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

Oliver Burghard and Reinhard Klein

Oct. 5th 2015

Where do we need geodesic distances ?

1. Karcher means

$$\mathrm{argmin}_x \sum w_i d^2(p_i, x)$$

[1] Huang et al.: Non-rigid registration under isometric deformations. CGF'08
[2] Kokkinos et al.: Intrinsic shape context descriptors for deformable shapes, CVPR'12

Where do we need geodesic distances ?

1. Karcher means

$$\mathrm{argmin}_x \sum w_i d^2(p_i, x)$$

2. Intrinsic k-means

$$\mathrm{argmin}_{\pi, p_i} \sum d(p_{\pi(x)}, x)$$

[1] Huang et al.: Non-rigid registration under isometric deformations. CGF'08

[2] Kokkinos et al.: Intrinsic shape context descriptors for deformable shapes, CVPR'12

Where do we need geodesic distances ?

1. Karcher means

$$\text{argmin}_x \sum w_i d^2(p_i, x)$$

2. Intrinsic k-means

$$\text{argmin}_{\pi, p_i} \sum d(p_{\pi(x)}, x)$$

3. Consistency of Correspondences[1]

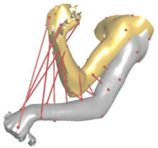

[1] Huang et al.: Non-rigid registration under isometric deformations. CGF'08

[2] Kokkinos et al.: Intrinsic shape context descriptors for deformable shapes, CVPR'12

Where do we need geodesic distances ?

1. Karcher means $\quad \text{argmin}_x \sum w_i d^2(p_i, x)$
2. Intrinsic k-means $\quad \text{argmin}_{\pi, p_i} \sum d(p_{\pi(x)}, x)$
3. Consistency of Correspondences[1]

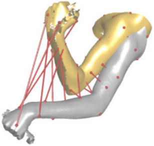

4. Exponential Maps[2]

[1] Huang et al.: Non-rigid registration under isometric deformations. CGF'08

[2] Kokkinos et al.: Intrinsic shape context descriptors for deformable shapes, CVPR'12

Where do we need geodesic distances ?

What do those use cases have in common?

1. All-pair distances for any pairs
2. Small distances/local distances are most important
3. Good approximations are sufficient
4. Distance computations take a significant fraction of execution time

We are not the first to care about distance fields

Mesh with n vertices, distances up to length ρ, m vertices are closer to query point then ρ. Current methods:

1. Window propagation[3]: $\underline{\mathcal{O}(m^2)}$, exact, optimal, can be local
2. Fast Marching[4]: $\underline{\mathcal{O}(m \log(m))}$, approximation, can be local
3. Geodesics in heat[5]: $\mathcal{O}(n)$ + preprocessing, approximation, <u>always global</u>
4. All-pairs-shortest-paths[6]: $\mathcal{O}(1)$ + preprocessing, approximative, <u>not continuous</u>
5. Our method: $\mathcal{O}(1)$ + preprocessing, approximative, *continuous*

---

[3] Chen et al.: Shortest paths on a polyhedron, SCG'90

[4] Kimmel et al.: Computing geodesic paths on manifolds, PNAS'98

[5] Crane et al.: Geodesics in heat, ToG'13

[6] Xin et al.: Constant-time all-pairs geodesic distance query on triangle meshes, I3D'12

Constant-time all-pairs geodesic distance query on triangle meshes

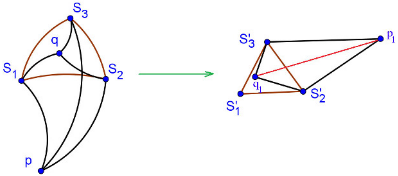

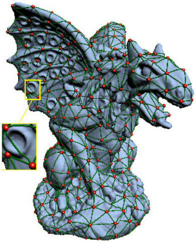

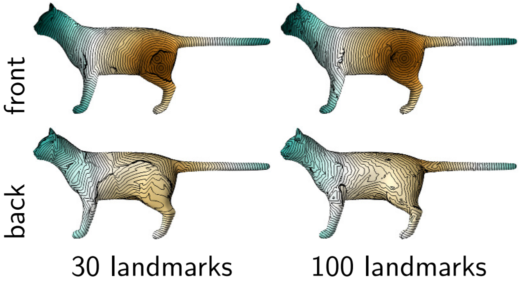

With few landmarks non-continuity when changing coarse triangles has huge impact on the quality of the result.

Is approximation of geodesic distances really so hard?

What can we infer from landmarks onto other distances?

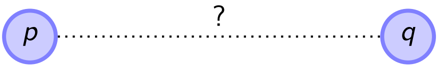

Is approximation of geodesic distances really so hard?

What can we infer from landmarks onto other distances?

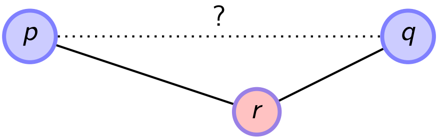

Is approximation of geodesic distances really so hard?

What can we infer from landmarks onto other distances?

