Convex combination maps over triangulations, tilings, and tetrahedral meshes

Michael S. Floater and Valérie Pham-Trong

Abstract: In a recent paper by the first author, a simple proof was given of a result by Tutte on the validity of barycentric mappings, recast in terms of the injectivity of piecewise linear mappings over triangulations. In this note, we make a short extension to the proof to deal with arbitrary tilings. We also give a simple counterexample to show that convex combination mappings over tetrahedral meshes are not necessarily one-to-one.

AMS subject classification: Primary: 05C10, 05C85, Secondary: 65D17, 58E20.

Key words: triangulation, tiling, tetrahedral mesh, convex combination, discrete maximum principle, parameterization, planar embedding.

Short title: Convex combination maps.

1. Introduction

The idea of convex combination maps over triangulations has developed recently as a tool for parameterization of triangulations [**4**] and morphing of both triangulations and polygons [**7, 8**]. Such maps are piecewise linear and generalize Tutte's barycentric mappings; see [**6**] and [**11**].

It is important in practice to be able to guarantee the injectivity of the map and a sufficient condition is that the image of the triangulation boundary is convex. The first proof of this (for barycentric mappings) was given by Tutte [**11**] and takes an abstract graph-theoretic viewpoint. However, as pointed out recently by Colin de Verdière, Pocchiola, and Vegter, Tutte's theory is complicated even from the point of view of graph theory, and this motivated their simpler proof in [**2**]. Since then, an elementary proof which is independent of graph theory, was given in [**6**] for the case that the graph is a triangulation. The idea in [**6**] was to use the discrete maximum principle for convex combination functions in a similar way to the use of the continuous maximum principle for harmonic functions used by Kneser [**9**] in his proof of the Radó-Kneser-Choquet theorem [**10, 9, 1**] for harmonic mappings.

In this paper, we first show how the proof in [**6**] can be extended without too much difficulty to guarantee the validity of convex combination maps over arbitrary tilings (playing a similar role to the 3-connected graphs dealt with by Tutte). Note that these tilings include rectangular grids, which occur frequently in applications.

We also give a simple counterexample to show that convex combination mappings over tetrahedral meshes are not necessarily one-to-one.

2. Convex combination maps over tilings

By a face F we will understand the region enclosed by a convex polygon in the plane.

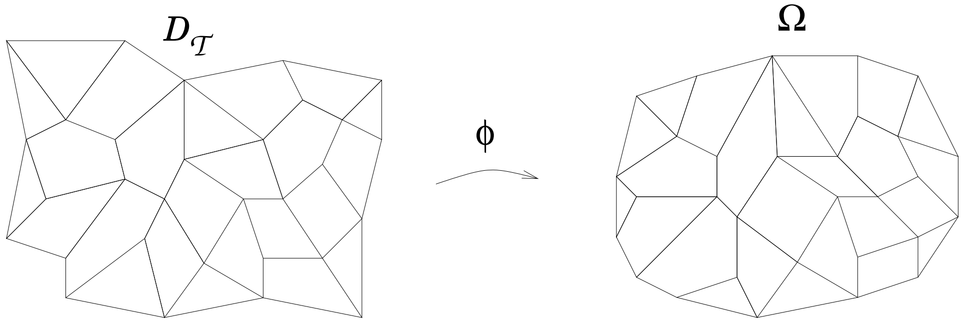

Figure 1. Convex combination map over a tiling

Definition 2.1. *Let $\mathcal{T}$ be a finite set of faces and let $D_{\mathcal{T}} = \bigcup_{T \in \mathcal{T}} T$. We will call $\mathcal{T}$ a* `tiling` *if*

(i) the intersection of any pair of faces is either empty, a common vertex, or a common edge, and
(ii) the edges in $\mathcal{T}$ belonging to only one face form a simple closed polygon $\partial D_{\mathcal{T}}$, the boundary of $D_{\mathcal{T}}$.

