

Reconstructing Ellipsoids from Projections¹

WILLIAM C. KARL, GEORGE C. VERGHESE, AND ALAN S. WILLSKY

Massachusetts Institute of Technology, Room 35-421, Cambridge, Massachusetts, 02139

Received October 1, 1992; revised November 16, accepted December 7, 1993

In this paper we examine the problem of reconstructing a (possibly dynamic) ellipsoid from its (possibly inconsistent) orthogonal silhouette projections. We present a particularly convenient representation of ellipsoids as elements of the vector space of symmetric matrices. The relationship between an ellipsoid and its orthogonal projections in this representation is linear, unlike the standard parameterization based on semiaxis length and orientation. This representation is used to completely and simply characterize the solutions to the reconstruction problem. The representation also allows the straightforward inclusion of geometric constraints on the reconstructed ellipsoid in the form of inner and outer bounds on recovered ellipsoid shape. The inclusion of a dynamic model with natural behavior, such as stretching, shrinking, and rotation, is similarly straightforward in this framework and results in the possibility of dynamic ellipsoid estimation. For example, the linear reconstruction of a dynamic ellipsoid from a single lower-dimensional projection observed over time is possible. Numerical examples are provided to illustrate these points. © 1994 Academic Press, Inc.

1. INTRODUCTION

Ellipsoids arise in many disciplines as simple yet effective object models. Such ellipsoidal models are used both directly to capture shape and indirectly as bounding approximations. In [1] an ellipse is used as a simple parameterized model in an attempt to recover object eccentricity and orientation from low signal-to-noise ratio tomographic data. In other medical areas ellipsoids are used to model both the shape and volume or area of anatomical parts, such as the heart, spine, and blood vessels [2–5]. In [6, 7] the state of a system is assumed confined to an unknown n -dimensional ellipsoid and the goal is essentially to reconstruct this ellipsoid from observations of its lower-dimensional projections. In [8] a group of closely spaced targets in space is observed through a number of passive sensors. The cluster of targets is modeled as an

ellipsoid and the observations as projections of it over time. The desire is to find the evolution of the shape of the ellipsoid. All these problems share the common goal of reconstructing an ellipsoid from (possibly inconsistent) observations of its lower-dimensional silhouette or shadow projections. In addition, many of these problems also involve a sequential element, in that the observations are extended in time or space.

This paper focuses on the orthogonal projection and reconstruction of centered n -dimensional ellipsoids, where by “orthogonal projection” we mean the generation of a silhouette shape through orthographic projection. Such silhouette projection observations arise in many ways in object reconstruction problems. In the realm of robotics, they can arise from repeated grasps or probes by a gripper [9, 10]. In low dose tomography the line integral observations may yield little more than shadow information [11–13], thus fitting into the silhouette framework above. Even when this is not the case, a preliminary step of projection support extraction coupled with object boundary estimation may be useful or desirable [11, 14]. This approach has proven particularly helpful in reflection tomography arising in laser range data [15]. A similar approach involving preextraction of boundary information was successfully used in [16, 17] in the estimation of cardiac ejection fraction.

In Section 2 we associate a centered ellipsoid with an underlying positive semidefinite (PSD) symmetric matrix. This symmetric matrix becomes our representation of the ellipsoid. While this association is not new, it has not been used, to our knowledge, for ellipsoid reconstruction. In Section 3 the connection between such an ellipsoid representation and that of its orthogonal projection is examined. In particular, we show that this relationship is *linear*, in contrast to that obtained for the commonly used parameterization based on ellipsoid semiaxis length and orientation angle. Since the set of symmetric matrices forms a vector space, a natural isomorphism exists between ellipsoids and points or vectors in this space, as discussed in Section 4. Next, the inverse problem of reconstructing a centered ellipsoid from a series of its orthogonal projections is treated in Section 5. Our repre-

¹ This work supported by the National Science Foundation under Grant MIP-9015281, the Army Office of Sponsored Research under Grant DAAL03-92-G-0115, and the Office of Naval Research under Grant N00014-91-J-1004. Address for correspondence: William C. Karl, MIT, Room 35-421, Cambridge, MA 02139.

sentations allow us to cast this problem in a standard linear estimation form and thus to provide a concise characterization of its solution. The representation of ellipsoids as symmetric matrices also provides a natural way to include certain constraints in the problem formulation, in the form of inner and outer bounds on recovered shape. Such constraints can represent our prior knowledge of possible shapes for a problem. The inclusion of a geometrically natural dynamic model into the linear estimation framework developed in Section 5 is straightforward, allowing the possibility of dynamic ellipsoid shapes. The generation and reconstruction from projections of such dynamic ellipsoids are examined in Section 6. Finally in Section 7 we provide some numerical experiments to illustrate the developments of the paper. Our focus in this paper is the presentation of an especially convenient framework for ellipsoid representation and manipulation and not on the statistical issues raised by any particular problem, which we leave for other works.

2. SYMMETRIC MATRICES

Here we develop the tie between ellipsoids and PSD symmetric matrices. A common way of representing an ellipsoid is by its semiaxis lengths together with their orientation angles with respect to the coordinate axes [3–5, 18]. While this representation is intuitive, in that these parameters directly capture elements of the ellipsoid geometry, it proves inconvenient for reconstruction as it leads to complicated nonlinear relationships between the original ellipsoid representation and that of its projections. Instead, we use an alternative representation of an ellipsoid as a symmetric matrix, which then leads to a simple linear relationship between the ellipsoid and its orthogonal projection. Any $n \times n$ symmetric PSD matrix X can be taken to represent an n -dimensional ellipsoid \mathcal{E} centered on the origin, comprising the set of points given by

$$\{z | z^T u \leq h(u) = \sqrt{u^T X u}, \forall u^T u = 1, u, z \in R^n\}. \quad (1)$$

Here $h(u) = \sqrt{u^T X u}$ is the (reduced) *support function* of the ellipsoid, [7, 19]. Conversely, for any ellipsoid \mathcal{E} centered at the origin, a unique symmetric PSD matrix can be found such that the description (1) characterizes the set of points of the ellipsoid. Thus, we can represent any ellipsoid \mathcal{E} by its corresponding PSD symmetric matrix X . Note that the ellipsoid is degenerate (i.e. has zero extent) in directions associated with vectors u in the null space of X . If X is positive definite, so its inverse exists, then an equivalent definition of the ellipsoid is given by

$$\{z | z^T X^{-1} z \leq 1, z \in R^n\}. \quad (2)$$

The geometric properties of the ellipsoid \mathcal{E} are reflected in algebraic properties of the corresponding symmetric




FIG. 1. Problem definition.

matrix X in a *natural* way. If λ_i, v_i are the eigenvalues and eigenvectors of the matrix X , then the principal axes of the ellipsoid are in the directions given by the v_i and the corresponding semi-axis lengths are given by the $\sqrt{\lambda_i}$. Thus going back and forth from an ellipsoid to its matrix representation is a simple matter. Further, any linear coordinate transformation of the form $\tilde{z} = Lz$ is reflected by a change of the corresponding representing matrix X to a matrix X_L given by [7]

$$X_L = LXL^T. \quad (3)$$

Such transformations can include, for example, rotation and stretching. For convenience in what follows, we will often not distinguish between the ellipsoid \mathcal{E} and the matrix X that represents it. In the following we restrict consideration to ellipsoids centered at the origin. Note that there is no loss of generality in this assumption if the center is known (as it is here) since, given a noncentered ellipsoid, we may always translate our coordinate system to the ellipsoid origin, thereby recovering the centered case.

3. ELLIPSOID PROJECTION

We now consider the orthogonal projection of the n -dimensional ellipsoid \mathcal{E} onto an m -dimensional subspace \mathcal{S} to obtain an m -dimensional shadow object, see Fig. 1 (note that both \mathcal{E} and \mathcal{S} actually pass through the origin, but are separated in the figure for clarity). This projected object will itself be an m -dimensional ellipsoid \mathcal{E}_s in the subspace \mathcal{S} . Let X be the $n \times n$ symmetric matrix representing \mathcal{E} and let Y denote the corresponding $m \times m$ matrix representing the ellipsoid \mathcal{E}_s in \mathcal{S} . If C is a matrix whose m columns form an orthonormal basis for the projection subspace \mathcal{S} , then the transformation from the coordinates z of points of the original space to their coordinates \tilde{z} in the projection is given by the relation $\tilde{z} = C^T z$. Applying (3) we find that the relationship between the original ellipsoid, specified by X , and its projection, specified by Y in the projection subspace, is simply given by

the equation²

$$Y = C^T X C. \quad (4)$$

In particular, this relationship between X and Y is *linear* in the elements of the matrix X . This relationship is much simpler than the nonlinear one existing between the commonly used semiaxes length and orientation parameters of the ellipsoid and those of its projection.

