IMPROVING JUNCTION DETECTION BY SEMANTIC INTERPRETATION

Published at VISAPP'07

Sinan Kalkan
Bernstein Centre for Computational Neuroscience, Univ. of Göttingen, Germany
sinan@bccn-goettingen.de

Shi Yan, Florian Pilz
Medialogy Lab, Aalborg Univ. Copenhagen, Denmark
{syan06, fpi}@imi.aau.dk

Norbert Krüger
Cognitive Vision Group, Univ. of Southern Denmark, Denmark
norbert@mip.sdu.dk

Keywords: Junction Detection, Junction Positioning, Junction Representation

Abstract: Every junction detector has a set of thresholds to make decisions about the junctionness of image points. Low-contrast junctions may pass such thresholds and may not be detected. Lowering the thresholds to find such junctions will lead to spurious junction detections at other image points. In this paper, we implement a junction-regularity measure to improve localization of junctions, and we develop a method to create semantic interpretations of arbitrary junction configurations at improved junction positions. We propose to utilize such a semantic interpretation as a *feedback mechanism* to filter false-positive junctions. We show results of our proposals on natural images using Harris and SUSAN operators as well as a continuous concept of intrinsic dimensionality.

1 INTRODUCTION

Junctions are utilized in computer vision and image processing for tasks that especially require finding correspondences between different views of the same scene, mainly due to their *distinctiveness*, *seldomness* and *stability*.

Correct localization of junctions[1] is crucial because even small errors in localization lead to wrong interpretations of the scene (Rohr, 1992). Nevertheless, it is shown in (Deriche and Giraudon, 1993; Rohr, 1992) that energy-based junction detection methods smooth out junctions and face the problem of wrong localization.

Junctions also have the property of being *interpretable*: *i.e.*, you can construct a meaningful interpretation about how the junction is formed, as proposed in (Parida et al., 1998; Rohr, 1992). Such a semantic interpretation (*SI*) can be utilized in rigid body motion estimation, depth estimation, feature matching etc. and should be more robust than a single junctionness measure in identification of junctions and in correspondence finding.

Junction detectors, no matter what the underlying methods are, have to make a decision about the junctionness of image points. The decision is made by a set of automatically or manually set thresholds (on a set of measures) that determine the *sensitivity* of the algorithm to contrast (in most of the cases, a high threshold means low sensitivity and vice versa). On the other hand, a method that utilizes a junction detector requires the detector to be *complete*: *i.e.*, the detector should be able to detect all the junctions that represent the image.

The relation between sensitivity and completeness presumably looks like as plotted in figure 1(a). Increasing the sensitivity increases not only the completeness of a detector[2] but also increases the amount of false-positives, or 'spuriousness', of the detector as illustrated in figure 1(b). These observations suggest that spuriousness and completeness are two competing objectives that make the problem of junction detection a multi-objective optimization problem, and it is known that a multi-objective optimization problem with competing objectives does not have a global optimum, but a set of optimal solutions which are called

[1] In this paper, corners are considered to be a special case of junctions, and the term 'corner' is avoided.

[2] Exact shape of this relation might be different in real world; however, the authors claim that completeness should be still an increasing function of sensitivity in any case.

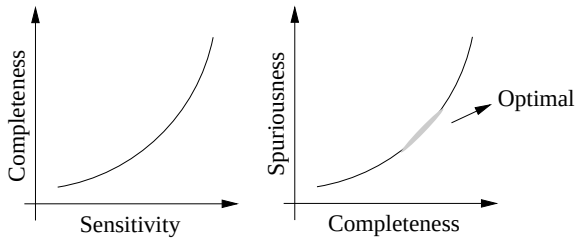

(a) (b)

Figure 1: The relations between sensitivity, completeness and spuriousness.

Pareto-optimal (see, figure 1(b) and *e.g.*, (Coello, 1999)).

Junction detection algorithms face this 'problem'[3] because detecting junctions in real images is an ill-posed problem at the level of feature-processing due to the fact that identifying *accurate* and *complete* boundaries and junctions of the objects requires an object-recognition step which is supposed to happen at a higher level in a vision system.

In this paper, we show that junction detectors can be followed by a 'semantic interpretation' step as a feedback mechanism to achieve a better completeness-to-spuriousness ratio. We achieve this by (1) increasing the sensitivity of the junction detectors (by decreasing their thresholds), (2) improving the positioning of the detection step using a regularity or intersection consistency step and then (3) extracting the semantic interpretation of the junctions to filter spurious junctions.

