

Planar Catadioptric Stereo: Geometry and Calibration *

Joshua Gluckman and Shree K. Nayar

Department of Computer Science

Columbia University

New York, NY 10027

Abstract

By using mirror reflections of a scene, stereo images can be captured with a single camera (catadioptric stereo). Single camera stereo provides both geometric and radiometric advantages over traditional two camera stereo. In this paper, we discuss the geometry and calibration of catadioptric stereo with two planar mirrors and show how the relative orientation, the epipolar geometry and the estimation of the focal length are constrained by planar motion. In addition, we have implemented a real-time system which demonstrates the viability of stereo with mirrors as an alternative to traditional two camera stereo.

1 Introduction

Optical systems consisting of a combination of refracting (lens) and reflecting (mirror) elements are called *catadioptric* systems [9]. Stereo is one area of computer vision which can benefit from such systems. By using two or more mirrored surfaces, a stereo view can be captured by a single camera (catadioptric stereo). This has the following advantages over traditional two camera stereo.

- **Identical System Parameters:** Lens, CCD and digitizer parameters such as blurring, lens distortions, focal length, spectral response, gain, offset, pixel size, etc. are identical for the stereo pair. Having identical system parameters facilitates stereo matching.
- **Ease of Calibration:** Because only a single camera and digitizer are used, there is only one set of intrinsic calibration parameters. Furthermore, we will show that the extrinsic calibration parameters are constrained by planar motion. Together these constraints reduce the total number of calibration parameters from 16 in traditional stereo to 10 in our case.
- **Data Acquisition:** Camera synchronization is not an issue because only a single camera is used. Stereo data can easily be acquired and conveniently stored with a standard video recorder without the need to synchronize multiple cameras.

*This work was supported in part by DARPA's VSAM Image Understanding Program, under ONR contract N00014-97-1-0553.

Previously, several researchers have demonstrated the use of both curved and planar mirrors to acquire stereo data with a *single* camera. Curved mirrors have been primarily used to capture a wide field of view. Systems using spherical and convex mirrors have been suggested by [19] and [24]. Nene and Nayar [20] enumerated the class of catadioptric stereo configurations which preserve a single viewpoint. The epipolar geometry of curved mirrors was studied in [25].

As pointed out by several researchers [26] [21] [18], it is possible to reconstruct a scene by imaging the scene reflection in a rotating planar mirror. However, these systems require more than one image and therefore a static scene. Mitsumoto *et al.* [17] previously described a stereo method which images an object and its reflections in a set of planar mirrors. Here, the mirrors were used to obtain occlusion free images of the object. Recently, Shashua suggested using catadioptric stereo for non-rigid stereo platforms [23].

A few researchers have demonstrated the use of a single camera and two or more planar mirrors to acquire stereo data in a single image [5] [14] [10]. However, the geometry and calibration of these systems has yet to be explored. In this paper, we will discuss several calibration issues in regard to *single* camera stereo with *planar* mirrors, including relative orientation, epipolar geometry, and determination of the focal length. These results will provide a theoretical foundation for catadioptric stereo. In addition, we will describe a real-time catadioptric stereo system which demonstrates the viability of stereo with mirrors as an alternative to traditional two camera stereo.

2 Geometry and Calibration

Figure 1 depicts the geometry of a catadioptric system with two planar mirrors. A scene point P is imaged as if seen from two different viewpoints v and v' . The location of the two virtual pinholes is found by reflecting the camera pinhole c about each mirror. Reflecting the optical axis of the camera about the mirrors determines the optical axes and thus the orientations of the two virtual cameras. The focal length of each virtual camera is equal to f , the focal length of the real camera. Therefore, the locations and orientations of the two virtual cameras are determined by the orientations and distances of the two mirrors with respect

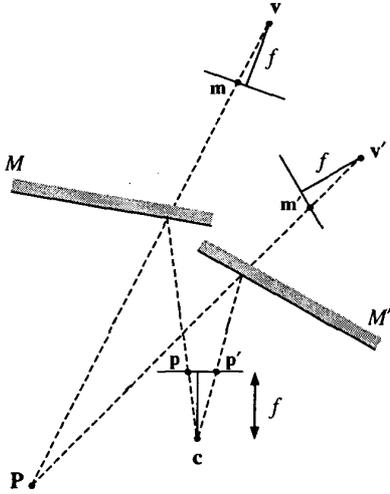


Figure 1: Stereo image formation with a single camera and two planar mirrors. A scene point P reflected by mirrors M and M' is imaged as if seen from two different viewpoints v and v' .

to the pinhole and optical axis of the camera.