Note that we could equivalently represent our projected ellipsoid \mathcal{E}_y by the matrix $\bar{Y} = C Y C^T$ rather than Y . The latter can always be recovered from the former, since C has full column rank. Whereas the $m \times m$ matrix Y represents a nondegenerate ellipsoid (one with no semiaxes of zero length) with respect to the subspace \mathcal{S} , the singular $n \times n$ matrix \bar{Y} represents a degenerate ellipsoid in the original space. The advantage of this alternate form for the projection is that the relationship (4) between the ellipsoid matrix and its projection then becomes

$$\bar{Y} = \bar{C} X \bar{C}, \quad (5)$$

where $\bar{C} = C C^T$ is now a true projector onto the subspace \mathcal{S} , i.e., is symmetric and satisfies $\bar{C}^2 = \bar{C}$. A consequence is that the projection \bar{Y} is invariant under different choices of basis C for the projection subspace \mathcal{S} . In what follows we will continue to use the representation Y because of its more transparent connection to physically measured quantities through (3), yet all of our results may be phrased in terms of the coordinate independent projections \bar{Y} . Such an approach is used in [21].

Finally, the special case of 1-dimensional projections is of interest for its connection to support measurements, which arise in a variety of applications from the robotics to medical imaging [10–12, 22–24]. The (reduced) support function $h(u)$ of an object is a scalar function of the direction specified by the vector u [19]. It gives a measure of the extent of an object in the direction u . If the subspace of projection \mathcal{S} is 1-dimensional, then C in (4) is a unit vector and the ellipsoid resulting from orthogonal projection is a line segment bounded by the support values $h(C)$ and $h(-C)$. Thus 1-dimensional shadows or projections correspond precisely to a pair of support observations in opposite directions. For 1-dimensional projections our observation, as given by the right hand side of (4), reduces to the expression for the squared support function $h_x^2(C)$ of the ellipsoid \mathcal{E} in the direction C , so that $Y = h_x^2(C)$. A consequence of this relationship is that our linear ellipsoid estimates of Section 5 based on Y suggest natural and convenient estimation schemes involving (squared)

support measurements. Of course, we may use these techniques even when the underlying object is not ellipsoidal, using the best fitting ellipsoid to obtain orientation and eccentricity information about an object. Such connections are explored in more detail in [11, 12, 14, 18, 21].

4. MATRIX-VECTOR REPRESENTATION

While the representation of an ellipsoid as a symmetric matrix X is convenient in that the projection relationship (4) is linear, manipulating such quantities can be somewhat cumbersome. The most common way to recast (4) in standard matrix-vector form results in the following equivalent relationship [25, 26]:

$$\text{vec}(Y) = (C^T \otimes C^T) \text{vec}(X), \quad (6)$$

where $\text{vec}(X)$ is the vector obtained by stacking the elements of X columnwise, and $A \otimes B = [a_{ij}B]$ is the Kronecker product of A and B , formed by taking all products of entries of A with B . While this approach does allow easy manipulation of the underlying quantities, it suffers from the problem that it is redundant. Not all the elements in the vectors $\text{vec}(Y)$ and $\text{vec}(X)$ are independent, since the matrix arguments are symmetric. Of course this redundancy can be eliminated by using the symmetry condition to reduce the dimension of the vectors $\text{vec}(Y)$ and $\text{vec}(X)$ from their original m^2 and n^2 elements to only the $m(m+1)/2$ and $n(n+1)/2$ independent elements, respectively [27–29]. The corresponding rows of $(C^T \otimes C^T)$ are eliminated as well, reducing its size. While this reduction is conceptually straightforward and may always be done, it destroys the special structure present in (4) and obscures the relationships between the quantities involved.

Instead of the above approach, we obtain a matrix-vector product which, while equivalent to (4), is *naturally induced* by the properties of the set of symmetric matrices. This transformation will allow us a simple characterization of the solvability of the inverse problem of determining X from a series of observations of the form (4). To this end, note that the set of $n \times n$ symmetric matrices together with the inner product $\langle A, B \rangle \equiv \text{tr}(A^T B)$ defines an $n(n+1)/2$ -dimensional Euclidean space.³ In particular, let \mathcal{X} denote the $n(n+1)/2$ -dimensional such space containing the original matrix X and let \mathcal{Y} denote the corresponding $m(m+1)/2$ -dimensional space containing the projection Y . In these spaces each point represents a symmetric matrix and conversely each symmetric matrix corresponds to a unique point. In particular, let x and y be the vector representations of the matrices $X \in \mathcal{X}$ and

² Interestingly, the result (4) is also the same algebraic relationship as is found between the curvature Hessian of a smooth surface at a point and that of its projection [20].

³ This inner product induces the Frobenius norm on a matrix $\langle A, A \rangle^{1/2} = \|A\|_F$.

$Y \in \mathcal{Y}$ with respect to corresponding orthonormal basis sets $\{M_j^{(n)}\}$ and $\{M_i^{(m)}\}$, so that:

$$(x)_j = \langle X, M_j^{(n)} \rangle \quad (7)$$

$$(y)_i = \langle Y, M_i^{(m)} \rangle \quad (8)$$

for $1 \leq j \leq n(n+1)/2$, $1 \leq i \leq m(m+1)/2$, where $(\cdot)_i$ denotes the i -th component of the argument. We may then view the *vectors* x and y as representations of the ellipsoid and its projection, respectively. For given bases these vectors are unique, nonredundant, and have a simple and clear relationship to the corresponding symmetric matrices. Note in particular that since the basis sets are orthonormal we have, e.g., for any $X_1, X_2 \in \mathcal{X}$ that $\langle X_1, X_2 \rangle = x_1^T x_2$. The basis sets $\{M_j^{(n)}\}$ and $\{M_i^{(m)}\}$ can be obtained, for example, by Gram-Schmidt orthogonalization of any spanning set of symmetric matrices. A particularly convenient choice of basis set is given by what we term the *standard symmetric basis*:

$$M_l^{(n)} = \begin{cases} e_i e_i^T & \text{if } l = (i-1)(2n+2-i)/2 + 1, \\ & 1 \leq i \leq n \\ (e_i e_j^T + e_j e_i^T)/\sqrt{2} & \text{if } l = i(2n+1-i)/2 - n + j, \\ & 1 \leq i < j \leq n, \end{cases} \quad (9)$$

where e_i is the i -th standard basis vector composed of all zeros except for a 1 in the i -th location. In this basis each entry of x is proportional to a single entry of the matrix X .

Now the matrix relating x and y is easily obtained from (4) and the definitions (7) and (8). In particular, if we denote this matrix by \tilde{C} , it is straightforward to show that

$$(\tilde{C})_{ij} = \langle M_i^{(m)}, C^T M_j^{(n)} C \rangle \quad (10)$$

for $1 \leq j \leq n(n+1)/2$, $1 \leq i \leq m(m+1)/2$, where $(\cdot)_{ij}$ denotes the ij -th component of the argument. We may now represent our original relation (4) equivalently as:

$$y = \tilde{C}x, \quad (11)$$

where the vectors y and x have *natural* interpretations as symmetric matrices. Finally, note that if we had used the projector form (5) rather than (4) for our projection definition, the resulting matrix \tilde{C} would itself also be a projector [21].