The intersection-consistency, or regularity measure that is implemented in this paper is based on the observation that the position of a junction is defined by the intersection of its edges (Parida et al., 1998). Our measure is similar to the $\mathcal{R}()$ function in equation 5 of (Parida et al., 1998) and the regularity function $S()$ in equation 4 of (Förstner, 1994). Both of these functions are based on the local image gradient whereas our method utilizes another edgeness measure called intrinsic dimensionality (see section 2.1 and (Krüger and Felsberg, 2003) for details).

For semantic interpretation of junctions, we propose representing junctions in terms of their constituents (*i.e.*, the edges that form the junctions) and how they form the junctions (*i.e.*, the directions of the constituent edges). There have already been studies related to the representation of the junctions (see, *e.g.*, (Simoncelli and Farid, 1996; Hahn and Krüger, 2000; Parida et al., 1998; Baker et al., 1998; Rohr, 1992)): In (Simoncelli and Farid, 1996), steerable wedge filters are developed for analyzing the orientation maps of edges and junctions, without creating an explicit representation of these features; in (Parida et al., 1998), by assuming that the number of junctions is known, a junction model is fitted to the data by minimizing an energy function; in (Baker et al., 1998), parameters of junctions with just two edges (*i.e.*, corners) are extracted by using dimensionality reduction techniques. In (Rohr, 1992), assuming the number of edges is known, a junction is extracted as a composition of L-junctions by fitting a parametric model to the image data. In (Hahn and Krüger, 2000), corners are detected, and their representations are created using Hough lines, and these corners are merged to create junction representations. In current paper, we employ a simple method that extracts the representation of a junction by analyzing the clusters in its orientation histogram. While doing so, *our method does not make any assumption* about the junction and is able to create representations of any junction configuration.

The contributions of this paper are (1) proposal of a new method for creating representations of junctions and (2) pinpointing a common problem in all junction detectors (namely, CS-problem) and (3) proposing a way to improve junction detectors with respect to this problem. We test our improvements on natural images, using three different junction detectors: SUSAN, Harris operators and the intrinsic dimensionality. The aim of this paper is not to compare the performance of these methods but propose a feedback mechanism to improve them. For a comparison of a set of interest point and junction detectors, the interested reader is directed to (Schmid et al., 2000).

Biological vision systems are claimed to be equipped with feedback mechanisms for disambiguation at several steps, especially in early cognitive vision (see, *e.g.*, (Bayerl and Neumann, 2004)). Current work is considered to be such a feedback mechanism for better detection of junctions.

This paper is organized as follows: In section 2, the main junction detection approaches and the continuous definition of intrinsic dimensionality are introduced. In sections 3 and 4, the new methods for improving the positioning and the semantic interpretation of junctions are proposed, respectively. In section 5, we present and discuss our results, concluding the paper in section 6.

[3]In the rest of the paper, CS-problem will denote the completeness-spuriousness problem.

2 JUNCTION DETECTION ALGORITHMS

In this section, we briefly describe the main approaches for junction detection without any claim of being complete (see, *e.g.*, (Schmid et al., 2000; Deriche and Giraudon, 1993; Smith, 1997) for more detailed reviews) and give a short presentation of intrinsic dimensionality (see (Harris and Stephens, 1988) and (Smith and Brady, 1997) for information about Harris and SUSAN operators, respectively).

Since the first attempts around late 1970s, there have been quite a number of works on the detection of junctions. The methods can be roughly divided into three main categories:

- **Contour-based:** These methods involve extracting an edge representation and then processing the maxima curvature or the linking of the edges to find the junctions (see, *e.g.*, (Asada and Brady, 1986; Deriche and Giraudon, 1990; Horaud and Veillon, 1990)).
- **Signal-based:** These methods involve finding the junctions by directly using the image intensities. The second order derivatives of intensities, or the Hessian matrix (Beaudet, 1978; Dreschler and Nagel, 1982), autocorrelation function of the image patch (Moravec, 1980; Harris and Stephens, 1988; Förstner, 1994) are the main tools used by such approaches.
- **Template-based:** These methods detect junctions that match certain templates (Rohr, 1992; Parida et al., 1998).

2.1 Continuous Concept of Intrinsic Dimensionality (*iD*)

In image processing, the *iD* was introduced by (Zetzsche and Barth, 1990) and was used to formalize a *discrete distinction* between edge-like and junction-like structures. This corresponds to a classical interpretation of local image structures in computer vision.

Homogeneous, edge-like and junction-like structures are respectively classified by *iD* as *intrinsically zero dimensional (i0D)*, *intrinsically one dimensional (i1D)* and *intrinsically two dimensional (i2D)*.

When looking at the spectral representation of a local image patch (see figure 2(a,b)), we see that the energy of an i0D signal is concentrated in the origin (figure 2(b)-top), the energy of an i1D signal is concentrated along a line (figure 2(b)-middle) while the energy of an i2D signal varies in more than one dimension (figure 2(b)-bottom).

