Rectification with unconstrained stereo geometry *

Andrea Fusiello
Dipartimento di Matematica e Informatica
Università di Udine, Italia
fusiello@dimi.uniud.it

Emanuele Trucco
Department of Computing and Electrical Engineering
Heriot-Watt University, UK
mtc@cee.hw.ac.uk

Alessandro Verri
Dipartimento di Fisica
Università di Genova, Italia
verri@infm.ge.it

Abstract

We present a linear rectification algorithm for general, unconstrained stereo rigs. The algorithm requires the two projection matrices of the original cameras, and enforces explicitly all constraints necessary and sufficient to achieve a unique pair of rectified projection matrices. We report tests proving the correct behaviour of our method, as well as the negligible decrease of the accuracy of 3-D reconstruction if performed from the rectified images directly. To maximise reproducibility and usefulness, we give a working, 22-line Matlab code, and a URL where code, example data and a user guide can be found. Stereo reconstruction systems are very popular in vision research and applications, hence the usefulness of a general and easily accessible rectification algorithm.

1 Introduction and motivations

Given a pair of stereo images, *rectification* determines a transformation of each image plane such that pairs of conjugate epipolar lines become collinear and parallel to one of the image axes. The important advantage of rectification is that computing correspondences, a 2-D search problem in general, is reduced to a 1-D search problem, typically along the horizontal raster lines of the rectified images [6, 3, 7, 8, 13]. The rectified images can be thought of as acquired by a new stereo rig, obtained by rotating the original cameras. This is indeed the basis for computing the *rectifying transformation*, as well as the perspective projection matrices rectifying the images (*rectifying projection matrices*). This paper presents

*revised February 11, 1998

a novel algorithm rectifying a calibrated stereo rig of unconstrained geometry and mounting general cameras. The only input required is the pair of projection matrices of the two cameras (but not their individual intrinsic or extrinsic parameters), which can be estimated using one of the many existing calibration methods [5, 14, 2, 10, 12]. The output is the pair of rectified projection matrices, which can be used to compute the rectified images. Reconstruction can be performed directly from the rectified images and projection matrices. Given the importance of rectification as a module for stereo systems, and the shortage of easily reproducible, easily accessible and clearly stated algorithms, we have made a "rectification kit" (code, examples, data and user's manual) available on line. Our work improves and extends [1], in which a matrix satisfying the constraints sufficient to guarantee rectification is hand-crafted, not derived, and the constraints necessary to guarantee a unique solution are left unspecified. Instead, we enforce explicitly *all* the constraints necessary and sufficient to derive a unique rectification matrix, and obtain the latter as the solution of the resulting system of simultaneous equations. Some authors report rectification under restrictive assumptions; for instance, [9] assumes a very restrictive geometry (parallel vertical axes of the camera reference frames). Recently, [11] have introduced an algorithm which performs rectification with general stereo geometry given a weakly calibrated stereo rig (fundamental matrix and three conjugate pairs). We do require strong calibration, which however can be achieved in many practical situations and by several algorithms.

This paper is organised as follows. Section 2 introduces our notations and summarises some necessary mathematics of perspective projections. Section 3 expresses the rectifying image transformation in terms of projections matrices. Section 4 and 5 derive the algorithm for computing the rectifying projection matrices. Section 6 gives the compact (22 lines), working MATLAB code for our algorithm, and indicates where to find our rectification kit on line. Section 7 reports tests on synthetic and real data. Section 8 is a brief discussion of our work.

