Quantifying shape using the medial axis

Erin Wolf Chambers
Saint Louis University

Based on collaborations with Tao Ju, David Letscher, and Chris Topp

Work supported by NSF grants IIS-1319573 and DBI-1759807

SAINT LOUIS UNIVERSITY

Higher purpose. Greater good.™

The medial axis

- The medial axis was first introduced by Blum in 1967:
  - The set of points with more than one closest point on the boundary
  - Can also be thought of as the set of quench sites of a fire started on the boundary of the shape which burns inward at uniform speed

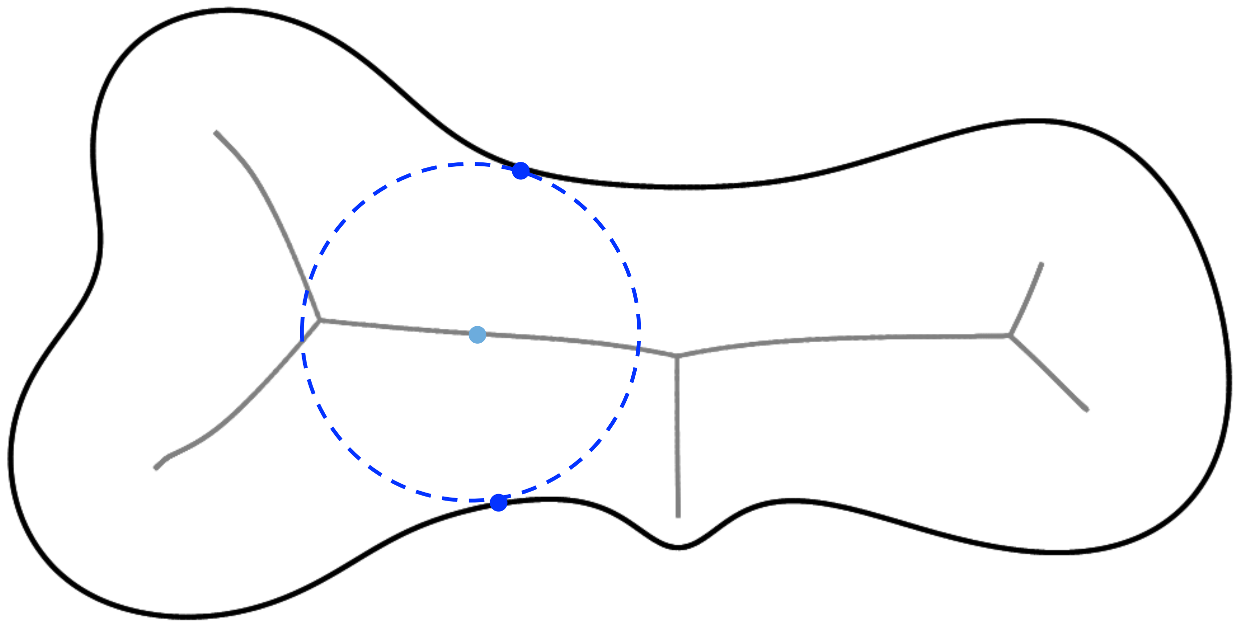

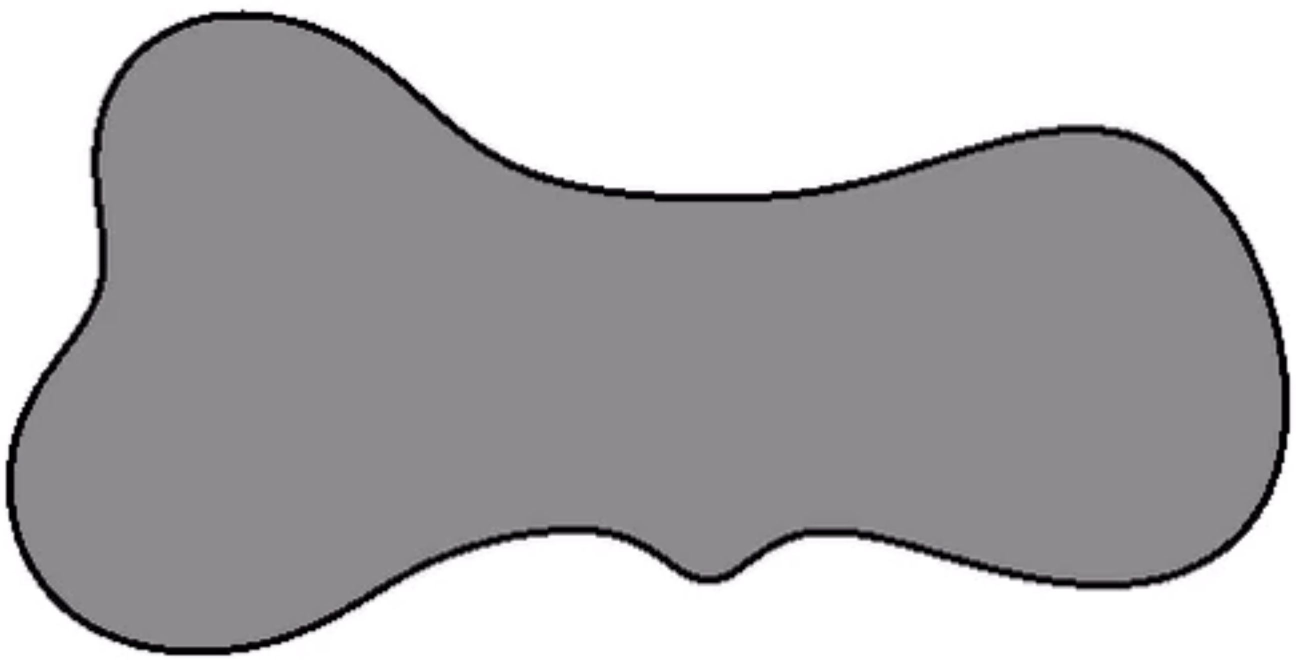

Some good properties

- Co-dimension at least 1
- Has the same topology as the original shape [Lieutier 2003]
- Central to the shape

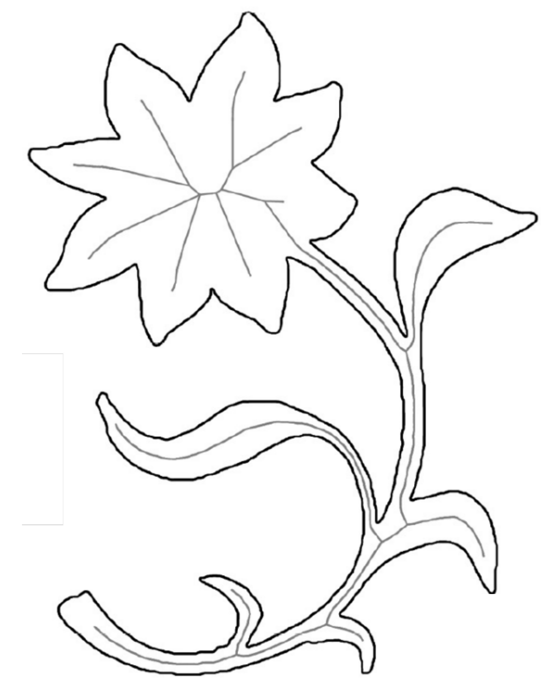

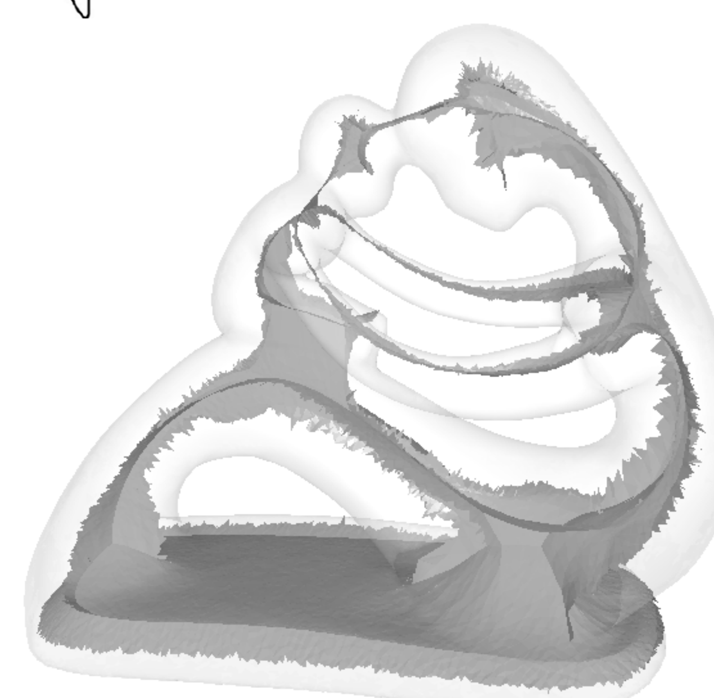

And some bad

- Very sensitive to boundary perturbations
- Can be difficult to compute in 3d
- Can have portions of different dimensions (so not always co-dimension 1)

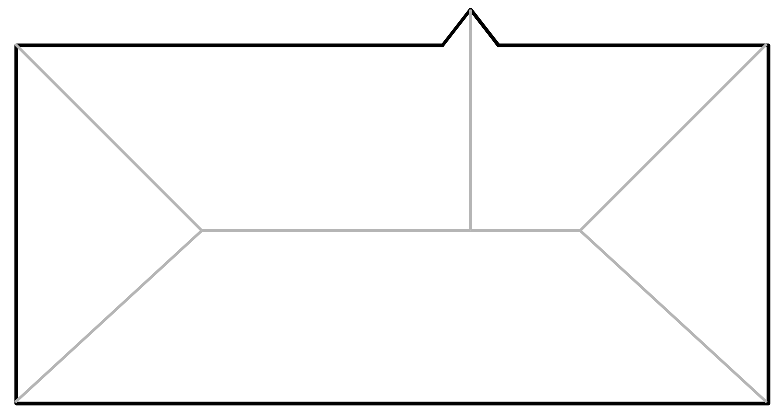

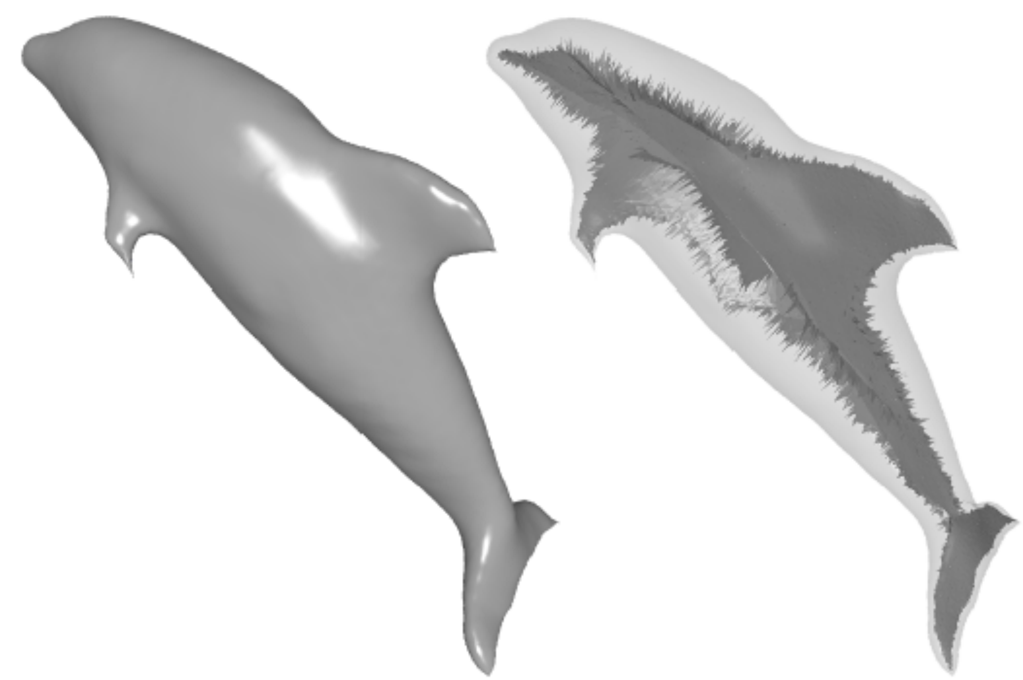

Our goal

- Use the medial axis as a basis for comparing two shapes.

- Main approach: compute a skeleton of a shape, as well as some relevant measures on the skeleton

    - Also has applications in shape modeling and shape segmentation

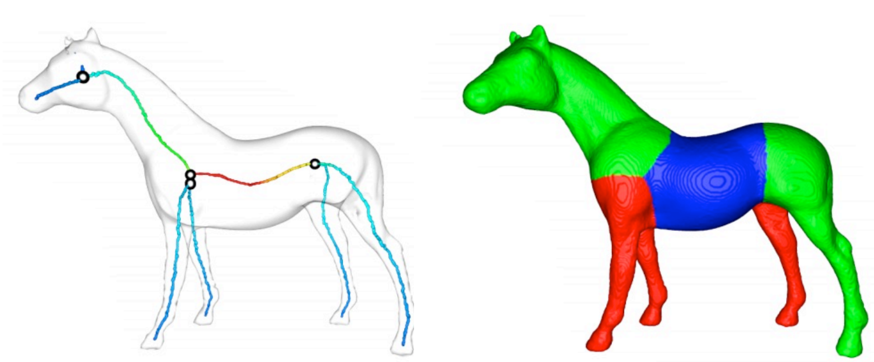

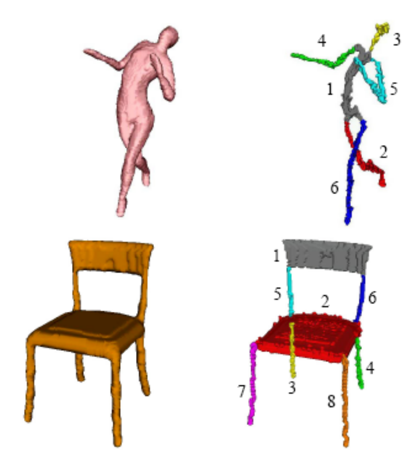

Significance measures

Significance measures can be used to prune the medial axis to retain only "significant" portions of it.

A few examples in prior work:

- Object angle [Attali 96, Amenta 01, Dey 04, Foskey 05, Sud 05]

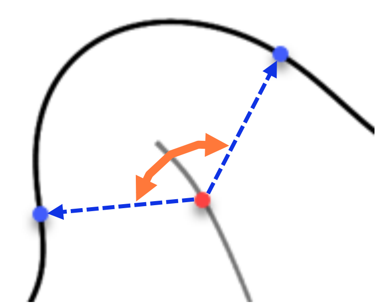

- Circumradius, or distance between the 2 nearest points on the boundary [Chazal 04, Chaussard 09] - used for the lambda medial axis

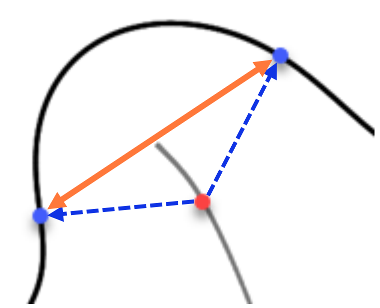

Back to skeletons

However, pruning based on these measures does not maintain the topology (as with object angle), or can cut off significant portions of the skeleton (as with circumradius).

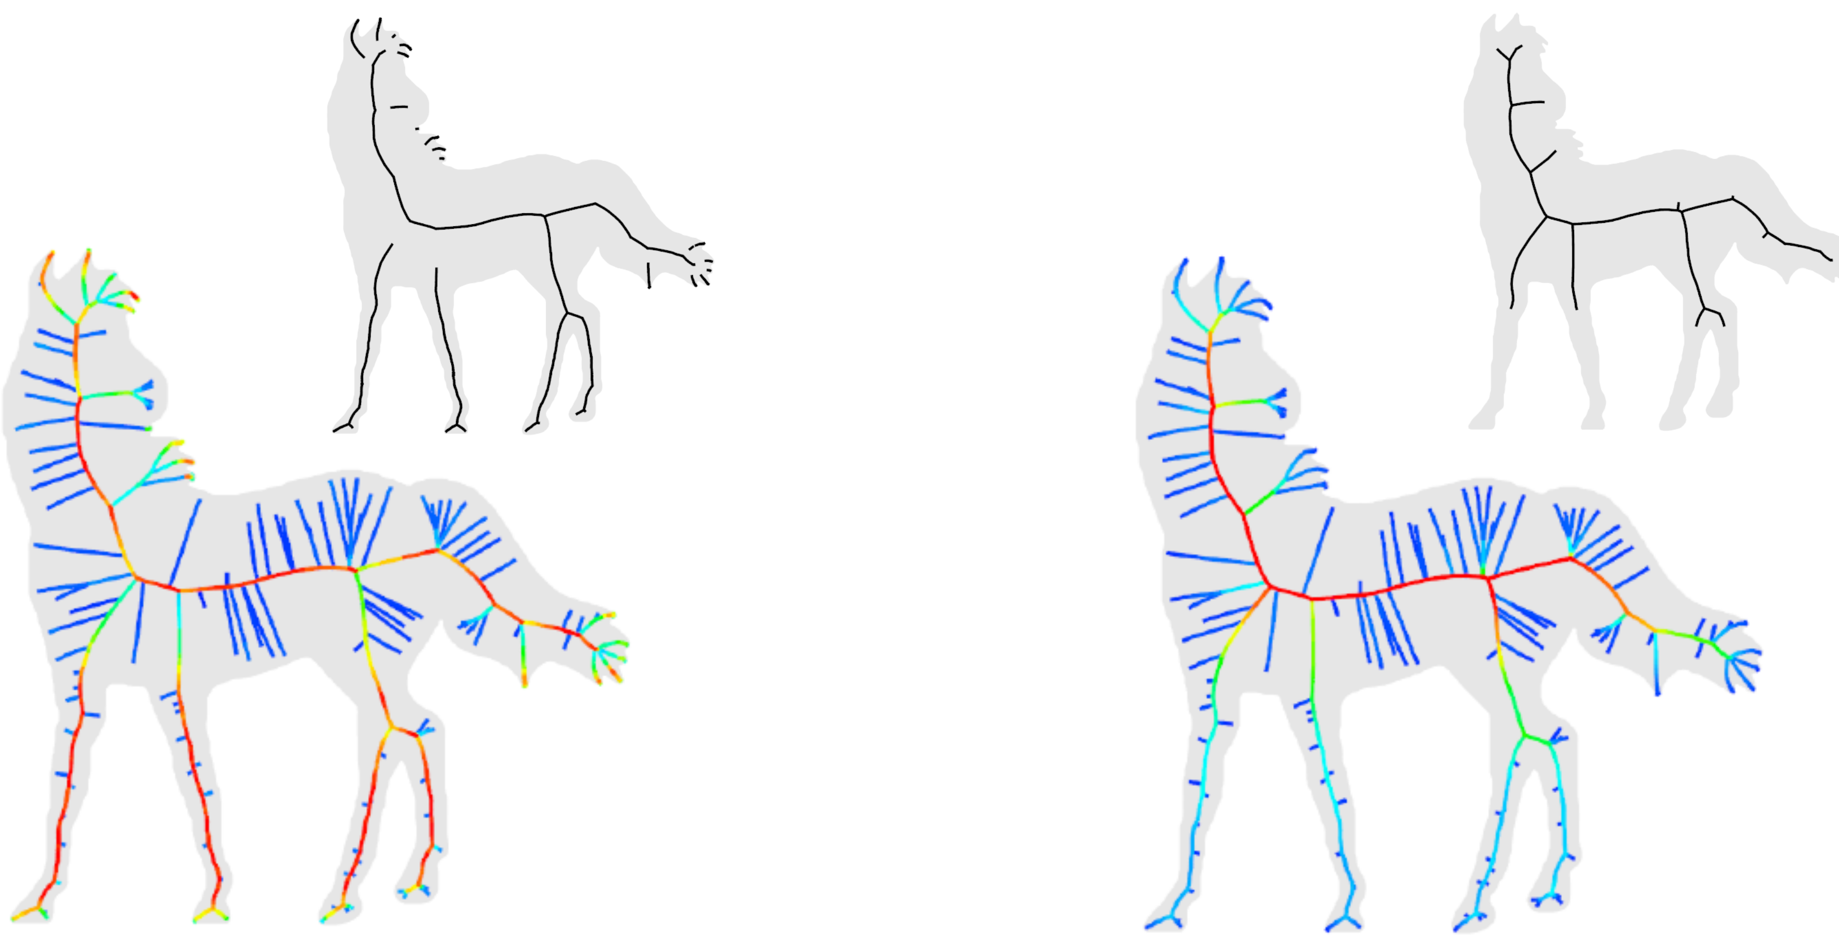

Potential residue

- The potential residue [Ogniewicz 1992] at a medial axis point is the shortest distance on the boundary between the two nearest boundary points to x.
- This function captures global features nicely, and can be generalized to 3d (the medial geodesic function [Dey-Sun 2006]).

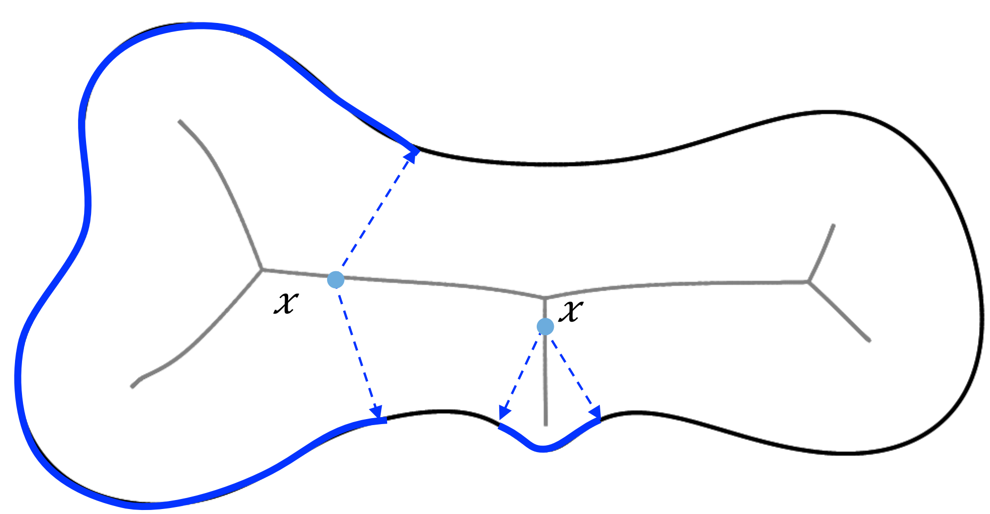

Medial geodesic function

- The natural generalization of potential residue to 3d is called the medial geodesic function [Dey-Sun 2006].
- While it has been implemented, the main drawback is the speed of computation: geodesic queries are relatively slow on arbitrary 3d objects.

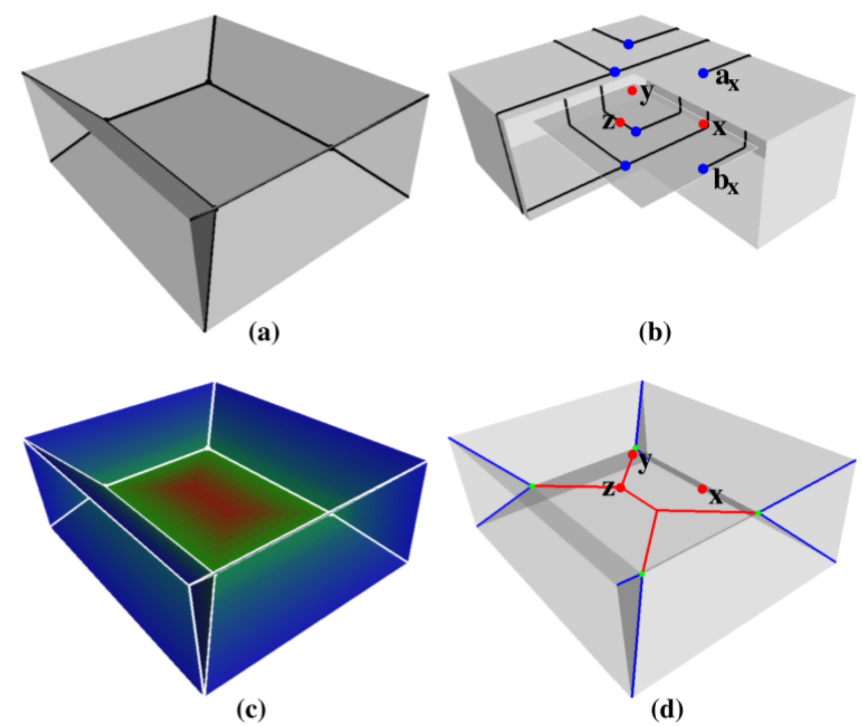

Erosion thickness

- One other significance measure in 2d is erosion thickness [Shaked 1998].
- This is defined as how much the shape erodes as a result of pruning the medial axis.

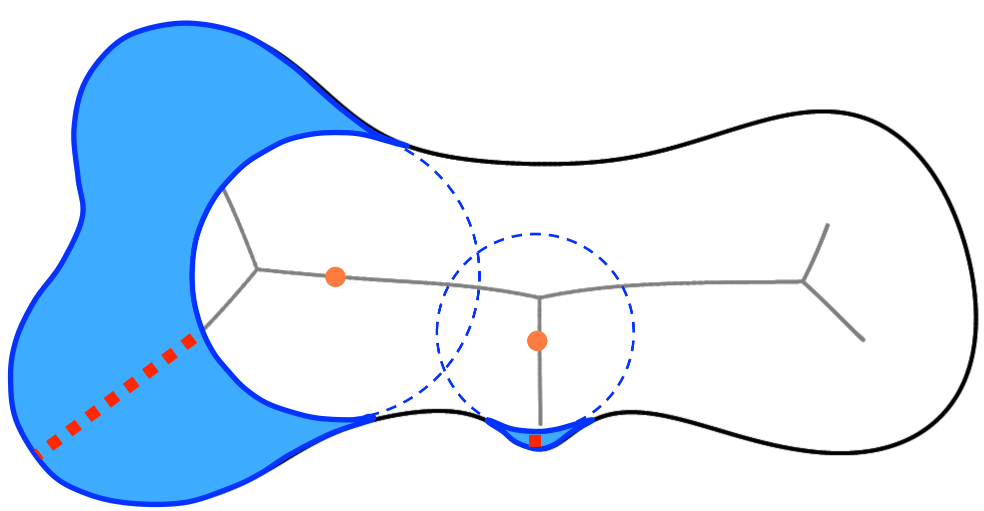

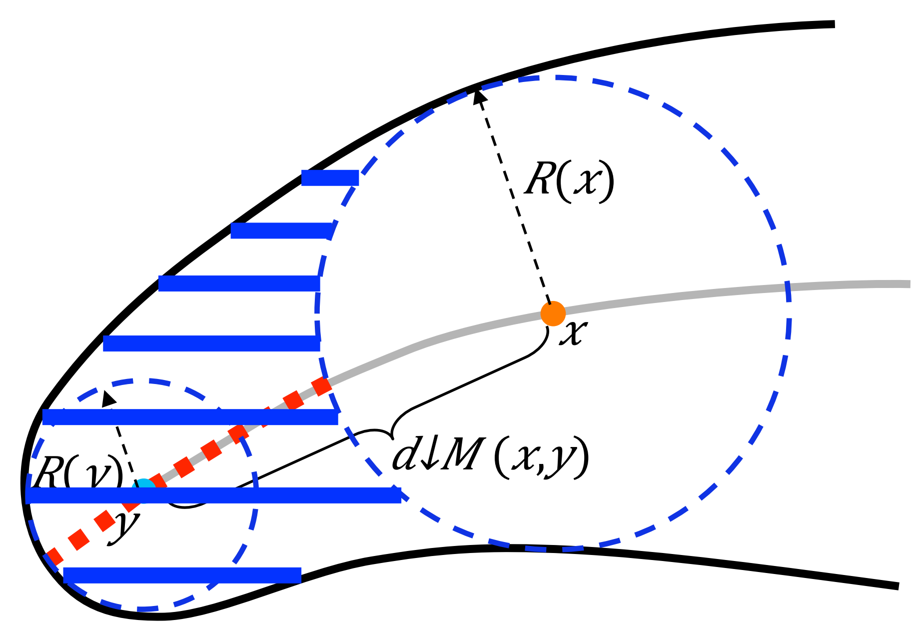

Erosion thickness comparison

Erosion thickness seems more robust to noise, although potential residue is also good if the noise is "random" - only particular examples cause issues.

