Abstract

In the polytope membership problem, a convex polytope K in $\mathbb{R}^d$ is given, and the objective is to preprocess K into a data structure so that, given a query point $q \in \mathbb{R}^d$, it is possible to determine efficiently whether $q \in K$. We consider this problem in an approximate setting and assume that d is a constant. Given an approximation parameter $\epsilon > 0$, the query can be answered either way if the distance from q to K's boundary is at most ϵ times K's diameter. Previous solutions to the problem were on the form of a space-time trade-off, where logarithmic query time demands $O(1/\epsilon^{d-1})$ storage, whereas storage $O(1/\epsilon^{(d-1)/2})$ admits roughly $O(1/\epsilon^{(d-1)/8})$ query time. In this paper, we present a data structure that achieves logarithmic query time with storage of only $O(1/\epsilon^{(d-1)/2})$, which matches the worst-case lower bound on the complexity of any ϵ-approximating polytope. Our data structure is based on a new technique, a hierarchy of ellipsoids defined as approximations to Macbeath regions. As an application, we obtain major improvements to approximate Euclidean nearest neighbor searching. Notably, the storage needed to answer ϵ-approximate nearest neighbor queries for a set of n points in $O(\log(n/\epsilon))$ time is reduced to $O(n/\epsilon^{d/2})$. This halves the exponent in the ϵ-dependency of the existing space bound of roughly $O(n/\epsilon^d)$, which has stood for 15 years (Har-Peled, 2001).