2 Notation and basics

Following [4], we consider a stereo pair composed of two pinhole cameras, each modelled by its *optical center* $\mathbf{c}$ and its *retinal plane* (or *image plane*) $\mathcal{R}$. In each camera, a point $\mathbf{w}$ in 3-D space is projected into an image point $\mathbf{m}$, which is the intersection of the line $\mathbf{wm}$ with $\mathcal{R}$. The transformation from $\mathbf{w}$ to $\mathbf{m}$ is modelled by the linear transformation $\tilde{\mathbf{P}}$ in projective (or homogeneous) coordinate:

$$\tilde{\mathbf{m}} = \tilde{\mathbf{P}}\tilde{\mathbf{w}}, \tag{1}$$

where

$$\tilde{\mathbf{m}} = \begin{pmatrix} U \\ V \\ S \end{pmatrix} \quad \tilde{\mathbf{w}} = \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix} \quad \mathbf{m} = \begin{pmatrix} U/S \\ V/S \end{pmatrix} \quad (\text{if } S \neq 0). \tag{2}$$

The points $\mathbf{w}$ for which $S = 0$ define the *focal plane* and are projected to infinity. The projection matrix $\tilde{\mathbf{P}}$ can be decomposed into the product $\tilde{\mathbf{P}} = \mathbf{T}_i\mathbf{T}_e$. $\mathbf{T}_e$ maps from world to camera coordinates and depends on the extrinsic parameters of the

stereo rig only; $\mathbf{T}_i$, which maps from camera to pixel coordinates and depends on the intrinsic parameters only, has the following form:

$$\mathbf{T}_i = \begin{pmatrix} -fk_u & 0 & u_0 \\ 0 & -fk_v & v_0 \\ 0 & 0 & 1 \end{pmatrix} \tag{3}$$

where f is the focal length in millimeters, k_u, k_v are the scale factors along the u and v axes respectively (the number of pixels per millimiter), and $\alpha_u = -fk_u$, and $\alpha_v = -fk_v$ are the focal lengths in horizontal and vertical pixels, respectively. If we write the projection matrix as

$$\tilde{\mathbf{P}} = \left(\begin{array}{c|c} \mathbf{q}_1^\top & q_{14} \\ \mathbf{q}_2^\top & q_{24} \\ \mathbf{q}_3^\top & q_{34} \end{array}\right) = (\mathbf{P}|\tilde{\mathbf{p}}), \tag{4}$$

we see that the plane $\mathbf{q}_3^\top \mathbf{w} + q_{34} = 0$ ($S = 0$) is the focal plane, and the two planes $\mathbf{q}_1^\top \mathbf{w} + a_{14} = 0$ and $\mathbf{q}_2^\top \mathbf{w} + a_{24} = 0$ intersect the retinal plane in the vertical ($U = 0$) and horizontal ($V = 0$) axis of the retinal coordinates, respectively.

The *optical center*, $\mathbf{c}$, is the intersection of the three planes introduced in the previous paragraph; therefore $\tilde{\mathbf{P}} \begin{pmatrix} \mathbf{c} & 1 \end{pmatrix}^\top = \mathbf{0}$, and $\mathbf{c} = -\mathbf{P}^{-1}\tilde{\mathbf{p}}$. The *optical ray* associated to an image point $\mathbf{m}$ is the line $\mathbf{cm}$, i.e. the set of points $\{\mathbf{w} : \tilde{\mathbf{m}} = \tilde{\mathbf{P}}\tilde{\mathbf{w}}\}$. The equation of this ray can be written in parametric form as $\mathbf{w} = \mathbf{c} + \lambda \mathbf{P}^{-1}\tilde{\mathbf{m}}$.

3 The rectification transformation

We now show that, if $\tilde{\mathbf{P}}_n$ is the projection matrix which rectifies one of the two views, the linear transformation (in projective coordinates) that maps the retinal plane of $\tilde{\mathbf{P}}_o = (\mathbf{P}_o|\tilde{\mathbf{p}}_o)$ onto the rectified retinal plane is given by the matrix $\mathbf{P}_n\mathbf{P}_o^{-1}$. For any 3-D point $\mathbf{w}$ we can write

$$\begin{cases} \tilde{\mathbf{m}}_o = \tilde{\mathbf{P}}_o \tilde{\mathbf{w}} \\ \tilde{\mathbf{m}}_n = \tilde{\mathbf{P}}_n \tilde{\mathbf{w}}. \end{cases} \tag{5}$$

