





Topological modeling of metallic foams

Ihab Sabik¹ , Paul Hans Kamm^{2,3} ,
Francisco García-Moreno^{2,3,*}  and Frank H Lutz^{1,†} 

¹ Institut für Mathematik, Technische Universität, Berlin, Germany

² Structure and Properties of Materials, Technische Universität, Berlin, Germany

³ Institute of Applied Materials, Helmholtz-Zentrum, Berlin, Germany

E-mail: garcia-moreno@helmholtz-berlin.de

Received 2 May 2025; revised 17 August 2025

Accepted for publication 26 September 2025

Published 17 October 2025



CrossMark

Abstract

Understanding material's micro- and macrostructure is key to linking inner structure information to large-scale physical properties and function. An improvement of the microstructure during production processes can lead to materials with enhanced physical properties or even to the development of new classes of materials. In this paper, we develop an 'accurate' topological model to describe the dual neighborhood structure of foams and other cellular materials (e.g. plant cell clusters or polycrystalline materials like steel, rocks, and ceramics) obtained by x-ray tomography, adhering to Plateau's rules: four pores meet at a vertex, and three meet along an edge. While the degree of non-convexity in the pores is *a priori* unknown, our approach is designed for structures without extreme geometric irregularities (e.g. tangled pores). It provides meaningful combinatorial-numerical validation in such cases but may not extend reliably to more complex geometries. To illustrate our method, we analyze a sample of aluminum foam comprising 1911 pores. The resulting abstract combinatorial foam structure partially aligns with that of the associated power diagram—sharing 51.45% of the tetrahedra in the dual complex—but also reveals significant deviations.

Keywords: topological modelling, metallic foams, x-ray tomography

* Author to whom any correspondence should be addressed.

† Deceased.



Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

1. Introduction

In its broadest sense, a foam is a structure comprising gas-filled pockets, commonly referred to as pores or bubbles, enclosed by a continuous phase, whether it be liquid or solid. When the surrounding medium is liquid, the gas-filled pockets are commonly known as *bubbles*, while in a solid medium, they are termed *pores*. The *liquid volume fraction* (or *solid volume fraction*) is defined as the ratio of the liquid phase (or solid phase) within the foam.

Foams exhibit characteristic structural patterns governed by Plateau's laws, which describe how the interfaces between bubbles or pores meet in equilibrium. In three dimensions, three films meet along an edge at angles of 120° , and four edges meet at a vertex with tetrahedral symmetry. These geometric rules, first formulated by Plateau [87] based on empirical observations and later proved by Taylor [99], are fundamental to understanding both the physical and combinatorial structure of foams.

In the realm of 3D, the term *liquid foam* is specifically applied when the liquid volume fraction falls within the range of 0 to 0.36. Once the liquid fraction surpasses 0.36, the system is categorized as a *bubbly liquid*, where the bubbles essentially move independently of each other (see [38, 48]). Similarly, the term *solid foam* is employed when the solid volume fraction spans from 0 to 0.5, and the system is referred to as a *porous solid* beyond that point; refer to [14]. This distinction arises from the significantly larger critical film thickness of metallic foams (50–80 μm) [61, 95] compared to liquid foams ($< 1\text{--}2 \mu\text{m}$) [48]. Consequently, the solid volume fraction in metallic foams can typically span a broad range.

Foams can be *monodisperse*, with pores (or bubbles) of equal volume, or *polydisperse*, with pores of varying size. A foam is considered effectively monodisperse if the dispersity (the ratio of standard deviation to mean of the pore diameter distribution) is less than 5% [47].

Different types of foams, such as polymeric [79] or metallic foams [39], also find applications in industry. Metallic closed-cell foams are typically produced by solidifying the molten phase [11], resulting in a solid foam [41].

1.1. Structure and formation of metallic foams

The structure of metallic foams is notably more intricate than that of liquid foams, presenting geometric challenges due to its dependence on various formation process parameters. Film stabilization in the liquid state is facilitated by solid particles [9] or oxide networks [61]. The movement of bubbles during structural rearrangements and coalescence is impeded by the stabilizing influence of the solid phase. The evolving bubble films undergo oxidation, dependent on the alloy composition, resulting in a corrugated skin during foam solidification [45].

Furthermore, processing parameters, including alloy composition, viscosity, surface tension, the blowing agent used, surrounding atmosphere, pressure, temperature profile, etc play crucial roles in determining the final structure [40]. In instances where active gas sources are employed (as in foams produced via the powder metallurgical route), some nucleated neighboring bubbles may not receive gas due to the high Laplace pressure of small bubbles. Consequently, gas preferentially diffuses to larger bubbles. The inner surface of these small bubbles oxidizes over time, remains stable, and results in the formation of micropores or satellite pores [81]. Following the observations of Ohgaki *et al* [86] and Mukherjee *et al* [81], we categorize pores with an equivalent diameter $< 9.5 \mu\text{m}$ as micropores and exclude them from the analysis. These factors collectively contribute to the structural evolution of liquid metal foams, with the foam pores typically becoming *non-convex* after solidification [8, 10].

1.2. Geometric representations and diagrams

In scenarios involving very small bubbles and a high liquid volume fraction, the bubbles essentially ‘appear’ as points in an ambient space. Here, the Voronoi decomposition of space captures the proximity information of individual bubbles. Notably, the independent bubbles in a bubbly liquid tend to assume the shape of round balls due to the isoperimetric theorem (see [98]). In such cases, the Apollonius diagram [20, 33, 50, 92, 103] or the Laguerre–Voronoi (power) diagram [1, 54, 69, 74, 89, 104], both generalizing the standard Voronoi diagram, are employed for space decomposition of the round bubbles. The Apollonius diagram, defined via Euclidean distance, provides a more accurate description of bubbly liquid geometry, particularly as the liquid fraction approaches 0.36 (see figure 7). For a recent implementation to compute the Apollonius diagram in 3D, see [103].

1.3. Application context and industrial relevance

The primary focus of this study is on metallic foams, specifically aluminum foams, and their combinatorial representation. Metallic foams, characterized by a multitude of gas pores, are lightweight engineering materials with remarkable physical and mechanical properties such as low density, high stiffness, and energy absorption properties [84]. These materials are increasingly used in transportation and energy applications for weight reduction and crashworthiness [12]. Figure 1(left) illustrates an aluminum foam sample with 1911 pores, serving as a showcase example discussed in detail throughout this paper, and figure 1(right) presents our cellular description of it. The fabrication process of this foam sample is explained in section 4.

The production of metallic foam involves the introduction of gas bubbles into a metallic liquid, leading to solidification. This is achieved through methods such as gas injection or chemical reactions [13]. The resulting material comprises a metallic framework with voids. Unlike other anisotropic materials like honeycombs, commercially available metal foams typically display isotropic load distribution due to random pore distribution and a polydispersed Gaussian pore size distribution [57]. Recent advancements allow for the production of aluminum-alloy foams with nearly monodisperse pore size distribution, catering to specific applications [78, 85, 101]. In certain instances, achieving short-range order between pores results in anisotropic properties [102].

1.4. Topological problem and research objectives

A key property of real-world foams is their low *average face degree*, or the average number of neighbors per pore. While this value can be arbitrarily large in a general 3D geometric decomposition (see [72]), studies by material scientists [75] and mathematicians [66] show that in simple real-world foam configurations, the average face degree—and consequently, the average vertex degree in the dual triangulation—typically ranges from 11 to 16, with most values between 12.5 and 14.5.

The main objective of this study is to develop a topological model for describing the dual neighborhood structure of cellular materials—specifically, metallic foams reconstructed from x-ray tomography. We aim to create a model that adheres to Plateau’s rules and captures meaningful structural information even when pores are not convex. Section 3 provides the necessary groundwork for our computations, and section 4 delves into the production process of the foam sample with 1911 pores. Following that, in section 5, we elucidate our topological modeling of simple regular CW complexes, which serve as space decompositions for metallic foams, particularly in cases involving general, non-convex pores. In sections 6 and 7, we discuss the




Figure 1. Left: X-ray tomogram of an aluminum foam sample with 1911 segmented pores. Each pore is shown in a distinct color solely for visual differentiation; the color encodes no geometric or topological information. Right: Our combinatorial-topological reconstruction of the sample as a 3D cellular complex. Each reconstructed cell is assigned the same color as its corresponding pore in the tomogram to visually link input and output. Axis ticks are in voxel units, where each voxel measures $41 \mu\text{m}$ per side.

(partial) reconstruction of a triangulated closed manifold from its graph. The main sections 8 and 9 are then devoted to a step-by-step description of our approach.

Although our primary case study involves metallic foams, the proposed model is applicable to a broader class of materials exhibiting similar cellular structures—including polycrystalline materials (e.g. steel, rocks, ceramics) and plant cell clusters—as long as they conform to the combinatorial Plateau rules.

2. Related work

In computational geometry, an extensive body of literature is dedicated to computing generalized, abstract Voronoi diagrams (and their dual Delaunay decompositions) for various 2D objects beyond points or disks. This includes disjoint line segments [19], disjoint convex polygons [76], pseudo-circles sets [55], and disjoint smooth convex objects [18]. For sets of ellipses in 2D, [32] provides an exact computation of their Voronoi diagram.

Efficient algorithms that yield exact results for objects in higher-dimensional spaces \mathbb{R}^d are available for sets of points (Voronoi diagrams, [28]), sets of discs (power diagrams, [7]), as well as for sets of line segments [18]. However, allowing non-convex objects introduces challenges and artifacts, such as regions with pinch points, pockets, or disconnected regions; refer to [4, 59] for a systematic study. Curved Voronoi diagrams of objects are closely related to their medial axis [18].

In 3D, the computation of bisector surfaces between neighboring bodies is a well-explored topic [31, 68, 106]. In general, regions bounded by bisector surfaces may not be homeomorphic to three-dimensional balls; the boundary surfaces might exhibit pinch points or other complications (see [59] for 2D illustrations). When dealing with concrete objects in 3-space, the evaluation of surface reconstruction algorithms poses challenges. As highlighted by Cazals and Giesen in [25], certain algorithms provide theoretical guarantees under specific conditions. However, their behavior is unspecified when these conditions are not met. Evaluating their performance on real data becomes a complex task.

A notable approach for general convex objects, beyond round balls or ellipsoids, involves considering the *set Voronoi cell*. This cell comprises points in space closer to the boundary

surface of a given body than to the surface of any other body. An algorithm in [93] computes approximations of the set Voronoi decomposition for convex 3D bodies, converging to the medial axis bisector surfaces as the resolution approaches infinity. Nevertheless, for non-convex bodies, without controlled and locally adapted resolution, this approach may inaccurately model the neighborhood structure, potentially leading to non-connected regions.

In practical scenarios, a foam occupies a bounded region in 3-space. However, for a cohesive description and to mitigate boundary effects, we assume that a foam decomposes a closed three-dimensional manifold M^3 . Two immediate choices for the underlying 3-manifold are the three-dimensional sphere S^3 or the three-dimensional torus T^3 . The latter is commonly used in simulations, representing a cuboid with periodic boundary conditions to model (flat) Euclidean space. On the other hand, the former is more suitable for analyzing measured data, where the outer region of a bounded foam in \mathbb{R}^3 is regarded, through inverse stereographic projection, as a foam bubble in S^3 at infinity. Topologically, a foam cluster (including the outer region) can be described as a *simple* regular CW complex (refer to the [appendix](#) and [definition 1](#) for formal definitions) decomposing the three-dimensional sphere S^3 (or some closed 3-manifold M^3).

Natural regular CW decompositions are associated with foams at specific volume fractions. As the volume fraction approaches 1, the foam's corresponding clipped Voronoi diagram [105], along with the outer region, forms a simple regular CW decomposition of S^3 . This assumes that the small pores are in general position. Similarly, for volume fractions within the ranges 0–0.2 and 0.36–1, the Apollonius and Laguerre–Voronoi diagrams for disks of varying sizes in ‘general position’ result in simple regular CW decompositions. However, there is currently no established model for foams with volume fractions between 0.2 and 0.36.

In this intermediate range, the standard approach for simulations and analyzing concrete data involves replacing an individual (generally non-convex) pore of a given volume with an *equivalent sphere* of the same volume centered at the barycenter of the original pore [62, 63]. Subsequently, the associated 3D Laguerre–Voronoi diagram or the 3D Apollonius diagram is computed (with an algorithm for the latter described only recently [103]). While this approach serves as a rough approximation of the corresponding cellular decomposition for the foam, the non-convexity of individual pores can lead to swift inaccuracies in local and global results, such as deviations in the numbers of neighbors of a pore and consequently the overall average face number.

In this work, we introduce a novel topological approach to model the (dual) *neighborhood structure* of a solid foam featuring (generally) non-convex pores within the solid volume fraction range of 0.2 to 0.36. This serves as a foundational step to extract additional combinatorial and geometric information from a foam sample, including determining the average face number and statistics on the geometric shapes and combinatorial types of individual associated cells.

To achieve this, we compute the neighborhood graph of the foam pores, which forms a generalized triangulation (see [23] for definitions). Its dual, a regular CW complex, decomposes space into fundamental building blocks. Each foam pore is approximated by the dual of a triangulation vertex, enabling the assignment of a combinatorial polytope-potentially with 2-gonal faces-to each pore.

It's crucial to emphasize that the output is *not* a uniquely determined decomposition, but instead belongs to a class of decompositions that effectively capture the neighborhood structure of the foam, up to local fluctuations and possible defects. This methodology draws parallels with the idea of selecting *a* Delaunay triangulation from the set of all Delaunay triangulations that triangulate a Delaunay decomposition of a point set not in general position.

3. Combinatorial topology foundations

We begin with some basic definitions. A triangulation is *neighborly* if every pair of vertices is connected by an edge, making its edge graph complete. In a triangulation, a set of vertices $X = \{x_1, \dots, x_{n+1}\}$ with $n \geq 3$ forms an *empty n -simplex* if every n -subset of X is a simplex in the triangulation, while X itself is not. A triangulation is *flag* if it contains no empty simplices. For equivalent definitions, see [6] and [73].

Recall that a *CW complex* is a topological space built by successively gluing cells of increasing dimension (0-cells, 1-cells, 2-cells, etc) along their boundaries in a coherent way. A *regular CW complex* is one where the closure of each cell is homeomorphic to a closed ball, and its boundary intersects the complex in a union of lower-dimensional cells. A *strongly regular CW complex* further requires that the attaching maps be homeomorphisms onto their images, ensuring that every face relation is uniquely determined and that no identifications occur along the boundary. This added regularity allows us to associate unambiguously a dual adjacency graph to the structure, which is essential for modeling the connectivity of foam pores. Refer to the [appendix](#) for technical definitions.

Our primary focus is on constructing the neighborhood structure of the pores in a nearest-neighbor fashion. This involves incrementally adjusting a threshold c_1 , which in turn determines which pores are considered neighbors based on their distance being less than c_1 . Ideally, this process yields a dual graph that forms the edge graph of a dual generalized triangulation, formally called its *1-skeleton*. However, this path involves resolving various problems and overcoming challenges in a heuristic manner.

Firstly, reconstructing a triangulation from its 1-skeleton is generally not possible. Extreme cases include neighborly triangulations, where reconstruction is impossible, and flag triangulations like Vietoris–Rips complexes, which guarantee reconstruction.

Warning: The Hauptvermutung [80, 91] states that triangulations of a given compact 3-manifold are all connected by bistellar flips. Consequently, there is no clear distinction between triangulations that can be reconstructed from their 1-skeleton and those where this is not possible. This implies that our approach is doomed to fail for arbitrary configurations of pores.

This consideration brings us back to real-world foams and other cellular materials. Having a low average face degree implies that the triangulations we consider are far from being neighborly and often exhibit flag patches, enabling partial local reconstruction. This reconstruction serves as the backbone for a triangulation that is then completed by assessing further local geometric information on the pores.

Secondly, when building the neighborhood graph, errors arise, including *missing edges* and *excessive edges*, as discussed in section 7. Missing edges create holes in the triangulations, which we address by allowing diagonal edges up to a second threshold, $c_2 > c_1$ to close quadrilateral holes for instance. In some cases, choices on which diagonal edges to pick result in triangulations that differ locally.

This ambiguity does not impair material analysis, since we are primarily interested in global topological and combinatorial statistics—such as average face degree, genus, or pore connectivity—rather than a uniquely defined triangulation. Consequently, we define our output as a member T of a class of triangulations $\mathbb{T}(c_1, c_2; E, F)$, where E denotes resolvable *exceptions* and F refers to unresolved *failure cases*.

This flexible, combinatorially-driven approach is preferable in practice for several reasons:

- It accommodates real-world noise and defects in tomographic data;
- It supports partial reconstruction where local conditions allow;
- It yields statistically stable outputs even when local triangulation choices vary;

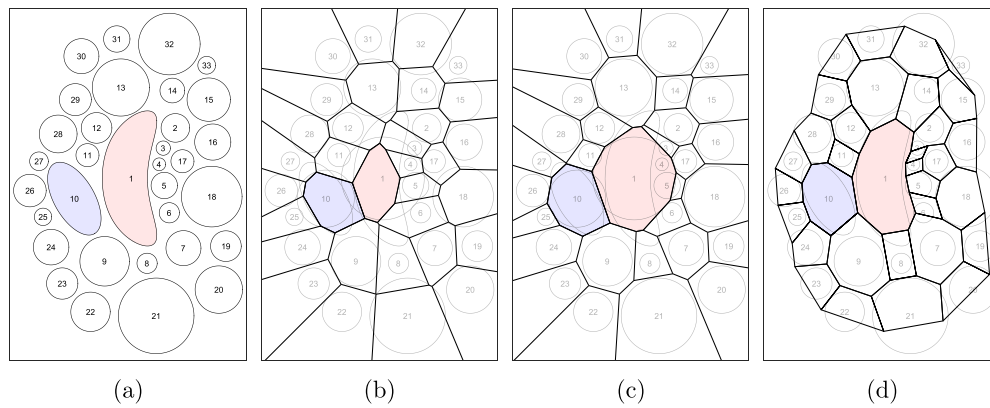


Figure 2. (a) A configuration of 33 two-dimensional foam pores, including 31 discs, one ellipse, and one non-convex pore. (b) Voronoi diagram of the barycenters of the pores. (c) Laguerre–Voronoi diagram associated with 33 area-equivalent disks centered at the barycenters. (d) ‘Combinatorially correct’ cellular description. Two adjacent pores are colored consistently across all panels to trace their structural representation. All other pores are shown in neutral tones. Colors serve only to track these two specific pores and do not reflect topological or geometric properties. A scale bar is not relevant in this schematic illustration.

- And it reflects the structural redundancy and resilience seen in actual foams.

Definition 1. A *simple* regular CW decomposition of a closed 3-manifold M^3 is a regular CW complex homeomorphic to M^3 , where cells meet in threes around edges and in fours around vertices.

This definition encodes the combinatorial aspect of Plateau’s rules. The boundary of a single cell in a simple regular CW complex is a *simple* decomposition of the two-dimensional sphere into p -gons with $p \geq 2$, such that three polygonal boundary pieces meet at every vertex. It’s important to note that if all individual cells are simple, the cluster need not be simple. For instance, if unit cubes tile some portion of size $k \times k \times k$ of the integer grid in \mathbb{R}^3 , and assuming periodic boundary conditions, then the entire decomposition of the 3-torus into k^3 cubes is not simple, even though each cube is simple.

Example 2. Figure 2 illustrates (a) 33 pores, comprising 31 discs, 1 ellipse, and one non-convex shape bounded by the level curve $x^2 + (y^2 - x)^2 = 0.1$ of Coste’s example [17]. The figure also presents (b) the Voronoi diagram of the barycenters of the pores, (c) the Laguerre–Voronoi diagram of the equivalent disks replacing the pores, and (d) a ‘combinatorially correct’ cellular description of the pore configuration. The non-convex pore # 1 in (a) is represented by an 8-gon in (b), a 10-gon in (c), and correctly as a 12-gon in (d). The MATLAB extension ‘Power Diagrams’ by McCollum [77] was employed to compute the power diagram and, with zero weights, derive the standard Voronoi diagram in both figures 2 and 3.

In example 2, the identification of neighboring pores is intuitive. However, in 3D, especially for non-convex pores, defining ‘representing neighborhood structures’ becomes more intricate and forms a fundamental aspect of our methodology.

Handling non-convex shapes necessitates certain adjustments. Instead of utilizing geometric polytopal complexes (as in Laguerre–Voronoi modeling) to represent pore configurations,

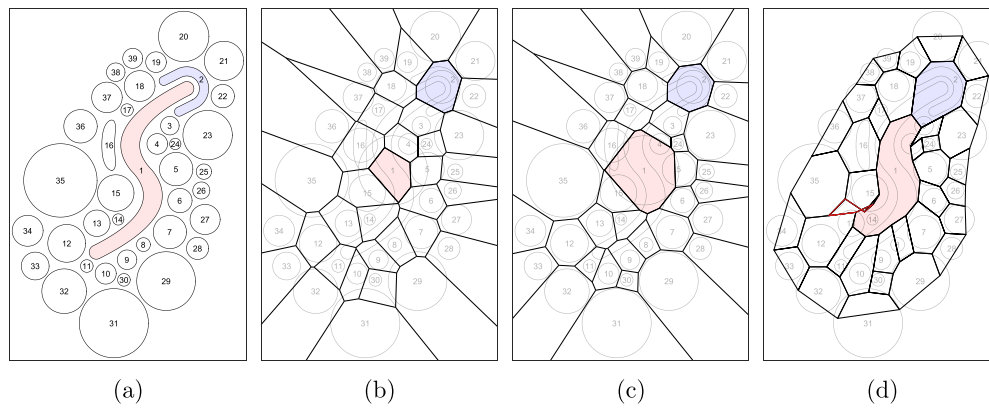


Figure 3. (a) A configuration with 39 pores, (b) the corresponding Voronoi diagram, (c) the corresponding Laguerre–Voronoi diagram, and (d) a ‘combinatorially correct’ cellular description. Two adjacent pores are highlighted consistently across all panels to illustrate structural correspondence; colors are purely illustrative and carry no geometric or topological meaning. A scale bar is not relevant in this schematic illustration.

we opt for describing the neighborhood structures of these configurations through *abstract* simple regular CW complexes. Some cells within these associated CW complexes have digonal faces (similar to an Apollonius modeling) and do not possess the combinatorial type of a polytope; that is, the cells need not be polytopal complexes, but rather regular CW-complexes (consult the [appendix](#) for definitions). Nevertheless, substantial portions of the representing CW complexes for foam pores consist of abstract polytopal cells.