5. RECONSTRUCTION

Now we are in a position to consider the inverse problem of reconstructing an ellipsoid \mathcal{E} from observation of a set of its (possibly inconsistent or noisy) orthogonal

projections onto the subspaces \mathcal{S}_i . From the discussion in Section 2, we may represent the desired ellipsoid by the symmetric matrix X and its projections onto the subspaces \mathcal{S}_i by the corresponding symmetric matrices Y_i . If C_i are matrices whose columns form orthonormal bases for the subspaces of projection \mathcal{S}_i (assumed known), then the relationship between the ellipsoid matrix X and its projections Y_i is given by (4). Our problem then is to determine the $n \times n$ positive semi-definite, symmetric matrix X , given observations of the form:

$$Y_i = C_i^T X C_i \quad (12)$$

for $1 \leq i \leq q$, where the matrices C_i have orthonormal columns.

Now we may use the relation (11) to express each of the observations as

$$y_i = \tilde{C}_i x \quad (13)$$

where, as for (11), x is the representation of the X , y_i is the representation of Y_i , and \tilde{C}_i is the matrix representation of the operator $C_i^T(\cdot)C_i$ as defined in (10). Stacking up the individual observation vectors y_i into a single vector, we obtain the overall relation:

$$\begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_q \end{bmatrix} = \begin{bmatrix} \tilde{C}_1 \\ \tilde{C}_2 \\ \vdots \\ \tilde{C}_q \end{bmatrix} x$$

or

$$y = Cx, \quad (14)$$

where y and C are defined in the natural way from the stacked observations. Thus, without a semidefiniteness constraint on the reconstruction, the problem of recovering the original ellipsoid \mathcal{E} is equivalent to finding the unknown vector x representing the desired ellipsoid, given the observations y and the projection geometry specified in C . The formulation (14) is a standard one in linear estimation. In general, however, a semi-definiteness constraint is needed on the matrix represented by x ; such issues are discussed in more detail in the next section.

The reader should note that our assumed observations are the projected *ellipsoids* $\mathcal{E}_{\mathcal{S}_i}$, or equivalently the matrices Y_i . Thus, the above formulation has an implicit step of *ellipsoid extraction* from the projections. For many situations this assumption should not pose a significant difficulty. Much work exists, for example, on extracting ellipses from planar data [30–36]. For the case of 1-dimen-

sional projections in particular, fitting the ellipse corresponds to nothing more than extracting the region of support of a line segment. Finally, we may always view a projection of any dimensionality as a group of (noisy) 1-dimensional projections instead of a single higher dimensional one, thus reducing the problem to the 1-dimensional case. This insight essentially reduces the observation problem to one of boundary point determination. In particular, we could perform the preliminary step of extracting the observed ellipsoids from each projection by using this technique of fitting (lower dimensional) ellipsoids to their 1-dimensional support data. We could also directly use all such sets of 1-dimensional observations from all projections simultaneously to directly estimate the desired ellipsoid. Note that whatever method is used to extract the ellipsoid observations from raw data, any noise in the data will manifest itself as perturbations in the parameters of the corresponding observed projected ellipsoid. Since the focus of the present paper is not ellipsoid extraction, however, we will simply assume that we are given a set of such ellipsoid observations (which thus correspond to PSD matrices).

5.1. Unconstrained Reconstruction

In this section we consider the solution of (12) without a PSD constraint on the resulting answer. First, we consider reconstruction from a *consistent* or noise-free set of observations. In this case, a PSD solution exactly matching the data always exists and thus the constraint is not needed. The formulation of (14) allows us to easily characterize the unique solutions of (12) in such a case. In particular, the inverse problem (12) has a unique solution if and only if the matrix \mathbf{C} of (14) has rank equal to $n(n+1)/2$ (i.e. full column rank). This solution, if it exists, is given by

$$x = \mathbf{C}^L \mathbf{y}, \quad (15)$$

where \mathbf{C}^L is any left inverse of \mathbf{C} .

We can now use this result to obtain conditions for reconstruction of \mathcal{E} that are stated directly in terms of the projection subspaces \mathcal{G}_i . For example, if the projection subspaces are restricted to be hyperplanes (so that $m = n - 1$), then three distinct such projections are necessary and sufficient to uniquely recover \mathcal{E} . Such results are obtained by using the definition (10) of \mathbf{C} coupled with a counting argument on the number of independent rows of \mathbf{C} . Other statements of this kind are, of course, possible, as discussed in [21]. The linear projection relationship (12) and isomorphic relationship between symmetric matrices and their representations makes such calculations straightforward, if tedious.

Inconsistent observations. Now consider the case of performing reconstruction based on a set of inconsistent

observations of the form:

$$Y_i = \mathbf{C}_i^L \mathbf{X} \mathbf{C}_i + \mathbf{W}_i, \quad (16)$$

where the \mathbf{W}_i are symmetric perturbations from ideal values. Such observations will in general result in an inconsistent set of equations of the form (14). In this case we may seek, for example, the unconstrained linear least squared error (LS) solution to the set. This estimate \mathbf{X}_{LS} is obtained as the solution of:

$$\mathbf{X}_{LS} = \arg \min_{\mathbf{X}} \sum_{i=1}^q \|\mathbf{Y}_i - \mathbf{C}_i^L \mathbf{X} \mathbf{C}_i\|_F^2 \quad (17)$$

The corresponding vector x_{LS} representing the solution matrix is the solution to:

$$\min_x \|\mathbf{y} - \mathbf{C}x\|_2^2 \quad (18)$$

The solution is obtained by choosing $\mathbf{C}^L = \mathbf{C}^+$, the Moore-Penrose inverse of \mathbf{C} , in (15). Thus $x_{LS} = \mathbf{C}^+ \mathbf{y}$ is the desired LS estimate without a semi-definiteness constraint on the solution.

To correspond to an ellipsoid, the matrix \mathbf{X} , represented by the vector x , must be positive semi-definite. The LS estimate given by $x_{LS} = \mathbf{C}^+ \mathbf{y}$ has no such constraint to guarantee this PSD property of the solution. If, however, a PSD matrix is obtained as the LS estimate without such a constraint, then clearly it is also the LS estimate subject to such a constraint. For many problems the observations are clean enough that the PSD nature of the solution is maintained anyway and no further effort is needed. In fact, it is possible to state sufficient conditions for this to be true, in the form of bounds on the allowed perturbations \mathbf{W}_i in the ideal observations as a function of the singular values of the matrices \mathbf{C} and \mathbf{X} . In particular, if the underlying matrix \mathbf{X} is a PSD matrix, the LS estimate will also be PSD if the following condition is satisfied [21]:

$$\sigma_{\min}^{-1}(\mathbf{C}) \sqrt{\sum_{i=1}^q \|\mathbf{W}_i\|_F^2} \leq \lambda_{\min}(\mathbf{X}), \quad (19)$$

where $\lambda_{\min}(\cdot)$ denotes the minimum eigenvalue of the argument, $\sigma_{\min}(\cdot)$ the minimum singular value, and the \mathbf{W}_i are the differences in the observations from their ideal values as given by (4).

Let us interpret this condition in terms of the underlying geometric quantities. In general, we would like to have the quantities $\lambda_{\min}(\mathbf{X})$ and $\sigma_{\min}(\mathbf{C})$ large and the quantities \mathbf{W}_i small. The term under the square root involving the \mathbf{W}_i can be taken as a measure of the overall "noise power" in the observations. Recalling the tie between an ellipsoid \mathcal{E} and the eigenvalues of the corresponding PSD symmet-

ric matrix X representing it, the term $\lambda_{\min}(X)$ can be seen to be the squared length of the smallest semiaxis of the ellipsoid. This length might be thought of as a measure of the closeness of the ellipsoid to degeneracy. Finally the term $\sigma_{\min}(\mathbf{C})$ reflects the nearness to singularity of the matrix \mathbf{C} . If this quantity is small, the columns of \mathbf{C} are nearly dependent. Since \mathbf{C} captures the effect of observation geometry, such a situation reflects the fact that our observations are nearly linearly dependent, as might happen if we were to use a set of projections on subspaces very close to each other (see, e.g., [37] for more detail on the nearness of subspaces). In summary, if the smallest aspect of the underlying ellipsoid is large relative to the perturbations in the ideal observations, and if our set of observations are well placed, we should be able to use the unconstrained LS estimate without the necessity of a PSD constraint on the solution.

