

# Epipolar Geometry Estimation via RANSAC Benefits from the Oriented Epipolar Constraint

Ondřej Chum Tomáš Werner Jiří Matas

Center for Machine Perception, Department of Cybernetics  
Faculty of Electrical Engineering, Czech Technical University in Prague  
166 27 Prague 6, Technicka 2, Czech Republic

## Abstract

*The efficiency of epipolar geometry estimation by RANSAC is improved by exploiting the oriented epipolar constraint. Performance evaluation shows that the enhancement brings up to a two-fold speed-up. The orientation test is simple to implement, is universally applicable and takes negligible fraction of time compared with epipolar geometry computation.*

## 1. Introduction

Establishing<sup>1</sup> correspondences in two views of a rigid scene is a critical part of a number of computer vision problems. It is generally accepted that local matching cannot avoid producing incorrect correspondences (outliers) and these have to be pruned out by imposing the epipolar constraint. Due to the presence of outliers, epipolar geometry estimators must be robust. The Random Sample Consensus – RANSAC [2] and related robust hypothesize-and-verify methods [15, 1] have become the methods of choice for outlier removal in epipolar geometry estimation [4] regardless of the features matched [12, 16, 13, 9, 7, 11].

For “real” cameras (physical devices), only points in front of the camera are visible. This is modeled in the framework of the oriented projective geometry [14], where cameras form images by projecting along *half-lines* emanating from a projection center. Points in two views taken by a camera satisfy, besides the epipolar constraint, some additional constraints [5, 17]. The constraints have been used before for outlier removal after the epipolar geometry was recovered [5], but not directly in the process of epipolar geometry estimation.

In this paper we show that the use of oriented constraints within RANSAC brings significant computational savings.

<sup>1</sup>OCh and JM were supported by grants EU IST-2001-32184, GACR 102/02/1539, Kontakt ME 678, STINT Dur IG2003-2 062, and the Austrian Ministry of Education CONEX GZ 45.535. TW was supported by CAK LN00B096.

The approach has no negative side-effects, such as limited applicability or poor worst-case performance, and thus should become a part of any state-of-the-art RANSAC implementation of epipolar geometry estimation.

In RANSAC, epipolar geometry is obtained by repeating a hypothesize-and-verify loop. If the hypothesized epipolar geometry violates the oriented constraint, the verification step does not have to be carried out. Since the orientation test takes negligible time compared to both the epipolar geometry computation and the verification, a speed-up may be achieved virtually for free. The overall impact of the orientation constraint on RANSAC depends on the fraction of hypothesized models that can be rejected, without verification, solely on the basis of failing the orientation constraint test. We empirically measure this fraction in a number of real scenes, both in a wide-baseline and narrow-baseline stereo settings. The performance evaluation carried out shows that the the impact of exploiting the orientation constraint on RANSAC running time is in many circumstances significant.

The rest of the paper is structured as follows. First, the derivation of the oriented epipolar constraint is reviewed in Section 2. RANSAC with the orientation constraint is introduced in Section 3. Next, in Section 4, performance of the improved RANSAC constraint is evaluated. Two quantities are measured: the fraction of hypothesis that fail the orientation test and the time saved as a consequence of inserting the orientation test into the RANSAC loop. Section 5 concludes the paper.

## 2 Oriented Epipolar Constraint

Let<sup>2</sup> a camera with  $3 \times 4$  projection matrix  $P$  observe a rigid scene. An image point represented by a homogeneous 3-vector  $\mathbf{x}$  is a projection of a scene point represented by a homogeneous 4-vector  $\mathbf{X}$  if and only if  $\mathbf{x} \sim P\mathbf{X}$  [4].