Definition 3. A *pseudo-realization* of an abstract simple regular three-dimensional CW complex associated with a 3D configuration of pores selectively represents cells that are abstract polytopes. This representation introduces vertices positioned at the average of the four barycenters of the pores where their representing cells meet. Subsequently, connections between neighboring vertices are established using straight line segments.

For boundary vertices, where three boundary pores $\mathcal{A}, \mathcal{B}, \mathcal{C}$ meet, choose a representation by computing the average of the barycenters of $\mathcal{A}, \mathcal{B}, \mathcal{C}$ plus the difference vector obtained by subtracting the average of $\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{D}$ from the average of $\mathcal{A}, \mathcal{B}, \mathcal{C}$, where \mathcal{D} is a mutual neighbor of $\mathcal{A}, \mathcal{B}, \mathcal{C}$ towards the interior (usually, \mathcal{D} is the unique mutual neighbor). Symbolically, this can be expressed as

$$2 \cdot \text{bary}(\{\text{bary}(\mathcal{A}), \text{bary}(\mathcal{B}), \text{bary}(\mathcal{C})\}) - \text{bary}(\{\text{bary}(\mathcal{A}), \text{bary}(\mathcal{B}), \text{bary}(\mathcal{C}), \text{bary}(\mathcal{D})\}).$$

To illustrate informative 2D examples, this definition can be adapted accordingly to two dimensions. Refer to figures 2(d) and 3(d) for the corresponding two-dimensional pseudo-realizations.

Example 4. Figure 3 displays a configuration featuring (a) 39 pores (36 disks and 3 non-convex pores), along with (b) the Voronoi diagram, (c) the Laguerre–Voronoi diagram, and (d) a ‘combinatorially correct’ cellular description. Key observations include:

- In (a), pores # 1 and # 2 are neighboring, yet they appear separated in the corresponding Voronoi and Laguerre–Voronoi diagrams.

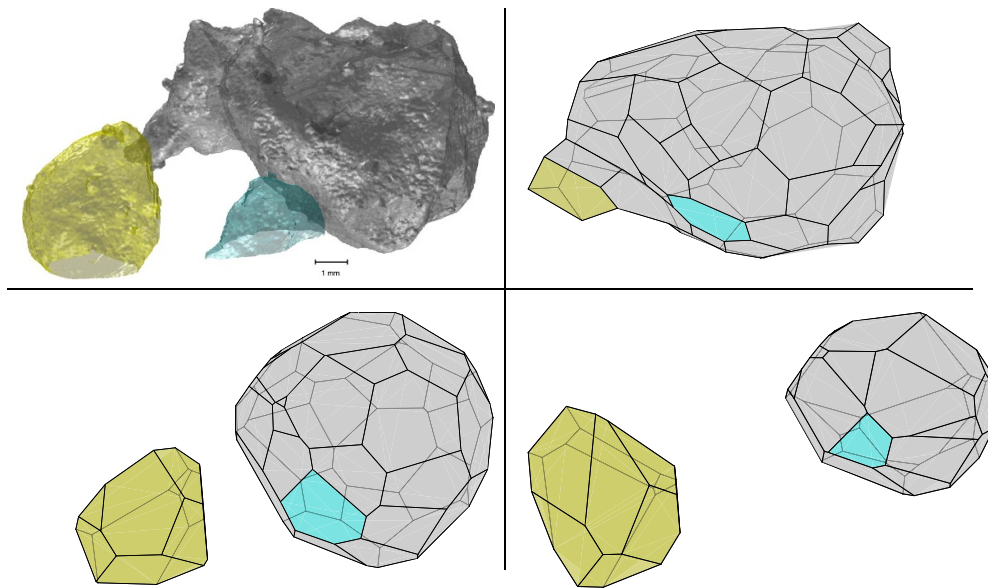


Figure 4. Top left: The second-largest pore #356 (gray) with two neighbors (yellow and blue). Top right: Its pseudo-realization from our reconstruction, which has 59 facets, including two highlighted facets (yellow and blue) corresponding to intersections with the pseudo-realizations of the yellow and blue pores. Bottom left: The Laguerre regions of the gray and yellow pores, which are not identified as neighbors; the Laguerre cell of the gray pore has 49 facets, and the blue facet (6-gon) marks the intersection with the blue pore's Laguerre cell. Bottom right: The Voronoi regions of the gray and yellow pores, also not identified as neighbors; the Voronoi cell of the gray pore has 32 facets, with the blue facet (5-gon) indicating the intersection with the blue pore's Voronoi cell.

- The geometric indentation of pore # 1 into pore # 2 is irrelevant for the overall neighborhood structure depicted in (d).
- The pseudo-realization (d) reveals an overlap in the pentagonal polygon representing pore # 13 within the cellular neighborhood structure of (a). The edges of this polygon are highlighted in red.

The pseudo-realization of the foam's neighborhood structure provides an overall view of the complex's structure, even though local voids (resulting from cells with digons being omitted) or folds may exist. For instance, figure 11 illustrates two multi connected pores and their pseudo-realized two-dimensional simple polytopal boundary spheres. The pseudo-realization of the simple boundary 2-sphere corresponding to pore #448 is appropriately represented as a surface on the visible side but exhibits two folds on the back side (not visible in the figure). Further comparisons between our pseudo-realization and the corresponding Voronoi and Laguerre cells of two selected pores are shown in figures 4 and 5. Both figures demonstrate that our reconstruction identifies more neighbors for ill-shaped pores.

While a more precise visualization might involve connecting the vertices of cells with curved segments rather than straight lines, *our goal and output in this work are abstract CW complex decompositions*. We provide the pseudo-realization (of the entire complex or individual cells) merely for illustrative purposes.




Figure 5. Elongated pore # 1154 (left) with the pseudo-realization of our reconstruction (middle) featuring 34 facets, while the associated Laguerre region (right) has only 13 facets.

In time-evolving simulations of cellular materials, employing tools like the surface evolver [21], two prevalent models for representing the underlying geometry, combinatorics, and topology are (abstract) polytopal complexes or simple regular CW complexes [94]. This includes simulations of foam complexes [34, 35] or polycrystalline metals [67, 72]. Such simulations often entail dynamic alterations in cell structures, involving processes like merging, shrinking, or disappearing. The incidence structure of the chosen model, whether polytopal or simple regular CW complex, proves instrumental in tracking combinatorial changes and managing topological transitions.

If, instead of starting with a predefined setup (as for the simulations), we want to assess real foams, we immediately face the following question.

Problem 5. *Given a configuration of pores, how can we accurately determine the incidence structure of a corresponding cell decomposition?*

Theorem 6. *Let S be a simple regular CW decomposition of a closed topological 3-ball $B \subseteq \mathbb{R}^3$ into m cells C_1, \dots, C_m . Let $B_1, \dots, B_m \subseteq B$ be m closed and pairwise non-intersecting topological 3-balls contained in the interior $\overset{\circ}{B}$ of B . Then there exists a self-homeomorphism f of B such that $B_{\pi(1)} \subseteq f^{-1}(\overset{\circ}{C}_1), \dots, B_{\pi(m)} \subseteq f^{-1}(\overset{\circ}{C}_m)$ for any permutation π of the set $\{1, \dots, m\}$.*

Proof. First, we shrink each of the balls B_1, \dots, B_m via a homeomorphism of B to standard (round) balls $B_{1,\epsilon}, \dots, B_{m,\epsilon}$ of a suitably small radius $\epsilon > 0$. Second, using additional homeomorphisms, we adjust the coordinate values of the centers of the shrunk balls, allowing us to freely and independently move them within $\overset{\circ}{B}$. Let f be the combined homeomorphism, and f^{-1} its inverse image.

Thus, for a given permutation π , we can ensure that $f(B_{\pi(1)}) \subseteq \overset{\circ}{C}_1, \dots, f(B_{\pi(m)}) \subseteq \overset{\circ}{C}_m$, which implies $B_{\pi(1)} \subseteq f^{-1}(\overset{\circ}{C}_1), \dots, B_{\pi(m)} \subseteq f^{-1}(\overset{\circ}{C}_m)$. \square

Example 7. Consider the six balls $B_1, \dots, B_6 \subseteq \mathbb{R}^2$ in figure 6(left), and a simple cell decomposition of the entire ball $B \subseteq \mathbb{R}^2$ into six cells C_1, \dots, C_6 (middle). In the right part of the figure, this cell decomposition is pulled back, where $\pi(i) = i$ for $i = 1, \dots, 6$.

Theorem 6 establishes that, from a topological perspective, any simple regular CW decomposition, provided it has the correct number of cells, is applicable for a given configuration of pores. In the realm of physics, the conventional approach for ‘determining’ the combinatorial neighborhood structure of a 3D foam cluster involves substituting individual, potentially non-convex pores with ‘equivalent spheres’ of equivalent volume (typically quantified by the number of voxels in a CT scan of the cluster). In a subsequent step, the associated 3D Laguerre–Voronoi (power) diagram is computed to generate a polyhedral representation of the cluster. In certain scenarios, opting for the computation of the 3D Apollonius diagram may be more appropriate. This is particularly true as the model relies on Euclidean distances to assess pore proximity. However, it is worth noting that the resulting bisector surfaces between pores are




Figure 6. A collection of pores (left) and an associated space decomposition (right) that conforms to a predefined cellular decomposition (middle).




Figure 7. Four disks with corresponding Laguerre–Voronoi diagram (left) and Apollonius diagram (right), along with their associated dual triangulations.

algebraic surfaces rather than flat polygonal faces. This characteristic introduces challenges when it comes to visualizing individual pores and the overall cluster.

In figure 7, we observe that for a 2D example, the choice of the model can lead to *distinct identifications* of neighboring pores. The physics of a system often suggests a natural model. For instance, consider the gravitational force between two homogeneous balls in empty space; in this scenario, a well-defined bisector surface marks the boundary between the two regions of attraction for the two bodies (corresponding to the Apollonius diagram in this example). This bistellar flip, representing the transition between the two dual triangulations in figure 7, is also known as the T1 transformation. The T1 transformation is a process where two neighboring bubbles swap a common edge to optimize local energy, and it is commonly employed in foam studies to dynamically adjust the bubble configuration.

In the absence of a specific describing force and moving beyond simplified models using, for example, disks or ellipses in 2D or balls or ellipsoids in 3D, especially when considering non-convex pores, we aim to define a ‘neighborhood structure’ for a cluster of objects based on their proximity. Naturally, pores that are ‘far apart’ should not be considered neighbors. This consideration rules out decompositions covered by theorem 6 that involve pairs of far apart pores as neighbors.

Consequently, our desired outcome is a (member of a) *class of cellular decompositions* rather than a unique decomposition. This class encompasses all decompositions that globally describe the cluster of pores but allows for local fluctuations, as illustrated in figure 7. This approach offers two significant advantages:

- Independence from *a priori* assumptions about the specific geometric shapes of the pores. There is no requirement for the pores to be round balls, ellipsoids, or convex objects, as often mandated by existing models in the literature.
- Independence, particularly in the case of measured data, from the specific voxelization and resolution of the pores, with only minimal preprocessing considerations (refer to section 5.1).

We exemplify our approach by examining a specific sample of an aluminum foam with 1911 pores. The resulting abstract combinatorial foam structure aligns partially with that of the associated power diagram but also exhibits notable differences.

4. Fabrication of an explicit aluminum foam sample

The study sample with 1911 pores was produced by means of gas injection into a melt: Alloy slabs of pre-fabricated metal matrix composite Duralcan F3S20S alloy (aluminum alloy 359 + 20 wt% SiC particles) were molten in a dedicated furnace (see figure 8) and diluted to reduce the particle content to 11 wt% SiC. The temperature of the melt was measured *in situ* using a thermocouple immersed directly within the melt. Further, gas was injected into the melt through a rotating, thin, and bent cannula to obtain small bubbles with a narrow size distribution. The introduction of synthetic air, comprising 20.5 vol% O₂, was initiated subsequent to the attainment of the target melt temperature of 670 °C. For cannula tip velocities below 0.4 m·s⁻¹, a pressure-regulated system was employed to maintain a consistent pressure of 1300 mbar. Conversely, for higher cannula tip velocities, a gas flow control device manufactured by Voegtlin, Germany, was utilized, operating within a flow rate range of 20 to 30 mln·min⁻¹. The obtained foam was conserved by switching off the furnace and natural cooling. Refer to [85] for a full description of the process.

X-ray tomography of the metal foam under investigation was performed using a laboratory setup consisting of a microfocus x-ray source (spot size 5 μm, at 100 kV and 100 μA), a flat panel detector (120 mm × 120 mm, pixel size 50 μm), both from Hamamatsu Photonics, Japan, and a turntable from XHuber, Germany. The sample was rotated 360°, 1000 projections were made and reconstructed into a volume with a voxel size of 25 μm. The setup is described in detail in [85].

A median filter is used to remove noise from the reconstructed volumes and the pores are segmented by a simple threshold. Since the material wall thickness between the gas phases lies in part within the range of spatial resolution, some pores appear to be connected after segmentation. In order to separate these pores properly from each other, a watershed transformation is carried out and the separated regions are consecutively numbered.

5. Topological modeling of clusters of pores

5.1. Modeling of individual pores

We are not overly concerned about the specific shapes of individual pores, as long as the entire input constitutes a uniform foam having a volume fraction in the range 0.2–0.36. In such cases, individual pores generally exhibit shapes akin to topological balls (not necessarily convex). These pores are permitted to have small or even larger holes. At the level of voxelized input, the boundaries of the pores need not be perfect spheres and can display various local defects. Allowing such a general input enhances the resilience of our approach to local deformations of the pores. This is possible because we primarily process combinatorial structure (with some limited geometry), which facilitates a robust implementation.




Figure 8. Left: Gas injection setup consisting of a furnace with a rotary feedthrough and a motor; Right: Diagram of the curved cannula, illustrating the offset of the cannula tip from the rotation axis. Reprinted from [85], Copyright (2021), with permission from Elsevier.

Example 8. Figure 9(left) presents three pores from the study sample. In this illustration, the blue pore exhibits a large hole, resulting in a (deformed) toroidal shape. While identifying and removing tiny holes is straightforward through persistent homology computations, larger holes might persist as features. However, for determining the neighborhood structure of the pores in this example, the presence of the hole is inconsequential. This is because the green pore is a neighbor of the red pore, regardless of whether the hole is filled. In alternative scenarios, this may not hold true, necessitating the severance of the connection between the red and green pores, as outlined in definition 9.

Definition 9. A *valid input* for a foam is a configuration of pores such that

- Each pore can be thickened into a topological ball in such a way that the interiors of all resulting balls are pairwise disjoint,
- The pores are not linked, and
- The foam is uniform with a volume fraction between 0.2 and 0.36.

Example 10. In the case of a liquid foam, our definition would encompass scenarios where a toroidal bubble is paired with a dumbbell-shaped bubble filling the hole, as discussed in [97, 98].

In definition 9, the nature of a pore is not yet specified. A practical input from measurements involves a 3D block, such as a 3D CT scan, containing $a \times b \times c$ voxels. Through a watershed algorithm, each voxel is classified as either a metal (liquid) voxel or an empty (gas) voxel. Subsequently, empty voxels are grouped into collections, forming gas bubbles of *strongly-connected* voxels. Here, two voxels are strongly connected if there exists a path connecting them, with the centers of the voxels as vertices and an edge representing a shared square face.




Figure 9. Left: Toroidal pore # 1376 (blue) alongside pores # 1186 (red) and # 1496 (green). Right: Elongated pore # 1154 (red) with numerous neighbors relative to its volume; see figure 10(right). This configuration features a *short excessive edge* between pore # 1245 (light-blue) and pore # 890 (yellow).

Strong connectedness in image processing is often referred to as 6-connectedness, indicating that in a 3D grid, a single cube has 6 neighboring cubes. This stands in contrast to 18-connectedness and 26-connectedness, where neighboring cubes can share an edge or a vertex, respectively. For computational aspects, refer to [37, 49, 88].

Definition 11. A *pore* is a strongly connected component of gas voxels.

For a pore, its *boundary* comprises the square faces of the pore voxels that are a face of only one of the voxels in the pore. Generally, the boundary may not be a cubical surface and can exhibit various defects.

It's possible for two distinct pores to be connected (26-connected or even 18-connected) but not 6-connected. In such cases, especially when one of the pores is composed of only a few voxels, it could be an artifact from the watershed process. Another source for small bubbles, aside from artifacts, is the presence of *satellite bubbles*, which may occur during the foam formation process [81]. While satellite pores that form a 'bridge' between larger pores can be physically relevant, isolated ones are typically negligible.

In our dataset of 1911 pores, each voxel has a length of $41 \mu\text{m}$. We conducted a preprocessing step to eliminate all strongly connected voxel clusters with fewer than 4189 voxels, effectively excluding small pores whose equivalent spheres have a radius of at most 10 voxels (equivalent to a diameter of 0.82 mm). Figure 10 illustrates the volume distribution of the 1911 pores, ranging from the smallest pore with 4300 voxels to the largest with 5151 369 voxels. Table 1 provides the volumes and the number of neighbors (as determined by our reconstruction algorithm) for the largest pores and other selected pores from the sample.

5.2. Modeling of a cluster of pores

We revisit problem 5, focusing on decomposing a given 3-ball B containing a collection of pores. In topology, several options exist for such decompositions, ranging from polytopal complexes (e.g. clipped Voronoi diagrams or clipped Laguerre–Voronoi diagrams) [105] to specific or general CW complexes, and even handle-body decompositions [46]. In our specific context, the natural choice, consistent with decompositions for dry foams and Voronoi diagrams, is to consider the components as 3-balls and the decomposition of B as a CW complex decomposition.

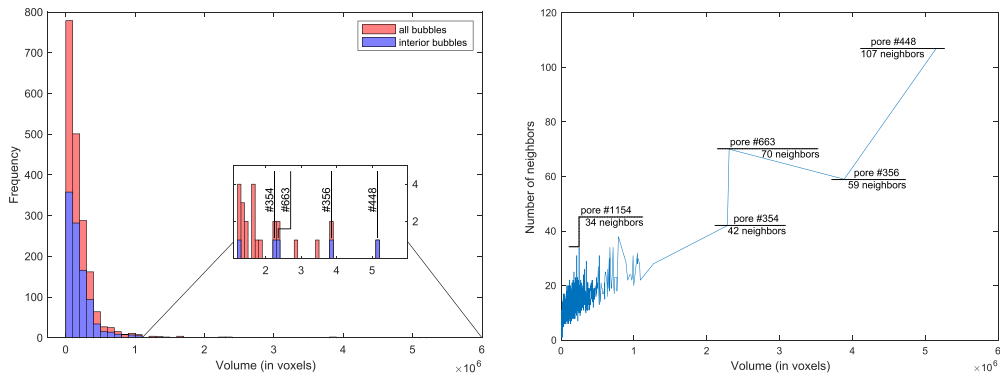


Figure 10. Left: Volume distribution of the 1911 pores in the study sample; Right: Correlation between the number of neighbors (determined by our reconstruction) and the pore volume.

Table 1. Volumes and number of neighbors of the largest and other selected pores in the sample.

Pore index	Number of voxels	Number of reconstructed neighbors
# 448	5151 369	107
# 356	3886 091	59
# 663	2306 342	70
# 354	2282 999	42
# 1154	246 658	34
# 856	143 387	7

In a general CW complex decomposition of B into 3-cells B_1, \dots, B_m , boundary identifications among the balls B_i , where $i = 1, \dots, m$, are permissible. This flexibility allows for modeling structures like a toroidal pore, as illustrated in figure 9, where a ball is glued to itself to create a 3-torus. However, when adhering to the conditions laid out in definition 9—specifically, linked pores and the ability to thicken pores to disjoint 3-balls—we can confine ourselves to the subclass of regular CW complexes. In contrast, subclasses such as strongly regular CW complexes or polytopal complexes prove too restrictive for modeling diverse real-world data.

Example 12. In figure 11, we showcase two pores (# 448 and # 663) from our study sample. In our reconstruction, the representing 3-cells exhibit a distinctive feature: they intersect in three shared polygonal 2-faces, highlighted in yellow. This contrasts with the usual requirement of only one intersection in polytopal or strongly regular CW complexes. In essence, these two cells are *multi-connected*, a characteristic evident in the dual complex through the presence of a multi-edge.

The combinatorial aspect of Plateau’s rules constrains the modeling to a *simple* regular CW complex, where the maximal faces of the dual cell complex consist of tetrahedra that are glued together along their boundaries.

Definition 13. A *generalized triangulation* of a 3-manifold M (possibly with boundary) is a regular CW complex decomposition of M into tetrahedra.




Figure 11. Our reconstruction (right) reveals a multi connection between the pores on the left, featuring pores #448 (blue) and #663 (red).




Figure 12. Left: Arrangement of five pores #591 (black, *a*), #550 (red, *b*), #347 (yellow, *c*), #615 (green, *d*), and #336 (blue, *e*) within the cluster. Right: Our reconstruction of pore #856 (black) exhibits 7 facets, while the corresponding Laguerre region features 19 facets (see figure 13). The neighboring pores are: #448 (orange), #875 (blue), #893 (red), #911 (magenta), #924 (cyan), #929 (purple), and #1137 (green).

In a generalized triangulation dual to a simple decomposition, the tetrahedra that are glued together have no identifications on their individual boundaries. Nevertheless, the set of vertices permits the presence of multiple edges, triangles, and tetrahedra.

Example 14. The dual generalized triangulation corresponding to the collection of five pores *a*, *b*, *c*, *d*, and *e* in figure 12(left) comprises four tetrahedra: two instances of *abcd* and two of *abce*. Additionally, there are two copies of the triangles *bcd* and *bce*. If the black pore *a* were removed from the pore cluster, the two copies of the tetrahedron *abcd* would merge into the single triangle *bcd*, and the two copies of *abce* would merge into the single triangle *bce*.

