

Efficient Distance Computation Using Best Ellipsoid Fit

Elon Rimon

Robotics Laboratory
Stanford University, Department of Computer Science

Stephen P. Boyd

Information Systems Laboratory
Stanford University, Department of Electrical Engineering

Abstract

Knowledge of the distance between a robot and its surrounding environment is vital for any robotic system. The robot must obtain this information rapidly in order to plan and react in real-time. Our technique first surrounds the robot links and the obstacles by optimal ellipsoids, and then computes the clearance of the links from the obstacles with a generalized distance function. This approach offers an attractive alternative to the widely used technique of computing the distance via polyhedral representation of the robot and the obstacles. In particular, our approach offers a drastic reduction in the complexity of the data structures: each polyhedron, typically represented by a list of its features and their adjacency graph; is replaced by a minimum-volume ellipsoid represented by its center and a symmetric matrix whose dimension is either two or three (the workspace dimension). Moreover, while the computation time of the distance between polyhedra is often a function of their geometrical complexity, computation time in the ellipsoidal case is essentially *constant*; and becomes even more rapid when it is computed repeatedly along the robot's trajectory.

Our method consists of the following two algorithms: The first computes the optimal ellipsoid surrounding a convex polyhedron. The second is an analytic formula for the *free margin* about one ellipsoid with respect to another, that is computed as a standard eigenvalue problem. An efficient incremental version of the latter algorithm is then proposed. This system has been implemented and preliminary simulation results are provided throughout the paper.

1 Introduction

The technique of bounding sets with minimum-volume ellipsoids seems to be applicable in other areas, such as pattern recognition and machine vision. The main concern of this paper, however, is to demonstrate the effectiveness of the ellipsoid representation for geometrical reasoning in the context of robotics. Specifically, this paper is concerned with the robot collision-detection problem, that consists of computing a quantity that reflects, as a function of the geometrical data, the amount of clearance between the robot and its environment. Knowledge of this distance is of central importance for planning collision-free paths [8], and its rapid computation is essential in the low-level control, where the gradient vector-field of the distance is used to guard the robot from collision [11]. A similar need also arises in many computer graphics applications, especially in physical simulations [22].

Unfortunately, the distance depends on the various geometrical features of the robot and the obstacles; and this introduces a major computational bottleneck. In many practical applications, however, the robot links tend to be elongated objects that can be effectively surrounded by ellipsoids. Supposing that the obstacles are described by union of convex polyhedra — there are efficient algorithms to decompose a polyhedron into union of convex polyhedra — we surround each convex polyhedron by an ellipsoid as well. The problem of computing the distance between polyhedral links and obstacles is thus replaced by the problem of computing the distance between pairs of ellipsoids. We shall see that a function related to the distance, the *free margin* about an ellipsoid, can be rapidly computed

in constant time, independent of the geometrical complexity of the original polyhedra.

More generally, this work is part of a larger program of research, whose purpose is to develop a geometric modeling system based on a catalogue of shapes expressed as Boolean combinations of linear and quadratic inequalities. Such a catalogue would be able to approximate all possible shapes and, unlike the purely polyhedral representation, seems to include shapes whose boundary is smooth (continuous normal). The catalogue must come with a set of operations that render it useful for robotic applications. To mention some of the most important ones: automatic construction of the shapes from sensory data, rapid collision detection, and the computation of a measure of the distance between pairs of shapes.

In the context of the latter problem, it can be easily verified that computing the Euclidean distance between two shapes described by intersection of linear and quadratic polynomials, such that each quadratic polynomial describes a convex region, is a *convex optimization problem* (see e.g., [2]). Being such, the Euclidean distance can be effectively computed up to desired accuracy in time linear in the number of geometrical features. In robotics, however, efficient iterative methods are not good enough. Most of the reactive controllers use the gradient vector-field (or subgradient when it is not differentiable) of the distance to guide the robot. It is therefore advantageous to obtain a closed-form expression for the gradient. Moreover, there is a basic need to explicitly parametrize the location of the configuration-space “obstacles” — the forbidden regions in the robot's configuration space — in terms of the geometrical data i.e., the inside-outside relation between shapes parametrized by the shapes' geometrical data. To the best of our knowledge the free margin function constitutes, for the first time in the context of robotics, an analytic formula for ellipsoids. The formula has the form of an eigenvalue problem, and we also give a closed-form formula for the gradient vector-field.

1.1 Organization of the Paper

This section continues with a brief account of the related literature. In Section 2 we describe an efficient algorithm for the minimal-volume ellipsoid surrounding a convex polyhedron. This also turns out to be a convex optimization problem, for which efficient ϵ -accurate algorithms that require time proportional to $\log(1/\epsilon)$ and linear in the number of vertices are known. We shall use the standard *ellipsoid algorithm*¹, whose features, as well as the class of convex optimization problems, are briefly described in Appendix A. The ellipsoid algorithm, although simple and efficient, is not the only known algorithm for solving convex optimization problems. In fact, interior-point algorithms recently developed by Nesterov and Nemirovsky [17] are much faster and may be a better choice for high-dimensional versions of this problem.

