Tutte Embeddings of Tetrahedral Meshes

Marc Alexa[1] ⓘ

Abstract

Tutte's embedding theorem states that every 3-connected graph without a K_5- or $K_{3,3}$-minor (i.e., a *planar* graph) is embedded in the plane if the outer face is in convex position and the interior vertices are convex combinations of their neighbors. We show that this result extends to simply connected tetrahedral meshes in a natural way: for the tetrahedral mesh to be embedded if the outer polyhedron is in convex position and the interior vertices are convex combination of their neighbors it is sufficient (but not necessary) that the graph of the tetrahedral mesh contains no K_6 and no $K_{3,3,1}$, and all triangles incident on three boundary vertices are boundary triangles.

1 Introduction

Every planar graph has a straight line embedding in the plane [5]. Tutte's celebrated embedding theorem [17] provides a constructive proof for 3-connected planar graphs: if the outer polygon of the graph is in convex position and the interior vertices are convex combinations of their neighbors then the realization is an embedding (and every face is convex). This means, we can compute an embedding by placing the vertices of the boundary polygon so that it is convex and then solving a linear system for the positions of the interior vertices. This procedure has become an invaluable tool in computer graphics, geometry processing, and CAGD, where it is used to construct

---

Marc Alexa
marc.alexa@tu-berlin.de

[1] Faculty of Computer Science & Electrical Engineering, TU Berlin, Sekretariat MAR 6-6, Marchstr. 6, 10587 Berlin, Germany

mappings for triangulations [6]. It has been generalized to other surface geometries and topologies [1, 2, 9].

A particularly useful generalization of Tutte embeddings would be the three-dimensional case and, in particular, tetrahedral meshes. Embedding a given tetrahedral mesh with fixed convex boundary would enable bijective piecewise linear mappings between different domains represented by the same boundary triangulation. Alas, Tutte's result fails to generalize to 3D in a simple way for arbitrary polyhedral complexes [4]. It has been observed to fail even for well behaved tetrahedral meshes in practice [3, Fig. 2]. We have conducted simple experiments on small tetrahedral meshes that suggests the convex combination approach will fail for "most" randomly chosen convex combination weights. Note that in most practical case one starts with an embedded tetrahedral mesh. If the mesh is already realized with convex boundary, there clearly exists a convex combination map that would generate the realization (just express every interior vertex as the convex combination of its neighbors). Our experiments suggest that the chance of finding such weights by random sampling is small.

The smallest possible counterexample is composed of two vertices inside a tetrahedron, with the two vertices connected to each other and all four boundary vertices [7]. The graph of this tetrahedral mesh is the complete graph on six vertices K_6. In view of the fact that Tutte's original work showed that 3-connected graphs may be embedded in the plane if (and only if) they have no K_5 and $K_{3,3}$ as a minor, one might ask if the observation that a K_6 is a counterexample for the tetrahedral case has any meaning. Note that K_6 is a forbidden minor in the class of *linklessly embeddable* graphs, which have been termed the natural three-dimensional analogue of planar graphs [14]. An important difference to the planer case, of course, is that a tetrahedral mesh with vertices and edges forming a K_6 *can* be embedded, meaning that there exist some convex combination weights that lead to an embedding (unlike in the case of non-planar graphs). We may still ask if tetrahedral meshes that are linklessly embeddable (concretely, have no K_6 and $K_{3,3,1}$ as a minor, see Sect. 2) guarantee Tutte embeddings. Indeed, this turns out to be the case.

Apart from excluded minors, however, there is one more obstruction. It has long been observed in the continuous domain that the harmonic extension from a boundary homeomorphism fails to be injective in 3D [12]. In fact, arbitrarily small perturbations of the identity map on the boundary are sufficient to cause the loss of injectivity [10]. For Tutte's embedding this has been elucidated by Floater[1] as follows:

> [...] there are very many ways to map, 1-1, one convex boundary into another (imagine creating lots of twists and turns). Won't that create foldover in the 3D embedding inside? So I guess the boundary mapping needs to be restricted.

We find that this problem materializes if an interior triangle in the tetrahedral mesh is fixed on the boundary, but the remaining vertices of its incident tetrahedra are not. Then "pulling" boundary vertices allows moving the interior vertices on either side of the fixed triangle. Interestingly, some proofs of Tutte's theorem require no chords in the graph. In 2D, this restriction can be lifted by observing that chords simply divide

---

[1] The quote is from private email exchange on the topic of tetrahedral Tutte embeddings.

the outer polygon into two smaller polygons. In 3D, however, a triangle on boundary vertices is generally not dividing the mesh into two components.

