

A faster dual algorithm for the Euclidean minimum covering ball problem

Marta Cavaleiro · Farid Alizadeh

June 30, 2017

Abstract Dearing and Zeck [5] presented a dual algorithm for the problem of the minimum covering ball in \mathbb{R}^n . Each iteration of their algorithm has a computational complexity of at least $\mathcal{O}(n^3)$. In this paper we propose a modification to their algorithm that, together with an implementation that uses updates to the QR factorization of a suitable matrix, achieves a $\mathcal{O}(n^2)$ iteration.

Keywords minimum covering ball · smallest enclosing ball · 1-center · minmax location · computational geometry

1 Introduction

Consider a given set of points $\mathcal{P} = \{p_1, \dots, p_m\}$ in the Euclidean space \mathbb{R}^n . Let $\|\cdot\|$ denote the Euclidean norm. The problem of finding the hypersphere

$$B(x, r) = \{y \in \mathbb{R}^n : \|y - x\| \leq r\}$$

with minimum radius that covers \mathcal{P} , which we will refer to as the minimum covering ball (MB) of \mathcal{P} , can be formulated as

$$MB(\mathcal{P}) := \begin{array}{ll} \min r^2 & \\ \text{s.t. } \|p_i - x\|^2 \leq r^2, & i = 1, \dots, m. \end{array} \quad (1)$$

We will use the notation $MB(\mathcal{P})$ both to refer to the problem of the minimum covering ball of a set \mathcal{P} and, depending on the context, also to the corresponding optimal ball.

The MB problem, reported to date back to the 19th century [18], is an important and active problem in computational geometry and optimization. Applications are wide and include facility location ($n = 2$ and $n = 3$), computer graphics (mostly $n = 3$), machine learning (usually with n very large).

Problem (1) can easily be converted in a quadratic program (QP) and solved using off-the-shelf QP solvers. Gärtner and Schönherr [10] developed a generalization of the simplex method for QP with the goal of targeting geometric QPs, with one of the main applications being the MB problem. The MB problem is also a second-order-cone program (SOCP) instance, so interior point methods may also be

Marta Cavaleiro
Rutgers University, MSIS Dept. & RUTCOR
100 Rockafellar Rd, Piscataway, NJ 08854. E-mail: marta.cavaleiro@rutgers.edu

Farid Alizadeh
Rutgers University, MSIS Dept. & RUTCOR
100 Rockafellar Rd, Piscataway, NJ 08854. E-mail: farid.alizadeh@rutgers.edu

used. Some work has also been done using the SOCP formulation while exploiting the special features of the MB problem [13,22].

The algorithmic complexity of the MB problem was first given by Megiddo in 1983, when he gave the first deterministic algorithm that solves the MB problem in linear time on the number of points, when the dimension is fixed [15,16].

The MB problem is an *LP-type* problem [6], sharing common properties with linear programs, in particular its combinatorial nature: the optimal ball is determined by at most $n + 1$ points of \mathcal{P} that lie on its boundary. Such set of points is usually called a *support set*. Many algorithms that search for such set have been developed. Welzl [19] presented a randomized algorithm that searches for a support set, solving the problem in expected linear time for fixed dimension. This algorithm was improved by Gärtner [9], however only dimensions $n \leq 25$ could be handled in practice. Later Fischer and Gärtner [7] proposed an algorithm with a pivoting scheme resembling the simplex method for LP based on previous ideas from [12], even adapting Bland's rule to avoid cycling. Their algorithm can handle problems in larger dimensions ($n \leq 10,000$). Using related ideas, Dearing and Zeck [5] developed a dual algorithm for the MB problem. This algorithm is the subject of this paper, and as we will prove, it can also deal with larger dimensions.

The concept of ϵ -core sets, proposed by Bădoiu et al. [3], introduced a new direction of research in approximate algorithms for the MB problem. A remarkable property is the existence of a ϵ -core set of size at most $1/\epsilon$ and independent of n [2,13]. Several algorithms focused on finding ϵ -core sets have been proposed [3,2,13,20,17,14], being in general able to deal with large dimensions in useful time.

Streaming algorithms, that only allow one pass over the input points, have also been studied. Zarrabi-Zadeh and Chan [21] gave a 3/2-approximation algorithm, and later an algorithm by Agarwal and Sharathkumar [1] was able to achieve a 1.22-approximation factor [4].

In this paper we propose a modification to the algorithm proposed by Dearing and Zeck [5] that makes it faster. Their algorithm looks for a support set, by solving a sequence of subproblems $MB(\mathcal{S})$, with $\mathcal{S} \subseteq \mathcal{P}$ affinely independent, until all points of \mathcal{P} are covered. At each iteration, set \mathcal{S} is updated by either adding a point that is not yet covered, or by replacing an existing point in \mathcal{S} by it. Problem $MB(\mathcal{S})$ is solved using a directional search procedure, and during this step possibly more points are removed from \mathcal{S} .

It is possible to implement the algorithm as presented in [5] taking advantage of the QR updates of a suitable matrix, which can be done in quadratic time on n . However, they would need to be done as many as $\mathcal{O}(n)$ times, resulting in an iteration having $\mathcal{O}(n^3)$ computational complexity. We modify the directional search procedure in such a way that, together with an implementation using QR updates, achieves a $\mathcal{O}(n^2)$ iteration.

The paper is organized as follows: In section 2 we review Dearing and Zeck's algorithm for the MB problem. In section 3 we present a modification of the algorithm that makes its directional-search step faster. In section 4 we show how the algorithm can be implemented using QR updates of a certain matrix, which together with the results from section 3, result in $\mathcal{O}(n^2)$ iteration. Finally, in section 5 we present some computational results that show the practical impact of our work.

