Informative Views and Active Recognition

Tal Arbel, Frank P. Ferrie, and Peter Whaite

TR-CIM-94-16 Dec. 1994

3D Vision Group
Centre for Intelligent Machines
McGill University
Montréal, Québec, Canada

Submitted to *5th International Conference on Computer Vision*, Nov. 1994

Postal Address: 3480 University Street, Montréal, Québec, Canada H3A 2A7
Telephone: (514) 398-6319 Telex: 05 268510 FAX: (514) 283-7897
Email: cim@cim.mcgill.ca

Informative Views and Active Recognition

Tal Arbel, Frank P. Ferrie, and Peter Whaite

Abstract

In this paper we introduce a method for distinguishing between informative and uninformative viewpoints as they pertain to an active observer seeking to identify an object in a known environment. The method is based on a generalized inverse theory using a probabilistic framework where assertions are represented by conditional probability density functions. Consequently, the method also permits the assessment of the beliefs associated with a set of assertions based on data acquired from a particular viewpoint. The importance of this result is that it provides a basis by which an external agent can assess the quality of the information from a particular viewpoint, and make informed decisions as to what action to take using the data at hand. What is important about the method is that it provides a formal recipe for representing and combining all prior knowledge in order to obtain the required density functions (which we refer to as belief distributions).

To illustrate the theory we show how the characteristics of belief distributions can be exploited in a model-based recognition problem, where the task is to identify an unknown model from a database of known objects on the basis of parameter estimates. This leads to a sequential recognition strategy in which evidence is accumulated over successive viewpoints (at the level of the belief distribution) until a definitive assertion can be made. Experimental results are presented showing how the resulting algorithms can be used to distinguish between informative and uninformative viewpoints, rank a sequence of images on the basis of their information (e.g. to generate a set of characteristic views), and sequentially identify an unknown object.

Résumé

Dans cet article, nous pr'esentons une m'ethode qui permet de distinguer les points de vue informatifs et non-informatifs d'un objet tels que perçus par un observateur actif qui cherche à identifier un objet dans un environnement connu. La m'ethode repose sur une généralisation de la théorie inverse utlisée en probabilité où les hypothèses sont représentés par des fonctions de densité de probabilité. Conséquemment, la méthode permet aussi l'estimation de la confiance associée à un ensemble d'hypothèses basés sur les données obtenues d'un certain point de vue. L'importance de ce résultat est qu'il procure une base par laquelle un agent externe peut estimer la qualité de l'information provenant d'un point de vue et en conséquence prendre une décision éclairée quant à l'action à réaliser. Ce qui est important à propos de cette méthode est qu'elle procure un formalisme pour représenter et combiner l'information acquise de manière à obtenir la fonction de densité de probabilité recherchée.

Pour illustrer la théorie, nous montrons comment les caractéristiques des fonctions de distribution de la confiance peuvent être exploitées dans le cadre d'un problème de reconnaissance basée sur un modèle. La tâche consiste à identifier un modèle à partir d'une base de données d'objets représentés paramétriquement. Ceci débouche sur une stratégie de reconnaissance séquentielle par laquelle les évidences sont accumulées sur plusieurs vues (au niveau des distributions de confiance) jusqu'à ce qu'une hypothèse définitive puisse être établie. Des résultats expérimentaux démontrent comment l'algorithme peut être utilisé pour: distinguer entre les vues informatives et non-informatives, classer une séquence d'images sur la base de leur information (i.e. pour générer un ensemble de vues caractéristiques) et identifier séquentiellement un object inconnu.

1. Introduction

Consider an active agent charged with the task of roaming the environment in search of some particular object. It has an idea of what it is looking for, at least at some generic level, but resources are limited so it must act purposefully when carrying out its task [1]. In particular, the agent needs to assess what it sees and quickly determine whether or not the information is useful so that it can evolve alternate strategies, the next place to look for example. Key to this requirement is the ability to make and quantify assertions while taking into account prior expectations about the environment. In this paper we show how the problem be cast in probabilistic terms from the point of view of inverse theory [24]. Assertions are represented by conditional probability density functions, which we refer to as *belief distributions*, that relate the likelihood of a particular hypothesis given a set of measurements. What is particularly important about our methodology is that it yields a precise recipe for generating these distributions, taking into account the different sources of uncertainty that enter into the process. Based on this result we show how the resulting distributions can be used to (i) assess the quality of a viewpoint on the basis of the assertions it generates and (ii) sequentially recognize an unknown object by accumulating evidence at the probabilistic level.

Uncertainty serves to condition prior expectations such that the shape of the resulting belief distribution can vary greatly, becoming very delta-like as the interpretation tends towards certainty. In contrast, ambiguous or poor interpretations consistently tend towards very broad or flat distributions [2]. We exploit this characteristic to define the notion of an *informative viewpoint*, i.e. a view which gives rise to assertions that have a high probability according to their associated belief distribution. There are at least two applications for this result. First, in the case of an active observer, viewpoints can be chosen so as to maximize the distribution associated with an object of interest. This does not specify *how* to choose an informative viewpoint[1], but can be used as a figure of merit for a particular choice. Second, in the case of an off-line planner, it is often advantageous to be able to pre-compute a set of characteristic views to aid in recognition [14, 15, 22, 10, 11, 16, 7]. A good strategy here would be to select the n best views of an object ranked according to its belief distribution.

The other important aspect of the paper is the recognition problem itself, and here we consider the case of a model-based approach. Specifically, model-based recognition focuses on matching an unknown model, which is computed on-line from sensory data, with a predetermined model computed off-line and residing in a database of known objects [4]. What differentiates approaches is largely a matter of the kinds of models used to represent objects in the scene and how models are matched. Our interest is in three-dimensional object recognition in which objects are represented by parametric

[1] Strategies for gaze planning are usually operationally defined [26, 28].

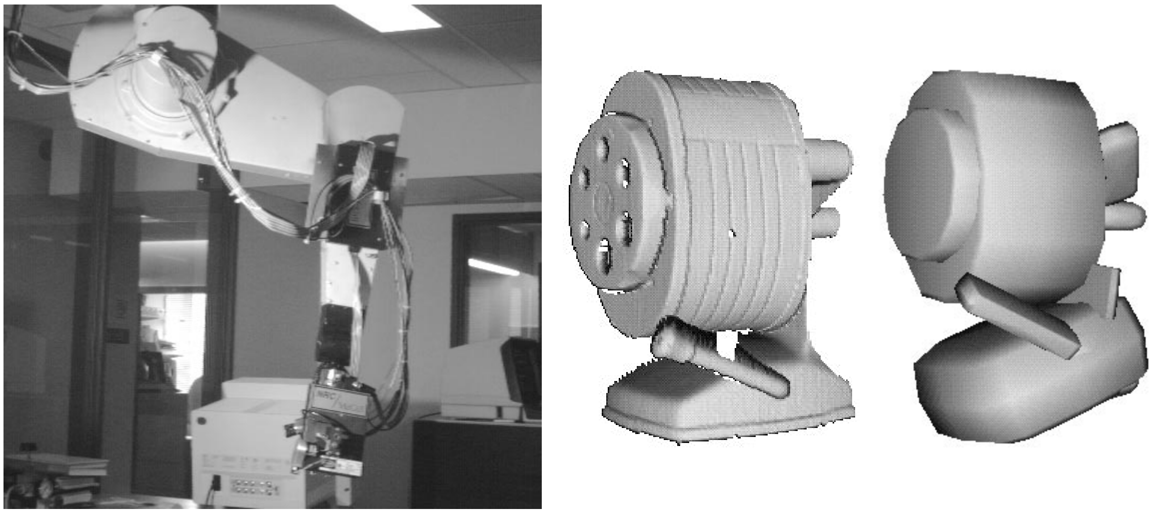

(a) (b) (c)