As we will show, the two restrictions mentioned above are sufficient to prove a version of Tutte's embedding theorem for tetrahedral meshes. The result is as follows:

Theorem 1.1 *Given a tetrahedral mesh with the following properties:*

- *The boundary is simply connected.*
- *A triangle incident on three boundary vertices is on the boundary.*
- *The graph is 4-connected.*
- *The graph has no K_6 or $K_{3,3,1}$ as a minor.*

If the vertex positions of the mesh are realized in $\mathbb{R}^3$ so that (a) *the boundary triangles form a strictly convex polyhedron and* (b) *each interior vertex is a strictly convex combination of its neighbors then the mesh is embedded.*

We hasten to point out that this result has little direct consequence on the practice of using the commonly generated tetrahedral meshes for creating PL mappings: almost all of them have a K_6-minor and, consequently, a convex combination mapping will likely not be an embedding. We discuss possible practical consequences in Sect. 5.

2 Tetrahedral Meshes and Linkless Embeddings

A tetrahedral mesh T is a simplicial complex consisting of vertices, edges, triangular faces and tetrahedral cells. We assume the tetrahedral mesh is a *topological ball*, meaning it has a simply connected interior. The boundary forms the graph of a polyhedron, a planar 3-connected graph with triangular faces [16]. Boundary faces are incident on one cell, interior faces are incident on two cells. In the following we assume that interior triangles satisfy the assumption of Theorem 1.1, i.e., they are incident on at most two boundary vertices.

The *star* of a vertex v is formed by the simplices incident on v. The *link* is the boundary of the star. We make the following observation about the connectivity of links:

Lemma 2.1 *The link L_v of a vertex v is 3-connected.*

Proof If v is an interior vertex, L_v is 3-connected by Steinitz's theorem [16]. Let v be a boundary vertex and V_b be the boundary vertices connected to v. The link L_v is a planar triangulation with boundary V_b. Every edge in L_v forms a triangle with v. A chord would induce a triangle that is not on the boundary of T, but all its vertices are boundary vertices. So there is no chord. This implies that the triangulation L_v is 3-connected [11, 3.2]. □

Every interior vertex has degree at least 4. Boundary vertices with degree 3 create *ears*: tetrahedra with three faces on the boundary connected to the remaining tetrahedral mesh through a single face. Since we ask that there are no interior faces incident on three boundary vertices, the tetrahedral meshes we consider contain no ears, and also boundary vertices have degree at least 4.
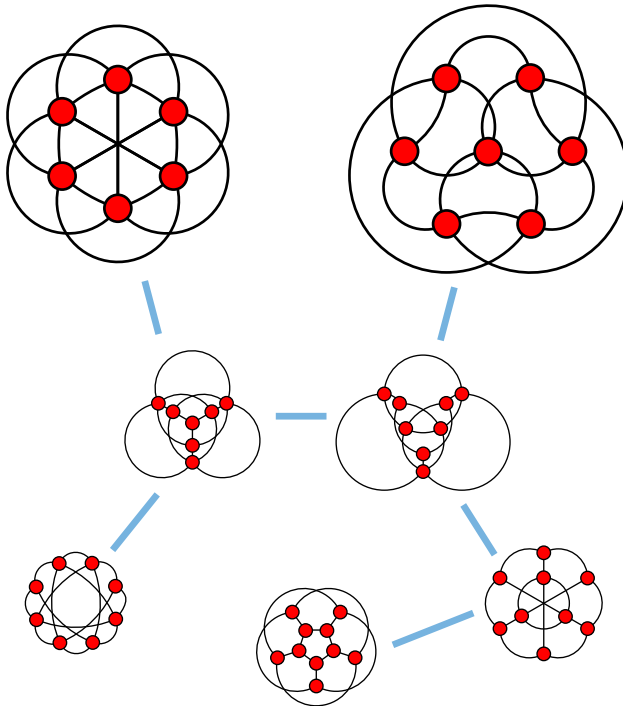
Fig. 1 The seven graphs of the Petersen family. The larger ones in the top row, K_6 and $K_{3,3,1}$, are necessary for the characterization of tetrahedral meshes that can be embedded by convex combination mappings. A graph is linklessly embeddable if it has no minor in this family. Drawing adopted from David Epstein (public domain)

