GRAPHICAL MODELS AND IMAGE PROCESSING
Vol. 58, No. 3, May, pp. 233–245, 1996
ARTICLE NO. 0020

Tomographic Reconstruction of Polygons from Knot Location and Chord Length Measurements[1]

LORI BELCASTRO

Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, Massachusetts 02139

WILLIAM C. KARL

Department of Electrical, Computer, and Systems Engineering, Boston University, 44 Cummington Street, Boston, Massachusetts 02215

AND

ALAN S. WILLSKY

Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, Massachusetts 02139

Received January 3, 1995; revised January 16, 1996; accepted January 30, 1996

In this work, we develop statistically based algorithms to reconstruct binary polygonal objects from sparse and noisy tomographic-based observation data. Traditional approaches to the reconstruction of geometric objects from projection data often lead to highly nonlinear estimation problems. To avoid the difficulties associated with such nonlinear problems, we first examine the problem of reconstruction of an object based on knot location measurements, i.e., measurements of the locations of abrupt change in the projections. The ties between this problem and that of multitarget radar tracking enable us to develop a sequential hypothesis-testing algorithm requiring only the solution of a series of linear estimation problems. In particular, data association hypotheses are generated, under each of which the inversion is linear. The complexity of the association possibilities are kept in check through the use of constraints on the reconstruction imposed by the tomography problem. The solution of this first problem is then used as an initialization to a more complete reconstruction which, while utilizing all the projection data, is nonlinear. We demonstrate that the estimates provided by the first, efficient algorithm are of good quality on their own, and, when combined with a fully nonlinear inversion, produce excellent object estimates. © 1996 Academic Press, Inc.

1. INTRODUCTION

In tomographic imaging, many algorithms have been developed to reconstruct a multidimensional function from its projections. Widely used methods such as filtered back-projection and Fourier methods [1, 2] are used to reconstruct high resolution images in a variety of applications [1]. Although these methods produce high quality images, they require a large number of projections and a relatively high signal-to-noise ratio (SNR). In applications where number, spacing, or SNR of the projections are severely constrained a high resolution image is virtually impossible.

In this paper, we develop algorithms that utilize *a priori* knowledge about the nature of underlying object to extract geometric information from sparse and noisy projection measurements. This approach follows from work in geometric-based algorithms [3–8] which develop parametric reconstruction algorithms to obtain geometric information from data. The primary objective of these approaches is to achieve *focusing* of the information in sparse and noisy data to directly determine the shape of the underlying object, rather than reconstructing all of the pixels of the associated image. In this spirit, we assume that the underlying object is a binary polygon and estimate the $2N_v$ parameters that define the N_v vertices of the polygon. Thus, the sparse data are focused on the parameters that are instrumental to the reconstruction of the object.

For a binary object, the projection data are a collection of projected thickness or chord length measurements of the object. The reconstruction of the vertices of a binary

[1] This work was supported by ARO under Contract DAAL03-92-G-0115, ARPA Grant F49620-93-1-0604, and an Office of Naval Research Fellowship.

1077-3169/96 $18.00

polygon from such chord length measurements is a nonlinear estimation problem [7], and thus subject to the same difficulties of all such nonlinear problems (e.g., extraneous local minima). To avoid these difficulties, we will first develop vertex estimates based on a second, related set of measurements which may be directly available or can be derived from the chord measurements. These measurements are of the *locations of the projected vertices* of the object in the projection data, which we term "knot-location" measurements, since they correspond to knots in the projection function. Such observations have been the subject of interest in their own right [9]. The use of such data allows us to exploit algorithms developed for multitarget tracking together with the geometric constraints to obtain an efficient estimation algorithm requiring only the solution of a sequence of *linear* problems. The results of this first, efficient, estimation algorithm are then used as the basis for solution of the full nonlinear problem of estimating the underlying polygon based on *both* knot and chord length measurements (including the case of chord length measurements alone).

This paper is organized as follows. First, the tomographic problem considered in this paper and in particular the role that knot locations play in it are described. Then this knot-based tomographic reconstruction problem is related to the multitarget tracking problem in radar. Using some of the approaches generally used in a radar context, an estimation algorithm to reconstruct the vertices of a binary polygon based upon knot location measurements is developed. Next, both chord and knot location measurements are incorporated into a full nonlinear estimation algorithm that uses our efficient knot-based reconstruction algorithm to provide the initialization. Reconstruction results for both algorithms are presented and the effects of noise, chord length measurements, and limited angle measurements are analyzed. Finally, conclusions and future work are summarized.

2. BACKGROUND

2.1. The Radon Transform

The goal of tomographic imaging is the reconstruction of a two-dimensional (2D) image, $f(x, y)$, from the complete set of 1D projections, $g(t, \theta)$,

$$g(t, \theta) = \int_{R^2} f(x, y)\, \delta(t - [x, y] \cdot e(\theta))\, dx\, dy, \qquad (1)$$

where $e(\theta) = [\cos(\theta), \sin(\theta)]^{\mathrm{T}}$ is a unit vector in direction θ. The function $g(t, \theta)$ is termed the 2D Radon transform of f and defines the projection of f along any line $L(t_i, \theta_j)$ at angle θ_j and displacement t_i, as shown in Fig. 1. For convenience, term the collection of all such projections at a given angle θ_j as the projection at that angle. The *support values* at θ_j are the two extrema of t that define the set of non-zero values of $g(t, \theta_j)$ (e.g., z_1 and z_4 in Fig. 1). Note that projections taken 180° apart are reflections about t so data is only required over a 180° angular range. Table 1 gives a list of symbols used throughout the paper.

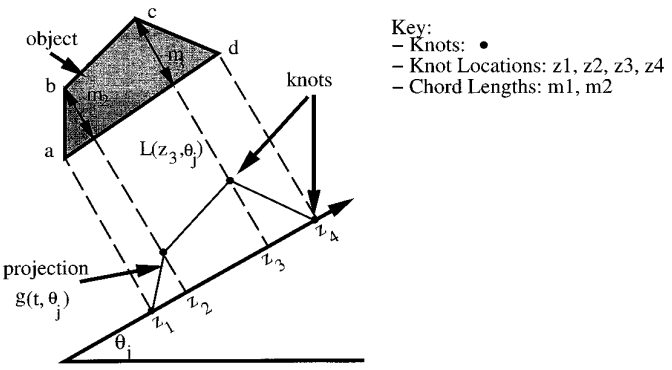

FIG. 1. Relationship of object, projection, knots, and chords.

2.2. Binary Polygonal Objects

The focus for the remainder of this paper will be on binary polygonal objects. The projection of a binary polygonal object is a piecewise-linear spline function, as shown

TABLE 1
List of Symbols

Symbol	Description
$f(x, y)$	2D object to be reconstructed
$g(t, \theta)$	2D Radon transform of $f(x, y)$
N_v	Number of vertices in the object
N_t	Number of chord length measurements in each view
N_θ	Number of angular views of projection data
N_q	Number of hypotheses retained at each stage of the Knot-Based Reconstruction Algorithm
m_{θ_j}	Vector of observed chord length measurements in view θ_j
M^{N_θ}	Overall vector of observed chord length measurements in all N_θ views
z_{θ_j}	Vector of observed knot location measurements in view θ_j
Z^{N_θ}	Overall vector of observed knot location measurements in all N_θ views
X^{N_θ}	Overall vector of observed knot and chord measurements for all θ_j
Λ_j	Noise covariance matrix for vector noise process j
V	Vertices of the underlying object $f(x, y)$
$\hat{V}$	Estimate of the vertices of the underlying object $f(x, y)$
I_j	jXj identity matrix
q_i^k	ith data association hypothesis for the first k views
q^k	Set of all possible data association hypothesis for the first k views

in Fig. 1 for the case of a 4-gon. In the projection, each knot location, or position of abrupt change in the slope of the spline function $g(t, \theta_j)$ [10], corresponds to the location of the projection of one of the vertices of the object. In each projection of the polygon the number of knots is equal to the number of vertices N_v of the object (except in degenerate cases). The chord length of a binary object is the thickness of the object along the line $L(t_i, \theta_j)$, thus the magnitude of the projection $g(t_i, \theta_j)$ at a particular t_i is precisely this length. Therefore, the projection of a binary object is simply a collection of chord lengths of the object. This relationship between chord length, knot location, and the underlying vertices provides the basis of the algorithms described in this paper.

