Use of Belief Networks for Modeling Indoor Environments*

R. Liscano
Ramiro.Liscano@nrc.ca

S. Elgazzar
elgazzar@iit.nrc.ca

Institute for Information Technology
National Research Council
Ottawa, Ont. K1A 0R6
CANADA

A. K. C. Wong
akcwong@watnow.uwaterloo.ca
Dept. of Systems Design
University of Waterloo
Waterloo, Ont. N2K 3G1
CANADA

Abstract

This article introduces an approach, based on Bayesian Networks, for the grouping of 3-D surfaces extracted from data obtained by a laser ranging sensor. A methodology for the specification of the network is presented along with an approach for determining the conditional probabilities. Determination of the conditional probabilities is based on a compatibility function that measures the uncertainty in the quality of fit of the data to a model of the features in the scene. Several compatibility functions for the grouping of 3-D surfaces are presented. These are coplanarity, parallel, planarity, and proximity. These compatibility functions are used with a Bayesian Network in determining belief values of possible groupings among the surfaces, in particular grouping for continuous surfaces and corners. This operation is a form of perceptual grouping of three dimensional data and is akin to the previous studies in perceptual grouping for two dimensional images.

1 Introduction

Research in the domain of modeling using 3-D data has primarily focussed on the extraction of 3-D surfaces and volumetric primitives for the purpose of either object recognition or creating more precise models from 3-D sensory data of machined parts [12, 4]. These type of objects can easily be carried and placed in a controlled environment and scanned using a high resolution active sensor. This is significantly different from the modeling of large indoor environments where it is necessary to bring the sensor to the environment, changing the characteristics of the sensed data dramatically. The result of this is that nearly all scans taken in these environments consist of sparse data and it becomes necessary to develop algorithms that can hypothesize the existence of surface continuity and intersections among surfaces. This may appear to be as simple as relaxing the tolerances used for computing the surface models from the sensory data, but it is more complicated than that. Objects occluded from the sensor and missing data make it necessary to use knowledge about the environment to hypothesize the existence of more complex surfaces.

Because of the larger domain in which the sensor is operating in, research in the modeling of indoor environments has primarily focussed on the incremental synthesis of sensor views and/or position estimation of the sensor [16, 15, 1]. For these systems to become viable tools for Computer Aided Design (CAD) it is necessary to develop approaches that hypothesize the formation of more composite features from the surfaces. Previous attempts in this domain [10, 14, 6] have integrated intensity data with range data to help define the boundaries of surfaces extracted from the 3-D data, and then used a set of heuristics to decide what surfaces should be joined. In most circumstances these heuristics are a set of rules with predefined thresholds that determine if the surfaces should be joined. In this article a Bayesian Network is proposed to manage the uncertainty associated with such decisions. A Bayesian Network offers a unified approach to the specification of relationships among surfaces as well as a method for computing a belief value in the existence of a compound feature given the evidence from the sensory data.

This article proposes definitions of proximity, planarity, and coplanarity of planar 3-D surfaces. This is akin to the idea of extending the 2-D perceptual organization rules to 3-D sensory data, except that it is not possible to associate these measurements to

*NRC Number 40174.

human perception since our sensing is primarily two-dimensional. These definitions are 3-D surface adjacency relationships that can be used to aid in the grouping of the surfaces and to hypothesize on how these surfaces are related among each other. For example, for sparse 3-D data the adjacency relationship between other surfaces is not directly evident, due in principal to missing data, but sufficient information may exist to hypothesize that the surfaces are neighbours. A Bayesian Network (BN) is used to compute belief values for the grouping of the surfaces into composite features. The network defines the causal relationship from composite features to the elements that form the composite features while the conditional probabilities are computed from geometrical measurements of proximity, planarity, and coplanarity between the surfaces.

Bayesian Networks have been used in the domain of computer vision for several applications. Some particular ones worth citing are, object recognition [8], multi-agent vision systems [7], and perceptual grouping for 2-D images [13]. The basis of these Bayesian Networks are all similar in that they take the view that a physical object and its relation to other objects will cause certain types of features to be detected by the sensor. They differ in the manner that the feature sets have been defined and how the conditional probabilities are computed. The Bayesian Network defined in this article focuses on the grouping of features into compound geometric features. This is similar in concept to the Perceptual Inference Network (PIN) developed by Sarkar and Boyer [13], except it has been developed for the grouping of 3-D surfaces instead of 2-D edges and the network has been designed by invoking a causality relation among the feature sets. The result is the elimination of duplicate similar nodes that can occur when designing the network based solely on the accumulation of evidence.

