Predicting an Upper Bound on SAR ATR Performance

MICHAEL BOSHRA, Member, IEEE

BIR BHANU, Fellow, IEEE
University of California, Riverside

We present a method for predicting a tight upper bound on performance of a vote-based approach for automatic target recognition (ATR) in synthetic aperture radar (SAR) images. In such an approach, each model target is represented by a set of SAR views, and both model and data views are represented by locations of scattering centers. The proposed method considers data distortion factors such as uncertainty, occlusion, and clutter, as well as model factors such as structural similarity. Firstly, we calculate a measure of the similarity between a given model view and each view in the model set, as a function of the relative transformation between them. Secondly, we select a subset of possible erroneous hypotheses that correspond to peaks in similarity functions obtained in the first step. Thirdly, we determine an upper bound on the probability of correct recognition by computing the probability that every selected hypothesis gets less votes than those for the model view under consideration. The proposed method is validated using MSTAR public SAR data, which are obtained under different depression angles, configurations, and articulations.

Manuscript received June 7, 1999; revised August 28, 2000; released for publication April 27, 2001.

IEEE Log No. T-AES/37/3/08559.

Refereeing of this contribution was handled by L. M. Kaplan.

This work was supported by DARPA/AFOSR Grant F49620-97-1-0184; the contents and information do not reflect positions or policies of the U.S. Government.

Authors' address: Center for Research in Intelligent Systems, University of California, B 232, Bourns Hall, Riverside, CA 92521, E-mail: ({michael, bhanu}@cris.ucr.edu).

I. INTRODUCTION

Model-based automatic target recognition (ATR) has received considerable attention during the last two decades [1, 2, 4, 18]. It involves comparing features extracted from scene data with those of targets in a model database. The features obtained by processing sensory data are corrupted by distortion factors, such as uncertainty (error in estimating feature locations and attributes), occlusion (missing features), and clutter (spurious features). Performance of recognition depends on these data factors, in addition to model factors such as the number of model targets, and the degree of similarity between them. The model-similarity factor plays a fundamental role in determining the recognition performance. Intuitively, a model target that is distinct from all other models can be correctly recognized, even in the presence of high levels of data distortion. On the other hand, if the target is similar to some other model targets, then there is a high likelihood of it being "confused" with one of the similar targets, even in the presence of moderate distortion levels. Accordingly, we can say that recognition performance is *inversely* proportional to the degree of model similarity.

We consider the above-mentioned data and model factors within a framework for predicting an upper bound on performance of a vote-based approach for ATR using synthetic aperture radar (SAR) data. In such an approach, each model target is represented by a set of SAR views that sample its signature at a variety of azimuth angles and a specific depression angle. The features used for recognition are scattering centers, which are local image peaks. These features have been commonly used for SAR ATR (e.g., [3, 9, 11, 20]). Recognition is performed by matching locations of scattering centers extracted from a data view with those of the model views. Each model hypothesis (a view of a specific target and associated location) is evaluated based on the number of model features that are consistent with data features (votes). The hypothesis that has the maximum number of votes is selected as the valid one.

The proposed performance prediction method can be outlined as follows. We are interested in determining an upper bound on the probability of correct recognition (PCR) of a given model view, in the presence of data distortion. Our approach for determining such a bound is to select a subset of possible erroneous hypotheses, and then compute the probability that each one of them gets less votes than those for the "distorted" model view. In order to obtain a reasonably *tight* upper bound, the selected hypotheses should satisfy the following condition: In case of recognition failure, there should be a good chance that the votes for at least one of the selected hypotheses reach or exceed those for the model view. Furthermore, in order to ensure that

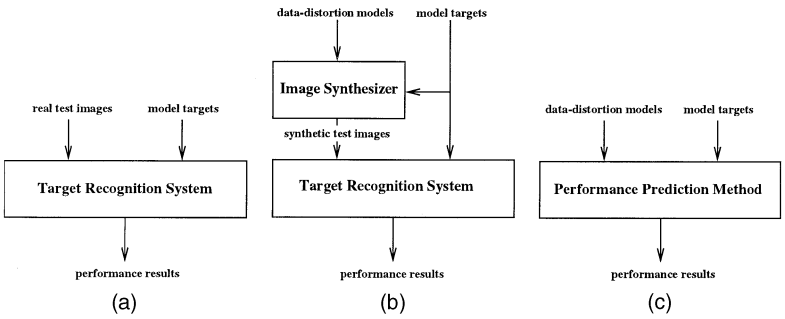

Fig. 1. Illustration of performance prediction approaches. (a) Empirical using real data. (b) Empirical using synthetic data. (c) Formal.

the computations are *mathematically tractable*, the assumption of statistical vote independence for the selected hypotheses should be reasonable. We obtain a subset of erroneous hypotheses satisfying both of these conditions by utilizing a quantitative measure of model similarity. In particular, we first compute the similarity between the given model view and each of the views in the model database, as a function of the relative transformation between them. Then, we select erroneous hypotheses that correspond to *peaks* in these similarity functions. The proposed method is validated using MSTAR public SAR data, obtained under different depression angles, configurations, and articulations.

The remainder of the paper is organized as follows. The next section reviews relevant efforts for predicting SAR ATR performance, and highlights our contributions. Section III formally defines the performance prediction problem considered in this research. The proposed similarity measure is described in Section IV. The method used to predict an upper bound on performance is presented in Section V. Its validity is demonstrated in Section VI by comparing experimentally determined PCR plots with predicted upper bounds. Finally, conclusions are drawn in Section VII.

II. RELEVANT RESEARCH AND OUR CONTRIBUTIONS

The approaches for ATR performance prediction can be classified into the following three categories.

1) *Empirical Approaches Using Real Data*: Performance is determined through experimentation using real data, as shown in Fig. 1(a). For example, Novak et al. [17] evaluated the performance of an ATR system using real polarimetric SAR images of camouflaged/uncamouflaged targets, and clutter. The estimated performance metrics included the probabilities of detection and false alarms, in addition to the confusion matrix for four classes (three targets and clutter). Catlin et al. [8] presented a method for estimating the function relating the probability of correct identification to aspect and depression angles using SAR data. Each point on the function was determined experimentally by passing a number of real images to an ATR system. The resulting function was modeled using either linear regression or an artificial neural network. Bhanu and Jones [3] used real SAR data, from the MSTAR public domain, to characterize the performance of a SAR ATR system that utilized scattering-center locations and magnitudes. The data, obtained under different depression angles, configurations, and articulations, were used to estimate receiver operating characteristic (ROC) curves relating the probability of correct identification to the probability of false alarm.

2) *Empirical Approaches Using Synthetic Data*: Performance is determined through experimentation using synthetic data (e.g., using Monte Carlo simulation). The data are obtained by generating images of model targets, and artificially introducing distortion to them (see Fig. 1(b)). Irving et al. [16] used Monte Carlo simulation to estimate classification performance using peak features represented by their locations. Performance was estimated as a function of factors such as number of target classes, peak positional uncertainty, and clutter rate. Horowitz and Brendel [13] estimated fundamental classification error considering two target views, and a Rayleigh distribution as a noise model. The classification error was determined using Monte Carlo trials. Diehl et al. [10] described a method to determine a number of performance metrics, such as the confusion matrix as a function of image resolution. These metrics were estimated experimentally using synthetic target images, which were degraded according to a noise model. Jones and Bhanu [14] empirically estimated a number of metrics for evaluating the performance of a scattering-center-based ATR system using XPATCH SAR data. These metrics included the probability of correct identification, and the confusion matrix (for four targets). The model data were distorted by changing model articulations (e.g., changing the orientation of a tank turret), eliminating model features, adding clutter features, and perturbing feature locations.

3) *Formal Approaches*: A statistical framework is developed to directly predict performance for a given set of statistical data-distortion models (see

Fig. 1(c)). Grenander et al. [12] presented a method to predict fundamental performance of template-based ATR for a given image noise model (a probabilistic likelihood function) using Hilbert–Schmidt estimator. The expected error of this estimator (Hilbert–Schmidt bound) forms a lower bound on the error of any other estimator. Irving et al. [15] presented a formal method for predicting performance of SAR target detection. Their method, which considered a target detection approach involving peak features, was based on statistical modeling of both data distortion factors, and model target views.

Most ATR performance prediction approaches belong to the first two categories. There is a lack of formal methods that predict performance based on the amount of information provided by sensory data [18]. Such approaches can provide a valuable insight into the roles played by different data and model factors (such as those mentioned in Section I) in determining the recognition performance.

The prediction method presented here belongs to the third category, the formal approaches. Compared with previous relevant work, the method presented here is unique in the following aspects.

1) An upper bound on performance of vote-based SAR ATR, using point features, is predicted.

2) Data uncertainty, occlusion, clutter and model similarity are considered simultaneously for performance prediction.

3) Real SAR data (public MSTAR) are used for validating the prediction method.