We know that the equation of the optical ray associated to $\mathbf{m}_0$ is

$$\mathbf{w} = \mathbf{c}_o + \lambda \mathbf{P}_o^{-1}\tilde{\mathbf{m}}_o; \tag{6}$$

hence

$$\begin{gathered} \tilde{\mathbf{m}}_n = \tilde{\mathbf{P}}_{n1} \begin{pmatrix} \mathbf{c}_o + \lambda \mathbf{P_o}^{-1}\tilde{\mathbf{m}}_o \\ 1 \end{pmatrix} = \tilde{\mathbf{P}}_{n1} \begin{pmatrix} \mathbf{c}_o \\ 1 \end{pmatrix} + \tilde{\mathbf{P}}_{n1} \begin{pmatrix} \lambda \mathbf{P_o}^{-1}\tilde{\mathbf{m}}_o \\ 0 \end{pmatrix} = \\ \tilde{\mathbf{P}}_{n1} \begin{pmatrix} \mathbf{c}_n \\ 1 \end{pmatrix} + \mathbf{P}_n\mathbf{P}_o^{-1}\tilde{\mathbf{m}}_o. \end{gathered} \tag{7}$$

Assuming that rectification does not alter the optical center ($\mathbf{c}_n = \mathbf{c}_o$), we obtain

$$\tilde{\mathbf{m}}_n = \mathbf{P}_n\mathbf{P}_o^{-1}\tilde{\mathbf{m}}_o. \tag{8}$$

This is a clearer and more compact result than the one reported in [1], in which $\mathbf{P}_o^{-1}$ is not written as the inverse of a projection matrix.

4 Constraining the rectifying projection matrices

To compute the rectifying projection matrices $\tilde{\mathbf{P}}_{n1}$ and $\tilde{\mathbf{P}}_{n2}$, we set up a linear, homogeneous system of equations formed by the constraints sufficient to guarantee rectification [1], and incorporate explicitly quadratic constraints on the entries of $\tilde{\mathbf{P}}_{n1}$ and $\tilde{\mathbf{P}}_{n2}$ to ensure a nontrivial solution. The constrained system and its equations are detailed in this section. In the following, we shall write $\tilde{\mathbf{P}}_{n1}$, $\tilde{\mathbf{P}}_{n2}$ as follows:

$$\tilde{\mathbf{P}}_{n1} = \left(\begin{array}{c|c} \mathbf{a}_1^\top & a_{14} \\ \mathbf{a}_2^\top & a_{24} \\ \mathbf{a}_3^\top & a_{34} \end{array}\right) \qquad \tilde{\mathbf{P}}_{n2} = \left(\begin{array}{c|c} \mathbf{b}_1^\top & b_{14} \\ \mathbf{b}_2^\top & b_{24} \\ \mathbf{b}_3^\top & b_{34} \end{array}\right). \tag{9}$$

Scale factor. Projection matrices are defined up to a scale factor. The common choice to block the latter is $a_{34} = 1$ and $b_{34} = 1$ [1], unfortunately, brings about two problems: first, the intrinsic parameters become dependent on the choice of the world coordinate system [5]; second, the resulting projection matrices do not satisfy the conditions guaranteeing that meaningful calibration parameters can be extracted from their entries of the matrices [4], that is (for example for $\tilde{\mathbf{P}}_{n1}$),

$$||\mathbf{a}_3|| = 1 \qquad (\mathbf{a}_1 \wedge \mathbf{a}_3)^\top (\mathbf{a}_2 \wedge \mathbf{a}_3) = 0. \tag{10}$$

To obviate the problems mentioned, we enforce

$$||\mathbf{a}_3|| = 1 \qquad ||\mathbf{b}_3|| = 1. \tag{11}$$

The second equation in (10) and its equivalent for $\tilde{\mathbf{P}}_{n2}$ are actually implied by the system defining our algorithm (proof omitted for reasons of space).