Before proceeding, note that we may also interpret a set of inconsistent observations (16) in a *stochastic* setting, obtaining a Bayesian or maximum-likelihood estimate for the matrix X . The linearity of the observation equation greatly facilitates such an approach. In particular, if we choose to model the noise in the observations W_i as being jointly Gaussian random variables in each entry,⁴ then the linearity of the problem ensures that the noise in the elements of y will also be Gaussian processes (not independent in general). Thus it is a straightforward matter to obtain statistically optimal (unconstrained) estimates. In general, the noise model and thus cost function we would use would depend, e.g., on the previous processing step which extracted the observed ellipsoid observations. Note that recursive implementations of both the LS and statistical solutions are also straightforward using the formulation given in (16). In the present paper our interest is in development of the matrix based ellipsoid reconstruction framework and not on such statistical issues. As a result, we take a deterministic view and treat any deviations W_k of the actual observations from ideal values as unknown perturbations, though such stochastic interpretations are easily accommodated.

5.2. Constrained Reconstruction

While in many instances an unconstrained, LS-type solution as presented in (17) is adequate, there may be situations when, due to incomplete, noisy observations, this estimate is *not* positive semidefinite. In addition, there may be situations where, because of prior information, we wish to impose constraints on the reconstructed ellipsoid in the form of bounds on its shape. Such bounds might reflect our prior knowledge of the minimum or maximum breadth or of the orientation of an object we wish to

estimate. Our formulation of the ellipsoid recovery problem as one of symmetric matrix estimation suggests a natural notion of solution bound and allows us to apply existing methods and algorithms for such constrained matrix problems here.

To this end, consider the requirement that the reconstructed ellipsoid matrix X lie in the matrix interval given by $\bar{X} \geq X \geq \underline{X}$. By such matrix inequalities we mean that the matrices $(X - \underline{X})$ and $(\bar{X} - X)$ are positive semi-definite. In particular, the PSD constraint is recovered if we choose $\underline{X} = 0$ and $\bar{X} = \infty \cdot I$. Such semi-definite interval matrix constraints, in turn, serve to naturally capture geometric constraints on the underlying ellipsoid. Specifically any ellipsoid satisfying such a condition will be contained within the outer extreme ellipsoid corresponding to \bar{X} and will contain the inner extreme ellipsoid associated with \underline{X} . The planar case is illustrated in Fig. 2, where the allowed area is the white region and some elements on the boundary of the interval set are shown as ellipses touching the boundary of one extreme or the other. This geometric interpretation follows from the definition of an ellipsoid given in (2) and the positive semi-definiteness of the quantities $(X - \underline{X})$ and $(\bar{X} - X)$ [21].

Now consider the special case of the constraint set obtained when the extreme matrices are given by a scalar times the identity, so that $\underline{X} = \underline{\alpha}I$ and $\bar{X} = \bar{\alpha}I$. The corresponding extreme ellipsoids then become nested spheres. This case corresponds to putting simple eigenvalue constraints on the reconstructed matrix X . Such constraints are nondirectional, since they do not favor one ellipsoid orientation over another. This fact is reflected in the central symmetry of the extreme ellipsoids. If $\underline{\alpha} = 0$ and we let $\bar{\alpha} \rightarrow \infty$ we again recover the PSD constraint. Since the inner bound is just the origin, note that the (possibly degenerate) ellipsoids corresponding to such a set of PSD matrices may be little more than a line. As expected, these ellipsoids, containing the origin, exist at the boundary of the set, corresponding to singular matrices X .

The algebraic problem of reconstructing a symmetric matrix under such semidefinite interval constraints as we




FIG. 2. Illustration of interval matrix geometry.

⁴ Note, however, that to correspond to observations of ellipsoids the matrices Y_i themselves must be PSD.

have been discussing is treated in detail in [21], where algorithms yielding a solution to the problem are given. The set of PSD matrices $X \geq 0$ forms a convex cone, so that the sets $(X - \underline{X}) \geq 0$ and $(\bar{X} - X) \geq 0$ are simply shifted and perhaps flipped versions of this PSD cone. When these constraints are active in a problem (i.e. affect its solution), then the solution matrix must lie on the boundary of the cone constraint set. The algorithms described in [21] iteratively form polygonal approximations to the constraint cone which are increasingly refined in the vicinity of the solution point. These polygonal approximations, obtained as the intersection of half-spaces, are defined by sets of linear inequalities and, in the limit, exactly capture the local constraint cone shape. These inequalities may be viewed as defining an infinite-dimensional linear programming problem in the limiting case. A pseudocode version of such an algorithm is given in the Appendix. In general these algorithms are not finite, but they do converge quickly. Reconstructions using such algorithms are demonstrated in Section 7. The special case of constrained reconstructions based on support function observations is discussed in [21].

6. DYNAMIC PROBLEMS

In this section we treat the problem of generating and estimating a *dynamically evolving* ellipsoid. These issues are a direct extension of our work on the static case in Section 2. We demonstrate a particular symmetric evolution equation and show how we may tailor the geometric characteristics of the ellipsoid evolution through choice of the dynamic matrix. This connection, based on our representation of the ellipsoid as a symmetric matrix together with the relation (3), is particularly simple. Our isomorphism between symmetric matrices and vectors then allows us to easily represent this evolution in a standard state-space form. Following this examination of the generation problem, we treat the inverse problem of *estimating* a dynamic ellipsoid from observations of its projections. Such dynamic ellipsoid problems appear, for example, in regard to tracking beating hearts in a series of images [16, 17, 38], following moving clouds of particles [6], and even in the tracking of blood vessel cross-sections [3, 4], where the ellipse evolution is *spatial*. Taken together, our framework provides both a simple evolution mechanism, with direct control of geometrically meaningful quantities of interest, coupled with a linear observation equation through (4). These characteristics stand in contrast to dynamic formulations based on parameterizations of the underlying ellipsoid in terms of orientation and axis-length, which result in nonlinear observation equations [3–5, 18] and, further, do not generalize easily to dimensions greater than 3.

6.1. Generation

We may animate an ellipsoid by imposing a dynamic relationship on its set of parameters. In particular, evolution of the elements of the representing matrix X will yield a corresponding dynamically evolving ellipsoid. We need only ensure that the resulting series of matrices X_k have the desired geometric behavior and remain symmetric and positive semidefinite. We thus seek an evolution structure on X_k that is simple to implement, easy to understand, yields interesting dynamical behavior, yet maintains symmetry and positive semidefiniteness. It is especially simple to obtain such desired behavior given this symmetric matrix ellipsoid representation.

Dynamic model. Because of its simplicity and the relationship (3), we use the following linear dynamic model to capture ellipsoid evolution, coupled with our projection relation (4) for observations:

$$\begin{aligned} X_{k+1} &= A_k^T X_k A_k + B_k \\ Y_k &= C_k^T X_k C_k + W_k, \end{aligned} \quad (20)$$

where the driving matrix B_k and the observation noise matrix W_k are assumed symmetric to maintain symmetry of the matrix state X_k and observation Y_k . In particular, note that the case of 1-dimensional projections corresponds to observations of the squared support function.

Positive semi-definiteness of X_k for all k is assured if X_0 is positive semidefinite and the matrices B_k are also. One possibility is to generate the B_k as the square of a matrix, $B_k = D_k^T D_k$ where D_k is an arbitrary matrix of proper dimensions. This approach will assure that B_k and hence X_k remain PSD, but the entries of X_k are now not simple functions of the elements comprising D_k , as they involve products between terms. In a stochastic setting, then, since the entries of X_k are nonlinear functions of the random variables comprising D_k , simple choices for the distribution of D_k may not translate into simple distributions of X_k .⁵ In contrast, directly choosing the independent entries of B_k as independent random processes yields linear relationships between these quantities, but then the positive semi-definiteness of X_k is difficult to guarantee, though it is likely if the drive is “small,” as discussed in connection with (19). Again, since our primary focus in this work is not statistical, we avoid this difficulty in what follows by simply assuming that B_k is a *known* deterministic PSD driving matrix. We leave statistical developments of the framework for other work, e.g., as in [16, 17] where

⁵ It is interesting to note however that if the elements of D_k are chosen from independent Gaussian distributions the resulting B_k will be Wishart [39], which is well understood (though it seems X_k will still not be in general). We thank a reviewer for bringing this to our attention.

a dynamic equation is coupled with an observation noise model in a statistical setting to generate estimates of cardiac ejection fraction.

