Improving RANSAC for Fast Landmark Recognition

Pablo Márquez-Neila[†], Jacobo García Miró[†], José M. Buenaposada[‡], Luis Baumela[†]

[†]Dep. Inteligencia Artificial
Facultad Informática
Universidad Politécnica de Madrid

[‡]Dep. Ciencias de la Computación
E.T.S.I. Informática
Universidad Rey Juan Carlos

http://www.dia.fi.upm.es/~pcr

Abstract

We introduce a procedure for recognizing and locating planar landmarks for mobile robot navigation, based in the detection and recognition of a set of interest points. We use RANSAC for fitting a homography and locating the landmark. Our main contribution is the introduction of a geometrical constraint that reduces the number of RANSAC iterations by discarding minimal subsets. In the experiments conducted we conclude that this constraint increases RANSAC performance by reducing in about 35% and 75% the number of iterations for affine and projective cameras, respectively.

1. Introduction

Vision is possibly the richest perception modality for mobile robot navigation. One fundamental problem in robotics is locating the robot in the environment. This is usually solved by recognizing and locating a set of known visual landmarks. These landmarks may have been previously captured and stored in a database [3] or automatically detected and located within a Simultaneous Localization and Mapping strategy [4]. More recently, the use of interest point detectors [8] and advanced classification techniques [11] have opened the possibility of building systems which detect and recognize a large set of landmarks.

In this paper we use an interest point detector and a classification procedure for detecting landmarks and estimating their relative orientation with the robot. Our landmarks are planar objects with texture suitable for extracting a set of corner features. We use RANSAC for fitting a homography to the set of features extracted from an image. From the homography we estimate the relative orientation between the camera and the landmark. In order to achieve a low computational cost in the landmark detection process, we use some well-known techniques for reducing the number of RANSAC iterations. We introduce a geometrical constraint that must be satisfied by all correct RANSAC minimal subsets. We discard all minimal RANSAC subsets that do not satisfy the constraint. For each discarded subset we save the computations required to fit a model to it and search for inliers to the fitted model, which are the two most expensive operations in RANSAC.

The organization of the paper is as follows. In section 2 we briefly review some previous results related to our problem. The landmark recognition procedure is described in section 3. In section 4 we describe the techniques used for improving the performance of RANSAC and introduce the geometrical constraint. Finally in sections 5 and 6 we describe some experimental results and draw conclusions.

2. Related work

Our work is mostly related to two fields of research: landmark recognition for localization and RANSAC improvements.

Landmark recognition and localization is the process that searches for known landmarks in unknown images and estimates their poses with respect to the camera coordinate system. Existing works handle this as a feature point matching problem, using local descriptors for feature points [4], or training classifiers over the space of possible feature points [6]. Such classifiers may be based on simple methods, as k-nearest neighbor, or more sophisticated ones, as randomized trees [5]. Our work concentrates on landmark recognition. Therefore, the landmark is assumed known at the beginning of our procedure. If this is not the case, landmark detection methods [4] may find suitable landmarks by examining the environment before their localization.

RANSAC [2] is a procedure used to robustly fit a model to data in presence of outliers. Some work has been performed on possible RANSAC improvements. MLESAC [10, 9] does not count inliers given a fixed threshold as RANSAC does, but takes a probabilistic approach to as-

sess the validity of a hypothesis. *Randomized RANSAC* [1] tries to increase efficiency randomizing the RANSAC hypothesis evaluation step. All these works introduce general RANSAC improvements, which are applicable in all cases. However, none of those improvements takes advantage of some specific conditions that hold in RANSAC homography estimations. The main contribution of this paper is the introduction of a geometrical restriction appearing in the estimations of homographies.

3. Planar landmark recognition

Our approach to landmark recognition is closely related to the work presented in [6]. We use a two-stage algorithm to recognize and locate the landmark.

In the first stage, we present to the system a frontal view of the planar landmark, known as the *landmark model*. The system extracts salient image patches from the model and trains a classifier $\hat{Y}$. This classifier manages to recognize the extracted model image patches in a wide range of scales, pose and illumination conditions. This is an off-line stage and, as such, it has to be done only once.

The second stage performs on-line. Landmark recognition is carried out by applying the classifier $\hat{Y}$ over prominent patches of new images presented to the system. If the landmark is present, classifier $\hat{Y}$ will detect most of the feature points. With the locations of the recognized key-points we estimate the precise pose of the landmark on the observed image and compute the relative location between camera and scene. Unfortunately many key-points are mismatched. We use an improved version of RANSAC [2] to efficiently and robustly fit the landmark model to the image and discard matching outliers.

