Junctions: Detection, Classification, and Reconstruction

Laxmi Parida, Davi Geiger, *Member, IEEE*, and Robert Hummel

Abstract—Junctions are important features for image analysis and form a critical aspect of image understanding tasks such as object recognition. We present a unified approach to *detecting* (location of the center of the junction), *classifying* (by the number of wedges—lines, corners, three-junctions such as *T* or *Y* junctions, or four-junctions such as *X*-junctions), and *reconstructing* junctions (in terms of radius size, the angles of each wedge and the intensity in each of the wedges) in images. Our main contribution is a modeling of the junction which is complex enough to handle all these issues and yet simple enough to admit an effective dynamic programming solution. Broadly, we use a template deformation framework along with a gradient criterium to detect radial partitions of the template. We use the minimum description length principle to obtain the optimal number of partitions that best describes the junction. Kona [27] is an implementation of this model. We (quantitatively) demonstrate the stability and robustness of the detector by analyzing its behavior in the presence of noise, using synthetic/controlled apparatus. We also present a qualitative study of its behavior on real images.

Index Terms—Junctions, corners, feature detection, low-level vision, minimum description length (MDL) principle, energy minimization.

1 Introduction

A critical component of most recognition systems is a stable, representative feature extraction from images. One of the key features used in recognition is junctions: *T*-junctions, *Y*-junctions, *X*-junctions, and so on. These junctions are also critical for stereo vision modules or motion modules, since these are places where occlusions can be identified. Such points, for example, coincide with the images of trihedral vertices of an object. These are critical features for recognition as suggested by [6], [7], [33]. They are also critical features for motion as suggested by Walach experiments [34] and by Movshon and Adelson [1]. A study of the role of junctions in stereo is presented in Malik [21].

There have been basically two different views/approaches for detecting junctions: edge detection followed by grouping of edges to form junctions [25], [24], [4], [2], [13], and, treating a junction as a template matching phenomenon [8], [12], [15]. In the former, it is assumed that the presence (or absence) of a junction is determined by "grouping" the intensity gradients near a hypothesized junction. Usually one is interested in examining large gradients in the direction perpendicular to the hypothesized Radial line. Experiments in this framework are limited, and even the richest ones shown in [25] are interesting but not exhaustive. In the latter approach, it is assumed that a (suitably small) local neighborhood is sufficient to detect a junction. The basic idea is to fit a junction-model to the input signal in a neighborhood. This involves minimizing an energy function which gives a measure of the "distance" of the junction-model from the input signal.

We use the template deformation framework to develop a "junction detector," to find corners (two-junctions), tri-corners (tri-junctions), quad-corners (quad-junctions), etc., defined as points where two or more homogeneous surface patches are located within a neighborhood of the point. Our approach is to use a combination of the two paradigms: grouping of edges (via Dynamic Programming) and fitting templates. We use a template deformation framework, using the minimum description length (MDL) principle, and we include the gradient criterium in order to detect the radial partitions of the template as a grouping mechanism. The task is to find the minimum number of wedges that best describes the junction. Note that as we increase the number of wedges, the junction description gets more accurate; hence, the task is to use the MDL principle to obtain an optimal number of such wedges. The minimum length encoding principle was first suggested by J. Rissanen [28], [30], [31], and is very close in spirit to the Kolmogorov sufficient statistic, as discussed in [3].

Kona[1] [27] is an implementation of our model. To test the stability/robustness of the detector, we have experimentally analyzed its behavior for the location and junction parameters against noise. We used a synthetic/controlled apparatus for a quantitative study. We have also qualitatively studied the stability/robustness of detector for real stereo images.

Our model is in contrast to other approaches which we briefly review here. In [13], in order to find corners, a local operator is first applied to detect possible junction locations. Then, inside an area around this, the edge lines are detected and depending on this number (of edge lines), a parametric model is used to fit the data is fit in a small region (25×25 pixels).

• *The authors are with the Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, U.S.A.*
E-mail: {parida, geiger, hummel}@cs.nyu.edu.

Manuscript received 24 Feb. 1997; revised 13 May 1998. Recommended for acceptance by K. Bowyer.
For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number 106851.

1. The word for "corner" in Hindi is *Kona*.

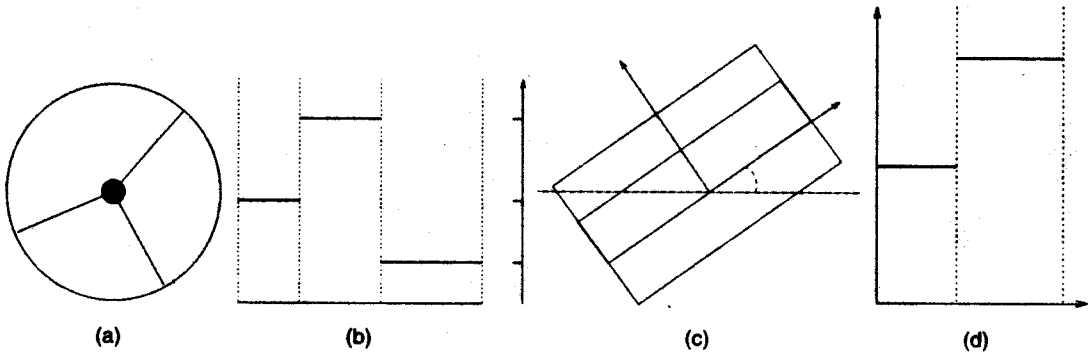

Fig. 1. Piecewise constant features. (a)-(b) A junction detector is shown on the top row: three-junction template with edges at angles θ_1, θ_2, and θ_3 with intensities T_1, T_2, and T_3. (c)-(d) A bar detector is shown: It is characterized by an angle α which gives the inclination of the bar and the intensities T_1 and T_2 at "heights" Y_1 and Y_2 from the center of the location.

The idea of performing local feature detection by projecting image data onto a subspace is fundamental in [8], [10]. Basically, the input is orthogonally projected onto a finite dimensional subspace of the Hilbert space of functions. An energy function (which is the L^2 norm of the the difference of the input and the fitted function) is minimized in this finite dimensional space. The two main issues are finding an orthonormal basis that spans a good finite dimensional subspace and minimizing the energy function. This approach can give closed form solutions for edges [8], [10], and lines [9]. A generalization to junctions is suggested in [32], using steerable filters, and in [18], through the use of principal component analysis (PCA). A possible drawback for the PCA is that it does not have an explicit model of the junctions (to be adapted).

The authors in [15], [12] employ an iconic model of the junctions, giving forms similar to our $\mathcal{R}$ (see (5) in Section 2). See Section 2 for a comparison.

In [16], corners and junctions (which are modeled as two adjacent corners) are represented by functions (models) that are blurred with a Gaussian (or an exponential filter) where the authors use a closed-form solution. In general, numerical methods are used to obtain parameters that minimize the distance to the input data using an L^2 norm. This is also the case in [17].

The paper is organized as follows: Section 2 describes the junction model. Section 3 studies detailed issues of the optimization process to estimate the junction parameters. Section 4 presents the results with a study of the stability of the detector. Section 5 concludes the paper.

2 The Junction Model

We model a junction as a region of an image where the values are piecewise constant in wedge-shaped regions emanating radially from a central point, covering a small disk centered at the point and omitting a (much) smaller disc centered at this point (see Fig. 1). The parameters of a junction consist of:

1) the radius of the junction-disk,
2) the center location,
3) the number of radial line boundaries,
4) the angular direction of each such boundary, and
5) the intensity within each wedge.

The radius of the disk addresses the "scale" issue, and the location of the center is a kind of "interest operator" [11] that determines the position where the feature is located in a region, possibly predefined. Another corner-detector [35] performs a similar role of detecting locations that have a corner. This uses a predetermined window/mask on each pixel and compares the intensity of every pixel with that of the center of the mask to determine whether there is a corner at the location.

We can formulate the junction detection problem as one of finding the parameter values that yield a junction that best approximates the local data using minimum description, and declaring local minima of the error as junctions. The best-fit parameter values provide attributes of the detected junction.

Let T denote the piecewise constant function/template. It has N angles and N intensities if N is the number of constant pieces. Further, let I denote the input signal.

Define the energy function, at a point (i, j) on the image as follows:

$$E = \mathcal{D} + \lambda \mathcal{G}, \tag{1}$$

where $\lambda \geq 0$.

The first term, $\mathcal{D}$, is a measure of the distance of the fitted function from the *data* using the L^2 norm:

$$\mathcal{D} = \int_0^\infty \int_0^{2\pi} |I(r,\theta) - T(\theta)|^2 g(r) r dr d\theta, \tag{2}$$

where $g(r)$ is an appropriate modulating function that goes to zero for large r, thus defining the template size.

We note that the input to the system is an image I^{raw} that can be filtered by a Gaussian mask to yield a smooth image I. Smoothing I^{raw} is needed due to the discretization of the image lattice and how it affects the junction detection. Our analysis of this effect is shown in Section 4.2.

The second term, $\mathcal{G}$, is a measure of the distance of the *gradient* using the L^2 norm.

$$\mathcal{G} = \int_0^\infty \int_0^{2\pi} |\nabla I(r,\theta) - \nabla T(\theta)|^2 g^*(r) r dr d\theta \tag{3}$$

where $g^*(r)$ is an appropriate modulating function, not necessarily the same as $g(r)$.