A closed-form formula for the free margin about one ellipsoid with respect to another is presented in Section 3. Specifically, let N be the dimension of the ambient space ($N = 2$ or 3 in our case), and let $\mathcal{E}(x_i, X)$ be the ellipsoid with center x_i and shape described by a positive-definite symmetric matrix X (a condition written as

$X > 0$),

$$\mathcal{E}(x_i, X) = \{x \in \mathbb{R}^N : (x - x_i)^T X (x - x_i) \leq 1\}.$$

Let the two ellipsoids be \mathcal{E}_1 and \mathcal{E}_2 . First, a formula for the point x^* in \mathcal{E}_2 at which the ellipsoidal level-surfaces surrounding \mathcal{E}_1 touch \mathcal{E}_2 for the first time is computed. Then the *free margin* about \mathcal{E}_1 with respect to \mathcal{E}_2 is computed in terms of x^* such that

$$\text{margin}(\mathcal{E}_1, \mathcal{E}_2) \begin{cases} < 0 & \text{iff } \mathcal{E}_1 \text{ overlaps } \mathcal{E}_2; \\ = 0 & \text{iff } \mathcal{E}_1 \text{ touches } \mathcal{E}_2; \\ > 0 & \text{otherwise.} \end{cases} \quad (1)$$

Geometrically, $\text{margin}(\mathcal{E}_1, \mathcal{E}_2)$ is the (signed) distance between \mathcal{E}_1 and \mathcal{E}_2 as determined by the metric associated with the matrix of \mathcal{E}_1 . When \mathcal{E}_1 is moving along a trajectory it becomes a function of its current configuration — \mathcal{E}_1 's position and orientation. Under this interpretation the condition $\text{margin}(\mathcal{E}_1, \mathcal{E}_2) \leq 0$ characterizes the configuration-space obstacle due to \mathcal{E}_2 i.e., configurations of \mathcal{E}_1 that involve intersection with \mathcal{E}_2 .

The formula for x^* involves the minimal eigenvalue of a $2N \times 2N$ matrix whose entries are expressed in terms of the geometrical data. The standard *QR method* is then used to find the minimal eigenvalue. In Section 4 we describe how to accelerate the computation along the robot trajectory by exploiting the previous computation. As long as the trajectory points are sufficiently close by, the minimal eigenvalue is computed with the faster *inverse iteration method*. We will make precise the notion of “sufficiently close by”, and will show in the process that the free-margin function is an *analytic function* of the geometrical data. This last result makes $\text{margin}(\mathcal{E}_1, \mathcal{E}_2)$ very attractive in actual implementations. In fact, we shall present a closed-form expression for its gradient vector-field, which is readily computable in terms of the geometrical data.

1.2 Related Literature

The topic of surrounding complicated shapes by simpler ones is considered in the computational geometry literature. For example, [16] discusses the problem of surrounding a polyhedron by minimum-volume box. Although such a box is an attractive alternative for the minimum-volume ellipsoid, it is currently not clear which approach is more effective. In fact, both approaches require the same number of parameters to represent their shape, and they seem to complement each other as effective geometrical approximation. The selection of the most effective approximating shape is the topic of research now in progress.

The appeal of ellipsoids as effective means for shape representation has been recognized in the machine vision literature for quiet some time (see e.g., [18]). In the context of computing the ellipsoidal approximation, Post [19] has proposed an exact algorithm that computes the minimum-volume ellipsoid in time proportional to m^2 independent of the ambient dimension N , where m is the number of vertices. Our algorithm uses standard convex programming techniques and solves the problem within ϵ accuracy in time $mp(N) \log(1/\epsilon)$, where $p(N)$ is a polynomial function which is a constant in our case.

The topic of closed-form formulas for the forbidden regions in configuration-space is relatively unexplored. It seems that general algebraic decision-methods can compute such formulas for polynomially bounded shapes (using, for instance, the multivariate resultant [4]), but we are not aware of any practical implementation of them. Specific closed-form formulas are known for the following cases: a polyhedron moving in the presence of polyhedral obstacles [12, Chap 3]; a convex rigid-body moving with fixed orientation amidst convex obstacles [1]; and specific planar articulated chains [5]. This paper presents such a formula for the ellipsoidal case, where an ellipsoidal object is moving in the presence of ellipsoidal obstacles. Characterization of the intersection between general quadratic shapes is also discussed in the computer graphics literature (see e.g., [14]).

¹The fact that the topic of this paper is ellipsoids and that convex optimization problems are solved by an algorithm based on n -ellipsoids is coincidental.

Computation of the distance between polyhedral shapes has long history in robotics (see e.g., [2, 8, 9, 15]), and the technique presented here complements these results for ellipsoidal shapes. In particular, an algorithm recently proposed by Lin and Canny for polyhedra [15], computes the distance between two moving convex polyhedra by tracking the closest two features. The computational effort of their algorithm is essentially constant, except at instances where the identity of the closest two features changes. At these singular events the new closest two features are found in time roughly linear in the number of geometrical features. In contrast, the computational effort of our ellipsoid approach is always constant and does not require the substantial bookkeeping required to manage the polyhedral features. Further, the free-margin function for ellipsoids is an *analytic function* of the geometrical data (refer to Corollary 4.1), while the polyhedral distance is not even differentiable. Of course, these gains come on the expense of using approximate shapes and a “distance” function that is not the Euclidean distance.

