APPROXIMATION OF SMOOTH SURFACES BY POLYHEDRAL SURFACES WITH HIDDEN VERTICES

UROP+ FINAL PAPER, SUMMER 2016.
CARLOS CORTEZ
MENTOR : ALEKSANDR BERDNIKOV

Project suggested by: Aleksandr Berdnikov/Emmy Murphy
September 1, 2016

Abstract. In $\mathbb{R}^3$, one can construct polyhedral surfaces such that, for some point p not on the surface, none of the vertices of the surface is visible from p. For a given compact surface K and a point we study the relation between the set of partial-linear embeddings of K with hidden vertices into $\mathbb{R}^3$ and the set of embeddings of K whose image is a smooth manifold with at least one principal curvature pointing towards p at each point that is visible from p. In particular, we establish results that suggest that elements of any of these spaces can be C^0-approximated by elements of the other one. Considering that these approximations may be extended to work parametrically, we conjecture that the space of partial linear mappings with vertices hidden from a point p is weakly homotopy equivalent to the space of mappings whose image has at least one principal curvature pointing towards p at all visible points, endowed with the C^0-topology.

1. Introduction

1.1. Motivation and results.

In $\mathbb{R}^2$, given a point p and a polygonal curve, one can always find a *visible* vertex, meaning, the segment between the vertex and p does not intersect the curve elsewhere. Perhaps surprisingly, this is no longer true when considering polyhedral surfaces in $\mathbb{R}^3$[1].

In the spirit of the h-principle, one may expect to find a natural obstruction to smooth surfaces being pointwise-approximated by polyhedral surfaces with hidden vertices and establish necessary and sufficient conditions for this to be possible. In section 2, we show that such an obstruction is given by having two principal curvatures pointing away from p at a point with a neighborhood all visible from p. We also establish that when this is not the case, the smooth surface can be at least locally approximated by polyhedral surfaces with hidden vertices.

In section 3, we work in the opposite direction and given a polyhedral surface with hidden vertices, we present a way to approximate it by a smooth surface with at least one principal curvature pointing towards the observer at each point.

In section 4, we sketch a generalization of the construction from section 3 establishing the approximation in the parametric case. We also speculate on the strategy one is likely to follow to prove that the constructed approximation establishes a weak homotopy equivalence between the space of polyhedral surfaces in $\mathbb{R}^3$ with hidden vertices and the considered class of smooth surfaces in $\mathbb{R}^3$.

1.2. Preliminary definitions.

We begin with a brief discussion of the objects we will be studying. Often we will assume without loss of generality that the point p appearing below is the origin.

Date: September 1, 2016.

We consider only compact smooth orientable 2-manifolds and smooth mappings of them into $\mathbb{R}^3$. We always choose the coorientation of a surface S so that p lies on the region of $\mathbb{R}^3 \setminus S$ to which the negative normal vector points.

Definition 1.1. *For a manifold K and an embedding $f : K \to \mathbb{R}^3$ of it into $\mathbb{R}^3$, we say a point $x \in \mathbb{R}^3$ is* visible *from p if the segment between p and the image of x lies completely in the closure of one of the two components of $\mathbb{R}^3 \setminus f(K)$. Otherwise, we say x is* hidden *from p. We say p is well-visible if there is a neighborhood U around it all of whose points are visible.*

Definition 1.2. *Consider a PL-manifold K. We say that a partially linear embedding $g : K \to \mathbb{R}^3$ is* shy *if there exists a point $p \in \mathbb{R}^3$ such that all 0-simplices v of the image $\mathcal{P} = g(K)$ are hidden from p. When we need to specify the point p, we say that g (or its image $\mathcal{P}$) is* shy with respect to *p.*

Definition 1.3. *For a smooth connected compact surface S and a point p, we denote by S_p the subset of points in S that are visible from p. We denote by S^- the subset of points x of S_p that have one or two negative principal curvatures when we choose the co-orientation of a chart around x to be such that p is on the negative side of the surface (we will always use this co-orientation). We say that a surface S is* visibly non-convex from *p (or* of visible non-convexity*) if $S_p \subseteq S^-$.*

Definition 1.4. *We say an embedding $g : L \to \mathbb{R}^3$ is a C^0-approximation of error ϵ of an embedding $f : K \to \mathbb{R}^3$ if there is a homeomorphism $h : g(L) \to f(K)$ with $||h(x)-x|| < \epsilon$ for all x. In addition, we say f is C^0-approximated by a family of mappings if for all $\epsilon > 0$, the family contains mappings that are C^0-approximations of error ϵ of f.*

Definition 1.5. *An ϵ-perturbation of a manifold K is the image of a homeomorphism h such that $d(h(x), x) < \epsilon$ for all $x \in K$.*

2. Approximating visibly non-convex surfaces with shy PL-mappings

Our first result shows the importance of the property of visible non-convexity for C^0-approximation by shy PL-manifolds.

Theorem 2.1. *Let $f : K \to \mathbb{R}^3$ be a smooth embedding of a 2-manifold with image S and let $x \in f(K)$ be a well-visible point such that both principal curvatures at x are positive (with respect to the agreed coorientation). Furthermore, assume $p \notin T_xS$. Then S cannot be C^0-approximated by shy PL-mappings.*

Proof. Using the hypothesis of two positive principal curvatures, we can choose a connected neighborhood U of x such that ∂U and p lie in different half-spaces of T_xS. Well-visibility of x lets us choose U such that all points in U are visible from p. Furthermore, choose U and $\epsilon' > 0$ such that, if $B(\epsilon, \partial U)$ denotes all points at distance less than ϵ' of $\partial \bar{U}$, $S \setminus B(\epsilon', \partial U)$ consists of two connected components such that the one which is a subset of U is visible from p for all $2\epsilon'$-perturbations of the other component.

For the sake of contradiction, assume there is a C^0-approximation of S of error ϵ by a shy PL-mapping of image $\mathcal{P}$ with ϵ small enough so that:

- $U \setminus B(\epsilon, \partial U) \neq \emptyset$
- $\epsilon < \frac{1}{2} \inf_{u \in \overline{B(\epsilon, \partial U)} \cap S} d(u, T_xS)$, where d denotes signed Euclidean distance to the plane T_xS, with sign chosen so that $d(p, T_xS) < 0$
- $\epsilon \leq \epsilon'$

If the approximation is given by a homeomorphism $h : \mathcal{P} \to S$, define the set $\mathcal{P}_U = h^{-1}(\bar{U})$, which has $\partial \mathcal{P}_U = h^{-1}(\partial U)$. Now, consider $y_0 \in \mathcal{P}_U$ such that

$$d(y_0, T_xS) = \sup_{y \in \mathcal{P}_U} d(y, T_xS).$$

Since we are maximizing a continuous function over a polyhedron, we can assume y to be a vertex or a point in the boundary. However, the latter cannot happen due to the second assumption in the choice of ϵ. We claim that y is visible from p, thus $\mathcal{P}$ is not shy.

In fact, we have $d(x, T_xS) = 0$, so $d(y_0, T_xS) \geq -\epsilon$ (we are using signed distances with the convention that coorientation is chosen so that the positive normal vector at x points to the region of $\mathbb{R}^3 \setminus S$ that p does not lie in) and thus $y_0 \notin B(\epsilon, \partial U)$:

- the choice of y_0 ensures that no other point of $\mathcal{P}_U$ lies on the segment $[y_0, p]$
- the second condition on ϵ ensures that $h^{-1}(B(\epsilon, \partial U))$ has all its points at a distance greater than ϵ of T_xS, thus not affecting the visibility of y_0
- the third condition on ϵ ensures that no point in $h^{-1}(S \setminus (U \cup B(\epsilon, \partial U)))$ affects the visibility of y_0

Hence, the theorem is proven.