Note that

$$\nabla I = \frac{\partial I(r,6)}{\partial r}\mathbf{e}_r + \frac{1}{r}\frac{\partial I(r,6)}{\partial \theta}\mathbf{e}_\theta$$

$$\nabla T(\theta) = \frac{1}{r}\frac{\partial T(\theta)}{\partial 6}\mathbf{e}_\theta$$

where $\mathbf{e}_r$ and $\mathbf{e}_\theta$ are the orthonormal vectors in the r and θ direction, respectively, evaluated at (r, θ).

Separating the angular and the radial terms, $\mathcal{G}$ can be written as

$$\mathcal{G} = \mathcal{A}\mathbf{e}_\theta + \mathcal{R}\mathbf{e}_r,$$

where

$$\mathcal{A} = \int_0^\infty \int_0^{2\pi} \frac{1}{r^2}\left(\frac{\partial I}{\partial \theta} - \frac{\partial T}{\partial \theta}\right)^2 g^*(r) r dr d\theta, \tag{4}$$

$$\mathcal{R} = \int_0^\infty \int_0^{2\pi} \left(\frac{\partial I}{\partial r}\right)^2 g^*(r) r dr d\theta, \tag{5}$$

Taking $g^*(r) = r^2 G_\sigma(r)$ where $G_\sigma(r)$ is the Gaussian function with standard deviation σ, our function $\mathcal{R}()$ is the same as the regularity measure $S()$ for junctions as in (4) of [15]. Also, it is the same as ω_d defined in (31) in [12], which is the measure of spirality omitting the term normal to the radial direction. However, in [12], a general method is presented to detect circular symmetry; thus our function $\mathcal{R}$ could be considered as a special case.

2.1 On the Exact Form of Template T and Energy E

This section gives a more rigorous definition of T and E as a *continuous* function of two variables. The reader may skip this section without loss of continuity.

Let the image be given by $I(x, y)$. T is a template that partitions the image in N regions, an "N-junction" template. Between the partition lines (wedges), say p and $p + 1$, we assign a constant value T_p for the template. Let us further consider a template for the edges that is zero everywhere except along the partition lines, $\theta = \theta_p$. We can write both templates as

$$T(\theta) \equiv \left\{\left(T_p, \theta_p\right); p \in (1, 2, \ldots, N)\right\} = \sum_{p=1}^{N} T_p u\left(\theta - \theta_p\right) u\left(\theta_{p+1} - \theta\right)$$

$$T'(\theta) = e(\theta) \equiv \left\{\left(e_p, \theta_p\right); p \in (1, 2, \ldots, N)\right\} = \sum_{p=1}^{N} e_p \delta\left(\theta - \theta_p\right) \tag{6}$$

where $\theta_{N+1} \equiv \theta_1$, $\delta(x)$ is the Dirac delta function, and $u(x)$ is the step function (or Heaviside function), i.e.:

$$\int_{-\infty}^{\infty} \delta(x)dx = 1 \quad \delta(x) = \begin{cases}\infty & x = 0 \\ 0 & x \neq 0\end{cases} \text{ and}$$

$$u(x) = \int_{-\infty}^{x} \delta(x)dx = \begin{cases}1 & x \geq 0 \\ 0 & x < 0\end{cases}$$

To fit a junction template $T(\theta)$ to the image $I(x, y)$ and to fit an edge template $e(\theta)$ to the image gradient along the angle θ (perpendicular to the radial line), an error function based on the L^2 norm weighted by a function $g(r)$ is used. Using the polar coordinate system,

$$r = \sqrt{x^2 + y^2},\ \theta = \arctan\frac{y}{x},\ \text{and}\ x = r\cos\theta,\ y = r\sin\theta$$

the image gradient is

$$\nabla I = \frac{\partial I(r,6)}{\partial r}\mathbf{e}_r + \frac{1}{r}\frac{\partial I(r,6)}{\partial 6}\mathbf{e}_\theta.$$

Thus, the energy E at a point (i, j) is:

$$E = \int_0^\infty \int_0^{2\pi} \{[I(r,\theta) - T(\theta)]^2 g(r) + \lambda\left|\frac{1}{r}\frac{\partial I(r,\theta)}{\partial \theta} - e(\theta)\right|^2 g^*(r)\} r dr d\theta \tag{7}$$

where $g(r)$ and $g^*(r)$ are functions that weights the importance of the piecewise constant fit according to the distance from the center $r = 0$.

3 Energy Minimization

Recall the energy equation:

$$E = \lambda\mathcal{R} + (\mathcal{D} + \lambda\mathcal{A}) = \lambda\mathcal{R} + \overline{E} \tag{8}$$

$\mathcal{R}$ is independent of the junction template and is used to filter out the image locations. $\overline{E}$ is minimized to obtain the most appropriate junction parameters. We now discuss the details of computing $\mathcal{R}$ (scale and location) and $\overline{E}$ (junction parameters).

3.1 Scale and Location (on $\mathcal{R}$)

We have considered $g(r) = 1/r$. Notice that for λ to be used in (1), $g^*(r) = r^2 g(r) = r$ is a choice that makes $\mathcal{D}$ and $\mathcal{G}$ have the same unit (intensity^2/length). (Thus, λ is unitless as expected.)

3.1.1 Estimating R_0 and R_1 (scale)

We further refine $g(r)$ as follows:

$$g(r) = \begin{cases}0 & r < R_0 \\ \frac{1}{r} & R_0 \leq r \leq R_1 \\ 0 & r > R_1\end{cases}$$

Notice that we introduce a "hole" of size $R_0 > 0$ and R_1 is the size of the window. To compute R_0 and R_1, we study $\mathcal{R}'$, which is $\mathcal{R}$ for r, the radius being considered. Fig. 2 shows the plot of the relative values of $\mathcal{R}^r$, e^r, which is

$$e^r = \frac{\mathcal{R}^r}{\mathcal{R}^{r+1}}. \tag{9}$$

Notice that $\mathcal{R}$ has comparatively larger values close to the origin, which we remove by introducing a "hole." We look for a range of r values where e^r is less than a threshold, τ_w. This leads to obtaining two values that correspond to R_0 and R_1. See Fig. 2 and Fig. 3 for examples.

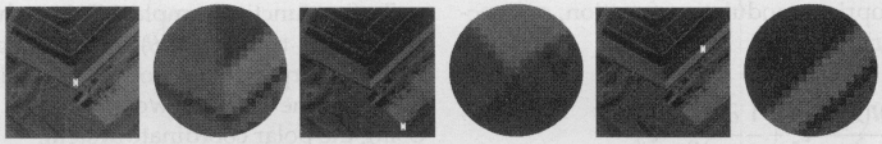

(a) (b) (c) (d) (e) (f)

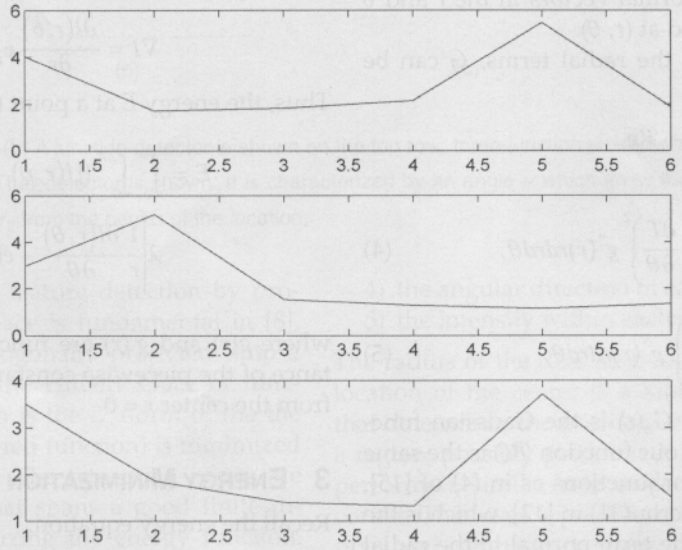

(g)

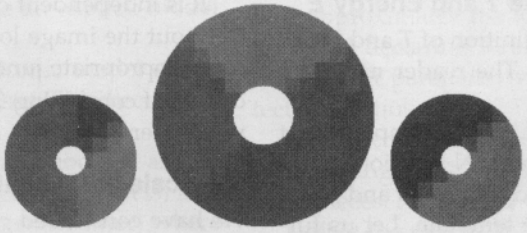

(h) (i) (j)

Fig. 2. Estimating R_0 and R_1, the size of the hole and the size of the window, respectively. The relative values of $\mathcal{R}$, e^r is plotted against the radius r. The threshold value of $\tau_w = 2.1$ was used to obtain the values of R_0 and R_1 as shown. Notice that when we pick the lower bound (R_0), we take one unit less than the picked interval (for instance, the range for the first graph is $r = 2 \ldots 4$ and $R_0 = 1$, $R_1 = 4$). We do this since we sample the values of $\mathcal{R}$ only at integral values and $\mathcal{R}^r = \mathcal{R}^s$, $r - 1 < s \leq r$ where s is a rational. (a)-(f) Marked images. (a)-(b) Location 1 (46,50). (c)-(d) Location 2 (68,81). (e)-(f) Location 3 (71,27). (g) e_r versus r plots for the three locations. (h)-(j) The computed windows. (h) Location 1 (46,50), $R_0 = 1$, $R_1 = 4$ pixels. (i) Location 2 (68,81), $R_0 = 2$, $R_1 = 8$ pixels. (j) Location 3 (71,27), $R_0 = 1$, $R_1 = 4$ pixels.

