



Quasi-Euclidean Epipolar Rectification

Pascal Monasse

[article](#) [demo](#) [archive](#)

published • 2011-09-13

→ BibTeX

reference • Pascal Monasse, *Quasi-Euclidean Epipolar Rectification*, Image Processing On Line, 1 (2011). http://dx.doi.org/10.5201/ipol.2011.m_qer

Communicated by Luis Alvarez

Demo edited by Agustin Salgado

• [Pascal Monasse](#) monasse@imagine.enpc.fr

Content

- Overview
- References
- The Fusiello-Irsara method
- Online Demo
- Algorithm
- Implementation
- Examples
- Source code
- Acknowledgments

Overview

The binocular stereo pipeline performs the following task: from a pair of images of a scene captured by two cameras (or the same one) at different positions, compute the distance map to one camera. Under particular conditions (called the *rectified case*) corresponding points are at same ordinate in both images and the distance is the inverse of an affine transform (whose coefficients depend on camera parameters) of the abscissa difference, called the *disparity*. In general, the output is simply the disparity map.

If the cameras satisfy the pinhole model assumption, an adequately chosen combination of rotations and adjustment of cameras' internal parameters yield the rectified case. This amounts to applying homographies to the images. Finding and applying these homographies is called *epipolar rectification*. The input is a discrete set of corresponding points in images. Since there are more degrees of freedom than constraints, several methods exist, each trying to minimize the change applied to images according to its own measure. One method to achieve this and more discussion can be found in [this book](#).

The method of [reference 2](#), expanded in [reference 3](#), assumes both cameras are the same (thus they must have the same size) but only partial knowledge of camera's internal parameters (uncalibrated case): square pixels, unknown focal length and principal point at image center. It then simulates appropriate pure rotations of each view, which can be done when the correct focal length is estimated.

References

1. Richard Hartley and Andrew Zisserman (2003). Multiple View Geometry in Computer Vision. Cambridge University Press. ISBN 978-0-521-54051-3
2. Andrea Fusiello and Luca Irsara, [Quasi-Euclidean uncalibrated epipolar rectification](#), 19th International Conference on Pattern Recognition (ICPR), Tampa, Florida, 2008. DOI: 10.1109/ICPR.2008.4761561.
3. Andrea Fusiello and Luca Irsara, [Quasi-Euclidean epipolar rectification of uncalibrated images](#), Machine Vision and Applications, Springer Berlin/Heidelberg, 2010. DOI: 10.1007/s00138-010-0270-3.

The Fusiello-Irsara method

Background

The epipolar constraint equation is written

$$X_l^T F X_r = 0,$$

with X_l and X_r 3-vectors of homogeneous coordinates of corresponding points and F the 3x3 fundamental matrix. In the rectified case, F has the special form (up to a scale factor)

$$F = [e_1]_{\times} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{pmatrix}.$$

In that case, epipoles are at infinity in horizontal direction since $Fe_1 = F^T e_1 = 0$, and corresponding points have same ordinate since

$$(x_l \ y_l \ 1) F (x_r \ y_r \ 1)^T = 0 \Leftrightarrow (x_l \ y_l \ 1) (0 \ -1 \ y_r)^T = 0 \Leftrightarrow y_l = y_r.$$

Any decomposition of a fundamental matrix F into

$$F = H_l^T [e_1]_{\times} H_r,$$

with H_l and H_r invertible 3x3 matrices, yields left and right homographies rectifying the images since then

$$X_l^T F X_r = 0 \Leftrightarrow (H_l X_l)^T [e_1]_{\times} (H_r X_r) = 0.$$

Rotations

The rotation of matrix R of a camera around its focus produces a homography of the image of matrix

$$H = KRK^{-1},$$

with K the upper-triangular matrix built from camera parameters, which is reasonably supposed written as:

$$K = \begin{pmatrix} f & 0 & x_C \\ 0 & f & y_C \\ 0 & 0 & 1 \end{pmatrix},$$

with (x_C, y_C) the principal point, assumed to be the image center, and f the unknown focal length.