We denote by $V = V(\mathcal{T})$ and $E = E(\mathcal{T})$ the sets of vertices and edges in $\mathcal{T}$. We take all notation regarding paths from [**6**]. In the case that all faces are triangles, the tiling is a *triangulation* in the sense of [**6**].

We will call a pair of boundary vertices (v, w) of $\mathcal{T}$ a *dividing pair* if they share a face, but do not share a boundary edge. This generalizes the concept of the dividing edge used in [**6**]. For example, the tiling in Figure 2 contains one dividing pair, (v_1, v_2). Following [**6**], we will say that a tiling $\mathcal{T}$ is `strongly connected` if it contains no dividing pairs.

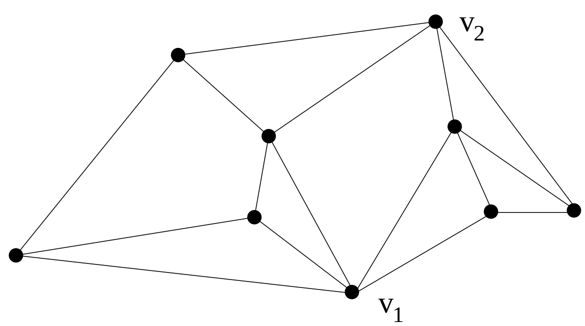

Figure 2. A dividing pair

In [**6**] we defined convex combination functions and mappings over the whole domain $D_{\mathcal{T}}$ by taking them to be linear in each triangle. To avoid unnecessary ambiguities when working with a general tiling $\mathcal{T}$, we will simply define such functions and mappings over *the vertices* V of $\mathcal{T}$. Let N_v denote the set of neighbours of a vertex v.

Definition 2.2. *Let $\phi : V \to \mathbb{R}$ be a function such that for every interior vertex v of $\mathcal{T}$, there exist weights $\lambda_{vw} > 0$, for $w \in N_v$, such that*

$$\sum_{w \in N_v} \lambda_{vw} = 1, \tag{2.1}$$

and

$$\phi(v) = \sum_{w \in N_v} \lambda_{vw} \phi(w). \tag{2.2}$$

Then we will call ϕ *a* `convex combination function`*.*

Similarly, we will understand a *convex combination mapping* to be any *mapping* $\phi : V \to \mathbb{R}^2$ satisfying the same two equations (2.1) and (2.2). The point $\phi(v)$ must then lie in the convex hull of the neighbouring points $\phi(w)$, $w \in N_v$. It is easy to show that such convex combination mappings satisfy a discrete maximum principle, which trivially extends the one for triangulations in [**6**]. Note that any mapping $\phi : V \to \mathbb{R}^2$ defines a unique piecewise linear mapping from the boundary $\partial D_{\mathcal{T}}$ to $\mathbb{R}^2$.

Suppose, furthermore, that we 'triangulate' our given tiling $\mathcal{T}$, i.e., we form a triangulation $\mathcal{T}'$ by adding edges which partition each face of $\mathcal{T}$ into triangles. Then any mapping $\phi : V \to \mathbb{R}^2$ uniquely extends to a continuous piecewise linear mapping $\phi' : D_{\mathcal{T}} \to \mathbb{R}^2$, linear over each triangle of $\mathcal{T}'$. Note that different triangulations of $\mathcal{T}$ will in general result in different extended mappings ϕ', though all will agree at the boundary $\partial D_{\mathcal{T}}$. Our main result is the following generalization of Theorem 4.1 of [**6**].

Theorem 2.3. *Suppose* $\mathcal{T}$ *is a strongly connected tiling and that* $\phi : V \to \mathbb{R}^2$ *is a convex combination mapping which maps* $\partial D_{\mathcal{T}}$ *homeomorphically into the boundary* $\partial\Omega$ *of some (closed) convex region* $\Omega \subset \mathbb{R}^2$. *If* $\mathcal{T}'$ *is any triangulation of* $\mathcal{T}$ *then the associated piecewise linear mapping* $\phi' : D_{\mathcal{T}} \to \mathbb{R}^2$ *is one-to-one.*

All the steps of the proof are completely analogous to those in [**6**] except the crucial part, namely Lemma 4.4 of [**6**]. This we now prove in the more general setting of tilings.

