Computer Aided Geometric Design ••• (••••) ••••••

Contents lists available at ScienceDirect

Computer Aided Geometric Design

www.elsevier.com/locate/cagd

Single-sized spheres on surfaces (S4)

Henriette Lipschütz [a], Martin Skrodzki [b], Ulrich Reitebuch [a], Konrad Polthier [a]

[a] *Institute of Mathematics, Freie Universität, Berlin, Germany*
[b] *RIKEN iTHEMS, Wako, Japan*

ARTICLE INFO

Article history:
Available online xxxx

Keywords:
Surface interpolation
Computational differential geometry
Object modeling

ABSTRACT

Surface representations play a major role in a variety of applications throughout a diverse collection of fields, such as biology, chemistry, physics, or architecture. From a simulation point of view, it is important to simulate the surface as good as possible, including the usage of a wide range of different approximating elements. However, when it comes to manufacturing, it is desirable to have as few different building blocks as possible, as these can then be produced cost-efficiently.

Our paper adds a procedure to be used in the simulation of natural phenomena as well as within the designers' creative toolbox. It represents a surface via a collection of equally sized spheres. In the first part of the paper, we give a mathematically precise definition of the underlying problem: How to cover as much as possible of a surface by single-sized spheres. This leads to questions regarding the extremal intersection area of spheres with reasonably well-behaved surfaces, for which we present upper and lower bounds. From these, we deduce how many spheres of a certain, fixed radius can be used at least and at most when interpolating a surface.

Following these theoretical results, we compare a depth-first, a breadth-first, and a heuristic algorithm for the generation of surface coverings by single-sized spheres. As opposed to the mathematical description, we show that our algorithms also work for surfaces with boundary elements or sharp features such as edges or corners. We prove the applicability of our algorithm by a multitude of experiments and compare our procedure to ellipsoidal and multi-sized sphere methods.

1. Introduction

A large number of structures appearing in nature is made from a small number of building blocks. Examples are honeycombs or the surface of diatoms, see Fig. 1b taken from Redmayne (1877), which often show densely packed spheres with little variation in size. Another interesting occurrence is the compound eye of insects, which are covered by nearly regular hexagons, each containing a lens of disk-shaped cross section, see Fig. 1a taken from Nation kingdom (2008). A three-dimensional arrangement of spherical objects can be observed in plants. For instance, on the microscopic level, pollen form such sphere packings, see Fig. 1c taken from Pixeltoo (2005).

These marvelous feats of evolution in nature compel with both their simplicity as well as their geometric properties. Underlying symmetries make these arrangements aesthetically pleasing and mathematically interesting. Therefore, they serve as rich inspiration for human-built constructions, for instance in the realm of architecture. Drawing direct inspiration from

E-mail address: Henriette.Lipschuetz@fu-berlin.de (H. Lipschütz).

https://doi.org/10.1016/j.cagd.2021.101971
0167-8396/

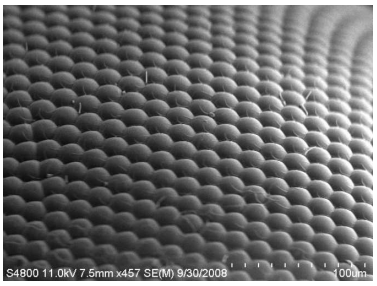

(a) Eye of a fly, taken from [1].

(b) Diatom (*triceratium favus*), taken from [2].

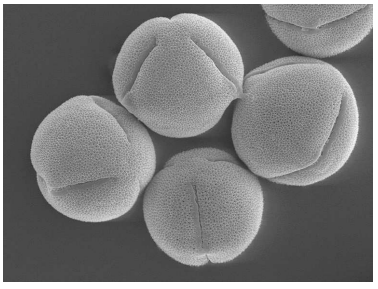

(c) Pollen (*ricinus communis sabguineus*), taken from [3].

Fig. 1. Real-world inspirations for spherical arrangements from biology.

these examples, the architect Buckminster Fuller designed the *Fly's Eye Dome*, see Fig. 2a taken from Wmpearl (2019), in 1965. A more complex example is the *Selfridges Building* in Birmingham, see Fig. 2b taken from Wedd (2007). Constructed in 2003 by the architecture company *Future Systems*, its sprayed concrete facade consists of a freeform surface with a circle packing, not unlike the packing of the diatom.

Lately, the computational architecture community paid attention to the creation of these specific freeform surfaces. Namely, several approaches have been presented to automatically generate circle packing meshes, such as the one on the Selfridges Building, from a given triangle mesh. These optimization-based procedures do succeed in the creation of circle packings on freeform surfaces, see Fig. 2c, but the resulting circular elements are not of a single size. From a construction point of view, this is by no means optimal as the production of a variety of building blocks is costly and the assembly bears the potential for misplaced parts. Finally, the currently existing algorithms can cover freeform surfaces with circles or ellipses, but not necessarily with intersection-free three-dimensional building blocks, such as spheres.

In this paper, we reduce the number of building blocks to a single sphere whose radius arises from the geometric properties of the surface. We investigate the amount of area of a given surface that can theoretically be covered by an arrangement of these spheres. Further, we present an algorithm that automatically builds a corresponding sphere cover on a given input surface. The optically appealing output is illustrated by several examples, which also provide quantitative measures that show the algorithm's relevance.

2. Related work

Several classical problems are concerned with the placement of circles or spheres on objects or in space. For instance, it is a well-known result from Tóth (1943), that a triangular arrangement of circles in the plane is the densest (ir-)regular placement of circles, covering $\frac{\pi\sqrt{3}}{6} \approx 0.9069$ of the plane. A similar result holds for sphere packings in three-dimensional Euclidean space, for which a recent result by Hales et al. (2017), establishes the hexagonal closed packing (HCP) as densest (ir-)regular packing with a covering of $\frac{\pi}{3\sqrt{2}} \approx 0.7405$. While our approach does consist of placing spheres in $\mathbb{R}^3$, we do not aim at a packing, but at interpolating a given surface. Said surface is not necessarily planar, thus we can also turn to covering ratios of the plane in terms of experimental comparisons.

Given a topological disk in $\mathbb{R}^2$, bounded by locally convex splines curves, the work of Machchhar and Elber (2017), provides a two-phase algorithm for filling the domain with circles. After choosing an appropriate radius, an initial set of circles is placed along the boundary of the domain; then, the inside is filled with circles touching at least two previously present ones. In the second phase, the circle positions are perturbed to enable placement of further circles. The algorithm

(a) Buckminster Fuller's *Fly's Eye Dome*, taken from [4].

(b) *Selfridges Building*, taken from [5].

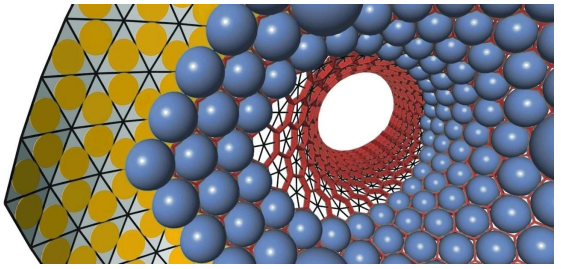

(c) Example of circle packing mesh with multi-sized spheres, taken from [6, Fig. 3].

Fig. 2. Architectural constructions inspired by disk and sphere packings in biology.

presented in this paper follows an approach similar to the first phase, but applied to placement of spheres on surfaces. Furthermore, instead of a perturbation phase, our algorithm employs a heuristic for placing the spheres.

A packing of circles on a given surface can be obtained e.g. by lifting a circle packing in the texture domain via a conformal map (see Kharevych et al., 2006, Fig. 16). The circles are, however, distorted in size by the mapping and thus resemble an Apollonian circle packing with its variety of disk sizes, see Bourke (2006). Furthermore, such lifts, while obeying a geodesic distance between the circle centers, do not necessarily obey a Euclidean distance between the centers when measured in the embedding space. The latter would cause spheres to collide, which is undesired in our setting. Closer resemblance of surface-interpolating spheres can be found in several natural phenomena that have been studied in the corresponding scientific disciplines.

Expanding the scope from a surface to a three-dimensional domain, several algorithms are available to compute a sphere-packing (e.g., Stoyan and Yaskov, 2013 or Sosin et al., 2021). In our context, these algorithms could be applied to the offset-volume of the surface and thus yield an interpolating sphere-packing on the input surface. However, the available methods employ costly optimization steps, while our algorithm computes a direct solution.