2 A Bayesian Network for the Grouping of Surfaces

This particular implementation of a Bayesian Network is based on the modeling of the formation of composite geometrical features from other fundamental features computed from the sensory data. The Bayesian Network allows the encoding of the expected formations that the sensor may detect when viewing a composite formation and infer from those features a belief value of the existence of the composite feature.

2.1 The Network's Structure

A Bayesian Network is a directed acyclic graph with the nodes representing a hypothesis of the existence of a proposition and the arcs signifying the causal relationship from one proposition to the next. This results in a simple model of the universe but one that can be useful for perceptual grouping, as was demonstrated in the Perceptual Inference Network [13].

In this particular Bayesian Network the nodes represent the existence of geometric formations resulting from the grouping of the 3-D sensory data and the arcs connecting the nodes represent a causal relationship among the formations. For example, a formation of 2 coplanar surfaces causes the existence of 2 parallel planar surfaces. This direction of thinking is generally the reverse of that used for sensory data perception, because in that domain the focus is on the accumulation of evidence (sensory data) towards the modeling of that data.

The steps for creating the structure of the network are the following:

- To decide on what geometric formations are desired to be represented by the network and the scope of the variable that will represent that formation. Currently only discrete variables are allowed and in most cases a Boolean one suffices, i.e. this represents the existence of the formation as TRUE or FALSE.
- To decide on the correlation among the geometric formations and on the hierarchy that represents which formations have direct causal effects on others. A directed acyclic graph should be formed which represents the cause-effect relations among the geometric formations.

Each parent node in the network can be considered as a hypothesis of a composite feature that resulted in the evidence represented by the children nodes. A similar interpretation of the structure of the network is to consider the parent nodes of the network as composite geometrical features formed by imposing constraints on the features defined by their children.

The following example defines a Bayesian Network for the detection of corners and continuity among a set of planar surfaces extracted from 3-D sensory data. The process of determining the existence of corners and continuity also classifies the surfaces as coplanar, parallel, planar, and proximal. These are the parent nodes of the Bayesian Network shown in figure 1.

Evidence from the sensory data flows from the end nodes to the root nodes commencing with detection of 3-D data points, to the formation of planar surfaces, then into parallel surfaces, and coplanar surfaces, and finally into a pair of continuous surfaces or a corner.

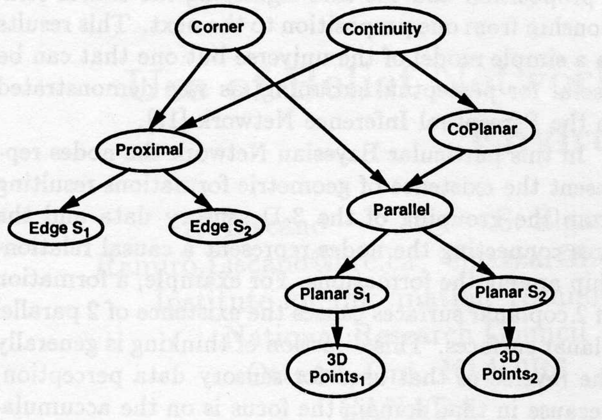

Figure 1: Example of a Bayesian Network for the detection of corners and continuity in planar surfaces.

In a similar fashion edges of the surfaces is used to determine the proximity among the surfaces which has a direct correlation on the belief value of the formation of a corner or surface continuity. Each node also stores the conditional probabilities (joint probability) associated with the formation of that node given the evidence available from the children nodes.

2.2 The Conditional Probabilities

It is still necessary to decide on a set of conditional probabilities for each node, of which 2 approaches are possible: the conditional probabilities can represent the state of the nodes conditioned on their parents (causal), or the state of the nodes conditioned on their children (effect). In this example, grouping surfaces for 3-D modeling, the conditional probabilities are more intuitive if represented as effect than causal. These conditional probabilities then represent the constraints imposed on the features that form the parent node. When new evidence appears at a node, it is represented as a conditional probability and assimilated and propagated through the network using the approach described by Pearl [11].

The conditional probabilities are developed from a set of functions, known as compatibility functions, that measure how well a set of features match to the assumed parent geometrical feature. They are in effect a measure of evidence of the existence of a composite geometrical feature, for example the formation of a corner from a set of surfaces. This is similar in concept to the system developed by Levitt et al. [8] in that the belief values of the nodes in the Bayesian Network represents a probability of the existence of that feature or combination of features. It differs in that this Bayesian Network is not designed to reflect the hierarchical structure imposed by manufacturing parts, but more on the features perceived by the sensor.

We would now like to define a mapping between the compatibility function, $f_u(\mathrm{e})$, and a value for the conditional probability $P(\mathbf{u}|\mathrm{e})$, where $\mathbf{u}$ represents a model for the compound feature and e represents the evidence accumulated from the sensor. This mapping should exhibit the following properties:

- It should be bounded between the interval $(1, 0)$.
- It should be a decreasing monotonic function.

