Wedge-based Corner Model for Widely Separated Views Matching

Robert Laganière and Étienne Vincent

School of Information Technology and Engineering
University of Ottawa, Ottawa, Canada, K1N 6N5
laganier,evincent,@site.uottawa.ca

Abstract

A novel interest point detector is presented in this paper. It uses a wedge model that characterizes corners by their orientation and angular width. The detector is compared to two popular corner detectors: Harris and SUSAN. It is also shown how widely separated views can be matched, by using the information provided by the detector to approximate local affine transformations between corners.

1. Introduction

The robust detection of interest points constitutes a fundamental step in the characterization and matching of images. Points of interest usually correspond to patterns of significant intensity variation in more than one direction. Many detectors are found in the literature. To evaluate their performance, different criteria may be considered:

1) *Accuracy* in localization, or the ability to consistently detect a given image pattern, at the exact same location, in spite of minor variability in the pattern appearance.

2) *Robustness*, or insensitivity to noise. Detection on noisy images can produce false positives, miss some feature points, or not localize them properly.

3) *Sensitivity*, that is the ability to detect points of interest in low contrast conditions. Most often, some parameters determines a tradeoff between sensitivity and robustness.

4) *Stability*, or the ability to continue to detect the same features under different geometrical transformations (especially perspective deformations), and under different conditions of illumination. A good measure of stability is the repeatability rate (see Subsection 4.2).

5) *Controllability*, or the number, and sensitivity of the parameters controlling the behavior of the detector. The effect of each parameter should be specific and predictable enough to allow an easy tuning of the detector.

6) *Richness* of the information provided. When a detector returns various characteristics of interest points, and not just a strength measure, the additional information can be exploited in the task to follow, such as the normalization of patterns to be matched (as in Section 5).

7) *Variability* in the characteristics of detected points. This ensures that several points of interest are detected, regardless of the nature of the image under analysis, and that the points are distinguishable from each other.

This paper presents a novel approach to interest point detection, which will be evaluated using the quality attributes presented above. In addition, it is shown how the information provided by this detector can improve the reliability of sparse matching, in the context of widely separated views.

2. Interest Point Detection

Many widely used feature point detectors use image derivatives to identify high curvature points. The Harris operator [3] which uses the autocorrelation matrix of the first derivatives, is such a detector. The eigenvalues of this matrix correspond to the gradient magnitudes in the two principal directions. Points of interest are then the ones where both these magnitudes are high.

In other approaches, detection is based on the direct comparison of intensity values inside small predefined windows. The SUSAN operator [2] fits in this category. It compares the central pixel with its neighbors to define a univalue segment assimilating nucleus (USAN), i.e. points inside a circular region having similar brightness. The size and the symmetrical axis of this nucleus then determine the presence of a feature point.

Finally, some detectors find image elements matching a predefined model [4, 5]. Such models represent ideal corners and detection consists in fitting control parameters to the underlying image intensity pattern.

3. The Proposed Detector

Our proposed corner detector combines the two later approaches above, by detecting corner shapes from intensity patterns. The goal was to devise a stable detector, suitable for feature matching across different views. This detector relies on a simple model of corners, consisting of a wedge (the corner), having its origin on the center of a circular neighborhood (the background) as illustrated in Fig. 1. This idealized corner can be described by two parameters: an angular position (θ), and an angular width (φ). This model is similar to the one used in [4].

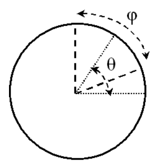

Figure 1. The corner model.

To identify corners, each pixel location should be examined, comparing its surrounding circular neighborhood with the ideal corner model. We propose to proceed as follow:
1) The circular window's intensity mean and variance are computed. The variance must be over a predefined threshold, to limit the sensitivity of the corner finder.
2) The circular area around the potential corner is segmented into background and foreground, as described in subsection 3.1.
3) The extracted foreground is fitted to the ideal corner model, by finding the best approximations for φ and θ (see subsection 3.2). The corner strength is obtained by comparing the segmented area to the parameterized corner model.

Applying this procedure to all pixels results in a corner map where each interest point is given a corner strength value, and two delimiting angles. The final set of corners is obtained by imposing a threshold on the corner strength values. This should be preceded by non-maxima suppression, to avoid corner clusters.

3.1. Segmenting a Point's Neighborhood

To obtain simple foreground/background segmentations of circular areas around potential corners, the surrounding pixels are classified as having an intensity above or below the mean intensity in that neighborhood. To improve stability, this is done using a sigmoidal function, rather then a hard threshold. The group with the fewest pixels becomes the foreground, and the other group forms the background. The center pixel must belong to the foreground, in order to be further considered.

Similarly, the SUSAN corner detector relies on a simple segmentation. Corners are then identified from the properties of the segmented areas, but without comparing them to a preestablished corner model.