Lemma 2.4. *Let* $\mathcal{T}'$ *and* ϕ' *be as in Theorem 2.3 and suppose* $T_1 \cup T_2$ *is a quadrilateral in* $\mathcal{T}'$. *If* $\phi'|_{T_1}$ *is one-to-one, then* $\phi'|_{T_1 \cup T_2}$ *is one-to-one.*

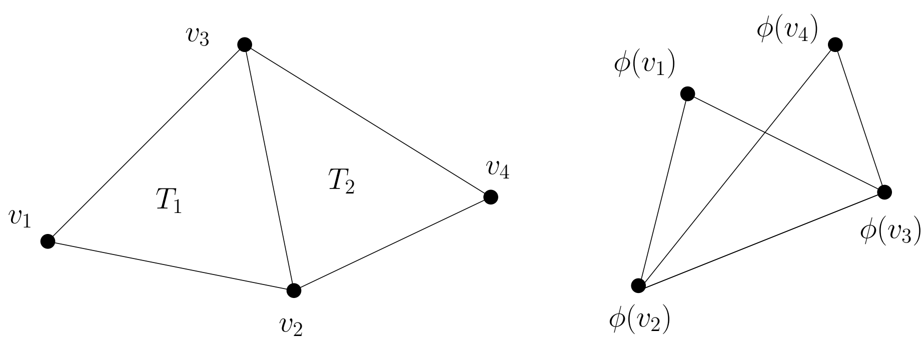

Figure 3. Mapping a quadrilateral

Proof: Suppose the triangles are $T_1 = [v_1, v_2, v_3]$ and $T_2 = [v_2, v_3, v_4]$ with common edge $[v_2, v_3]$, as in Figure 3. Since $\phi(T_1)$ is non-degenerate, the point $\phi(v_1)$ lies on one side of the infinite straight line L passing through the points $\phi(v_2)$ and $\phi(v_3)$. Our task is to show that the point $\phi(v_4)$ lies on the side of L opposite to $\phi(v_1)$. Let $ax_1 + bx_2 + c = 0$ be the equation of L and define the function $h : V \to \mathbb{R}$ as

$$h(v) = a\phi_1(v) + b\phi_2(v) + c.$$

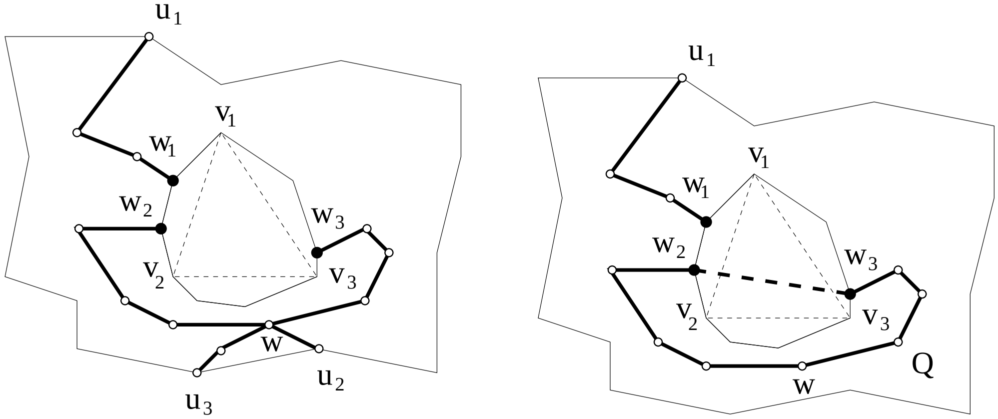

Figure 4. Rising and falling paths, case (i)

Then h is a convex combination function satisfying $h(v_2) = h(v_3) = 0$ and we may assume without loss of generality that $h(v_1) > 0$, and the task is now to show that $h(v_4) < 0$.