By *tetrahedral graph* we mean the graph induced by the vertices and edges of T. Linkless embeddings of graphs, intuitively, are realizations of the graph in $\mathbb{R}^3$ so that there are no two cycles that are linked, i.e., that cannot be homotopically deformed so that the two cycles are topological disks. Sachs [14] showed that the Petersen graphs (see Fig. 1) are intrinsically linked and minimal. He suspected that all linklessly embeddable graphs can be characterized as those without a Petersen graph as a minor. This was eventually proved by Robertson et al. [13].

Of the seven graphs in the Peterson family, we only need the complete graph K_6 and, to show that certain degeneracies cannot occur, the complete tripartite graph $K_{3,3,1}$. This seems quite natural, as the other graphs have vertices with degree less than 4, so cannot serve as minimal counter-examples for tetrahedral Tutte embeddings.

The proof of the main statement is inspired by earlier proofs and their adaptations of Tutte [17], Geelen [8], Spielman [15], and the one focusing on PL mappings by Floater [6].

3 Realization by Convex Embedding, and Degeneracies

The graph T is *realized* by assigning coordinates to the vertices, i.e., $V \mapsto \mathbb{R}^{3 \times |V|}$. We denote the coordinate of vertex v as $\mathbf{x}(v) \in \mathbb{R}^3$. We assume the vertices on the boundary are realized so that the boundary faces form a convex polyhedron.

Every interior vertex is realized as a convex combination of its neighbors:

$$\mathbf{x}(v) = \sum_{(v,w) \in T} b_{(v,w)} \mathbf{x}(w), \qquad \sum_{(v,w) \in T} b_{(v,w)} = 1, \quad b_{(v,w)} > 0.$$

This implies that $\mathbf{x}(v)$ lies strictly in the interior of the convex hull of its neighbors. In the following we want to show that (a) the convex hull of the neighbors cannot degenerate, i.e., be contained in a common plane; and (b) that $\mathbf{x}(v)$ lies strictly inside the boundary polyhedron for interior vertices v.

We start by recalling that a plane through $\mathbf{q} \in \mathbb{R}^3$ with normal vector $\mathbf{n} \in \mathbb{R}^3$, $\mathbf{n}^\mathsf{T} \mathbf{n} = 1$, is defined as the set

$$P_{\mathbf{n},\mathbf{q}} = \{\mathbf{x} \in \mathbb{R}^3 : \mathbf{n}^\mathsf{T}(\mathbf{x} - \mathbf{q}) = 0\}.$$

The *positive open half-space* is the set of points on the side of the plane in the direction $\mathbf{n}$, i.e., $P_{\mathbf{n},\mathbf{q}}^+ = \{\mathbf{z} \in \mathbb{R}^3 : \mathbf{n}^\mathsf{T}(\mathbf{z} - \mathbf{q}) > 0\}$. The negative open half-space is defined analogously. We make the following simple observation for interior vertices:

Lemma 3.1 *Let v be an interior vertex. Consider the plane $P_{\mathbf{n},\mathbf{x}(v)}$. If v has a neighbor realized in $P_{\mathbf{n},\mathbf{x}(v)}^+$ then it has at least one neighbor in $P_{\mathbf{n},\mathbf{x}(v)}^-$ (and vice versa).*

Proof This follows by contradiction from $\mathbf{x}(v)$ being in the interior of the convex hull of its neighbors. □

If v is a boundary vertex, it is possible that that all neighbors are contained in $P_{\mathbf{n},\mathbf{x}(v)}$ and either $P_{\mathbf{n},\mathbf{x}(v)}^+$ or $P_{\mathbf{n},\mathbf{x}(v)}^-$. In this case we call v *extreme* along $\mathbf{n}$, because the convexity of the boundary polyhedron implies that there is no other vertex v' with $\mathbf{n}^\mathsf{T}(v') > \mathbf{n}^\mathsf{T}(v)$. We also observe that interior vertices are never extreme:

Lemma 3.2 *An interior vertex is not realized on the boundary.*

Proof Assume vertex v_i is realized on a boundary element b—this may be a vertex, edge, or face. Let V_i be the vertices realized in b that are connected to v_i by a path realized in b. All vertices in V_i are realized in a common plane $P_{\mathbf{n},\mathbf{v}_i}$ and $\mathbf{n}$ can be chosen so that $\mathbf{n}^\mathsf{T} \mathbf{v}_i \geq \mathbf{n}^\mathsf{T} \mathbf{v}$, $v \in T$, because the boundary polyhedron is convex. We may delete the at most three boundary vertices incident on b and V_i remains connected to a vertex v' outside V_i because we assume T to be 4-connected. The vertex v' is not realized in $P_{\mathbf{n},\mathbf{v}_i}$ and $\mathbf{n}$, so by Lemma 3.1 there exists a vertex in V_i with neighbors realized in both half-spaces $P_{\mathbf{n},\mathbf{x}(v)}^+$, which is a contradiction. □

We say that a path $v_0, v_1, \ldots$ is *non-decreasing* w.r.t. $\mathbf{n}$ if the realizations of the vertices along the path satisfy $\mathbf{n}^\mathsf{T}(v_0) \leq \mathbf{n}^\mathsf{T}(v_1) \leq \ldots$ We observe that any vertex is connected to an extreme boundary vertex by a non-decreasing path.

Lemma 3.3 *For any vertex v_0 and any direction $\mathbf{n}$ there is a non-decreasing path $v_0, v_1, \ldots, v_b$ satisfying $\mathbf{n}^\mathsf{T}(v_0) \leq \mathbf{n}^\mathsf{T}(v_0) \leq \ldots \leq \mathbf{n}^\mathsf{T}(v_b)$, with v_b on the boundary and extreme.*

Proof Consider vertex v_k along the path, which may be an interior or boundary vertex. If v_k has a neighbor v' satisfying $\mathbf{n}^\mathsf{T}\mathbf{x}(v') > \mathbf{n}^\mathsf{T}\mathbf{x}(v_k)$ set $v_{k+1} = v'$ and continue with v_{k+1}.

Now assume we have found a vertex v such that $\mathbf{n}^\mathsf{T}\mathbf{x}(v) \geq \mathbf{n}^\mathsf{T}\mathbf{x}(v')$ for all vertices v' adjacent to v. If v is on the boundary we are done. If v is interior, Lemma 3.1 forces $\mathbf{n}^\mathsf{T}\mathbf{x}(v) = \mathbf{n}^\mathsf{T}\mathbf{x}(v')$, meaning all neighbors of v lie in the plane $P_{\mathbf{n},\mathbf{x}(v)}$. Let V be the set of vertices realized in $P_{\mathbf{n},\mathbf{x}(v)}$ that are connected to v by a path realized in $P_{\mathbf{n},\mathbf{x}(v)}$. Because the boundary polyhedron is not flat there are vertices in T not in V. But T is connected so there must be a vertex $v' \in V$ connected to a vertex not in $P_{\mathbf{n},\mathbf{x}(v)}$. If v' is interior, Lemma 3.1 implies that v' has a neighbor v'' satisfying $\mathbf{n}^\mathsf{T}(v'') > \mathbf{n}^\mathsf{T}\mathbf{x}(v') = \mathbf{n}^\mathsf{T}\mathbf{x}(v_i)$. Because v is connected to v' through vertices contained in $P_{\mathbf{n},\mathbf{x}(v_i)}$ there is non-decreasing path through V to v'', and the process continues. This shows that the process has to end in a boundary vertex v_b, which is extreme because it satisfies $\mathbf{n}^\mathsf{T}\mathbf{x}(v_b) \geq \mathbf{n}^\mathsf{T}\mathbf{x}(v')$ for all neighbors v'. □

Corollary 3.4 *Let $P_{\mathbf{n},\mathbf{q}}^+$ be an open half-space and $V_{\mathbf{n},\mathbf{q}}^+$ be the vertices contained in it, i.e., $v \in V_{\mathbf{n},\mathbf{q}}^+$ implies $\mathbf{x}(v) \in P_{\mathbf{n},\mathbf{q}}^+$. The graph induced by $V_{\mathbf{n},\mathbf{q}}^+$ is connected.*

Proof Let $V_{\mathbf{n}}^b$ contain all boundary vertices v_b that maximize $\mathbf{n}^\mathsf{T}\mathbf{x}(v_b)$. This set consists of either a single vertex, or two vertices incident on a common edge, or three vertices incident on a common face. Note that $V_{\mathbf{n}}^b$ is connected. Any vertex in $V_{\mathbf{n},\mathbf{q}}^+$ is connected to $V_{\mathbf{n}}^b$ by a non-decreasing path, which is entirely contained in $P_{\mathbf{n},\mathbf{q}}^+$. □

We now show that the neighbors of a vertex cannot degenerate to a flat configuration if the graph is linklessly embeddable.

