

Global Localization and Relative Positioning Based on Scale-Invariant Keypoints

Jana Košecká, Fayin Li and Xialong Yang

Computer Science Department, George Mason University, Fairfax, VA 22030

Abstract

The localization capability of a mobile robot is central to basic navigation and map building tasks. We describe a probabilistic environment model which facilitates global localization scheme by means of location recognition. In the exploration stage the environment is partitioned into locations, each characterized by a set of scale-invariant keypoints. The descriptors associated with these keypoints can be robustly matched despite changes in contrast, scale and viewpoint. We demonstrate the efficacy of these features for location recognition, where given a new view the most likely location from which this view came from is determined. The misclassifications due to dynamic changes in the environment or inherent appearance ambiguities are overcome by exploiting location neighborhood relationships captured by a Hidden Markov Model. We report the recognition performance of this approach in an indoor environment consisting of eighteen locations and discuss the suitability of this approach for a more general class of recognition problems. Once the most likely location has been determined we demonstrate how to robustly compute the relative pose between the representative view and the current view.

1 Introduction and Related Work

The acquisition of unknown environment models, navigation and pose maintenance belong to the essential capabilities of a mobile robots. The approaches for vision-based model acquisition and localization typically strived to obtain either metric or topological models. The topological models were commonly induced by visibility regions associated with the artificial landmarks. Artificial landmarks simplified the issues of landmark recognition and enabled reliable estimation of the robot's pose with respect to a landmark (1; 2). The techniques which tried to bypass the choice of artificial landmarks have been mainly motivated by approaches used for object

Email address: {kosecka,fl,xyang}@cs.gmu.edu (Jana Košecká, Fayin Li and Xialong Yang).

recognition. One of the main concerns of these methods is the choice of image representation, which could guarantee some amount of invariance with respect to variations in pose, illumination and scale and be robust to partial occlusion. The image representations proposed in the past comprised of descriptors computed locally at selected image locations or globally over the entire image.

The image locations were selected using various saliency measures and were characterized by rotationally or affine invariant feature descriptors (3; 4; 5), which enabled effective matching of overlapping and possibly widely separated views. Alternative global descriptors were derived from local responses of filters at different orientations and scales (6), multi-dimensional histograms (7; 8) or orientation histograms (9). These types of global image descriptors integrated the spatial image information and enabled classification of views into coarser classes (e.g. corridors, open areas), yielding only qualitative localization. In case of omni-directional views representations in terms of eigenviews obtained by principal component analysis were applied successfully both for topological and metric model acquisition, thanks to small variations of the image appearance within a location and rotationally invariant image representations (10; 11). The problem of building a metric model and simultaneous localization (SLAM) using solely visual sensing has been demonstrated successfully in case of smaller, single room environments using single video stream (12) or trinocular stereo (13). The use of local point features for both metric and topological localization was proposed by (14). In both of these instances the odometry readings were used in connection with the visual measurements.

In order to enable the map building and localization in large scale environments solely by means of visual sensing, suitable representations of the environment at different spatial scales and associated means of localization need to be sought. The advantages of such representations have been pointed out previously by (15) both from the perspective of model building, localization as well as navigation given the model. These types of hybrid models have been already explored previously using ultrasound sensing (16). In the presented work, we tackle the problem of both global localization and relative positioning in large scale environments using only visual sensing.

1.1 Approach

In our approach, the final environment model is represented in terms of individual locations, each characterized by a set of representative views. Within the location we will endow the model with a local geometry relative to a set of representative views. The global localization stage is motivated by the recent advances in object recognition using local scale invariant features and adopts the strategy for localization by means of location recognition. The image sequence acquired by a robot

during the exploration is first partitioned to individual locations. The locations correspond to the regions in the space across which the features can be matched successfully. Each location is represented by a set of model views and their associated scale-invariant features. In the first localization stage, the current view is classified as belonging to one of the locations using the standard voting approach. In the second stage we exploit the knowledge about neighborhood relationships between individual locations captured by a Hidden Markov Model (HMM). The main contribution of this stage is the instantiation of the Hidden Markov Model in the context of this problem and demonstration of an improvement in the overall recognition rate. This topological structure is essential particularly in the case of large scale environments which often contain uninformative regions, violating the continuity of the mapping between the environment appearance and camera pose. In such case imposing a discrete structure on the space of continuous observations enables us to overcome these difficulties while maintaining a high recognition rate. Once the most likely view has been determined we will demonstrate how to compute the relative pose of the robot with respect to the model view in the absence of the focal length of the camera. This second stage will then enable local metric localization given the model.

