Computational Geometry: Theory and Applications 132 (2026) 102219

Contents lists available at ScienceDirect

Computational Geometry: Theory and Applications

journal homepage: www.elsevier.com/locate/comgeo

Packing d-dimensional balls into a $d+1$-dimensional container ☆

Helmut Alt [a], Sergio Cabello [b,c,*], Otfried Cheong [d], Ji-won Park [e], Nadja Seiferth [a]

[a] Institut für Informatik, Germany
[b] University of Ljubljana, Slovenia
[c] Institute of Mathematics, Physics and Mechanics, Slovenia
[d] SCALGO, Denmark
[e] Université de Lorraine, CNRS, France

ARTICLE INFO

Article history:
Received 6 August 2024
Received in revised form 5 September 2025
Accepted 7 September 2025
Available online 11 September 2025

Keywords:
Packings
Translations
Unit disks
Unit balls
Volume

ABSTRACT

In this article, we consider the problems of finding in $d+1$ dimensions a minimum-volume axis-parallel box, a minimum-volume arbitrarily-oriented box and a minimum-volume convex body into which a given set of d-dimensional unit-radius balls can be packed under translations. The computational problem is neither known to be NP-hard nor to be in NP. We give a constant-factor approximation algorithm for each of these containers based on a reduction to finding a shortest Hamiltonian path in a weighted graph, which in turn models the problem of stabbing the centers of the input balls while keeping them disjoint. We also show that for n such balls, a container of volume $O(n^{\frac{d-1}{d}})$ is always sufficient and sometimes necessary. As a byproduct, this implies that for $d \geqslant 2$ there is no finite size $(d+1)$-dimensional convex body into which all d-dimensional unit-radius balls can be packed simultaneously.

1. Introduction

Packing a set of geometric objects in a non-overlapping way into a minimum-size container is an intriguing problem. Because of its practical significance, many variants of the problem have been investigated, see the surveys [1,14] and the references therein. Even a simple variant of packing a set of rectangles into a rectangular container turns out to be NP-complete [10]. In many cases, not much is known about the true complexity of the problem.

Constant-factor approximation algorithms of polynomial running time have been found for many variants of packing problems, in particular for finding minimum-size rectangular or convex containers for a set of convex polygons under translations [2], that is, the objects may be translated but rotations are not allowed. Also, approximation algorithms for rigid motions (translations and rotations) are known in this case, see for instance [16].

☆ A preliminary version was presented at WALCOM 2019.

* Corresponding author.
E-mail addresses: alt@mi.fu-berlin.de (H. Alt), sergio.cabello@fmf.uni-lj.si (S. Cabello), otfried@scalgo.com (O. Cheong), ji-won.park@inria.fr (J.-w. Park), nadja.seiferth@fu-berlin.de (N. Seiferth).

https://doi.org/10.1016/j.comgeo.2025.102219

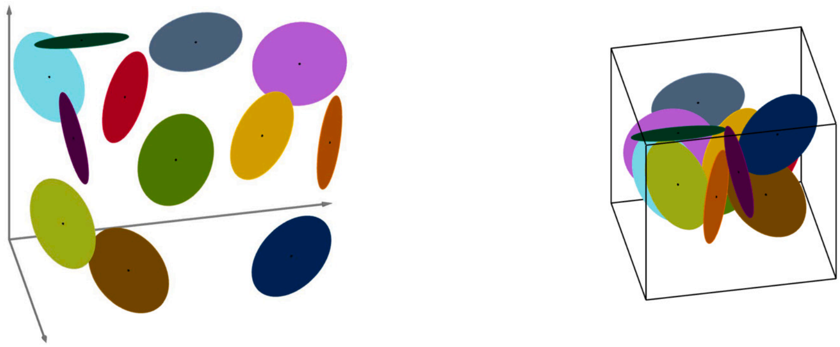

Fig. 1. Input and solution for the DiskPackingAABox-problem.

In three dimensions, most previous results are concerned with "regular" packing problems where objects to be packed are axis-parallel boxes, see for instance [7,12]. In particular, approximation algorithms for packing rectangular cuboids or convex polyhedra into minimum-volume rectangular cuboids or convex containers under rigid motions are known [3]. Whether this is possible under only translations remains open.

In this paper, we provide several results for packing disks of unit radius under translations into minimum-volume containers of different types. In fact, we also handle a higher-dimensional analogue, as follows. Let $d \geqslant 1$ be fixed. We define a *unit hyperdisk* to be a d-dimensional unit-radius ball in $\mathbb{R}^{d+1}$. The task is to pack unit hyperdisks under translations in containers of dimension $d+1$ that may be axis-parallel boxes (see Fig. 1 for an example with $d=2$), boxes that are arbitrarily oriented, or arbitrary convex bodies. In all cases, the objective is to minimize the $(d+1)$-volume of the container. We describe constant-factor approximation algorithms for these problems. Our approximation factors are high and grow quickly with the dimension, but it is of theoretical interest that the problem, whose decision version is neither known to be NP-hard nor to be in NP, can be approximated within a constant factor. Our techniques also show that for n unit hyperdisks, a container of volume $O(n^{\frac{d-1}{d}})$ is always sufficient and sometimes necessary. As a byproduct, this implies that for $d \geqslant 2$ there is no finite size $(d+1)$-dimensional convex container into which all unit hyperdisks can be packed simultaneously.

The input to our problem is a set of n unit hyperdisks in $\mathbb{R}^{d+1}$. Nearly all the hyperdisks we consider are unit, and we often drop the adjective. Since we are allowed to freely translate the hyperdisks, we can assume each hyperdisk to be centered at the origin, so that it is fully defined by its normal vector. Two hyperdisks are *parallel* if their normal vectors are multiples of each other (in particular, a vector and its negative define parallel hyperdisks). A set of *distinct hyperdisks* is a set where no two hyperdisks are parallel.

Two hyperdisks *overlap* if there is a point that lies in the relative interior of both. If they intersect but do not overlap, then they *touch*. By *non-overlapping*, we mean a placement of hyperdisks such that no two hyperdisks overlap, but where hyperdisks are allowed to touch. In other words, we treat hyperdisks as if they consist of their relative interior only. With this, we do not need to consider placements of the hyperdisks where they are arbitrarily close but remain disjoint.

Our main problem is as follows (see Fig. 1):

Definition 1. Given a set of distinct unit hyperdisks in $\mathbb{R}^{d+1}$, find a container of minimum volume in $\mathbb{R}^{d+1}$ such that all hyperdisks can be translated into the container without overlap. The problem is called

- DiskPackingAABox when the container is an axis-parallel box;
- DiskPackingBox when the container is a box of arbitrary orientation;
- DiskPackingConvex when the container is a convex body.

In all cases we also want an actual packing of the hyperdisks inside the container.

For example, the special case where $d=2$ is about packing 2-dimensional unit disks in $\mathbb{R}^3$, while for $d=1$ we talk about packing unit segments in the plane. It is clear that the optimal values of DiskPackingAABox, DiskPackingBox and DiskPackingConvex are monotonically decreasing.

We will approximately solve these packing problems by arranging subsets of hyperdisks such that their centers lie on a line and they are packed as densely as possible. Let O be a given ordering of a set of hyperdisks, and let $s \in \mathbb{R}^{d+1}$ be a vector. We define the *length* of O with respect to direction s as follows: let ℓ be a line with direction vector s. Place the hyperdisks such that their centers lie on ℓ and appear in the ordering O when traversing the line in direction s, and such that two consecutive hyperdisks touch but do not overlap. Then the *length of the ordering* O with respect to s is the distance from the center of the first hyperdisk to the center of the last hyperdisk. Fig. 2 illustrates the definition for $d=1$, where the hyperdisks are line segments.