One application of this algorithm that benefits from it having a faster iteration is a branch and bound approach to solve the problem of the minimum k -covering ball (that seeks the ball with smallest radius that contains at least k of m given points in \mathcal{P}). At each node of the search tree, Dearing and Zeck's algorithm can be employed to solve the corresponding subproblem. Since it starts with the solution of the parent node, it usually needs very few iterations to solve each subproblem. And since a very large number of subproblems need to be solved, having a fast iteration is essential. This application will be the subject of a future publication by the authors of this paper.

2 The dual algorithm by Dearing and Zeck

The dual problem of (1) is

$$\begin{aligned} \max_{\pi \in \mathbb{R}^m} \quad & \sum_{i=1}^m \pi_i \|p_i\|^2 - \left\| \sum_{i=1}^m \pi_i p_i \right\|^2 \\ \text{s.t.} \quad & \sum_{i=1}^m \pi_i = 1 \\ & \pi_i \geq 0, \quad i = 1, \dots, m, \end{aligned} \tag{2}$$

and the optimal primal and dual solutions satisfy the following complementary slackness conditions

$$\sum_{i=1}^m \pi_i (p_i - x) = 0 \quad \text{and} \quad \pi_i (r^2 - \|x - p_i\|^2) = 0, \quad i = 1, \dots, m. \tag{3}$$

An important consequence of duality and the complementary slackness conditions, and a well known fact, is the following lemma

Lemma 1 *Consider a ball $B(x, r)$ that covers \mathcal{P} , and let $\mathcal{S} \subseteq \mathcal{P}$ be the set of points on the boundary of $B(x, r)$. Then $B(x, r) = MB(\mathcal{P})$ if and only if $x \in \text{conv}(\mathcal{S})$.*

As a consequence, the optimal ball is determined by at most $n + 1$ affinely independent points of \mathcal{P} . Moreover, it is easily proved that the center is the intersection of the bisectors of the facets of $\text{conv}(\mathcal{S})$, the convex hull of \mathcal{S} .

We now define *support set* for the $MB(\mathcal{P})$ problem, an analogous concept to the one of basis for linear programming.

Definition 1 (Support set) An affinely independent subset $\mathcal{S} \subseteq \mathcal{P}' \subseteq \mathcal{P}$ is a *support set* of $MB(\mathcal{P}')$ if $MB(\mathcal{S}) = MB(\mathcal{P}')$ and \mathcal{S} is minimal, in the sense that it does not exist a $\mathcal{S}' \subset \mathcal{S}$ s.t. $MB(\mathcal{S}') = MB(\mathcal{P}')$.

The following lemma is yet another consequence from duality and complementary slackness:

Lemma 2 *\mathcal{S} is a support set of $\mathcal{P}' \subseteq \mathcal{P}$ if and only if x , the center of $MB(\mathcal{P}')$, is in $\text{ri conv}(\mathcal{S})$, the relative interior of $\text{conv}(\mathcal{S})$.*

The dual algorithm developed by Dearing and Zeck [5] solves the problem $MB(\mathcal{P})$ by finding a support set of $MB(\mathcal{P})$ the following way: at the beginning of each iteration we have the solution of $MB(\mathcal{S})$, where \mathcal{S} is a support set of $MB(\mathcal{S})$; a point $p \in \mathcal{P}$ that is not yet covered by $MB(\mathcal{S})$ is then selected, and problem $MB(\mathcal{S} \cup \{p\})$ is solved by iteratively removing points from the set $\mathcal{S} \cup \{p\}$ until a support of $MB(\mathcal{S} \cup \{p\})$ is found. The radius of the covering ball strictly increases at each iteration and the algorithm stops when all points have been covered. As a consequence, the algorithm is finite [5, Theo. 3.2].

We now review the algorithm in more detail. We denote by x and r the center and radius of the ball at each iteration, and by \mathcal{S} the corresponding support set.

As \mathcal{S} and x are updated throughout each iteration, the algorithm maintains the following invariants

- \mathcal{S} is affinely independent,
- $x \in \text{conv}(\mathcal{S})$.

As a consequence, dual feasibility and complementary slackness conditions are both always satisfied.

2.1 Initialization

The routine starts with a support set $\mathcal{S} \subseteq \mathcal{P}$ of $MB(\mathcal{S})$, and x and r the solution to $MB(\mathcal{S})$. If such data is not available, the algorithm picks any two points $\{p_{i_1}, p_{i_2}\} \in \mathcal{P}$, and solves $MB(\mathcal{S})$ with $\mathcal{S} = \{p_{i_1}, p_{i_2}\}$:

$$x = \frac{p_{i_1} + p_{i_2}}{2} \quad \text{and} \quad r = \|x - p_{i_1}\|.$$

2.2 Iteration

At the beginning of each iteration, we have the minimum covering ball of a support set $\mathcal{S} = \{p_{i_1}, \dots, p_{i_s}\} \subseteq \mathcal{P}$, whose center and radius are x and r respectively. Moreover $x \in \text{ri conv}(\mathcal{S})$.

1. *Optimality check*: If all points of \mathcal{P} are covered by $MB(\mathcal{S})$ then $MB(\mathcal{S}) = MB(\mathcal{P})$. Otherwise, the algorithm picks a point $p \in \mathcal{P}$ that is not yet covered.

2. *Update \mathcal{S}* : If $\mathcal{S} \cup \{p\}$ is affinely independent then $\mathcal{S} = \mathcal{S} \cup \{p\}$, and the two invariants are maintained for the current \mathcal{S} and x .