3.1. Classifier training

As explained above, the classifier must learn to match a set of image patches in different conditions of poses and illumination. This could be achieved by providing a large number of sample images, each one presenting the landmark with a different orientation and translation. However, this is not practical (or even possible) in most real world situations and it would require a tedious work for the user.

Instead, the user has to provide only one landmark model to our system, as the one shown in Fig. 1. Thanks to the planarity assumption, images of the landmark in different poses are synthesized automatically applying homographies to the original plane image. We systematically compute homographies for different orientation angles about the principal axes (X, Y, Z) and translations along the Z-axis of the camera reference system. Some of these synthetically generated views of the landmark are shown in Fig. 2.

After having computed views for different landmark poses, we proceed to extract salient image patches from

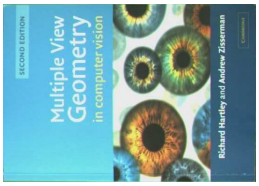

Figure 1. Example of a landmark model presented to the system.

Figure 2. Automatically generated views of the landmark.

each view. Specifically, our system applies Harris corner detector, although any other feature point detector, as SIFT [7], could have been used as well. We obtain a set of patches $\mathbf{p}_i \in \mathcal{P}$ of prominent parts of the images in the view set, where $\mathcal{P}$ is the set of all possible patches of a given size. Most of these image patches are views of the same landmark key-point k_j in different orientations and scales, as can be seen in Fig. 3. Only key-points with a high probability of being found in most images of the view set are of interest. Key-points not meeting this requirement and their related image patches are discarded. The remaining patches $\mathbf{p}_i$ associated with specific landmark key-point k_j are labeled with the name of this key-point, and the set of tuples $(\mathbf{p}_i, k_j)$ obtained is the dataset used to train the classifier.

Our classifier is a mapping $\hat{Y} : \mathcal{P} \rightarrow K$ which, given a image patch, estimates its associated key-point. $K = \{k_0, k_1, \ldots, k_N\}$ is the set of valid landmark key-points and the special key-point k_0, which denotes points that do not belong to the landmark, *i.e.*, an unknown key-point. The classifier uses the k-nearest neighbor algorithm with a threshold to classify new patches (although other methods, as randomized trees [5], could have been used as well). If, given a new patch $\mathbf{q}$, the distance from $\mathbf{q}$ to its nearest neighbor is larger than a fixed threshold, then $\mathbf{q}$ is assigned to key-point k_0. Otherwise, $\mathbf{q}$ is classified to the most frequent key-point of its k nearest neighbors. In our

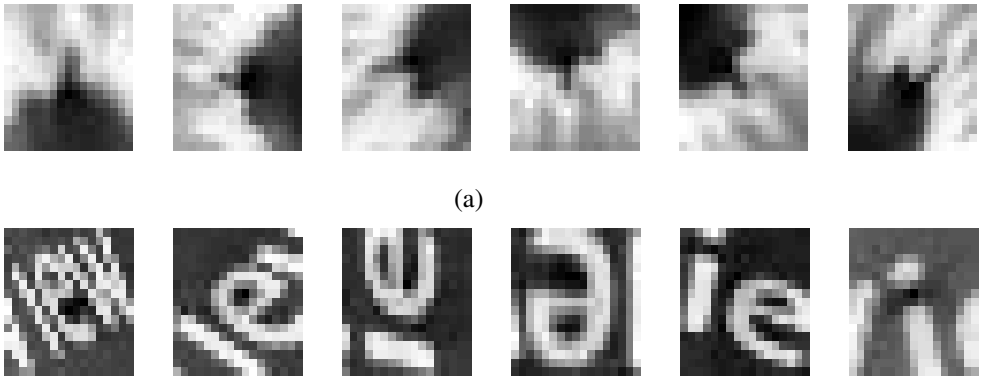

(a)

(b)

Figure 3. Some image patches of the same landmark key-point in different orientations and scales.

experiments, we have fixed $k = 5$.

The dimension of the patches' space $\mathcal{P}$ is quite large, even for patches of small size (*e.g.*, for patches of 17×17 pixels, the space has 289 dimensions). We use Principal Component Analysis to reduce the dimensionality of this space.

Finally, invariance to illumination changes is achieved by normalizing image patches from the view set so that their intensities have the same maximum and minimum values. Each patch is independently normalized before training the classifier, so that complex illumination conditions (as partial shadows) are correctly handled.