We begin by observing that the triangle T_1 is contained in some face F of $\mathcal{T}$. The boundary of the face F is a closed path which we will call R. Notice that the three vertices v_1, v_2, v_3 partition the path R into three open paths: R_{12} connecting v_1 to v_2; R_{23} connecting v_2 to v_3; and R_{31} connecting v_3 to v_1. Starting from v_1, let w_1 be the first vertex of R_{12} which either (a) is a boundary vertex or (b) has a different h value to the next vertex in R_{12}: see Figure 4. Clearly we will either have $w_1 = v_1$ or w_1 will be some vertex of R_{12} between v_1 and v_2. We will never have $w_1 = v_2$ since $h(v_1) \neq h(v_2)$. Note that in any case $h(w_1) = h(v_1) > 0$. Next we construct, in exactly the same manner as in Lemma 4.4. of [**6**], a rising path P_1 from w_1 to some boundary vertex u_1 of $\mathcal{T}$. Since h is strictly increasing along P_1, we have that $h(u_1) > 0$. In the case that w_1 is itself a boundary vertex, we have $u_1 = w_1$ and P_1 is the null path w_1.

In an analogous manner, starting from v_2, let w_2 be the first vertex of R_{12} which either (a) is a boundary vertex or (b) has a different h value to the next vertex in R_{12}. Note that $h(w_2) = h(v_2) = 0$. Similar to the construction of the rising path P_1 we construct a falling path from w_2 to a boundary vertex u_2. Along this falling path, h is strictly decreasing, and in particular $h(u_2) < 0$ if w_2 is an interior vertex and $h(u_2) = 0$ otherwise.

Similarly, starting from v_3, let w_3 be the first vertex of R_{13} which either (a) is a boundary vertex or (b) has a different h value to the next vertex in R_{13}. Let P_3 be a falling path starting at w_3 and ending at a boundary vertex u_3.

Note that if either falling path P_2 or P_3 passes through the vertex v_4 then we will immediately have $h(v_4) < 0$ and so we are finished. So we will assume from now on that v_4 does not belong to P_2 or P_3.

Next let $\hat{\mathcal{T}}$ be the tiling formed from $\mathcal{T}$ by splitting the face F into two faces by adding the segment $e = [w_2, w_3]$, so that e becomes an edge of $\hat{\mathcal{T}}$.

Next observe that the rising path P_1 is clearly distinct to both the falling paths P_2 and P_3. However, it is possible that P_2 and P_3 meet at some common vertex. We will therefore consider the two cases (i) P_2 and P_3 meet at some common vertex (see Figure 4) and (ii) P_2 and P_3 never intersect (see Figure 5).

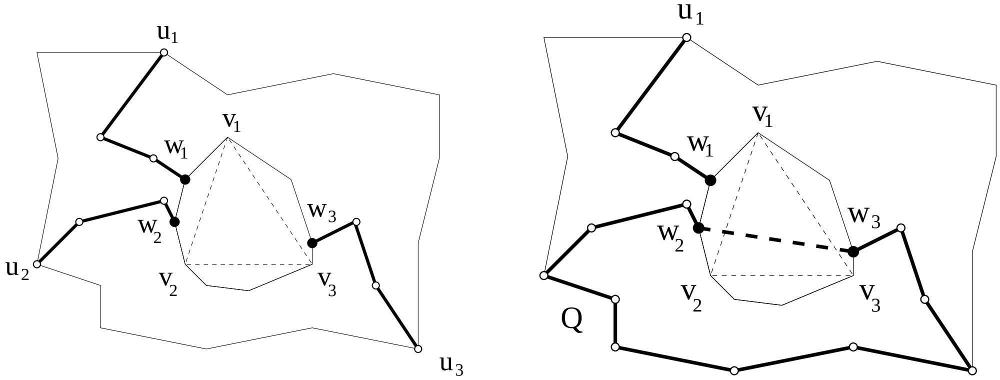

Figure 5. Rising and falling paths, case (ii)