If $\mathcal{S} \cup \{p\}$ is not affinely independent, a point $p_{i_k} \in \mathcal{S}$ is dropped from \mathcal{S} , in such way that $\mathcal{S} \setminus \{p_{i_k}\} \cup \{p\}$ is affinely independent and $x \in \text{conv}(\mathcal{S} \setminus \{p_{i_k}\} \cup \{p\})$. Consider s the cardinality of \mathcal{S} . Point p_{i_k} to leave \mathcal{S} is calculated as follows:

- Let π_1, \dots, π_s solve $\sum_{j=1}^s \pi_j p_{i_j} = x$ and $\sum_{j=1}^s \pi_j = 1$;
- Let $\omega_1, \dots, \omega_s$ solve $\sum_{j=1}^s \omega_j p_{i_j} = -p$ and $\sum_{j=1}^s \omega_j = -1$;
- p_{i_k} is such that $\frac{\pi_k}{-\omega_k} = \min_{j=1, \dots, s} \left\{ \frac{\pi_j}{-\omega_j} : \omega_j < 0 \right\}$.

\mathcal{S} is now updated: $\mathcal{S} = \mathcal{S} \setminus \{p_{i_k}\} \cup \{p\}$.

3. *Solution of $MB(\mathcal{S})$* : The optimal center of $MB(\mathcal{S})$ can occur in two possible locations:

(a) Either the center is in the interior of $\text{conv}(\mathcal{S})$, in which case it is the intersection of the bisectors of the facets of $\text{conv}(\mathcal{S})$ with $\text{aff}(\mathcal{S})$, the affine space generated by the points in \mathcal{S} . In this case \mathcal{S} is the support set of $MB(\mathcal{S})$.

(b) Or the center is on one of the facets of $\text{conv}(\mathcal{S})$, implying that the support set of $MB(\mathcal{S})$ is a proper subset of \mathcal{S} .

To find the optimal center of $MB(\mathcal{S})$, the algorithm uses a directional search procedure that starts at the current center x , and proceeds along a direction d . During this search the algorithm either immediately finds the optimal solution to $MB(\mathcal{S})$, or identifies a point from \mathcal{S} that is not part of the support set of $MB(\mathcal{S})$, removing it from \mathcal{S} , and performing a new directional search next. The details of this directional search are described below:

i. *The direction d* : d , the direction along which the line search is performed, satisfies the following:

(a) It is on the $(s-1)$ -dimensional subspace generated by the points in \mathcal{S} , $\text{Sub}(\mathcal{S})$, where s is the cardinality of \mathcal{S} , that is

$$u_j^T d = 0, \quad j = 1, \dots, n - s + 1,$$

with $\{u_j\}_j$ a basis for $\text{Null}(\mathcal{S})$, the null space of $\text{Sub}(\mathcal{S})$;

(b) It is parallel to the intersection of the bisectors of the facets of the polytope $\text{conv}(\mathcal{S} \setminus \{p\})$, or, equivalently, it is orthogonal to $\text{Sub}(\mathcal{S} \setminus \{p\})$, so

$$(p_{i_j} - p_{i_s})^T d = 0, \quad p_{i_j} \in \mathcal{S} \setminus \{p_{i_s}, p\};$$

(c) It “points towards” p , in the sense that the distance to p from any point on the ray $\ell^+ = \{x + \alpha d : \alpha \geq 0\}$ decreases as α increases:

$$(p - p_{i_s})^T d = 1.$$

ii. *Calculating the next iterate:* The ray ℓ^+ intersects both the intersection of the bisectors of the facets of $\text{conv}(\mathcal{S})$, let α_b correspond to that point, and one (or the intersection of several) of those facets, let α_f correspond to that point. Two cases are possible:

Case 1: $\alpha_b < \alpha_f$, that is, the intersection with the bisectors occurs first. If this is the case, the solution to $MB(\mathcal{S})$ is the point $x + \alpha_b d$. The algorithm goes back to Step 1 with $x = x + \alpha_b d$ and the support set \mathcal{S} .

Case 2: $\alpha_b \geq \alpha_f$. In this case the opposite point $p_{i_i} \in \mathcal{S}$ to the (or one of the) intersected facet(s) is not part of the support set of $MB(\mathcal{S})$, being therefore removed from \mathcal{S} : $\mathcal{S} = \mathcal{S} \setminus \{p_{i_i}\}$. Note that p_{i_i} can never be p . The algorithm now returns to the beginning of Step 3 with the new \mathcal{S} and $x = x + \alpha_f d$, for a new directional search to solve $MB(\mathcal{S})$.

For a full description of the algorithm and its correctness we refer the reader to [5].

3 Directional search via orthogonal projections

In [5], the authors find α_f , the point where the ray ℓ^+ intersects the boundary of $\text{conv}(\mathcal{S})$, by calculating the intersection of the ray with each one of the facets (as many as $n + 1$). This is the central reason why, even using efficient updates to the QR factorization of a matrix, their iteration could not have a better complexity than $\mathcal{O}(n^3)$, since such updates would need to be done $\mathcal{O}(n)$ times.

We now show how one can find α_f without having to check each facet. The idea consists on projecting $\text{conv}(\mathcal{S})$ and the ray ℓ^+ orthogonally onto $\text{aff}(\mathcal{S} \setminus \{p\})$. Recall that ℓ^+ is perpendicular to $\text{aff}(\mathcal{S} \setminus \{p\})$, so its projection will be a single point, the projection of x . In order to find the intersected facet, we find in which two projected facets of $\text{conv}(\mathcal{S})$ the projection of x fell into. We then calculate the intersection of the ray with the two facets, and the one with smallest α that is non-negative, corresponds to the facet of $\text{conv}(\mathcal{S})$ intersected by the ray.

