

Rectification for Any Epipolar Geometry

Daniel Oram
Advanced Interfaces Group
Department of Computer Science
University of Manchester
Manchester, M13, UK
oramd@cs.man.ac.uk

Abstract

This paper proposes a new approach to rectification, a process whereby a stereo pair of images is re-sampled so as to make imposing the two view geometric constraint simple. This is normally performed by applying a single linear transformation per image, a method which has the drawback that some camera motions produce rectified images which are heavily distorted or unbounded. More recent attempts have used nonlinear transformations to allow any camera motion, but as a consequence distort images so that matching features no longer look the same in both images. This work provides a hybrid linear/nonlinear method that greatly reduces this problem, and simplifies the technique. The technique also provides further improvements by selecting the rectifying transformation so as to minimise perspective effects between the images.

1 Introduction

Rectification is a process used to facilitate the analysis of a stereo pair of images by making it simple to enforce the two view geometric constraint. For a pair of related views the *epipolar geometry* provides a complete description of relative camera geometry. Once the epipolar geometry has been determined it is possible to constrain the match for a point in one image to lie on a line (the *epipolar line*) in the other image and vice-versa. The process of rectification makes it very simple to impose this constraint by making all matching epipolar lines coincident and parallel with an image axis. Many stereo algorithms assume this simplified form because subsequent processing becomes much easier if differences between matched points will be in one direction only.

In the past, stereo images were primarily rectified using optical techniques [11], but more recently these have been replaced by software techniques. These model the geometry of optical projection by applying a single linear transformation to each image, effectively rotating both cameras until the image planes are the same (see [4, 3] for a small sample). Such techniques are often referred to as *planar rectification*. The advantages of this linear approach are that it is mathematically simple, fast and preserves image features such as straight lines.

Unfortunately, planar rectification is not general and if there is a large forward component in the camera movement it may produce unbounded, large or badly warped images. In the past this was not a problem because stereo vision was usually performed using stereo rigs with near parallel cameras. However, recent advances in uncalibrated stereo




Figure 1: The epipolar geometry

vision (e.g. [1, 5]), have focused more on hand held sequences in which forward movement can frequently be present.

To deal with these restrictions, a cylindrical rectification technique was proposed in [10] which used a separate transformation for each epipolar line. However, the technique was complex, omitted many implementation details and worked largely in 3D. A later work, [8] overcame most of these problems, by using the tools of oriented projective geometry to perform a similar nonlinear rectification without using 3D.

In this work, a new general rectification technique is presented which further improves on these techniques. Firstly, it uses existing matches between the images, e.g. those used to calculate the epipolar geometry, to determine a rectification which minimises perspective distortion effects such as motion parallax between the images. Secondly, by using the same nonlinear transformation for both images, it overcomes a problem with existing general techniques where the application of different nonlinear transformations to each image results in matching image features being warped differently in each image. Finally it also simplifies the approach of [8] and handles a number of other minor problems overlooked by previous work such as sub pixel coordinates and infinite epipoles.

2 Preliminaries

Throughout this text, column vectors will be referred to with bold lower case letters, such as \mathbf{x} , and row vectors as transposed column vectors \mathbf{x}^T . Matrices will be represented by an upper case letter e.g. A . Also, the notation \approx will be used to indicate equality of vectors or matrices subject to multiplication by a non-zero scale factor. Priming will be used to indicate quantities in the second image of a pair e.g. \mathbf{m}' . Familiarity with basic projective geometry and homogeneous notation will be assumed.

2.1 Epipolar Geometry

Figure 1 gives a pictorial representation of the central projection of a point M onto two image regions I_1 and I_2 by two cameras with centres of projection O and O' . Given only I_2 , all that can be determined about M from its image \mathbf{m}' is that M must lie on the infinite line defined by O' and M (the back projection of \mathbf{m}'). As a direct consequence of this,




Figure 2: Oriented epipolar geometry. Matching for the point M can be seen to be restricted to a half epipolar line from the epipole. The point N illustrates that if the point matches to the other half of the epipolar line in $I1$ it must be behind $I2$.

given $I1$ as well, the matching projection of M i.e. \mathbf{m} must lie on the projection of the line $\langle O', M \rangle$ in $I1$.

The projection of $\langle O', M \rangle$ is also known as an epipolar line and it can be seen that all points m'_k in $I2$ will generate a pencil of epipolar lines in $I1$ all containing the point e . The point e is known as the epipole and is the image of the point O' , the only point common to all back projected lines in $I2$.

