Vision, Modeling, and Visualization (2015)
D. Bommes, T. Ritschel, and T. Schultz (Eds.)

Simple, Robust, Constant-Time Bounds on Surface Geodesic Distances using Point Landmarks

Oliver Burghard and Reinhard Klein

Bonn University

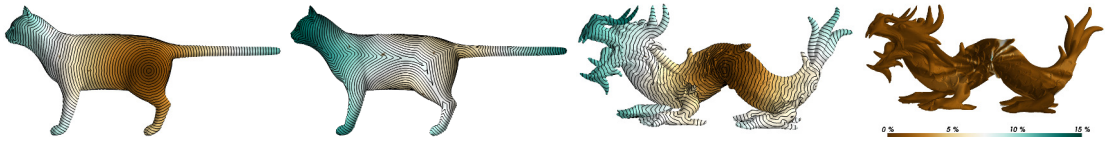

Figure 1: *The lower bound d_{min} is a good distance approximation (a), such as the upper bound d_{max} on the backside (b; 30 landmarks). The quality of d_{min} on the Stanford Dragon (c) can be seen in its low relative error (d; 100 landmarks).*

Abstract

In this paper we exploit redundant information in geodesic distance fields for a quick approximation of all-pair distances. Starting with geodesic distance fields of equally distributed landmarks we analyze the lower and upper bound resulting from the triangle inequality and show that both bounds converge reasonably fast to the original distance field. The lower bound has itself a bounded relative error, fulfills the triangle equation and under mild conditions is a distance metric. While the absolute error of both bounds is smaller than the maximal landmark distances, the upper bound often exhibits smaller error close to the cut locus. Both the lower and upper bound are simple to implement and quickly to evaluate with a constant-time effort for point-to-point distances, which are often required by various algorithms.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Line and curve generation

1. Introduction

Geodesic distances on surfaces are an important tool providing intrinsic information derived from the metric. Even though there has been much research in approximating geodesic distances their calculation can take a significant time in current processing.

To motivate our approximation notice that distance fields at different points typically share a lot of common information (see [XYH12] for a discussion and our later analysis). The triangle inequality estimates lower and upper distance bounds between two points based on distances to a third point. Starting with a reasonable set of landmark points we derive a lower and upper bound on all-pair geodesic distances. We analyze these bounds and show that they are accurate, simple to implement and efficient to compute.

The lower bound has itself a bounded relative error, so that it can be used as an approximation for geodesic distances. The upper bound often exhibits smaller absolute errors close to the cut locus, which makes them better suited for certain class of applications. And because the difference of lower and upper bound is limited by the maximal distance of landmark points, so is the absolute approximation error by both bounds.

Under the mild condition that no point has equal distances to all landmarks (which should not happen for more than 3 landmarks) the lower bound is a distance metric ($d_{min}(p,p) = 0$, triangle equation and $p \neq q \Rightarrow d_{min}(p,q) > 0$). As some efficient methods for calculating geodesic do not assure the triangle equation (e.g. [CWW13]), our lower bound might be used as an approximation instead.

DOI: 10.2312/vmv.20151253

Our distance fields are efficient to compute in the sense that there is a constant effort required for calculating bounds on the distance of two points. Often algorithms depend on pairwise distances only instead of global distance fields (e.g. Karcher means or Voronoi regions). Such algorithms will be typically much faster with our approach, than with global distance fields (see [CWW13, XWL*15]).

2. Related work

Calculating geodesic distance fields there are two different classes of algorithms, exact and approximate ones. Exact algorithms [MMP87, CH90, XW09, XWL*15] often utilize that single-source distances (on a piece-wise linear mesh) equal to a set of quadratic functions on an edge (called windows). Similar to Dijkstra on graphs, they distribute windows between triangles over adjacent edges. Best algorithms have a complexity of $O(n^2)$ [CH90,XWL*15], which is minimal [XWL*15] and thus their complexity is optimal.

Exact distances are often not critical for applications as long as errors are small. Indeed frequently used piece-wise linear surfaces often are an approximation of a continuous surfaces themselves. Geodesic distances solve the Eikonal PDE: $\|\nabla d_p(x)\| = 1$, so that [KS98] approximate geodesics as solutions to this equation. Their approach hast a complexity of $O(n \log n)$). Predicting the geodesic gradient from heat diffusion [CWW13] speeds up approximation further. Ignoring a one-time matrix factorization, it reduces the complexity to $\mathcal{O}(n)$ per distance field, which is trivially optimal for an entire distance field.

Yet this complexity is not optimal when approximating a point to point distances. Typically applications require only few distances and not the whole distance field, e.g. calculating Karcher means [Kar77], intrinsic Voronoi regions [XW10] or non-rigid registration [HAWG08]. After preprocessing a constant time approximation, i.e. not depending on the number of mesh vertices such as our algorithm, would be optimal.

A different set of algorithms has this constant complexity for point to point distances. They define intrinsic distances by embedding a manifold into some euclidean space and back-projecting the distance metric [LRF10, CL06, QH07]. This construction guarantees a distance metric with constant time point to point distances. Still no embedding has been found so that distances are assured to approximate geodesic distances (indeed exact preservation of geodesic distance is often impossible).