Because the form of both the dynamic and observation equation in (20) are identical to (4), we may immediately express them as the following equivalent vector equation using the relationship (11):

$$\begin{aligned} x_{k+1} &= \bar{A}_k x_k + b_k \\ y_k &= \bar{C}_k x_k + w_k, \end{aligned} \quad (21)$$

where the vectors x_k , y_k , b_k , and w_k are the representations of the corresponding matrices in (20) with respect to consistent symmetric basis sets, while \bar{A}_k and \bar{C}_k are matrix representations of the operators $A_k^T(\cdot)A_k$ and $C_k^T(\cdot)C_k$ with respect to the same basis sets. Thus, instead of the direct equations (20), we may equivalently generate the matrices X_k using the standard state space equations of (21). Note again that the entries of b_k and w_k are linear combinations of the respective entries of B_k and W_k , so that, in particular, Gaussian entries in B_k and W_k give rise to Gaussian entries in b_k and w_k .

The form (21) is convenient because of the great amount of existing work on such equations. In particular, the observability and controllability of the underlying ellipsoid X follows immediately from the properties of the matrices \bar{A}_k and \bar{C}_k together with standard results of linear system theory [40]. An example of the tie between such algebraic properties and the geometric properties of the underlying ellipsoid problem will be seen below when we consider a particular class of dynamic matrices \bar{A}_k . Finally, we note that the symmetry in the problem may be exploited to obtain a *square root* algorithm for the evolution of (20), as discussed in [21]. Such algorithms are often used because of their speed and numerical reliability.

Shaping parameters. Let us investigate how the choice of the dynamic matrices in (20) affects the shape of the corresponding ellipsoid. For ease of visualization we consider the planar case, though the same arguments hold in arbitrary dimensions. Geometrically, we may think of changing the ellipsoid by applying a linear transformation to the underlying coordinate system, for example scaling, stretching, and rotation transformations. The relationship (3) then describes how this transformation will affect the underlying ellipsoid matrix. We may apply such transformations over time by simply choosing A_k in the dynamic relationship (20) as the transpose of the desired transformation.

Consider the following family of transformations of the underlying space:

$$\begin{aligned} \tilde{z} &= \begin{bmatrix} \cos(\phi) & -\sin(\phi) \\ \sin(\phi) & \cos(\phi) \end{bmatrix} \begin{bmatrix} \alpha & 0 \\ 0 & 1/\alpha \end{bmatrix} \begin{bmatrix} \tau & 0 \\ 0 & \tau \end{bmatrix} \\ &\quad \times \begin{bmatrix} \cos(\psi) & -\sin(\psi) \\ \sin(\psi) & \cos(\psi) \end{bmatrix} z. \end{aligned} \quad (22)$$

This transformation can capture *any* linear transformation of the plane (to see this note that the SVD of an arbitrary matrix may always be put into this form). It rotates objects by ψ , magnifies them by a factor of τ , stretches them in the z_1 direction by α and shrinks them in the z_2 direction by $1/\alpha$, and then rotates them by the angle ϕ . This transformation was used in a static setting with $\psi = 0$ in [1] to capture a rich class of object profiles. In particular, it was used in the estimation of object size, eccentricity, and orientation. The generalization to higher dimensions is straightforward, with the scaling term replaced by a multiple of the identity, the stretching term becoming a diagonal matrix of determinant 1, and the rotations becoming orthogonal matrices.

The effect on an ellipsoid of applying such a transformation to the coordinate system is given in (3). In particular, we may impose this class of transformations dynamically if we choose our matrices A_k as follows:

$$\begin{aligned} A_k &= \begin{bmatrix} \cos(\psi_k) & \sin(\psi_k) \\ -\sin(\psi_k) & \cos(\psi_k) \end{bmatrix} \begin{bmatrix} \tau_k & 0 \\ 0 & \tau_k \end{bmatrix} \begin{bmatrix} \alpha_k & 0 \\ 0 & 1/\alpha_k \end{bmatrix} \\ &\quad \times \begin{bmatrix} \cos(\phi_k) & \sin(\phi_k) \\ -\sin(\phi_k) & \cos(\phi_k) \end{bmatrix} \end{aligned} \quad (23)$$

Different choices of the parameters of these A_k will result in the application of the corresponding transformations to the ellipse at time point k . For example, suppose we used an A_k with $\alpha_k = 1$, $\tau_k = 1$, $\psi_k = 0$, and $\phi_k = \pi/8$. The resulting ellipse will not change its shape but only rotate by $\pi/8$ radians every step. We show every other step of such a simulation, yielding a tumbling ellipse in Fig. 3.⁶

An interesting subclass of these transformations is obtained by considering only those that preserve the volume of the ellipse. In cell tracking applications such a constraint might reflect incompressibility or conservation of mass of a cell undergoing deformational forces in a particular direction. Since the volume of an ellipsoid is a constant times the square root of the determinant of its defining matrix X [41], any transformation preserving the determinant of X achieves the desired goal. For the class of dynamic matrices under consideration, this restriction

⁶ Note that all ellipses are actually at the origin, but have been separated in the figure for clarity. We will display such evolving ellipses this way in what follows.




FIG. 3. Rotating ellipse.

corresponds to requiring that the scaling term τ be set to 1. An example of such a case with $\alpha_k = \frac{3}{4}$, $\psi_k = 0$, $\phi_k = \pi/8$, and $\tau_k = 1$ is given in Fig. 4, where we have shown every other step of the sequence. This choice of parameters corresponds to compressing along the first coordinate axis by a factor of $\frac{3}{4}$ and stretching along the second coordinate axes by $\frac{4}{3}$ in addition to rotation by $\pi/8$ at each step (the envelope of ellipses traced out in this fashion itself appears to be an ellipse).

6.2. Estimation

In this section we consider the problem of estimating the state of the dynamically evolving ellipsoid (20). First, we recast this problem in terms of the formulation given in (14), obtaining a batch method of solution. Some comments on the connection between geometry and solvabil-




FIG. 4. Deforming ellipse.

ity are made. The inclusion of constraints, including a PSD constraint, for this case is straightforward from our work so far. We then briefly point out the possibility of recursive solutions to the problem.

Batch methods. Certainly one way of solving the problem of estimating the state of the ellipsoid in (20) is to stack up our observations y_k using the form of the equations given in (21), lumping the dynamics and known input into the output matrices and observations, respectively. Doing this operation yields a batch formulation of the problem. In particular, we obtain the following equivalent linear equation for the initial state:

$$\mathbf{y} = \mathbf{C}\mathbf{x}_0 + \mathbf{w}, \quad (24)$$

where the matrices \mathbf{y} and \mathbf{C} are now given by

$$\mathbf{C} = \begin{bmatrix} \tilde{\mathbf{C}}_0 \\ \tilde{\mathbf{C}}_1 \tilde{\mathbf{A}}_0 \\ \tilde{\mathbf{C}}_2 \tilde{\mathbf{A}}_1 \tilde{\mathbf{A}}_0 \\ \vdots \\ \tilde{\mathbf{C}}_r \tilde{\mathbf{A}}_{r-1} \cdots \tilde{\mathbf{A}}_0 \end{bmatrix} \quad (25)$$

$$\mathbf{y} = \begin{bmatrix} y_0 \\ y_1 \\ \vdots \\ y_r \end{bmatrix}$$

$$- \begin{bmatrix} 0 \\ \tilde{\mathbf{C}}_1 b_0 \\ \vdots \\ \tilde{\mathbf{C}}_r (b_{r-1} + \tilde{\mathbf{A}}_{r-1} b_{r-2} + \cdots + (\tilde{\mathbf{A}}_{r-1} \cdots \tilde{\mathbf{A}}_1) b_0) \end{bmatrix}$$

and the observation perturbation vector \mathbf{w} is given by

$$\mathbf{w} = \begin{bmatrix} w_0 \\ w_1 \\ \vdots \\ w_r \end{bmatrix}.$$

To find a unique LS estimate of the initial state \mathbf{x}_0 requires that the matrix \mathbf{C} given in (25) have full column rank. In this case, the unconstrained LS estimate of the initial ellipsoid state is given by $\hat{\mathbf{x}}_0 = \mathbf{C}^+ \mathbf{y}$, where \mathbf{C}^+ is again the Moore-Penrose inverse of \mathbf{C} . The corresponding unconstrained LS estimate at any other time is obtained by using this estimate of the initial state as an initial condition to the equation (20) or (21).