We can now define the DiskStabbing-problem as follows (see Fig. 3):

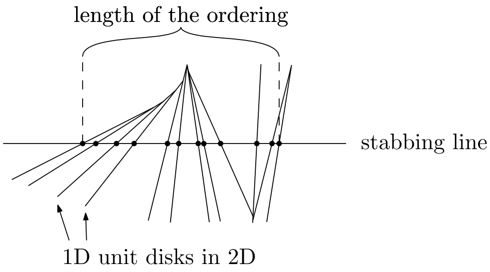

Fig. 2. The length of an ordering of 1-dimensional hyperdisks.

Fig. 3. Input and solution for the DiskStabbing-problem with $d = 2$.

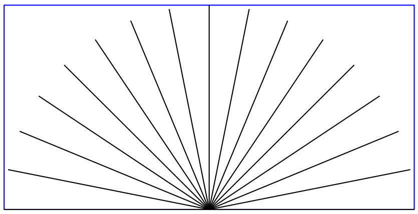

Fig. 4. All unit-length segments can be packed into a container of area 2.

Definition 2 *(*DiskStabbing*).* Given a set of distinct unit hyperdisks in $\mathbb{R}^{d+1}$, and an additional vector $s \in \mathbb{R}^{d+1}$ defining the direction of a line, find an ordering of the hyperdisks that minimizes the length with respect to s.

In Section 2, we study the geometry of the DiskStabbing-problem and define a metric on the family of unit hyperdisks. To show that we indeed have a metric, we will make use of the fact that we are considering *unit* hyperdisks; for hyperdisks of different sizes, the same approach does not define a metric. An important consequence of this metric is that each solution to the DiskStabbing-problem gives a non-overlapping packing of the unit hyperdisks. Note that this is not obvious because the definition of DiskStabbing only requires that consecutive hyperdisks in the ordering are non-overlapping, but in principle non-consecutive hyperdisks in the ordering could still overlap. This property will allow us to prove that an optimal solution to DiskStabbing implies an approximate solution to DiskPackingAABox. We prove this first for hyperdisks with similar normal vectors in Section 3, then for the general case in Section 4.

In Section 5 we then show how to approximately solve DiskStabbing by (approximately) computing a shortest Hamiltonian path in a complete weighted graph, implying a constant-factor approximation for DiskPackingAABox.

In Section 6 we consider the DiskPackingBox and DiskPackingConvex problems, where we reuse several of the ideas employed for the case of axis-parallel containers.

For $d = 1$, all (infinitely many) unit-length line segments can be packed into a rectangle of area 2 (Fig. 4), and even smaller containers are possible. In Section 7 we show that a similar result does not hold for $d \geqslant 2$: there is no bounded-volume convex container into which all d-dimensional hyperdisks can be packed. More precisely, we consider a family $\mathcal{D}$ of n hyperdisks in $\mathbb{R}^{d+1}$, a stabbing direction s, and an angle $\phi < \pi/2$, such that the normal of each disk in $\mathcal{D}$ makes angle at most ϕ with s. We show that there is then a stabbing of $\mathcal{D}$ in direction s of length $O(n^{\frac{d-1}{d}})$ (and therefore also a packing of this volume). We also construct such a family $\mathcal{D}$ where this bound is tight.

The reader may wish to contrast this with the *Potato Sack Theorem* of Auerbach et al. [4], see Kosiński [13] for a proof and the survey by Fejes Tóth [8] for a general discussion. It states that when rigid motions are allowed, any sequence of convex bodies in $\mathbb{R}^d$ can be packed into a box of finite volume, as long as there is a common bound on the diameter of the bodies and their total volume is bounded.

In Section 8 we briefly consider packing other fixed shapes like squares and mention some open problems.

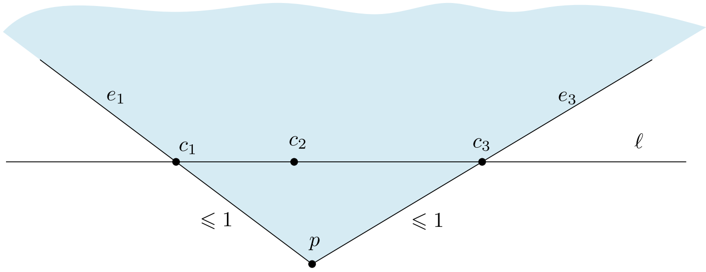

Fig. 5. Proof of Theorem 3, first step. The point p is not necessarily an endpoint of the segments e_1 and e_3. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

2. A metric on unit hyperdisks

Given a vector $s \in \mathbb{S}^d$, we define the *s-distance* $d_s(D_1, D_2)$ between two hyperdisks D_1 and D_2 with centers c_1 and c_2 as the length of the ordering D_1, D_2 with respect to s. In other words, it is the distance $|c_1c_2|$ when the hyperdisks touch and the ray $\overrightarrow{c_1c_2}$ has direction s. When D_1 and D_2 are parallel, we define $d_s(D_1, D_2) = 0$.

Theorem 3. *For any $s \in \mathbb{S}^d$, the distance d_s is a metric on the set of d-dimensional unit hyperdisks with normals not orthogonal to s (and where parallel hyperdisks are considered equivalent).*

Proof. If D_1 and D_2 are not parallel, it is clear that $d_s(D_1, D_2) > 0$. Otherwise, $d_s(D_1, D_2) = 0$ by definition.

Symmetry holds since a point reflection keeps both hyperdisks identical (their normal vectors are negated), but reverses the order in which a line with direction s meets them.

It remains to prove the triangle inequality $d_s(D_1, D_3) \leqslant d_s(D_1, D_2) + d_s(D_2, D_3)$. It clearly holds when any two of the three hyperdisks are parallel, so we will assume for a contradiction that there are three distinct hyperdisks D_1, D_2, D_3 with centers c_1, c_2, c_3 such that $d_s(D_1, D_3) > d_s(D_1, D_2) + d_s(D_2, D_3)$. We place c_1, c_2, c_3 in this order on a line ℓ with direction s such that $|c_1c_2| = d_s(D_1, D_2)$ (so the hyperdisks touch but do not overlap) and that $|c_1c_3| = d_s(D_1, D_3)$ (so these hyperdisks touch in a point p and do not overlap). Since $|c_2c_3| = d_s(D_1, D_3) - d_s(D_1, D_2) > d_s(D_2, D_3)$, the hyperdisks D_2 and D_3 do not intersect at all.

The point p does not lie on the line ℓ because the normals of the hyperdisks D_1 and D_3 are not orthogonal to s. We consider the situation in the (two-dimensional) plane h containing the line ℓ and the point p. Since $c_i \in h$ and the normal of D_i is not orthogonal to s, the intersection $e_i := D_i \cap h$ is a segment of length two centered at c_i, for $i \in \{1, 2, 3\}$. We choose a coordinate system for h where ℓ is the x-axis and p lies below ℓ, see Fig. 5. By assumption, the segment e_2 has length two, its midpoint is c_2, and it does not intersect e_1 or e_3.

We first observe that if $\measuredangle pc_1c_2$ or $\measuredangle pc_3c_2$ is obtuse, then the entire triangle c_1pc_3 lies inside a circle of radius 1 around c_2 because in this case $|c_2p| \leqslant \max\{|c_1p|, |c_3p|\} \leqslant 1$ and $|c_1c_2| \leqslant \max\{|c_1p|, |c_3p|\} \leqslant 1$, so e_2 would have to intersect e_1 or e_3, a contradiction. It follows that $\measuredangle pc_1c_2$ and $\measuredangle pc_3c_2$ are acute, and e_2 is therefore contained in the wedge bounded by the rays pc_1 and pc_2 (light blue in Fig. 5).