□

Remark 1. If we consider the weakened visibility hypothesis "x is visible from p" and remove the $p \notin T_xS$ condition, we cannot longer use the local structure of positive curvatures around x to outrule the existence of a C^0-approximation by shy PL-mappings. Indeed, it is easy to construct a surface for which there is a neighborhood U around x in which only x is visible, by creating other three points x_1, x_2, x_3, all in the segment $[p, x]$ and such that $\cap_{i=1,2,3} T_{x_i}S$ is the line through p and x. Then S can be approximated by PL-mappings that are shy in a neighborhood around x by simply enforcing.

Furthermore, one can prove that surfaces of visible non-convexity can be locally approximated by shy PL-mappings around generic points.

Theorem 2.2. *For a smooth embedding of a compact surface $f : K \to \mathbb{R}^3$ with image S of visible non-convexity with respect to a point $p \notin f(K)$ and any $x \in S$ such that $p \notin T_xS$, one can find a neighborhood U around x such that $\bar{U}$ can be C^0-approximated to arbitrary precision by a PL-mapping of image $\mathcal{P}$ and whose only visible vertices lie on $\partial\mathcal{P}$. Furthermore, the simplices with a visible vertex can be assumed to be arbitrarily small.*

Proof. **Case 1:** x is hidden.

In this case the claim is trivial as we can take a neighborhood $U \subset S$ of x such that all of its points are hidden and C^0-approximate U by any PL-manifold.

Case 2: x is visible

First take a neighborhood U of x small enough so that the stereographic projection $\pi : U \to T_xS$ from p onto the plane T_xS is injective. By the hypothesis of one negative curvature, for $y \in U$ close enough to x, there is a cone of directions $\mathcal{C}_y$ centered at $\pi(y)$ such that the preimages of rays in those directions have a negative normal curvature. It will however be better to interpret this cone as a subset of S^1, forgetting the center $\pi(y)$. This cone varies continuously with respect to y, so one can find a neighborhood $V \subseteq U$ of x small enough so that $\mathcal{C}_V = \cap_{y \in V} \mathcal{C}_y$ is a non-trivial cone (i.e. contains an open set of S^1). Consider any two directions in the same connected component of $\mathcal{C}_V \subseteq S^1$ and take two small vectors w_1, w_2 in these directions. Consider the lattice L spanned by w_1 and w_2 on T_xS.

Consider the sets of segments:

$$L_1 = \{[aw_1 + bw_2, (a+2)w_1 + bw_2] : a, b \in \mathbb{Z}, a + b \equiv 0 (\text{mod } 2)\}$$
$$L_2 = \{[aw_1 + bw_2, aw_1 + (b+2)w_2] : a, b \in \mathbb{Z}, a + b \equiv 1 (\text{mod } 2)\}$$

Define also $\tilde{L_i}$ as a set of segments with endpoints $\pi^{-1}(\partial l)$ for any $l \in L_i$. Then, by the hypothesis of w_1, w_2 pointing in directions contained in $\mathcal{C}_V$, all points in U, in particular those in $\pi^{-1}(L)$, have negative normal curvature in the directions tangent to $\pi^{-1}(w_i)$. This implies that all of $(\tilde{L_1} \cup \tilde{L_2})$ lies on the set of non-positive co-orientation with respect to S, with only the

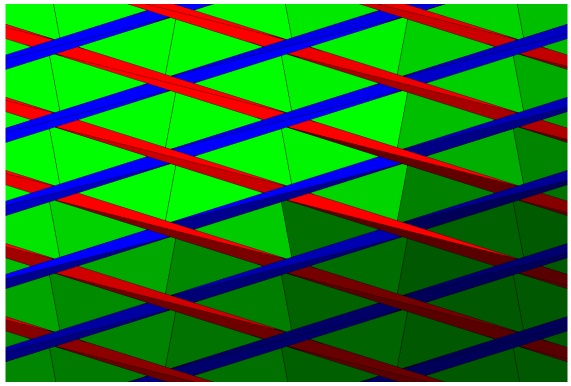

(A) Top-view of the approximation by a shy PL-mapping, i.e., as it is seen from p

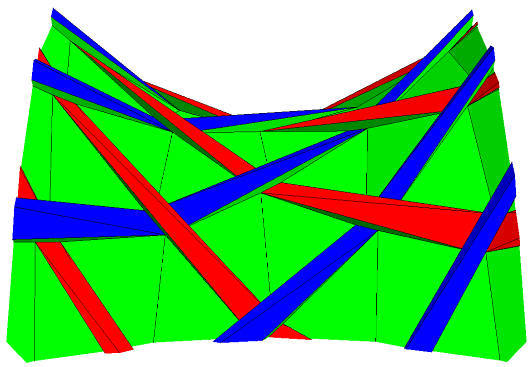

(B) Side-view of the approximation by a shy PL-mapping

FIGURE 1. All vertices of the PL-mapping are vertices of at least one red or blue triangle. All vertices of a red triangle are hidden by a blue triangle and vice-versa. The red and blue triangles arise as the slightly widened faces around the subsets L_1 and L_2 of the lattice L as in the ensuing proof. Although these faces are rectangles in the proof, each of which corresponds to two adjacent triangles in the pictures, both procedures work–the latter is simply slightly better for the pictures. The green faces fill the gaps in a well-established fashion to complete the PL-surface.

vertices lying on S itself. Now, each vertex in L_1 lies inside a segment of L_2 (e.g. aw_1+bw_2 lies on $(aw_1+(b-1)w_2, aw_1+(b+1)w_2)$), so each vertex of $\tilde{L_1}$ (inside $\pi^{-1}(U)$) is hidden by a segment of $\tilde{L_2}$ (e.g. $\pi^{-1}(aw_1+bw_2)$ is hidden by the interior of $[\pi^{-1}(aw_1+(b-1)w_2), \pi^{-1}(aw_1+(b+1)w_2)] \in \tilde{L_2}$ as all of this interior lies on the domain of negative co-orientation) and vice versa.

Finally, choose a small δ and replace each segment in $\tilde{L_1} \cup \tilde{L_2}$ by a rectangle of width δ in such a way that adjacent segments are replaced by rectangles whose short sides are parallel and a slight distance apart (much smaller than δ, say η), each segment becomes a middle line of the rectangle replacing it, and the widths of the rectangles are perpendicular to the rays from p to the middle line. This new structure has four vertices at a distance smaller than $\frac{1}{2}\delta + \eta$ from each of those in $\tilde{L_1} \cup \tilde{L_2}$, but now they are hidden (since the widths of a rectangle and the rectangle above it are not parallel, as they come from different $\tilde{L}_i$, we can choose η to be small enough so that the upper rectangle, of width δ, covers the short side of the lower rectangle, also of width δ, even after the displacement by η). The structure can be completed into a PL-manifold without adding new vertices, similar to Figure 1, and, if w_1, w_2 were chosen small enough, it will be a C^0-approximation of error ϵ of $\bar{U}$ with vertices only corresponding to points in δU.