Before we proceed into the details, consider the following notation:

- $\mathcal{S} = \{p_1, \dots, p_{s-1}, p\}$ and $\mathcal{S}' = \mathcal{S} \setminus \{p\}$;
- \mathcal{C} , the polytope $\text{conv}(\mathcal{S})$, and $\partial\mathcal{C}$ its boundary;
- $F_j = \text{conv}(\mathcal{S} \setminus \{p_j\})$, the facet of \mathcal{C} opposed to point $p_j \in \mathcal{S}$, $j = 1, \dots, s - 1$;
- $F_0 = \text{conv}(\mathcal{S} \setminus \{p\})$, the facet of \mathcal{C} opposed to p ;
- $\ell = \{x + \alpha d : \alpha \in \mathbb{R}\}$ and $\ell^+ = \{x + \alpha d : \alpha \geq 0\}$;
- α_j , $j = 0, \dots, m$, be the intersection of ℓ with facet F_j .

Recall that, at the beginning of Step 3ii, we have:

- \mathcal{S} is an affinely independent set, and therefore $p \notin \text{aff}(\mathcal{S}')$;
- d is a direction in $\text{Sub}(\mathcal{S})$, orthogonal to $\text{aff}(\mathcal{S}')$, that points towards p and passes through the intersection of the bisectors of the facets of $\text{conv}(\mathcal{S})$;
- x , the current solution, is in $\text{conv}(\mathcal{S})$.

Let x' and p' be the orthogonal projections of x and p onto $\text{aff}(\mathcal{S}')$, respectively, and let

$$x' = \sum_{j=1}^{s-1} \pi_j p_j \text{ s.t. } \sum_{j=1}^{s-1} \pi_j = 1 \quad \text{and} \quad p' = \sum_{j=1}^{s-1} \omega_j p_j \text{ s.t. } \sum_{j=1}^{s-1} \omega_j = 1, \quad (4)$$

be their unique representations as affine combinations of points of \mathcal{S}' . An important observation is that the projection of all points of ℓ onto $\text{aff}(\mathcal{S}')$ coincides with x' . The two Lemmas that follow will be useful later.

Lemma 3 $x' \in \text{conv}(\mathcal{S}' \cup \{p'\})$.

Lemma 3 is a direct consequence of the linearity of the projection operator and the fact that $x \in \text{conv}(\mathcal{S})$.

Lemma 4 If $\pi_j < 0$ then $\omega_j < 0$.

Proof From Lemma 3 we know that there exists $\beta \geq 0$ and $\beta' \geq 0$ such that

$$x' = \sum_{j=1}^{s-1} \beta_j p_j + \beta' p' \quad \text{and} \quad \sum_{j=1}^{s-1} \beta_j + \beta' = 1.$$

Therefore,

$$x' = \sum_{j=1}^{s-1} \beta_j p_j + \beta' \sum_{j=1}^{s-1} \omega_j p_j = \sum_{j=1}^{s-1} (\beta_j + \beta' \omega_j) p_j.$$

Let $\delta_j = \beta_j + \beta' \omega_j$. Note that $\sum_{j=1}^{s-1} \delta_j = 1$. Since \mathcal{S} is an affinely independent set, the representation of x' as an affine combination of the points in \mathcal{S} is unique, therefore we must have $\pi_j = \delta_j = \beta_j + \beta' \omega_j$ for all j . Thus, if $\pi_j < 0$ then $\omega_j < 0$. Note that $\pi_j = 0$ does not imply $\omega_j \leq 0$.

Consider the general case where the intersection of line ℓ with $\partial\mathcal{C}$ is two points z_1 and z_2 . Let F_{k_1} be one of the facets containing z_1 and F_{k_2} be one containing z_2 . Point z_1 and/or z_2 may be on the intersection of several facets, but knowing one of them suffices. The projection of z_1 and z_2 onto $\text{aff}(\mathcal{S}')$ is x' , therefore x' will be written as a unique convex combination of the projections of the points that form F_{k_1} and also as a unique convex combination of the projections of the points that form F_{k_2} , and only of those and no other projected facets. In the particular cases when ℓ intersects $\partial\mathcal{C}$ on a single point or on an infinite number of points, x' will still be written as a convex combination of the projected points of one of the intersected facets. Lemma 5 proves this fact. Note that the intersection always exists since $x \in \ell \cap \mathcal{C}$. Theorem 1 finds those at most two unique convex representations of x' , finding consequently the intersected facets of $\partial\mathcal{C}$ by ℓ .

Lemma 5 Line ℓ intersects F_k , the facet of \mathcal{C} opposed to $p_k \in \mathcal{S}'$, if and only if $x' \in \text{conv}(\mathcal{S}' \setminus \{p_k\} \cup \{p'\})$.

Proof We only need to prove that if $x' \in \text{conv}(\mathcal{S}' \setminus \{p_k\} \cup \{p'\})$ then ℓ intersects F_k , since the opposite is trivial as a consequence of the linearity of the projection operator.

Let $B = [p_2 - p_1, \dots, p_{s-1} - p_1]$. Then $p' = B(B^T B)^{-1} B^T (p - p_1) + p_1$. Firstly, we prove that there exists a $\gamma > 0$ s.t. $d = \gamma(p - p')$. Clearly, $p - p' \in \text{Sub}(\mathcal{S})$, and $p - p'$ is orthogonal to $\text{Sub}(\mathcal{S}')$ since

$$B^T (p - p') = B^T (p - p_1) - B^T (p - p_1) = 0.$$

Moreover, d and $p - p'$ have the same direction since (recall that $(p - p_{s-1})^T d > 0$)

$$(p - p_{s-1})^T (p - p') = \|p - p'\|^2 + (p' - p_{s-1})^T (p - p') = \|p - p'\|^2 > 0.$$

Thus we conclude that there exists a $\gamma > 0$ such that $d = \gamma(p - p')$.