We can therefore move c_1 to the left and c_3 to the right while keeping c_2 and p unchanged, until $|c_1p| = |pc_3| = 1$. This can only grow the wedge, so it will still contain e_2. The triangle c_1pc_3 is isosceles, as shown in Fig. 6. Let u be the upper, v the lower endpoint of e_2. Clearly v lies inside the triangle c_1pc_3. Without loss of generality, we assume that e_2 has a negative slope, so that v is the *right* endpoint of e_2.

Consider now the segment e_2' parallel to e_2 centered at c_1, with upper endpoint u' and lower endpoint v'. The points p, u', v' all lie on the circle C of radius 1 around c_1, see Fig. 6.

Since v lies above p, so does v', and therefore v' lies on the counterclockwise arc of C from p to ℓ. Let now u_1 be the point on e_1 on the horizontal line $u'u$, and let v_i be the point on e_i on the horizontal line $v'v$, for $i \in \{1, 3\}$, see Fig. 6. Since v' lies to the right of p, we have $|v'v_3| < |v'v_1| = |u'u_1|$. However, $|u'u| = |v'v|$, so $v \in \triangle c_1pc_3$ implies $|v'v| < |v'v_3|$, so $|u'u| < |v'v_3| < |u'u_1|$, and u lies outside the wedge bounded by the rays pc_1 and pc_2, a contradiction. $\square$

As mentioned in the Introduction, Theorem 3 implies that any solution to the DiskStabbing-problem for unit hyperdisks gives a non-overlapping set of hyperdisks. In fact, this has been the essence of its proof.

Consider now a configuration of two hyperdisks D_1 and D_2 that realize the s-distance $d_s(D_1, D_2)$ for a given direction vector s. We let n_i denote the normal of hyperdisk D_i, c_i its center, and h_i the hyperplane supporting D_i. Let g be the $(d-1)$-dimensional affine subspace $h_1 \cap h_2$ (for $d = 2$ this is a line with direction vector $n_1 \times n_2$). The two hyperdisks must touch in a point $p \in g$. Let $\bar{h}_i$ denote the half-hyperplane of h_i delimited by g that contains c_i. Note that the intersection of g and the hyperdisk D_i is either the point p or a $(d-1)$-dimensional ball of positive radius.

We can distinguish three cases that cover all possibilities (but (i) and (ii) do not exclude each other):

(i) D_1 lies entirely in $\bar{h}_1$. In this case D_1 is tangent to g in p, $|pc_1| = 1$ and $|pc_2| \leqslant 1$;

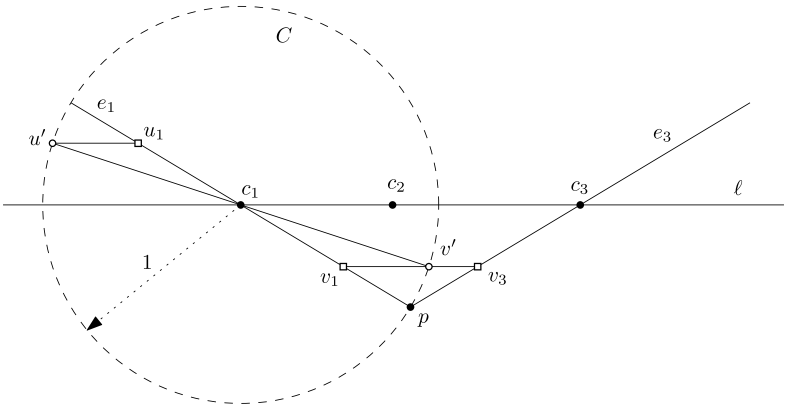

Fig. 6. Proof of Theorem 3 (final step).

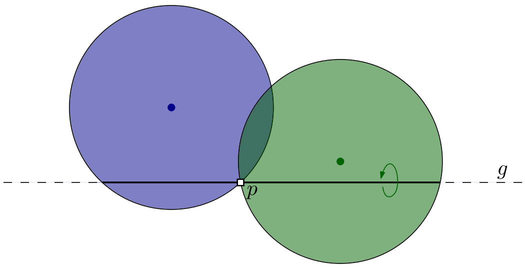

Fig. 7. Example illustrating case (iii) for $d = 2$, which is embedded in $\mathbb{R}^3$. The disk on the right is rotated around the line g in $\mathbb{R}^3$ and then both disks touch only at p.

(ii) D_2 lies entirely in $\bar{h}_2$. In this case D_2 is tangent to g in p, $|pc_2| = 1$ and $|pc_1| \leqslant 1$;

(iii) D_1 does not lie entirely in $\bar{h}_1$ and D_2 does not lie entirely in $\bar{h}_2$. In this case $|c_1 p| = |pc_2| = 1$ and g intersects both hyperdisks in $(d-1)$-dimensional balls of positive radius that have disjoint interiors, and therefore are tangent at p.

The last case may be harder to visualize. In $\mathbb{R}^{d+1}$ such case can be constructed as follows, see Fig. 7 for $d = 2$. Take two hyperdisks in $\mathbb{R}^d$ that intersect, let p be a point where the boundaries of the hyperdisks intersect, let g be a $(d-1)$-dimensional affine subspace through p that does not contain other intersection points of the hyperdisks and that goes through the interior of both hyperdisks, and rotate one of the hyperdisks around g. The point p remains the touching point of the hyperdisks and g intersects the interior of both hyperdisks.

We have the following lemmas.

Lemma 4. *Let D_1 and D_2 be two unit hyperdisks whose normal vectors form an angle ξ. Then $d_s(D_1, D_2) \geqslant \sin\xi$.*

Proof. The angle formed by $\bar{h}_1$ and $\bar{h}_2$ is either ξ or $\pi - \xi$. By possibly negating n_2, we can assume the former. The claim is then obvious in case (i), as c_1 has distance $\sin\xi$ from the hyperplane h_2, and case (ii) is symmetric.

In case (iii) we consider the isosceles triangle $\triangle c_1 p c_2$ and set $\gamma := \measuredangle c_1 p c_2$. We will show $\gamma \geqslant \xi$, so we have $|c_1 c_2| = 2\sin(\gamma/2) \geqslant 2\sin(\xi/2) \geqslant \sin\xi$.

To see why $\gamma \geqslant \xi$, let $g^\perp$ be the 2-dimensional affine subspace orthogonal to g and through p. Let c_i' be the orthogonal projection of c_i on g, and let $c_i^\perp$ be the orthogonal projection of c_i on $g^\perp$. See Fig. 8. Since we are using orthogonal projections on orthogonal affine subspaces through p, we have $c_i - c_i^\perp = c_i' - p$. Moreover, the vectors $(c_i - c_i')$ and $(c_j' - p)$ are orthogonal for $i, j \in \{1, 2\}$. We will also use the fact that the orthogonal projection of a segment cannot be longer than the original segment. Therefore $|c_i^\perp p| \leqslant |c_i p| = 1$ and $|c_1' c_2'| \leqslant |c_1 c_2|$.

Since $D_1 \cap g$ and $D_2 \cap g$ are interior disjoint $(d-1)$-dimensional balls of positive radius tangent at p and contained in the $(d-1)$-dimensional subspace g, their centers have to be collinear with p. Noting that c_i' is the center of $D_i \cap g$, we get that the points c_1', p and c_2' are collinear and in that order. Therefore $(c_1' - p) \cdot (c_2' - p) < 0$.