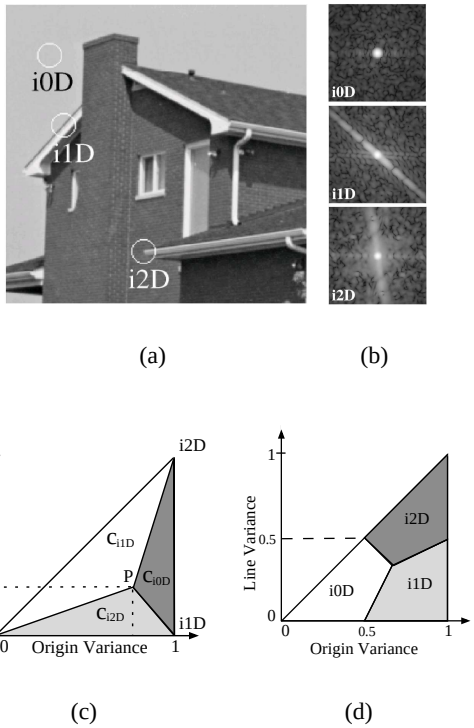

Figure 2: Illustration of *iD* (Sub-figures (a,b,c) taken from (Felsberg and Krüger, 2003)). (a) Three image patches for three different intrinsic dimensions. (b) The 2D spatial frequency spectra of the local patches in (a), from top to bottom: i0D, i1D, i2D. (c) The topology of *iD*. Origin variance is variance from a point, i.e., the origin. Line variance is variance from a line, measuring the junctionness of the signal. c_{iND} for $N = 0, 1, 2$ stands for confidence for being i0D, i1D and i2D, respectively. Confidences for an arbitrary point P is shown in the figure which reflect the areas of the sub-triangles defined by P and the corners of the triangle. (d) The decision areas for local image structures.

Recently, it has been shown (Krüger and Felsberg, 2003; Felsberg and Krüger, 2003) that the structure of the *iD* can be understood as a triangle that is spanned by two measures: origin variance and line variance. Origin variance describes the deviation of the energy from a concentration at the origin while line variance describes the deviation from a line structure (see figure 2(b) and 2(c)); in other words, origin variance measures non-homogeneity of the signal whereas the line variance measures the junctionness. The corners of the triangle then correspond to the 'ideal' cases of *iD*. The surface of the triangle corresponds to signals that carry aspects of the three 'ideal' cases, and the distance from the corners of the triangle indicates the similarity (or dissimilarity) to *ideal* i0D, i1D and i2D signals.

As shown in (Krüger and Felsberg, 2003; Felsberg and Krüger, 2003), this triangular interpretation allows for a *continuous formulation* of *iD* in terms of 3 confidences assigned to each discrete case. This is achieved by first computing two measurements of origin and line variance which define a point in the triangle (see figure 2(c)). The bary-centric coordinates (see, e.g., (Coxeter, 1969)) of this point in the triangle directly lead to a definition of three confidences that add up to one. These three confidences reflect the volume of the areas of the three sub-triangles which are defined by the point in the triangle and the corners of the triangle (see figure 2(c)). For example, for an arbitrary point P in the triangle, the area of the sub-triangle i0D-*P*-i1D denotes the confidence for i2D as shown in figure 2(c). That leads to the decision areas for i0D, i1D and i2D as seen in figure 2(d).

For the example image in figure 2, computed *iD* is shown in figure 3.

3 IMPROVING LOCALIZATION

Our approach is to detect junctions (using Harris, Susan or *iD*), and then to compute a junction regularity measure, called intersection-consistency (*IC*), in the neighborhood of the detected junctions. The new *improved* position of a junction is determined by the local maximum of *IC* in the 3x3-neighborhood.

IC is measured by checking whether the pixels in the image patch point towards the center of the patch or not. Pointing towards the center of the local image patch is measured by the distance between the center $\mathbf{p}_c$ and the line going through the pixel. The line is defined according to the position of the pixel $\mathbf{p}$ and the computed orientation information $\theta_{\mathbf{p}}$. The weighted average of these distances then defines the intersection consistency at $\mathbf{p}_c$:

$$ic(\mathbf{p}_c) = \int [c_{i1D}(\mathbf{p})]^2 [1 - d(l^{\mathbf{p}}, \mathbf{p}_c)/d(\mathbf{p}, \mathbf{p}_c)] d\mathbf{p}, \quad (1)$$

where $\mathbf{p}$ is the index of the pixels in image patch P; $c_{i1D}(\mathbf{p})$ is the confidence for i1D of pixel $\mathbf{p}$; $l^{\mathbf{p}}$ is the line going through pixel $\mathbf{p}$ with a slope defined according to the orientation $\theta_{\mathbf{p}}$; $d(l^{\mathbf{p}}, \mathbf{p}_c)$ is the distance between $l^{\mathbf{p}}$ and $\mathbf{p}_c$; and, $d(\mathbf{p}, \mathbf{p}_c)$ is the distance between $\mathbf{p}$ and $\mathbf{p}_c$. Note that $d(l^{\mathbf{p}}, \mathbf{p}_c)$ is normalized with $d(\mathbf{p}, \mathbf{p}_c)$ because we would like to give equal weights to votes of every pixel whether they are close to the center or not.