<sup>2</sup>In this section,  $\mathbf{a} \sim \mathbf{b}$  denotes equality of two vectors up to a non-zero scale and  $\mathbf{a} \stackrel{\pm}{\sim} \mathbf{b}$  equality up to a positive scale. Vector product of two 3-vectors is  $\mathbf{a} \times \mathbf{b}$ . Symbol  $[\mathbf{a}]_{\times}$  denotes the matrix such that  $[\mathbf{a}]_{\times} \mathbf{b} = \mathbf{a} \times \mathbf{b}$ . Matrix pseudoinverse is denoted  $P^{+}$  and vector norm  $\|\mathbf{a}\|$ .

Following the classical (i.e. unoriented) projective geometry, the homogeneous quantities  $\mathbf{x}$ ,  $\mathbf{X}$ , and  $\mathbf{P}$  represent the same geometric objects if multiplied by a non-zero scale. E.g., the homogeneous vectors  $\mathbf{x} = (x, y, 1)^\top$  and  $-\mathbf{x} = (-x, -y, -1)^\top$  represent the same image point with affine coordinates  $(x, y)^\top$ .

It has been noticed that the *oriented projective geometry* [14, 6] is a more appropriate model for multiple view geometry as it can represent ray orientations. In oriented geometry, vectors  $\mathbf{x}$  and  $-\mathbf{x}$  represent two different image points, differing by whether the corresponding scene point lies in front of or behind the camera. The (oriented) relation between scene point  $\mathbf{X}$  and its image  $\mathbf{x}$  is  $\mathbf{x} \stackrel{\pm}{\sim} \mathbf{P}\mathbf{X}$ , as opposed to unoriented relation  $\mathbf{x} \sim \mathbf{P}\mathbf{X}$ .

The orientation of image points is known from the fact that all visible points lie in front of the camera [6]. Formally, image points lie on the positive side of the image line at infinity  $\mathbf{l}_\infty$ . For the usual choice  $\mathbf{l}_\infty = (0, 0, 1)^\top$ , the correctly oriented homogeneous vector representing an image point with affine coordinates  $(x, y)^\top$  is  $(x, y, 1)^\top$  or its positive multiple.

Let two cameras with projection matrices  $\mathbf{P}$  and  $\mathbf{P}'$  observe a rigid scene. It is well-known [4] that there exists a  $3 \times 3$  *fundamental matrix*  $\mathbf{F}$  of rank 2 such that any pair  $\mathbf{x} \leftrightarrow \mathbf{x}'$  of corresponding image points satisfies the *epipolar constraint*

$$\mathbf{x}'^\top \mathbf{F} \mathbf{x} = 0. \quad (1)$$

The oriented version of the epipolar constraint is [17]

$$\mathbf{e}' \times \mathbf{x}' \stackrel{\pm}{\sim} \mathbf{F} \mathbf{x}. \quad (2)$$

It is implied by the following lemma (we omit the proof).

**Lemma 1** Any  $3 \times 4$  full-rank matrices  $\mathbf{P}$  and  $\mathbf{P}'$  and a 4-vector  $\mathbf{X}$  satisfy

$$\mathbf{e}' \times (\mathbf{P}'\mathbf{X}) = \mathbf{F}(\mathbf{P}\mathbf{X}), \quad (3)$$

where  $\mathbf{F} = [\mathbf{e}']_\times \mathbf{P}'\mathbf{P}^+$ ,  $\mathbf{e}' = \mathbf{P}'\mathbf{C}$ , and  $\mathbf{C}$  is uniquely determined by equations  $\mathbf{P}\mathbf{C} = 0$  and  $\det(\mathbf{P}^\top | \mathbf{C}) = \|\mathbf{C}\|^2$ .

Note that (1) is invariant to the change of signs of  $\mathbf{x}$ ,  $\mathbf{x}'$  and  $\mathbf{F}$  whereas (2) is not. Therefore (2) implies (1) (multiply (2) by  $\mathbf{x}'^\top$  from the left), but not *vice versa*. Thus, the oriented epipolar constraint is stronger than the unoriented one.