The weight function $a(x)$ which is defined in (37) in [12] is similar to our $g(r)$ function.

The use of a hole is also supported by psychophysical experiments [34] that suggest humans may also utilize such a method. The experiments suggested by Wuerger et al. [34] have shown that different motion effects can be produced depending on the junction movement. In a Wallach type of experiment, junctions occur due to the intersection of lines, and junctions move due to the movement of the lines. The results of the experiments in [34] seem to suggest that junctions are detected although there is a small gap (R_0) between the lines.

3.1.2 Estimating the Location

We do the following to select the best location in an image region.[2] We compute $\mathcal{R}^r = \mathcal{R}/r$ for all the points in the region (with not necessarily the same window size). The one with the minimum value defines the location. We illustrate this in Fig. 4.

3.2 Reconstructing the Junction (on $\overline{E}$)

In this section, we explain how the junction parameters (number of wedges, wedge angles, and wedge intensities) are estimated. Recall that the last section discussed the es-

2. Notice that $\mathcal{R}$ has unit (intensity)^2 length; thus $\mathcal{R}$ has unit (intensity)^2.

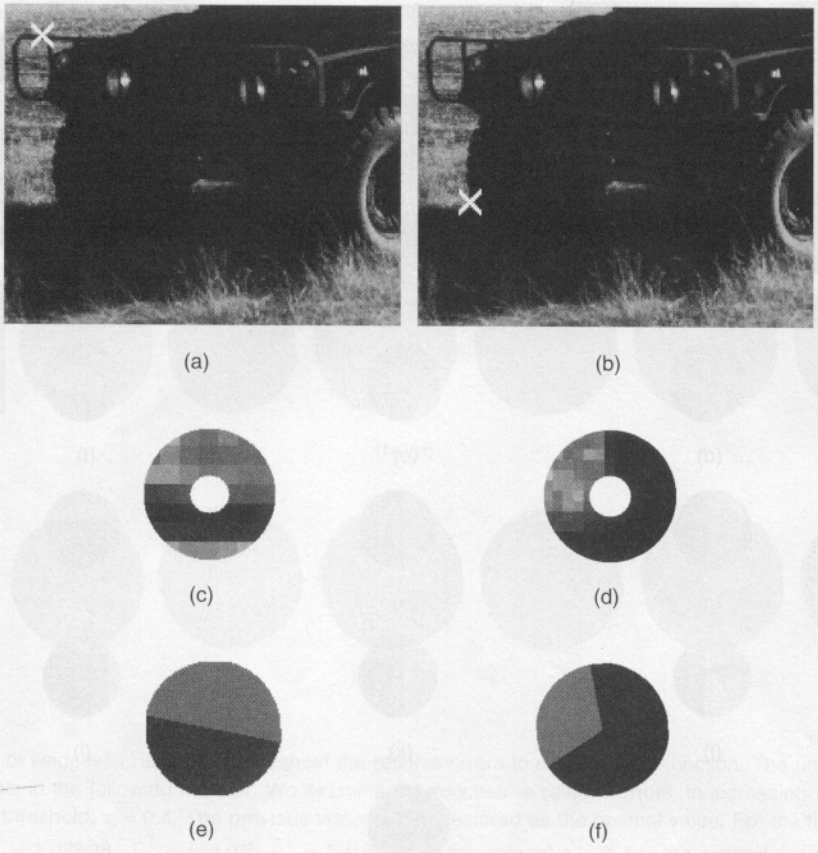

Fig. 3. An example to show the dynamic computation of R_0 and R_1 at different locations on the image. We have used $\tau_w = 2.1$. (a) Position on the image. (b) Position on the image. (c) Adaptive window at (a), ($R_0 = 1$, R_1=3 pixels). (d) Adaptive window at (b), ($R_0 = 2$, $R_1 = 6$ pixels). (e) The best fit to (c). (f) The best fit to (d).

timation of the radius (size) of the junction and the size of the hole.

We can carry out an appropriate numerical integration to obtain the value of the first term $\mathcal{R}$. The second factor, $\overline{E}$, is used to estimate the junction template. The unknown factors are: N, the number of intersecting lines (or wedges) at the junction,[3] $\{\theta_p\}$, $\{T_p\}$, $p = 1 \ldots N$, where N is the number of wedges, θ_ps are the angles where the partitions occur, T_ps are the intensity values.

We can write down $\overline{E}$ as

$$\overline{E} = \mathcal{F} + \mathcal{V},$$

where, $\mathcal{F}$ is fixed, that is, it does not depend on the unknown parameters, whereas $\mathcal{V}$ does. Recall $g^*(r) = r^2 g(r)$. With some straightforward manipulations, we obtain:

$$\mathcal{F} = \int_0^\infty \int_0^{2\pi} \left[I^2(r,\theta) + \lambda \left(\frac{\partial I(r,\theta)}{\partial \theta} \right)^2 \right] g(r) r dr d\theta$$

$$\mathcal{V} = \int_0^\infty \int_0^{2\pi} \left[T^2(\theta) - 2T(\theta) I(\theta) + \lambda \left(\frac{\partial T^2}{\partial \theta} - 2 \frac{\partial I^2}{\partial \theta} \frac{\partial T}{\partial \theta} \right) \right] g(r) r dr d\theta$$

$\mathcal{F}$ can be approximated numerically. Also, $\mathcal{V}$ can be approximated as $\tilde{\mathcal{V}}$.

$$\tilde{\mathcal{V}} = \sum_{p=1}^{N} [(\theta_{p+1} - \theta_p)(-2T_p \bar{I}_{\theta_p, \theta_{p+1}} + T_p^2 C) + \lambda(-2T_p'(\partial_\theta \tilde{I})_{\theta_p} + C^* T_p'^2)]$$

where

$$C = \int_0^\infty g(r) r dr = R_1 - R_0, \tag{10}$$

$$\tilde{I}(\theta) = \int_0^\infty I(r,\theta) g(r) r dr, \tag{11}$$

$$C^* = \int_0^\infty g^*(r) r dr = \int_0^\infty r^2 dr = \frac{R_1^3 - R_0^3}{3}$$

$$(\partial_\theta \tilde{I})_{\theta_p} \equiv \frac{\partial \tilde{I}(\theta)}{\partial \theta}\Big|_{\theta_p} = \frac{\partial \int_0^\infty I(r,\theta) g(r) r dr}{\partial \theta}\Big|_{\theta_p}$$

$$\bar{I}_{\theta_p \theta_{p+1}} = \Delta\theta \sum_{j=j_p}^{j_{p+1}} \tilde{I}(j\Delta\theta) \tag{12}$$

Further, $\tilde{I}(\theta)$, $(\partial_\theta \tilde{I})_{\theta_p}$ can be approximated numerically. For the sake of brevity, we have omitted the details of the derivations here, and present it in Appendix A.

3. N for homogeneous region is one, for line and corner is two, for junctions like T-, Y-junction is three, four for X-junction, and so on.