Now, suppose $x' \in \text{conv}(\mathcal{S}' \setminus \{p_k\} \cup \{p'\})$, that is,

$$x' = \sum_{j \neq k} \beta_j p_j + \beta' p' \quad \text{with} \quad \sum_{j \neq k} \beta_j + \beta' = 1 \quad \text{and} \quad \beta' \geq 0, \beta_j \geq 0, j = 1, \dots, s-1.$$

Observe that for any α we have

$$x' + \alpha d = \sum_{j \neq k} \beta_j p_j + \beta' p' + \alpha \gamma (p - p') = \sum_{j \neq k} \beta_j p_j + (\beta' - \alpha \gamma) p' + \alpha \gamma p.$$

Let $\alpha' = \frac{\beta'}{\gamma}$. We have that $x' + \alpha' d \in \text{conv}(\mathcal{S}' \setminus \{p_k\} \cup \{p\})$ since $\sum_{j \neq k} \beta_j + \alpha' \gamma = 1$. This implies that there exists an α such that $x + \alpha d \in F_k$, that is, ℓ intersects F_k .

Now we show in Theorem 1 how to calculate the intersection of ℓ^+ and $\partial\mathcal{C}$, that is α_f . Recall that

$$\alpha_f = \min_{j=1,\dots,s-1} \{\alpha_j : \alpha_j \geq 0\},$$

and α_j , the intersection of ℓ with any F_j , can be calculated by finding a basis $\{w_1, \dots, w_{n-(s-1)+1}\}$ for the null space of $\text{Sub}(\mathcal{S}' \setminus \{p_k\} \cup \{p\})$, and

$$\alpha_j = \frac{(p-x)^T w_i}{d^T w_i}, \text{ for any } i = 1, \dots, n-s+2, \text{ s.t. } d^T w_i \neq 0. \quad (5)$$

Theorem 1 Consider the representations (4) of x' and p' . Let F_k , the facet opposed to $p_k \in \mathcal{S}'$, be the one first intersected by ℓ^+ , and let α_k be the value of α at which the intersection occurs. To find p_k and α_k , there are two possible cases:

- Case 1. If there is a $k = 1, \dots, m$ such that $\pi_k = \omega_k = 0$, then p_k is the point opposed to the facet intersected first, and $\alpha_f = \alpha_k = 0$.
- Case 2. Suppose case 1 does not hold. First, find k_1 such that

$$\frac{\pi_{k_1}}{\omega_{k_1}} = \min_{j=1,\dots,s-1} \left\{ \frac{\pi_j}{\omega_j} : \pi_j \geq 0, \omega_j > 0 \right\}, \quad (6)$$

then, find α_{k_1} as in (5). Let $\mathcal{J} := \{j : \pi_j \leq 0, \omega_j < 0\}$. If $\mathcal{J} \neq \emptyset$, find k_2 such that

$$\frac{\pi_{k_2}}{\omega_{k_2}} = \max_{j=1,\dots,s-1} \left\{ \frac{\pi_j}{\omega_j} : \pi_j \leq 0, \omega_j < 0 \right\}, \quad (7)$$

and find α_{k_2} as in (5). If $\mathcal{J} = \emptyset$ simply consider $\alpha_{k_2} = -\infty$.

The facet first intersected by ℓ^+ , F_k , is such that

$$k = \arg \min_{j=k_1, k_2} \{\alpha_j : \alpha_j \geq 0\},$$

and $\alpha_f = \alpha_k$.

Proof Case 1. Suppose there is a $k \in \{1, \dots, m\}$ such that $\omega_k = \pi_k = 0$. Then, since $x = x' + \delta d$, for some $\delta \geq 0$, and the fact that there is a $\gamma > 0$ such that $d = \gamma(p - p')$ (see the proof of Lemma 5) we have:

$$x = x' + \delta\gamma(p - p') = \sum_{\substack{j=1 \\ j \neq k}}^{s-1} \pi_j p_j + \delta\gamma \left(p - \sum_{\substack{j=1 \\ j \neq k}}^{s-1} \omega_j p_j \right) = \sum_{\substack{j=1 \\ j \neq k}}^{s-1} (\pi_j - \delta\gamma\omega_j) p_j + \delta\gamma p,$$

and this representation of x as a convex combination of \mathcal{S} is unique, since \mathcal{S} is affinely independent. Consequently, $\pi_j - \delta\gamma\omega_j \geq 0$ and $\delta\gamma \geq 0$, concluding that $x \in \text{conv}(\mathcal{S}' \setminus \{p_k\} \cup \{p\}) \equiv F_k$. Therefore ℓ^+ intersects F_k at $\alpha_f = \alpha_k = 0$.

Remark: The case $\pi_k = \omega_k = 0$ needed to be treated separately, since in such case F_k is perpendicular to F_0 , and so formula (5) could not be applied (we would have 0/0).

Case 2. Since $x' \in \text{conv}(\mathcal{S}' \cup \{p'\})$, from the proof of Lemma 4, there exists $\beta' \geq 0$ such that

$$x' = \sum_{j=1}^{s-1} \beta_j p_j + \beta' p' = \sum_{j=1}^{s-1} (\pi_j - \beta' \omega_j) p_j + \beta' p'. \quad (8)$$

Formula (8) gives all possible ways to represent x' as a convex combination of $\mathcal{S}' \cup \{p'\}$ as a function of β' . We are now interested in knowing the minimum and maximum values of β' , β'_{min} and β'_{max} respectively.

Suppose $\pi_j \geq 0$ for all $j = 1, \dots, s-1$. We have that $x' \in \text{conv}(\mathcal{S}' \cup \{p'\})$ and so $\beta_{min} = 0$. It is easy to see that $\beta_{max} = \frac{\pi_{k_1}}{\omega_{k_1}}$ as in (6), and that any $\beta \in [0, \beta_{max}]$ yields $\pi_j - \beta\omega_j \geq 0$. We then conclude that, when all $\pi_j \geq 0$ we have $x' \in \text{conv}(\mathcal{S}' \setminus \{p_{k_1}\} \cup \{p'\})$, and so from Lemma 5 we have that ℓ intersects F_{k_1} . Note that β_{max} may be 0. When $\beta_{max} > 0$, observe that there is no other $\beta \in]0, \beta_{max}[$ such that $\pi_j - \beta\omega_j = 0$, so there is no other way to write x' as a convex combination of p' and $s-2$ points of \mathcal{S}' .