2.3. Chord Length Measurements

We represent the vertices (x_i, y_i) of the underlying object (and thus the polygonal object itself) by the following length $2N_v$ vector V containing their coordinates:

$$V = [x_1\, y_1 \,|\, \cdots \,|\, x_{N_v}\, y_{N_v}]^T \tag{2}$$

The vector of chord length measurements, $m(\theta_j)$ are then modeled by

$$m(\theta_j) = f(V, \theta_j) + v(\theta_j), \quad v(\theta_j) \sim N(0, \Lambda_v), \tag{3}$$

where $m(\theta_j) = [m(t_1, \theta_j), m(t_2, \theta_j), \ldots, m(t_{N_t}, \theta_j)]^T$ with $t_1 \le t_2 \le \cdots \le t_{N_t}$ is a vector of the N_t ordered chord length measurements at angle θ_j, $f(V, \theta_j)$ is a nonlinear vector function relating the ordered, noise-free chord measurements at a particular angle θ_j to the vertices of the object and $v(\theta_j)$ is a vector Gaussian noise process independent from angle to angle and shift to shift. For this work we will assume that $\Lambda_v = \sigma_v^2 I_{N_t}$, where I_N is an $N \times N$ identity matrix. Each of the N_θ views can now be combined (stacked) to form a single overall measurement equation,

$$\begin{aligned} M^{(N_\theta)} &= [m^T(\theta_1) \,|\, \cdots \,|\, m^T(\theta_{N_\theta})]^T \\ &= F(V) + \Upsilon, \quad \Upsilon \sim N(0, \Lambda_\Upsilon), \end{aligned} \tag{4}$$

where $F(V)$, Υ, and the covariance Λ_Υ are defined in the obvious way.

Because the function $F(V)$ is a nonlinear function of the vertices, reconstructing V based on (4) is difficult. Instead of immediately attempting solution of this nonlinear problem, we will first develop methods based on knot location measurements because they can be extracted directly from the projection data [9, 10], are geometrically related to the parameters of interest (the vertices), and will allow us to develop estimates based on solving a series of *linear* estimation problems.

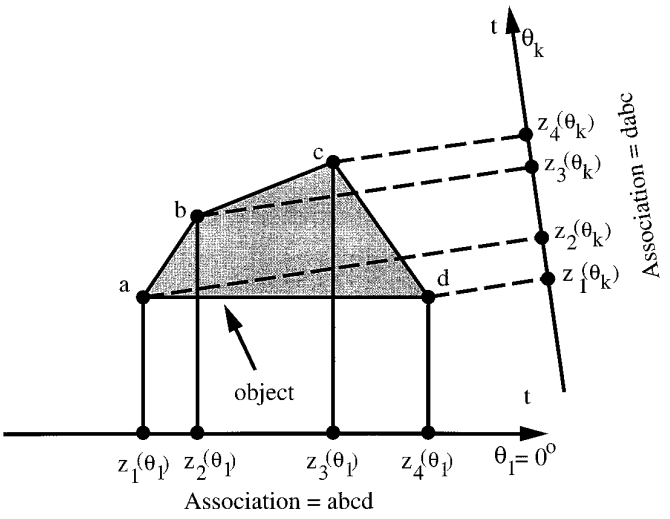

FIG. 2. Initial knot association and notation.

2.4. Knot Location Measurements

The knot locations for the projection at angle θ_j can be arbitrarily, and without loss of generality, labeled in increasing order as $z_1(\theta_j), \ldots, z_{N_v}(\theta_j)$, as shown in Fig. 2. Each knot location corresponds to a particular vertex in the object (e.g., in Fig. 2 $z_1(\theta_1)$ corresponds to vertex a), but without knowledge of the underlying object, the correspondence of the set of knot locations $z_1(\theta_j), z_2(\theta_j), z_3(\theta_j), \ldots$ to the underlying vertices $a, b, c, \ldots$ of the object is unknown. In particular, the function $h(V, \theta_j)$ relating the vertices to the noise-free knot locations at angle θ_j is given by

$$h(V, \theta_j) = \text{sort}\,(E(\theta_j)\, V) \tag{5}$$

where sort$(\cdot)$ sorts the argument in ascending order, $E(\theta_j) = I_{N_v} \otimes e^T(\theta_j)$ with $\otimes$ representing the Kronecker product,[2] and $e(\theta) = [\cos(\theta), \sin(\theta)]^T$ a unit vector in direction θ. From the definition of $E(\theta)$, we see that $E(\theta)V$ is simply the projection of the vertices at angle θ. The observations are a nonlinear function of the vertex locations due to the sort operation, which destroys the associations of the knot locations to the object vertices.

Based on this discussion, we will use the following model for the noisy knot location observations at each view θ_j:

$$z(\theta_j) = h(V, \theta_j) + n(\theta_j), \quad n(\theta_j) \sim N(0, \Lambda_{n(\theta_j)}). \tag{6}$$

Here $z(\theta_j) = [z_1(\theta_j), z_2(\theta_j), \ldots, z_{N_v}(\theta_j)]^T$ with $z_1(\theta_j) \le z_2(\theta_j) \le \cdots \le z_{N_v}(\theta_j)$ is the vector of observed, ordered knot location measurements at angle θ_j, $h(V, \theta_j)$ is the

[2] $A \otimes B = [A_{ij}B]$.

nonlinear vector function given in (5), and $n(\theta_j)$ is a vector Gaussian noise process, independent from angle to angle and knot to knot. We assume that $\Lambda_{n(\theta_j)} = \text{diag}[\sigma^2_{n_i}(\theta_j)]$. Thus while the noise is assumed independent, the variance $\sigma^2_{n_i}(\theta_j)$ on each knot location measurement $z_i(\theta_j)$ is thus allowed to be different to reflect the belief that knot observations extracted from projections using methods such as [10] will, in general, possess differing degrees of uncertainty. If the measurements at each view are combined (stacked), the following overall knot measurement equation is obtained:

$$\begin{aligned} Z^{(N_\theta)} &= [z^{\mathrm{T}}(\theta_1) \mid z^{\mathrm{T}}(\theta_2) \mid \cdots \mid z^{\mathrm{T}}(\theta_{N_\theta})]^{\mathrm{T}} \\ &= H(V) + \mathcal{N}, \quad \mathcal{N} \sim N(0, \Lambda_{\mathcal{N}}). \end{aligned} \tag{7}$$

Here $H(V)$, $\mathcal{N}$, and the covariance $\Lambda_{\mathcal{N}}$ are defined in the obvious way. Note that if we were given the association of the knot location data to the object vertices in each view, the relationship between the knot observations $Z^{(N_\theta)}$ and the unknown vertices V would be linear. The challenge arises because we do not know these data associations *a priori*. To estimate the vertices of the object from noisy measurements of these knot locations we thus approach the problem through a two-step procedure by first estimating the association of knot locations to the vertices of the object in each projection through a hypothesis testing method and then solving the resulting linear estimation problem.

2.5. Combined Chord and Knot Measurements

The two observation equations (4) and (7) can be combined to form an overall measurement equation,

$$X^{(N_\theta)} = \begin{bmatrix} Z^{(N_\theta)} \\ M^{(N_\theta)} \end{bmatrix} = \begin{bmatrix} H(V) \\ F(V) \end{bmatrix} + \begin{bmatrix} \mathcal{N} \\ \Upsilon \end{bmatrix}, \quad \begin{bmatrix} \mathcal{N} \\ \Upsilon \end{bmatrix} \sim N(0, \Lambda), \tag{8}$$

where Λ captures the joint covariance of the noises. We make the simplifying assumption that the noise from the chord length and knot location measurements are *uncorrelated* with each other. In practical situations where a knot extraction process might derive knot location measurements from the original projection data, there would undoubtedly actually be a correlation between the knot location uncertainty and the chord length uncertainty. As we will show in Section 4.3, however, this simplifying assumption appears quite good and does not have a significant impact on reconstruction results. Finally, note that $\Lambda_{\mathcal{N}}$, Λ_Υ, and hence Λ are diagonal matrices because of our assumptions of noise independence.