FIGURE 1. (a) Mobile laser range-finding system used to construct object models. (b) Laser range-finder image of a pencil sharpener rendered as a shaded image. (c) An articulated, part-oriented model of the sharpener using superellipsoid primitives; 8 superellipsoids are used, one corresponding to each of the parts of the object.

shape descriptors (i.e. models) such as superellipsoids [6, 5, 20, 12], deformable solids [9, 19], and algebraic surfaces [23]. In our context, models are constructed through a process of *autonomous exploration* [26, 27, 28] in which a part-oriented, articulated description of an object is inferred through successive probes with a laser range-finding system. Figure 1a shows the set-up used to perform experiments — a two-axis laser range-finder mounted on the end-effector of an inverted PUMA-560 manipulator. For any particular viewpoint, such as the one shown in Figure 1b, a process of bottom-up shape analysis leads to an articulated model of the object's shape (Figure 1c) in which each part is represented by a superellipsoid primitive [12]. Associated with each primitive is a covariance matrix $\mathbf{C}$ which embeds the uncertainty of this representation and which can be used to plan subsequent gaze positions where additional data can be acquired to reduce this uncertainty further [26, 27]. A system which automatically builds object models based on this principle is reported in [28, 17].

Many approaches have been advocated for the problem of *matching* models. The majority of these employ various metrics to measure the distance between models in the appropriate parameter spaces, e.g., Mahalanobis distance [13], dot product [19] to mention but a few. These strategies rarely include both the uncertainties in the parameters of the measured models and the ambiguities of the representations in the database. However, when fitting a model to data that is noisy, there is an inherent lack of uniqueness in the parameters that describe the model. In these cases it is impossible to make a definitive statement as to which model fits the data best [26]. For this reason, rather than choose external constraints that would force the choice of one model over another, it would be more instructive to embed the uncertainty in the chosen description into the representation. This is precisely the approach that we have taken in computing the belief distribution.

Our methodology is based on a probabilistic inverse theory first introduced by Tarantola in [24]. Earlier work has shown how this theory can be used to methodically synthesize belief distributions corresponding to each model hypothesis, $\mathcal{H}_i$, given the parameters corresponding to the unknown model, $\mathcal{M}$, computed from the current measurement D_j, i.e. $P(\mathcal{H}_i|\mathcal{M}_{D_j})$ [2]. This procedure explicitly accounts for uncertainties arising from the estimation of the unknown model parameters, database model parameters, and prior expectations on the frequency of occurrence for each of the database entries. In this case the solution reduces to the classical Bayesian solution, similar to the result obtained by Subrahmonia et al. [23] - the primary difference being in the techniques used to obtain the solution. In addition, they (and many others [4, 8]) are interested in the constructing a discriminant that makes an absolute identification of the measured object. We argue that making assessments about identity from single measurements can be erroneous. We are more interested in communicating the belief in the identification as feedback to the recognition procedure, into order to further reduce the ambiguity in an active strategy.

The key idea behind a sequential recognition strategy is that of improving interpretation by accumulating evidence. But at what level of representation should this evidence be accumulated? The autonomous exploration procedure that we use to generate the set of database models, for example, sequentially constructs a complete 3-D representation at the level of surface geometry [28]. One could follow a similar approach at the recognition phase, i.e. recalculate each belief distribution as the explorer adds new data to its representation of the unknown object. Unfortunately this would be computationally prohibitive, largely due to the expense of data fusion [21]. A better approach would be to process each view independently and avoid the fusion problem at the data level by seeking instead to combine information at the level of the belief distribution. That is, given two data sets D_j and D_{j+1} corresponding to different viewpoints we seek a conjunction of $P(\mathcal{H}_i|\mathcal{M}_{D_j})$ and $P(\mathcal{H}_i|\mathcal{M}_{D_{j+1}})$ that is equivalent to $P(\mathcal{H}_i|\mathcal{M}_{D_j+D_{j+1}})$. An active agent would then be able to gather evidence until the composite belief distribution associated with a particular hypothesis exceeds a prescribed figure of merit.

Although the theory formally defines conjunction, such an operation requires knowing how a change in viewpoint conditions the respective belief distributions.

However, as we will show later on, if the maximum likelihood hypothesis [2] is largely invariant over a sequence of trials, then a robust interpretation can be made by tabulating the votes for each one and picking the hypothesis with the highest score. We also show that this invariance can be maximized by using the structure of the belief distribution to filter out uninformative hypotheses.

The remainder of the paper is organized as follows. We begin in Section 2 with an introduction to the general inverse theory, and an explanation of how to apply the theory to the problem of recognizing parametric models. In Section 3, we describe how to distinguish between informative and uninformative viewpoints, and introduce an incremental recognition scheme in Section 4. In Section 5, we describe a series of experiments illustrating the application of the method to characterize informative views within the context of recognizing both single-part objects, and parts of multiple-part objects. As well, we perform a series of incremental recognition experiments that test the accumulation of evidence from sequential viewpoints. Finally, we conclude in Section 6 with a summary of the results and a pointer to future applications.

2. The Inverse Problem Theory

The recognition problem requires us to infer from measurements of an unknown object that model which most closely represents it in a database of known objects. Like all inverse problems, the recognition problem is ill posed in that, i) several models can give rise to identical measurements and, ii) experimental uncertainty gives rise to uncertain measurements. As a result it is not possible to identify the unknown object uniquely. There are various ways of conditioning ill posed problems, but these all require strong, and often implicit, a priori assumptions about the nature of the world. As a result a method may work well only in specific cases and because of the hidden implicit nature of the conditioning assumptions, cannot be easily modified to work elsewhere.

For this reason we have adopted the very general inverse problem theory of Tarantola [24]. In it the sources of knowledge used to obtain inverse solutions are made explicit, so if conditioning is required, the necessary assumptions about that knowledge are apparent and can be examined to see if they are realistic. The theory uses probability density functions to represent the following sources of knowledge:

1. Knowledge given by a theory which describes the physical interaction between models $\mathbf{m}$ and measurements $\mathbf{d}$, denoted $\theta(\mathbf{d}, \mathbf{m})$,
2. Knowledge about the model from measurements, denoted $\rho_D(\mathbf{d})$.
3. Information from unspecified sources about the kinds of models which exist in the world (namely that there are a discrete number of them). We denote this knowledge $\rho_M(\mathbf{m})$. Knowledge like this is a powerful constraint and can be used to eliminate many of the unconstrained solutions.

[2] This refers to the hypothesis that the correct answer is the one with the highest belief.

Given this knowledge the theory tells us how it should be combined, but leaves any decision about its usefulness up to the tasks that require it. For example, when attempting to recognize objects we would ideally want the unknown model be identified correctly all the time. Because of experimental uncertainties this can never happen, and there is always the possibility that an object will be identified incorrectly. Only the task can know if the likelihood of errors is acceptable.

This raises the interesting question of what we should do if the level of errors is not acceptable. Because the sources of knowledge are explicit they are not only visible to the operational tasks, but are also potentially open to manipulation by them. In principal it should be possible for the task to condition or actively acquire the a priori knowledge required to make the solution acceptable. We have already demonstrated that what we call autonomous exploration functions well at the model building level [25, 28] and we now intend, with the aid of this theory, to incorporate feedback from the recognition task as well.