Lemma 3.5 *Let v be a vertex in a tetrahedral graph T. If v and all its neighbors are realized in a common plane $P_{\mathbf{n},\mathbf{v}}$ then T has $K_{3,3,1}$ as a minor.*

Proof Let V be the set of vertices contained in P that are connected to v by a path realized in $P_{\mathbf{n},\mathbf{x}(v)}$. We distinguish two types of vertices in V: vertices V_i whose neighbors are all in V (such as v); and vertices V_b, which have at least one neighbor not in V. Note that vertices in V_i cannot be boundary vertices.

Let V^+, resp. V^- be the vertices contained in $P_{\mathbf{n},\mathbf{x}(v)}^\pm$. Both are connected (Corollary 3.4). Since v is interior, Lemma 3.2 implies that none of the vertices in V_b is extreme along $\mathbf{n}$. So all vertices in V_b are connected to vertices in both V^+ as well as V^-: for interior vertices this follows from Lemma 3.1 and for boundary vertices from the realization of the boundary as a convex polyhedron.

Pick a vertex $v_b \in V_b$ that is connected to a vertex $v_i \in V_i$. Such pair must exist because V is connected. Consider the link L_{v_b} of v_b. It contains v_i as well as a vertices in $V^\pm$. The vertices in $V_i \cap L_{v_b}$ cannot be connected by an edge to the vertices in $(V^+ \cup V^-) \cap L_{v_b}$, so there must be a set of vertices $v_b^k \in$ in $V_b \cap L_{v_b}$ separating the two sets. This set contains at least three vertices, because L_{v_b} is 3-connected (Lemma 2.1).
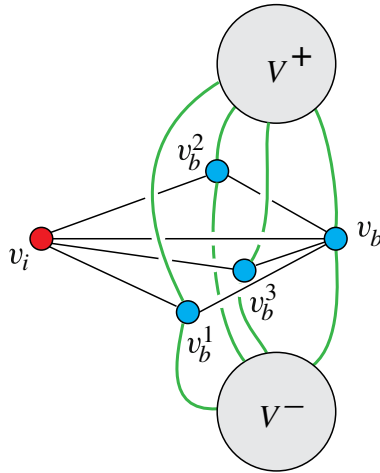
Fig. 2 Illustration for Lemma 3.5: Vertex v_i and all its neighbors are in a common plane. Vertices v_b^* are in the same plane but connected to vertices in both half-spaces. The vertices form a $K_{3,3,1}$

Now notice that V_i, V^+, V^- are connected to all $v_b^{1,2,3}$, and v_b is connected to V_i, V^+, V^- as well as $v_b^{1,2,3}$. This is a $K_{3,3,1}$ (see illustration in Fig. 2). □

By showing that the neighbors of a vertex are not co-planar, we have also shown that they are not co-linear or degenerate to a single point. Moreover, we have established that for any plane $P_{\mathbf{n},\mathbf{x}(v)}$, vertex v will have neighbors in both half-spaces $P_{\mathbf{n},\mathbf{x}(v)}^{\pm}$. This observation allows us to sharpen Lemma 3.3 to strictly increasing path $v_0, v_1, \ldots$ satisfying $\mathbf{n}^\mathsf{T}\mathbf{x}(v_i) < \mathbf{n}^\mathsf{T}\mathbf{x}(v_j)$ for $i < j$.

Corollary 3.6 *For any direction $\mathbf{n}$ every interior vertex v has a strictly increasing path $v = v_0, v_1, \ldots, v_b$ that ends in a boundary vertex v_b.*

Proof The proof is analogous to Lemma 3.3 except that Lemmas 3.1 and 3.5 now guarantee that every interior vertex v_k has a neighboring vertex v_{k+1} satisfying $\mathbf{n}^\mathsf{T}\mathbf{x}(v_{k+1}) > \mathbf{n}^\mathsf{T}\mathbf{x}(v_{k+1})$. □

4 Local and Global Injectivity

The fact that interior vertices are strictly in the interior of the boundary polyhedron (Lemma 3.2) establishes the injectivity for tetrahedra incident on boundary triangles:

Corollary 4.1 *Consider a boundary face f, its incident tetrahedron t and the vertex v in t not in f. All elements of t have positive (signed) measure.*

Proof Since v is strictly inside the boundary polyhedron, t has positive volume. If any of the triangles or edges had zero area or lengths, the volume of the tetrahedron would be zero, so they are all strictly positive. □
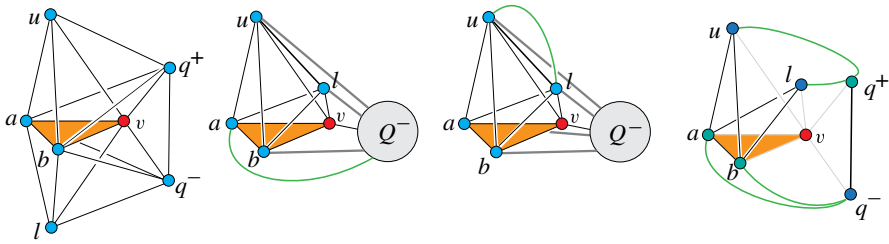
Fig. 3 Illustration for Lemma 4.3: The interior face $f = (a, b, v)$ (in orange) is incident on two tetrahedra spanned by vertices u and l. The leftmost illustration shows a proper realization. The other illustrations correspond the cases 1, 2, and 3 (from left to right)

Let us now consider two tetrahedra incident on an interior triangle. We want to show that if one of the tetrahedra is realized injectively then the other one must be as well. This local consistency will then imply global injectivity.

We need the following observation about *planar triangulations*, i.e., a planar graph whose faces all have degree 3, except possibly the outer face.

Proposition 4.2 *In a planar triangulation without chords, the graph induced by the interior vertices is connected. Contracting the interior vertices results in a triangulation of the boundary polygon with a single interior vertex connected to all boundary vertices.*

Proof We assume the planar triangulation has at least two interior vertices. Assume that a and b are interior vertices not connected by a path of interior edges. Since there are no chords, the triangulation is 3-connected [11, 3.2], so a and b are connected to three boundary vertices $v_b^k, k = 1, 2, 3$. Add vertex e outside the boundary and connect it to v_b^k with non-crossing paths. This creates a $K_{3,3}$, which is impossible, because the triangulation (including the paths to the exterior vertex) is planar.

Since the interior vertices are connected they can be contracted to a single vertex without altering the boundary. Each boundary vertex is connected to a least one interior vertex, because there are no chords, so there are no ears. □

Lemma 4.3 *Let f be an interior triangle with incident tetrahedra t_u, t_l and u, l the vertices in t_u, t_l not in f. Let $P_{\mathbf{n}, \mathbf{x}(v)}$ be the plane through f and assume $\mathbf{x}(u) \in P_{\mathbf{n}, \mathbf{x}(v)}^+$. If $\mathbf{x}(l) \notin P_{\mathbf{n}, \mathbf{x}(v)}^-$ then T contains K_6 or $K_{3,3,1}$ as a minor.*

Proof Assume $\mathbf{x}(l)$ is contained $P_{\mathbf{n}, \mathbf{x}(v)} \cup P_{\mathbf{n}, \mathbf{x}(v)}^+$. As f is an interior face it has at least one interior vertex. Let v be the interior vertex of f and a, b the other two. The link L_v of v is a polyhedral graph. It contains the triangles (a, b, u) and (l, b, a). The remainder of L_v is a triangulation bounded by the cycle (a, u, b, l). For illustrations see Fig. 3.

The vertex q^-: As v has neighbors realized in $P_{\mathbf{n}, \mathbf{x}(v)}^+$ it must have at least one neighbor in $P_{\mathbf{n}, \mathbf{x}(v)}^-$ (by Lemma 3.1). Every neighbor of v is in L_v, but none of a, u, b, l are realized in $P_{\mathbf{n}, \mathbf{x}(v)}^-$. So there must be at least one other vertex in L_v. Among the

vertices $q \in L_v$ realized in $P^-_{\mathbf{n},\mathbf{x}(v)}$ we pick one vertex and call it q^-. Note that all vertices $q \in L_v$ are connected to v.