2.1. Simulation of natural phenomena

In chemistry, in order to stabilize emulsion droplets, it is important to investigate dense covers of spherical shaped particles on said droplets, see Burke et al. (2015). In contrast to deterministically grown structures, the particles are spread after fixing a number of infinitesimal starting positions on an ellipsoid. Afterwards, these particles begin to grow into elastic spheres which subsequently start to collide. Similar to the procedure in this paper, the work of Burke et al. (2015), deduces a sphere radius r from curvature properties of the underlying surface. However, their method is limited to ellipsoidal input surfaces.

In the context of statistical physics, disk packings create polycrystalline textures with irregular grain boundaries and linear shear fractures, see Lubachevsky and Stillinger (1990). As in the work discussed above, disks are grown from a number of randomly chosen starting points, which are moved by an optimization to a completely jammed state. This is in contrast to the procedure described in this paper, where sphere positions are not altered once they have been chosen. Furthermore, the optimization approach of Lubachevsky and Stillinger (1990) is computationally very expensive and only applicable to planar domains. We provide a comparison with our approach on such a domain in Section 5.1.

2.2. Interpolating surfaces

A first approach towards interpolation of general surfaces by a collection of circular elements has been pursued in the work of Shimada and Gossard (1995). While they successfully presented an algorithm to generate circle- and sphere-packings for given two- and three-dimensional domains respectively, their elements intersect each other and are therefore not suited for a construction context. First steps in this direction are taken by Schiftner et al. (2009) in the context of freeform architecture. Fundamental is the introduction of a so-called *circle packing mesh* (CP mesh). It denotes a triangle mesh for which the incircles of two triangles with a common edge have the same contact point on that edge, such that the circles form a packing. For the generation of these meshes from a triangle mesh, the authors propose an optimization algorithm. Performing the optimization requires the solution of a nonlinear least-squares problem and as the CP mesh property is just encoded in the target functional, in general, only approximations to CP meshes are obtained. Furthermore, the circle sizes vary over the approximated CP mesh and while the incircles touch, the inspheres can intersect in the final output. However, the authors also introduce a descriptor to measure the deviation from their CP mesh approximation to an actual CP mesh, which we will use in our experiments in Section 5.4.

Building on this optimization-based construction of CP meshes, the work of Pottmann et al. (2015) relaxes the outcome from strict circular cells to anisotropic cell packings. Once more, the starting point is a given triangle mesh, which is now optimized such that its edges are of nearly equal length when measured in the local metric. Ultimately, this results in only minimal changes of the target functional, while the resulting mesh is now comprised of inscribed ellipses instead of incircles that touch tangentially on the edges of the mesh. This additional degree of freedom allows for e.g. incorporating the curvature of the input mesh within the anisotropic construction (see Pottmann et al., 2015, Figure 9). Yet, because of the underlying optimization, the inscribed ellipses are not guaranteed to be free of collision, neither are the induced ellipsoids. Furthermore, it is still costly to solve the optimization.

In contrast to the two works discussed above, Xu et al. (2018) aim at the creation of an ellipses packing structure on a given freeform surface. To do so, again, an optimization is performed to guide the input triangle mesh to an ellipse packing mesh. The target functional of the optimization is similar to the one used by Schiftner et al. (2009) but is altered to account for the ellipsoidal shapes. Similar to the work of Pottmann et al. (2015), the radii and the orientations of the ellipsoids are determined using a positive definite matrix built from the principal curvature directions and the principal curvatures. As discussed by the authors in Appendix A of Xu et al. (2018), even when explicitly enforcing the ellipsoidal packing constraint, the resulting ellipsoids can still intersect each other. Thus, aside from the costly nonlinear optimization problem, while generating a set of interpolating ellipses, a non-intersecting packing is not guaranteed.

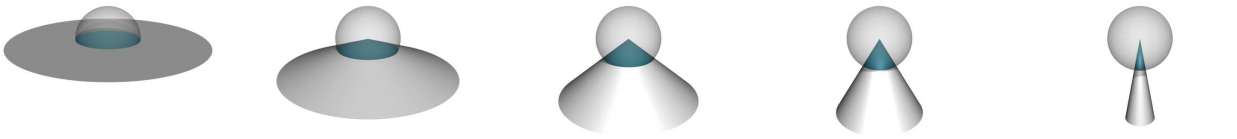

(a) Area πr^2. (b) Area $0.9397\pi r^2$. (c) Area $0.766\pi r^2$. (d) Area $0.5\pi r^2$. (e) Area $0.1734\pi r^2$.

Fig. 3. Cone intersected by a sphere of radius r at the apex.

Our work contributes to the literature and approaches regarding CP meshes as follows.

- Derivation of theoretical bounds as to what part of a given surface can be covered at least and at most when using spheres of a single, derived radius r.
- Presentation of a deterministic, sequential algorithm for the generation of sphere covers that interpolate a given triangulated input surface without optimization.
- Our algorithm can handle arbitrary input with boundary components as well as sharp features.

3. Theoretical foundation

We are interested in distributing equally measured spheres on a given input surface. For this part of the paper, we assume the input surface to be a two-manifold. To connect the surface and the sphere radius, consider the following observations. Since all spheres are to be of equal measurements, they are fully characterized by their radius $r \in \mathbb{R}_{>0}$. Placing a sphere on the surface $\mathcal{M}$ such that the center v of the sphere lies on $\mathcal{M}$, the sphere cuts out a piece of $\mathcal{M}$. Depending on r, the sphere might contain the whole surface which is not meant to be the case. Therefore, a reasonable range for r has to be determined. For r sufficiently small, the sphere intersects with the surface such that it cuts out one or several pieces of $\mathcal{M}$. These pieces are considered as *covered*. To maintain the preliminary ansatz of a single building block, the spheres are supposed to touch each other most, but not to intersect and to cover as much area of $\mathcal{M}$ as possible.

To deduce a choice of properties $\mathcal{M}$ has to fulfill to find a reasonable covering by spheres as described above, consider the following examples of two-dimensional surfaces embedded into $\mathbb{R}^3$:

Example 1. Consider a right circular cone $\mathcal{K}$ of height $h \in \mathbb{R}_{\geq 0}$ and radius $R \in \mathbb{R}_{>0}$. When placing a sphere of radius r on the apex of $\mathcal{K}$, the cut out area varies depending on the ratio of h and R, between nearly zero and πr^2 as an upper bound, see Fig. 3.

To prevent this—possibly very small—intersection area for some spheres, we should impose a limit on the curvature of the input manifold $\mathcal{M}$ such that each placed sphere roughly covers a similar area of the surface.

Example 2. Consider a thin, planar strip embedded in $\mathbb{R}^3$ with a width $w \ll r$. Then, a sphere of radius r will cover only about $2r \cdot w \ll \pi r^2$.

To prevent this—very small—intersection area at boundary components, we require the input manifold $\mathcal{M}$ to be closed, i.e. to be free of boundaries.

Example 3. The *clothoid* or *Euler spiral* is a planar curve fulfilling the property that its osculating circle is antiproportional to its arc length, see Fig. 4a. Let Z be the cylinder over the clothoid, see Fig. 4b. Consider the intersection of Z with a sphere

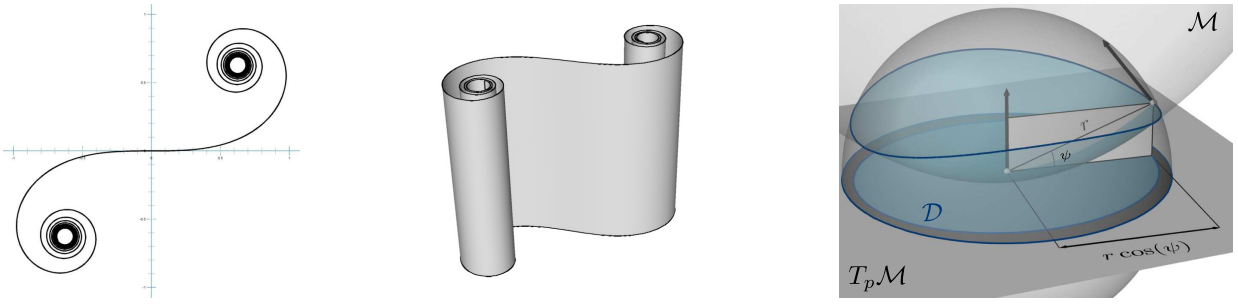

(a) The clothoid curve. (b) Cylinder over the clothoid. (c) Trigonometric relation.