The distances between the center of the local image patch and the lines through the pixels is weighted by $(c_{i1D})^2$ because the computed orientation information is defined only for edge-like structures, and *IC* by definition involves intersection consistency of edge-like structures.

$ic(P)$ value will be high (1) if the image patch has only one edge which goes through the center of the patch or (2) if the image patch has a junction whose intersection point is located at the center of the patch.

For comparison, $\mathcal{R}()$ function of (Parida et al., 1998) and $S()$ function of (Förstner, 1994) are given in equations 2 and 3 respectively:

$$\mathcal{R} = \int_0^{\infty} \int_0^{2\pi} (\frac{\delta I}{\delta r}) g^*(r) r dr d\theta, \quad (2)$$

where (r, θ) is the polar coordinate of points relative to the center of the image patch, and $g^*(r)$ is a modulating function and can be set to $rG_\sigma(r)$ where G_σ is the Gaussian.

$$S(p, \sigma) = \iint d^2(p,q) \parallel \nabla g(q) \parallel^2 G_\sigma(p-q) dq. \quad (3)$$

where p is the center of the image patch; q denotes image points in the image patch; $d(p,q)$ is the distance of center point p to the line defined by q; and, $\nabla g(q)$ is the intensity gradient (g_x, g_y).

4 SEMANTIC INTERPRETATION OF JUNCTIONS

In this section, we describe how we estimate the semantic interpretation (*SI*) of a junction. This estimation process does not make any assumptions on the configuration of the edges (interested reader is directed to (Waltz, 1975) for different classes of junctions and their properties). The *SI* of a junction that does not fall into any meaningful junction category can be used to find false-positives.

In our *SI*, a junction is represented by a set of rays $r_1 ... r_n$ corresponding to the set of n edges that intersect at the junction. Each ray r_i represents a specific edge i, defined as a half-line expanding from the intersection point in a certain direction $\tilde{\theta}_i$. We introduce another parameter c_i for the confidence of the edge which can be used as a weight when the *SI* of a junction is utilized. With these parameters, we can define the semantic interpretation $SI(\mathcal{J})$ of a junction $\mathcal{J}$ as follows:

$$SI(\mathcal{J}) = \{r_1, ..., r_n\} = \left\{ \left(c_1, \tilde{\theta}_1\right), ..., \left(c_n, \tilde{\theta}_n\right) \right\}. \quad (4)$$

$\tilde{\theta} \in [0, 2\pi)$ is the junction-relative orientation; the image relative orientation $\theta_p \in [0, \pi)$ of a pixel p at (x_p, y_p) is transformed to junction-relative orientation for junction $\mathcal{J}$ at (x,y) as follows:

$$\tilde{\theta}_p = \begin{cases} \theta_p, & \text{if } tan^{-1}[(x-x_p)/(y-y_p)] < \pi, \\ \theta_p + \pi, & \text{if } tan^{-1}[(x-x_p)/(y-y_p)] \geq \pi. \end{cases} \quad (5)$$

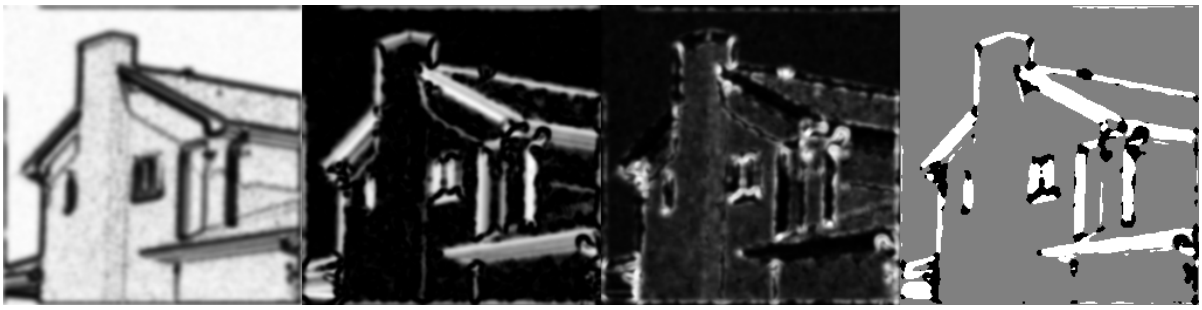

Figure 3: Computed *iD* for the image in figure 2, black means zero and white means one. From left to right: $c_{\text{i0D}}, c_{\text{i1D}}, c_{\text{i2D}}$ and highest confidence marked in gray, white and black for i0D, i1D and i2D, respectively.