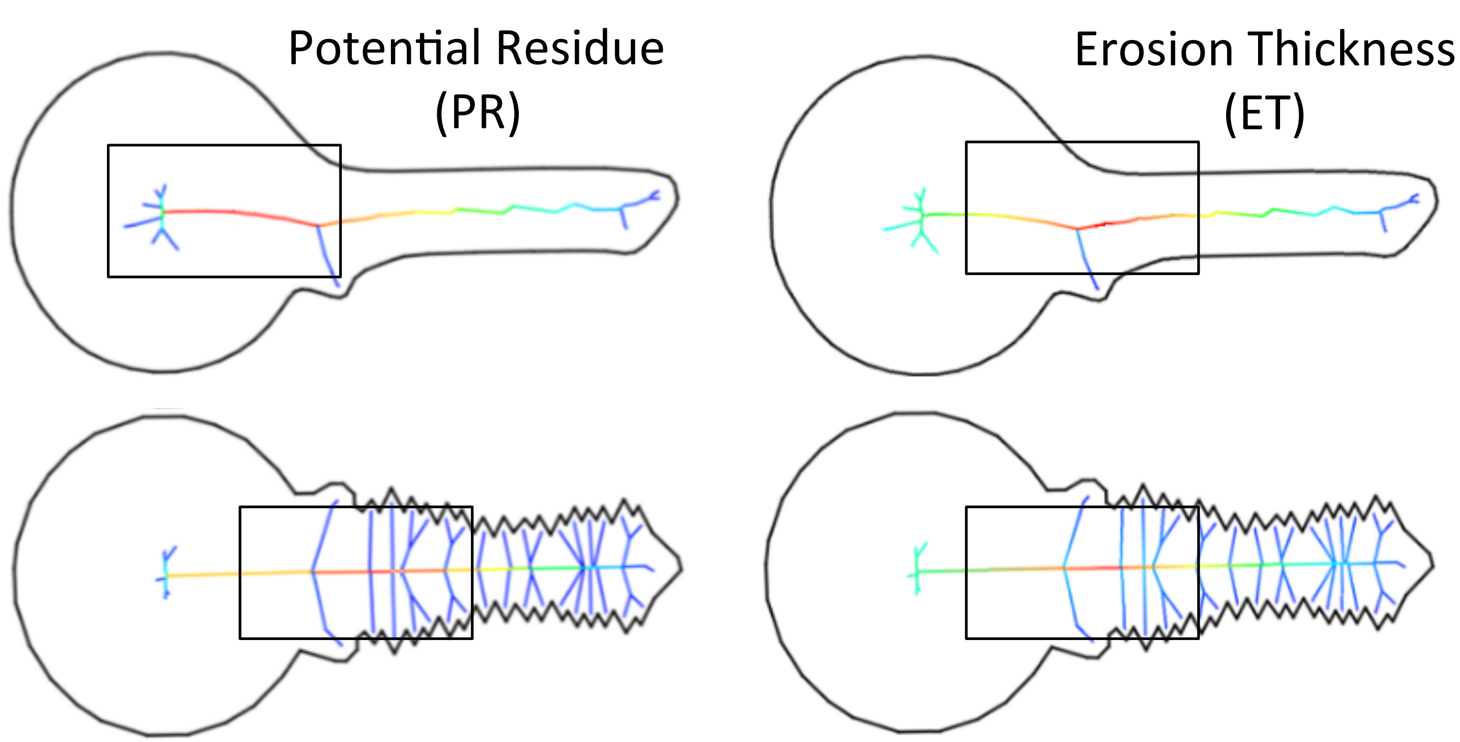

Erosion thickness: downsides

- However, erosion thickness is limited to 2d, as there is not an immediate way to generalize to 3d.

- In addition, there is no explicit definition.

  - It is computed using an iterative pruning process, and hence it is much harder to prove mathematical properties about the quality of the pruning.

The burn time function

In [Liu et al 2011], we define the burn time of a point on the medial axis: the time arrival time of a fire front that is started at all medial axis boundary points, and which dies at interior junctions of the medial axis.

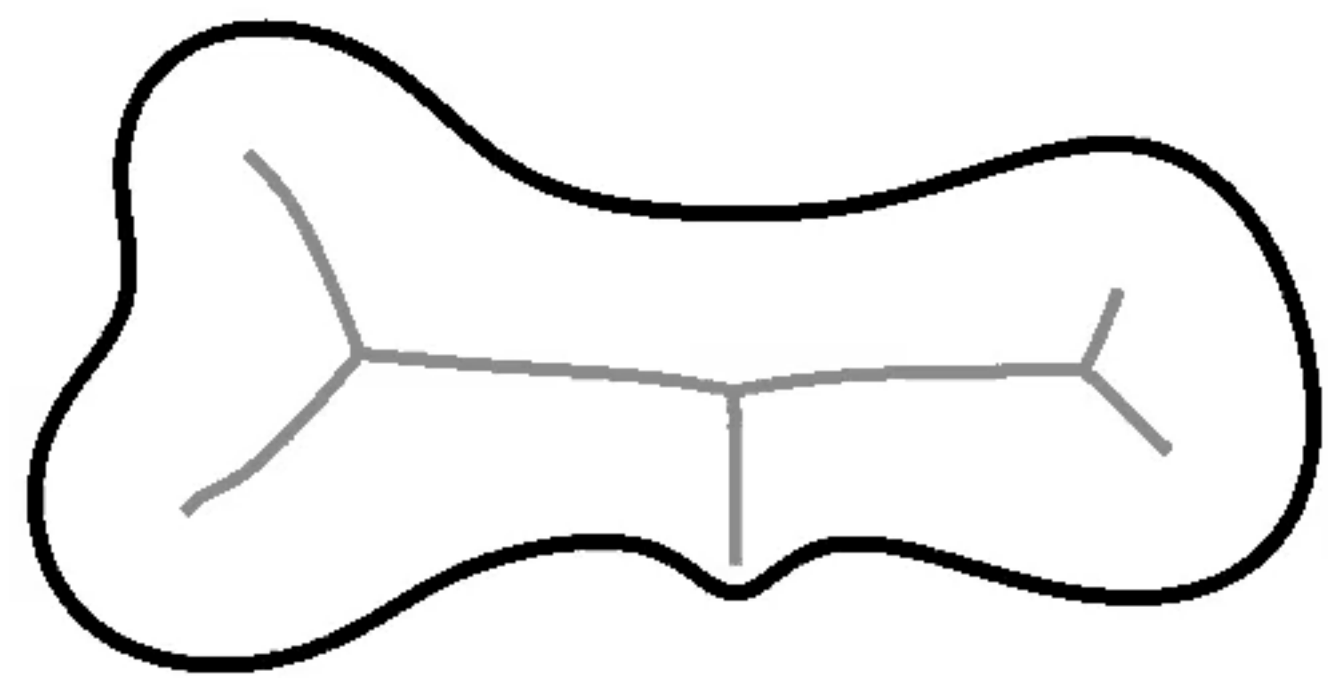

The burn time function

This burn time function (which we originally called the extended distance function) gives a natural way to classify important features in the medial axis, as well as how “central” a point is.

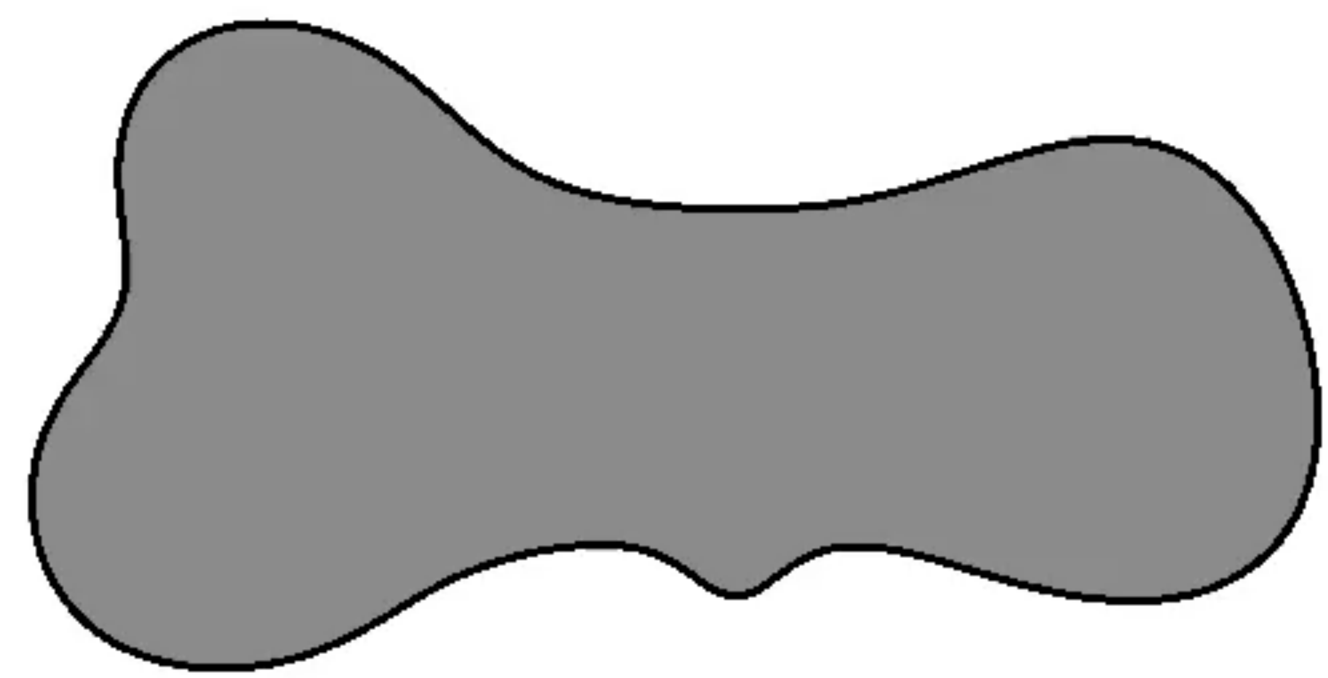

Exposing trees

- Some definitions are needed to formalize the 2d intuition and to generalize it to higher dimensions

- An exposing tree for x is a finite tree contained on the medial axis, where all leaves are on the boundary and the tree must branch when crossing a non-manifold vertex of the medial axis

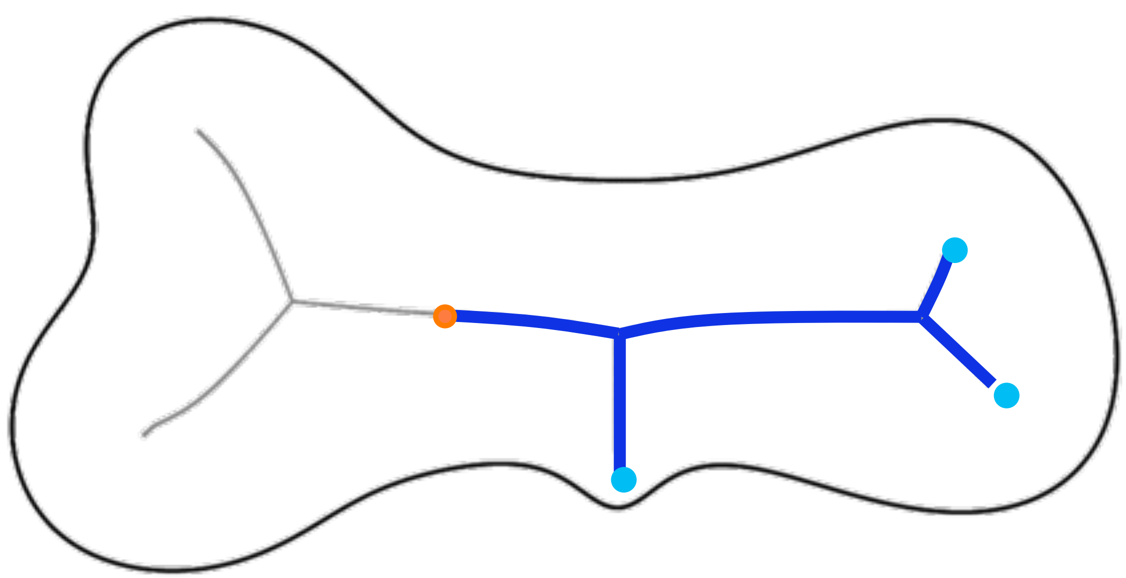

Burn time

- The length of an exposing tree is the longest root to leaf path plus local feature size at the leaf

- The burn time of a point is the minimum over all trees T of length(T).

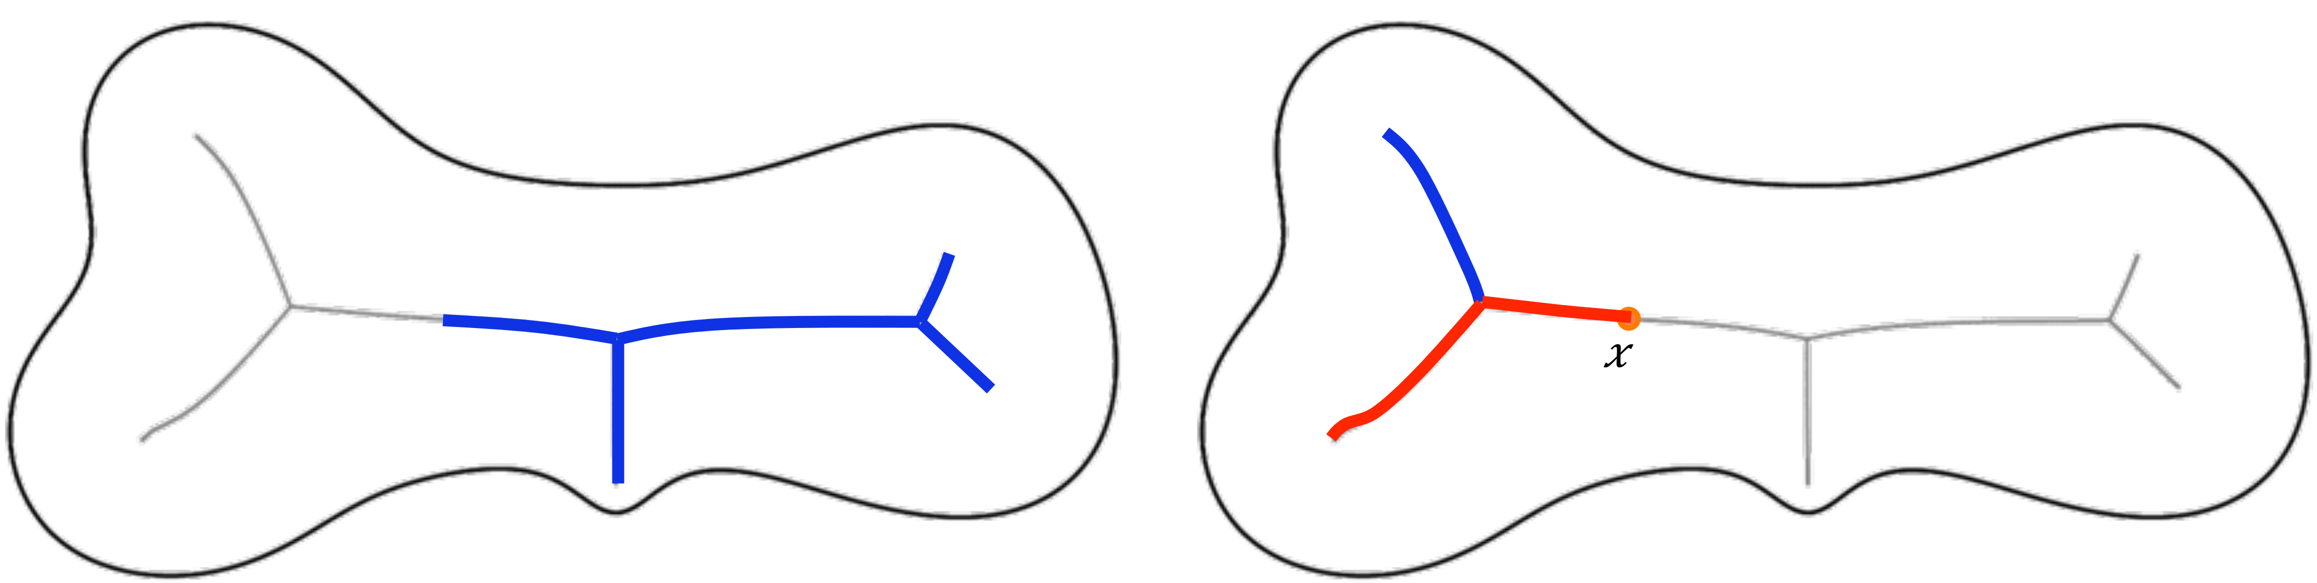

Finiteness

In 2d, we prove that burn time exists and is finite everywhere except the maximally closed sub complex [Liu et al 2011].

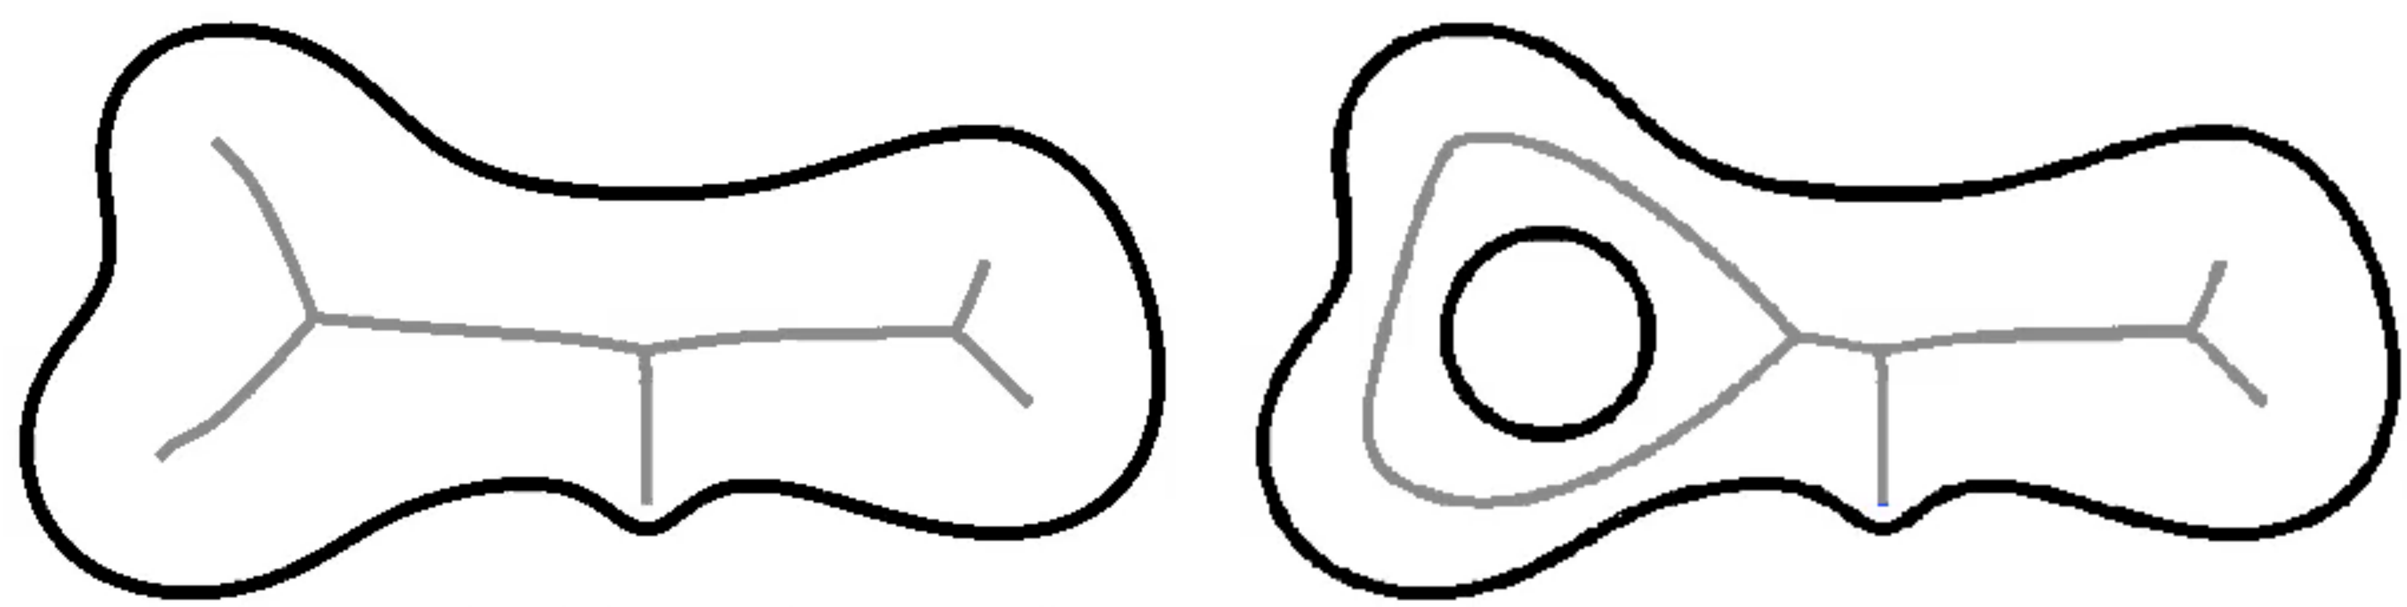

Properties of burn time in 2d

- In [Liu et al], we also prove several nice properties of this function in 2d on simply connected shapes:
  - It is continuous except at branch points; is is upper semi-continuous everywhere.
  - It has no local minima, so is a good tool for finding center points of 2d shapes.

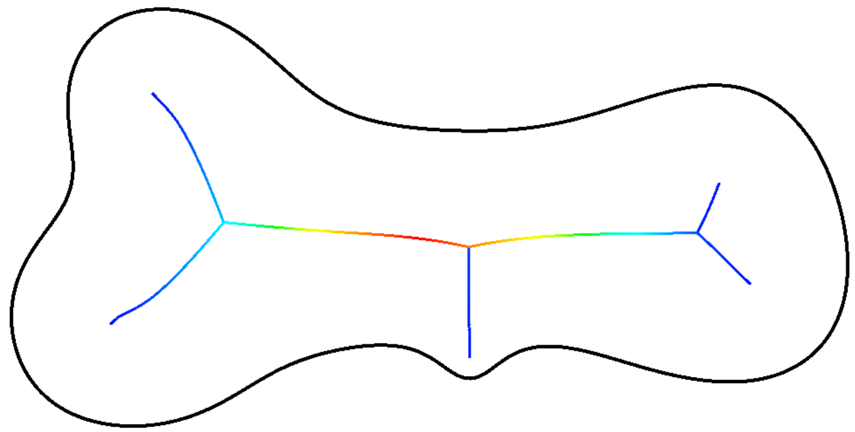

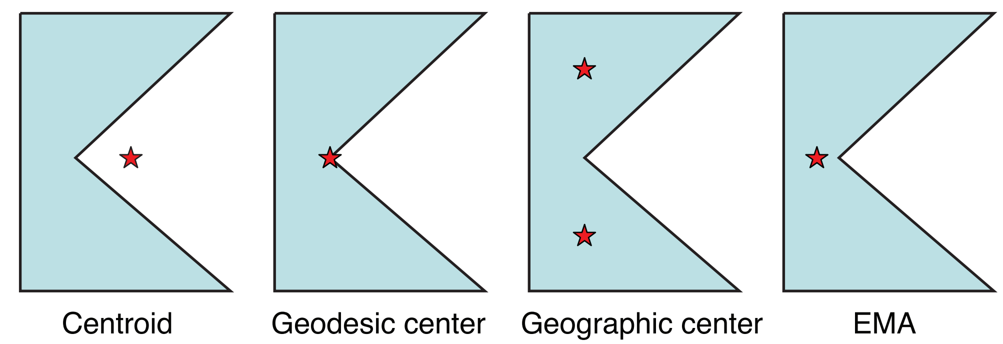

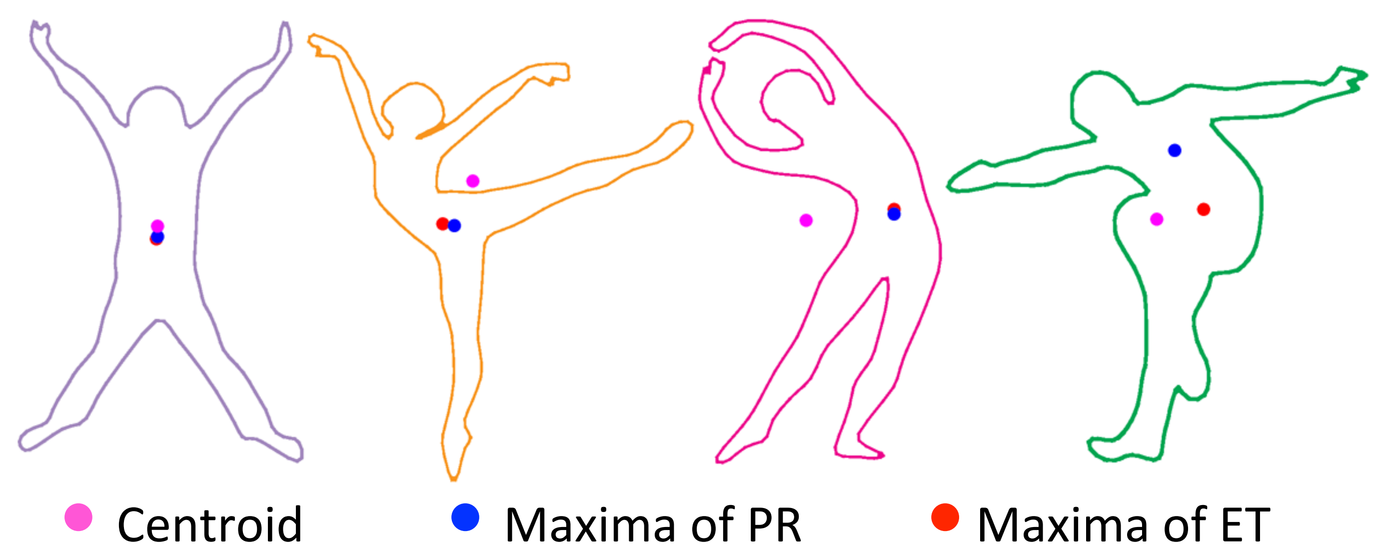

Erosion thickness

Burn time in 2d gives an alternate way to define erosion thickness:

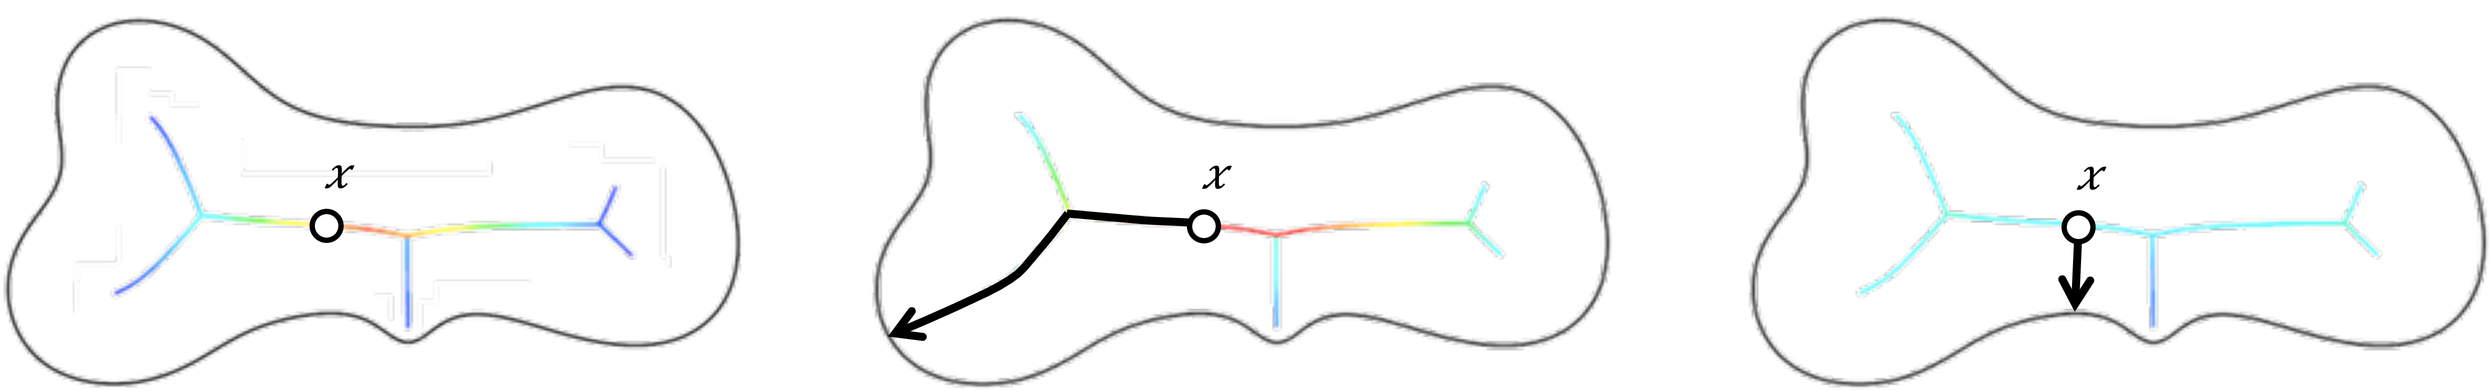

Erosion Thickness $ET(x)$ = Burn Time $BT(x)$ − Radius $R(x)$

"Tubularity" "Length" "Thickness"

Shape alignment application

This can also be used for shape alignment, and is particularly good for articulated shapes:

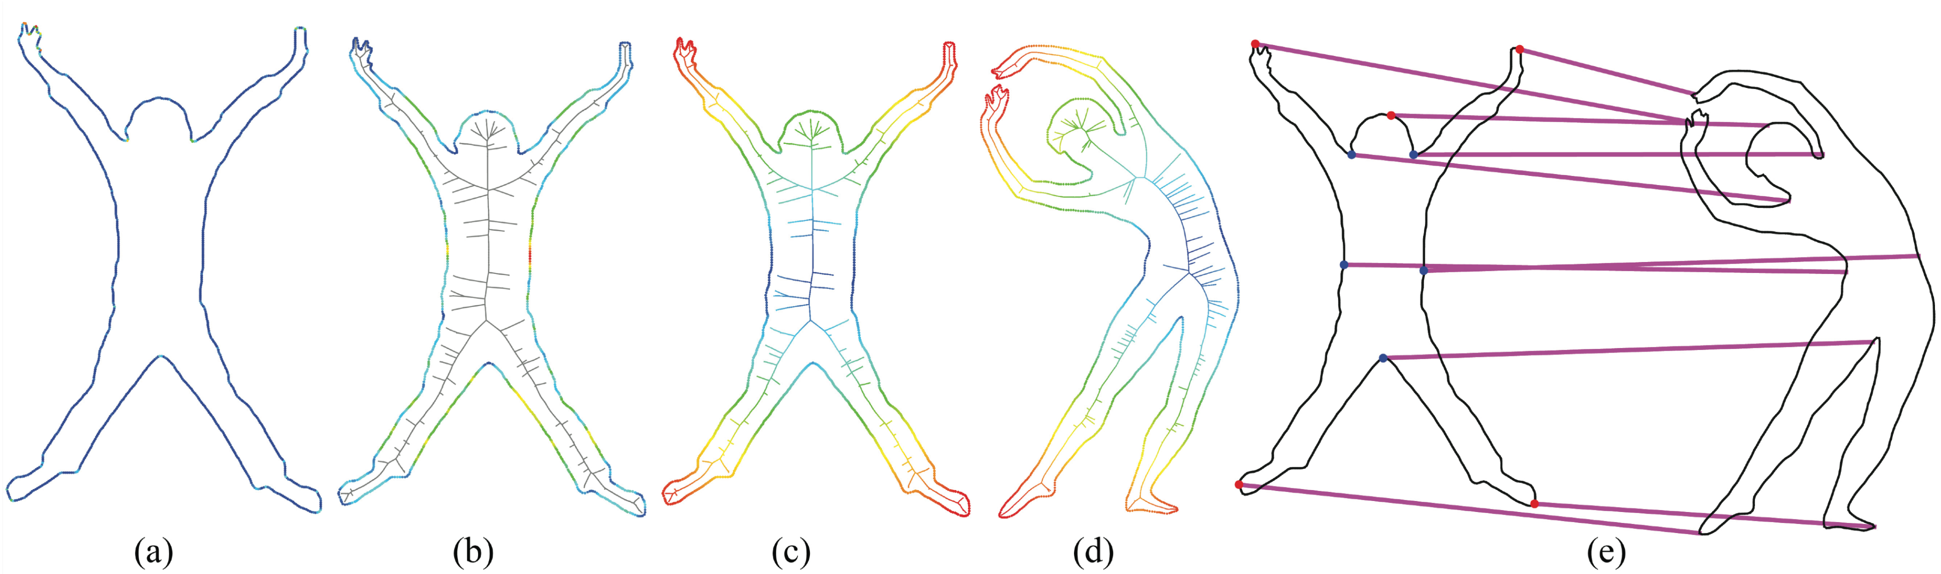

local curvature

local feature size

erosion thickness

matching based on erosion thickness

Comparison in 2d: pruned medial axis

Distance

Object angle

Potential residue

Burn time

One extension: weighted EDF (w-EDF)

In 2d, EDF (or burn time) considers simply the length of the longest tube that can be fit in the shape

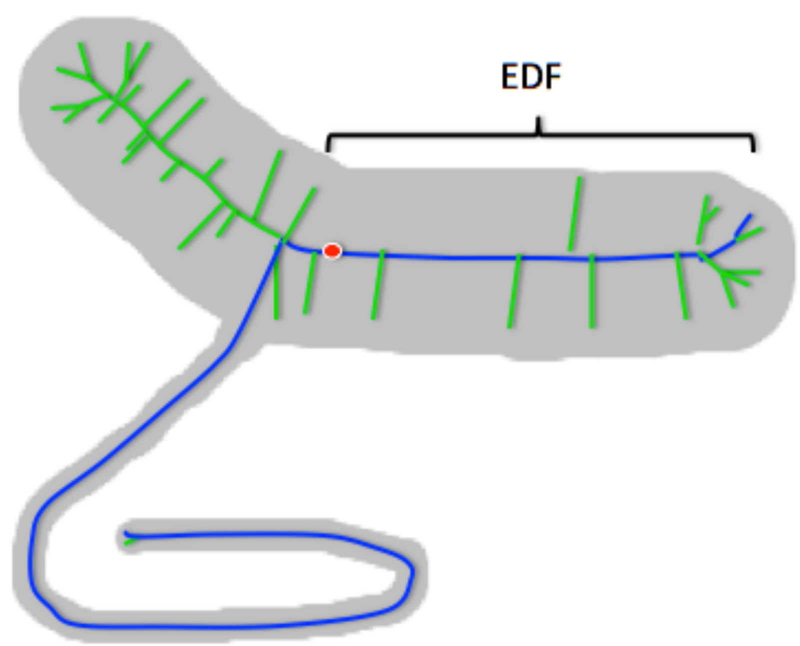

$$EDF(x) = \sup_{f|x \in f} r_f(x)$$

Weighted EDF

W-EDF [Leonard-Morin-Hahmann-Carlier 2016] is a natural extension which weights by area, instead of length:

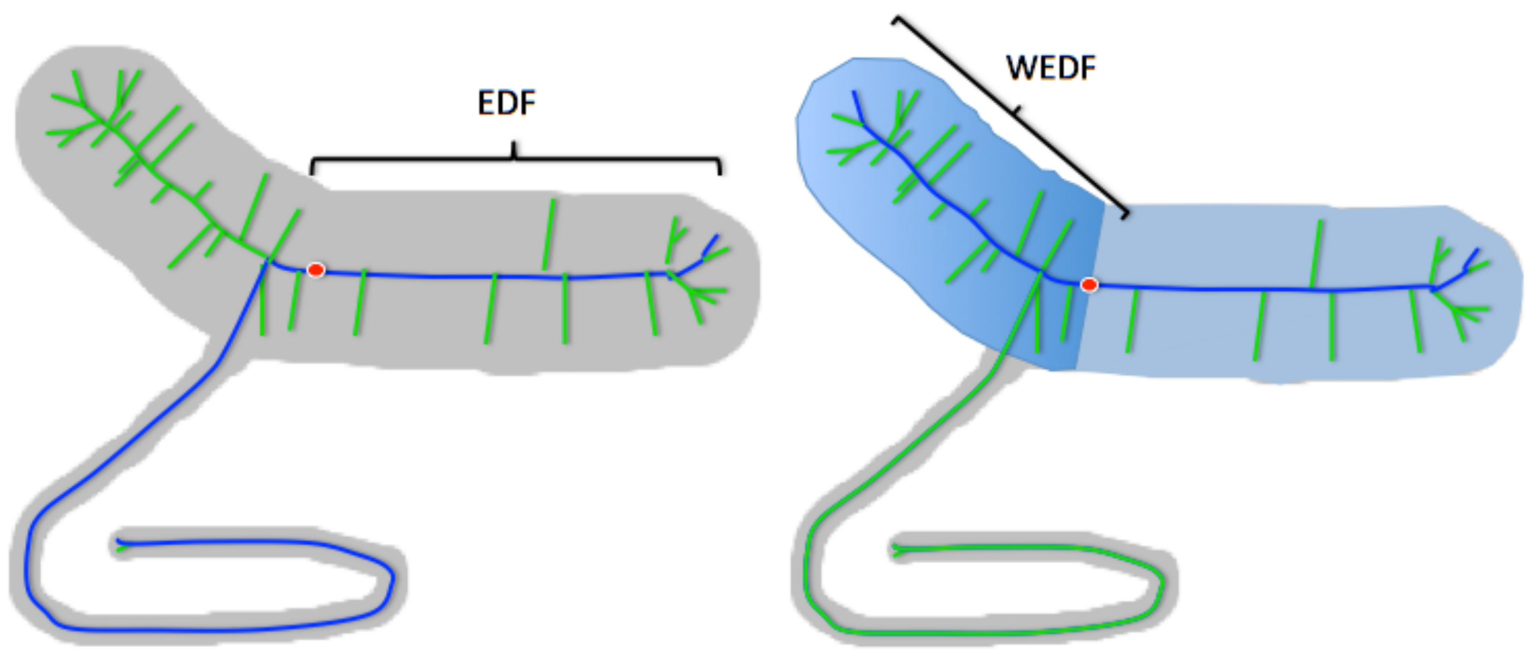

$$EDF(x) = \sup_{f|x \in f} r_f(x) \qquad WEDF(x) = \sup_{f|x \in f} A_f(x).$$

w-EDF motivation

The goal is to identify major parts of an input shape, separating features (or “details”) from the core shape.

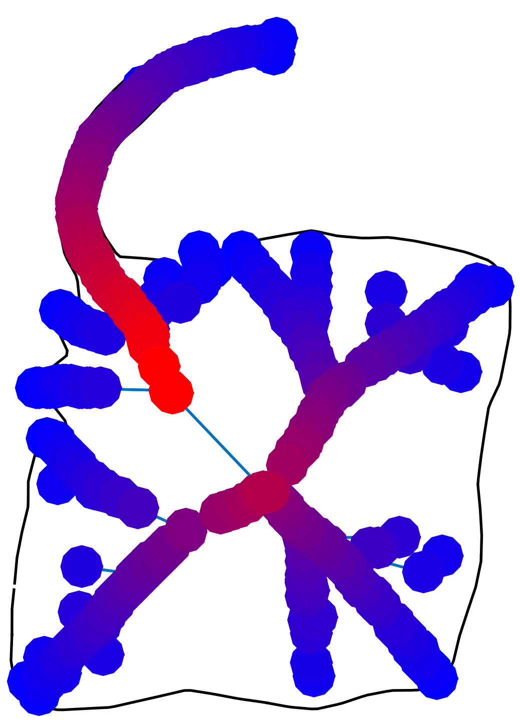

EDF

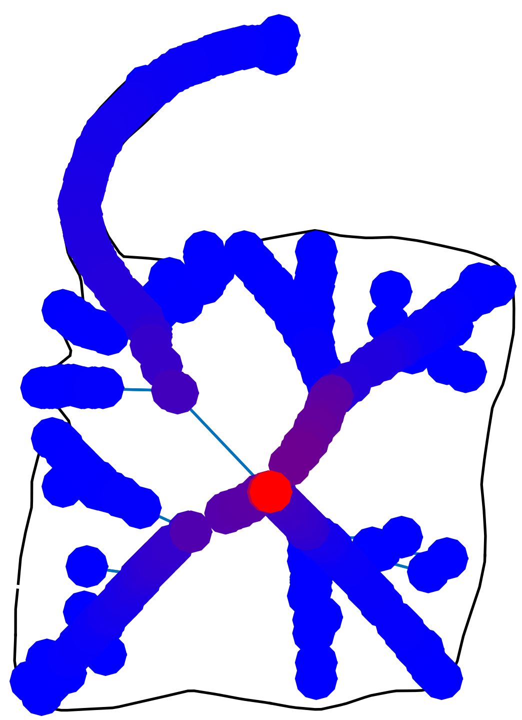

w-EDF

W-EDF decompositions

- Their algorithm continues a “part” across a branch only if the adjacent branch has a very different value.
  - This adds noise robustness.

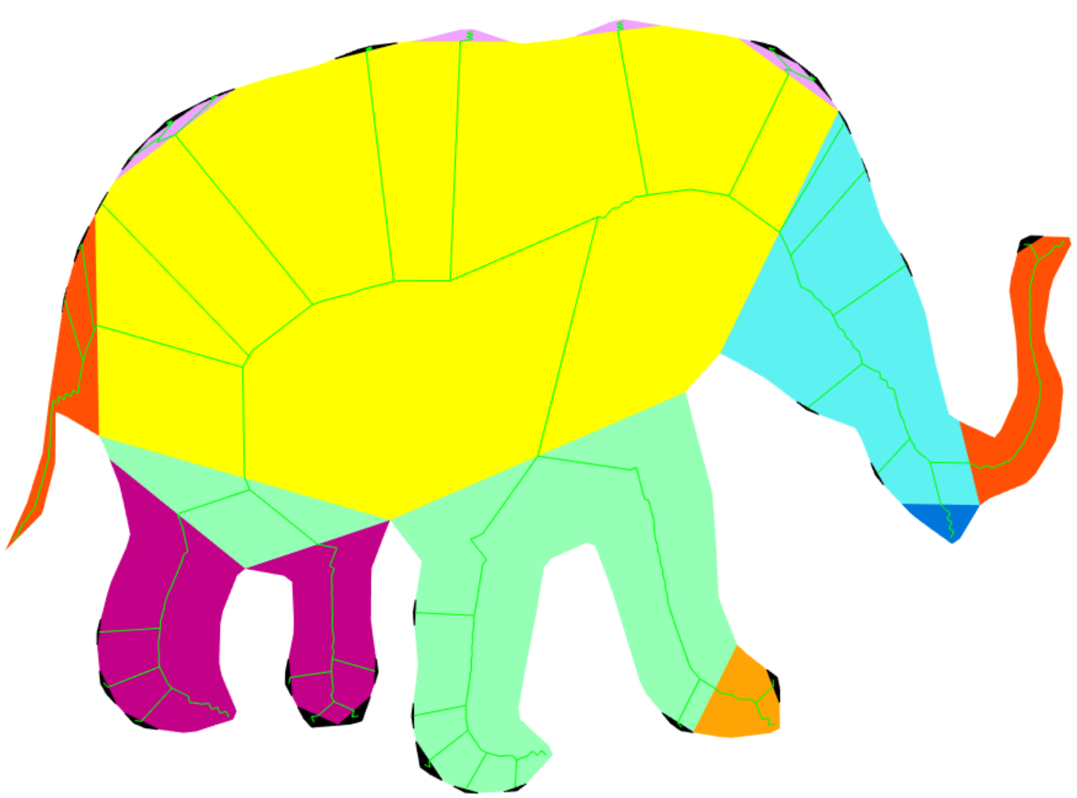
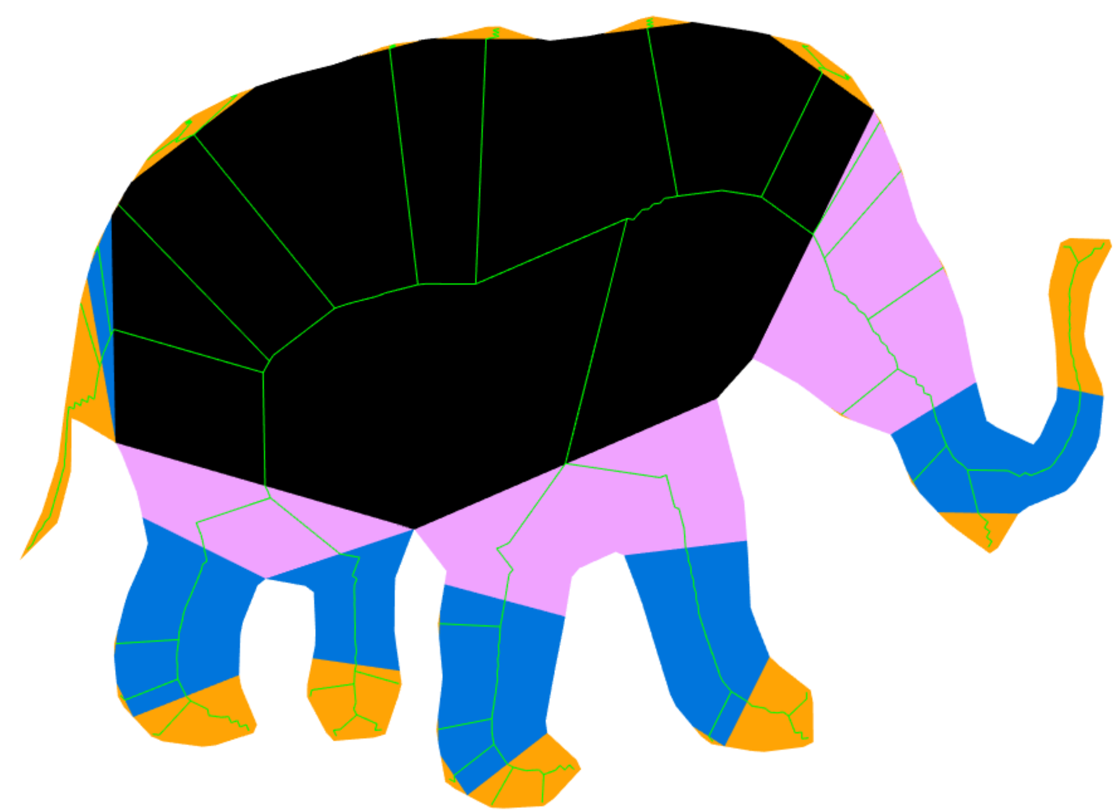

Articulated images

This w-EDF decomposition also turns out to be particularly robust against articulation (when the same shape moves around):

Moving to 3 dimensions

- In 3d, most standard medial axis approximations yield piecewise flat cellular complexes where the local geometry consists of sheets glued along singular curves
- Generically, there are 6 local pictures possible [Giblin Kimia]
- Intuitively, burning will still start at the boundary of the medial axis and proceed inward, but crossing the singular curves is more complex.

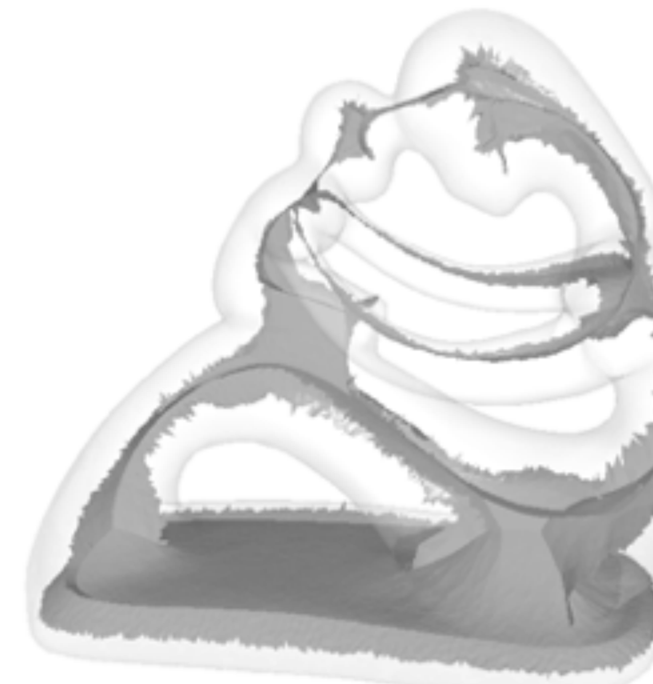

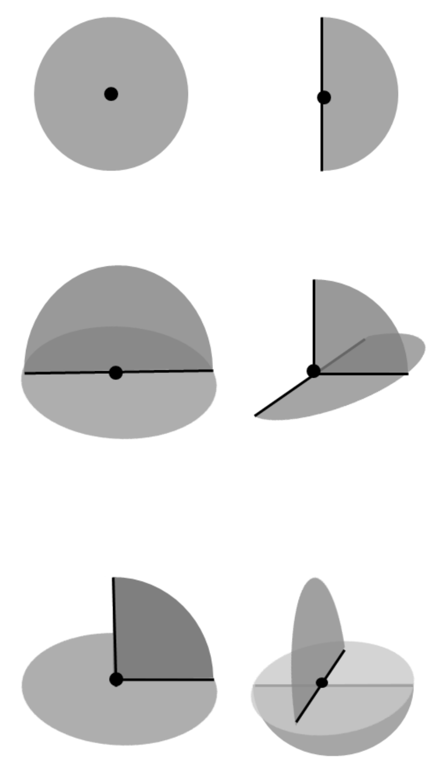

Exposing sets

- We say a point x is exposed in its local neighborhood by a set of adjacent sheets if there is no disk neighborhood remaining:

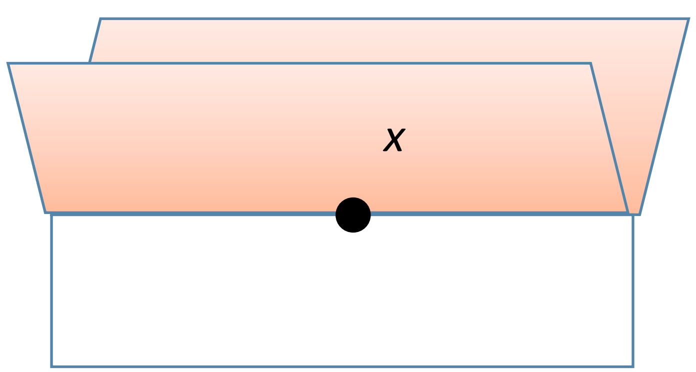

Any two adjacent sheets will expose x.

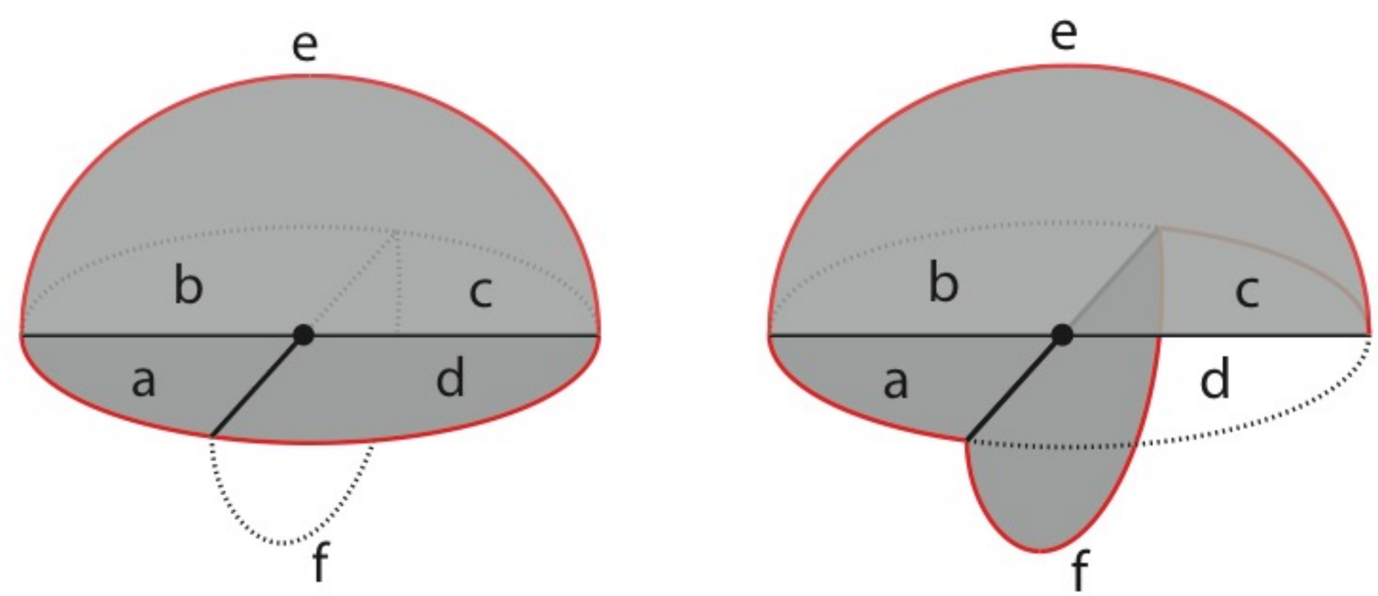

Here, removing only one sheet or removing sheets b and d will not expose the center, as there is still a disk surrounding it (shown in red).