In relation to rapid distance computation, some researchers have suggested to trade computation with memory, by computing the distance beforehand for all possible configurations on a suitably discretized configuration space (see e.g., [13]). Note that this computation must use the aforementioned inside-outside functions to be effective. Unfortunately, the required memory grows exponentially with the dimension of the configuration space, and this approach becomes impractical for more than few degrees of freedom. For low-dimensional stationary environments, however, it offers very rapid (discretized) distance computation that can be made to be independent of the geometrical complexity of the environment.

2 The Optimal Ellipsoid

It was shown quite some time ago that for any compact set (closed and bounded) with non-empty interior \mathcal{P} there exists a unique ellipsoid \mathcal{E} of minimal volume containing it [10]. This ellipsoid is called the Löwner-John ellipsoid of \mathcal{P} , or simply the **L-J ellipsoid**. Of course, the L-J ellipsoid contains the convex hull of \mathcal{P} and for this reason we shall restrict our attention to convex sets \mathcal{P} . The L-J ellipsoid has a remarkable property that \mathcal{P} contains the ellipsoid obtained from \mathcal{E} by shrinking it from its center by a factor equal to the dimen-




Figure 1. The N -dimensional L-J ellipsoid is obtained by intersecting the $(N + 1)$ -dimensional L-J ellipsoid centered at the origin with the plane at height $x_{N+1} = 1$.

sion of the ambient space. This establishes an upper bound on the distance of the surface of \mathcal{E} from \mathcal{P} , and consequently indicates that the L-J ellipsoid is always an intuitively acceptable approximation.

We describe now an efficient ϵ -accuracy algorithm for what we shall call the **L-J problem** — to compute the L-J ellipsoid containing a given convex polyhedron \mathcal{P} . We do this in two steps. First, following [17], the L-J problem is shown to be a *convex optimization* problem. Then the standard *ellipsoid algorithm* used to solve such problems is described in the context of our problem. Appendix A contains a short account of the convex optimization problems and the ellipsoid algorithm.

Let N be the dimension of the ambient space ($N = 2$ or 3 in our case), and denote its coordinates by (x_1, \dots, x_N) . Also, let \mathcal{E}_0^{N+1} be an $(N+1)$ -dimensional ellipsoid centered at the origin of \mathbb{R}^{N+1} . The idea is to embed \mathcal{P} in a space of dimension $N+1$, in the plane at height $x_{N+1} = 1$, and then to compute the minimum-volume \mathcal{E}_0^{N+1} containing the embedded \mathcal{P} . It turns out that the resulting \mathcal{E}_0^{N+1} determines the N -dimensional L-J ellipsoid by simply intersecting \mathcal{E}_0^{N+1} with the plane $x_{N+1} = 1$, as shown in Figure 1. This fact is mentioned in [17, pp 229].

The latter problem — of computing the minimum-volume \mathcal{E}_0^{N+1} containing \mathcal{P} — is a convex optimization problem. To see this, consider the volume of an $(N+1)$ -ellipsoid, $\mathcal{E}(x, X)$, given by

$$\text{volume}(\mathcal{E}) = \frac{\beta_{N+1}}{\sqrt{\det X}}, \quad (2)$$

where X is \mathcal{E} 's $(N+1) \times (N+1)$ symmetric positive-definite matrix, and β_{N+1} is the volume of the unit ball in \mathbb{R}^{N+1} , but we will not need β_{N+1} . Since (2) is equivalent to the equation

$$\log(\text{volume}(\mathcal{E})) = \log \beta_{N+1} - \frac{1}{2} \log(\det X),$$

the L-J problem is equivalent to minimizing $-\log(\det X)$ subject to the constraints that X be symmetric positive-definite, and that \mathcal{E}_0^{N+1} contain the polyhedron \mathcal{P} embedded at height $x_{N+1} = 1$. In general, an ellipsoid contains a convex polyhedron if and only if it contains its vertices, v_1, \dots, v_m . Thus the L-J problem becomes

$$\min\{-\log(\det X)\} \quad (3)$$

subject to

$$X > 0 \quad \text{and} \quad \left(\begin{array}{c} v_i^T \\ 1 \end{array} \right) X \left(\begin{array}{c} v_i \\ 1 \end{array} \right) \leq 1 \quad \text{for } i = 1, \dots, m. \quad (4)$$

Let the optimization variables be the distinct entries of the symmetric matrix X , and let M be the number of these entries, $M = \frac{1}{2}(N+1)(N+2)$. The problem (3)-(4) is convex if the objective function (3) and the constraints in (4) are *convex functions* in terms of the entries of X . Indeed, it is shown in [21] that $-\log(\det X)$ is convex in the region $X > 0$. The constraint $X > 0$ in (4) can be written as

$$\lambda_{\max}(-X) < 0 \quad (5)$$

($\lambda_{\max}(-X)$ is the largest eigenvalue of the matrix $-X$), and is shown in [21] to be convex. Similarly, the constraint $\left(\begin{array}{c} v_i^T \\ 1 \end{array} \right) X \left(\begin{array}{c} v_i \\ 1 \end{array} \right) \leq 1$ is a linear inequality in the entries of X and is also convex. Having shown that (3)-(4) is a convex optimization problem, we can now apply the ellipsoid algorithm.