2.1. The Inverse Solution. The theory postulates that our knowledge about a set of parameters is described by a probability density function over the parameter space. This requires us to devise appropriate density functions in order to represent what we know about the world. The solution to the inverse problem then becomes quite straight forward — it is simply a matter of combining the sources of information. The logical operation of *conjunction* is appropriate, i.e. the solution to the inverse problem is given by the theory AND the measurements AND any a priori information about the models. Tarantola extends the notion of logical conjunction to define the conjunction of two states of information [24, pages 29–31]. With this definition we can therefore combine the information from the joint prior probability density function $\rho(\mathbf{d}, \mathbf{m})$ and the theoretical probability density function $\theta(\mathbf{d}, \mathbf{m})$ to get the a posteriori state of information

$$\sigma(\mathbf{d}, \mathbf{m}) = \frac{\rho(\mathbf{d}, \mathbf{m})\ \theta(\mathbf{d}, \mathbf{m})}{\mu(\mathbf{d}, \mathbf{m})} \tag{1}$$

where $\theta(\mathbf{d}, \mathbf{m}) = \theta(\mathbf{d}|\mathbf{m})\ \mu_M(\mathbf{m})$ and $\rho(\mathbf{d}, \mathbf{m}) = \rho_D(\mathbf{d})\ \rho_M(\mathbf{m})$ over the joint space $M \times D$. The so called non-informative probability density $\mu(\mathbf{d}, \mathbf{m}) = \mu_D(\mathbf{d})\mu_M(\mathbf{m})$ represents the reference state of information in much the same way that noise is used when measuring information in terms of signal to noise ratios. The formulation of appropriate non-informative densities is a complex issue, but for our purposes we will assume that all the non-informative densities are uniform over their respective spaces.

Accordingly, (1) is more general that the equations obtained through traditional approaches, but degenerates to them in specific cases. Under the conditions mentioned, the solution is identical to the Bayesian solution [24, page 61] where the a posteriori information about the model parameters is given by the marginal probability density function:

$$\sigma(\mathbf{m}) = \rho_M(\mathbf{m}) \int_D \frac{\rho_D(\mathbf{d})\ \theta(\mathbf{d}|\mathbf{m})}{\mu_D(\mathbf{d})}\ d\mathbf{d}. \tag{2}$$

2.2. The Part Recognition Problem. In the system we have constructed, range measurements are taken, surfaces are reconstructed, segmented into parts, and individual models are fit to each part. We will treat *the whole system as a measuring instrument.* Given some model $\mathbf{m}$ in the scene, range measurements are taken and from these an *estimate* of the model $\mathbf{d}$ is obtained, which we call a *measurement of the model* in the scene.

2.2.1. *Information Obtained from Physical Theories.* We first formulate an appropriate distribution to represent what is known about the physical theory that predicts estimates of the model parameters given a model in the scene. Such a theory is too difficult to formulate mathematically given the complications of our system. We therefore build an empirical theory through a process called the *training* or learning stage. Here, Monte Carlo experiments are run on measures of a known model exactly as in traditional statistical pattern classification methods. The conditional probability density function $\theta(\mathbf{d}|\mathbf{m})$ is calculated for each model $\mathbf{m}$ by assuming a multivariate normal distribution. Therefore, the equation for $\theta(\mathbf{d}|\mathbf{m})$ is:

$$\theta(\mathbf{d}|\mathbf{m}) = N(\mathbf{d} - \mathbf{m}, \mathbf{C}_T) \tag{3}$$

where N is the multivariate normal distribution, with a covariance matrix, $\mathbf{C}_T$, describing estimated modelling errors for a model $\mathbf{m}$.

2.2.2. *Information Obtained from Measurements.* Much of the knowledge we have about a problem comes in the form of experimental measurements. In our system [28], we obtain an estimate of the observed model parameters $\mathbf{d}_{obs}$, and also an estimate of their uncertainty in the covariance operator $\mathbf{C}_d$. The assumption we make is that the multivariate normal distribution $N(\mathbf{d} - \mathbf{d}_{obs}, \mathbf{C}_d)$ represents our knowledge of the measurements. The probability density function representing this information is the conditional probability density function $\nu(\mathbf{d}_{obs}|\mathbf{d})$, such that:

$$\nu(\mathbf{d}_{obs}|\mathbf{d}) = \rho_D(\mathbf{d})/\mu_D(\mathbf{d}) = N(\mathbf{d} - \mathbf{d}_{obs}, \mathbf{C}_d) \tag{4}$$

2.2.3. *A Priori Information on Model Parameters.* In the current context, there are a discrete number of reference models, $\mathbf{m}_i, i = 1 \ldots M$. The probability density function used to convey this knowledge is

$$\rho_M(\mathbf{m}) = \sum_i P(\mathbf{m}_i)\ \delta(\mathbf{m} - \mathbf{m}_i), \tag{5}$$

where $P(\mathbf{m}_i)$ is the a priori probability that the i^{th} model occurs.

2.2.4. *Solution to the Inverse Problem.* Substituting the probability density functions (3), (4), and (5) into (2) gives us the final equation for the a posteriori probability density function

$$\sigma(\mathbf{m}) = \sum_i P(\mathbf{m}_i) N(\mathbf{d}_{obs} - \mathbf{m}_i, \mathbf{C}_D)\ \delta(\mathbf{m} - \mathbf{m}_i). \tag{6}$$

where $\mathbf{C}_D = \mathbf{C}_d + \mathbf{C}_T$. This density function is comprised of one delta function for each model in the database. Each delta function is weighted by the *belief* $P(\mathbf{m}_i)N(\mathbf{d}_{obs} - \mathbf{m}_i, \mathbf{C}_D)$ in the model $\mathbf{m}_i$. The final distribution represents the "state of knowledge" of the parameters of $\mathbf{m}_i$. The beliefs in each of the reference models are computed by convolving the normal distributions in (3) and (4). The advantage of the method is that rather than establish a final decision as to the exact identity of the unidentified object, it communicates the degree of confidence in assigning the object to each of the model classes. It is then up to the interpreter to decide what may be inferred from the resulting distribution.

The methodology introduced applies to the recognition of any parametric primitive. For our purposes, superellipsoid models were chosen because of the range of shapes they can represent as well as their computational simplicity. However, representations based on superquadrics pose a number of problems due to degeneracies in shape and orientation. In solving this problem, work has being done in representing objects by multi-modal distributions, where each mode contains information about a possible equivalent form. Discussion of this process is beyond the scope of this paper.

2.3. Example – Recognizing a Sphere. Figure 2 illustrates the kinds of results we get by applying the theory to a typical recognition problem. Here, the reference models were produced by training on models created with data acquired by scanning the objects all around their surfaces (i.e. complete 3D data). The reference models, consisting of a smaller sphere, a large sphere, and a lemon, can be seen in Figure 2a. The larger sphere was then measured from a single viewpoint, and the resulting model is shown in Figure 2b. The system's ability to distinguish the larger sphere from both the smaller sphere and the lemon was then tested. The result is the belief distribution found in Figure 2c. One can see that the system has a significantly higher degree of confidence in the hypothesis that the measured model was a large sphere.