If the same reasoning is followed through for the projection of the line $\langle O, M \rangle$ in $I2$ then a matching epipolar line is found. Corresponding epipolar lines can be seen as being the intersection of the *epipolar plane* with the images where the epipolar plane is defined by the camera centres O, O' and the world point M . It can be seen that as M varies the epipolar plane rotates around the baseline forming a pencil of epipolar planes known as the epipolar pencil.

The epipolar geometry can be represented in the projective case by the fundamental matrix. For details of this and how it may be estimated the reader is referred to any of the numerous works on the subject, for example [7].

2.2 Oriented Projective Geometry

The concept of oriented projective geometry was first applied to the structure from motion problem in [2], and refined in [6], where it was shown that it was possible to improve a standard projective reconstruction by distinguishing between points in front of and behind the camera. This was achieved using the constraint that all points known to be visible in an image must be in front of the camera and expressed in the reconstruction of the 3D points by enforcing the convention that points in front of the camera project $\lambda \mathbf{m} = P\mathbf{X}$ to a positive scale factor λ and those behind to a negative scale factor.

2.2.1 Orientation and Half Epipolar Lines

The notion of oriented projective geometry can also be applied to the epipolar geometry, so that a point in one image matches to a half epipolar line in the other image instead of a full epipolar line. This is illustrated in figure 2, where two world points, N and M on

the same epipolar plane Π are projected through the camera centres \mathbf{C} and \mathbf{C}' to form matching image points $\mathbf{m} \leftrightarrow \mathbf{m}'$ and $\mathbf{n} \leftrightarrow \mathbf{n}'$. Since by definition the camera centres and epipoles lie on the same line (the baseline) matches are restricted to being between points on the same side of the base line. If the correspondences are not on the same side of the baseline, the world point will be behind one of the cameras, as illustrated by the point \mathbf{N} . Note that as a consequence orientation of the epipolar geometry is only relevant to matching if the epipole is within the image.

2.3 Homographies Compatible with a Fundamental Matrix

In this section, homographies H providing a planar projective mapping between image points $\mathbf{m} \leftrightarrow \mathbf{m}'$ will be considered i.e. $\mathbf{m} \simeq H\mathbf{m}'$. Such homographies can be used to perform linear planar rectification by applying one to each image so as to make all epipolar lines coincident and parallel (see [3]). However, because making epipolar lines parallel can result in unbounded and badly warped images (see [3]), only homographies which make epipolar lines coincident, but not necessarily parallel will be considered here. These will be termed compatible homographies.

More concisely, given a fundamental matrix F for an image pairing and a match between the images $\mathbf{m} \leftrightarrow \mathbf{m}'$, a compatible homography H will transfer \mathbf{m}' so that the resultant point $H\mathbf{m}'$ lies on the corresponding epipolar line $F\mathbf{m}'$. The result of applying this homography to all the points in one image will by definition be a pair of images which have coincident epipolar lines and hence the same epipole, i.e. $He' \simeq e$ and $H^{-1}e \simeq e'$. The set of homographies that are consistent with the geometry of a particular image pairing can be obtained from the fundamental matrix as (see [3]):

$$H \approx [e']_{\times} F - e' \mathbf{a}^T \quad (1)$$

where \mathbf{a} is an arbitrary 3 vector such that $\det H \neq 0$. Note that this means there is a 3 parameter set of homographies that are compatible with the fundamental matrix. Since compatible homographies can be considered as transforming points on a world plane between two different images these 3 parameters can be considered to represent a plane in the scene. In image terms, these 3 extra parameters amount to defining a one dimensional projective transformation that is applied along all the epipolar lines in an image.

Previous generalised rectification techniques made no attempt to select the free parameters \mathbf{a} in equation 1 using any principled manner. In fact, different versions of these free parameters were applied to each epipolar line in each image. Subsequently, distortion arises which causes features to not look the same in both images for reasons other than perspective or photometric effects. In this work, the use of a compatible homography reduces these problems because the same parameters are applied to matching epipolar lines. However, it does not solve the problem, because different parameters are still applied to each epipolar line.

One final point of note is that a compatible homography is a point to point mapping and so must enforce orientation. This is because assuming the matches are correct, the mapping must map between the correct half epipolar lines because no points will match to incorrect half epipolar lines. Subsequently, there is no need to explicitly enforce orientation as in the method of [8].