Given a convex polyhedron $\mathcal{P} \subset \mathbb{R}^N$ described by its vertices v_1, \dots, v_m , the ellipsoid algorithm computes a matrix X that minimizes (3) up to ϵ accuracy,

$$0 \leq (-\log(\det X)) - (-\log(\det X^*)) \leq \epsilon,$$

where X^* is the true minimum. It is shown in [21] that the N -ellipsoid obtained by intersecting the resulting ϵ -optimal $(N+1)$ -ellipsoid with the plane $x_{N+1} = 1$ is also ϵ -optimal. Hence the original L-J problem is solved within a specified relative error of $\epsilon^{-\epsilon}$. Let us describe now the algorithm and some of its details.

First, we will need to compute subgradients of $\lambda_{\max}(-X)$. Let $\mathcal{E}(0, X)$ be an $(N+1)$ -ellipsoid. Using the definition of subgradient given in Appendix A, the subgradient of $\lambda_{\max}(-X)$ is a matrix G satisfying the inequality

$$\lambda_{\max}(-Z) \geq \lambda_{\max}(-X) + \text{tr}(G^T(Z - X)),$$

for all symmetric matrices Z (tr denotes the trace). Let v be a unit-magnitude eigenvector of X corresponding to its maximal eigenvalue. It can be easily verified that the inequality

$$\lambda_{\max}(-Z) \geq v^T[-Z]v = v^T[-X]v - v^T(Z - X)v,$$

valid for all symmetric matrices Z , implies that the desired G is simply the symmetric matrix

$$G = -vv^T.$$

In the following, $(X_k, A_k) \in (\mathbb{R}^M, \mathbb{R}^{M \times M})$ is the center and matrix of the k^{th} ellipsoid in the ellipsoid algorithm and, for simplicity, X_k also represents the corresponding matrix of the $(N+1)$ -ellipsoid $\mathcal{E}(0, X_k)$. For simplicity, as well, we replace $\lambda_{\max}(-X_k)$ by the equivalent expression $\lambda_{\min}(X_k)$. Last, we will need to use the stack notation: if $A \in \mathbb{R}^{n \times n}$, then A^s denotes the $n^2 \times 1$ vector obtained by stacking the columns of A over each other. More details about the algorithm can be found in Appendix A and in [3].

$X_1, A_1 \leftarrow$ an initial M -ellipsoid that contains the minimum;
 $k \leftarrow 0$;
 repeat {
 $k \leftarrow k + 1$;

compute $\lambda_{\min}(X_k)$ and $(v_i^T, 1)X_k \left(\begin{array}{c} v_i \\ 1 \end{array} \right)$ for $i = 1, \dots, m$;

if $(\lambda_{\min}(X_k) \leq 0)$ { /* X_k is not positive definite */

compute eigenvector v for $\lambda_{\min}(X_k)$;

$h_k = (-vv^T)^s$; /* compute a subgradient */

$\tilde{g} \leftarrow h_k / \sqrt{h_k^T A_k h_k}$;

} else {

if $((v_i^T, 1)X_k \left(\begin{array}{c} v_i \\ 1 \end{array} \right) > 1$ for some $i = i_0$) { /* X_k is infeasible */

$h_k = \left(\left(\begin{array}{c} v_{i_0} \\ 1 \end{array} \right) (v_{i_0}^T, 1) \right)^s$; /* compute a subgradient */

$\tilde{g} \leftarrow h_k / \sqrt{h_k^T A_k h_k}$;

} else { /* X_k is feasible */

$g_k = \nabla(-\log(\det X_k)) = (-X_k^{-1})^s$; /* compute a gradient */

$\tilde{g} \leftarrow g_k / \sqrt{g_k^T A_k g_k}$;

}

$X_{k+1} \leftarrow X_k - \frac{1}{n+1} A_k \tilde{g}$;

$A_{k+1} \leftarrow \frac{n^2}{n^2-1} \left(A_k - \frac{2}{n+1} A_k \tilde{g} \tilde{g}^T A_k \right)$;

} until $(\lambda_{\min}(X_k) > 0)$ and (no i_0 exists) and $(\sqrt{g_k^T A_k g_k} \leq \epsilon)$.

It is shown in [21] that the initial X_1, A_1 can be conveniently chosen in terms of the radii of two balls in \mathbb{R}^N , one that is contained in the polyhedron \mathcal{P} and one that contains it.




Figure 2. The L-J ellipsoid of two convex polygons

The ellipsoid algorithm computes the L-J ellipsoid to a relative accuracy of $e^{-\epsilon}$ in $mp(M) \log(\frac{1}{\epsilon})$ steps, where $p(M)$ is a constant polynomial (see Appendix A). We have implemented a two-dimensional ($N = 2$ hence $M = 6$) version of this algorithm on a DEC5000 machine. A typical computation takes 600 iterations and runs for about 2 seconds. Two numerical examples are shown in Figure 2.

3 The Free Margin Function

Given two ellipsoids in \mathbb{R}^N , $\mathcal{E}_1 = \mathcal{E}(x_1, P)$ and $\mathcal{E}_2 = \mathcal{E}(x_2, Q)$, we would like to find the point x^* in \mathcal{E}_2 such that

$$(x^* - x_1)^T P(x^* - x_1) \leq (x - x_1)^T P(x - x_1), \quad (6)$$

for all $x \in \mathcal{E}_2$. Geometrically, x^* is the point in \mathcal{E}_2 that is the closest to \mathcal{E}_1 with respect to a distance function whose equidistance level-sets are the ellipsoidal surfaces surrounding \mathcal{E}_1 . We consequently define the *free-margin* of \mathcal{E}_1 about \mathcal{E}_2 is defined to be,

$$\text{margin}(\mathcal{E}_1, \mathcal{E}_2) \triangleq (x^* - x_1)^T P(x^* - x_1) - 1$$

(\triangleq denotes a definition).