### 3. RANSAC with Oriented Constraint

The standard seven-point algorithm [4] is used to hypothesize the fundamental matrix. The two-dimensional space of  $3 \times 3$  matrices satisfying (1) for the seven sampled correspondences is found by QR factorization rather than SVD, as suggested in [10]. Each fundamental matrix is then tested whether it satisfies the oriented constraint (2). This is the

	corrs	inliers	$\varepsilon$ [%]
Juice	447	274	61.30
Shelf	126	43	34.13
Valbonne	216	42	19.44
Great Wall	318	68	21.38
Leuven	793	379	47.79
Corridor	607	394	64.91

Table 1: The number of correspondences (‘corrs’), inliers (‘inliers’) and the fraction of inliers (‘ $\varepsilon$ ’) in the experiments.

only step in which the new algorithm differs from the standard one. The test can be performed very efficiently, requiring only 27 – 81 floating point operations (i.e. multiplications and additions). If the orientation test is passed, the support of the fundamental matrix is computed as the number of correspondences with Sampson’s distance [4] below threshold.

## 4. Experiments

The RANSAC algorithm with the oriented constraints was tested on six standard image pairs, including wide-baseline stereo (experiments Juice, Shelf, Valbonne and the Great Wall) and narrow-baseline stereo (Leuven and Corridor).

**Obtaining tentative correspondences.** By a tentative correspondence we mean a pair of points  $\mathbf{x} \leftrightarrow \mathbf{x}'$ , where  $\mathbf{x}$  is from the first image and  $\mathbf{x}'$  is from second image. The set of tentative correspondences contains both inliers and outliers.

The tentative correspondences for the wide-baseline experiments were obtained automatically by matching normalized local affine frames [8]. Only mutually best candidates were selected as tentative correspondences.

In the narrow-baseline experiments, the Harris operator [3] was used to detect interest points. Point pairs with mutually maximal normalised cross-correlation of rectangular windows around interest points were kept as tentative correspondences. The size of the used windows were 15 and 7 pixels for Leuven and Corridor, respectively. The proximity constraint, ensuring that the coordinates of corresponding points would not differ by more than 100 and 30 pixels respectively, was also used. The numbers of tentative correspondences and the numbers of inliers for each experiment are summarized in Table 1. The image pairs, with inliers and outliers superimposed, are depicted in Figure 1.

**The fraction of rejected models.** The number of hypotheses that can be rejected on the basis of the orientation constraint was measured. The results of the experiment are summarized in Table 2. A sample of seven correspondences was selected at random 500,000 times. The seven-point algorithm produces 1 to 3 fundamental matrices satisfying the unoriented epipolar constraint (1). The total number of the models is given in the ‘models’ column of Table 2. The

	models	rejected	passed [%] %
Juice	1,221,932	760,354	37.77
Shelf	1,233,770	1,094,533	11.29
Valbonne	1,256,648	1,176,042	6.41
Great Wall	1,274,018	1,181,084	7.29
Leuven	1,194,238	336,515	71.82
Corridor	1,187,380	120,916	89.82

Table 2: The number of fundamental matrices generated by RANSAC (‘models’) over 500,000 samples, the number of models rejected by the orientation constraint (‘rejected’) and the percentage of models that passed the test (‘passed’). Note that regardless of the setting, there is always approximately 2.4 fundamental matrices per sample on average.

	standard	oriented	speed-up [%]
Juice	1.4	0.9	35.08
Shelf	53.6	45.1	15.82
Valbonne	930.8	584.8	37.17
Great Wall	1109.0	599.3	45.96
Leuven	18.5	15.1	18.52
Corridor	1.2	1.2	5.77

Table 3: Time (in ms) spent in standard and oriented versions of RANSAC and the relative speed-up (right column).

number of models that are rejected by the orientation constraint (2) is shown in the ‘rejected’ column.

The fraction of rejected models varies widely. What affects the fraction of hypothesis that can be rejected solely based on the orientation constraint? This question goes far beyond the scope of this paper. From the results of the experiments we observed that more models were rejected in the wide-baseline setting than in the narrow-baseline one. We believe this is due to different distribution of outliers which is caused by limited correspondence search window in the narrow-baseline case. The fraction of rejected models is proportional to the fraction of outliers among the tentative correspondences.