Fig. 4. Clothoid and a cylinder over it as example of a problematic surface. Trigonometric relations on the tangent plane.

of fixed radius r such that its center c lies on Z. Depending on c, the intersection varies between a slightly deformed disk to a rolled up strip in the inside of the sphere having arbitrarily large area and probably some unconnected intersections.

This final example shows that a sphere can cover a large area of the surface, while the covered part is fragmented into several disconnected layers. As we only want to count that part of the surface as covered that is connected to the sphere center, we have to avoid this situation. Before formulating a corresponding strong condition, we recall some definitions. In the following, $\|p\|$ denotes the Euclidean norm of $p \in \mathbb{R}^3$.

Definition 1. Let $\mathcal{M}$ be a closed $\mathcal{C}^2$-manifold embedded in $\mathbb{R}^3$ and let $\operatorname{len}(f) = \int_0^1 |f'(t)|\, dt$ denote the length of a curve $f \in \mathcal{C}^1([0,1], \mathcal{M})$ in $\mathcal{M}$. The *geodesic distance* $d_{\mathcal{M}}$ of $m, m' \in \mathcal{M}$ is defined as

$$d_{\mathcal{M}}(m, m') := \inf\left\{\operatorname{len}(f) \mid f \in \mathcal{C}^1([0,1], \mathcal{M}) :\ f(0) = m \wedge f(1) = m'\right\}.$$

To guarantee the geodesic distance to be finite, we consider only connected manifolds or the connected components of $\mathcal{M}$ individually. To avoid an unconnected intersection of $\mathcal{M}$ and the spheres, we connect the radius of the spheres to the *reach* of $\mathcal{M}$. Therefore, define the *medial axis* and *reach* of $\mathcal{M}$ as follows:

Definition 2. The *medial axis* of a $\mathcal{C}^2$-manifold $\mathcal{M}$ embedded in $\mathbb{R}^3$ is given as

$$\mathcal{A}(\mathcal{M}) := \left\{ p \in \mathbb{R}^3 \mid \exists m_1 \neq m_2 \in \mathcal{M} : \min_{m' \in \mathcal{M}} d(p, m') = d(p, m_1) = d(p, m_2) \right\}.$$

Definition 3. Let a $\mathcal{C}^2$-manifold $\mathcal{M}$ embedded in $\mathbb{R}^3$ be given. Then, the *reach* ρ of $\mathcal{M}$ is defined as

$$\rho := \inf\{\|a - m\| \mid a \in \mathcal{A}(\mathcal{M}) \wedge m \in \mathcal{M}\}.$$

Now, we can summarize our strong conditions on $\mathcal{M}$ as:

Assumption 1. In the following, we assume $\mathcal{M}$ to be a (1) compact, (2) connected, (3) $\mathcal{C}^2$-manifold (4) embedded in $\mathbb{R}^3$, (5) without boundary, and (6) of strictly positive, finite reach ρ.

As discussed above, the densest covering of the plane with equally sized spheres is given by a regular triangle grid. Therefore, the numbers of neighboring spheres as well as the amount of covered surface will be our means of comparison in Section 5.1. In the following, we derive further bounds.

3.1. Upper and lower bounds on the covered area

Let $\mathcal{M}$ be a manifold fulfilling the properties from Assumption 1 and let $B(p)$ denote a ball of fixed radius $r \in \mathbb{R}_{>0}$ centered at $p \in \mathcal{M}$. Lemma 3 in Boissonnat et al. (2019) states that for any $p, q \in \mathcal{M}$ fulfilling $\|p - q\| < 2\rho$ the following estimation holds

$$\|p - q\| \leq d_{\mathcal{M}}(p, q) \leq 2\rho \arcsin\left(\frac{\|p - q\|}{2\rho}\right).$$

Lemma 6 in Boissonnat et al. (2019) states that for $p, q \in \mathcal{M}$, the following holds for the angle $\sphericalangle\left(T_p\mathcal{M}, T_q\mathcal{M}\right)$ between the tangent planes of $\mathcal{M}$ at p and q, denoted by $T_p\mathcal{M}$ and $T_q\mathcal{M}$ respectively,

$$\sphericalangle\left(T_p\mathcal{M}, T_q\mathcal{M}\right) \leq \frac{d_{\mathcal{M}}(p, q)}{\rho}, \tag{1}$$

hence Equation (1) implies for the normal vectors N_p in p and N_q in q that

$$\varphi := \sphericalangle\left(N_p, N_q\right) \leq 2\arcsin\left(\frac{\|p - q\|}{2\rho}\right) \Rightarrow \|p - q\| \geq r(\varphi) := 2\rho \sin\left(\frac{\varphi}{2}\right). \tag{2}$$

Choose $\varphi_{\max} < \frac{\pi}{2}$ since then $\langle N_p, N_{p'} \rangle > 0$ for $p' \in \mathcal{M} \cap B(p)$ for all $r \leq r(\varphi_{\max})$.

After choosing appropriate coordinates, $\mathcal{M} \cap B(p)$ can be represented as the graph of some function $f(u, v)$ over the (topological) disk $\mathcal{D} = \pi_{T_p\mathcal{M}}(\mathcal{M} \cap B(p))$ such that $N_p = e_3$ where e_3 denotes the third vector of the standard basis. Then, using a re-parametrization of φ given by the new coordinates, the area of $\mathcal{M} \cap B(p)$ can be estimated by (using relation $\star$: $\cos(\varphi_{\max}) = 1 - 2\sin^2\left(\frac{\varphi_{\max}}{2}\right)$):

$$\text{area}\,(\mathcal{M}\cap B(p)) = \int_{\mathcal{D}} \frac{1}{|\cos(\varphi)|}\,du\,dv \overset{\substack{\cos(\varphi_{\max})\leq\cos(\varphi),\\ \text{area}(\mathcal{D})\leq r^2\pi}}{\leq} \frac{1}{|\cos(\varphi_{\max})|} r^2\pi \overset{\star,(2)}{\leq} \left|\frac{1}{1-\frac{r^2}{2\rho^2}}\right| r^2\pi. \tag{3}$$

On the other hand, the projection $\mathcal{D}$ of $\mathcal{M}\cap B(p)$ to the tangent plane $T_p\mathcal{M}$ at p gives a lower bound for the area of the cut surface since $\text{area}(\mathcal{D}) \leq \pi r^2$:

$$\text{area}\,(\mathcal{M}\cap B(p)) \overset{\text{proj.}}{\geq} \text{area}(\mathcal{D}) \overset{\substack{\text{Fig. 4c /w}\\ \psi\leq\varphi_{\max}}}{\geq} \cos^2(\varphi_{\max})\, r^2\pi \overset{\star,(2)}{\geq} \left(1-\frac{r^2}{2\rho^2}\right)^2 r^2\pi. \tag{4}$$

3.2. Hard bounds on the number of spheres

We can now use the bounds deduced above to state a lower and upper number of non-intersecting spheres that can be put on the input manifold $\mathcal{M}$.

Theorem 1 *(Upper bound number of spheres). Given a manifold $\mathcal{M}$ with requirements from Assumption 1, for $r \ll \rho$, we can fit at most $\left\lfloor \left(1-\frac{r^2}{2\rho^2}\right)^{-2} \frac{\text{area}(\mathcal{M})}{\pi r^2} \right\rfloor$ spheres of radius r on $\mathcal{M}$ such that they are at most touching, but not intersecting.*

Proof. Given a manifold $\mathcal{M}$ with requirements from Assumption 1, assume r is small compared to ρ such that for a given ball $B(v_i)$ of radius r with center $v_i \in \mathcal{M}$, we have due to Equation (4) that

$$\text{area}\,(\mathcal{M}\cap B(v_i)) \geq \left(1-\frac{r^2}{2\rho^2}\right)^2 r^2\pi.$$

Furthermore, as all ball centers are placed with Euclidean distance at least $2r$ from each other, two balls $B(v_i)$ and $B(v_j)$ do not intersect. Therefore, the area of $\mathcal{M}$ covered by $B(v_i)$ is not included in any other cover $\mathcal{M}\cap B(v_j)$, $j \neq i$. Thus, we can fit at most

$$\left\lfloor \frac{\text{area}(\mathcal{M})}{\left(1-\frac{r^2}{2\rho^2}\right)^2 \pi r^2} \right\rfloor$$

spheres on $\mathcal{M}$. □

Conversely, we can also prove the following lower bound on the number of spheres.