Clearly, $\text{margin}(\mathcal{E}_1, \mathcal{E}_2)$ is positive when \mathcal{E}_1 and \mathcal{E}_2 are disjoint, zero when they touch, and negative when their interiors overlap. Note, however, that $\text{margin}(\mathcal{E}_1, \mathcal{E}_2)$ is not symmetric i.e.,

$$\text{margin}(\mathcal{E}_1, \mathcal{E}_2) \neq \text{margin}(\mathcal{E}_2, \mathcal{E}_1),$$

and that it resembles the actual Euclidean distance only when \mathcal{E}_1 and \mathcal{E}_2 are close to each other. Ideally, one would like to compute the Euclidean distance, but we are not aware of any closed-form formula for it. In contrast, the computation of the free-margin function turns out to be equivalent to an eigenvalue problem, that can be solved by traditional methods.

First let us apply a coordinate transformation that will make \mathcal{E}_2 look like a unit ball,

$$\bar{x} \triangleq Q^{\frac{1}{2}}(x - x_2) \quad \text{or} \quad x = Q^{-\frac{1}{2}}\bar{x} + x_2. \quad (7)$$

In the new coordinates our problem becomes:

$$\min\{(\bar{x} - c)^T C(\bar{x} - c)\} \quad \text{such that} \quad \|\bar{x}\|^2 \leq 1, \quad (8)$$

where

$$C \triangleq Q^{-\frac{1}{2}} P Q^{-\frac{1}{2}} \quad \text{and} \quad c \triangleq Q^{\frac{1}{2}}(x_1 - x_2)$$

(C is positive definite). The ellipsoid \mathcal{E}_2 has thus become a unit ball. For our purposes we may assume that the center c is always outside the unit ball. This implies that the quadratic polynomial in (8) must attain its minimum on the boundary of the unit ball, where $\|\bar{x}\|^2 = 1$. For simplicity, we shall hereafter replace \bar{x} by x .

Using Lagrange multiplier, a necessary condition for x^* to be a solution of (8) is

$$\lambda x^* = C(x^* - c) \quad \text{for some scalar } \lambda; \quad (9)$$

or, equivalently,

$$x^* = [C - \lambda I]^{-1} b,$$

where

$$b \triangleq C c,$$

and I is the $N \times N$ identity matrix. Substituting x^* into the constraint $\|x\|^2 = 1$,

$$b^T [C - \lambda I]^{-2} b = 1, \quad (10)$$

yields a $2N$ -degree polynomial in λ .

It turns out that the minimal (real) root of the polynomial (10) solves the problem. This fact is evident from the following identity:

Theorem 1 ([6]) If (x_1, λ_1) and (x_2, λ_2) are solutions of (9)-(10), then

$$\begin{aligned} (x_2 - c)^T C(x_2 - c) - (x_1 - c)^T C(x_1 - c) \\ = \frac{\lambda_2 - \lambda_1}{2} \|x_1 - x_2\|^2. \end{aligned}$$

Thus the problem is solved if we can compute the minimal (real) root of the polynomial (10).

Using a method developed by Gander, Golub, and Matt [7], the problem is transformed into an eigenvalue problem via the introduction of two new variables, $y \in \mathbb{R}^N$ and $z \in \mathbb{R}^N$, as follows

$$y \triangleq [C - \lambda I]^{-2} b \quad \text{and} \quad z \triangleq [C - \lambda I]^{-1} b. \quad (11)$$

Expressing equations (10) and (11) in terms of (y, z) and λ , yields the following system of equations

$$\begin{aligned} b^T y &= 1 \\ [C - \lambda I] z &= b \\ [C - \lambda I] y - z &= 0. \end{aligned} \quad (12)$$

Substituting $b^T y = 1$ into the right side of the second equation yields,

$$[C - \lambda I] z = [bb^T] y.$$

Combining the third equation in (12) with the last equation yields,

$$\begin{bmatrix} C & -I \\ -bb^T & C \end{bmatrix} \begin{pmatrix} y \\ z \end{pmatrix} = \lambda \begin{pmatrix} y \\ z \end{pmatrix}, \quad (13)$$

a standard eigenvalue problem. Let λ^* be the minimal (real) eigenvalue of (13). It is shown in [21] that λ^* is exactly the minimal root of the $2N$ -degree polynomial (10). It follows that the point x^* , and consequently $\text{margin}(\mathcal{E}_1, \mathcal{E}_2)$, can be computed in terms of λ^* as follows:

Theorem 2 Given two ellipsoids $\mathcal{E}_1 = \mathcal{E}(x_1, P)$ and $\mathcal{E}_2 = \mathcal{E}(x_2, Q)$, let λ^* be the minimal eigenvalue of (13). Then the point $x^* \in \mathcal{E}_2$ is given by

$$x^* = [C - \lambda^* I]^{-1} b, \quad (14)$$

where $C = Q^{-\frac{1}{2}} P Q^{-\frac{1}{2}}$ and $b = Q^{-\frac{1}{2}} P(x_1 - x_2)$; and the free-margin function, $\text{margin}(\mathcal{E}_1, \mathcal{E}_2)$, satisfies equation (1) above.

