

Manifold Homotopy via the Flow Complex

Bardia Sadri*

Abstract

It is known that the critical points of the distance function induced by a dense sample P of a submanifold Σ of \mathbb{R}^n are distributed into two groups, one lying close to Σ itself, called the *shallow*, and the other close to medial axis of Σ , called *deep* critical points. We prove that under (uniform) sampling assumption, the union of stable manifolds of the shallow critical points have the same homotopy type as Σ itself and the union of the stable manifolds of the deep critical points have the homotopy type of the complement of Σ . The separation of critical points under uniform sampling entails a separation in terms of distance of critical points to the sample. This means that if a given sample is dense enough with respect to two or more submanifolds of \mathbb{R}^n , the homotopy types of all such submanifolds together with those of their complements are captured as unions of stable manifolds of shallow versus those of deep critical points, in a filtration of the flow complex based on the distance of critical points to the sample. This results in an algorithm for homotopic manifold reconstruction when the target dimension is unknown.

1 Introduction

Surface reconstruction is the problem of producing from a discrete sample of a surface Σ a concisely represented surface $\tilde{\Sigma}$ that closely approximates Σ and shares its topology, provided that the sample is dense enough. Due to its many applications, this problem has a rich literature spanning several disciplines. In computational geometry a great deal of attention has been given to algorithms that guarantee topological and geometric accuracy of the output under assumptions on the density of the sample; see e.g. [7, 3, 5, 4, 6, 2, 8, 9] or [13] for a survey. Traditionally, “topological equivalence” has been interpreted as *homeomorphism* or even the stronger notion of *ambient isotopy*. This in particular requires the reconstructed object $\tilde{\Sigma}$ to also be a manifold and of the same dimension as the target surface Σ . In this paper, we relax this interpretation to *homotopy equivalence* (see [26] for definitions). Thus we seek to *homotopically reconstruct* the target manifold which consists of finding a concisely represented (here polyhedral) subset of \mathbb{R}^n that is within small Hausdorff distance to the original manifold and shares its homotopy type. We emphasize here that the outcome of a homotopic reconstruction of a sampled manifold need not be a manifold itself.

An intuitive idea that has inspired several reconstruction algorithms [27, 1, 9] is to interpret the reconstructed surface as the zero level set of a (signed) distance function which evaluates to (near) zero in all of the sampled points. The distance function to the target surface itself is clearly one such function. When the sample is sufficiently dense, the distance to the sample closely approximates the distance to the target surface. Thus one may use the distance function induced by the sample as the starting point for building the distance function that leads to the reconstructed surface. Study of distance functions as a natural approach to surface reconstruction has lead to the examination

*Department of Computer Science, University of Toronto, Toronto, Ontario, Canada. The majority of work on this paper was carried out when the author was at the University of Illinois at Urbana-Champaign.

of deeper properties of such functions such as their singularities, gradients, or steepest ascent flows [32, 20].

The *flow complex* was introduced by Giesen and John [22] as a tool for geometric modeling though much of the mathematical foundations behind the flow complex were well-explored prior to that (see [24] and references therein). In essence, for a discrete set $P \subset \mathbb{R}^n$, the *cells* of the flow complex of P is a cell complex that partitions the entire space into a number of cells each of which is a stable set (aka *stable manifold*) under a *flow map* ϕ_P that results from the integration of a vector field v_P that generalizes the gradient of the distance to P [28]. Each cell of the flow complex is the stable manifold of (set of all points that flow into) a critical point of the distance function induced by P . In [21], it was noted empirically that if P is a dense sample of a surface Σ , then the flow complex of P contains a subcomplex that *approximates* Σ .

Prior to [21], flow methods were employed for surface reconstruction (e.g. [20]) but the first of such algorithms with geometric and topological guarantees was found by Dey et al. [18] who proved a sharp separation of critical points of the distance function induced by surface samples into two groups. The points in the first group, called the *shallow critical points* lie close to the surface itself, and those in the second group, called the *deep critical points* lie close to the *medial axis* of the surface. They further showed that for surfaces in \mathbb{R}^3 , the boundary of the union of stable manifolds of inner or outer deep critical points is homeomorphic to the original surface, provided that the sample is dense enough and meets extra regularity conditions.

Contributions. We prove important topological properties about the flow complex induced by dense samples of submanifolds of \mathbb{R}^n of arbitrary dimension and codimension. These properties in a way generalize reconstruction result of [18], albeit with certain reservations. On the downside, we strengthen the adaptive sampling assumption of [18] to the *uniform* sampling where the sample is a subset of the manifold with bounded Hausdorff distance to it. Moreover, the notion of topological equivalence is weakened from homeomorphism to homotopy equivalence, thus our results translate to algorithms for homotopic reconstruction. On the upside, we prove that the union of stable manifolds of the shallow critical points approximates the manifold and captures its homotopy type while that of the deep ones does the same for the complement of the manifold. Plus, we show that this works for any closed submanifold of a Euclidean space of any dimension not just for (codimension-1) surfaces. Capturing the homotopy type of the complement in addition to that of the manifold results in a much stronger topological guarantee. For example, all closed curves, knotted or not, have the same homotopy type and are in fact homeomorphic, but it is the homotopy types of their complements that allow us to distinguish knotted curves from each other or from the unknotted ones. The homotopy equivalence of union of stable manifolds of shallow critical points to the target manifold simply follows from a sequence of known results on the homotopy types of flow shapes, alpha shapes, and union of balls [17, 29, 19, 11]. The other homotopy equivalence, i.e. between the union of stable manifolds of deep critical points and the complement of the manifold, is the core result of this paper and its *raison d’etre*. Standard distance flow arguments as those used in [28, 15, 23] fail in this case; see Section 5. We thus use a novel proof method that successively applies such flow arguments on a family of intermediate sets that are indexed by the indices of shallow critical points.

In many scientific settings, a subspace of interest from a high dimensional space is expected to be a manifold or lie close to one. Naturally occurring data often have much lower inherent dimension than the space in which they live. Inferring the topology of such subspaces based on a collected sample poses a version of the manifold reconstruction problem in which the dimension of the target manifold is not known. In fact, it can be the case that the given sample is a dense sample for multiple submanifolds of



various dimensions of the larger space. For example, in the figure on the right, a sample of the curve can also be regarded as a sample for the torus in which the curves is contained and given the sample alone as input, either the curve or the torus can be returned as a reconstruction of the target manifold. As a result, in recent years, there has been growing interest in manifold reconstruction algorithms that are not supplied with a target dimension (See e.g. [25]).

For uniform samples, the separation of critical points, which is determined in terms of their distance from the target manifold, translates into a separation in terms of the distance from the sample itself. In other words, if one sorts the critical points in the order of their distance to the sample, the shallow critical points make a prefix of this ordering. Thus if one filters the flow complex by putting together the stable manifolds, i.e. cells associated to, critical points in all prefixes of this ordered sequence, one is guaranteed to reach in this filtration a shape homotopy equivalent to any manifold that is represented densely enough by the given sample. As mentioned above, the union of stable manifolds of the remaining critical points then captures the homotopy type of the complement of that manifold. Notice that this filtration is independent of the manifold and is simply a function of the given sample. Thus in the above example of the curve on the torus, if the given sample is dense for both the curve and the torus, the above filtration results homotopic reconstruction of both the curve and the torus, as well as their complements, in different stages.