□

Of course, one would like to be able to C^0-approximate the whole of any surface of visible non-concavity by a single shy PL-mapping. There are two obstacles in this direction: constructing the local PL-mappings around points with $p \in T_xS$ and merging the local PL-mappings into a single global one. We present two partial results regarding the first obstacle. First, notice that the set $\{x \in S' | p \in T_xS'\}$ corresponds exactly to the singular points under a radial projection π_{S^2} to the unit sphere centered at p ($x \mapsto \frac{x}{||x||}$ when $p = 0$). The following proposition shows there exists an arbitrarily small perturbation S' of S in which the set of singular values under this projection to a field of view has a well-characterized structure. Secondly, we present a picture that explains some intuition on how to construct a local PL-mapping around these points.

Proposition 2.3. *For a smooth embedding $f : K \to \mathbb{R}^3$ of visible non-convexity of a compact surface K and any $\delta > 0$, there exists an embedding $g : K \to \mathbb{R}^3$ with image $g(K) = S'$ also of visible non-convexity that satisfies $||g(y) - f(y)|| < \delta$ for all $y \in K$ and such that the set of singularities of the projection onto a unit sphere centered at p is a 1-submanifold of S' consisting of only two types of points: 'fold' and 'cusp' points.*

Proof. The existence of S' relies on Whitney's Singularity theory [2], which allows us to control the singularities that appear when projecting a surface into a field of view from p. More precisely, by translating p and the image of f, we can assume $p = 0$. Since $p \notin f(K)$, we can think of f as a map into $\mathbb{R}^3 \setminus \{0\} \simeq \mathbb{R}_{\geq 0} \times S^2$. Write $f = f_r \times f_\theta$ for the corresponding radial and angular component functions of f. Then $f_\theta : K \to S^2$ is a smooth map between compact surfaces, so Whitney's singularity theory asserts it can be C^∞-approximated to arbitrary precision by a map g_θ with only two types of singularities (i.e. points where the differential does not have full rank), which in addition form a 1-submanifold $M \subseteq K$:

- fold points: points in K such that local coordinates can be chosen in the domain and image so that the point is the origin and the map is $(x, y) \mapsto (x^2, y)$
- cusp points: points in K such that local coordinates can be chosen in the domain and image so that the point is the origin and the map is $(x, y) \mapsto (x^3 + xy, y)$

Now, df_x was assumed to have full rank for all $x \in K$ and K was assumed to be compact, so there exists a smooth function g_θ sufficiently close (in $C^3(K, \mathbb{R}^3)$) to f_θ such that $g = f_r \times g_\theta : K \to \mathbb{R}^3$ satisfies that:

- dg_x is also full-rank for all $x \in K$
- g is injective
- the image of g is a surface of visible non-convexity. This is possible as the set S^- of points with negative curvature is open in S and contains the closed set S_p of points visible from p. Then we can find an open S_0 such that $S_p \subseteq S_0$ and $\bar{S_0} \subseteq S^-$ and request that only points in S_0 become visible under the perturbation and that, since the approximation is C^3 close, there remains at least one negative curvature at all points in the image of S_0

The first two conditions ensure g is an embedding. Finally, notice that $g_\theta = \pi_{S^2} \circ g$ is the projection onto the unit sphere, so the singular points of g_θ are precisely the desired 1-submanifold of 'fold' and 'cusp' points.

□

Lastly, the following figure shows how to construct a shy PL-mapping that approximates S' around a neighborhood of a cusp point or fold point, alas the construction has not been formalized.

FIGURE 2. In the figure, the construction behaves similar to that of Theorem 2.2 away from the singular points. However, unlike the construction of Theorem 2.2, along the fold points and single cusp point, the vertices of the red faces are covered by the subsequent one, rather than by a blue face.

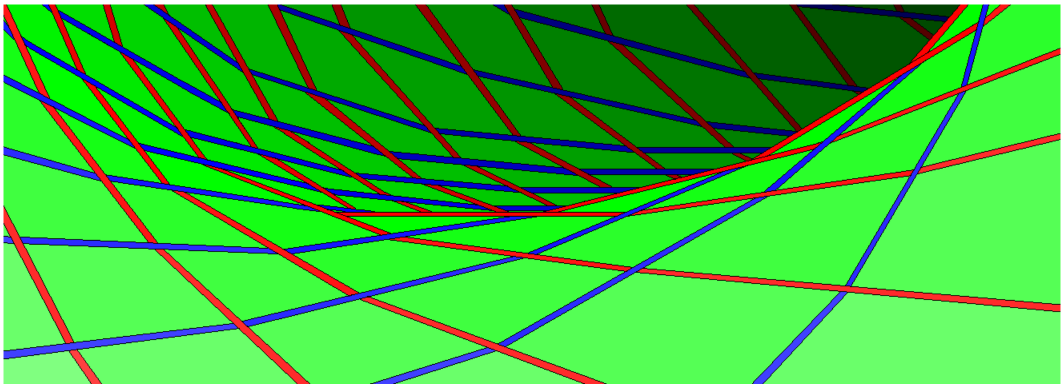

3. Approximating shy PL-mappings with visibly non-convex surfaces

Here we prove a sort of converse of the results of the previous section. The key idea of the proof is to slightly push each face in the direction of positive co-orientation so as to make them visibly non-convex, then smooth out the edges while preserving this property, and ultimately note that the smoothing at the vertices allows a lot of flexibility since they are not visible, although a standard procedure is not described.

Theorem 3.1. *Consider a PL-embedding $g : K \to \mathbb{R}^3$ with image $\mathcal{P}$ that is shy with respect to a point p. Then there exist radii $r_v < R_v$ for each vertex $v \in K$ such that:*

- *$g|_{K\setminus g^{-1}(\cup_v B(v,r_v))}$ can be C^0-approximated by mappings whose images are surfaces of visible non-convexity from p*
- *each of these surfaces intersects the boundaries of each of the $B(v, R_v)$ transversally and the intersection is diffeomorphic to a circle*
- *for each v, all points of $B(v, R_v)$ are hidden from p with respect to the approximation constructed*

Remark 2. Ideally, the approximation would be done for the whole of g, in which case the theorem would read "any PL-embedding g of a compact surface can be C^0-approximated by a smooth map whose image is a surface of visible non-convexity". This result follows from theorem 3.1 if one assumes the intuitive claim that the image S of the approximation given by the theorem can be 'completed' with smooth caps diffeomorphic to D^2 that lie inside each $B(v, r_v)$ and glue smoothly with $S \setminus \cup_v B(v, r_v)$. Furthermore, for the purpose of section 4, it would be important that such construction can be made to behave continuously for parametrized PL-embeddings.

Proof. Fix $\epsilon > 0$ for which we will construct a C^0-approximation with error ϵ.

The choices that follow are so that the sets $N_e \cap \mathcal{P}$ (to be defined) intersect only in small neighborhoods around the vertices of $\mathcal{P}$ all of whose points are hidden and such that the N_e vary continuously under homotopies of PL-mappings of g.

For each vertex v let $r_v > 0$ be the greatest positive number such that $B(v, 2r_v) \cap \mathcal{P}$ is hidden, such that any other vertex is at a distance at least $3r_v$, and such that any point in the interior of a face not incident to v is at a distance at least $2r_v$.

For each edge e, consider the closed subsegment $e^* = e \setminus (B(v, r_v) \cup B(w, r_w))$ where v, w are the vertices to which e is incident. Assign $d_e > 0$ to each edge such that

$$3d_e = \min_{x \in e^*, y \in f^* \neq e^*} ||x - y||$$

i.e. $3d_e$ is the minimum distance between e^* and any other such truncated edge. In particular, $d_e < r_v$ whenever v is incident to e.