3. RECONSTRUCTION ALGORITHMS

Given our observation data, as defined by (7) or (8), we seek estimates of the polygonal object (i.e., its vertices) based on maximum likelihood (ML) estimation techniques. In particular, we seek

$$\hat{V} = \arg\max_V p_{y|V}(Y \mid V), \tag{9}$$

where $p_{y|V}(Y \mid V)$ denotes the conditional probability density of the observations Y given the vector V of vertices of the object and $Y = Z$ or $Y = X$, depending on our focus, as described next. In addition to assuming that the object is a binary polygon, we will also assume that the object is convex with a known number of vertices N_v. While convexity is not essential, it is convenient in determining a unique connection of the vertices to form the underlying object. In the absence of convexity, the chord length measurements could be used to effectively disambiguate among possible vertex connections (e.g., by evaluating all possible connections of the vertices).

3.1. Knot-Based Reconstruction Algorithm

Vertex Reconstruction and Multitarget Tracking

The reconstruction of the vertices V of a binary polygonal object from knot location measurements Z is very closely linked to the multitarget tracking problem of radar. In particular, note that the data of the tomography problem can be viewed as arising from measurements of a rotating object taken at a single, fixed location. By viewing the problem in this way, we can interpret the vertex locations of the polygon as targets and the corresponding knot location measurements as radar reports. In this framework, the association of knot location observations to object vertices is analogous to the association of radar reports to targets.

The similarity of these problems allows estimation techniques developed in the radar context to be applied here. A number of these algorithms are based on an adapted form of a hybrid state estimation problem [11–14], which propose a simultaneous solution to a discrete-state estimation problem (i.e., data association) and a continuous-state estimation problem (i.e., target location estimation). There is also an important difference between the tomography and the radar problem that we exploit to our advantage. In contrast to the radar problem, "target" maneuverability and dynamics are constrained by the fact that the targets or vertices define a rigid object. These constraints can be used to drastically limit the number of possible knot-to-vertex correspondences and thus reduce the number of data association hypotheses that must be considered.

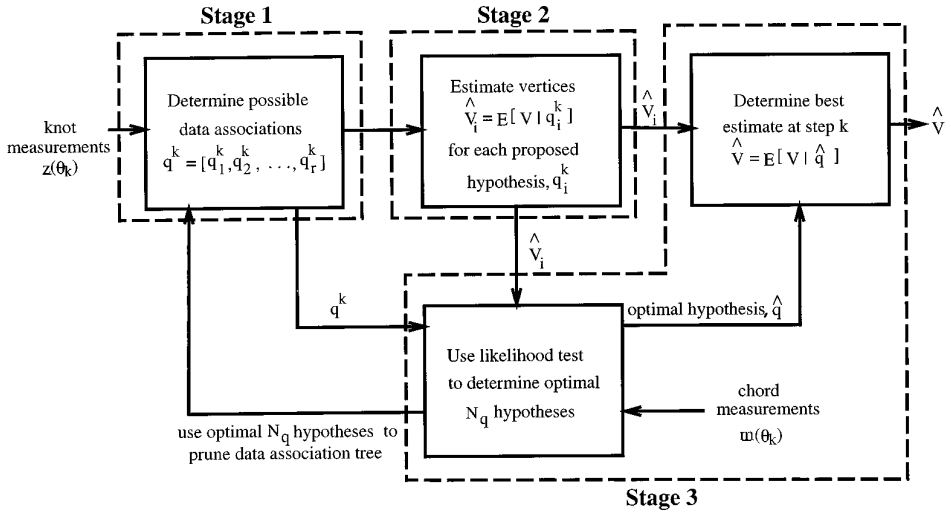

FIG. 3. Block diagram of KBA.

Overview of the Knot-Based Reconstruction Algorithm

A block diagram of the algorithm developed in this section, which we term the knot-based reconstruction algorithm (KBA), is depicted in Fig. 3. Basically, the algorithm is a three stage process, based on the hybrid-state estimation algorithm used in multitarget tracking. In the first stage, the possible associations of the current knot location measurements to the object vertices are enumerated and merged with the existing set of such hypotheses carried forward from the previous steps. The second stage then generates ML estimates of the vertex location for each association hypothesis under the assumption that it is correct; a linear problem. In the third stage the likelihood of each of the data association hypothesis is evaluated given the current set of measurements of both the knot locations *and* chord lengths. The chord lengths are not required to rank the hypotheses, but because they provide additional information, are available anyway, and the computational cost of using them in the analysis in this way is minimal, we include them. In particular, note that all that is required is the *evaluation* of a nonlinear function for this purpose, rather than minimization of it. Finally, to reduce complexity, the likelihood ratios are used to *prune* the set of hypotheses prior to the incorporation of the next view. We describe each step in more detail next.

Part I: Data Association. In the first stage of the algorithm we update the current set of data association hypotheses by incorporating the possible associations between the current view of ordered knot location observations and the vertices of the underlying polygon. Without loss of generality, we let the first view correspond to $\theta_1 = 0°$ and associate the ordered knot locations $z_1, z_2, z_3, \ldots$ in this first view to vertices we arbitrarily label as $a, b, c, \ldots$, as shown in Fig. 2 for a 4-gon. In general, such associations will be represented by sequences of letters corresponding to the vertices of the underlying object, e.g., *abcd* for the first view of the example of Fig. 2.

For convenience, we use the *target tree* approach described by Kurien and co-workers [11, 13] to represent the possible data associations in a hypothesis tree. Each level of the tree, arising from a given angular view, contains the possible associations of the ordered knot observations of that view with the object vertices. As one moves along a branch of the tree from the root to a leaf, the data associations are listed in the order that the angular views are introduced into the algorithm (which does not necessarily correspond to increasing angular order). A path from the root to a leaf of the tree thus provides a complete association of all ordered knot observations in all views to the vertices in the object, and thus corresponds to one possible complete data association hypothesis, q_i^k, up to the kth view. The set of all such hypotheses up to the kth measurement is denoted by $q^k = [q_1^k, q_2^k, \ldots]$.

In general, for an object with N_v vertices there are $N_v!$ total possible knot to vertex association hypotheses that must be considered *for each view*. Without further constraint, each of these possibilities must be considered in combination with all the possibilities produced by the previous views and thus a combinatorial explosion occurs as the number of views increases. This number can be greatly reduced, however, by using the following constraints for associations considered as an increasing function of continuous angle θ:

Data Association Rules:

1. *Maximum Number of Associations:* The total angular view range [0, 180°) can be split into at most N_v

$(N_v - 1)/2 + 1$ separate regions, within which the knot-to-vertex association is fixed. Thus there are *at most* N_v $(N_v - 1)/2 + 1$ distinct knot-to-vertex associations.

2. *Uniqueness of Associations:* The association within each distinct region in Rule 1 is unique. Thus identical associations must be contiguous as a function of angle (e.g., if the associations at $\theta = \theta_1$ and $\theta = \theta_2$ are the same, then so are any associations for $\theta \in [\theta_1, \theta_2]$).

3. *Reciprocal Associations:* An association and its reciprocal cannot both appear, except for the association of the initial and final association regions, since $g(t, \theta_j) = g(-t, \theta_j + 180°)$.

These rules are only a subset of the possible rules that can be derived to determine the allowable data associations given the geometry of this problem. They are meant only to provide a simple and efficient means of reducing the total number of discrete data association hypotheses that must be considered. A precise definition of the space of valid data associations is still an open question.

While we assume for simplicity that the number vertices is known, in general this may not be true. In this case, the knot location data may contain what are termed ''missed detections'' or ''false alarms'' which correspond to missing data measurements and extraneous data measurements, respectively. To account for such cases, the algorithm can be extended to include the global hypothesis techniques used in the track-oriented approach of multitarget tracking [11].