3 General Rectification

The rectification method presented here comes in two stages. First, a compatible homography is selected so as to minimise distortion due to perspective effects in some supplied set of matches, and then applied to one image to make all matching epipolar lines coincident. Epipolar lines are then made parallel to an image axis by parameterisation of both images with polar coordinates centred on the epipole. Note that, because a compatible homography has been used on one image, the same nonlinear epipolar alignment process can be used for both images and so problems of inconsistent image warping avoided. On input, the rectification process expects to be provided with two rectangular images as well as a fundamental matrix and a set of point matches such as those used to calculate the fundamental matrix.

3.1 Determining a Compatible Homography

This section will address the method by which a compatible homography is obtained. This method uses known matches and the fundamental matrix between the images to attempt to find a compatible homography that minimises inter image distortion due to perspective effects.

To do this, an attempt is made to find the homography compatible with the fundamental matrix F that transfers points \mathbf{m}'_i in image two as close as possible to their matches in image one \mathbf{m}_i . This will find the best fitting plane for all the observed points (remember a homography transfers points on a world plane between images), and so if the image is then warped to make the found plane exhibit no parallax, perspective effects should be reduced. Assuming the point matches have been identified subject to a Gaussian distributed error, the following least-squares criterion should be minimised for n points:

$$\min_H \sum_{i=1}^{i=n} d_E^2(\mathbf{m}_i, ([\mathbf{e}']_{\times} F - \mathbf{e}' \mathbf{a}^T) \mathbf{m}'_i)$$

where d_E is Euclidean distance and the compatible homography is parameterised as in equation 1. Replacing Euclidean distance with algebraic distance, and $[\mathbf{e}']_{\times} F$ with H , the result is two linear equations $k \in (1, 2)$ in terms of \mathbf{a} per point match $\mathbf{m} \leftrightarrow \mathbf{m}'$:

$$(\mathbf{m}'_k \mathbf{e}'_3 - \mathbf{e}'_k) \mathbf{m}^T \mathbf{a} = (\mathbf{m}'_k \mathbf{h}_3^T - \mathbf{h}_k) \mathbf{m}$$

where subscripts indicate the n th item in a vector and \mathbf{h}_n the n th row of H . Stacking these equations gives a linear system of the form $X \mathbf{a} = \mathbf{b}$ which can be solved using any standard linear least-squares technique.

Whilst this linear algorithm is effective, it does come with the major problem that, even if the matches conform to the epipolar geometry, they can still be incorrect because the epipolar geometry only constrains matches to lie on a line. Consequently, some form of robust solution must be found. One approach is to use a random sampling method to minimise a robust Huber function $\rho(x)$ of the residuals x (σ is the robust standard deviation), i.e:

$$\rho(x) = \begin{cases} x^2 & x^2 < 3.84\sigma \\ 3.84\sigma & x^2 \geq 3.84\sigma \end{cases} \quad \sigma = 1.4826 \left[1 + \frac{5}{n-p} \right] \text{median}_i |r_i|$$

for n observations and a parameter space of dimension p , $p = 3$ for this case (see [9] for full details). This can be achieved by taking m minimal samples ($m = 300$ in my implementation) of 3 points and using them to find a solution \mathbf{a}_n . The solution \mathbf{a}_n , which minimises the Huber function just outlined is then accepted as the best solution. Outliers are rejected using a 95% confidence limit and the robust standard deviation σ of the best solution i.e. $x^2 \geq 3.84\sigma$. Finally, the calculation can be repeated using the linear method.

Even if no point matches are outlying, the robust approach is still advisable because depth variation may mean that some points, although correctly matched will be so far off the best fit plane they will skew the selected plane very badly. This might occur if there is a very large amount of depth variation in an image or a very dominant plane. This is based on the assumption that it is desirable to find the best fit plane in general that will result in the most points matching well, and not to allow small insignificant image regions to skew the fit.

3.2 Unbounded Images

Since a linear transformation is being used, it is possible for the compatible homography to result in an unbounded image. This occurs in the extremely unlikely situation of the epipole in first image being infinite, and the epipole in the second image being within the image. In this case, the compatible homography will cause the epipole in the second image to be mapped to infinity, causing an unbounded second image. Fortunately, such a degeneracy is easily handled by swapping the images so that the homography maps points from image one (infinite epipole) to image two (epipole in image).