We note that prediction of a lower bound on performance is reported elsewhere by Boshra and Bhanu [6].

III. PROBLEM DEFINITION

In this section, we present a formal definition of the performance prediction problem considered in this work.

Our SAR ATR approach can be defined as follows. We are given the following.

1) Model Database $\mathcal{MD}$: This database consists of a set of model views, belonging to model targets under consideration. Each model view, $\mathcal{M}_i \in \mathcal{MD}$, is represented by the 2-D locations of a set of scattering centers, $\{F_{ik}\}$. These locations are assumed to be discretized at some resolution.

2) Data View $\hat{\mathcal{M}}$: This view belongs to one of the model targets. As is the case with model views, $\hat{\mathcal{M}}$ is represented by discretized 2-D locations of a set of scattering centers, $\{\hat{F}_l\}$.

3) Transformation Space $\mathcal{T}$: This space defines possible transformations between data and model views. Notice that $\mathcal{T}$ is discretized, since we are dealing with discretized features. In our SAR ATR approach, we assume that $\mathcal{T}$ is the space of 2-D translations [14]. The prediction method, however, is applicable to other transformations as well.

The objective is to identify the model view that corresponds to $\hat{\mathcal{M}}$. The recognition process can be outlined as follows. Let $\mathcal{M}_i^{\hat{\tau}}$ be a *hypothesis* of model view $\mathcal{M}_i$ at location $\hat{\tau} \in \mathcal{T}$ with respect to data view $\hat{\mathcal{M}}$. Recognition involves matching $\hat{\mathcal{M}}$ with $\mathcal{M}_i^{\hat{\tau}}$, for all $\mathcal{M}_i \in \mathcal{MD}$, and $\hat{\tau} \in \mathcal{T}$, and selecting the best matching hypothesis. The criterion used for matching is the number of model features that are *consistent* with data features. Each consistent model feature is said to "vote" for the associated hypothesis. The decision on whether a pair of data and model features is consistent depends on the description of the positional uncertainty of the data features. In the following, we define both the data uncertainty model considered in this work, and the notion of feature consistency.

Data Uncertainty: The uncertainty associated with data feature $\hat{F}_l$ is a probabilistic description of the location of F_{ik}, the model feature that $\hat{F}_l$ corresponds to. We assume that this uncertainty is described by a uniform distribution. Accordingly, it can be represented by an *uncertainty region* $R_u(\cdot)$. In particular, we have $F_{ik}^{\hat{\tau}} \in R_u(\hat{F}_l)$, where $F_{ik}^{\hat{\tau}} \in \mathcal{M}_i^{\hat{\tau}}$. Furthermore, we assume that the uncertainty distributions associated with the data features are independent. Notice that the data uncertainty model considered here does *not* include the relative data/model target transformation. This is because, as mentioned above, our recognition mechanism compensates for it (through examining possible relative transformations).

Feature Consistency: A pair of data and model features, $\hat{F}_l$ and $F_{ik}^{\hat{\tau}}$, is said to be consistent if $\hat{F}_l$ can be interpreted as an uncertain measurement of $F_{ik}^{\hat{\tau}}$. It can be easily shown that the condition for consistency is $F_{ik}^{\hat{\tau}} \in R_u(\hat{F}_l)$, or equivalently $\hat{F}_l \in \bar{R}_u(F_{ik}^{\hat{\tau}})$, where $\bar{R}_u(\cdot)$ is the reflection of $R_u(\cdot)$ about the origin. We refer to $\bar{R}_u(\cdot)$ as a *consistency region*. In most practical situations, region $R_u(\cdot)$ is symmetric about the origin (e.g., a circle, ellipse or square). In such a case, we have $R_u(\cdot) = \bar{R}_u(\cdot)$.

From the above two definitions, it can be shown that the number of votes for hypothesis $\mathcal{M}_i^{\hat{\tau}}$ given data view $\hat{\mathcal{M}}$, VOTES$(\mathcal{M}_i^{\hat{\tau}}; \hat{\mathcal{M}})$, can be expressed as follows:

$$\text{VOTES}(\mathcal{M}_i^{\hat{\tau}}; \hat{\mathcal{M}}) = |\{F_{ik}^{\hat{\tau}} : F_{ik}^{\hat{\tau}} \in \mathcal{M}_i^{\hat{\tau}} \text{ and } \exists \hat{F}_l \in \hat{\mathcal{M}} \text{ s.t. } \hat{F}_l \in \bar{R}_u(F_{ik}^{\hat{\tau}})\}|. \quad (1)$$

Next, we define our performance prediction problem. Let us assume that we are given a data view $\hat{\mathcal{M}}_i$ which is a distorted version of model view $\mathcal{M}_i$. In case of recognition failure, $\hat{\mathcal{M}}_i$ will be misinterpreted

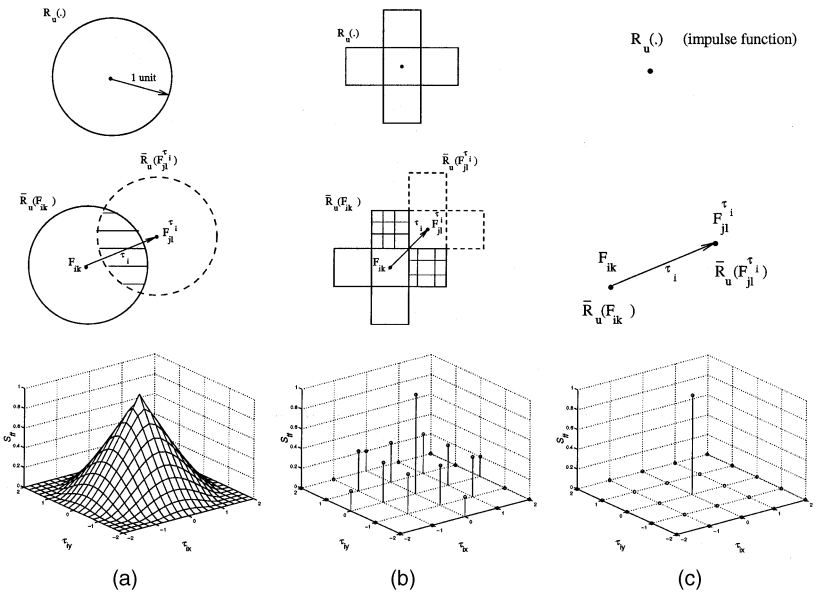

Fig. 2. Illustration of feature/feature similarity for three uncertainty regions. (a) Circle of unit radius. (b) Discrete four-neighbor region. (c) Point region (implying that there is no positional uncertainty). Note that τ_{ix}, and τ_{iy} represent the components of τ_i along the x and y axes, respectively (for sake of simplicity, we assume $F_{ik} = F_{jl}^{\tau_i}|_{\tau_i=0}$).

as one of a set of possible erroneous hypotheses, $\mathcal{H}_i$. This set can be formally defined as

$$\mathcal{H}_i = \{\mathcal{M}_j^{\tau_i} : |\mathcal{M}_j^{\tau_i} \cap R_c| > 0,\ \forall \mathcal{M}_j \in \mathcal{MD},\ \text{and}\ \tau_i \in \mathcal{T}\} - \{\mathcal{M}_i\}$$

where $\mathcal{M}_j^{\tau_i}$ is a hypothesis of model view $\mathcal{M}_j$ at location $\tau_i \in \mathcal{T}$ with respect to view $\mathcal{M}_i$, and R_c is the *clutter region*, which is a region surrounding $\mathcal{M}_i$ that bounds the locations of spurious features in $\hat{\mathcal{M}}_i$. The condition for the correct recognition of $\hat{\mathcal{M}}_i$ as $\mathcal{M}_i$ (or equivalently $\mathcal{M}_i = \mathcal{M}_i^{\tau_i}|_{\tau_i=0}$, where 0 is the origin of $\mathcal{T}$) is that the number of votes for $\mathcal{M}_i$ is greater than the number of votes for *every* possible erroneous hypothesis. Accordingly, we can express the probability of correctly recognizing $\mathcal{M}_i$ in distorted view $\hat{\mathcal{M}}_i$ as

$$\Pr[\mathcal{M}_i;\hat{\mathcal{M}}_i] = \Pr[\forall \mathcal{M}_j^{\tau_i} \in \mathcal{H}_i : V_i > V_j^{\tau_i};\hat{\mathcal{M}}_i] \qquad (2)$$

where V_i is the number of votes for $\mathcal{M}_i$ (i.e., V_i = VOTES($\mathcal{M}_i;\hat{\mathcal{M}}_i$)), and $V_j^{\tau_i}$ is the number of votes for hypothesis $\mathcal{M}_j^{\tau_i}$ (i.e., $V_j^{\tau_i}$ = VOTES($\mathcal{M}_j^{\tau_i};\hat{\mathcal{M}}_i$)). Our objective is to determine an upper bound on $\Pr[\mathcal{M}_i;\hat{\mathcal{M}}_i]$.

