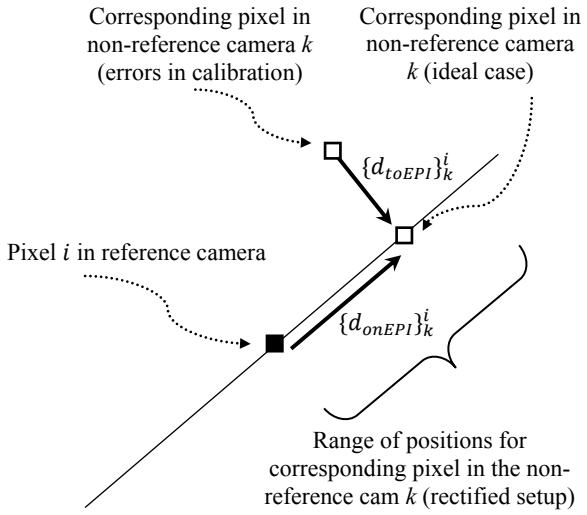

Figure 2: World pixels imaged on epipolar lines

DOI: 10.2352/ISSN.2470-1173.2016.13.IQSP-009

Our method is described in a few steps:

Step 1: We use robust feature point detectors and matching across all cameras in the array. We proceed to detect inaccurate calibration by measuring the magnitudes of the orthogonal geometric components $\{d_{toEPI}\}_k^i$ referred to above. For scenes known to be at infinity, the magnitude of the components on epipolar lines $\{d_{onEPI}\}_k^i$ should also be zero, otherwise the locations of the projected world points on epipolar lines are proportional to the baselines in the camera array.

Step 2: We hypothesize a disparity at each feature point by using the redundancy inherent in the array camera, and assuming a certain type of perturbation from nominal calibration.

In our case, we compute the estimated disparity in pixels $\widetilde{d_i}$ for each interest point i from all valid correspondence vectors using weighted averaging of the measured pixel distances on epipolar line $\{d_{onEPI}\}_k^i$ by the known baseline multipliers b_k as shown in Figure 1:

$$\widetilde{d_i} = \frac{1}{M}\sum_{k\in Match(i)} \frac{\{d_{onEPI}\}_k^i}{b_k}. \quad (1)$$

Here, we denote *Match(i)* the set of non-reference cameras identifying a matching interest point i in the reference camera and M as the number of non-reference cameras with an identified match.

This hypothesized disparity minimizes the overall distance in pixels between the set of observed disparities scaled by the baseline multipliers $\{d_{onEPI}\}_k^i/b_k$ in a least squared sense. This is because

$$\widetilde{d_i} = \arg\min_{d_i} \sum_{k\in Match(i)} \left(\frac{\{d_{onEPI}\}_k^i}{b_k} - d_i\right)^2. \quad (2)$$

The hypothesized disparity in our case is consistent with solving for the least perturbation - in the direction from epipolar lines - from nominal calibration.

Step 3: At this point, we are in a position to compute a correction vector field based on the measured locations of feature points, and the hypothesized disparity $\widetilde{d_i}$. This is done for each non-reference camera in the array, and essentially quantifies the amount of residual shift from the nominal calibration. This vector field is computed only for the selected feature points. We first compute the residual vector field on epipolar direction $\{r_{onEPI}\}_k^i$ by subtracting $\widetilde{d_i}$ from $\{d_{onEPI}\}_k^i$ while keeping its orthogonal direction $\{r_{toEPI}\}_k^i$, equal to $\{d_{toEPI}\}_k^i$ component. We project this vector field to XY coordinates into $\{r_X\}_k^i$, $\{r_Y\}_k^i$, which is the sparse correction vector field that we use. We then use surface interpolation to fit a smooth surface to this sparse correction vector field.

Step 4: Finally, we use this smooth, dense correction vector field as our calibration residual-vector field for each non-reference camera in the array, to apply a geometric correction to each such camera, with respect to the reference camera.

Our approach finds a solution that corrects the geometric calibration by finding a perturbation vector field that is consistent to the chosen criterion described above. While this may not necessarily be the actual solution, it is a plausible solution that produces good results for the type of depth map application needed which is super-resolution. For other types of applications, hypothesizing for a particular disparity at each pixel may be done by choosing a different criterion.

The above correction technique does not necessarily have to be performed for every image that the camera captures. Rather, a detector that quantifies a loss of calibration by monitoring the magnitudes $\{d_{toEPI}\}_k^i$ can be utilized to trigger the processing for calibration recovery described in this paper. This can be done in an online monitoring mode of operation.

Results

We show that using our method we can recover the geometric calibration that may be lost due to various external factors. Using super-resolution as the main application for registration, we illustrate the effectiveness of our algorithm both visually as well as objectively using SSIM scores [11]. We demonstrate that once calibration is lost due to various external factors, the loss in image quality is significant, making recovery of the calibration critical.