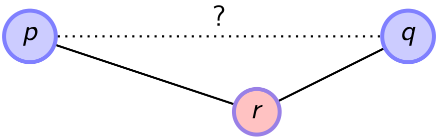

Triangle equation gives an upper bound:

$$d(p, q) \leq d(p, r) + d(r, q)$$

Is approximation of geodesic distances really so hard?

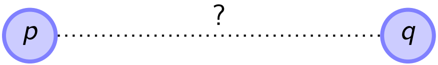

Is approximation of geodesic distances really so hard?

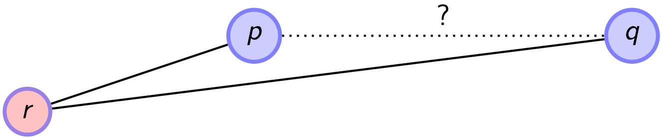

Is approximation of geodesic distances really so hard?

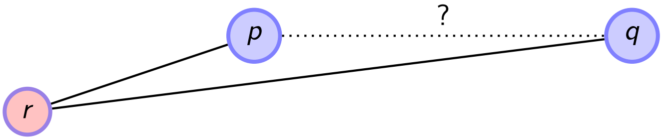

Triangle equation gives an *lower* bound:

$$d(q, r) \leq d(p, r) + d(p, q)$$
$$d(p, q) \geq d(q, r) - d(p, r)$$

What can we know from the triangle equation *alone*?

Therefore one can infer *lower* and *upper* bounds from a single landmark

$$|d(p,r) - d(r,q)| \quad \leq \quad d(p,q) \quad \leq \quad d(p,r) + d(r,q)$$

What can we know from the triangle equation *alone*?

Therefore one can infer *lower* and *upper* bounds from a single landmark

$$|d(p,r) - d(r,q)| \quad \leq \quad d(p,q) \quad \leq \quad d(p,r) + d(r,q)$$

And for a set of R landmarks:

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

(Maximal lower bound and minimal upper bound)

How do they look like?

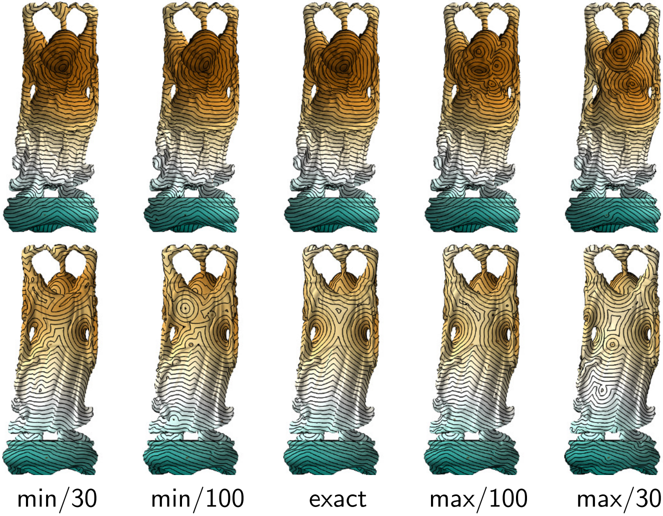

Properties

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

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

Strictly positive iff for $p \neq q$ there is r with $d(p, r) \neq d(q, r)$. For a single landmark conflicting points build a subset $\{x|d(x, r) = d(p, r) \forall r \in R\}$, that typically is a sub-manifold of co-dimension 1. On real world data the intersection of such sets is $\{p\}$ for $|R| > 3$.

- d_{min} is mostly a distance metric

A closer look at involved errors

For the absolute errors $e_{min}(p, q) := d(p, q) - d_{min}(p, q)$ and $e_{max}(p, q) := d_{max}(p, q) - d(p, q)$ we have (due to the triangle equation):

$$e_{min}(p, q) \leq e_{min}(p', q) + 2\, d(p, p')$$
$$e_{max}(p, q) \leq e_{max}(p', q) + 2\, d(p, p')$$

If $V_{min}(p)$ and $V_{max}(p)$ are the sets where d_{min} and d_{max} are exact:

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

How do V_{min} and V_{max} look like?

V_{min} and V_{max} on an Euclidean line

Euclidean line

After fixing p and r, the line is divided at r into two, at one part d_{min} is exact (blue), on the other d_{max} (orange). Adding more landmarks those areas will overlap and there are regions where both are precise.

Note: V_{min} is simply the union of all V^r_{min} and analog V_{max}.

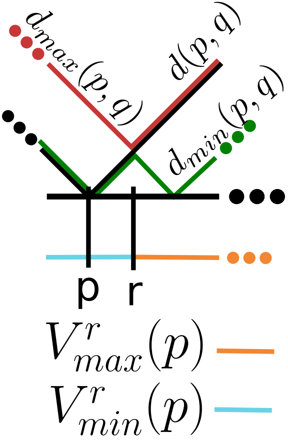

V_{min} and V_{max} on a circle

Euclidean circle

Things become more difficult as *topology* has its influences.
Let R^+ denote the maximal shortest path starting in r including p, R^- be the opposite maximal shortest path on the same geodesic. Define $P^{+/-}$ analog and P as $P^+ \cup P^-$. Then (paper)

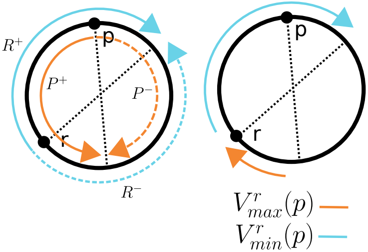

$$V^r_{min}(p) = R^+ \cap P = R^+$$
$$V^r_{max}(p) = R^- \cap P$$

V_{min} and V_{max} on a surface

Surface

At a surface there are *topological* and *tangential* errors.