[XYH12] proposed a method for constant time geodesic distance approximation with a similar motivation such as our method. From equally spaced landmarks they construct a coarse intrinsic Delaunay triangulation on the surface with precomputed pair-wise distances on the landmarks. Mesh distances are then interpolated utilizing this triangulation by projecting quadrilaterals into the euclidean plane, as side lengths and one diagonal are known.

Their method is similar to ours and delivers good approximations to the exact distances. Yet their approximations might be lead to less smooth, even non-continuous distance fields as they depend on the induced coarse triangulation (which changes non-continuously). They have no bounds on the approximation error and their approximated distance might not be a distance metric. See Sec. 4 for a comparison.

The paper is structured as follow: First we introduce our lower and upper bounds on geodesic distances. Then we analysis their properties and error bounds. Finally we show qualitative and quantitative evaluations.

3. Landmark induced distance bounds

To motivate our method notice that on a surface all-pair geodesic distances share much redundant information. For example, each distance field from p amounts to the information of all shortest paths starting in p. Therefore with all-pair distances all geodesics, which shortest paths are subsets of, can be reconstructed and vice-versa.

A distance field is typically required or desired to adhere to the triangle inequality, that is the shortest path from p to q must get longer if we additionally require that it passes some other point r:

$$d(p,q) \leq d(p,r) + d(r,q) \tag{1}$$

Subtraction of $d(p,r)$ and changing q and r gives a reformulation with a *lower* and *upper* bound on distance $d(p,q)$ induced by a distance field from some landmark point r:

$$|d(p,r) - d(r,q)| \leq d(p,q) \leq d(p,r) + d(r,q) \tag{2}$$

Unifying bounds induced by a set of landmark points $R = \{r_1, \ldots, r_k\}$ we gain the following bounds:

Definition 1 The minimal and maximal induced distances of the landmarks R over the distance metric $d : \mathcal{M} \times \mathcal{M} \to \mathbb{R}_0^+$ are:

$$d_{min}(p,q) := \max_{r \in R} |d(r,p) - d(r,q)| \tag{3}$$

$$d_{max}(p,q) := \min_{r \in R} d(r,p) + d(r,q) \tag{4}$$

The actual landmark r that gave rise to a maximum of $d_{min}(p,q)$ or a minimum of $d_{max}(p,q)$ is called the *inducing* landmark. We add an upper r index to denote the distance w.r.t. a single landmark and will use this also on upcoming definitions: d^r_{min}, d^r_{max}.

For completeness we quickly recapitulate properties defining a distance metric:

Definition 2 A pseudo distance metric is a map $d : \mathcal{M} \times$

	pos	sym	ident	strict pos	tri. ineq.
d_{min}	✓	✓	✓	mostly	✓
d_{max}	✓	✓	×	✓(mostly)	×

Figure 2: *Properties emerging from the definitions of d_{min} and d_{max} (see Theorem 1). Properties in brackets are valid for landmark distances that do not emerge from a proper distance metric (see Sec. 3.1).*

$\mathcal{M} \mapsto \mathbb{R}$ with $(\forall p,q,r \in \mathcal{M})$

$$d(p,q) \geq 0 \quad \text{(pos)} \quad (5)$$
$$d(p,q) = d(q,p) \quad \text{(sym)} \quad (6)$$
$$d(p,p) = 0 \quad \text{(id)} \quad (7)$$
$$d(p,q) \leq d(p,r) + d(r,q) \quad \text{(tri. ineq.)} \quad (8)$$

A distance metric is a pseudo metric with $(\forall p,q \in \mathcal{M})$

$$d(p,q) > 0 \text{ if } p \neq q \quad \text{(strict pos)} \quad (9)$$

As the next Theorem and Tab. 2 shows d_{min} and d_{max} share several desirable properties:

Theorem 1 $d_{min}(p,q)$ is a pseudo distance metric and a proper distance metric iff there are no points with equal distances to *all* landmarks. d_{max} is symmetric and strictly positive. (Please find all proofs in the appendix)

After defining the absolute approximation error as

$$e_{min}(p,q) := d(p,q) - d_{min}(p,q) > 0 \quad (10)$$
$$e_{max}(p,q) := d_{max}(p,q) - d(p,q) > 0 \quad (11)$$

there are the following relations:

$$d_{min}(r,\cdot) = d_{max}(r,\cdot) = d(r,\cdot) \quad (12)$$
$$e_{max}(p,q) + e_{min}(p,q) \leq \min_{r \in R, z \in \{p,q\}} d(z,r) \quad (13)$$
$$e_{min}(p,q) \leq e_{min}(p',q) + 2d(p,p') \quad (14)$$
$$e_{max}(p,q) \leq e_{max}(p',q) + 2d(p,p') \quad (15)$$

On the landmark points the approximation is exact (Eq. 12). Otherwise it is bounded by the maximal inter landmark distance (Eq. 13), which assures that adding landmarks decreases the absolute errors. For example placing landmarks on a torus in a regular grid, n grid points will lead to a landmark distance in $\mathcal{O}(\sqrt{1/n})$. Thus for twice the precision 4 times the landmarks are needed.