Position of the optical centers. The optical centers of the rectified projections, $\mathbf{c}_1$ and $\mathbf{c}_2$, must be the same as those of the original projections:

$$\tilde{\mathbf{P}}_{n1} \begin{pmatrix} \mathbf{c}_1 \\ 1 \end{pmatrix} = \mathbf{0} \qquad \tilde{\mathbf{P}}_{n2} \begin{pmatrix} \mathbf{c}_2 \\ 1 \end{pmatrix} = \mathbf{0}. \tag{12}$$

Eq. (12) gives six linear constraints:

$$\left\{\begin{array}{l} \mathbf{a}_1^\top \mathbf{c}_1 + a_{14} = 0 \\ \mathbf{a}_2^\top \mathbf{c}_1 + a_{24} = 0 \\ \mathbf{a}_3^\top \mathbf{c}_1 + a_{34} = 0 \\ \mathbf{b}_1^\top \mathbf{c}_2 + b_{14} = 0 \\ \mathbf{b}_2^\top \mathbf{c}_2 + b_{24} = 0 \\ \mathbf{b}_3^\top \mathbf{c}_2 + b_{34} = 0. \end{array}\right. \tag{13}$$

Common focal plane. The two rectified projections must share the same focal plane, i.e.

$$\mathbf{a}_3 = \mathbf{b}_3 \qquad a_{34} = b_{34}. \tag{14}$$

Alignment of conjugate epipolar lines. The vertical coordinate of the projection of a 3-D point onto the rectified retinal plane must be the same in both image, i.e:

$$\frac{\mathbf{a}_2^\top \mathbf{w} + a_{24}}{\mathbf{a}_3^\top \mathbf{w} + a_{34}} = \frac{\mathbf{b}_2^\top \mathbf{w} + b_{24}}{\mathbf{b}_3^\top \mathbf{w} + b_{34}}. \tag{15}$$

Using Eq. (14) we obtain the constraints

$$\mathbf{a}_2 = \mathbf{b}_2 \qquad a_{24} = b_{24}. \tag{16}$$

Notice that the equations written to this point are sufficient to guarantee rectification (indeed some authors stop here, e.g. [1]), *but not a unique solution*: the orientation of the rectified retinal plane and the intrinsic parameters are still free. Our algorithm constrains these quantities explicitly, as follows.

Orientation of the rectified retinal plane. We choose the rectified focal planes to be parallel to the intersection of the two original focal planes, i.e.

$$\mathbf{a}_3^\top (\mathbf{f}_1 \wedge \mathbf{f}_2) = 0, \tag{17}$$

where $\mathbf{f}_1$ and $\mathbf{f}_2$ are the third rows of $\mathbf{P}_{o_1}$ and $\mathbf{P}_{o_2}$ respectively. Notice that the dual equation $\mathbf{b}_3^\top (\mathbf{f}_1 \wedge \mathbf{f}_2) = 0$ is redundant thanks to Eq. (14).

Orthogonality of the rectified reference frames. The intersections of the retinal plane with the planes $\mathbf{a}_1^\top \mathbf{w} + a_{14} = 0$ and $\mathbf{a}_2^\top \mathbf{w} + a_{24} = 0$ correspond to the v and u axes, respectively, of the retinal reference frame. As this reference frame to be orthogonal, the planes must be perpendicular, hence, using Eq. (16),

$$\mathbf{a}_1^\top \mathbf{a}_2 = 0 \qquad \mathbf{b}_1^\top \mathbf{a}_2 = 0. \tag{18}$$

Principal points. Given a full-rank 3×4 matrix satisfying constraints (10), the principal point (u_0, v_0) is given by [4]:

$$u_0 = \mathbf{a}_1^\top \mathbf{a}_3 \qquad v_0 = \mathbf{a}_2^\top \mathbf{a}_3. \tag{19}$$