Example 15. Through manual inspection, we verified that our reconstruction algorithm successfully identified all 7 neighbors of pore #856, depicted in black in figure 12(right). The neighbor count for the associated Laguerre region is 19. This discrepancy is primarily attributed to the presence of the substantial neighboring pore #448 (orange): the equivalent sphere of this pore does not obstruct the equivalent sphere of pore #856 from detecting other spheres

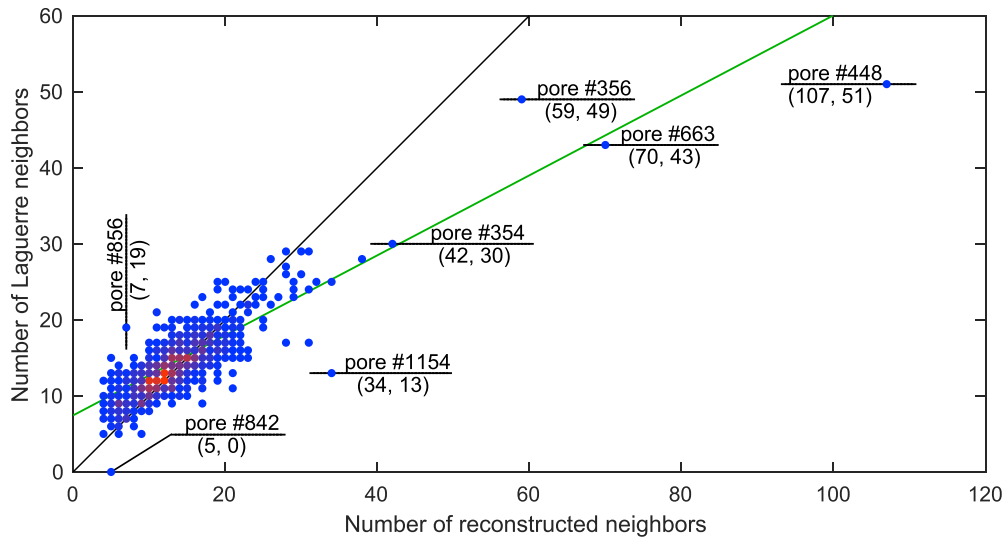


Figure 13. Scatter plot depicting the correlation between the pore's neighbor count in our reconstruction and the corresponding neighbor count in the Laguerre diagram. Color intensity reflects data point frequency, with deeper red indicating higher occurrence. The black line represents the reference $y = x$, while the green line shows the fitted trendline obtained via a first-degree polynomial regression.

beyond pore #448. Consequently, the Laguerre region corresponding to pore #856 becomes adjacent to regions representing some of those spheres. This underscores a notable drawback of using equivalent spheres to represent pores.

The deviations between our topological reconstruction and the power diagram arise from fundamental differences in the modeling assumptions. Power diagrams impose convexity by construction, assigning each cell to a region defined by weighted distance from a generator point. This geometric rigidity cannot accommodate topological complexities such as non-convex or multiply connected pores. Our combinatorial model, by contrast, is sensitive to such features, and therefore captures structural relationships that a power diagram necessarily omits. Figure 14 provides a quantitative comparison of the number of facets across different geometric representations.

The 1-skeleton of the dual generalized triangulation forms a graph, with a vertex corresponding to each pore. Edges, potentially multi-edges, connect pores that are (multi-)neighbors. The main focus is the 'reconstruction' of the dual generalized triangulation, achieved through the neighborhood graph and additional local geometric information. The primal simple regular CW complex decomposition is subsequently directly derived from this triangulation.

As per theorem 6, any simple regular CW complex decomposition with m cells is topologically suitable for a collection of m pores without additional information. However, in practical scenarios, pores situated far from each other are not considered neighbors. Thus, we 'register' pores as neighbors if their distance is less than a specified threshold c_1 (refer to section 8.2 for the heuristic distance computation between two pores). In figure 15(left), we visualize the distance segments between pores—pairs of voxels corresponding to distinct pores, connected by a segment if their distance is less than c_1 . The resulting neighborhood graph, displayed




Figure 14. Normalized histogram comparing the number of facets in Voronoi cells, Laguerre cells, and our reconstructed cells. Voronoi cells exhibit the highest average face count (14.79), followed by Laguerre cells (13.97), and reconstructed cells (13.32). Distributions are overlaid for visual comparison; share is normalized over interior pores only.

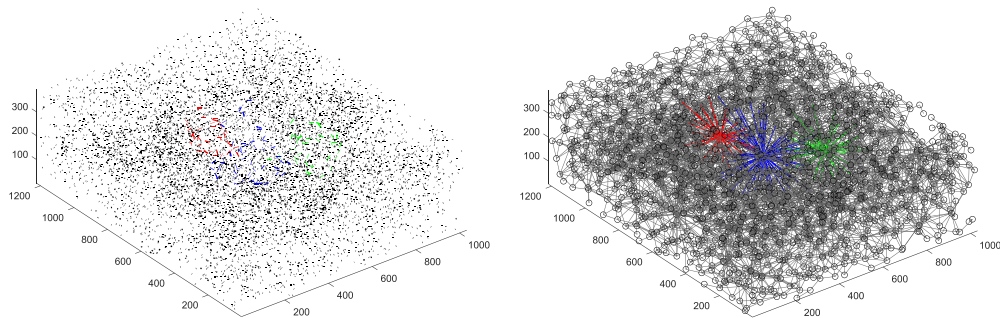


Figure 15. Initial adjacencies of pores # 663, # 448, # 356 (left) are represented by red, blue, and green distance segments, while the corresponding edges of the dual graph are shown on the right.

in figure 15(right), features barycenters of pores as vertices and edges connecting neighbors determined by the distance threshold.

The selection of the threshold c_1 significantly impacts the edges represented in the neighborhood graph. A small threshold may lead to *missing edges*, overlooking connections between pores. Conversely, a large threshold may incorrectly link pores as neighbors, even when obstructed by other pores; see figure 9(right). These inaccuracies result in what can be referred to as *excessive edges*.

Example 16. In figure 16(left), the Delaunay triangulation is shown for four points representing discs of equal radius. In the middle and on the right, the disk centers form a square, and




Figure 16. Delaunay triangulation (left), missing edge (middle), and excessive edge (right).




Figure 17. Example of foam pores with dual graphs in the shape of an octahedron, a pentagonal bipyramid, and a hexagonal bipyramid.

one disk is replaced by an ellipse of equal area. The ellipse has a vertical orientation in the middle image and a horizontal orientation on the right.

In our methodology, pores are considered neighbors when their distance is less than the threshold c_1 . In the middle image, if c_1 is at most the length of the diagonal of the square, the vertical diagonal becomes a missing edge, which would be added first to complete the triangulation as c_1 increases. On the right, if c_1 exceeds the length of the diagonal, both diagonals are included, but the vertical one is excessive. This is because the ellipse and the right disk are closer to each other than the top and bottom discs, so they get connected earlier for smaller values of c_1 .

Example 17. In figure 9(right), an excessive edge is registered between the light-blue and the yellow pore, despite the presence of the red pore between them. Missing edges occur, for example, in configurations like those in figure 17, which involve pores in octahedral positions or in the positions of bipyramids over a pentagon or a hexagon, respectively (see [65]).

5.3. Dual graphs of cellular manifolds and reconstructibility

Let K^d be a triangulation of a closed d -dimensional manifold M^d in the form of an abstract simplicial complex. A combinatorial approach to defining a ‘dual complex’ for K^d involves the *dual block complex* $(K^d)^*$ [83] (see also [42, 51]). For each vertex $v \in K^d$, the *closed dual block* $\bar{D}(v)$ is the union of all simplices containing v in the barycentric subdivision $\text{sd}(K^d)$ of K^d and their faces. The main difficulty is that $\bar{D}(v)$ might not be a topological cell, i.e. it may not be homeomorphic to the d -dimensional ball B^d .

Consider the example where Σ^{d-2} is a *combinatorial homology* $(d-2)$ -sphere. The double suspension theorem of Edwards [30] and Cannon [24] implies that the double suspension $\text{susp}(\text{susp}(\Sigma^{d-2}))$ of Σ^{d-2} is homeomorphic to S^d . If $d \geq 5$ and Σ^{d-2} is not homeomorphic to S^{d-2} , $\text{susp}(\text{susp}(\Sigma^{d-2}))$ is a non-PL triangulation of S^d [15]. Consequently, the collection of (open) dual blocks for the non-trivial original faces of a triangulation K^d of a d -manifold M^d may not form a CW complex decomposition of M^d .

If K^d is a triangulation of a combinatorial d -manifold—meaning a triangulation in which the link of each vertex is a combinatorial sphere (see [22])—the blocks $\bar{D}(v)$ of vertices v form proper topological cells. However, for general manifolds with $d \geq 4$, the $(d-1)$ -dimensional vertex links may fail to be polytopal spheres, i.e. they might not be realizable as boundary complexes of simplicial d -polytopes. This presents an obstruction to defining a dual polytopal complex for K^d .

Definition 18. The *dual graph* $G(\mathcal{C})$ of a simple regular CW decomposition \mathcal{C} of a compact d -manifold M^d is the graph on the maximal cells of \mathcal{C} as vertices, with an edge connecting two cells $C, D \in \mathcal{C}$ whenever they intersect in a common $(d-1)$ -dimensional cell.

Dancis [26] established that a compact homology d -manifold M^d , triangulated as a simplicial complex K^d , is uniquely determined by its $(\lfloor \frac{d}{2} \rfloor + 1)$ -skeleton. For even d , if the $d/2$ -homology of K^d vanishes or, in the case of orientability, is a finite group, then K^d can be reconstructed from its $d/2$ -skeleton. Given our focus on reconstructing triangulations of manifolds from their graphs, Dancis' result predominantly confines reconstructibility to triangulations of the 2-sphere S^2 .

In fact, triangulations of all other surfaces, whether orientable or non-orientable, may contain a complete subgraph K_r , $r \geq 6$, in their 1-skeleta or even be neighborly (see [52, 90]), hindering reconstruction (locally or globally).

Example 19. For $d \geq 3$, the d -dimensional boundaries $\partial C_{d+1}(n)$ of the $(d+1)$ -dimensional cyclic polytopes $C_{d+1}(n)$ on n -vertices have a complete 1-skeleton, thus yield neighborly triangulations of the d -dimensional sphere S^d . Reconstructions of these triangulations from their graphs are not possible. Using the spheres $\partial C_{d+1}(n)$ as summands in connected sums gives non-reconstructible triangulations for every PL manifold M^d . For $d=3$, Walkup [100] proved that every closed 3-manifold admits neighborly triangulations.

In contrast to the aforementioned negative results, several significant subclasses of triangulations do allow for reconstruction from their 1-skeleta:

- Flag triangulations of manifolds are reconstructible.
- Barycentric subdivisions of arbitrary triangulations are reconstructible since barycentric subdivisions are flag.
- Elementary subdivisions, where every facet is subdivided, are reconstructible, although these triangulations are not flag.

The reconstruction of a flag simplicial complex K from its 1-skeleton involves representing it as the clique complex of that 1-skeleton, i.e. $K = \text{Clique}(\text{Skel}_1(K))$. For a general simplicial complex K , it holds that $K \subseteq \text{Clique}(\text{Skel}_1(K))$. In the extreme case where $\text{Skel}_1(K)$ is the complete graph K_n on n vertices, the resulting clique complex is the $(n-1)$ -dimensional simplex, $\text{Clique}(\text{Skel}_1(K)) = \Delta_{n-1}$.

As conjectured by Perles and proven by Blind and Mani [16], any simple d -polytope can be reconstructed from its 1-skeleton; refer to Kalai [53] for a simple proof. However, our inputs

are not the graphs of simple polytopes but the dual graphs of simple regular CW decompositions of compact 2D or 3D manifolds, serving as the 1-skeleta of the dual triangulations.

6. 2D reconstruction

The crux of Dancis' reconstruction of a triangulation K^d of a d -manifold M^d from its k -skeleton is determining whether a subcomplex N^k of K^d , isomorphic to the boundary of a $(k+1)$ -simplex, bounds a $(k+1)$ -simplex in K^d , for $k \geq \lceil \frac{d}{2} \rceil$; see [26, main lemma 10]. In the context of a triangulation K^2 of the 2-sphere S^2 , this entails deciding for every 3-clique t in the graph $\text{Skel}_1(K^2)$ whether it constitutes a triangle of K^2 or not.

6.1. Reconstruction of triangulations of S^2

The following lemma is immediate and also follows from the work of Dancis [26].

Lemma 20. *Let K^2 be a triangulation of S^2 , and let t be a triangle of $\text{Clique}(\text{Skel}_1(K^2))$ with a vertex set $V(t)$. Then, the graph $\text{Skel}_1(K^2) - V(t)$ is connected if and only if t is a triangle of K^2 ; otherwise, t is an empty triangle.*

Proof. Let t be a triangle in $\text{Clique}(\text{Skel}_1(K^2))$. The boundary ∂t of t triangulates a 1-sphere in K^2 and, according to Jordan's curve theorem, divides K^2 into two discs. If t is not a triangle of K^2 , then the two discs contain interior vertices, each separated by the 1-sphere ∂t , resulting in $\text{Skel}_1(K^2) - V(t)$ being disconnected. To establish that if t is a triangle of K^2 , then $\text{Skel}_1(K^2) - V(t)$ is connected, we leverage the fact that removing a triangle from K^2 produces a triangulated disk with boundary ∂t . According to Steinitz' theorem [96], the graph of this disk is 3-connected. \square

By applying lemma 20 to a triangulation K^2 of S^2 , we can readily determine which triangles in the clique complex $\text{Clique}(\text{Skel}_1(K^2))$ belong to K^2 and which do not.

6.2. Dual graph based certification for 2D triangulations

Later in this paper, we address situations where an input graph G does not precisely represent the 1-skeleton of a 3-sphere triangulation K^3 . These variations manifest as either *missing edges* or *excessive edges*. During an 'incremental reconstruction' of K^3 from G , we identify and iteratively address missing edges or excessive ones.

To begin, we replace the utility of lemma 20 (used in the reconstruction of K^2 of the 2-sphere S^2 from its 1-skeleton) with local checks. These checks either

- *certify* a 3-clique $t \in \text{Clique}(\text{Skel}_1(K^2))$ as a triangle of K^2 ,
- *decertify* t by marking it as an empty triangle, or
- label t as *undecided*.

In a triangulation K^2 of a surface M^2 , every edge e lies in exactly two triangles of K^2 and, consequently, in at least two 3-cliques of $\text{Skel}_1(K^2)$. If e lies in exactly two 3-cliques of $\text{Skel}_1(K^2)$, these two 3-cliques must be triangles of K^2 , thus certified by the edge e . The link of a vertex v in K^2 is a triangulated 1-sphere with $|V(\text{link}_{K^2}(v))| = \text{deg}(v) \geq 3$. In the case of $\text{deg}(v) = 3$, the three 3-cliques of $\text{Skel}_1(K^2)$ containing v must be triangles of K^2 . Consequently, every vertex of degree 3 certifies the three triangles it is contained in.

Algorithm 1. 2D Certification Algorithm.**Input:** The 1-skeleton $\text{Skel}_1(K^2)$ of a triangulation K^2 of a surface M^2 **Output:** Lists C , D , and U of certified, decertified, and undecided 3-cliques of $\text{Skel}_1(K^2)$

0. **Initialization.** We initialize $C = D = \emptyset$ and $U = \{t \in \text{Clique}(\text{Skel}_1(K^2)) : |t| = 3\}$. For every edge e of $\text{Skel}_1(K^2)$, let $C(e)$, $D(e)$, and $U(e)$ be the respective subsets of C , D , and U of cliques that contain the edge e . An edge e of $\text{Skel}_1(K^2)$ is *unprocessed* if it lies in no certified 3-clique, is *touched* if it lies in exactly one certified 3-clique, and *processed* if it lies in two certified 3-cliques of $\text{Skel}_1(K^2)$.
1. **Germes for certification.**
 - (i) Every vertex v of degree 3 certifies the three 3-cliques it is contained in as triangles of K^2 .
 - (ii) Every edge e that is contained in exactly two 3-cliques certifies these as triangles of K^2 .
2. **Certification and decertification via edges e .**
 - (i) If $|C(e)| = 2$ and $|U(e)| \geq 1$, then the 3-cliques in $U(e)$ can not be triangles of K^2 and thus can be marked as decertified.
 - (ii) If $|C(e)| = |U(e)| = 1$, then the only triangle t in $U(e)$ has to be a triangle of K^2 and thus can be removed from $U(e)$ and added to $C(e)$. As a consequence, the other two edges of t become touched if they were unprocessed and processed if they were touched before, respectively.
 - (iii) If $|C(e)| = 0$ and $|U(e)| = 2$, then the two triangles in $U(e)$ have to be triangles of K^2 and thus can be removed from $U(e)$ and added to $C(e)$. As a consequence, the other edges of the two triangles become touched if they were unprocessed and processed if they were touched before, respectively.
 - (iv) Repeat Steps (i)–(iii) until no further progress is made.
3. **Decertification via splitting.** Pick a 3-clique $t \in U$ (in some order, only once) and check if t splits $\text{Skel}_1(K^2)$ into two components; if so, then t is an empty triangle of K^2 and can be marked as decertified, i.e. t is removed from U and added to D , then Steps 2 (i)–(iv) are repeated (after updating $U(e)$ and $D(e)$ for each of the three edges e of t). Process all remaining 3-cliques of U until no further progress is made.
4. **Output.** Return the lists C , D , and U .

Remark 21. The 2D Certification Algorithm accurately certifies triangles of K^2 in the initial Steps 1(i)+(ii). Steps 2(i)–(iii) leverage the condition that each edge in a triangulated surface belongs to precisely two triangles. Step 3 aids in identifying separating 3-cliques.

For a triangulation K^2 of a general surface S , a separating 3-clique in K^2 implies an empty triangle. However, when $S = S^2$, a 3-clique is separating if and only if it is an empty triangle, as per lemma 20. Thus, In the case of $S = S^2$, Step 3 alone suffices to distinguish between triangles and empty triangles.

Nevertheless, we intentionally execute Steps 1–2 first. This strategy is particularly beneficial for general surfaces $S \neq S^2$, where these initial steps can yield positive certificates for specific surface patches.

Theorem 22. For a triangulation K^2 of the 2-sphere S^2 , the 2D Certification Algorithm outputs a list of certified triangles corresponding to the triangles of K^2 . The decertified 3-cliques represent the empty triangles, and the list of undecided 3-cliques is empty.

Proof. If the triangulation K^2 is flag (i.e. has no empty triangles), then all triangles are certified in Step 1 (ii). Otherwise, according to lemma 20, each empty triangle in K^2 allows a split of K^2 into two components, denoted as P^2 and Q^2 , such that $K^2 = P^2 \# Q^2$, representing the connected sum of the two components. This repeated splitting along empty triangles results in components with no empty triangles. The total resulting pieces can be organized in a binary

tree structure, with each piece as a node; two nodes at the same level are connected summands of their lowest common ancestor.

The leaves Z_i , with $i \in \{1, \dots, s\}$, fall into two categories: either $Z_i = \partial\Delta_3$ or Z_i is flag. For an empty triangle t in K^2 , we have $K^2 = Z_j \# Q^2$, where $j \in \{1, \dots, s\}$ and Z_j has no empty triangles. In the subsequent steps, all triangles of Z_j , except for t , are certified either in Step 1 (i) if $Z_i = \partial\Delta_3$ or in Step 1 (ii) if Z_j is flag. The empty triangle t gets decertified later, either in Step 2 (i) or later in Step 3. Consequently, faces of the neighboring piece Z_k to Z_j are certified through the repetition of Step 1 (ii). This approach ensures that triangles in the leaf pieces are certified in the initial round. Following that, the inner pieces are sequentially certified until all triangles in K^2 are certified, and all empty triangles are decertified, leaving no undecided triangles. \square

Theorem 23. *For a flag triangulation K^2 of a surface S , the 2D Certification Algorithm returns the list of certified triangles as the triangles of K^2 , with empty lists for both decertified and undecided triangles.*

Proof. In a flag triangulation K^2 of a surface S , there are no vertices of degree 3—apart from the tetrahedron triangulating the 2-sphere. The 2D Certification Algorithm terminates with Step 1 (ii) since every edge is contained in exactly two 3-cliques of K^2 . \square

In general, a triangulation of a surface need not be reconstructible from its graph.

Example 24. Let K^2 be the unique neighborly vertex-minimal 6-vertex triangulation of the real projective plane \mathbb{RP}^2 (up to combinatorial isomorphism, see [64]). In this case, $\text{Skel}_1(K^2)$ is the complete graph K_6 , and $\text{Clique}(\text{Skel}_1(K^2))$ is the 5-simplex Δ_5 on 6 vertices. The 2D Certification Algorithm marks all triangles of Δ_5 as undecided.

Example 25. Let K^2 be a flag triangulation of a surface M^2 , and let L^2 be a neighborly triangulation of a surface $N^2 \neq S^2$. For the connected sum $K^2 \# L^2$, the 2D Certification Algorithm certifies all triangles of K^2 and returns the triangles of L^2 as undecided (except for the triangle at which the connected sum is taken, which is marked as decertified). Hence, although the entire triangulation may not be fully reconstructible, certain patches can still be reconstructed successfully.

A *locally planar triangulation* (in the most general sense) of a surface $M^2 \neq S^2$ is a triangulation of M^2 for which its *edge-width*, the minimal length of a non-contractible cycle, is at least four [2, 3]. In topological graph theory, locally planar triangulations are often considered with larger edge-width [27].) Such triangulations generalize flag triangulations in the sense that there can be empty triangles, but every such splits the surface, a property detectable by the 2D Certification Algorithm.

Theorem 26. *A locally planar triangulation K^2 of a surface $M^2 \neq S^2$ is reconstructible from its 1-skeleton.*

Example 27. Let K^2 be a neighborly triangulation of a surface M^2 other than the 2-sphere. In this case, K^2 is not locally planar. If we consider the triangulation L^2 obtained by subdividing each triangle of K^2 , then L^2 is also not locally planar. Nonetheless, L is reconstructible from its 1-skeleton.