Two important remarks are in order:

1. The significance of the results of this paper is primarily theoretical as in practice the flow complex is expensive to compute. Numerical issues can affect the structural accuracy of this complex and its exact computation has only been implemented in \mathbb{R}^3 [12]. This is in contrast to much more robustly manageable and more efficiently computable structures such as offset surfaces, alpha-shapes, or witness complexes (See e.g. [16]). Nevertheless, distance functions have been repeatedly used as the basis of many reconstruction algorithms and we believe that the stable manifolds of the flow induced by these functions capture much of the structure of these functions and elucidate their role in manifold reconstruction.
2. Although the results presented in this paper are stated for uniform noise-free samples of smooth submanifolds of \mathbb{R}^n , they all generalize to considerably broader settings at the cost of adjustments of constants: noisy (but uniform) samples can be accommodated using the results of [15]. Furthermore, the assumption of the target shape being a smooth manifold can be dramatically relaxed to allow arbitrary compact subsets of \mathbb{R}^n with positive μ -reach in which case the sample can be taken as any finite (μ, κ) -approximation of the shape for an appropriate choice of κ (See [14] for the definitions as well as a more general separation result). We omit these generalizations from this manuscript and leave them for the full-version of this paper.

The structure of the paper is as follows: In Section 2 we formally introduce the flow map ϕ_P induced by a point set P as well as the resulting flow complex. In Section 3 we prove a slightly different version of the critical points separation result of [18] for uniform samples of submanifolds of \mathbb{R}^n of arbitrary dimension. In Section 4 we show that the union of stable manifolds of the shallow critical points capture the homotopy type of the manifold. Then in Section 5 we prove the corresponding result for the deep critical points and the complement of the manifold. Concluding remarks are provided in Section 6.

2 Background and Preliminaries

Let P be closed nonempty subset of \mathbb{R}^n . The complement of P is the open set $P^c = \mathbb{R}^n \setminus P$. For any point $x \in P^c$, let $h_P(x) = \inf_{y \in P} \|x - y\|$ be the *distance function* defined by P and let $A_P(x) = \{y \in P : \|x - y\| = h_P(x)\}$.

While the distance function h_P is not smooth, it induces a vector field v_P over P^c which behaves like the gradient of h_P in the sense that $v_P(x) \neq 0$ if and only if there is a unique direction of steepest ascent for h_P at x in which case the direction of this steepest ascent is given by $v_P(x)$ (see [24] for more general statement and details). The vector $v_P(x)$ at a point x is characterized by $v_P(x) = \frac{x - d_P(x)}{h_P(x)}$, where $d_P(x)$, called *driver of x* is the center of the smallest enclosing ball of $A_P(x)$, or equivalently, the closest point in $\text{conv } A_P(x)$, the *convex hull* of $A_P(x)$, to x . The *critical points* of h_P are those points x for which $v_P(x) = 0$, or equivalently, $x = d_P(x) \in \text{conv } A_P(x)$.

Lieutier [28] proved that if P^c is bounded, then Euler schemes defined by v_P on P^c uniformly converge and this results in a *flow map* $\phi_P : \mathbb{R}^+ \times P^c \rightarrow P^c$ (where \mathbb{R}^+ is the set of non-negative reals) which he also proved to be continuous in both variables (in fact some of these results on distance functions were discovered earlier in higher generality on semi-concave functions [30]). Intuitively, $\phi_P(t, x)$ is the point y that is reached from following the vector field v_P for time interval of length t , starting at x , by infinitesimal movements proportional to the magnitude of v_P . The map ϕ_P has the classical properties of a flow map, namely $\phi_P(0, x) = x$, $\phi_P(s + t, x) = \phi_P(s, \phi_P(t, x))$, and for any point x and any $t \geq 0$, $v_P(\phi_P(t, x))$ is the *right-derivative* of $\phi_P(t, x)$. Lieutier also proved that h_P along any flow orbit, i.e. $t \mapsto h_P(\phi_P(t, x))$ is increasing and in addition satisfies

$$h_P(\phi_P(t, x)) = h_P(x) + \int_0^t \|v_P(\phi_P(\tau, x))\|^2 d\tau. \quad (1)$$

The special case where P is finite is of particular interest to us and the rest of this section goes over special properties of the flow maps in this case. Let $\text{Vor } P$ and $\text{Del } P$ respectively denote the Voronoi and Delaunay complexes induced by P . For any point $x \in \mathbb{R}^n$, we represent by $V_P(x)$ the *lowest dimensional face* of $\text{Vor } P$ that contains x , and by $D_P(x)$ the face in $\text{Del } P$ dual to $V_P(x)$. The set $A_P(x)$ is the vertex set of $D_P(x)$ and $d_P(x)$ becomes the closest point on $D_P(x)$ to x . It can be verified that all points in the relative interior of the same Voronoi face have the same driver. Since the affine hulls of a Voronoi face and its dual are orthogonal with total dimension n , they intersect in exactly one point. Thus if $V_P(x)$ and $D_P(x)$ intersect, then this intersection consists of a single critical point which is the driver of x . All critical points (except for the maximum at infinity) are characterized the same way (as intersection points of duals). Following [22], we make a *general position assumption* that all pairs of Voronoi and Delaunay objects that are dual to and intersect each other, do so in their relative interiors. The *index* of a critical point c is defined as the dimension of $D_P(c)$.

For a given flow map ϕ_P , the *flow orbit* of a regular point x , denoted $\phi_P(x)$ is defined as $\phi_P([0, +\infty), x)$. For a set T we use $\phi_P(T)$ for $\bigcup_{x \in T} \phi_P(x)$. Notice that by this definition $T \subseteq \phi_P(T)$.

For a critical point c of h_P , the set of all points x whose flow orbit converges to c is called the *stable manifold* of c and denoted by $\text{Sm}(c) = \{x : \phi_P(+\infty, x) = c\}$. Although there is no flow out of a critical point c , we study the orbits of points very close to c . Some of these points flow into c while other flow away from it. We define the *unstable manifold* $\text{Um}(c)$ of a critical point c , as the set of all points into which points arbitrarily close to c flow. Formally, $\text{Um}(c) = \bigcap_{\varepsilon > 0} \phi_P(B(c, \varepsilon))$, where $B(c, \varepsilon)$ denotes the open ball of radius ε centered at c . In other words, the unstable manifold of c consists of c and all the integral lines that start infinitesimally close to c .

Proposition 2.1 *Let P be finite. For a critical point c of h_P , $\text{Um}(c) = \phi_P(V_P(c))$.*

A set T is said to be *flow-tight* for ϕ_P if $\phi_P(T) = T$. Stable and unstable manifolds of critical points and their union and intersections are flow tight. Let \mathcal{C}_P be the set of critical points of h_P induced by P (including the critical point at infinity). The *(stable) flow complex* of P , denoted $\text{Sfc } P$ is the collection of stable manifolds of all critical points in \mathcal{C}_P . Generically, the cell associated to an index k critical point is a topological open k -ball. Moreover, if for critical points $c, c' \in \mathcal{C}_P$, $c \in \partial \text{Sm}(c')$, then $\text{Sm}(c) \subset \partial \text{Sm}(c')$. The following lemma states an important structural property of the stable and unstable flow complex that follows from the correctness of the algorithms for computing these complexes [31, 10].