There are several of these type of mappings and in this particular implementation we propose the declining S-Curve function, shown in figure 2 and represented mathematically as,

$$P(\mathbf{u}|\mathrm{e}) = \begin{bmatrix} 1 & \rightarrow & f_u(\mathrm{e}) = 0 \\ 1 - 2(f_u(\mathrm{e})/\gamma)^2 & \rightarrow & 0 \leq \beta \\ 2((f_u(\mathrm{e}) - \gamma)/\gamma)^2 & \rightarrow & \beta \leq \gamma \\ 0 & \rightarrow & f_u(\mathrm{e}) \geq \gamma \end{bmatrix}$$

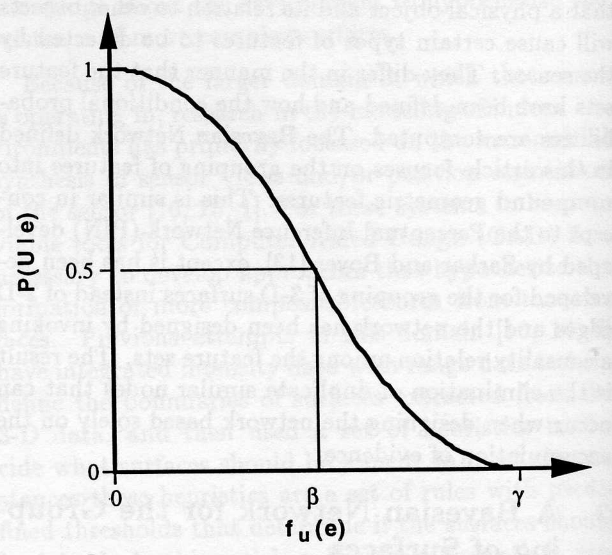

Figure 2: A declining S-curve.

The compatibility function $f_u(\mathrm{e})$ in most cases is a linear function that when equal to the value 0 signifies a perfect match with the model $\mathbf{u}$. Mapping the compatibility function into an S-Curve function offers several advantages to the user:

- It separates the process of defining the compatibility function from the assignment of a conditional probability value;

- It introduces a subjective measure into the definition of the conditional probability that can act as a method for defining the resolution of a match of the geometric features with the model;

This added flexibility in determining $P(\mathbf{u}|\mathbf{e})$ is akin to the concept of specifying a resolution to the perception process. The S-Curves are the primary interface to the user and correspond to the manner in which a user can decide the bounds on the compatibility functions. The values for β and γ are determined subjectively or through experimentation. In the examples given in this article they have been determined subjectively.

3 Compatibility Functions

In this section we will introduce several compatibility functions that measure the quality of the planar surfaces, how parallel the surfaces are to each other, the coplanarity of the parallel surfaces, and the proximity of two surfaces . These compatibility functions are enough to determine belief values that the surfaces form a corner or are a continuation of the same surface.

3.1 Planar Surfaces

Let us assume that the data is normally distributed about a planar surface and therefore its distribution function can be represented by the following equation,

$$p(d_i|\mathbf{u}) \quad = \quad \frac{1}{\sqrt{2\pi}\sigma_i} exp\left(\frac{1}{2}\frac{d_i^2}{\sigma_i^2}\right),$$

where $\mathbf{u}$ is the model for the planar surface and d_i is the distance of point i from the surface $\mathbf{u}$. σ_i is the variance of the Gaussian noise of the individual points i, which in our situation is unknown, and we set to unity.

A commonly used compatibility function for determining a surface of fit to the data points is given by,

$$f_{plan}(d_i) \quad = \quad \frac{1}{n_s}\sum_i d_i^2.$$

where n_s is the number of points in surface s.

3.2 Parallel and Coplanar Surfaces

The compatibility functions for parallel surfaces and coplanar surfaces involve the comparison of two planar surfaces. A measure of parallelism between two planar surfaces can be derived by evaluating the length of the vector computed from the cross product of the normals to the surfaces. This leads to the following geometrical compatibility function,

$$f_{prll}(\mathbf{s}_i, \mathbf{s}_j) \quad = \quad |\mathrm{N}_{\mathbf{s}_i} \times \mathrm{N}_{\mathbf{s}_j}|,$$

where $\mathrm{N}_{\mathbf{s}_i}$ and $\mathrm{N}_{\mathbf{s}_j}$ are the normals corresponding to the planar surfaces $\mathbf{s}_i$ and $\mathbf{s}_j$.