6.3. 2D certification for decompositions with boundary

In practical scenarios, a cluster of pores typically occupies a confined region in Euclidean space. To prevent degeneracy in 2D, we make the assumption that the pores forming the cluster




Figure 18. A collection of 2D pores and their adjacency list.

cover a two-dimensional disk. Including the outer region as an extra pore guarantees a simple regular CW decomposition of the 2-sphere S^2 .

In figure 18, a cluster comprising nine pores, with the outer region represented as the 10th pore, is shown alongside the dual graph of the pores stored as an adjacency list. In this instance, each 3-clique of the dual graph corresponds to a triangle in the dual triangulation. This allows us to reconstruct the cell structure of the pore cluster from the dual adjacency graph. The resulting dual triangulation, in conjunction with the primal cellular decomposition, is reconstructed as an abstract combinatorial object.

7. 3D reconstruction

In the context of voxelized tomography data for metallic foam samples, the dual graphs constructed for static thresholds often contain missing and excessive edges, categorized as ‘errors’; refer to section 8 for details. Before addressing these complications in an adaptive manner, we initially explore the (partial) reconstruction of triangulations, treated as abstract simplicial complexes, for closed 3-manifolds in section 7.1. Subsequently, in section 7.2, we incorporate local geometric information to reconstruct geometric triangulations of triangulated 3-balls B^3 within Euclidean 3-space \mathbb{R}^3 .

7.1. Dual graph based certification for 3D triangulations

Let K^3 be a triangulation of a closed 3-manifold M^3 . For every edge $e \in K^3$, the link of e , denoted as $\text{link}_{K^3}(e)$, forms a cycle. We have the inclusions:

$$\text{link}_{K^3}(e) \leq \text{Skel}_1(\text{link}_{\text{Clique}(\text{Skel}_1(K^3))}(e)) \leq \text{link}_{\text{Clique}(\text{Skel}_1(K^3))}(e).$$

The 1-complex $L(e) := \text{Skel}_1(\text{link}_{\text{Clique}(\text{Skel}_1(K^3))}(e))$ consists of the cycle $\text{link}_{K^3}(e)$, possibly along with extra chords or vertices and edges, and may not be connected in general.

If K^3 is a neighborly triangulation with the vertex set $V(K^3)$, then the clique complex of $\text{Skel}_1(K^3)$ is a $(|V(K^3)| - 1)$ -simplex, i.e. $\text{Clique}(\text{Skel}_1(K^3)) = \Delta_{|V(K^3)|-1}$. In this case, a combinatorial reconstruction of K^3 from its graph is not possible.

Lemma 28. *If K^3 is a flag triangulation of a closed 3-manifold, then*

- $\text{Clique}(\text{Skel}_1(K^3)) = K^3$,

- $L(e) = \text{link}_{K^3}(e)$,
- there are no k -cliques in $\text{Skel}_1(K^3)$ of size $k \geq 5$,
- the list of facets of K^3 is given by all 4-cliques in $\text{Skel}_1(K^3)$, and
- every 3-clique in $\text{Skel}_1(K^3)$ is a triangle of K^3 that is contained in exactly two tetrahedra of K^3 .

Theorem 29. Let K^3 be a triangulation of a closed 3-manifold M^3 , and let q be a 5-clique in the 1-skeleton of K^3 . Then either $K^3 \cong \partial\Delta_4$ or at most four out of the five 4-subcliques of q are tetrahedra of K^3 . In case exactly four 4-subcliques of q are tetrahedra of K^3 , one of the vertices of q has degree 4 in $\text{Skel}_1(K^3)$.

Proof. Let $N^3 \leq K^3$ be the subcomplex formed by the four-element subsets of q that appear as tetrahedra in the triangulation K^3 . If all five four-element subsets of q are contained in N^3 , then N^3 is simplicially isomorphic to the boundary of a 4-simplex $\partial\Delta_4$. Consequently, $\partial\Delta_4$ is a closed, full-dimensional submanifold of M^3 , implying $M^3 \cong \partial\Delta_4$ and hence $K^3 \cong N^3 \cong \partial\Delta_4$.

If $N^3 \neq \partial\Delta_4$ and exactly four 4-subcliques of q appear in N^3 , then there is exactly one vertex, denoted v , that appears in each of those four 4-subcliques. The link of v in K^3 consists of the 3-element subsets formed by removing v from the four 4-subcliques, which form the boundary of a 3-simplex Δ_3 . Since the link of v is a 2-sphere, it follows by the same reasoning above that the link in K^3 consists solely of these four 3-element subsets, implying that v has degree 4 in $\text{Skel}_1(K^3)$. \square

Corollary 30. Let $K^3 \neq \partial\Delta_4$ be a triangulation of a closed 3-manifold M^3 . Additionally, let q be a 5-clique in the 1-skeleton of K^3 where each vertex of q has a degree of at least 5 in the 1-skeleton of K^3 . Then, at most three out of the five 4-subcliques of q are tetrahedra of K^3 .

For a triangulation K^3 of a closed 3-manifold M^3 and a vertex $v \in K^3$, the complex $\text{link}_{K^3}(v)$ is a triangulated 2-sphere, and we have the inclusions

$$\text{link}_{K^3}(v) \leq \text{Skel}_2(\text{link}_{\text{Clique}(\text{Skel}_1(K^3))}(v)) \leq \text{link}_{\text{Clique}(\text{Skel}_1(K^3))}(v).$$

Definition 31. A triangulated r -sphere contained in an s -dimensional simplicial complex L^s (with $s \geq r$) is *Hamiltonian* if it contains all the vertices of L^s .

Trivially, $\text{link}_{K^3}(v)$ is a Hamiltonian 2-sphere in $\text{link}_{\text{Clique}(\text{Skel}_1(K^3))}(v)$, and enumerating all triangulated Hamiltonian 2-spheres in $L(v) := \text{Skel}_2(\text{link}_{\text{Clique}(\text{Skel}_1(K^3))}(v))$ is possible and yields all candidate spheres for $\text{link}_{K^3}(v)$ [60, 70].

The 3D Certification Algorithm correctly certifies and decertifies 4-cliques in $\text{Skel}_1(K^3)$, identifying tetrahedra or missing 3-faces. Similar to the two-dimensional case, it may conclude with some undecided 4-cliques. As previously mentioned, the algorithm's performance may be suboptimal for general triangulations, but we focus on triangulations with an average vertex degree ranging from 11 to 16, ensuring computational feasibility. For triangulations dual to space decompositions induced by foam structures, there are patches in which the algorithm is effective, allowing for a partial reconstruction of K^3 from its 1-skeleton $\text{Skel}_1(K^3)$. Figure 19 illustrates a high-level flowchart of algorithm 2.

Step 1 (i) of the algorithm, checking if the link of an edge has a unique cycle, is a straightforward test. For Step 1 (ii), we iteratively collapse edges that are not included in any cycle until no further collapses are possible, followed by the removal of all vertices with degree 0. The order of Steps 2–5 is interchangeable. In our implementation, whenever C , D , and U are updated, we aim to leverage this information locally for incident vertices, edges, and 4-cliques.

For Steps 2 and 3, we employ a brute-force approach, given that the links in the triangulation we are interested in (both one- and two-dimensional) typically contain few cycles. Step 4

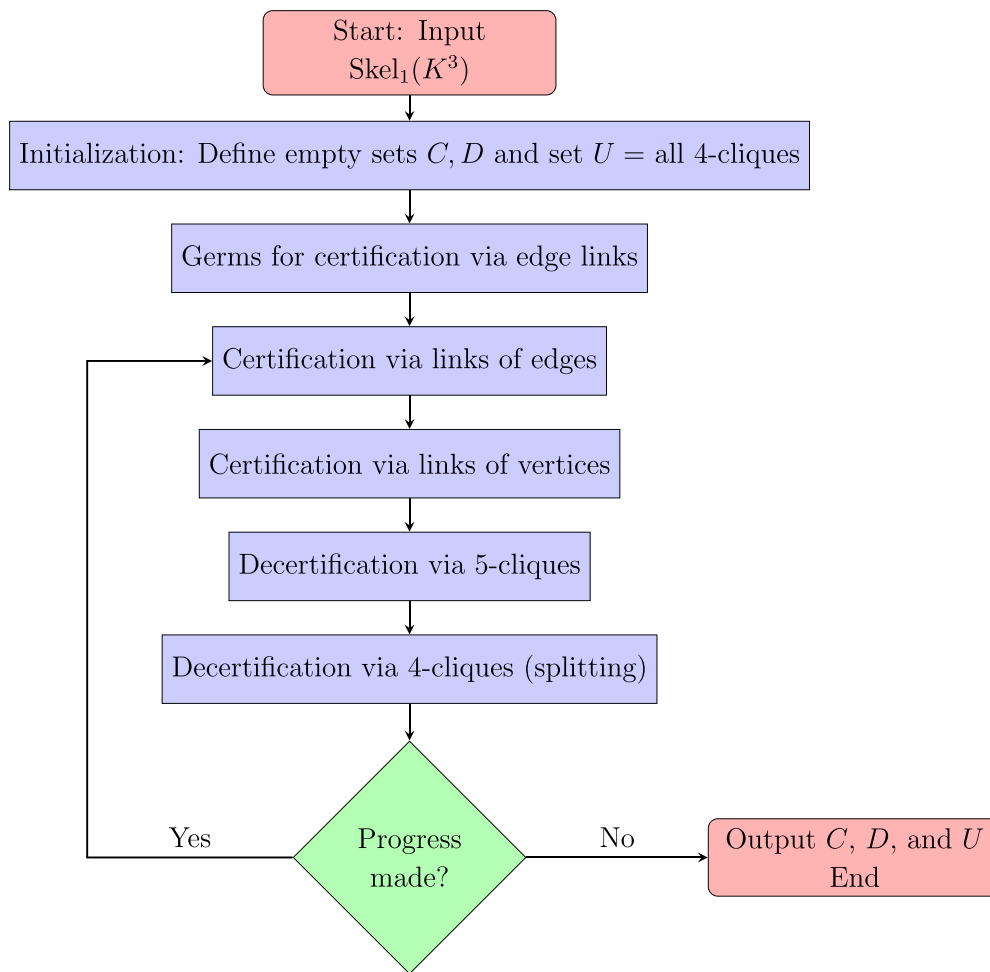


Figure 19. Flowchart of algorithm 2.

is executed when the information about 4-cliques in a 5-clique is updated. Step 5 is often implicitly covered by previous steps.

If a vertex v in K^3 has a degree of 4, this vertex certifies the four tetrahedra it is part of. This certification is implicitly verified by Step 1 (i) as a trivial scenario where the edges incident to v have a degree of 3.

In the situation where the link of a vertex v in K^3 contains a 3-clique t that splits the link, the 4-clique $v * t$ is not a face of K^3 . Such 4-cliques are appropriately decertified in Step 3.

Example 32 (Cross configuration). Consider a flag triangulation K^3 of a closed 3-manifold M^3 . Let abc and abd be two adjacent triangles that both appear in the link of a vertex v and also in the link of a vertex $w \neq v$. If we add the edge cd to $\text{Skel}_1(K^3)$, the resulting graph represents, in general (without considering combinatorial isomorphisms in specific examples), the 1-skeleton of two distinct triangulations of M^3 . Both triangulations include the additional tetrahedron $abcd$. In one case, the vertex v sees the triangles abc and abd , while in the other case, v sees the triangles acd and bcd . Conversely, for the vertex w , the scenarios are reversed.

Algorithm 2. 3D Certification Algorithm, Part I (Combinatorial Conditions).

Input: The 1-skeleton $\text{Skel}_1(K^3)$ of a triangulation K^3 of a closed 3-manifold M^3

Output: Lists C , D , and U of certified, decertified, and undecided 4-cliques of $\text{Skel}_1(K^3)$

- Initialization.** Initialize $C = D = \emptyset$ and $U = \{q \in \text{Clique}(\text{Skel}_1(K^3)) \mid |q| = 4\}$. For each edge e of $\text{Skel}_1(K^3)$, define $C(e)$, $D(e)$, and $U(e)$ as the respective subsets of C , D , and U of cliques that contain the edge e . Assign values 1 or 0 to edges $f \in L(e)$ based on whether $e * f$ lies in C or D , respectively. No value is assigned to an edge $f \in L(e)$ with $e * f \in U$. Define $L_1(e) := \{f \in L(e) \mid f \text{ has value } 1\}$ and $L_0(e) := \{f \in L(e) \mid f \text{ has value } 0\}$.

Similarly, for a vertex v and a triangle $t \in L(v)$, assign values 1 or 0 to t based on whether $v * t$ lies in C or D . Define the sets $L_1(v)$ and $L_0(v)$ as above.

- Germs for certification and decertification via links of edges.**

- For every edge e of K^3 , if $L(e)$ contains a unique cycle Z , then move $e * f$ to C for every $f \in Z$, and move $e * g$ to D for every $g \in L(e) - Z$.
- For every edge e of K^3 , let T_e be the union of all cycles in $L(e)$. Assign 0 to every $g \in L(e) - T_e$ and move $e * g$ to D .

This step is illustrated in figure 20.

- Certification and decertification via links of edges e .**

- Let E be the set of edges obtained as the intersection over all cycles in $L(e) - L_0(e)$ that contain all the edges of $L_1(e)$. Assign 1 to every $g \in E$ and move $e * g$ to C .
- Let E be the union of the edges of all cycles in $L(e) - L_0(e)$ that contain all the edges of $L_1(e)$. Assign 0 to every $g \in L(e) - L_0(e) - E$ and move $e * g$ to D .

- Certification and decertification via links of vertices v .** Similar to Step 2, for links of vertices we have $L_1(v) \subseteq \text{link}_{K^3}(v)$. Inspect the union and the intersection of Hamiltonian 2-spheres in $L(v)$ to certify and decertify 4-cliques of $\text{Skel}_1(K^3)$.

- Decertification via 5-cliques.** If a 5-clique has no vertex of degree 4 in $\text{Skel}_1(K^3)$ and three of its five 4-cliques have been certified, then (by theorem 29) decertify the remaining two 4-cliques.

- Decertification via 4-cliques (splitting).** Pick a 4-clique $q \in U$ and check if q splits $\text{Skel}_1(K^3)$ into two components. If so, move q to D .

- Iteration and output.** Repeat Steps 2–5 until no further progress is made. Then return the lists C , D , and U of certified, decertified, and undecided tetrahedra, respectively.




Figure 20. Top row: Different configurations of edge links $L(e)$ along with assigned values to edges. Bottom row: Values assigned in red are deduced after applying Step 2 of algorithm 2.

Example 32 illustrates that local modifications to the graph of a triangulation K^3 of a closed 3-manifold M^3 can pose combinatorial challenges, impeding the reconstructibility of the triangulation from its 1-skeleton $\text{Skel}_1(K^3)$.

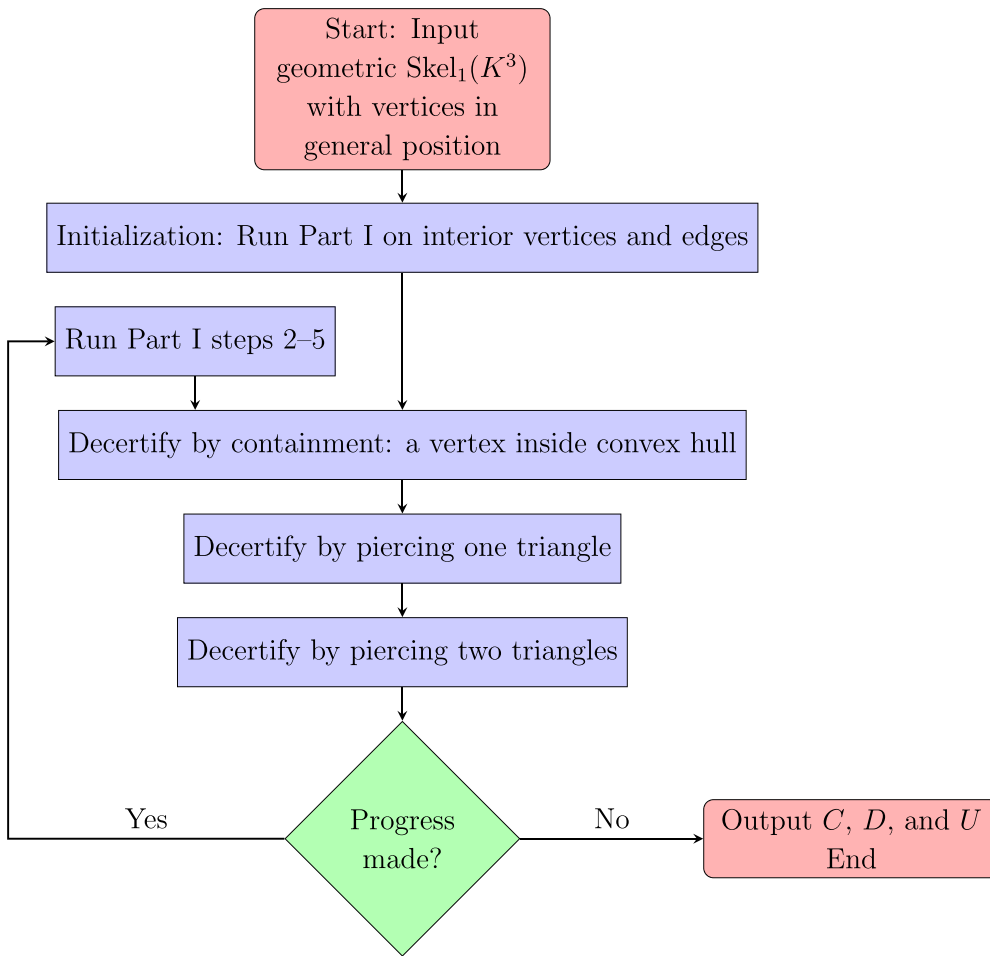


Figure 21. Flowchart of algorithm 3.

7.2. Dual graph based certification for geometric 3D triangulations

In this section, we integrate geometric information into the 3D Certification Algorithm.

Figure 21 illustrates a high-level flowchart of algorithm 3. In general, this algorithm may not succeed in fully reconstructing the geometric triangulation of a 3-ball from its 1-skeleton. Consider, for instance, a simplicial 3-polytope P with vertex set V . By selecting a point in the interior of P and forming the join of that vertex with the boundary (triangles) of P , a triangulation K^3 of P is obtained. Removing any tetrahedron from K^3 results in the same 1-skeleton as K^3 itself.

Theorem 33. *Let X be a finite set of points in \mathbb{R}^3 in general position with $|X| \geq 4$, and K^3 be a geometric triangulation of the point set X , i.e. K^3 is a triangulation of $\text{conv}(X)$. Then K^3 is reconstructible from $\text{Skel}_1(K^3)$.*

Proof. A 4-clique q of K^3 is decertified if $\text{conv}(q)$ contains another point of X or if $\text{conv}(q)$ is pierced by an edge of $\text{Skel}_1(K^3)$.

Algorithm 3. 3D Certification Algorithm, Part II (Geometric Conditions).

Input: The 1-skeleton of a geometric triangulation K^3 of a three-dimensional ball B^3 in \mathbb{R}^3 , with the vertices of K^3 in general position

Output: Lists C , D , and U of certified, decertified, and undecided 4-cliques of $\text{Skel}_1(K^3)$

0. **Initialization.** Run algorithm 2 on $\text{Clique}(\text{Skel}_1(K^3))$, with the individual steps restricted to interior vertices and interior edges. This produces lists C , D , and U of certified, decertified, and undecided 4-cliques, respectively.
1. **Geometric decertifications.** Let a, b, c, d be distinct vertices forming a 4-clique q in $\text{Clique}(\text{Skel}_1(K^3))$.
 - (i) **Decertification by containment.** If there is a vertex v of K^3 in $\text{conv}(a, b, c, d)$, then v decertifies q , and q is moved to D .
 - (ii) **Decertification by piercing one triangle.** Let v be a vertex outside $\text{conv}(a, b, c, d)$ and av an edge of $\text{Skel}_1(K^3)$. If av and $\text{conv}(b, c, d)$ non-trivially intersect, av decertifies $\text{conv}(b, c, d)$ and q is moved to D .
 - (iii) **Decertification by piercing two triangles.** Let v, w be vertices outside $\text{conv}(a, b, c, d)$ and vw an edge of $\text{Skel}_1(K^3)$. If vw non-trivially intersects $\text{conv}(a, b, c)$ and $\text{conv}(a, b, d)$ of $\text{conv}(a, b, c, d)$, vw decertifies these triangles and q is moved to D .
2. **Iteration and output.** Repeat Steps 0 (without the initialization Step 0 of algorithm 2) and 1 until no further progress is made. Then return the lists of certified, decertified, and undecided tetrahedra.

Since $\text{conv}(q) \subseteq \text{conv}(X)$, $\text{conv}(q)$ must be covered by the tetrahedra of the triangulation K^3 of $\text{conv}(X)$. If $\text{conv}(q)$ neither contains other vertices of X nor is pierced by an edge of $\text{Skel}_1(K^3)$, then q might still lie within another tetrahedron, leading to the decertification of that tetrahedron. Alternatively, q could be covered by at least two other tetrahedra. In the latter case, one of the triangles of the covering tetrahedra is pierced by q , leading to the decertification of those tetrahedra.

The only remaining possibility is that q is a tetrahedron of K^3 , and thus is covered by itself. Hence, after executing algorithm 3, K^3 is reconstructed by considering all undecided tetrahedra to be certified. \square

Corollary 34. Set S be a set of points in \mathbb{R}^3 in general position. Then the Delaunay triangulation $\text{DT}(S)$ of S is reconstructible from its 1-skeleton $\text{Skel}_1(\text{DT}(S))$.

Example 35. As a test case for algorithms 2 (combinatorial) and 3 (geometric), we generated a Poisson–Voronoi diagram for 1000 points with dual Delaunay triangulation K^3 and a face vector $f = (1000, 7427, 12785, 6357)$.

In this sample, 927 vertices are interior, and 73 lie on the boundary of the convex hull of the 1000 points. The graph $\text{Skel}_1(K^3)$ has 7427 edges and contains 9104 4-cliques, of which 6357 are tetrahedra of K^3 . Additionally, $\text{Skel}_1(K^3)$ has 1406 5-cliques, 25 6-cliques, and no 7-cliques.