Theorem 2 *(Lower bound number of spheres). Given a manifold $\mathcal{M}$ with requirements from Assumption 1, for $2r \ll \rho$, we can fit at least $\left\lceil \left|1-\frac{2r^2}{\rho^2}\right| \frac{\text{area}(\mathcal{M})}{\pi(2r)^2} \right\rceil$ balls of radius r on $\mathcal{M}$ such that they are at most touching, but not intersecting.*

Proof. Given a manifold $\mathcal{M}$ with requirements from Assumption 1, consider the following procedure for the construction of balls $\bar{B}(v_j)$ of radius $2r$. Start with an arbitrary point $p_1 \in \mathcal{M}$ and place the center v_1 of a ball $\bar{B}(v_1)$ at p_1. Then, repeat the following procedure until $\mathcal{M}$ is completely covered: Choose a next point $p_j \in \mathcal{M}$ that is not covered by all currently constructed spheres $\bar{B}(v_1), \ldots, \bar{B}(v_{j-1})$. Place center v_j of $\bar{B}(v_j)$ at p_j. Finally, obtain a set of $2r$-balls $\bar{B}(v_1), \ldots, \bar{B}(v_k)$, $k \in \mathbb{N}$, such that $\mathcal{M} \subset \bigcup_{j=1}^{k} \bar{B}(v_j)$.

Note that two of these balls can intersect each other. Thus, to cover $\mathcal{M}$ completely, we need at least $\left\lceil \left|1-\frac{2r^2}{\rho^2}\right| \frac{\text{area}(\mathcal{M})}{\pi(2r)^2} \right\rceil$ balls because $\text{area}\left(\mathcal{M}\cap\bar{B}(v_j)\right) \leq \left|\frac{1}{1-\frac{2r^2}{\rho}}(2r)^2\pi\right|$ by Equation (3). Consider the balls $\bar{B}(v_1), \ldots, \bar{B}(v_k)$ as constructed above, then we have $\left\lceil \left|1-\frac{2r^2}{\rho^2}\right| \frac{\text{area}(\mathcal{M})}{\pi(2r)^2} \right\rceil \leq k$.

We will now place k balls of radius r on $\mathcal{M}$ that are neither touching nor intersecting by placing the center v_i of an ball $B(v_i)$ at the center v_i of each $2r$-ball $\bar{B}(v_i)$ as constructed above for $1 \leq i \leq k$. Naturally, all centers v_i lie on $\mathcal{M}$. Furthermore, by construction, $\|v_i - v_j\| \geq 2r$ therefore balls $B(v_i)$ and $B(v_j)$ are non-intersecting for all $i, j \in \{1, \ldots, k\}$, $i \neq j$. □

The theoretical bounds from Theorems 1 and 2 provide us with estimates as to how many spheres can be used to cover an as-large-as-possible area of the input manifold $\mathcal{M}$. However, these results are not constructive, i.e. they do not provide us with a concrete procedure how to find the center points of the covering spheres. Therefore, in the following section,

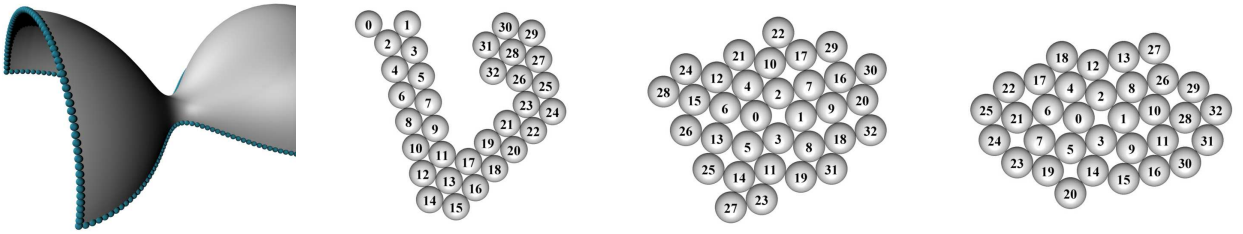

(a) Init. Boundary. (b) Depth-First. (c) Breadth-First. (d) Heuristic.

Fig. 5. Graph growing process for the different approaches.

we will present three versions of a sequential and deterministic algorithm to compute said centers. The theoretical results presented here can be related to the experiments as listed in Section 5, as shown in Fig. 9d.

4. Algorithmic implementation

In this section, we will discuss an algorithmic implementation that explicitly constructs a sphere cover. Given a finite, connected two-manifold $\mathcal{M}$ via a corresponding triangulation $\mathcal{T}$ of bounded positive *discrete reach*.[1] We construct a graph $G = (V, E)$ with a set of vertices V and a set of edges E. Additionally, the graph is embedded into the given triangulation $\mathcal{T}$, i.e. all vertices $v_i \in V$ are placed on the vertices, edges, or the interior of triangles from $\mathcal{T}$. Furthermore, all edges $e_{ij} = (v_i, v_j)$ have length $\|v_i - v_j\| = 2r$ when measured as distance of their embeddings in the embedding space $\mathbb{R}^3$. The construction is performed sequentially, starting from an initial graph G_0.

4.1. Initialization

The initial graph $G_0 = (V_0, E_0)$ is found as follows. For a closed surface $\mathcal{M}$, we start with a pair of vertices $v_0, v_1 \in \mathcal{T}$ with $4r \geq \|v_0 - v_1\| \geq 2r$ such that there exists a point $p \in \mathcal{T}$ with $\|p - v_0\| = \|p - v_1\| = 2r$. For a surface $\mathcal{M}$ with boundary, we initialize the graph with a set of points along each boundary component, see Fig. 5a, distributed with $2r$ distances, possibly leaving a gap somewhere.

4.2. Growing the graph

Starting from the initial graph G_0, we iteratively find a point $p \in \mathcal{T}$ which has distance $2r$ to two vertices $v_i, v_j \in V$ and distance $\|p - v_k\| \geq 2r$ for all $v_k \in V$. Given such a point p, it is added as a new vertex v to V, while the edges (v_i, v) and (v_j, v) are added to the edge set E. This process is iterated until there is no pair of vertices v_i, v_j such that a point p at appropriate distance can be found. Note that with each new vertex, we add two new edges, so we have $2(|P| - |P_0|) = |E| - |E_0|$. In total, for closed surfaces (starting with $P_0 = 2$ and $E_0 = 0$), we get $4 \cdot |P| - 8 = 2 \cdot |E|$, thus the average valence of vertices in G is $\frac{2 \cdot |E|}{|P|} = 4 - \frac{8}{|P|} \approx 4$, by construction. A priori, there are manifold choices for points p to be added as new vertices. We consider the following three strategies for identifying the next point p to be added.

4.2.1. Depth-first

The depth-first approach constructs as-straight-as-possible triangle strips on $\mathcal{T}$. That is, when searching for a pair of vertices v_i, v_j to find a new point p to be added as next vertex, the procedure starts with as-high-as-possible vertex indices i, j, see Fig. 5b. The resulting strip of vertex triangles continues until it comes too close to an already covered region. In this case, it continues by constructing a new triangle strip, starting as close as possible to the end of the last triangle strip. By construction, the covering strips are as themselves as dense as possible. However, there are only few connections between the triangle strips, which causes large crevasses on the surface, see Fig. 6a.

4.2.2. Breadth-first

The breadth-first approach resembles a disk-growing around the initial vertices v_i, v_j. This is achieved by taking as-low-as-possible indexed vertices v_i, v_j into account when searching for a vertex pair v_i, v_j to create a new vertex v to be inserted, see Fig. 5c. Compared to the depth-first approach, the surface is covered much more homogeneously. However, the breadth-first approach still creates visible seams. These arise wherever the growing process comes close to already covered regions, see Fig. 6b.

[1] Following Definition 3, we consider the reach to be the distance of the medial axis to the surface. In our setting, the *discrete reach* is given as the distance of the medial axis as computed by Tagliasacchi et al. (2012), to the triangulation $\mathcal{T}$.