Coplanar surfaces are a more restricted case of parallel surfaces in which the two surfaces are in fact the same surface if the boundaries were removed and the surfaces extended. To determine coplanarity the surfaces must be parallel to each other and the angle between the normals of the surfaces and the line joining the center points of the two surfaces is approximately 90°. A geometrical compatibility function based on the dot product between one of the surface normals and the vector joining the two centre points of the surfaces can be defined as an added constraint to the parallel constraint. We propose the additional following equation as a measure of coplanarity,

$$f_{copl}(\mathbf{s}_i, \mathbf{s}_j) \quad = \quad < \mathrm{N}_{\mathbf{s}_i}, C_{\mathbf{s}_j} - C_{\mathbf{s}_i} >,$$

where $C_{\mathbf{s}_j}$ and $C_{\mathbf{s}_i}$ correspond to the location of the centers of the planar surfaces $\mathbf{s}_j$ and $\mathbf{s}_i$ respectively.

3.3 Surface Proximity

Defining a proximity compatibility function for 3-D surfaces is challenging, because the definition of proximity is not unique for this type of data. Defining a proximity measure requires one to decide on a viewing direction to the data and on a definition of distance between the surfaces. This is ideal for sensory data that still maintains its scanning sequence because the viewing direction is implicitly maintained in its representation as a 2-D image of 3-D points. The proposed approach defines a compatibility function among the surfaces by measuring the distance between the boundary of two surfaces. Our approach is similar to Fan's et al. [5] proximity measure for 2-D lines except it is applied to 3-D data. In this case the lines are replaced by surfaces and the gap by the area between the edges of the surfaces. With this in mind we propose the following measure of proximity,

$$f_{prox}(\mathbf{e}_i, \mathbf{e}_j, \mathbf{s}_i, \mathbf{s}_j) \quad = \quad \frac{n_g(\mathbf{e}_i, \mathbf{e}_j)}{n_{\mathbf{s}_i} + n_{\mathbf{s}_j}},$$

where $n_g(\mathbf{e}_i, \mathbf{e}_j)$ is a function that computes the size of the gap between the planar surfaces based on the edges of the surfaces $\mathbf{s}_i$ and $\mathbf{s}_j$, and $n_{\mathbf{s}_i}$ and $n_{\mathbf{s}_j}$ correspond to the number of points in surfaces $\mathbf{s}_i$ and $\mathbf{s}_j$ respectively. The procedure for determining the size of the gap has been described in greater detailed by Liscano et al. [9].

3.4 Corners and Continuous Surfaces

At this point there does not exist any other added constraints in hypothesizing the formation of corners or continuous surfaces except for combining the evidence from the children nodes. Therefore, it is not necessary to specify a compatibility function and the conditional probabilities for the formation of these compound features can be directly specified as is done in section 4.

4 Quantifying the Belief Network

Before the network can compute the belief values for the nodes it is necessary to compute the conditional probabilities and instantiate the root nodes with evidence. In this section, a detailed example will be given demonstrating how the network values are instantiated and the belief values are propagated.

From section 2 the network appears like that in figure 3 except that in this figure the arrows are shown pointing in the direction of evidence and the necessary conditional probabilities have been specified along the links.

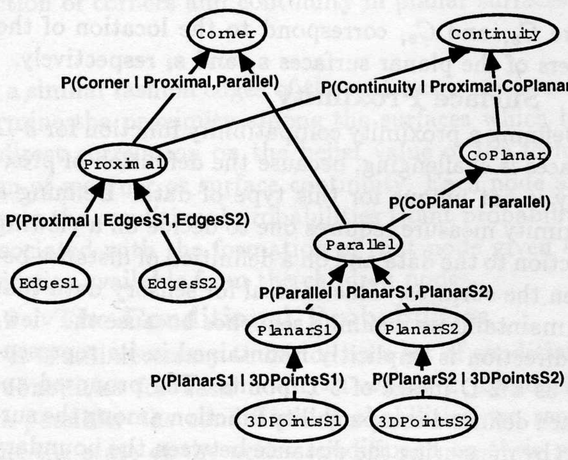

Figure 3: Flow of evidence in the Bayesian Network for the detection of corners and surface continuity.

The variables for the network are Boolean variables which signify the existence of a particular feature or grouping of features. In this particular case the only known variables are the 3-D points the edges extracted from the intensity data of the surfaces. The rest of the variables are treated as unknowns, and when the network is executed belief values of the existence of the compound features are computed.