Lemma 2.2 *If for $c \in \mathcal{C}_P$, $\text{ind } c = k$, then every critical point $c' \in \partial \text{Sm}(c)$ has index less than k , provided that $\text{Sm}(c)$ does not intersect the $(n - k - 1)$ -skeleton of $\text{Vor } P$. Under the same assumption, if $c \in \partial \text{Um}(c')$, then $\text{ind } c' < \text{ind } c$.*

All but a measure-0 set of points P satisfy the requirement that $\text{Sm}(c)$ must stay clear from faces of $\text{Vor } P$ of dimension $n - k - 1$ or smaller (see [31]).

Terminology. For the rest of this paper, by a *manifold* we refer to a C^2 -smooth closed submanifold Σ of \mathbb{R}^n of arbitrary codimension. The medial axis $M(\Sigma)$ of Σ consists of points in space with 2 or more closest points in Σ . The *reach* of Σ is the minimum distance of any point of $M(\Sigma)$ to Σ . The C^2 -smoothness of Σ implies that its reach is strictly positive. Any point $x \notin M(\Sigma)$, has a unique closest point \hat{x} in Σ . The half-line bounded at \hat{x} through x hits $M(\Sigma)$ for the first time at a point \tilde{x} (or at infinity).

A point set $P \subset \Sigma$ is a *uniform ξ -sample* of Σ if $\forall x \in \Sigma \exists p \in P : \|x - p\| \leq \xi$. For a given parameter $r \geq 0$, the union of balls $\bigcup_{p \in P} B(p, r)$ is denoted by $B^{(r)}(P)$. The α -*shape* of P of parameter r , denoted $K^{(r)}(P)$ is the underlying space of *restriction* of $\text{Del } P$ to $B^{(r)}(P)$ (see [19]). The *flow shape* of P for parameter r , denoted $F^{(r)}(P)$ is the union of stable manifolds of critical points at distance $\leq r$ from P (see [17]).

3 Shallow versus deep critical points

For any point $x \in \mathbb{R}^n \setminus (\Sigma \cup M(\Sigma))$ let $\mu(x) = \|\tilde{x} - \hat{x}\|$. If \tilde{x} is at infinity, then $\mu(x) = \infty$. Otherwise, the ratio $0 < \frac{\|x - \hat{x}\|}{\|\tilde{x} - \hat{x}\|} < 1$, is a relative measure of how close to Σ or $M(\Sigma)$ the point x is. It turns out [18, 15] that when a (possibly noisy) sample P of Σ satisfies some density requirements, then critical points of h_P are distributed, according to the above measure, into two distinguishable groups, one lying very close to Σ and the other to $M(\Sigma)$. We state an essentially weaker version of the lemma that is formulated for uniform samples. A proof can be found in the Appendix. Variants of this result for noisy samples, or for surfaces of positive μ -reach can be found in [15] and [14].

Theorem 3.1 *Let P be an $\varepsilon\tau$ -sample of a manifold Σ of reach τ with $\varepsilon \leq 1/\sqrt{3}$. Then for every critical point c of h_P , either $\|c - \hat{c}\| \leq \varepsilon^2\tau$, or $\|c - \hat{c}\| \geq (1 - 2\varepsilon^2)\tau$. In the former case we call c a shallow critical point and in the latter a deep critical point.*

The following Corollary is a technical improvement of Proposition 7.1 in [15] or Lemma 3.3 in [29].

Corollary 3.2 *Let P be an $\varepsilon\tau$ -sample of a manifold Σ of reach τ with $\varepsilon \leq 1/\sqrt{3}$. Then, for every shallow critical point c of h_P , $h_P(c) \leq \sqrt{5/3} \cdot \varepsilon\tau$, and for every deep critical point c' of h_P , $h_P(c') \geq (1 - 2\varepsilon^2)\tau$.*

Proof. Let c be a shallow critical point of h_P and let $\lambda = \|c - \hat{c}\|/\tau$. Since c is shallow, $\lambda \leq \varepsilon^2$. Since $\mu(c) \geq \tau$, by Lemma A.1, $h_P(c) \leq \ell(\varepsilon, \lambda)\tau$. Therefore

$$\begin{aligned} h_P(c) &\leq \sqrt{\lambda^2 + \varepsilon^2(1 + \lambda)} \cdot \tau \leq \sqrt{\varepsilon^4 + \varepsilon^2(1 + \varepsilon^2)} \cdot \tau \\ &= \sqrt{1 + 2\varepsilon^2} \cdot \varepsilon\tau \leq \sqrt{\frac{5}{3}} \cdot \varepsilon\tau. \end{aligned}$$

On the other hand, if c' is a deep critical point, by Theorem 3.1, $\|c' - \hat{c}'\| \geq \mu(c') - 2\varepsilon^2\tau \geq (1 - 2\varepsilon^2)\tau$. The proof follows from the fact that $h_P(c') \geq \|c' - \hat{c}'\|$. \square

For any $0 \leq \delta < 1$, the δ -tubular neighborhood of a manifold Σ of reach τ is defined as the set $\Sigma_\delta = \{x \in \mathbb{R}^n : \|x - \hat{x}\| \leq \delta\tau\}$. Notice that $M(\Sigma) \subset \Sigma_\delta^c$.

The following statement is classical. A proof is supplied in the Appendix for completeness.

Lemma 3.3 *For any $0 \leq \delta < 1$, $cl\Sigma_\delta^c$ is homotopy equivalent to Σ^c . In fact, the former is a strong deformation retract of the latter.*

The following lemma is an adaptation of a similar lemma from [23] for uniform samples. The proof is provided in the Appendix.

Lemma 3.4 *Let P be an $\varepsilon\tau$ -sample of a manifold Σ of reach τ with $\varepsilon \leq 1/(1 + \sqrt{2})$. Then, $cl\Sigma_\delta^c$ is flow-tight under the flow ϕ_P , for any $\frac{\varepsilon^2}{1-\varepsilon} < \delta < 1 - \varepsilon - \frac{\varepsilon^2}{1-\varepsilon}$. In particular this is true for $\delta = 1/2$.*

The above lemma implies that union of stable manifolds of shallow critical points is contained in Σ_δ for $\delta = \varepsilon^2/(1 - \varepsilon)$ thus providing the Hausdorff distance guarantee for our homotopic reconstructions. Alternatively, one can replace Σ_δ with a union of balls of an appropriately small radius centered at all sample points and show that the complement of the this union is flow tight under ϕ_P (See e.g. [15]).

4 Homotopy Type of the Manifold

In this section we show that in a dense enough sample of a submanifold of \mathbb{R}^n , the union of stable manifolds of the shallow critical points has the same homotopy type as the manifold itself. This statement follows from the following sequence of results.