The solution $\mathbf{C}^+ \mathbf{y}$ provides an answer to the unconstrained problem. It is a simple matter to obtain *con-*

strained (e.g., PSD) estimates of the initial state, given our previous development. We need only combine the constrained reconstruction methods discussed in Section 5.2 with the formulation of (24). If the A_k are invertible, as they will be for any reasonable choice of parameters in (23), then constrained reconstructions at other times may be found by simply writing (24) in terms of the state at the desired point in time rather than x_0 . The solutions of such problems are straightforward, given our previous development.

Naturally, stochastic interpretations and recursive solutions of the state estimation problem represented by (20) are also possible, leading in a straightforward way to statistically optimal shape estimates. In particular, recursive solutions to the unconstrained problem obtained through Kalman filtering [42] are simple to implement using the formulation of (21), as is done in [16, 17]. At present, however, we know of no optimal recursive solution to the problem of estimating a *constrained* initial matrix for the system (20). Recall though that if the unconstrained estimate lies within the constraint set, then it must also be the optimal constrained estimate. Of course, we may take the ad hoc approach of projecting the unconstrained estimate onto the constraint set when such an estimate is desired. Such a procedure is suboptimal but may yield reasonable results for many cases.

Observability. The existence of a unique solution to (24) (or a time shifted version of it if our interest is at other than $k = 0$) requires full rank of the matrix C in (25). This matrix will be recognized as the observability matrix of the linear dynamical system given by (20). Thus our rank requirement for solution of (24) is really nothing more than a statement of observability of the corresponding dynamical system (20). Such observability is straightforward to check for a given problem using (25) and often reflects geometric properties of the problem.

For example, suppose our lower-dimensional views are fixed so that $\tilde{C}_k = \tilde{C}$ (nonsquare) and that we choose the dynamic matrices A_k in (23) such that $\phi = \psi = 0$ and $\alpha = 0$, corresponding to uniform shrinking of the ellipsoid with no rotation or stretching. The corresponding matrices \tilde{A}_k of (21) then become $\tilde{A}_k = \tau I$, where I is the identity. The observability matrix C for this case is given by

$$\begin{bmatrix} \tilde{C} \\ \tau \tilde{C} \\ \tau^2 \tilde{C} \\ \vdots \\ \tau^{r-1} \tilde{C} \end{bmatrix}$$

which is clearly rank deficient. Thus we cannot reconstruct an ellipsoid that is uniformly shrinking from a sin-

gle, fixed lower-dimensional viewpoint. The geometrical difficulty is that we get no information about the ellipsoid perpendicular to the projection subspace. The presence of a rotation or stretching term would yield such information, as would changing our view point by allowing C_k to be a rotation matrix.

7. EXPERIMENTS

In this section we present the results of several simulations to demonstrate the developments of the paper. We limit ourselves to planar examples and 1-dimensional (support-type) observations here for ease of visualization, though the demonstrated procedures and methods are valid in arbitrary dimensions. Our goal in these experiments is not to form statistically optimal estimates (which we leave to other works) but simply to demonstrate the potential use of the developments of the paper. In particular, we have not matched our estimators to the perturbations (i.e., noise model) applied to the ideal observations in any of the experiments, but rather taken a deterministic perspective and minimized the sum of the squared residuals to obtain an unweighted LS estimate.

First, we examine the case of reconstructing a static ellipse from a series of its perturbed projections. We compare unconstrained to matrix-interval-constrained (and in particular PSD-constrained) reconstructions, showing how the addition of constraints, reflecting prior knowledge, may aid a reconstruction. In particular, we demonstrate how such algebraic constraints manifest themselves geometrically. Next, we examine the reconstruction of dynamic ellipses. The class of dynamic matrices defined in (23) is used. We show reconstructions demonstrating how knowledge of the evolution of an ellipse (in the form of the dynamic equation in (20)) allows us to reconstruct it from knowledge of only a single spatial projection over time.

Assumptions. In the numerical experiments of this section it is assumed that the center of the ellipse is known and that all data has been translated to the origin, as throughout this section. For convenience all examples use 2-dimensional ellipsoids (ellipses) and 1-dimensional (scalar) projection observations. Recall that our observations are taken to be the *projected ellipses* themselves, as represented by Y_k of (20). As discussed in Section 3, such 1-dimensional ellipse projections are closely related to the support or extent $h(C)$ of the ellipse in the direction C . In particular, our ellipse projections Y_k equal the squared ellipse support length $h^2(C_k)$. Thus we are effectively demonstrating the reconstruction of ellipses from support data in this 1-dimensional case. Though we will take as our observations the values Y_k , in keeping with




FIG. 5. Underlying ellipse.

the model (20), for ease of visualization in what follows we will actually plot the equivalent perturbed support values $\sqrt{Y_k}$, corresponding to the noisy projection extent.

To correspond to a lower-dimensional ellipse observation in the 1-dimensional case, each (scalar) Y_k must be nonnegative. To guarantee this in our experiments, we arbitrarily choose our corresponding observation perturbations W_k from one-sided log-normal distributions [43]. In particular, each ideal observation $C_k^T X_k C_k$ is corrupted by samples of a zero-mean log-normal distribution of a specified variance (which will depend on the experiment). This distribution is chosen for each ideal sample so that its left support point is at $-C_k^T X_k C_k$, thus ensuring that Y_k will be nonnegative for any level of noise variance. We use this model as a convenient way to perturb the ideal samples while keeping them nonnegative and intend no, necessarily practical, statistical interpretation to be drawn from it.

7.1. Static Ellipse Reconstruction


The underlying ellipse used throughout this section is shown in Fig. 5. Its semiaxes are of lengths 1 and 2, with the major axis inclined at $\pi/4$ radians to the first coordinate axis. The corresponding matrix X is given by:

$$X = \frac{1}{2} \begin{bmatrix} 5 & 3 \\ 3 & 5 \end{bmatrix} \quad (26)$$

We will reconstruct this ellipse based on observation of a series of its perturbed projections Y_k . We use 5 equally spaced projections (support samples squared) of the ellipse (26) taken over 0 to π . In particular, the associated observation matrices C_k are given by $C_k = [\cos(\theta_k), \sin(\theta_k)]^T$ where $\theta_k = (k-1)\pi/5$. We examine both a small and large perturbation case.

Small perturbation case. For this experiment we add zero-mean log-normal noise of variance .3 (described above) to the ideal observations $C_k^T X C_k$ defined in (12) to obtain a “signal-to-noise ratio” (SNR) of about 5. The SNR in this case is defined to be $\sqrt{\sum \|C_k^T X C_k\|^2} / (\sigma \sqrt{n})$, where n is the total number of observations and σ is the standard deviation of applied noise. The corresponding perturbed support observations are shown together with the underlying ellipse in Fig. 6. Since the axis lengths are 2 and 1, the level of observation corruption is relatively low in this example.

In Fig. 7 three different reconstructions are shown. In the upper left of the figure the unconstrained LS solution obtain by using (15) is shown. The upper right shows the corresponding PSD-constrained reconstruction. This reconstruction is identical to that obtained without the PSD constraint, showing that for this small perturbation case the PSD constraints were not active. The bottom-right plot shows inner and outer bounds of an interval-constrained reconstruction (the outer bounding ellipse leaves the plot area). The constraints are arbitrary and only meant to show how such constraints may be included. The corresponding reconstruction is shown in the lower left hand corner. This estimate is also identical to the unconstrained one in this instance, again demonstrating that in this slightly perturbed case the unconstrained, linear estimate performs adequately. If we had no directional information as to the orientation of the ellipse we could replace the elliptical bounds by circles, with no such inherent directional bias. Recall that such circular constraints correspond to *eigenvalue* constraints on the reconstructed matrix.

FIG. 6. Perturbed observations and true ellipse (SNR ≈ 5).




FIG. 7. Reconstructions.





FIG. 9. Reconstructions.