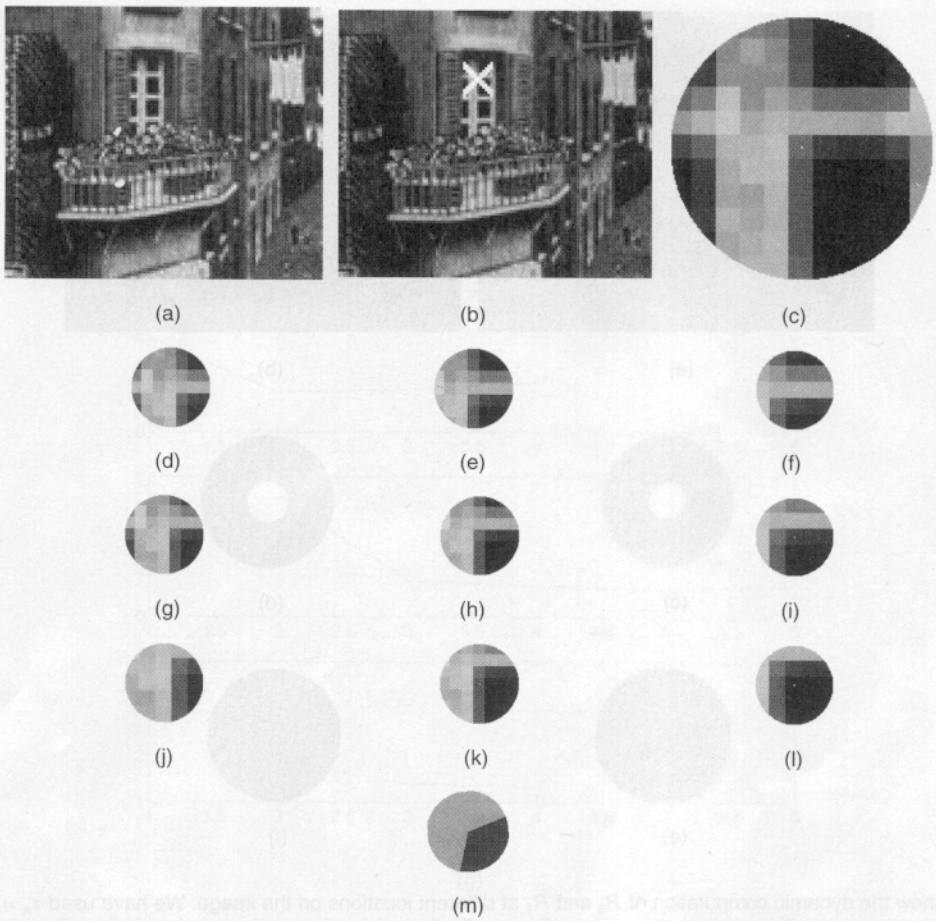

Fig. 4. Estimating the location of a junction: The use of $\mathcal{R}^r$, to locate the center of the junction. $L_{x,y}$ indicates the x and y coordinates on the image. For convenience, the input data is normalized, that is, the intensities are scaled to lie in [0,1]. The unit of each value is (intensity)^2. Note that the location $L_{52,26}$ has the minimum value in the neighborhood. ($\tau_w = 2.1$). (a) Input image. (b) $L_{52,26}$ marked. (c) Region around $L_{52,26}$. (d) $L_{51,25}$, $\mathcal{R}^r = 1.24$. (e) $L_{52,25}$, $\mathcal{R}^r = 1.28$. (f) $L_{53,25}$, $\mathcal{R}^r = 1.02$. (g) $L_{51,26}$, $\mathcal{R}^r = 1.52$. (h) $L_{52,26}$, $\mathcal{R}^r = 0.66$. (i) $L_{53,26}$, $\mathcal{R}^r = 1.02$. (j) $L_{51,27}$, $\mathcal{R}^r = 1.12$. (k) $L_{52,27}$, $\mathcal{R}^r = 1.24$. (l) $L_{53,27}$, $\mathcal{R}^r = 1.16$. (m) The junction at location $L_{52,26}$.

3.2.1 Estimating Wedge Angles and Intensities

Although the energy equation E looks fairly complex, it has a remarkably simple and natural interpretation. Once the role of each factor of the energy equation is ascertained, some computations (like evaluating $\mathcal{F}$, for instance) can be omitted without changing the solution. To recapitulate, the factors are used thus:

1) $\mathcal{R}$ is used for scale, i.e., to determine the size of the window, R_1. This is also used to obtain the exact location of the junction in a neighborhood.
2) $\overline{E}$ is used to obtain the junction parameters.

The $\tilde{I}(\theta)$ (11) can be viewed as integrating the intensity along a radial line. Thus, the two-dimensional image is projected on to a one-dimensional coordinate θ, appropriately discretized.

Let $\tilde{I}(\theta_i)$ be defined for $\theta_1, \theta_2, \ldots, \theta_d$, and, $(\theta_{i+1} - \theta_i) = 2\pi/d, \forall\ i$.[4] For a p-junction, $T_1, T_2, \ldots, T_p$, are the template intensities and the wedge boundaries are at

$$\theta_{k_1}, \theta_{k_2}, \ldots, \theta_{k_p}, \theta_{k_1} \le \theta_{k_2} \ldots \le \theta_{k_p}.$$

Assuming we know the θ_ps, we can obtain T_ps by setting $\frac{\partial \mathcal{V}}{\partial T_p} = 0, \forall p$. When $\lambda = 0$, $\frac{\partial \mathcal{V}}{\partial T_p} = f(\theta_p)$, for some function $f()$. Thus fixing θ_ps, the T_ls can be shown to be the following:

$$T_l = \frac{\sum_{j=k_l}^{k_{l+1}} \tilde{I}(\theta_j)}{k_{l+1} - k_l}.$$

In other words, T_l is a piecewise constant fit which is the average value of the data in that region. The energy for the fit is

$$\overline{E}^p = \sum_{l=1}^{p} \sum_{j=k_l}^{k_{l+1}} \left(\tilde{I}(\theta_j) - T_l\right)^2.$$

When $\lambda \neq 0, \overline{E}^p$ has some extra terms. It can be verified that T_l is the same as before, and, $\overline{E}^p$ is:

4. If $i = d$, $i + 1$ is defined to be one.

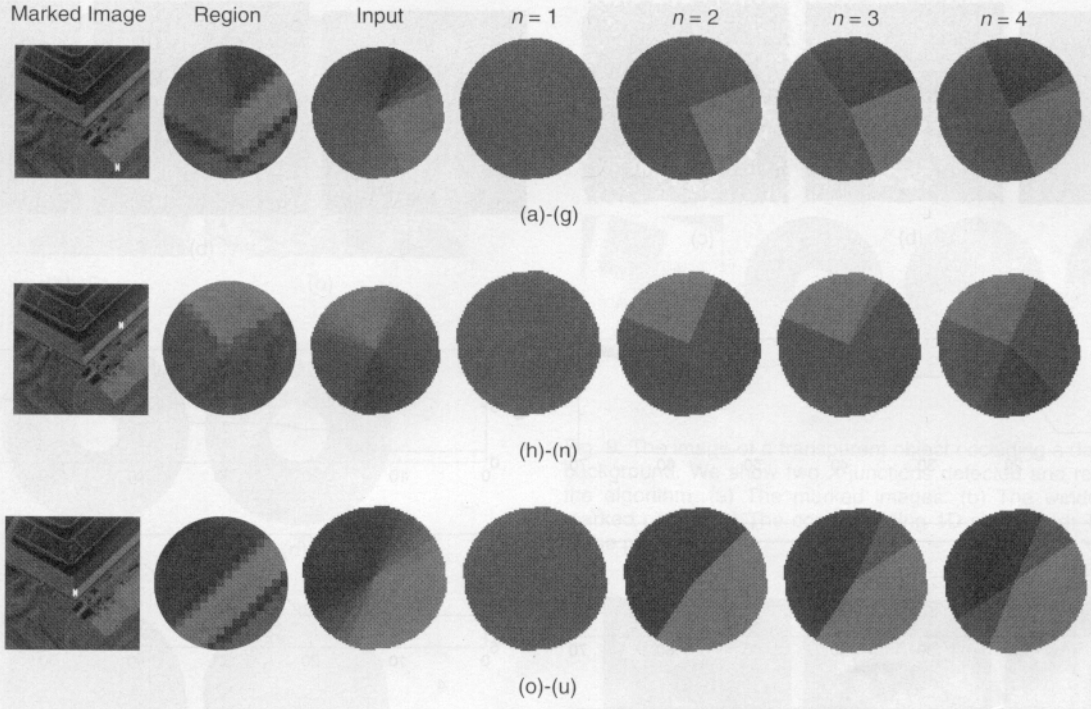

Fig. 5. Estimating the number of wedges: τ_c is used to threshold the relative errors to declare an N-junction. The unit of $\overline{E}^n$ is intensity^2 * length. The thresholding process works in the following manner: We examine the successive relative errors, in increasing n, until we reach an n where the relative error exceeds the threshold, $\tau_c = 0.4$. The previous value of n is declared as the optimal value. For the first example (top row), the E^n values are $\overline{E}^1 = 2{,}814.93$, $\overline{E}^2 = 1{,}088.28$, $\overline{E}^3 = 654.07$, $\overline{E}^4 = 504.52$, and the optimal n is 3. For the second example (middle row), the E^n values are $\overline{E}^1 = 2{,}174.61$, $\overline{E}^2 = 567.28$, $\overline{E}^3 = 508.28$, $\overline{E}^4 = 408.21$, and the optimal n is 2. For the third example (bottom row), the E^n values are $\overline{E}^1 = 7{,}154.97$, $\overline{E}^2 = 1{,}178.82$, $\overline{E}^3 = 942.57$, $\overline{E}^4 = 670.92$, and the optimal n is 2. (a)-(g) Location 1 (46,50). (h)-(n) Location 2 (68,81). (o)-(u) Location 3 (71,27).

$$\overline{E}^p = \sum_{l=1}^{p}\sum_{j=k_l}^{k_{l+1}}\left(\tilde{I}\left(\theta_j\right) - T_l\right)^2 + \lambda\sum_{i=k_1}^{k_p}\left(\tilde{I}\left(\theta_{k_i}\right) - \tilde{I}\left(\theta_{k_i-1}\right)\right)^2.$$

We compute the θ_ps by exploring all possible set of θ_ps. We summarize the dynamic program

$$C_{jl}^p = \begin{cases} \text{cost of fitting } T_p \text{ to } \theta_j, \theta_{j+1}, \ldots, \theta_l & l \geq j \\ \text{cost of fitting } T_p \text{ to } \theta_l, \theta_{l+1}, \ldots, \theta_1 \ldots \theta_j & l < j \end{cases}$$

$$\varepsilon_{sj}^p = \begin{cases} C_{sj}^1 & p = 1 \\ \min_{i<j}\left(\varepsilon_{si}^{p-1} + C_{(i+1)j}^p\right) & \text{otherwise} \end{cases}$$

$$\overline{E} = \min_{s=1,m} \varepsilon_{sm}^N$$

This dynamic programming solution has been implemented in Kona.

$\mathcal{F}$ is a constant (fixed), depending only on the image, and does not influence the junction parameters. The scenario where $\mathcal{F}$ could play a role is at the time of comparing the energy values in a neighborhood. But we use $\mathcal{R}$ to filter the neighborhood so $\mathcal{F}$ may be ignored without any damage.