Lemma 4.1 [29] *Let Σ be a manifold of reach τ and let P be an $\varepsilon\tau$ -sample of Σ for any $\varepsilon \leq \frac{1}{2}\sqrt{3/5}$. Then $B^{(r)}(P)$ deformation retracts (and is in particular homotopy equivalent) to Σ , for any $2\varepsilon\tau < r < \sqrt{3/5} \cdot \tau$.*

Lemma 4.2 [19] *For any $r \geq 0$, $B^{(r)}(P)$ and the α -shape $K^{(r)}(P)$ are homotopy equivalent.*

Lemma 4.3 [17, 11] *For any r , the flow shape $F^{(r)}(P)$ and the α -shapes $K^{(r)}(P)$ are homotopy equivalent.*

Theorem 4.4 *Let Σ be a manifold of reach τ and let P be an $\varepsilon\tau$ -sample of Σ for $\varepsilon \leq \frac{1}{2}\sqrt{3/5}$. Then Σ is homotopy equivalent to the union U of stable manifolds of shallow critical points of h_P .*

Proof. For a critical point c of h_P , by Corollary 3.2 $h_P(c) \leq \sqrt{5/3} \cdot \varepsilon\tau$ if c is shallow and $h_P(c) \geq (1 - 2\varepsilon^2)\tau$ if c is deep. For $\varepsilon < \frac{1}{2}\sqrt{3/5}$ the latter bound is strictly greater than the former and therefore there is a positive value r for which $h_P(c) < r$ for every shallow critical point c and $h_P(c') > r$ for every deep critical point c' . Thus the flow shape $F^{(r)}(P)$ is precisely the union of stable manifolds of shallow critical points of h_P with respect to Σ . Lemmas 4.1, 4.2, 4.3 now imply that this union is homotopy equivalent to Σ . \square

5 Homotopy Type of the Complement of the Manifold

In this section we prove that the union of stable manifolds of deep critical points has the homotopy type of Σ^c using the continuity of the flow map ϕ_P . The technique is inspired from the work of Lieutier [28]. A proof can be found in [31] (Theorem 4.20, page 111).

Theorem 5.1 *Let P be a finite set of points in \mathbb{R}^n . If for sets $Y \subset X \subset \mathbb{R}^n$, X and Y are both flow-tight for ϕ_P , i.e. $\phi_P(X) = X$ and $\phi_P(Y) = Y$, and if $X \setminus Y$ is bounded, and, finally, if there is a constant $c > 0$ for which $\|v_P(x)\| \geq c$ for all $x \in X \setminus Y$, then X and Y are homotopy equivalent.*

A difficulty in using the above theorem is that ϕ_P is proven in [28] to be continuous on P^c as long as it is a bounded set. This can be overcome by clipping the space with a very large ball, thus letting $P_0 = P \cup B^c$ where B is a very large ball satisfying $P \subset \frac{1}{5}B$. It can then be verified that within $\frac{1}{2}B$, ϕ_P and ϕ_{P_0} agree which is enough for what we want to prove. In the sequel \mathcal{C}_Σ denotes the set of shallow critical points of P where P is an $\varepsilon\tau$ -sample of a manifold Σ of reach τ . The value of ε is determined later. For shorthand, we write S for Σ^c as well S_δ for Σ_δ^c .

Lemma 5.2 *Let c be a critical point of h_P and let $U \subseteq \mathbb{R}^n$ be a flow-tight set for ϕ_P with $c \notin U$. Let $V = \text{rel int } V_P(c)$. For $r \geq 0$, let $V_r = V \cap B(c, r)$. Then for every $r \geq 0$, if $U \cap B(c, r) \subset V$, then $U \setminus V_r$ is flow-tight for ϕ_P and U and $U \setminus V_r$ have the same homotopy type.*

Proof. We build a deformation retraction from U to $U \setminus V_r$. Since c is a critical point and V is the relative interior of the lowest-dimensional Voronoi face that contains c , c is the driver of the points in V . Consequently, if $x \neq c$ is a point in $V \cap U$, $d_P(x) = c$ and since U is flow-tight for ϕ_P , we have

$$\{x + t(x - c) : t \geq 0\} \cap \text{cl } V \subseteq U.$$

Now, define the map $\rho_r : V_r \rightarrow \partial V_r$ (where ∂V_r is defined relative to the affine hull of V_r) as $\rho_r(x) = x + t(x - c)$ for the smallest $t \geq 0$ such that $\rho_r(x) \in \partial V_r$. In other words, $\rho_r(x)$ is the point at which the ray shot from x in the direction $x - c$ hits the boundary of V_r . Since V_r is convex (it is the intersection of V and $B(c, r)$ which are both convex), it is easy to see that the map ρ_r is continuous (it is a central projection from a point in a convex set to the boundary of the convex set) and because of the assumption that $U \cap B(c, r) \subset V$ this implies that the retraction map $\rho_r^* : U \rightarrow U \setminus V_r$ defined below is also continuous.

$$\rho_r^*(x) = \begin{cases} \rho_r(x) & x \in V_r, \\ x & x \in U \setminus V_r. \end{cases}$$

We now define the map $R_r : [0, 1] \times U \rightarrow U$ as

$$R_r(t, x) = \begin{cases} (1 - t)x + t\rho_r(x) & x \in V_r, \\ x & x \in U \setminus V_r. \end{cases}$$

which gives us a straight line homotopy from the identity map of U to the retraction ρ_r .

If $U \cap B(c, r) \subset V$, then for any $y \in V_r \cap U$, only points in V_r can flow into y . In other words, $y = \phi_P(t, x)$ for some $t \geq 0$ and $x \in U$ implies that $x \in V_r$. Therefore, all flow lines that are affected by removal of V_r from U start in V_r . But we saw above that each such flow line loses an initial segment in $U \setminus V_r$. Thus $U \setminus V_r$ is flow tight for ϕ_P . \square

Theorem 5.3 *Let $\varepsilon \leq \frac{1}{2}\sqrt{3/5}$. Let $\tilde{S} = \bigcup_{c \in \mathcal{C}_\Sigma} \text{Sm}(c)$ be the union of stable manifolds of all deep critical points of h_P with respect to Σ . Then \tilde{S} is homotopy equivalent to S .*




Figure 1: Illustration of the steps in the proof of Theorem 5.3: a) The grayed region represents $\tilde{S} = U_0$. The solid circles are sample points and the empty circles are the shallow index-1 critical points. The Voronoi diagram is shown in thin gray lines and the dotted curve shows the boundary of S_δ . b) A small neighborhood of shallow index-0 critical (sample) points is removed from U_0 to result U'_0 . c) U'_0 is shown to be homotopy equivalent to U_1 which is S_δ plus the unstable manifold of shallow critical points of index 1 or higher, clipped to \tilde{S} . d) A small neighborhood of shallow index-1 critical points are deleted from U_1 to result U'_1 , etc. e) We finally arrive at S_δ .

Proof. Let $U_\Sigma = \bigcup_{c \in \mathcal{C}_\Sigma} \text{Um}(c)$ be the union of unstable manifolds of all shallow critical points and let $U = (S_\delta \cup U_\Sigma) \cap \tilde{S}$ for $\delta = 1/2$. We show that $U \simeq S_\delta$ which proves this theorem since $S_\delta \simeq S$ by Lemma 3.3. First notice that U is a flow-tight set. This is because S_δ and U_Σ are both flow-tight and so is \tilde{S} .