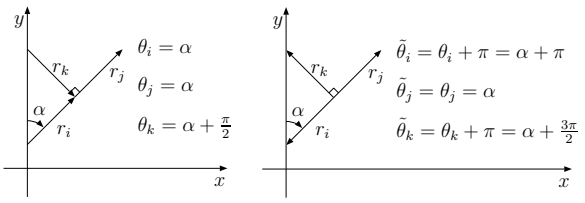

(a) (b)

Figure 4: Image- and junction-relative representations of the directions of the edges of an example junction. (a) Image-relative directions. (b) Junction-relative directions.

In figure 4, the image-relative and junction-relative orientations of two edges are shown for a junction.

The junction-relative orientation $\tilde{\theta}_i$ for each ray r_i is extracted by finding the dominant orientations in the neighborhood N of the junction $\mathcal{J}$. We can construct the set of pixels in N that point towards the center of the junction $\mathcal{J}$ as follows:

$$\tilde{\Theta}_c = \left\{ \tilde{\theta}_p \mid p \in N \text{ and } d(l_p, (x,y)) < T \right\}, \quad (6)$$

where l_p is the line defined by the pixel p with orientation $\tilde{\theta}_p$.

The number of rays and their orientations are determined by the clusters in the histogram $H_l(\tilde{\Theta}_c)$ where l is the index of the bins. The set of clusters $\{C_m\}$ is the set of H_l (1) where the first derivative $\delta H_l / \delta l$ changes sign, and (2) where the energy (*i.e.*, the number of elements in the bin) is above a threshold. Figure 5 shows the rays extracted from an example junction.

A junction is marked as false-positive if $n < 2$ or $n = 2$ and $\tilde{\theta}_1 \simeq \tilde{\theta}_2$.

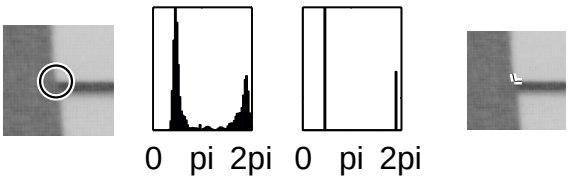

Figure 5: Illustration of the *SI* of a junction. From left to right: the junction marked with a circle; the distribution of junction-relative orientation; detected ray orientations; estimated *SI*.

Table 1: The parameters used in the experiments.

Algorithm	**Low sensitivity**	**High sensitivity**
SUSAN	$brightness > 20$	$brightness > 13$
Harris	$E > 1000$	$E > 300$
iD	$c_{i2D} > c_{i1D}$ & $c_{i2D} > c_{i0D}$	$c_{i2D} > 0.3$

5 RESULTS AND DISCUSSIONS

SUSAN implementation is taken from the author of (Smith and Brady, 1997), and Harris implementation is taken from (Noble, 1989). The parameters of SUSAN, Harris and *iD* are provided in table 1.

In figure 6, the results of the three junction detection methods are presented for several image patches extracted from real images. For each example and each method, original detection results, improved positions and the *SI* are plotted.

The examples demonstrate that junction detection methods face the problem of wrong localization as pointed out in (Deriche and Giraudon, 1993; Rohr, 1992). Moreover, it is very likely that the methods produce false positives especially visible in the case of SUSAN and *iD*. However, we show in figure 6 that the effect of the positioning problem can be decreased, and false positives can be removed by using the *SI* of the junctions.