Now consider that there exists $\pi_j < 0$. Then $x' \notin \text{conv}(\mathcal{S}')$ and therefore $\beta_{min} > 0$. Since β' must be positive, in order to have a convex combination in (8), β' must satisfy the following conditions:

$$\beta' \geq \frac{\pi_j}{\omega_j}, \quad \forall j : \omega_j < 0, \pi_j \leq 0 \Rightarrow \beta' \geq \frac{\pi_{k_2}}{\omega_{k_2}}, \quad (9)$$

$$\beta' \leq \frac{\pi_j}{\omega_j}, \quad \forall j : \omega_j > 0, \pi_j \geq 0 \Rightarrow \beta' \leq \frac{\pi_{k_1}}{\omega_{k_1}}. \quad (10)$$

For k_1 and k_2 as in (6) and (7) respectively. The conditions above are feasible since there must be such a β' , and because of Lemma 4. This allow us to conclude that $\frac{\pi_{k_2}}{\omega_{k_2}} \leq \frac{\pi_{k_1}}{\omega_{k_1}}$ and $\beta_{min} = \frac{\pi_{k_2}}{\omega_{k_2}}$ and $\beta_{max} = \frac{\pi_{k_1}}{\omega_{k_1}}$. Finally we conclude that $x' \in \text{conv}(\mathcal{S}' \setminus \{p_{k_1}\} \cup \{p'\})$ and $x' \in \text{conv}(\mathcal{S}' \setminus \{p_{k_2}\} \cup \{p'\})$, with $k_1 \neq k_2$. So ℓ intersects both facets F_{k_1} and F_{k_2} .

4 The Implementation

Recall the QR factorization of some matrix A of size $n \times m$ into the product of matrices Q and R , where $Q_{n \times n}$ is an orthogonal matrix and $R_{n \times m}$ is upper triangular. For our purposes we will consider $n \geq m$. If A has full column rank, then the diagonal of R is non-zero, and Q can be partitioned in $[V_{n \times m} \ U_{n \times (n-m)}]$ such that the columns of V form a basis of $\text{Span}(A)$, the column range of A , and the columns of U form a basis to $\text{Null}(A)$, the null space of A . Different algorithms are available to find a QR factorization of a matrix, but in terms of computational work, and for a general matrix $n \times n$, they all need $\mathcal{O}(n^3)$ steps [11, §5.2]. However, the QR factorization of A can be “recycled” and used to calculate the QR factorization of a matrix obtained from A by either rank-one changes, appending a row or column to A , or deleting a row or column from A [11, §12.5]. This is accomplished by using *Givens rotations*, and, in the case when $m = n$, such procedure needs $\mathcal{O}(n^2)$ steps.

We now describe in detail how one can implement the algorithm taking advantage of the QR factorization updates. At the beginning of each iteration, we have the QR factorization of the $n \times s$ matrix S , whose columns are the points in \mathcal{S} , the support set found in the previous iteration:

$$S = [p_{i_1} \ p_{i_2} \ \dots \ p_{i_s}].$$

Let Q_S and R_S be the matrices of the QR factorization of S , which were inherited from the previous iteration. The first iteration is the only time a QR factorization from scratch is performed.

4.1 Update \mathcal{S} procedure

The first step of this phase is to check whether $\mathcal{S} \cup \{p\}$ is affinely independent. If it is not, the next step is to find $p_{i_k} \in \mathcal{S}$, such that $\mathcal{S} \setminus \{p_{i_k}\} \cup \{p\}$ is affinely independent.

Consider matrices B and \bar{B} as follows:

$$B = \begin{bmatrix} p_{i_1} & \dots & p_{i_s} \\ 1 & \dots & 1 \end{bmatrix}, \quad \bar{B} = \begin{bmatrix} p_{i_1} & \dots & p_{i_s} & p \\ 1 & \dots & 1 & 1 \end{bmatrix}. \quad (11)$$

B is full column rank since \mathcal{S} is affinely independent. When \mathcal{S} has $n+1$ points we automatically know that $\mathcal{S} \cup \{p\}$ is affinely dependent. Consider then $s \leq n$, which implies that \bar{B} has at least as many

columns as rows. $\mathcal{S} \cup \{p\}$ is affinely independent if and only if \bar{B} is full column rank. This can be checked by looking at the element in position $(s+1, s+1)$ of $R_{\bar{B}}$, from the QR factorization of $\bar{B} = Q_{\bar{B}}R_{\bar{B}}$: if it is zero or close to zero (to account for precision errors) then $\mathcal{S} \cup \{p\}$ is affinely dependent, otherwise it is affinely independent.

Matrix B can be obtained by inserting a row of ones in S , and \bar{B} by then inserting the column $\begin{pmatrix} p \\ 1 \end{pmatrix}$ in B , thus the corresponding QR factorizations can be efficiently computed from the QR factorization of S and of B respectively.

If $\mathcal{S} \cup \{p\}$ is affinely independent, we leave S as is. Otherwise, the linear systems are solved

$$B\pi = \begin{pmatrix} x \\ 1 \end{pmatrix}, \quad \text{and} \quad B\omega = \begin{pmatrix} -p \\ -1 \end{pmatrix}, \quad (12)$$

by reducing them to linear systems with upper triangular matrices, using the QR factorization of B :

$$R_B\pi = Q_B^T \begin{pmatrix} x \\ 1 \end{pmatrix}, \quad R_B\omega = Q_B^T \begin{pmatrix} -p \\ -1 \end{pmatrix}, \quad (13)$$

which can be solved using *Back Substitution* and performed in $\mathcal{O}(n^2)$ [11, §3.1]. Now \mathcal{S} is updated, $\mathcal{S} = \mathcal{S} \setminus \{p_{i_k}\} \cup \{p\}$, for some p_{i_k} , and at this point we get the new Q_S and R_S , the QR factorization of the new matrix S , obtained from the previous one by removing the k -th column. We do not add p yet to S , we leave that to the very end of the iteration, when p is finally covered.