Recall that the index of a critical point c , $\text{ind } c$, is $\dim D_P(c)$, i.e. the dimension of the Delaunay face dual to the lowest dimensional Voronoi face of $\text{Vor } P$ that contains c . Equivalently $\text{ind } c = |A_P(c)| - 1$. Let \mathcal{C}_Σ^i , $i = 0, \dots, n$ denote the set of shallow critical points of index i . Thus $\mathcal{C}_\Sigma = \mathcal{C}_\Sigma^0 \cup \dots \cup \mathcal{C}_\Sigma^n$. We define for every $0 \leq i < n$ the set $U_i \subset \mathbb{R}^n$ as

$$U_i = \left(S_\delta \cup \bigcup_{j \geq i} \bigcup_{c \in \mathcal{C}_\Sigma^j} \text{Um}(c) \right) \cap \tilde{S}.$$

In other words, U_i is the restriction to \tilde{S} of the union of the reduced shape S_δ plus the unstable manifolds of shallow critical points of index i or higher. Notice that since $\mathcal{C}_\Sigma \cap S_\delta = \emptyset$ (by Theorem 3.1) and because S_δ is flow tight for ϕ_P , every flow path that reaches S_δ belongs to the stable manifold of some critical point in S_δ . Therefore $S_\delta \subset \tilde{S}$. Also, notice that $U_0 = \tilde{S}$ and $U_n = S_\delta$ because shallow critical points are not included in \tilde{S} and the unstable manifold of a critical point of index n , i.e. a maximum is the critical point itself.

To complete the proof, we show by induction on i that all U_i 's, $i = 0, \dots, n$, are homotopy equivalent and this will prove the Theorem. To this end, first observe that all U_i 's are flow-tight for ϕ_P by definition. For a base case, we show that U_0 and U_1 are homotopy equivalent. Observe that every point $x \in U_0 \setminus U_1$ belongs to the unstable manifold of some critical point in \mathcal{C}_Σ^0 , i.e. a minimum which is simply a sample point in P . Since U_0 and U_1 are flow tight and $U_0 \setminus U_1$ is bounded (by being contained in Σ_δ), it suffices to show that $\|v_P(x)\| > c$ for some $c > 0$, regardless of the choice of $x \in U_0 \setminus U_1$. Recall that $v_P(x) = (x - d_P(x))/h_P(x)$. If $V_P(x) \cap D_P(x) = \emptyset$, the distance between the driver $d_P(x)$ which is contained in $D_P(x)$ and x is bounded from below with the distance between $V_P(x)$ (which contains x) and $D_P(x)$. Let ζ denote the smallest distance between any pair of dual Delaunay and Voronoi objects that do not intersect. Since $x \in \Sigma_\delta$ which is a bounded set and because $P \subset \Sigma_\delta$, we get

$$\|v_P(x)\| = \frac{\|x - d_P(x)\|}{h_P(x)} \geq \frac{\zeta}{\text{diam } \Sigma_\delta}.$$

If, on the other hand, $V_P(x) \cap D_P(x) = \{c\}$, then $x \in \text{Um}(c)$. Since $x \in U_0 \setminus U_1$, c has to have index 0 and therefore $c = d_P(x)$ and $\|v_P(x)\| = 1$.

Thus we assume that $U_0 \simeq \dots \simeq U_i$ and prove that $U_i \simeq U_{i+1}$. We do this in two stages. First we construct a set U'_i by removing a neighborhood of every index- i shallow critical point in such a way that U'_i is still flow tight for ϕ_P . We then show that $U_i \simeq U'_i \simeq U_{i+1}$ (See Figure 1). The construction of U'_i uses Lemma 5.2. We thus remove for each shallow critical point c a neighborhood $B(c, r_c)$ from U_i for which $r_c > 0$ is chosen small enough so that $B(c, r_c) \cap U_i \subset V_P(c)$. This is the case unless the unstable manifold of some other shallow critical point c' reaches arbitrarily close to c and is not contained in $\text{Um}(c)$. Lemma 2.2 now implies that in this case $\text{ind } c' < \text{ind } c$. But the unstable manifolds of shallow critical points of index less than $i = \text{ind } c$ are not included in U_i . Thus by Lemma 5.2 we can remove a neighborhood of every shallow critical point of index i from U_i to get a set U'_i that is flow-tight for ϕ_P and is homotopy equivalent to U_0 .

Next we show that $U'_i \simeq U_{i+1}$. For this we use Theorem 5.1. Since U'_i and U_{i+1} are both flow-tight for ϕ_P , all we need to do is to find a lower bound for $\|v_P(x)\|$ for points $x \in U'_i \setminus U_{i+1}$. For any such point x , the driver $d_P(x)$ is contained in $D_P(x)$. There are two cases to consider; depending on whether $V_P(x)$ and $D_P(x)$ intersect or not.

If $V_P(x) \cap D_P(x) = \emptyset$ then as argued above $\|v_P(x)\| \geq \zeta / (\text{diam } \Sigma_\delta)$. On the other hand if $V_P(x)$ and $D_P(x)$ do intersect, their intersection will (by definition) be a critical point c which coincides with $d_P(x)$. In this case $x \in \text{Um}(c)$.

Notice that c cannot be a deep critical point since these critical points and their unstable manifolds are contained in $S_{1-2\epsilon^2} \subset S_\delta$ which is flow-tight for ϕ_P . Thus c is a shallow critical point. But in that case c must have index $\leq i$ since unstable manifolds of shallow critical points of index $i + 1$ and higher are included in U_{i+1} while $c \in U'_i \setminus U_{i+1}$. Therefore

$$\|v_P(x)\| = \frac{\|x - c\|}{h_P(x)} \geq \frac{\text{dist}(U'_i, \bigcup_{j=0}^i \mathcal{C}_\Sigma^j)}{\text{diam } \Sigma_\delta} > 0.$$

Theorem 5.1 now implies that $U'_i \simeq U_{i+1}$. The proof follows from the observation that $\Sigma_\delta \subset U_\Sigma$ and therefore $\tilde{S} \subset S_\delta \cup U_\Sigma$ implying that $U = \tilde{S}$. \square

The following corollary immediately follows from Corollary 3.1, Theorem 4.4, and Theorem 5.3. In essence, it enables us to reconstruct all the manifold a given sample densely represents (see Figure 2).

Corollary 5.4 *Let $\Sigma_1, \dots, \Sigma_s$ be manifolds of various dimensions for all of which the same sample P is an $\epsilon\tau_i$ -sample where τ_i is the reach of Σ_i , $i = 1, \dots, s$. If c_1, \dots, c_m are the set of critical points of h_P ordered such that $h_P(c_1) < \dots < h_P(c_m)$, then for each i , there is a j_i such that $\bigcup_{j \leq j_i} \text{Sm}(c_j) \simeq \Sigma_i$ and $\bigcup_{j > j_i} \text{Sm}(c_j) \simeq \Sigma_i^c$.*

6 Conclusions

In this paper, we proved that the separation of critical points of a smooth submanifold of \mathbb{R}^n leads to a very natural way of homotopic reconstruction of the submanifold and its complement. Combined with the fact that this separation can be formulated in terms of the distance from the sample, one obtains a way of homotopic reconstruction of all submanifolds that are represented densely enough by the input sample.