Geometrically, $\text{margin}(\mathcal{E}_1, \mathcal{E}_2)$ determines the smallest ellipsoid with center x_1 and matrix P that touches \mathcal{E}_2 .

Two planar ($N = 2$) examples are shown in Figure 3. We have computed the minimal eigenvalue using the *QR algorithm*. This algorithm requires roughly $4(2N)^3$ operations [21]. The QR algorithm was used without exploiting the specifics of our matrix, and the average time for one distance computation was 2.5 msec (on a DEC5000 machine). In the next section we describe how to track only the minimal eigenvalue, and consequently achieve a considerable efficiency gain.




Figure 3. The closest point, x^* , according to the generalized distance determined by the ellipse surrounding the robot link

4 Incremental Computation of the Generalized Distance

We have seen that the computation of the generalized distance with the QR method is efficient. It computes, however, all the eigenvalues of the matrix

$$M \triangleq \begin{bmatrix} C & -I \\ -bb^T & C \end{bmatrix},$$

while only the *minimal* eigenvalue is needed.

In robotics, as well as in computer-graphics animation, the distance is typically computed along a trajectory i.e., the matrix M becomes $M(x(k))$, where $x(k)$ is the robot's k^{th} configuration. The computation time can be substantially reduced by tracking only the minimal eigenvalue along the trajectory. The *Inverse Iteration method* [20, pp 394] is suitable for this task. It is initialized with an estimate for the minimal eigenvalue, $\hat{\lambda}$, and for the corresponding eigenvector, \hat{v} ; and works as follows:

```

x(0) ← v̂;
k ← 0;
A ← [M - λ̂I]-1;
repeat {
k ← k + 1;
x(k) ← Ax(k - 1);
normalize x(k);
} until ( ||x(k) - x(k - 1)|| ≤ ε )
λ* = λ̂ + ||x(k - 1)||2 / (x(k) · x(k - 1)).

```

The idea behind this method is simple. Let λ^* and v^* be the minimal eigenvalue and the corresponding eigenvector of M . For any value of $\hat{\lambda}$, v^* is also an eigenvector of $M - \hat{\lambda}I$, with eigenvalue $\lambda^* - \hat{\lambda}$. Hence if $\hat{\lambda}$ is closer to λ^* than to the other eigenvalues of M , then the error $\|x(k) - x(k - 1)\|$ converges exponentially to zero. The number of steps K required for the error to become less than ϵ satisfies

$$K \leq c \log\left(\frac{1}{\epsilon}\right),$$

where c is a constant that depends on the initial estimate $\hat{\lambda}$. Note that K grows slowly with the accuracy ϵ . The constant c depends on $\hat{\lambda}$ via the ratio $|\hat{\lambda} - \lambda^*| / |\hat{\lambda} - \lambda(M)|$ for $\lambda \neq \lambda^*$, and we shall hereafter make the rather gross simplification that c is approximately unity.

For this method to be of practical use, one must characterize the distance between λ^* and the other eigenvalues of M . The following theorem asserts that λ^* is the only eigenvalue in the left-hand side of the complex plane. We shall hereafter call λ^* the *minimal* eigenvalue of M . To the best of our knowledge this fact was previously unknown.

Theorem 3 *The minimal eigenvalue of M , λ^* , is negative real whenever the center c of the ellipsoid $\mathcal{E}(c, C)$ is outside the unit ball. Moreover, all the other eigenvalues of M satisfy*

$$\text{Re}\{\lambda(M)\} \geq \lambda_1 > 0,$$

where λ_1 is the minimal eigenvalue of C ($C > 0$).

The theorem, whose proof is given in [21], asserts that λ^* is always isolated in the complex plane. This, in turn, affords a conclusion, stated in the following corollary, that λ^* is a real analytic function of the geometrical data. We will use in the corollary the following notation. Let R be an $N \times N$ rotation matrix that diagonalizes C , and let Λ be the resulting diagonal matrix,

$$\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n) \triangleq R^T C R.$$

and let

$$\bar{b} \triangleq R^T b.$$

Corollary 4.1 *λ^* , and hence margin($\mathcal{E}_1, \mathcal{E}_2$), are real analytic functions of the geometrical data. Moreover, a formula for the gradient of λ^* is given by*

$$\frac{\partial \lambda^*(\bar{b}, \Lambda)}{\partial b_i} = -\frac{1}{\alpha} \frac{\bar{b}_i}{(\lambda_i - \lambda^*)^2}$$

and

$$\frac{\partial \lambda^*(\bar{b}, \Lambda)}{\partial \lambda_i} = \frac{1}{\alpha} \frac{\bar{b}_i^2}{(\lambda_i - \lambda^*)^3},$$

where

$$\alpha \triangleq \sum_{i=1}^N \frac{\bar{b}_i^2}{(\lambda_i - \lambda^*)^3}.$$

Theorem 3 guarantees that there is a fixed-size disc of radius larger than λ_1 in the complex plane from which all initial guesses $\hat{\lambda}$ will converge to λ^* . A practical criterion to detect correct convergence is the attainment of ϵ -convergence to some negative real number in less than K steps, where $K = \log(\frac{1}{\epsilon})$.