3.2. Key-point recognition for incoming images

Since the classifier has been trained for different key-point projective transformations, no further preprocessing to achieve invariance is needed. For each incoming image, the Harris corner detector finds salient patches on it. Then the classifier identifies each patch as a known key-point or as a point not belonging to the landmark. Correspondences between the locations of key-points in the landmark model image and in the incoming image are then used to estimate the homography between both planes. Obviously, the classifier is subject to error, leading to wrong key-point correspondences that may affect negatively the homography estimation. We use RANSAC as a robust estimation method in presence of outliers.

3.3. Robust landmark pose estimation with RANSAC

Given $\mathcal{C}$, a set of correspondences between key-points in the model and in the input image, in which there may be outliers, RANSAC computes the homography that best fits the data, ignoring, to some extent, outliers.

RANSAC operation is simple and well known. Iteratively, it takes random subsets from the original set $\mathcal{C}$ and fits a model (in this case, a homography) to them. The number of elements in each subset is the minimum for fitting the model (specifically, four for estimating a projectivity and three for an affinity). This is what we call a *minimal subset*. Then, the model is tested against all other correspondences in $\mathcal{C}$. Correspondences that fit the model are considered as hypothetical inliers. Those that do not fit are considered as hypothetical outliers. After a fixed number $\mathcal{I}$ of iterations, the model with highest number of hypothetical inliers is selected.

Algorithm 1 RANSAC algorithm for homography estimation.

1: **for** $h = 1$ to $h = \mathcal{I}_{\max}$ **do**
2: Take randomly a subset S_h of hypothetical inliers from the correspondence set $\mathcal{C}$.
3: Fit a homography M_h to the hypothetical inliers.
4: Count the number of inliers u_h for M_h (*i.e.*, the number of correspondences in $\mathcal{C}$ their error is smaller than predefined threshold θ).
5: **if** $u_h > u^*$ **then**
6: Holds current homography M_h as the best one.
7: $u^* \leftarrow u_h$.
8: **end if**
9: **end for**

Given a proportion p of inliers in the dataset, the probability P of finding a correct hypothesis after $\mathcal{I}$ RANSAC iterations is given by [2]:

$$P = 1 - (1 - p^m)^{\mathcal{I}}, \tag{1}$$

where m is the size of the minimal subset ($m = 4$ for estimating projective homographies and $m = 3$ for affine deformations). Therefore, given a desired confidence level P, the theoretical number of necessary RANSAC iterations is

$$\mathcal{I}_{\max} = \frac{\log(1 - P)}{\log(1 - p^m)}. \tag{2}$$

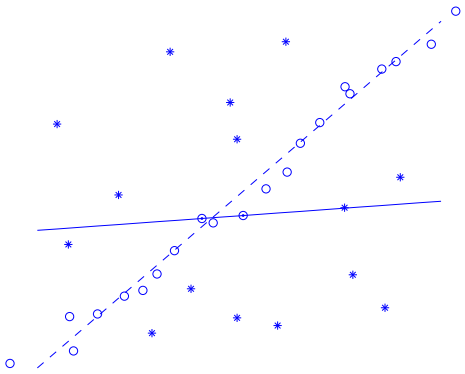

Figure 4. Example of noisy inliers which lead to a wrong hypothesis. ∘'s are noisy inliers, and ∗'s mark outliers. Indicated inliers lead to a wrong model (continuous line). The correct model is given by the dashed line.

In our experiments we have found this estimation to be over-optimistic for noisy inliers, since a minimal subset of correct inliers may not lead to a valid hypothesis, as shown in Fig. 4.

So, the actual number $\mathcal{I}^*$ of RANSAC iterations is usually larger than the theoretical value $\mathcal{I}_{\max}$ and condition $\mathcal{I}^* \geq \mathcal{I}_{\max}$ always holds for a given confidence level.

4. Improved RANSAC

In order to improve the efficiency of RANSAC we aim to reduce the actual number $\mathcal{I}^*$ of iterations by introducing some improvements to the algorithm. The first two improvements presented here are based on the work of Tordoff and Murray [9]. We also introduce a geometric constraint that must be met by inliers in a minimal subset. This constraint serves as a test to be applied to all minimal subsets before the homography is fitted.