Additionally, let $2c_e$ be the minimum distance between e' and any point that lies on a face other than the two that intersect at e'.

Let $\delta_e = \min(d_e, c_e, \epsilon)$. Including ϵ in the minimum is not necessary for this proof, but for the purposes of Section 4, we require δ_e to decrease as ϵ does.

Write N_e for the intersection of the tubular neighborhood of radius δ_e around e and $\mathcal{P}$. Let $\mathcal{P}_0 = \mathcal{P} \setminus \cup_e N_e$, i.e. a set resulting from removing an open neighborhood around the 1-skeleton of $\mathcal{P}$. The conditions above are chosen so that:

- δ_e, and therefore N_ϵ, varies continuously when the embedding g does
- the choice of r_v ensures the $B(v, r_v)$ are disjoint and that they do not intersect faces non-incident to v
- the choice of d_e ensures that for non-adjacent edges e_1, e_2, it happens that $N_{e_1} \cap N_{e_2} = \emptyset$
- the choice of d_e also ensures that, for adjacent edges e_1, e_2 incident to a vertex v, it holds that $N_{e_1} \cap N_{e_2} \subseteq B(v, r_v)$.
- the choice of c_e ensures that if N_e intersects a face it is either a face that has e as an edge, or it is a face incident to an endpoint v of e and in which case the intersection lies in $B(v, r_v)$

In particular, the N_e pairwise intersect only in small non-visible disjoint neighborhoods around the vertices. We thus see that for a face Δ with edges $e(\Delta)$ we have $\Delta \setminus \cup_{e \in \mathcal{P}} N_e = \Delta \setminus \cup_{e \in e(\Delta)} N_e$, so in the construction of $\mathcal{P}_0$, at each face we have only removed a small tubular neighborhood of each of its edges. Hence, $\mathcal{P}_0$ is a collection consisting of one non-empty proper open connected subset of each face of $\mathcal{P}$.

To construct a C^0-approximation $F : \mathcal{P} \to \mathbb{R}^3$ of error ϵ, we first define it on $\mathcal{P}_0$, then around the truncated edges e^*, and finally in the balls at the vertices $B(v, r_v)$.

For each face Δ consider any point $x_0 \in \Delta \cap \mathcal{P}_0$ and let $\hat{n}$ be the normal unit vector to Δ in the positive coorientation. Let P be the plane on which Δ lies and let $r_\Delta = \max_{e \text{ an edge of } \Delta} r_e$. For $x \in P$ define

$$f_{\Delta,\delta,\lambda}(x) = x + (\delta - \lambda||x - x_0||^2)\hat{n}$$

where δ, λ are such that $\delta = \sup_{x \in B(\Delta, r_\Delta) \cap P} \lambda||x - x_0||^2 = \frac{1}{100} \min\left(\epsilon, \inf_{y \in \Delta \cap \mathcal{P}_0, z \in \Delta', \Delta \neq \Delta'} ||y - z||\right)$. The image of this function on P is a paraboloid and in particular

$$||x - f_{\Delta,\delta,\lambda}(x)|| \leq \epsilon/100$$

for $x \in B(r_3)$. We define the C^0-approximation on $\mathcal{P}_0$ to be

$$F|_{\mathcal{P}_0}(x) = f_{\Delta_x,\delta,\lambda}(x)$$

where Δ_x is the face on which x lies. The choice of δ, λ ensures $F_{\mathcal{P}_0}$ is injective and that it is a C^0-approximation of error ϵ.

Before proceeding with the construction, let's show the following

Claim 3.2. *At each point in the image of $F|_{\mathcal{P}_0}$, both principal curvatures are negative.*

Proof. Since each connected component of $\mathcal{P}_0$ is a subset of a plane and we can assume $p = 0$ without loss of generality, it suffices to show that if S is a plane not passing through the origin with a normal vector $\hat{n}$, orientation chosen so that $\hat{n}$ is pointing to the positive side and $d > 0$, then $f_{\Delta,\delta,\lambda}(\Delta)$ has two negative curvatures at each point for any choices of $x_0 \in \Delta \subseteq S$, $\delta > 0$, $\lambda > 0$.

Consider an orthonormal parametrization ψ of the plane that Δ lies on with $\psi(u,v) = x_0 + u\hat{s} + v\hat{t}$ where $\{\hat{s}, \hat{t}, \hat{n}\}$ is an orthonormal basis. Let $\phi = f_{\Delta,\delta,\lambda} \circ \psi|_{\Delta}$, a parametrization of the image of our construction. We compute the coefficients of the first and second fundamental forms:

$$\phi_u = \hat{s} - 2\lambda u\hat{n}, \phi_v = \hat{t} - 2\lambda v\hat{n}$$

$$\phi_{uu} = -2\lambda\hat{n}, \phi_{uv} = 0, \phi_{vv} = -2\lambda\hat{n}$$

$$E = 1 + 4\lambda^2 u^2, F = 4\lambda^2 uv, G = 1 + 4\lambda^2 v^2$$

$$e = -2\lambda, f = 0, g = -2\lambda$$

Hence the mean and Gaussian curvatures satisfy:

$$H = \frac{eG - 2fF + gE}{2(EG - F^2)} = \frac{-2\lambda(2 + 4\lambda^2(u^2 + v^2))}{2(1 + 4\lambda^2(u^2 + v^2))} < 0$$

$$K = \frac{eg - f^2}{EG - F^2} = \frac{4\lambda^2}{1 + 4\lambda^2(u^2 + v^2)} > 0$$

and since $H = \frac{1}{2}(\kappa_1 + \kappa_2)$ and $K = \kappa_1\kappa_2$ where κ_1, κ_2 are the principal curvatures, we conclude both of them are negative.

□

It then remains to construct $F|_{\cup_e N_e}$ in such a way that F is smooth and that visible non-convexity holds for its image.

We continue to define F around each edge. Namely, let π_e be the orthogonal projection onto e and let

$$V(e) = \{x \in N_e : x \notin B(v, r_v) \text{ for any } v\}$$

We define $F|_{V(e)}$ for all e at this step.

For each edge, let Δ_1, Δ_2 be two faces of $\mathcal{P}$ with $\Delta_1 \cap \Delta_2 = e$. Also, write $\Delta_i^* = \Delta_i \cap \mathcal{P}_0$. Define a coordinate system with orthonormal basis $\hat{t}, \hat{z}, \hat{n}$ such that:

- $(0, 0, 0)$ is the midpoint of e
- $\hat{n}$ is orthogonal to e, bisects the angle between Δ_2 and Δ_2, and points towards the region of positive coorientation
- $\hat{z}$ lies on e
- $\hat{t}$ is orthonormal to the previous two. Without loss of generality assume the direction of $\hat{z}$ was chosen so that we can choose $\hat{t}$ such that simultaneously $\hat{z} \times \hat{t} = \hat{n}$ and Δ_2 has positive t-coordinate.

Let P_i be the plane containing Δ_i. The functions $f_{\Delta_i,\delta,\lambda} : P_i \to \mathbb{R}^3$ define two paraboloids each containing one of the $F|_{\mathcal{P}_0}(\Delta_i^*)$. We create an interpolation between these paraboloids to define $F(V(e))$.