The Fusiello-Irsara method looks for rotation matrices R_l and R_r and focal length f such that

$$E(x_l, y_l, x_r, y_r) = X_l^T K^{-T} R_l^T K^T [e_1]_{\times} K R_r K^{-1} X_r = 0,$$

which can be simplified by removing innermost two instances of K since

$$K^T [e_1]_{\times} K = f^2 [e_1]_{\times}.$$

This equality also shows that in the rectified case, we can always change the internal parameters of both cameras to K_n , amounting to homography

$$K_n K^{-1},$$

and remain rectified.

Also if R_x is any rotation around x-axis, we have

$$R_x^T [e_1]_{\times} R_x = [e_1]_{\times},$$

so we can ignore the x-axis rotation in R_l for example. Finally, the unknown matrices can be parameterized by

$$R_l = R_z(\theta_{lz})R_y(\theta_{ly}), \quad R_r = R_z(\theta_{rz})R_y(\theta_{ry})R_x(\theta_{rx}), \quad K = K(f = 3^g(w+h)),$$

so that expected value of g is in the interval $[-1,1]$, the same order of magnitude as angles.

The Sampson error associated to algebraic error E is

$$E_s^2 = E^T (JJ^T)^{-1} E,$$

with J the 1×4 matrix of partial derivatives of E wrt the 4 variables:

$$J = ((FX_r)_1 \quad (FX_r)_2 \quad (F^T X_l)_1 \quad (F^T X_l)_2),$$

leading to

$$E_s(X_l, X_r)^2 = \frac{E(X_l, X_r)^2}{\| [e_3]_{\times} F^T X_l \|^2 + \| [e_3]_{\times} F X_r \|^2}.$$

The method looks for the set of parameters

$$(\theta_{ly} \quad \theta_{lz} \quad \theta_{rx} \quad \theta_{ry} \quad \theta_{rz} \quad g)$$

minimizing the sum of Sampson errors over matching pairs by Levenberg-Marquardt method.

Jacobian computation

The iterative error minimization requires computation of the Jacobian matrix. Noting any variable p , we have

$$\frac{\partial E_s}{\partial p} = \frac{X_l^T F' X_r}{D} - N \frac{\overline{F^T X_l^T F'^T X_l} + \overline{F X_r^T F' X_r}}{D^3},$$

with

$$F = (R_l K^{-1})^T [e_1]_{\times} (R_r K^{-1}), \quad F' = \frac{\partial F}{\partial p}, \quad N = X_l^T F X_r, \quad D = \sqrt{\| \overline{F^T X_l} \|^2 + \| \overline{F X_r} \|^2}$$

and the overline operator defined as

$$\overline{(a \quad b \quad c)^T} = (a \quad b)^T.$$

The partial derivatives wrt any of the 5 angles are easy to compute, as each one is involved in only one rotation matrix:

- $\frac{\partial F}{\partial \theta_{ly}} = (R_z(\theta_{lz}) R'_y(\theta_{ly}) K^{-1})^T [e_1]_{\times} R_r K^{-1}$
- $\frac{\partial F}{\partial \theta_{lz}} = (R'_z(\theta_{lz}) R_y(\theta_{ly}) K^{-1})^T [e_1]_{\times} R_r K^{-1}$
- $\frac{\partial F}{\partial \theta_{rx}} = (R_l K^{-1})^T [e_1]_{\times} R_z(\theta_{rz}) R_y(\theta_{ry}) R'_x(\theta_{rx}) K^{-1}$
- $\frac{\partial F}{\partial \theta_{ry}} = (R_l K^{-1})^T [e_1]_{\times} R_z(\theta_{rz}) R'_y(\theta_{ry}) R_x(\theta_{rx}) K^{-1}$
- $\frac{\partial F}{\partial \theta_{rz}} = (R_l K^{-1})^T [e_1]_{\times} R'_z(\theta_{rz}) R_y(\theta_{ry}) R_x(\theta_{rx}) K^{-1},$

...

with

$$R_x(\theta) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\theta & -\sin\theta \\ 0 & \sin\theta & \cos\theta \end{pmatrix}, \quad R'_x = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \cos(\theta + \frac{\pi}{2}) & -\sin(\theta + \frac{\pi}{2}) \\ 0 & \sin(\theta + \frac{\pi}{2}) & \cos(\theta + \frac{\pi}{2}) \end{pmatrix}$$

being its derivative, and similarly for R_y and R_z .