Eq. 14 and Eq. 15 allow limiting the error with distances to well approximated sets, which leads to much better convergence. Let $V_{min}(p) := \{q | d_{min}(p,q) = d(p,q)\}$ and $V_{max}(q) := \{q | d_{max}(p,q) = d(p,q)\}$ be the sets of points where d_{min} respective d_{max} are exact. Then the distance of either p or q bounds the absolute error:

$$e_{min}(p,q) \leq 2d(q, V_{min}(p)) \quad (16)$$
$$e_{max}(p,q) \leq 2d(q, V_{max}(p)) \quad (17)$$

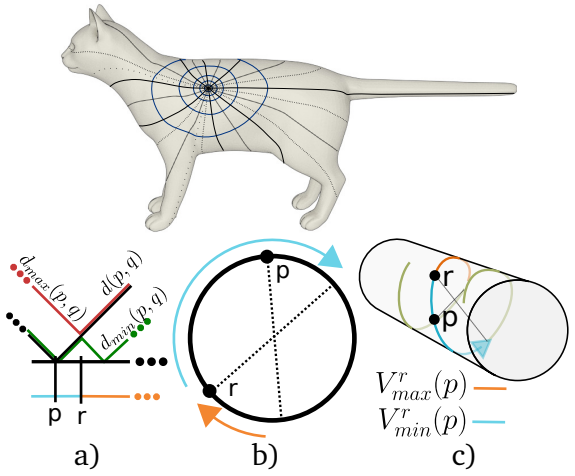

Figure 3: *Top: Illustration of a few maximal shortest paths induced from a single landmark point. Bottom: (a) The euclidean line is split up by a single landmark between two regions - one where d_{min} is exact and one where d_{max} is exact. b) On the circle there due to topology influence there is a third region where none is exact. c) The geodesic spanned by r and p shown with same regions (V_{min}, V_{max}). Lika in b due to topology there are regions on the geodesic where neither is exact.*

Now for estimating the error one has to determine the set $V_{min}(p)$ and $V_{max}(p)$, i.e. all q where d_{min} or d_{max} are exact. Because $e_{min}(p,q) = \min_{r \in R} e^r_{min}(p,q)$ and analog $e_{max} = \min_{r \in R} e^r_{max}(p,q)$ determining $V^r_{min}(p)$ and $V^r_{max}(p)$ is sufficient (remember that upper indices are restrictions to single landmarks).

We assume that we have a smooth surface and for simplicity assume that there is a single shortest path between two points. For geodesic distances and some landmark r, the induced distance d^r_{min} is exact iff q is located on the shortest path of p and r, or p is located on the shortest path of q and r. d_{max} is exact iff r is located on the shortest path connecting p and q.

Shortest paths starting in r either intersect only in r or one is the subset of the other. This partial ordering gives rise to maximal shortest paths. A few of those maximal shortest paths are visualized in Fig. 3. Let R^+ denote the maximal shortest path starting in r including p, R^- be the opposite maximal shortest path located on the same geodesic. Similar let P be the union of maximal shortest path containing r and its opposite. Then:

$$V^r_{min}(p) = R^+ \cap P = R^+ \qquad V_{min}(p) = \bigcup_r V^r_{min} \quad (18)$$
$$V^r_{max}(p) = R^- \cap P \qquad V_{max}(p) = \bigcup_r V^r_{max} \quad (19)$$

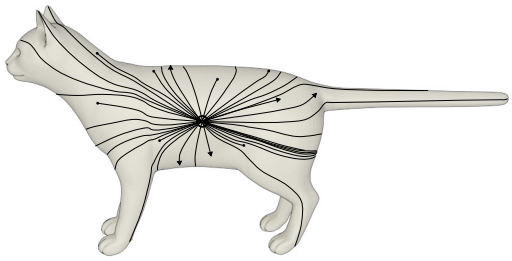

Figure 4: *Visualization of $V_{min}(p)$. For a single point p we collect the maximal induced shortest paths from all landmark points (exactly one curve for each landmark). The error $e_{min}(p,q)$ is bound by twice the distance of q to any of these curves. The curves of d_{min} have smaller lengths than the maximal shortest paths starting in p, what we described as the topological error (ends are marked with small arrows).*

In Fig. 3 we see 3 different domains and fixed p and r. In the first example the euclidean line is partitioned by V_{min} and V_{max} into two segments and for every point on the line, either d_{min} or d_{max} is exact. Any subset of a geodesic (the euclidean line) is a shortest path. Top right we see a closed circle demonstrating the topological influence in comparison with the line. There is a region where neither d_{min} nor d_{max} is exact. Finally bottom left we see a geodesic on a smooth surface showing topological and tangential error.

An illustration of V_{min} can be seen in Fig. 4. Equally spread landmarks create curves quickly becoming dense everywhere. Errors get smaller, the closer p and q, which in our tests resulted in a bounded relative error of d_{min} as well. This dense field of lines leads to a decrease of absolute errors of d_{min} and d_{max} as well as the relative error of d_{min} ($e_{min}(p,q)/d(p,q)$).

Approximation errors can be classified into two categories. For a good approximation one needs a landmark r inducing maximal shortest paths P so that P is close to q. We call this first class of errors tangential errors. Additionally r must be located on P in such a way that errors on P can be inferred and this second class of errors we call topological errors, as it does not appear in euclidean domains. Moving a landmark r along P will change the topological error, moving r so that P changes, changes the tangential error.