3.2.2 Estimating the Number of Wedges

We estimate the number of wedges by measuring the rate of increase of energy by thresholding the ratio of the energy measurement, $\overline{E}^n$. More precisely, the (optimal) number of wedges, N, is computed by thresholding the relative error, r^n,

$$r^n = \frac{\overline{E}^{n+1}}{\overline{E}^n}.$$

Notice that by definition, r^n is unitless. Although, in principle, we are looking for the minimum r^n, in practice, we terminate the computation when r^n drops below a predefined threshold, say τ_c, as n is increased. (See Fig. 5 for an example.) Note that as the number of parameters increases, $\overline{E}^n$ decreases, i.e., $r^n < 1$.

To summarize, the following steps are involved in detecting junctions on a large region of an image:

1) Compute $\mathcal{R}$, the measure of radial variation, and R_0, the size of the hole and R_1, the size of the template at every point.
2) Filter the locations using a threshold on $\mathcal{R}$.
3) In a neighborhood of a filtered location, pick the one with minimum $\mathcal{R}$, and remove all other locations within a radius of R_1 of this. Repeat this for all the filtered locations.
4) Compute the junction parameter for all the filtered locations.

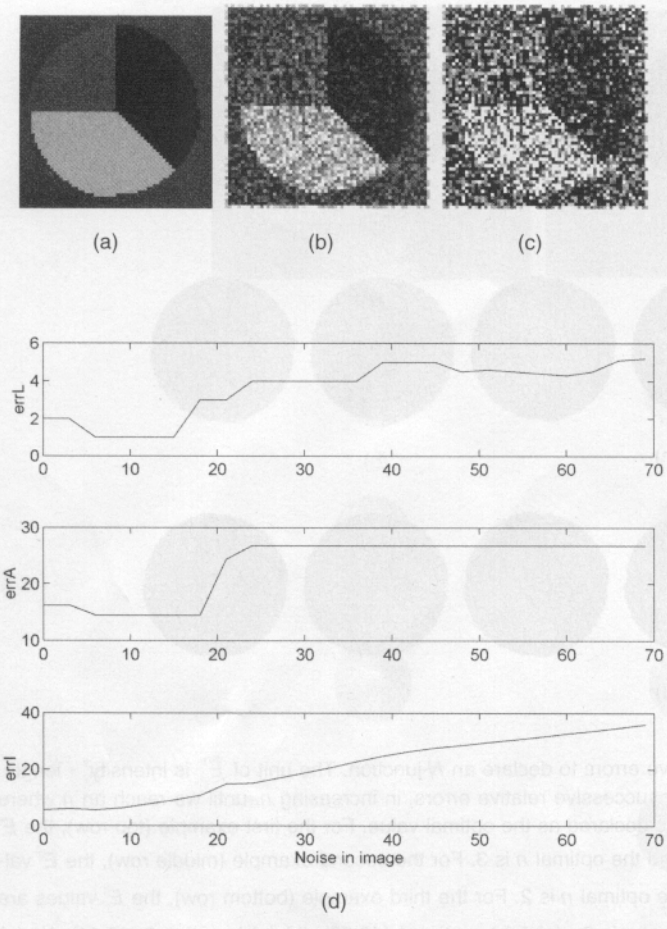

Fig. 6. Test of stability: The range of intensities in the image is 0–255. A noise with a Gaussian distribution, with standard deviation σ, is introduced randomly at each pixel. The first image in the top left has $\sigma = 0$ (that is, no noise) and the angles of the three-corner at the center, (x, y), are $a_{1f} = 90^\circ$, $a_{2f} = 180^\circ$ and $a_{3f} = 315^\circ$ with intensities $i_{1f} = 120$, $i_{2f} = 200$ and $i_{3f} = 40$. σ is varied from 0 to 69 to obtain 24 images, three of which are shown in the top row for illustration. The three-corners around the center, (x_σ, y_σ), of each image is extracted and the following errors are computed: error in location, $errL_\sigma = |x_\sigma - x| + |y_\sigma - y|$, error in intensities, $(errI_\sigma)^2 = \sum_{j=1}^{3} (i_{j\sigma} - i_{jf})^2/3$, error in angles, $(errA_\sigma)^2 = \sum_{j=1}^{3} (a_{j\sigma} - a_{jf})^2/3$. Finally, $errL_\sigma$ vs $\{\sigma\}$, $errA_\sigma$ vs $\{\sigma\}$, and, $errI_\sigma$ vs $\{\sigma\}$ are plotted. See the text for further explanations. (a) $\sigma = 0.0$. (b) $\sigma = 30.0$. (c) $\sigma = 69.0$. (d) Noise in image.

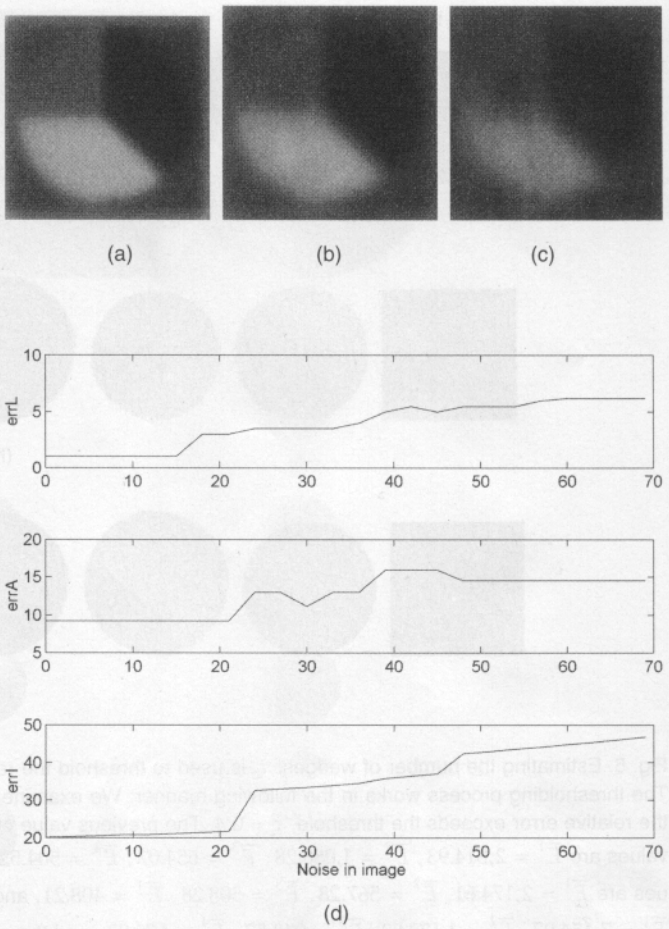

Fig. 7. Test of stability over smoothed noisy images: We use the same set of (noisy) images as used in Fig. 6, except that these images are now smoothed. We increase the smoothing factor as the image gets noisier. The first eight images ($\sigma = 0, 3, \ldots, 21$) are smoothed using a $\sigma_s = 2.0$, the next eight ($\sigma = 24, 27, \ldots, 45$) uses $\sigma_s = 3.0$ and the last set ($\sigma = 48, 51, \ldots, 69$) uses $\sigma_s = 4.0$. We plot the errors as in Fig. 6. As expected we are able to get rid of the high errors for the very sharp images (low values of σ). (a) $\sigma = 0.0$, $\sigma_s = 2.0$. (b) $\sigma = 30.0$, $\sigma_s = 3.0$. (c) $\sigma = 69.0$, $\sigma_s = 4.0$.

Kona is an implementation of this junction model. It is programmed in C. A version in C++ and a parallel implementation of Kona is also available.

4 Results of Experiments

We carried out a series of experiments on synthetically generated images (see Fig. 6 and Fig. 7) to test the stability of the algorithm in the presence of noise in the image, and to further understand the role of smoothing. In these experiments, we suppressed the use of threshold values using a fixed $R_1 = 4$ and fixed $N = 3$, i.e., forcing the algorithm to pick up the optimal three-corners. We then carried the experiments on real images to test the general performance of the algorithm.

4.1 Stability of the Algorithm

Fig. 6 shows the result of one such experiment where we look for three-corners in the center region of the image. It shows that when the image edges are sharp ($\sigma = 0 \ldots 3$ in the figure), the error in the angles and the intensities are slightly higher. The errors are least in the range of $\sigma = 6 \ldots 20$, and increase further on.

Note that the (Manhattan) distance of the location of the three-corner is less than six pixels; the error in the angles is bounded by 30° at the very worst and the intensity differs by 40 units in the worst situation. It also shows that, in the best situation the angle differs, on an average, by 15° from the true answer.

4.2 Effect of Image Smoothing

We present in Fig. 7 the results of one such experiment. The images used here are the smoothed version of the ones shown in Fig. 6. Note the reduction in error around the sharp images or low values of σ.