Now, the derivative wrt g is

$$\frac{\partial F}{\partial g} = (R_l K^{-1}(g))^T [e_1]_{\times} R_r K^{-1}(g) + (R_l K^{-1}(g))^T [e_1]_{\times} R_r K^{-1}(g)$$

with

$$K^{-1}(g) = \begin{pmatrix} 1/f & 0 & -w/(2f) \\ 0 & 1/f & -h/(2f) \\ 0 & 0 & 1 \end{pmatrix} \quad K^{-1}'(g) = -\log 3 \begin{pmatrix} 1/f & 0 & -w/(2f) \\ 0 & 1/f & -h/(2f) \\ 0 & 0 & 0 \end{pmatrix}$$

Fixing free parameters

We take advantage of the invariance properties of the rectified configuration, wrt to rotation around baseline (x -axis) and change of internal parameters, to fix the center point of each image through application of rectifying homography.

First, we apply an x -rotation so as to keep ordinate of central pixel fixed, meaning replacing the homographies by:

- $H'_l = K R_x(\alpha) R_l K^{-1}$
- $H'_r = K R_x(\alpha) R_r K^{-1}$

with the angle

$$\alpha = \text{atan} \frac{(H_l C)_y - y_c}{f} = \text{atan} \frac{(H_r C)_y - y_c}{f}.$$

Finally, we adjust independently the abscissa of left and right camera to keep the center point at center. Noting

$$K_n(x) = \begin{pmatrix} f & 0 & x \\ 0 & f & h/2 \\ 0 & 0 & 1 \end{pmatrix}$$

and knowing the resulting abscissa x' of center point through H'_l , based on matrix $K=K_n(w/2)$, we need to translate by $w/2-x'$, yielding the final homographies:

- $H''_l = K_n(w - (H'_l C)_x) R_x(\alpha) R_l K^{-1}$
- $H''_r = K_n(w - (H'_r C)_x) R_x(\alpha) R_r K^{-1}$.

Online Demo

The online demo of the algorithm has additional steps to make it more widely testable, as the demonstrated algorithm itself takes as input only corresponding points between images and the size of images and outputs homographies:

1. Extract key points with SIFT algorithm in each image;

2. Match key points based on descriptors;
3. Remove outlier correspondences by a *contrario* RANSAC algorithm (ORSA) looking for a consensus on fundamental matrix;
4. Apply Fusiello-Irsara epipolar rectification algorithm described above;
5. Apply homographies to both images.

Remember that input images must have the same size. A common failure case is when one epipole is in the image. Another problem is when geometrical distortion is noticeable in images (this happens especially with wide-angle lenses). In that case, it is preferable to correct this distortion (use for example [this IPOL algorithm](#)).

The image rectangle can be mapped to a large quadrilateral if the epipole is outside but close to image boundary. To avoid potentially very long transfer time over the network, we do not show the rectified image containing the full quadrilateral. Instead, we show only the portion of same size as original image and with same center point, since this point has been fixed by adjustment of parameters of homography.

Here is an example of algorithm output:

```

01 sift:: 1st image: 473 keypoints
02 sift:: 2nd image: 496 keypoints
03 sift:: matches: 146
04 seed: 1286188619
05 Remove 19/146 duplicate matches
06 Optimized stochastic mode (ORSA).
07   nfa=-130.795 size=111 (niter=1)
08   nfa=-147.135 size=102 (niter=6)
09   nfa=-156.207 size=106 (niter=10)
10   nfa=-190.564 size=103 (niter=16)
11   nfa=-207.249 size=107 (niter=32)
12   nfa=-207.436 size=105 (niter=424)
13   nfa=-211.225 size=106 (niter=475)
14 best matching found: 106 points log(nfa)=-211.225 (500 iterations)
15 F= [ -3.42803e-09 4.35331e-07 -0.000178794; -5.00755e-07 -4.76413e-08 0.00
123017; 0.000225377 -0.00117995 -0.0373931 ]
16 Geometric error threshold: 0.869643
17 LM iterations: 8 f=760.855
18 Final rectification error: 0.206727 pix
19 Disparity: -90 -26