3.2. Fitting the Model

The similarity of segmented areas to ideal wedge models (W_θ^φ), is measured as the sum of their absolute differences:

$$c_s(x,y) = \sum_{(i,j)\in\mathcal{C}_{xy}} |W_\theta^\varphi(i,j) - sig(I(x+i,y+j) - \bar{I}(x,y))| \quad (1)$$

Where $\mathcal{C}_{xy}$ is the circular window around the pixel under consideration.

Some θ and φ approaching a minimum for this pseudo Hamming distance can be found using the following scheme:
1) The circular area around the potential corner is subdivided into small wedges,$W_{n\Delta\theta}^{\varphi_{min}}$ having a width of φ_{min} and rotated around the circular area by increments of $\Delta\theta$.
2) Corner coverage is computed for each of these wedges, as the sum of the segmented value over all its pixels:

$$c_c(x,y) = \sum_{(i,j)\in W_{n\Delta\theta}^{\varphi_{min}}} sig(I(x+i,y+j) - \bar{I}(x,y)) \quad (2)$$

A wedge will be considered to potentially lie on the corner if its corner coverage is greater than a predetermined threshold value (c_{min}). So far, the process is similar to the rotated wedge filtering presented in [11], where it is used to compute intensity mean values, from which one-dimensional angular derivative are computed.
3)The wedge with the highest coverage is selected, and considered to belong to the foreground of the corner.
4)From this initial wedge, all adjacent wedges with a coverage greater than c_{min} are retained. The corner model is the one formed by the union of all these selected elementary wedges. The angle spanned by this set of adjacent wedges must belong to a range $\varphi_{min} < \varphi < \varphi_{max}$.

The union of all the wedges selected in the above procedure determines W_θ^φ (for equation (1)), φ being the angle spanned by the union, and θ being its bisector.

4. Experimental Comparisons

To test the proposed corner detector, experiments that compare it to SUSAN, Harris were conducted.

4.1. Evaluating Corner Location

The first set of experiments measures the *robustness* and *accuracy* of each detector. A synthetic image was used, where the position of each corner (the vertices of geometrical shapes) is known.

Figure 2 shows the feature points found by three different detectors on this simple test image. Note the variability of contrast conditions at different corner locations, caused by

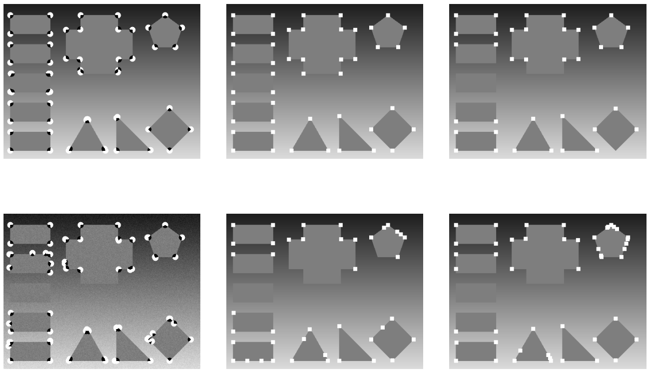

Figure 2. Corner detected using (from left to right), our detector, SUSAN, Harris. The second row shows detection when Gaussian noise of variance 50 was added.

the smooth gray level transition of the background. This allows the appreciation of the respective *sensitivity* of the operators, which have been tuned to obtain the best possible results.

For each vertex in this image, the closest detected corner was found, and the distance between the two points was measured. These measurements were taken for all three detectors and for noisy versions of the same image. Results are presented in Figure 3, where the graphs show, for different noise levels, the number of points within 1,2,3 or 4 pixels of a detected corner. In Figure 2, the behavior of the detectors is also shown, in the case where a gaussian noise of variance 50 was added to the original synthetic image. These experiments show SUSAN to be a more accurate detector. However, as the level of noise is increased, our wedge-based detector leads to superior results.

4.2. Evaluating Repeatability

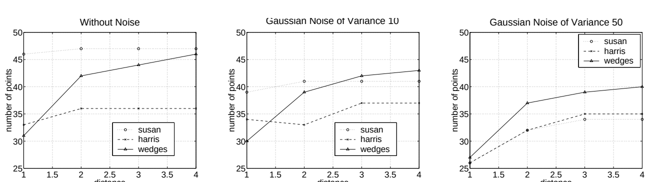

Figure 3. Evaluating the accuracy of different corner detectors.

The second set of experiments measures the *stability* of the different detectors, which is the most important attribute in many applications. This is done by evaluating the repeatability rate [1], i.e. the number of features that are consistently detected in two views of a scene. To do so, we must be able to determine, for each point in an image, where its corresponding point is located in another image.

It is well known that images of a planar surface are related by a homographic transformation. This also holds for images where the difference in viewpoint corresponds to a pure rotation. Two image pairs, corresponding to these two cases, are shown in Figure 4. The homographies between these images were computed, it then becomes possible to verify, for each detected interest point in one image, wether its corresponding point is also detected in the other image.

Figure 4 also shows the computed repeatability rates for the corner detectors. These repeatability rates were computed for different corner acceptance thresholds. In this case, SUSAN clearly demonstrates the worst performances. The two other operators presents a similar behavior.