The main result of this paper generalizes to allow noisy samples, i.e. discrete point sets which are within a small enough Hausdorff distance of the manifold. Using the original separation result of




Figure 2: An example illustrating Corollary 5.4. Left: a torus knot is sampled densely. The yellow curve is the reconstructed manifold, consisting of stable manifolds of shallow critical points (with respect to the curve). The scattered white points are deep critical points whose stable manifolds are not included. Right: the sample taken from the curve is dense enough for the torus surface to reconstruct it as a union of stable manifolds of shallow critical points (with respect to the torus). As can be seen the complements of these manifolds have the right homotopy type. In particular, the complement of the torus knot shows that the given curve is knotted.

[18], Theorem 5.3 can be generalized to allow adaptive samples, where the sampling density varies with the local feature size. The proof essentially remains unchanged modulo using the adaptive analogue of 3.4 proven in [23]. For adaptive noisy samples, one can achieve a similar result using a corresponding separation theorem of [15]. However, we know of no analogue for Theorem 4.4 or Corollary 5.4 under adaptive sampling. The obstacle in this case is that the union of stable manifolds of shallow critical points may fail to be a *flow shape*, meaning that the shallow and deep critical points are not separated in terms of their distance to the sample. Consequently, the homotopy equivalence of flow shapes and alpha shapes (Lemma 4.3) may cease to hold.

References

- [1] Marc Alexa, Johannes Behr, Daniel Cohen-Or, Shachar Fleishman, David Levin, and Cláudio T. Silva. Point set surfaces. In *IEEE Visualization*, 2001.
- [2] Nina Amenta and Marshall W. Bern. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.*, 22:481–504, 1999.
- [3] Nina Amenta, Marshall W. Bern, and David Eppstein. The crust and the β -skeleton: Combinatorial curve reconstruction. *Graphical Models and Processing*, 60(2):125–135, 1998.
- [4] Nina Amenta, Sunghee Choi, Tamal K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *Internat. J. Comput. Geom. Appl.*, 12(1-2):125–141, 2002.
- [5] Nina Amenta, Sunghee Choi, and Ravi Krishna Kolluri. The power crust, unions of balls, and the medial axis transform. *Comput. Geom. Theory Appl.*, 19(2-3):127–153, 2001.
- [6] Nina Amenta, Thomas J. Peters, and Alexander Russell. Computational topology: Ambient isotopic approximation of 2-manifolds. *Theo. Comp. Sci.*, 305(1-3):3–15, 2003.
- [7] Jean-Daniel Boissonnat. Geometric structures for three-dimensional shape representation. *ACM Trans. Graph.*, 3(4):266–286, 1984.

- [8] Jean-Daniel Boissonnat and Frédéric Cazals. Smooth surface reconstruction via natural neighbour interpolation of distance functions. *Comput. Geom. Theory Appl.*, 22(1-3):185–203, 2002.
- [9] Jean-Daniel Boissonnat and Steve Oudot. Provably good sampling and meshing of Lipschitz surfaces. In *Proc. 22nd Annu. ACM Sympos. Comput. Geom.*, pages 337–346, 2006.
- [10] Kevin Buchin, Tamal K. Dey, Joachim Giesen, and Matthias John and. Recursive geometry of the flow complex and topology of the flow complex filtration. *Comput. Geom. Theory Appl.*, 40:115–157, 2008.
- [11] Kevin Buchin and Joachim Giesen. Flow complex: General structure and algorithms. In *Proc. 16th Canad. Conf. Comput. Geom.*, pages 270–273, 2005.
- [12] Frédéric Cazals. Robust construction of the extended three-dimensional flow complex. Research Report 5903, INRIA, 2006.
- [13] Frédéric Cazals and Joachim Giesen. Delaunay triangulation based surface reconstruction. In Jean-Daniel Boissonnat and Monique Teillaud, editors, *Effective Computational Geometry for Curves and Surfaces*. Springer-Verlag, 2006.
- [14] Frédéric Chazal, David Cohen-Steiner, and André Lieutier. A sampling theory for compacts in Euclidean space. In *Proc. 22nd Annu. ACM Sympos. Comput. Geom.*, pages 319–326, 2006.
- [15] Frédéric Chazal and André Lieutier. Topology guaranteeing manifold reconstruction using distance function to noisy data. In *Proc. 22nd Annu. ACM Sympos. Comput. Geom.*, pages 112–118, 2006.
- [16] Frédéric Chazal and Steve Yann Oudot. Towards persistence-based reconstruction in euclidean spaces. In *SCG '08: Proceedings of the twenty-fourth annual symposium on Computational geometry*, pages 232–241, 2008.
- [17] Tamal K. Dey, Joachim Giesen, and Matthias John. Alpha-shapes and flow shapes are homotopy equivalent. In *Proc. 35th Annu. ACM Sympos. Theory Comput.*, pages 493–502, 2003.
- [18] Tamal K. Dey, Joachim Giesen, Edgar A. Ramos, and Bardia Sadri. Critical points of the distance to an epsilon-sampling of a surface and flow-complex-based surface reconstruction. In *Proc. 21st Annu. ACM Sympos. Comput. Geom.*, pages 218–227, 2005.
- [19] H. Edelsbrunner. The union of balls and its dual shape. *Discrete Comput. Geom.*, 13:415–440, 1995.
- [20] Herbert Edelsbrunner. Surface reconstruction by wrapping finite point sets in space. *Discrete & Computational Geometry*, 32:231–244, 2004.
- [21] Joachim Giesen and Matthias John. Surface reconstruction based on a dynamical system. *Computer Graphics Forume*, 21:363–371, 2002.
- [22] Joachim Giesen and Matthias John. The flow complex: A data structure for geometric modeling. In *Proc. 14th ACM-SIAM Sympos. Discrete Algorithms*, pages 285–294, 2003.
- [23] Joachim Giesen, Edgar A. Ramos, and Bardia Sadri. Medial axis approximation and unstable flow complex. In *Proc. 22nd Annu. ACM Sympos. Comput. Geom.*, pages 327 – 336, 2006.
- [24] Karl Grove. Critical point theory for distance functions. *Symposia in Pure Mathematics*, 54(3):357–385, 1993.
- [25] Leonidas J. Guibas and Steve Oudot. Reconstruction using witness complexes. In *Proc. 18th ACM-SIAM Sympos. Discrete Algorithms*, pages 1076–1085, 2007.
- [26] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2001.
- [27] Hugues Hoppe, Tony DeRose, Tom Duchamp, John Alan McDonald, and Werner Stuetzle. Surface reconstruction from unorganized points. In *Proc. SIGGRAPH '92*, pages 71–78, 1992.
- [28] André Lieutier. Any bounded open subset of \mathbb{R}^n has the same homotopy type as its medial axis. *Computer-Aided Design*, 36(11):1029–1046, 2004.