In section 3 several compatibility functions were defined that correspond to the conditional probabilities at each node. For each pair of surfaces in a scene there exists an instantiation of the Bayesian Network along with the corresponding conditional probabilities. These conditional probabilities must be determined at a node for all the states the node has and conditioned on all possibilities of the evidence coming into that node. Each node has 2 states, True or False, so therefore the number of conditional probabilities is 2^{1+N_c} where N_c is the number of children nodes associated with a particular node. For most of these states the conditional probabilities can be set to 0 since they are not relevant. For example, for the node *Parallel*, eight conditional probabilities must be specified, but those associated with the non-existence of planar surfaces are not relevant so this reduces the conditional probability for the formation of parallel surfaces to the following two,

$$P(\mathbf{Parallel}|\mathbf{PlanarS}_i, \mathbf{PlanarS}_j) = \mathcal{S}(f_{prll}(\mathbf{s}_i, \mathbf{s}_j); 0, \beta, \gamma)$$
$$P(\neg\mathbf{Parallel}|\mathbf{PlanarS}_i, \mathbf{PlanarS}_j) = 1 - \mathcal{S}(f_{prll}(\mathbf{s}_i, \mathbf{s}_j); 0, \beta, \gamma)$$

This same argument can be applied to the computation of conditional probabilities for the other nodes.

The detection of corners and continuity among surfaces is unique in that there does not exist any compatibility function. In this situation the conditional probabilities have to be directly specified in the following manner,

$$P(\mathbf{Corner}|\mathbf{Proximal}, \neg\mathbf{Parallel}) = 1.0$$
$$P(\mathbf{Continuity}|\mathbf{Proximal}, \mathbf{CoPlanar}) = 1.0$$

All other conditional probabilities for the detection of corners and continuity among surfaces are set to 0.

The ultimate desire is to compute the belief in the formation of the compound features represented by the nodes in the network, in particular the formation of corners and the continuity of the surfaces. This belief value, as defined by Pearl [11], is the conditional probability of the existence of the compound feature conditioned on particular states of the other feature nodes.

5 Experimental Results

This section covers the details of an experimental implementation of the Bayesian Network for the detection of corners and continuous surfaces utilizing data acquired from a compact laser camera called BIRIS [2]. The output of the camera consists of one acquisition of 256 range and intensity values along a projected plane of light. The system has been calibrated up to a range of 3 m and will function up to a 5 m range by extrapolating from a calibration table. To acquire more data than that of a single acquisition the BIRIS sensor was mounted onto a pan and tilt unit. The ability to tilt the sensor is crucial in being able to acquire more data, since the field of view of the sensor is fairly limited. When the Biris Laser Scanner is panned at a constant speed for a fixed tilt angle the result is a rectangular image of dimension 256 x N_{aq}, where N_{aq} is the number of acquisitions taken during the panning sequence. This process can be repeated for several different tilt angles so that a number of registered scans can be acquired. This article focuses primarily on the grouping of segments in a single scan.

Figure 4 (a) is an example of a typical scan taken by the BIRIS Laser Scanner of part of a laboratory room at the Institute for Information Technology. This image size is 256 x 256 pixels and was taken using a tilt angle of 0° and a pan angle of 60°. Any dark pixel values (null intensity values) are locations where it was not possible to register a reliable signal either because objects in that vicinity are beyond the range of the sensor or the energy of the returned signal was too weak to register. Each pixel location in the image also corresponds to a 3-D location of that point in space, the result is an image with registered intensity and range values. To get a better understanding of this, figure 4 (b) is an isometric view of the 3-D points acquired from the sensor.

Note that the intensity image in figure 4 (a) is inverted to that of the 3-D data points shown in figure 4 (b). This is caused by the fact that the sensor was scanned in a counterclockwise direction, from right to left, and the data has been displayed from first acquisition to last. This has no bearing on the computed results.

5.1 Planar Surfaces and Boundaries

We will describe briefly in this section the approach taken for the computation of the planar surfaces from the sensory data along with their respective boundaries. The approach for computing a planar model from the data has been previously discussed in greater detail in Boulanger's [3] article, and only a brief description is offered here. The approach taken for estimating the boundary has been outlined in Liscano et al. [9].

Planar surfaces are extracted from the 3-D points using an algorithm based on a hierarchical segmentation procedure which starts with small planar regions constrained by the detection of depth discontinuities and groups these smaller regions into larger regions until the accumulated error of the grouping operation is beyond a predefined threshold value. The algorithm uses a Bayesian decision theory to determine if the planar neighbouring regions should be grouped into a larger region.

The results of performing the segmentation on the range data in figure 4 are shown in figure 5 as both an intensity image and an isometric view of the 3-D data. Again similar to figure 4, the intensity and 3D view are inverted. The individual segments in the intensity image are shown using separate gray values to represent each segment. The advantage of maintaining the points in an image is that one can take advantage of the already established ordering in the 2-D image.