If $c_i' = c_i$, then $p = c_i^\perp$ and thus $|c_1 c_2| \geqslant |c_1' c_2'| \geqslant |pc_i'| = 1 \geqslant \sin\xi$. Therefore, we consider only the case where $c_1' \neq c_1$ and $c_2' \neq c_2$, which means that $p \neq c_1^\perp$ and $p \neq c_2^\perp$. We then have $\measuredangle c_1^\perp p c_2^\perp = \xi$ because rotating D_2 an angle of ξ around g would place D_2 in the hyperplane supporting D_1. We have

$$(c_1^\perp - p) \cdot (c_2^\perp - p) = |c_1^\perp p| \cdot |c_2^\perp p| \cdot \cos\xi \leqslant |c_1 p| \cdot |c_2 p| \cdot \cos\xi = \cos\xi,$$

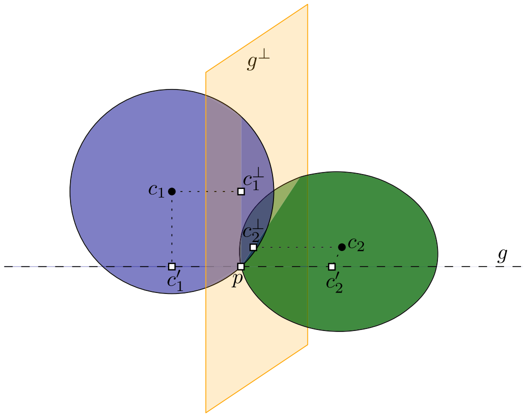

Fig. 8. Proof of Lemma 4 for case (iii).

and therefore

$$\begin{aligned}\cos\gamma &= (c_1 - p)\cdot(c_2 - p)\\ &= \big((c_1 - c_1') + (c_1' - p)\big)\cdot\big((c_2 - c_2') + (c_2' - p)\big)\\ &= (c_1 - c_1')\cdot(c_2 - c_2') + (c_1' - p)\cdot(c_2' - p) \qquad \big[\text{because } (c_i - c_i')\cdot(c_j' - p) = 0\big]\\ &< (c_1^\perp - p)\cdot(c_2^\perp - p) + 0\\ &\leqslant \cos\xi. \quad \square\end{aligned}$$

Lemma 5. *Let D_1 and D_2 be two unit hyperdisks whose normal vectors form an angle ξ, and such that their normals make an angle of at most $\phi < \pi/2$ with the direction s. Then $d_s(D_1, D_2) \leqslant \frac{\sin\xi}{\cos\phi}$.*

Proof. Again we can assume that the angle between $\bar{h}_1$ and $\bar{h}_2$ is ξ. This means that c_2 has distance $t \leqslant \sin\xi$ from h_1. Let c_2' be the point at distance t on the ray from c_1 with direction n_1. Since the ray c_1c_2 has direction s, then $\measuredangle c_2'c_1c_2 \leqslant \phi$, and it follows that $t/|c_1c_2| = \cos(\measuredangle c_2'c_1c_2) \geqslant \cos\phi$. This implies $|c_1c_2| \leqslant t/\cos\phi$ and the claim follows. $\square$

3. Stabbing helps to pack unit hyperdisks with similar normal vectors

We define the angle $\phi_0 = \arccos(1/\sqrt{d+1})$. (For $d = 2$, $\phi_0 \approx 54.7°$.) This is the angle between the diagonal and an edge of the $(d+1)$-dimensional hypercube. In this section we consider a set $\mathcal{D}$ of unit hyperdisks whose normal vector makes an angle of at most ϕ_0 with the positive x_{d+1}-axis, that is, with the vector $(0, 0, \dots, 0, 1)$. We will show that the DiskStabbing and DiskPackingAABox problems are related by proving that an optimal solution to DiskStabbing for $\mathcal{D}$ with respect to direction $s = (0, 0, \dots, 0, 1)$ provides a constant-factor approximation to DiskPackingAABox for $\mathcal{D}$. In the next section we will then extend this to arbitrary sets of unit hyperdisks.

Let us first define E to be the maximum *extent* of all the hyperdisks. Formally, let E be the smallest axis-parallel box that contains all the hyperdisks if we place them with their center at the origin, and let $E_1, E_2, \dots, E_{d+1}$ be the dimensions of E. We have

$$2 \geqslant E_i \geqslant 2\sin(\pi/2 - \phi_0) = 2\cos\phi_0 = 2/\sqrt{d+1}, \qquad \text{for } 1 \leqslant i \leqslant d.$$

Let now M denote a minimum-volume axis-parallel box into which $\mathcal{D}$ can be packed—that is, an optimal solution to DiskPackingAABox—and denote its dimensions by $M_1, M_2, \dots, M_{d+1}$. Its volume is $\textsc{opt} = M_1M_2\dots M_{d+1}$.

When we project M onto a hyperplane normal to the x_{d+1}-axis, we obtain a d-dimensional box of size $M_1 \times M_2 \times \dots \times M_d$. We place a grid of $n_1 \times n_2 \times \dots \times n_d$ points inside this box with grid size $\mu = 1/\sqrt{d(d+1)}$, starting at the point at distance $1/\sqrt{d+1} + \mu/2$ from the boundary, and stopping when we have passed the same distance from the opposite boundary. See the left side of Fig. 9. We have

$$n_i = \left\lceil \frac{M_i - \frac{2}{\sqrt{d+1}}}{\mu} \right\rceil < \frac{M_i - \frac{2}{\sqrt{d+1}} + \mu}{\mu} < \frac{M_i}{\mu}, \qquad \text{for } 1 \leqslant i \leqslant d.$$

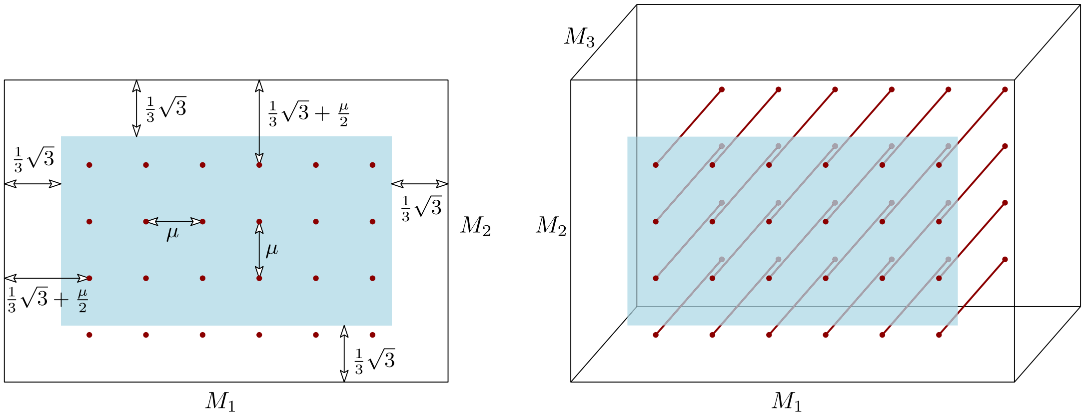

Fig. 9. Left: Placing a grid on a d-dimensional face of the optimal container M, for $d = 2$. Right: the vertical segments stabbing the hyperdisks.

We observe that since the hyperdisks lie entirely in M, the projection of a hyperdisk center on the box must have distance at least $1/\sqrt{d+1}$ from the box boundary (that is, it must lie in the shaded region of Fig. 9), so it has distance at most $\frac{\mu}{2}\sqrt{d}$ to the nearest grid point.

For each of the $n_1 \times \cdots \times n_d$ grid points, we consider the *vertical* segment of length M_{d+1} obtained by intersecting M with a line parallel to the x_{d+1}-axis through the grid point. See the right side of Fig. 9. By our construction, every disk center has distance at most $\frac{\mu}{2}\sqrt{d}$ to one of these segments. This means that such a segment stabs the hyperdisk in a point q at distance at most $\frac{\mu}{2}\frac{\sqrt{d}}{\sin(\pi/2-\phi_0)} = \frac{\mu}{2}\sqrt{d(d+1)} = \frac{1}{2}$ from the disk center. The smaller disk of radius $1/2$ around q lies entirely inside the original disk, and we replace each original disk with this smaller disk. With this, we have obtained n hyperdisks of radius $1/2$ that are packed inside the same box and such that each of them is pierced through the center by one of the vertical segments through the grid points.