2 Image Representation

In order to obtain image representation which captures the essential appearance of the location and is robust to occlusions and changes in image brightness we adopt the representation in terms of local scale-invariant features. The use of local features and their associated descriptors in the context of object recognition has been demonstrated successfully by several researchers in the past (4). In particular we examine the effectiveness of scale-invariant (SIFT) features proposed by D. Lowe (17).

2.1 Scale-Invariant Features

The SIFT features correspond to highly distinguishable image locations which can be detected efficiently and have been shown to be stable across wide variations of viewpoint and scale. Such image locations are detected by searching for peaks in the image $D(x, y, \sigma)$ obtained by taking a difference of two neighboring images in the scale space

$$D(x, y, \sigma) = (G(x, y, k\sigma) - G(x, y, \sigma)) * I(x, y) = L(x, y, k\sigma) - L(x, y, \sigma).$$

The image scale space $L(x, y, \sigma)$ is first build by convolving the image with the Gaussian kernel with varying σ , such that at particular σ , $L(x, y, \sigma) = G(x, y, \sigma) *$




Fig. 1. The circle represents the keypoint’s location, with the radius proportional to the keypoint’s scale.

$I(x, y)$. Candidate feature locations are obtained by searching for local maxima and minima of $D(x, y, \sigma)$. In the second stage the detected peaks with low contrast or poor localization are discarded. A more detailed discussion about enforcing the separation between the features, the sampling of the scale space and improvements in the feature localization can be found in (17; 18). Once the location and scale have been assigned to candidate keypoints, the dominant orientation is computed by determining the peaks in the orientation histogram of its local neighborhood weighted by the gradient magnitude. The keypoint descriptor is then formed by computing local orientation histograms (with 8 bin resolution) for each element of a 4×4 grid overlaid over 16×16 neighborhood of the point. This yields a 128 dimensional feature vector which is normalized to unit length in order to reduce the sensitivity to image contrast and brightness changes in the matching stage. Figure 1 shows the keypoints found in the example images in our environment. In the reported experiments the number of features detected in an image of size 480×640 varies between 10 to 1000. In many instances this relatively low number of keypoints is due to the fact that in indoor environments many images have a small number of textured regions. Note that the detected SIFT features correspond to distinguishable image regions and include both point features as well as regions along line segments.

3 Environment Model

The initial environment model is obtained in the exploration stage. Given a temporally sub-sampled sequence acquired during the exploration (images were taken approximately every 2-3 meters), The path along which the training sequence was taken visited all locations (some of them twice) and is outlined in Figure 2. In this data set the heading direction was in most cases aligned with the principal directions of the world coordinate frame. The transitions between individual locations are determined depending on the number of features which can be successfully matched between the successive frames. These are depicted in Figure 3 for a sequence captured by a still digital camera along the path which visited all eighteen

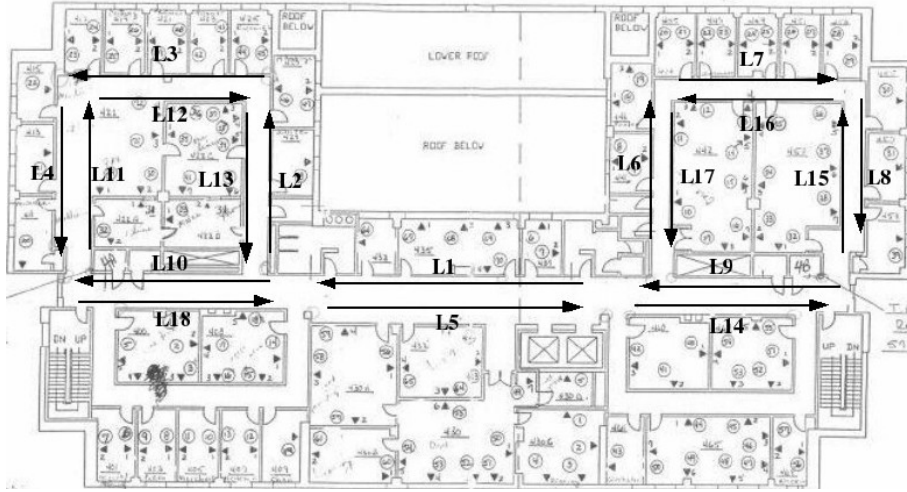


Fig. 2. The map on the fourth floor of our building. The arrows correspond to the heading of the robot and the labels represent individual locations.