Let $\pi_i : P_i \to P_{zt}$ be the projection from each P_i onto the z-t-plane, which is a diffeomorphism. We write $f_i(z,t) = f_{\Delta_i,\delta,\lambda}(\pi_i^{-1}(z,t))$ for convenience. Consider a smooth function $g(z,t) : P_{zt} \to [0, 1]$ satisfying

- the value of g depends only on t, i.e. $\frac{\partial g}{\partial z} = 0$
- g is decreasing as a function of t, i.e. $\frac{\partial g}{\partial t} \leq 0$
- $g(z,t) = 1$ for $(z,t) \in \pi_1(\Delta_1^*)$ and $g(z,t) = 0$ for $(z,t) \in \pi_2(\Delta_2^*)$
- $g(z, 0) = \frac{1}{2}$
- $g(z,t) \in (0, 1)$ only if $||(z,t) - \pi_i^{-1}(z,t)|| \leq \epsilon/100$

In particular, the choice of g becomes unique if we impose that it must be the rescaling of of maximal support that satisfies the conditions above of the function $t \mapsto \frac{1}{e}\exp(-\frac{1}{1-(\frac{t+1}{2})^2})$ for $-1 \leq t < 1$ and $t \mapsto 0$ otherwise. Having a unique such choice is not important for the current proof, but will be essential for the ensuing discussion.

We define the interpolation between the paraboloids by

$$G_e(z,t) = g(z,t)f_1(z,t) + (1-g(z,t))f_2(z,t)$$

which is smooth as all the functions involved are smooth.

Let π_{zt} be the projection onto the z-t-plane (so $\pi_i = \pi_{zt}|_{P_i}$). Finally, for $x \in \Delta_1^* \cup V(e) \cup \Delta_2^*$ we define

$$\tilde{F}_e(x) = G_e(\pi_{zt}(x))$$

. This function agrees with the one for F on points in $\mathcal{P}_0$ as for $x \in \Delta_1^*$ we have $g(\pi_{zt}(x)) = g(\pi_1(x)) = 1$ so

$$\tilde{F}_e(x) = G_e(\pi_{zt}(x)) = f_1(\pi_{zt}(x)) = f_{\Delta_1,\delta,\lambda}(x) = F|_{\mathcal{P}_0}(x)$$

and similarly on Δ_2^*. Let $\mathcal{P}_1 = \cup_e V(e)$. Since the $V(e)$ are disjoint, this allows us to extend F to all of $\mathcal{P}_0 \cup \mathcal{P}_1$ by setting $F(x) = \tilde{F}_e(x)$ for $x \in V(e)$.

We furthermore show the

Claim 3.3. *$F|_{\mathcal{P}_0\cup\mathcal{P}_1}$ is an embedding, a C^0-approximation with error $\epsilon/10$, and is visibly non-concave.*

We first show that no point in the current domain of F is displaced by more than ϵ. If $x \in \mathcal{P}_0$, $||x - F(x)|| < \epsilon/100$. If $x \in \mathcal{P}_1 \cap \Delta_1 \setminus \mathcal{P}_0$ and $||\pi_1(x) - x|| > \epsilon/100$, then $g(x) \in \{0,1\}$, so $F(x) = f_{\Delta_1,\delta,\lambda}(x)$, so $||x - F(x)|| < \epsilon/100$. If $x \in \mathcal{P}_1 \cap \Delta_1 \setminus \mathcal{P}_0$ and $||\pi_1(x) - x|| > \epsilon/100$, then $||x - f_{\Delta_1,\delta,\lambda}(x)|| < \epsilon/100$ and

$$\begin{aligned}||x - f_{\Delta_2,\delta,\lambda}(x)|| &\leq ||x - \pi_1(x)|| + ||\pi_1(x) - \pi_2^{-1}(\pi_1(x))|| + ||\pi_2^{-1}(\pi_1(x)) - f_{\Delta_2,\delta\lambda}(\pi_2^{-1}(\pi_1(x)))|| \\ &< \epsilon/100 + \epsilon/100 + \epsilon/100 = 3\epsilon/100\end{aligned}$$

where the bound on the last term comes from the choice of λ and since $\pi_2^{-1}(\pi_1(x)) \in B(\Delta_2, r_{\Delta_2}) \cap P_2$. Thus in this case

$$||x - F(x)|| \leq ||x - f_{\Delta_1,\delta,\lambda}(x)|| + ||x - f_{\Delta_2,\delta,\lambda}(x)|| < 6\epsilon/100$$

Hence, $||x - F|_{\mathcal{P}_0\cup\mathcal{P}_1}(x)|| < \epsilon/10$ for all $x \in \mathcal{P}_0 \cup \mathcal{P}_1$.

We can now show $F|_{\mathcal{P}_0\cup\mathcal{P}_1}$ is an embedding. For injectivity, if two points $y,\ y' \in \mathcal{P}_0 \cup \mathcal{P}_1$ satisfied $F(y) = F(y')$ then the choice of λ and δ (which ensure that our construction displaces points in distances that are much smaller than those between non-adjacent faces) and the inequality $||x - F|_{\mathcal{P}_0\cup\mathcal{P}_1}(x)|| < \epsilon/10$, imply that for any faces on which y and y' lie must share an edge (this statement includes the case in which y and y' lie on the same face). Call the edge e and name the faces adjacent to e as Δ_1, Δ_2 as in the construction of G_e. It suffices then to show that G_e is injective. This follows from looking at transversal sections. Namely, for each value of z, let W_z be the plane orthogonal to the z-axis at that value. Then for $x \in P_{zt}$ we see

$$x \in W_z \Leftrightarrow f_{\Delta_i,\delta,\lambda}(\pi_i^{-1}(x)) \in W_z \text{ for } i = 1,2 \Leftrightarrow G_e(x) \in W_z$$

, which in addition implies $\frac{\partial}{\partial z}\langle G_e(x), \hat{z}\rangle = 1$ for all $x \in P_{zt}$. We have

$$\frac{\partial}{\partial t}G_e(z,t) = g(z,t)\frac{\partial}{\partial t}f_1(z,t) + (1-g(z,t))\frac{\partial}{\partial t}f_2(z,t) + (f_1(z,t) - f_2(z,t))\frac{\partial}{\partial t}g(z,t)$$

As $\lambda \to 0$, we know $\frac{\partial}{\partial t}f_i(z,t) \to \frac{\partial}{\partial t}\pi_i^{-1}(z,t)$ and we also know $\langle \frac{\partial}{\partial t}\pi_i^{-1}(z,t), \hat{t}\rangle = 1$ and, in particular, is independent of λ. On the other hand, as $\lambda \to 0$, $f_1(z,t) - f_2(z,t) \to \pi_1^{-1}(z,t) - \pi_2^{-1}(z,t)$, so $\langle f_1(z,t) - f_2(z,t), \hat{t}\rangle \to 0$. Finally, since $g(z,t)$ was chosen independently of λ, we see that for $\lambda \to 0$, the equation above and the considerations in this paragraph imply

$$\left\langle \frac{\partial}{\partial t}G_e(z,t), \hat{t} \right\rangle \to g(z,t) + (1-g(z,t)) = 1$$

which is a positive constant. Hence, if λ was chosen small enough, it holds that $\frac{\partial}{\partial t}\langle G_e(x), \hat{t}\rangle > 0$ for all $x \in \pi_{zt}(\Delta_1 \cup \Delta_2)$.

Together with $\frac{\partial}{\partial z}\langle G_e(x), \hat{z}\rangle = 1$, this implies the injectivity of G, thus of $F_{\mathcal{P}_0 \cup \mathcal{P}_1}$.