We can now place all $n_1 n_2 \cdots n_d$ vertical segments behind each other and obtain a hyperdisk stabbing of length $n_1 n_2 \cdots n_d M_{d+1}$ for $\mathcal{D}$, where the hyperdisks have now radius $1/2$. Simply enlarging this stabbing by factor 2 results in a solution to DiskStabbing for $\mathcal{D}$ of length

$$\begin{aligned} 2 \cdot n_1 n_2 \cdots n_d \cdot M_{d+1} &< 2 \cdot \frac{M_1}{\mu}\frac{M_2}{\mu}\cdots\frac{M_d}{\mu} \cdot M_{d+1} \\ &= \frac{2}{\mu^d} \cdot M_1 M_2 \cdots M_{d+1} = \frac{2}{\mu^d}\textsc{opt} \\ &= 2\sqrt{\big(d(d+1)\big)^d}\textsc{opt} < 2(d+1)^d\textsc{opt}. \end{aligned}$$

Theorem 6. *Let $\mathcal{D}$ be a family of unit hyperdisks whose normals make an angle of at most $\phi_0 = \arccos(1/\sqrt{d+1})$ with the x_{d+1}-axis. An optimal solution to* DiskStabbing *for $\mathcal{D}$ with respect to the x_{d+1}-direction has length at most $2(d+1)^d \cdot \textsc{opt}$, where* opt *is the volume of an optimal solution to* DiskPackingAABox *for $\mathcal{D}$. The* DiskStabbing*-solution thus provides a $((2^{d+1}(d+1)^d + 1))$-approximation for* DiskPackingAABox.

Proof. Since an optimal solution to DiskPackingAABox implies a stabbing of length at most $2(d+1)^d \cdot \textsc{opt}$, an optimal stabbing has at most this length, implying the first claim.

The smallest axis-parallel box enclosing the stabbing is a valid packing of the hyperdisks. Here we make use of the triangle inequality from Theorem 3: if consecutive hyperdisks do not overlap, then all hyperdisks are non-overlapping. If L_{d+1} is the length of the stabbing, its volume is at most $E_1 E_2 \cdots E_d (L_{d+1} + E_{d+1})$ (the E_{d+1}-term is needed because we measure the length of a stabbing only from the first to the last center). Since $E_i \leqslant 2$ for $1 \leqslant i \leqslant d$ and $L_{d+1} \leqslant 2(d+1)^d \cdot \textsc{opt}$, this is at most

$$2^d \cdot 2(d+1)^d \cdot \textsc{opt} + E_1 E_2 \cdots E_{d+1} \leqslant (2^{d+1}(d+1)^d + 1) \cdot \textsc{opt},$$

where we made use of the fact that $\textsc{opt} \geqslant E_1 E_2 \cdots E_{d+1}$. $\square$

4. Stabbing helps to pack general unit hyperdisks inside an axis-parallel box

We now address the general DiskPackingAABox problem. Given a set $\mathcal{D}$ of unit hyperdisks, we partition it into $d+1$ subsets $\mathcal{D}_1, \mathcal{D}_2, \ldots, \mathcal{D}_{d+1}$, depending on which coordinate axis the normal vector makes the smallest angle with (this corresponds simply to the highest absolute coordinate value in the normal vector). We then consider an optimal stabbing for $\mathcal{D}_i$ with respect to the x_i-axis, for $1 \leqslant i \leqslant d+1$. Note that the normals for each subfamily make an angle of at most ϕ_0 with its corresponding axis.

Theorem 7. *Optimal stabbings for $\mathcal{D}_i$ with respect to the x_i-axis, for $1 \leqslant i \leqslant d+1$, can be used to provide a $3\big(3d+3\big)^{d+1}$-approximation for* DiskPackingAABox *for $\mathcal{D}$.*

Proof. Again we define E (of dimension $E_1 \times E_2 \times \cdots \times E_{d+1}$) to be the *extent* of all the hyperdisks (in the entire set $\mathcal{D}$). If $\mathcal{D}_i = \emptyset$ for all but one index i, then we are done by Theorem 6, so we assume this is not the case. This implies that for all extent dimensions we have

$$2/\sqrt{d+1} \leqslant E_i \leqslant 2.$$

Since $\text{OPT} \geqslant E_1 E_2 \cdots E_{d+1}$, we have $\text{OPT} \geqslant 2^{d+1}/(d+1)^{(d+1)/2}$, and therefore $(d+1)^d \text{OPT} \geqslant 1$.

Consider now an index $i \in \{1, 2, \ldots, d+1\}$ where $\mathcal{D}_i \neq \emptyset$. Let ℓ^* be the length of an optimal stabbing for the set $\mathcal{D}_i$ in the direction of the x_i-axis, and let s be a segment of length ℓ^* parallel to the x_i-axis that stabs the hyperdisks in $\mathcal{D}_i$. We partition s into $\lceil \ell^* \rceil$ pieces of length at most one. For each piece, we take the smallest axis-parallel box containing the hyperdisks whose centers lie on the piece. Each box has dimensions at most $E_1 \times \cdots \times E_{i-1} \times (1+E_i) \times E_{i+1} \times \cdots \times E_{d+1}$, so it fits into a hypercube of side length 3. The number of these hypercubes needed for $\mathcal{D}_i$ is

$$\lceil \ell^* \rceil \leqslant \ell^* + 1 \leqslant 2(d+1)^d \cdot \text{OPT} + 1 \leqslant 3(d+1)^d \cdot \text{OPT}.$$

Repeating this for all $d+1$ sets D_i, we find that we can pack $\mathcal{D}$ into

$$(d+1) \cdot 3 \cdot (d+1)^d \cdot \text{OPT} = 3 \cdot (d+1)^{d+1} \cdot \text{OPT}$$

hypercubes of side length 3, with a total volume of

$$3^{d+1} \cdot 3 \cdot (d+1)^{d+1} \cdot \text{OPT} = 3 \cdot \big(3d+3\big)^{d+1} \cdot \text{OPT}.$$

We can stack the hypercubes along any coordinate direction to obtain an axis-parallel box packing $\mathcal{D}$ of the same volume. Note that, at the cost of an additional constant factor of, say, 2^{d+1}, we could even pack $\mathcal{D}$ into a hypercube. This is for example obtained by doubling each dimension of the hypercube container whenever it becomes too small to continue placing the cubes of side length 3. □

5. Approximating the optimal stabbing and packing inside an axis-parallel box

It remains to show how to actually solve the DiskStabbing problem.

Theorem 8. *A $\frac{3}{2}$-approximation for* DiskStabbing *can be computed in polynomial time.*

Proof. Given a set of n unit hyperdisks by their normal vectors and a direction vector s, we construct the complete weighted graph G on n vertices $\{1, 2, \ldots, n\}$, where the weight of the edge (i, j) is $d_s(D_i, D_j)$. The Hamiltonian paths in this graph correspond one-to-one to the orderings of the hyperdisks, and the length of the path in the graph is equal to the length of the ordering with respect to s. It follows that the optimal ordering is given by the shortest Hamiltonian path in G.

Christofides' approximation algorithm applies to our problem, since by Theorem 3 the triangle inequality holds in G. We can therefore compute a Hamiltonian path for G whose length is at most $3/2$ times the optimal [11]. □

We then have the following corollary.

