Computing Minimum Volume Enclosing Axis-Aligned Ellipsoids*

Piyush Kumar† E. Alper Yıldırım‡

September 22, 2006

Abstract

Given a set of points $\mathcal{S} = \{x^1, \ldots, x^m\} \subset \mathbb{R}^n$ and $\epsilon > 0$, we propose and analyze an algorithm for the problem of computing a $(1+\epsilon)$-approximation to the the minimum volume axis-aligned ellipsoid enclosing $\mathcal{S}$. We establish that our algorithm is polynomial for fixed ϵ. In addition, the algorithm returns a small core set $\mathcal{X} \subseteq \mathcal{S}$, whose size is independent of the number of points m, with the property that the minimum volume axis-aligned ellipsoid enclosing $\mathcal{X}$ is a good approximation of the minimum volume axis-aligned ellipsoid enclosing $\mathcal{S}$. Our computational results indicate that the algorithm exhibits significantly better performance than that indicated by the theoretical worst-case complexity result.

Key words: Axis-aligned ellipsoids, enclosing ellipsoids, core sets, approximation algorithms.

Extended Abstract

Real time computer graphics and computer gaming call for the computation of simple bounding regions for rapid culling in the rendering process and for rapid determination that two objects are not intersecting during the collision detection process. An axis-aligned ellipsoid is one such bounding region with low storage complexity that can support fast intersection tests. Currently, in computer graphics, this is done by computing the axis-aligned bounding box of the object first and then computing the minimum volume axis-aligned ellipsoid (MVAE) enclosing the resulting box. While such a scheme usually enables one to quickly compute a reasonably good approximation of the MVAE enclosing the object under consideration, the volume of the resulting ellipsoid may be significantly larger than that of the optimal ellipsoid for certain geometric objects. For instance, if the object is almost spherical, the simple procedure outlined above would return an ellipsoid whose volume may exceed the volume of the MVAE enclosing the object by a factor of as much as $n^{n/2}$, where n is the dimension of the object. Therefore, this procedure may lead to false positive results in collision detection, which provides one of our motivations to study the minimum volume enclosing axis-aligned ellipsoid problem.

Another application of this problem in higher dimensions is in machine learning, where the kernel approach is widely used. Kernel functions are functions that live in low dimensional spaces but behave like inner products in high dimensions. Computational geometers have used a similar approach since almost the advent of the field. In this approach, the main idea is based on *linearization*, for which kernel functions provide an implicit way. For instance, kernel functions are used in support vector machines to separate nonlinearly separable data via calculating a hyperplane in a different space. In case of enclosing shapes, it is easy to calculate the linearization given a kernel function and then apply optimization algorithms to compute enclosing shapes in the explicitly linearized space. The problem with this approach is that the time complexity is quadratic or more in the size of the input. For minimum enclosing balls (MEBs), the kernel version of the problem can be solved very efficiently. The problem with calculating the minimum volume enclosing ellipsoid (MVEE) in kernel space is that the core set size dependence of MVEEs is quadratic (or higher) in dimension so the number of support vectors produced with such an algorithm is too large for good generalization bounds (Ref. 3). Keeping this in mind, a natural question emerges: Is there a bounding shape whose quality is between MEB and MVEE and that can be efficiently computed? Axis-aligned ellipsoids are clearly an answer to this question but efficient algorithms to compute enclosing MVAEs in higher dimensions have not been studied yet. Moreover, for machine learning applica-

*Full paper link: http://www.optimization-online.org/DB_HTML/2006/06/1419.html

†Department of Computer Science, Florida State University, Tallahassee, FL 32306-4530, USA. e-mail: piyush@acm.org

‡Department of Industrial Engineering, Bilkent University, Bilkent, Ankara 06800, Turkey. e-mail: yildirim@bilkent.edu.tr. This author was supported in part by NSF through CAREER Grant DMI-0237415.

tions, it is also important to identify a small subset of the input points with the property that the MVAE enclosing this subset is a good approximation of the MVAE enclosing the original set of points. These considerations form a basis for studying minimum volume axis-aligned enclosing ellipsoids.