Large perturbation case. Now we repeat the experiment, but with a larger perturbation to our ideal observations. This time we add zero-mean log-normal noise of variance 9 to the 5 ideal projection observations, for an SNR of about 1. In Fig. 8 the resulting noisy support data are shown. The amount of perturbation to the ideal observations is quite large this time, being on the order of the semi-axis lengths.

In Fig. 9 the same three reconstructions provided before are shown. In the upper left of the figure the unconstrained LS solution is displayed. The unconstrained estimated matrix for this example is not positive semidefinite, having

FIG. 8. Noisy observations and true ellipse (SNR ≈ 1).

eigenvalues at -2.5 and 19.2 . The corresponding shape really does not make sense to draw and certainly does not resemble an ellipse. The curve that is displayed is actually a hyperbola (a different conic section), obtained because we used the ellipse definition given in (2) to produce these plots. The upper-right plot shows the corresponding PSD-constrained reconstruction. This time the PSD reconstruction yields the degenerate ellipse given by a line, corresponding to the fact that the constrained solution must lie on the boundary of the PSD set and thus possess a zero eigenvalue and semiaxis. While yielding a PSD matrix, the effect is not particularly appealing geometrically. Again, the bottom-right plot shows the inner and outer bounds for an interval constrained reconstruction. The corresponding reconstruction is shown in the lower-left hand corner. This estimate appears to be the best, demonstrating the beneficial effect of accurate prior knowledge.

7.2. Dynamic Ellipse Reconstruction

Here we demonstrate the reconstruction of a *dynamic* ellipse from corrupted lower-dimensional projection data. The generation of such an ellipse was discussed in Section 6.1, where a particular class of dynamic matrices was defined and demonstrated. For this example we have arbitrarily chosen a constant dynamic matrix $A_k = A$ with the parameters defined in Section 6.1 chosen as $\tau = .8$, $\alpha = .9$, $\psi = 0$, and $\phi = \pi/15$. We use a periodic driving term B_k in (20). This drive is shown in Fig. 10 through half of its cycle (recall that all ellipses in the following are actually centered at the origin, but are displayed




FIG. 10. Driving term for the dynamic ellipse.

shifted for clarity). The initial condition for the simulation is given by the matrix X in (26). The corresponding ellipse state for this choice of drive and dynamic matrix is shown for every other time point in Fig. 11. One can imagine such constructions as providing a model for a beating heart, for example. Such a model (actually based on periodic A_k rather than B_k) is used in [16, 17] to estimate ejection fraction from a series of noisy projection images of the heart.

At each time point the ellipse is projected onto the first coordinate axis, yielding a corresponding output matrix in equation (20) of $C_k = [1 \ 0]^T$. Zero-mean log-normal noise of variance 4 is then added to each of the resulting ideal projections $C_k^T X_k C_k$ (support measurements squared) to yield our corrupted observations, as discussed earlier. The average SNR for this case, defined as the average of




FIG. 11. Ellipse state.




FIG. 12. Dynamic reconstruction of ellipses of Fig. 11 together with noisy observations.

the (scalar) observations at all times $C_k^T X_k C_k$ divided by the standard deviation of the noise, is about 1. Using these corrupted observations, the original ellipse is reconstructed using a least squares criterion. This reconstruction is shown in Fig. 12 at every other time point as a solid line. The true ellipses are shown as dotted lines. Also shown in the figure as solid vertical lines are the support measurements corresponding to the noisy observations used in the reconstruction. As can be seen from the figure, despite measurements that are quite corrupted the estimate tracks the dynamic ellipse after about seven time steps. Note that no single frame of data, even if it were ideal, contains enough information to be able to reconstruct the ellipse. The dynamic equation allows us to integrate the data from multiple frames and hence resolve the inherent ambiguity.

In the above example we assume that we have perfect knowledge of the dynamics and drive and the reconstructions seem quite good. For interest, in Fig. 13 we show the effect of adding $N(0, .06)$ noise to the independent elements of our assumed dynamic matrix and drive terms during the reconstruction, resulting in an average error in these matrices of about 10%. This mimics the effect of imperfect knowledge of the dynamics and drive on the reconstruction. The errors in our knowledge of the parameter matrices defining the dynamics is quite substantial, yet the reconstruction still appears quite good, suggesting the robustness of the approach in a stochastic setting. Clearly, any detailed such sensitivity analysis will depend

on the specific problem context and its goals (e.g., see [16, 17] for such a treatment).

8. CONCLUSIONS

In this paper we have examined the problem of reconstructing an ellipsoid from its projections. We presented an approach based on a representation of ellipsoids as elements of the vector space of symmetric matrices. This representation led in turn to a particularly simple relationship between an ellipsoid and its orthogonal silhouette projection. This approach allowed us to simply and precisely characterize the solutions of the associated reconstruction problem. The inclusion of constraints in the form of bounds on the reconstructed ellipse is straightforward in our framework, leading to semidefinite interval constrained symmetric matrix estimation problems. Inclusion of a dynamic element with natural geometric interpretations is similarly easy, leading to the possibility of dynamic shape generation and estimation. Examples of the above points were illustrated through numerical examples. Application of this framework to problems of statistical estimation is straightforward given our development and is left for future work.

APPENDIX: CONSTRAINED RECONSTRUCTION ALGORITHM

Here we present an outline for an algorithm to compute:

$$\min_{X \in \bar{X}} \sum_{i=1}^q \|Y_i - C_i^T X C_i\|_F^2 \quad (27)$$

where

$$\bar{X} = \{X | \bar{X} \geq X \geq \underline{X}\} \quad (28)$$

defines the matrix interval set as described in Section 5. As shown in [21], the set \bar{X} is the intersection of two cones and can be represented as the intersection of an infinite number of halfspaces:

$$\bar{X} = \bigcap_{N \in \text{PSD}^{(1)}} \mathcal{H}\mathcal{S}(N, \langle N, \underline{X} \rangle) \cap \mathcal{H}\mathcal{S}(-N, -\langle N, \bar{X} \rangle) \quad (29)$$

where $\mathcal{H}\mathcal{S}(N, d)$ denotes the closed halfspace $\{X | \langle X, N \rangle \geq d\}$ of normal N and distance d from the origin and $\text{PSD}^{(1)}$ denotes the set of rank 1 PSD matrices. Thus the normals to the bounding halfspaces are obtained from the rank 1 PSD matrices and the distances of these halfspaces from the origin are found from the bounding matrices \underline{X} and \bar{X} . With this notation the algorithm is as follows:




FIG. 13. Reconstruction of dynamic ellipse with imperfect knowledge of dynamics and drive.

ALGORITHM 1 (successive halfspace).

Step 1. Choose an initial approximating set of rank 1 PSD normals $\{N_{i0}\}$. Set $k = 0$.

Step 2. Solve the following to obtain $X(k)$:

$$X(k) = \arg \min_{X \in \bar{X}_k} \sum_{i=1}^q \|Y_i - C_i^T X C_i\|_F^2 \quad (30)$$

where

$$\bar{X}_k = \bigcap_{N \in \{N_{ij}\}_k} \mathcal{H}\mathcal{S}(N, \langle N, \underline{X} \rangle) \cap \mathcal{H}\mathcal{S}(-N, -\langle N, \bar{X} \rangle).$$

Step 3. For the lower bound \underline{X} , find all eigenvalues $\underline{\lambda}_j$ and corresponding eigenvectors \underline{v}_j of $[X(k) - \underline{X}]$ such that $\underline{\lambda}_j \leq 0$.

Step 4. For the upper bound \bar{X} , find all eigenvalues $\bar{\lambda}_j$ and corresponding eigenvectors \bar{v}_j of $[\bar{X} - X(k)]$ such that $\bar{\lambda}_j \leq 0$. If there are no such $\underline{\lambda}_j$ or $\bar{\lambda}_j$ the solution is optimal: STOP.

Step 5. Update the approximating set \bar{X}_k by augmenting the associated set $\{N_{ij}\}_k$ with rank 1 normals corresponding to the forbidden eigenvectors found in Steps 3 and 4: $\{N_{ij}\}_{k+1} = \{N_{ij}\}_k \cup \{\underline{v}_j \underline{v}_j^T\} \cup \{\bar{v}_j \bar{v}_j^T\}$.

Step 6. Set $k = k + 1$. Goto Step 2.