Surface boundaries are extracted by detecting high curvature points along the edges of the surfaces in the image plane and joining these points with straight lines. Some degree of filtering is required along the boundary points. Also, edges with many consecutive high curvature points are approximated by a straight line. The result is an approximation for the boundary of the surface that follows closely the edges that appeared straight and clips off edges that are too ragged. This is not a problem, since ragged edges are primarily a consequence of unreliable data.

These special high curvature points are used as control points in defining a polygonal surface which can be used as a method for representing a planar surface in 3-D space. Extracting the boundaries of the surfaces results in far more fragmentation than when the planar surfaces are extracted. This is due primarily to the complete separation of the surfaces by defining two boundaries between some surfaces where at one point there was only continuous data points. Figure 6 (a) shows the extracted boundaries associated with each of the segments. Figure 6 (b) is the same image but the polygonal corner points are highlighted in white while the edge is in gray.

At this point the fundamental features have been extracted and it is possible to compute the conditional probabilities from the compatibility functions described in section 3. The experimental results are presented in the following section.

5.2 Instantiating the Network

The system implemented is not yet fully automated: the conditional probabilities were computed by applying the compatibility functions from sections 3 to the sensory data and manually transferred these values into the Bayesian Network. The following examples demonstrate results from the network applied to the sensory data depicted in figures 5 and 6. For each compatibility function the parameters for the S-Curve have to be determined before the conditional probabilities can be computed. The S-Curve parameters used were the following:

Compatibility	β	γ
proximity	0.25	0.50
coplanarity	10	20
parallel	10	20
planar	0.578	1.156

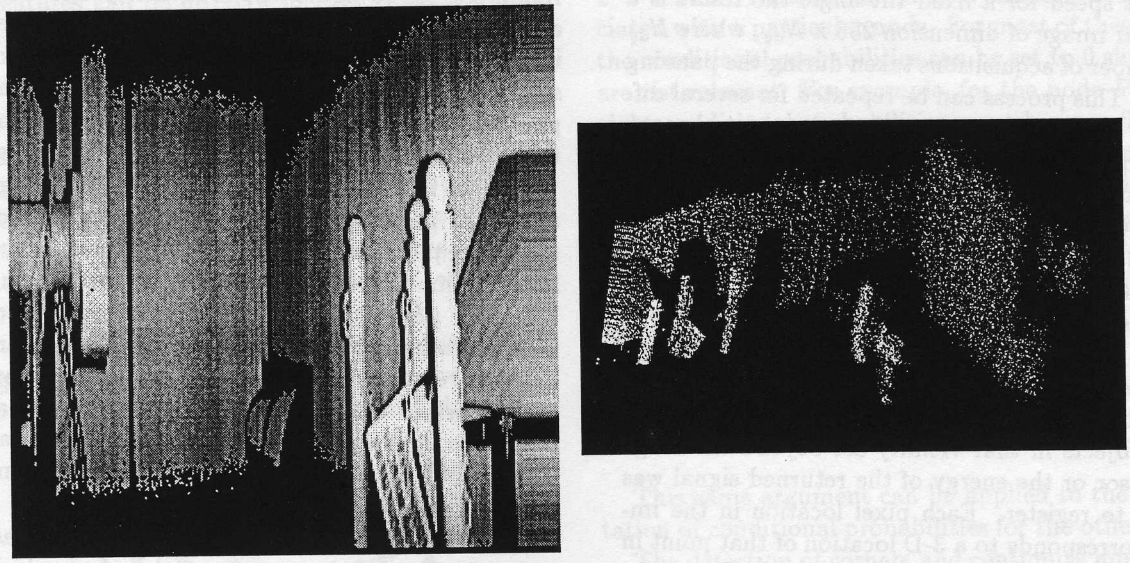

(a)

(b)

Figure 4: Intensity image and 3-D data of a scan from the BIRIS sensor.

These values were subjectively selected but the selection follows a certain reasoning. The proximity value of 0.25 represents the desire to consider surfaces with gaps smaller than 1/4 the sum of the areas of the 2 neighbouring surfaces as proximal, anything beyond 0.5 is not proximal. Surfaces are considered parallel if the surface normals are within 10° with respect to each other, and they are not parallel if beyond 20°. Coplanarity is similar to parallel surfaces except the angular measurement is not between the surface normals but among one surface normal and the vector defined by the surfaces' centroids. The parameter values for the quality of the surface being considered as a plane was determined by computing the average in the variance of the data points for all the surfaces.