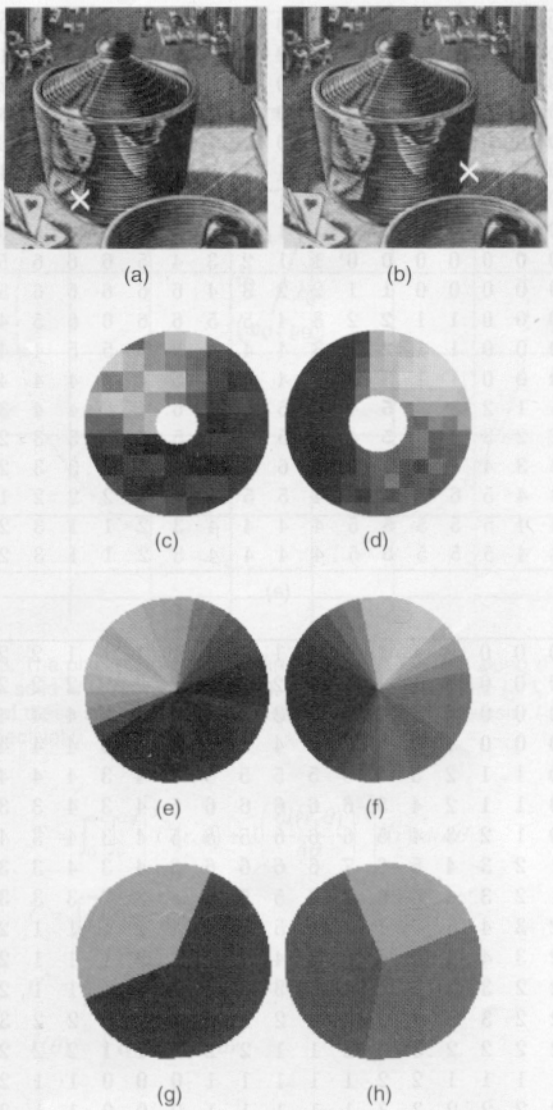

Fig. 8. Examples of a sharp corner (column 1) and a T-junction on an image (column 2). We have used $\tau_w = 2.1$ to get R_0 and R_1 values, and τ_c=0.4 to get the N-junctions. (a) Position on the image. (b) Position on the image. (c) Adaptive window for (a), (R_0 = 1, R_1=4 pixels). (d) Adaptive window for (b), (R_0 = 2, R_1=6 pixels). (e) Projection of (c) to 1D signal. (f) Projection of (d) to 1D signal. (g) Best-fit to (e). (h) Best-fit to (f).

It is worth noting that even for sharp junctions (without noise added), the algorithm performs better when smoothing is introduced. It is possible that smoothing the image just along radial directions would be more appropriate, to avoid smoothing along wedges. In our experiments, this possible effect was not noticed, while the substantial speed/efficiency of filtering on x and y (and not along r) encouraged us to use an overall Gaussian blur.

4.3 Results on Real Images

The threshold for deciding N of the junction, τ_c, was set to 0.4, and the threshold for deciding R_0 and R_1, τ_w, was set to 2.1, for all the images we used.[5] Also, we chose a value of $\lambda = 1$ for these experiments.[6] We observed that using the number of discretization units of the angle as 16 has worked well: Eight was found to be too coarse, and 32 did not result in any significant higher accuracy.[7]

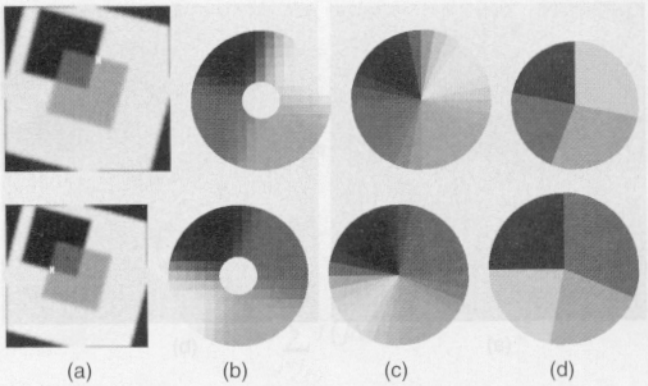

Fig. 9. The image of a transparent object occluding a dark object in the background. We show two X-junctions detected and reconstructed by the algorithm. (a) The marked images. (b) The window around the marked points. (c) The corresponding 1D signals. (d) The 4-junctions at the points.

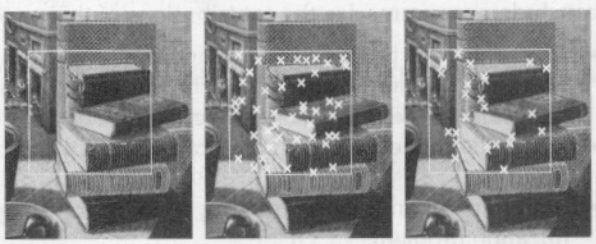

Fig. 10. Example of an image showing the results of the junction detector Kona. The 2-junctions have been filtered which has removed the large angle corners (> 160 degrees). Also, the low-contrast 3-junctions have been filtered out. The values of the various control parameters are: $\tau_c = 0.4$ (to obtain N the number of wedges), $\tau_w = 2.1$ (to obtain the size of hole, R_0, and the radius, R_1, of the template), number of radial lines = 32, $\lambda = 1$. (a) Marked input image. (b) Two-junctions. (c) Three-junctions.

Fig. 8 shows the function/template that is fit to different points on the images. Fig. 9 shows an image where the X-junctions have been detected, and Fig. 10 shows the results of Kona on a region of the image. After the filtering, it computes the parameters for about sixteen junctions in a minute, on a Sun Sparc Station.

We demonstrate the stability of the junction detector by running it on a pair of stereo images as shown in Fig. 11. (The results of this are intended for use by stereo algorithms.) The analysis of the result of this is presented in Fig. 12. We measure the stability by first finding a correspondence between location (x_l, y_l) in the left image and (x_r, y_r) in the right stereo image and, then, comparing the number of junctions in a small neighborhood of the two locations.

5. Recall that τ_c and τ_w are thresholds for relative errors, hence, are unitless.

6. Under ideal circumstances, $\mathcal{R} = 0$, and further, the only nonzero values of $\mathcal{A}$ are along the wedge partition. Thus the value of λ controls the amount of "contrast" in the neighboring wedge intensities of the reconstructed junction—the higher the value, the higher the desired contrast.

7. We choose a multiple of four because of the four quadrants—using any arbitrary number could do as well, in principle.

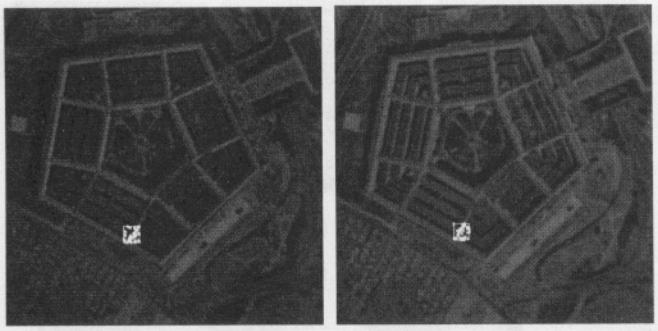

(a) (b)

Fig. 11. Kona is run on a 25 × 25 size window (at location (190,350)) on a pair of stereo images. It detects junctions, with significant contrast, on similar locations on both the images. The analysis of this is presented in a simplified form in Fig. 12. (a) Left Image. (b) Right Image.

0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1
0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	2	2
0	0	0	0	0	0	**0**	**0**	**0**	**0**	**0**	**0**	**1**	**1**	**2**	**2**	**2**	**3**	**3**
0	0	0	0	0	0	**0**	0	0	0	1	1	2	2	3	3	3	3	4
0	0	0	0	0	0	**0**	0	1	1	2	2	3	3	4	4	4	4	5
0	0	0	0	0	0	**0**	0	1	1	2	2	3	4	5	5	5	5	6
0	0	0	0	0	0	**0**	0	1	1	2	3	4	5	6	6	6	6	6
0	0	0	0	0	0	**0**	0	1	1	2	3	4	5	6	6	6	5	5
0	0	0	0	0	0	**1**	1	2	2	3	4	6	6	6	6	6	5	5
0	0	0	0	1	1	**2**	2	3	4	5	5	6	6	6	6	5	4	5
0	0	0	0	1	1	**2**	2	3	4	4	4	5	5	5	5	4	4	4
0	0	0	0	1	1	**3**	3	3	4	4	4	5	4	4	4	4	4	4
1	1	1	2	3	3	**5**	4	4	5	5	5	6	4	4	4	4	3	3
2	2	2	3	4	3	**5**	4	4	5	5	4	5	3	3	3	3	2	3
2	3	3	4	5	4	**6**	6	5	6	6	5	6	4	3	3	3	2	4
3	4	4	5	6	5	**5**	5	4	5	5	4	4	3	2	2	2	1	3
3	4	4	5	5	5	**5**	5	4	4	4	4	3	2	1	1	3	2	3
3	4	4	5	5	5	**5**	5	4	4	4	4	3	2	1	1	3	2	3

(a)