Case (i)

Let w be the first vertex P_2 and P_3 have in common, starting from w_2 and w_3 respectively. Then the sub-path Q_2 of P_2 from w_2 to w and the sub-path Q_3 of P_3 from w_3 to w, together with the edge $[w_2, w_3]$ form a closed path Q in $\hat{\mathcal{T}}$; see Figure 4. The faces inside Q form a subtiling $\tilde{\mathcal{T}}$ of $\hat{\mathcal{T}}$. We have established that the closed path Q does not pass through either v_1 or v_4. Moreover either v_1 is inside Q and v_4 outside or vice versa. For if the triangle T_2 is contained in the face F as in Figure 6a, then the line segment $[v_1, v_4]$ clearly crosses Q precisely once, as in Figure 6b. Otherwise, T_2 belongs to a different face to T_1 and v_2 and v_3 are neighbours, as in Figure 6c, and in this case the curve consisting of the two line segments $[v_1, v_2]$ and $[v_2, v_4]$ clearly crosses Q precisely once; see Figure 6d.

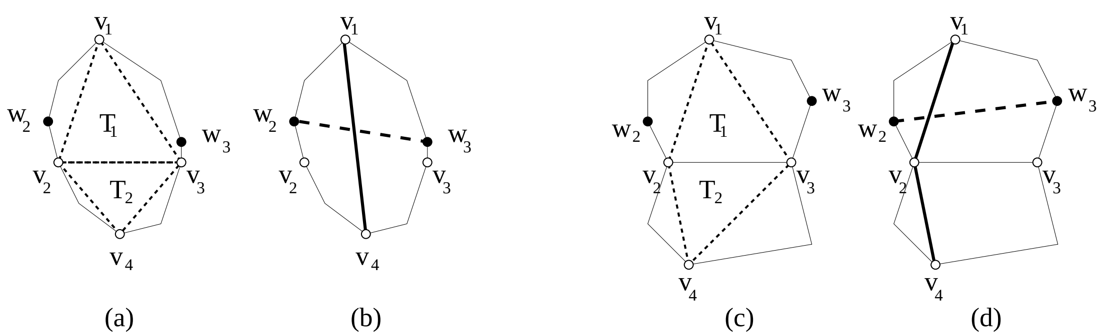

Figure 6. Faces containing T_1 and T_2

However, the path Q cannot enclose v_1, for then Q would have to cross P_1. Therefore Q must enclose v_4.

Now, by a simple generalization of lemma 2.2 of [**6**], we can stablish that v_4 must be connected in $\tilde{\mathcal{T}}$ to at least three vertices of its boundary Q. Thus v_4 must be connected in $\tilde{\mathcal{T}}$ to some vertex v in Q other than w_2 and w_3. Due to the construction of the falling paths, we have $h(v) < 0$, and then by the Discrete Maximum Principle for h in $\tilde{\mathcal{T}}$ we deduce that $h(v_4) < 0$.

Case (ii)

In this case all three paths P_1, P_2, P_3 are distinct, and therefore u_1, u_2, u_3 are distinct and we have $h(u_1) > 0$, $h(u_2) \leq 0$, and $h(u_3) \leq 0$. Moreover, due to the assumption that $\mathcal{T}$ is strongly connected, the vertex pair (w_2, w_3) cannot be a dividing pair, which means that at least one of its endpoints, say w_2, is an interior vertex. Thus we can assume that $u_2 \neq w_2$ and $h(u_2) < 0$.