3. Determining Which Viewpoints are Informative

In [3], it has been shown that recognition based on complete information produced perfect results in all cases. Since complete information is not always available, and potentially expensive to acquire, recognition schemes based on single viewpoints are required. We will show that recognition based on one view is not always reliable. In fact, the degree of reliability depends upon the amount of information available. For example, some viewpoints capture enough of the unique characteristics of the object to sufficiently distinguish it from the others in the database. We will refer

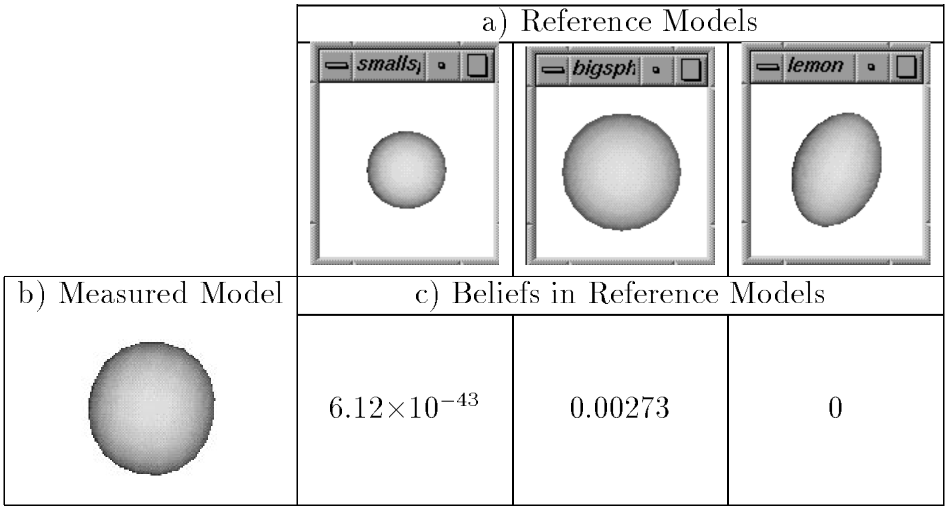

FIGURE 2. Recognizing a sphere. (a) The reference models are: a smaller sphere, a larger sphere, and a lemon. (b) The measured unknown model. (c) The belief distribution.

to these viewpoints as *informative viewpoints*. Other viewing positions, where it is impossible to say which object in the database the unknown is closest to, are called *uninformative viewpoints*. By determining if a viewpoint is informative or not, we can establish if further sampling is necessary to be able to recognize the object well.

An important result of the inverse solution is that rather than establish an absolute identity for the unknown object, the system produces a belief in each of the models in the database. We can then use the belief distribution to tell us whether a particular viewpoint is informative or not. This is easily accomplished by denoting views with a clear winner, in terms of a significantly higher belief in one model than the others, as informative views. From these positions, the system is able to capture the attributes of the model that distinguish it from the others. The important contribution of this work is to be able to recognize these viewpoints, and use them in the synthesis of object identity.

We would also like to use the beliefs for the converse, i.e. to label a viewpoint as uninformative. This indicates that results from the current viewing position do not tell us much about the object's identity. This situation occurs when the unnormalized belief in each of the models is very low (or zero). Here, it is impossible to say which reference model the unknown might correspond to. This situation occurs when the distribution of the unknown model does not significantly overlap with any of the reference distributions. There are two possible reasons for this to occur. The first is the case where the distribution of the measured model is very wide due to large uncertainties in its parameters. The result is low beliefs in all the reference models in the database. This case occurs when the scanning has occurred from a viewpoint where insufficient data was collected. The second case occurs when there is a breakdown in some of the prior assumptions. In this case, the issue was not one of insufficient data. Here, the parameters determined from that particular viewpoint differ significantly from any of the models in the database. The resulting distribution could actually be quite sharp, but simply does not overlap with any of the reference model distributions. In this case, it could be that the linearity assumption breaks down, implying that perhaps the assumption of a normal distribution is not valid. Zero belief cases exist when the values of the a posteriori probability density functions are extremely low. Due to numerical underflow, the procedure produces beliefs of zero for each of the reference models. Figure 3 illustrates the difference between informative and uniformative viewpoints for the case of a cylinder. Here, one can see that the system is able to distinguish the cylinder from a block with great ease, if the cylinder is measured from an informative viewpoint. However, if measured from an uninformative viewpoint, there is little confidence in either model. In this case, the beliefs are in fact below the numerical precision of the system, and therefore become zeros.

The problem of distinguishing between the two kinds of states becomes one of determining the threshold, below which one can safely state that the beliefs are in fact insignificant. It is obvious that cases where the beliefs in all the models are zero are uninformative. However, this threshold depends on the numerical precision of the system. In this sense, it is chosen externally (and is, therefore, a random cutoff point). We therefore feel justified in raising this threshold to one that excludes other low confidence states. The expectation is that this will eliminate false positive states, as they are thought to occur with low belief. One can determine this cutoff point empirically, by observing the belief distributions from different viewpoints, and noting if there is a clear division between the clear winner states and the low confidence states. A bi-modal distribution would indicate that a application of a predefined threshold can easily distinguish between these states.

4. Incremental Recognition

Provided that the low belief states have been identified, we wish to make a statement about the remaining beliefs. Even though the majority of the cases can be clearly divided into informative and uninformative states, there are still ambiguous cases where a "significant" belief in more than one model exists. Because of these situations, it becomes apparent that evidence from more than one viewpoint is needed. The question becomes: how do we accumulate evidence from different views, when the evidence is in the form of a conditional probability density function? The immediate response is given by the theory (Section 2) which formally defines the operation of conjunction of information, i.e. the belief distributions. To state this more formally, we denote belief distributions corresponding to each model

Database Models

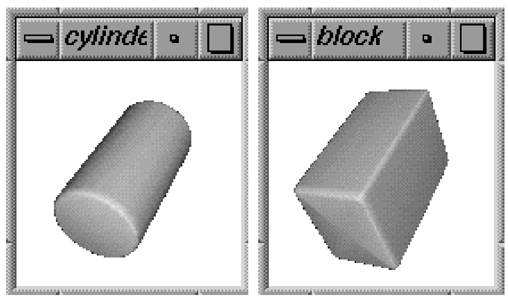

Measured Model	View 1	View 2	View 3	View 4
	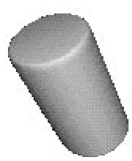	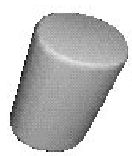	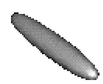	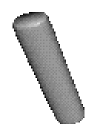
Belief in cylinder	2.237	0.009181	0.0	0.0
Belief in block	0.0	0.0	0.0	0.0
	a) Informative		b) Uninformative	

At the top of this figure are the two reference models in the data base: the cylinder and the square block. Beneath these are measured models of the *cylinder* obtained after scanning its surface from 4 different viewing positions. Below each model one can find the unnormalized belief distributions obtained when attempting to recognize each of the measured models.

FIGURE 3. (a) Informative and (b) Uninformative Views of a Cylinder.

hypothesis, $\mathcal{H}_i$, given the parameters of the unknown model, $\mathcal{M}$, computed from the measurement, D_j, by $P(\mathcal{H}_i|\mathcal{M}_{D_j})$. Then, given two data sets D_j and D_{j+1} corresponding to different viewpoints we seek a conjunction of $P(\mathcal{H}_i|\mathcal{M}_{D_j})$ and $P(\mathcal{H}_i|\mathcal{M}_{D_{j+1}})$ that is equivalent to $P(\mathcal{H}_i|\mathcal{M}_{D_j+D_{j+1}})$. An active agent would then gather sufficient evidence in this fashion until the composite belief distribution associated with a particular hypothesis exceeds a predefined level of acceptability.

The problem with combining belief distributions from different viewpoints is that they are not normalized. Therefore, relative values between the views are meaningless. The normalizing factor is some unknown function of viewpoint, and is difficult to obtain analytically. In the long version of this paper, we show that the reason for the difficulty lies in that this factor can be shown to be a function of $\int_D d\mathbf{d}$. The issue of how to define this space is a difficult one to address. In order to do so, a commitment to a permissible region of observed parameters must be established *prior* to experimentation. In this sense, this problem is analogous to finding correspondence between range data from different viewpoints in that the relationship between the views is unknown. Because of the difficulty of finding this factor, the beliefs are not normalized. As a result, it becomes difficult to match a belief of 500, for example, from one view, with a value of 50 from another. Each of these values may reflect the strongest possible belief from their respective views, however it is difficult to compare them in a sensible fashion. As well, in situations where there is a belief of 50 in one model and 40 in another, it becomes impossible to establish a clear winner.