Exposing sets (cont.)

- Exposing requires the point to be able to be "burned away", which is why there can be no closed disk surrounding it.

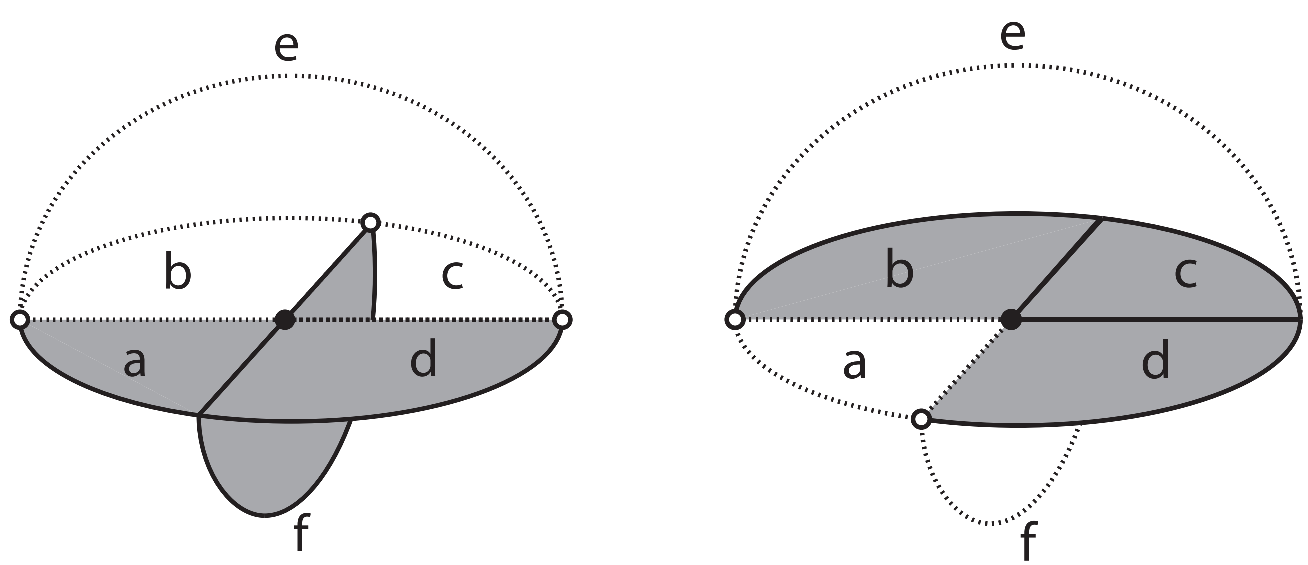

Here, sheets b,c and e expose the center point, or sheets a, e, and f.

- This is key when developing a more combinatorial and closed form definition of burn time.

Burning 3d medial axes

- An exposing tree for a point x in again a finite tree on the medial axis, rooted at x and with leaves at the boundary. All edges must be contained in the 2-manifold regions of the medial axis.
- However, when the tree crosses singular curves, the root of that subtree must be exposed by the sheets the subtree lies on.
- The longest path in the tree again gives the length of the tree, and burn time of a point is the infimum over all possible trees.

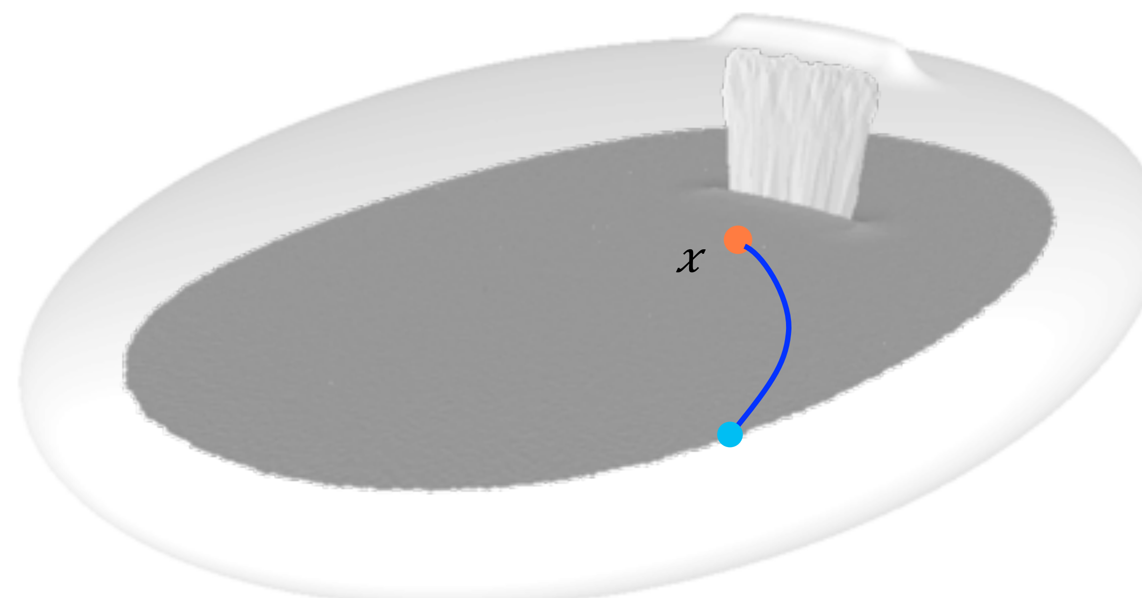

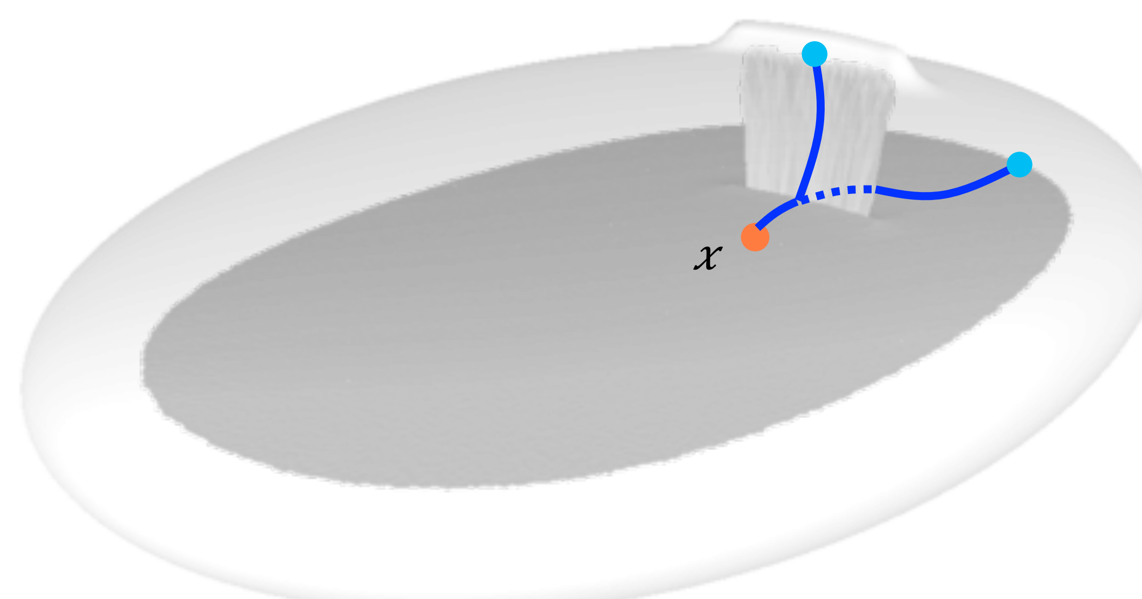

An example of burn time

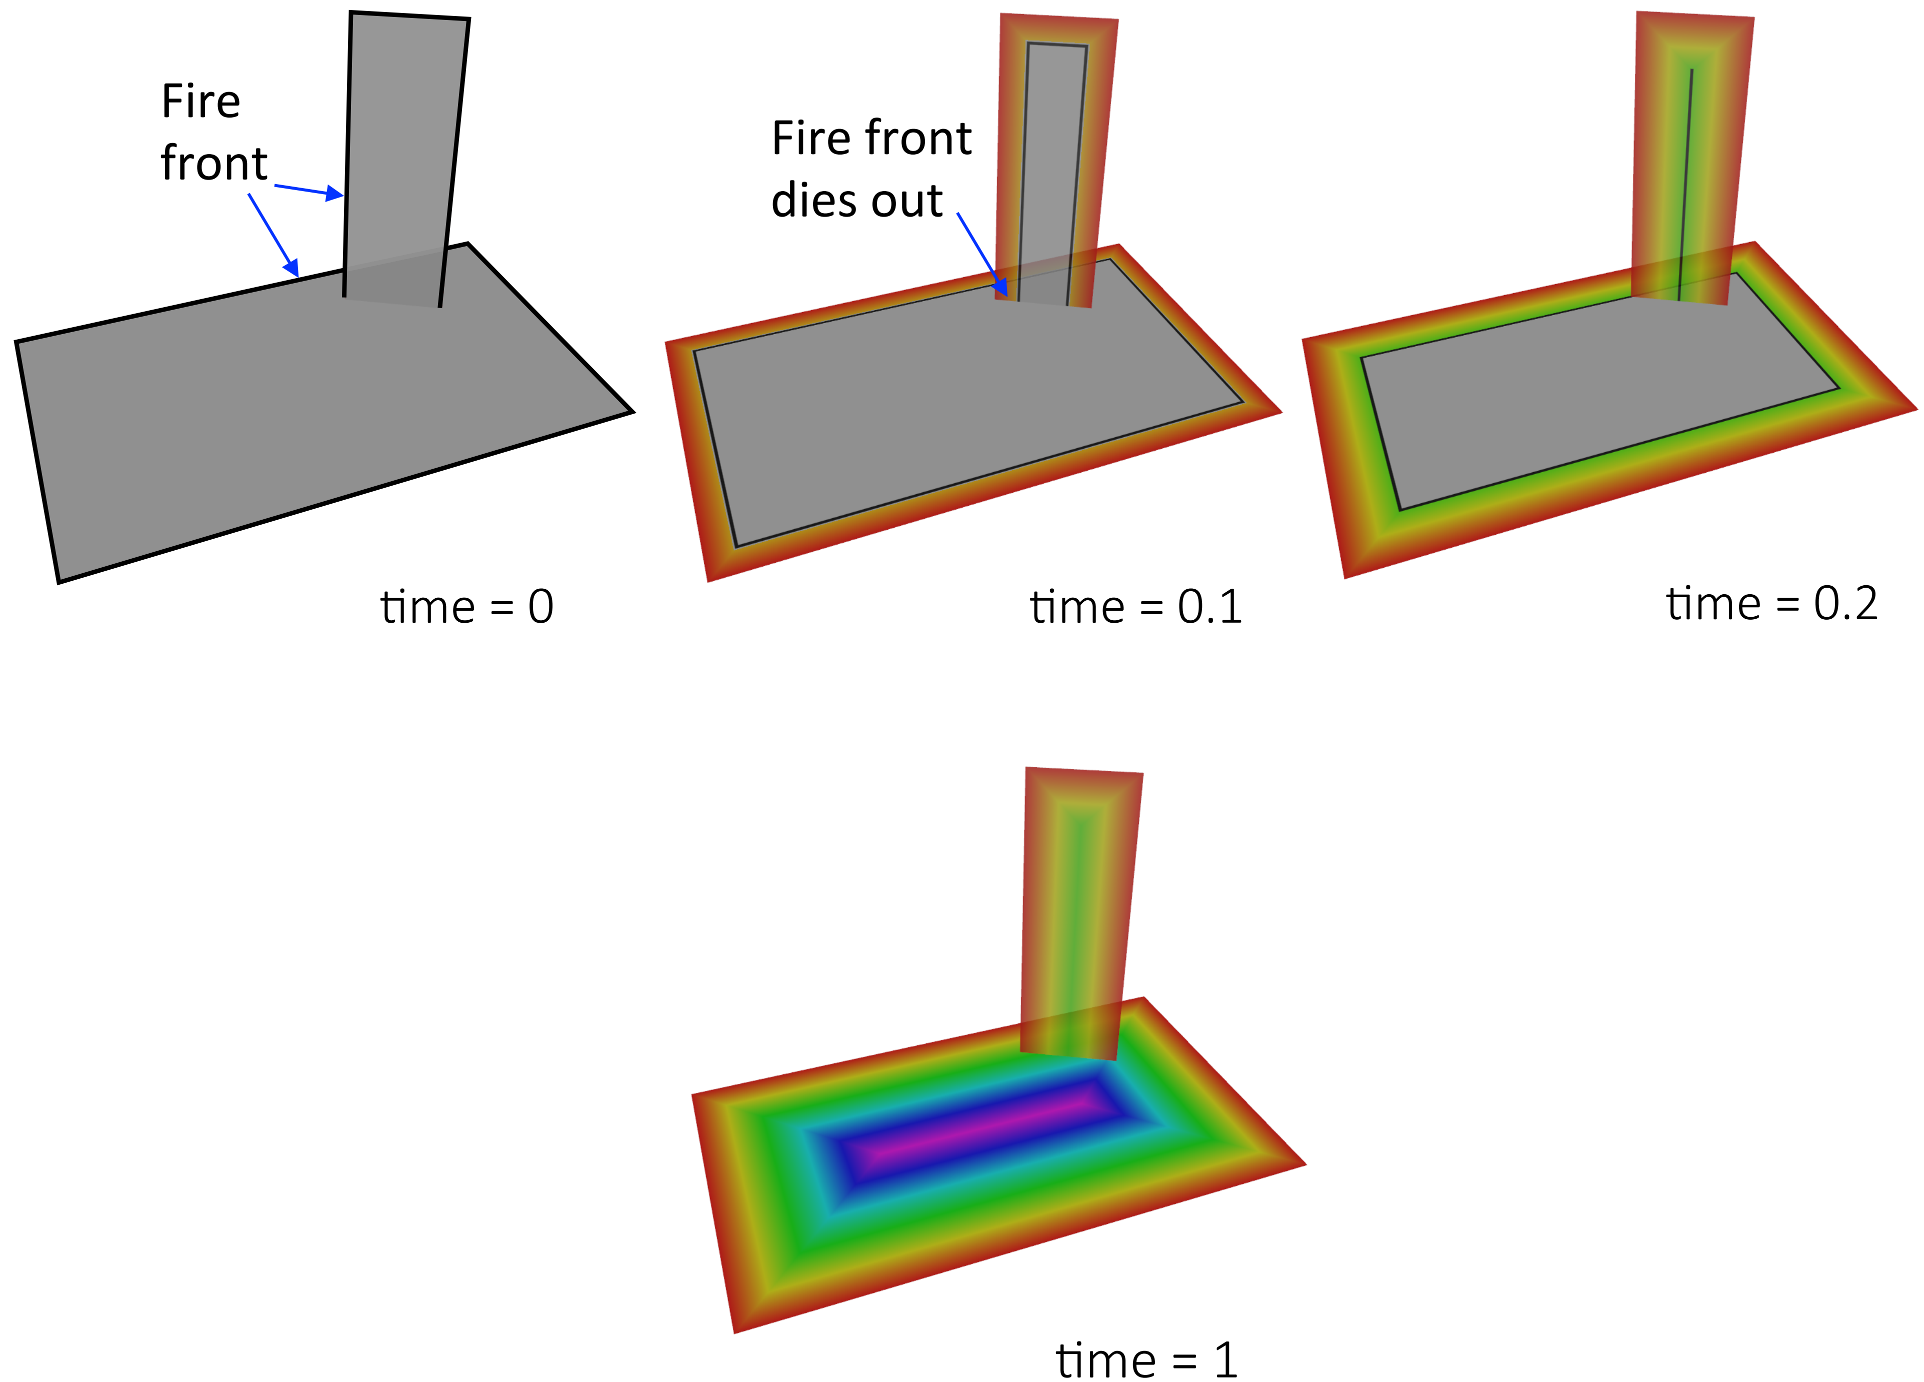

Another example

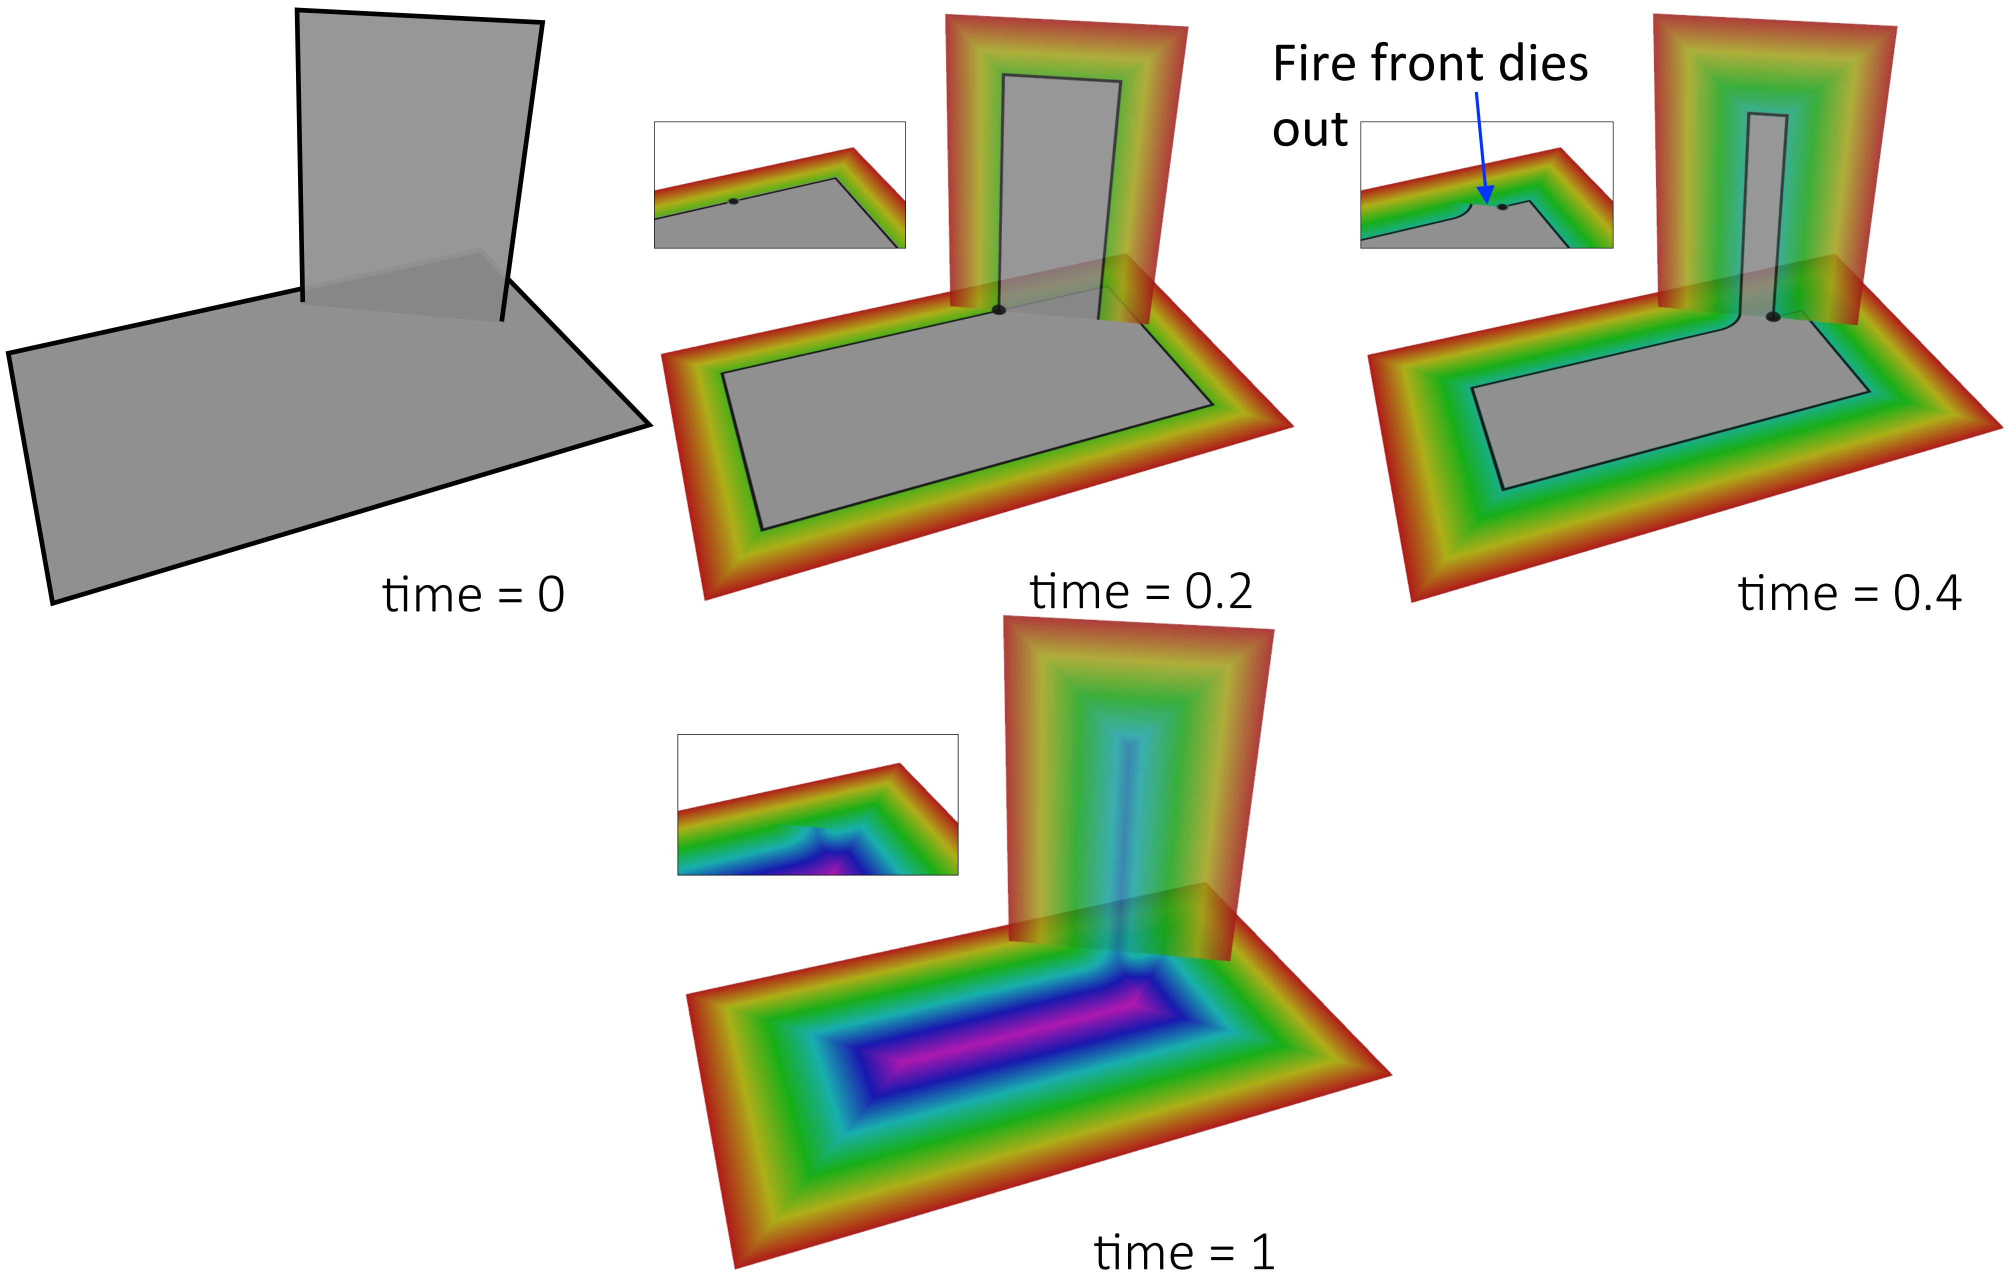

Burn time properties

- We formalize a definition of burn time and prove the following properties of burn time (analogous to the 2d results from earlier work) [Yan et all 2016]:
  - Burn time is upper semi-continuous on singular regions.
  - Burn time is 1-Lipschitz and continuous on manifold regions.
  - Burn time is finite away when not on the maximally closed sub complex of M.

