Junction Detection and Semantic Interpretation Using Hough Lines

Marco Hahn, Norbert Krüger

Lehrstuhl für kognitive Systeme
Institut für Informatik,
Christian–Albrechts–Universität zu Kiel
Preusserstrasse 1-9, 24105 Kiel, Germany
mah{nkr}@ks.informatik.uni-kiel.de

November 28, 1999

Abstract

We present a method for junction detection and semantic interpretation that uses an image and a set of extracted lines as input data. A unified model is used to work on both tasks. The model allows the handling of ambiguous information in images. Our method is fast since we do only evaluate 1D parts of the image.

1 Introduction

Junction detection and classification can be divided into three subtasks. The first one is finding junction hypotheses in images. The next one is about determining a semantic interpretation of a junction. Such an interpretation contains the number of edges meeting at the junction and the orientation of the edges. The last subtasks involves at least two images and the information on the junctions contained in each image. Knowledge of the intrinsic and extrinsic camera parameters allows determining 3D junction location and 3D edge orientation.

A complete assessment of the semantic information contained in a given region of an image is usually not possible, but requires the use of information of the global context of the image. We decided to develop methods that do not make final decisions, but rather extract a set of local interpretations which allow a complete interpretation at later processing stages. Thus, we derive only junction hypotheses and possible assessments of these. If ambiguities exists, these are not eliminated.

Images are a two dimensional entity. Searching for junctions in images is often done by evaluating squared or circular areas. The methods presented, however, work 1D entities, i. e. lines. Therefore their processing time is far shorter and global information can be used. Here we concentrate on junction finding and 2D classification, i. e. determining the number of edges of a junction and their orientation. The 3D interpretation is described in [Hah99]. This paper covers the following topics: junction detection in images and classifying found 2D junctions.

There are two main types of 2D junction detection methods [DG93]: First, there are methods that search for edges and look for intersections of these. The use of chain codes [Jäh97] is one way. Places of maximum curvature in the codes are probable junctions. Techniques of the second type do directly process gray scale images. They either analyze gradient and curvature of gray level surface, like Kitchen and Rosenfeld [KR82] did, or are based on heuristics like the interest operator of Moravec [Mor77]. However, the above-mentioned methods do only determine the location of a junction, they do not deliver any semantic interpretation. The methods presented in this paper use the intersections of lines contained in the image to generate junction hypotheses but to evaluate these hypotheses the image gradient is used.

Methods for 2D junction classification have been proposed by Rohr [Roh92], Michaelis and Sommer [MS94], and others. Rohr uses a parameterized model of a junction and fits it to the gray level surface near a given junction location. His model of the neighborhood consists of a square patch containing several planes of constant gray level. Michaelis and Sommer [MS94] use filters that are steerable in orientation and scale. The complex response of the quadrature pair contains information on the existence and the type of an event. It is possible to distinguish between sharp and smeared events and between lines and edges.

Both methods may use a neighborhood ranging from only a few to several dozen pixels. All three methods require knowledge of the location of a junction, they only determine the semantics of a junction.

The methods presented in this paper use a unified model to tackle the problems of junction detection and classification in 2D images and computing edge orientation in 2D space. A preprocessing stages obtains a list of lines, which is used to generate a set of hypotheses. The junction detection is based on the intersections of lines that partially correlate to edges contained in the image, but does not directly use the gray scale image information. The semantic assessment of the junctions is based both on the list of lines and on the gray scale data. The interpretations are weighted by a confidence measure and are not absolute. Misinterpretations due to processing errors at early stages are therefore correctable at later stages. As long as the input data, the list of lines, is correct and complete, all junction types may be recognized. Due to the 1D processing of parts of the image, the methods work much faster compared to other mentioned techniques. Both problems, junction detection and classification, are handled within the same framework. The other junction classification algorithms presuppose found junctions, while our method finds and evaluates junctions.

The methods are described in the order of their application to images. First, we explain the generation of 2D junction hypotheses. These hypotheses are subjected to a filtering and a merging process. After these steps, junctions are usually evaluated without any knowledge of the underlying gray scale information. Using sets of junctions derived from two images, it is shown in [Hah99] how 3D junctions and edges can be obtained. Thereafter, the results of experiments are given. Finally we discuss the results and possible extensions of the presented methods. Figure 1 shows the processing order and the processed data between different parts of the method presented in this paper.

2 Detecting and Evaluating Junction Hypotheses

Detecting Junction Hypotheses: The detection of junction hypotheses proceeds in two steps. The first step extracts intersections of lines that are not yet sufficient to describe arbitrarily complex junctions. In the second step, we survey the near surroundings of low level junction hypotheses for other hypotheses and merge them if certain preconditions are met. The resulting higher level junction hypotheses can then represent junctions consisting of any number of straight edges.