For this reason, we have chosen not to choose a "winner" in ambiguous situations, and state that all positive beliefs indicate equally likely hypotheses. We illustrate this philosophy by binarizing the conditional probability density function values at each view, such that all beliefs above the threshold become ones. In this fashion, we have divided the possible results to include:

(1) *Informative states*: states with one clear winner (a single positive value).
(2) *Uninformative states*: states without a clear winner. This includes:
  a) *Ambiguous states*: states with more than one possible winner (more than one single positive value).
  b) *Undetermined states*: states with no winners (all zero values).

It is important to note that ambiguous states are, in fact, undetermined states that lie above the chosen threshold. In theory, careful choice of cutoff level should eliminate these states as well (without eliminating a large number of informative states). Figure 4 illustrates these different states in the case of a square block. Here, the system is asked to identify a square block from different views, and correctly distinguish it from a similar rounder one. This example indicates that the results match human intuition. The clear winners, or informative states, in Figure 4a indicate that the system is able to identify the block despite wide variations in its three dimensions. The ambiguous cases (Figure 4b) occur when the resulting models are rounder in shape. Here, the system has trouble differentiating between the models. In fact, these models resemble the rounded block more than the square one. In the third case (Figure 4c), the system does not have significant belief in any of the models. Intuitively, one can see that these models are not similar to either reference model.

Using this method of representation, rather than base conclusions on maximum likelihood methods from independent viewpoints, methods that combine evidence from single viewpoints would consider all models whose beliefs are above a threshold to be equally significant. In accordance with Marr's "Principle of Least Commitment" [18], all possible hypotheses, rather than just one are communicated to the external processes.

By normalizing our confidence values in this manner, combining them from different viewpoints becomes straightforward. At any point, one can histogram the binarized beliefs accumulated thus far. The idea is that if the evidence in the correct hypothesis prevails in a largely view-invariant manner, then after a sequence of trials, a clear winner should emerge. In addition, the confidence in the incorrect models should become insignificant. Figure 5 illustrates an attempt at sequentially recognizing the square block at $40°$ increments. As in the previous example, the square and round blocks are used as reference models. The raw beliefs are binarized by imposing a threshold of 10^{-13}. Notice that the ambiguous case quickly becomes insignificant with the increase of evidence in the correct model. After only 9 iterations, the clear winner emerges, casting all doubt aside.

5. Experiments and Results

In this set of experiments, we wish to illustrate how one can use the belief distributions to distinguish between informative and uninformative viewpoints. As well, we wish to show that evidence, in the form of the belief distributions, can be accumulated from sequential viewpoints. The expectation is that a clear winner will emerge rapidly after a small number of views.

5.1. Characterizing Informative Viewpoints by External Threshold. It has already been shown that the system can successfully recognize an instance of any object in the database with perfect results, provided that all its surfaces are accessible, independently of viewpoint and sampling order [3]. It has also been shown that recognition from single-views retains most of the selectivity of the previous case, displaying high discriminatory powers, with the occasional false-positive cases. The hypothesis has been that, if the majority of the incorrect cases occur with low beliefs, by raising the threshold, one could eliminate most of them. Through experimentation, we wish to test the hypothesis on a set of real objects, and see if by raising the cutoff, we are able to transform those ambiguous cases that would normally result in false-positive indications into uninformative viewpoints. In this manner, rather than make

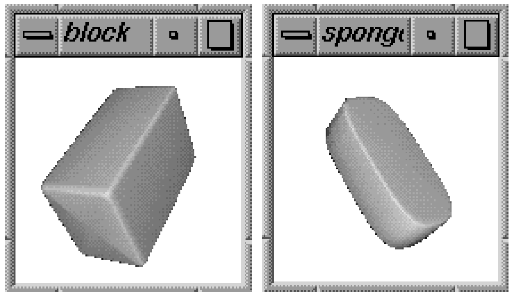

Measured Model	Belief in Block		Belief in Round Block	
	Unnormalized	Binarized	Unnormalized	Binarized
	0.2	1	0	0
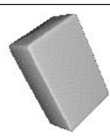	0.007	1	0	0
	2.0×10^{-13}	1	5.8×10^{-6}	1
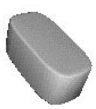	3.4×10^{-13}	1	0.002	1
	0	0	0	0
	0	0	0	0

Above are the two reference models: a block and a rounded block. In the left column of the table are the models of the block measured from informative (first pair), ambiguous (middle pair) and undetermined (last pair) viewpoints. To their right, one can find the unnormalized, and binarized (threshold of 10^{-13}) belief distributions obtained when attempting to recognize each of the measured models.

FIGURE 4. Informative, Ambiguous, and Undetermined States for the Block.

View Angle	Measured Model	Belief in Block		Belief in Round Block	
		Unnormalized	Binarized	Unnormalized	Binarized
0°		2.0×10^{-13}	1	5.8×10^{-6}	1
40°		0	0	0	0
80°	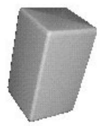	0.2	1	0	0
120°	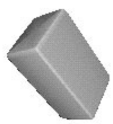	0.03	1	0	0
160°		0	0	0	0
200°		0.1	1	0	0
240°		0	0	0	0
280°		0.03	1	0	0
320°		0.001	1	0	0
		Final Score	6		1

Displayed above are the 9 models resulting from sequentially measuring the square block at 40° increments. From left to right, one can see the viewing angle, the measured model, the unnormalized and binarized (threshold of 10^{-13}) belief distribution resulting from attempting to recognize each of the measured models. The final distribution is the histogram of the binarized distributions.

FIGURE 5. Incremental Recognition of a Block.

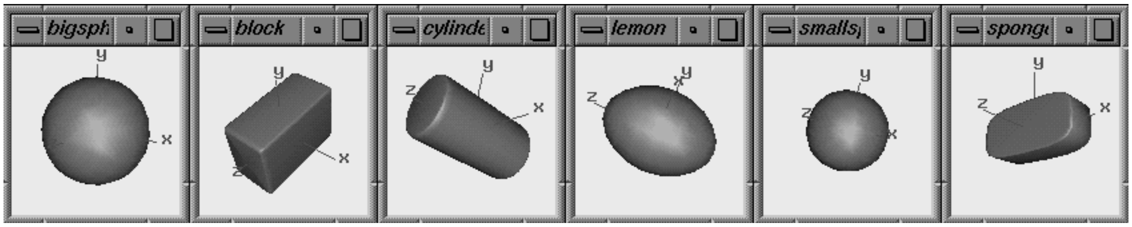

BS B C L SS RB

Displayed above are the reference objects that result from training on complete surface data: a big sphere (BS), a block (B), a cylinder (C), a lemon (L), a smaller sphere (SS), and a rounded block (RB).

FIGURE 6. Six representatives that result from training.

assessments about the object's identity based only on the maximum likelihood results from these views, external processes can establish that they do not provide significant information and further sampling becomes essential.

Six objects were used to examine the recognition procedure: two spheres (rad = $20mm$, rad = $25mm$), a block, a cylinder, a lemon, and a block with rounded edges. The objects were selected because they consisted of single parts that conformed well to superellipsoids. They varied in size and shape, so as not to be clustered together too tightly in five-dimensional feature space. However, their distributions overlapped sufficiently enough in several dimensions so that the recognition procedure was challenged in its discrimination task.

Training (Section 2.2.1) automatically produced object class representatives, by measuring the object numerous times. Each individual model was created by scanning the object from several views using a laser range-finder, then a superellipsoid model was fit to the data, and the resulting parameters stored. For the purposes of creating a stable database for recognition, it was established that three views of each object, $120°$ apart, were sufficient to constrain the fitting procedure. Each sample was scanned from a random scanning position, producing 24 samples of each object. Figure 6 illustrates the six representative models of each object that result from training.