2.1 Relative Orientation

In traditional stereo, the two cameras can be placed in any configuration, and therefore the relative orientation between the cameras is described by 6 parameters (3 for rotation and 3 for translation). For catadioptric stereo, the relative orientation between the two virtual cameras, regardless of the configuration of the mirrors, is restricted to planar motion (the direction of translation must lie in the plane normal to the axis of rotation). This constraint reduces the number of degrees of freedom of relative orientation from 6 to 5 (3 for rotation and 2 for translation in a plane).

To derive this result we consider the rigid transformation D between the two reflected viewpoints v and v' . Each virtual viewpoint is related to the camera center c as,

$$v = D_1 c$$

and

$$v' = D_2 c,$$

where D_1 and D_2 are reflection transformations. Then, the relative orientation D is simply,

$$D = D_2 D_1^{-1}.$$

Representing the two mirrors as planes with normals n_1 and n_2 and distances d_1 and d_2 measured from the camera center c , the reflection transformations for the two mirrors are given by

$$D_1 = \begin{bmatrix} \mathbf{I} - 2n_1 n_1^T & 2d_1 n_1 \\ \mathbf{0} & 1 \end{bmatrix}$$

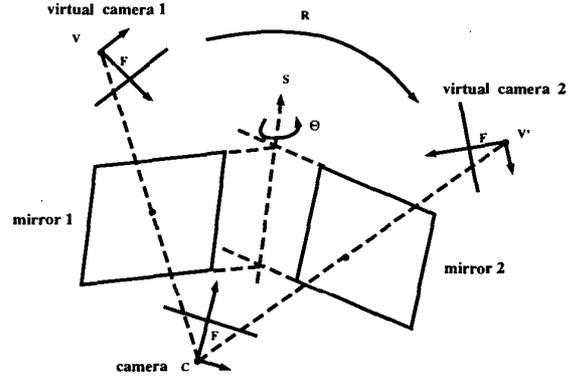


Figure 2: The relative orientation between the two virtual cameras is described by a pure rotation about the axis S (called a screw axis) which is the intersection of the planes containing the two mirrors. A pure rotation about this axis constrains the translation between the virtual cameras to lie in a plane perpendicular to the screw axis, and is thus called planar motion. This planar motion constraint holds true for any locations and orientations of the two mirrors.

and

$$D_2 = \begin{bmatrix} \mathbf{I} - 2n_2 n_2^T & 2d_2 n_2 \\ \mathbf{0} & 1 \end{bmatrix}.$$

Since the inverse of a reflection transformation is itself, the relative orientation of the two virtual cameras is simply,

$$D = D_2 D_1 = \begin{bmatrix} \mathbf{R} & t \\ \mathbf{0} & 1 \end{bmatrix}$$

where

$$\mathbf{R} = \mathbf{I} + 4(n_1 \cdot n_2)n_1 n_2^T - 2n_1 n_1^T - 2n_2 n_2^T,$$

and

$$t = 2d_1 n_1 - (2d_1(n_1 \cdot n_2) + 2d_2)n_2. \quad (1)$$

Note that $(n_1 \times n_2) = \mathbf{R}(n_1 \times n_2)$ and therefore, the rotation matrix \mathbf{R} has a rotational axis of $n_1 \times n_2$. From (1) the direction of translation lies in the plane defined by n_1 and n_2 . Therefore, the rotational axis is normal to the plane containing the direction of translation. This type of motion is termed *planar motion*.

Figure 2 shows how the relative orientation between the two virtual cameras is constrained by planar motion. This constraint arises from the fact that the virtual cameras are related by a pure rotation about the axis S (called a *screw axis*) which is the intersection of the planes containing the two mirrors. A pure rotation about S constrains the motion between the virtual cameras to lie in a plane perpendicular to S . We will be referring to the plane which contains

the camera centers (both real and virtual) as the *plane of motion*. This plane is orthogonal to the screw axis and its intersection with the image plane is termed the *horizon line* of the plane of motion.