We set the two principal points to $(0, 0)$ and use Eqs. (14) and (16) to obtain the constraints

$$\begin{cases} \mathbf{a}_1^\top \mathbf{a}_3 = 0 \\ \mathbf{a}_2^\top \mathbf{a}_3 = 0 \\ \mathbf{b}_1^\top \mathbf{a}_3 = 0. \end{cases} \tag{20}$$

Focal lengths in pixels. The horizontal and vertical focal lengths in pixels, respectively, are given by

$$\alpha_u = ||\mathbf{a}_1 \wedge \mathbf{a}_3|| \qquad \alpha_v = ||\mathbf{a}_2 \wedge \mathbf{a}_3||. \tag{21}$$

By setting the values of α_u and α_v, for example, to the values extracted from $\tilde{\mathbf{P}}_{o1}$, we obtain the constraints

$$\begin{cases} ||\mathbf{a}_1 \wedge \mathbf{a}_3||^2 = \alpha_u^2 \\ ||\mathbf{a}_2 \wedge \mathbf{a}_3||^2 = \alpha_v^2 \\ ||\mathbf{b}_1 \wedge \mathbf{a}_3||^2 = \alpha_u^2, \end{cases} \tag{22}$$

which, by virtue of the equivalence $||\mathbf{x} \wedge \mathbf{y}||^2 = ||\mathbf{x}||^2||\mathbf{y}||^2 - (\mathbf{x}^T\mathbf{y})^2$ and Eq. (20), can be rewritten as

$$\begin{cases} ||\mathbf{a}_1||^2||\mathbf{a}_3||^2 = \alpha_u^2 \\ ||\mathbf{a}_2||^2||\mathbf{a}_3||^2 = \alpha_v^2 \\ ||\mathbf{b}_1||^2||\mathbf{a}_3||^2 = \alpha_u^2. \end{cases} \tag{23}$$

5 Solving for the rectifying projection matrices

We organise the constraints introduced in the previous section in the following four systems:

$$\begin{cases} \mathbf{a}_3^\top \mathbf{c}_1 + a_{34} = 0 \\ \mathbf{a}_3^\top \mathbf{c}_2 + a_{34} = 0 \\ \mathbf{a}_3^\top (\mathbf{f}_1 \wedge \mathbf{f}_2) = 0 \\ ||\mathbf{a}_3|| = 1 \end{cases} \quad \begin{cases} \mathbf{a}_2^\top \mathbf{c}_1 + a_{24} = 0 \\ \mathbf{a}_2^\top \mathbf{c}_2 + a_{24} = 0 \\ \mathbf{a}_2^\top \mathbf{a}_3 = 0 \\ ||\mathbf{a}_2|| = \alpha_v \end{cases} \quad \begin{cases} \mathbf{a}_1^\top \mathbf{c}_1 + a_{14} = 0 \\ \mathbf{a}_1^\top \mathbf{a}_2 = 0 \\ \mathbf{a}_1^\top \mathbf{a}_3 = 0 \\ ||\mathbf{a}_1|| = \alpha_u \end{cases} \quad \begin{cases} \mathbf{b}_1^\top \mathbf{c}_2 = -b_{14} \\ \mathbf{b}_1^\top \mathbf{a}_2 = 0 \\ \mathbf{b}_1^\top \mathbf{a}_3 = 0 \\ ||\mathbf{b}_1|| = \alpha_u \end{cases} \tag{24}$$

Each system is a linear homogeneous system subject to a quadratic constraint, that is,

$$\mathbf{A}\mathbf{x} = 0 \text{ subject to } ||\mathbf{x}'|| = k, \tag{25}$$

where $\mathbf{x}'$ is a vector composed by the the first three components of $\mathbf{x}$, and k is a real number. The four systems above are solved in sequence; the solution of each can be computed using the decomposition of $\mathbf{A}$ suggested by [4], or casting the problem as a *generalised eigenvector problem.*