Now let Q be the closed path in $\hat{\mathcal{T}}$ consisting of the paths P_2 and P_3, the edge $[w_2, w_3]$, and the part of the boundary of $\mathcal{T}$ connecting u_2 and u_3 which does not contain u_1; see Figure 5. As in case (i), the faces inside Q form a subtiling $\tilde{\mathcal{T}}$ of $\hat{\mathcal{T}}$. Due to the convexity of the image of $\partial D_{\mathcal{T}}$, and since $h(u_2) < 0$ and $h(u_3) \leq 0$, we have $h(v) < 0$ for all boundary vertices v of $\mathcal{T}$ in Q, other than u_2 and u_3. Therefore, $h(v) < 0$ for every vertex v in Q other than w_2 and w_3. Since the vertex v_4 is enclosed by Q, the Discrete Maximum Principle for h in $\tilde{\mathcal{T}}$ implies that $h(v_4) < 0$. ∎

3. Convex combination maps over tetrahedral meshes

In this second part of the paper we will show through a counterexample that convex combination maps over volumes partitioned into tetrahedra are not necessarily one-to-one even when the polyhedral boundary of the volume is mapped to a convex polyhedron. We note that Colin de Verdière, Pocchiola, and Vegter [**2**] have also constructed a counterexample to the generalization of Tutte's embedding theorem in $\mathbb{R}^3$. However, their approach is graph-theoretic and appears to us to be much more complicated. Generalizations of Tutte's theory to $\mathbb{R}^3$ have also been studied in [**3**].

By a tetrahedron we mean the convex hull $T = [v_1, v_2, v_3, v_4]$, of four non-coplanar points v_1, v_2, v_3, v_4 in $\mathbb{R}^3$.

Definition 3.1. *Let $\mathcal{T}$ be a finite set of tetrahedra and let $D_{\mathcal{T}} = \bigcup_{T \in \mathcal{T}} T$. We will call $\mathcal{T}$ a* `tetrahedral mesh` *if*

(i) the intersection of any pair of tetrahedra is either empty, a common vertex, a common edge, or a common face, and
(ii) the faces in $\mathcal{T}$ belonging to only one tetrahedron form a simple polyhedron $\partial D_{\mathcal{T}}$, homeomorphic to a sphere.

Similar to tilings we let $V = V(\mathcal{T})$, $E = E(\mathcal{T})$, and $F = F(\mathcal{T})$ be the sets of vertices / edges / faces in $\mathcal{T}$ and we call vertices / edges / faces contained in $\partial D_{\mathcal{T}}$ *boundary* vertices / edges / faces and otherwise *interior* ones.

Similar to Definition 2.2 we call $\phi : V \to \mathbb{R}^3$ a *convex combination map* if it satisfies equations (2.1) and (2.2), for all interior vertices v. Clearly ϕ has a unique extension to a continuous piecewise linear mapping $D_{\mathcal{T}} \to \mathbb{R}^3$ (i.e. linear over each tetrahedron), which we can also call ϕ.

To describe the counterexample, we first specify the tetrahedral mesh $\mathcal{T}$, and second describe the convex combination map. The tetrahedral mesh $\mathcal{T}$ has six vertices and eight tetrahedra $\mathcal{T} = \{T_1, \ldots, T_8\}$.
We begin by letting v_1, v_2, v_3, v_4 be any non-coplanar points in $\mathbb{R}^3$; see Figure 7. They form a tetrahedron $T_{1234} = [v_1, v_2, v_3, v_4]$ which forms the volume $D_{\mathcal{T}}$. We then choose