As we have seen, single camera stereo with two planar mirrors constrains the external calibration parameters to planar motion, reducing the number of extrinsic parameters from 6 for two camera stereo to 5 for catadioptric stereo. Because only a single camera is used, the intrinsic parameters (focal length, pixel size, image center, skew) are exactly the same for the two stereo views, reducing the number of intrinsic parameters from 10 to 5. Together, these constraints place restrictions on the epipolar geometry.

2.2 Epipolar Geometry

Epipolar geometry is a description of the geometric relationship between a pair of stereo images. It is represented by the fundamental matrix \mathbf{F} and is the minimal information necessary to determine the epipolar lines [3]. For a pair of image correspondences \mathbf{p} and \mathbf{p}' , \mathbf{F} introduces the following well-known epipolar constraint:

$$\mathbf{p}'^T \mathbf{F} \mathbf{p} = 0. \quad (2)$$

\mathbf{F} is also known as the uncalibrated version of the essential matrix \mathbf{E} described by Longuet-Higgins [13], because

$$\mathbf{F} = \mathbf{A}'^{-T} \mathbf{E} \mathbf{A}^{-1}, \quad (3)$$

where \mathbf{A}' and \mathbf{A} are matrices representing the internal calibration parameters of the stereo cameras. Both \mathbf{F} and \mathbf{E} are rank 2 matrices. For an arbitrary stereo pair the rank 2 constraint is the only constraint on the fundamental matrix.

When the intrinsic parameters are identical for both views and the relative orientation is limited to planar motion, an additional constraint is imposed on the fundamental matrix. From a result due to Maybank [16], it is known that the symmetric part of the essential matrix, $\mathbf{E} + \mathbf{E}^T$, is rank 2 for planar motion. When the intrinsic parameters remain constant ($\mathbf{A}' = \mathbf{A}$), it is simple to show that the symmetric part of the fundamental matrix is also rank 2, providing the following additional constraint on the fundamental matrix,

$$\det(\mathbf{F} + \mathbf{F}^T) = 0. \quad (4)$$

This constraint reduces the number of free parameters in the fundamental matrix from 7 to 6 and has been used by [2] and [29] to help constrain the self-calibration problem.

When estimating the fundamental matrix from image correspondences it is useful to have a parameterization of \mathbf{F} which implicitly enforces (4). We can derive such a parameterization by considering the image projection of the screw axis.

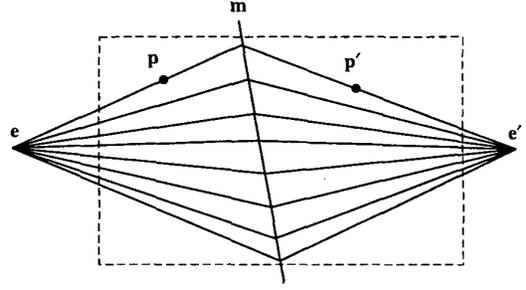


Figure 3: The epipolar geometry due to planar motion. When motion is constrained to lie in a plane, all corresponding epipolar lines must intersect at \mathbf{m} the image projection of the screw axis. Therefore, the two epipoles \mathbf{e} and \mathbf{e}' and the line \mathbf{m} completely determine the epipolar geometry. Because the plane of motion contains the two virtual camera centers, the horizon line of the plane of motion is the line containing the two epipoles.

The screw axis (see figure 2) remains fixed with respect to the coordinate systems of the two virtual cameras, therefore its image \mathbf{m} is identical for both the left and right stereo views. This implies that corresponding epipolar lines must intersect on \mathbf{m} . The resulting epipolar geometry is depicted in figure 3. As shown in this figure, the epipolar line of a point \mathbf{p} is the line containing epipole \mathbf{e}' and the intersection of \mathbf{m} with the line through epipole \mathbf{e} and point \mathbf{p} . If \mathbf{p} and \mathbf{p}' are corresponding points then

$$\mathbf{p}' \cdot (\mathbf{e}' \times (\mathbf{m} \times (\mathbf{e} \times \mathbf{p}))) = 0, \quad (5)$$

and from (2) the fundamental matrix has the form¹

$$\mathbf{F} = [\mathbf{e}']_{\times} [\mathbf{m}]_{\times} [\mathbf{e}]_{\times}. \quad (6)$$

With the help of a symbolic algebra package, we have confirmed that the parameterization given in (6) does indeed enforce the planar motion constraint (4). Other parameterizations of the fundamental matrix for planar motion are also possible, see for instance [29].