As mentioned in section 1, junction detectors have

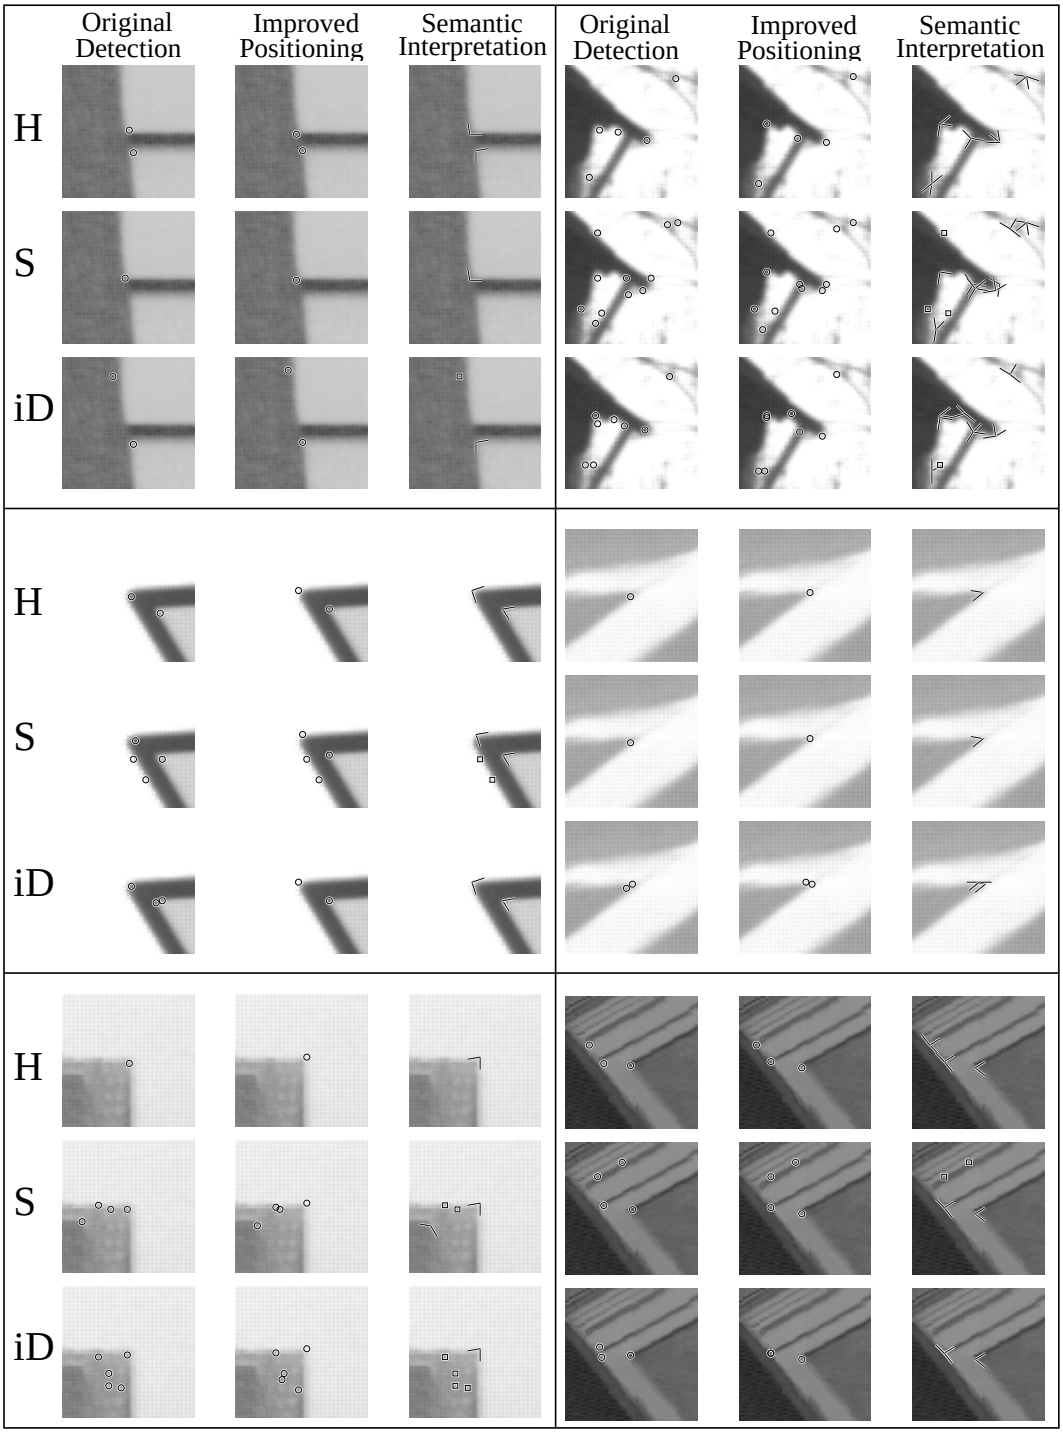

Figure 6: A set of example junctions and the results of junction detectors and the results of *IC* and *SI*on these results. For each example, S, H and iD denote SUSAN, Harris and *iD* respectively. In each example, the first column shows the original detections of the algorithms; the second column shows the effect of improved positioning via *IC*; and, the third column shows the estimated *SI* and how it can be used to get rid of spurious junctions. Spurious junctions that are estimated with *SI* are marked in small squares.