4.1. Probabilistic correspondence subset selection

When minimal subset S_h is randomly selected in each RANSAC iteration, the probability of taking correspondence v_i is the same for all elements in $\mathcal{C}$. However, not all correspondences are equally probable. Using a k-nearest neighbor classifier, it is possible to infer a discrete probability distribution π over the possible correspondences v_i. Then the performance of RANSAC may be improved by choosing correspondences according to $\pi(v_i)$. That is, when minimal subsets of four correspondences are selected, a number $0 < \rho \leq 1$ is drawn uniformly for each one. Correspondence v_i is chosen if $\Pi_{i-1} < \rho \leq \Pi_i$, where $\Pi_j = \sum_{i=0}^{j} \pi(v_i)$, and then added to the hypothetical inliers set if it is not already there.

4.2. Weighted inliers counting

When evaluating the quality of the estimated homography M_h, only correspondences with a low error (lower than a threshold θ) are considered. The quality of M_h is then the number of estimated inliers. However, each of these correspondences v_i do not fit equally well to the model M_h, with some of them having lower errors than others. Using the above quality measure, homographies with a large number of bad inliers are considered better than those with a lower number of very good inliers.

This approach tries to solve this drawback. Given the error ε_h^i of correspondence v_i with homography M_h, we can compute a weight associated to it as

$$w(\varepsilon_h^i) = \begin{cases} 1 - \dfrac{\varepsilon_h^i}{\theta} & \text{if } \varepsilon_h^i \leq \theta \\ 0 & \text{otherwise} \end{cases} . \tag{3}$$

Then, quality of homography M_h is given by

$$u_h = \sum_{v_i \in \mathcal{C}} w(\varepsilon_h^i). \tag{4}$$

This leads to a more accurate estimation of homography quality, since correspondences having lower errors will be more relevant than those correspondences with high values of ε_h^i.

4.3. Geometric constraint for RANSAC

Not all minimal subsets lead to valid homographies. We may improve the performance of RANSAC by checking whether a minimal subset is valid, before actually fitting model M_h to them (step 3 in algorithm 1). This would avoid estimating the homography and searching for inliers to the fitted model, which are by far the two most expensive operations of RANSAC.

The geometric constraint that we introduce here must be met by all minimal subsets S_h under affine projection. Let I_0 be the landmark model given to the system as explained above, and I_t be the incoming image at instant t. Let $R_0 = \{\mathbf{r}_1^0, \mathbf{r}_2^0, \ldots, \mathbf{r}_N^0\} \subset \mathbb{R}^2$ be the set of N two-dimensional coordinates of key-points extracted in the off-line stage of our algorithm from I_0, and $R_t = \{\mathbf{r}_1^t, \mathbf{r}_2^t, \ldots, \mathbf{r}_{N'}^t\} \subset \mathbb{R}^2$ be the set of N' two-dimensional coordinates of feature points extracted from I_t using the Harris corner detector. A correspondence $v_i \in \mathcal{C}$ is a tuple $(\mathbf{r}_j^0, \mathbf{r}_k^t)$ which maps an element of R_0 to an element of R_t. A minimal subset is defined as a set $S_h = \{v_a, v_b, v_c\}$ of three correspondences.

Given any set S_h of three correspondences

$$\begin{aligned} v_a &= (A', A),\ A' \in R_0,\ A \in R_t, \\ v_b &= (B', B),\ B' \in R_0,\ B \in R_t, \\ v_c &= (C', C),\ C' \in R_0,\ C \in R_t, \end{aligned}$$

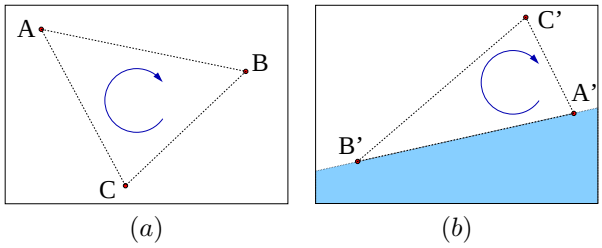

Figure 5. Geometric constraint that must be met by S_h in each iteration of RANSAC. (a) Points A, B, C in image I_t. (b) Set S_h of hypothetical inliers gives their corresponding points A', B', C' in landmark model image I_0. Under affine deformations, point C' must not be located in marked region, since A', B', C' must have the same relative order than A, B, C. If this constraint does not hold, set S_h should be discarded.

it holds that relative order of points A', B', C' and that of points A, B, C must be the same. Formally, using the vectors $\phi = B - A$, $\psi = C - A$ and their corresponding vectors ϕ' and ψ', we express the rule as:

$$sgn\left(\begin{vmatrix}(\phi)_x & (\phi)_y\\ (\psi)_x & (\psi)_y\end{vmatrix}\right) = sgn\left(\begin{vmatrix}(\phi')_x & (\phi')_y\\ (\psi')_x & (\psi')_y\end{vmatrix}\right), \quad (5)$$

where $|\cdot|$ is the determinant function, $(\mathbf{r})_x$ and $(\mathbf{r})_y$ refer to the coordinates x and y of vector $\mathbf{r}$, respectively, and $sgn(x)$ is the sign function defined as usual. This condition is graphically depicted in Fig. 5. All sets of correspondences which do not hold this geometric constraint should be discarded, since they would surely lead to an invalid homography.