Fig. 3. The number of keypoints matched between consecutive views for the sequence comprised of 18 locations (snapshot was taken every 2-3 meters) captured by a digital camera (left); the number of keypoints detected in each frame (right).

locations. The transitions between individual locations are marked by the peaks in the graph, which are automatically detected. The training sequence is partitioned into 18 different locations, each represented by an image sub-sequence separated by two peaks. Different locations in our model correspond to hallways, sections of corridors and meeting rooms approached at different headings. The transitions between the locations typically occur either at places where navigation decisions have to be made or when the appearance of the location changes suddenly. In the current experiment, the environment is mostly comprised of a network of rectangular corridors and hallways which are typically traversed with four possible headings (N, S, W, E). The deviations from these headings can be handled as long as there is a sufficient overlap between the model views acquired during the exploration and current views. The number of views per location vary between 8 to 20 depending on the appearance variation within the location. In order to obtain a more compact representation of each location a number of representative views is chosen per location, each characterized by a set of SIFT features. The sparsity of the model is directly related to the capability of matching SIFT features in the presence of



Fig. 4. Examples of representative views of 14 out of 18 locations.

larger variations in scale. The number of representative views varied between one to four per location and was obtained by regular sampling of the views belonging to individual locations. Examples of representative views associated with individual locations are depicted in Figure 4. The number of features in the model can be further reduced by keeping only the features, which discriminate each location well. This issue is addressed in more detail in (19).

4 Location recognition

The environment model obtained in the previous section consists of a database of model views ¹. The i -th location in the model, with $i = 1, \dots, N$ is represented by n views I_1^i, \dots, I_n^i with $n \in \{1, 2, 3, 4\}$ and each view is represented by a set of SIFT features $\{S_k(I_j^i)\}$, where k is the number of features. In the initial stage we tested the location recognition by using a simple voting scheme.

4.1 Voting Approach

Given a new query image Q and its associated keypoints $\{S_l(Q)\}$ a set of corresponding keypoints between Q and each model view I_j^i , $\{C(Q, I_j^i)\}$, is first computed. The correspondence is determined by matching each keypoint in $\{S_l(Q)\}$ against the database of $\{S_k(I_j^i)\}$ keypoints and choosing the nearest neighbor based on the Euclidean distance between two descriptors. We only consider point matches with high discrimination capability, whose nearest neighbor is at least 0.6 times closer than the second nearest neighbor. More detailed justification behind the

¹ It is our intention to attain a representation of location in terms of views (as opposed to some abstract features) in order to facilitate relative positioning tasks in the later metric localization stage.

representative	train	test	test
# of views	#1 (250)	#2 (134)	#3 (130)
one	84%	46%	44%
two	97.6%	68%	66%
four	100%	82%	83%

Table 1

Recognition performance for one training and two test sequences in terms of % of correctly classified views as a function of number of representative views.

choice of this threshold can be found in (17). In the subsequent voting scheme we determine the location whose keypoints were most frequently classified as nearest neighbors. The location where the query image Q came from is then determined based on the number of successfully matched points among all model views

$$C(i) = \max_j |\{C(Q, I_j^i)\}| \text{ and } [l, num] = \max_i C(i)$$

where l is the index of location with maximum number num of matched keypoints. Table 1 shows the location recognition results as a function of number of representative views per location on the training sequence of 250 views and two test sequences of 134 and 130 images each. All three sequences were sparse, with images taken 2-3 meters apart. The two test sequences were taken at different days and times of day, exhibiting larger deviations from the path traversed during the training. Despite a large number of representative views per location relatively poor performance on the second and third test sequence was due to several changes in the environment between the training and testing stage. In 5 out of 18 locations several objects were moved or misplaced. Examples of dynamic changes can be seen in Figure 5. The poorer performance due to dynamic changes is not surprising, since the most discriminative SIFT features often belong to objects some of which are not inherent to particular locations. In the next section we describe how to resolve these issues by modelling the neighborhood relationships between individual locations.