Part II: Estimation of the Vertices Given the Data Associations. Given the updated list of data association hypotheses from Part I of the algorithm, this stage of the algorithm generates ML estimates of the vertices for each hypothesis, assuming it is correct. This is a simple linear problem. In particular, under hypothesis q_i^k, (7) can be written as

$$\begin{aligned} Z^{(k)} &= \begin{bmatrix} P_1(q_i^k) & & \bigcirc \\ & \ddots & \\ \bigcirc & & P_k(q_i^k) \end{bmatrix} \begin{bmatrix} E(\theta_1) \\ \hline \vdots \\ \hline E(\theta_k) \end{bmatrix} V + \mathcal{N} \\ &= P(q_i^k)\, E^{(k)}\, V + \mathcal{N}, \quad \mathcal{N} \sim N(0, \Lambda_{\mathcal{N}}), \end{aligned} \tag{10}$$

where $P_j(q_i^k)$ is a permutation matrix for view j which depends on the i-th association hypothesis q_i^k for the k views, $E(\theta_j)$ is defined, cf. (5), and $\mathcal{N}$ and $P(q_i^k)$ and $E^{(k)}$ are aggregate matrices for the k-th step that are defined in the obvious way. Note that $E^{(k)}\, V$ is the vector of the projected vertices at all the k views and $P(q_i^k)$ is the permutation matrix that rearranges this complete set of projected vertices to match the knot associations contained in the hypothesis q_i^k. Given hypothesis q_i^k, this equation is linear with respect to the vertices, so the ML estimate of V is given by

$$\begin{aligned} \hat{V}(Z^{(k)} \mid q_i^k) = (E^{(k)^{\mathrm{T}}} P^{\mathrm{T}}(q_i^k)\, \Lambda_{\mathcal{N}}^{-1}\, P(q_i^k)\, E^{(k)})^{-1} \\ E^{(k)^{\mathrm{T}}} P^{\mathrm{T}}(q_i^k)\, \Lambda_{\mathcal{N}}^{-1}\, Z^{(k)}, \end{aligned} \tag{11}$$

where we have made the dependence of the estimate on the data $Z^{(k)}$ and hypothesis q_i^k explicit. Further, since we assume that the noise on each knot location is independent, $\Lambda_{\mathcal{N}}$ is a diagonal matrix, and the problem separates into N_v separate estimation problems—one for each vertex. In summary, an ML estimate of V is obtained for each of the association hypotheses from Part I.

Part III: Hypotheses Evaluation and Pruning. While the knot-to-vertex association constraints presented earlier serve to greatly reduce the number of possible data associations, there are still a large number of remaining possible hypotheses. Thus, *pruning techniques* developed in the multitarget target literature are also used to limit the number of possible hypotheses. In particular, the quality of each proposed hypothesis is found by evaluating the conditional log-likelihood of the *complete* knot and chord data up to the kth view $X^{(k)}$, conditioned on the current hypothesis q_i^k and assuming the corresponding ML estimate $\hat{V}(Z^{(k)} \mid q_i^k)$ of the vertices based on the knot observations alone is correct,

$$\log[p_{x|q_i^k, \hat{V}}(X^{(k)} \mid q_i^k, \hat{V}(Z^{(k)} \mid q_i^k))], \tag{12}$$

where $\hat{V}(Z^{(k)} \mid q_i^k)$ is defined, cf. (11), and $X^{(k)}$ is complete data vector; cf. (8). The primary departure of this log-likelihood from the standard Generalized Log-Likelihood is that the unknown nonrandom quantity, V, is estimated from a subset $Z^{(k)}$ of the total measurements $X^{(k)}$ (i.e., from the knot location measurements but not from the chord measurements).

The hypotheses q_i^k are now ranked based on the log-likelihoods of (12). The top N_q hypotheses are retained and the remainder are pruned from the tree and thus further consideration. Finally, the top ranked hypothesis $\hat{q}^k$ is taken as the optimal hypothesis at stage k and the associated ML vertex estimate $\hat{V}(Z^{(k)} \mid \hat{q}^k)$ is taken as the optimal estimate. Experimental results of this algorithm follow in Section 4.

3.2. Knot-and-Chord-Based Reconstruction Algorithm

Here we develop a nonlinear algorithm to estimate the vertices of a binary polygonal object from the full set of noisy measurements $X^{(N_\theta)}$ of both knot locations and chord lengths. Like the KBA, the estimate is obtained by using ML estimation techniques. This algorithm, which we term the knot-and-chord-based algorithm (KCBA), minimized the cost criterion of (9) from measurements of both

knot locations *and* chord lengths.[3] This cost is equivalent to the following weighted nonlinear least-squares error (WNLSE) criterion,

$$\begin{aligned}\hat{V}_{ML} &= \arg\min_{V} \left(\begin{bmatrix} Z \\ M \end{bmatrix} - \begin{bmatrix} H(V) \\ F(V) \end{bmatrix}\right)^{\mathrm{T}} \\ &\quad \Lambda^{-1} \left(\begin{bmatrix} Z \\ M \end{bmatrix} - \begin{bmatrix} H(V) \\ F(V) \end{bmatrix}\right) \qquad (13) \\ &= \arg\min_{V}: (Z - H(V))^{\mathrm{T}} \Lambda_{\mathcal{V}}^{-1} (Z - H(V)) \\ &\quad + (M - F(V))^{\mathrm{T}} \Lambda_{\Upsilon}^{-1} (M - F(V)),\end{aligned}$$

where the dependence of the observations on N_θ is suppressed for clarity. Being nonlinear, this problem is plagued by multiple local minima, etc., and must generally be solved iteratively. To ameliorate such difficulties, a good initial guess is required, which is provided by the KBA. In summary, the steps of the algorithm are: (i) Use the KBA to generate an initial guess V_0 from knot measurements and (ii) solve the knot *and* chord based ML problem (13) using any iterative technique starting from this initial value. Note that by setting $\Lambda_{\mathcal{V}}^{-1} = 0$ in (13) we will obtain the ML estimate based on chord length measurements alone.

4. RECONSTRUCTION RESULTS

This section provides three sets of numerical experiments: experiments using the KBA, experiments using the KCBA, and for completeness, experiments with a complete end-to-end system. Although the algorithms are valid for the reconstruction of n-gons, for brevity we limit our reconstruction results to 4-gons. The principal algorithmic impact of including more vertices is that there are more data association hypothesis to consider at each stage of the reconstruction. In addition, the inclusion of many vertices suggests the desire for detailed structural information about the underlying object. Given that our focus is on limited and noisy data, the reconstruction detail implied by such a large number of parameters is not really warranted anyway.

For the first two sets of experiments in Sections 4.1 and 4.2, the knot and chord measurements we use are independently generated from the underlying object and perturbed with additive independent Gaussian noise. In an actual tomographic reconstruction scenario, the noisy projection data (or chord length measurements) might be put through a separate knot extraction algorithm, such as a procedure similar to that in [10], which would produce knot location measurements and their corresponding noise statistics. These data, along with the original chord projection data, would then be used as the input to the reconstruction algorithms developed in this paper. In Section 4.3 we present a final set of experiments showing the results of an end-to-end system designed to capture the spirit of such a scenario.

[3] This algorithm was implemented via the Matlab function FMINS based on the simplex search method of Nelder and Mead.

Performance Measures

In the experiments that follow, the reconstructions are performed by minimizing the L^2 norm on the measurements (i.e., in the Radon or projection space). Ultimately, however, it is the difference between the true underlying shape and our reconstructed shape that is of direct interest in quantifying the quality of the reconstructed object. The measure we use is the Hausdorff distance [15] $\Delta^{\mathrm{H}}(S, \hat{S})$ between two convex sets S and $\hat{S}$, defined as

$$\Delta^{\mathrm{H}}(S, \hat{S}) = \inf\{\varepsilon \mid S \subset \hat{S}^{(\varepsilon)}, \hat{S}^{(\varepsilon)} \subset S\}, \qquad (14)$$

where $S^{(\varepsilon)} = \{S \mid d(s, S) \leq \varepsilon\}$ is the ε-neighborhood of the set S and $d(s, S)$ is the minimum distance between the point s and set S. Thus the Hausdorff distance is a measure of the largest distance by which the two sets differ. In order to make reconstruction evaluations with this measure comparable, we will use a normalized version, we term the percent Hausdorff error and define as

$$\%\ \text{Hausdorff error} = \frac{\Delta^{\mathrm{H}}(\mathcal{S}, \hat{\mathcal{S}})}{\Delta^{\mathrm{H}}(\mathcal{S}, \emptyset)} \times 100, \qquad (15)$$

where $\mathcal{S}$ is the true object, $\hat{\mathcal{S}}$ the estimate, and $\emptyset$ is the single point at the origin.

4.1. Knot-Based Algorithm Experiments

Here we present results of the KBA developed. For these experiments our chord measurements are generated according to (4). The variance σ_m^2 of the actual applied additive noise depended on the experiment and was chosen to yield a given equivalent average signal-to-noise ratio (SNR) of the chord length observations, defined as

$$\text{SNR (dB)} = 10 \log_{10} \frac{\|F(V)\|^2}{N_t N_\theta \sigma_m^2}, \qquad (16)$$

where $F(V)$ are the noise-free chord lengths and $N_t N_\theta$ is the total number of data points. In all our experiments we assume that this variance is known, so that $\sigma_v^2 = \sigma_m^2$.