IV. MODEL SIMILARITY

In this section, we define a quantitative measure of the structural similarity between a pair of model views represented by locations of point features (e.g., scattering centers in SAR images).

The proposed similarity measure is dependent upon the amount of uncertainty in the data. This agrees with the intuitive observation that as views get "blurred," it becomes more difficult to distinguish between them, which is equivalent to saying that they become more "similar" to each other. In the following, we introduce a sequence of definitions that lead to a formal definition of model similarity.

Feature/Feature Similarity: The similarity between a pair of model features, $F_{ik} \in \mathcal{M}_i$ and $F_{jl}^{\tau_i} \in \mathcal{M}_j^{\tau_i}$, denoted by $S_{ff}(F_{ik},F_{jl}^{\tau_i})$, is defined as the probability that an uncertain measurement of F_{ik} is consistent with $F_{jl}^{\tau_i}$. Formally,

$$S_{ff}(F_{ik},F_{jl}^{\tau_i}) = \frac{\text{AREA}(\bar{R}(F_{ik}) \cap \bar{R}(F_{jl}^{\tau_i}))}{\text{AREA}(\bar{R}(F_{ik}))}$$

where AREA(R) is the area of region R. Notice that $S_{ff}(F_{ik},F_{jl}^{\tau_i})$, which lies in the range $[0,1]$, is proportional to the extent of overlap between $\bar{R}(F_{ik})$ and $\bar{R}(F_{jl}^{\tau_i})$. We sometimes refer to feature pairs with overlapping/nonoverlapping consistency regions as *similar/dissimilar* features, respectively. Fig. 2 shows $S_{ff}(F_{ik},F_{jl}^{\tau_i})$ as a function of τ_i for a sample of three uncertainty/consistency regions, assuming that $\mathcal{T}$ is the space of 2-D translations.

View/Feature Similarity: The similarity between a view $\mathcal{M}_i$ and a feature $F_{jl}^{\tau_i} \in \mathcal{M}_j^{\tau_i}$, is defined as the probability that an uncertain measurement of *any* feature in $\mathcal{M}_i$ is consistent with $F_{jl}^{\tau_i}$. Formally, view/feature similarity, $S_{vf}(\mathcal{M}_i,F_{jl}^{\tau_i})$, can be defined as follows:

$$S_{vf}(\mathcal{M}_i,F_{jl}^{\tau_i}) = 1 - \prod_k (1 - S_{ff}(F_{ik},F_{jl}^{\tau_i})).$$

View/Hypothesis Similarity: The structural similarity between a view $\mathcal{M}_i$ and a view hypothesis

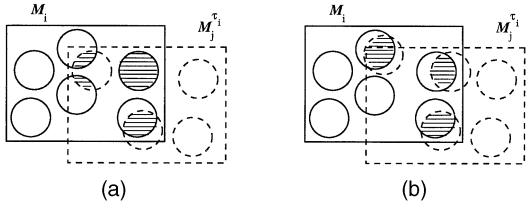

Fig. 3. View/hypothesis similarity between $\mathcal{M}_i$ and $\mathcal{M}_j^{\tau_i}$ (solid and dashed circles denote consistency regions associated with the features of $\mathcal{M}_i$ and $\mathcal{M}_j^{\tau_i}$, respectively). (a) Structural similarity: notice that $S_j^{\tau_i} \in [1,3]$, and $E(S_j^{\tau_i})$ is about 2. (b) Uniform structural similarity: notice that feature/feature similarity is about $\frac{2}{3}$ for each of the three similar pairs of features, and so distribution of $S_j^{\tau_i}$ is approximated by $B_{S_j^{\tau_i}}(s;3,\frac{2}{3})$.

$\mathcal{M}_j^{\tau_i}$, is defined as the number of features in $\mathcal{M}_j^{\tau_i}$ consistent with uncertain measurements of features in $\mathcal{M}_i$. Obviously, view/hypothesis similarity, denoted as $S_{vh}(\mathcal{M}_i,\mathcal{M}_j^{\tau_i})$ or simply $S_j^{\tau_i}$, is a random variable, which is bounded as follows:

$$\min(S_j^{\tau_i}) = |\{F_{jl}^{\tau_i} : F_{jl}^{\tau_i} \in \mathcal{M}_j^{\tau_i} \text{ and } S_{vf}(\mathcal{M}_i,F_{jl}^{\tau_i}) = 1\}|$$

$$\max(S_j^{\tau_i}) = |\{F_{jl}^{\tau_i} : F_{jl}^{\tau_i} \in \mathcal{M}_j^{\tau_i} \text{ and } S_{vf}(\mathcal{M}_i,F_{jl}^{\tau_i}) > 0\}|.$$

Fig. 3(a) shows an example of view/hypothesis similarity and associated bounds.

In order to simplify the estimation of view/hypothesis similarity, and subsequently the prediction of a performance bound in the next section, we make the following reasonable assumptions about $\mathcal{M}_i$, $\mathcal{M}_j^{\tau_i}$ and the similarity between their features. 1) There is no overlap between the consistency regions of the features that belong to each of $\mathcal{M}_i$ and $\mathcal{M}_j^{\tau_i}$, 2) there is a one-to-one correspondence between similar features in $\mathcal{M}_i$ and $\mathcal{M}_j^{\tau_i}$, and 3) the similarity between every pair of similar features in $\mathcal{M}_i$ and $\mathcal{M}_j^{\tau_i}$ is a constant value, which is the average view/feature similarity. These assumptions result in a "uniform" view of the structural similarity between $\mathcal{M}_i$ and $\mathcal{M}_j^{\tau_i}$. As an example, Fig. 3(b) illustrates the uniform similarity corresponding to the view/hypothesis pair shown in Fig. 3(a). Based on the uniform-similarity assumption, we can express the expected value of $S_j^{\tau_i}$ as

$$E(S_j^{\tau_i}) = \sum_l S_{vf}(\mathcal{M}_i,F_{jl}^{\tau_i}).$$

In addition, we can represent the probability distribution function (pdf) of $S_j^{\tau_i}$ by the following binomial distribution:

$$\Pr[S_j^{\tau_i} = s] = B_{S_j^{\tau_i}}(s;N_j^{\tau_i},Q_j^{\tau_i})$$

where

$$N_j^{\tau_i} = \max(S_j^{\tau_i}) \tag{3}$$

$$Q_j^{\tau_i} = \frac{E(S_j^{\tau_i})}{N_j^{\tau_i}} \tag{4}$$

$B_X(x;n,p) = C(n,x)p^x(1-p)^{n-x}$, and $C(a,b) = (a!/(a-b)!b!)$. Notice that the similarity between $\mathcal{M}_i$ and $\mathcal{M}_j^{\tau_i}$ is quantified by the tuple $(N_j^{\tau_i},Q_j^{\tau_i})$.

View/View Similarity: The similarity between a pair of views $\mathcal{M}_i$ and $\mathcal{M}_j$, is defined as the view/hypothesis similarity $S_j^{\tau_i}$, for all $\tau_i \in \mathcal{T}$. Thus, the view/view similarity is a probabilistic function. As an illustration, Figs. 4(a) and 4(b) show a pair of simple model views and the corresponding expected similarity function, respectively. Peaks in this function

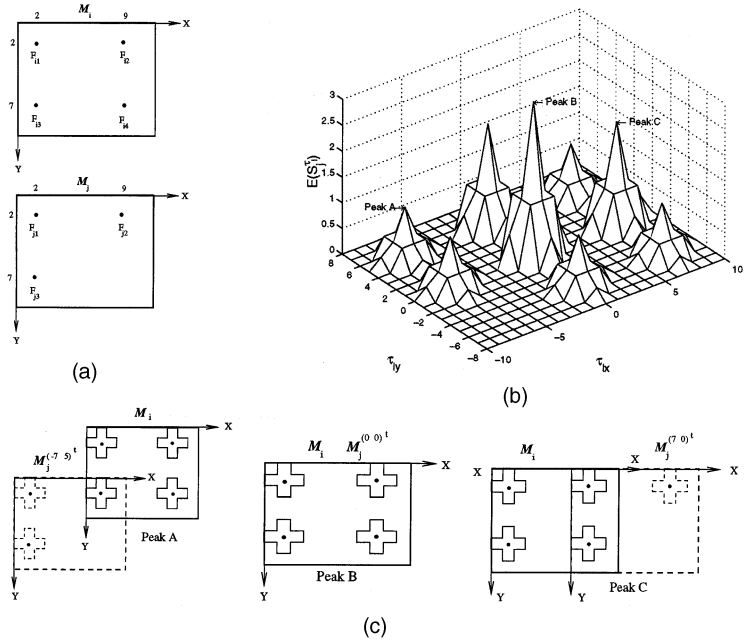

Fig. 4. Illustration of view/view similarity assuming four-neighbor discrete consistency region, and discrete 2-D translation space. (a) Pair of model views $\mathcal{M}_i$ and $\mathcal{M}_j$. (b) Corresponding expected-similarity function $E(S_j^{\tau_i})$. (c) Sample of three peak hypotheses.