5.3 Results

The Bayesian Network was instantiated on all pair of neighbouring surfaces, and figure 7 shows the results of running the network using surfaces 16 and 17 from figure 5 (a). The histograms in the nodes represent the belief value of the existence of the feature represented by the node. The following table displays computed belief values in the formation of Corners (Cor) and Continuous (Con) surfaces:

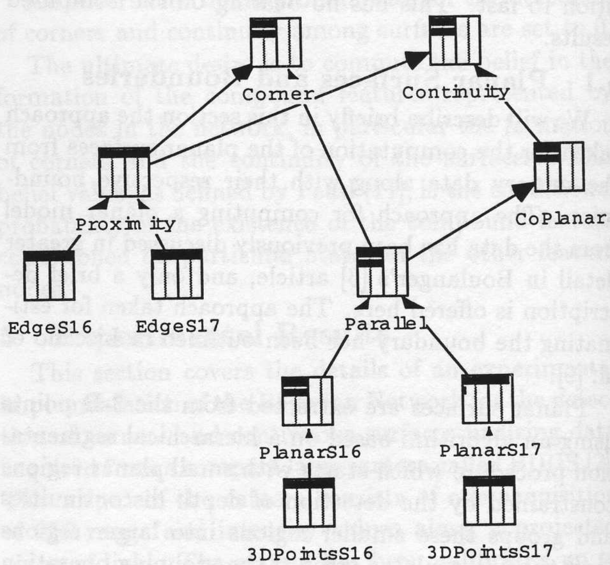

Figure 7: Results of the Bayesian Network using surfaces S_{16} and S_{17}.

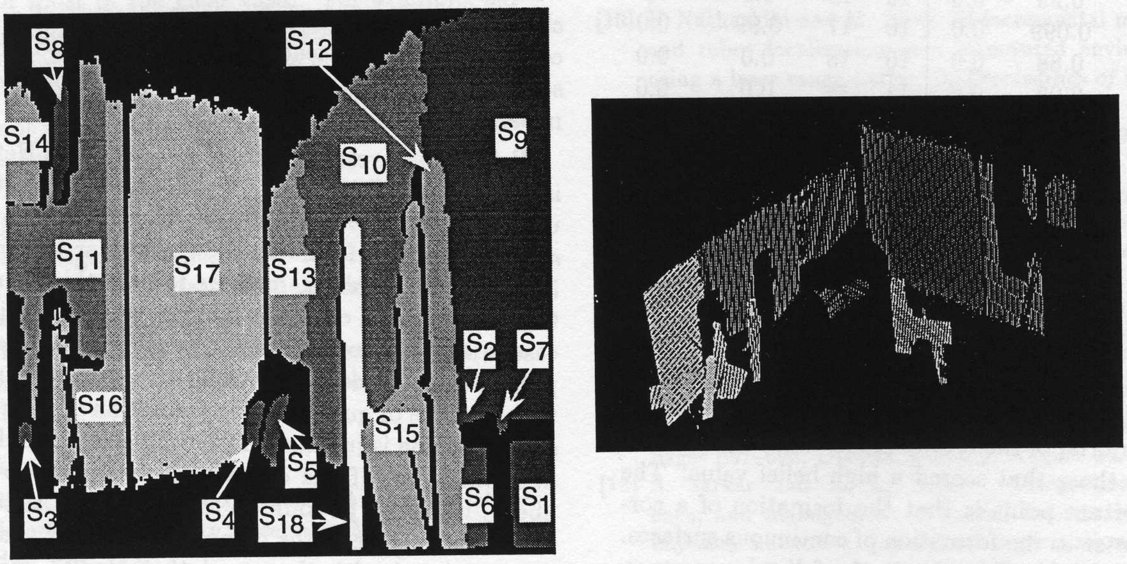

(a) (b)

Figure 5: Surfaces extracted from the 3-D Points, shown as an intensity image an isometric view.

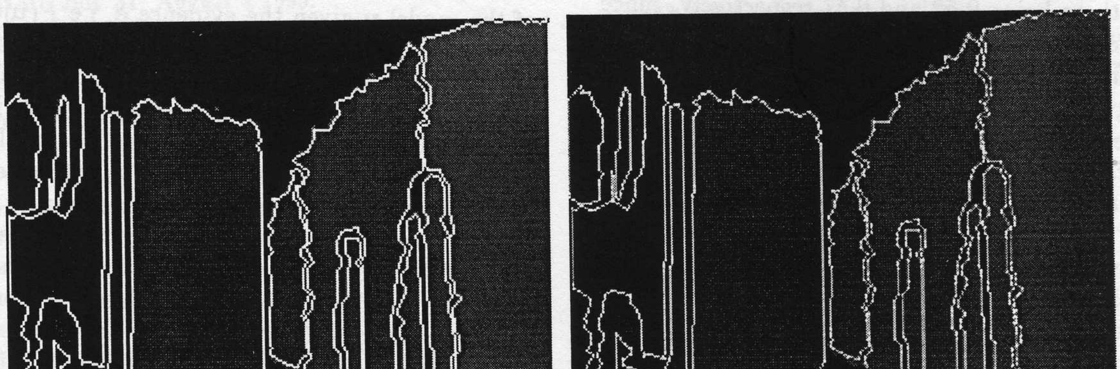

(a) (b)

Figure 6: Boundaries extracted from the surfaces.