Our knot observations are generated according to (7). The variance $\sigma_{z_i}^2(\theta_j)$ of the noise actually applied to the noise-free knot i in view j is chosen according to

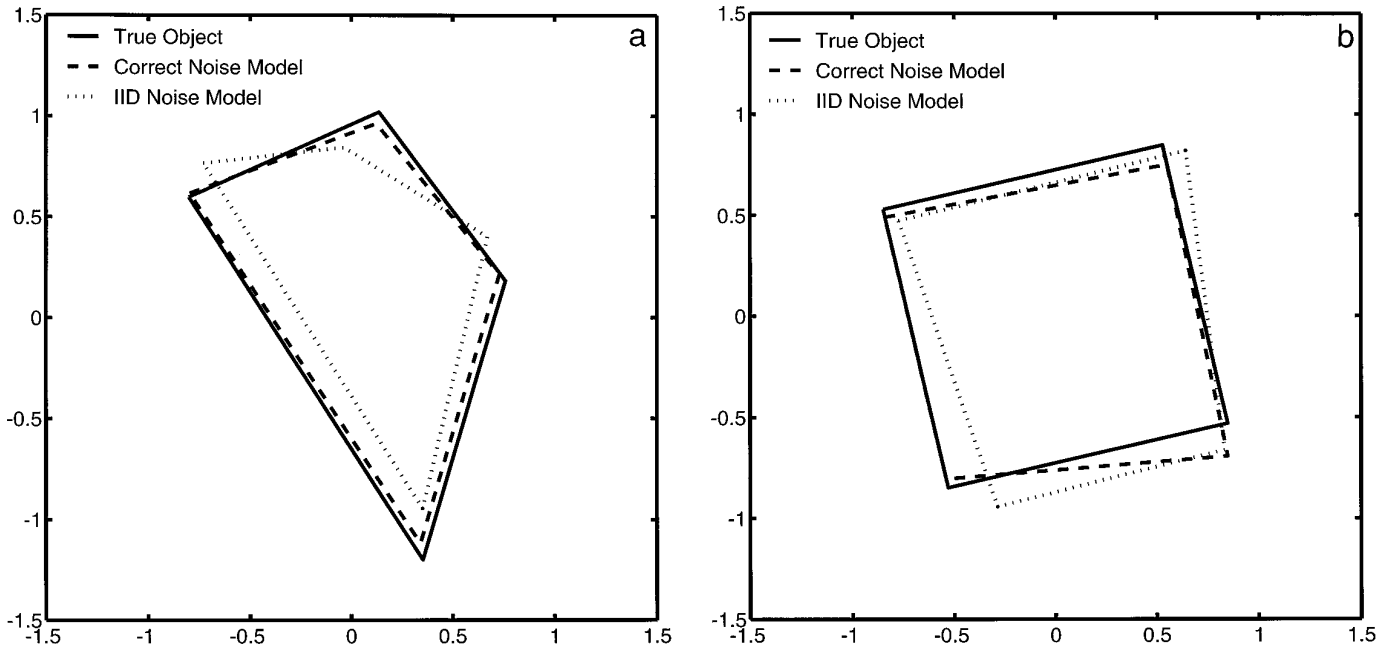

FIG. 4. Sample reconstructions of KBA, $N_\theta = 27$ views, $N_q = 10$ retained hypotheses, 10 dB SNR on chords. (a) Kite reconstruction. (b) Square reconstruction

$$\sigma^2_{z_i}(\theta_j) = \frac{K^2}{|\Delta s_i(\theta_j)|^2 M_i(\theta_j)}\,\sigma^2_m, \tag{17}$$

where σ^2_m is the variance of the noise actually applied to the noise-free chord length measurements from (16), $|\Delta s_i(\theta_j)|$ is the magnitude of the change in slope of the noise free projection at knot i in view j, $M_i(\theta_j)$ is the separation between adjacent knots at location i in view j, and K is a global scaling constant. For interior knots, $M_i(\theta_j)$ is defined as the magnitude of the distance between the knots adjacent to knot i, while for a boundary knot, we use twice the distance to the adjacent interior knot. Thus $M_i(\theta_j)$ is a measure of local knot separation. The value of the scaling constant K, which can be viewed as setting the overall relative strength of knot noise verses chord noise, is somewhat arbitrarily chosen so that the average standard deviation of knot location noise terms is 1.5 times greater than the standard deviation of the chord measurement noise terms. This choice of applied noise variance for the knot observations attempts to capture the most important features we would expect from a knot extraction routine. In a complete implementation this noise model would follow from the particular knot extraction algorithm. Our goal here is to focus on the use of knot location and chord length measurements to reconstruct the underlying object and the behavior of this part of the process, so we have intentionally simplified our knot noise model.

In the examples of this section we perform experiments where we assume that this actual applied noise variance is known, so that $\sigma^2_{n_i}(\theta_j) = \sigma^2_{z_i}(\theta_j)$ and the model variance in (6) is correct. We also consider a simplified case where the model variance is set to $\sigma^2_{n_i}(\theta_j) = (1.5\ \sigma_m)^2$, so the model assumes an identically distributed noise on the knots. This more simplistic model, which we refer to as the IID model, assumes no information is available from the knot extraction algorithm about the relative goodness of one knot measurement compared to another, so it offers a form of worst case scenario for the algorithms.

Sample Reconstructions

For all of the reconstructions shown, the true object is depicted by solid lines while the reconstructions are represented by dashed or dotted lines. Two objects are considered, a square and a ''kite'' with noisy chord length and knot location data generated so that the SNR on the chord length data is 10dB. Projection data were taken at 27 equally spaced angles over the interval [0°, 180°) with five chord measurements in each view uniformly spaced over $t = [-1, 1]$. Note that this results in views where some of the chord measurements are zero (i.e., the object is not in the field of view) and conversely, in views where the chord measurements are confined to the interior of the object (for the kite object). Thus, each object had a total of 108 knot location measurements (used to estimate the vertices and prune the hypotheses) and 135 chord measurements (used to prune the hypotheses). Finally, only $N_q = 10$ hypotheses are retained at each stage of the algorithm.

Figure 4 shows sample reconstructions of two objects.

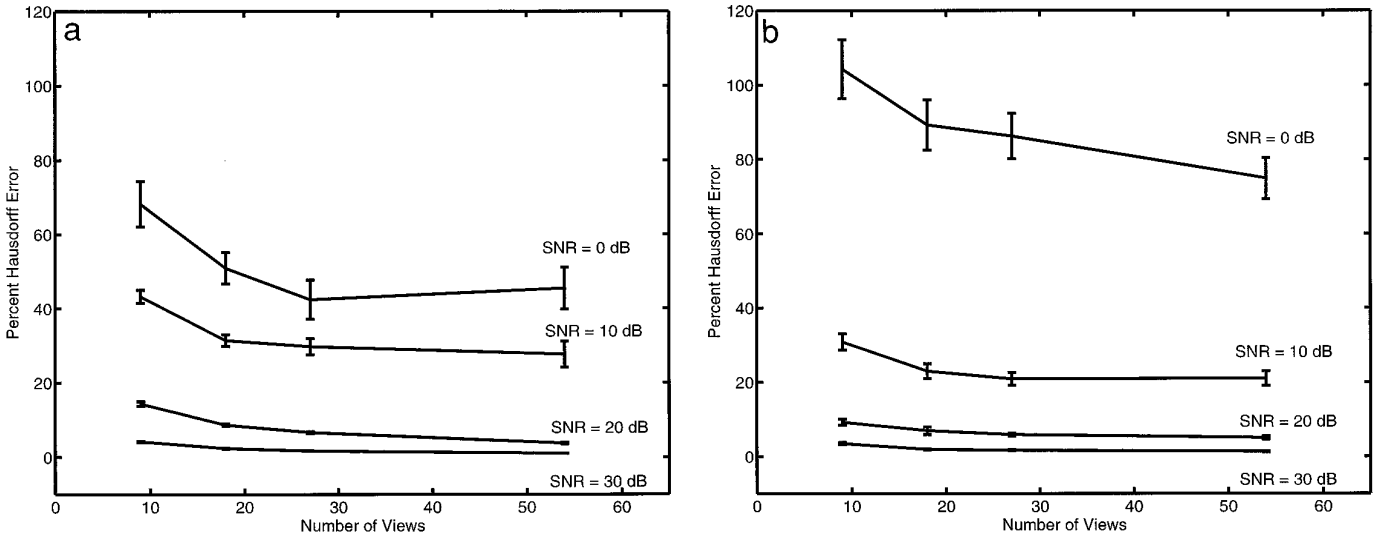

FIG. 5. Performance as a function of number of views used in the KBA. (a) Hausdorff error, correct model. (b) Hausdorff error, IID model.