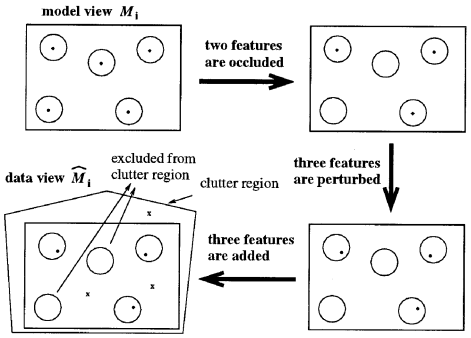

Fig. 5. Simple example showing sequence of distortion processes that transform a model view $\mathcal{M}_i$ to a corresponding data view $\hat{\mathcal{M}}_i$. Circles shown represent consistency regions associated with features of $\mathcal{M}_i$. Clutter features are shown as crosses to distinguish them from perturbed features of $\mathcal{M}_i$ that are observed in $\hat{\mathcal{M}}_i$.

correspond to view hypotheses that have a higher degree of similarity with $\mathcal{M}_i$ than the neighboring ones. A sample of three *peak hypotheses* is shown in Fig. 4(c). In the next section, we use peak hypotheses for predicting an upper bound on ATR performance.

V. PERFORMANCE PREDICTION METHOD

In this section, we present a method for predicting an upper bound on recognition performance, in the presence of data distortion.

A. Modeling of Data Distortion

A model view $\mathcal{M}_i$ is assumed to be distorted to form a data view $\hat{\mathcal{M}}_i$ in the following sequence (refer to the example shown in Fig. 5).

1) A group of O model features are *occluded*. We assume that occlusion is uniform, i.e., each subset of features of size O is equally likely to be occluded as any other subset of the same size.

2) Unoccluded model features are uniformly *perturbed* within associated consistency regions, $\bar{R}_u(\cdot)$.

3) Resulting (occluded and perturbed) model view is *cluttered* by adding C spurious features to form $\hat{\mathcal{M}}_i$. We assume that the distribution of these features is uniform within a clutter region R_c, which encompasses the original model view (e.g., convex hull, bounding box). In order to control the number of votes for $\mathcal{M}_i$, we exclude the consistency regions associated with the occluded features of $\mathcal{M}_i$ from the original clutter region (see Fig. 5). The resulting region R_c' is referred to as the *effective clutter region*.

B. Predicting an Upper Bound on Performance

The probability of correctly recognizing $\mathcal{M}_i$, given $\hat{\mathcal{M}}_i$, depends on the votes for $\mathcal{M}_i$, V_i, and those for the erroneous hypotheses in $\mathcal{H}_i$, $\{V_j^{\tau_i}\}$ (refer to (2)). According to our modeling of the data distortion process, described in Section VA, it is easy to obtain V_i:

$$V_i = |\mathcal{M}_i| - O. \tag{5}$$

On the other hand, $V_j^{\tau_i}$ is a random variable that depends on the degree of similarity between $\mathcal{M}_i$ and $\mathcal{M}_j^{\tau_i}$, and the number of clutter features that happen to be consistent with features in $\mathcal{M}_j^{\tau_i}$. The pdf of $V_j^{\tau_i}$ is derived in the next section.

Determining $\Pr[\mathcal{M}_i;\hat{\mathcal{M}}_i]$, defined in (2), is a very complicated task, because the random variables $\{V_j^{\tau_i}\}$ are not independent. This is especially true for adjacent hypotheses, due to the spatial overlap between their consistency regions. The complexity of this task is the reason that we are resorting to the estimation of performance bounds, rather than attempting to estimate the actual performance.

We present three possible upper bounds on $\Pr[\mathcal{M}_i;\hat{\mathcal{M}}_i]$, which vary in their tightness and reliance on assumptions. A recognition error occurs when the number of votes for *any* of the hypotheses in $\mathcal{H}_i$ reaches or exceeds the number of votes for $\mathcal{M}_i$. This event occurs with a specific probability for every hypothesis in $\mathcal{H}_i$. The maximum of these probabilities forms a *lower* bound on the probability of recognition error. This directly leads to the following upper bound on $\Pr[\mathcal{M}_i;\hat{\mathcal{M}}_i]$ (refer to (2) and (5)):

$$\Pr[\mathcal{M}_i;\hat{\mathcal{M}}_i] < 1 - \max_{\tau_i,j} \Pr[V_j^{\tau_i} \geq |\mathcal{M}_i| - O;\hat{\mathcal{M}}_i] \tag{6}$$

where $\mathcal{M}_j^{\tau_i} \in \mathcal{H}_i$. Obviously, this bound is not expected to be tight, since it considers only a single hypothesis in $\mathcal{H}_i$.

In order to obtain a tighter bound, one possible approach is to select a subset of hypotheses in $\mathcal{H}_i$, and assume that the votes for those that belong to the *same* view are independent. In our work, we consider hypotheses that correspond to *peaks* in the expected similarity function (refer to Section IV). The set of peak hypotheses associated with model view $\mathcal{M}_i$, $\mathcal{H}_{pi}$, can be defined as follows:

$$\mathcal{H}_{pi} = \{\mathcal{M}_j^{\tau_i} : \mathcal{M}_j^{\tau_i} \in \mathcal{H}_i, \text{ and } \tau_i \text{ is a peak in } E(S_j^{\tau_i})\}.$$

The rationale behind the choice of peak hypotheses can be stated as follows. 1) Peak hypotheses are more similar to the given model view than the neighboring hypotheses, and so, a distorted model view is generally more likely to be misinterpreted as a peak hypothesis than an off-peak neighboring one. 2) Peak hypotheses are not very close to each other, thus making the vote-independence assumption reasonable. Based on this approach, we can obtain the following upper bound:

$$\Pr[\mathcal{M}_i;\hat{\mathcal{M}}_i] < \min_j \prod_{\tau_i} (1 - \Pr[V_j^{\tau_i} \geq |\mathcal{M}_i| - O;\hat{\mathcal{M}}_i]) \tag{7}$$

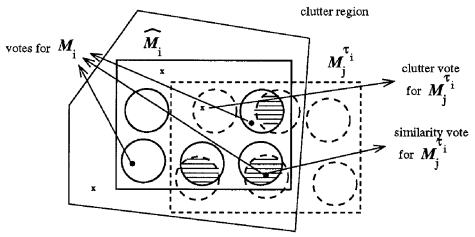

Fig. 6. Illustration of votes for model view $\mathcal{M}_i$ and similarity and clutter votes for view hypothesis $\mathcal{M}_j^{\tau_i}$. Solid and dashed circles represent consistency regions of features that belong to $\mathcal{M}_i$ and $\mathcal{M}_j^{\tau_i}$, respectively, and crosses represent clutter features. Empty solid circles correspond to occluded features of $\mathcal{M}_i$. Features that lie within solid (dashed) circle vote for $\mathcal{M}_i$ ($\mathcal{M}_j^{\tau_i}$).

where $\mathcal{M}_j^{\tau_i} \in \mathcal{H}_{pi}$. Obviously, the above bound is expected to be considerably tighter than the one defined in (6), since it considers a representative subset of hypotheses of a model view (peak ones), not just a single hypothesis as in (6).

A bound that is tighter than the one defined in (7) can be obtained by considering a representative subset of hypotheses that belong to *all* model views, not just a single view as in (7). In order to achieve this goal, we further make the assumption that votes for peak hypotheses belonging to *different* views are independent. Based on such an assumption, we obtain the following upper bound:

$$\Pr[\mathcal{M}_i;\hat{\mathcal{M}}_i] < \prod_{j,\tau_i}(1-\Pr[V_j^{\tau_i} \geq |\mathcal{M}_i| - O;\hat{\mathcal{M}}_i]) \quad (8)$$

where $\mathcal{M}_j^{\tau_i} \in \mathcal{H}_{pi}$. While this assumption is clearly invalid in some extreme cases, such as when we have identical or very similar views, we argue that it is reasonable in most practical SAR ATR scenarios. This is because, even in the presence of similar targets, their corresponding SAR views are still expected to be significantly different, due to the extreme sensitivity of the SAR imaging process to slight geometric variations in target shape.

C. Determining the PDF of Hypothesis Votes

In this section, we outline the derivation of the pdf of $V_j^{\tau_i}$, the number of votes for hypothesis $\mathcal{M}_j^{\tau_i}$ given distorted data view $\hat{\mathcal{M}}_i$.