3.1. Arbitrary landmark distance fields

The input of the d_{min} and d_{max} is strictly speaking not a distance metric, but consists of $|R|$ different distance fields ($d(r,\cdot) := d_r(\cdot)$). For the given input, there might not exist a distance metric reconstructing input distances. This might be due to numerical errors or might be because the input distances were not derived from a distance metric in the first place. In the following we reason about effects on the bounds.

There are various reasons why arbitrary distance fields might not be compatible with any distance metric: Distances might not be symmetric ($d_r(r') \neq d_{r'}(r)$), triangle equation might not hold between two distance fields, distances might not be 0 at the landmarks ($d_r(r) \neq 0$), distance might be 0 elsewhere ($d_r(p) = 0, p \neq r$) or distance might be negative ($d_r(p) < 0$). Interestingly d_{min} will still be a (pseudo) distance metric:

Theorem 2 Given arbitrary distance fields $d_r : \mathcal{M} \to \mathbb{R}$ with $d_r(r) = 0 \forall r \in R$ as input and define d_{min} and d_{max} as:

$$d_{min}(p,q) := \max_{r\in R} |d_r(p) - d_r(q)| \tag{20}$$

$$d_{max}(p,q) := \min_{r\in R} |d_r(p)| + |d_r(q)| \tag{21}$$

Then most results of Theorem1 stay valid: $d_{min}(p,q)$ is a pseudo distance metric and a proper distance metric iff there are no points with equal distances to *all* landmarks. d_{max} is symmetric and positive.
On landmark points there is:

$$d_{max}(r,\cdot) \leq d_r(\cdot) \leq d_{min}(r,\cdot) \tag{22}$$

and for two points $p,q \in \mathcal{M}$ the inequality $d_{max}(p,q) < d_{min}(p,q)$ holds if and only if there exists two landmarks r_1,r_2 where the triangle inequality can not be fulfilled for p,q,r_1,r_2 (not necessarily pairwise different).

The reason why the exactness of d_{min} depends only on the triangle equation is simply that violations of identity, positivity and symmetry lead to triangle equation violations.

Let $d_{min}[d_r]$ and $d_{max}[d_r]$ denote the distance fields emerging from the distances d_r at the fixed landmarks. Then $d_{min}[d_{min}[d_r]]$ equals $d_{min}[d_r]$, because $d_{min}[d_r]$ is a distance metric, that will be exactly reproduced (Eq. 12). This is generally not true for $d_{max}[d_{max}[d_r]] = d_{max}[d_r]$.

As a simple example we inspect a triangle with edge lengths $1,2,4$, that violates the triangle inequality. All three vertices should be landmark points. Then d_{min} and d_{max} are $2,3,4$ and $1,2,3$ respectively. As guaranteed d_{min} adheres to the triangle inequality, but in this simple case also $\$d_{\max}$. Because the triangle equation was initially violated for all edges we have $d_{max} \leq d_{min}$ everywhere.

There could be 'better' distance metric approximations. For example, we could define the optimal least squares approximation with a least squares energy

$$\text{argmin}_{d \text{ is a distance metric}} \sum_{r\in R} \int (d(r,x) - d_r(x))^2 dx$$

which defines a quadratic program. d_{min} is not optimal, but for our example above $1\frac{1}{3}, 2\frac{1}{3}, 3\frac{2}{3}$ would be.

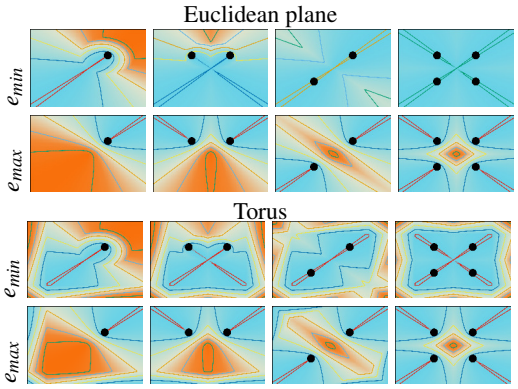

Figure 5: *e_{min} and e_{max} on different domains with up to 4 landmarks (black dots). Top rows depict an unbounded euclidean space and the bottom rows a bounded torus (i.e. warping around left-to-right and top-to-bottom). Please see the text for further discussion. [Coordinates: -3 to 3, landmarks on a circle of radius $\sqrt{2}$, colors from 0 (light blue) to 3 (light brown) with 6 equal spaces contour lines].*

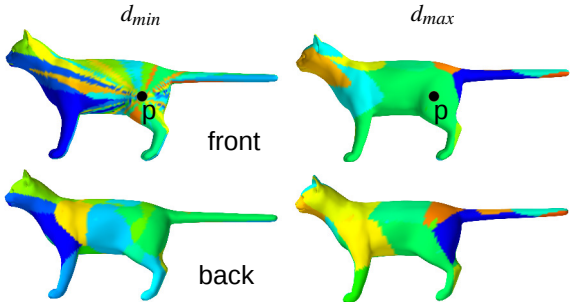

Figure 6: *Regions colored based on the landmark inducing distances to p. Regions should resemble Voronoi regions around lines of $V_{min}(p)$ and $V_{max}(p)$.*