In each case two reconstructions are shown: the dashed line represents the reconstruction generated assuming the correct applied noise model on the knot measurements while the dotted line represents the reconstruction assuming the IID knot-noise model. Figure 4a shows the kite object. The Hausdorff percent errors are 6.5 and 20.4%, respectively. Similarly, Fig. 4b shows a sample reconstruction of the square. The Hausdorff percent errors are 15.79 and 19.22%, respectively.

Monte-Carlo Simulations

Monte-Carlo simulations were used to test the performance of the KBA. All simulations are performed using the kite test object. The Monte-Carlo simulations consisted of 100 independent reconstructions for each scenario. In each plot error bars denote the 95% confidence intervals of the sample mean values that result from the 100 runs of the algorithm. Data are generated under the same conditions but, depending on the experiment, we will vary the SNR, number of retained hypotheses, etc.

Reconstruction Error as a Function of SNR and Number of Views. Figure 5 summarizes the effect on the KBA of varying the SNR and number of views. The number of chord measurements per view N_t was 5 in all reconstructions. Figures 5a and 5b present the Hausdorff errors for both the correct and the IID knot-noise model, respectively. The most dramatic decrease in error occurs when the SNR is increased from 0 to 10 dB. The KBA is unable to resolve the knot-to-vertex data associations in the high noise scenario. At a given SNR all of the reconstructions behave similarly as a function of the number of views even though the relative error levels vary as a function of the SNR. The results that assume an IID knot location noise model are slightly worse than those that assume the correct applied noise model (as expected).

Reconstruction Error as a Function of Limited Angle Projection Data. The reconstruction performance of the KBA over limited angular ranges was also examined. A full set of projection data (i.e., chord measurements and knot location measurements) was generated over the specified angular range with five chord measurements per view at a SNR of 10dB. Next, reconstructions were performed on *subsets* of the projection data containing 18 uniformly spaced views over limited angular ranges. For each angular range, 100 independent Monte-Carlo reconstructions were performed with $N_q = 15$ hypotheses retained at each step of the KBA.

As the results in Fig. 6a demonstrate, the reconstruction errors decrease significantly as the angular range is increased. This result is attributed to poor triangulation geometry over the smaller angular ranges. In addition, above a 45° range the error appears to level off.

In addition, a second set of Monte-Carlo simulations was done over the range [0°, 45°). For this case, 100 reconstructions were performed for cases of 4, 5, 9, 18, and 27 views. The remainder of the parameters (chord length, SNR, etc.) were set exactly as before. Figure 6b shows that the limited angle reconstruction errors decrease as a function of the number of views used to perform the reconstruction. The results of Fig. 6 are significant in that they show that the KBA behaves robustly under limited angle scenarios.

4.2. Knot-and-Chord-Based Algorithm Experiments

In this section we present experiments with the KCBA. Recall this algorithm uses the output of the KBA to initial-

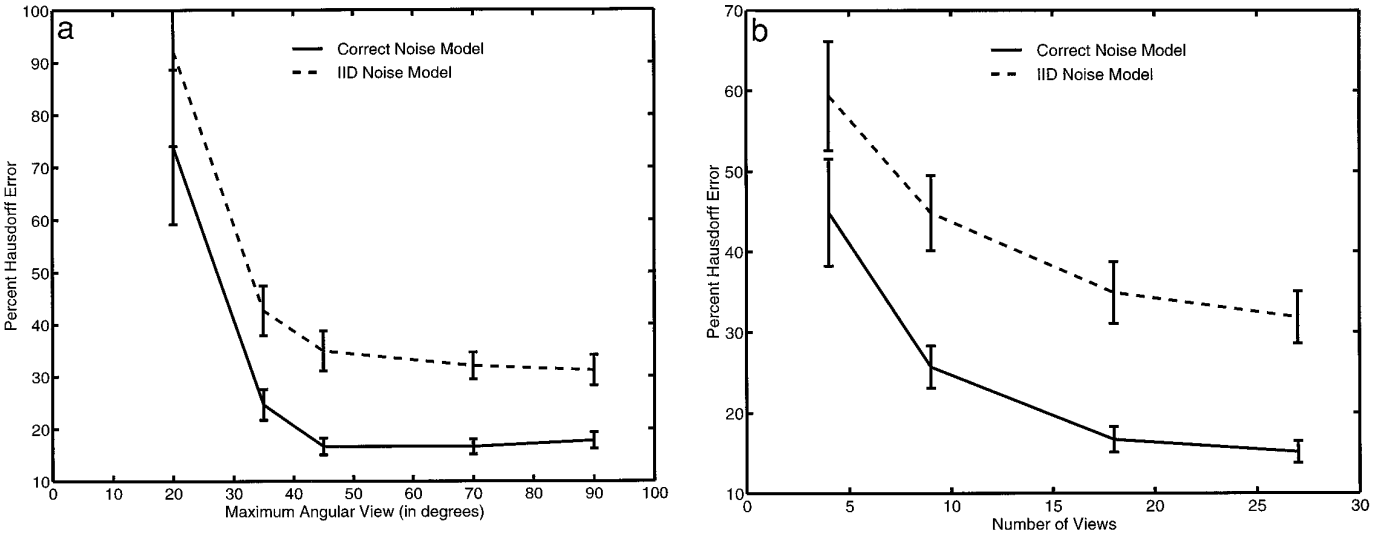

FIG. 6. Performance as a function of limited angular views using the KBA. (a) Varying angular range. (b) Number of views.

ize a minimization of a nonlinear cost function. The data was generated in the same way described in Section 4.1. Further, while the knot data were generated using the variance given in (17), the assumed knot-variance model for the reconstruction was the IID knot-noise model $\sigma^2_{n_i}(\theta_j) = (1.5\ \sigma_m)^2$.

Sample Reconstructions

For all of these results the true object is depicted by solid lines, the initial guess is represented by the dash-dotted lines, and the final nonlinear reconstruction is denoted by dotted lines. The sample reconstructions of this section use the same data set used to generate the sample reconstructions of Section 4.1. Figure 7 shows sample reconstructions of the kite using the KBA to generate the initial guess. The percent Hausdorff error decreases from 20.4% for the KBA to 8.9% for the KCBA showing significant improvement from the initial estimates provided by the KBA.

Monte-Carlo Simulations

Monte-Carlo simulations were also used to test the KCBA. All simulations were performed using the kite object. The Monte-Carlo simulations consist of 50 independent reconstructions for each reconstruction scenario. In each plot error bars denote the 95% confidence intervals of the sample mean values. The initial estimates of the vertices were generated using the KBA with the IID assumed knot location noise model. Finally, unless otherwise stated, $N_q = 15$ hypotheses were retained at each step, five uniformly spaced chord measurements over the range $t = [-1, 1]$ were taken per view, and the SNR on the chord length measurements is set to 10 dB.

Reconstruction Error Versus Quantity of Chord Measurements. The performance of the algorithm as a function of the number of chord measurements was tested. Each reconstruction was based on 18 uniformly spaced views for a total of 72 knot location measurements. The variance on the noise added to the chord length measurements was set to a constant value throughout, so as the number of measurements per view is varied, the SNR will also change. This variance σ^2_m was chosen to yield a SNR

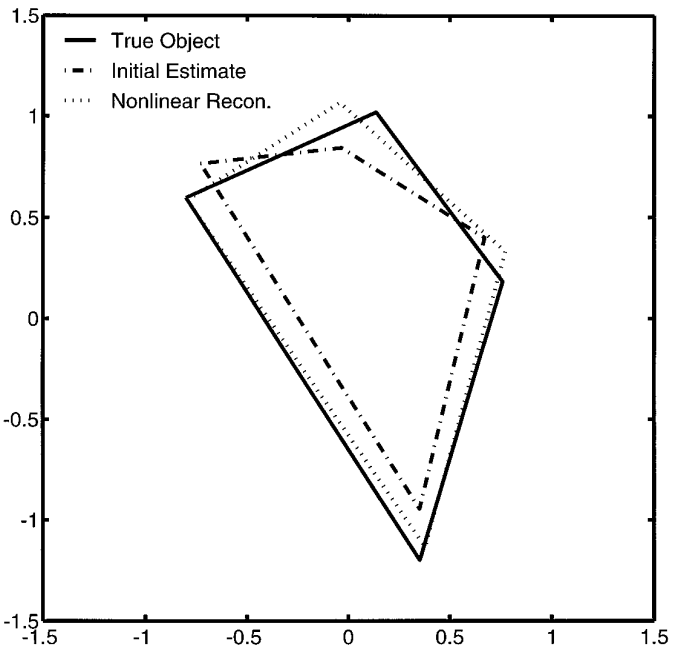

FIG. 7. Sample reconstructions of KCBA, $N_\theta = 27$ views, $N_q = 10$ hypotheses retained, 10 dB SNR on chords.