In the first set of experiments, recognition was performed on thirty-six single-view samples of each object. Here, data was collected at $40°$ intervals along 4 different great circle routes. Using maximum likelihood as the basis for recognition, i.e. choosing the model with the highest confidence value, the results shown in Figure 7.

By examining the belief distributions, one can see that the two false positive cases occurred. In these cases, the system identified the block as being the rounded block, despite the fact that the resulting distribution overlapped with the distribution of the reference block as well. The belief in both models was quite low, indicating that the system is quite uncertain about the identification. In fact, in many cases, a false-positive identification is associated with low beliefs. This suggests that if the threshold for undetermined states were raised, the incorrect identifications would

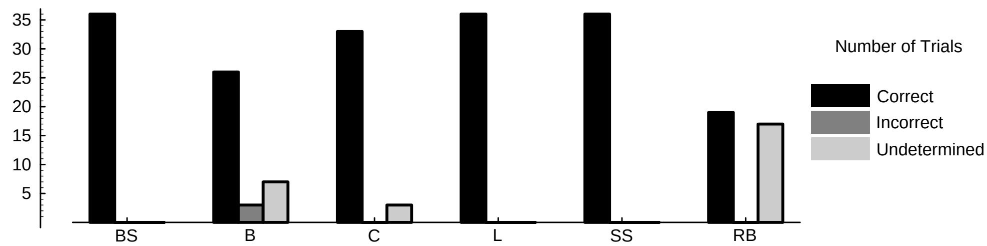

FIGURE 7. Matching samples taken from single viewpoints.

become undetermined states.

In order to justify raising this threshold, the beliefs resulting from the experiment described above were plotted on a logarithmic scale graph. The hope in observing these results was that the scatter of the beliefs was bi-modal. This would imply that a distinct separation between informative and uninformative cases exists, permitting the application of a threshold to distinguish between the two. The results can be found in the plot in Figure 8.

The results illustrate clustering effect in the beliefs. The first large cluster indicates that the highest degree of confidence lies in the correct model hypotheses. Beneath this group, is a scatter of beliefs in the incorrect model. The degree of evidence of these hypotheses varies from model to model. This second large cluster occurs for beliefs in models that lie below the numerical precision of the system (denoted the "U" level). The distinct bi-modality of the results justifies the application of an external threshold differentiating between the high confidence informative views and the low confidence uninformative views. In addition, they indicate that the value of this threshold is not critical. For example, for the Big Sphere model, the cutoff point can lie anywhere from 10^{-5} to 10^{-60} (above the "U" level). However, the desire is to choose this threshold so as to eliminate the majority of false positive cases. Although the plot does not illustrate the maximum likelihood results, making it impossible to tell where false positive indications occur, one can see that by placing the cutoff above the scatter of incorrect hypotheses, one can ensure a minimal amount of incorrect maximum likelihood indications. Furthermore, one can see from the results that one does not necessarily need to choose a universal threshold level for all the models. By examining the difference in the Big Sphere and the Rounded Block distributions, one can see that choosing individual cutoff levels would render the results more accurate. For maximal efficiency, these levels can be computed off-line prior to experimentation, and then used in the recognition stage.

For the purposes of testing the hypothesis that an external cutoff would divide the results into informative and uninformative cases (and eliminate the majority of false-positive cases), the threshold for undetermined states was uniformly raised to 0.00001. Figure 9 shows the results of imposing this threshold on the belief distributions. One

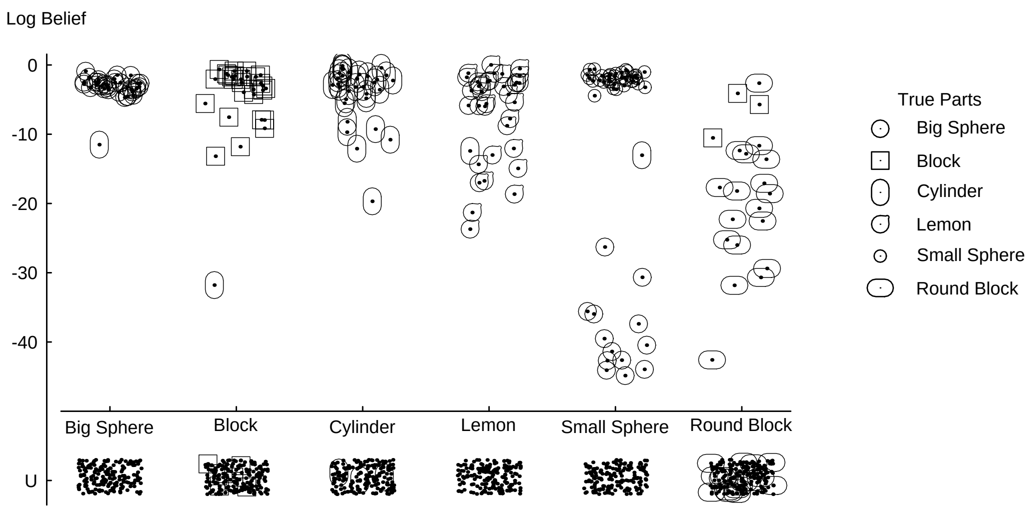

Above are the results from attempting to recognize 36 different single-view samples of each of the models in the database. The beliefs in the different models are represented by different symbols, each symbol indicating the true model used during that trial.

The level of numerical underflow of the system is represented by a "U" on the y axis. Because so many trials fall into this category they are marked with a simple point, *except* when the belief is for the true model used in the trial.

By observing the log of the beliefs, one can see the bi-modality in the results.

FIGURE 8. Log of beliefs in the Big Sphere, Block, Cylinder, Lemon, Small Sphere, and Round Block.

can see that all but one incorrect state (B) has become undetermined. However, several correct identifications have become undetermined as well. This is to be expected since setting this threshold causes all uncertain identifications to be removed. We therefore make the empirical observation that, by raising the threshold, states that are not undetermined are accompanied by a high accuracy in recognition.

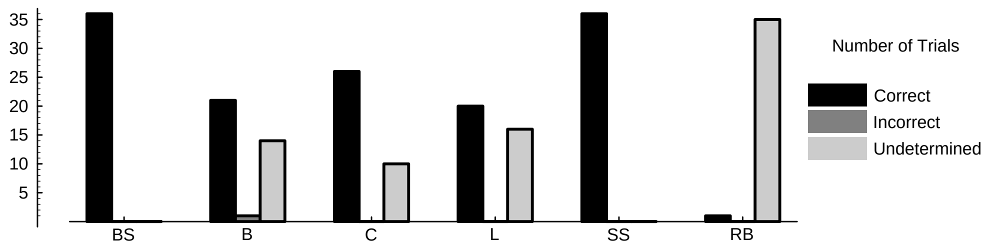

FIGURE 9. Matching samples taken from a single viewpoint while imposing a threshold of 0.00001.

Since these object are less complex than most found in the real world, we are more interested in recognizing objects that consist if several articulated parts. Our current focus is "recognition by parts", whereby measured objects are segmented into their constituent parts, each of which is compared to the parts in the database. The task of recognizing these parts is much more challenging than recognizing single-part objects due to problems of self-occlusion and segmentation.

A toy potato-head consisting of two ears, two eyes, a nose and a head was chosen for the purposes of testing the part recognition algorithm on complex objects. Figure 10a displays the actual potato-head toy used in the experiment. Most of constituent parts were superellipsoids, however the head's shape was tapered. The object was chosen because its parts were similar to each other as well as to the reference spheres making discrimination a challenging task. Ten samples were used in the training procedure. Each sample was produced by scanning the object from several viewpoints in an exploration sequence. The reference model resulting from training can be found in Figure 10b.