```

Lines 01 and 02 show the number of SIFT points extracted in each image. Line 03 shows the number of SIFT correspondences. Lines 04 to 16 are output from the ORSA algorithm. Line 04 shows the seed used for random number generation and is useful only for debugging purpose. Line 05 shows the number of remaining correspondences when duplicates are removed: some SIFT points may be found to have several main directions in their descriptor, which creates duplicates and can fool the ORSA match independence assumption. Lines 07 to 13 show each time a more meaningful group of correspondences is found. Line 14 sums up the best consensus. Line 15 shows the fundamental matrix. Line 16 displays the geometric threshold that was selected by ORSA to discriminate inliers and outliers. Here, every SIFT point must have its corresponding SIFT point at less than 0.87 pixel from the epipolar line. Lines 17 and 19 show the result of the rectification itself. The number of Levenberg-Marquardt iterations is here 8 and the found focal length is displayed. Line 18 shows the resulting mean Sampson error. Finally, line 19 gives the minimum and maximum disparity of SIFT points in the rectified images.

Algorithm

The algorithm is a Levenberg-Marquardt iterative minimization of the Sampson error, starting with all 6 unknowns at 0. It stops when one of the following conditions is satisfied:

1. the RMSE is below 0.1 pixel.
2. the relative error change from one iteration to the next is below 1E-3.
3. 300 iterations have been performed.

The first case is considered a success and the last one a failure. The second case may be either: the global minimum is reached but the data inaccuracy prevents the error from dropping below 0.1 (success) or a local minimum is reached (failure).

(success) or a local minimum is reached (failure).

Implementation

Eliminating null equations

During Levenberg-Marquardt minimization (see [reference 1](#)), one has to solve a system of 6 equations

$$HX = J^T E, \text{ with } H = J^T J + \lambda \text{diag}(J^t J),$$

with J the Jacobian matrix of the error term, of size $N \times 6$ with N the number of correspondences. The solution X is the update to apply to the current set of parameters. One of these 6 equations may happen to be the trivial one

$$(0 \ 0 \ 0 \ 0 \ 0 \ 0) X = 0.$$

This is the case when one column of J is null. This happens in particular at the first iteration for the g -derivative, when all 5 angles are 0, since

$$F' = K^{-T} [e_1]_{\times} K^{-1} + K^{-T} [e_1]_{\times} K^{-1} = k [e_1]_{\times} \Rightarrow \frac{\partial E_s}{\partial g}(0, 0, 0, 0, 0, g) = \frac{kN}{D} - N \frac{kI}{L}$$

Notice that the diagonal elements of $J^T J$ are the square norms of column vectors of J . So, to avoid division by zero or by a negligible quantity, we find the maximum diagonal entry M of H and put in `m_nullCols` the indices of diagonal elements lower than kM , with $k=1E-9$. We remove these lines and columns from the linear system to solve in the method

```
void MinLM::compress(matrix<flnum>& JtJ, vector<flnum>& B).
```

The solution of the smaller system is then completed with 0 at removed indices in the method

```
void MinLM::uncompress(vector<flnum>& B),
```

since these elements correspond to variables with null derivative, which must not be changed.

Homography application with anti-aliasing filter

Applying directly the computed homographies to images can result in aliasing artifacts. Although less severe than in the case of a zoom out for example, because we simulate rotations in our case, it is nonetheless important to attenuate the aliasing. This happens when there is a compression in one direction: at each point of the original image, the homography can be approximated by its 2×2 Jacobian matrix. When one of the singular values is below 1, compression occurs and anti-aliasing filter is recommended. For this, we have to know the maximum unidirectional compression factor over the image. Under normal conditions (image of line at infinity of original image does not meet the frame of the result image), the maximum compression (lowest singular value of Jacobian) is reached at one of the four corners of the original image. The procedure is the following:

1. At each corner of the original image frame, compute 2×2 Jacobian matrix of homography H and its minimum singular value. Note z the minimum of these over all 4 corners.
2. If $z < 1$, compose to the left H with a zoom in of factor $s=1/z$; apply the resulting homography to the image, yielding an image of size `sw_sh`; convolve with a Gaussian filter of standard deviation $0.8(s*s-1)^{1/2}$; downsample the smoothed image by a factor s .
3. If $z \geq 1$, regular transform is applied without any anti-aliasing filter.