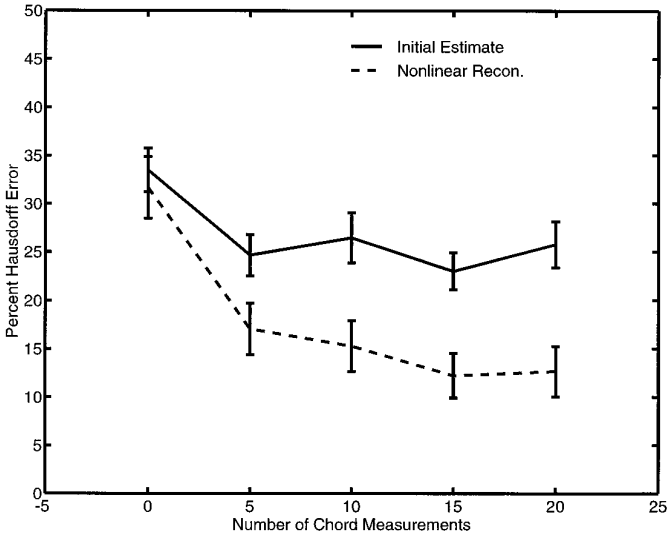

FIG. 8. Performance as a function of number of chords used in the KCBA.

of 10 dB for the case of 10 chord measurements per view and the noise on the knot locations was set with (17) using this variance.

The results are given in Fig. 8. The solid line represents reconstruction errors of the initialization while the dashed line represents final reconstruction error obtained from the algorithm. As the number of chord length measurements increases, the difference between the reconstruction errors of the initial estimate and those of the final nonlinear reconstruction become more pronounced. This occurs because the KCBA uses the information in the chord measurements to obtain the estimate of the vertices while the KBA (used to generate the initial estimate) only uses chord length to prune the possible discrete data association hypotheses.

Reconstruction Error as a Function of SNR and Number of Views. This section characterizes the effect of SNR and number of views on the algorithm. Figure 9 presents the results for both the initial estimate (generated from the KBA) and the final nonlinear reconstruction for four cases: (a) SNR = 0 dB, (b) SNR = 10 dB, (c) SNR = 20 dB, and (d) SNR = 30 dB. As expected, each error curve decreases as the number of views is increased and as the SNR is increased. Thus, the sample reconstructions of Fig. 7 are representative of the kind of error reduction possible with the KCBA (note that in the previous section, the reconstruction errors are more than halved compared to the sample reconstruction of Section 4.1). As for the reconstruction results at a SNR of 0 dB, the errors are also consistent with the previous results. The initial estimate provided to the nonlinear optimization algorithm is too far away from the true vertices of the object, supporting the claim that a good initial guess is important for the nonlinear optimization routine. Additionally, the knot location and chord length measurements may be too noisy to resolve this problem. As a result, the algorithm produces results that have reconstruction errors on the same order of magnitude as the initial estimate.

4.3. End-to-End System Experiments

For completeness, we also perform some experiments based on an end-to-end system. We directly use the noisy projection data as the starting point in the experimental reconstructions, extracting our knot location observations from these data. To extract the knot locations in each projection, we find the ML estimate of the linear spline which best fits the projection data. Further, we do not assume that we know the number of knots in any particular view *a priori*. Instead, we use the Akiake criterion [16] to determine this value. For each view θ_j we generate a dense set of noisy projection data as described in Section 4.1. We then take as the knot observations for the view at angle θ_j the knot locations which minimize the following criterion:

$$\arg\min_{P,\hat{N}_v(\theta_j)} \frac{1}{2}\|m(\theta_j) - \hat{m}(P, \theta_j)\|^2_{\Lambda_v^{-1}} + 2\hat{N}_v(\theta_j). \qquad (18)$$

Here $m(\theta_j)$ are the noisy chord length measurements at angle θ_j, $\hat{m}(P, \theta_j)$ are the corresponding chord estimates obtained from a linear spline defined by the vector P, of length $\hat{N}_v(\theta_j)$ which contains the estimated knot locations and corresponding chord lengths at these points. Thus P defines the linear spline fitting projection θ_j, and gives us the knots locations in that view. Rather than searching over all possible values of $\hat{N}_v(\theta_j)$, we consider only the two possibilities $\hat{N}_v(\theta_j) = 3$ or $\hat{N}_v(\theta_j) = 4$ for tractability. Since the problem is nonlinear and our focus is not on the knot extraction algorithm, we initialize the minimization (18) with the true projection spline parameters to avoid local minima and focus on the subsequent stages of the algorithm.

The Monte-Carlo simulations consisting of 50 independent reconstructions of the kite test object were used to test the performance of the end-to-end system. The estimates were generated with the KCBA using the KBA as the initialization. Projection data were taken at 18 equally spaced angles over the interval [0°, 180°) with equally spaced chord measurements over the range $t = [-1, 1]$ in each view. The noisy chord length data were generated such that there were 50 measurements in each view. These measurements were used to detect the knot location measurements as outlined above. A subset of these noisy chord length data (five equally spaced measurements per view) were then used along with the extracted knot observations as input to the KCBA. The SNR on this subset of chord