Erosion thickness in 3d

- We also give the first extension of erosion thickness to 3d using burn time [Yan et al 2016].
  - Recall the 2d picture:

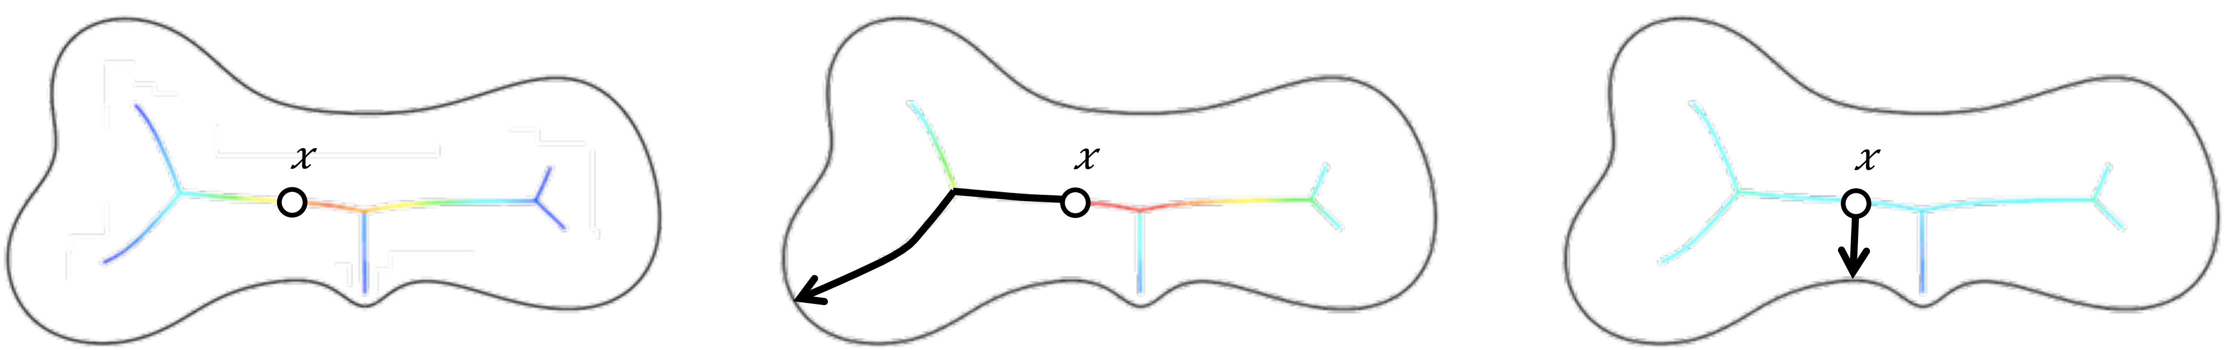

Erosion Thickness $ET(x)$	=	Burn Time $BT(x)$	−	Radius $R(x)$
"Tubularity"		"Length"		"Thickness"

Erosion thickness in 3d

We can define a similar function using burn time in 3 dimensions as well, capturing similar types of features:

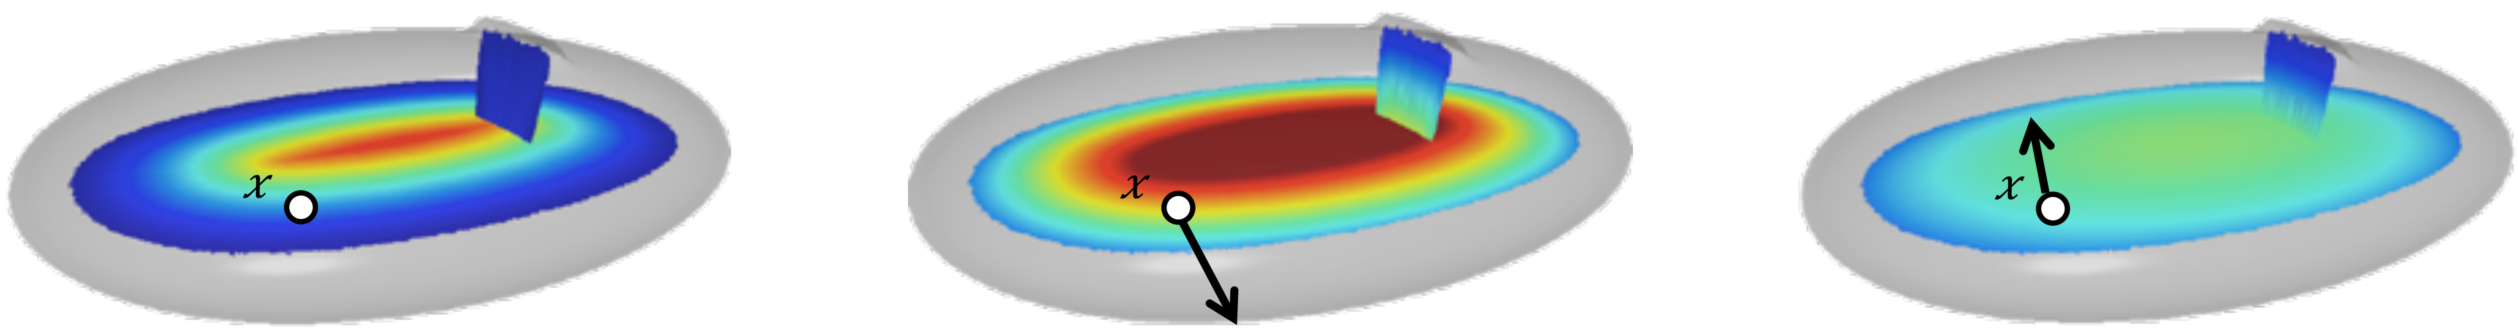

Erosion Thickness	=	Burn Time	−	Radius
$ET(x)$		$BT(x)$		$R(x)$
“Plate-likeness”		“Width”		“Thickness”

Computing burn time

- In 3d, computing burn time exactly is difficult
  - The only algorithm we have is based on computing geodesics [MMP 1987], but is not guaranteed to terminate [Sykes 2016].
- Instead, we use approximate Dijkstra on a refinement of the input mesh.

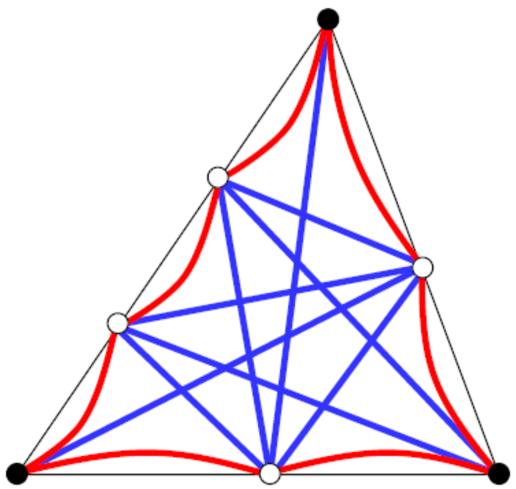

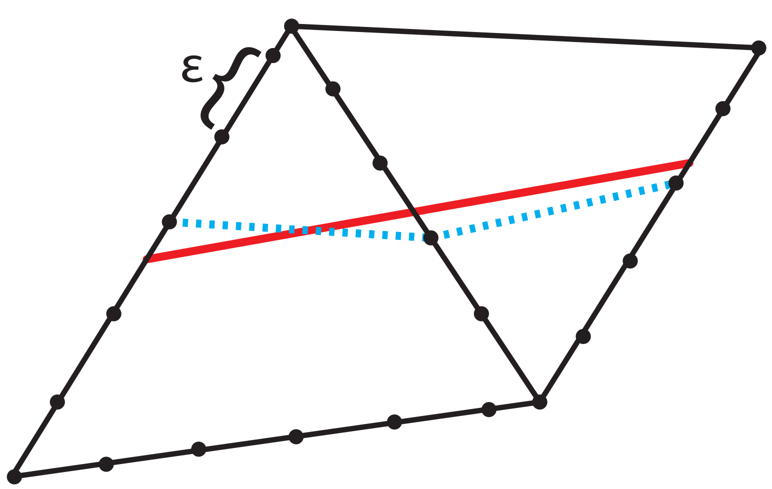

Approximation guarantee

- Theorem [Yan et al 2016]: Let M be a medial axis of a piecewise linear manifold whose triangulation, T, has F flat faces. Let w be the longest distance between any two Steiner points. Then our approximation algorithm gives a value within 2wF of the actual burn time.
  - Essentially, we get constant error per face we cross.
  - The proof is very similar to prior work on approximating geodesics on meshes [Lanthier et al 1997].

Final Result

- In the end, we can get a guaranteed approximation to burn time, given a fine enough refinement of the mesh.

- (Code available at: https://yajieyan.github.io/project/et/)

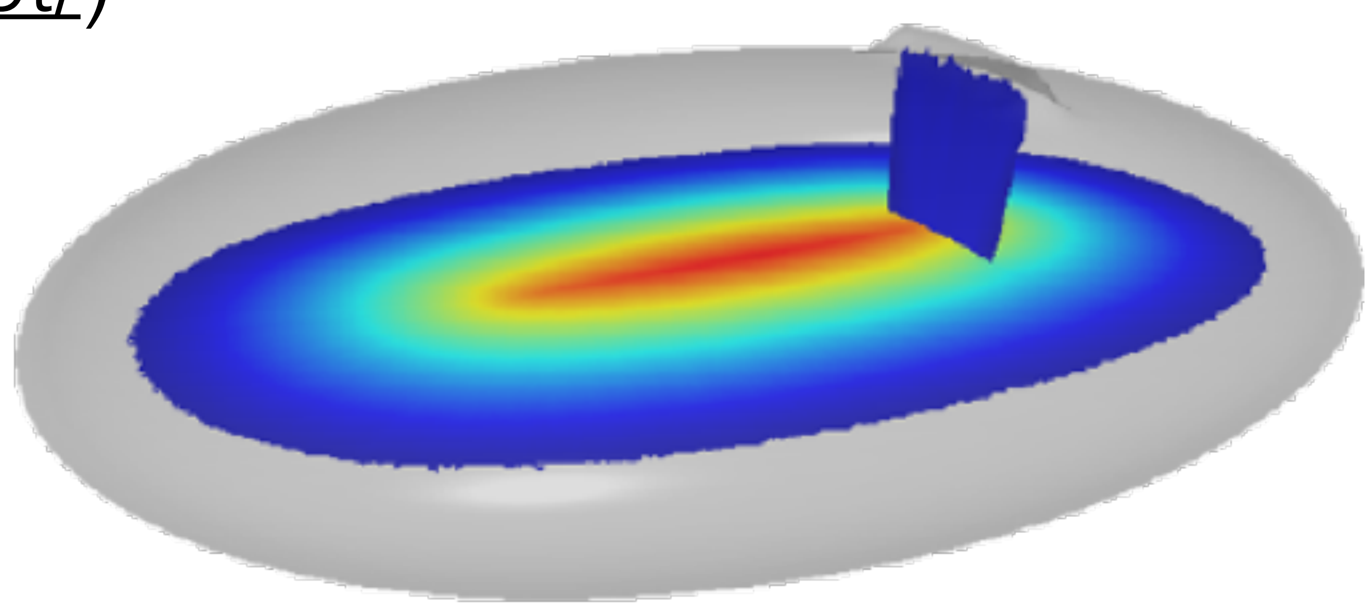

Final result

Erosion thickness

We can then approximate erosion thickness as well:

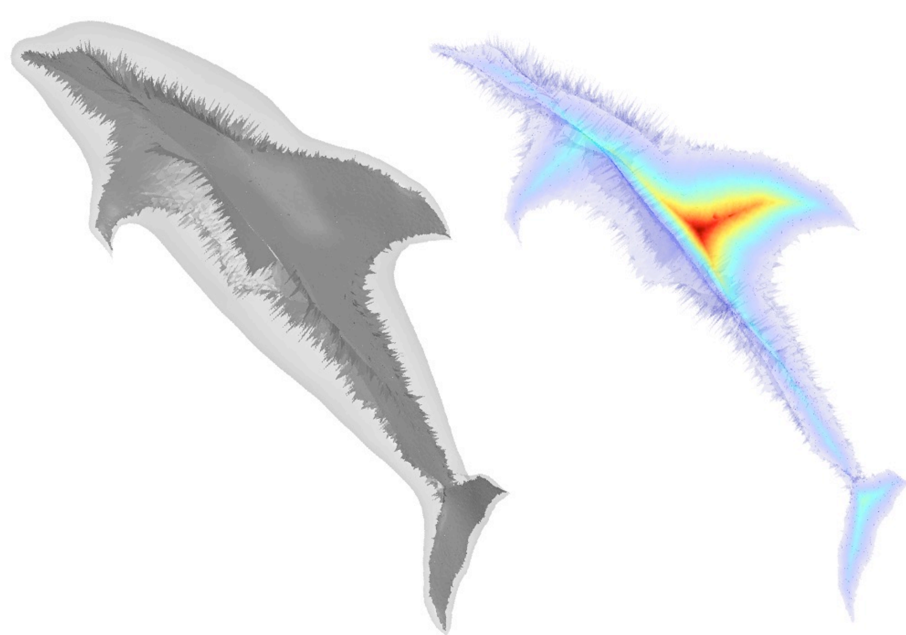

Original dolphin

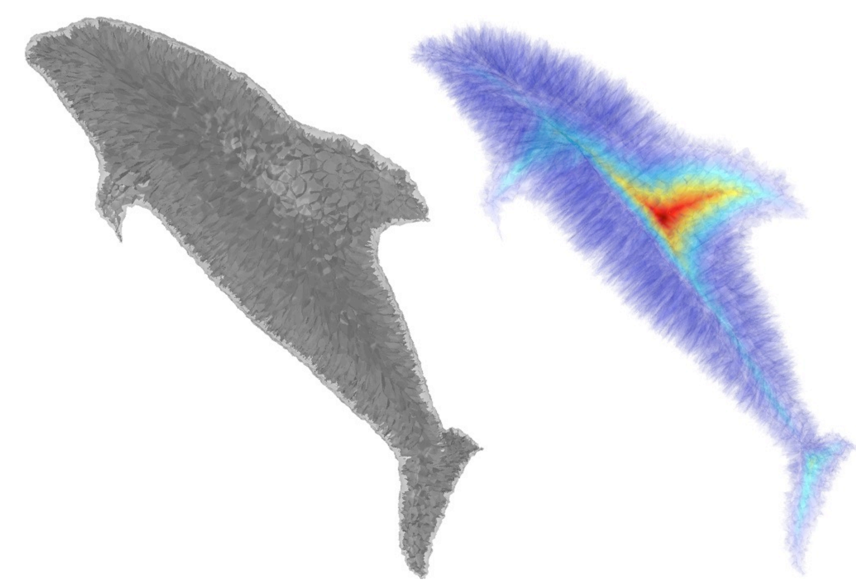

Noisy dolphin

Comparisons

As in 2d, our results yield nicer shape descriptors than other options:

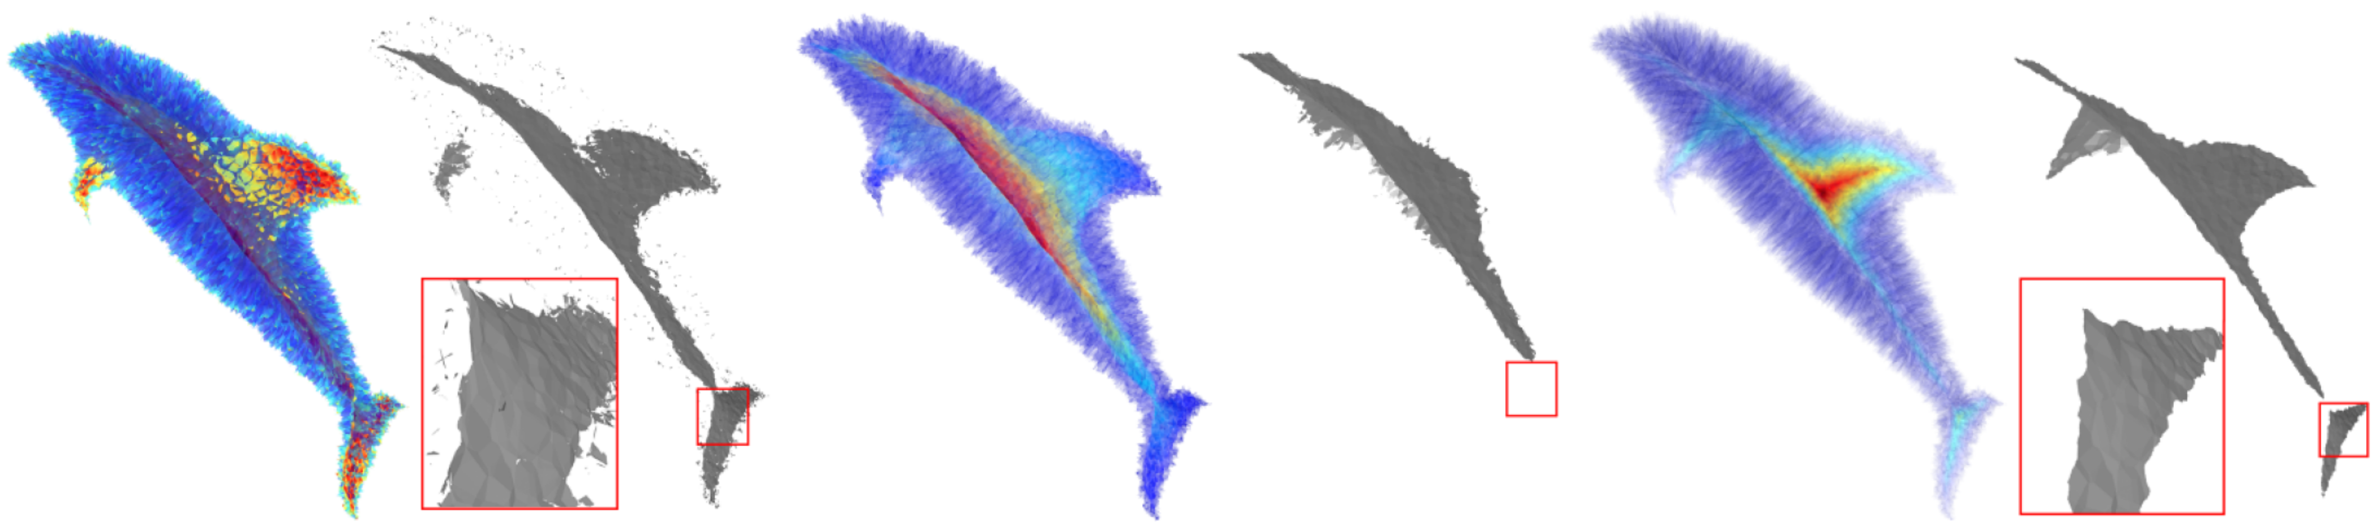

Object angle Circumradius Erosion thickness

Comparison with MGF

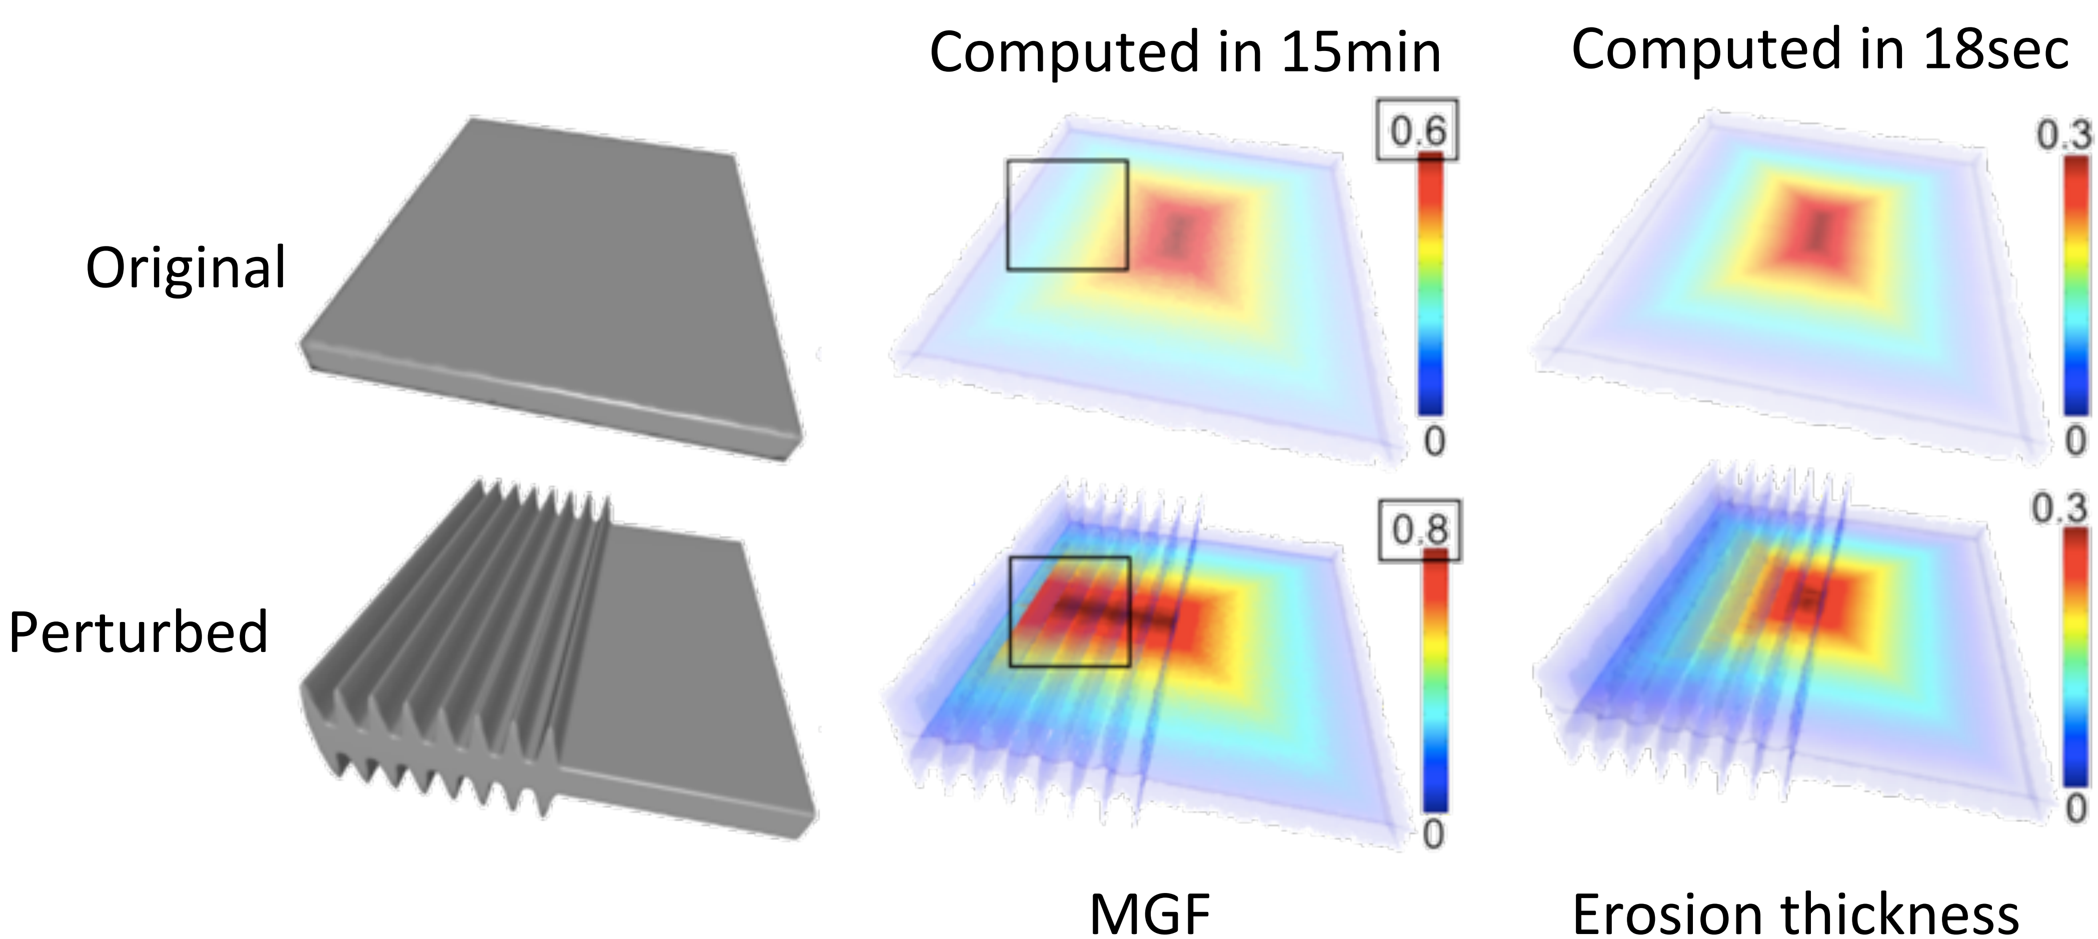

Skeletons in 3d

However, using erosion thickness alone to prune does not give a good skeleton:

Shape + MA

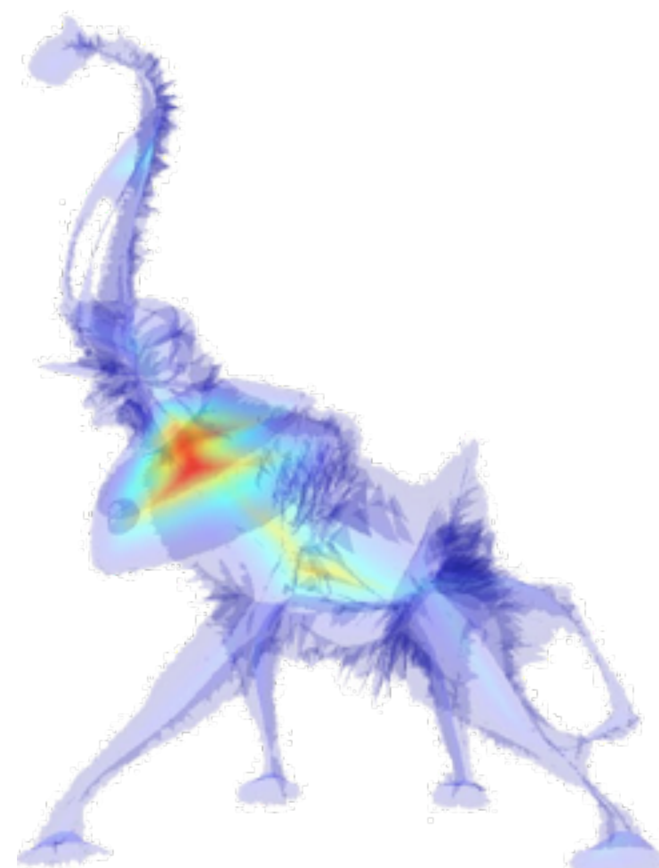

ET

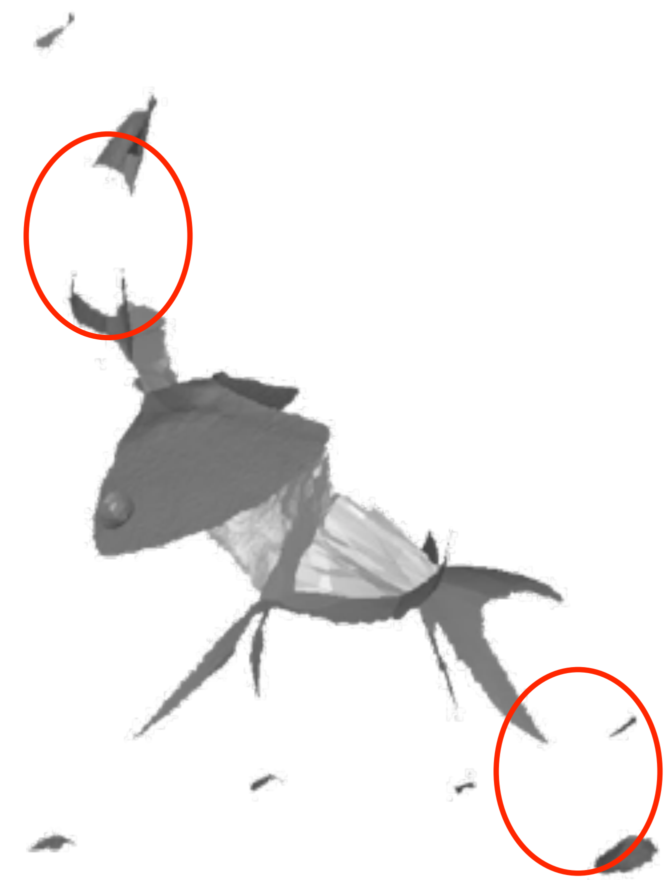

Naïve thresholding

Skeletons in 3d

- Similarly, using just quench sites of the burn time function (where it is not differentiable) is not enough to get a good skeleton.
    - The result will not have the correct topology.
- In a sense, we want to get all ridges of the burn time function; this will reconnect the pieces.