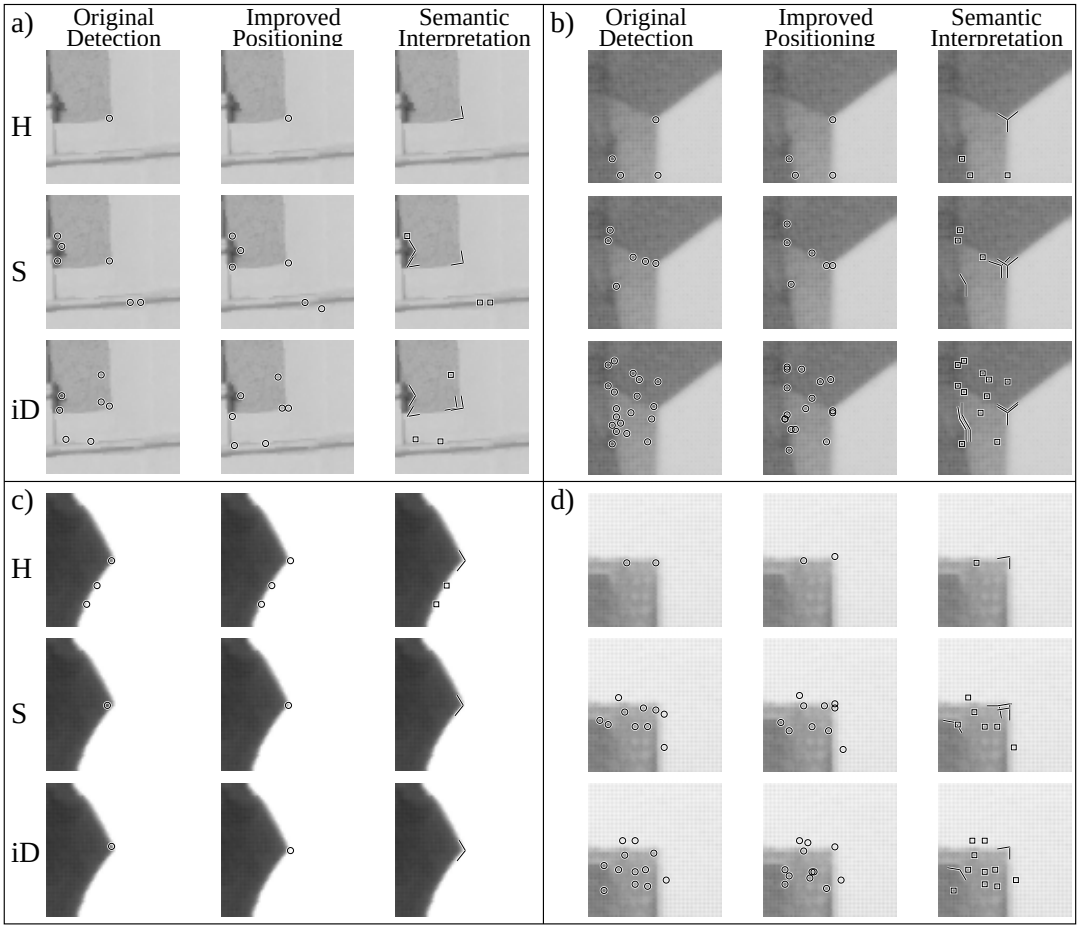

Figure 7: The effect of high 'sensitivity' on the performance of junction detectors. Junction detectors can now detect low contrast junctions that they miss with low sensitivity. H and S denote Harris and SUSAN respectively. For each subfigure, the first column shows original detection results, the second the results of improved positioning with *IC* and the third the *SI*. Spurious junctions that are estimated with *SI* are marked in small squares.

a level of sensitivity that cannot be made universal; *i.e.*, it is not possible to adjust the parameters of a junction detector in order to detect every junction without producing a big ratio of false positives.

In figure 7, we show a set of examples for SUSAN, Harris and *iD* with low thresholds (table 1). The thresholds have been decreased so that the detectors can detect the junctions that they have missed with their default parameter values (*e.g.*, the junction in the center of figure 7(a)). From 7(a), we see that increasing sensitivity of a detector can help in detecting low contrast however important junctions. On the other hand, figure 7 shows that all methods produce spurious results when the sensitivity is high, especially in the case of SUSAN and *iD*. However, by making use of *IC* and *SI*, it is possible to get rid of most of the spurious junctions and detect a wider range of junctions with more accuracy even for high sensitivity levels.

6 CONCLUSION

In this paper, we proposed two methods for improving the detection and the representation of junctions: (1) an operator called *intersection consistency* that measures how consistent a junction is with its neighborhood, and (2) a way to create semantic interpretation of junctions.