4. Evaluation

To get a first qualitative idea of the bounds, Fig. 5 contains plots of d_{min} and d_{max} in an euclidean plane (top) and on the torus (bottom). There are up to 4 landmarks (black points) and we infer bounds for the distances from the origin o to the plane. First observation is that errors are indeed bound by twice the distance from the origin to the closest landmark ($2 \cdot \sqrt{2}$ in our case) which follows from Eq. 13. The error is bound for each point by twice its distance to the closest position without error ($V_{min}(o)$ and $V_{max}(o)$). In the euclidean plane the distances to $V_{min}(o)$ and e_{min} decrease quickly, which is not true for d_{max}. The torus additionally exhibits topological error, which leads to worse lower bounds d_{min}, but affects d_{max} less.

One interesting insight from the Euclidean case is that relative errors of d_{min}, i.e. e_{min}/d are bound (if there is at least one landmark). Let α_r be the minimal angle between the shortest path from p to q and some path in $V^r_{min}(p)$. Then

$$e_{min}(p,q)/d(p,q) \leq 1 - \cos\alpha_r$$

The same is true on a smooth manifold for some small neighborhood around p. But then the relative error is also globally bound. In our experiments the largest relative errors appeared locally, so that the finer the directions of the tangent space are sampled by shortest paths $V^r_{min}(p)$, the smaller is the maximal relative error of d_{min} (see Fig. 1d and Fig. 10).

For evaluation on real world data we need to decide on landmark points. In our tests we chose farthest point sampling, which worked quite well. We choose a random point first and then iteratively add the point with maximal distance to all previously chosen. Distance calculations were done with [CWW13] so that for k landmarks and n points the runtime is $\mathcal{O}(kn)$, excluding the once required matrix factorization (tests with exact geodesics [MMP87] led to similar results).

After having landmark points spread we can evaluate d_{min} and d_{max} on actual meshes. Fig. 6 shows a cat, where two points share the same color if distances to a fixed p are induced by the same landmark. The exactness of $d_{min}(p,q)$ depends on the minimal distance from q to one of the sets $V^r_{min}(p)$. Thus we expect regions of same colors to resemble Voronoi regions of $V^r_{min}(p)$. Same is true for d_{max} and $V^r_{max}(p)$. In agreement with our previous writing they change frequently for d_{min}, less so for d_{max}.

A quantitative analysis of the bounds, their distances and absolute and relative errors were done on 3 different models in Fig. 7. First 1000 landmark points and their distances were calculated with furthest point sampling. Then the last 100 were chosen as test points, on whose distance fields the bounds were compared to the exact distances. The graph e_{min} for example contains the mean value of e_{min} as measured from the test points to all others. The graph is twice logarithmic, so that exponential functions become straight lines whose slope is the exponent. Two guide lines were added showing the functions $\mathcal{O}(\sqrt{1/n})$ and $\mathcal{O}(1/n)$ to which the other plots can be set in relation. In agreement to our theoretical considerations, absolute errors e_{min} and e_{max} and the relative error e_{min}/d are decreasing similar to $\mathcal{O}(1/n)$, while the distance to the closest landmark point decreases only with order $\mathcal{O}(\sqrt{1/n})$, which is thus not the reason for good convergence. The bounds can be further visually inspected in Figures 8, 9 and 10.

From the related work the work of Xin et al. [XYH12] is most significant, as they approximate geodesic distances in constant time as well. For a fair comparison we chose for their algorithm the same landmark points as for our approach, from which they then build a coarse triangulation to infer distances. They deliver a good approximation, w.r.t. the

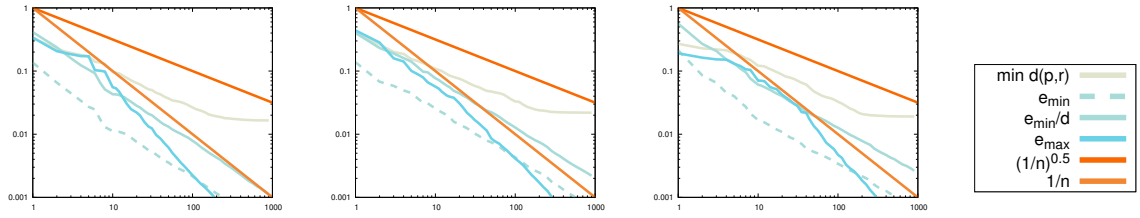

Figure 7: *Absolute and relative errors in a logarithmic plot over the number of points. In agreement with our model errors decrease approximative with $\mathcal{O}(1/n)$ in contrast the distance to the closest landmark decreases only with order $\mathcal{O}(\sqrt{1/n})$. Straight lines represent $\mathcal{O}(1/n)$ and $\mathcal{O}(\sqrt{1/n})$ for reference.*