5 Modelling spatial relationships between locations

We propose to resolve these difficulties by incorporating additional knowledge about neighborhood relationships between individual locations. The rationale behind this choice is that despite the presence of ambiguities in recognition of individual views the temporal context should be instrumental in resolving them. The use of temporal context is motivated by the work of (6) which addresses the place recognition problem in the context of wearable computing application. The temporal context is determined by spatial relationships between individual locations and

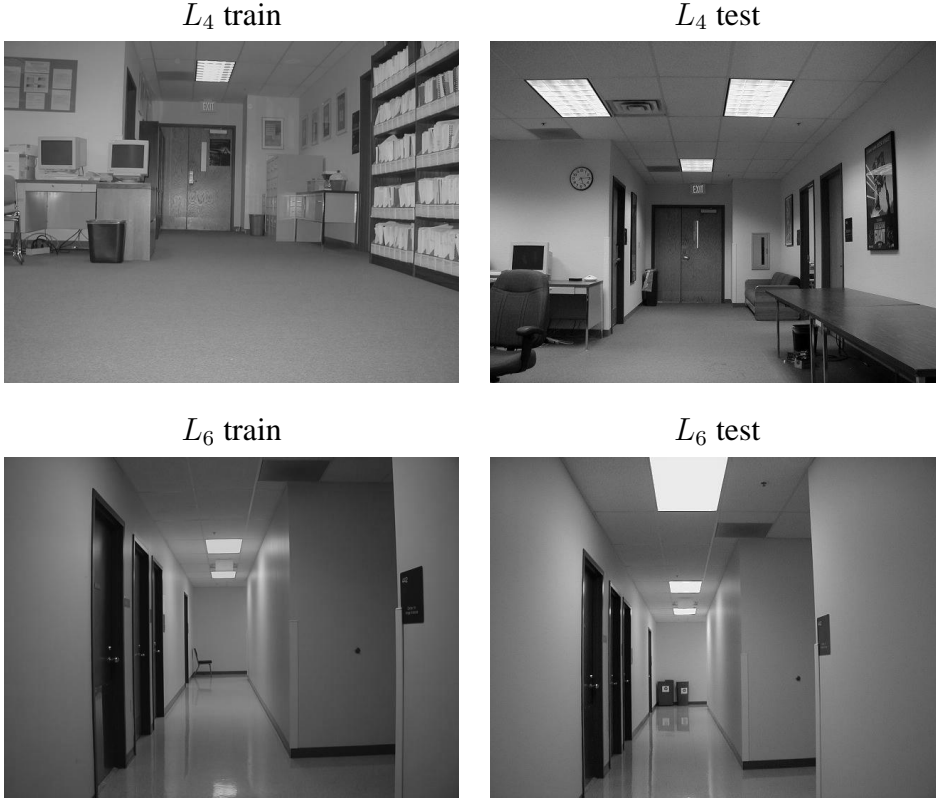


Fig. 5. Changes in the appearance of location L_4 and L_6 between the training and testing. In the top image pair the bookshelve was replaced by a table and couch and in the bottom pair recycling bins were removed.

is modelled by a Hidden Markov Model (HMM). In this model the states correspond to individual locations and the transition function determines the probability of transition from one state to another. Since the locations cannot be observed directly each location is characterized by an observation o_t and its associated likelihood $p(o_t|L_t = l_i)$. The most likely location is at each instance of time obtained by maximizing the conditional probability $P(L_t = l_i|o_{1:t})$ of being at time t and location l_i given the available observations up to time t . The location likelihood can then be estimated recursively using the following formula

$$P(L_t = l_i|o_{1:t}) \propto p(o_t|L_t = l_i)P(L_t = l_i|o_{1:t-1}) \quad (1)$$