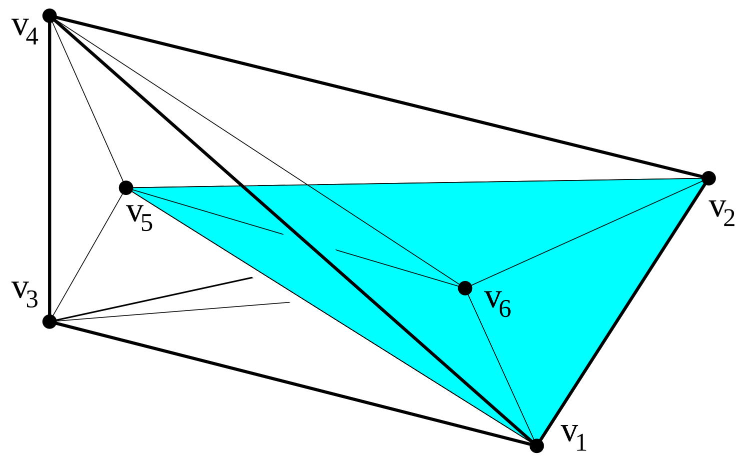

Figure 7.

any arbitrary point v_5 in the interior of T_{1234}. We can represent v_5 uniquely in barycentric form as

$$v_5 = \sum_{i=1}^{4} \lambda_i v_i, \qquad \sum_{i=1}^{4} \lambda_i = 1, \tag{3.1}$$

where $\lambda_i > 0$, $i = 1, \ldots, 4$. Notice that v_5 naturally defines a partition of T_{1234} into four tetrahedra;

$$T_{1234} = T_{5234} \cup T_{1534} \cup T_{1254} \cup T_{1235}.$$

Two of these tetrahedra do not contain F_{125}, the shaded face in Figure 7, and we let

$$T_1 := T_{5234}, \qquad T_2 := T_{1534}.$$

We split each of the other two tetrahedra, T_{1254} and T_{1235}, into three by choosing any point v_6 in the interior of the triangle F_{125}, and we set

$$T_3 = T_{6254}, \qquad T_4 = T_{1654}, \qquad T_5 = T_{1264},$$

and

$$T_6 = T_{6235}, \qquad T_7 = T_{1635}, \qquad T_8 = T_{1236},$$

Clearly, $\mathcal{T} := \{T_1, \ldots, T_8\}$ is indeed a tetrahedral partition of T_{1234}, for we have

$$\begin{aligned} T_{1234} &= T_1 \cup T_2 \cup T_{1254} \cup T_{1235} \\ &= T_1 \cup T_2 \cup (T_3 \cup T_4 \cup T_5) \cup (T_6 \cup T_7 \cup T_8), \end{aligned}$$

and the only intersections, if any, among $T_1, \ldots, T_8$ are common faces, edges, or vertices. We can clearly represent v_6 uniquely in the barycentric form

$$v_6 = \mu_1 v_1 + \mu_2 v_2 + \mu_5 v_5, \qquad \mu_1 + \mu_2 + \mu_5 = 1, \tag{3.2}$$

with $\mu_1, \mu_2, \mu_5 > 0$.

Next, we re-express v_6 as a 'strict' convex combination of all its neighbours $v_1, \ldots, v_5$, i.e., a combination in which all five coefficients are strictly positive. For example, we can express v_6 as

$$\begin{aligned} v_6 &= \mu_1 v_1 + \mu_2 v_2 + \mu_5 \left(\frac{1}{2} v_5 + \frac{1}{2} \sum_{i=1}^{4} \lambda_i v_i \right) \\ &= (\mu_1 + \mu_5 \lambda_1/2) v_1 + (\mu_2 + \mu_5 \lambda_2/2) v_2 \\ &\quad + (\mu_5 \lambda_3/2) v_3 + (\mu_5 \lambda_4/2) v_4 + (\mu_5/2) v_5 = \sum_{i=1}^{5} \hat{\mu}_i v_i. \end{aligned} \tag{3.3}$$

Similarly, for any α, we can express v_5 as

$$\begin{aligned} v_5 &= \alpha \sum_{i=1}^{4} \lambda_i v_i + (1-\alpha) \frac{v_6 - \mu_1 v_1 - \mu_2 v_2}{\mu_5} \\ &= (\alpha \lambda_1 - (1-\alpha)\mu_1/\mu_5) v_1 + (\alpha \lambda_2 - (1-\alpha)\mu_2/\mu_5) v_2 \\ &\quad + \alpha \lambda_3 v_3 + \alpha \lambda_4 v_4 + ((1-\alpha)/\mu_5) v_6 = \sum_{i=1}^{4} \hat{\lambda}_i v_i + \hat{\lambda}_6 v_6, \end{aligned} \tag{3.4}$$

and all the coefficients $\hat{\lambda}_i$ are strictly positive if we take $0 < \alpha < 1$ and α close enough to 1.