S_i	S_j	Cor	Con	S_i	S_j	Cor	Con
1	6	0.57	0.0	10	13	0.0395	0.8805
1	9	0.99	0.0	10	15	0.0	0.0
4	5	0.099	0.0	10	17	0.99	0.0
4	17	0.88	0.0	10	18	0.0	0.0
5	10	0.98	0.0	11	14	0.0	0.0
5	13	1.0	0.0	11	16	0.0	0.0
5	17	0.57	0.0	12	15	0.0	0.0
6	9	1.0	0.0	13	17	0.99	0.0
9	10	1.0	0.0	15	18	0.0	0.0
9	12	0.57	0.0	16	17	0.52	0.47
10	12	0.6897	0.0	17	18	0.0	0.0

Note that not all surface pairs are represented, only the neighbouring surfaces are considered and surfaces that are significant in size. For example surface 2 was quickly eliminated from any further computations.

The majority of the results are intuitively correct in particular those that scored a high belief value. The first important point is that the formation of a corner is counter to the formation of continuous surfaces. The belief values reflect this in the following manner. When the belief of the formation of a corner is high the belief in the formation of a continuous surface is low, and the converse is true. Also surfaces that may not form continuous surfaces or be good candidates for corners are shown to have both low corner and continuous belief values. In some cases the information is ambiguous, surfaces S_{16} and S_{17} have similar values for the formation of a corner and continuous surfaces these being 0.52 and 0.47 respectively. Since the network has no hidden nodes it is possible to try and determine the cause of such close values. In this example, one can see from figure 7 that the principal cause came about from the low value in the belief that surface S_{16} is a planar surface and this is due to the quality of fit of the points to the surface. Most surfaces reflected consistency in their results. For example, surfaces S_{13} and S_{17} are corners, S_{10} and S_{13} are continuous and therefore S_{10} and S_{17} should be corners and are so.

6 Discussion and Conclusions

This article presented an example of the use of a Bayesian Network for the grouping of 3-D surfaces into either corners or continuous planar surfaces. An approach for the specification of the Bayesian Network was presented that used the idea of grouping simple features into compound features by imposing an added constraint to the geometry of the surfaces at each level of the network. Although the concept of using constraints for object recognition and feature grouping is not unique it is more common to apply a constraint satisfaction solution to this problem than to represent it a causal network. Also an approach for the computation of the conditional probabilities was presented, which uses compatibility functions that measure how close a set of features is to the formation of a more complex compound feature. This was tested on 3-D sensory data captured from a typical indoor environment.

Using causal relations to define a Bayesian Network is not as easy at it seems. For example, the network used in this article treats the edges extracted from the surfaces and the formation of the planar surfaces as independent variables. In some respect this is true in that the evidence comes from different sources, in one case it is intensity data and in the other the actual 3-D points. Also the edge of a surface does not contribute to a measure of quality in the formation of the planar surface so one is justified in treating them as independent variables. If one examines how the edges come about, it is easy to come to the conclusion that they are a consequence of the planar surfaces and therefore are correlated with them and that results in a link from the planar surface variable to the respective edge variable of that surface. This may be a possible alteration to the Bayesian Network used in this article, but leads to another complication in the determination of the conditional probabilities.

Even though the network was designed with causal relations in mind, the conditional probabilities have been specified as evidential knowledge of the existence of the model $\mathbf{u}$ given the evidence $\mathbf{e}$, i.e $P(\mathbf{u}|\mathbf{e})$. Another approach is to define causal conditional probabilities. This has the advantage in that it is easier to introduce new evidence from other sources into the network by specifying the conditional probability that the new evidence came from the hypothesized compound feature instead of having to compute a probability distribution conditioned on the evidence. It unfortunately complicates matters when one is considering the joining of two features into one compound feature, similar to that of the joining of 2 surfaces as one parallel entity. If one now tries to compute a conditional probability in the causal direction, this leads to, instead of having one equation that compares the angle between the two surfaces, it is necessary to determine two equations for each surface, that measure the probability of the particular surface resulting from the existence of a parallel surface formation. This is currently under investigation as an extension to the current work.

Typical to any systems that manage uncertainty there is no right or wrong answer and results have to be compared to human judgment. Comparing the net-