where the observation likelihood $p(o_t|L_t = l_i)$ characterizes how likely is the observation o_t at time t to come from location l_i . The choice of observation likelihood depends on the available image representation and the matching criterion. When local descriptors are used as observations, several such choices have been proposed in the context of probabilistic approaches to object recognition (20; 21). The proposed likelihood functions account for the feature density and spatial relationships between features and have been shown to improve overall recognition rate. In the context of global image descriptors the locations were modelled in terms of Gaussian mixture models proposed in (6). Since the location recognition problem is

notably simpler than the object recognition problem due to the background clutter not being so prominent ², we use slightly different form of the probabilistic model. The essential features of this probabilistic model, which we describe next, is the selection of discriminative features in the matching stage and integration of the evidence they provide for individual locations. Given a set of features in the query image $\{S_k(Q)\} = \{g_k^Q\}$, we first define the so called strangeness parameter α_k^i , which characterizes the discrimination capability of k -th feature, with respect to i -th location

$$\alpha_k^i = \frac{\min_{g_j \in S_i} (\|g_k^Q - g_j\|)}{\min_{g_j \notin S_i} (\|g_k^Q - g_j\|)}. \quad (2)$$

α_k^i is the ratio of minimal intra-distance within the class and minimal inter-distance to features from other classes. If α_k^i is greater than 1, the feature g_k^Q is not contributing to classification of Q as label l_i . The smaller the α_k^i is the more discriminative is the feature for the purpose of classifying Q as i -th location. For the computation of the likelihood we select only top r features from the query image, ranked by their strangeness. The likelihood of feature g_k^Q belonging to location $L = l$ (denoted here as L_l) is defined as

$$P(g_k^Q | L_l) = P(\alpha_k^l | L_l) \propto \exp(-\frac{(\alpha_k^l)^2}{2\sigma^2}). \quad (3)$$

Since we do not know how many features among $\{g_k^Q\}$ belong to location l , we need to integrate the evidence over all possible hypotheses and all measurements. A hypothesis in our case indicates that a subset $h_j \neq \emptyset$ of selected features $\{g_k^Q\}$ is classified as location l . Assuming the selected features are independent, we can now compute the probability of a single hypothesis h_j conditioned on location l

$$P(h_j | L_l) = \prod_m P(\alpha_m^l | L_l) \prod_n (1 - P(\alpha_n^l | L_l)), \quad (4)$$



where index m ranges over features which belong to location l with certain probability, n ranges over features which do not belong to location l and $m + n = r$ is the number of selected features, i.e. the length of a hypothesis. Then the probability a selected feature set $\{g_k^Q\}$ belonging to location l (assuming that the features are independent) can be computed as

$$P(\{g_k^Q\} | L_l) = \sum_{h_j} \prod_m P(\alpha_m^l | L_l) \prod_n (1 - P(\alpha_n^l | L_l)). \quad (5)$$



Using this probability to model the likelihood $P(\{g_k^{Q(l)}\} | L_l)$, it can be shown that $P(\{g_k^Q\} | L_l)$ can be simplified to

$$P(\{g_k^Q\} | L_l) \propto (1 - \prod_{k=1}^{M_l} (1 - \exp(-\frac{\alpha_k^l{}^2}{2\sigma^2}))). \quad (6)$$

² The probabilistic models used in the object recognition, must also account for the fact that large number of detected features comes from background and not the object.



Seq. 2 with and without HMM



Seq. 3 with and without HMM

Fig. 6. Classification results for Sequence 2 and Sequence 3 with and without considering the spatial relationships. The black circles correspond to the labels of most likely locations.

More detail justification for this likelihood model as well as the comparison with previously used approaches can be found in (19).

In order to explicitly incorporate the location neighborhood relationships, the second term of equation (1) can be further decomposed

$$P(L_t = l_i | o_{1:t-1}) = \sum_j^N A(i, j) P(L_{t-1} = l_j | o_{1:t-1}) \quad (7)$$

where N is the total number of locations and $A(i, j) = P(L_t = l_i | L_{t-1} = l_j)$ is the probability of two locations being adjacent. In the presence of a transition between two locations the corresponding entry of A was assigned a unit value and in the final stage all the rows of the matrix were normalized. The results of location recognition employing this model are in Figure 6. The recognition rate for Sequence 2 was 96.3% and for Sequence 3 it was 95.4%. The location label assigned to each image is the one with the highest probability. While in both cases some images were misclassified the overall recognition rates are an improvement compared to the rates

reported in Table 1. Despite the classification errors in Sequence 2, the order of visited locations was correctly determined. For Sequence 3, where we exhibited some intentional deviations between the path taken during training and testing, the classification of location 14 was incorrect. The effect of HMM model can be examined by making all the probabilities in the transition matrix A uniform essentially neglecting the knowledge of location neighborhood relationships. The assigned location labels for this case are in the right column of Figure 6, with noticeably degraded recognition performance.