We restrict the combinatorial steps of algorithm 2 to interior edges and vertices. As a result, 4078 4-cliques of $\text{Skel}_1(K^3)$ are certified as tetrahedra of K^3 (i.e. 64.15% of the tetrahedra), and 538 4-cliques get decertified. Approximately 25.45% of the links of interior edges in $\text{Clique}(\text{Skel}_1(K^3))$ have a unique cycle and are entirely reconstructed in Step 1 (i) of algorithm 2. Eventually, 8.41% of the links of interior vertices are reconstructed.

Algorithm 3 decertifies the remaining 4-cliques that are not tetrahedra of K^3 , leading to a full reconstruction of K^3 .

8. Adaptive determination of the graph G of a dual generalized 3D triangulation

The input in the following is a collection of pores as specified by definition 11. Each pore is a strongly connected union of gas voxels.

8.1. Preprocessing

In our showcase example of 1911 pores, each pore contains at least 4300 voxels. Smaller pores were removed in a preprocessing step for two reasons. Firstly, very tiny conglomerates of strongly connected voxels with up to 10 voxels were removed as noise. Secondly, satellite pores, as discussed in the introduction, were eliminated. In the showcase example, the voxel size is $41 \mu\text{m}$, and an equivalent ball with a volume of 4300 voxels has a radius of approximately 0.41 mm .

8.2. Computation of an upper bound to the distance of two pores

Given two non-convex pores as distinct strongly connected collections of voxels \mathcal{C} and \mathcal{D} , their *voxelized distance*, denoted as $d(\mathcal{C}, \mathcal{D})$, is the minimum distance between any voxel in \mathcal{C} and any voxel in \mathcal{D} ,

$$d(\mathcal{C}, \mathcal{D}) = \min \{ d(x, y) \mid x \in \mathcal{C}, y \in \mathcal{D} \},$$

where the distance $d(x, y)$ between two voxels x and y is the distance between their centers.

However, computing this exact distance is computationally expensive. Instead, we use an upper bound, $d_{\text{up}}(\mathcal{C}, \mathcal{D})$, determined as follows. We employ a method inspired by the Newton–Raphson method, starting from a voxel $x_0 \in \mathcal{C}$ and traversing a *zigzag path* $x_0 y_1 x_1 y_2 x_2 \dots y_k x_k$ of voxels (see figure 22). Here,

- $y_i \in \{z \in \mathcal{D} \mid d(x_{i-1}, z) \leq d(x_{i-1}, y) \forall y \in \mathcal{D}\}$ is a voxel in \mathcal{D} with minimal distance to voxel $x_{i-1} \in \mathcal{C}$, and
- $x_i \in \{z \in \mathcal{C} \mid d(z, y_i) \leq d(x, y_i) \forall x \in \mathcal{C}\}$.

In case of multiple choices, we randomly select the voxels x_i and y_i .

In our computations, we set $k = 3$ and randomly choose four starting voxels $x_0 \in \mathcal{C}$, including (in an abuse of notation) the barycenter of \mathcal{C} . Additionally, we exchange the roles of \mathcal{C} and \mathcal{D} . The upper bound $d_{\text{up}}(\mathcal{C}, \mathcal{D})$ is the minimum distance over the distances $d(x_3, y_3)$ for the considered zigzag paths between the pores \mathcal{C} and \mathcal{D} . Thus, $d_{\text{up}}(\mathcal{C}, \mathcal{D})$ serves as an upper bound to $d(\mathcal{C}, \mathcal{D})$. While computing $d(\mathcal{C}, \mathcal{D})$ would take $O(|\mathcal{C}| \cdot |\mathcal{D}|)$ time, we compute $d_{\text{up}}(\mathcal{C}, \mathcal{D})$ in $O(|\mathcal{C}| + |\mathcal{D}|)$ time.

8.3. Determination of boundary pores

The metallic foam sample discussed in this work is cuboid-shaped, measuring $1033 \times 1221 \times 398$ voxels with a nearly square base and small height (see figure 1). Notably, the sample includes a large pore with voxels on both the top and bottom sides, prompting the question of distinguishing between pores on the boundary and in the interior.

To represent pores, we use their barycenters. The boundary of the convex hull of the finite set of barycenters $X \in \mathbb{R}^3$ can indicate which pores are on the boundary (with barycenters as boundary vertices). However, this approach has drawbacks, such as distant pores becoming neighbors on the boundary due to long edges of the convex hull, while pores with barycenters just inside the convex hull are considered interior pores.

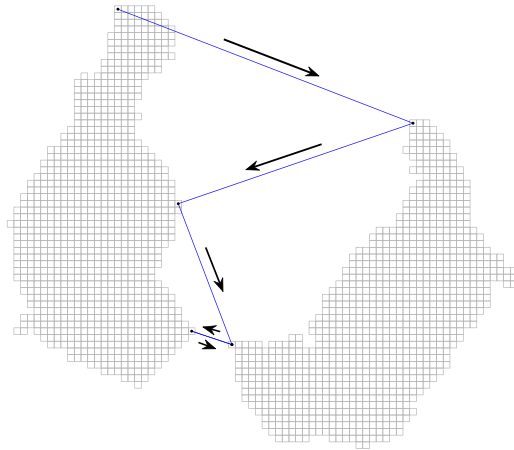


Figure 22. A zigzag path to obtain an upper bound $d_{\text{up}}(\mathcal{C}, \mathcal{D})$ of the voxelized distance between two pores.




Figure 23. The alpha complex for the barycenters of the pores of the foam sample. Axis ticks are in voxel units, where each voxel measures $41 \mu\text{m}$ per side.

For a more refined boundary description, we employ alpha complexes, introduced by Edelsbrunner *et al* in 1983 [29]. Assuming points in X are in general position, for each $x \in X$, we consider an r -ball $B_r(x)$. The *alpha complex* $\text{Alpha}(r, X)$ is the nerve complex of the union of sets $U_r(X) = \{B_r(x) \cap V_x \mid x \in X\}$, where V_x is the Voronoi region of $x \in X$. In extreme cases, $\lim_{r \rightarrow 0} \text{Alpha}(r, X) = X$, while $\lim_{r \rightarrow \infty} \text{Alpha}(r, X)$ yields the Delaunay triangulation of X with the convex hull of X as the boundary.

For our purposes, we stop incrementing r when $\text{Alpha}(r, X)$ becomes a triangulated 3-ball (figure 23), which occurs when $r = 88$ voxel side lengths for this foam sample, where a voxel side length equals $41 \mu\text{m}$. For similar samples, values for r ranged between 73 and 135, partly depending on CT scan resolution. In our sample, 917 out of a total of 1911 pores are determined as boundary pores, while 994 are interior pores.

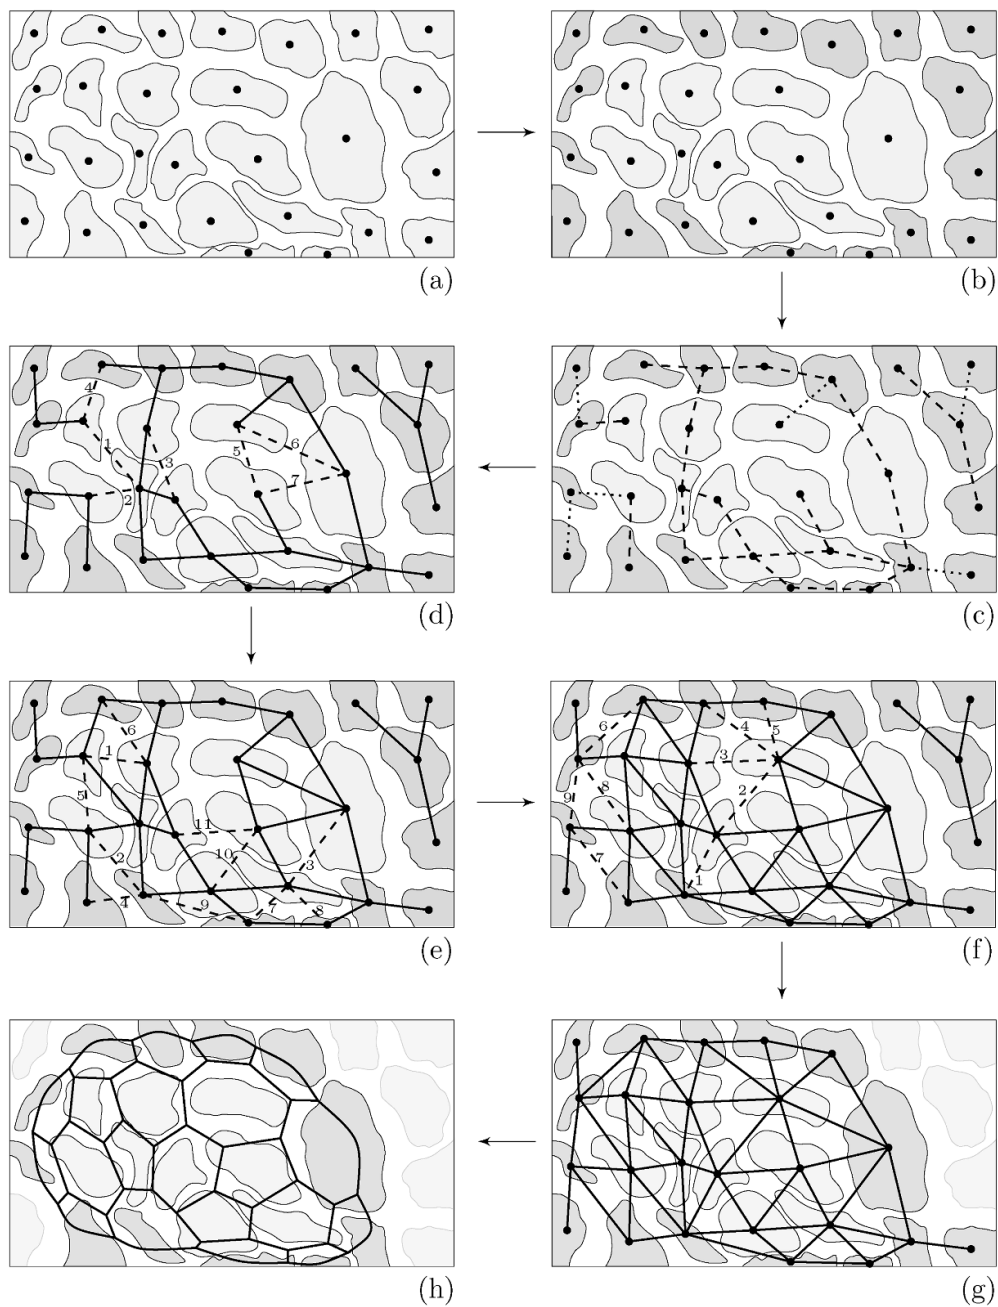


Figure 24. A sample of 2D pores and the main steps to reach a cellular decomposition.

8.4. 2D sketch of the main steps

Figure 24 provides a 2D overview of the primary steps executed in 3D to derive a dual triangulation and a corresponding primal cell decomposition. In figure 24(a), a set of pores is depicted along with their barycenters.

Step (a)→(b): Transitioning from figures 24(a) and (b), the alpha shape computation identifies boundary pores (shown in darker gray).

Step (b)→(c): Utilizing d_{up} , we establish connections (dashed edges) between pairs of pores \mathcal{C} and \mathcal{D} when $d_{\text{up}}(\mathcal{C}, \mathcal{D}) < c_1$. Isolated interior pores are connected to their nearest neighbor with dotted edges.

Step (c)→(d): In the triangulation of a 2D disk (or a 3D ball), every interior vertex should have a degree of at least 3 (or at least 4, respectively). Transitioning from figures 24(c) and (d), vertices with a degree less than 3 have additional edges greedily added, starting with the shortest each time. The order in which edges are added is indicated by the labels 1 to 7.

Step (d)→(e): For each vertex, its current vertex link is determined. If it is not a closed cycle, the next shortest edge that reduces the number of connected components of a link is added globally. In this step, edges of length at most c_2 are added. This addition might lead to some few excessive edges, the removal of which requires an additional step in 3D. This step is not shown here as it lacks a 2D counterpart.

Step (e)→(f): Further edges of length at most c_3 are added in this step to close remaining holes in the links, along with addressing any potentially occurring additional excessive edges.

Step (f)→(g): Figure 24(g) displays the resulting set of edges. Through certification and decertification, a dual triangulation is determined. Components consisting solely of boundary pores are removed as outliers.

Step (g)→(h): Spikes of boundary pores are collapsed, resulting in a triangulated disk (or ball in 3D). The primal cell decomposition is presented in figure 24(h).

Our choice of a greedy approach in Steps (c)→(d), (d)→(e), and Step (e)→(f) for constructing an abstract graph requires justification. In 2D, a closely related geometric problem is the *minimum-weight triangulation (MWT)* problem, proven to be NP-hard by Mulzer and Rote [82]. The MWT problem involves finding a triangulation of a finite set of points $S \subseteq \mathbb{R}^2$ in general position with the smallest total edge length. Each candidate triangulation of S triangulates the convex hull of S and has the same fixed number of edges.

In 3D, our scenario involves a finite set of pores, and we aim to represent them as vertices in an abstract graph, essentially the 1-skeleton of a three-dimensional dual generalized triangulation. Unlike the fixed number of edges in the 2D MWT problem, in 3D, candidate triangulations may vary in the number of edges, and the graphs can include multiple edges.

Mulzer and Rote [82, figure 2] show that different optimal triangulations can exist for the 2D MWT problem, such as those obtained by retriangulating a 2D strip at its two ends and propagating the triangulation through the strip. However, for the 2D MWT problem, the β -skeleton [58]—a graph consisting of particular edges—is part of every optimal triangulation for a suitable choice of β [56, 82].

In Step (b)→(c), we introduce the concept of a ‘backbone network’, composed of edges that should belong to every ‘generic’ triangulation dual to a ‘homogeneous’ metallic foam sample. In our showcase sample, the edges of length less than c_1 constitute 57% of the final triangulation, representing a significant portion. Therefore, the greedy approach to fill in the holes in Steps (d)→(e) and (e)→(f) has a high likelihood of deviating from ‘an optimal triangulation’ only by local fluctuations.

8.5. Adaptive computation of nearest neighbors

If the homogeneous foam sample is significantly larger than the average or maximum pore diameter, pores that are spatially distant from each other *do not* qualify as neighbors in the foam. Therefore, potential neighboring pores for a given pore include all pores located within a specified distance from that pore. This scenario is addressed in the literature, treating it as an

extension of the nearest neighbor searching problem; for a comprehensive survey, refer to [5]. A common strategy involves dividing the space into a regular grid of cubes and checking points in nearby cubes to determine whether they fall within a specified range from an initial point.

In our foam sample, spanning 1221 voxels in length, the pores exhibit an average diameter of 63.5 voxels, with a minimum diameter of 20.2 voxels and a maximum of 214.3 voxels. Given the substantial size and variability, the grid approach was not feasible, and we opted for a straightforward iteration over all pairs of pores. However, at the voxel level, we only calculate $d_{\text{up}}(\mathcal{C}, \mathcal{D})$ when the distance between the barycenters of two pores is constrained by the sum of their equivalent ball radii plus ten times the threshold distance c_1 , i.e.

$$d(\text{bary}(\mathcal{C}), \text{bary}(\mathcal{D})) \leq r(\mathcal{C}) + r(\mathcal{D}) + 10c_1.$$

We iteratively adapt the threshold parameter c_1 , starting from $c_1 = 2$ voxels. (Note that the minimum distance between two voxels from distinct strongly connected components is at least $\sqrt{2} \times$ the voxel side length.) For each candidate pair of pores \mathcal{C} and \mathcal{D} , we evaluate whether $d_{\text{up}}(\mathcal{C}, \mathcal{D}) \leq c_1$. If this condition is met, we register \mathcal{C} and \mathcal{D} as neighbors.

After constructing the adjacency graph for a given c_1 , we compute the average degree deg_{av} over all interior vertices. If $\text{deg}_{\text{av}} < 6$, we increment c_1 (by one voxel or a resolution-dependent constant) and repeat the process. The procedure halts as soon as $\text{deg}_{\text{av}} \geq 6$ for the first time.

This stopping criterion is not meant to match a theoretical average precisely but is instead chosen to ensure that the initial graph is sufficiently dense. In our experience, graphs reaching this density threshold contain enough correct adjacencies to allow subsequent combinatorial steps (e.g. edge certification via local topology) to identify and recover missing edges or prune excessive ones.

To fulfill Step (b)→(c) of section 8.4, we additionally connect each isolated (interior or boundary) pore \mathcal{C} to its nearest neighbor \mathcal{D} . The resulting graph G represents each pore \mathcal{C} as vertex $c \in V(G)$. Before connecting the isolated pores, the average vertex degree was 7.702, increasing to 7.715 afterward.

8.6. Degree requirement for edges

Further incrementing c_1 until achieving an average vertex degree between 11 and 16 would be disadvantageous, as it introduces unwanted excessive edges that locally impede the graph from representing the 1-skeleton of a triangulation. Instead, we strategically continue adding edges where needed.

Consider an edge cd in G where at least one of the corresponding pores \mathcal{C} or \mathcal{D} is an interior pore. In the hypothetical case where G forms the 1-skeleton of a triangulated 3-ball as a simplicial complex K , the link of every interior edge e in K is a cycle of length at least 3. If the link of cd in $\text{Clique}(G)$ has fewer than three vertices, it indicates a potential oversight of mutual neighbors to cd , unless c and d share exactly two neighbors e and f , where ef forms a double edge in the link of cd within a simple regular CW decomposition. Although rare, this scenario may occur, such as when a small pore is enclosed by two large pores (see figure 12). In the general case, we examine neighbors of c and d that are not mutual neighbors. Let b be a vertex representing a pore \mathcal{B} , adjacent to only one of c or d in G . For all such vertices, we pick the one, say b_{min} , that minimizes the distance to the other vertex c or d . If this distance is smaller than $2c_1$, we register b_{min} as a mutual neighbor to both c and d by adding the edge $b_{\text{min}}c$ or $b_{\text{min}}d$ to G , respectively.

As outlined in Step (c)→(d) of section 8.4 for the 2D situation, we iterate over all edges in G with a degree less than 3, selecting the overall shortest edge that, upon addition, increases




Figure 25. Two examples of vertex links $L_G^c(v)$ of a vertex v in the clique complex $\text{Clique}(G)$.

the degree of a low-degree edge. This process is repeated until no candidate edges with lengths at most $2c_1$ remain.

It’s worth noting that during this step, the criterion of adding an edge only if the distance between two pores is at most $2c_1$ may result in retaining some edges with a degree less than 3. The degree of such edges is either increased later by subsequent steps, or their link is recognized as a digon.

8.7 Detection of holes in vertex links

Let K^3 be a triangulation of a closed 3-manifold, and v be a vertex of K^3 . Then, $\text{link}_{K^3}(v)$ is a triangulated 2-sphere, and $\text{link}_{K^3}(v) \leq \text{link}_{\text{Clique}(\text{Skel}_1(K^3))}(v)$. The complexes $\text{link}_{K^3}(v)$ and $L_{K^3}^c(v) := \text{link}_{\text{Clique}(\text{Skel}_1(K^3))}(v)$ have the same number of vertices, but $L_{K^3}^c(v)$ can have additional edges, triangles, or higher-dimensional faces, for instance, when $\text{Skel}_1(K^3)$ forms a complete graph.

Consider an edge $e = ab$ in $L_{K^3}^c(v)$ that does not belong to $\text{link}_{K^3}(v)$. In this case, vertices a and b are connected by e in K^3 , but e lies within the interior of $K^3 - \text{star}_{K^3}(v)$. If abc is an empty triangle in $\text{link}_{K^3}(v)$, then abc is part of $\text{Clique}(\text{Skel}_1(K^3))$. Additionally, if abc and abd are triangles in $\text{link}_{K^3}(v)$ and cd is an edge in $\text{Clique}(\text{Skel}_1(K^3))$, then acd and bcd form triangles in $L_{K^3}^c(v)$, but they are not triangles in $\text{link}_{K^3}(v)$.

Consider the graph G obtained from section 8.6, with v being a vertex of G , and $L_G^c(v) := \text{link}_{\text{Clique}(G)}(v)$. Each vertex v of G is geometrically represented by the barycenter of the corresponding pore. However, $L_G^c(v)$ might contain ‘holes’ due to missing edges, preventing it from forming a triangulated 2-sphere; refer to figure 25 for visual examples of such links. In the following, we aim to heuristically close these holes, taking care to avoid false fill-ins.

Consider an edge $e = ab$ in $L_G^c(v)$. The edge e is *loose* if it is not part of any triangle in $L_G^c(v)$, and *free* if it is included in exactly one triangle of $L_G^c(v)$. Loose and free edges are natural candidates for bounding holes. However, as illustrated in figure 25(left), the link $L_G^c(v)$ of a vertex v might exhibit holes bounded by cycles containing edges that belong to two (or more) triangles of $L_G^c(v)$, as exemplified by the edges ab , ag and bc in figure 25(left).

Definition 36. An edge e in $L_G^c(v)$ is *exposed* if it is either loose, free, or if it is part of a unique inclusion-maximal clique A in $L_G^c(v)$ such that $\text{conv}(v, e)$ forms a boundary triangle of

Algorithm 4. Hole Detection in Vertex Links.

Input: Graph G on barycenters of pores as its vertex set V , with edges corresponding to neighboring pores (as identified so far)

Output: Updated graph G on V with added edges to close holes in vertex links

1. **Vertex loop.** Loop over vertices v of G corresponding to interior pores and consider their vertex links $L_G^c(v) = \text{link}_{\text{Clique}(G)}(v)$.
2. **Cycle computation.** For each link, compute all admissible cycles on the set of exposed edges.
3. **Candidates gathering.** For each admissible cycle, consider all cycle chords and chords connecting collapsed vertices to cycle vertices.
4. **Edge addition.** Select the overall shortest candidate with a length of at most $6c_1$, add it to G , and update the affected vertex links accordingly.
5. **Iteration.** Repeat until no further candidate edge of length $\leq 6c_1$ is left.

the convex hull $\text{conv}(v, A)$. In other words, the edge e lies on the shadow of $\text{conv}(A)$ as seen from v .