Figure 3: Illustration of the proof of Lemma A.1

- [29] P. Niyogi, S. Smale, and S. Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete & Computational Geometry*, 2006.
- [30] A. Petrunin. Semiconcave functions in alexandrov geometry. In *Surveys in Differential Geometry*, volume XI: Metric and Comparison Geometry. International Press of Boston, 2007.
- [31] Bardia Sadri. *Surface and Medial Axis Topology Through Distance Flows Induced by Discrete Samples*. PhD thesis, University of Illinois at Urbana-Champaign, 2006.
- [32] Dirk Siersma. Voronoi diagrams and morse theory of the distance function. In O.E. Barndorff-Nielsen and E.B.V. Jensen, editors, *Geometry in Present Day Science*, pages 187–208. World Scientific, 1999.

A Some Missing Proofs

The following lemma encapsulates some required calculations needed in the proof of Theorem 3.1.

Lemma A.1 *Let τ be the reach of a submanifold Σ of \mathbb{R}^n and let P be an $\varepsilon\tau$ -sample of Σ . Let $x \in \mathbb{R}^n \setminus (\Sigma \cup M(\Sigma))$ be a point satisfying $\|x - \hat{x}\| = \lambda\tau$ for some $0 < \lambda < 1$ and let $\mu_0 = \mu(x)/\tau$. Let B_μ be the ball of radius $\mu(x)$ centered at \hat{x} and let B_τ be the ball of radius τ centered at a point $x_\tau = \hat{x} + \tau(\hat{x} - x)/\|\hat{x} - x\|$ and of radius τ . Define the set $L(x) \subset \mathbb{R}^n$ as*

$$L(x) = \overline{B}(x, \ell(\varepsilon, \lambda)\tau) \setminus (B_\mu \cup B_\tau),$$

where $\ell(\varepsilon, \lambda) = \sqrt{\lambda^2 + \varepsilon^2(1 + \lambda)}$. Then,

1. $A_P(x) \subset L(x)$. In particular, the points in $A_P(x)$ are within distance $\ell(\varepsilon, \lambda)\tau$ from x .
2. $L(x) \subset \overline{B}(x, r_{\mu_0}(\varepsilon, \lambda)\tau) \subseteq \overline{B}(x, r(\varepsilon, \lambda)\tau)$ where $r_{\mu_0}(\varepsilon, \lambda) = \varepsilon\sqrt{\frac{1+\lambda}{1-\lambda/\mu_0}}$ and $r(\varepsilon, \lambda) = \varepsilon\sqrt{\frac{1+\lambda}{1-\lambda}}$.
3. $L(x)$ is contained in a cone with apex x , axis $\hat{x} - x$, and half-angle ψ where $\psi = \psi(\varepsilon, \lambda) = \arcsin(r(\varepsilon, \lambda)/\ell(\varepsilon, \lambda))$.

Proof of Lemma A.1. The interiors of both B_μ and B_τ are disjoint from Σ . By the sampling assumption, there is a sample point in the ball $B_\varepsilon = \overline{B}(\hat{x}, \varepsilon\tau)$. Thus the closest point in P to x is within distance ℓ from x , where ℓ is the distance between x and y , where y is a farthest point from p in the set $L_0 = B_\varepsilon \setminus (B_\mu \cup B_\tau)$. Notice that ℓ is an upper bound for $h_P(x)$. Figure 3 shows a planar section of this setting. In the figure, the region L_0 is shaded with the darkest gray. Let B_ℓ denote the ball $\overline{B}(x, \ell)$. Since P does not intersect B_μ or B_τ , $A_P(x)$ must be contained in the region

$$L_1 = L(x) = B_\ell \setminus (B_\mu \cup B_\tau).$$

Let z be a point in this region farthest away from \hat{x} and let $r = \|\hat{x} - z\|$. Let γ be the angle between $y - \hat{x}$ and the hyper-plane tangent to Σ at \hat{x} and normal to $x - \hat{x}$. It can be easily seen from Figure 3 (right) that $\sin \gamma = \varepsilon/2$.

To simplify notation, let us take τ as unit length. Since the angle $\angle(y - \hat{x}, x - \hat{x})$ is $\pi/2 + \gamma$ we have using the cosine rule

$$\begin{aligned}\ell^2 &= \|x - z\|^2 = \|x - y\|^2 \\ &= \lambda^2 + \varepsilon^2 - 2\varepsilon\lambda \cos(\pi/2 + \gamma) \\ &= \lambda^2 + \varepsilon^2 - 2\varepsilon\lambda \sin \gamma \\ &= \lambda^2 + \varepsilon^2(1 + \lambda).\end{aligned}$$

Now by applying the cosine rule to the triangle $xz\tilde{x}$, we have for the angle $\beta = \angle(x - \tilde{x}, z - \tilde{x})$

$$\begin{aligned}\cos \beta &= \frac{\mu_0^2 + (\mu_0 - \lambda)^2 - \ell^2}{2\mu_0(\mu_0 - \lambda)} \\ &= \frac{\mu_0^2 + (\mu_0 - \lambda)^2 - \lambda^2 - \varepsilon^2(1 + \lambda)}{2\mu_0(\mu_0 - \lambda)} \\ &= 1 - \frac{\varepsilon^2}{2\mu_0^2} \cdot \frac{1 + \lambda}{1 - \lambda/\mu_0}.\end{aligned}$$

If we rewrite the above equality as

$$\cos \beta = 1 - 2 \left(\frac{\varepsilon}{2\mu_0} \sqrt{\frac{1 + \lambda}{1 - \lambda/\mu_0}} \right)^2,$$

and observe on the figure that $\sin(\beta/2) = (r/2)/\mu_0$, we can use the identity $\cos \beta = 1 - 2 \sin^2(\beta/2)$ to obtain,

$$r = \varepsilon \cdot \sqrt{\frac{1 + \lambda}{1 - \lambda/\mu_0}}.$$

To complete the proof, we need to only show that the angle $\beta' = \angle(\hat{x} - x, z - x)$ is smaller than $\psi(\varepsilon, \delta)$ as given in the statement of the Lemma. From the figure $\sin \beta' = h/\ell$ where h is the distance between z and the line supporting the segment $x\hat{x}$. Since $h \leq r$, $\sin \beta' \leq r/\ell = \sin \psi$. \square

Proof of Theorem 3.1. For simplicity we scale the distances so as to have $\tau = 1$. Refer Figure 3 in the proof of Lemma A.1 (Appendix A). First observe that if \tilde{c} is at infinity, the open halfspace whose boundary is tangent to Σ at \tilde{c} and contains c , is disjoint from Σ and therefore from P . This immediately implies that c cannot be a critical point. Thus we assume that \tilde{c} is at finite distance from \tilde{c} .