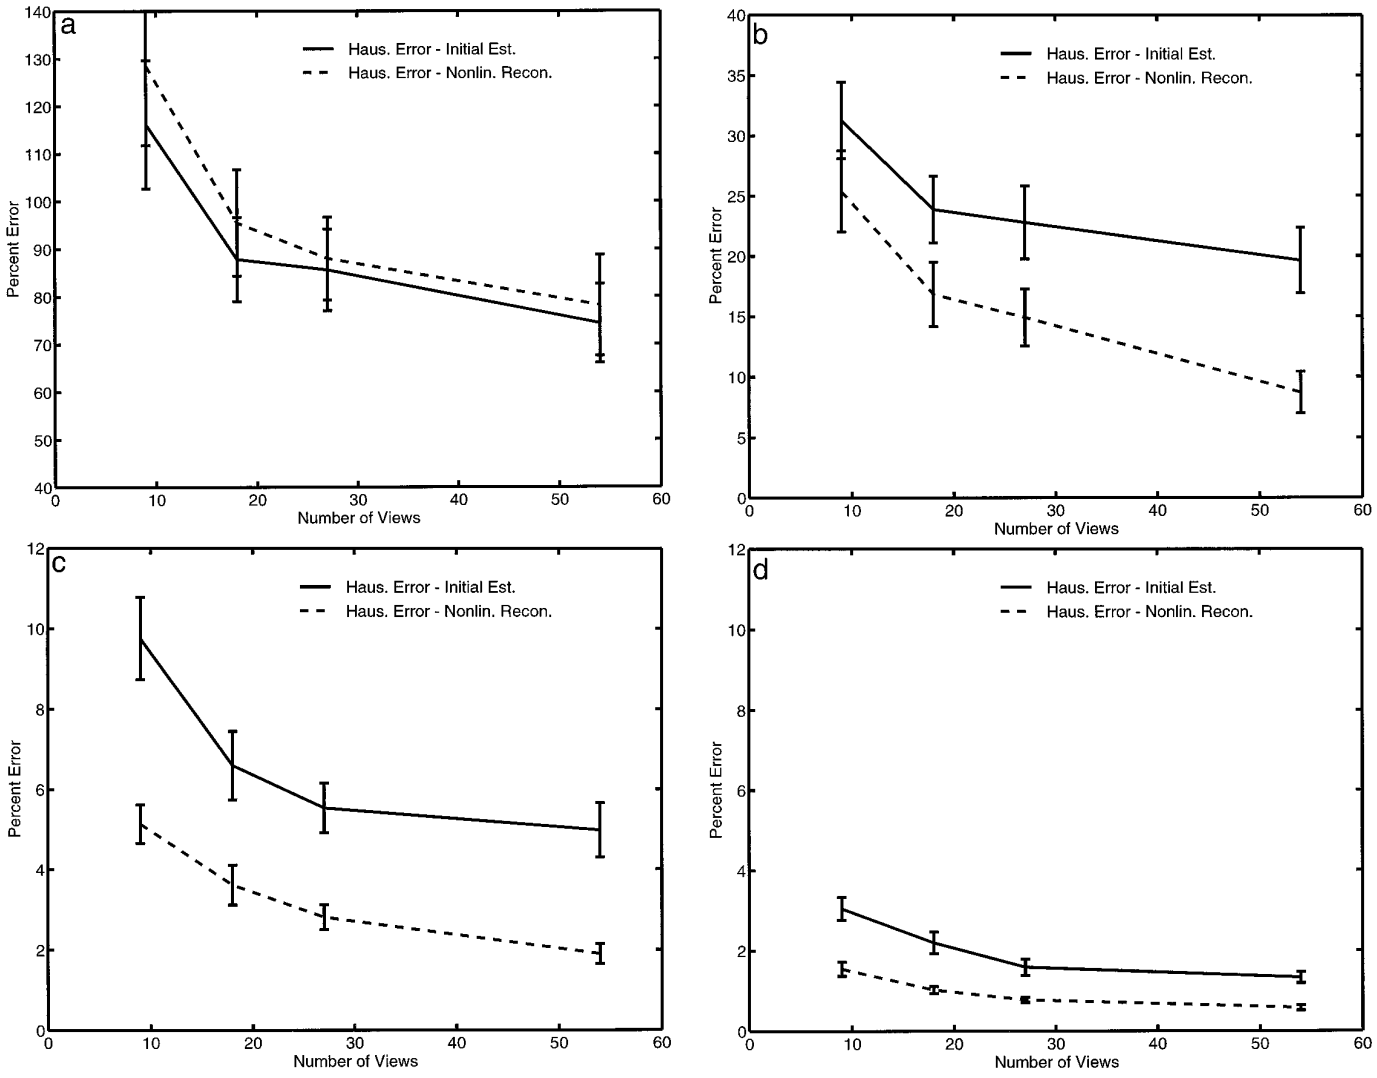

FIG. 9. Performance as a function of number of views used in the KCBA at different SNRs. (a) SNR = 0 dB. (b) SNR = 10 dB. (c) SNR = 20 dB. (d) SNR = 30 dB.

length measurements was set to 10 dB and, for simplicity, the assumed noise model on the knot locations used the variance of method (18), which was empirically determined by Monte-Carlo simulation.

Reconstruction results are presented in Fig. 10. Even though the SNR on the chord length measurements was the same as the cases when we used simulated knot location data, the reconstruction errors are smaller. Therefore, the knot extraction process used in this end-to-end system produced significantly better data than was assumed in the simulated knot location data used in the previous reconstruction results.

5. CONCLUDING REMARKS

In this paper, we have developed two statistically-based algorithms to reconstruct convex binary polygonal objects from sparse and noisy measurements of projection data. Specifically, our techniques concentrated on reconstructions that use the geometric-based information found in measurements of chord lengths and knot locations to estimate the vertices of a polygonal object. We have shown that these algorithms can be used to produce quality reconstructions in the face of limited and noisy projection data.

The first algorithm, the KBA, focuses on techniques that are traditionally used in multitarget tracking problems to determine the position of targets. By solving a simultaneous discrete-valued hypothesis testing problem and continuous-valued ML estimation problem, we avoided the direct nonlinear estimation of the vertices.

The second algorithm developed in this paper, the KCBA, directly uses *both* knot location and chord length data to estimate the vertices of the object. Thus, the KCBA extends the work of the KBA by incorporating both the

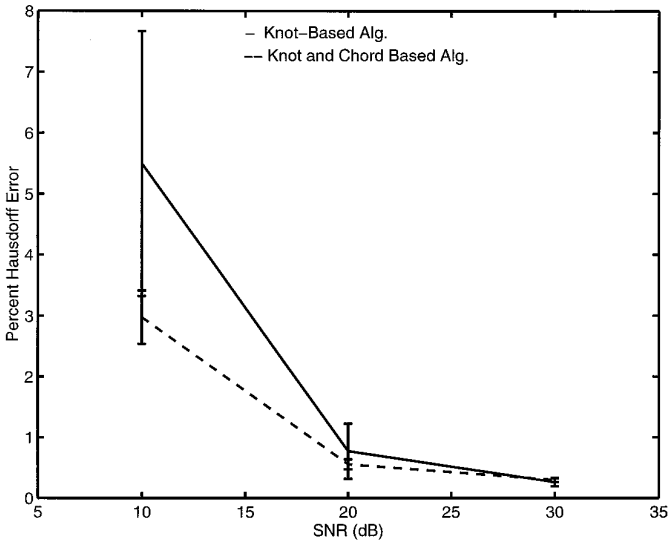

FIG. 10. Performance as a function of SNR using the end-to-end system.

chord length and knot location measurements into the estimate of the vertices.

REFERENCES

1. G. T. Herman, *Image Reconstruction from Projections*, Academic Press, New York, 1980.
2. R. A. Brooks and G. Di Chiro, Principles of computer assisted tomography (cat) in radiographic and radioisotropic imaging, *Phys. Med. Biol.* **21**(5), 1976, 689–732.
3. Y. Bresler and A. Macovski, Three-dimensional reconstruction from projections with incomplete and noisy data by object estimation, *IEEE Trans. Acoust. Speech Signal Process.* **ASSP-35**(8), Aug. 1987, 1139–1152.
4. D. J. Rossi and A. S. Willsky, Object shape estimation from tomographic measurements: A performance analysis, *Signal Process.* **18**(1), Sept. 1989, 63–87.
5. J. L. Prince and A. S. Willsky, Convex set reconstruction using prior shape information. *CVGIP Graphical Models Image Process.* **53**(5), Sept. 1991, 413–427.
6. W. C. Karl, G. C. Verghese, and A. S. Willsky, Reconstructing ellipsoids from projections. *CVGIP Graphical Models Image Process.* **56**(2), Mar. 1994, 124–139.
7. P. Milanfar, W. C. Karl, and A. S. Willsky, Reconstructing binary polygonal objects from projections: A statistical view, *CVGIP Graphical Models Image Process.* **56**(5), 1994, 371–391.
8. S. Jaggi, W. C. Karl, and A. S. Willsky, Dynamic estimation of left-ventricular ejection fraction, in *Proc. IS&T/SPIE Symposium on Electronic Imaging Science and Technology: Biomedical Image Processing and Biomedical Visualization, San Jose, CA, Jan. 1993* (R. S. Acharya and D. B. Goldgof, Eds.), Vol. 1905, pp. 206–217.
9. I. Polyak, A. S. Willsky, and W. C. Karl, Robust knot detection and spline approximation using wavelet transform extrema and multitarget tracking, in *Proceedings of the IEEE-SP International Symposium on Time-Frequency and Time-Scale Analysis, Oct. 1994*, pp. 290–293.
10. A. M. Mier-Muth and A. S. Willsky, *A Sequential Method for Spline Approximation with Variable Knots*, MIT Electronic Systems Laboratory ESL-P-759, MIT, 1977.
11. T. Kurien, A. L. Blitz, R. B. Washburn, and A. S. Willsky, Optimal manuever detection and estimation in multiobject tracking, in *Proceedings of the Sixth MIT/ONR Workshop on Command and Control, Dec. 1983*, pp. 164–171.
12. Donald B. Reid, An algorithm for tracking multiple targets, *IEEE Trans. Automat. Control* **24**(6), Dec. 1979.
13. T. Kurien, Issues in the design of practical multitarget tracking algorithms, in *Multitarget-Multisensor Tracking: Applications and Advances* (Y. Bar-Shalom, Ed.), Chap. 3, Artech House, Norwood, MA, 1992.
14. H. A. P. Blom and Y. Bar-Shalom, The interacting multiple model algorithm for systems with markovian switching coefficients, *IEEE Trans. Automat. Control* **33**(8), Aug. 1988, 780–783.
15. M. Berger, *Geometry I*, Springer-Verlag, Berlin/New York, 1987.
16. J. Rissanen, Modeling by shortest data description, *Automatica* **14,** 1978, 465–471.