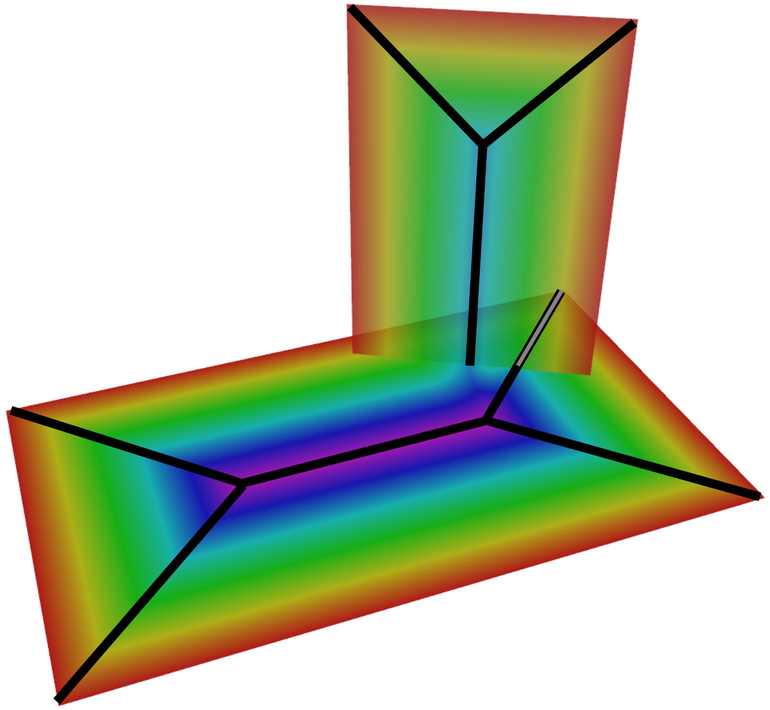

Approximate skeleton

- Given our approximation of burn time, we develop a discrete algorithm that traces ridges by pruning the dual graph of our refined mesh.

- We keep edges of the graph which contribute to the burn time of some vertex, and take the dual:

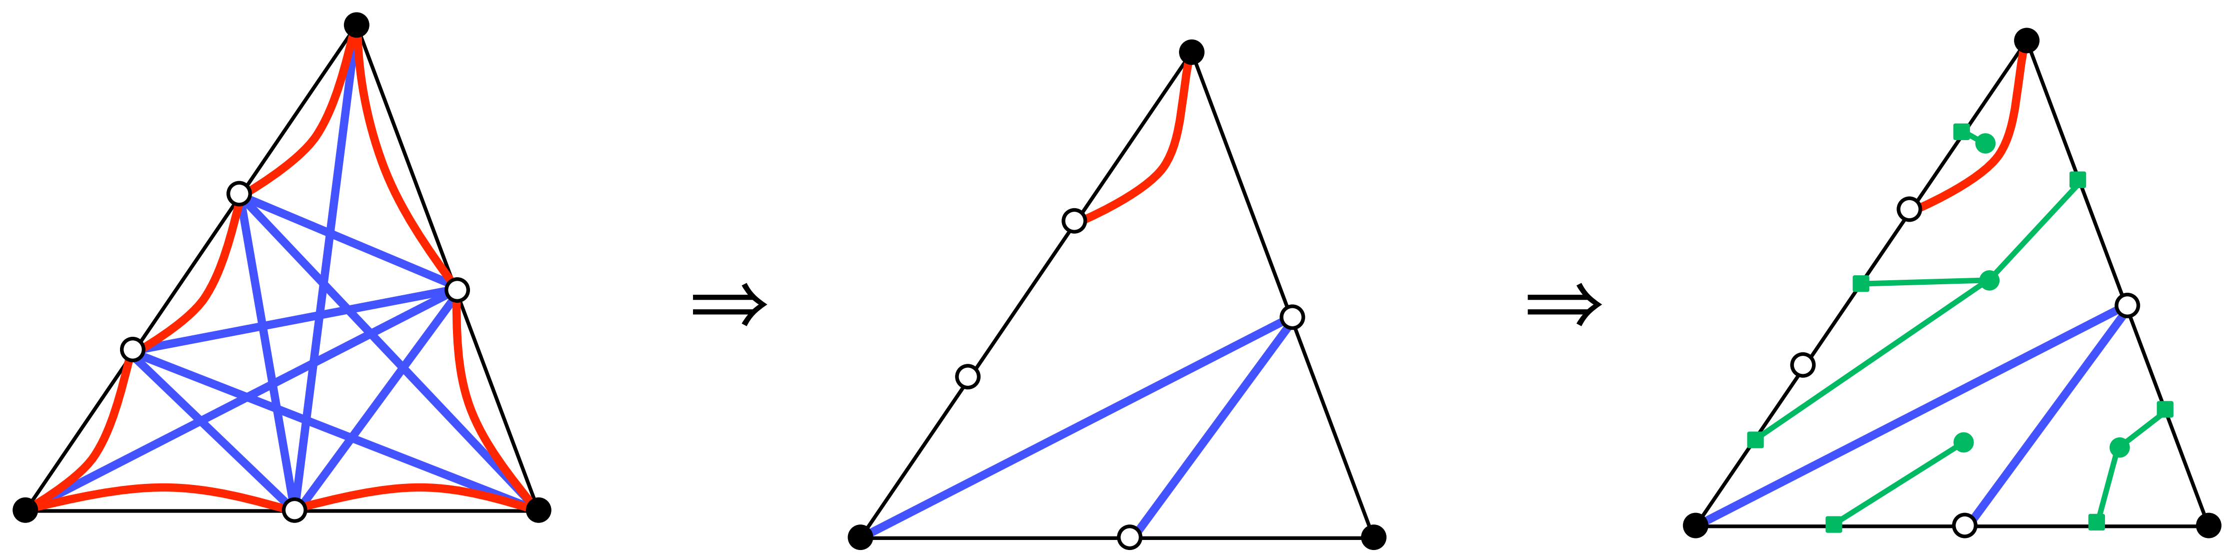

Approximate skeleton

The resulting graph is guaranteed to be homotopy equivalent to the medial axis, since it is a retract of the original shape.

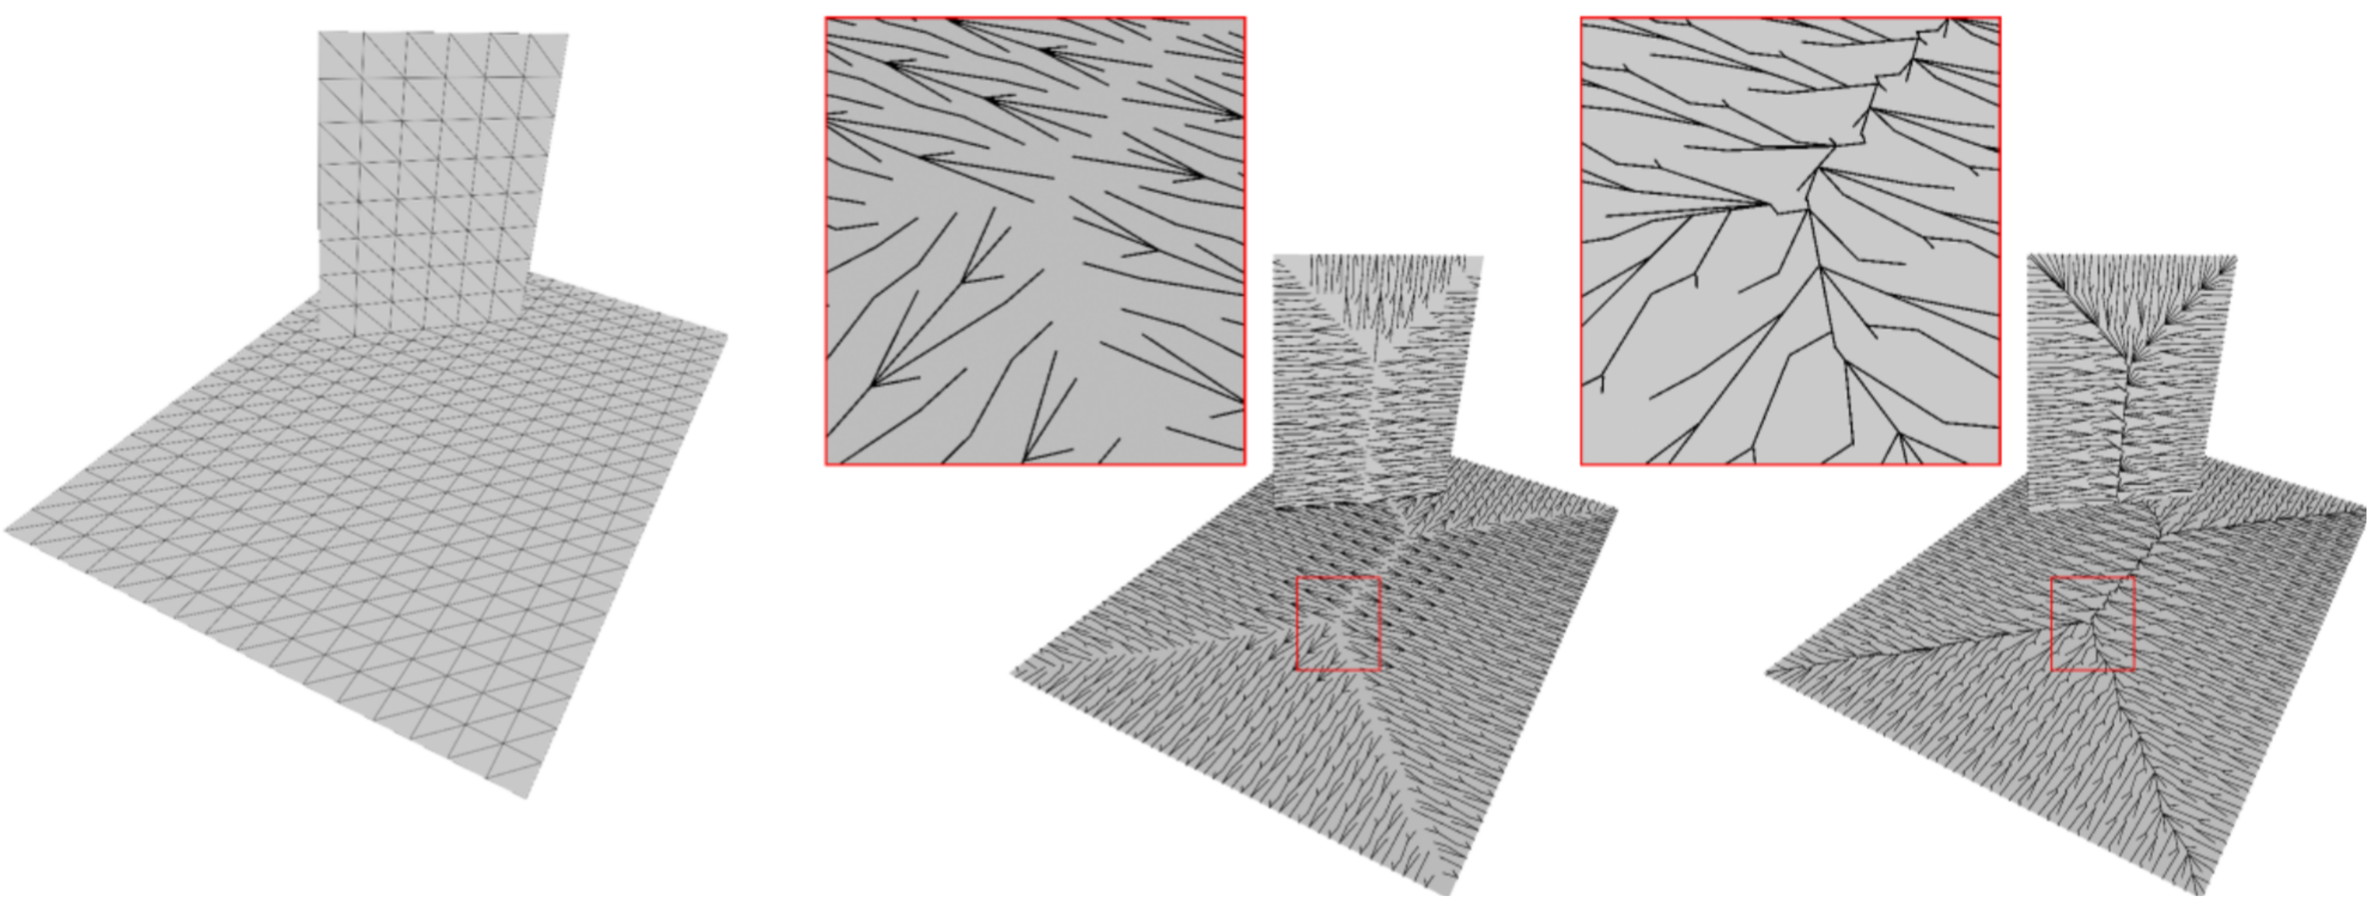

Adding significance measures:

Burn time (or erosion thickness) helps to identify the core portions of this approximate medial curve:

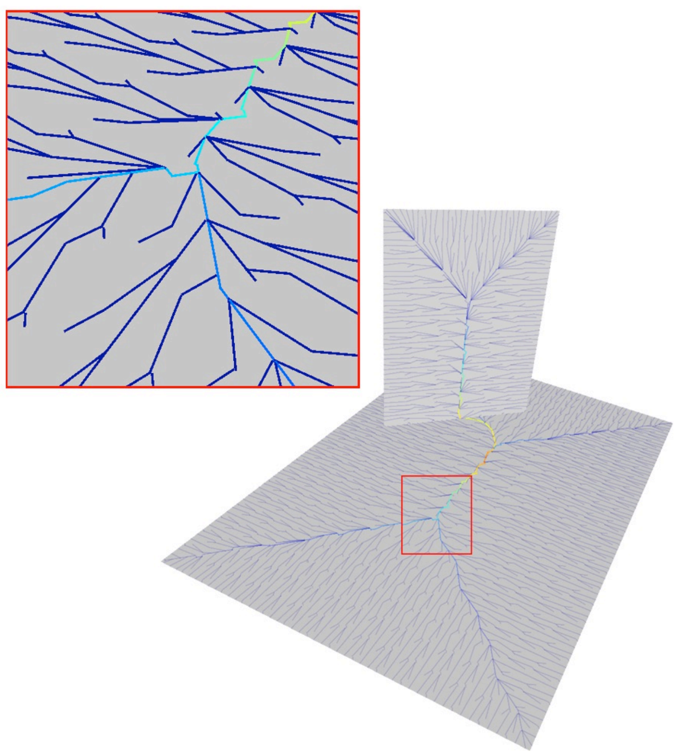

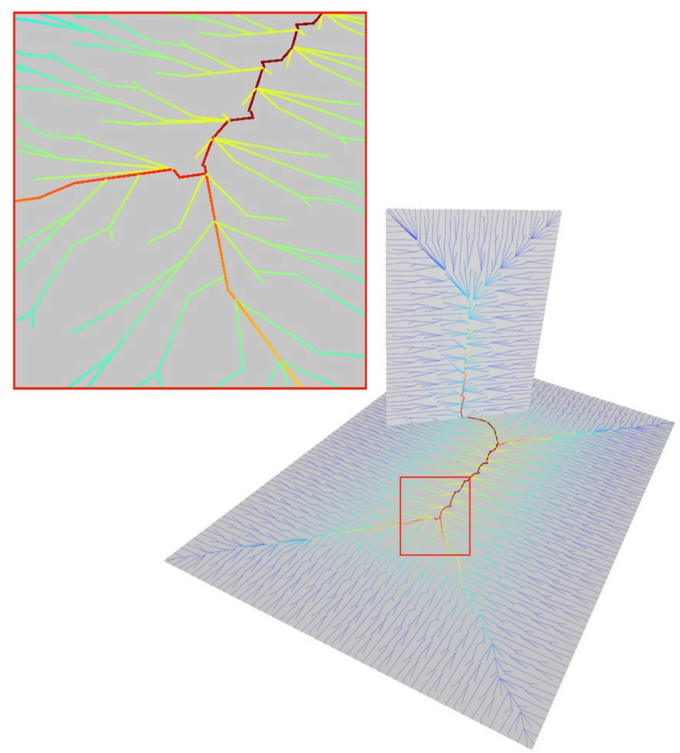

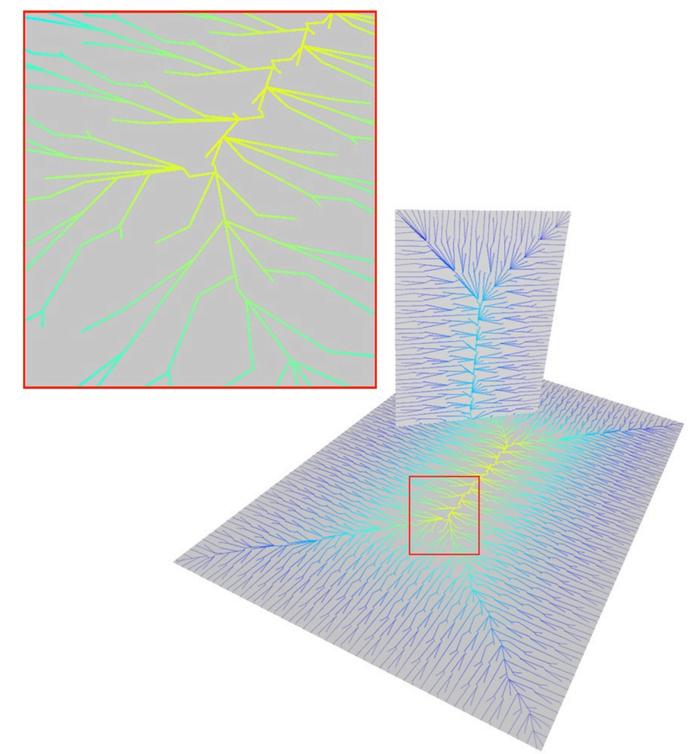

Erosion thickness on medial curve = Burn time on medial curve - Burn time on medial axis

Adding significance measures:

We can then use burn time or erosion thickness to prune:

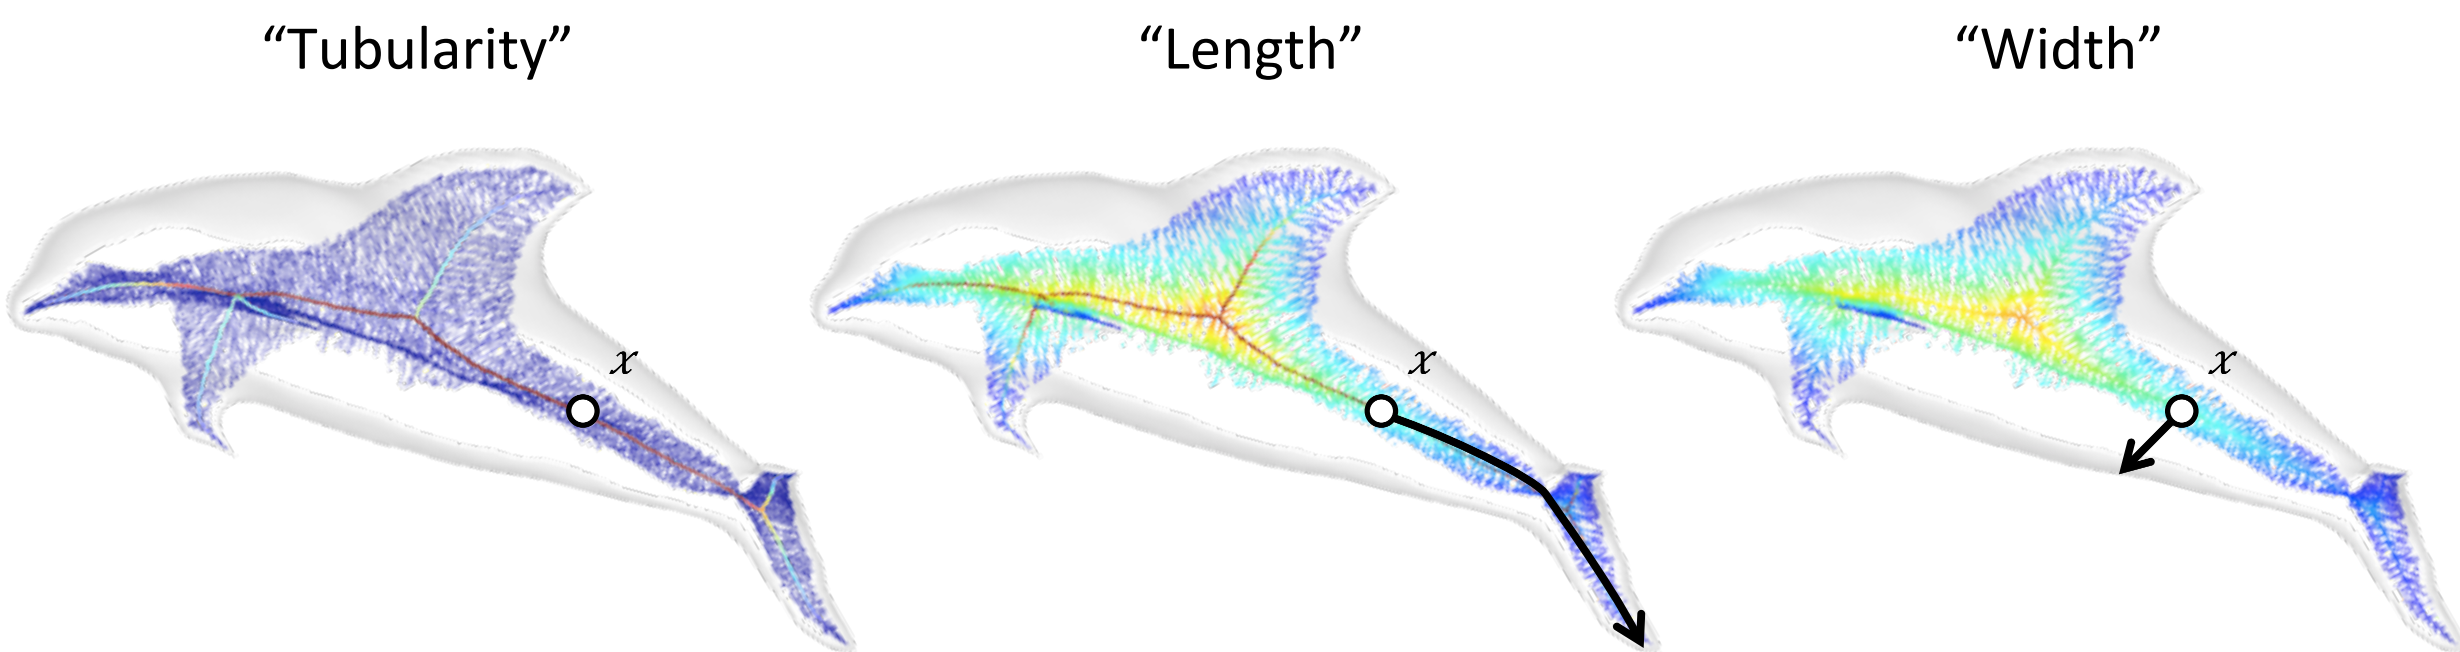

Erosion thickness on medial curve = Burn time on medial curve - Burn time on medial axis

Skeleton results

Resulting skeletons:

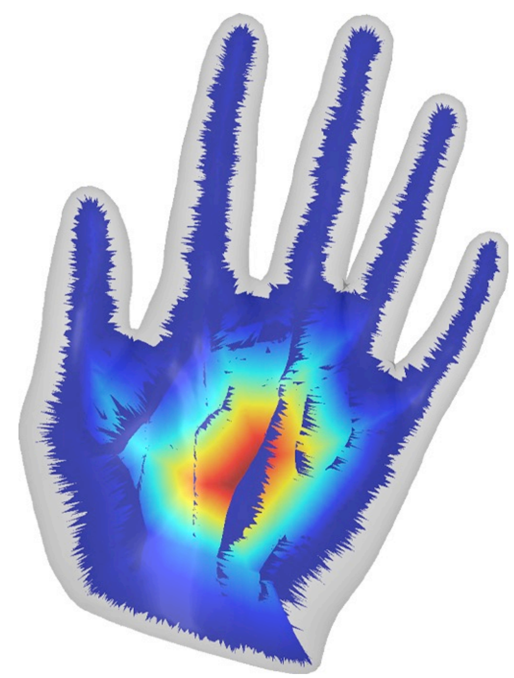

ET on medial axis

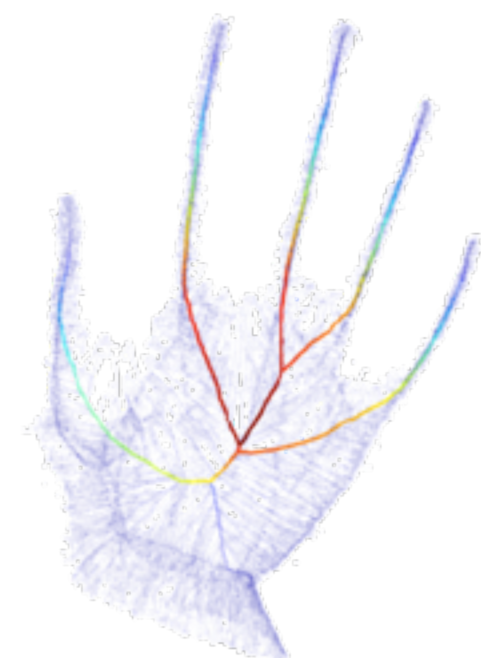

ET on medial curve

Comparisons

- The scale axis transform has a similar purpose, but does not maintain topology.

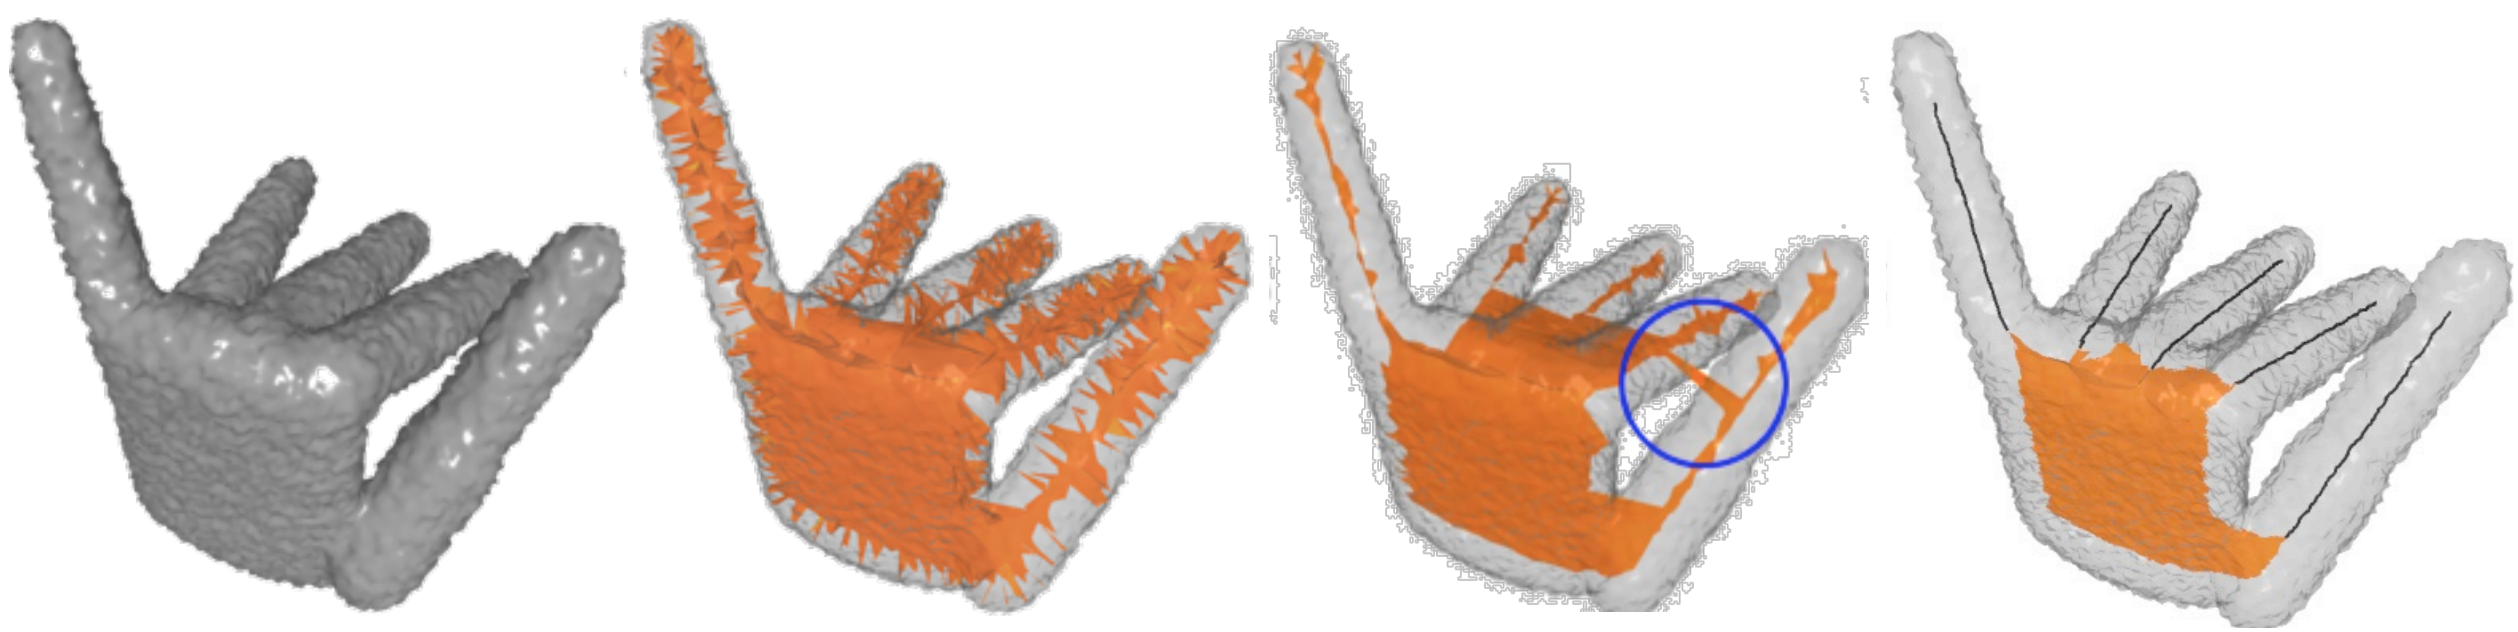

Noisy shape　SAT: small scale　SAT: large scale　ET-based

Hybrid skeletons

We also use this to develop hybrid skeletons, capturing significant 2d sheets from the medial axis.