Injectivity of the differential is then evident: we know from the previous paragraph $\frac{\partial}{\partial z}\langle G_e(x), \hat{z}\rangle = 1$, $\frac{\partial}{\partial t}\langle G_e(x), \hat{z}\rangle = 0$ and $\frac{\partial}{\partial t}\langle G_e(x), \hat{t}\rangle > 0$. Hence dF_x has rank two, so it is injective.

Finally, visible non-concavity has already been shown at points in $F(\mathcal{P}_0)$. Although not all of $F(\mathcal{P}_1)$ is necessarily visible, we show that at all its points we have one negative principal curvature. For $x \in \mathcal{P}_1$, consider the edge e such that $x \in V(e)$ and without loss of generality assume $\{\hat{z}, \hat{t}, \hat{n}\}$ coordinates have been chosen so that $x \in \Delta_1$. Since the principal curvatures are the eigenvalues of the second fundamental form, it suffices to show there is a curve on $F(\mathcal{P}_0)$ such that the second fundamental form evaluated on the tangent direction to the curve at x is negative. Consider the image of a line in $V(e)$ parallel to e. Such a line is (locally) parametrized as $\alpha(\tau) = (z_0 + \tau, t_0, n_0)$ where $x = (z_0, t_0, n_0)$. Let $\gamma(\tau) = F \circ \alpha(\tau)$ and write $c(\tau) = (u(\tau), v(\tau))$. We have

$$II(u', v') = \langle \gamma''(\tau), \mathbf{N}\rangle$$

where $\mathbf{N}$ is the normal to the surface at x.

On the one hand, letting $x_i = (z_i, t_i, n_i)$ be the point with respect to which $f_{\Delta_i,\delta,\lambda}$ was defined, we have

$$\begin{aligned} f_{\Delta_i,\delta,\lambda} \circ \alpha(\tau) &= x + (\delta - \lambda ||(z_0 + \tau - z_i, t_0 - t_i, n_0 - n_i)||^2)\hat{n}_i \\ \Rightarrow \frac{d}{d\tau}(f_{\Delta_i,\delta,\lambda} \circ \alpha(\tau)) &= -2\lambda(\tau + z_0 - z_i)\hat{n}_i \\ \Rightarrow \gamma'(\tau) &= -2\lambda g(\tau)(\tau + z_0 - z_1)\hat{n}_1 - 2\lambda(1 - g(\tau))(\tau + z_0 - z_2)\hat{n}_2 \\ \Rightarrow \gamma''(\tau) &= -2\lambda(g(\tau)\hat{n}_1 + (1 - g(\tau))\hat{n}_2) \end{aligned} \tag{1}$$

To estimate $\mathbf{N}$, we consider that $G_e : P_{zt} \to \mathbb{R}^3$ gives a parametrization of $F(\mathcal{P}_0 \cup \mathcal{P}_1)$ in some neighborhood of x. For shorthand, write $G_z = \frac{\partial}{\partial z}G_e(z,t)$ and $G_t = \frac{\partial}{\partial z}G_e(z,t)$. Also, write $b_i = \frac{\partial}{\partial t}\pi_i^{-1}(z,t) = \hat{t} + (-1)^{i+1}h\hat{n}$, for some $h \in \mathbb{R}$. As usual, as $\lambda \to 0$, we have $f_i|_{\pi_{zt}(\Delta_1 \cup \Delta_2)} \to \pi_i^{-1}|_{\pi_{zt}(\Delta_1 \cup \Delta_2)}$ uniformly as C^∞ maps. In particular, we have convergence on the functions and on their first derivatives. Hence, considering that g depends only on t so $\frac{\partial}{\partial z}g(z,t) = 0$, we may write

$$\begin{aligned} G_z &= (f_1(z,t) - f_2(z,t))\frac{\partial}{\partial z}g(z,t) + (g(z,t)\frac{\partial}{\partial z}f_1(z,t) + (1 - g(z,t))\frac{\partial}{\partial z}f_2(z,t)) \\ &= (g(z,t)\frac{\partial}{\partial z}\pi_1^{-1}(z,t) + (1 - g(z,t))\frac{\partial}{\partial z}\pi_2^{-1}(z,t)) + \epsilon_1(z,t) \\ &= \hat{z} + \epsilon_1(z,t) \end{aligned}$$

where $\epsilon_1(z,t)$ is a small error vector that converges uniformly to 0 on $\pi_{zt}(\Delta_1 \cup \Delta_2)$ when $\lambda \to 0$. Similarly

$$\begin{aligned} G_t &= (f_1(z,t) - f_2(z,t))\frac{\partial}{\partial t}g(z,t) + (g(z,t)\frac{\partial}{\partial t}f_1(z,t) + (1 - g(z,t))\frac{\partial}{\partial t}f_2(z,t)) \\ &= (\pi_1^{-1}(z,t) - \pi_2^{-1}(z,t))\frac{\partial}{\partial t}g(z,t) + (g(z,t)\frac{\partial}{\partial t}\pi_1^{-1}(z,t) + (1 - g(z,t))\frac{\partial}{\partial t}\pi_2^{-1}(z,t)) + \epsilon_2(z,t) \\ &= \rho\frac{\partial}{\partial t}g(z,t)\hat{n} + \hat{t} + (2g(z,t) - 1)h\hat{n} + \epsilon_2(z,t) \end{aligned}$$

where $\epsilon_2(z,t)$ is a small error vector that converges uniformly to 0 on $\pi_{zt}(\Delta_1 \cup \Delta_2)$ when $\lambda \to 0$ and ρ is the real number such that $\pi_1^{-1}(z,t) - \pi_2^{-1}(z,t) = \rho\hat{n}$.

Hence, for some scalar $\sigma > 0$, we have

$$\sigma\mathbf{N} = G_t \times G_z = \hat{n} - (\rho\frac{\partial}{\partial t}g(z,t) + (2g(z,t) - 1)h)\hat{t} + \epsilon_3(z,t)$$

where $\epsilon_3(z,t)$ is a small error vector that converges uniformly to 0 on $\pi_{zt}(\Delta_1 \cup \Delta_2)$ when $\lambda \to 0$.

We know $n_i \cdot b_i = 0$ and that n_i has positive n-coordinate, so $n_i = \xi((-1)^i h\hat{t} + \hat{n})$ for $\xi = \frac{1}{\sqrt{h^2+1}} > 0$. Thus

$$\gamma''(\tau) = -2\lambda\xi(\hat{n} + (1 - 2g(z,t))h\hat{t})$$

Finally,

$$\langle \gamma''(\tau), \mathbf{N}\rangle = -\frac{2\lambda\xi}{\sigma}((1 + (1 - 2g(z,t))^2 h^2 - (1 - 2g(z,t))h\rho\frac{\partial}{\partial t}g(z,t)) + \epsilon_4(z,t)$$

where $\epsilon_4(z,t) \in \mathbb{R}$ and converges to 0 uniformly over $\pi_{zt}(\Delta_1 \cup \Delta_2)$ when $\delta \to 0$. Hence, assuming δ was chosen to be small enough, as the terms other than ϵ_4 do not depend on λ, we would obtain $\langle \gamma''(\tau), \mathbf{N}\rangle < 0$ if we show $(1 - 2g(z,t))h\rho\frac{\partial}{\partial t}g(z,t) < 0$.