Corollary 9. *A $4\big(3d+3\big)^{d+1}$-approximation to* DiskPackingAABox *can be computed in polynomial time.*

Proof. In the proof of Theorem 7, we stab the hyperdisks of $\mathcal{D}_i$ using a segment of length $\ell \leqslant \frac{3}{2}\ell^*$, so the number of hypercubes of side length 3 needed for $\mathcal{D}_i$ is

$$\lceil \ell \rceil \leqslant \ell + 1 \leqslant \tfrac{3}{2}\ell^* + 1 \leqslant \tfrac{3}{2}\big(2(d+1)^d \cdot \text{OPT}\big) + 1 \leqslant 4(d+1)^d \cdot \text{OPT}.$$

Repeating this for all $d+1$ sets D_i, the number of hypercubes of side length 3 packing $\mathcal{D}$ is at most

$$(d+1)4(d+1)^d \cdot \text{OPT} = 4(d+1)^{d+1} \cdot \text{OPT},$$

with a total volume of $4 \cdot \big(3d+3\big)^{d+1} \cdot \text{OPT}$. □

6. Packing inside a convex container or a box with arbitrary orientation

We now turn our attention to packing unit hyperdisks under translations inside a convex container of minimum volume or a box of minimum volume, where we can freely select the orientation of the box. Again, our aim is to get a constant-factor approximation.

6.1. Approximating a convex body by a box

We will need a geometric lemma that shows that finding a box of (approximately) minimum volume where we can pack the hyperdisks suffices to get a constant-factor approximation to the (approximately) smallest convex container. We phrase the lemma in $\mathbb{R}^d$, even though later we will use it in $\mathbb{R}^{d+1}$. Similar arguments have been used in the literature before, we include the lemma for completeness. The factor $d^{3d/2}$ is not tight. In two dimensions, for instance, it is known that factor 2 is possible [15].

Lemma 10. *Any convex body K in $\mathbb{R}^d$ is contained in a box of volume at most $d^{3d/2} \cdot \mathrm{vol}(K)$.*

Proof. Let $B(r)$ denote the ball in $\mathbb{R}^d$ centered at the origin with radius r.

It is a classical result in convexity that there is a unique ellipsoid E of maximum volume contained in K. Moreover, scaling E around its center by d results in an ellipsoid that contains K. This ellipsoid is called the *Löwner-John ellipsoid*; see for example the books by Barvinok [5, Chapter 5, Section 2] or Ben-Tal and Nemirovski [6, Theorem 4.9.1]. Therefore, we have $\mathrm{vol}(E) \leqslant \mathrm{vol}(K) \leqslant d^d\,\mathrm{vol}(E)$. We choose a coordinate system with origin at the center of E and aligned with the axes of E. Scaling along each coordinate axis, we can transform E into the unit-radius ball $B(1)$. Let K' be the image of K under this scaling.

We then have $B(1) \subseteq K' \subseteq B(d)$. Inside the ball $B(1)$ we inscribe the axis-parallel cube Q_{in} of side length $2/\sqrt{d}$, around the ball $B(d)$ we circumscribe the axis-parallel cube Q_{out} of side length $2d$. Therefore, we have obtained $Q_{\mathrm{in}} \subseteq K' \subseteq Q_{\mathrm{out}}$ such that

$$\mathrm{vol}(Q_{\mathrm{out}}) = (2d)^d = d^{3d/2} \cdot \left(\frac{2}{\sqrt{d}}\right)^d = d^{3d/2} \cdot \mathrm{vol}(Q_{\mathrm{in}}) \leqslant d^{3d/2} \cdot \mathrm{vol}(K').$$

Applying the inverse scaling along each coordinate axis, such that $B(1)$ is transformed back into E, we obtain the desired box because an affine transformation keeps ratios of volumes invariant. □

6.2. The algorithm

We will need another lemma that allows us to replace any box by an axis-parallel box, assuming only that the original box is not too thin. We assume $d \geqslant 2$.

Lemma 11. *Assume there is a box K of arbitrary orientation where $\mathcal{D}$ can be packed, and assume that each side of K has length at least $1/2$. Then there is an axis-parallel box K' where $\mathcal{D}$ can be packed with*

$$\mathrm{vol}(K') \leqslant (2d+4)^{d+1} \cdot \mathrm{vol}(K).$$

Proof. We cover K with hypercubes of side length 1 that are aligned with K. The total number of these hypercubes is at most $2^{d+1}\,\mathrm{vol}(K)$, since we need to round up each side length to the nearest integer, at most doubling the side length.

Consider now one such small hypercube Q. Its diagonal has length $\sqrt{d+1}$, therefore the projection of Q on any axis has length at most $\sqrt{d+1} < d$. We can therefore enclose Q in an axis-parallel hypercube Q' of side length $d+2$, such that every unit hyperdisk whose center lies in Q is completely contained in Q'.

It follows that $\mathcal{D}$ can be packed into the union of the hypercubes Q'. The total volume of all the hypercubes Q' is

$$(d+2)^{d+1} \cdot 2^{d+1} \cdot \mathrm{vol}(K) = (2d+4)^{d+1}\,\mathrm{vol}(K),$$

and we can stack the hypercubes Q' along any coordinate direction to obtain an axis-parallel box of this volume. □

Theorem 12. *A $4\big(3d+3\big)^{2d+2}$-approximation to* DiskPackingBox *can be computed in polynomial time.*

Proof. We select an arbitrary disk D_0 from $\mathcal{D}$ and choose a coordinate system where the normal of D_0 is the x_{d+1}-axis. We then use Corollary 9 to compute an axis-parallel box with respect to this coordinate system where $\mathcal{D}$ can be packed.

Let K_0 be the box computed by the algorithm, and let K^* be a minimum-volume box (in an arbitrary orientation) where $\mathcal{D}$ can be packed. Let $\textsc{opt}_B = \mathrm{vol}(K^*)$ be the optimal volume.

If each side of K^* has length at least 1/2, then by Lemma 11 there is an axis-parallel (with respect to our coordinate system based on D_0) box K' of volume $(2d+4)^{d+1} \cdot \text{OPT}_B$ where $\mathcal{D}$ can be packed. Corollary 9 guarantees that

$$\text{vol}(K_0) \leqslant 4(3d+3)^{d+1}\,\text{vol}(K') \leqslant 4(3d+3)^{d+1} \cdot (2d+4)^{d+1} \cdot \text{OPT}_B < 4(3d+3)^{2d+2} \cdot \text{OPT}_B.$$

It remains to consider the case where K^* has a side of length smaller than 1/2. We will denote both the direction and the length of this side by s^*. From $s^* < 1/2$ follows that the normal of every hyperdisk in $\mathcal{D}$ makes an angle at most $\vartheta = \arcsin \frac{s^*}{2}$ with the direction s^*. We have $\vartheta < \arcsin \frac{1}{4} \approx 14.48° < \phi_0/2$. (Recall that we defined the angle $\phi_0 = \arccos(1/\sqrt{d+1}) \geqslant \arccos(1/\sqrt{3}) \approx 54.74°$.)

Let now s_0 be the direction of the normal of D_0, that is, the x_{d+1}-axis in our chosen coordinate system. Then the normal of any hyperdisk in $\mathcal{D}$ makes an angle of at most $2\vartheta < \phi_0$ with s_0. It follows that the algorithm of Corollary 9 will compute a single disk stabbing for all disks in $\mathcal{D}$ with respect to the direction s_0, and will then fit a box to this disk stabbing.

Consider again direction s^*. Since the normal of every disk makes an angle of at most $\vartheta < \phi_0$ with s^*, using a suitable coordinate system and Theorem 6 implies that an optimal disk stabbing with direction s^* has length at most $2(d+1)^d \text{OPT}_B$. Let $D_1, \ldots, D_n$ be the ordering of such an optimal disk stabbing for direction s^*. (D_0 is one of these disks, possibly in the middle.) By Lemma 4, we have $d_{s^*}(D_i, D_{i+1}) \geqslant \sin \xi_i$, where ξ_i is the angle between the normals of D_i and D_{i+1}. On the other hand, by Lemma 5, we have $d_{s_0}(D_i, D_{i+1}) \leqslant \sin \xi_i / \cos(2\vartheta) \leqslant d_{s^*}(D_i, D_{i+1})/\cos(2\vartheta)$. Since $\cos(2\vartheta) = 1 - 2\sin^2\vartheta = 1 - 2(\frac{s^*}{2})^2 \geqslant \frac{7}{8}$, we have $d_{s_0}(D_i, D_{i+1}) \leqslant \frac{8}{7} d_{s^*}(D_i, D_{i+1})$, so the ordering $D_1, \ldots, D_n$ gives us a disk stabbing for direction s_0 of length at most $\frac{16}{7}(d+1)^d \text{OPT}_B$.