6 Pose Estimation

Once the most likely location and best matched view has been found we can compute the relative displacement between the current view and model view.

The current view and the matched model view are related by a rigid body displacement $g = (R, T)$ represented by a rotation $R \in SO(3)$ and translation $T = [t_x, t_y, t_z]^T \in \mathbb{R}^3$. Provided that the camera is calibrated, g can be estimated from the epipolar geometry between the two views. This recovery problem can be further simplified taking into account the fact that the motion of the robot is restricted to a plane. Here we outline an algorithm for this special case and demonstrate how to recover the displacement in case of unknown focal length. The case of general motion and unknown focal length was studied by (22) and the solution for the case of planar motion case has been proposed by (?) in the context of uncalibrated stereo. Here we demonstrate a slightly different, more concise solution to the problem. Consider the perspective camera projection model, where 3D coordinates of point $\mathbf{X} = [X, Y, Z]^T$ are related to their image projections $\mathbf{x} = [x, y, 1]^T$ by an unknown scale λ ; $\lambda \mathbf{x} = \mathbf{X}$. In case the camera is calibrated the two views of the scene are related by $\lambda_2 \mathbf{x}_2 = R \lambda_1 \mathbf{x}_1 + T$, where $(R, T) \in SE(3)$ is a rigid body transformation and λ_1 and λ_2 are the unknown depths with respect to individual camera frames. After elimination of the unknown scales from the above equation, the relationship between the two views is captured by the so called epipolar constraint

$$\mathbf{x}_2^T \hat{T} R \mathbf{x}_1 = \mathbf{x}_2^T E \mathbf{x}_1 = 0, \quad (8)$$

where $E = \hat{T}R$ is the essential matrix³. In case of planar motion, assuming translation in $x - z$ plane and rotation around y -axis by an angle θ , the essential matrix

³ \hat{T} denotes a 3×3 skew symmetric matrix associated with vector T .

has the following sparse form

$$E = \begin{bmatrix} 0 & -t_z & 0 \\ t_z c\theta + t_1 s\theta & 0 & t_z s\theta - t_1 c\theta \\ 0 & t_x & 0 \end{bmatrix} = \begin{bmatrix} 0 & e_1 & 0 \\ e_2 & 0 & e_3 \\ 0 & e_4 & 0 \end{bmatrix} \quad (9)$$

where $s\theta(c\theta)$ denote $\sin\theta(\cos\theta)$ respectively. Given at least four point correspondences, the elements of the essential matrix $[e_1, e_2, e_3, e_4]^T$ can be obtained as a least squares solution of a system of homogeneous equations of the form (8). Once the essential matrix E has been recovered, the four different solutions for θ and $T = \pm[t_x, 0, t_z]$ can be obtained (using basic trigonometry) directly from the parametrization of the essential matrix (9). The physically correct solution is then obtained using the positive depth constraint.

In the context of robotic applications, it is often of interest to be able to estimate the relative camera pose in the absence of intrinsic camera parameters. Even in the case when the camera is calibrated ahead of time, the focal length parameter is the one most likely to be changed, especially in the context of active vision systems. This is the case we examine next. In the case of unknown focal length the two views are related by so called fundamental matrix F

$$\tilde{\mathbf{x}}_2^T F \tilde{\mathbf{x}}_1 = 0 \text{ with } \mathbf{x} = K^{-1}\tilde{\mathbf{x}}. \quad (10)$$