Using (6) and a set of image correspondences, \mathbf{F} can be determined by searching the parameter space of \mathbf{e} , \mathbf{e}' and \mathbf{m} while minimizing a suitable cost function such as the sum of distances of corresponding points from their epipolar lines. This process requires non-linear minimization and thus initial estimates of \mathbf{e} , \mathbf{e}' and \mathbf{m} are needed.

Initial estimates of \mathbf{e} , \mathbf{e}' and \mathbf{m} can be extracted from an estimate of \mathbf{F} obtained by the linear 8-point algorithm [7]. \mathbf{e} and \mathbf{e}' can be extracted from the left and right null space of \mathbf{F} . Using the following equations, \mathbf{m} can be obtained from the eigenvectors of the symmetric part of fundamental

¹ $[\cdot]_{\times}$ is the matrix form of the cross product operator such that for any vectors \mathbf{p} and \mathbf{q} , $(\mathbf{p} \times \mathbf{q}) = [\mathbf{p}]_{\times} \mathbf{q}$.

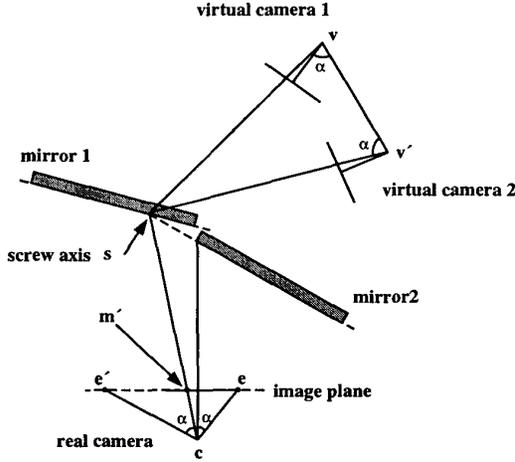


Figure 4: In the plane of motion the two virtual camera centers will be the same distance from the screw axis. As a result, the angles formed by the virtual camera centers v and v' and the screw axis s are equivalent. This constraint can be used to recover the focal length provided that m' is not the image center.

matrix $F_s = F + F^T$ [1]. Letting λ_1, λ_2 and $\mathbf{n}_1, \mathbf{n}_2$ be the positive and negative eigenvalues and eigenvectors of F_s , we have either

$$\mathbf{m} = \sqrt{\lambda_1} \mathbf{n}_1 + \sqrt{-\lambda_2} \mathbf{n}_2 \quad (7)$$

or

$$\mathbf{m} = \sqrt{\lambda_1} \mathbf{n}_1 - \sqrt{-\lambda_2} \mathbf{n}_2. \quad (8)$$

The ambiguity can be resolved by noting that one of these expressions is equivalent to $\mathbf{e} \times \mathbf{e}'$ and \mathbf{m} is the other one.

As we have shown, catadioptric stereo with planar mirrors introduces an additional constraint on the fundamental matrix which reduces the number of parameters to estimate from 7 to 6. Next, we will discuss recovering the focal length from a single catadioptric stereo image.

2.3 Recovering the Focal Length

With knowledge of the fundamental matrix, the scene geometry can be reconstructed up to an unknown projective transform [3]. To obtain a Euclidean reconstruction from a stereo pair, it is necessary to determine the internal camera parameters. With video cameras, it is often the case that the aspect ratio is known, the skew is zero, and the image center is roughly the center of the image; therefore, Euclidean reconstruction amounts to determining the focal lengths of the cameras. Through the Kruppa equations [30], the fundamental matrix places two quadratic constraints on the internal calibration parameters. As demonstrated by Hartley [6], these two constraints are sufficient to solve for the focal lengths when the other internal parameters are known.

For catadioptric stereo, we have only one unknown focal length f and we can solve for f from the Kruppa equations

$$\mathbf{F} \omega \mathbf{F}^T = x [\mathbf{e}']_x \omega [\mathbf{e}']_x, \quad (9)$$

where

$$\omega = \begin{bmatrix} f^2 & 0 & 0 \\ 0 & f^2 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

and x is an unknown scale factor (\mathbf{F} and \mathbf{e}' are projective quantities and thus only known up to a scale factor).