This implies that the algorithm, which computes a 3/2-approximation to the optimal disk stabbing in direction s_0, returns a disk stabbing of length at most $\frac{24}{7}(d+1)^d \text{OPT}_B \leqslant 4(d+1)^d \text{OPT}_B$. The box computed by the algorithm therefore has volume at most

$$\text{vol}(K_0) \leqslant 2^d 4(d+1)^d \text{OPT}_B + \text{vol}(E_0),$$

where E_0 is the extent of $\mathcal{D}$ in our chosen coordinate-system. The sides of E_0 orthogonal to s_0 have length at most 2, the side of E_0 in direction s_0 has length at most $2\sin(2\vartheta)$ (it is determined by a disk whose normal makes the largest angle with direction s_0), so

$$\text{vol}(E_0) \leqslant 2^{d+1} \sin 2\vartheta = 2^{d+1} \cdot 2 \sin\vartheta \cos\vartheta = 2^{d+1} \cdot s^* \cos\vartheta.$$

On the other hand, K^* has one side of length s^*, and all other sides have length at least $2\cos\vartheta = 2\cos(\arcsin \frac{s^*}{2}) > 2\cos(\arcsin \frac{1}{4}) = \frac{1}{2}\sqrt{15} > 1$. Thus

$$\text{vol}(K^*) \geqslant (2\cos\vartheta)^d \cdot s^* = 2\cos\vartheta\,(2\cos\vartheta)^{d-1} \cdot s^* > 2s^* \cos\vartheta \geqslant \frac{\text{vol}(E_0)}{2^d}.$$

It follows that $\text{vol}(E_0) \leqslant 2^d \cdot \text{OPT}_B$, and we finally have

$$\text{vol}(K_0) \leqslant 2^d(4(d+1)^d + 1) \cdot \text{OPT}_B < 4(3d+3)^{2d+2} \cdot \text{OPT}_B. \quad \square$$

By Lemma 10, the smallest box of arbitrary orientation where we can pack $\mathcal{D}$ is a $(d+1)^{(3d+3)/2}$-approximation to the smallest convex body where we can pack $\mathcal{D}$. Thus, by returning the box computed in Theorem 12 we obtain the following corollary.

Corollary 13. *A $36 \cdot 9^d (d+1)^{(7d+7)/2}$-approximation to* DiskPackingConvex *can be computed in polynomial time.*

7. The volume needed to stab or pack *n* hyperdisks

While in two dimensions all (infinitely many) unit-length line segments can be packed into a single fixed rectangle (Fig. 4), this is not true in higher dimensions. We will show that volume $\Omega(n^{\frac{d-1}{d}})$ is sometimes needed and volume $O(n^{\frac{d-1}{d}})$ is always sufficient to stab n unit hyperdisks in $\mathbb{R}^{d+1}$. This will imply similar lower bounds for packing inside a convex container and similar upper bounds for packing inside an axis-parallel box.

Theorem 14. *There is a family $\mathcal{D}$ of n unit hyperdisks in $\mathbb{R}^{d+1}$ whose normals make an angle of at most $\phi_0 = \arccos(1/\sqrt{d+1})$ with the positive x_{d+1}-axis, such that, for each subfamily $\mathcal{D}' \subset \mathcal{D}$, any stabbing of $\mathcal{D}'$ in any direction s needs length $\Omega(|\mathcal{D}'|/n^{1/d})$.*

Proof. The claim is obviously true for $d = 1$, so we assume $d \geqslant 2$. We construct a set $\mathcal{D}$ of n hyperdisks as follows: we pick a d-dimensional hypercube with side length c in the hyperplane normal to the x_{d+1}-axis that is centered at the origin, for some constant $c > 0$. We partition the hypercube into grid cells with side length $\varepsilon = c/n^{1/d}$, and project each grid point in the positive x_{d+1}-direction onto the unit-radius sphere, see Fig. 10. These n points define the normal vectors for our set $\mathcal{D}$ of hyperdisks. When c is small enough (depending on d), all normals make an angle of less than ϕ_0 with the x_{d+1}-axis.

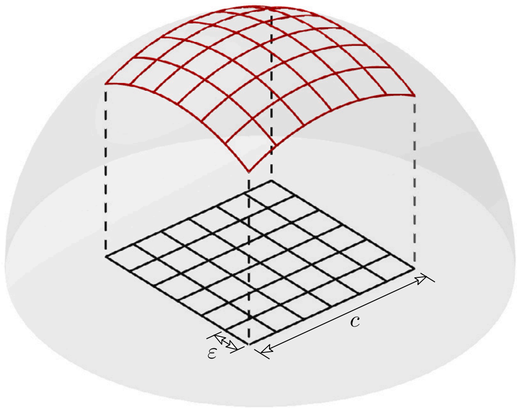

Fig. 10. Projecting a grid onto the unit sphere.

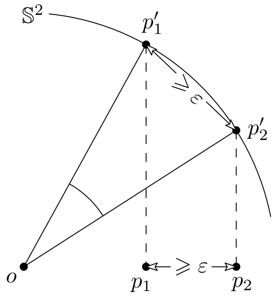

Fig. 11. Two grid points p_1, p_2 and their projection onto $\mathbb{S}^2$.

We observe next that the angle made by any two of the normal vectors is at least ε. Indeed, since any two grid points p_1, p_2 have distance at least ε, so do their projections p'_1 and p'_2. The angle $\measuredangle p'_1 o p'_2$ is the length of the great circle arc between p'_1 and p'_2 on the unit sphere $\mathbb{S}^d$, so it is longer than $|p'_1 p'_2|$, see Fig. 11.

Using Lemma 4 and the inequality $\sin x \geqslant x/2$ for all $x \in [0, \pi/2]$, this implies that for any two hyperdisks $D_1, D_2 \in \mathcal{D}$ and any direction s we have $d_s(D_1, D_2) \geqslant \sin\varepsilon \geqslant \varepsilon/2$. Therefore, for any $\mathcal{D}' \subseteq \mathcal{D}$, any stabbing of $\mathcal{D}'$ has length at least $(|\mathcal{D}'| - 1) \cdot \frac{\varepsilon}{2} = \Omega(|\mathcal{D}'|/n^{1/d})$. □

The arguments made in Section 3 imply the following consequence of Theorem 14.

Theorem 15. *There is a family $\mathcal{D}$ of n unit hyperdisks in $\mathbb{R}^{d+1}$ such that any convex container where we can pack $\mathcal{D}$ under translations has volume $\Omega(n^{\frac{d-1}{d}})$. In particular, there is no convex container of bounded volume into which all unit hyperdisks in $\mathbb{R}^{d+1}$ can be packed under translations.*

Proof. Consider the set $\mathcal{D}$ of hyperdisks of Theorem 14. Let OPT_C be the volume of an optimal convex container for $\mathcal{D}$. We have to show that $\text{OPT}_C = \Omega(n^{\frac{d-1}{d}})$.