The votes for $\mathcal{M}_j^{\tau_i}$ come from two different sources: 1) *similarity*: due to the structural similarity between $\mathcal{M}_j^{\tau_i}$ and $\mathcal{M}_i$, and 2) *clutter*: due to the random coincidence between clutter features and consistency regions of features in $\mathcal{M}_j^{\tau_i}$ (an example of similarity and clutter votes is shown in Fig. 6). Accordingly, we can express $V_j^{\tau_i}$ as

$$V_j^{\tau_i} = V_{sj}^{\tau_i} + V_{cj}^{\tau_i} \quad (9)$$

where $V_{sj}^{\tau_i}$ and $V_{cj}^{\tau_i}$ are the numbers of similarity and clutter votes for $\mathcal{M}_j^{\tau_i}$, respectively. Let $N_{oj}^{\tau_i}$ be a random variable denoting the number of the *unoccluded* features in $\mathcal{M}_i$ that are similar to features in $\mathcal{M}_j^{\tau_i}$ (e.g., in Fig. 6, $N_{oj}^{\tau_i} = 2$). The pdf of $V_{sj}^{\tau_i}$ can be expressed as

$$\Pr[V_{sj}^{\tau_i} = v_s;\hat{\mathcal{M}}_i] = \sum_{n_o}\Pr[N_{oj}^{\tau_i} = n_o;\hat{\mathcal{M}}_i]\Pr[V_{sj}^{\tau_i} = v_s;\ N_{oj}^{\tau_i} = n_o,\hat{\mathcal{M}}_i]. \quad (10)$$

From (9) and (10), we can express the pdf of $V_j^{\tau_i}$ as

$$\Pr[V_j^{\tau_i} = v;\hat{\mathcal{M}}_i] = \sum_{n_o}\Pr[N_{oj}^{\tau_i} = n_o;\hat{\mathcal{M}}_i] \times \sum_{v_s}\Pr[V_{sj}^{\tau_i} = v_s;\ N_{oj}^{\tau_i} = n_o,\hat{\mathcal{M}}_i] \times \Pr[V_{cj}^{\tau_i} = v - v_s;\ N_{oj}^{\tau_i} = n_o, V_{sj}^{\tau_i} = v_s,\hat{\mathcal{M}}_i].$$

In order to simplify the process of computing the pdf of $V_j^{\tau_i}$, we consider the uniform view of similarity described in Section IV. It can be shown that $N_{oj}^{\tau_i}$ is described by the following hypergeometric distribution:

$$\Pr[N_{oj}^{\tau_i} = n_o;\hat{\mathcal{M}}_i] = H_{N_{oj}^{\tau_i}}(N_j^{\tau_i} - n_o;\ O, N_j^{\tau_i}, |\mathcal{M}_i| - N_j^{\tau_i}) \quad (11)$$

where $H_X(x;n,a,b) = C(a,x)C(b,n-x)/C(a+b,n)$. The conditional pdf of $V_{sj}^{\tau_i}$ is the following binomial distribution:

$$\Pr[V_{sj}^{\tau_i} = v_s;\ N_{oj}^{\tau_i} = n_o,\hat{\mathcal{M}}_i] = B_{V_{sj}^{\tau_i}}(v_s;n_o,Q_j^{\tau_i}). \quad (12)$$

Notice that the expected number of similarity votes, $n_o Q_j^{\tau_i}$, is directly proportional to the degree of similarity between $\mathcal{M}_i$ and $\mathcal{M}_j^{\tau_i}$, which is represented by the tuple $(N_j^{\tau_i}, Q_j^{\tau_i})$ (defined in (3) and (4), respectively). This provides a quantitative justification of our claim, at the beginning of Section I, that performance is inversely proportional to the degree of model similarity. One simple approximation of the conditional pdf of $V_{cj}^{\tau_i}$ can be outlined as follows. The effective clutter region R_c', defined in Section VA, is split into two subregions, R_k' and $(R_c' - R_k')$, such that a clutter feature falling within the first (second) subregion does (does not) vote for $\mathcal{M}_j^{\tau_i}$. It can be shown that R_k' is the union of the consistency regions of the features of $\mathcal{M}_j^{\tau_i}$ inside R_c' that are *not* contributing similarity votes. The conditional pdf of $V_{cj}^{\tau_i}$ can be approximated by the following binomial distribution:

$$\Pr[V_{cj}^{\tau_i} = v_c;\ N_{oj}^{\tau_i} = n_o, V_{sj}^{\tau_i} = v_s,\hat{\mathcal{M}}_i] \approx B_{V_{cj}^{\tau_i}}\left(v_c;C,\frac{\text{AREA}(R_k')}{\text{AREA}(R_c')}\right) \quad (13)$$

where

$$\text{AREA}(R_c') = \text{AREA}(R_c) - \text{AREA}(R_u(\cdot))O, \quad \text{and}$$

$$\text{AREA}(R_k') = \text{AREA}(R_u(\cdot))(|\mathcal{M}_j^{\tau_i} \cap R_c| - v_s - Q_j^{\tau_i}(N_j^{\tau_i} - n_o)).$$

The above pdf is simplified since it does not consider factors such as the following. 1) In some cases, clutter features have to be at a minimum distance from each other. For example, if point features extracted from a SAR image correspond to local peaks, then no feature can lie in the eight-neighbor region of another one. 2) Clutter features that fall within the same consistency region contribute only a single vote. A more accurate approximation of the pdf of $V_{cj}^{\tau_i}$, which considers these factors, is presented in [6] (along with more details about determining the pdf of hypothesis votes).

D. Algorithm

In this section, we present the algorithm used to determine an upper bound on average PCR, for a given a model database $\mathcal{MD}$. The algorithm consists of two steps. In the first step, the similarity information involving all model views and associated peak hypotheses, $\mathcal{M}_i/\mathcal{H}_{pi}$ for all i, is calculated and accumulated in a *peak similarity histogram* (PSH). In the second step, the PSH is used, along with the distortion models, to predict an upper bound on PCR. These two steps are described as follows.

1) *Construction of PSH*: From the analysis presented in the previous section, it can be observed that the pdf of $V_j^{\tau_i}$ depends on four view-dependent parameters only: $|\mathcal{M}_i|$, $|\mathcal{M}_j^{\tau_i} \cap R_c|$ (*size* of the section of $\mathcal{M}_j^{\tau_i}$ that lies inside clutter region R_c), and the similarity parameters $(N_j^{\tau_i}, Q_j^{\tau_i})$. Accordingly, only these parameters need to be stored when calculating the similarity between model views. In particular, for each view/peak-hypothesis pair $\mathcal{M}_i/\mathcal{M}_j^{\tau_i}$, where $\mathcal{M}_j^{\tau_i} \in \mathcal{H}_{pi}$, entry $(|\mathcal{M}_i|, |\mathcal{M}_j^{\tau_i} \cap R_c|, N_j^{\tau_i}, \lfloor N_j^{\tau_i} Q_j^{\tau_i} + 0.5 \rfloor)$ in the 4-D PSH is incremented by one.

2) *Computation of PCR Bound*: We consider the upper bound defined in (8). Let

$$T(a,b,c,d) = \Pr[V_j^{\tau_i} \geq |\mathcal{M}_i| - O; \hat{\mathcal{M}}_i]$$

such that $a = |\mathcal{M}_i|$, $b = |\mathcal{M}_j^{\tau_i} \cap R_c|$, $c = N_j^{\tau_i}$, and $d = \lfloor N_j^{\tau_i} Q_j^{\tau_i} + 0.5 \rfloor$. The upper bound on average PCR for model database $\mathcal{MD}$, the main result of this paper, can be approximated as follows:

$$\mathrm{PCR}(\mathcal{MD}) < \prod_a \prod_b \prod_c \prod_d (1 - T(a,b,c,d))^{\mathrm{PSH}(a,b,c,d)/|\mathcal{MD}|}.$$

VI. EXPERIMENTAL VALIDATION

In this section, we validate the proposed performance prediction method by comparing actual PCR plots, as a function of occlusion and clutter rates, with predicted upper bounds.

1) *Model Databases*: We have selected two model databases consisting of military targets from the MSTAR public data domain. The first database $\mathcal{MD}_1$ consists of three targets, T72, BMP2, and BTR70, while the second database $\mathcal{MD}_2$ consists of two targets, T72 and ZSU. Each target is represented by a number of SAR views, which sample its signature at a variety of azimuth angles and a specific depression angle. The two databases are described in Table I. Each model view is processed to extract eight-neighborhood peaks, which are the scattering centers used for recognition. The strongest 30 scattering centers in magnitude are chosen to represent each model view. Fig. 7 shows a sample of three model views (from $\mathcal{MD}_1$) along with extracted scattering centers superimposed on corresponding regions of interest (ROIs).