The number of operations required by the inverse-iteration method is shown in [21] to be $5N^3 + (2N)^2 \log(\frac{1}{\epsilon})$. A comparison of this count with that of the QR method, with the substitution $\log(\frac{1}{\epsilon}) = K$, shows that the number of steps K required for an efficiency gain of two is

$$K = 3N \quad \text{where } N = 2, 3.$$

The accuracy of the solution obtained after K such steps is $\epsilon = 10^{-3N}$.

We have experimented with the incremental algorithm on a planar scene, in which an ellipsoidal link navigates in the presence of one stationary ellipsoidal obstacle. The link executes a biased random-walk, and at each step the free margin about the ellipsoidal link is computed with the incremental method. The random-walk was repeated for several randomly chosen polygonal link and obstacle pairs. A numerical example is shown in Figure 4, in which the rotational increment is one degree and the translational increment is one cell in a 128×128 grid. The average number of iterations required to attain $\epsilon = 1e - 06$ accuracy was 7, and the average time for one distance computation was 1 msec (on a DEC5000 machine)—shorter by a factor of about 2.5 than the QR-method discussed in the previous section. The incremental method ceased to converge correctly for rotational increment of about five degrees and translational increment of about five cells.




Figure 4. The closest point, marked by $*$, is traced as the ellipse surrounding the robot link moves around an obstacle (the ellipse surrounding the obstacle is not shown)

4.1 Conclusion

We have proposed in this paper a “complete” system: First a polyhedral robot and environment are approximated by ellipsoids. Then the free margin about each of the ellipsoidal links is computed in constant time per ellipsoidal obstacle. Hence, in terms of the ellipsoidal representation, the free margin about an n -link robot in an environment described by union of m convex polyhedra takes $O(n \cdot m)$ to compute. We have also shown how to accelerate the computation by exploiting the previous one along the robot's trajectory. Our ellipsoidal approach compares favorably with the polyhedral one, since there is no need to deal with the geometrical features of the underlying polyhedra the free-margin function takes essentially constant time to compute.

We have also presented in this paper an analytic formula for the free-margin function, and for its gradient vector-field. This, in turn, expands the catalogue of shapes for which a closed-form formula for the forbidden regions in configuration-space is known. This new formula also advances our larger program of research, concerned with setting up a system that will admit arbitrary Boolean combinations of linear and quadratic inequalities.

Last, it is worthwhile to note in the context of bounding shapes by ellipsoids that similar convex programming techniques can be used to efficiently compute the ellipsoid of maximal volume contained in a given convex polyhedron. This could be used to model the outer boundary of a robot work-cell. Moreover, by maintaining a pair of such ellipsoids, one surrounding the polyhedron and one contained in it, an estimate for the tightness of the ellipsoidal approximation can be effectively computed. The intersection between two polyhedra could then be first checked for the enclosing ellipsoids, for if they do not intersect than the underlying polyhedra are disjoint. Otherwise their interior ellipsoids are checked for intersection, and their intersection would imply that the underlying polyhedra intersect.

A Convex Programming

A real-valued function $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$ is a **convex function** if

$$\phi(sx_1 + (1-s)x_2) \leq s\phi(x_1) + (1-s)\phi(x_2)$$

for all $x_1, x_2 \in \mathbb{R}^n$ and $0 \leq s \leq 1$. A **convex optimization** (or **convex programming**) problem is to compute x^* that minimizes $\phi(x)$, subject to the constraint that x be in \mathcal{K} , where ϕ is a convex function and $\mathcal{K} \subset \mathbb{R}^n$ is a convex region. One standard algorithm used to solve this problem is the **ellipsoid algorithm**. It requires that \mathcal{K} be described by a convex function $\psi : \mathbb{R}^n \rightarrow \mathbb{R}$ in the form

$$\mathcal{K} = \{x \in \mathbb{R}^n : \psi(x) \leq 0\}.$$

The algorithm produces a sequence of points $x_k \in \mathbb{R}^n$ that converge to x^* . It needs to compute at the k^{th} step a separating plane passing through x_k for one of the two convex regions

$$\{x : \phi(x) \leq \phi(x_k)\} \quad \text{or} \quad \{x : \psi(x) \leq \psi(x_k)\}.$$

The separating plane is not necessarily unique, since the boundary of the region may have a sharp corner at x_k . In such situations the separating-plane's normal becomes a subgradient. More precisely, if $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$ is convex, but not necessarily differentiable, $g \in \mathbb{R}^n$ is a **subgradient** of ϕ at x if

$$\phi(z) \geq \phi(x) + g^T(z - x) \quad \text{for all } z \in \mathbb{R}^n.$$

The ellipsoid algorithm is initialized with an n -ellipsoid containing the minimizer x^* (e.g. a large n -ball containing \mathcal{K}). At the k^{th} step the current center of the ellipsoid, x_k , is compared against the constraint function ψ . If the constraint is violated ($\psi(x_k) > 0$), a separating plane passing through x_k for the region $\{x : \psi(x) \leq \psi(x_k)\}$ is computed. Otherwise, a separating plane passing through x_k for the region $\{x : \phi(x) \leq \phi(x_k)\}$ is computed. Clearly, one side of the resulting plane contains the entirety of \mathcal{K} in the first case, and contains the minimizer x^* in the other. In both cases the $(k+1)^{\text{th}}$ ellipsoid is computed as the *minimum-volume* ellipsoid that contain the intersection of the k^{th} ellipsoid with the half space determined by the separating plane. A closed-form formula for such an ellipsoid is known,