When estimating projectivity homographies, in which four correspondences are needed, one might test for this condition applying eq. (5) over first three correspondences (v_a, v_b and v_c), over all four possible sets of three correspondences, or over some of these four possible sets.

The proposed geometric constraint may be used in conjunction with probabilistic S_h selection and weighted inliers counting, as explained above, to obtain an optimized version of RANSAC listed in algorithm 2.

Algorithm 2 Improved RANSAC for homography estimation.

for $h = 1$ to $h = \mathcal{I}_{\max}$ **do**
  Take a subset S_h of hypothetical inliers according to $\pi(v_i)$ as explained in section 4.1.
  if S_h passes geometric test in section 4.3 **then**
    Fit a homography M_h to the hypothetical inliers S_h.
    Compute the homography quality u_h using eq. (4).
    if $u_h > u^*$ **then**
      Holds current homography M_h as the best one.
      $u^* \leftarrow u_h$.
    end if
  end if
end for

5. Experiments

Our scheme for landmark recognition has been tested in several ways. First, we made a general experiment to test the framework as a whole. A landmark model served to train a k-nearest neighbor classifier in the off-line stage. Then, in the online stage, we used a video sequence with landmark images taken in different poses. Pose estimation results are displayed in Fig. 6.

As we saw in previous sections, the real number of needed RANSAC iterations $\mathcal{I}^*$ is usually larger than theoretical number $\mathcal{I}_{\max}$. Our geometric constraint aims to reduce the real number of RANSAC iterations and bring it nearer to $\mathcal{I}_{\max}$. We have noticed that the inlier proportion p and the Gaussian noise added to the positions of prominent points in I_t have a relevant influence in the performance of RANSAC. The lower the inlier proportion and the higher the noise level, the higher the number of needed RANSAC iterations.

In following experiments, we try to assess the impact of our geometric restriction in the performance of RANSAC under different conditions of inlier proportion and noise level. In each experiment we proceed as follows. First, we set an inlier proportion and a noise level, and the set of correspondences $\mathcal{C}$ is generated using these parameters. Then, we run RANSAC 5000 times using the obtained set of correspondences $\mathcal{C}$. Each of these executions is halted at the iteration in which a good homography (*i.e.*, a homography that fits to at least 85 percent of inliers in $\mathcal{C}$) is found. From each execution we record the required number of iterations using standard RANSAC and the required number of iterations using our geometric restriction. We also compute the number of theoretical RANSAC iterations with eq. (2).

Experiments 1–3 have a fixed proportion of inliers ($p = 0.6$) with different levels of Gaussian noise added to the image points $\mathbf{r}_i^t \in R_t$: experiment 1 adds Gaussian noise with standard deviation $\sigma = 2$ pixels to the extracted image points; experiment 2 uses $\sigma = 3$ pixel and experiment 3 does not add noise ($\sigma = 0$ pixels). Experiment 4 applies Gaussian noise with $\sigma = 2$ pixels, but it has a reduced proportion of outliers ($p = 0.7$), while we increase the proportion of outliers ($p = 0.5$) in experiment 5.

As listed above, our geometric constraint works by pruning subsets of invalid correspondences under affine transformations. When applied under projective transformations, we have introduced three alternatives to use our restriction in each iteration: (a) check one subset $\{v_a, v_b, v_c\} \subset S_h$ of three arbitrary correspondences; (b) check all possible four subsets of three correspondences of each S_h and (c) check