The input for our 2D methods is a list of lines. Parts of these lines correspond to edges in the image, while other parts are unrelated to the respective image information. The lines are detected using the orientation selective Hough transform [Bal81] with Ackermann's [Ack00] method for finding and extracting the maxima in the Hough transform of an image. (See figure 1, (1).)

Since junctions are intersections of edges, we use the list of Hough lines to determine locations of possible junctions. Each Hough line is intersected with all other Hough lines. If the intersection is within image boundaries, we use the location as a junction hypothesis. Such a junction hypothesis has four edge hypotheses, which we independently assess with a method described below. We call the result of the assessment a confidence. A confidence is always in the unit interval $[0, 1]$, with 1 for a certain edge and 0 for an extremely unlikely edge. Our low level junctions are thus represented by the following feature vector: $(\mathbf{x}; c_1^+, c_1^-; c_2^+, c_2^-)$, where $\mathbf{x}$ is the location, and c^+ and c^- are the confidences

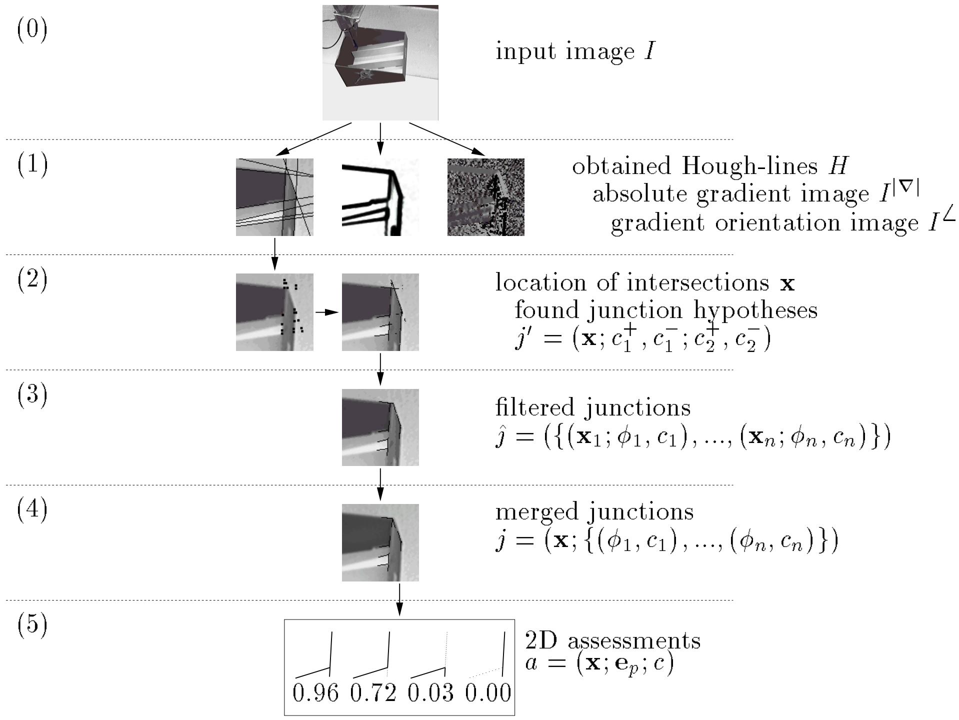

Figure 1: Order of usage of the methods described in this paper. The gradient is computed for the input image and a set of Hough lines is detected. Intersections between all Hough lines are determined and the gray scale information along the Hough lines is used to characterize four edge candidates. The edge candidates are filtered and merged into general junctions that can be assessed.

---

in both directions. The indices 1 and 2 denote both intersecting lines. This step corresponds to part (2) of figure 1.

Planck's Formula: To determine the edge confidences, we assess how the image information confirms the edge hypothesis along each Hough line. The whole half-line need not be assessed: For various reasons are edges blurred near junctions [Bey91]. Furthermore it becomes ever more unlikely that ever more remote parts of the images still contain the edge. Thus we decided to apply a weighting function to the assessments along the Hough line. If we put the origin of the function to the location of the junction hypothesis, the function should be very close to zero for a few pixels, thereafter sharply rise, have a wide maximum, and finally fall slowly down to zero again. We decided to use a weighting function that is a generalization of a formula first used by Planck to describe the thermal properties of a black body. Since the confidences are normalized to the unit interval, the weighting function must integrate to 1. For the normalization we include the first factor, while the second determines the shape of the graph:

$$p(x) = \frac{1}{b!\, c^{b+1}} \cdot \frac{x^b}{e^{x/c}} \tag{1}$$

We can adapt the parameters b and c to change the location and width of the part significantly different from zero. In Figure 2 the effects of different parameters b and c are shown.