Near this degeneracy, problems will also occur with large images. In order to deal with this, the whole technique is modified by swapping the images the homography transfers between so that points are transferred to the image with an epipole closest to the image centre. For simplicity from this point onwards this swapping will be assumed and the image which has the compatible homography applied to it will be considered to be image two of the pair.

3.2.1 Applying to The Image

After determination, the compatible homography can be used to warp all points in the second image to the first image plane as $H\mathbf{m}'_i$, thus making all epipolar lines coincident, giving both images the same epipole and orienting the epipolar geometry. The nonlinear mapping of epipolar lines to the rectified image using polar coordinates will then be exactly the same for both images.

3.3 Rectifying the images

After the compatible homography has been applied to one image, rectification can proceed so as to make epipolar lines parallel with the x axis. This is achieved by parameterising all image points in terms of polar coordinates centred on the epipole. Subsequently, each rectified point is described by a y coordinate given by an associated polar angle $-\pi \leq \theta \leq \pi$, and an x coordinate given by the distance of the rectified point from the epipole.

Recalling section 2.2, it should be noted that it is necessary to consider only positive distances from the epipoles. Points at negative and positive distances do not belong to the

same half epipolar line and so cannot match to the same half epipolar line in the other image.

Before reparameterisation can proceed, it is first necessary to find the common bounds of the rectified images in polar coordinates. This amounts to identifying the range of epipolar lines common to both images as well as the maximum and minimum distance of points from the epipole for both images. Once bounded, the rectified images can then be built up line by line, with the distance between consecutive epipolar lines selected individually so as to avoid pixel compression. The output image is then created as the region that bounds the reparameterised image.

3.3.1 Finding the common region

Before finding the epipolar lines common to both images, it is best to first identify the extreme epipolar lines for both images. If it is assumed that the input images are rectangular, maximal epipolar lines are guaranteed to pass through the image corners and the angle range spanned by the image is guaranteed to be at most π radians. Consequently, the maximal corners can be found by determining the polar angle for each corner, and selecting the maximum and minimum corners so that the total angle range is less than π radians. Because of the restriction on the angle range, this can easily be achieved by first normalising all the angles, so that one angle is 0 radians, and then using normal minimum and maximum.

Since a compatible homography is available, all the image corners from image two can be transferred to image one prior to finding the maximal corners. The common region is then found as the second maximum and second minimum angle, such that the image spans less than π radians. In effect, this means that given minimum angles n, n' and maximum angles x, x' for both images, the second minimum sn and maximum sx are given by:

$$sn = \begin{cases} MIN(n, n') & \|n - n'\| > \pi \\ MAX(n, n') & \|n - n'\| \leq \pi \end{cases} \quad sx = \begin{cases} MAX(x, x') & \|x - x'\| > \pi \\ MIN(x, x') & \|x - x'\| \leq \pi \end{cases}$$

This scheme will fail if an epipole is within an image, because in that case the relevant image will cover an angle range of 2π radians. Fortunately, this is an easily handled anomaly. If the epipole is within one image, then minimum and maximum angles can simply be set to the bounds of the other image. If the epipole is within both images, the maximum and minimum angles can be set to $-\pi$ and π .

3.3.2 Selecting the Epipolar Lines To Rectify

The next step is to build a table that will be used to transfer epipolar lines to and from different scan lines in the rectified image. To do this, the process starts from one extreme epipolar line, assigns it the rectified line $y = 0$ and associates it with the relevant angle. Subsequent epipolar lines $y = n$ are then found by taking a small angle step from the previous epipolar line so that there is no pixel compression within the region of the epipolar line intersecting the image. The worst case pixel will always be situated at the furthest distance from the epipole, i.e. the image edge opposite to the epipole. Figure 3 shows how the angle step θ can be calculated very simply as $\theta = \arctan\left(\frac{1}{d}\right)$ where d can be found by intersecting the epipolar line with the image.




Figure 3: Determining the minimum distance between consecutive epipolar lines so as to avoid pixel loss

Note that when this table is built up, each epipolar line from both images can be unrectified and intersected with the image. From this, the maximum and minimum distance from the epipole can be found. Subsequently, the maximum and minimum distance anywhere in both images can be found, and both output images completely bounded.