The fundamental matrix F is in this special planar, partially calibrated case related to the essential matrix E as follows

$$F = K^{-T} E K^{-1} \text{ with } K = \begin{bmatrix} f & 0 & 0 \\ 0 & f & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (11)$$

where f is the unknown focal length. The remaining intrinsic parameters are assumed to be known. In the planar motion case the matrix $F = [0, f_1, 0; f_2, 0, f_3; 0, f_4, 0]$ can be recovered from the homogeneous constraints of the form (10) given a minimum of four matched points. The extraction of the unknown motion parameters and the focal length f however is not straightforward, since the translation and the focal length appear in the parametrization of the matrix F in a multiplicative way. We propose to use additional constraints provided by the so called Kruppa's equations (23). It can be easily verified that a fundamental matrix F between the two views and the unknown intrinsic parameter matrix K satisfy the following constraint

$$F K K^T F^T = \lambda^2 \hat{e} K K^T \hat{e}^T \quad (12)$$

where $e = \frac{KT}{\|KT\|}$ is the epipole and λ is the unknown scale of the fundamental matrix. In our previous work (23) we have shown that for the special case of

planar motion the above equation is satisfied if and only if $\lambda = 1$. Since F and $e = [-f_1, 0, f_4]^T$ can be estimated, the renormalized equation (12) yields following useful constraint on intrinsic parameters K

$$F K K^T F^T = \hat{e} K K^T \hat{e}^T. \quad (13)$$

Given the planar motion case, the middle entries of matrices on the left and right side of equation (13) yield a constraint on the focal length and the entries of the fundamental matrix

$$f_2^2 f^2 + f_3^2 = f_4^2 f^2 + f_1^2.$$

The solution for the focal length can then be directly obtained from the above equation as

$$f = \sqrt{\frac{f_1^2 - f_3^2}{f_2^2 - f_4^2}}. \quad (14)$$

Once f is computed, the relative displacement between the views can be obtained by the method outlined for the calibrated case. Additional care has to be taken in assuring that the detected matches do not come from a degenerate configuration. We have used RANSAC algorithm for the robust estimation of the pose between two views, with slightly modified sampling strategy. Figure 7 shows two examples of relative positioning with respect to two different representative views. The initial estimate of the motion and focal length is further refined by nonlinear minimization, where the total reprojection error of all the matched points is minimized


$$E(R, T, f) = \min \sum_{i=1}^n \|\mathbf{x}^i - \pi([KR, KT]\mathbf{X}^i)\|^2 + \|\mathbf{x}_r^i - \pi(\mathbf{X}^i)\|^2, \quad (15)$$

where K is partially known matrix of intrinsic parameters, \mathbf{x}^i and \mathbf{x}_r^i are the matched SIFT features between the current view and the most likely reference view and \mathbf{X}_i are 3D coordinates of points expressed with respect to the reference view. Note that 3D structure of the scene is estimated as well. The focal length estimates obtained for these examples are $f = 624.33$ and $f = 545.30$. The relative camera pose for individual views is represented in the figure by a coordinate frame.


7 Conclusions and Future Work

We have demonstrated the suitability and the discrimination capability of the scale-invariant SIFT features in the context of location recognition and global localization task. Although the matching and location recognition methods can be accomplished using an efficient and simple voting scheme, the recognition rate is affected by dynamic changes in the environment and inherent ambiguities in the appearance of individual locations. We have shown that these difficulties can be partially resolved by exploiting the neighborhood relationships between the locations captured by Hidden Markov Models.

Relative positioning in Location 1



Relative positioning in Location 2



Location 1



Location 2

Fig. 7. Relative positioning experiments with respect to the representative views. Bottom: Query views along the path between the first view and the representative view for two different locations. Top: Recovered motions for two locations.

Since the notion of location is not defined precisely and is merely inferred in the learning stage the presented method enables only qualitative global localization in terms of individual locations. Following the global localization we compute the relative pose of the robot with respect to the closest reference view (24) found in the matching stage. This enables us to achieve metric localization with respect to the reference view, which can be followed by relative positioning tasks. More extensive experiments as well as integration with the exploration and navigation strategies on-board of mobile robot platform are currently underway.

8 Acknowledgements

The authors would like to thank D. Lowe for making available the code for detection of SIFT features. This work is supported by NSF grant IIS-0118732.