Now, since we assumed $x \in \Delta_1$, we have $t \leq 0$, so $g(z,t) \geq \frac{1}{2}$ which gives $(1 - 2g(z,t)) \leq 0$. We also have by hypothesis $\frac{\partial}{\partial t}g(z,t) \leq 0$. Finally, there are three cases:

- If the internal angle at e is smaller than π, then $h > 0$ and $\rho < 0$
- If the internal angle at e is π, then $h = \rho = 0$
- If the internal angle at e is greater than π, then $h < 0$ and $\rho > 0$

Hence, in all cases it holds that $(1 - 2g(z,t))h\rho\frac{\partial}{\partial t}g(z,t) \leq 0$, which implies that the curvature of $F|_{(\mathcal{P}_0 \cup \mathcal{P}_1)}$ in the direction of γ is negative and completes the proof of the claim. □

To finish the proof of the theorem, we need to choose R_v appropriately. Since $B(v, 2r_v)$ is hidden by $\mathcal{P}_0 \cup \mathcal{P}_1$, we know that given ϵ was small enough (for instance that ϵ-perturbation of $\mathcal{P}$ keeps all of $B(v, \frac{3}{2}r_v)$ hidden) and since r_v is independent of ϵ, there is R_v slightly larger than r_v such that $B(v, R_v)$ is hidden by any approximation of $\mathcal{P}_0 \cup \mathcal{P}_1$ of error ϵ. We may take $R_v = 1.1r_v$ for well-definedness. Furthermore, for $\lambda = 0$ and any radius $r_v \leq r \leq 2r_v$ it holds that $\partial B(v,r) \cap \mathcal{P}_0 \cap \mathcal{P}_1$ is a transversal intersection homeomorphic to a circle. In particular, for $r = R_v$, the intersection lies in the interior of $\mathcal{P}_0 \cap \mathcal{P}_1$, so for small enough $\lambda > 0$, it holds that $\partial B(v, R_v) \cap F(\mathcal{P}_0 \cap \mathcal{P}_1)$ is a transversal intersection diffeomorphic to a circle, and we are done.

□

Notice that if we restrict to PL-embeddings with PL-structure given only by triangular faces (always possible as all surfaces can be triangulated), we can choose x_0 canonically as the center of mass of each face. Then $F|_{\mathcal{P}_0 \cup \mathcal{P}_1}$ as constructed in the proof is completely determined given a choice of triangulation T of $\mathcal{P}$, and of parameters ϵ and λ (that are small enough with respect to the conditions outlined along the proof). For this reason, we write $F_{T,\epsilon,\lambda}$ for the corresponding map obtained from the construction with these parameters in what follows.

4. Behavior of the constructions for parametrized mappings

In this section we describe how the suggested extensions of the previous results may piece together to establish a stronger relation between shy PL-embeddings of a surface and embeddings with smooth image of visible non-convexity.

4.1. The space |*ShyPL-K*| of shy PL-mappings.

The ultimate purpose of the investigations has been to try to show that for any fixed compact 2-manifold K the space of its shy PL-embeddings into $\mathbb{R}^3$ (call this set *ShyPL-K*, but we do not specify a topology as we will see some issues in choosing the right topological space to study) is weakly homotopy equivalent to the space of its embeddings into $\mathbb{R}^3$ whose image is a surface of visible non-convexity (call it *VNC-K* and consider the C^0-topology on it). The usual way to establish such a weak homotopy equivalence is by constructing a map between the two spaces and showing that the induced maps in homotopy groups are injective and surjective. However, in the present case, it is unclear that there could exist a generic map from *ShyPL-K* to *VNC-K* (or how one could construct such a map explicitly). To resolve this, the intention is to define such a map by choosing an appropriate $\epsilon > 0$ for each point in a space of mappings and defining the map at such point to be one of the constructed C^0-approximations of error ϵ in such a way that they vary continuously.

In section 3 we partially achieve this, at least when we restrict to a given triangulation of K as follows. First, for T a triangulation of K, we write *ShyPL-K*$_T \subseteq$ *ShyPL-K* for the subset of mappings that are partial-linear with respect to T. The topology in *ShyPL-K*$_T$ we consider is the C^0 one, namely, the one induced by the metric $d(f, f') = \int_K ||f(x) - f'(x)||$. Now, choose a small $\epsilon > 0$ continuously for each $f \in$ *ShyPL-K*$_T$. Then the r_v and R_v from the construction in section 3 are determined (and depend continuously on f).

Then, theorem 3.1 provides a construction of mappings arbitrarily C^0-close to f. More precisely, it proves there exists some $\Lambda(f, \epsilon) > 0$ such that for any $0 < \lambda < \Lambda(f, \epsilon)$, we can construct a map $F_{T,\epsilon,\lambda} : \mathcal{P} \setminus B(v, r_v) \to \mathbb{R}^3$ that is C^0-close to the identity on $\mathcal{P} \setminus B(v, r_v)$. If, in addition we assume as suggested by the remark next to theorem 3.1 that the construction can be canonically extended to all of $\mathcal{P}$, we would have a mapping $F_\lambda : \mathcal{P} \to \mathbb{R}^3$ such that $F \circ f$ is C^0-close to f. Although not proven, it seems the variation of Λ with respect to f and ϵ behaves in a way that may allow us to choose λ in a continuous manner for each $f \in$ *ShyPL-K*$_T$. Then, we would have obtained a well-defined continuous map from *ShyPL-K*$_T$ into *VNC-K*.

At this point it may be tempting to construct a map from *ShyPL-K* into *VNC-K* as follows: for each $f \in$ *ShyPL-K*, consider the coarsest triangulation T of K with respect to which f is partial-linear and define the map from *ShyPL-K* to be that given by the map in *ShyPL-K*$_T$. However, there are two issues with this:

- the triangulations with respect to which f is partial-linear are only partially ordered, so there may not be a single coarsest one
- this approach will not give a continuous map with respect to the C^0-topologies on *ShyPL-K* and *VNC-K* as the construction does not behave well when a new edge appears during a folding of an initially flat face. Namely, if f and f' are partial linear with respect to T and T', respectively, and the latter is a subtriangulation of the former, the obtained F_λ and F'_λ in *ShyPL-K*$_T$ and *ShyPL-K*$_{T'}$ may not become C^0-close when f and f' do regardless of the choice of λ.

Both issues are related in that they arise from the fact that the construction in Theorem 3.1 is not determined by f, ϵ and λ, but is heavily dependent on the triangulation of $\mathcal{P}$.

Hence, if we are to use the construction in section 3 to study the relation between *ShyPL-K* and *VNC-K*, we need to be careful on how to define the space of shy PL-mappings and its topology. Namely, we consider the geometric realization of the simplicial set that assigns to each ordinal $[n]$ the set of chains $(f, T_1) \to \ldots \to (f, T_n)$ where T_j is a subtriangulation of T_i whenever $i \leq j$ and f is partial-linear with respect to all T_i. Abusing the notation for geometric realization, call this space $|$*ShyPL-K*$|$. Geometrically, this creates a n-simplex for each chain (which is convenient as although all points in the simplex correspond to the same embedding f, we can continuously vary the image in *VNC-K* as we move along the simplex to account for the introduction of new edges) and glues them together in a coherent manner (namely subchains of a longer chain appear as a subsimplices of a larger simplex in $|$*ShyPL-K*$|$).