TABLE I
Description of Model Databases

Database	Depression Angle	Target (Serial No.)	No. of Views	Total No. of Views
$\mathcal{MD}_1$	17°	T72 (#812)	231	
		BMP2 (#c21)	233	
		BTR70 (#c71)	233	
				697
$\mathcal{MD}_2$	30°	T72 (#a64)	288	
		ZSU (#d08)	288	
				576

2) *Test Sets*: We have chosen five sets of test views. The first two sets $\mathcal{TS}_1$ and $\mathcal{TS}_2$ are obtained by introducing controlled distortion to model databases $\mathcal{MD}_1$ and $\mathcal{MD}_2$, respectively. Distortion is introduced in a manner similar to that described in Section VA (see also Fig. 5). Details of this process are as follows.

a) The chosen data distortion parameters are $O \in \{9, 10, \ldots, 18\}$, $C = O$, $R_u(\cdot)$ = four-neighbor region, and R_c = convex hull of the features representing the model view under consideration (note that, in Fig. 7, we show a sample of ROIs, not convex hulls). Notice that since $C = O$, test views have the same number of features as the model ones (30 features).

b) The locations of clutter features are restricted such that no feature lies within the eight-neighbor region of another one. This constraint is introduced to simulate the process of scattering-center extraction.

c) For every O/C value, each model view is distorted four times. Accordingly, for a model database of size N and a specific O/C value, the size of corresponding test subset is $4N$.

The remaining three test sets involve real distortion. The third set $\mathcal{TS}_3$ is a variant of $\mathcal{MD}_1$ obtained under a different depression angle (15°, versus 17° for model views). The fourth set $\mathcal{TS}_4$ is a variant of $\mathcal{MD}_1$ obtained under different configurations (e.g., different number of barrels, missing tow cable) at the same depression angle

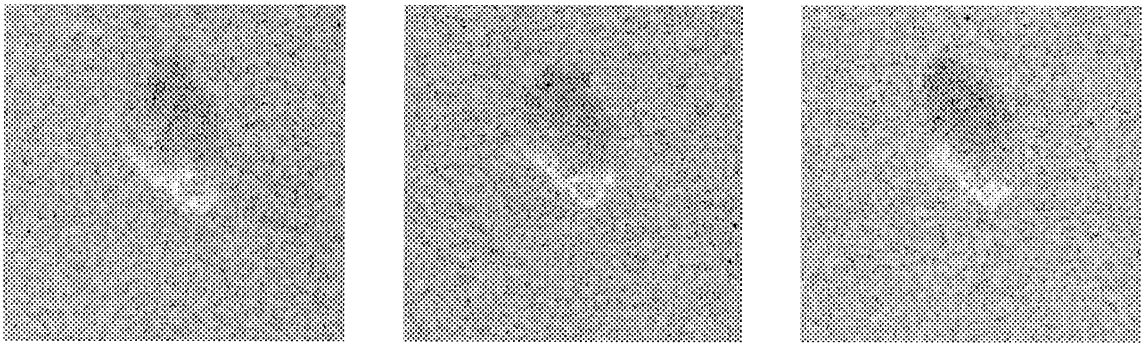

(a) (b) (c)

Fig. 7. Examples of three MSTAR model views at azimuth angle = 132°, and corresponding ROIs with extracted scattering centers superimposed on them. (a) $\mathcal{MD}_1$/T72/132°. (b) $\mathcal{MD}_1$/BMP2/132°. (c) $\mathcal{MD}_1$/BTR70/132°.

TABLE II
Description of Test Sets

Set	Distortion	Reference Database	Depression Angle	Target (Serial No.)	No. of Views	Total No. of Views
$\mathcal{TS}_1$	controlled	$\mathcal{MD}_1$	17°	T72 (#812)	$10 \times 4 \times 231$	
				BMP2 (#c21)	$10 \times 4 \times 233$	
				BTR70 (#c71)	$10 \times 4 \times 233$	
						$10 \times 4 \times 697$
$\mathcal{TS}_2$	controlled	$\mathcal{MD}_2$	30°	T72 (#a64)	$10 \times 4 \times 288$	
				ZSU (#d08)	$10 \times 4 \times 288$	
						$10 \times 4 \times 576$
$\mathcal{TS}_3$	depression-angle change	$\mathcal{MD}_1$	15°	T72 (#812)	193	
				BMP2 (#c21)	194	
				BTR70 (#c71)	194	
						581
$\mathcal{TS}_4$	configuration differences	$\mathcal{MD}_1$	17°	T72 (#132)	231	
				BMP2 (#9563)	233	
						464
$\mathcal{TS}_5$	articulation changes	$\mathcal{MD}_2$	30°	T72 (#a64)	124	
				ZSU (#d08)	118	
						242

(17°). The fifth set $\mathcal{TS}_5$ is a variant of $\mathcal{MD}_2$ obtained under articulation changes (e.g., rotation of a turret or a gun) at the same depression angle (30°). More details about these data sets can be found in [19]. The quantification of scattering-center location invariance with respect to change in depression angles, configuration differences and articulation, and their persistency with azimuth are given in the paper by Bhanu and Jones [3]. Each test view is represented by the strongest 30 scattering centers. We assume that $R_u(\cdot)$ and R_c are the same as in the first two sets. For each test view, the corresponding model-view/location is estimated by finding the best match between the test view and model ones within $\pm 3°$ in azimuth. If no model view exists within such a range, the test view is compared with the model view that is closest in azimuth. The resulting test/model match directly determines the amount of distortion (occlusion and clutter) in the test view. Since test and model views are represented by the same number of features, test-view distortion is described by a single O/C value. Table II summarizes the chosen test sets.

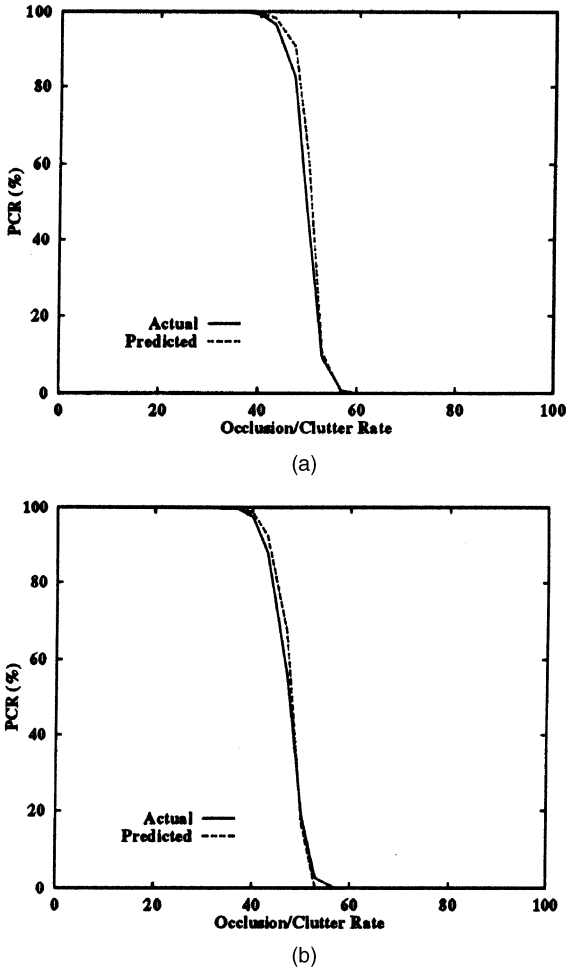

Fig. 8. Actual PCR plots versus predicted upper bounds for controlled-distortion sets. (a) $\mathcal{TS}_1$. (b) $\mathcal{TS}_2$.

3) *Determination of Actual and Predicted PCR Plots*: An uncertainty-accommodating recognition system has been implemented for experimentally determining PCR, as a function of O/C. This system examines *all* the relevant problem space, which is four-dimensional (target, azimuth, and translations along the range and cross-range directions). Accordingly, its recognition performance is optimal, for the given vote-based criterion (defined in (1)). Each test set is passed to the recognition system. Results are accumulated in a histogram whose bins correspond to O/C values. This histogram is directly used to determine the actual PCR plot. The predicted PCR plot (an upper bound) is determined as explained in Section VD. Notice that, in our experiments, the shape of the clutter region is view dependent (the convex hull). Accordingly, the area of the clutter region, AREA(R_c), needed to calculate the conditional pdf of clutter votes (refer to Section VC) is not of a constant value. In our case, AREA(R_c) is substituted by the average of convex hull areas corresponding to views of the appropriate model database.

4) *Results*: Figs. 8(a) and 8(b) show actual and predicted PCR plots for the two test sets involving controlled distortion, $\mathcal{TS}_1$ and $\mathcal{TS}_2$, respectively. From these figures, we observe that our method succeeds in predicting tight upper bounds on ATR performance. In particular, the knees of the actual and predicted plots are very close to each other in both cases. This demonstrates that the proposed method accurately determines the occlusion/clutter rate at which performance starts to degrade. Furthermore, we observe that the actual and predicted plots have the same shape, and are close to each other. This demonstrates the overall accuracy of our method.