We utilize two types of calibration perturbations in this paper. The first type comprises using wrong calibration data which may correspond to field effects (such as shock), with or without additional temperature effects. In general, these perturbations tend to be more severe, and less monotonic compared to the second type of calibration perturbations used next. This latter situation comprises using calibration files that no longer correspond to the temperature at which the array camera operates (since camera was calibrated at a given, e.g., room temperature). For example in the case of using polymer lenses, and plastic holders, temperature will have an effect on the geometry of the array camera, which needs to be taken into account.

For our first experiment, we show results for the case where a large scale perturbation of geometry has occurred from the state at which the initial offline calibration was performed. In Figure 3 and Figure 5 we show:

a) The high resolution image quality / corresponding depth map quality when using the appropriate offline calibration. These are our 'reference' results;

b) The high resolution image quality / depth map quality as being significantly affected by artifacts when we use the original offline calibration, while a perturbation of it has occurred;

c) The high resolution image quality / depth map quality as being significantly improved compared to b), by performing our automatic calibration correction, correcting the original offline calibration.

Overall, results with our automatic geometric calibration compare visually very well to the reference results in a). We show additional super-resolution results in Figure 8. In Figure 4 we illustrate some example vector fields for the uncorrected and corrected case. For the uncorrected case, the reason for a residual correction vector field is twofold: a) there is always a small, residual calibration error and b) estimation errors between the computed hypothesized disparity in step 2 and the actual disparity.

For objective evaluation we report the results in Figure 6 for various natural scenes. The SSIM measurements are performed on image content that has been masked spatially to limit the evaluation to textures and its immediate neighborhood. This is done so that the results are not perturbed by just noise variations in

DOI: 10.2352/ISSN.2470-1173.2016.13.IQSP-009

the smooth (texture-less) areas, and because the loss of geometry calibration will manifest itself in the effects observed on texture rendering in the output image. For the purpose of spatial masking, we used a simple Sobel edge detector tuned to capture fine textures in the image while avoiding the detection of random structures in the noise present in the images. The resulting mask was then slightly dilated spatially using a morphological operator with a square structuring element of size 3×3 pixels. The mask detection was obtained from the reference (well-calibrated) high resolution images. In essence, we would like to measure the perturbations in the rendered texture of the incorrectly calibrated, or recovered-calibration-camera, with respect to their original position/appearance in the correctly calibrated camera.

For our second experiment, we illustrate the effect of temperature on geometric calibration for our array camera modules manufactured using plastic lenses and the successful recovery of geometric calibration using our method. Compared to the previous experiment, the calibration loss is more gradual, and increases as the temperature deviates from that of the initial offline calibration. Note that temperature rise is almost unavoidable due to heat generated as the sensor is operated over a period of time (example, when capturing video).

Figure 7 shows SSIM results of using calibration corresponding to incorrect operating temperature, for the case corresponding to incorrect calibration and recovered calibration respectively. The nominal (correct) calibration was performed at approximately 35 °C (thermistor reading near camera module) which was also the approximate temperature at which the image was taken. The first reported result at 40 °C is naturally very close even for the incorrect calibration case.

Conclusion

We present a method to automatically detect and correct for errors in geometric calibration in camera arrays. Our method exploits the multiple camera calibrated stereo setup and is fast and efficient. We are aware of automatic correction for geometric calibration in the literature for stereo (two-camera) setups, but our method is significantly simpler and operates in multi-camera arrays, by exploiting redundancies in the multi-camera setup. Experimentally, we demonstrate that our method provides a solution to a critical problem faced by modern camera array systems.

References

[1] K. Venkataraman, D. Lelescu, J. Duparre, A. McMahon, G. Molina, P. Chatterjee, R. Mullis, S. Nayar, “PiCam: An Ultra-Thin High Performance Monolithic Camera Array,” in ACM Transactions on Graphics (Proc. SIGGRAPH Asia), 32(5), 2013.

[2] R. Hartley and A. Zisserman, "Multiple View Geometry in Computer Vision", 2nd ed, Cambridge University Press, 2004.

[3] R. Szeliski, “Computer Vision, Algorithms and Applications”, Springer, Springer-Verlag, 2010.

[4] T. Moons, L. Van Gool, and M. Vergauwen, “3D Reconstruction from multiple images. Foundations and Trends in Computer Graphics and Computer Vision”, 4 (4), 2010.

[5] R.Y. Tsai, “A Versatile Camera Calibration Technique for High Accuracy 3D Machine Vision Technique using off the Shelf TV Cameras and Lenses”, IEEE Journal of Robotics and Automation, RA-3(4), pp 323-344, 1987.

[6] Z. Zhang, “A Flexible New Technique for Camera Calibration”, IEEE Trans PAMI, 22(11), pp 1330-1334, 2000.

[7] C.-R. Huang, C.-S. Chen, P.-C. Chung, "An improved algorithm for two-image camera self-calibration and Euclidean structure recovery using absolute quadric”, Pattern recognition, 37, pp 1713-1722, 2004.

[8] Z. Zhang, R. Deriche, O. Faugeras, Q.-T. Luong, "A robust technique for matching two uncalibrated images through the recovery of the unknown epipolar geometry," Artif. Intell. 78, pp 87–119, 1995.