From this, and the previous experiment, it is seen that the proposed detector offers a good compromise between the *accuracy* of SUSAN, and the *stability* of Harris.

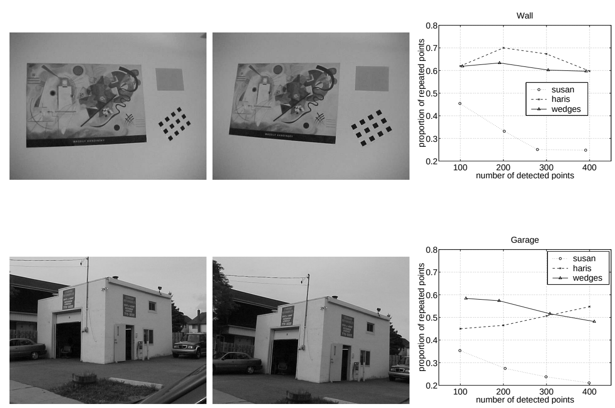

Figure 4. Two image pairs, with graphs of their corresponding repeatability rates

5. Matching Points of Interest

Matching is an important task in computer vision. Feature point detection is often used to select points for sparse matching between images. When the change in viewpoint is small, robust matching can be accomplished through normalized correlation [6]. However, if views are more widely separated, the simple correlation of windows around interest points is not a meaningful similarity measure, as image patterns around the points can undergo significant deformation. Thus, a matching scheme which is invariant to perspective deformation induced by changes in viewpoint is required. The information about corner shapes provided by our detector will achieve this.

The solution consists in applying a local transformation between pairs of image patches, before evaluating their similarity, as in [8, 7]. A good approximation of perspective deformation consists in planar affine transformations [9]. These can be represented with 6 DOF, 3×3 non-singular matrices that transforms the (x, y) coordinates of an image patch as follow:

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & t_1 \\ a_{21} & a_{22} & t_2 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \qquad (3)$$

The patches that are to be matched consist in circular windows around the points. The translation between them are known, this leaves 4 DOF, which are estimated from the parameters of the wedge corner models. First, it is assumed that the difference in scale can be neglected. Then, the rotation angle is given by the difference in angular position between the two detected wedges, θ and θ'. Finally, the non-isotropic component of the affinity is assumed to apply in the bisector direction of the rotated wedge, with a ratio given by the angular widths of the two wedges, $\varphi : \varphi'$. Matching is then accomplished by correlating the intensity values found in the vicinity of the corner in the first view, and the corresponding intensity values in the second view, as given by the affine transformation. All matches with correlation scores higher than a some threshold are kept.

This approach still produces many false matches. However, it will produce far fewer than a direct correlation scheme, in the case of widely separated views. A robust estimator of two-view geometric transformations can filter out these false matches [10].

Figure 5 shows a pair of images on which our simple matching procedure was applied. The resulting candidate match set contained 300 point pairs, and was used to estimate, through a RANSAC scheme, the homography between the two views. A mosaic was constructed using the found homography. In contrast, when 300 matches were obtained by direct normalized correlation, no valid homography could be obtained, no matter how many RANSAC iterations were used.

Figure 5. Homography estimation (the Mosaic was generated with the help of Rafy Terzian and Stéphanie Deguire).

References

[1] C. Schmid, R. Mohr, C. Bauckhage, Comparing and Evaluating Interest Points, *ICCV*, pp. 230-235, 1998.

[2] S.M. Smith, J.M. Brady, A New Approach to Low Level Image Processing *Tech. Rep.*, Oxford U., 1995.

[3] C. Harris, M. Stephens, A Combined Corner Detector, *Alvey Vision Conf.*, pp. 147-151, 1988.

[4] K. Rohr, Modelling and Identification of characteristic intensity variations, *Image Vision Computing*, 10, pp. 66-76, 1992.

[5] T. Blaszka, R. Deriche, Recovering and Characterizing Image Features using an Efficient Model Based Approach, *Tech. Rep. RR-2422*, INRIA, 1994.

[6] P. Torr, A. Zisserman, Robust Computation and Parameterization of Multiple View Relations, *ICCV*, pp. 727-732, 1998.

[7] P. Pritchett, A. Zisserman, Wide Baseline Stereo Matching, *ICCV*, pp. 754-760, 1998.

[8] I. Zoghalmi, O. Faugeras, R. Deriche, Using Geometric Corners to Build a 2D Mosaic from a set of Images, *CVPR*, pp. 420-425, 1999.

[9] A. Baumberg, Reliable Feature Matching Across Widely Separated Views, *CVPR*, pp. 774-781, 2000.

[10] G. Roth, A. Whitehead, Using Projective vision to find Camera Positions in an Image Sequence, *Vision Interface*, pp. 225-232, 2000.

[11] W. Yu, K. Daniilidis, G. Sommer, Rotated Wedge Averaging Method for Junction Characterization, *CVPR*, pp. 390-395, 1998.