Because of Lemma 10 there is a box M of volume at most $(d+1)^{3(d+1)/2} \cdot \text{OPT}_C$ where we can pack $\mathcal{D}$. We break $\mathcal{D}$ into $d+1$ groups $\mathcal{D}_1, \ldots, \mathcal{D}_{d+1}$ depending on which of the $d+1$ directions of the edges of M minimize the angle with the normal. Because of the pigeonhole principle, one of the groups, say $\mathcal{D}_1$, has at least $n/(d+1)$ hyperdisks. Let s be the corresponding direction of the edge of M minimizing the angle with the hyperdisks of $\mathcal{D}_1$. Because of Theorem 14, any stabbing of $\mathcal{D}_1$ has length at least $\Omega((n/(d+1))/n^{1/d}) = \Omega(n^{\frac{d-1}{d}})$. This is in particular true for the stabbing in the direction s. Using Theorem 6 (in the second inequality of the forthcoming chain), we get

$$(d+1)^{3(d+1)/2} \cdot \text{OPT}_C \geqslant \text{vol}(M) \geqslant \frac{\Omega(n^{\frac{d-1}{d}})}{2(d+1)^d} = \Omega(n^{\frac{d-1}{d}}),$$

which implies that $\text{OPT}_C = \Omega(n^{\frac{d-1}{d}})$. □

We now prove the reverse direction for stabbing.

Theorem 16. *Let $\mathcal{D}$ be a family of n unit hyperdisks in $\mathbb{R}^{d+1}$ whose normals make an angle of at most $\phi_0 = \arccos(1/\sqrt{d+1})$ with the positive x_{d+1}-axis. Then there is a stabbing for $\mathcal{D}$ of length $O(n^{\frac{d-1}{d}})$.*

Proof. Let $S \subset \mathbb{S}^d$ be the set of directions of the unit normals of $\mathcal{D}$. There is then a tour that visits the points of S one by one, and such that the sum of angles between consecutive points is bounded by $O(n^{\frac{d-1}{d}})$. This can be seen by first arguing that a minimum spanning tree of this cost exists [9], replacing that by an Euler tour, and finally shortcutting points that have been visited before. The claim now follows immediately from Lemma 5. $\square$

Using Theorem 7, this immediately gives us the following upper bound, which is now for the weaker class of container, axis-parallel boxes.

Corollary 17. *Let $\mathcal{D}$ be a family of n unit hyperdisks in $\mathbb{R}^{d+1}$. Then there is an axis-parallel box of volume $O(n^{\frac{d-1}{d}})$ into which $\mathcal{D}$ can be packed.*

8. Other objects and open problems

Our approximation algorithms can be extended to any arbitrary fixed shape A of dimension d, provided that A can be enclosed by some hyperdisk D_{out} (that is, it is bounded) and contains another hyperdisk D_{in} (that is, it has nonempty relative interior). More precisely, if we are given a finite set of congruent copies of A in $d+1$ dimensions, we can approximate the minimum-volume container into which the set can be packed by translations.

This can be done by just applying our algorithm to the corresponding set of copies of D_{out}. Since this gives a constant-factor approximation of the optimal packing of the D_{out}'s, it also gives an approximation of the optimal packing of the D_{in}'s. Observe however, that the approximation factor is multiplied by r^{d+1}, where r is the ratio between the radii of D_{out} and D_{in}. Since the optimal packing of the A's provides some packing of the D_{in}'s its container must be at least as large, from which we obtain an approximation for the A's.

The approximation factor obtained this way depends on the shape of A. For standard shapes such as squares ($r = \sqrt{2}$), equilateral triangles ($r = 2$) etc., we can directly compute it from the approximation factors of the algorithms.

We have given a constant-factor approximation for a special case of objects, unit-radius hyperdisks in $\mathbb{R}^{d+1}$. It remains an open problem whether an optimal packing of hyperdisks of different radii can be efficiently approximated. It is also unclear what happens, for example, when packing $(d-1)$-dimensional unit balls in $\mathbb{R}^{d+1}$ under translations. Finally, approximating the packing of arbitrarily oriented boxes or convex polyhedra seems to be much more difficult.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors thank Günter Rote for helpful comments and a simplification in the proof of Theorem 3.

N. Seiferth was partially supported by a fellowship within the FITweltweit program and H. Alt by the Johann-Gottfried-Herder program, both of the German Academic Exchange Service (DAAD). N. Seiferth was also partially supported by the German Science Foundation within the collaborative DACH project *Arrangements and Drawings* as DFG Project MU3501/3-1. J. Park was supported in part by starlab project (IITP-2015-0-00199) and NRF-2017M3C4A7066317. Research funded in part by The Slovenian Research and Innovation Agency (P1-0297, J1-2452, N1-0218, N1-0285). Research funded in part by the European Union (ERC, KARST, project number 101071836). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

Data availability

No data was used for the research described in the article.

References

[1] Helmut Alt, Computational aspects of packing problems, Bull. Eur. Assoc. Theor. Comput. Sci. 118 (2016).
[2] Helmut Alt, Mark de Berg, Christian Knauer, Approximating minimum-area rectangular and convex containers for packing convex polygons, J. Comput. Geom. 8 (1) (2017) 1–10.
[3] Helmut Alt, Nadja Scharf, Approximating smallest containers for packing three-dimensional convex objects, in: 27th Annu. Internat. Sympos. Algorithms Comput. (ISAAC), 2016, pp. 11:1–11:14.
[4] Herman Auerbach, Stefan Banach, Stanisław Mazur, Stanisław Ulam, Problem 10.1, in: R.D. Mauldin (Ed.), The Scottish Book, vol. 74, Birkhäuser, 1981.

[5] Alexander Barvinok, A Course in Convexity, Graduate Studies in Mathematics, vol. 54, American Mathematical Soc., 2002.
[6] Aharon Ben-Tal, Arkadi Nemirovski, Lectures on modern convex optimization, in: MOS-SIAM Series on Optimization, Society for Industrial and Applied Mathematics, 2001.
[7] Florian Diedrich, Rolf Harren, Klaus Jansen, Ralf Thöle, Henning Thomas, Approximation algorithms for 3D orthogonal knapsack, J. Comput. Sci. Technol. 23 (5) (2008) 749–762.
[8] Gábor Fejes Tóth, Packing and covering properties of sequences of convex bodies, in: Lagerungen: Arrangements in the Plane, on the Sphere, and in Space, Springer International Publishing, 2023, pp. 293–302.
[9] Leonard Few, The shortest path and the shortest road through n points, Mathematika 2 (2) (1955) 141–144.
[10] Robert J. Fowler, Mike Paterson, Steven L. Tanimoto, Optimal packing and covering in the plane are NP-complete, Inf. Process. Lett. 12 (3) (1981) 133–137.
[11] J.A. Hoogeveen, Analysis of Christofides' heuristic: some paths are more difficult than cycles, Oper. Res. Lett. 10 (5) (1991) 291–295.
[12] Klaus Jansen, Lars Prädel, A new asymptotic approximation algorithm for 3-dimensional strip packing, in: Viliam Geffert, Bart Preneel, Branislav Rovan, Julius Stuller, A. Min Tjoa (Eds.), SOFSEM 2014: Theory and Practice of Computer Science - 40th International Conference on Current Trends in Theory and Practice of Computer Science, Nový Smokovec, Slovakia, January 26-29, 2014, Proceedings, in: Lecture Notes in Computer Science, vol. 8327, Springer, 2014, pp. 327–338.
[13] A. Kosiński, A proof of an Auerbach-Banach-Mazur-Ulam theorem on convex bodies, Colloq. Math. 4 (2) (1957) 216–218.
[14] Guntram Scheithauer, Zuschnitt- und Packungsoptimierung: Problemstellungen, Modellierungstechniken, Lösungsmethoden, Vieweg+Teubner Verlag, 2008.
[15] Otfried Schwarzkopf, Ulrich Fuchs, Günter Rote, Emo Welzl, Approximation of convex figures by pairs of rectangles, Comput. Geom. 10 (1998) 77–87.
[16] Léonard von Niederhäusern, Packing polygons, Master's thesis, EPFL Lausanne, FU Berlin, 2014.