The the potato-head was then measured from 32 independent viewing positions. Recognition was performed on each of these samples in turn, using a database consisting of the parts of the potato-head with the single-part objects used earlier as distractors. The results of the using maximum likelihood on the beliefs can be seen in Figure 11.

The results indicate that the system is able to successfully recognize instances of articulated parts of a complex object with only partial information available. Figure 11a illustrates that the system is able to maintain its selectivity even though very little information is available from single viewpoints, especially with the added effects of self-occlusion.

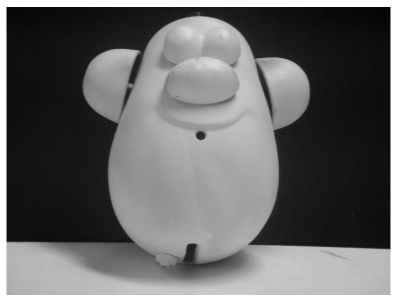

a) Original potato-head toy.

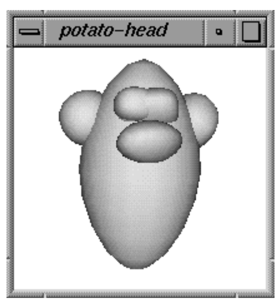

b) Reference potato-head model created by training.

FIGURE 10. Potato-head: a) real object and b) reference model.

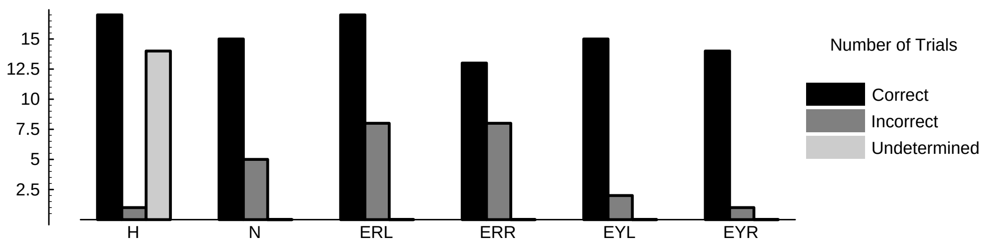

a) Threshold = computational underflow

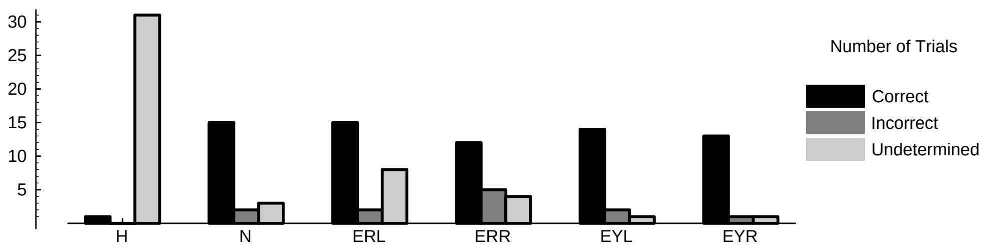

b) Threshold = 0.00001

Displayed above are the tables describing the belief distributions of the potato-head measured from single view-points. The parts of the potato-head are: a head (H), a nose (N), a left ear (ERL), a right ear (ERR), a left eye (EYL), and a right eye (EYR). Here, labelling one eye as the other, or one ear as the other was considered to be a correct identification. Zero values are defined by a) numerical underflow of system and b) a threshold of 0.00001.

FIGURE 11. Matching samples of the potato-head taken from single viewpoints.

The incorrect states arose due to the similarity of the eyes to the smaller sphere and to the nose. As well, the ears were extremely close to the bigger sphere in size and shape. However, as hypothesized, the beliefs in the incorrect model were very small, and most were eliminated by imposing a threshold. This is shown in Figure 11b.

In most cases, the external threshold retained most of the correct states, confirming that the system had high confidence in the correct identifications. The exception was the case of the head, where low beliefs caused almost all of the correct identifications to become undetermined states. This is because the head is tapered, breaking the assumption that the objects *can* in fact be accurately modeled by the current superellipsoid representation which does not describe tapering. Thus, because different superellipsoid models can be used to describe the head equally well, the system had trouble identifying the one in the database, leading to undetermined states.

5.2. Accumulating Evidence from Sequential Viewpoints. The experiments described suggest the possibility of an incremental recognition procedure. It is based in the following observations obtained empirically over successive trials:

i) Viewpoints that provide very little information, or *uninformative* views, generally can be detected by their low confidence levels (beliefs). Because of the bi-modality of the belief spread, these can be discovered by application of a threshold. Detection of such events is a clear indicator that further sampling is required.

ii) *Informative* views are generally accompanied by high beliefs, but with the possibility of a false-positive indication. These can also be detected by threshold application.

iii) The likelihood of successive false-positive indications is very small. First, this is a consequence of the high selectivity of the reference distributions which result in low frequencies of false-positive indications in the first place (e.g. Figure 9). Second, it is unusual for observer motion to result in similar viewpoints in two successive views (general position assumption).

To explore the possibility of an incremental scheme, an experiment was performed whereby evidence from single-views was accumulated. Here, the method described earlier was employed, whereby the system binarized the beliefs above the predefined threshold at each view. Evidence at each stage was computed by histogramming the binarized beliefs accumulated thus far. In order to simplify the experiment, only the single-part objects were used in this set of experiments. Table 1 displays the result of accumulating evidence after 36 single-view iterations. Table 1a illustrates the results when the zero states were established by the numerical limitations of the system, whereas in b, a threshold of 0.00001 was imposed externally. One can see from these results that, after several iterations, choosing a winner based on a maximum likelihood scheme on the accumulated beliefs gave the correct answer in

	BS	B	C	L	SS	RB
BS	36	0	0	0	13	0
B	0	28	0	0	0	3
C	1	1	33	1	1	0
L	0	0	0	36	0	0
SS	0	0	0	0	36	0
RB	0	0	0	0	0	18

a) Threshold = computational underflow

	BS	B	C	L	SS	RB
BS	36	0	0	0	0	0
B	0	21	0	0	0	1
C	0	0	26	0	0	0
L	0	0	0	20	0	0
SS	0	0	0	0	21	0
RB	0	0	0	0	0	1

b) Threshold = 0.00001

Displayed above are the tables describing the accumulation of evidence from 36 single-view experiments. Each row describes the histogram of the binarized belief distributions for a particular measured model. The columns refer to the reference models. Zero values are defined by a) numerical underflow of system and b) a threshold of 0.00001.

TABLE 1. Histogram of binarized belief distributions after single-view iterations.

all cases. The false-positive cases became insignificant due to insufficient evidence. In fact, Table 1b illustrates that hardly any evidence in incorrect models remained after applying the threshold of 0.00001. However, in the case of the rounded block, the majority of the evidence in the correct model was also eliminated, indicating that perhaps this choice of threshold was too high in this case. Its belief values were, in fact, significantly lower than the rest of the objects. In their case, this choice of threshold seems to be appropriate in that it removes the false-positive cases, while maintaining a high degree of confidence in the correct hypotheses. This justifies using independent threshold levels for each of the models in the database.

6. Discussion and Conclusions

In this paper, we have introduced a method for distinguishing between informative and uninformative viewpoints and for assessing the beliefs associated with a particular set of assertions based on this data. The importance of this result is that it provides a basis by which an external agent can assess the quality of the information from a particular viewpoint, and make informed decisions as to what action to take using the data at hand. Our approach was based on a generalized inverse theory [24] using a probabilistic framework where assertions are represented by conditional probability density functions (belief distributions). The importance of the method is that it provides a formal recipe for representing and combining all prior knowledge in order to obtain these distributions. We have illustrated how to apply the theory to solve a 3-D model-based recognition problem and have shown how the resulting belief distributions can be used to assess the quality of the interpretation. An important characteristic of the resulting belief distributions is that they are bi-modal, simplifying the problem of determining how to distinguish between informative and uninformative viewpoints.