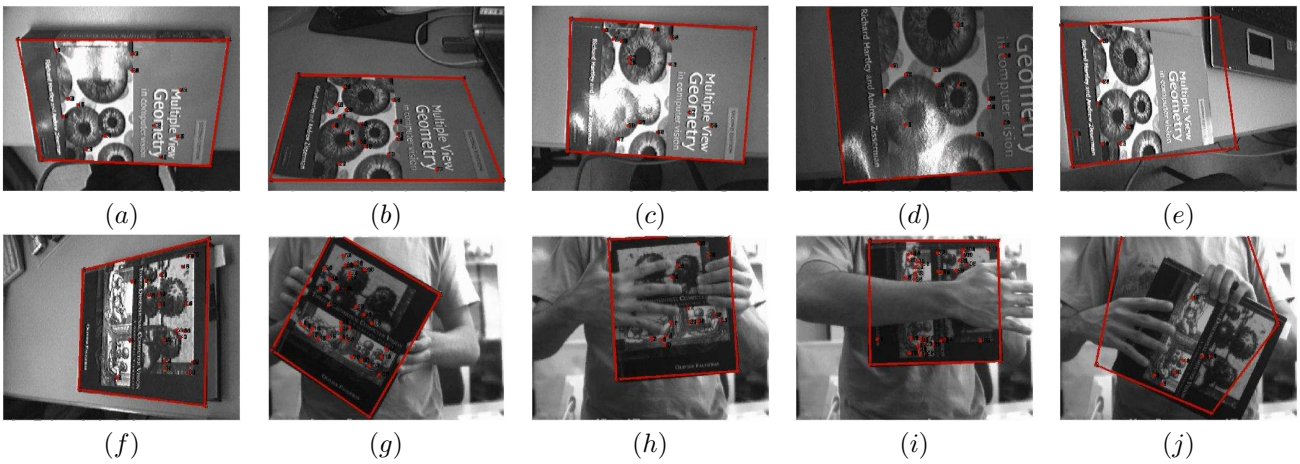

Figure 6. Recognition of two different landmarks in different poses. Points represent extracted feature points that were recognized by the classifier. Lines are the estimation of the landmark edges after fitting a homography with RANSAC. (a), (b), (f), (g) Landmark correctly recognized in different poses. (c), (d), (h), (i) Our method is robust to partial occlusions and different illumination conditions. (e), (j) Wrong estimations.

	Inlier prop.	**Noise level**	**Transform.**
Exp. 1	$p = 0.6$	$\sigma = 2$	Affine
Exp. 2	$p = 0.6$	$\sigma = 3$	Affine
Exp. 3	$p = 0.6$	$\sigma = 0$	Affine
Exp. 4	$p = 0.7$	$\sigma = 2$	Affine
Exp. 6	$p = 0.5$	$\sigma = 2$	Affine
Exp. 6	$p = 0.6$	$\sigma = 2$	Projective

Table 1. Experiment configurations. Each experiment has different a value of inlier proportion p (first column in the table), Gaussian noise level (second column) and type of transformation (third column).

only some of these subsets. We analyze the ability of our method under non-affine transformations with alternatives (a) and (b) in experiment 6(a) and 6(b), respectively.

Table 1 summarizes the configurations of experiments, and Fig. 7 displays the results. This Figure gives valuable information. First, the actual number of RANSAC iterations is higher than the predicted theoretical value, when noisy input inliers are present. In all experiments, our cheap constraint detects and prunes near half of S_h. This leads to a reduction in the number of iterations to be completed in each RANSAC execution by around 35% in affine transformations. Specifically, in experiments 1–5 this reduction is, on average for all runs, 37.09%, 35.89%, 35.88%, 29.71%, and 39.75%, respectively. Experiment 3 leads to an interesting result. Under noise-free input data, the actual number of RANSAC iterations matches that given by the theoretical prediction made with eq. (2). Since our approach filters out some invalid correspondences, the improved number of iterations is lower than the theoretical value computed with (2).

It should be noted that we can apply our constraint under projective transformations obtaining good results. Experiment 6(a) shows that our first alternative (one test per iteration) provides a reduction of iterations of 39.34%, which is not far from results in affinity cases. However, second alternative (four tests per iteration) increases the ratio of filtered subsets up to 74.36%.

Time improvement due to our method depends on the size of correspondence set $\mathcal{C}$. In experiment 6(b), we got a time reduction of 45.41% with 116 correspondences using our unoptimized implementation.

6. Conclusions and future work

We have presented an efficient method for landmark recognition. It consists of an off-line stage, in which the landmark is learned, and an on-line stage, in which the landmark is recognized and its pose is estimated. We construct a k-nearest neighbor classifier for detecting landmark keypoints, and use it to find possible landmark points in new incoming images. We then use RANSAC for fitting a homography to the estimated correspondences between model key-points and detected feature points in incoming images. Our major contribution in this work is a geometric constraint for a faster RANSAC. It aims to filter out minimal sets of correspondences that are not going to lead to valid homographies. This constraint lowers by a 35% the number of required RANSAC iterations in the affine case, and up to a 75% in the projective case.