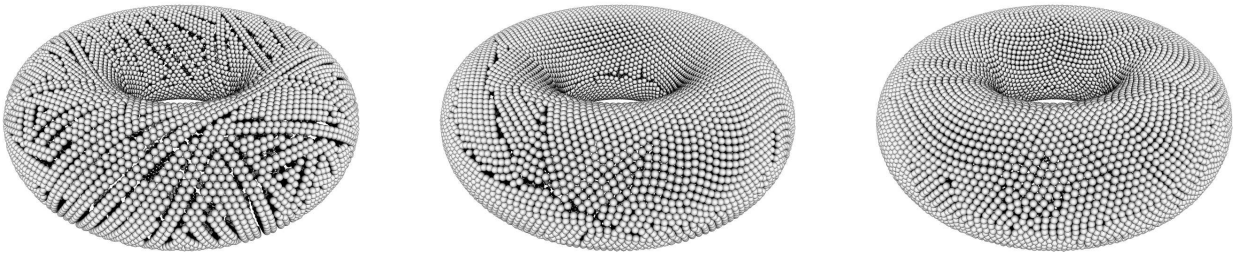

(a) Depth-First. (b) Breadth-First. (c) Heuristic.

Fig. 6. Illustrations of the three different graph growing models as discussed in Section 4.2.

4.2.3. Heuristic approach to avoid visible seams

To further improve the visual quality of our algorithm, we propose a heuristic. The idea is to prevent the formation of longer seams as they occur in the breadth-first approach. In order to do so, we maintain a linked list along each boundary component of the surface area covered so far. Thereby, we can easily deduce the distance of two boundary vertices v_i, v_j when traversing along this boundary component. Now, a new point p is selected as vertex, created from v_i, v_j via the following priorities. First, we try to find a triple (p, v_i, v_j) such that v_i and v_j are in different boundary components of the growing graph-covered area, see Fig. 7d. This can happen for surfaces of higher genus or surfaces with more than one boundary component. Second, we try to connect two points v_i, v_j in the current boundary by a new vertex v such that the distance between v_i and v_j—when measured as a discrete distance in the boundary of the currently created graph—is maximal compared to other pairs of boundary vertices, see Fig. 7b. If the best pair is not unique, we prioritize by breadth-first growing, see Fig. 5d.

Compared to the depth-first and breath-first approach, the heuristic successfully avoids the creation of large crevasses or visible seams. The corresponding surface area is better distributed over the entire graph, causing several *monovacancies* as observed by Lubachevsky and Stillinger (1990) and Fig. 6c.

4.3. Design-level comparison to previous algorithms

The previously proposed algorithms, (see Schiftner et al. (2009); Lubachevsky and Stillinger (1990); Pottmann et al. (2015); Xu et al. (2018)), all incorporate solving an optimization problem. While Lubachevsky and Stillinger (1990), obtain the starting positions from a random sampling, the other listed algorithms, (see Schiftner et al. (2009); Pottmann et al. (2015); Xu et al. (2018)), deduce them from the input mesh. Afterwards, all four algorithms employ iterative solvers to minimize the cost of the current configuration.

In contrast to these schemes, our algorithm does not need any optimization procedure. Following the greedy strategy of our approach, each placed vertex center is final and does not have to be moved. This provides us with fast solutions whose computed graphs G are always guaranteed CP meshes by construction.

The algorithm of Lubachevsky and Stillinger (1990) is tailored to create circle packings on planar domains. The other algorithms, (see Schiftner et al. (2009); Pottmann et al. (2015); Xu et al. (2018)), optimize for a circle or ellipses packing on a given input mesh. However, whether the optimization does find non-intersection circles/ellipses and whether the induced spheres/ellipsoids are touching, but non-intersecting, largely depends on the quality of the input mesh. Our algorithm can create a true CP mesh from any input connected triangulation of strictly positive discrete reach. Note that at this point we relax the conditions imposed on $\mathcal{M}$ in Section 3. That is, we allow for surfaces with boundary and sharp features—such as edges or corners. In the following, we will add to this high-level comparison by providing quantitative results.

5. Experiments

In this final section, we perform three sets of experiments on a variety of geometric models. As explained in Section 4.1, our algorithm can be initialized either with two *center spheres* placed somewhere on the surface, but at most distance $4r$ apart, or by placing a number of initial spheres along the boundary components of the geometry. We will distinguish these

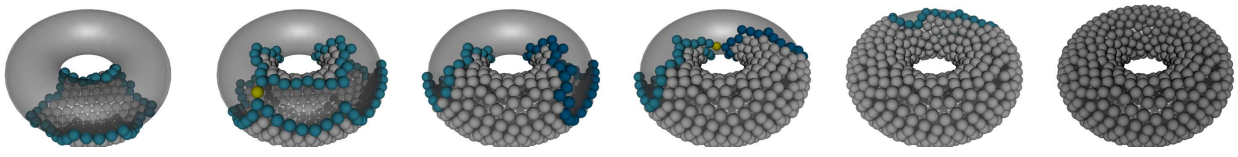

(a) Start. (b) Split front. (c) Two fronts. (d) Join fronts. (e) Continue. (f) Final cover.

Fig. 7. Heuristic disk growing process.

Table 1
Legend of abbreviations used in Section 5.

N	final number of spheres	D	depth-first	ic	init center
TAE	Triangle Average Edge	B	breadth-first	ib	init boundary
ER/TAE	Edge RMS to TAE	H	heuristic	area	area($\mathcal{T}$)

two cases as *init center (ic)* or *init boundary (ib)*. Furthermore, Sections 4.2.1 to 4.2.3 depicted three different versions (*Depth-First (D)*, *Breadth-First (B)*, *Heuristic (H)*) of our algorithm. To measure their performance, we will run all three on several geometries.

The input to the algorithm is comprised of the sphere radius r and an input geometry via its triangulation $\mathcal{T}$. In order to relate r to the size of the geometry, we list the diagonal of the axis-aligned bounding box of the geometry. Furthermore, we give the area of the input surface as the sum of the triangle areas of $\mathcal{T}$.

Output quantities arising in the experiments are the *final number N of spheres* that have been constructed to cover the input surface. We provide an estimate—using area($\mathcal{T}$)—as to what percentage of this area is covered by the resulting sphere cover. Furthermore, we provide the quantity *TAE*, the triangle average edge length, which is computed as follows. After creating the graph G, we triangulate all cells of G on $\mathcal{T}$ and thus obtain another triangulation $\mathcal{T}'$. The value *TAE* gives the average edge length in $\mathcal{T}'$. We use this quantity to relate the *edge root mean square error (ER)* to. That is to say, the value *ER/TAE* states the expected deviation of an edge length in $\mathcal{T}'$ from the average edge length of $\mathcal{T}'$. A legend to the abbreviations of these quantities is given in Table 1.

5.1. Growing square

From the results presented in Section 2, we know that the optimal packing of circles in a plane is a triangular lattice that covers approximately 90.69% of the plane. In the work of Lubachevsky and Stillinger (1990), starting from an initial sample of 2,000 randomly chosen points, the authors grow spheres and handle the collisions that subsequently occur. After handling $2 \cdot 10^4$ pairwise collisions, they achieve a covering fraction of about 56.28%.

We run our algorithm on a square with different initialization and vary the sphere radius. The results are shown in Table 2. Note that independent of the chosen initialization or sphere radius, the lowest covering rate obtained is 66%, about 10% better than the results of Lubachevsky and Stillinger (1990), without having to handle such a large number of collisions.

Even when increasing the runtime and handling $2 \cdot 10^5$ pairwise collisions, the algorithm of Lubachevsky and Stillinger (1990), arrives at a covering rate of about 73.94%. Still, this is only better than five of the 28 experiments we ran. Only with

Table 2
Running our proposed algorithm on a square with different initialization and varying sphere radius r. See Table 1 for a legend of the column headings.