Though f can be determined in this manner, we would like to explore additional constraints on the focal length which arise from the planar motion. It turns out there are two such constraints. The first results from the fact that the plane of motion is perpendicular to the plane which contains the screw axis and the camera center. The plane of motion projects to the horizon line ($\mathbf{e} \times \mathbf{e}'$) and the plane containing the screw axis and the camera center projects to \mathbf{m} (the image of the screw axis). The 3D angle between the visual planes of two image lines \mathbf{x} and \mathbf{y} is given by [27],

$$\cos \theta = \frac{(\mathbf{x}^T \omega \mathbf{y})}{\sqrt{(\mathbf{x}^T \omega \mathbf{x})(\mathbf{y}^T \omega \mathbf{y})}}. \quad (10)$$

From this, we can define the following constraint between f and the lines $\mathbf{l} = (\mathbf{e} \times \mathbf{e}')$ and \mathbf{m} ,

$$0 = \frac{(\mathbf{l}^T \omega \mathbf{m})}{\sqrt{(\mathbf{l}^T \omega \mathbf{l})(\mathbf{m}^T \omega \mathbf{m})}}. \quad (11)$$

A second constraint can be derived from the image points \mathbf{e}, \mathbf{e}' , and the point $\mathbf{m}' = \mathbf{l} \times \mathbf{m}$, which is the intersection of the image of the screw axis \mathbf{m} and the horizon line \mathbf{l} . From figure 4, we can see that the angle formed between the image rays through \mathbf{e} and \mathbf{m}' is equal to the angle formed by \mathbf{e}' and \mathbf{m}' . Using a relationship similar to (10) but for image rays [27], we can express this as

$$\frac{(\mathbf{e}^T \omega^{-1} \mathbf{m}')}{\sqrt{(\mathbf{e}^T \omega^{-1} \mathbf{e})(\mathbf{m}'^T \omega^{-1} \mathbf{m}')}} = \frac{(\mathbf{e}'^T \omega^{-1} \mathbf{m}')}{\sqrt{(\mathbf{e}'^T \omega^{-1} \mathbf{e}')(\mathbf{m}'^T \omega^{-1} \mathbf{m}')}} \quad (12)$$

When using these equations to recover the focal length, care must be taken to avoid degenerate configurations. In particular, when \mathbf{m} passes through the image center, (12) will not lead to a solution for f . We can ensure \mathbf{m} does not pass through the image center by displacing the mirrors as in figure 4. Equation (11) can not be used when \mathbf{m} and \mathbf{l} are perpendicular. Avoiding this configuration is more difficult, it requires displacing the mirrors and tilting the camera upward or downward with respect to the mirrors.

2.4 Experiments: Recovering the Focal Length

To test the accuracy of recovering the focal length from a single catadioptric image, we performed a series of experiments comparing the focal length obtained from the

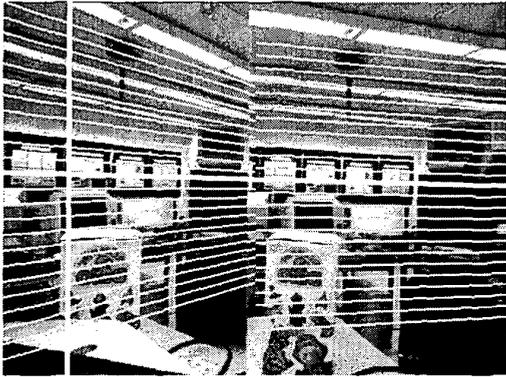


Figure 5: The recovered epipolar geometry of a catadioptric image. The vertical thick line is \mathbf{m} the estimated image of the screw axis, where the corresponding epipolar lines meet. The horizontal thick line is the line connecting the two epipoles, the horizon line of the planar motion.

angle constraint (12) to the focal length obtained from Tsai calibration [28].

We took a series of 10 catadioptric images using a Sony XC-75 camera with a Computar 4mm pinhole lens (no radial distortions are present). For each image the mirrors were placed in a configuration similar to figure 4 in order to avoid \mathbf{m} passing through the image center. Throughout the sequence we varied the angle between the mirrors and used several different scenes.