In practice, our approach can be used for landmark-based localization with mobile platforms. Any kind of textured, planar landmark can be used. It may be used for recogniz-

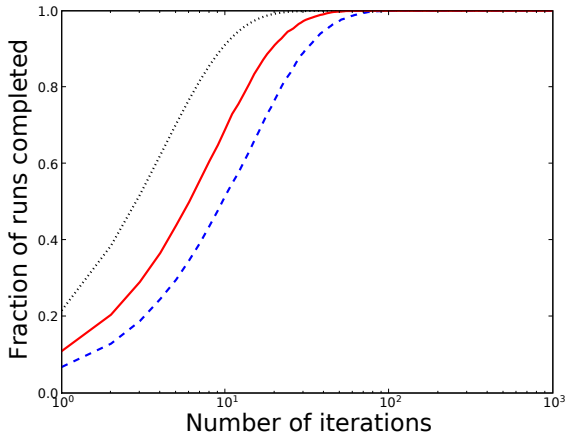

Experiment 1

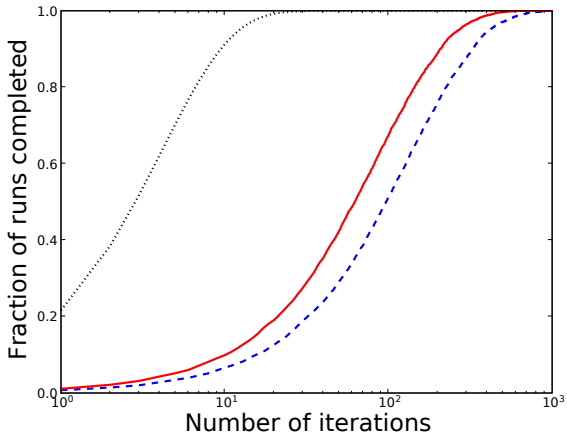

Experiment 2

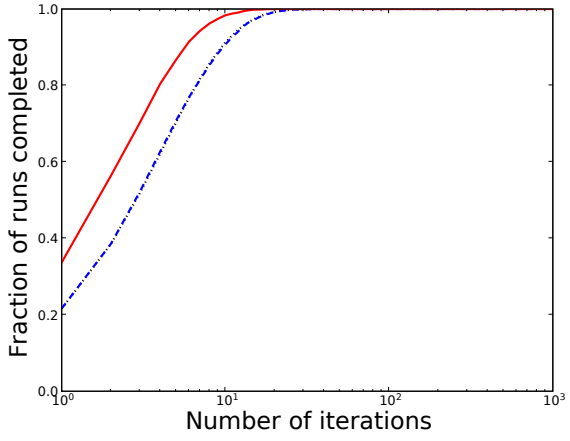

Experiment 3

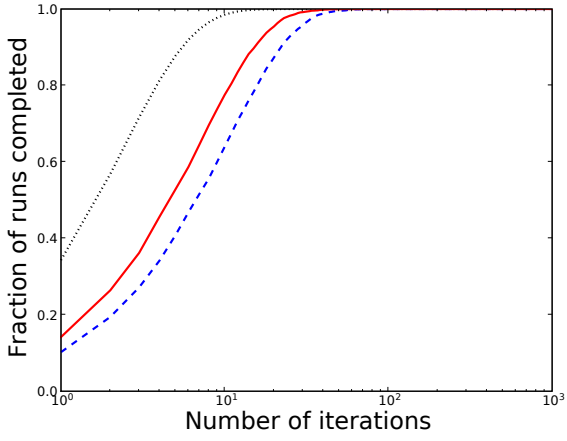

Experiment 4

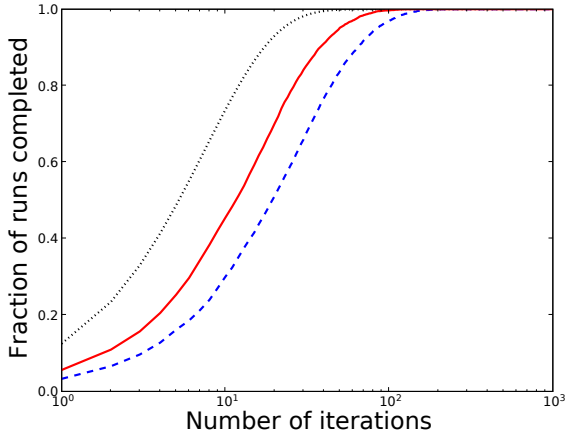

Experiment 5

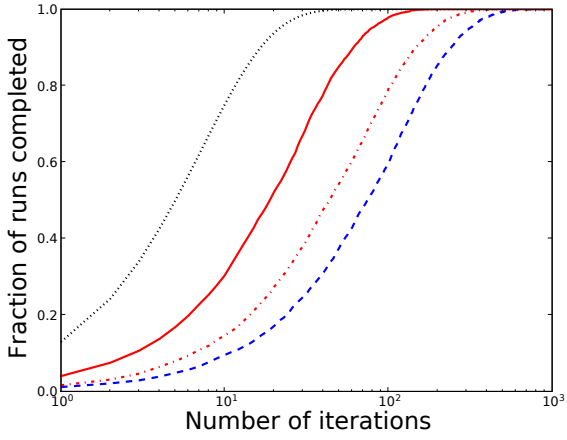

Experiment 6

Figure 7. RANSAC was run 5000 times and each run stopped when the first good solution is found. Each figure shows the cumulative proportion of runs completed by a particular iteration number, plotted against that iteration number. The dotted (····) curve shows theoretical values given by eq. (2); the dashed (- - -) curve shows actual results obtained using the standard RANSAC algorithm; finally the continuous (—) curve shows results obtained using our geometrical restriction. In last figure, dash-dot curve (- · - · -) represents experiment 6(a) outcome (one test per iteration), and continuous curve (—) shows experiment 6(b) outcome (four tests per iteration).

ing several different planar landmarks by running more than one instance of the algorithm.

We are working in several directions to improve the geometric constraint introduced in the paper. First, simplifying its application in non-orthographic images so that an only, more restrictive condition is needed, instead of repeating the same affinity condition several times. Second, searching for other similar geometric constraints applicable to pose and parameter estimation of different types of surfaces.

Acknowledgments

The authors gratefully acknowledge funding from the Spanish *Ministerio de Educación y Ciencia* under contract `TRA2005-08592-C02-02`. Pablo Márquez-Neila was funded by the *Programa Personal Investigador de Apoyo* from the *Comunidad de Madrid*.

References