Model	Alg.	Radius r	BB diag.	Area	N	$N\pi r^2$ / area	TAE	ER/TAE
Square (ic)	H	0.125	1.7678	1.5625	21	66.0%	0.25	0.0%
Square (ic)	H	0.0625	1.5910	1.2656	70	67.9%	0.1250	0.0%
Square (ic)	H	0.03125	1.5026	1.1289	293	79.6%	0.0625	10.2%
Square (ic)	H	0.015625	1.4584	1.0635	1074	77.5%	0.03125	14.2%
Square (ic)	H	0.0078125	1.4363	1.0315	4275	79.4%	0.015625	14.5%
Square (ic)	H	0.00390625	1.4253	1.0157	16902	79.8%	0.0078125	14.9%
Square (ic)	H	0.001953125	1.4197	1.0078	67337	80.1%	0.0042270	15.1%
Square (ic)	B	0.125	1.7678	1.5625	22	69.1%	0.25	0.0%
Square (ic)	B	0.0625	1.5910	1.2656	73	70.8%	0.125	0.0%
Square (ic)	B	0.03125	1.5026	1.1289	296	80.4%	0.063080	8.8%
Square (ic)	B	0.015625	1.4584	1.0635	1161	83.7%	0.031545	8.9%
Square (ic)	B	0.0078125	1.4363	1.0315	4600	85.5%	0.015726	6.6%
Square (ic)	B	0.00390625	1.4253	1.0157	18506	87.3%	0.0078804	6.9%
Square (ic)	B	0.001953125	1.4197	1.0078	73074	86.9%	0.003985	8.2%
Square (ib)	H	0.125	1.7678	1.5625	24	75.4%	0.2860	26.4%
Square (ib)	H	0.0625	1.5910	1.2656	78	75.6%	0.1435	22.0%
Square (ib)	H	0.03125	1.5026	1.1289	284	77.2%	0.07121	20.1%
Square (ib)	H	0.015625	1.4584	1.0635	1084	78.2%	0.03551	19.5%
Square (ib)	H	0.0078125	1.4363	1.0315	4221	78.5%	0.01769	19.0%
Square (ib)	H	0.00390625	1.4253	1.0157	16681	78.7%	0.008823	18.7%
Square (ib)	H	0.001953125	1.4197	1.0078	66216	78.7%	0.004415	18.8%
Square (ib)	B	0.125	1.7678	1.5625	22	69.1%	0.2995	28.4%
Square (ib)	B	0.0625	1.5910	1.2656	78	75.6%	0.1399	25.4%
Square (ib)	B	0.03125	1.5026	1.1289	302	82.1%	0.06608	17.0%
Square (ib)	B	0.015625	1.4584	1.0635	1166	84.1%	0.03271	14.7%
Square (ib)	B	0.0078125	1.4363	1.0315	4570	85.0%	0.01626	13.9%
Square (ib)	B	0.00390625	1.4253	1.0157	17712	83.6%	0.008222	14.4%
Square (ib)	B	0.001953125	1.4197	1.0078	69939	83.2%	0.004130	14.0%

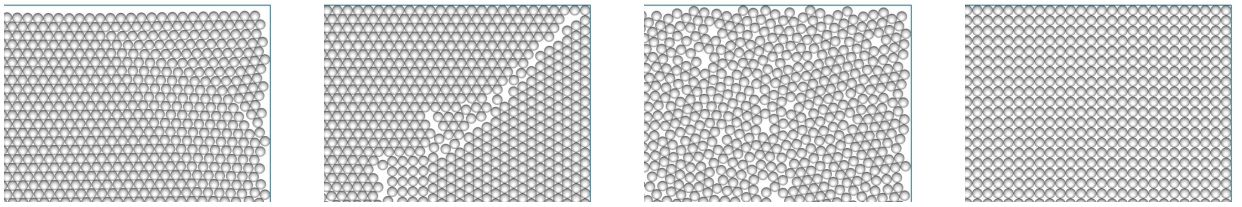

(a) Breadth-first, ic. (b) Breadth-first, ib. (c) Heuristic, ic. (d) Heuristic, ib.

Fig. 8. Zoom-in on a corner of a sphere cover on the square, border of the square shown as a blue line. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

another order of magnitude, running for $2 \cdot 10^6$ pairwise collisions, the covering rate obtained by Lubachevsky and Stillinger (1990), reaches about 86.81%. This places it roughly on par with the best covering rate obtained in our experiments. It was created by running the breadth-first approach with initial spheres in the center of the square and yielded a covering rate of 87.3%, admittedly with a larger number of spheres, namely 18,506. However, even with a comparable number of spheres, i.e. at $N \approx 2,000$, our experiments still provide coverage of close to or even above 80%. Note that the approach of Lubachevsky and Stillinger (1990), was able to surpass this number slightly, reaching a covering rate of 88.14% after handling $2.6 \cdot 10^7$ collisions.

A final point to be taken from this set of experiments is the difference the initialization makes. In Fig. 8, we show a corner of the square, covered with spheres of radius $r \approx 3.9 \cdot 10^{-3}$. In Figs. 8a and 8c, only two spheres are initially placed at the center of the square, from which the algorithm grows the covering. In contrast, the other two cases show the result for fixing spheres along the boundary of the square. Note that while fixing the boundary, the breadth-first approach still creates a visible seam, which is effectively prevented using the heuristic approach.

5.2. Growing sphere

For our second experiment, we turn to placing spheres on an actually curved surface in $\mathbb{R}^3$. In order to be able to perform computations precisely, we turn to a sphere cover on the unit sphere. Our sphere model is a discretization of the sphere, obtained from an icosahedron that has been subdivided seven times and the resulting vertices have been projected onto the unit sphere. With a surface area of approximately $0.99998 \cdot 4\pi$, this construction yields a reliable discretization of the unit sphere.

The graphs in Fig. 9 show the covering rate ξ of the radius r spheres, i.e. what proportion of the unit-spheres they cover. This rate is shown depending on the reciprocal of the radius r. The background colors are related to (from bottom to top) different covers of the plane:

- dark gray, hexagonal cover, $\xi < 0.604998$,
- gray, pentagonal cover (irregular), $0.604998 < \xi < 0.7090829$,
- light gray, square lattice, $0.7090829 < \xi < 0.785398$,
- white, hexagonal closest packing, $0.785398 < \xi < 0.906899$,
- dark gray, better cover than in the planar case, $0.906899 \leq \xi$.

The algorithm uses two initial spheres of radius r as input, we place them at $\left(\pm r, 0, \sqrt{1-r^2}\right)$, which produces a distance of $2r$ between these two and puts them symmetrically with regard to the 'north pole' of the covered sphere. From these two starting points, we run the three different versions of our algorithm.

Note that the covering rate ξ does not grow monotonically for any of the three approaches as given in Fig. 9. The depth-first approach actually passes a local maximum at about $r = \frac{1}{12}$ before settling at roughly $\xi \approx 72\%$. Thereby, this covering of the sphere barely surpasses an irregular pentagonal cover of the plane and performs about as well as a square lattice in the plane. Both, the breadth-first and the heuristic approach, perform better. Starting with some variation for large radii r, the covering rate ξ grows monotonically and settles at about $\xi \approx 82\%$ for the breadth-first and at about $\xi \approx 83\%$ for the heuristic approach. Therefore, both perform about as well as covering on a sphere as a hexagonal closest packing in the

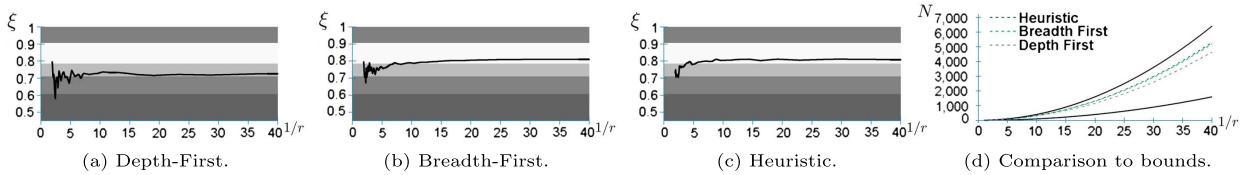

(a) Depth-First. (b) Breadth-First. (c) Heuristic. (d) Comparison to bounds.

Fig. 9. Each of the three graphs represents the covering ratio ξ of spheres of radius r when placed on a unit-sphere. The last graphs compare the performance of the three approaches with regard to the theoretical bounds derived in Section 3.

Table 3
Statistics about running our proposed algorithm on the Fertility model with different initialization positions twelve times. See Table 1 for a legend of the column headings.

Model	Alg.	Radius r	BB diag.	Area	N	$N\pi r^2$ / area	TAE	ER/TAE
Fertility (Min)	H	2.0	256.562	5979.6064	15,511	81.49%	2.1365	12.5%
Fertility (Avg)	H	2.0	256.562	5979.6064	15,550	81.69%	2.1417	12.7%
Fertility (Max)	H	2.0	256.562	5979.6064	15,594	81.93%	2.1461	12.9%

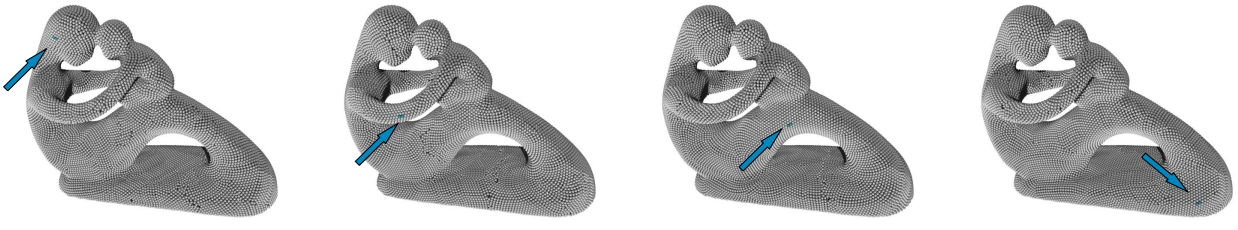

(a) Start: Head. (b) Start: Arm. (c) Start: Leg. (d) Start: Base.