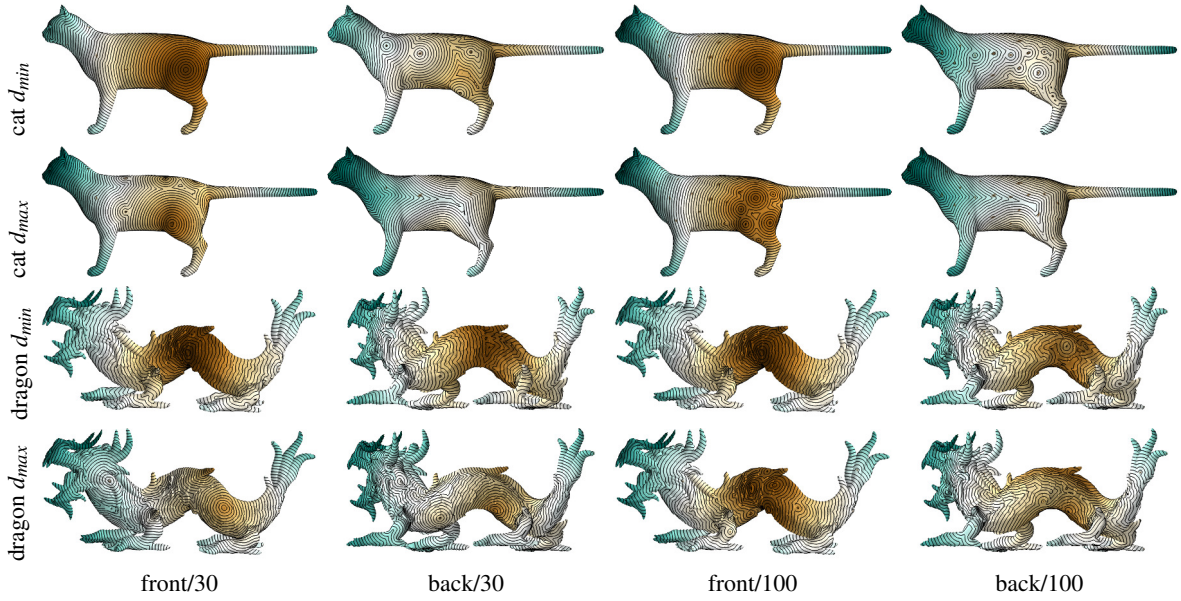

Figure 8: *Visualization of boundaries on the Tosca Cat and Stanford Dragon for a single query point (30/100 landmarks).*

absolute approximation error. Yet through the influence of the triangulation, their approximation is not continuous. Further their approximation might not result in a distance metric. For large Gaussian curvatures the real distances might deviate largely for their approximation, while our method always gives assured bounds. Finally, the implementation of our algorithm is of a remarkable simplicity.

5. Future work

It could be well worth, investigating ideas for better landmark placing. We present the following simple Theorem, that might help relating potential landmarks to the resulting approximation error.

Theorem 3 Let $p,q,s \in \mathcal{M}$, further $d(p,s) \leq d(p,q)$ and $d(q,s) \leq d(p,q)$ and h_s be the shortest distance of s to any

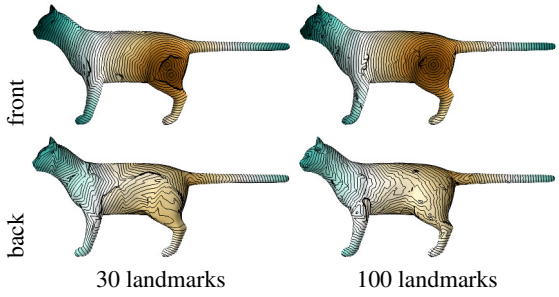

Figure 11: *Comparison to [XYH12]. They calculate constant time all-pairs distances as well. We utilize the same sample points as in our results (furthest point sampling). Note the discontinuities and see the text for discussions.*