The path C: If l is realized in $P^+_{\mathbf{n},\mathbf{x}(v)}$ there exists a path C from l to u realized in $P^+_{\mathbf{n},\mathbf{x}(v)}$ by Lemma 3.4. If l is realized in $P_{\mathbf{n},\mathbf{x}(v)}$ it has a neighbor realized in $P^+_{\mathbf{n},\mathbf{x}(v)}$ by Lemmas 3.5 and 3.1 (and because f is not on the boundary). In this case there exists a path C from l to u such that $C \setminus \{l\}$ is realized in $P^+_{\mathbf{n},c}$. We now consider three cases based on the intersection of C and L_v.

Case 1: Vertices u and l are connected by an edge. This creates the two cycles (a, l, u) and (b, u, l) in L_v. The triangulation of one them contains q^-. Assume this is b, u, l—the other case is analogous. The triangulation of (b, u, l) is connected by Proposition 4.2 and we contract it to Q^-, which is connected to u, l, and b. A path to a is constructed from decreasing paths starting in a and any vertex in Q^-. All vertices in this path, except for a are realized in $P^-_{\mathbf{n},\mathbf{x}(v)}$, so the path cannot contain any of b, u, l, or v. This means a, u, b, l together with v and Q^- form a K_6.

Case 2: The edge (u, l) is not present and $C \cap L_v = \{u, l\}$, i.e., the path C has no intersection with L_v except for the endpoints. In this case contract the interior vertices of the triangulation in the cycle (a, u, b, l). The triangulation has no chords, because (u, l) is not present, so it is connected (Proposition 4.2) and we contract it into Q^-. All boundary vertices a, u, b, l of the triangulation in the cycle are connected to Q^- (Proposition 4.2). As above, a, u, b, l together with v and Q^- form a K_6.

Case 3: The edge (u, l) is not present and C intersects L_v in at least one vertex q^+. The name q^+ indicates that it is realized in $P^+_{\mathbf{n},c}$ because $C \setminus l$ is. This implies that q^+ is distinct from q^-. Because $q^+ \in C$, it is connected to both u and l. We construct decreasing paths from a, b, and q^-, establishing a path from q^- to a and to b. Each of a and b has a neighbor in $P_{\mathbf{n},\mathbf{x}(v)}$, so a is not in the path from q^- to b and vice versa. Lastly, q^+ and q^- are connected in $L_v \setminus \{a, u, b, l\}$, because the triangulation inside (a, u, b, l) has no chords and is connected (Proposition 4.2). Now $\{u, l, q^-\}$ and $\{a, b, q^+\}$ form a $K_{3,3}$. Since all of the six vertices are in L_v, together with v they form a $K_{3,3,1}$. □

This establishes that if one of the two tetrahedra incident on an interior face is non-degenerate and correctly oriented, the other one is as well. Since the tetrahedra incident on the boundary are non-degenerate and correctly oriented (Corollary 4.1) and the dual graph of T is connected it follows that all tetrahedra are non-degenerate and correctly oriented.

The global injectivity (i.e., every point in the interior of the boundary is inside exactly one simplex) can be established by a homotopy argument (similar to the 2D case [8, 15]): for any point $\mathbf{q}$ in the interior of the boundary consider a half-line l, originating at $\mathbf{q}$, not intersecting any vertex or edge. Such line exists because the shadows of the vertices and edges on the sphere of directions consists of finitely many at most one-dimensional subsets. Start with any point $\mathbf{x}$ on l outside the boundary and move towards $\mathbf{q}$. The number of tetrahedra that contain $\mathbf{x}$ can only change if $\mathbf{x}$ crosses a face. As $\mathbf{x}$ crosses the boundary face, the number of tetrahedra containing $\mathbf{x}$ changes from 0 to 1. Note that no interior vertex, edge, or face can intersect any

boundary face (by Corollary 4.1). In the interior, the number of tetrahedra containing **x** stays constant as **x** crosses interior faces by Lemma 4.3. This is the case even if l intersects more than one face in the same point. However, the fact that the number of tetrahedra containing any interior point is always one rules out the possibility that two faces intersect, because in the vicinity of the intersection more than one tetrahedron would contain the point.