Tangential errors stem from the distance of point q to P.

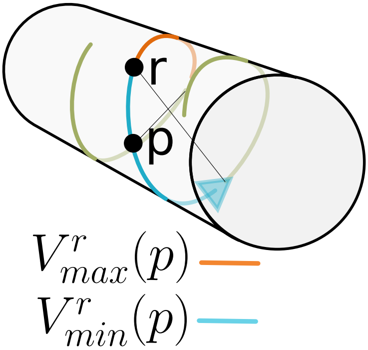

Maximal shortest paths vs V_{min}

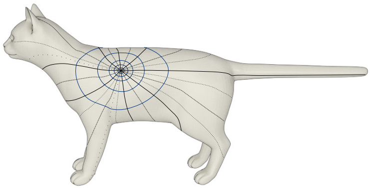

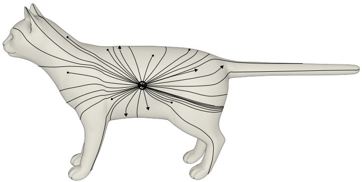

Left is an illustration of the maximal shortest paths starting in some query point. If those were known, distances were known. Instead only a subset of those shortest paths is known. The subset are the paths induced from the landmarks (leading to tangential errors) and on those paths only a subset is known (topological errors).

Example: Absolute errors on flat domains

In the euclidean plane d_{min} is precise for few landmarks, as V_{min} gets denser. d_{max} keeps a large error far from the landmarks.
On the torus the topological error keeps d_{min} from converging close to the cut locus, where d_{max} converges still.

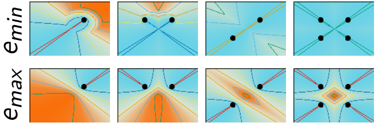

Euclidean plane

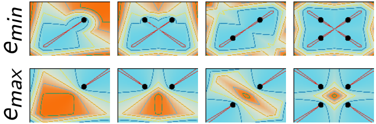

Torus

Reviewing the results

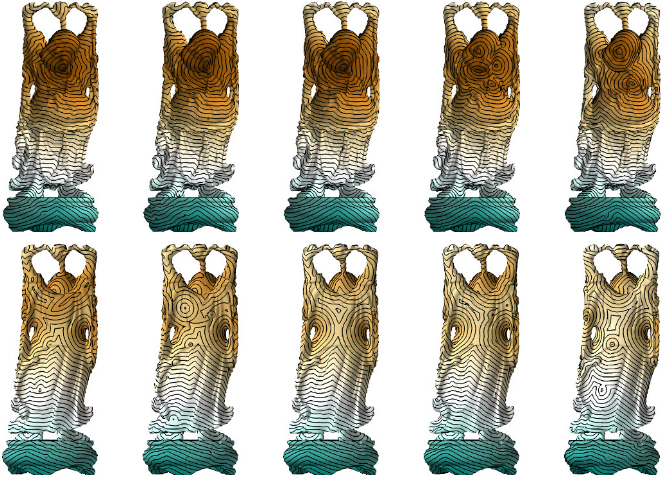

min/30 min/100 exact max/100 max/30

Quantitative approximation errors

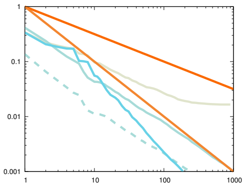

Tosca Cat

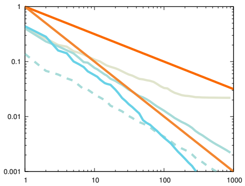

Happy Buddha

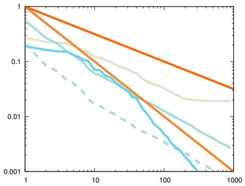

Stanford Dragon

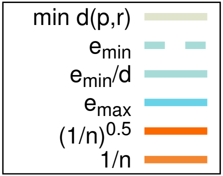

Thank you for your attention

Questions?

Relative errors

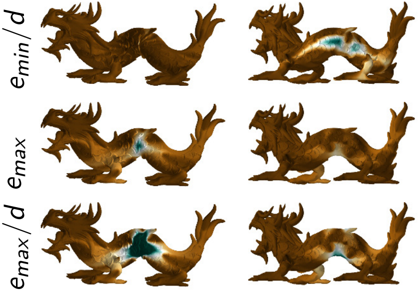

Color scale for relative plots: 0%-35%. Note, that e_{max}/d is unbounded. Relative errors of d_{min} are bounded on compact meshes (see paper).

Future work

1. Utilize multiple landmarks for each point to have
   - better convergence
   - more stability
   - while retaining $\mathcal{O}(1)$ per query
2. Optimize landmark positions