We have also demonstrated that some viewpoints can give rise to ambiguous information, where the system has confidence in more than one hypothesis. Similar to the motivation behind *autonomous exploration* in the model-building phase [28], ambiguous views have spawned the development of an incremental recognition scheme, where we seek information from a new viewpoint to reduce the overall ambiguity. We have shown how evidence, in the form of the belief distributions, can be accumulated from a sequence of views. The experiments have demonstrated that the maximum likelihood hypothesis is largely viewpoint-invariant, implying that merging votes for the different hypotheses over a sequence of views should lead to a clear winner. Because the beliefs are not normalized, we have given equal weighting to all hypotheses by binarizing the values above a threshold. We have illustrated that by histogramming the binarized beliefs and picking the highest score of the result, we choose the correct winner in all cases.

An important contribution of the method is its potential for a wide variety of applications. For example, an active recognition agent can choose viewpoints that will maximize the belief distribution associated with an object of interest. We have not specified *how* to choose this viewpoint, but the method can be used to determine if the particular choice leads to a sufficient level of information. Another important application of the methodology is a strategy for off-line computation of a pre-computed set of characteristic views. One can rank these views according to the belief distributions, and then store the n best views. Predefining these views speeds up on-line computations by directing the active agent's attention to informative views, thereby reducing the search space of viable hypotheses. These and other topics are currently under investigation in our laboratory.

References

1. E. Aloimonos, Y. Purposive, qualitative, active vision. *CVGIP: Image Understanding*, 56(1):3–129, 1992. special issue.
2. T. Arbel, P. Whaite, and F. P. Ferrie. Recognizing volumetric objects in the presence of uncertainty. Technical Report TR-CIM-94-03, Center for Intelligent Machines, McGill

University, Montréal, Québec, Canada, Mar. 1994. Available via ftp at `ftp.cim.mcgill.ca` in `/pub/3d/papers/tr-cim-94-04.ps.gz`.

3. T. Arbel, P. Whaite, and F. P. Ferrie. Recognizing volumetric objects in the presence of uncertainty. In *Proceedings 12th International Conference on Pattern Recognition*, pages 470–476, Jerusalem, Israel, Oct 9-13 1994. IEEE Computer Society Press.
4. F. Arman and J. Aggarwal. Model-based object recognition in dense-range images - a review. *ACM Computing Surveys*, 25(1), apr 1993.
5. R. Bajcsy and F. Solina. Three dimensional object recognition revisited. In *Proceedings, 1ST International Conference on Computer Vision*, London,U.K., June 1987. Computer Society of the IEEE, IEEE Computer Society Press.
6. A. H. Barr. Superquadrics and angle preserving transformations. *IEEE Computer Graphics and Applications*, 1(1):11–23, Jan. 1981.
7. K. Bowyer and C. Dyer. Aspect graphs: An introduction and survey of recent results. In *Close Range Photogrammetry Meets Machine Vision, Proc. of SPIE*, volume 1395, pages 200–208, 1990.
8. R. T. Chin and C. R. Dyer. Model-based recognition in robot vision. *Computing Surveys*, 18(1):67–108, mar 1986.
9. T. Darrell, S. Sclaroff, and A. Pentland. Segmentation by minimal description. In *Proceedings, 3RD International Conference on Computer Vision*, pages 112–116, Osaka,Japan, Dec. 1990. Computer Society of the IEEE, IEEE Computer Society Press.
10. D. Eggert and K. Bowyer. Computing the orthographic projection aspect graph of solids of revolution. In *PROC. of IEEE Workshop on the Interpretation of 3-D Scenes*, pages 102–108, Austin, Texas, November 27-29 1989. IEEE.
11. D. Eggert, K. Bowyer, C. Dyer, H. Christensen, and D. Goldgof. The scale space aspect graph. In *Proceedings, Conference on Computer Vision and Pattern Recognition*, pages 335–340, Champaign, Il., June 15-18 1992. IEEE.
12. F. P. Ferrie, J. Lagarde, and P. Whaite. Darboux frames, snakes, and super-quadrics: Geometry from the bottom up. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(8):771–784, Aug. 1993.
13. D. Keren, D. Cooper, and J. Subrahmonia. Describing complicated objects by implicit polynomials. Technical Report 102, Brown University LEMS, Laboratory for Engineering Man/Macine Systems, Division of Engineering, Brown University, Providence RI 021912 USA, 1992.
14. J. Koenderink. The singularities of the visual mapping. *Biological Cybernetics*, 24(1):51–59, 1976.
15. J. Koenderink. Internal representation of solid shape with respect to vision. *Biological Cybernetics*, 32(4):211–216, 1979.
16. D. Kriegman and J. Ponce. Computing exact aspect graphs of curved objects: Solids of revolution. In *PROC. of IEEE Workshop on the Interpretation of 3-D Scenes*, pages 116–122, Austin, Texas, November 27-29 1989. IEEE.
17. A. Lejeune and F. Ferrie. Partitioning range images using curvature and scale. In *PROC. IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 800–801, New York City, New York, June 15-17 1993.
18. D. Marr. *Vision*. W.H. Freeman & Co., San Francisco, 1982.
19. A. Pentland and S. Sclaroff. Closed form solutions for physically based shape modelling and recognition. In T. Kanade and K. Ikeuchi, editors, *IEEE Transactions on Pattern Analysis and Machine Intelligence: Special Issue on Physical Modeling in Computer Vision*, volume 13(7),

pages 715–729, July 1991.

20. N. S. Raja and A. K. Jain. Recognizing geons from superquadrics fitted to range data. *Image and Vision Computing*, April 1992.
21. G. Soucy. View correspondence using curvature and motion consistency. Master's thesis, Dept. of E.E., McGill Univ., 1992.
22. T. Sripradisvarakul and R. Jain. Generating aspect graph for curved objects. In *PROC. of IEEE Workshop on the Interpretation of 3-D Scenes*, pages 109–115, Austin, Texas, November 27-29 1989. IEEE.
23. J. Subrahmonia, D. B. Cooper, and D. Keren. Practical reliable bayesian recognition of 2D and 3D objects using implicit polynomials and algebraic invariants. LEMS 107, Brown University LEMS, Laboratory fo Engineering Man/Machine systems, Division of Engineering, Brown University, Providence, RI 02912, USA, 1992.
24. A. Tarantola. *Inverse Problem Theory: Methods for Data Fitting and Model Parameter Estimation.* Elsevier Science Publishing Company Inc., 52, Vanderbuilt Avenue, NewYork, NY 10017, U.S.A., 1987.
25. P. Whaite and F. Ferrie. Model building and autonomous exploration. In *SPIE - Intelligent Robots and Computer Vision XII: Active Vision and 3D Methods*, pages 73–85, Boston, Massachusetts, Sept. 8–9 1993.
26. P. Whaite and F. P. Ferrie. From uncertainty to visual exploration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(10):1038–1049, Oct. 1991.
27. P. Whaite and F. P. Ferrie. Active exploration: Knowing when we're wrong. In *PROC. Fourth International Conference on Computer Vision*, pages 41–48, Berlin, Germany, May 11-14 1993. Computer Society of the IEEE, IEEE Computer Society Press.
28. P. Whaite and F. P. Ferrie. Autonomous exploration: Driven by uncertainty. In *Proceedings, Conference on Computer Vision and Pattern Recognition*, pages 339–346, Seattle, Washington, June 21-23 1994. Computer Society of the IEEE, IEEE Computer Society Press.