Let e be an exposed edge that lies in a unique 4-clique A in $L_G^c(v)$. Geometrically, there are two possible cases for this. The shadow of A , as seen from v , is either a triangle or a square. We refer to the former as a *pyramid configuration* and the latter as a *cross configuration*. It is important to note that the distinction between pyramid and cross configurations may not always be accurate, especially for elongated, non-convex pores where the positions of the barycenters can vary with further stretching. However, these two configurations and their shadows have proven to be very helpful in our implementation.

In figure 25(left), the 4-clique $abfg$ forms a pyramid with shadow ab , ag , and bg , whereas the shadow of the 4-clique $bchi$ consists of the four edges bc , bh , ci , and hi . Only the shadow edges ab , ag , and bc of the two 4-cliques $abfg$ and $bchi$ are exposed, as the other ones are included in further triangles of $L_G^c(v)$.

If an edge e belongs to two inclusion-maximal 3-cliques in $L_G^c(v)$, it is not exposed and is not considered as a candidate edge for a hole. However, in the link $L_G^c(v)$ of v , the two triangles containing e could potentially be folded up, creating a hole along e that we do not detect at this stage. Our restriction to exposed edges has proven effective, and any missed holes were later identified and closed when G was extended through the addition of other edges.

Let H be the subgraph consisting of all exposed edges of $L_G^c(v)$. We successively collapse away free vertices (of degree 1) of H to obtain a graph H' . In H' , we enumerate all cycles of length at least four. If a cycle of H' has a chord, as in figure 25(right), we discard it. If the removal of a cycle of H' disconnects H , we discard it as well. If a cycle bounds a triangulated disk in $L_G^c(v)$, then we do not consider it. At this point, this condition prevents the closure of a vertex link in the case that this link consists of only one triangulated disk; this will be addressed later in section 8.12. The remaining cycles of H' are called *admissible*.

Now, we gather *candidate missing edges* to close an admissible hole. These include all chords connecting non-adjacent vertices of the hole. Additionally, if there are any trees of free vertices that were collapsed away, we consider all edges connecting a vertex in those trees to a non-adjacent vertex on the admissible cycle.

Next, we iterate over all the gathered candidate missing edges for every vertex v in G and each admissible cycle in $L_G^c(v)$. We select the shortest edge, provided its length is at most $6c_1$, and add it. Subsequently, we update the list of admissible holes and candidate missing edges accordingly, and add the shortest among them if its length remains within the limit of $6c_1$.

The aforementioned steps are outlined in algorithm 4.




Figure 26. (Left) The purple pore and the red pore are separated by the blue, yellow, and green pores. (Middle and Right) The purple pore and the red pore are neighbors.

8.8. Adding edges to close edge links

Consider an edge $e = uv$ in a triangulation K^3 of a closed 3-manifold. The link of e in K^3 is a triangulated 1-sphere, represented as a cycle, and this cycle is contained in the link of e in $\text{Clique}(\text{Skel}_1(K^3))$.

Now, assume that e is an edge of G with at least one interior vertex. We explore the case where $L_G^c(e) := \text{link}_{\text{Clique}(G)}(e)$ is acyclic, aiming to gradually add edges to G to ensure that $\text{link}_{\text{Clique}(G)}(e)$ eventually contains a cycle. Let $N(u)$ and $N(v)$ be the sets of neighbors of u and v in G , respectively. Additionally, define $N(u, v) := N(u) \cap N(v)$ as the set of vertices that are mutual neighbors to both u and v —in other words, $N(u, v)$ constitutes the vertices of the link $L_G^c(e)$ of e in $\text{Clique}(\text{Skel}_1(K^3))$.

Next, we gather candidate missing edges:

- ur for $r \in N(v)$, but $r \notin N(u, v)$, with $d(u, r) \leq 2c_1$,
- vs for $s \in N(u)$, but $s \notin N(u, v)$, with $d(v, s) \leq 2c_1$,
- rs for $r, s \in N(u, v)$, but rs is not an edge of $L_G^c(e)$, with $d(r, s) \leq 2c_1$.

We select the shortest edge from the above options, iterating over all edges e with an acyclic link, and add it to G . This process is repeated until no further edge of length at most $2c_1$ can be added.

Similar to section 8.6, the constraint of adding only edges with a length of at most $2c_1$ maintains the possibility of digonal edge links.

8.9. Adding degree-3 edges

In a triangulation K^3 of a closed 3-manifold, an edge $e = uv$ of degree 3 is a candidate edge for a bistellar flip that removes this edge; see [15]. In this step, we consider the opposite direction where we have two 4-cliques, $uabc$ and $vabc$, in G , as registered in the previous steps, for which uv is not an edge of G . In other words, the edges of the two 4-cliques form a bipyramid with apices u and v over a triangle abc . Figure 26 illustrates three such configurations, selected from our showcase sample, with the five pores: a in blue, b in yellow, c in green, u in red, and v in pink. In the left case, the purple pore and the red pore are separated by the blue, yellow, and green pores, whereas in the middle and right cases, the purple pore and the red pore are neighbors. In the latter two cases, the three edges ab , ac , and bc form a triangular hole, and uv is considered a missing edge. We add the edge uv to G whenever the distance between u and v is at most $3c_1$.

This step impacts the average degree of a vertex and, consequently, the average face degree of the primal simple regular CW decomposition. Thus, depending on the threshold $3c_1$, the average degree should be considered as an interval rather than a fixed number.

In the case where a 3-clique is contained in more than two 4-cliques, we check the distance for every pair of additional vertices. Similar to previous steps, no exact decision can be made regarding whether to add a respective edge uv , as non-convex pores can have spikes, etc which might lead to a wrong decision and perhaps excessive edges—a situation we need to address later in section 8.11.

8.10. Detection of further holes in vertex links

In the previous two steps, additional edges were introduced, potentially enabling the detection of further holes in the vertex links within $\text{Clique}(G)$. However, instead of merely re-executing the procedure outlined in section 8.7, we slightly relax the condition stipulated in that step, which mandates that a cycle should consist solely of exposed edges. If such a cycle is not found in a vertex link, we explore maximal paths in that vertex link composed exclusively of exposed edges. If the two endpoints of such a path are connected by an edge in the link, we anticipate this non-exposed edge to close the maximal path into a potential hole. Once again, we iterate over all vertex links and all potential chords of the identified cycles, selecting the shortest chord with a length of at most $5c_1$ to be added to G .

As highlighted earlier in section 8.7, if an edge e is part of two triangles within a vertex link, these two triangles could be folded, as seen from the vertex. Consequently, even though not strictly combinatorially exposed, e could be ‘geometrically exposed’ and aid in closing a hole.

8.11. Removal of excessive edges

Commencing with the graph G as outlined in section 8.5, the introduction of edges to G in sections 8.6 through 8.10 may introduce conflicts. In section 8.5, we initially incorporate *all* edges with lengths not exceeding the threshold c_1 , as well as edges linking isolated pores to G , without consideration of the relative positions of the pores—whether geometrically or combinatorially in the clique complex. This indiscriminate incorporation of edges in section 8.5 stems from the initial lack of sufficient information on the vertex links within the clique complex of G to judiciously determine the reasonability of including specific edges. The primary objective of this section, as well as section 9.4, is to pinpoint potential surplus edges and eliminate them.

In a flag triangulation K^3 of a closed 3-manifold, each 4-clique of $\text{Skel}_1(K^3)$ constitutes a tetrahedron of K^3 , and there are no cliques of larger size. For Delaunay triangulations dual to Poisson–Voronoi diagrams, as observed experimentally in section 1, k -cliques with $k \geq 7$ are infrequent. Consequently, in the graph G constructed thus far, it is prudent to examine all 6-cliques to identify potential conflicts caused by a single edge within such a clique, which might warrant its removal as an excessive edge. This examination will be conducted in section 8.15. However, it is plausible that we have overlooked some edges in G at this stage due to unresolved holes. As a precautionary measure, and before scrutinizing 6-cliques, we perform a more general check. In this check, we iterate over all vertices w of G and consider pairs (t, e) comprising triangles $t = abc$ and edges $e = uv$ in $L_G^c(w) = \text{link}_{\text{Clique}(G)}(w)$ that intersect geometrically. This intersection occurs when the triangle $\text{bary}(\mathcal{A})\text{bary}(\mathcal{B})\text{bary}(\mathcal{C})$ intersects the edge $\text{bary}(\mathcal{U})\text{bary}(\mathcal{V})$. The objective is to establish conditions ensuring that the three pores \mathcal{A} , \mathcal{B} , and \mathcal{C} act as separators for the pores \mathcal{U} and \mathcal{V} .

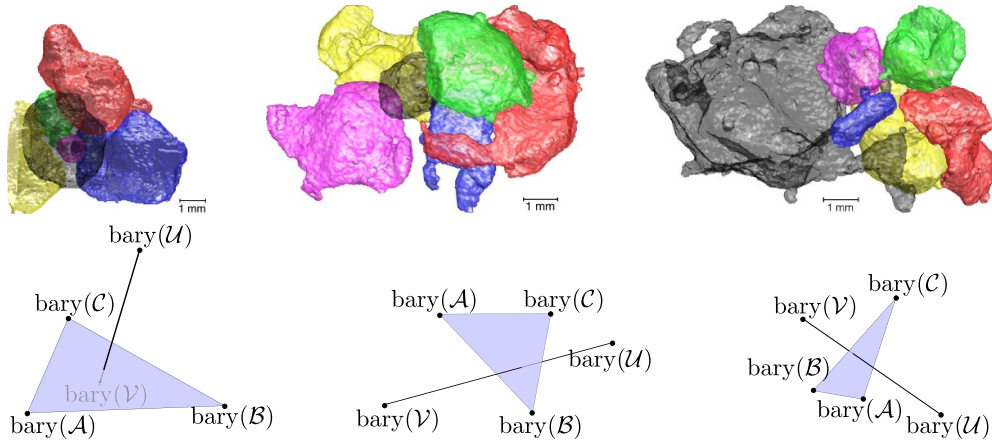


Figure 27. Three pairs of (edge, triangle) from Step 8.11.

Figure 27 illustrates three configurations, each consisting of six pores, from the showcase sample. In these configurations, the vertex- w pore is rendered in transparent gray, while the a -, b -, and c -pores are in yellow, blue, and green, and the u - and v -pores are in red and purple, respectively.

In figure 27(left), the a -, b -, and c -pores fail to act as separators for the u - and v -pores, even though the edge $\text{bary}(U)\text{bary}(V)$ intersects the triangle $\text{bary}(A)\text{bary}(B)\text{bary}(C)$, albeit marginally. To account for this, we declare an edge $e = uv$ in $L_G^c(w)$ to be excessive if there exists a triangle abc in $L_G^c(w)$ such that:

- The edge $\text{bary}(U)\text{bary}(V)$ intersects the triangle $\text{bary}(A)\text{bary}(B)\text{bary}(C)$, say at a point P .
- The barycentric coordinates of P in the segment $\text{bary}(U)\text{bary}(V)$, denoted as α and β , satisfy $0.3 \leq \alpha, \beta \leq 0.7$.
- The distance between $\text{bary}(U)$ and $\text{bary}(V)$ is greater than each of the distances $\text{bary}(A)\text{bary}(B)$, $\text{bary}(A)\text{bary}(C)$, and $\text{bary}(B)\text{bary}(C)$, i.e. uv is not short at a local scale.
- The distance between $\text{bary}(U)$ and $\text{bary}(V)$ is at least $2c_1$, i.e. uv is not short with respect to the threshold c_1 .

In figure 27(left), the edge uv is not removed because the second condition is not satisfied. However, in both the middle and right configurations, the edge uv is removed. In these cases, a pore with a tentacle connects two distant pores, prompting us to adjust G by disconnecting the red and purple pores, respectively.

To prevent the excessive removal of edges in this step and mitigate any bias introduced by the choice of the initial vertex w for which triangles and edges of its link are inspected, we globally collect candidate excessive edges uv and order them by their length $d(\text{bary}(U), \text{bary}(V))$. The longest such edge is then removed first, and this process is iterated until no more candidate excessive edges remain. The loop only considers pairs (t, e) where both t and e are faces of some link $L_G^c(w)$, focusing the search for excessive edges on a local configuration of pores.

8.12. Detection of further holes in vertex links: round 2

We revisit the procedure outlined in section 8.10, this time employing a larger threshold of $10c_1$ to address any remaining holes, and allowing the boundary of a triangulated disk to qualify as

a candidate hole. It is crucial to note that edges removed as excessive in any previous step are prohibited from being reintroduced.

8.13. Hamiltonicity of vertex-links and detection of digonal faces

Let K^3 be a triangulation of a closed 3-manifold, and consider a vertex v in K^3 . As mentioned earlier, the complex $\text{link}_{K^3}(v)$ is a triangulated 2-sphere, and it is a subcomplex of $\text{link}_{\text{Clique}(\text{Skel}_1(K^3))}(v)$. Specifically, for a graph G to represent the 1-skeleton of a triangulation of a closed 3-manifold, it is required that for every vertex v in G , the 2-skeleton $\text{Skel}_2(\text{link}_{\text{Clique}(G)}(v))$ must include a triangulated 2-sphere, and its vertex set must match the set of vertices in $\text{link}_{\text{Clique}(G)}(v)$.

Definition 37. A finite two-dimensional simplicial complex K^2 is termed *Hamiltonian* if it contains a triangulated 2-sphere with its vertex set being the vertices of K^2 . Such a sphere is referred to as a *Hamiltonian 2-sphere*.

This is an immediate relaxation of only requiring that K^2 contains a homology 2-cycle that ‘passes’ through all vertices. However, in the case where G is a complete graph, as we would have for the 1-skeleton of a neighborly triangulation of a closed 3-manifold, $L_G^c(v) = \text{link}_{\text{Clique}(G)}(v)$ is a simplex, and thus $\text{Skel}_2(L_G^c(v))$ is Hamiltonian but lacks non-trivial homology 2-cycles.

Additionally, our definition of ‘Hamiltonian’ differs from requiring that the 1-skeleton of a higher-dimensional complex should contain a Hamiltonian (edge-)cycle, as discussed in [36].

It follows directly from definition 37 that if the link complex $L_G^c(v)$ for some vertex v of G is not Hamiltonian, then G cannot be the 1-skeleton of a triangulated 3-manifold.

To check whether $L_G^c(v)$ is Hamiltonian, an exhaustive enumeration can be employed. Assuming n is the number of vertices of $L_G^c(v)$, a triangulated 2-sphere on n vertices has a face vector $f = (n, 3n - 6, 2n - 4)$. Therefore, we need to identify $2n - 4$ out of the triangles of $L_G^c(v)$ to form a 2-sphere. As the set of triangles of $L_G^c(v)$ is a subset of the set of all triangles on the n vertices, an adaptation of the enumeration procedure outlined in [71] will either return a Hamiltonian 2-sphere or determine that $L_G^c(v)$ is not Hamiltonian.

For larger links $L_G^c(v)$, this exact computation can be expensive, but it is not essential for our purposes. As shown in figure 28 for the graph G of our showcase sample, as computed in the previous sections, the number of triangles in a link essentially lies between $2n - 4$ and twice that number. However, there are two link outliers with $n = 5$, one with four and one with five triangles, falling below $2 \cdot 5 - 4 = 6$ triangles. For these two cases (and also for a few others with larger n), we assess below that they have digonal faces, allowing them to undercut the number $2n - 4$.

Since most links of vertices in the clique complex have between $2n - 4$ and $4n - 8$ triangles, it can be checked heuristically in most cases whether $L_G^c(v)$ contains a Hamiltonian 2-sphere or not using the following steps (0) and the reductions (i)–(iii):

- Step 0: If $L_G^c(v)$ has a vertex w for which its link in $L_G^c(v)$ contains no cycle, we declare the link $L_G^c(v)$ as *not good* and stop.
- Step i: If $L_G^c(v)$ has a vertex w for which its link in $L_G^c(v)$ is a 3-clique, we can delete w from $L_G^c(v)$, performing a bistellar flip that removes w . The resulting reduced complex then contains a Hamiltonian 2-sphere on $n - 1$ vertices if and only if $L_G^c(v)$ contains a Hamiltonian 2-sphere on n vertices.
- Step ii: If $L_G^c(v)$ contains a 4-clique $abcd$ in which two nonadjacent edges, say ab and cd , are each contained in exactly two triangles of $L_G^c(v)$, then we pick one of the edges, say ab ,




Figure 28. A scatter plot of the number of triangles in links of interior vertices in the clique complex. The solid line corresponds to $2n - 4$ triangles, which is the number that a triangulated 2-sphere with n vertices has. The dashed line represents twice that number.

as a free edge and remove it, effectively collapsing the tetrahedron $abcd$. Again, the reduced complex contains a Hamiltonian 2-sphere if and only if $L_G^c(v)$ contains a Hamiltonian 2-sphere.

We first perform (i) steps as long as possible, and then after each (ii) step, retry (i). If we find a (not necessarily unique) Hamiltonian 2-sphere this way, we regard the link of v as *good*.

If the exhaustive search steps (i) and (ii) fail to find a Hamiltonian 2-sphere in $L_G^c(v)$, we create a copy of the current link and continue with additional collapses:

- Step iii: If $L_G^c(v)$ contains a 4-clique $abcd$ where one (or more) of its edges, say ab , is contained in exactly two triangles (abc and abd) of $L_G^c(v)$, we once again take ab as a free edge to collapse away the tetrahedron $abcd$ by removing ab from $L_G^c(v)$. If the reduced complex contains a Hamiltonian 2-sphere, then $L_G^c(v)$ contains a Hamiltonian 2-sphere. However, the converse need not hold.

If a Hamiltonian 2-sphere is found, the link $L_G^c(v)$ of v is considered good. If not, we restart with the processed link after performing the (i) and (ii) steps and initiate a backtrack search for a Hamiltonian 2-sphere on the remaining triangles. A time limit of 20 s is imposed on this search. In our showcase example, there were only two cases where the time limit was reached: one for the largest pore with 107 neighbors (which, upon manual inspection, still had a hole in its link and was later identified as multiconnected) and another pore with 19 neighbors, where a Hamiltonian 2-sphere was not detected in its link within the time limit due to the high number of triangles.

If a Hamiltonian 2-sphere is not found after the steps (i) + (ii) + (iii) + backtrack search, or if a timeout occurs, the link $L_G^c(v)$ of v is considered not good, and these links are collected. The following additional check is then performed:

- Step iv: [Digon situation] If a vertex v has a link $L_G^c(v)$ that is not good, we iterate over all vertices in $L_G^c(v)$. If there exists a vertex u in $L_G^c(v)$ with only four or five neighbors in the entire graph G , we examine whether the link $L_G^c(u)$ has a cycle (e.g. if $L_G^c(u)$ is the boundary of a tetrahedron). If $L_G^c(u)$ does not have a cycle, we test whether the removal of the edge vu

makes the updated link $L_{G-uv}^c(v)$ Hamiltonian, and possibly also the links of other neighboring vertices to u . If so, we term the corresponding pore \mathcal{U} as a ‘football pore’ and temporarily remove u from G . In this case, $L_G^c(u)$ contains a multi-edge, and a correct reconstruction of \mathcal{U} would contain digonal facets.

For instance, in figure 12(left), the link of the blue pore, v , contains only one triangle, red–black–yellow, that includes the black pore. The black pore, u , has only four neighbors in G (red, yellow, blue, and green pores), and this configuration causes the link of the blue pore to be not good.

In our showcase sample, only three configurations with digonal faces were found in Step (iv) above. One involved a center black pore with four neighbors, and the other two had a center pore with five neighbors (two on one side and one on the other). Removing these three football pores resolved the link difficulties caused by digonal faces in our showcase sample. We store these removals to reinstate the respective pores later for the computation of the average face number. By focusing on vertices u with only four or five neighbors in G in our search, we might overlook more general configurations involving football pores. However, such configurations are expected to be rare, occurring mainly for small pores enclosed by very large pores.

For the remaining collected not-good links, the next step is to attempt adding edges to close further holes, thereby potentially making additional links Hamiltonian.

- Step v: [Addition of edges] If v is a vertex with a not-good link $L_G^c(v)$, it may be due to missing connections between v and a vertex w not yet in $L_G^c(v)$ or between two vertices r and s in $L_G^c(v)$ that should be connected. Candidate edges include those of the form vw , where w is a second neighbor of v (i.e. not in the link $L_G^c(v)$ but connected to v in G by a path of length two), and edges rs that connect not-yet-connected vertices of $L_G^c(v)$. The length of these edges must be at most $10c_1$, and their addition to G must make the link of v Hamiltonian.

All these candidate edges are collected for all vertices of G having not-good links, and the shortest one is selected and added to G . Subsequently, the graph G and affected vertex links are updated.

8.14. Adding long edges to close remaining holes

In the 2D sketch shown in figure 24(g), the link of the largest pore forms a path that might be closed by introducing a long edge between its end-vertices.

Revisiting the procedure outlined in section 8.7, we now examine candidate edges with lengths up to $20c_1$. The process involves iteratively selecting the shortest candidate edge, updating the configuration, and repeating as needed.

8.15. Removing not needed edges from 6-cliques

To streamline the subsequent computation of the dual triangulation, we undertake the removal of edges in this step, focusing on those that may not be strictly necessary. The process involves iterating over all 6-cliques and examining each edge vw within a 6-clique. We test whether the removal of this edge from G maintains the Hamiltonicity for the links $L_G^c(v)$, $L_G^c(w)$, and $L_G^c(u)$, where u is a mutual neighbor of v and w . From the resulting set of candidate edges, we select the longest one, eliminate it from G , and subsequently update both G and the impacted links.




Figure 29. Comparison of the link of a vertex v in the clique complex: after completing step 8.5 (left), after building the adjacency graph by step 8.15 (middle), and in the final reconstructed triangulation (right).

8.16. Overview and other foam samples

Different foam samples might necessitate adjustments to the thresholds kc_1 corresponding to various k values. However, for several samples akin to our showcase example, we obtained comparable results with the aforementioned choices.