6 Summary of the algorithm

Given the high diffusion of stereo in research and applications, we have endeavoured to make our algorithm as easily reproducible and usable as possible. To this purpose, we give below our working MATLAB code, which is simple and compact (22 lines). The comments provided should make it understandable without knowledge of MATLAB. Moreover, a "rectification kit" including code, data sets, and instructions on how to use the algorithm can be downloaded from `ftp://taras.dimi.uniud.it/pub/sources/rectif_m.tar.gz`.

```
function [T1,T2,Pn1,Pn2] = rectify(Po1,Po2)
% RECTIFY computes the rectified projection matrices
% Pn1, Pn2, and the rectifying transformation of
% the retinal plane T1 and T2 (in homog. coords.)
% Po1 and Po2 are the original projection  matrices.

% focal length (extp(a,b) is external product of vectors a,b)
au = norm(extp(Po1(1,1:3)', Po1(3,1:3)'));
av = norm(extp(Po1(2,1:3)', Po1(3,1:3)'));

% optical centers
c1 = - inv(Po1(:,1:3))*Po1(:,4);
c2 = - inv(Po2(:,1:3))*Po2(:,4);

% unit vectors of retinal planes
fl = Po1(3,1:3)';   fr = Po2(3,1:3)';
nn = extp(fl,fr);
```

```
% select the first 3 components
L = [1 1 1 0 ]';

% lagra(A,L,k) computes the minimum of the norm of Ax,
% with the constraint "norm of x(L) equal to k"

A = [ [c1' 1]' [c2' 1]' [nn'  0]' ]'
a3 = lagra(A,L,1) ;

A = [ [c1' 1]' [c2' 1]' [a3(1:3)' 0]' ]';
a2 = lagra(A,L,av) ;

A = [ [c1' 1]' [a2(1:3)' 0]' [a3(1:3)' 0]' ]';
a1 = lagra(A,L,au) ;

A = [ [c2' 1]' [a2(1:3)' 0]' [a3(1:3)' 0]' ]';
b1 = lagra(A,L,au);

% adjustment transformation
T = [
  1 0 0
  0 1 0
  0 0 1 ];

% rectifying projection matrices
Pn1 = T * [ a1 a2 a3 ]';     Pn2 = T * [ b1 a2 a3 ]';

% rectifying image transformations
T1 = Pn1(1:3,1:3)* inv(Po1(1:3,1:3));
T2 = Pn2(1:3,1:3)* inv(Po2(1:3,1:3));
```

This MATLAB code receives in input two 3×4 projection matrices, and returns two 3×3 image transformation matrices (to be applied to image points) and two 3×4 rectifying projection matrices (to be used to perform reconstruction from rectified images).

Reconstruction. Reconstruction is performed through Eq. (1). Given two conjugate points, $\mathbf{m}_1 = (u', v')^\top$ and $\mathbf{m}_2 = (u'', v'')^\top$, and the two projection matrices $\tilde{\mathbf{P}}_{n1}$ and $\tilde{\mathbf{P}}_{n2}$, we can write the overconstrained linear system

$$\mathbf{A}\mathbf{w} = \mathbf{y}, \tag{26}$$

where

$$\mathbf{A} = \begin{pmatrix} (\mathbf{a}_1 - u'\mathbf{a}_3)^\top \\ (\mathbf{a}_2 - v'\mathbf{a}_3)^\top \\ (\mathbf{b}_1 - u''\mathbf{b}_3)^\top \\ (\mathbf{b}_2 - v''\mathbf{b}_3)^\top \end{pmatrix} \qquad \mathbf{y} = \begin{pmatrix} -a_{14} + u'a_{34} \\ -a_{24} + v'a_{34} \\ -b_{14} + u''b_{34} \\ -b_{24} + v''b_{34} \end{pmatrix}, \tag{27}$$

and $\mathbf{w}$ gives the position of the 3-D point projected to the conjugate points. Notice that different sizes or centers can be obtained, if required, by pre-multiplying both rectifying projection matrices by a suitable matrix $\mathbf{T}$, set to identity in our MATLAB code, of the same structure as (3).