Fig. 10. Four different initial positions. The output of the algorithm remains stable under changing the starting positions.

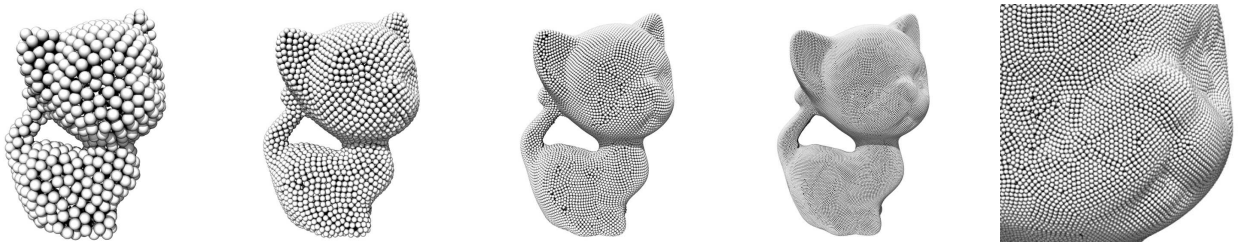

(a) r=0.025. (b) r=0.0125. (c) r=0.00625. (d) r=0.003125. (e) Zoom-in (d).

Fig. 11. The *Kitten* model processed with the heuristic version of our algorithm.

plane. Furthermore, Fig. 9d also relates the experimental results to the theoretical bounds derived in Section 3. We see that the number of spheres obtained from the breadth-first and the heuristic approach tend towards the theoretical maximum.

5.3. Dependence on the initial vertex placement

The next set of experiments is devoted to testing the stability of our algorithm with respect to the initial graph G_0. Following Section 4.1, surfaces with boundary components do not exhibit a choice in the initial vertex placement. Therefore, we turn to a closed model, the *Fertility* model.

For our experiments, we utilize the heuristic and fix a radius $r = 2.0$. Then, we run our algorithm on the fertility model twelve times, each run with a different starting position. See Fig. 10 for a selection of four different outcomes produced by the algorithm when the initial pair of points is placed on the head, arm, leg, and base of the model, respectively. The observed differences over the twelve runs are minimal, see Table 3. The number of spheres that are placed varies by some per mill between the minimum and maximum number obtained in the series. Correspondingly, the estimated covered area, the TAE, and the ER/TAE measure remain stable throughout the series. Therefore, we conclude that our algorithm is stable under the variation of the starting position.

5.4. Various surfaces

In the final part of this section, we highlight that the algorithm can indeed be applied to freeform surfaces, to geometries with sharp features as well as to those with boundary components. First, we would like to illustrate the visual effect obtained by varying the sphere radius. Therefore, we run our algorithm on the *Kitten* model with four different radii. The largest is chosen to be slightly smaller than the reach found in the Kitten's tail, the other three are successively halves of the one before. The result is provided in Fig. 11. As can be seen from the quantitative results in Table 4, the covering rate does not increase significantly. This is despite the fourfold number of spheres employed in each model compared to its predecessor. Similarly, the ER/TAE ratio stays roughly constant. Both effects are mainly due to the fact that we employ the heuristic version of the algorithm to obtain all four models.

Table 4
Quantitative results on a variety of closed, bounded, freeform, and CAGD models. See Table 1 for a legend of the column headings.

Model	Alg.	Radius r	Box diag.	Area	N	$N\pi r^2$ / area	TAE	ER/TAE
Kitten	H	0.025	1.3178	1.6668	676	79.6%	0.05372	13.2%
Kitten	H	0.0125	1.3178	1.6668	2764	81.4%	0.02675	12.5%
Kitten	H	0.00625	1.3178	1.6668	11157	82.1%	0.013341	12.3%
Kitten	H	0.003125	1.3178	1.6668	44746	82.4%	0.006673	12.2%
Bunny	H	0.001	0.2502	0.05821	15084	81.4%	0.0021414	12.6%
Bunny	B	0.001	0.2502	0.05821	14868	80.2%	0.002158	14.6%
Bunny	D	0.001	0.2502	0.05821	13428	72.5%	0.002326	64.7%
Fandisk	H	0.015	7.6156	60.6691	70851	82.5%	0.03182	11.7%
Fandisk	B	0.015	7.6156	60.6691	70365	82.0%	0.03191	13.4%
Fandisk	D	0.015	7.6156	60.6691	63378	73.8%	0.03642	198.5%
Rocker arm	H	0.002	1.1689	1.3180	86227	82.2%	0.004263	12.2%
Rocker arm	B	0.002	1.1689	1.3180	85076	81.1%	0.004292	13.8%
Rocker arm	D	0.002	1.1689	1.3180	77044	73.5%	0.004774	114.3%
Rotational	H	0.0234	13.6643	74.6075	36236	83.5%	0.04998	12.0%
Rotational	B	0.0234	13.6643	74.6075	36153	83.4%	0.04993	13.7%
Rotational	D	0.0234	13.6643	74.6075	31572	72.8%	0.05594	115.5%
Torus	H	0.05	8.7172	78.6545	8221	82.1%	0.1070	12.5%
Torus	B	0.05	8.7172	78.6545	8179	81.7%	0.1071	13.7%
Torus	D	0.05	8.7172	78.6545	7311	73.0%	0.1143	27.3%
Two tubes	H	0.049	12.0739	78.5807	8570	82.3%	0.1070	15.0%
Two tubes	B	0.049	12.0739	78.5807	9054	86.9%	0.1020	11.7%
Two tubes	D	0.049	12.0739	78.5807	8114	77.9%	0.1112	44.0%
Tower	H	0.15	30.3315	603.4286	7123	83.4%	0.3224	12.7%
Tower	B	0.15	30.3315	603.4286	6997	82.0%	0.3237	15.0%

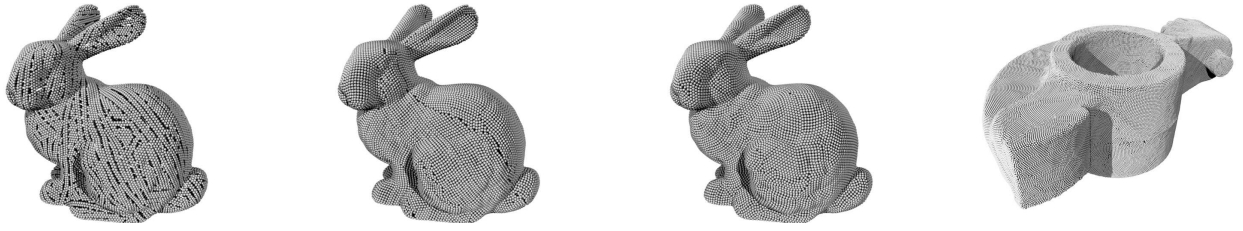

(a) Depth-First. (b) Breadth-First. (c) Heuristic. (d) Heuristic.

Fig. 12. The *Bunny* model, processed with all three versions of our algorithm, and the *Rockerarm* model.

In the *Torus* model, we have already seen the visual benefit of the heuristic approach, see Figs. 6 and 7. Proceeding with closed geometries that do not exhibit sharp features, we consider the *Bunny* and the *Rocker Arm* model. For both models, we can make the same general observation in terms of visual quality as we made for the torus. The number of visual crevasses is highest for the depth-first approach, the breadth-first approach shows a limited number of visible seams, and the heuristic efficiently distributes these defects on the entire geometry, see Fig. 12. The *Rocker Arm* model exhibits fairly straight curves and features, see Fig. 12d. However, the sphere cover is not aligned with these. For improved visual quality, spheres should initially be placed along lines of high curvature, which is left for future work.

Next, we present several freeform surfaces: a *rotational* model, the *Two tubes* model, and the *Tower* model, see Fig. 13. These freeform surfaces are comparable to the models shown in the previously published articles (see Schiftner et al. (2009); Pottmann et al. (2015); Xu et al. (2018)). Note that we achieve true CP meshes by construction of the graph G in all of the models shown in this experimental section. In terms of the descriptor δ_{bp} from Schiftner et al. (2009), or the equivalent measure δ_{ij} (see Xu et al., 2018, Equ. 16), while both report a deviation from the minimum 0, our true CP meshes always obtain a value of 0.