For each catadioptric image we found an initial estimate $\hat{\mathbf{F}}$ of the fundamental matrix and a set of correspondences using the robust method of [31] and publicly available at www.inria.fr/robotvis/personnel/zzhang/zzhang-end.html. We then enforced the planar motion constraint (4) by performing non-linear optimization using the parameterization defined in (6). The error criteria minimized was the sum of squared distances to epipolar lines and the Levenberg-Marquardt algorithm [22] was used to perform the minimization. After minimization equation (12) and the estimates of \mathbf{e} , \mathbf{e}' and \mathbf{m} were used to obtain an estimate of the focal length.

Figure 5 shows a typical scene and the recovered epipolar geometry and the following table shows the estimated focal lengths (in pixels) for each image as compared to the focal length obtained from Tsai calibration.

catadioptric images					Tsai
440	437	439	440	437	407
433	439	420	428	423	

The discrepancy of the focal lengths for the catadioptric images as compared to Tsai is possibly due to several factors: inaccurate localization of the epipoles, bias from assuming the location of the center of projection is the image

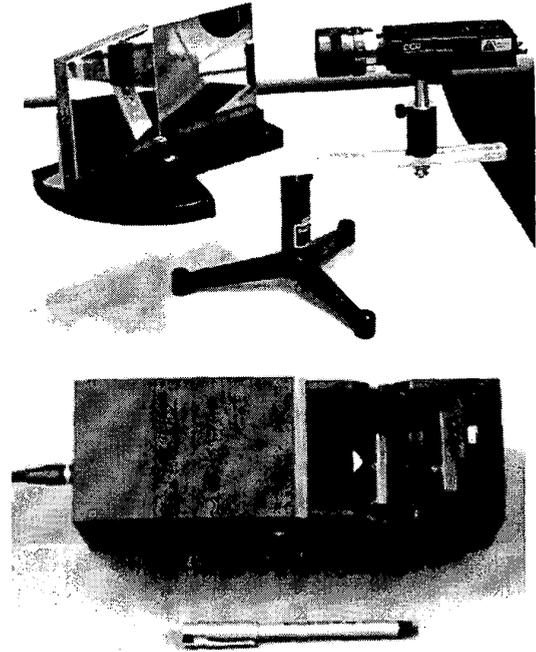


Figure 6: Catadioptric stereo sensors. (top) A single Sony XC-77 b/w camera and a 12.5mm Computar lens is used with two high quality front silvered Melles Griot 5" mirrors. The distance between the camera and mirrors can be altered, which changes the baseline of the stereo system. The angle between the mirrors can also be adjusted to control vergence and rotation between the stereo views. (bottom) This compact unit uses a single Sony XC-75 b/w camera and a 4mm Computar pinhole lens with 2" Melles Griot mirrors.

center, and/or inaccuracies in Tsai calibration.

3 Real-Time Implementation

Real-time stereo systems have been implemented by several researchers [4] [15] [11] [12]. All of these systems use two or more cameras to acquire stereo data. Here, we describe a real-time catadioptric stereo system which uses a single camera and only a PC to compute depth-maps in real-time. Figure 6 shows a picture of two catadioptric stereo sensors we have designed.

3.1 Calibration and Rectification

To achieve real-time performance it is necessary to have scanline correspondence between the stereo pair. This allows stereo matching algorithms to be implemented efficiently as described by Faugeras *et al.* [4]. Because catadioptric stereo requires rotated mirrors (if only two mirrors are used) and hence rotated views, we must rectify the stereo pair at run-time. To compute the rectification transform we first need to estimate the fundamental matrix.

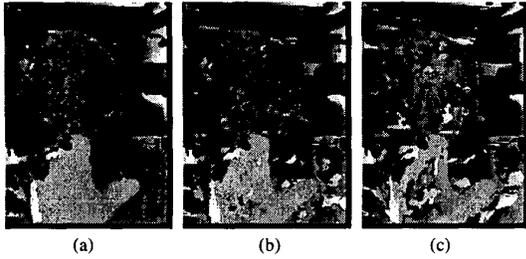


Figure 7: Comparison of three commonly used measures of similarity on an image taken by a catadioptric stereo sensor. (a) Depth map computed using sum of absolute differences. (b) Depth map computed using normalized cross-correlation. (c) Depth map computed using sum of absolute differences after a Laplacian of Gaussian operator was applied. For all three measures a 5×5 window was used and no thresholds were applied.