Naïve thresholding

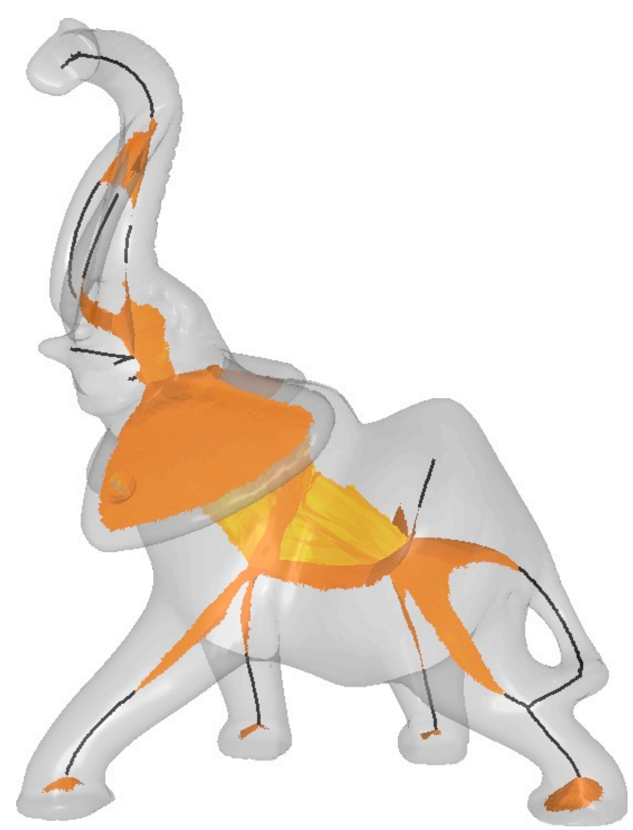

Curve+surface skeleton

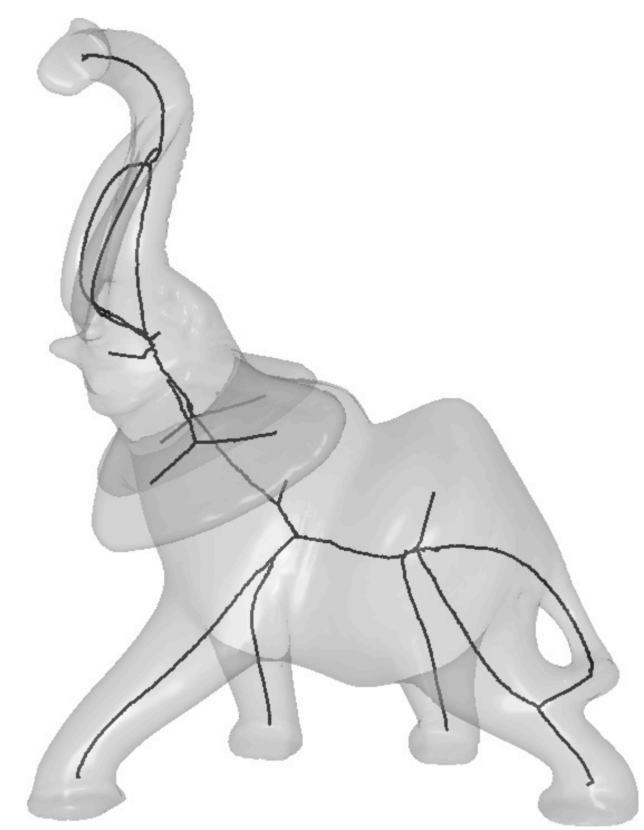

Curve only skeleton

One application: root systems

Root architecture is controlled genetically, and the shape of the root changes considerably

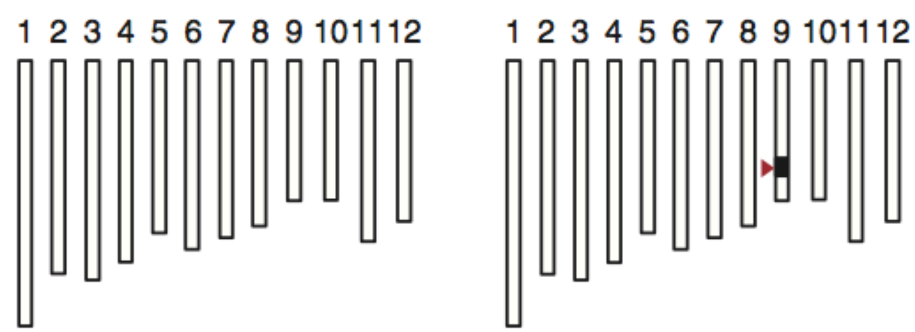

Uga et al. Nature Genetics 2013

Widely grown breeding line of rice

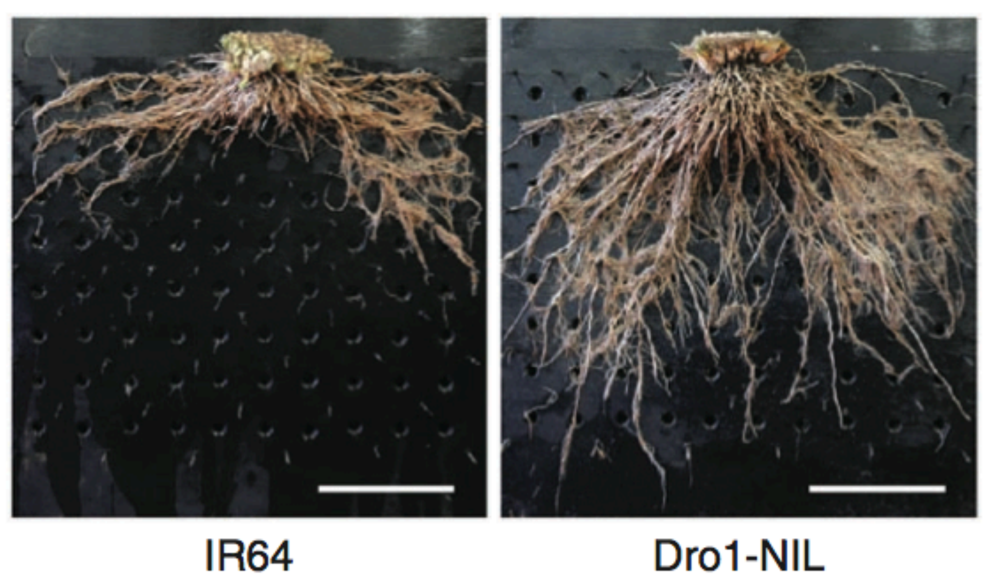

Effect on roots after 1 base pair in Dro-1 is changed

Shapes of root systems

Biologists care about isolating these shape genes because the root architecture drastically impacts crop production

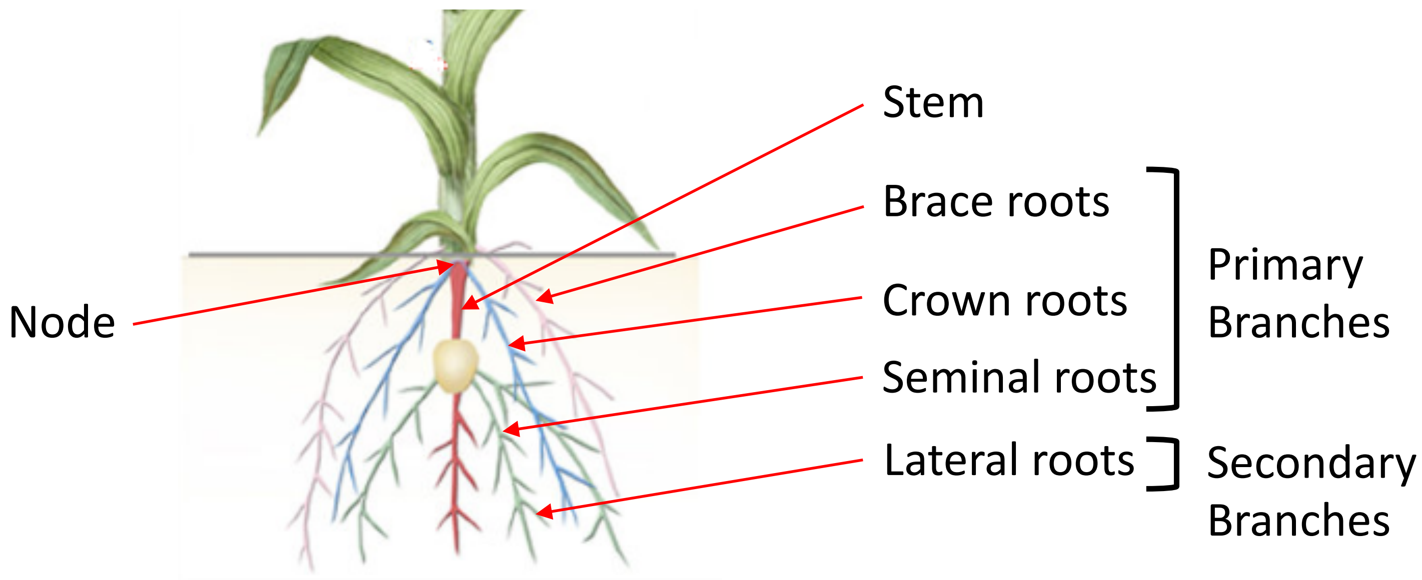

Scanning root systems

Recently, we have begun applying burn time to learn the shapes of root systems.

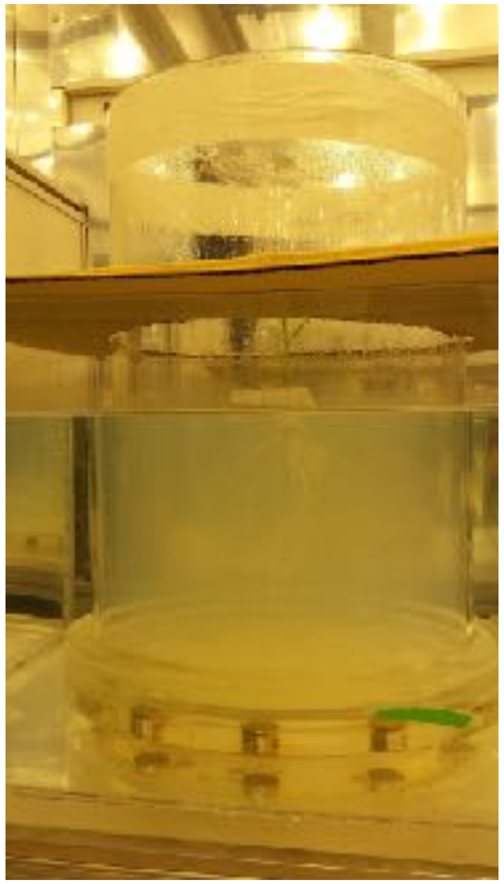

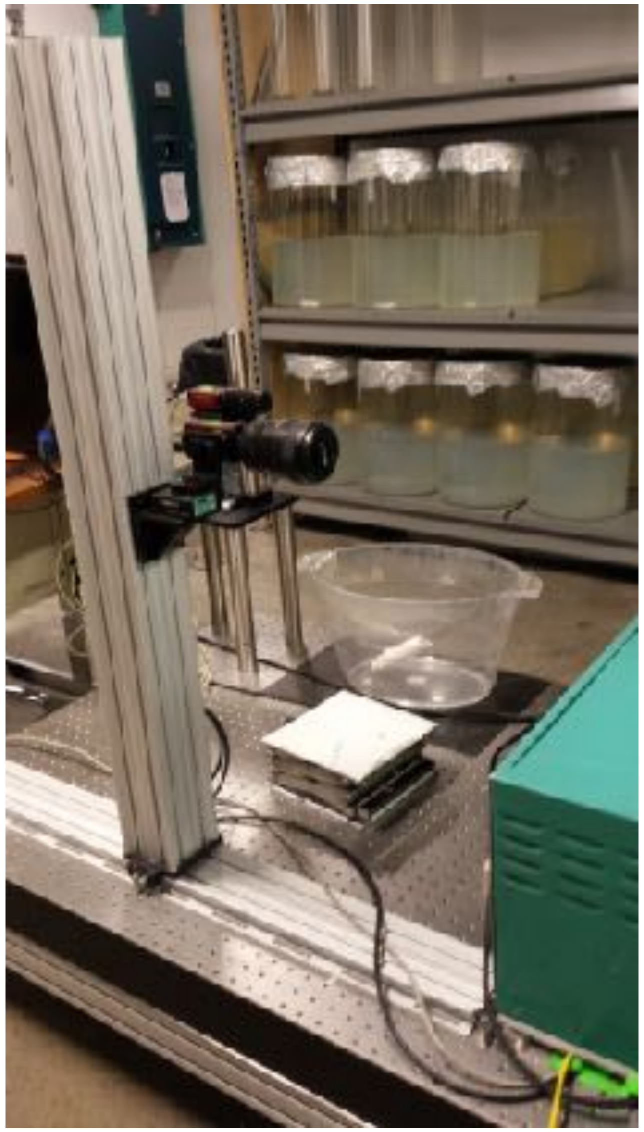

Analyzing the roots

For corn roots, they dig and clean the systems, then scan in a high resolution x-ray imaging system, resulting in large, high quality images.

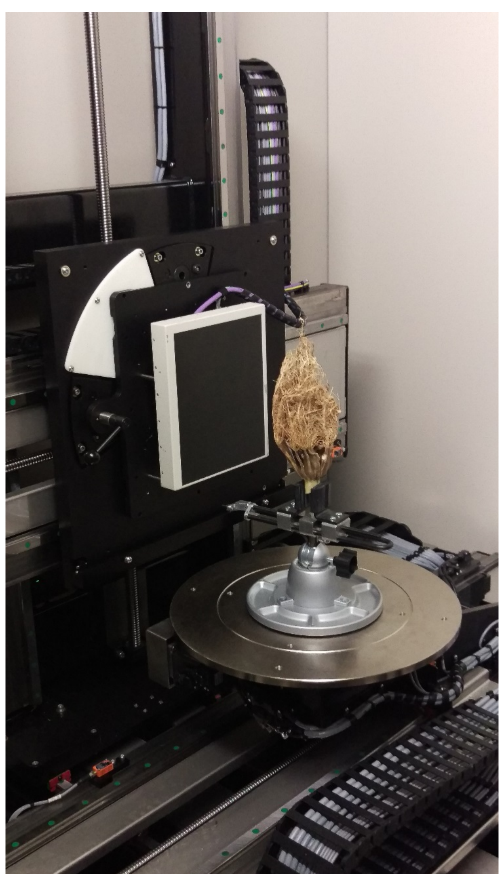

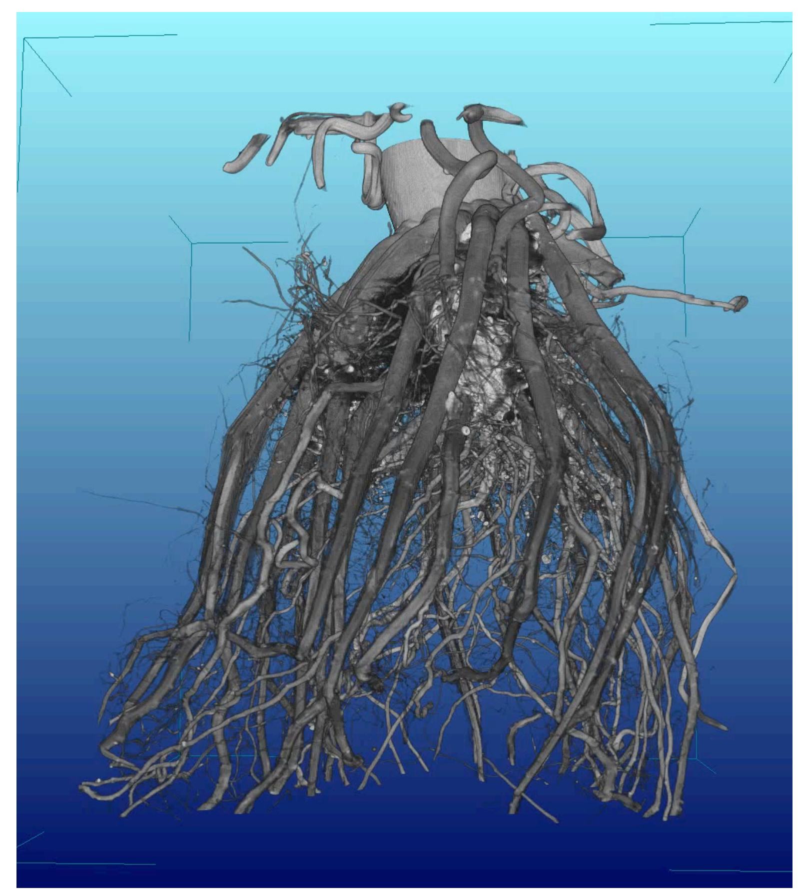

Root shapes

- We then have several challenges: denoising, restoring correct topology, and quantifying shape measures which capture the traits we are looking for.

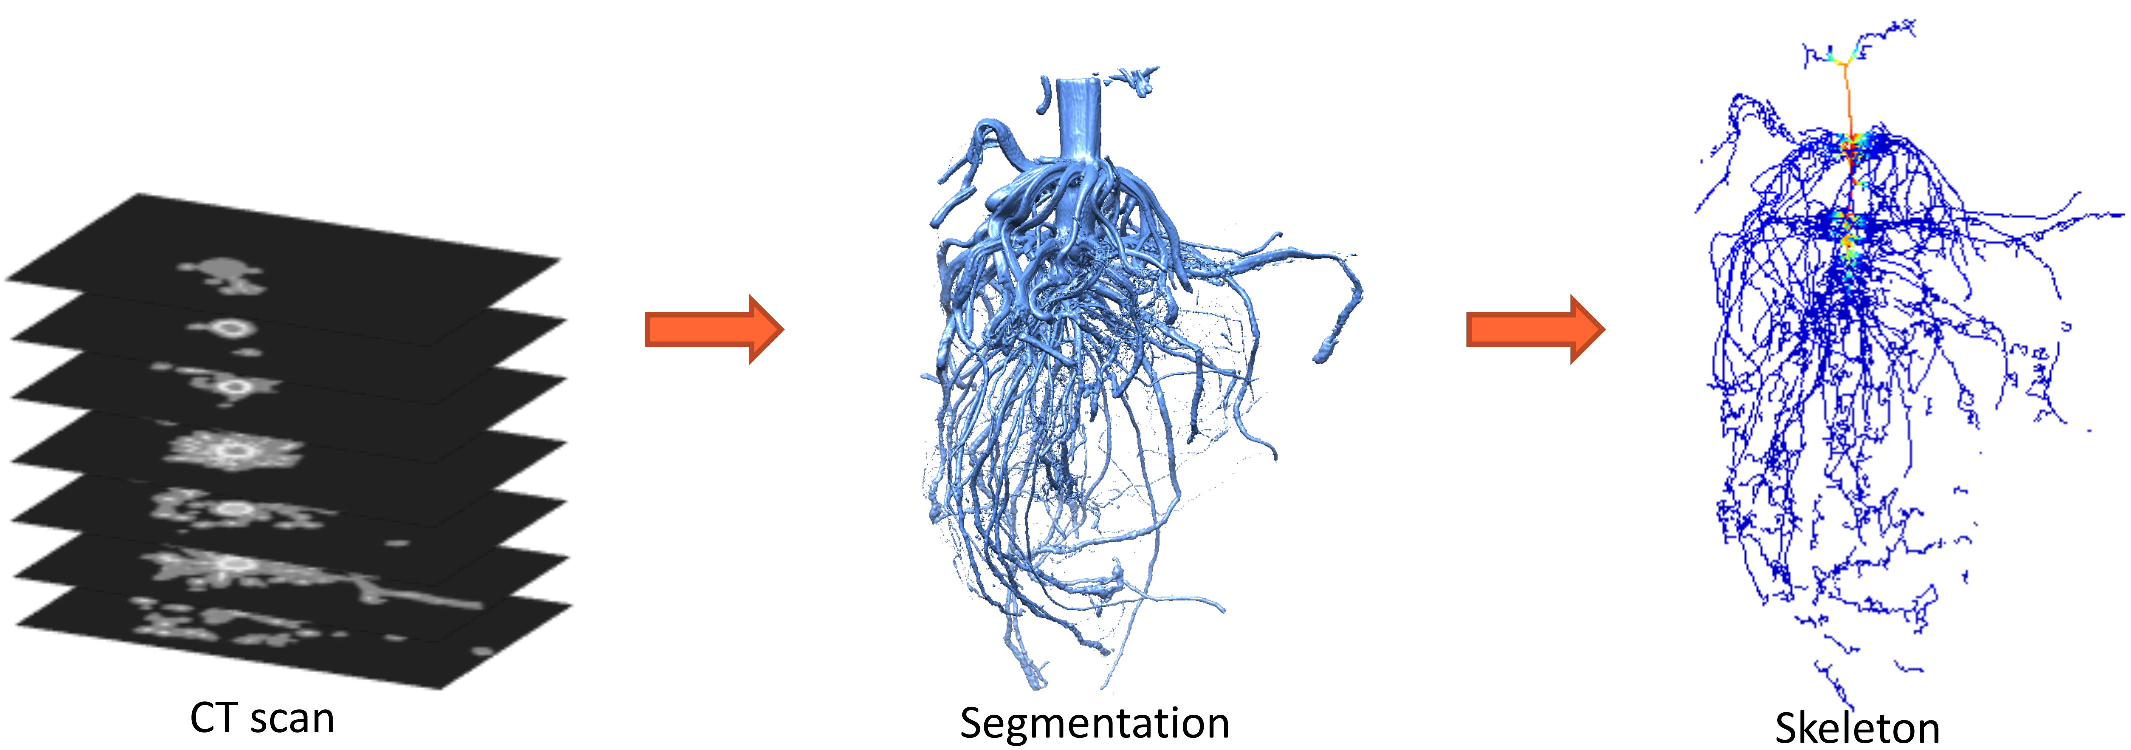

Corn Roots

- We use persistent homology to identify likely “noise”.
- Our first implementation does a naive simplification, and manages to remove 99% of the extra handles or holes.

Before

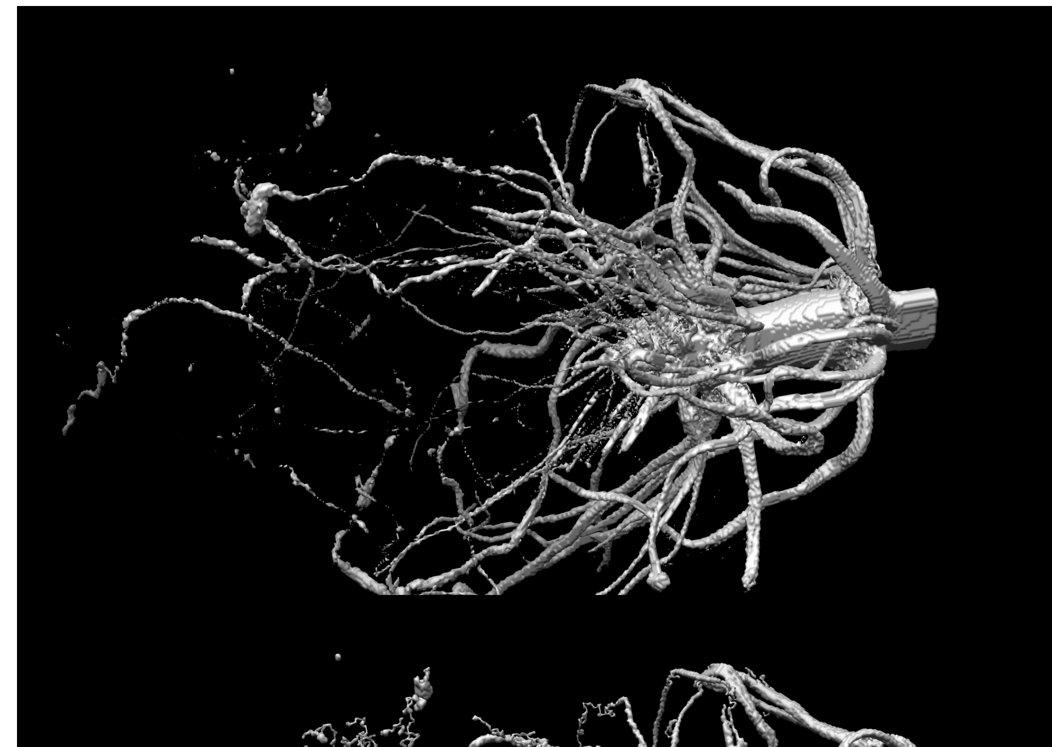

After

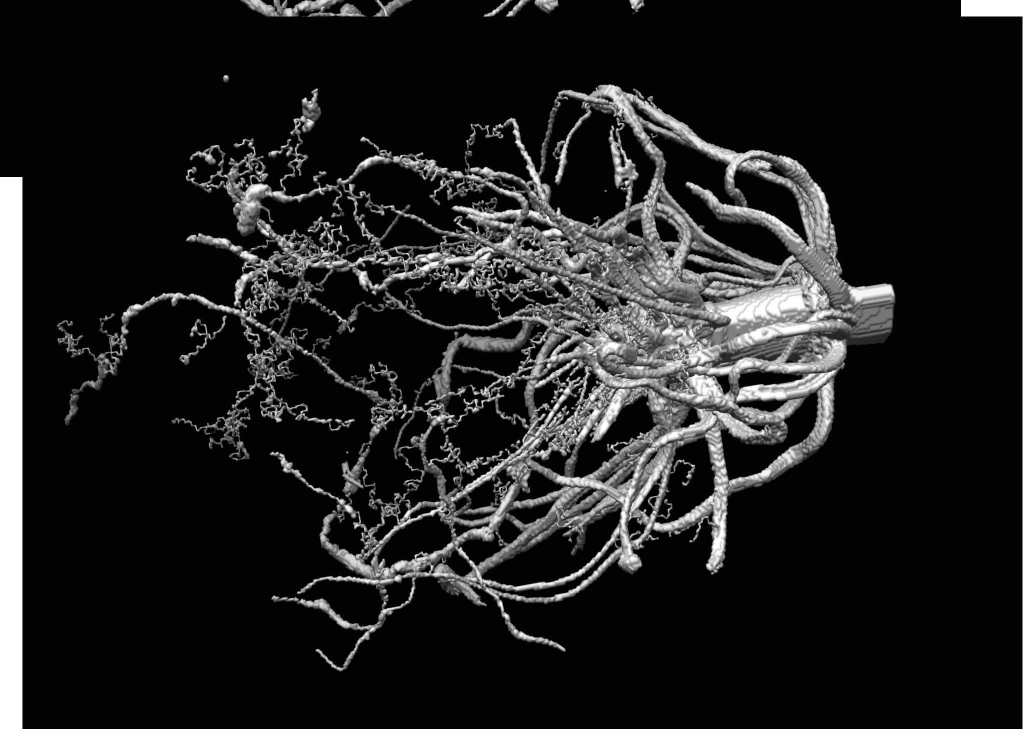

Code available at:
http://git.cs.slu.edu/public-repositories/shape-simplification-software

Repairing the topology

- Our approach is solving a variant of the Homological Simplification Problem: given a pair of simplicial complexes $S \subseteq N$, add simplicies to S to fill “noisy” holes or voids connecting erroneously disconnected components.
  - NP-Hard, even for 3-complexes embedded in 3d
- In our (heuristic) approach, we also remove simplicies to obtain the desired topology.