Assessing the Edge Hypotheses: The gradient information of an image can be used to find and characterize edges. The gradient orientation $I^{\angle}$ is perpendicular to the edge orientation and the maximum of the gradient magnitude $I^{|\nabla|}$ defines the location of the edge. We use these two features to assess the edges. The gradient magnitude is one of the three factors of the assessment.

The gradient orientation should be perpendicular to the edge orientation as given by the orientation of the underlying Hough line. Due to image noise, the orientation of the gradient is also noisy. We

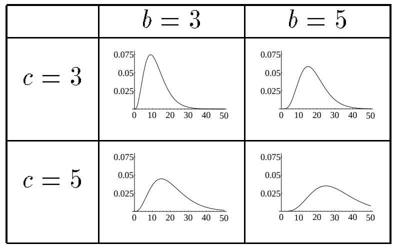

Figure 2: Shape of the weighting formula for different values of b and c.

therefore check whether it is within a certain range of the orientation of the Hough line. If this is the case, the gradient magnitude at the location is taken into account, otherwise not.

The last factor is the weighting function (1). If we put the intersection at the origin of the weighting function and use a parameterized access to the Hough line, we get the following weighting contribution of a single pixel at $\mathbf{x}(t)$:

$$c_{pixel}(t) = p(t) \cdot I^{|\nabla|}(\mathbf{x}(t)) \cdot \left\{ \begin{array}{ll} 1 & \text{iff } \left|\psi - \nabla_{\angle}(\mathbf{x}(t))\right| \leq \delta_{\angle} \\ 0 & \text{iff } \left|\psi - \nabla_{\angle}(\mathbf{x}(t))\right| > \delta_{\angle} \end{array} \right\} \qquad (2)$$

The full confidence is computed by summing the pixelwise evaluations along a certain stretch of the Hough line. We have now arrived at part (2) of figure 1.

The confidences of all edges found in a given image may not fully exploit the unit interval. After we normalized the confidences to the unit interval, we found that almost all real edges have confidences above 0.1, while almost artifacts have confidences below that threshold. Thus we call this level significant confidence and filter all edges that lie below this threshold.

3 Merging Junction Hypotheses

The junction hypotheses as extracted with the methods given so far cannot model arbitrarily complex junctions. Indeed, even very simple common junctions are impossible to describe with such a model: in general, the corner of a cube will result in a junction containing three edges with no one being collinear to any other. Such a junction will result in three Hough lines that have three *different* intersections. Furthermore, of the four edges originating in each junction, usually exactly two will have significant confidences. Each real-world edge will cause two significant edges to be detected. The merging process should handle these problems and result in a junction that contains three edges with significant confidences. To achieve this, the location information has to be re-assessed to yield a single location, and the set of edges has to be filtered and reorganized. (See figure 3.)

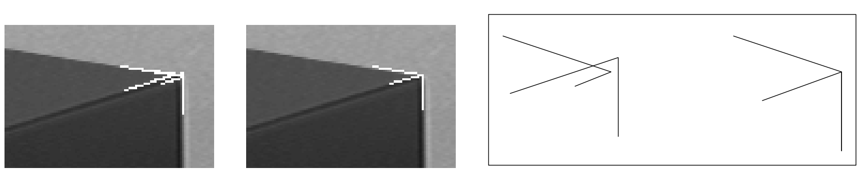

(a) (b) (c)

Figure 3: Example of a junction merge: (a) Two junction hypotheses with two localized edges each (b) A single merged junction hypothesis with three edges (c) Schematic view of the transformation and of the involved edges

Merging Junction Hypotheses: The merging process proceeds in two steps. First we transfer the location and angle information from the Hough lines to the edges. An edge is modeled as $(\mathbf{x}; \phi, c)$

(see figure 1, part (3)), with ϕ being the angle between edge and the horizontal. Junctions hypotheses become sets of localized edges: $\hat{\jmath} = (\{(\mathbf{x}_1; \phi_1, c_1), ..., (\mathbf{x}_n; \phi_n, c_n)\}, n \in \{2,3,4\})$, where n gives the number of edges in the set. Note that after merging all $\mathbf{x}_i$ are equal.