An estimate of the fundamental matrix is found using the method described in the previous sections. After computing the fundamental matrix, we find a rectification transform using the method of Hartley and Gupta [8]. Once computed, this transform is used to warp each incoming image at run-time. The brightness value of each pixel in the warped image is determined by back projecting to the input image through the rectification transform and bilinearly interpolating among adjacent pixels.

3.2 Stereo Matching

The underlying assumption of all stereo matching algorithms is that the two image projections of a scene patch are similar. The degree of similarity is computed using a variety of measures such as brightness, texture, color, edge orientation, etc. To minimize computations, most real-time systems use a measure of similarity based on image brightness. However, differences in focal settings, lens blur and gain control between the two cameras results in the two patches having different intensities. For this reason many methods, such as normalized cross-correlation, Laplacian of Gaussian, and normalized sum of squared differences, have been employed to compensate for camera differences [4] [15] [11] [12]. By using a single camera, catadioptric stereo avoids both the computational cost and loss of information which results from using these methods. As figure 7 shows, normalized cross-correlation and the Laplacian of Gaussian can degrade the performance of stereo matching due to loss of information and finite arithmetic. By removing differences in offset and gain, normalized cross-correlation and the Laplacian of Gaussian also remove shading information which is useful for matching.

One of the simplest measures of similarity between two image patches is the sum of absolute differences (SAD). Because we use only a single camera, SAD is a suitable

choice. SAD keeps the data size small and is easily implemented on SIMD (single instruction multiple data) processors such as those with MMX technology. Furthermore, SAD lends itself to efficient scanline correspondence algorithms.

In our implementation, stereo matches are found by using a standard window based search. The search is limited to an interval of 32 pixels along the epipolar line (scanline) of a 320×240 image. By using the SAD measure, scanline correspondence, and SIMD instructions we were able to achieve a throughput of approximately 20 frames per second on a 300Mhz Pentium II machine. An example catadioptric stereo image and computed depth map are shown in figure 8.

4 Conclusion

We have examined the geometry of stereo with two planar mirrors in an arbitrary configuration and shown that both the relative orientation and the epipolar geometry are constrained by planar motion. In addition, we have shown how the focal length can be extracted from a single catadioptric image. We have also implemented a real-time stereo system which demonstrates that high quality depth maps can be obtained when a single camera is used.

In this paper, we have investigated the geometric properties of catadioptric stereo. In the future we intend to examine the radiometric properties. Although single camera stereo eliminates inter-camera differences, intra-camera differences still remain. We intend to research pixel differences across the CCD and the $\cos^4(\alpha)$ decay in image irradiance. Both of these may result in different intensities at corresponding image points. However, through calibration these effects can be measured and removed, thus further improving the quality of the stereo matching.

In conclusion, we feel that the sensor used to acquire the stereo data is just as important as the algorithm used for matching. In this respect, catadioptric stereo offers a significant benefit by improving the quality of the stereo data at no additional computational cost.

References

- [1] M. Armstrong. *Self-calibration from image sequences*. PhD thesis, University of Oxford, 1996.
- [2] M. Armstrong, A. Zisserman, and R. Hartley. Self-calibration from image triplets. In *Proceedings of the 1996 European Conference on Computer Vision*, 1996.
- [3] O. Faugeras. What can be seen in three dimensions with an uncalibrated stereo rig? In *Proceedings of the 1992 European Conference on Computer Vision*, pages 563–578. Springer-Verlag, 1992.
- [4] O. Faugeras, B. Hotz, H. Mathieu, T. Vieville, Z. Zhang, P. Fau, E. Theron, L. Moll, G. Berry, J. Vuillemin, P. Bertin, and C. Proy. Real-time correlation-based stereo: algorithm, implementation and application. Technical Report 2013, INRIA Sophia Antipolis, 1993.
- [5] A. Goshtasby and W. Gruver. Design of a single-lens stereo camera system. *Pattern Recognition*, 26(6):923–937, 1993.
- [6] R. Hartley. Estimation of relative camera positions for uncalibrated cameras. In *Proceedings of the 1992 European Conference on Computer Vision*, 1992.