7 Experimental Results

We ran tests to verify that the algorithm performed rectification correctly, and also to check that the accuracy of the 3-D reconstruction did not decrease when performed from the rectified images directly.

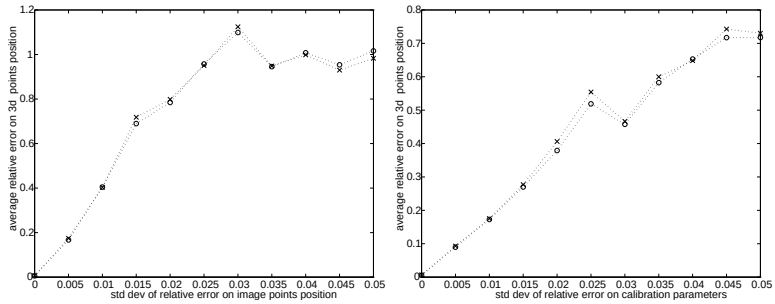

Figure 1: Reconstruction error vs noise levels in the image coordinates (left) and calibration parameters (right) for a synthetic stereo rig with camera translation $-[100\ 20\ 30]$ mm and rotation angles roll=19^o pitch=32^o and yaw=5^o. Crosses refer to reconstruction from rectified images, circles to reconstruction from unrectified images.

Correct rectification. The tests used both synthetic and real data. Each set of synthetic data consisted of a cloud of 3-D points and a pair of projection matrices; both points and matrices were chosen randomly, to cover a range of very different stereo geometries and camera parameters. For reasons of space, we do not show examples. Real-data experiments used several calibrated stereo pairs available from the INRIA-Syntim WWW site, which include the cameras' projection matrices [12]. Figure 2 shows an example. The right image of each pair shows three epipolar lines corresponding to the points marked by a cross in the left image. The output images are cropped to the size of the input images for display purposes only; the pixel coordinates of the rectified images are not constrained to lie in any special part of the image plane.

Accuracy of reconstruction. In order to evaluate the errors introduced by rectification on reconstruction, we compared the accuracy of 3-D reconstruction computed from original and rectified images. We used synthetic, noisy images of random clouds of 3-D points and random stereo geometries, as explained in the previous subsection. Imaging errors were simulated by perturbing the image coordinates, and calibration errors by perturbing the intrinsic and extrinsic parameters, both with additive, Gaussian noise. As an example, Figure 1 shows the average (over the set of points) relative error measured on 3-D point position, plotted against noise, for one of the several synthetic stereo rigs used for testing. Each point plotted is an average over 100 independent trials. The abscissa is the standard deviation of the relative error on coordinates of image point or calibration parameters. It can be seen that the accuracy does not suffer when reconstructing directly from the rectified images.

8 Discussion

Stereo matching is greatly simplified if the epipolar lines are parallel and horizontal in each image, i.e. if the images are rectified. We have developed a rectification algorithm which imposes explicitly all the constraints sufficient to obtain two

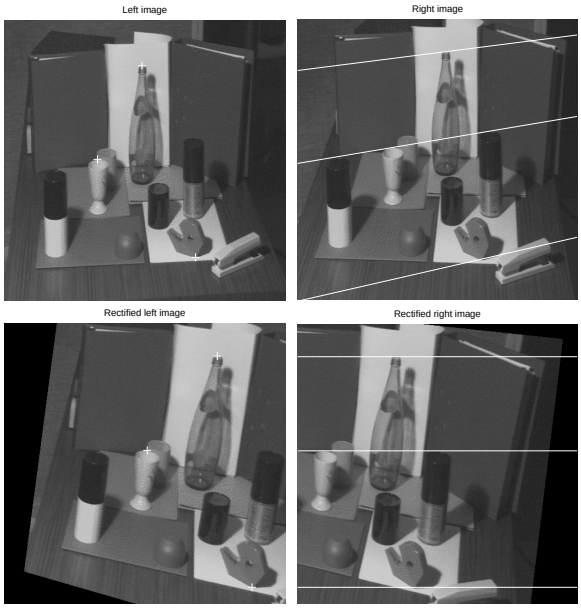

Figure 2: Original "Color" stereo pair (top) and rectified pair (bottom). The left pictures plot the epipolar lines corresponding to the point marked in the right pictures.