[9] M. Okotumi, T. Kanade, "A Multiple Baseline Stereo," IEEE Trans. PAMI, vol. 15, no. 4, pp. 353-363, 1993.

[10] B. Wilburn, N. Joshi, V. Vaish, E.-V. Talvala, E. Antunez, A. Barth, A. Adams, M. Horowitz, M. Levoy, "High Performance Imaging using a Camera Array," in ACM Trans Graph, 24 (3), pp. 765-776, 2005.

[11] Z. Wang, A. C. Bovik, H. R. Sheikh and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," IEEE TIP, 13 (4), pp. 600-612, 2004.

Author Biography

Florian Ciurea is currently Principal Imaging Scientist at Pelican Imaging, working on computational array cameras. Previously, he was with Sony US Research Center working on various computational imaging features for cameras. He has received his M.S. in Computer Science from "Politehnica" University in Bucharest, Romania and his Ph.D. in Computing Science from Simon Fraser University. He has been working on color and image processing on cameras for almost 15 years.

Dan Lelescu has close to 20 years of research experience in areas such as two- and higher-dimensional (i.e., plenoptic) image and video processing, imaging, computational cameras, content-based retrieval from image or video databases, and complex-system statistical modeling (e.g., mobility modeling in wireless LAN). He is currently Chief Imaging Scientist at Pelican Imaging., working on architectures and processing for array computational cameras. Dan earned his M.S. in Electrical Engineering from the Technical University "Politehnica" Timisoara, Romania, and his Ph.D. in Electrical Engineering and Computer Science from the University of Illinois, Chicago

Priyam Chatterjee received the B.Tech. degree in Information Technology from University of Kalyani, India, an M.Tech. in Electrical Engineering from Indian Institute of Technology, Bombay, India and the Ph.D. degree in Electrical Engineering from University of California, Santa Cruz, USA. He is currently Senior Scientist at Pelican Imaging. His primary research interests are in image and video restoration, depth and imaging pipelines for computational cameras.

Kartik Venkataraman is currently Founder and CTO at Pelican Imaging. He previously headed the Computational Camera Group at Micron Imaging (Aptina). He spearheaded the design of Extended Depth of Field (EDOF) imaging systems for mobile imaging and an end-to-end simulation environment for camera system architecture and module simulations that was adopted in parts of the mobile imaging ecosystem. Previously at Intel, Kartik was principally associated with investigating CPU design, 3D graphics, medical imaging and visualization together with Johns Hopkins Medical School. Venkataraman received his Ph.D. in Computer Science from UC, Santa Cruz, MS in Computer Engineering from Amherst, and B.Tech (Honors) in EE from the IIT, Kharagpur.

DOI: 10.2352/ISSN.2470-1173.2016.13.IQSP-009

Calibrated | Incorrect calibration | Recovered calibration

Figure 3: Super-resolution results as part of the PiCam imaging pipeline

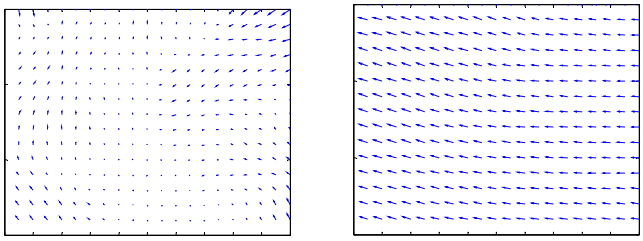

Figure 4: Vector plots of residual vector field for Calibrated vs Incorrect calibration (Array cam 10)

DOI: 10.2352/ISSN.2470-1173.2016.13.IQSP-009

Calibrated | Incorrect calibration | Recovered calibration

Figure 5: Wireframe depth maps (depth map computed in textured areas)

	R		G		B	
	Incorrect calibration	Recovered calibration	Incorrect calibration	Recovered calibration	Incorrect calibration	Recovered calibration
Image #1	0.79	0.97	0.88	0.97	0.61	0.92
Image #2	0.70	0.89	0.87	0.96	0.70	0.93
Image #3	0.72	0.90	0.88	0.96	0.69	0.91
Image #4	0.60	0.94	0.81	0.95	0.55	0.93
Image #5	0.77	0.96	0.88	0.97	0.62	0.88

Figure 6: SSIM scores with respect to the calibrated case (various images)

Temperature	R		G		B	
	Incorrect calibration	Recovered calibration	Incorrect calibration	Recovered calibration	Incorrect calibration	Recovered calibration
40 °C	0.95	0.97	0.97	0.98	0.97	0.98
49 °C	0.91	0.97	0.95	0.98	0.94	0.97
59 °C	0.81	0.96	0.91	0.98	0.82	0.97
69 °C	0.73	0.96	0.89	0.97	0.67	0.96

Figure 7: SSIM scores with respect to the calibrated case for various temperatures (same image)

DOI: 10.2352/ISSN.2470-1173.2016.13.IQSP-009

Calibrated image

Incorrect calibration

Recovered calibration

Figure 8: Additional super-resolution results