Linear systems (13) are guaranteed to be determined since \mathcal{S} is affinely independent and both x and p are in $\text{aff}(\mathcal{S})$. However, we may encounter degeneracies if some diagonal elements of R_B are close to zero. In order to avoid that, we adopt the same strategy as [8], and so, at the beginning of the iteration, we select a p that is not too close to the affine space of \mathcal{S} .

4.2 Solving MB(\mathcal{S}): calculating d

Consider that at the beginning of this phase $S = [p_{i_1} \dots p_{i_s}]$ and let C be the following matrix:

$$C = [p_{i_1} - p_{i_s} \dots p_{i_{s-1}} - p_{i_s} \ p - p_{i_s}]. \quad (14)$$

Let $C = Q_C R_C$ be the QR factorization of C , obtained by updating the QR factorization of $S = Q_S R_S$ twice, since C can be obtained from S by adding two rank one matrices

$$C = S + (p - p_{i_s})e_n^T - p_{i_s}1_n^T.$$

with $e_j \in \mathbb{R}^n$ is the vector with 1 in the j -th entry and all the other zero. C is full column rank since $\mathcal{S} \cup \{p\}$ is affinely independent, so we can partition $Q_C = [V \ U]$, where V is a $n \times s$ matrix whose columns are a basis to $\text{Span}(C)$, and U is a $n \times (n-s)$ matrix whose columns are a basis to $\text{Null}(C)$. Finally, to find d we need to solve the linear system

$$\begin{bmatrix} C^T \\ U^T \end{bmatrix} d = e_s \quad \iff \quad \left[\begin{array}{c|c} R_C^T & \\ \hline 0_{(n-s) \times s} & I_{n-s} \end{array} \right] Q_C^T d = e_s,$$

where the latter is a lower triangular system.

4.3 Solving MB(\mathcal{S}): calculating the next iterate

The first step now is to calculate x' and p' , the orthogonal projection of x and p respectively onto $\text{aff}(\mathcal{S}')$, where $\mathcal{S}' = \mathcal{S} \setminus \{p\}$. Let D be the following matrix:

$$D = [p_{i_1} - p_{i_s} \ \dots \ p_{i_{s-1}} - p_{i_s}].$$

And let Q_D and R_D be the matrices of the QR factorization of D , which can be obtained easily from the QR factorization of matrix C (14), since D is obtained from C by removing the last column. Let V be the matrix with the first $s-1$ columns of Q_D , which form an orthogonal basis to $\text{Sub}(\mathcal{S}')$. Consider $\text{aff}(\mathcal{S}') = p_{i_1} + \text{Sub}(\mathcal{S}')$. The projections x' and p' can be calculated the following way:

$$x' = VV^T(x - p_{i_1}) + p_{i_1}, \quad p' = VV^T(p - p_{i_1}) + p_{i_1}.$$

Now, two linear systems need to be solved

$$B\pi = \begin{pmatrix} x' \\ 1 \end{pmatrix} \quad \text{and} \quad B\omega = \begin{pmatrix} p' \\ 1 \end{pmatrix}, \quad (15)$$

where B is the same matrix as in (11) but with the points of \mathcal{S}' , so (15) are solved following the same steps used to solve (12).

After finding k_1 and k_2 , as in (6) and (7) respectively, we need to calculate α_{k_1} and α_{k_2} . To calculate α_{k_1} , we need a basis for the null space of $\text{Sub}(\mathcal{S}' \setminus \{p_{k_1}\} \cup \{p\})$, in order to apply formula (5). If $k_1 < s$, $\text{Sub}(\mathcal{S}' \setminus \{p_{i_k}\} \cup \{p\}) \equiv \text{Span}(F)$ with

$$F = [p_{i_1} - p_{i_s} \ \dots \ p_{i_{k_1-1}} - p_{i_s} \ p_{i_{k_1+1}} - p_{i_s} \ \dots \ p_{i_{s-1}} - p_{i_s} \ p - p_{i_s}].$$

Matrix F can be obtained from C (14) by deleting its k_1 -th column. On the other hand, if $k_1 = s$ then $\text{Sub}(\mathcal{S}' \setminus \{p_{i_k}\} \cup \{p\}) \equiv \text{Span}(F)$ with

$$F = [p_{i_1} - p \ \dots \ p_{i_{s-1}} - p],$$

and F can be obtained from C by deleting the last column and then adding the rank one matrix $(p_{i_s} - p)1_n^T$. Therefore, in both cases, the QR factorization of F can be obtained by updating the QR factorization of C . A basis for the null space of $\text{Sub}(\mathcal{S}' \setminus \{p_{k_1}\} \cup \{p\})$ is then formed by the last $n-s+1$ columns of Q_F . The value α_{k_2} is calculated in an analogous way.

The only time a QR factorization of a matrix is calculated from scratch is at the beginning of the algorithm. Then, at each iteration, a constant number of ‘‘QR updates’’ are performed to matrices with n rows and at most n columns, so such updates take $\mathcal{O}(n^2)$ steps. This results in an iteration that is done in quadratic time.