4.2. Construction of a continuous map $\mathcal{F} : |$*ShyPL-K*$| \to$ *VNC-K*.

We now exploit this advantage of having points representing not only the partial linear mappings f, but also the partial linear structure that we are considering on it. Namely, we describe how to construct a continuous map $\mathcal{F} : |$*ShyPL-K*$| \to$ *VNC-K* assuming that the remark after the statement of the theorem 3.1 has provided maps $F_{T,\epsilon,\lambda}$: *ShyPL-K* $\to$ *VNC-K* and that for each triangulation T we have chosen continuous parameters $\epsilon_T(f)$ and $\lambda_T(f)$ such that we have a continuous map $\mathcal{F}_T$: *ShyPL-K*$_T \to$ *VNC-K* satisfying $\mathcal{F}_T(f) = F_{T,\epsilon_T(f),\lambda_T(f)}(f)$. The key idea is that the introduced simplices let us interpolate between the maps at their vertices by first adjusting the λ parameter, then the ϵ parameter to be small enough so that we can take linear combinations of the approximations corresponding to various triangulations and, once they are small enough, taking linear interpolations of these approximations. Geometrically what this does is shrink the neighborhoods around the edges so that they do not interfere with each other across triangulations and thus their linear combinations possess a direction of negative curvature. We describe such construction recursively. However, we impose an additional condition: $\epsilon_{T_n}(f)$ has been chosen small enough so that for any T_i of which T_n is a subtriangulation and any edge e_0 of T_i, the associated neighborhood $N_{T_i}(e_0)$ around this edge as defined in the construction of Theorem 3.1 (if $\epsilon_{T_n}(f)$ is used instead of $\epsilon_{T_i}(f)$) does not intersect any of the neighborhoods $N_{T_n}(e)$ as e ranges over the edges $e \neq e_0$ of T_n.

For each $[x] \in |ShyPL\text{-}K|$ consider a pre-image x (under the map induced by the quotient topology), a simplex Δ^n to which x belongs and suppose it corresponds to the chain $(f, T_1) \to \ldots \to (f, T_n)$. Let $x = (k_1, \ldots, k_n) \in \Delta^n$ where the k_i are the barycentric coordinates of x with respect to the vertices $x_1, \ldots, x_n$ of Δ^n. Let m be the largest index such that $k_m \neq 0$. Then, we let $\mathcal{F}(x) = \sum_{i=1}^{m} t_i F_{T_i,\epsilon,\lambda}(f)$ where t_i, ϵ, and λ are described below, assuming that for $x' = (\frac{k_1}{1-k_m}, \ldots, \frac{k_{m-1}}{1-k_m}, 0, \ldots, 0)$ the corresponding parameters are ϵ', λ', t_i':

- If $\frac{1}{2} \leq k_m \leq 1$, we let, $\lambda = \lambda_{T_m}$ $\epsilon = \epsilon_{T_m}$, $t_i = t_i'(2-2k_m)$ for $1 \leq i \leq m-1$, and $t_m = 2k_m - 1$
- If $\frac{1}{4} \leq k_m \leq \frac{1}{2}$, let $\lambda = \lambda_{T_m}$, $\epsilon = (4k_m - 1)\epsilon_{T_m} + (2 - 4k_m)\epsilon'$, $t_i = t_i'$ for $1 \leq i \leq m-1$, and $t_m = 0$
- If $0 \leq k_m \leq \frac{1}{4}$, let $\lambda = 4k_m\lambda_{T_m} + (1 - 4k_m)\lambda'$, $\epsilon = \epsilon'$, $t_i = t_i'$ for $1 \leq i \leq m-1$, and $t_m = 0$

It is then immediate that this definition is independent of the representative x of the class $[x]$ that we used, as various representatives only differ by having longer barycentric coordinates with additional zeroes on the right.

Furthermore, the condition of visible non-convexity is easy to verify in this construction as the resulting surface is a linear combination of visibly non-convex mappings such that set of points that a linear combination is taken of, there is a fixed direction in which all of them have negative curvature. More precisely, see that for each $y \in K$ there is at most one edge e of T_n such that $f(y) \in N_{T_i}(e)$ for at least one index i for which this neighborhood is defined. Then, in all the mappings $F_{T_i,\epsilon,\lambda}$ either the image of $f(y)$ has negative curvatures in all directions or it has negative curvature in the direction of $F_{T_i,\epsilon,\lambda}(\gamma(\tau))$ where $\gamma(\tau)$ is as described in Theorem 3.1.

4.3. **Relation between parametric approximation and the conjectured weak homotopy equivalence.**

Lastly, we describe how one could expect to show that $\mathcal{F}$ is a weak homotopy equivalence between $|ShyPL\text{-}K|$ and $VNC\text{-}K$, i.e. show that $\mathcal{F}$ induces an isomorphism in homotopy groups, assuming that surfaces of visible non-convexity can also be approximated by shy PL-mappings, not only for a fixed mapping as suggested by Section 2, but parametrically, which allows one to define $\mathcal{G}$ below. In the construction of $\mathcal{F}$ we could have specified any arbitrary precision $\epsilon > 0$ so that $||\mathcal{F}(f) - f||_{C^0} \leq \epsilon$ for all f and we expect that for some appropriate choice of ϵ we could proceed as follows:

- **Surjectivity:** Given a spheroid $\mathcal{S} \subseteq VNC\text{-}K$, we may find $\delta > 0$ so that if $\mathcal{G} : \mathcal{S} \to ShyPL\text{-}K$ is an approximation parametrized by $\mathcal{S}$ with $||\mathcal{G}(f) - f||_{C^0} \leq \delta$ for all $f \in \mathcal{S}$, we have $||\mathcal{F} \circ \mathcal{G}(f) - f||_{C^0} \leq 2\epsilon$ for all $f \in \mathcal{S}$ (possible since $\mathcal{S}$ is compact). Although this only ensures that $\mathcal{S}$ and $\mathcal{F} \circ \mathcal{G}(\mathcal{S})$ are C^0-close, one could expect to construct a homotopy in $VNC\text{-}K$ between them by specifying a canonical procedure to 'unfold' the more complex $\mathcal{G} \circ \mathcal{F}(f)$ back into f while preserving visible non-convexity throughout.
- **Injectivity:** Given spheroids $\mathcal{F}(\mathcal{S}_1)$ and $\mathcal{F}(\mathcal{S}_2)$ in $VNC\text{-}K$ and a homotopy $H : S^n \times I \to VNC\text{-}K$ between them, we want to specify $\delta > 0$ and $\mathcal{G}$ as above so that $||\mathcal{G}(f_t) - f_t||_{C^0} \leq \delta$

for all $f_t \in H(S^n \times I)$. Then, if one creates homotopies H_i for $i = 1, 2$ between $\mathcal{G} \circ \mathcal{F}(\mathcal{S}_i)$ and $\mathcal{S}_i$ (possibly by unfolding edges into the original PL-structure), one can compose the homotopies H_1^{-1}, $\mathcal{G}(H)$, and H_2 to create a homotopy between $\mathcal{S}_1$ and $\mathcal{S}_2$, which shows injectivity.

5. Acknowledgements

The author would like to thank the UROP+ program at MIT for allowing him the opportunity to do this research. He would like to thank his faculty supervisor Prof. Emmy Murphy for suggesting the chosen direction of the research, as well as new possible ones as it progressed. He also thanks his mentor Aleksandr Berdnikov for guiding the research, teaching him some new mathematics, improving the quality of the paper, and, most importantly, for the proposal of the problem and for his outstanding geometric intuition. Finally, he thanks Dr. Slava Gerovitch for organizing the UROP+ program.

References

[1] Joseph O'Rourke. Art Gallery Theorems and Algorithms. Oxford University Press, 1987.

[2] Hassler Whitney. *On Singularities of Mappings of Euclidean Spaces. I. Mappings of the Plane into the Plane.* Annals of Mathematics Second Series, Vol. 62, No. 3 (Nov., 1955), pp. 374-410