unique projection matrices as the solution of four constrained, homogeneous linear systems; the resulting algorithm is quite simple. The correct behaviour of the algorithm has been demonstrated with both synthetic and real images. Interestingly enough, reconstruction can be performed directly from the disparities of the rectified images, using the rectified projection matrices. Our tests show that this process does not introduce appreciable errors compared with reconstructing from the original images. We believe that a general rectification algorithm, together with the material we have made available from the URL given in Section 6, can prove a useful resource for the research and application communities alike.

Acknowledgements

This work benefited from discussions with Bruno Caprile, and was partially supported by grants from the British Council-MURST/CRUI and EPSRC (GR/L18716). The real stereo pairs are available from INRIA-Syntim under Copyright (`http://www-syntim.inria.fr/syntim/analyse/paires-eng.html`).

References

[1] N. Ayache. *Artificial Vision for Mobile Robots: Stereo Vision and Multisensory Perception*. The MIT Press, 1991.

[2] B. Caprile and V. Torre. Using vanishing points for camera calibration. *International Journal of Computer Vision*, 4:127–140, 1990.

[3] I. J. Cox, S. Hingorani, B. M. Maggs, and S. B. Rao. A maximum likelihood stereo algorithm. *Computer Vision and Image Understanding*, 63(3):542–567, May 1996.

[4] O. Faugeras. *Three-Dimensional Computer Vision: A Geometric Viewpoint.* The MIT Press, Cambridge, 1993.

[5] O. Faugeras and G. Toscani. Camera calibration for 3D computer vision. In *Proceedings of the International Workshop on Machine Vision and Machine Intelligence*, Tokyo, Japan, February 1987.

[6] D. Geiger, B. Ladendorf, and A. Yuille. Occlusions and binocular stereo. *International Journal of Computer Vision*, 14(3):211–226, April 1995.

[7] S. S. Intille and A. F. Bobick. Disparity-space images and large occlusion stereo. In Jan-Olof Eklundh, editor, *European Conference on Computer Vision*, pages 179–186, Stockholm, Sweden, May 1994. Springer-Verlag.

[8] T. Kanade and M. Okutomi. A stereo matching algorithm with an adaptive window: Theory and experiments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(9):920–932, September 1994.

[9] D. V. Papadimitriou and T. J. Dennis. Epipolar line estimation and rectification for stereo images pairs. *IEEE Transactions on Image Processing*, 3(4):672–676, April 1996.

[10] L Robert. Camera calibration without feature extraction. *Computer Vision, Graphics, and Image Processing*, 63(2):314–325, March 1995. also INRIA Technical Report 2204.

[11] L. Robert, C. Zeller, O. Faugeras, and M. Hébert. Applications of non-metric vision to some visually-guided robotics tasks. In Y. Aloimonos, editor, *Visual Navigation: From Biological Systems to Unmanned Ground Vehicles*, chapter 5, pages 89–134. Lawrence Erlbaum Associates, 1997.

[12] Jean-Philippe Tarel and André Gagalowicz. Calibration de caméra à base d'ellipses. *Traitement du Signal*, 12(2):177–187, 1995.

[13] C. Tomasi and R. Manduchi. Stereo without search. In B. Buxton and R. Cipolla, editors, *European Conference on Computer Vision*, pages 452–465, Cambridge (UK), April 1996.

[14] R. Tsai. A versatile camera calibration technique for high-accuracy 3d machine vision metrology using off-the-shelf tv cameras and lenses. *IEEE Journal of Robotics and Automation*, 3(4):323–344, August 1987.