As shown in (Deriche and Giraudon, 1993; Rohr, 1992), energy-based junction detectors face the problem of wrong positioning. Using our intersection consistency, it is possible to achieve better positioning.

We also addressed the issue of having thresholds on the detection of junctions with respect to the 'completeness' and 'spuriousness' of the detection. By making use of our semantic interpretation as a feedback mechanism, we show that using semantic interpretation can improve detection performance of junc-

tion detectors.

7 ACKNOWLEDGEMENTS

This work is supported by the ECOVISION and Drivsco projects. We would like to thank Babette Dellen, Irene Markelić, Nicolas Pugeault and Florentin Wörgötter for their fruitful comments.

REFERENCES

Asada, H. and Brady, M. (1986). The curvature primal sketch. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(1):2–14.

Baker, S., Nayar, S. K., and Murase, H. (1998). Parametric feature detection. *Int. Journal of Computer Vision*, 27(1):27–50. March.

Bayerl, P. and Neumann, H. (2004). Unified mechanisms in the disambiguation and grouping of visual information in motion, stereo and monocular depth perception. *Perception Supplement*, 33:79.

Beaudet, P. (1978). Rotationally invariant image operators. In *Proc.4th Int. Joint Conf. Pattern Recognition*, pages 579–583. Kyoto, Japan.

Coello, C. A. C. (1999). A comprehensive survey of evolutionary-based multiobjective optimization techniques. *Knowledge and Information Systems*, 1(3):129–156.

Coxeter, H. (1969). *Introduction to Geometry (2nd ed.)*. Wiley & Sons.

Deriche, R. and Giraudon, G. (1990). Accurate corner detection: An analytical study. *ICCV 90, Osaka Japan*.

Deriche, R. and Giraudon, G. (1993). A computational approach for corner and vertex detection. *IJCV*, 10(2):101–124.

Dreschler, L. and Nagel, H. H. (1982). Volumetric model and 3d trajectory of a moving car derived from monocular tv frame sequences of a street scene. *Computer Graphics and Image Processing*, 20:199–228.

Felsberg, M. and Krüger, N. (2003). A probablistic definition of intrinsic dimensionality for images. *Pattern Recognition, 24th DAGM Symposium*.

Förstner, W. (1994). A framework for low level feature extraction. In *ECCV '94: Proceedings of the third European conference on Computer Vision (Vol. II)*, pages 383–394, Secaucus, NJ, USA. Springer-Verlag New York, Inc.

Hahn, M. and Krüger, N. (2000). Junction detection and semantic interpretation using hough lines. *International ICSC Symposium on Engineering of Intelligent Systems (EIS)*.

Harris, C. G. and Stephens, M. J. (1988). A combined corner and edge detector. In *Proc. Fourth Alvey Vision Conference, Manchester*, pages 147–151.

Horaud, R. and Veillon, F. (1990). Finding geometric and relational structures in an image. In *ECCV 90*, pages 374–384, New York, NY, USA. Springer-Verlag New York, Inc.

Krüger, N. and Felsberg, M. (2003). A continuous formulation of intrinsic dimension. *Proceedings of the British Machine Vision Conference*.

Moravec, H. (1980). Obstacle avoidance and navigation in the real world by a seeing robot rover. Technical Report CMU-RI-TR-3, Carnegie-Mellon University, Robotics Institute.

Noble, A. (1989). *Descriptions of Image Surfaces*. PhD thesis, Dept. of Engineering Science, Oxford University. p45.

Parida, L., Geiger, D., and Hummel, R. (1998). Junctions: detection, classification and reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(7):687–698.

Rohr, K. (1992). Recognizing corners by fitting parametric models. *Int. Journal of Computer Vision*, 9(3):213–230.

Schmid, C., Mohr, R., and Bauckhage, C. (2000). Evaluation of interest point detectors. *Int. Journal of Computer Vision*, 37(2):151–172.

Simoncelli, E. and Farid, H. (1996). Steerable wedge filters for local orientation analysis. *IEEE Trans Image Proc*, 5(9):1377–1382.

Smith, S. and Brady, J. (1997). SUSAN - a new approach to low level image processing. *Int. Journal of Computer Vision*, 23(1):45–78.

Smith, S. M. (1997). Reviews of optic flow, motion segmentation, edge finding and corner finding. Technical Report TR97SMS1, Oxford University.

Waltz, D. (1975). Understanding line drawings of scenes with shadows. In Winston, P. H., editor, *The Psychology of Computer Vision*. McGraw-Hill, New York.

Zetzsche, C. and Barth, E. (1990). Fundamental limits of linear filters in the visual processing of two dimensional signals. *Vision Research*, 30(7):1111–1117.