$$\mathcal{E}_{k+1} = \{x : (x - x_{k+1})^T X_{k+1}^{-1} (x - x_{k+1}) \leq 1\},$$

where

$$x_{k+1} = x_k - \frac{1}{n+1} X_k \bar{g}$$

and

$$X_{k+1} = \frac{n^2}{n^2-1} \left(X_k - \frac{2}{n+1} X_k \bar{g} \bar{g}^T X_k \right),$$

and \bar{g} is the normal to the separating plane. At each step the volume of the new ellipsoid is less than the volume of the previous one:

$$\text{volume}(\mathcal{E}_{k+1}) \leq e^{-1/2n} \text{volume}(\mathcal{E}_k),$$

by a factor that depends only on the dimension n of the ambient space. This affords in turn a conclusion that the number of steps K required to achieve ϵ -accuracy solution satisfies,

$$K \leq 2n^2 \log\left(\frac{c}{\epsilon}\right),$$

where c is constant. Note that this number grows slowly with both dimension n and accuracy ϵ ($n = 2$ or 3 and is constant in our case). Note, as well, that each step of the algorithm requires an evaluation of $\psi(x_k)$, an operation that is typically linear in the number of constraints used to describe \mathcal{K} . A more complete description can be found in [21] and [3, Chap. 14].

References

- [1] C. Bajaj and M. Kim. Generation of configuration space obstacles: Moving algebraic surfaces. *The International Journal of Robotics Research*, 9(1):92-112, February 1990.
- [2] J. E. Bobrow. A direct minimization approach for obtaining the distance between convex polyhedra. *The International Journal of Robotics Research*, 8(3):65-76, June 1989.
- [3] S.P. Boyd and C.H. Barratt. *Linear Controller Design*. Prentice Hall, NJ, 1991.
- [4] J. F. Canny. *The complexity of robot motion planning*. MIT Press, Cambridge, MA, 1988.
- [5] J.R. Dooley and J.M. McCarthy. Parameterized descriptions of the joint space obstacles for a 2r closed chain robot. In *IEEE International Conference on Robotics and Automation*, pages 1536-1541, Cincinnati, OH, May 1990.
- [6] W. Gander. Least squares with a quadratic constraint. *Numerische Mathematik*, 36:291-307, 1981.
- [7] W. Gander, G.H. Golub, and U. Matt. A constrained eigenvalue problem. *Linear Algebra and its Applications*, pages 815-839, 1989.
- [8] E. Gilbert and D. W. Johnson. Distance functions and their application to robot path planning in the presence of obstacles. *IEEE Journal of Robotics and Automation*, RA-1(1):21-30, March 1985.
- [9] E.G. Gilbert and C.P. Foo. Computing the distance between general convex objects in three-dimensional space. *IEEE Transactions on Robotics and Automation*, 6(1):53-61, Feb 1990.
- [10] F. John. Extremum problems with inequalities as subsidiary conditions (1948). In J. Moser, editor. *Fritz John Collected Papers*, chapter Vol 2, pages 543-560. Birkhauser, Boston, 1985.
- [11] O. Khatib. Real time obstacle avoidance for manipulators and mobile robots. *The International Journal of Robotics Research*, 5(1):90-99, Spring 1986.
- [12] J.-C. Latombe. *Robot Motion Planning*. Kluwer Academic Publishers, Boston, 1990.
- [13] J. Lengyel, M. Reichert, B.R. Donald, and D.P. Greenberg. Real-time robot motion planning using rasterizing computer graphics hardware. In *ACM SIGGRAPH*, pages 327-335, Las-Vegas, July 1990.
- [14] J. Levin. A parametric algorithm for drawing pictures of solid objects composed of quadratic surfaces. *Communications of the ACM*, 19(10):555-563, Oct 1976.
- [15] M. C. Lin and J.F. Canny. A fast algorithm for incremental distance calculation. In *IEEE International Conference on Robotics and Automation*, pages 1008-1014, Sacramento, CA, April 1991.
- [16] R. R. Martin and P.C. Stephenson. Containment algorithm for objects in rectangular boxes. In *Theory and Practice of Geometric Modeling*, pages 307-325. Springer Verlag, New York, 1989.
- [17] Y.E. Nesterov and A.S. Nemirovsky. *Interior Point Polynomial Methods in Convex Programming: Theory and Applications*. Springer Verlag, New York, 1992.
- [18] A. P. Pentland. Perceptual organization and the representation of natural form. In *Readings in Computer Vision*, pages 680-699. M. Kaufmann Publishers, Los Altos, CA, 1987.
- [19] M.J. Post. Minimum spanning ellipsoids. In *ACM Symposium on Theory of Computing*, pages 108-116. Washington, D.C., April 1984.
- [20] W.H. Press, B.P. Flannery, S.A. Teukolsky, and W.T. Vetterling. *Numerical Recipes in C*. Cambridge University Press, NY, 1988.
- [21] E. Rimón and S.P. Boyd. Efficient distance computation using best ellipsoid fit. Technical report, Information Systems Laboratory, Stanford University, Dec 1991.
- [22] S. Sclaroff and A. Pentland. Generalized implicit functions for computer graphics. In *ACM SIGGRAPH*, pages 247-250. Las Vegas, July 1991.