0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	2	2	2
0	0	0	0	0	1	1	1	1	2	2	2	2	2	2	2	2	2
0	0	0	0	0	1	2	2	2	3	3	4	4	3	3	4	4	4	3
0	0	0	0	1	2	3	3	3	4	4	4	4	3	3	4	4	4	3
0	0	1	1	2	3	4	5	5	5	5	5	5	4	3	4	4	4	3
0	0	1	1	2	4	5	6	6	6	6	6	5	4	3	4	3	3	2
0	0	1	2	3	4	5	6	6	6	5	5	5	4	3	4	3	4	3
0	1	2	3	4	5	6	7	6	6	6	6	5	4	3	4	3	3	2
0	1	2	3	4	5	6	7	6	5	5	5	4	3	2	3	3	3	2
1	2	3	4	5	7	7	7	6	5	5	4	2	2	1	1	1	2	2
1	2	3	4	4	6	6	6	5	4	4	4	2	2	1	1	1	2	2
1	2	2	3	3	5	5	4	3	3	3	3	1	1	1	1	1	2	2
1	2	2	3	3	4	4	3	2	2	2	2	1	1	1	2	2	3	3
1	2	2	2	2	3	3	2	1	1	2	2	1	1	1	2	2	2	2
1	1	1	1	1	2	2	1	1	1	1	1	0	0	0	1	1	2	2
2	2	2	2	2	3	3	1	1	1	1	1	0	0	0	1	1	2	2
1	1	2	2	2	2	2	1	2	1	1	1	1	1	1	1	1	2	2
1	1	2	2	2	2	2	1	2	1	1	1	1	1	1	1	1	2	2
2	2	3	3	3	3	2	1	2	1	1	1	1	1	1	1	1	2	2

(b)

Fig. 12. Kona is run on a 25 × 25 size window as shown in Fig. 11. The 19 × 19 matrix gives a count of the number of junctions Kona detected in a 7 × 7 window centered at that location. (Thus the (1,1) entry of the matrix corresponds to location (193,353) of the image and gives the number of junctions detected.) By inspection, we have indicated the alignment points of the pair of images by presenting the numbers in bold. It can be seen that the number of junctions (after aligning the two images), match up reasonably well (i.e., gives good correlation values). (a) Matrix corresponding to the left stereo image. (b) Matrix corresponding to the right stereo image.

5 Conclusions

To conclude, we summarize the contributions of this paper:

- A modeling of the junction that includes the gradient term in the model and supports the removal of a small disc at the center of the junction.
- Using the model to obtain the scale and location of a junction.
- A simple dynamic programming solution to the optimization problem.

This paper demonstrates the successful use of piecewise constant functions, using a minimum description length principle, to detect features like junctions and corners. Similar piecewise constant functions may be used for features like bar detectors, blobs, endpoints etc.

Kona is an implementation on this model. This detector is being currently tested in a multitude of tasks such as object recognition (based on geometric hashing), detection of illusory contours, and depth estimation from stereo pairs of images. It is also being tested in a novel approach to image compression using image diffusion and multijunctions.

Appendix A: Computing *E*

Exploiting the radial invariance of the templates $T(\theta)$ and $e(\theta)$, we can write (7) as follows:

$$E = \int_0^\infty \int_0^{2\pi} \left[I^2(r,\theta) - 2T(\theta) I(r,\theta) + T^2(\theta) \right] g(r) r dr d\theta$$

$$+\lambda \int_0^\infty \int_0^{2\pi} \left[\left(\frac{1}{r} \frac{\partial I(r,\theta)}{\partial \theta} \right)^2 - 2e(\theta) \frac{1}{r} \frac{\partial I(r,\theta)}{\partial \theta} + e^2(\theta) \right] g^*(r) r dr d\theta$$

$$= \int_0^\infty \int_0^{2\pi} \left[I^2(r,\theta) + \lambda \left(\frac{\partial I(r,\theta)}{\partial \theta} \right)^2 \right] g(r) r dr d\theta$$

$$+ \sum_{p=1}^{N} \lim_{\epsilon \to 0} \int_{\theta_{p^-}^{\theta_{p+1}^-}}^{\theta_{p+1}^-} \int_0^\infty$$

$$\left[T_p^2 - 2T_p I(r,\theta) + \lambda \left(re^2(\theta) - 2re(\theta) \frac{\partial I(r,\theta)}{\partial \theta} \right) \right] g(r) r dr d\theta$$

where we have assumed $g^*(r) = r^2 g(r)$. We can now exploit the radial invariance of the templates to perform the integrals over dr. Define the following:

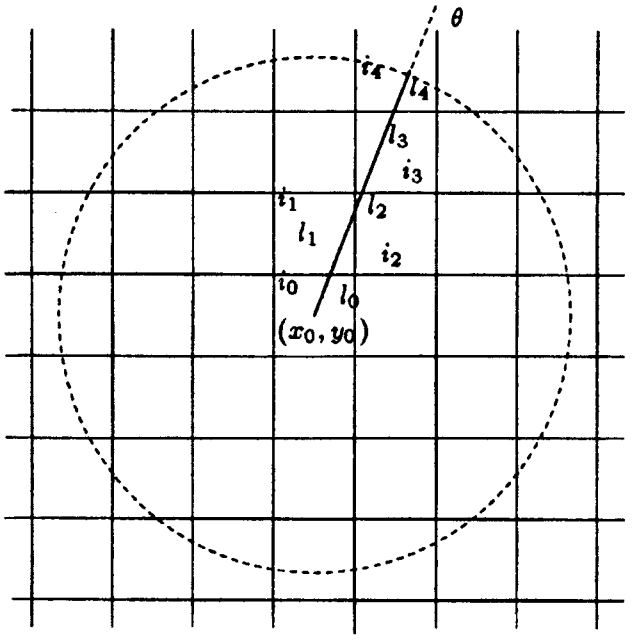

Fig. 13. The pixels shown as rectangular cells. The ray along θ shown by the solid line starting at (x_0,y_0) is split into segments l_0, l_1, ..., l_4. Each of these segment lies in a distinct pixel having intensity i_0, i_1, ..., i_4 respectively.

$$\mathcal{F} = \int_0^{\infty}\int_0^{2\pi}\left[I^2(r,\theta) + \lambda\left(\frac{\partial I(r,\theta)}{\partial\theta}\right)^2\right]g(r)rdrd\theta$$

$$C = \int_0^{\infty} g(r)rdr = R_1 - R_0$$

$$C^* = \int_0^{\infty} g^*(r)rdr = \int_0^{\infty} r^2 dr = \frac{R_1^3 - R_0^3}{3}$$

$$\tilde{I}(\theta) = \int_0^{\infty} I(r,\theta)g(r)rdr \qquad (13)$$

Note that

$$\int_0^{\infty}\frac{\partial I(r,\theta)}{\partial\theta}g(r)rdr = \frac{\partial}{\partial\theta}\int_0^{\infty} I(r,\theta)g(r)rdr = \frac{\partial\tilde{I}(\theta)}{\partial\theta}.$$

The radial integrals can be performed numerically. Thus,

$$E = \mathcal{F} + \sum_{p=1}^{N}\lim_{\epsilon\to 0}$$

$$\int_{\theta_{p^{-\epsilon}}}^{\theta_{p+1^{-\epsilon}}}\left[-2T_p\tilde{I}(\theta) + T_p^2 C + \lambda\left(-2re(\theta)\frac{\partial\tilde{I}(\theta)}{\partial\theta} + C^* e^2(\theta)\right)\right]d\theta$$

We can now perform the integrals over $d\theta$ to obtain the following:

$$E = \mathcal{F} + \sum_{p=1}^{N}$$

$$\left[\left(\theta_{p+1} - \theta_p\right)\left(-2T_p\bar{I}_{\theta_p,\theta_{p+1}} + T_p^2 C\right) + \lambda\left(-2re_p\left(\partial_\theta I\right)_{\theta_p} + C^* e_p^2\right)\right]$$

where we have used that $e(\theta) = \sum_{p=1}^{N} e_p\delta\left(\theta - \theta_p\right)$, and

$$\left(\partial_\theta I\right)_{\theta_p} = \frac{\partial\tilde{I}(\theta)}{\partial\theta}\Big|_{\theta_p} = \frac{\partial\int_0^{\infty} I(r,\theta)g(r)rdr}{\partial\theta}\Big|_{\theta_p}$$

$$e_p F\left(\theta_p\right) = \lim_{\epsilon\to 0}\int_{\theta_{p^{-\epsilon}}}^{\theta_{p+1^{-\epsilon}}} e(\theta)F(\theta)d\theta$$

$$\bar{I}_{\theta_p,\theta_{p+1}} = \frac{1}{\left(\theta_{p+1} - \theta_p\right)}\int_{\theta_p}^{\theta_{p+1}}\tilde{I}(\theta)d\theta$$

$$= \Delta\theta\sum_{j=j_p}^{j_{p+1}}\tilde{I}(j\Delta\theta)$$

where the integral $\int_{\theta_p}^{\theta_{p+1}}\tilde{I}(\theta)d\theta$ is adapted to the image lattice, by computing $\tilde{I}(\theta)$ at every interval $\Delta\theta$ and with $(\theta_{p+1} - \theta_p) = (j_{p+1} - j_p)\Delta\theta$.

Appendix B: Numerical Approximation of $\tilde{I}(\theta)$

Using the definition of $g(r) = 1/r$ and (11), we have

$$\tilde{I}(\theta) = \int_0^{\infty} I(r,\theta)g(r)rdr$$

$$= \int_{R_0}^{R_1} I(r,\theta)\left(\frac{1}{r}\right)rdr$$

To carry out the numerical approximation, we define a grid, as shown in Fig. 13, such that each cell in the grid corresponds to a pixel on the image. A ray centered at (x_0, y_0) of length R_1 is broken by the grid into k segments $l_1, l_2, \ldots, l_k$, where each segment lies in a rectangular cell (pixel) with intensity i_j, $j = 1, 2, \ldots, k$. Note that $R_1 = l_1 + \ldots + l_k$. Thus, we have the following approximation: $\tilde{I}(\theta) \approx \sum_{j=1}^{k} l_j i_j$.