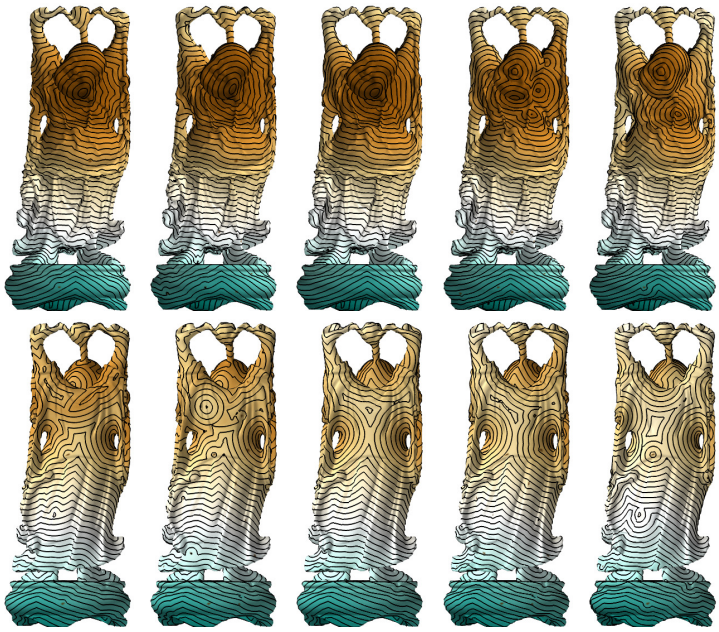

Figure 9: *Visualization of boundaries on the Happy Buddha.*

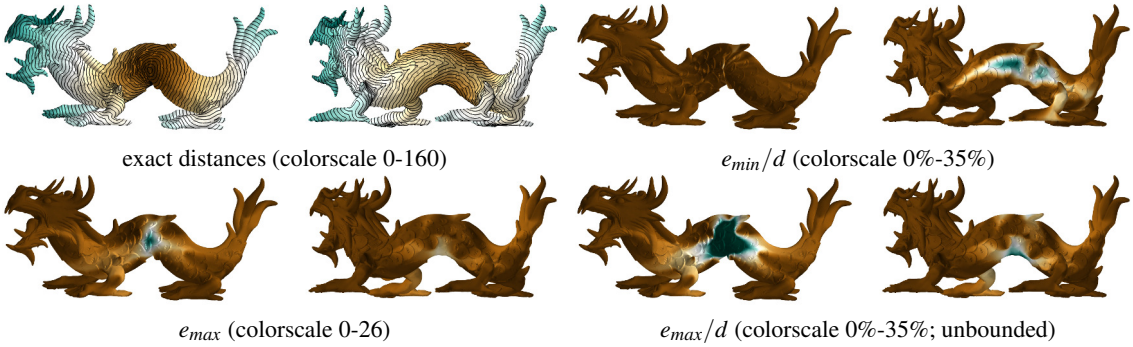

Figure 10: *Visualization of the errors for the Stanford dragon.*

shortest path connecting p and q. Then

$$e^q_{min}(p,s) = e^p_{min}(s,q) = e^s_{max}(p,q) \leq h_s \qquad (23)$$

When solving for distance metrics (for example with linear/quadratic programs), it might be interesting to represented these over finite distance fields as discussed here (Sec. 3.1).

Additionally, it would be interesting to investigate the information that general pairwise distances [LRF10,SRGB14] share and it would be interesting whether a generalization of our method is applicable in their setting.

6. Appendix

Theorem 1 For d_{min} *positivity*, *symmetry* and *identity* follow from the definition. *Strict positivity*: $p \neq q$ then $d_{min}(p,q) > 0$ iff there is a landmark r with $\tilde{d}_r(p) \neq \tilde{d}_r(q)$. If the contours of landmark distances (where distances are constant) are a 1-dimensional set, such as if they were geodesic distances, then the set of points sharing equal distance to all landmarks decreases in dimension with each additional landmark. So for a fixed point p the set $\{q|d_{min}(p,q) = 0\}$ is a set of curves for 1 landmark, a set of points for 2 landmarks and will be only p itself for more than 3 landmarks. *Triangle inequality*: For a single landmark the triangle inequality is assured: $|d(p,r) - d(q,r)| = |d(p,r) - d(r,x) + d(r,x) - d(q,r)| \leq |d(p,r) - d(r,x)| + |d(r,x) - d(q,r)|$. Let r be the landmark

maximizing the left side, then also for this landmark the inequality is fulfilled and the right side is only increased replacing single landmark distances with d_{min}. For d_{max} *positivity* and *symmetry* follow from the definition. *Strict positivity* follows from the strict positivity of the landmark distance metric. *Identity* is only true at landmark points, otherwise $d(p,p) > d(p,r) > 0$.

Eq. 10&11 $d_{min}(p,q) \leq d(p,q)$ and $d(p,q) \leq d_{max}(p,q)$ is a reformulation of the triangle inequality and true for distance metrics.

Eq. 12 $\forall r \in R$: $d_{max}(r,p) \leq d(r,p) + d(r,r) = |d(r,p) - d(r,r)| \leq d_{min}(r,p)$. Thus $d(r,\cdot) = d_{min}(r,\cdot) = d_{max}(r,\cdot)$.

Eq. 13 $e_{max}(p,q) - e_{min}(p,q) = d_{max}(p,q) - d_{min}(p,q) \leq d^r_{max}(p,q) - d^r_{min}(p,q) \leq (d(p,r) + d(q,r)) - |d(p,r) - d(q,r)| \leq (d(p,r) + d(q,r)) + \min(d(p,r) - d(q,r), d(q,r) - d(p,r)) \leq 2\min(d(q,r), d(p,r))$.

Eq. 14 First we look at a single landmark r. From $|d(p,r) - d(q,r)| \geq d(p',r) + d(p',p) - d(q,r) \geq d(p',r) - d(q,r) - d(p',p)$ and $|d(p,r) - d(q,r)| \geq d(q,r) - (d(p',r) + d(p',p))$ follows $|d(p,r) - d(q,r)| \geq |d(q,r) - d(p',r)| - d(p',p)$ and $d(p,q) - |d(p,r) - d(q,r)| \leq (d(p,q) + d(p',p)) - (|d(q,r) - d(p',r)| - d(p',p))$.

If we know $e^r_{min}(p,q) \leq e^r_{min}(p',q) + 2d(p',p)$ then there is one r' with $e^{r'}_{min}(p',q) = e_{min}(p',q)$ and $e_{min}(p,q) \leq e^{r'}_{min}(p,q) \leq e^{r'}_{min}(p',q) + 2d(p',p) \leq e_{min}(p',q) + 2d(p',p)$.