[1] O. Chum and J. Matas. Randomized RANSAC with $T_{d,d}$ test. In *Proc. of British Machine Vision Conference*, volume 2, pages 448–457, 2002.

[2] M. A. Fishler and R. C. Bolles. Random sample consensus: a paradigm for model tting with applications to image analysis and automated cartography. *Tech report 213, AI Center, SRI International*, 1980.

[3] J. Gaspar, N. Winters, and J. Santos-Victor. Vision-based navigation and environmental representations with an omnidirectional camera. *IEEE Trans. on Robotics and Automation*, 16(6):890–898, 2000.

[4] J. Hayet, F. Lerasle, and M. Devy. A visual landmark framework for mobile robot navigation. *Image and Vision Computing*, 25(8):1341–1351, August 2007.

[5] V. Lepetit and P. Fua. Keypoint recognition using randomized trees. *Transactions on Pattern Analysis and Machine Intelligence*, 28(9):1465–1479, 2006.

[6] V. Lepetit, J. Pilet, and P. Fua. Point matching as a classification problem for fast and robust object pose estimation. In *Proc. of Computer Vision and Pattern Recognition*, volume 2, pages 244–250, 2004.

[7] D. Lowe. Distinctive image features from scale-invariant keypoints. *Int. Journal of Computer Vision*, 60(2):91–110, November 2004.

[8] K. Mikolajczyk, T. Tuytelaars, C. Schmid, A. Zisserman, J. Matas, F. Schaffalitzky, T. Kadir, and L. Van Gool. A comparison of affine region detectors. *Int. Journal of Computer Vision*, 65(1-2):43–72, November 2005.

[9] D. W. Tordoff, Ben J.and Murray. Guided-MLESAC: Faster image transform estimation by using matching priors. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(10):1523–1535, 2005.

[10] P. Torr and A. Zisserman. MLESAC: A new robust estimator with application to estimating image geometry. *Computer Vision and Image Understanding*, 78:138–156, 2000.

[11] A. Torralba, K. Murphy, and W. Freeman. Sharing visual features for multiclass and multiview object detection. *IEEE Trans. Pattern Anal. Mach. Intell.*, 29(5):854–869, May 2007.