8.17. Computational complexity and scalability

The Hamiltonian 2-sphere detection step is a central component of our certification algorithm and involves identifying spanning 2-spheres in the links of vertices. This task is computationally challenging in general, as finding Hamiltonian cycles in graphs is NP-hard. However, in our setting, these tests are applied locally-typically on link graphs with 10 to 25 vertices-which keeps the problem tractable in practice.

Algorithm 4, which evaluates cycles in vertex links, is also computationally expensive. This step requires computing all cycles in certain induced subgraphs, a process that can grow exponentially in the worst case. Though heuristics and pruning help limit the number of candidate cycles, this step remains one of the primary bottlenecks in the pipeline.

For the aluminum foam dataset containing 1911 pores, the full reconstruction process-including all combinatorial and topological tests-was completed in approximately 19 h on a MacBook Pro with an M1 chip and 16 GB of RAM (single-threaded implementation in Python). Memory usage remained modest throughout since the algorithm relies solely on local combinatorial operations.

While the method performs reliably for medium-sized datasets (up to a few thousand pores), scalability to significantly larger systems (e.g. $>10\,000$ pores) would require further optimization. Potential improvements include: (1) parallelization of local computations, (2) memoization of frequently encountered link configurations, and (3) use of physical constraints (e.g. expected face degrees) to reduce the combinatorial search space.

Figure 29 depicts the progress of reconstructing the link of an interior vertex; left: the link of the vertex after recording the initial adjacencies in Step 8.5. This link is evidently non-Hamiltonian. Middle: the link in the clique complex after completing Step 8.15, presenting multiple choices for Hamiltonian cycles. Right: the link in the triangulation we construct in section 9.

In table 2, we provide the statistics of the adjacency graph after each step in this section, excluding the final step, which will be discussed in the next section.

Table 2. Progress of individual steps in constructing an adjacency graph.

Step	Number of added edges	Number of removed edges	Number of all edges	Average degree of interior vertices
8.5 Adaptive computation of nearest neighbors	6431	0	6431	7.72
8.6 Degree requirement for edges	2126	0	8557	10.58
8.7 Detection of holes in vertex links	2031	0	10 588	13.24
8.8 Adding edges to close edge links	483	0	11 071	13.62
8.9 Adding degree-3 edges	102	0	11 173	13.76
8.10 Detection of further holes in vertex links	95	0	11 268	13.77
8.11 Removal of excessive edges	0	56	11 212	13.70
8.12 Detection of further holes in vertex links (repeat)	69	0	11 281	13.78
8.13 Hamiltonicity of vertex-links and detection of digonal faces	17	0	11 298	13.79
8.14 Adding long edges to close remaining holes	2	0	11 300	13.80
8.15 Removing not needed edges from 6-cliques	0	12	11 288	13.78
9.4 Repairing remaining problematic vertex links	0	11	11 277	13.76

9. Building the generalized triangulation

Let G denote the graph introduced in the preceding section, specifically for our showcase sample. In the subsequent discussion, we maintain G as a constant (at least temporarily; only in section 9.4 will we eliminate additional edges from G as excessive edges). Through the processes outlined in the prior section, the majority of links among interior vertices in $\text{Clique}(G)$ are Hamiltonian, with only a few exceptions. These exceptions may arise due to factors such as a pore being multi-connected to a neighboring pore, digons in its link, or boundary effects.

The objective of this section is to establish a consistent generalized triangulation, denoted as K^3 , for the entire sample region, treating it as a 3-ball on the vertex set of G , with the graph G as the 1-skeleton of K^3 . However, given the heuristic nature of our approach and the limited assumptions on the input data, we may not achieve a consistent triangulation for every interior vertex. The subset of vertices for which we encounter difficulties in constructing their vertex stars is referred to as the *fail part* in our output.

9.1. An initial partial triangulation

Start with $L_G^c(v) = \text{link}_{\text{Clique}(G)}(v)$. The optimal scenario is when $L_G^c(v)$ forms a triangulated 2-sphere. In such instances, $L_G^c(v)$ is either the boundary of a tetrahedron or is flag*.




Figure 30. Left: Barycenter locations of the 109 pores with triangulated 2-sphere links in the clique complex; the eight colored dots highlight those with tetrahedral links. Middle: Locations of the 259 pores featuring flag* links in the final triangulation. Right: Locations of 88 (out of 1223) 5-cliques whose tunnels are changed.

For our featured sample, 109 vertices have their links determined this way, with eight vertices having the boundary of a tetrahedron as their link, and 101 having flag* links. The distribution of links by size is as follows:

4 : 8, 6 : 10, 7 : 8, 8 : 8, 9 : 8, 10 : 19, 11 : 13, 12 : 15, 13 : 6, 14 : 4, 15 : 4, 16 : 2, 18 : 1, 19 : 1.

Figure 30(left) illustrates the barycenter locations of these 109 pores. If two pores are neighbors, they are connected by an edge. The eight red dots represent pores with tetrahedral links. Figure 30(middle) displays a plot of the 259 pores with flag* links in the final triangulation, including the 101 flag* pores from the left image. Once again, pores are connected by an edge if they are neighbors. The distribution of these flag* pores is critical to our approach, forming a backbone network throughout the sample, enhancing the likelihood of successfully reconstructing links for additional neighboring pores. The larger holes in figure 30(left) correspond to the locations of the largest pores in the sample, with the three largest ones (as in figure 15) and their connecting edges to flag* pores indicated in red, green, and blue in figure 30(middle).

The pore with the flag* link comprising 19 vertices is highlighted in black in figure 31(left), accompanied by visualizations of its link (middle) and its pseudo-realization (right). The black pore has the pink pore as one of its neighbors; as determined by our reconstruction. The spike of the pink pore influences the positioning of the representing vertex (barycenter) for the pink pore in our visualization of the black pore's link, causing it to be situated notably to the right. Despite the somewhat non-convex, hummingbird-like shape of the pink pore, the reconstruction of this link is easily accomplished owing to its flag*-ness.

In the more general case, if $L_G^c(v)$ is not a triangulated 2-sphere but includes a unique Hamiltonian 2-sphere N^2 , we determine N^2 using algorithm 2. The resulting N^2 can be either a stacked flag* 2-sphere or a stacked sc* 2-sphere (refer to [72] for further details).

If $L_G^c(v)$ has more than one Hamiltonian 2-cycle, this may arise from a local cross configuration, as illustrated in example 32 and discussed in section 8.7. In the presence of a single cross, vertex v is part of a 5-clique where the cross forms a K_4 in $L_G^c(v)$, leading to two distinct Hamiltonian 2-cycles in $L_G^c(v)$ —each utilizing two of the four triangles of the 4-clique.

More generally, if vertex v has a Hamiltonian link $L_G^c(v)$ and is situated in one or several k -cliques where $k \geq 5$, then $L_G^c(v)$ is likely to contain more than one Hamiltonian 2-sphere.




Figure 31. Left: The black pore has the pink pore among its 19 neighbors. Middle: The link of the black pore with 19 vertices (depicted as a central thickened dot) showcases the barycenter positions of the pores. The right-most dot signifies the pink pore, while the red dots denote boundary pores. Right: Pseudo-realization of the simple polytope representing the black pore.

Table 3. Progress along the individual steps in constructing a triangulation.

Statistics on the graph G :	Number of certified 4-cliques	Properly reconstructed links of interior edges	Properly reconstructed links of interior vertices	Number of changed tunnels
- Number of interior vertices: 991				
- Number of interior edges: 8080				
- Number of 4-cliques: 10 878				
- Number of 5-cliques: 1223				
- Number of 6-cliques: 0				
Initial 109 interior vertices with complete links	1436	12.48%	11%	—
Combinatorial results of algorithm 2	4498	41.98%	17.66%	—
Geometric improvements by algorithm 3	4804	46.57%	20.79%	—
Algorithm 5: Intermediate result with default tunnels	8299	97.02%	80.93%	0
Algorithm 5: Final result with changed tunnels	8328	98.18%	90.72%	88
Step 9.3: Repairing remaining problematic edge links	8337	99.72%	97.66%	0
Step 9.4: Repairing remaining problematic vertex links	8337	99.81%	99.19%	0

We initiate algorithm 2 followed by algorithm 3 on the graph G representing our showcase sample. The initial 109 vertices v with $L_G^c(v)$ as a triangulated 2-sphere account for 11% of the interior vertices. These are part of the 17.66% of interior vertices for which $L_G^c(v)$ contains a unique Hamiltonian 2-sphere, as determined by algorithm 2. However, algorithm 3 demonstrates limited progress, merely enhancing the fraction of successfully reconstructed vertex links to 20.79% (refer to table 3). In contrast, if G represented the 1-skeleton of a geometric triangulation K^3 for a convex 3-ball, algorithm 3 would facilitate the complete reconstruction of the triangulation K^3 from its 1-skeleton. The subsequent section explores the reasons behind the suboptimal performance of the geometric checks in algorithm 3 within the context of the showcase sample.




Figure 32. A 5-clique whose tunnel needs to be changed.




Figure 33. (Left) Possible configuration where edge dv pierces the empty triangle abc and decertifies it. (Right) Possible configuration where edge dv fails to decertify triangle abc but instead wrongly decertifies a neighboring triangle.

9.2. 5-cliques and tunnels

In the configuration shown in figure 32(left), we observe five interconnected pores forming a K_5 in G . Their barycenters form a bipyramid over a triangle, with the red, dark blue, and green pores forming the base triangle and the yellow and light blue pores forming the apices (figure 32 right). Due to the non-convexity of the yellow pore, we have a ring formed by the green, light blue, and yellow pores, with the connection between red and dark blue pores forming a ‘tunnel’ through this ring.

This observation prompts us to modify the initial geometric tunnel de (with respect to the barycenters of the pores) through the empty triangle abc to the tunnel ab through the empty triangle cde . According to corollary 30, only three out of the five tetrahedra of K_5 can be present in a triangulation (unless one of the vertices of the K_5 is of degree 4 in G). In algorithm 3, step 1(ii) decertifies two of the five possible tetrahedra of a K_5 , which is effective for geometric triangulations but may lead to erroneous decisions in our specific situation.

Figure 33 illustrates a K_5 for a central vertex v , along with its link (in gray). Assume that the highlighted triangle abc is an empty triangle. Then, depending on the locations of the barycenters of the pores, the edge dv either pierces the triangle abc (figure 33 left), thereby decertifying it, or fails to decertify abc and instead erroneously decertifies a neighboring triangle (figure 33 right). If we make errors of this kind for multiple K_5 ’s, our reconstruction process may reach an impasse.

To make matters worse, we might mistakenly decertify a tetrahedron $abcd$ geometrically if for some pore e neighboring, say, pore d , the edge de intersects the triangle abc , even though the five pores $a, b, c, d,$ and e do not form a K_5 . See figure 34 for an example of such a configuration.

This presents us with a quandary. Algorithm 3 is an effective tool for reconstructing a triangulation from its graph in the case of geometric triangulations of convex 3-balls, such as




Figure 34. Example of a configuration where edge de intersects triangle abc , leading to the erroneous decertification of tetrahedron $abcd$ despite the pores not forming a K_5 .

Delaunay triangulations dual to Poisson–Voronoi diagrams. However, it fails when applied to the graph G , where the vertices are positioned at the barycenters of the pores. To resolve this issue, we propose a modified approach that omits steps 1 and 2 of algorithm 3, which involve general decertifications. Instead, we focus on 5-cliques (excluding those with a degree-4 vertex) for which corollary 30 establishes that at most three of their five 4-cliques can be present in a triangulation. For each 5-clique, we initially set its tunnel geometrically as its *default tunnel*. However, we allow the tunnels to be modified if this leads to local improvements for the involved vertex links.

In the event that an interior vertex v of G has degree 4 and its link $L_G^c(v)$ is Hamiltonian, meaning $L_G^c(v)$ is the boundary of a tetrahedron, then we decertify this tetrahedron as a special combinatorial case identified in Step 1(i) of algorithm 3. For any other 5-clique $abcde$ of G with vertices a, b, c, d , and e (barycenters of pores), the convex hull $\text{conv}(a, b, c, d, e)$ either forms a tetrahedron or a bipyramid over a triangle.

In cases where $\text{conv}(a, b, c, d, e)$ forms a bipyramid over a triangle, we select the corresponding geometric tunnel as the default tunnel. If $\text{conv}(a, b, c, d, e)$ is a tetrahedron (with all vertices having a degree of at least five), we randomly choose one of the ten possible tunnels combinatorially as the default tunnel. Each default tunnel choice leads to the decertification of two tetrahedra. We gather these tetrahedra and remove them from $\text{Clique}(G)$ to derive a reduced complex $R \subseteq \text{Clique}(G)$. For each vertex v in G , $\text{link}_R(v) \subseteq L_G^c(v)$ holds true. However, the selection of default tunnels might inadvertently puncture $\text{link}_R(v)$ due to incorrect decertifications, thus creating holes, even though the link $L_G^c(v)$ was originally Hamiltonian.

We identify all vertices v where $L_G^c(v)$ is Hamiltonian, but $\text{link}_R(v)$ is not, and place them in a set V_{non} . Subsequently, we attempt to close the holes for the vertices in V_{non} by adjusting tunnels. We organize the set V_{non} as a list, sorted in ascending order based on the weighted distance of a pore's barycenter from the overall barycenter of the cuboid in our showcase sample.

Consider a vertex v belonging to V_{non} with $\text{link}_R(v)$ being non-Hamiltonian. As discussed earlier in section 8.7, we identify *exposed vertices* of $\text{link}_R(v)$ lying in exposed edges of the link. We then examine all 5-cliques containing v and an exposed vertex of $\text{link}_R(v)$. For each such K_5 denoted as $vabcd$, we determine the intersection $A = \{v, a, b, c, d\} \cap V_{\text{non}}$. Since there are $\binom{5}{2} = 10$ combinatorial choices for a tunnel in a 5-clique, we have nine options, apart from the initially chosen tunnel, to change the tunnel.

For each of these nine tunnel change possibilities, we verify whether the chosen tunnel removal excludes v from V_{non} and, consequently, from A . Among these options, we select the one that never adds any of the vertices of the K_5 to A while minimizing the size of A . In other words, we adopt a greedy approach to tunnel changes: we opt to change a tunnel if

Algorithm 5. Compute a Partial Simplicial Complex Triangulation.

Input: Abstract graph G from section 8, where vertices represent pores of a metallic foam, subset G_{int} of interior vertices of G , and pore barycenters G_{bary}

Output: Lists C , D , and U of certified, decertified, and undecided 4-cliques of G

0. **Initialization.** Set C , D , and U as the output of algorithm 2 applied to G .
 1. **Default tunnel selection.** For each K_5 in G without a vertex of degree 4, choose its default tunnel, thereby decertifying two out of five 4-cliques in each such K_5 .
 2. **Algorithm rerun with default tunnels.** Rerun algorithm 2 on the reduced complex with the default tunnels until no further progress is made.
 3. **Collection of vertices with bad links.** Identify interior vertices v in G with Hamiltonian $L_G^c(v)$ but non-Hamiltonian links in the reduced complex.
 4. **Tunnel modification.** Change tunnels until no further progress is made.
 5. **Result return.** Return the updated lists C , D , and U .
-

- (i) the link of v is repaired,
- (ii) no new holes are introduced for other vertices in $\text{link}_R(v)$, and
- (iii) the net effect, measured as the reduction in the total number of vertices with non-Hamiltonian links in $\text{link}_R(v)$, is maximal among the available options.

This strategy prevents us from getting stuck in loops but may result in leaving the link of v unrepaired as we move on to the next vertex in V_{non} . However, the repair of v may occur later when it appears in the link of another vertex w . In such a case, v is removed from the set A' corresponding to w and also from V_{non} .

The remaining vertices of A' that remain unrepaired after changing a tunnel in a 5-clique containing w may have experienced some improvement, addressing one of their issues but not all. We gather all vertices of A' that were processed before w and prioritize checking them before moving on to the next vertex in the sorted list V_{non} following w .

As mentioned earlier, algorithm 3 encounters challenges when decertifying graphs G that do not represent the 1-skeleta of geometric triangulations. In the case of our graph G , the algorithm erroneously decertifies 37 tetrahedra not situated in 5-cliques. These inaccuracies persist when algorithm 2 is rerun, preventing the correction of these erroneous decertifications. Consequently, only 20.79% of interior vertex links are properly reconstructed, as detailed in table 3.

In contrast, when restricting geometric decertification to 5-cliques exclusively, combined with the selection of default tunnels and the subsequent rerun of algorithm 2, we achieve an 80.93% proper reconstruction of interior vertex links. Further improvement to 90.72% is accomplished by modifying 88 tunnels, where each change met the improvement criteria outlined earlier. The corresponding 5-cliques with changed tunnels in our showcase sample are depicted in figure 30(right), showing a sparse distribution across the sample.

9.3. Repairing remaining problematic edge links

Based on the current results following the execution of algorithm 5, as outlined in table 3, 98.18% of the interior edges now possess a triangulated 1-sphere as their link. In contrast, only 90.72% of the vertices have a triangulated 2-sphere as their link. For the remaining vertices without a triangulated 2-sphere link, the vertex-links are often approximate spheres, potentially featuring only a few vertices w in the link of v where the link of the corresponding edges vw is not a triangulated 1-sphere.

Our focus now shifts to rectifying these edge links that fall short of being triangulated 1-spheres. Given the manageable number of such edges, we employ a brute force approach to address these link issues.

The process begins with an edge $e = uv$ identified with a faulty link. Let V_e denote the set of mutual neighbors of vertices u and v , and G_e represent the subgraph induced by V_e in G . Additionally, C_e denotes the set of simple cycles in G_e , computed selectively based on the sparsity of G_e —a condition usually met for the foam structures in our input.

Among the cycles in C_e , we gather those that do not disconnect the links of u or v within the partial triangulation established so far. This subset of C_e serves as the pool of candidates for potential links of edge e in a proper triangulation. Iterating over these candidates, let C be one such cycle. Temporarily adjusting the triangulation so that C becomes the link of e (by adding and removing necessary tetrahedra for this purpose), we compute the difference between the edges in the neighborhood of e that have had their links repaired and those that have had their links compromised. This serves as a metric to monitor potential improvements. We select the cycle maximizing this metric with this maximum being positive, considering a similar metric for repaired vertex links as a tie-breaker (i.e. the difference between the vertices in the neighborhood of e that have had their links repaired and those that have had their links compromised), and modify the triangulation accordingly.

If no such cycle is found, we examine whether e qualifies as a multi-edge. To determine this, we temporarily remove edge e from the triangulation and compute the links of vertices u and v . If each link contains more than one hole, and these holes coincide, we categorize edge e as a multi-edge.

In our study sample, this step identifies one pair of multi-connected pores, illustrated in figure 11.

9.4. Repairing remaining problematic vertex links

After excluding multi-connected vertices and digons, 97.66% of interior vertices now exhibit proper links. Once again, we employ a brute force methodology to identify potential excessive edges in the links of the remaining interior vertices with non-proper links. For a given vertex v with a non-proper link, the candidates are edges uv with distant corresponding pores, $d_{\text{up}}(\mathcal{U}, \mathcal{V}) > c_1$, and the removal of edge uv repairs the link of v without impairing any links in its vicinity. Following the approach outlined in section 8, we eliminate the candidate edge with the most distant corresponding pores, update the list of candidates, and iterate until no further candidate edges remain.

10. Limitations and conclusion

After implementing Step 9.4, the majority of interior vertices in our study sample now exhibit proper links. However, certain cases necessitate manual investigation to identify local obstructions preventing them from achieving proper links. Figure 35 highlights three limitations of our reconstruction method:

- **Failure to identify excessive edges due to small pore distances:** The green and yellow pores are incorrectly identified as connected despite being separated by a thin layer of the dark blue pore.
- **Failure to detect multi-connections:** The dark blue and red pores are not correctly identified as having a multi-connection.




Figure 35. Limitations of the proposed reconstruction method.

- **Inaccurate pore distance estimation:** The estimated distance between the dark blue and light blue pores ($15.91 c_1$, as computed via the method outlined in section 8.2) is significantly different from the actual distance ($4.35 c_1$).

Typically, the first limitation is addressed in Step 9.4, while the second is captured in Step 9.3. However, the simultaneous occurrence of both in the same vicinity is the primary cause of the failure.

Another potential fail case might arise when, during computation and with the increase in c_1 and c_2 , we encounter a larger clique in the dual graph, i.e. a complete subgraph. Theoretically, this can happen, as the Delaunay triangulation of a point set can be neighborly [43]. However, cliques K_r with $r \geq 7$ are infrequent. We conducted Delaunay triangulations for 100 Poisson point processes in 3D, each consisting of 1000 points, and identified a 7-clique in only four of those samples.

Failure can also occur when a foam (or another cellular material) exhibits a genuine defect, meaning that, at a local level, the material loses its cellular structure. We identify and output these failure cases F . Further analysis is then required, which may involve enhancing our approach with additional heuristics (incurring extra computational cost) or physically assessing whether the material indeed has a genuine defect.

To assess the robustness of our method, we applied it to additional x-ray tomography datasets beyond the aluminum foam sample discussed in detail. These included real-world metallic foams with varying degrees of pore polydispersity and structural noise. Across these samples, the reconstruction success rate ranged from 97% to 100%. We also tested the method on synthetic foams generated via Laguerre–Voronoi tessellations, where the underlying ground-truth topology is known. In all synthetic cases, the reconstruction rate was 100%. These results confirm the generalizability and accuracy of our algorithm across both empirical and controlled settings.

In conclusion, a systematic analysis is essential to evaluate the geometric, topological structure, and combinatorial arrangement of foam pores, correlating these aspects with physical properties. The method presented in this paper approximates foam pores with combinatorial polytopes by leveraging neighborhood graphs and generalized triangulations, defining a ‘combinatorial profile’ of the foam. This enables meaningful computations of metrics such as the average face degree and the distribution of combinatorial pore types.

By providing a structured framework for analyzing foam topology, this approach bridges the gap between geometric reconstructions and combinatorial representations, offering a robust tool for studying disordered cellular materials. Understanding the combinatorial structure of foams is crucial for applications in material science, where properties such as mechanical strength and thermal conductivity are closely tied to pore connectivity. Moreover, this method

lays the groundwork for more advanced reconstructions by integrating additional geometric and physical data, making it a valuable step toward more precise modeling of real-world foams.