Figure 8: Stereo image and depth map. On the left is an image taken by a catadioptric stereo system and on the right is the depth map computed with the SAD measure and a 7×7 window.

- [7] R. Hartley. In defense of the 8-point algorithm. In *Proceedings of the 5th International Conference on Computer Vision*, pages 1064–1070, 1995.
- [8] R. Hartley and R. Gupta. Computing matched-epipolar projections. In *Proceedings of the 1993 Conference on Computer Vision and Pattern Recognition*, 1993.
- [9] E. Hecht and A. Zajac. *Optics*. Addison-Wesley, 1974.
- [10] M. Inaba, T. Hara, and H. Inoue. A stereo viewer based on a single camera with view-control mechanism. In *Proceedings of the International Conference on Robots and Systems*, July 1993.
- [11] T. Kanade, A. Yoshida, K. Oda, H. Kano, and M. Tanaka. A stereo machine for video-rate dense depth mapping and its new applications. In *Proceedings of the 1996 Conference on Computer Vision and Pattern Recognition*, 1996.
- [12] K. Konolige. Small vision systems: hardware and implementation. In *8th Int'l Symposium of Robotics Research*, Hayama, Japan, Oct. 1997.
- [13] H. Longuet-Higgins. A computer algorithm for reconstructing a scene from two projections. *Nature*, 293:133–135, 1981.
- [14] H. Mathieu and F. Devernay. Systeme de miroirs pour la stereoscopie. Technical Report 0172, INRIA Sophia-Antipolis, 1993. in French.
- [15] L. Matthies. Stereo vision for planetary rovers: stochastic modeling to near realtime implementation. *International Journal of Computer Vision*, 8(1):71–91, 1993.
- [16] S. Maybank. *Theory of reconstruction from image motion*. Springer-Verlag, Berlin, 1993.
- [17] H. Mitsumoto, S. Tamura, K. Okazaki, N. Kajimi, and Y. Fukui. 3d reconstruction using mirror images based on a plane symmetry recovery method. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(9):941–945, 1992.
- [18] D. Murray. Recovering range using virtual multicamera stereo. *Computer Vision and Image Understanding*, 61(2):285–291, 1995.
- [19] S. Nayar. Robotic vision system. United States Patent 4,893,183, Aug. 1988.
- [20] S. Nene and S. Nayar. Stereo with mirrors. In *Proceedings of the 6th International Conference on Computer Vision*, Bombay, India, January 1998. IEEE Computer Society.
- [21] Y. Nishimoto and Y. Shirai. A feature-based stereo model using small disparities. In *cvpr87*, pages 192–196, 1987.
- [22] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. *Numerical Recipes in C*. Cambridge University Press, 1992.
- [23] A. Shashua. Omni-rig sensors: what can be done with a non-rigid vision platform? In *Workshop on applications of computer vision*, 1998.
- [24] D. Southwell, A. Basu, M. Fiala, and J. Reyda. Panoramic stereo. In *Proceedings of the Int'l Conference on Pattern Recognition*, 1996.
- [25] T. Svoboda, T. Pajdla, and V. Hlavac. Epipolar geometry for panoramic cameras. In *Fifth European Conference on Computer Vision*, pages 218–232, 1998.
- [26] W. Teoh and X. Zhang. An inexpensive stereoscopic vision system for robots. In *Proc. Int'l Conference on Robotics*, pages 186–189, 1984.
- [27] B. Triggs. Autocalibration and the absolute quadric. In *Proceedings of the 1997 Conference on Computer Vision and Pattern Recognition*, 1997.
- [28] R. Tsai. An efficient and accurate camera calibration technique for 3d machine vision. 1986.
- [29] T. Vieville and D. Lingrand. Using singular displacements for uncalibrated monocular visual systems. Technical Report 2678, INRIA Sophia-Antipolis, 1995.
- [30] C. Zeller and O. Faugeras. Camera self-calibration from video sequences: the kruppa equations revisited. Technical Report 2793, INRIA, Sophia-Antipolis, France, 1996.
- [31] Z. Zhang, R. Deriche, O. Faugeras, and Q. Luong. A robust technique for matching two uncalibrated images through the recovery of the unknown epipolar geometry. *Artificial Intelligence Journal*, 78, 1995.