The second step generates our final 2D junction hypotheses $\mathbf{j} = (\mathbf{x}; \{(\phi_1, c_1), ..., (\phi_n, c_n)\})$ from this information. In this structure, the location information $\mathbf{x}$ is bound to the junction and the orientation and confidence values are bound to their respective edges. (See figure 1, part (4).)

Before the merging starts, we filter all those junctions that do not have significant confidences on at least two edges derived from different Hough lines. Obviously, junctions with zero or one significant edge do not correspond to real world junctions. Junctions that have only two significant edges in opposite directions of the same Hough line are of no use, either. They originate in intersections of two Hough lines of which one is describing the edge of an object while the other Hough line does not correspond to any significant image information in that part of the image that contains the intersection.

After this filtering, we sort all junctions by the sum of their edge confidences. The merge candidates are then evaluated in the order of decreasing summed confidence:

For each junction, we check which other junctions are within a certain distance. All such junctions are merged into a single one. The location of the merged junction is the weighted average of the location of the junction and its neighbors. The weighting factors are the summed edge confidences.

The edge information is the union of the edge sets of the junction and its neighbors. The new edge set is sorted and if it contains multiple edges with the same orientation, these are merged into a single one, thereby combining the different confidences into one. This is done with a steerable function that computes a weighted average from a set of input values, the OWA operator [Yag88], for which we adjust the steering parameter so that the operator acts more like a maximum function, since the lower confidences are probably due to non-optimal locations for the confidences evaluation.

Thus we have yielded our merged junction hypotheses (figure 1, part (4)). Figure 3 shows the result of a merge process.

Semantic Interpretation of the Found Junction Hypotheses: The edge information of the merged junction hypotheses can be used to characterize the junction hypotheses. It is possible, though, that one or more of the detected edges do not correspond to real edges. To allow for such uncertainties, our assessment uses the following strategy:

The edge set is divided into two disjunct sets, the pro-edges and the contra-edges. The pro-edges are supposed to be real edges supporting a certain assessment and the contra-edges are supposed to be artifacts contradicting it. If a certain assessments is a good one, the pro-edges will have high confidences, while the contra-edges will have low confidences. Again, each junction hypothesis must have at least two pro-edges.

We divide the edge set into all possible pairs of subsets. For each set of pro-edges $\mathbf{e}_p$, we apply the OWA operator to the confidences of the edges in the set with a behavior close to the minimum function (i. e. $z_{\mathbf{e}_p} < 1$), while for each set of contra-edges $\mathbf{e}_c$ we apply the OWA operator close to the maximum function ($z_{\mathbf{e}_c} > 1$). For a good assessment, the first value will be high and the second will be low. Thus the difference

$$c(\mathbf{e}_p, \mathbf{e}_c, z_{\mathbf{e}_p}, z_{\mathbf{e}_c}) = \max\{\Omega_{z_{\mathbf{e}_p}}(\mathbf{e}_p) - \Omega_{z_{\mathbf{e}_c}}(\mathbf{e}_c), 0\} \tag{3}$$

will be a good measure of the quality of an assessment. The maximum function ensures that very inappropriate divisions of the edge set do not yield a negative confidence.

The full assessment a contains the location of the junction hypothesis, the set of pro-edges, and a confidence as computed with equation (3): $a = (\mathbf{x}; \mathbf{e}_p; c(\mathbf{e}_p, \mathbf{e}_c, z_{\mathbf{e}_p}, z_{\mathbf{e}_c}))$. Similar to discarding edges with very low confidences, we discard assessments with very low confidences. Thus we have arrived at the last step of figure 1.

In most cases there is only one assessment with a high confidence, however ambiguous cases may yield two or more assessments with significant confidences, which may be resolved by later processing stages.

4 Experimental Results

Samples of Detected 2D Junction Hypotheses And of 2D Semantic Interpretations: Figure 4(a) shows the confidences for several edges of figure 4(b). For the junctions B, D and E there are more than one possible assessments. Note that the length of the marked edges corresponds to the confidence, i. e. longer lines represent edges with higher confidences. Missing edges are usually due to missing Hough lines. In table 1 the confidences for each assessment and the orientation and confidence for all marked edges are listed.