5 Computational results

In order to understand in practice the effect of the new directional procedure together with implementation described in Section 4, we implemented the algorithm in MATLAB. In order to compare it to the original version from Dearing and Zeck [5], we re-wrote their algorithm also using QR updates and implemented it too. Our experiments were conducted using MATLAB R2014a (version 8.3) on a PC with an Intel Core i5 2.30 GHz processor, with 4 GB RAM. Tables 1 and 2 show the average running times of the two versions of the algorithm on instances with 1000 and 10000 points, respectively, drawn uniformly at random from the unit cube.

We observed that for smaller dimensions the original algorithm is slightly faster, which is expected. It is with larger dimensions that we observe that the new version of the algorithm with the changes we proposed in Section 3 is considerably faster. This change naturally does not affect the number of iterations, but only the computational work of each iteration.

Table 1 Average time in seconds for datasets with $m = 1000$ points in variable dimension n uniformly sampled in an unit cube.

Problem		Time in seconds	
n	m	D&Z original	D&Z new
50	1000	0.15	0.32
100	1000	0.42	0.96
500	1000	32.52	20.13
1000	1000	140.45	57.50
5000	1000	8852.20	887.15

Table 2 Average time in seconds for datasets with $m = 10000$ points in variable dimension n uniformly sampled in an unit cube.

Problem		Time in seconds	
n	m	D&Z original	D&Z new
50	10000	2.48	2.95
100	10000	5.03	5.97
500	10000	70.04	42.55
1000	10000	267.68	114.88
5000	10000	17044.10	1463.20

References

1. P. K. Agarwal and R. Sharathkumar. Streaming algorithms for extent problems in high dimensions. *Algorithmica*, 72(1):83–98, 2015.
2. M. Badoiu and K. L. Clarkson. Smaller core-sets for balls. In *Proc. 14th Annual ACM-SIAM Symp. on Discrete Algorithms*, SODA '03, pages 801–802, Philadelphia, PA, USA, 2003. SIAM.
3. M. Badoiu, S. Har-Peled, and P. Indyk. Approximate clustering via core-sets. In *Proc. 34th Annual ACM Symp. on Theory of Computing*, STOC '02, pages 250–257, New York, NY, USA, 2002. ACM.
4. T. M. Chan and V. Pathak. *Streaming and Dynamic Algorithms for Minimum Enclosing Balls in High Dimensions*, pages 195–206. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.
5. P. Dearing and C. R. Zeck. A dual algorithm for the minimum covering ball problem in \mathbb{R}^n . *Oper. Res. Lett.*, 37(3):171–175, 2009.
6. M. Dyer, N. Megiddo, and E. Welzl. *Linear programming*, chapter 45. Chapman and Hall/CRC, Boca Raton, FL, 2nd edition, 2004.
7. K. Fischer and B. Gärtner. The smallest enclosing ball of balls: Combinatorial structure and algorithms. *Internat. J. Comput. Geom. Appl.*, 14(4–5):341–387, 2004.
8. K. Fischer, B. Gärtner, and M. Kutz. Fast smallest-enclosing-ball computation in high dimensions. In *Proc. Algorithms - ESA 2003: 11th Annual European Symp., Budapest, Hungary, September 16-19, 2003*, pages 630–641, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.
9. B. Gärtner. *Fast and Robust Smallest Enclosing Balls*, pages 325–338. Springer Berlin Heidelberg, Berlin, Heidelberg, 1999.
10. B. Gärtner and S. Schönherr. An efficient, exact, and generic quadratic programming solver for geometric optimization. In *Proc. 16th annual ACM Symp. on Computational Geometry (SCG)*, pages 110–118, 2000.
11. G. H. Golub and C. Van Loan. *Matrix Computations (3rd Ed.)*. Johns Hopkins University Press, Baltimore, MD, USA, 1996.
12. T. H. Hopp and C. P. Reeve. An algorithm for computing the minimum covering sphere in any dimension. *Technical Report NISTIR 5831, National Institute of Standards and Technology, Gaithersburg, MD, USA*, 1996.
13. P. Kumar, J. S. B. Mitchell, and E. A. Yildirim. Computing core-sets and approximate smallest enclosing hyperspheres in high dimensions. In *Proc. 5th Workshop on Algorithm Engineering and Experiments*, ALENEX03, pages 44–55. Springer Berlin Heidelberg, 2003.
14. T. Larsson and L. Källberg. Fast and robust approximation of smallest enclosing balls in arbitrary dimensions. In *Proc. 11th Eurographics/ACMSIGGRAPH Symp. on Geometry Processing*, SGP '13, pages 93–101, Aire-la-Ville, Switzerland, Switzerland, 2013. Eurographics Association.
15. N. Megiddo. Linear-time algorithms for linear programming in \mathbb{R}^3 and related problems. *SIAM J. Comput.*, 12(4):759–776, 1983.

16. N. Megiddo. Linear programming in linear time when the dimension is fixed. *J. ACM*, 31(1):114–127, Jan. 1984.
17. F. Nielsen and R. Nock. Approximating smallest enclosing balls with applications to machine learning. *Internat. J. Computat. Geom. Appl.*, 19(05):389–414, 2009.
18. J. J. Sylvester. A question in the geometry of situation. *Quarterly Journal of Pure and Applied Mathematics*, pages 1–79, 1857.
19. E. Welzl. *Smallest enclosing disks (balls and ellipsoids)*, pages 359–370. Springer Berlin Heidelberg, Berlin, Heidelberg, 1991.
20. E. A. Yildirim. Two algorithms for the minimum enclosing ball problem. *SIAM J. Optim.*, 19(3):1368–1391, 2008.
21. H. Zarrabi-Zadeh and T. M. Chan. A simple streaming algorithm for minimum enclosing balls. In *In Proc. 18th Canad. Conf. Computat. Geom*, pages 139–142, 2006.
22. G. Zhou, K.-C. Tohemail, and J. Sun. Efficient algorithms for the smallest enclosing ball problem. *Comput. Optim. Appl.*, 30(2):147–160, 2005.