**Running time.** The time saved by using the oriented epipolar constraint was measured. The standard and oriented versions of RANSAC were executed 500 times and their running-time was recorded (on a PC with K7 2200+ processor). To ensure that both methods draw the same samples, the generator of pseudo-random numbers was initialized by the same seed.



## 5. Conclusions

In this paper, RANSAC enhanced by the oriented epipolar constraint was experimentally evaluated. The application of the oriented constraint reduced the running time by 5% to 46%, compared with standard RANSAC. The effi-

ciency increase of RANSAC with the orientation constraint is achieved by reducing the number of verification steps. As a consequence, the more time-consuming the verification step is the higher relative speed-up is achieved. This applies not only to situations with large number of correspondences, but also to RANSAC-type algorithms that perform expensive verification procedures. An example of such an algorithm is MLESAC [15] which estimates the parameters of a mixture of inlier and outlier distributions. Since the evaluation of the orientation test takes virtually no time compared to the epipolar geometry computation, the epipolar geometry estimation via RANSAC should exploit the oriented epipolar constraint.

## References


- [1] O. Chum, J. Matas, and Š. Obdržálek. Enhancing ransac by generalized model optimization. In *Proc. of the ACCV*, volume 2, pages 812–817, January 2004.
- [2] M. Fischler and R. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *CACM*, 24(6):381–395, June 1981.
- [3] C. J. Harris and M. Stephens. A combined corner and edge detector. In *Proc. of Alvey Vision Conference*, pages 147–151, 1988.
- [4] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, Cambridge, UK, 2000.
- [5] S. Laveau and O. Faugeras. 3-D scene representation as a collection of images. In *ICPR '94*, pages 689–691, October 9–13 1994.
- [6] S. Laveau and O. Faugeras. Oriented projective geometry for computer vision. In *Proc. of ECCV*, volume I, pages 147–156, 1996.
- [7] J. Matas, O. Chum, M. Urban, and T. Pajdla. Robust wide baseline stereo from maximally stable extremal regions. In *Proc. of BMVC*, volume 1, pages 384–393. BMVA, 2002.
- [8] J. Matas, Š. Obdržálek, and O. Chum. Local affine frames for wide-baseline stereo. In *Proc. ICPR*, volume 4, pages 363–366. IEEE CS, Aug 2002.
- [9] K. Mikolajczyk and C. Schmid. An affine invariant interest point detector. In *Proc. ECCV*, volume 1, pages 128–142, 2002.
- [10] D. Nister. An efficient solution to the five-point relative pose problem. In *Proc. of CVPR*, volume II, pages 195–202, 2003.
- [11] D. Nister. Preemptive ransac for live structure and motion estimation. In *Proc. ICCV03*, volume I, pages 199–206, October 2003.
- [12] P. Pritchett and A. Zisserman. Wide baseline stereo matching. In *Proc. ICCV*, pages 754–760, 1998.
- [13] F. Schaffalitzky and A. Zisserman. Viewpoint invariant texture matching and wide baseline stereo. In *Proc. 8th ICCV, Vancouver, Canada*, July 2001.
- [14] J. Stolfi. *Oriented Projective Geometry: A Framework for Geometric Computations*. Academic Press, Inc., 1250 Sixth Avenue, San Diego, CA 92101, 1991.
- [15] P. H. S. Torr and A. Zisserman. MLESAC: A new robust estimator with application to estimating image geometry. *CVIU*, 78:138–156, 2000.
- [16] T. Tuytelaars and L. Van Gool. Wide baseline stereo matching based on local, affinely invariant regions. In *Proc. 11th BMVC*, 2000.
- [17] T. Werner and T. Pajdla. Oriented matching constraints. In *Proc. of BMVC*, pages 441–450. BMVA, September 2001.



Juice

Shelf

Valbonne



Great Wall

Leuven

Corridor

Figure 1: The experimental settings. Inliers and outliers are superimposed over the first and second images respectively.