3.4 Rectifying and Unrectifying points

Unfortunately, avoiding pixel compression means it is necessary to rectify and unrectify points using look-up tables. Each y coordinate of the rectified image can be associated with an angle to enable unrectification, and vice-versa to enable rectification. However, this makes it difficult to rectify any point and to unrectify points with a subpixel y coordinate. In order to perform such operations, it is necessary to interpolate the look-up tables. For this reason, it is best to represent unrectified epipolar lines by their polar angles and interpolate using the angles.

3.5 Resampling the image

The image can be resampled very efficiently. For each line of the first image, the maximum and minimum distance of points from the epipoles can be unrectified, to give an epipolar line segment. Pixel sized steps from one extent of this segment to the other can be taken, and the output row of the rectified image can be built up. The same scheme can be used for the other image, but with the different maximum and minimum distances for the relevant image. Then, as the epipolar line in image one is being worked along in pixel sized steps, it is transferred into image two using the compatible homography.

3.6 Infinite epipoles

Images with infinite epipoles will in fact not work with the above technique because all distances will become infinite, and all angles will be the same. Fortunately, they can easily be detected as an image with an effective angle range of 0 radians. Note that an infinite epipole in the second image is irrelevant because the use of a compatible homography means that points from the second image are transferred into the first image. Consequently, only an infinite epipole in the first image need be detected. For this case, the rectification can simply apply the compatible homography and rotate both images so that the epipole lies on the x axis. Although this is an exception case, it is extremely easy to detect and handle.




Figure 4: Forward movement image pair prior to rectification (left) and after (right)



Figure 5: Near parallel image pair prior to rectification (left) and after (right)

4 Examples

Figures 4 to 5 give a qualitative feel for the effects of rectification on some example images. For all the scenes, the fundamental matrix was first estimated techniques described in [7].

Figure 4 illustrates the rectification of an image pair produced with a camera undergoing a mainly forward movement, the sort that would result in an unbounded image if planar rectification were used. Notice how the epipole has been mapped to a line and how this has resulted in disparities which are purely horizontal making this a usable rectification. The other stereo pair in figure 5 was taken using a more conventional near parallel camera movement and can be seen to have produced very little image distortion.

5 Conclusion

This work has presented a simple and fast algorithm for rectification of any stereo image pair without the need for any calibration. Compared to previous techniques it has reduced inter image distortion caused by the nonlinear rectification as well as providing a simpler and more complete algorithm. It has also provided a means of minimising image distortion due to perspective effects if some pre-determined matches are provided (such as might have been used to calculate the epipolar geometry). Overall, the result is a more general technique that produces images which are easier to match.

References

- [1] A. W. Fitzgibbon and A. Zisserman. Automatic camera recovery for closed or open image sequences. In *Proc. ECCV*, pages 311–326, June 1998.
- [2] R. Hartley. Cheirality invariants. In *Proc. DARPA Image Understanding Workshop*, pages 743–753, 1993.
- [3] R. I. Hartley. Theory and practice of projective rectification. Technical Report 2538, INRIA, April 1995.
- [4] R. I. Hartley and R. Gupta. Computing matched epipolar projections. In *Proc. CVPR*, pages 549–555, New York, 1993.
- [5] R. Koch, M. Pollefeys, and L.V. Gool. Multi viewpoint stereo from uncalibrated video sequences. In *Proc. ECCV*, pages 55–71. Springer-Verlag, 1998.
- [6] S. Laveau and O. Faugeras. Oriented projective geometry for computer vision. In *Proc. ECCV*, volume 1064 of *LNCS*, pages 147–156, 1996.
- [7] Q. T. Luong and O. D. Faugeras. The fundamental matrix: Theory, algorithms, and stability analysis. In *IJCV*, volume 17, pages 43–75, January 1996.
- [8] M. Pollefeys, R. Koch, and L. Van. Gool. A simple and efficient rectification method for general motion. In *Proc. ICCV*, pages 496–501, Corfu, 1999.
- [9] P. J. Rousseeuw and A. M. Leory. *Robust Regression and Outlier Detection*. Jogn Wiley & Sons, New York, 1987.
- [10] R. Sébastien, J. Meunier, and J. C. Ingemar. Cylindrical rectification to minimize epipolar distortion. In *Proc. ICVPR*, pages 393–399, 1997.
- [11] C. C. Slama, editor. *Manual of Photogrammetry*. American Society of Photogrammetry, Falls Church, Va, fourth edition, 1980.