The discrete Gaussian kernel is cut at 4 times its standard deviation. Interpolation is done with splines of order 5.

Differences with authors' implementation

Fusiello and Irsara's Matlab implementation of their algorithm, available [here](#) does not correspond exactly to the text in their paper. we underline here the differences, both with the paper version and their implementation.

1. They write the epipolar constraint $X_r F X_l = 0$, whereas we write it $X_l F X_r = 0$. In other words, the fundamental matrices are transpose of each other. We find our convention more logical as subscript l (left) corresponds to the term on the left of the fundamental matrix.
2. The Matlab code contains fully developed formulas for the Jacobian matrix, probably gotten from some symbolic computation program. We use matrix calculus to compute it. Moreover, the Matlab code has the *wrong* Jacobian: it corresponds to the Jacobian matrix of the square errors. This is harmless in their implementation as they specify to use finite differences for the computation, not the closed formulas.
3. They decompose a rotation matrix with Euler angles in this order: $R_x R_z R_y$, we prefer the more natural $R_z R_y R_x$.
4. They cancel the x-rotation in the right term of F , we do it in the left term.
5. The paper claims using Levenberg-Marquardt (LM) minimization with *lsqnonlin* Matlab function. The way this function is called in their code uses the default "trust region reflective" algorithm. On some examples, we found this minimization to be less susceptible to get caught in a local minimum of the energy. However, we implemented it with the better known LM algorithm, as claimed in the paper. Modifying their Matlab code to use LM in *lsqnonlin*, we get results very similar to ours.
6. The Matlab code tries the minimization up to 20 times with random initialization of the focal length while the results are not satisfactory. We try only once with $g=0$.
7. The Matlab code has the option of keeping the center image point fixed. This is done by translating the principal point. We prefer using only an x-translation (to fix abscissa) of the principal point and an x-rotation (to fix ordinate).

Examples

We show some examples of successful rectification by the above algorithm. In most cases, 10 or fewer iterations of the Levenberg-Marquardt algorithm are necessary. To check visually the result, you can make your browser window large enough to get the rectified images side by side, then change its height so as to cut the images at some level; check that discernible features (blobs, corners) move along that line.

In each case, the first two images are the original ones and the last two the rectified ones.

Notice that different runs of the full pipeline with same data do not necessarily yield the exact same results as there is a random component in ORSA.

Beijing lion





107 correspondences are found and the final RMSE is 0.25 pixels after 6 iterations.

Cachan statue





964 correspondences are found and the final RMSE is 0.13 pixels after 5 iterations.

Cluny church





110 correspondences are found and the final RMSE is 0.27 pixels after 35 iterations.

Carcassonne





503 correspondences are found and the final RMSE is 0.19 pixels after 55 iterations.

Birthday cake





146 correspondences are found and the final RMSE is 0.69 pixels after 11 iterations. The final error is higher than in previous examples, due to some geometric distortion of the camera.

Source code

C++ source code of the pipeline (containing SIFT point detection and matching, ORSA, rectification) is provided in a compressed tar archive: [MissStereo.tar.gz](#). Tested under linux Ubuntu 10.04 with GNU g++ 4.4.3.

Acknowledgments

This work was supported by Centre National d'Etudes Spatiales (CNES) and by Agence Nationale de la Recherche ANR-09-CORD-003 (Callisto project). Different parts of the support code used in the demo were written by Toni Buades, Nicolas Limare, Lionel Moisan and Zhongwei Tang.

[feeds & twitter](#) • [sitemap](#) • [contact](#) • [privacy policy](#) • ISSN: 2105-1232 • DOI: 10.5201/ipol
supported by CMLA, ENS Cachan • DMI, Universitat de les Illes Balears • Fing, Universidad de la República • Gandi.net
© 2009-2014, IPOL Image Processing On Line & the authors