By Lemma A.1, $A_P(x) \subset B_\ell = B(x, \ell)$, where $\ell = \ell(\varepsilon, \lambda)$ as defined in Lemma A.1. On the other hand $A_P(x)$ is disjoint from $B_\mu = B(\tilde{x}, \mu)$ where $\mu = \mu(x)$. Let H be the hyperplane normal to $\tilde{x} - \hat{x}$ through x and let R be the radius of the ball of intersection between H and B_μ . The plane H is at distance $\|x - \tilde{x}\| = \mu(x) - \lambda$ from \tilde{x} . By the Pythagorean theorem

$$R^2 = \mu^2 - (\mu - \lambda)^2.$$

If the radius ℓ of B_ℓ is less than R , then $B_\ell \setminus B_\mu$ is strictly contained in the open half-spaces of $\mathbb{R}^n \setminus H$ that contains \hat{x} . Since $x \in H$, this implies that $x \notin \text{conv}(B_\ell \setminus B_\mu)$ which further entails that $x \notin \text{conv} A_P(x)$. Since $d_P(x) \in \text{conv} A_P(x)$, this would imply that $\angle(\tilde{x} - x, v_P(x)) < \pi/2$. In particular, x cannot be a critical point if $R > \ell$ or equivalently if

$$\mu^2 - (\mu - \lambda)^2 > \lambda^2 + \varepsilon^2(1 + \lambda).$$

Rearranging the above inequality gives us

$$2\lambda^2 + (\varepsilon^2 - 2\mu)\lambda + \varepsilon^2 < 0.$$

Solving for λ , we get $\lambda_{\min} < \lambda < \lambda_{\max}$, where $\lambda_{\min} = \frac{1}{2} \left(\mu - \varepsilon^2/2 - \sqrt{(\mu - \varepsilon^2/2)^2 - 2\varepsilon^2} \right)$ and $\lambda_{\max} = \frac{1}{2} \left(\mu - \varepsilon^2/2 + \sqrt{(\mu - \varepsilon^2/2)^2 - 2\varepsilon^2} \right)$. Since $\mu \geq 1$, $\varepsilon \leq 1/\sqrt{3}$ is sufficient to have $(\mu - \varepsilon^2/2)^2 - 2\varepsilon^2 \geq 0$. Thus for $\varepsilon \leq 1/\sqrt{3}$, both λ_{\min} and λ_{\max} are real. The assumption of $\varepsilon \leq 1/\sqrt{3}$ can be written as $3\varepsilon^2 \leq 1$ from which

$$2\varepsilon^2 + 1 \leq 2 - \varepsilon^2 \leq 2\mu - \varepsilon^2 = 2(\mu - \varepsilon^2/2).$$

Multiplying by $2\varepsilon^2 \geq 0$, gives us

$$2\varepsilon^2(2\varepsilon^2 + 1) \leq 4\varepsilon^2(\mu - \varepsilon^2/2).$$

By adding $(\mu - \varepsilon^2/2)^2$ to both sides and rearranging we get

$$(\mu - \varepsilon^2/2)^2 - 2\varepsilon^2 \geq ((\mu - \varepsilon^2/2) - 2\varepsilon^2)^2.$$

The smaller side being non-negative allows us to take square roots of both sides which by rearranging results $\lambda_{\min} \leq \varepsilon^2$. As for λ_{\max} , using the inequality $\sqrt{1-t} \geq 1-t$ for $0 \leq t \leq 1$

$$\begin{aligned} \lambda_{\max} &= \frac{1}{2}(\mu - \varepsilon^2/2) \left(1 + \sqrt{1 - \frac{2\varepsilon^2}{(\mu - \varepsilon^2/2)^2}} \right) \\ &\geq \frac{1}{2}(\mu - \varepsilon^2/2) \left(2 - \frac{2\varepsilon^2}{(\mu - \varepsilon^2/2)^2} \right) \\ &= \mu - \varepsilon^2/2 - \frac{\varepsilon^2}{\mu - \varepsilon^2/2} \geq \mu - \varepsilon^2/2 - \frac{\varepsilon^2}{1 - \varepsilon^2/2} \\ &\geq \mu - 2\varepsilon^2. \end{aligned}$$

Thus if $\varepsilon^2 < \lambda < \mu - 2\varepsilon^2$, the point x is separated from $\text{conv } A_P(x)$ and therefore x cannot be a critical point. To complete the proof, we note that $\mu \geq \tau = 1$. \square

Proof of Lemma 3.3. Consider the retraction map $r : \Sigma^c \rightarrow \text{cl } \Sigma_\delta^c$ given by

$$r(x) = \begin{cases} \hat{x} + \delta\tau \cdot (x - \hat{x})/\|x - \hat{x}\| & x \in \Sigma^c \setminus \text{cl } \Sigma_\delta^c \\ x & x \in \text{cl } \Sigma_\delta^c \end{cases}$$

The map r is continuous on $\Sigma^c \setminus \text{cl } \Sigma_\delta^c$ since $(x - \hat{x})/\|x - \hat{x}\|$ changes continuously with \hat{x} (because the surface is smooth), the map $x \mapsto \hat{x}$ is continuous because the only points of discontinuity of this map are medial axis points of which there are none in $\Sigma^c \setminus \text{cl } \Sigma_\delta^c$. The continuity of r on all of its domain follows from a gluing argument using the fact that the points on the boundary of Σ_δ^c are mapped to themselves both with the identity map and with the mapping $x \mapsto \hat{x} + \delta\tau \cdot (x - \hat{x})/\|x - \hat{x}\|$.

If we now define the map $R : [0, 1] \times \Sigma^c \rightarrow \text{cl } \Sigma_\delta^c$ as

$$R(t, x) = \begin{cases} (1-t)x + tr(x) & x \in \Sigma^c \setminus \text{cl } \Sigma_\delta^c \\ x & x \in \text{cl } \Sigma_\delta^c, \end{cases}$$

the map R is a straight-line homotopy from the identity of Σ^c to the retraction map r . \square

B Proof of Lemma 3.4

Lemma B.1 *Let x be a point satisfying $\|x - \hat{x}\| = \delta\tau$. Then, the angle α that $v_P(x)$ makes with $\tilde{x} - x$ is bounded by*

$$\arccos \left(\frac{2\delta(1 - \varepsilon - \delta) - \varepsilon^2}{2(1 - \delta)(\delta + \varepsilon)} \right),$$

provided that the argument of the arccos is between 0 and 1.

Proof. Let c be the point on the line segment $\overline{\hat{x}\tilde{x}}$ at distance τ from \hat{x} .

Let $B = B(c, \tau)$ and let $B' = \overline{B}(x, (\delta + \varepsilon)\tau)$. The driver $d_P(x)$ of x has to be contained in the convex hull of $B' \setminus B$. Let w be a point in the intersection of ∂B and $\partial B'$. The inner angle of the triangle cxw at x is at least $\pi - \alpha$. From the cosine rule we get

$$\begin{aligned} \cos(\pi - \alpha) &\leq \frac{(1 - \delta)^2\tau^2 + (\delta + \varepsilon)^2\tau^2 - \tau^2}{2(1 - \delta)(\delta + \varepsilon)\tau^2} \\ &= \frac{2\delta(\delta + \varepsilon - 1) + \varepsilon^2}{2(1 - \delta)(\delta + \varepsilon)}. \end{aligned}$$

It follows

$$\cos \alpha \geq \frac{2\delta(1 - \delta - \varepsilon) - \varepsilon^2}{2(1 - \delta)(\delta + \varepsilon)},$$

which implies the statement of the lemma. \square

Proof of Lemma 3.4. Take a point x in the boundary of S_δ . For the statement of the theorem to hold, it suffices to for the angle α as defined by Lemma B.1 to be smaller than $\pi/2$ or equivalently $\cos \alpha > 0$. Thus as long as $\delta < 1$, by the same Lemma, it suffices to have $2\delta(1 - \delta - \varepsilon) - \varepsilon^2 > 0$. \square