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: https://github.com/IhabSabik/reconstruction_foam.

Acknowledgments

We thank Junichi Nakagawa for highlighting this problem, Mike A Noack for assisting in the production of the metal foam sample used in this study, and Myfanwy E Evans for her valuable feedback. This research is carried out in the framework of the DFG funded Cluster of Excellence EXC 2046 MATH+: The Berlin Mathematics Research Center.

In Memory of Frank H Lutz, whose contributions were essential to this work.

Appendix. Cellular decompositions of manifolds

In topology, general cellular complexes are typically described as CW complexes (see [44]), with topological balls serving as their building blocks. Specifically, each cell c is homeomorphic to a k -dimensional (open) ball $B^k = \{x \in E^k \mid \|x\| < 1\}$ in a Euclidean space E^k for some k .

A CW complex X is a Hausdorff topological space decomposed into a disjoint union of open cells c , satisfying two key conditions:

- **Closure-finiteness:** Let the $(k-1)$ -skeleton X_{k-1} of X be the union of all cells of X of dimension at most $k-1$. Then, for any open k -cell c of X , there exists a continuous map $f_c : B^k = \{x \in E^k \mid \|x\| \leq 1\} \rightarrow X$ on the closed ball $\overline{B^k}$ such that $f_c(B^k) = c$ and $f_c(\overline{B^k})$, the closure of c , intersects only finitely many cells of X_{k-1} . The restriction $f_c|_{\partial\overline{B^k}}$ of f_c to the boundary $\partial\overline{B^k}$ of $\overline{B^k}$ is called the *attaching map* for c .
- **Weak topology:** A subset $Y \subseteq X$ is closed in X if and only if $Y \cap f_c(\overline{B^k})$ is closed for every cell c of X , meaning Y intersects the closure of each cell in a closed set.

CW complexes follow an inductive building procedure, commencing with a discrete set of points called *vertices*. Subsequently, 1-cells, called *edges* are attached, followed by 2-cells, and so forth. The continuous maps that incorporate cells need only adhere to the specified restrictions, often enabling economical decompositions of topological spaces as CW complexes. For instance, the d -dimensional sphere S^d can be decomposed into a CW complex with only two cells: one zero-dimensional cell and one d -dimensional cell, where the entire boundary of the latter is mapped to the single vertex under the attaching map.

Regular CW complexes are CW complexes for which $f : B^k \rightarrow f(\overline{B^k}) \subseteq X$ is a homeomorphism for every cell c of X , indicating that the closure of a cell c in X is a closed ball. In such complexes, two cells may intersect in disjoint sets, potentially of different dimensions. For example, in one-dimensional regular CW complexes viewed as graphs, parallel edges can exist.

Strongly regular CW complexes extend the regularity condition by ensuring that any two closed cells intersect in a single closed cell, which may be empty. These complexes can be combinatorially encoded through the incidence structure of their cells.

Polytopal complexes are a subset of strongly regular CW complexes, with the added requirement that the closure of every k -cell is realizable as a k -dimensional polytope. A polytopal complex is called *pure* if all of its inclusion-maximal faces share the same dimension.

Generalized triangulations of d -manifolds, whether with or without boundaries, are collections of d -simplices whose $(d - 1)$ -faces are affinely glued together in pairs, resulting in a complex that is homeomorphic to the d -manifold (see [23]). For instance, identifying both vertices of a single edge provides a generalized triangulation of the 1-sphere S^1 .

While CW complexes are permitted to have infinitely many cells and can even be of infinite dimension, for the purpose of our discussion, we exclusively focus on *finite* CW complexes characterized by a finite number of faces. When considering decompositions of Euclidean spaces E^d , we also allow for (possibly non-bounded) *polyhedral cells*, which are intersections of finitely many closed half-spaces.

ORCID iDs

Ihab Sabik  [0009-0009-0963-2841](https://orcid.org/0009-0009-0963-2841)

Paul Hans Kamm  [0000-0001-6353-768X](https://orcid.org/0000-0001-6353-768X)

Francisco García-Moreno  [0000-0001-6114-0240](https://orcid.org/0000-0001-6114-0240)

Frank H Lutz  [0009-0008-5375-4136](https://orcid.org/0009-0008-5375-4136)

References

- [1] Abendroth M, Werzner E, Settgast C and Ray S 2017 An approach toward numerical investigation of the mechanical behavior of ceramic foams during metal melt filtration processes *Adv. Eng. Mater.* **19** 1700080
- [2] Albertson M O and Hutchinson J P 1977 The independence ratio and genus of a graph *Trans. Am. Math. Soc.* **226** 161–73
- [3] Albertson M O and Stromquist W R 1982 Locally planar toroidal graphs are 5-colorable *Proc. Am. Math. Soc.* **84** 449–57
- [4] Alt H, Cheong O and Vigneron A 2005 The Voronoi diagram of curved objects *Discrete Comput. Geom.* **34** 439–53
- [5] Arya S and Mount D M 2017 Computational geometry: proximity and location *Handbook of Data Structures and Applications* 2nd ed (Chapman & Hall/CRC) p 63.1–63.22
- [6] Attali D, Lieutier A and Salinas D 2011 Efficient data structure for representing and simplifying simplicial complexes in high dimensions *Proc. 27th Annual Symp. on Computational Geometry (SoCG '11)* (Association for Computing Machinery) pp 501–9
- [7] Aurenhammer F 1987 Power diagrams: properties, algorithms and applications *SIAM J. Comput.* **16** 78–96
- [8] Avinash G, Harika V, Sandeepika C and Gupta N 2018 Pore size control in aluminium foam by standardizing bubble rise velocity and melt viscosity *IOP Conf. Ser.: Mater. Sci. Eng.* **338** 012010
- [9] Babcsán N, García-Moreno F and Banhart J 2007 Metal foams-high temperature colloids: part II: in situ analysis of metal foams *Colloids Surf. A* **309** 254–63
- [10] Banhart J 1999 Foam metal: the recipe *Europhys. News* **30** 17–20
- [11] Banhart J 2000 Manufacturing routes for metallic foams *JOM* **52** 22–27
- [12] Banhart J 2005 Aluminium foams for lighter vehicles *Int. J. Veh. Des.* **37** 114–25
- [13] Banhart J 2006 Metal foams: production and stability *Adv. Eng. Mater.* **8** 781–94
- [14] Baumeister J 1997 Überblick: Verfahren zur Herstellung von Metallschäumen *Beiträge zum Symposium Metallschäume* ed J Banhart (MIT Verlag) pp 3–14

- [15] Björner A and Lutz F H 2000 Simplicial manifolds, bistellar flips and a 16-vertex triangulation of the Poincaré homology 3-sphere *Exp. Math.* **9** 275–89
- [16] Blind R and Mani-Levitska P 1987 Puzzles and polytope isomorphisms *Aequationes Math.* **34** 287–97
- [17] Bochnak J, Coste M and Roy M-F 1998 *Real Algebraic Geometry (Ergebnisse der Mathematik und Ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)])* vol 36 (Springer) Translated from the 1987 French original, Revised by the authors
- [18] Boissonnat J-D, Wormser C and Yvinec M 2007 Curved Voronoi diagrams *Effective Computational Geometry for Curves and Surfaces, Mathematics and Visualization* (Springer) pp 67–116
- [19] Boissonnat J-D and Yvinec M 1998 *Algorithmic Geometry* (Cambridge University Press)
- [20] Bollobás B and Riordan O 2008 Percolation on random Johnson-Mehl tessellations and related models *Probab. Theory Relat. Fields* **140** 319–43
- [21] Brakke K A 1992 The surface evolver *Exp. Math.* **1** 141–65
- [22] Brehm U and Kühnel W 1987 Combinatorial manifolds with few vertices *Topology* **26** 465–73
- [23] Burton B A 2011 Simplification paths in the Pachner graphs of closed orientable 3-manifold triangulations (arXiv:1110.6080)
- [24] Cannon J W 1979 Shrinking cell-like decompositions of manifolds. Codimension three *Ann. Math.* **110** 83–112
- [25] Cazals F and Giesen J 2006 *Delaunay Triangulation Based Surface Reconstruction* (Springer) pp 231–76
- [26] Dancis J 1984 Triangulated n -manifolds are determined by their $[n/2] + 1$ -skeletons *Top. Appl.* **18** 17–26
- [27] DeVos M, Kawarabayashi K-I and Mohar B 2008 Locally planar graphs are 5-choosable *J. Comb. Theory B* **98** 1215–32
- [28] Dwyer R A 1991 Higher-dimensional Voronoi diagrams in linear expected time *Discrete Comput. Geom.* **6** 343–67
- [29] Edelsbrunner H, Kirkpatrick D G and Seidel R 1983 On the shape of a set of points in the plane *IEEE Trans. Inf. Theory* **29** 551–9
- [30] Edwards R D 1980 The topology of manifolds and cell-like maps *Proc. of the Int. Congress of Mathematicians (Helsinki, 1978)* pp 111–27
- [31] Elber G and Kim M-S 2000 A computational model for nonrational bisector surfaces: Curve-surface and surface-surface bisectors *Proc. Geometric Modeling and Processing 2000. Theory and Applications* pp 364–72
- [32] Emiris I, Tsigaridas E and Tzoumas G 2008 The predicates for the exact Voronoi diagram of ellipses under the Euclidean metric *Int. J. Comput. Geom. Appl.* **18** 567–97
- [33] Emiris I Z and Karavelas M I 2006 The predicates of the Apollonius diagram: algorithmic analysis and implementation *Comput. Geom. Theory Appl.* **33** 18–57
- [34] Evans M E, Kraynik A M, Reinelt D A, Mecke K and Schröder-Turk G E 2013 Networklike propagation of cell-level stress in sheared random foams *Phys. Rev. Lett.* **111** 138301
- [35] Evans M E, Schröder-Turk G E and Kraynik A M 2017 A geometric exploration of stress in deformed liquid foams *J. Phys.: Condens. Matter* **29** 124004
- [36] Ewald G 1973 Hamiltonian circuits in simplicial complexes *Geom. Dedicata* **2** 115–25
- [37] Franklin W and Landis E N 2006 Connected components on $1000 \times 1000 \times 1000$ datasets *16th Fall Workshop on Computational and Combinatorial Geometry*
- [38] Furuta Y, Oikawa N and Kurita R 2016 Close relationship between a dry-wet transition and a bubble rearrangement in two-dimensional foam *Sci. Rep.* **6** 37506
- [39] García-Moreno F 2016 Commercial applications of metal foams: their properties and production *Materials* **9** 85
- [40] García-Moreno F, Babcsan N and Banhart J 2005 X-ray radioscopy of liquid metalfoams: influence of heating profile, atmosphere and pressure *Colloids Surf. A* **263** 290–4
- [41] García-Moreno F, Tobin S T, Mukherjee M, Jiménez C, Solórzano E, Kumar G S V, Hutzler S and Banhart J 2014 Analysis of liquid metal foams through x-ray radioscopy and microgravity experiments *Soft Matter* **10** 6955–62
- [42] Glaser L C 1970 *Geometrical Combinatorial Topology. Vol. I (Van Nostrand Reinhold Mathematics Studies)* vol 27 (Van Nostrand Reinhold Co.)
- [43] Gonska B and Padrol A 2016 Neighborly inscribed polytopes and Delaunay triangulations *Adv. Geom.* **16** 349–60

- [44] Hatcher A 2005 *Algebraic Topology* (Cambridge University Press)
- [45] Heim K, Vinod-Kumar G, García-Moreno F, Rack A and Banhart J 2015 Stabilisation of aluminium foams and films by the joint action of dispersed particles and oxide films *Acta Mater.* **99** 313–24
- [46] Hempel J 2004 *3-Manifolds* (AMS Chelsea Publishing) Reprint of the 1976 original
- [47] Höhler R, Yip Cheung Sang Y, Lorenceau E and Cohen-Addad S 2008 Osmotic pressure and structures of monodisperse ordered foam *Langmuir* **24** 418–25
- [48] Hutzler S, Weaire D, Saugey A, Cox S and Peron N 2005 The physics of foam drainage *Proc. MIT European Detergents Conf., Würzburg* pp 191–206
- [49] Isenburg M and Shewchuk J 2009 Streaming connected component computation for trillion voxel images *Workshop on Massive Data Algorithmics*
- [50] Johnson W A and Mehl R F 1939 Reaction kinetics in processes of nucleation and growth *Trans. Am. Inst. Mining Metall. Eng.* **135** 416–42
- [51] Joswig M 2002 Projectivities in simplicial complexes and colorings of simple polytopes *Math. Z.* **240** 243–59
- [52] Jungerman M and Ringel G 1980 Minimal triangulations on orientable surfaces *Acta Math.* **145** 121–54
- [53] Kalai G 1988 A simple way to tell a simple polytope from its graph *J. Comb. Theory A* **49** 381–3
- [54] Kanaun S and Tkachenko O 2006 Mechanical properties of open cell foams: simulations by Laguerre tessellation procedure *Int. J. Fract.* **140** 305–12
- [55] Karavelas M I and Yvinec M 2003 The Voronoi diagram of planar convex objects *Algorithms—ESA 2003 (Lecture Notes in Computer Science)* vol 2832 (Springer) pp 337–48
- [56] Keil J M 1994 Computing a subgraph of the minimum weight triangulation *Comput. Geom.* **4** 18–26
- [57] Kenesei P, Kádár C, Rajkovits Z and Lendvai J 2004 The influence of cell-size distribution on the plastic deformation in metal foams *Scr. Mater.* **50** 295–300
- [58] Kirkpatrick D G and Radke J D 1985 A framework for computational morphology *Computational Geometry (Machine Intelligence and Pattern Recognition)* vol 2 (North-Holland) pp 217–48
- [59] Klein R 1988 Abstract Voronoi diagrams and their applications (extended abstract) *Computational Geometry and its Applications (Würzburg (Lecture Notes in Computer Science))* vol 333 (Springer) pp 148–57
- [60] Köhler E G and Lutz F H 2005 Triangulated manifolds with few vertices: vertex-transitive triangulations I (arXiv:math/0506520v1)
- [61] Körner C, Arnold M and Singer R F 2005 Metal foam stabilization by oxide network particles *Mater. Sci. Eng. A* **396** 28–40
- [62] Kraynik A 2006 The structure of random foam *Adv. Eng. Mater.* **8** 900–6
- [63] Kraynik A M, Reinelt D A and van Swol F 2004 Structure of random foam *Phys. Rev. Lett.* **93** 208301
- [64] Kühnel W and Banchoff T F 1983 The 9-vertex complex projective plane *Math. Intell.* **5** 11–22
- [65] Kumar V S and Kumaran V 2005 Voronoi neighbor statistics of hard-disks and hard-spheres *J. Chem. Phys.* **123** 074502
- [66] Kusner R and Sullivan J M 1996 Comparing the Weaire-Phelan equal-volume foam to Kelvin's foam *Forma* **11** 233–42
- [67] Lazar E A, Mason J K, MacPherson R D and Srolovitz D J 2012 Complete topology of cells, grains and bubbles in three-dimensional microstructures *Phys. Rev. Lett.* **109** 095505
- [68] Leymarie F F and Kimia B B 2001 The shock scaffold for representing 3D shape *Visual Form 2001* (Springer) pp 216–27
- [69] Liebscher A, Proppe C, Redenbach C and Schwarzer D 2013 Stochastic multiscale modeling of metal foams *Procedia IUTAM (IUTAM Symp. on Multiscale Problems in Stochastic Mechanics)* vol 6 pp 87–96
- [70] Lutz F H 1999 Triangulated manifolds with few vertices and vertex-transitive group actions. *Berichte aus der Mathematik [Reports from Mathematics] Dissertation* Verlag Shaker, Technische Universität
- [71] Lutz F H 2008 Enumeration and random realization of triangulated surfaces *Discrete Differential Geometry (Oberwolfach Semin)* vol 38 (Birkhäuser) pp 235–53
- [72] Lutz F H, Mason J K, Lazar E A and MacPherson R D 2017 Roundness of grains in cellular microstructures *Phys. Rev. E* **96** 023001
- [73] Lutz F H and Nevo E 2016 Stellar theory for flag complexes *Math. Scand.* **118** 70–82

- [74] Lyckegaard A, Lauridsen E M, Ludwig W, Fonda R W and Poulsen H F 2011 On the use of Laguerre tessellations for representations of 3D grain structures *Adv. Eng. Mater.* **13** 165–70
- [75] Matzke E B 1946 The three-dimensional shape of bubbles in foam—an analysis of the role of surface forces in three-dimensional cell shape determination *Am. J. Bot.* **33** 58–80
- [76] McAllister M, Kirkpatrick D and Snoeyink J 1996 A compact piecewise-linear Voronoi diagram for convex sites in the plane *Discrete Comput. Geom.* **15** 73–105
- [77] McCollum F 2023 Power diagrams, Matlab central file exchange
- [78] Mettan X et al 2021 Acoustic-pressure-assisted engineering of aluminum foams *Adv. Eng. Mater.* **23** 2100306
- [79] Mills N 2007 *Polymer Foams Handbook: Engineering and Biomechanics Applications and Design Guide* (Butterworth Heinemann)
- [80] Moise E E 1952 Affine structures in 3-manifolds: V. the triangulation theorem and hauptvermutung *Ann. Math.* **56** 96–114
- [81] Mukherjee M, García-Moreno F, Jiménez C, Rack A and Banhart J 2017 Microporosity in aluminium foams *Acta Mater.* **131** 156–68
- [82] Mulzer W and Rote G 2008 Minimum-weight triangulation is NP-hard *J. ACM* **55** 1–29
- [83] Munkres J R 1984 *Elements of Algebraic Topology* (Addison-Wesley Publishing Company)
- [84] Neu T R, Heim K, Seeliger W, Kamm P H and García-Moreno F 2024 Aluminum foam sandwiches: a lighter future for car bodies *JOM* **76** 2619–30
- [85] Noack M A, Bülk F, Wang N, Banhart J and García-Moreno F 2021 Aluminium foam with sub-mm sized cells produced using a rotating gas injector *Mater. Sci. Eng. B* **273** 115427
- [86] Ohgaki T, Toda H, Kobayashi M, Uesugi K, Niinomi M, Akahori T, Kobayash T, Makii K and Aruga Y 2006 In situ observations of compressive behaviour of aluminium foams by local tomography using high-resolution x-rays *Phil. Mag.* **86** 4417–38
- [87] Plateau J A F 1873 *Statique expérimentale et théorique des Liquides Soumis aux Seules Forces moléculaires* (Gauthier-Villars)
- [88] Randolph Franklin W, de Magalhaes S V G and Landis E N 2024 *Fast 3D Euclidean Connected Components* 1st edn (Association for Computing Machinery) pp 17–28
- [89] Redenbach C 2009 Microstructure models for cellular materials *Comput. Mater. Sci.* **44** 1397–407
- [90] Ringel G 1955 Wie man die geschlossenen nichtorientierbaren Flächen in möglichst wenig Dreiecke zerlegen kann *Math. Ann.* **130** 317–26
- [91] Rubinstein J H, Segerman H and Tillmann S 2018 Traversing three-manifold triangulations and spines *L'Enseign. Math.* **65** 155–206
- [92] Scaringe R P 2015 Apollonius solutions in R^d *CoRR* (arXiv:1512.08846)
- [93] Schaller F M, Kapfer S C, Evans M E, Hoffmann M J, Aste T, Saadatfar M, Mecke K, Delaney G W and Schröder-Turk G E 2013 Set Voronoi diagrams of 3D assemblies of aspherical particles *Phil. Mag.* **93** 3993–4017
- [94] Schick C 2004 A mathematical analysis of foam films *Doctoralthesis* Technische Universität Kaiserslautern
- [95] Stanzick H, Wichmann M, Weise J, Helfen L, Baumbach T and Banhart J 2002 Process control in aluminum foam production using real-time x-ray radiography *Adv. Eng. Mater.* **4** 814–23
- [96] Steinitz E 1922 Polyeder und Raumeinteilungen *Encyklopädie der Mathematischen Wissenschaften mit Einschluss Ihrer Anwendungen, Dritter Band: Geometrie, Iii.1* vol 2, ed W F Meyer and H Mohrmann (B. G. Teubner) ch III A B 12, pp 1–139
- [97] Sullivan J M 1999 Foams and bubbles: geometry and simulation *Int. J. Shape Model.* **05** 101–14
- [98] Sullivan J M 1999 The geometry of bubbles and foams *Foams and Emulsions* (Springer) pp 379–402
- [99] Taylor J E 1976 The structure of singularities in soap-bubble-like and soap-film-like minimal surfaces *Ann. Math.* **103** 489–539
- [100] Walkup D W 1970 The lower bound conjecture for 3- and 4-manifolds *Acta Math.* **125** 75–107
- [101] Wang N, Chen X, Li Y, Liu Z, Zhao Z, Cheng Y, Liu Y and Zhang H 2017 The cell size reduction of aluminum foam with dynamic gas injection based on the improved foamable melt *Colloids Surf. A* **527** 123–31
- [102] Wang N, Noack M A, Kamm P H, Banhart J and García-Moreno F 2022 Short-range ordered aluminum foams *Adv. Eng. Mater.* **24** 2100795
- [103] Wang P, Yuan N, Ma Y, Xin S, He Y, Chen S, Xu J and Wang W 2020 Robust computation of 3D Apollonius diagrams *Comput. Graph. Forum* **39** 43–55

- [104] Wejrzanowski T, Skibinski J, Szumbariski J and Kurzydłowski K 2013 Structure of foams modeled by Laguerre-Voronoi tessellations *Comput. Mater. Sci.* **67** 216–21
- [105] Yan D-M, Wang W, Lévy B and Liu Y 2010 Efficient computation of 3D clipped Voronoi diagram *Advances in Geometric Modeling and Processing* ed B Mourrain, S Schaefer and G Xu (Springer) pp 269–82
- [106] Zhao X, Wang H and Komura T 2014 Indexing 3D scenes using the interaction bisector surface *ACM Trans. Graph.* **33** 1–14