Next, we analyze the tightness of the predicted bounds. It can be observed that these bounds become slightly less tight when performance starts to degrade, and then they become very tight towards the end of the plot. This can be explained as follows. Let $\bar{\mathcal{H}}_i \subset \mathcal{H}_i$ be the subset of erroneous hypotheses that reach or exceed the votes for the correct one, in case of a recognition failure. It can be easily seen that the expected size of $\bar{\mathcal{H}}_i$ is directly proportional to the extent of distortion (occlusion/clutter). This implies that the probability that $\bar{\mathcal{H}}_i$ includes no peak hypotheses is inversely proportional to the extent of distortion. Accordingly, when performance starts to degrade (at the knee section of the plot), the bound defined in (8) is a strict upper bound. Its tightness, quantitatively defined as the ratio $(1 - \text{predicted\_bound})/(1 - \text{actual\_PCR})$, gradually increases with distortion, and eventually, the predicted plot almost coincides with the actual one. This tightness/distortion relationship explains the variation in the vertical distance between actual and predicted plots. Notice that this relationship does *not* imply that the vertical distance between predicted and actual performance monotonically decreases with distortion. We note here that, in Fig. 8(b), there is a hardly visible discrepancy between actual and predicted plots towards the bottom section of the plots. This is attributed to the fact that the assumption of statistical independence among peak hypotheses of different views (refer to Section VB) is not entirely accurate. However, as can be observed in both cases, this assumption is clearly very reasonable for use in SAR ATR performance prediction.

It is interesting to compare the performance plots corresponding to sets $\mathcal{MD}_1$ and $\mathcal{MD}_2$ (Figs. 8(a) and 8(b), respectively). It can be observed that the actual PCR for set $\mathcal{MD}_2$, as a function of the O/C rate, is slightly inferior to that for set $\mathcal{MD}_1$, although $\mathcal{MD}_2$ consists of a smaller number of views (refer to Table I). This is due to the fact that the average convex-hull area for $\mathcal{MD}_2$ is smaller than that for $\mathcal{MD}_1$ (482 for $\mathcal{MD}_2$, compared with 532 pixels for $\mathcal{MD}_1$). Since we are considering a fixed number of model features per view, a smaller convex hull leads to: 1) more clutter votes for the erroneous hypotheses corresponding to the views of $\mathcal{MD}_2$ (refer to (13)), and 2) an increase in the similarity between the model

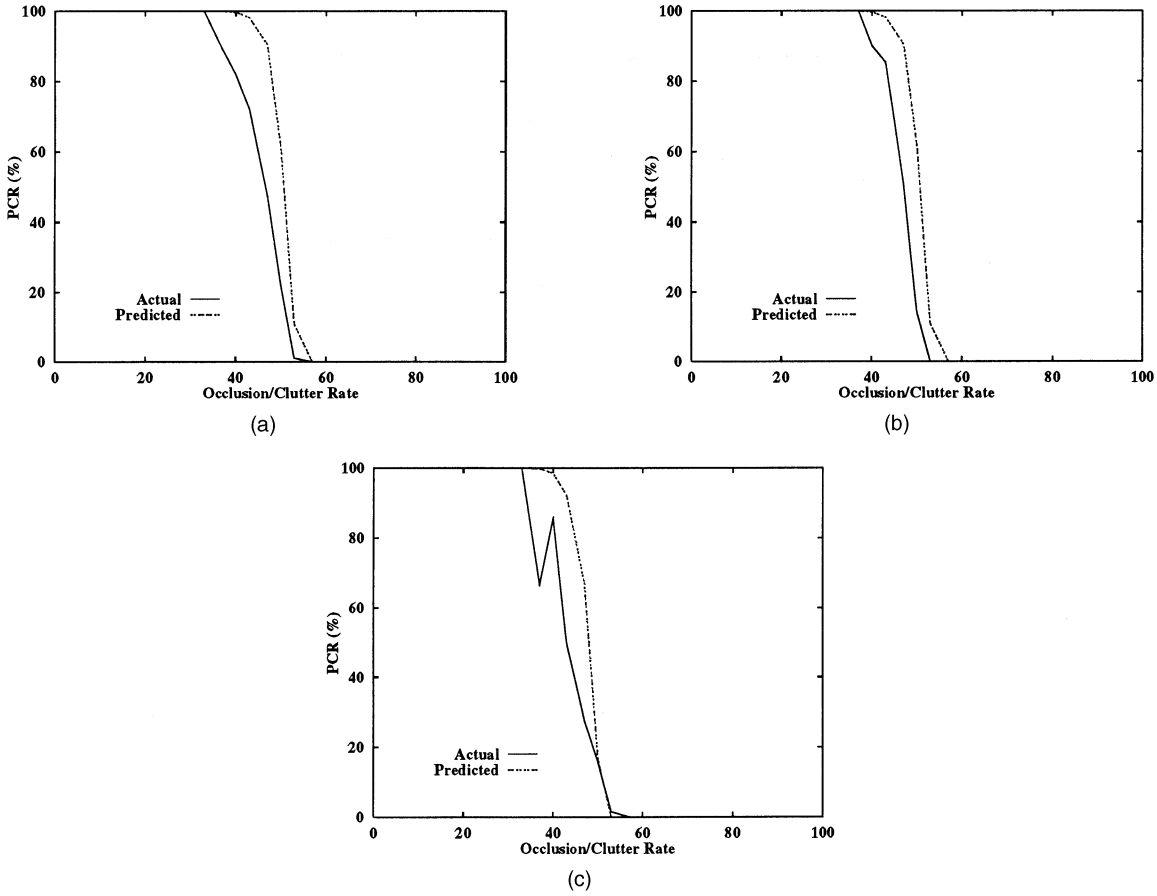

Fig. 9. Actual PCR plots versus predicted upper bounds for real-distortion sets. (a) $\mathcal{TS}_3$ (depression-angle change). (b) $\mathcal{TS}_4$ (configuration differences). (c) $\mathcal{TS}_5$ (articulation changes). (The break in actual plot in (c) is due to insufficient data at corresponding O/C value).

views, and accordingly more similarity votes in case of set $\mathcal{MD}_2$.

The results for the three test sets involving real distortion, $\mathcal{TS}_3$ to $\mathcal{TS}_5$, are shown in Figs. 9(a)–(c), respectively. From these figures, we observe that the proposed method succeeds in determining reasonable upper bounds on performance. We also observe that the bounds are not as tight as in the first case. This is because of differences between the actual distortion models and the assumed ones (i.e., uniform models for uncertainty, occlusion, and clutter). These differences are manifested in actual PCR performance that is slightly more optimistic in the controlled-distortion cases, than the corresponding real-distortion ones (compare the actual plots in Figs. 9(a)–(b) with that in Fig. 8(a), and the actual plot in Fig. 9(c) with that in Fig. 8(b)). Development of methods to learn more accurate distortion models for a given set of data is an interesting subject for future research. An initial work in this area has been presented in [7].

Finally, we note that the results presented in this section demonstrate robustness properties of the vote-based SAR ATR approach considered in this work. In all test cases, performance remains virtually perfect up to more than 30% occlusion/clutter rate. This demonstrates the tolerance of the vote-based approach to considerable levels of data distortion. Further performance improvements, in terms of tolerance to higher levels of distortion, can be obtained by incorporating additional scattering-center attributes, such as peak magnitude (e.g., [3]).

VII. CONCLUSIONS

A novel statistical method has been presented for predicting a tight upper bound on performance of vote-based SAR ATR. The method is based on a quantitative measure of structural similarity between model target views. It considers data distortion factors such as uncertainty, occlusion, and clutter. The validity of the method has been demonstrated by comparing experimentally determined PCR plots with predicted upper bounds in the presence of both controlled and real distortion. Test sets involving real distortion are obtained at different depression angle, configurations, and articulations. Future work involves the development of tighter upper bounds through learning the distortion models, as well as the prediction of other performance metrics such as the confusion matrix and the probability of false alarm.

Finally, we discuss extensions of the proposed method to other related SAR ATR tasks.

1) Consideration of scattering-center magnitudes involves increasing the dimensionality of the feature space from two to three. For example, the uncertainty/consistency region would become a 3-D volume (e.g., an ellipsoid, or a cylinder) in order to consider the uncertainty in magnitude, as well as location. The clutter region would also become a 3-D volume. Other than that, the method remains basically the same.

2) Consideration of Gaussian positional uncertainty, instead of uniform uncertainty, is a more complicated process, although the method remains conceptually the same. It mainly involves the following. 1) The voting criterion would become a weighted criterion that depends on both the distance between corresponding data and model features, and the standard deviation of the Gaussian pdf describing the positional uncertainty. 2) The similarity between a pair of model views would be approximated by the number of feature pairs that are "reasonably" close to each other, as well as the average distance between them. 3) The votes for a distorted version of a model view would become a random variable, instead of a constant in our case (refer to Section VB). 4) The distributions of the similarity and clutter (weighted) votes for an erroneous hypothesis would need to be modified to reflect the weighted-voting criterion (refer to Section VC).