References

- [1] Briggs, A., Scharstein, D., Abbott, S.: Reliable mobile robot navigation from unreliable visual cues. In: Fourth International Workshop on Algorithmic Foundations of Robotics, New Hampshire. (2000)
- [2] Taylor, C.J., Kriegman, D.: Vision-based motion planning and exploration algorithms for mobile robots. *IEEE Transaction on Robotics and Automation* **14** (1998) 417–427
- [3] Sims, R., Dudek, G.: Learning environmental features for pose estimation. *Image and Vision Computing* **19** (2001) 733–739
- [4] Lowe, D.: Object recognition from local scale invariant features. In: International Conference on Computer Vision. (1999) 1150–1157
- [5] Wolf, J., Burgard, W., Burkhardt, H.: Robust vision based localization for mobile robots using image-based retrieval system based on invariant features. In: IEEE International Conference on Robotics and Automation. (2002)
- [6] Torralba, A., Murphy, K., Freeman, W., Rubin, M.: Context-based vision system for place and object recognition. In: International Conference on Computer Vision. (2003) 273 – 282
- [7] Schiele, B., Crowley, J.L.: Object recognition using multidimensional receptive field histograms. *International Journal of Computer Vision* **36** (2000) 31 – 50
- [8] Ulrich, I., Nourbakhsh, I.: Appearance-based place recognition for topological localization. In: IEEE International Conference on Robotics and Automation, San Francisco, CA. (2000) 1023–1029
- [9] Košecká, J., Zhou, L., Barber, P., Duric, Z.: Qualitative image based localization in indoors environments. In: IEEE Proceedings of CVPR. Volume 2. (2003) 3–8
- [10] Artac, M., Jogan, M., Leonardis, A.: Mobile robot localization using an incremental eigenspace model. In: IEEE Conference on Robotics and Automation. (2002) 1025 – 1030
- [11] Gaspar, J., Winters, N., Santos-Victor, J.: Vision-based navigation and environmental representations with an omnidirectional camera. *IEEE Transactions on Robotics and Automation* (2000) 777–789
- [12] Davidson, A., Murray, D.: Simultaneous localization and map building using active vision. *IEEE Transactions on PAMI* **24** (2002) 865–880
- [13] Se, S., Lowe, D., Little, J.: Mobile robot localization and mapping with uncertainty using scale invariant visual landmarks. *International Journal of Robotics Research* (2002) 735–758
- [14] P. Rybski, S. Roumeliotis, M.G., Papanikopoulous, N.: Appearance-based minimalistic metric SLAM. In: Intl. Conference on Intelligent Robots and Systems IEEE/RSJ. (2003) 194–199
- [15] Kuipers, B., Byun, Y.T.: A robot exploration and mapping strategy based on a semantic hierarch of spatial representations. *Journal of Robotics and Autonomous Systems* (1991) 47–63
- [16] Tomatis, N., Nourbakhsh, I., Siegwart, R.: Hybrid simultaneous localization

- and map building: a natural integration of topological and metric. *Robotics and Autonomous Systems* (2003) 3–14
- [17] Lowe, D.: Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* **60** (2004) 91–110
- [18] Brown, M., Lowe, D.: Invariant features from interest point groups. In: *Proceedings of BMVC, Cardiff, Wales*. (2002) 656–665
- [19] Li, F., Košecká, J.: Probabilistic location recognition using reduced feature set. Technical Report GMU-TR-2005-3, George Mason University (2005)
- [20] Pope, A., Lowe, D.: Probabilistic models of appearance for 3-d object recognition. *International Journal of Computer Vision* **40** (2000) 149 – 167
- [21] Schmid, C.: A structured probabilistic model for recognition. In: *Proceedings of CVPR, Kauai, Hawaii*. (1999) 485–490
- [22] Sturm, P.: On focal length calibration from two views. In: *Proceedings of CVPR, Kauai*. Volume 2. (2001) 145–150
- [23] Ma, Y., Vidal, R., Košecká, J., Sastry, S.: Kruppa’s equation revisited: Degeneracy and renormalization. In: *ECCV*. (2000) 561–577
- [24] Xang, X., Košecká, J.: Experiments in location recognition using scale-invariant SIFT features. Technical Report GMU-TR-2004-2, George Mason University (2004)