5 Discussion

Whereas both K_5 and $K_{3,3}$ cannot be embedded in the plane, a tetrahedral mesh with K_6 as the underlying graph may well be embedded in $\mathbb{R}^3$. Whether a realization based on convex combinations is embedded depends on the weights. We have not been able to create a similar situation for a tetrahedral mesh with only a $K_{3,3,1}$-minor (and no K_6-minor). Or, in other words, we were unable to generate an example that would demonstrate the necessity of excluding $K_{3,3,1}$-minors. In fact, every graph without a K_6-minor turned out to be embedded by a any convex combination map. This leads to the question whether excluding $K_{3,3,1}$ is necessary for the proof.

There are certain bounds on the number k-cliques in a graph without K_t-minors [18]. For the maximal number of 4-cliques in a graph on n vertices without a K_6-minor we find $4n - 15$. Since every tetrahedron is a 4-clique a tetrahedral mesh on n vertices with more than $4n - 15$ tetrahedra has to contain a K_6-minor. Consequently, only for tetrahedral meshes with few elements there is hope that convex combination mappings are guaranteed to work. This may be considered in contrast to the speculation by Chilakamarri et al. [4] that 'high' connectivity is necessary for an extension of Tutte's theorem. If this is interpreted as the graph containing many edges, we find that rather the opposite is the case for tetrahedral meshes.

Moreover, already the bound on the number of 4-cliques suggests that most tetrahedral meshes used in practice are guaranteed to contain a K_6 and, based on our and others' experiments, unlikely embedded by 'random' convex combination maps. One route for future investigations may be the reduction of the number of tetrahedra using bistellar flips, or the reduction of the graph by edge contractions.

References

1. Aigerman, N., Lipman, Y.: Orbifold Tutte embeddings. ACM Trans. Graph. **34**(6), # 190 (2015)
2. Aigerman, N., Lipman, Y.: Hyperbolic orbifold Tutte embeddings. ACM Trans. Graph. **35**(6), # 217 (2016)
3. Campen, M., Silva, C.T., Zorin, D.: Bijective maps from simplicial foliations. ACM Trans. Graph. **35**(4), # 74 (2016)
4. Chilakamarri, K., Dean, N., Littman, M.: Three-dimensional Tutte embedding. Congr. Numer. **107**, 129–140 (1995)
5. Fáry, I.: On straight line representation of planar graphs. Acta Univ. Szeged. Sect. Sci. Math. **11**, 229–233 (1948)
6. Floater, M.S.: One-to-one piecewise linear mappings over triangulations. Math. Comput. **72**(242), 685–696 (2003)
7. Floater, M.S., Pham-Trong, V.: Convex combination maps over triangulations, tilings, and tetrahedral meshes. Adv. Comput. Math. **25**(4), 347–356 (2006)
8. Geelen, J.: On "How to draw a graph". University of Waterloo (2012). https://www.math.uwaterloo.ca/~jfgeelen/Publications/tutte.pdf
9. Gortler, S.J., Gotsman, C., Thurston, D.: Discrete one-forms on meshes and applications to 3D mesh parameterization. Comput. Aided Geom. Des. **23**(2), 83–112 (2006)
10. Laugesen, R.S.: Injectivity can fail for higher-dimensional harmonic extensions. Complex Variables Theory Appl. **28**(4), 357–369 (1996)
11. Laumond, J.-P.: Connectivity of plane triangulations. Inform. Process. Lett. **34**(2), 87–96 (1990)
12. Melas, A.D.: An example of a harmonic map between Euclidean balls. Proc. Am. Math. Soc. **117**(3), 857–859 (1993)
13. Robertson, N., Seymour, P., Thomas, R.: Sachs' linkless embedding conjecture. J. Comb. Theory Ser. B **64**(2), 185–227 (1995)
14. Sachs, H.: On a spatial analogue of Kuratowski's theorem on planar graphs—an open problem. In: Graph Theory (Łagów 1981). Lecture Notes in Mathematics, vol. 1018, pp. 230–241. Springer, Berlin (1983)
15. Spielman, D.A.: Tutte's Theorem: How to draw a graph. Lecture notes for Spectral Graph Theory, Yale University (2018). https://www.cs.yale.edu/homes/spielman/561/lect15-18.pdf
16. Steinitz, E.: Polyeder und Raumteilungen. In: Encyklopädie der Mathematischen Wissenschaften, vol. 3-1-2, # 12 (1922)
17. Tutte, W.T.: How to draw a graph. Proc. Lond. Math. Soc. **13**, 743–767 (1963)
18. Wood, D.R.: Cliques in graphs excluding a complete graph minor. Electron. J. Comb. **23**(3), # P3.18 (2016)