Further, if $R_0 > 0$, the segments are $l_m, l_{m+1}, \ldots l_k$, for some $1 \le m \le k$, and $R_1 = l_m + l_{m+1} + \ldots + l_k$, and $\tilde{I}(\theta) \approx \sum_{j=m+1}^{k} l_j i_j$.

Acknowledgments

This research has been supported by AFOSR grant #F49620-96-1-0159 and DARPA grant #MDA972-93-1-0011.

We are grateful to J. Morel and T. Binford for their comments and suggestions. We wish to thank the anonymous referees for their detailed comments that substantially improved the paper. Thanks to R. Chandrasekar for his careful reading of the paper.

References

[1] E.H. Adelson and J.A. Movshon, "Phenomenal Coherence of Moving Visual Patterns," *Nature*, vol. 300, pp. 523-525, 1982.

[2] V. Caselles, B. Coll, and J.M. Morel, *A Kanizsa Programme*, technical report, Univ. de les Illes Balears, Spain, 1996.

[3] T. Cover and J. Thomas, *Elements of Information Theory*. New York: John Wiley & Sons, Inc., 1991.

[4] W. Freeman and E. Adelson, "Junction Detection and Classification," *Proc. ARVO*, 1991.

[5] D. Geiger, K. Kumaran, and L. Parida, "Visual Organization for Figure/Ground Separation," *CVPR*, San Francisco, 1996.
[6] A. Guzman, "Decomposition of a Visual Scene Into Three-Dimensional Bodies," *Proc. AFIPS Fall Joint Computer Conf.*, 1968.
[7] D.A. Huffman, "A Duality Concept for the Analysis of Polyhedral Scenes," *Machine Intelligence*, vol. 6. Edinburgh, U.K.: Edinburgh Univ. Press, 1971.
[8] M.F. Hueckel, "An Operator Which Locates Edges in Digitized Pictures," *J. ACM*, vol. 18, pp. 113-125, 1971.
[9] M.F. Hueckel, "A Local Operator Which Recognizes Edges and Lines," *J. ACM*, vol. 20, pp. 634-647, 1973.
[10] R. Hummel, "Feature Detection Using Basis Functions," *Computer Graphics and Image Processing*, vol. 9, pp. 40-55, 1979.
[11] W. Forstner and E. Gulch, "A Fast Operator for Detection and Precise Location of Distinct Points, Corners, and Centres of Circular Features," *Proc. Intercommission Conf. Fast Processing of Photogrammetric Data*, pp. 281-305, Interlaken, Switzerland, 1987.
[12] J. Bigun, "A Structure Feature for Some Image Processing Applications Based on Spiral Functions," *Compter Vision, Graphics and Image Processing*, vol. 51, pp. 166-194, 1990.
[13] K. Rohr, "Recognizing Corners by Fitting Parametric Models," *Int'l J. Computer Vision*, vol. 9, no. 3, pp. 213-230, 1992.
[14] U. Weidner, "Parametric Information-Preserving Surface Restoration," *Computer Vision—ECCV*, pp. 218-224, 1994. (J.-O. Eklundh, ed., Lecture Notes in Computer Science, vol. 801.)
[15] W. Forstner, "A Framework for Low Level Feature Extraction," *Computer Vision—ECCV*, pp. 383-394, 1994. (J.-O. Eklundh, ed., Lecture Notes in Computer Science, vol. 801.)
[16] R. Deriche and T. Blaszka, "Recovering and Characterizing Image Features Using an Efficient Model Based Approach," *Proc. Computer Vision and Pattern Recognition*, pp. 530-535, New York, June 1993.
[17] G. Giraudon and R. Deriche, "On Corner and Vertex Detection," *Proc. Computer Vision and Pattern Recognition*, pp. 650-655, Hawaii, 1991.
[18] S. Nayar, S. Baker, and H. Murase, "Parametric Feature Detection," *CVPR*, San Francisco, 1996.
[19] K. Brunnstorm, T. Lindeberg, and J.O. Eklundh, "Active Detection and Classification of Junctions by Foveation With a Head-Eye System Guided by the Scale-Space Primal Sketch," *Trans. ISRN KTH/NA/P-91/31-SE, CVAP*, Royal Institute of Technology, 1992.
[20] B.M. Ter Haar Romey, L.M.J. Florack, J.J. Koenderink, and M.A. Viergever, "Invariant Third-Order Properties of Isophotes: T-Junction Detection," *Proc. Scandinavian Conf. Image Analysis*, 13-16 Aug. 1991.
[21] J. Malik, "On Binocularly Viewed Occlusion Junctions," *ECCV*, B. Buxton and R. Cipolla, eds. New York: Springer Verlag, 1996. (J.-O. Eklundh, ed., Lecture Notes in Computer Science, vol. 801.)
[22] J. Morel and S. Solimini, *Variational Methods in Image Segmentation*. Boston: Birkhauser, 1995, pp. 47-58.
[23] D. Mumford and T. Shah, "Optimal Approximation by Piecewise Smooth Functions and Associated Variational Problems," *Comm. on Pure and Applied Mathematics*, vol. 17, no. 5, pp. 557-685, July 1989.
[24] M. Nitzberg, D. Mumford, and T. Shiota, *Filtering, Segmentation, and Depth*. New York: Springer-Verlag, 1993.
[25] D. Beymer, "Junctions: Their Detection and Use for Grouping Images," master's thesis, Massachusetts Institute of Technology, 1989.
[26] D.J. Beymer, "Finding Junctions Using the Image Gradient," *Computer Vision and Pattern Recognition*, pp. 720-721, Lahania, Maui, Hawaii, 3-6 June 1991.
[27] L. Parida, D. Geiger, and B. Hummel, "Kona: A Multi-Junction Detector Using Minimum Description Length Principle," *Energy Minimization Methods in Computer Vision and Pattern Recognition*, pp. 51-65, May 1997. (M. Pelillo and E. Hancock, eds., Lecture Notes in Computer Science, vol. 1,223.)
[28] J. Rissanen, "A Universal Prior for Integers and Estimation by Minimum Description Length," *Annals Statistics*, vol. 11, pp. 416-431, 1983.
[29] Y. Lamdan and H.J. Wolfson, "Geometric Hashing: A General and Efficient Model-Based Recognition Scheme," *Second IEEE Int'l Conf. Computer Vision*, pp. 238-249, 1988.
[30] J. Rissanen, "Stochastic Complexity and Modeling," *Ann. Stat.*, vol. 14, pp. 1,080-1,100, 1986.
[31] J. Rissanen, "Stochastic Complexity (With Discussions)," *J. Royal Statistical Soc.*, vol. 49, pp. 223-239, 252-265, 1987.
[32] E. Simoncelli and H. Farid, "Steerable Wedge Filters for Local Orientation Analysis," *IEEE Trans. Image Processing*, Sept. 1996.
[33] D.L. Waltz, *Understanding Line Drawings of Scenes With Shadows: The Psychology of Computer Vision*. New York: McGraw-Hill, 1972.
[34] S. Wuerger, R. Shapley, and N. Rubin, "On the Visually Perceived Direction of Motion, Hans Wallach: 60 Years Later," *Vision Research*, vol. 25, pp. 1,317-1,367, 1996.
[35] S. Smith and J. Brady, "SUSAN—a New Approach to Low Level Image Processing," DRA technical report TR95SMS1c, Defense Research Agency, Farnborough, Hampshire, GU14 6TD, UK, 1995.

Laxmi Parida is a doctoral candidate in the Courant Institute of Mathematical Sciences, New York University, working in the areas of computer vision and computational biology. Her research interests also include geometric computing, design, and analysis of algorithms.

Davi Geiger holds a BS from PUC-Rio and a PhD from MIT, both in physics. At MIT in 1990, he worked at the Artificial Intelligence Laboratory, developing a thesis in statistical physics and computer vision. From 1990 to 1993, he was at Siemens Corporate Research in New Jersey, working in medical imaging problems. In 1993, he held a visiting position at the Isaac Newton Institute for Mathematics and Applications in Cambridge, England. He joined the Computer Science Department in the Courant Institute of Mathematical Sciences, New York University in 1994 and is currently an assistant professor in computer science and neural science at NYU. Dr. Geiger's research interests are in perception, recognition, and its applications. In particular, his study focuses on mechanisms that are driven by local computations and global solutions, such as the machinery of Markov random fields. He is supported by NSF, AFOSR, and DARPA and has been awarded the NSF CAREER award in 1998. He is a member of the IEEE.

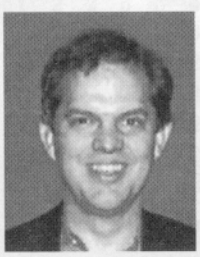

Robert Hummel received a BA in mathematics from the University of Chicago in 1975 and a PhD in mathematics from the University of Minnesota in 1980. From 1980 to 1997, he was at the Courant Institute of Mathematical Sciences, New York University, where he was on the faculty of the Computer Science Department. In 1992-1993, he was a member of the Defense Science Study Group at the Institute for Defense Analysis. From 1995 to 1997, he was a member of the Naval Studies Board. Since 1997, he has been a program manager at the Department of Defense's Defense Advanced Research Projects Agency, where he manages the MSTAR project on automated target recognition technology development.