We now construct a convex combination map $\phi : D_{\mathcal{T}} \to D_{\mathcal{T}}$ which is not one-to-one. We first choose the images of the boundary vertices to be the vertices themselves, i.e., we let $\phi(v_i) = v_i$, for $i = 1, \ldots, 4$. We then construct ϕ in such a way that

$$\phi(v_5) = v_6, \qquad \text{and} \qquad \phi(v_6) = v_5,$$

so that v_5 and v_6 'swap places'. Due to equations (3.3) and (3.4), these two equations are clearly satisfied by the map defined by the two convex combinations

$$\begin{aligned} \phi(v_5) &= \sum_{i=1}^{4} \hat{\mu}_i \phi(v_i) + \hat{\mu}_5 \phi(v_6), \\ \phi(v_6) &= \sum_{i=1}^{4} \hat{\lambda}_i \phi(v_i) + \hat{\lambda}_6 \phi(v_5). \end{aligned}$$

It is easy to see that $\phi : D_{\mathcal{T}} \to D_{\mathcal{T}}$ is not injective. For example, the image under ϕ of the tetrahedron $T_3 = T_{6254}$ is clearly T_3 itself (though with the opposite orientation). Meanwhile the image of the tetrahedron $T_4 = T_{1264}$ is clearly

$$\phi(T_4) = T_3 \cup T_4 \cup T_5,$$

which therefore *contains* $\phi(T_3)$. Thus there are at least two points (one in T_3 and one in T_4) which both map to any given point in the interior of T_3.

References

1. G. Choquet, Sur un type de transformation analytique généralisant la représentation conforme et définie au moyen de fonctions harmoniques, Bull. Sci. Math. **69** (1945), 156–165.
2. E. Colin de Verdière, M. Pocchiola, and G. Vegter, Tutte's barycenter method applied to isotopies, in *Electronic proceedings of the thirteenth Canadian Conference on Computational Geometry* (2001), (http://compgeo.math.uwaterloo.ca).
3. C. Ó. Dúnlaing, personal communication, April, 2002.
4. M. S. Floater, Parametrization and smooth approximation of surface triangulations, Comp. Aided Geom. Design **14** (1997), 231–250.
5. M. S. Floater, Parametric tilings and scattered data approximation, International Journal of Shape Modeling **4** (1998), 165–182.
6. M. S. Floater, One-to-one piecewise linear mappings over triangulations, Math. Comp. **72** (2003), 685–696.
7. M. S. Floater and C. Gotsman, How to morph tilings injectively, J. Comp. Appl. Math. **101** (1999), 117–129.
8. C. Gotsman and V. Surazhsky, Guaranteed intersection-free polygon morphing, Computers and Graphics **25-1** (2001), 67–75.
9. H. Kneser, Lösung der Aufgabe 41, Jahresber. Deutsch. Math.-Verein. **35**, (1926), 123–124.
10. T. Radó, Aufgabe 41, Jahresber. Deutsch. Math.-Verein. **35**, (1926), 49.
11. W. T. Tutte, How to draw a graph, Proc. London Math. Soc. **13** (1963), 743–768.

Michael S. Floater
Mathematics Institute
University of Oslo
P.O. Box 1053 Blindern
0316 Oslo, NORWAY
`michaelf@ifi.uio.no`

Valérie Pham-Trong
Masters EPITA
14–16 rue Voltaire
94276 Le Kremlin-Bicetre
FRANCE
`valerie.phamtrong@free.fr`