This algorithm iteratively refines its polygonal approximation to the constraint set by adding supporting hyperplanes in the vicinity of the solution. Note that (30) can be formulated as a standard linear inequality constrained least square problem [44] using the vector space formulations of Section 4.

REFERENCES

1. D. J. Rossi, *Reconstruction from Projections Based on Detection and Estimation of Objects*, Ph.D. thesis, Massachusetts Institute of Technology, Aug. 1982.
2. J. A. K. Blokland, A. M. Vossepoel, A. R. Bakker, and E. K. J. Paulwels, Delineating elliptical objects with an application to cardiac scintigrams, *IEEE Trans. Medical Imaging* **MI-6**(1), 1987, 57–66.
3. Y. Bresler and A. Macovski, Three-dimensional reconstruction from projections with incomplete and noisy data by object estimation, *IEEE Trans. Acoustic, Speech, and Signal Process.* **ASSP-35**, 1987, 1139–1152.
4. J. A. Fessler and A. Macovski, Object-based 3-D reconstruction of arterial trees from magnetic resonance angiograms, *IEEE Trans. Medical Imaging*, 10, Mar. 1991, 25–39.
5. P. Lipson, A. L. Yuille, D. O'Keefe, J. Cavanaugh, J. Taaffe, and D. Rosenthal, Automated bone density calculation using feature extraction by deformable templates, in *Proceedings, First IEEE Conference on Visualization in Biomedical Computing*, Atlanta, GA, May 22–25 1990, pp. 477–484.
6. F. C. Schweppe and H. K. Knudsen, The theory of amorphous cloud trajectory prediction, *IEEE Trans. Inform. Theory* **IT-14**, 1968, 415–427.
7. F. C. Schweppe, *Uncertain Dynamic Systems*, Prentice-Hall, Englewood Cliffs, NJ, 1973.
8. O. E. Drummond, S. S. Blackman, and K. C. Hell, Multiple sensor tracking of clusters and extended objects, in *Proceedings, 1988 Tri-Service Data Fusion Symposium*, Laurel, MD, May 19, 1988.
9. P. C. Gaston and R. Lozano-Perez, Tactile recognition and localization using object models: The case of polyhedra on the plane, *IEEE J. Pattern Anal. Machine Intell.* **6**(3), 1984, 257–266.
10. J. L. Schneider and T. B. Sheridan, An automated tactile sensing strategy for planar object recognition and localization, *IEEE J. Pattern Anal. Machine Intell.* **12**, Aug. 1990, 775–786.
11. J. L. Prince and A. S. Willsky, Convex set reconstruction using prior shape information, *CVGIP: Graphical Models Image Process.* **53**, Sept. 1991, 413–427.
12. J. L. Prince and A. S. Willsky, Reconstructing convex sets from support line measurements, *IEEE J. Pattern Analysis Machine Intell.* **12**, Apr. 1990, 377–389.
13. H. Stark and H. Peng, Shape estimation in computer tomography from minimal data, in *Pattern Recognition and Artificial Intelligence: Towards and Integration* (E. S. Gelsema and L. N. Kanal, Eds.), *Machine Intelligence and Pattern Recognition*, Vol. 7 pp. 185–200, North-Holland, Amsterdam, 1988.
14. D. J. Rossi and A. S. Willsky, Reconstruction from projections based on detection and estimation of objects—Parts I and II: Performance analysis and robustness analysis, *IEEE Trans. Acoustic, Speech, Signal Process.* **ASSP-32**(4) 1984, 886–906.
15. A. S. Lele, S. R. Kulkarni, and A. S. Willsky, Convex-polygon estimation from support-line measurements and applications to target reconstruction from laser-radar data, *J. Opt. Soc. Am. A*, **9**, Oct. 1992, 1693–1714.
16. S. Jaggi, W. C. Karl, and A. S. Willsky, Dynamic estimation of left-ventricular ejection fraction, in *Proceedings of the 1993 SPIE/SPSE Symposium on Electronic Imaging*, 1993. Special Session on Cardiac Image Processing.
17. S. Jaggi, W. C. Karl, and A. S. Willsky, Estimation of dynamically evolving ellipsoids with applications to medical imaging. To appear *IEEE Trans. Med. Img.*
18. J. L. Prince, *Geometric Model-Based Estimation from Projections*, Ph.D. thesis, Massachusetts Institute of Technology, Jan. 1988.
19. T. Bonnesen and W. Fenchel, *Theory of Convex Bodies*, BCS Associates, Moscow, ID, 1987.
20. W. C. Karl and G. C. Verghese, Curvatures of surfaces and their shadows, *Linear Algebra and its Appl.* **130**, 1990, 231–255.
21. W. C. Karl, *Reconstructing Objects from Projections*, Ph.D. thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, Jan. 1991.
22. B. K. P. Horn, *Robot Vision*, MIT Press, Cambridge, MA, 1986.
23. P. L. Van Hove, *Silhouette-Slice Theorems*, Ph.D. thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, September 1986.
24. A. S. Lele, Convex set reconstruction from support line measurements and its application to laser radar data, Master's thesis, Massachusetts Institute of Technology, April 1990.
25. A. Graham, *Kronecker Products and Matrix Calculus: With Applications*, Ellis Horwood Series in Mathematics and Its Applications, Halsted Press, Chichester, 1981.
26. J. R. Magnus and H. Neudecker, *Martix Differential Calculus, with Applications in Statistics and Econometrics*, Wiley New York, 1988.
27. J. R. Magnus and H. Neudecker, The elimination matrix: Some lemmas and applications, *SIAM J. Algebraic Discrete Meth.* **1**(4), 1980, 422–449.
28. W. J. Vetter, Vector structures and solutions of linear matrix equations, *Linear Algebra Appl.* **10**, 1975, 181–188.
29. F. J. H. Don, On the symmetric solutions of a linear matrix equation, *Linear Algebra Appl.* **93**, 1987, 1–6.
30. S. Tsuji and F. Matsumoto, Detection of ellipses by a modified hough transform, *IEEE Trans. Comput.* **C-27**(8), 1978, 777–781.
31. D. H. Ballard, Generalizing the hough transform to detect arbitrary shapes, *Pattern Recognit.* **13**(2), 1981, 111–112.
32. H. K. Yuen, J. Illingworth, and J. Kittler, Detection partially occluded ellipses using the hough transform, *Image Vision Comput.* **7**, Feb. 1989, 31–37.
33. F. L. Bookstein, Fitting conic sections to scattered data, *Comput. Graphics Image Process.* **9**, 1979, 56–71.
34. T. Nagata, H. Tamura, and K. Ishibashi, Detection of an ellipse by use of a recursive least-squares estimator, *J. Robotic Systems* **2**(2), 1985, 163–177.
35. B. B. Chaudhuri and G. P. Samanta, Elliptic fit of objects in two and three dimensions by moment of inertia optimization, *Pattern Recognit. Lett.* **12**, Jan. 1991, 1–7.
36. R. Safaee-Rad, K. C. Smith, B. Benhabib, and I. Tchoukanov, Application of moment and fourier descriptors to the accurate estimation of elliptical-shape parameters, *Pattern Recognit. Lett.* **13**, July 1992, 497–508.
37. G. H. Golub and C. F. V. Loan, *Matrix Computations*, Johns Hopkins Univ. Press, Baltimore, 1989.
38. S. Jaggi, *Estimation of Dynamically Evolving Ellipsoids with Applications to Cardiac Imaging*, Master's thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, Sept. 1992.

39. R. J. Muirhead, *Aspects of Multivariate Statistical Theory*, Wiley, New York 1982.
40. C. T. Chen, *Introduction to Linear System Theory*. Holt, Rinehart, & Winston, New York, 1970.
41. H. Guggenheimer, *Applicable Geometry: Global and Local Convexity*. Applied Mathematics Series, Krieger, Huntington, NY, 1977.
42. A. Gelb, (Ed.), *Applied Optimal Estimation*. MIT Press, Cambridge, MA, 1984.
43. A. Drake (Ed.), *Fundamentals of Applied Probability Theory*, McGraw-Hill, New York, 1967.
44. C. L. Lawson and R. J. Hanson, *Solving Least Squares Problems*. Prentice-Hall, Englewood Cliffs, NJ, 1974.