To further compare with the results of Xu et al. (2018), consider the columns *Radius* and *TAE* in Table 4. The first states the used value for the sphere radius r in this specific experiment, while the second gives the average edge length in the triangulation $\mathcal{T}'$ obtained from G as described in the first part of this section. If not only G, but also $\mathcal{T}'$ was a perfect CP mesh, *TAE* would be exactly twice the value of the radius. While this is true for the meshes obtained on the square and presented in Table 2, it is generally not true for the other models, see Table 4. The exhibited deviation is, however, well within the bounds reported by Xu et al. (2018), in their paper via histograms for the respective models.

Note that some of the ER/TAE values in Table 4 are very large for certain covers obtained by the depth-first approach. These are in fact artifacts of the meshing algorithm used to create the triangle mesh $\mathcal{T}'$ from G. Recall that the depth-first approach produces long, connected strips that are pairwise seldom interconnected. Therefore, they are separated by equally long cavities, see e.g. Fig. 12a. The utilized meshing algorithm creates extremely long triangles within these cavities, that

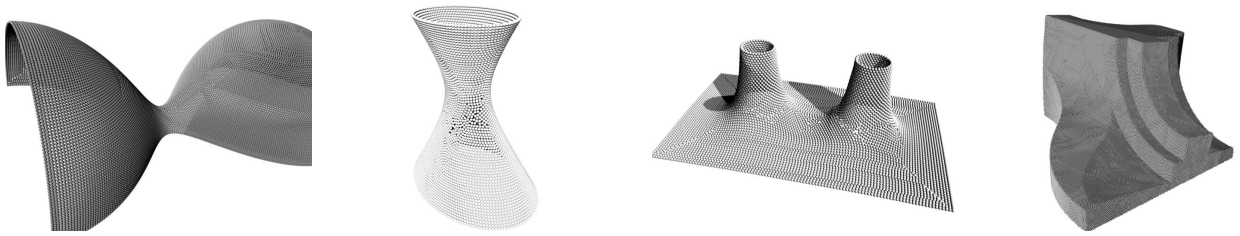

(a) Rotational model. (b) Tower model. (c) Two tubes model. (d) Fandisk model.

Fig. 13. Freeform surfaces and a surface including sharp features, all covered via the heuristic approach.

cause the large quantitative deviations. Similarly, the data for running the depth-first approach on the *Tower* model is not given, as we created a sphere cover, which led to a disconnected graph and thus threw off the triangulation algorithm.

The final geometry we consider is the *Fandisk* model. It is shown in Fig. 13d. As opposed to the other models, this geometry incorporates sharp corners and edges and does thereby not satisfy Assumption 1. However, our algorithm is still capable of handling the model and produces a suitable sphere covering.

6. Conclusion

In this article, we have discussed the creation of sphere cover meshes on surfaces. For surfaces satisfying the conditions of Assumption 1, we were able to prove upper and lower bounds on the area of the surface covered by a sphere that is placed on the surface and has radius depending on the surface's reach, see Equations (3) and (4). From these two bounds, in Theorem 1, we deduced an upper bound on the number of spheres that can be packed on the surface and in Theorem 2, we deduced a corresponding lower bound. As these theoretical descriptions are not constructive, we also provide an algorithm for the explicit creation of sphere covers on any triangulated surface. Our algorithm has three different modes of construction: depth-first, breadth-first, and heuristic. In our experiments, we showed that the heuristic—while being slightly more costly to implement—works best in terms of visual quality and quantitative results.

As opposed to previously published algorithms, (see Schiftner et al. (2009); Pottmann et al. (2015); Xu et al. (2018)), our algorithm does not incorporate the solution of a nonlinear optimization problem. Also, unlike the approach of Lubachevsky and Stillinger (1990), we do not need to handle a large number of collisions between the packed spheres. While the optimization-based approaches do not come with a guaranteed intersection-free circle packing, our algorithm constructs an intersection-free sphere cover by design. The obtained covering rates are comparable or better to those of Lubachevsky and Stillinger (1990), as are the other quantities, when compared to the results of the other listed works (see Schiftner et al. (2009); Xu et al. (2018)). However, the results of our algorithm could improve further by using optimization. For instance, the output could be used as input for the optimization step of previous works. It remains future work to investigate the potential of this direction.

Additionally, we aim at better visual representations of geometries that incorporate high-curvature regions. This can—for instance—be realized by placing a set of initial spheres along feature lines of the input. Furthermore, we plan to investigate sharp cases for the bounds presented in the theoretical section.

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.

References

Boissonnat, J.-D., Lieutier, A., Wintraecken, M., 2019. The reach, metric distortion, geodesic convexity and the variation of tangent spaces. J. Appl. Comput. Topol. 3, 29–58.
Bourke, P., 2006. An introduction to the Apollonian fractal. Comput. Graph. 30 (1), 134–136.
Burke, C.J., Mbanga, B.L., Wei, Z., Spicer, P.T., Atherton, T.J., 2015. The role of curvature anisotropy in the ordering of spheres on an ellipsoid. Soft Matter 11 (29), 5872–5882.
Hales, T., Adams, M., Bauer, G., Dang, T.D., Harrison, J., Le Truong, H., Kaliszyk, C., Magron, V., McLaughlin, S., Nguyen, T.T., et al., 2017. A Formal Proof of the Kepler Conjecture. Forum of Mathematics, Pi, vol. 5. Cambridge University Press, p. e2.
Kharevych, L., Springborn, B., Schröder, P., 2006. Discrete conformal mappings via circle patterns. ACM Trans. Graph. 25 (2), 412–438.
Lubachevsky, B.D., Stillinger, F.H., 1990. Geometric properties of random disk packings. J. Stat. Phys. 60 (5–6), 561–583.
Machchhar, J., Elber, G., 2017. Dense packing of congruent circles in free-form non-convex containers. In: Geometric Modeling and Processing 2017. Comput. Aided Geom. Des. 52–53, 13–27.
Nation kingdom, 2008. available online at https://commons.wikimedia.org/wiki/File:FLY_EYE.jpg.
Pixeltoo, 2005. available online at https://commons.wikimedia.org/wiki/File:Pollen_Ricinus_communis_sanguineus.jpg.
Pottmann, H., Jiang, C., Höbinger, M., Wang, J., Bompas, P., Wallner, J., 2015. Cell packing structures. Comput. Aided Des. 60, 70–83.
Redmayne, J.T., 1877. available online at https://commons.wikimedia.org/wiki/File:Diatom_-_Triceratium_favus.jpg.

Schiftner, A., Höbinger, M., Wallner, J., Pottmann, H., 2009. Packing circles and spheres on surfaces. In: ACM SIGGRAPH Asia 2009 Papers. ACM, pp. 1–8.
Shimada, K., Gossard, D.C., 1995. Bubble mesh: automated triangular meshing of non-manifold geometry by sphere packing. In: Proceedings of the Third ACM Symposium on Solid Modeling and Applications, SMA '95. Association for Computing Machinery, New York, NY, USA, pp. 409–419.
Sosin, B., Rodin, D., Sliusarenko, H., Bartoň, M., Elber, G., 2021. The construction of conforming-to-shape truss lattice structures via 3d sphere packing. Comput. Aided Des. 132, 102962.
Stoyan, Y., Yaskov, G., 2013. Packing congruent spheres into a multi-connected polyhedral domain. Int. Trans. Oper. Res. 20 (1), 79–99.
Tagliasacchi, A., Alhashim, I., Olson, M., Zhang, H., 2012. Mean Curvature Skeletons. Computer Graphics Forum, vol. 31. Wiley Online Library, pp. 1735–1744.
Tóth, L.F., 1943. Über die dichteste Kugellagerung. Math. Z. 48, 676–684.
Wedd, I., 2007. available online at https://pixabay.com/photos/building-architecture-selfridges-572490/.
Wmpearl, 2019. available online at https://commons.wikimedia.org/wiki/File:Fly%27s_Eye_Dome_by_Buckminster_Fuller,_Crystal_Bridges_Museum.JPG.
Xu, Q.-C., Deng, B., Yang, Y.-L., 2018. Ellipsoid Packing Structures on Freeform Surfaces. Computer Graphics Forum, vol. 37 (7. Wiley Online Library, pp. 87–95.