Eq. 15 From $d(p,q) \geq d(p',q) - d(p,p')$ follows $(d(p,r) + d(q,r)) - d(p,q) \leq (d(p,p') + d(p',r) + d(q,r)) - (d(p',q) - d(p,p'))$. Transfer onto multiple landmarks is analog to above.

Thm.2 Thm.1 stays valid as all utilized requirements w.r.t. d_{min} are still valid. **Eq. 22** $\forall r \in R$: $d_{max}(r,p) \leq |d(r,p)| + |d(r,r)| = |d(r,p) - d(r,r)| \leq d_{min}(r,p)$ **Eq. 23** We assume that distance are induced from a distance metric. Then if $a < e + f$ and $b < e + f$ follows $e^A_{min}(B,C) = e^B_{min}(A,C) = e^C_{max}(A,B) = a + b - e - f \leq 2h$.

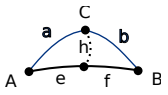

cont. Thm.2 For two points p and q distances $d_{min}(p,q)$ and $d_{max}(p,q)$ are induced by two landmarks r_1 and r_2 (where they are minimal/maximal). $d_{min}(p,q)$ is a lower bound and $d_{max}(p,q)$ is an upper bound on distances between p and q induced by the triangle inequality from r_1 and r_2. Thus $d_{max}(p,q) < d_{min}(p,q)$ iff landmarks contradict the triangle inequality in p and q.

References

[CH90] CHEN J., HAN Y.: Shortest paths on a polyhedron. In *Proceedings of the Sixth Annual Symposium on Computational Geometry* (New York, NY, USA, 1990), SCG '90, ACM, pp. 360–369. doi:10.1145/98524.98601. 2

[CL06] COIFMAN R. R., LAFON S.: Diffusion maps. *Applied and Computational Harmonic Analysis 21*, 1 (2006), 5–30. Special Issue: Diffusion Maps and Wavelets. doi:10.1016/j.acha.2006.04.006. 2

[CWW13] CRANE K., WEISCHEDEL C., WARDETZKY M.: Geodesics in heat: A new approach to computing distance based on heat flow. *ACM Trans. Graph. 32*, 5 (Oct. 2013), 152:1–152:11. doi:10.1145/2516971.2516977. 1, 2, 5

[HAWG08] HUANG Q.-X., ADAMS B., WICKE M., GUIBAS L. J.: Non-rigid registration under isometric deformations. *Computer Graphics Forum 27*, 5 (2008), 1449–1457. doi:10.1111/j.1467-8659.2008.01285.x. 2

[Kar77] KARCHER H.: Riemannian center of mass and mollifier smoothing. *Communications on Pure and Applied Mathematics 30*, 5 (1977), 509–541. doi:10.1002/cpa.3160300502. 2

[KS98] KIMMEL R., SETHIAN J. A.: Computing geodesic paths on manifolds. *Proceedings of the National Academy of Sciences 95*, 15 (1998), 8431–8435. arXiv:http://www.pnas.org/content/95/15/8431.full.pdf. 2

[LRF10] LIPMAN Y., RUSTAMOV R. M., FUNKHOUSER T. A.: Biharmonic distance. *ACM Trans. Graph. 29*, 3 (July 2010), 27:1–27:11. doi:10.1145/1805964.1805971. 2, 7

[MMP87] MITCHELL J. S. B., MOUNT D. M., PAPADIMITRIOU C. H.: The discrete geodesic problem. *SIAM J. Comput. 16*, 4 (Aug. 1987), 647–668. doi:10.1137/0216045. 2, 5

[QH07] QIU H., HANCOCK E.: Clustering and embedding using commute times. *Pattern Analysis and Machine Intelligence, IEEE Transactions on 29*, 11 (Nov. 2007), 1873–1890. doi:10.1109/TPAMI.2007.1103. 2

[SRGB14] SOLOMON J., RUSTAMOV R., GUIBAS L., BUTSCHER A.: Earth mover's distances on discrete surfaces. *ACM Trans. Graph. 33*, 4 (July 2014), 67:1–67:12. doi:10.1145/2601097.2601175. 7

[XW09] XIN S.-Q., WANG G.-J.: Improving chen and han's algorithm on the discrete geodesic problem. *ACM Trans. Graph. 28*, 4 (Sept. 2009), 104:1–104:8. doi:10.1145/1559755.1559761. 2

[XW10] XIN S.-Q., WANG G.-J.: Applying the improved chen and han's algorithm to different versions of shortest path problems on a polyhedral surface. *Comput. Aided Des. 42*, 10 (Oct. 2010), 942–951. doi:10.1016/j.cad.2010.05.009. 2

[XWL*15] XU C., WANG T., LIU Y., LIU L., HE Y.: Fast wavefront propagation (fwp) for computing exact geodesic distances on meshes. *Visualization and Computer Graphics, IEEE Transactions on 21*, 7 (July 2015), 822–834. doi:10.1109/TVCG.2015.2407404. 2

[XYH12] XIN S.-Q., YING X., HE Y.: Constant-time all-pairs geodesic distance query on triangle meshes. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* (New York, NY, USA, 2012), I3D '12, ACM, pp. 31–38. doi:10.1145/2159616.2159622. 1, 2, 5, 6