In this paper, we propose and study an algorithm that computes an approximate minimum volume axis-aligned ellipsoid enclosing a given set of m points in $\mathbb{R}^n$. More precisely, given $\mathcal{S} := \{x^1, \ldots, x^m\} \subset \mathbb{R}^n$ and $\epsilon > 0$, our algorithm computes an axis-aligned ellipsoid $\mathcal{E} \subset \mathbb{R}^n$ that satisfies

$$\mathcal{S} \subseteq \mathcal{E}, \quad \text{Vol}\, \mathcal{E} \leq (1+\epsilon)\ \text{Vol}\, \mathcal{E}^*, \tag{1}$$

where $\mathcal{E}^*$ denotes the MVAE enclosing $\mathcal{S}$ and $\text{Vol}(\cdot)$ denotes the volume. An axis-aligned ellipsoid $\mathcal{E}$ satisfying (1) is said to be a $(1+\epsilon)$-approximation to the MVAE enclosing $\mathcal{S}$.

Our algorithm is mainly motivated by the algorithm developed earlier by the authors of this paper for the minimum volume enclosing ellipsoid problem (Ref. 2) (henceforth the KY algorithm), which, in turn, improves upon the algorithm of Khachiyan (Ref. 1) by using a simple initial volume approximation. In particular, we establish that our algorithm computes a $(1+\epsilon)$-approximation to the MVAE enclosing $\mathcal{S}$ in

$$O\left(mn^2(\log n + n^2[(1+\epsilon)^{2/n} - 1]^{-1})\right) \tag{2}$$

arithmetic operations (cf. Theorem **??**), which is linear in m, the number of points in $\mathcal{S}$, and is polynomial for fixed $\epsilon > 0$. In particular, the overall complexity reduces to $O(mn^5/\epsilon)$ for $\epsilon \in (0,1)$, which implies that our algorithm is especially well-suited for instances of the problem for which $m \gg n$ and for moderately small values of ϵ.

Despite the underlying similarity between the KY algorithm and the algorithm proposed in this paper, our theoretical complexity analysis here is slightly more involved than that of (Ref. 2) and relies on somewhat different tools. In contrast with the KY algorithm, the one-dimensional line search problem that arises in each iteration of our algorithm does not have a closed form solution. We circumvent this difficulty by using an approximate solution, which, in turn, leads to a worse complexity result given by (2) than that of the KY algorithm. On the other hand, we also establish that our theoretical analysis in general cannot be improved by demonstrating several examples

Similar to the KY algorithm, our algorithm in this paper also computes a subset $\mathcal{X} \subseteq \mathcal{S}$ with the property that

$$\text{Vol}\, \mathcal{E}^*_{\mathcal{X}} \leq \text{Vol}\, \mathcal{E}^* \leq \text{Vol}\, \mathcal{E} \leq (1+\epsilon)\text{Vol}\, \mathcal{E}^*_{\mathcal{X}} \leq (1+\epsilon)\text{Vol}\, \mathcal{E}^*,$$

where $\mathcal{E}^*_{\mathcal{X}}$ denotes the MVAE enclosing $\mathcal{X}$. It follows from (1) that the ellipsoid $\mathcal{E}$ returned by our algorithm is simultaneously a $(1+\epsilon)$-approximation to the MVAE enclosing $\mathcal{X}$ and to that enclosing $\mathcal{S}$. In addition, $|\mathcal{X}| = O\left(n(\log n + n^2[(1+\epsilon)^{2/n} - 1]^{-1})\right)$, which is independent of $|\mathcal{S}| = m$. Following the earlier literature, we call $\mathcal{X}$ an ϵ-core set (or a core set) of $\mathcal{S}$ to signify that $\mathcal{X}$ provides an approximate and compact representation of $\mathcal{S}$. To the best of our knowledge, this establishes the first core set result for the minimum volume enclosing axis-aligned ellipsoid problem. In comparison with the KY algorithm, the theoretical estimate of the size of the core set for this problem is considerably larger. Similarly to the overall complexity result, this is a byproduct of a more pessimistic theoretical analysis, which, in general, cannot be improved.

In an attempt to highlight the potential discrepancy between the worst-case theoretical complexity result and the practical behavior of the algorithm, we implemented two different versions in MATLAB. While the first version numerically computes an "exact" solution to the line search problem mentioned above, the second one is an exact implementation of our algorithm using only the approximate solution of the line search problem. Our computational experiments reveal that the former version is usually much faster than the latter. These results indicate that the overall computation time and the size of the core set in practice tend to be much smaller than the corresponding worst-case estimates for our algorithm.

References

1. L. G. Khachiyan. Rounding of polytopes in the real number model of computation. *Mathematics of Operations Research*, 21:307–320, 1996.
2. P. Kumar and E. A. Yıldırım. Minimum volume enclosing ellipsoids and core sets. *Journal of Optimization Theory and Applications*, 126(1):1–21, 2005.
3. J. Shawe-Taylor and N. Cristianini. *Kernel methods for pattern analysis*. Cambridge University Press, 2004.