3) The prediction theory developed here can be combined with adaptive target recognition algorithms to optimize the performance of complete ATR systems, with respect to desired PCR and probability of false alarm, in practical situations [5].

REFERENCES

[1] Bhanu, B. (1986)
Automatic target recognition: State of the art survey.
IEEE Transactions on Aerospace and Electronic Systems, **22**, 4 (1986), 364–379.

[2] Bhanu, B., Dudgeon, D. E., Zelnio, E. G., Rosenfeld, A., Casasent, D., and Reed, I. S. (1997)
Introduction to the special issue on automatic target detection and recognition.
IEEE Transactions on Image Processing, **6**, 1 (1997), 1–6.

[3] Bhanu, B., and Jones, G., III (2000)
Recognizing target variants and articulations in synthetic aperture radar images.
Optical Engineering, **39**, 3 (2000), 712–723.

[4] Bhanu, B., and Jones, T. (1993)
Image understanding research for automatic target recognition.
IEEE Aerospace and Electronics Systems Magazine, **8**, 10 (1993), 15–23.

[5] Bhanu, B., Lin, Y., Jones, G., and Peng, J. (2000)
Adaptive target recognition.
Machine Vision and Applications, **11**, 6 (2000), 289–299.

[6] Boshra, M., and Bhanu, B. (2000)
Predicting performance of object recognition.
IEEE Transactions on Pattern Analysis and Machine Intelligence, **22**, 9 (2000), 956–969.

[7] Boshra, M., and Bhanu, B. (2000)
Validation of SAR ATR performance prediction using learned distortion models.
In *Proceedings of SPIE Conference on Algorithms for Synthetic Aperture Radar Imagery VII*, 4053, Orlando, FL, 2000.

[8] Catlin, A., Myers, L., Bauer, K., Rogers, S., and Broussard, R. (1997)
Performance modeling for automatic target recognition systems.
In *Proceedings of SPIE Conference on Algorithms for Synthetic Aperture Radar Imagery IV*, 3070, Orlando, FL, 1997, 185–193.

[9] Chiang, H., Moses, R. L., and Potter, L. C. (2000)
Classification performance prediction using parametric scattering feature models.
In *Proceedings of SPIE Conference on Algorithms for Synthetic Aperture Radar Imagery VII*, 4053, Orlando, FL, 2000.

[10] Diehl, V., Benedict-Hall, G., and Heydemann, C. (1998)
Upper bound calculations of ATR performance for Ladar sensors.
In *Proceedings of SPIE Conference on Algorithms for Synthetic Aperture Radar Imagery V*, 3370, Orlando, FL, 1998, 602–613.

[11] Ettinger, G. J., Klanderman, G. A., Wells, W. M., and Grimson, W. E. L. (1996)
A probabilistic optimization approach to SAR feature matching.
In *Proceedings of SPIE Conference on Algorithms for Synthetic Aperture Radar Imagery III*, 2757, 1996, 318–329.

[12] Grenander, U., Miller, M. I., and Srivastava, A. (1998)
Hilbert-Schmidt lower bounds for estimators on matrix lie groups for ATR.
IEEE Transactions on Pattern Analysis and Machine Intelligence, **20**, 8 (1998), 790–802.

[13] Horowitz, L. L., and Brendel, G. F. (1997)
Fundamental SAR ATR performance predictions for design tradeoffs.
In *Proceedings of SPIE Conference on Algorithms for Synthetic Aperture Radar Imagery IV*, 3070, Orlando, FL, 1997, 267–284.

[14] Jones, G., III, and Bhanu, B. (1999)
Recognition of articulated and occluded objects.
IEEE Transactions on Pattern Analysis and Machine Intelligence, **21** 7 (1999), 603–613.

[15] Irving, W. W., Washburn, R. B., and Grimson, W. E. L. (1997)
Bounding performance of peak-based target detectors.
In *Proceedings of SPIE Conference on Algorithms for Synthetic Aperture Radar Imagery IV*, 3070, 1997, 245–257.

[16] Irving, W. W., Wissinger, J., and Ettinger, G. (1997)
Analysis of classification performance with peak-location feature.
Technical report, Alphatech, Nov. 1997.

[17] Novak, L. M., Owirka, G. J., and Netishen, C. M. (1993)
Performance of a high-resolution polarimetric SAR automatic target recognition system.
Lincoln Laboratory Journal, **6**, 1 (1993), 11–24.

[18] Ratches, J. A., Walters, C. P., Buser, R. G., and Guenther, B. D. (1997)
Aided and automatic target recognition based upon sensory inputs from image forming systems.
IEEE Transactions on Pattern Analysis and Machine Intelligence, **19**, 9 (1997), 1004–1019.

[19] Ross, T., Worrell, S., Velten, V., Mossing, J., and Bryant, M. (1998)
Standard SAR ATR evaluation experiments using the MSTAR public release data set.
In *Proceedings of SPIE Conference on Algorithms for Synthetic Aperture Radar Imagery V*, 3370, Orlando, FL, 1998, 566–573.

[20] Wissinger, J., Washburn, R., Morgan, D., Chong, C., Friedland, N., Nowicki, A., and Fung, R. (1996)
Search algorithms for model-based SAR ATR.
In *Proceedings of SPIE Conference on Algorithms for Synthetic Aperture Radar Imagery III*, 2757, 1996, 279–293.

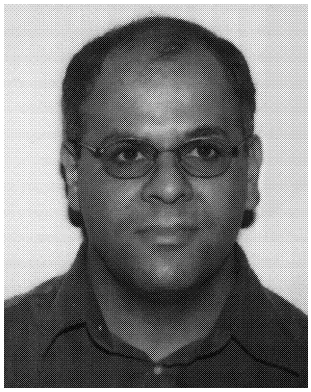

Michael Boshra (S'95—M'97) received the B.Sc. and M.Sc. degrees in computer science from the University of Alexandria, Egypt, in 1988 and 1992, respectively, and the Ph.D. degree in computing science from the University of Alberta, Canada, in 1997. He is currently a technical staff member at AuthenTec Inc., in Melbourne, FL, working on the development of algorithms for fingerprint identification. Prior to joining AuthenTec, he was a post-doctoral fellow in the Center for Research in Intelligent Systems (CRIS) at the University of California, Riverside from 1997 to 1999. From 1989 to 1992, he worked as a research assistant at the National Research Center, Giza, Egypt. His current research interests include fingerprint identification, object recognition, performance prediction, and sensor-data integration. He is a member of ACM.

Bir Bhanu (S'72—M'82—SM'87—F'96) received the S.M. and E.E. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, the Ph.D. degree in electrical engineering from the Image Processing Institute, University of Southern California, Los Angeles, and the M.B.A. degree from the University of California, Irvine. He also received the B.S. degree (with Honors) in electronics engineering from the Institute of Technology, BHU, Varanasi, India, and the M.E. degree (with Distinction) in electronics engineering from Birla Institute of Technology and Science, Pilani, India.

Currently Dr. Bhanu is the Director of the Center for Research in Intelligent Systems (CRIS) at the University of California, Riverside where he has been a Professor and Director of Visualization and Intelligent Systems Laboratory (VISLab) since 1991. Previously, he was a Senior Honeywell Fellow at Honeywell Inc. in Minneapolis, MN. He has been on the faculty of the Department of Computer Science at the University of Utah, Salt Lake City, UT, and has worked at Ford Aerospace and Communications Corporation, CA, INRIA-France and IBM San Jose Research Laboratory, CA. He has been the principal investigator of various programs for DARPA, NASA, NSF, AFOSR, ARO and other agencies and industries in the areas of learning and vision, image understanding, pattern recognition, target recognition, navigation, image databases, and machine vision applications. He is the co-author of books on Computational Learning for Adaptive Computer Vision (Forthcoming), Genetic Learning for Adaptive Image Segmentation (Kluwer 1994), and Qualitative Motion Understanding (Kluwer 1992). He has received two outstanding paper awards from the Pattern Recognition Society and has received industrial awards for technical excellence, outstanding contributions and team efforts. He has been the guest editor of several IEEE transactions and other journals, and on the editorial board of various journals. He holds 10 U.S. and international patents and over 200 reviewed technical publications in the areas of his interest. He has been General Chair for IEEE Workshops on Applications of Computer Vision. Chair for the DARPA Image Understanding Workshop, General Chair for the IEEE Conference on Computer Vision and Pattern Recognition, and Program Chair for the IEEE Workshops on Computer Vision Beyond the Visible Spectrum.

Dr. Bhanu is a Fellow of AAAS and IAPR.