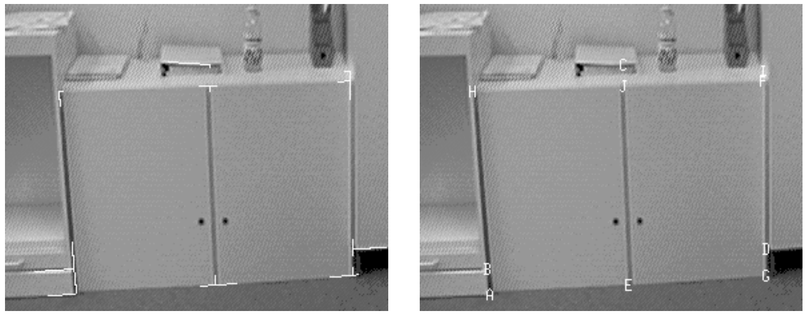

(a) (b)

Figure 4: Example of a semantic interpretation: (a) Found junction hypotheses (b) Label of the junctions; for assessments see table 1

Junction	Confidence of the assessment	1st edge		2nd edge		3rd edge	
		ϕ_1 (degree)	c_1	ϕ_2 (degree)	c_2	ϕ_3 (degree)	c_3
A	0.50	95	0.45	184	0.65	-	-
B	0.56	95	0.64	184	0.91	-	-
	0.40	275	0.15	95	0.64	184	0.91
C	0.32	272	0.11	177	0.95	184	0.91
D	0.26	4	0.79	90	0.23	270	0.17
	0.20	4	0.79	90	0.23	-	-
E	0.30	4	0.53	92	0.22	184	0.35
	0.17	4	0.53	184	0.35	-	-
F	0.33	182	0.28	270	0.46	-	-
G	0.20	4	0.16	90	0.17	184	0.52
H	0.21	275	0.38	2	0.15	-	-
I	0.17	177	0.14	270	0.29	-	-
J	0.19	272	0.17	2	0.19	182	0.26

Table 1: Assessment of junction hypotheses of figure 4. For each junction all assessments with significant confidences are shown. Each assessment consists of two or three edges, for which the orientations ϕ and confidences c of are shown.

5 Summary

We presented a method that does junction detection and classification, together with edge evaluation, in a unified model. The model puts special emphasis on preserving ambiguities. Thus later processing stages have more freedom in assessing an image. They can better combine several features from different parts of the image. Since we work on lines and not 2D parts of the input data, we achieve a short run time behavior. Our model is easily extended to 3D entities and allows determination of edge orientations in 3D [Hah99]. In current work, we aim to use the junction classifications together with other visual entities such as line segments, color, and texture to represent, recognize and manipulate objects.

Acknowledgement

We would like to thank Marcus Ackermann and Bodo Rosenhahn whose work at the software library KiViGraP was very helpful for our simulations. For technical support we would like to thank Henrik Schmidt.

References

[Ack00] Ackermann, Marcus. *3D-Rekonstruktion mit Hilfe des Bewegungszusammenhangs*. Diplomarbeit, Institut für Informatik und Praktische Mathematik, Christian-Albrechts-Universität, Kiel, 2000.

[Bal81] Ballard, Dana H. *Generalizing the hough transform to detect arbitrary shapes*. Pattern Recognition, 13(2), 1981.

[Bey91] Beymer, David J. *Finding junctions using the image gradient*. Technischer Bericht A.I.Memo No. 1266, Massachusetts Institute of Technology, 1991.

[DG93] Deriche, Rachid und Giraudon, Gerard. *A computational approach for corner and vertex detection*. International Journal of Computer Vision, 10(2):101–124, 1993.

[Hah99] Hahn, Marco. *Semiglobale Verfahren zur Generierung von Eckpunkthypothesen in 2D und 3D*. Diplomarbeit, Institut für Informatik und Praktische Mathematik, Christian-Albrechts-Universität, Kiel, 1999.

[Jäh97] Jähne, Bernd. *Digitale Bildverarbeitung*. Springer-Verlag, Berlin, vierte Auflage, 1997.

[KR82] Kitchen, L. und Rosenfeld, A. *Gray-level corner detection*. Pattern Recognition Letters, S. 95–102, 1982.

[Mor77] Moravec, Hans P. *Towards automatic visual obstacle avoidance*. In *Proceedings 5th International Joint Conference on Artificial Intelligence*, S. 584. Cambridge, MA, 1977.

[MS94] Michaelis, Markus und Sommer, Gerald. *Junction classification by multiple orientation detection*. In Eklundh, Jan-Olof, Hg., *European Conference on Computer Vision*, Nummer 801 in Lecture Notes in Computer Science, S. 101–108. Springer-Verlag, 1994.

[Roh92] Rohr, Karl. *Recognizing corners by fitting parametric models*. International Journal of Computer Vision, 9(3):213–230, 1992.

[Yag88] Yager, Ronald R. *Ordered weighted averaging aggregation operators in multi-criteria decision making*. IEEE Transactions on Systems, Man and Cybernetics, 18:183–190, 1988.