Zooming in

- However accurate the topology, the geometry of our repair needs to be improved!

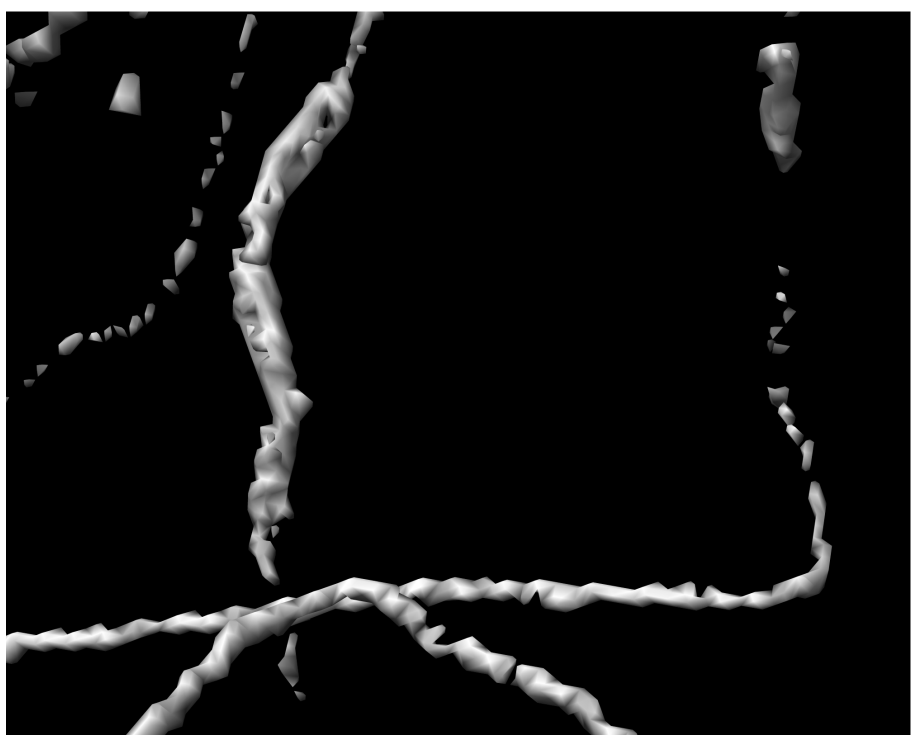

Before

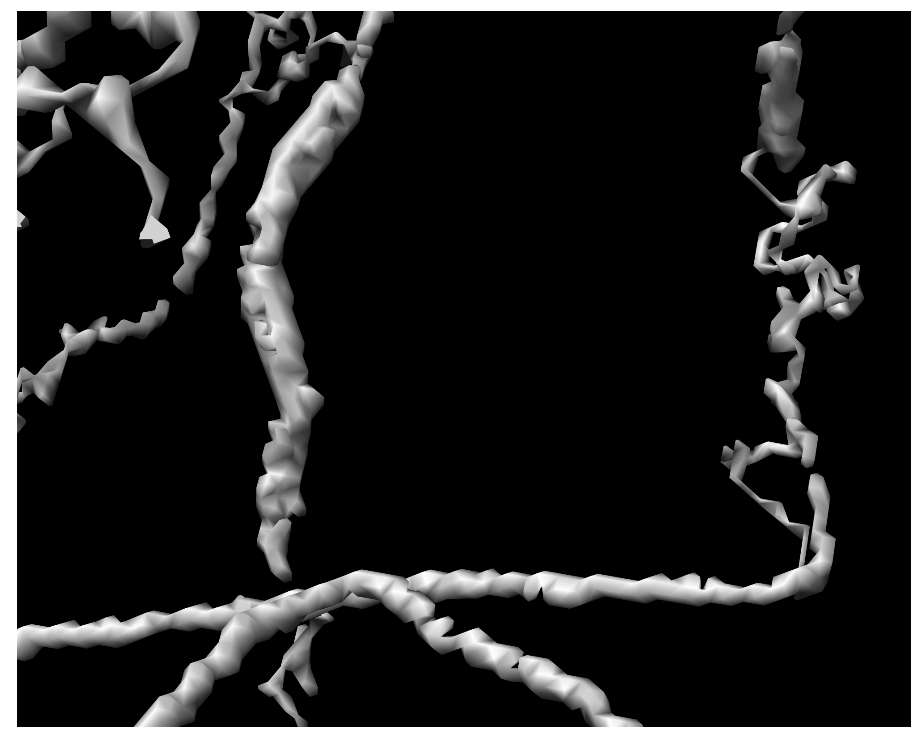

After

Improving the geometry

- Our goal now is to use the burn axis skeleton and shape information to reconstruct more accurate skeletons, and use these to calculate shape features for genetic analysis.

Preliminary tool

Tao Ju’s lab is working on an interface for our algorithms.

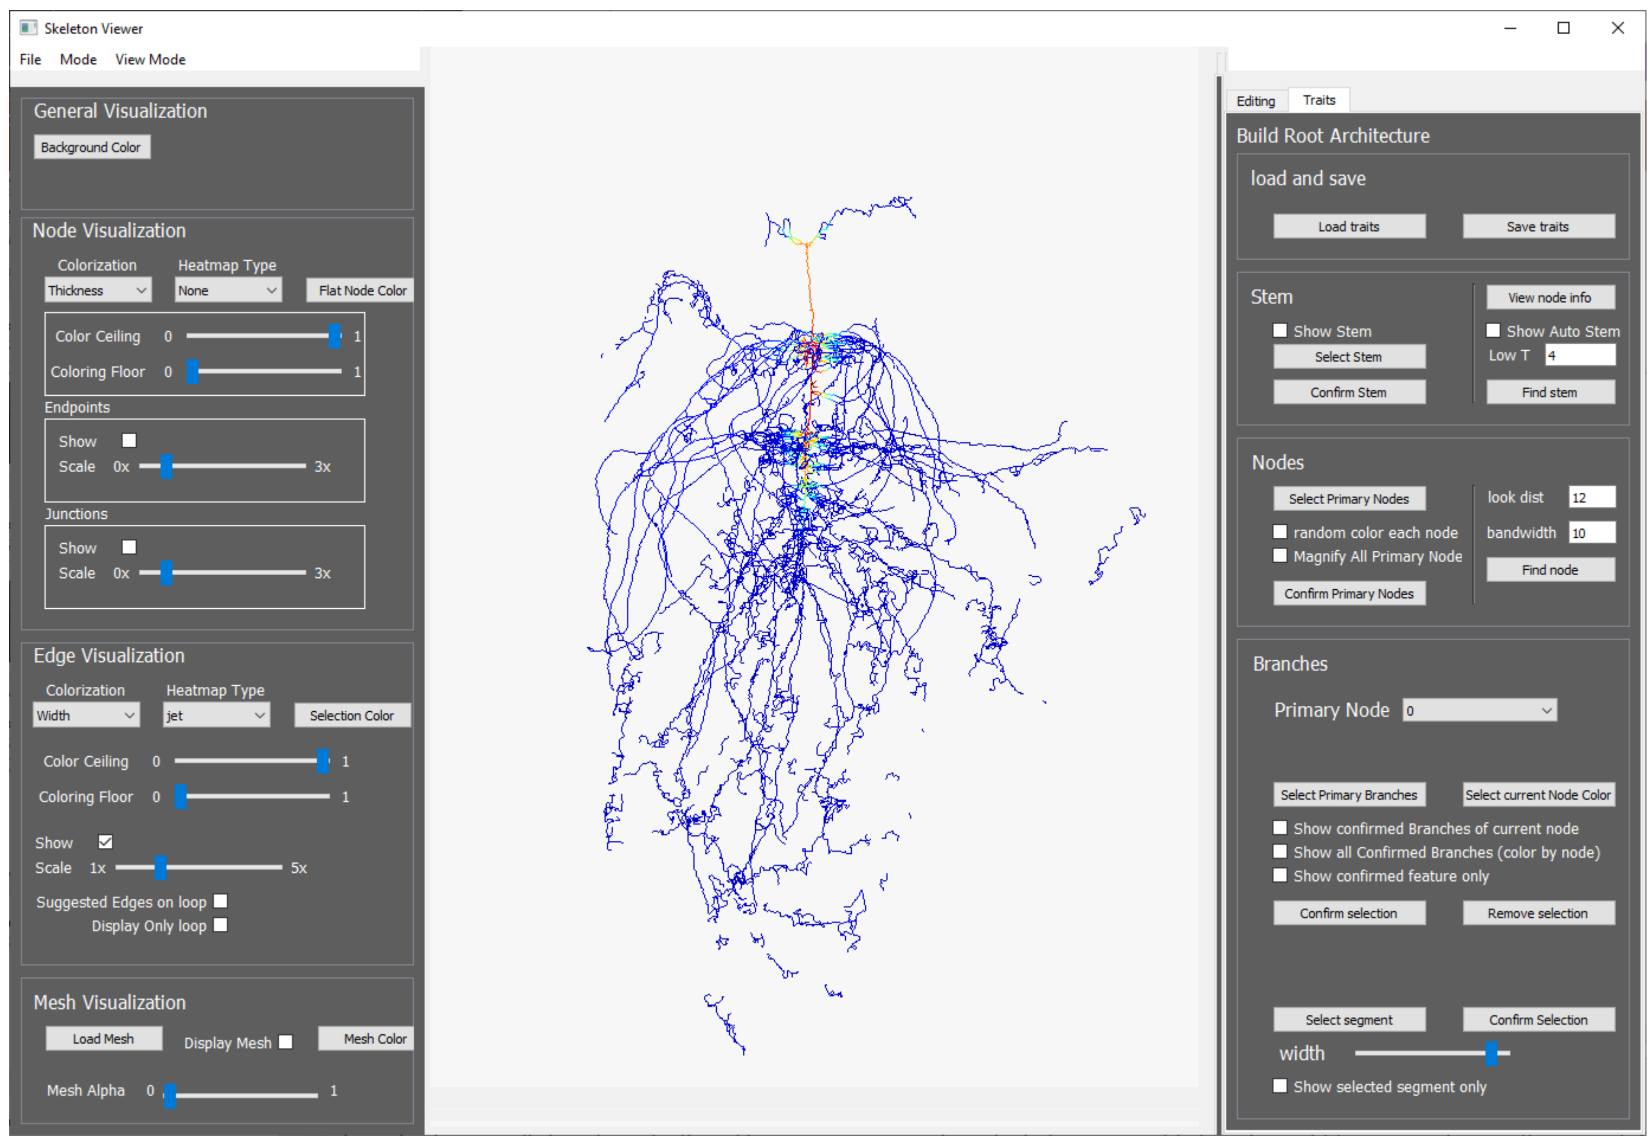

Preliminary tool

- Our preliminary tool allows hand identification of the primary nodes and branches, and then gives measurements for relevant features involving these.

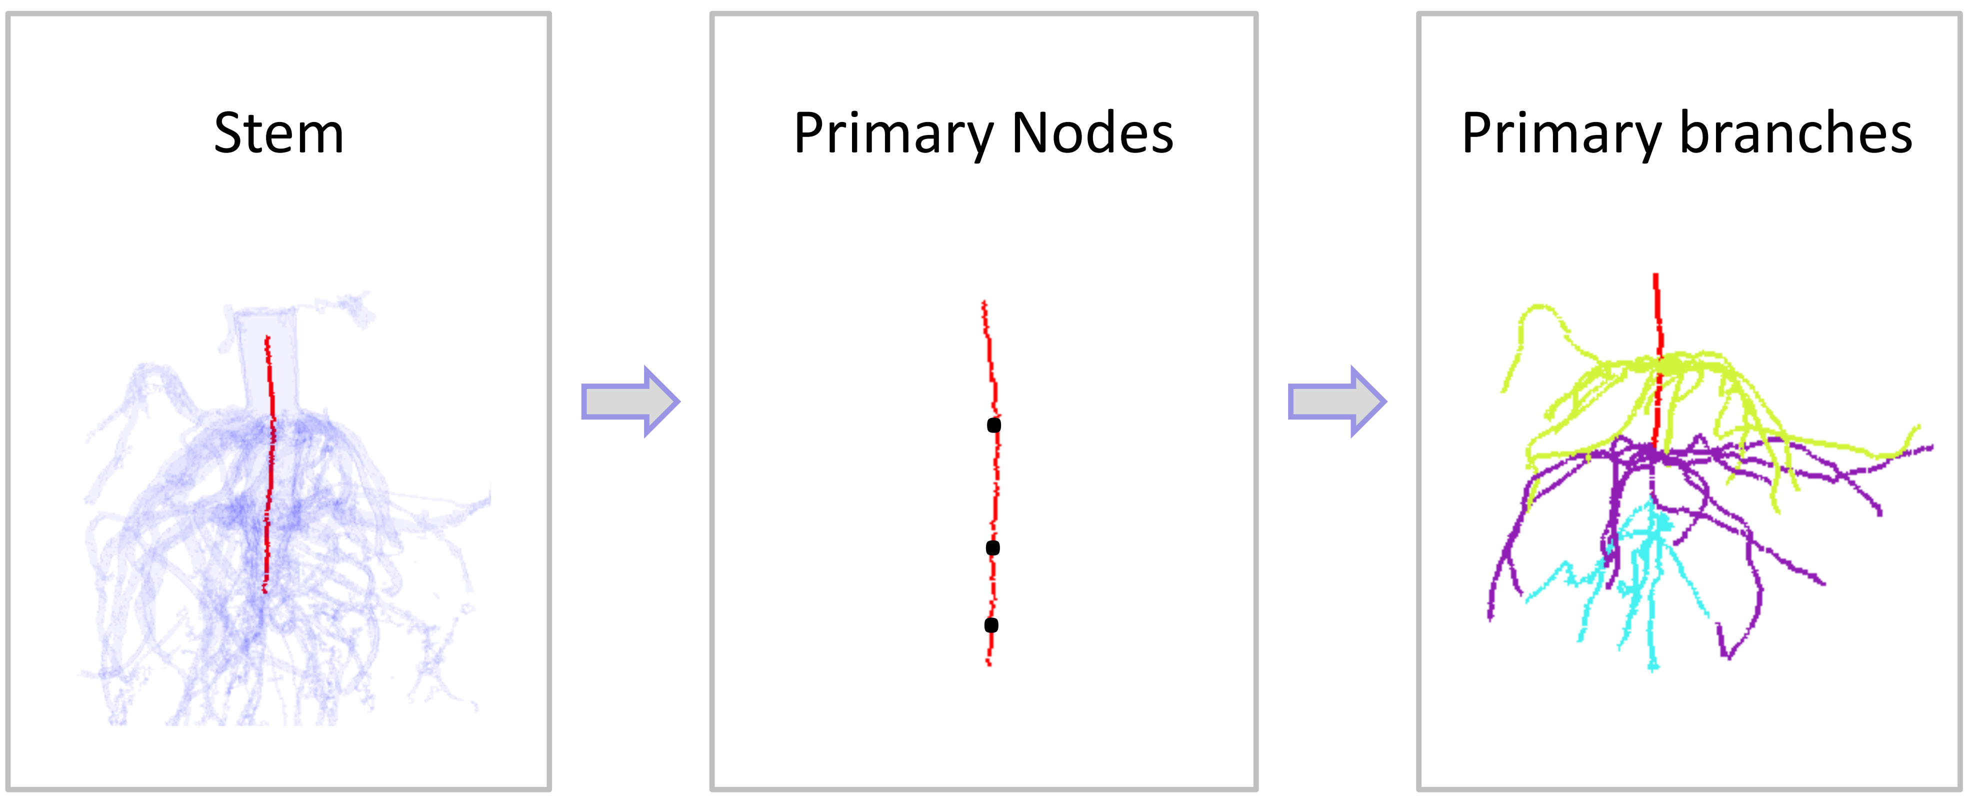

Auto-identification

Green: ground truth stem. White: stem inferred by software.

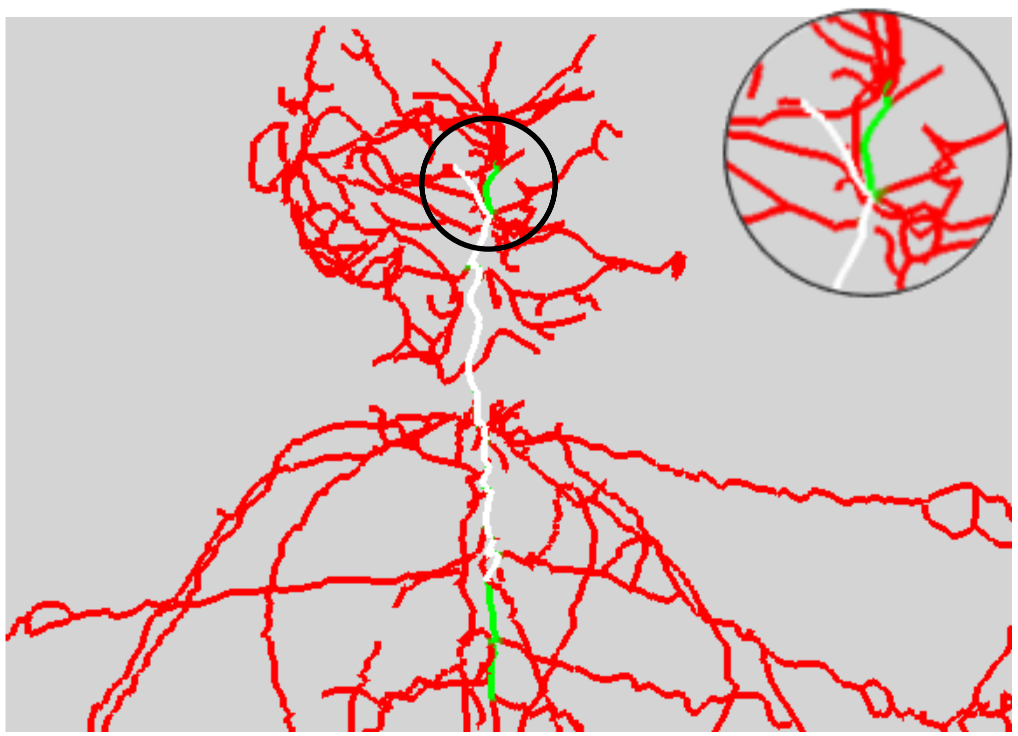

4 week root

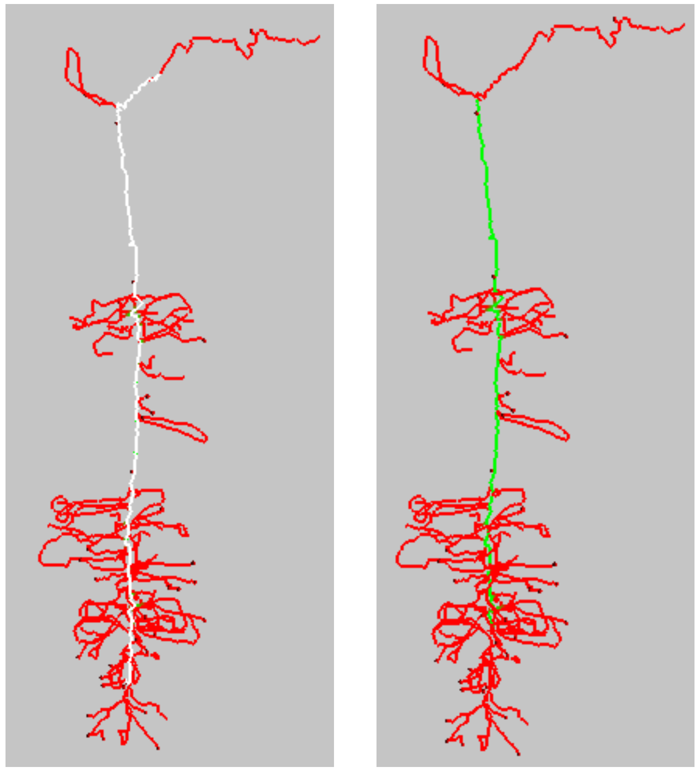

6 week root

Still to do: auto identification on more complex root systems

Four week corn root

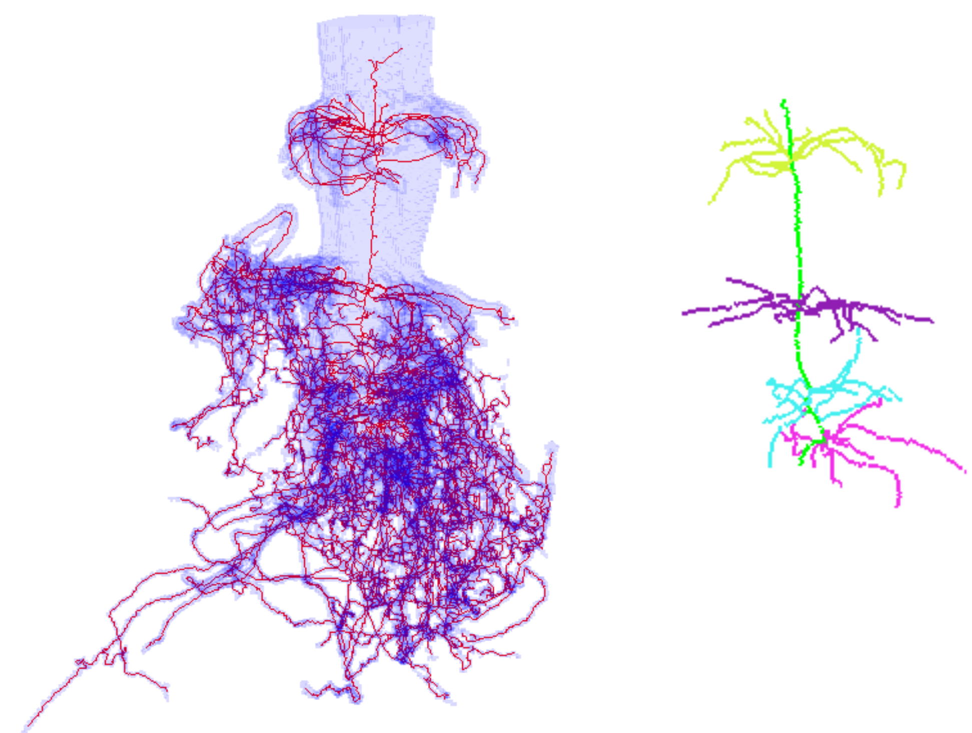

Flowering (mature) corn root

Plus more accurate shape repair tools

- Dirt in more complex roots continues to be a problem, as are “holes” in the roots

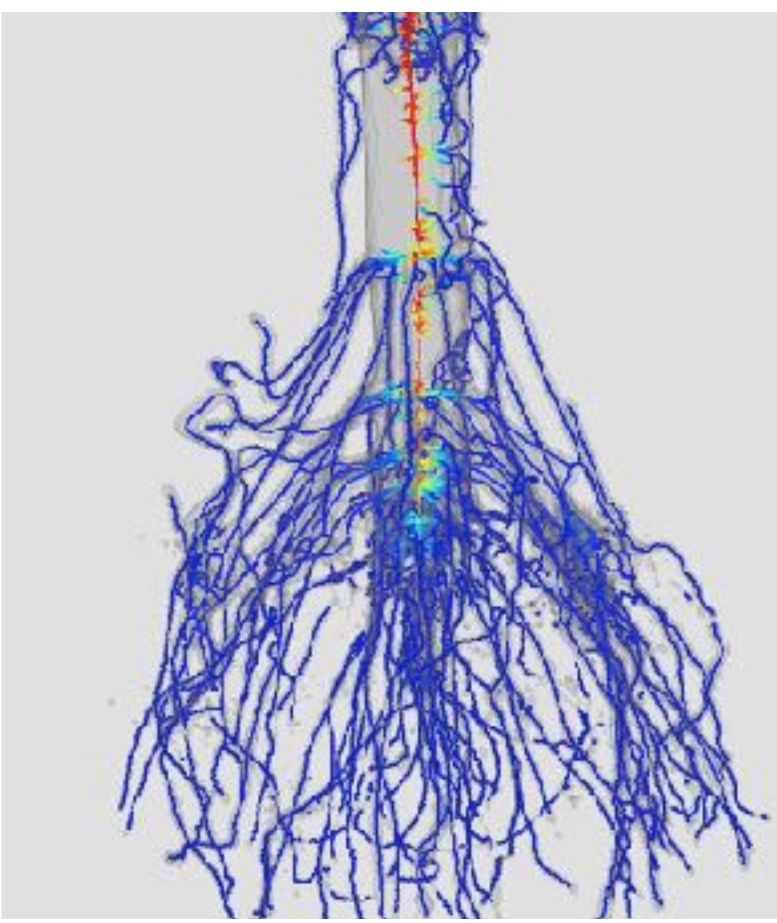

Root identification tool

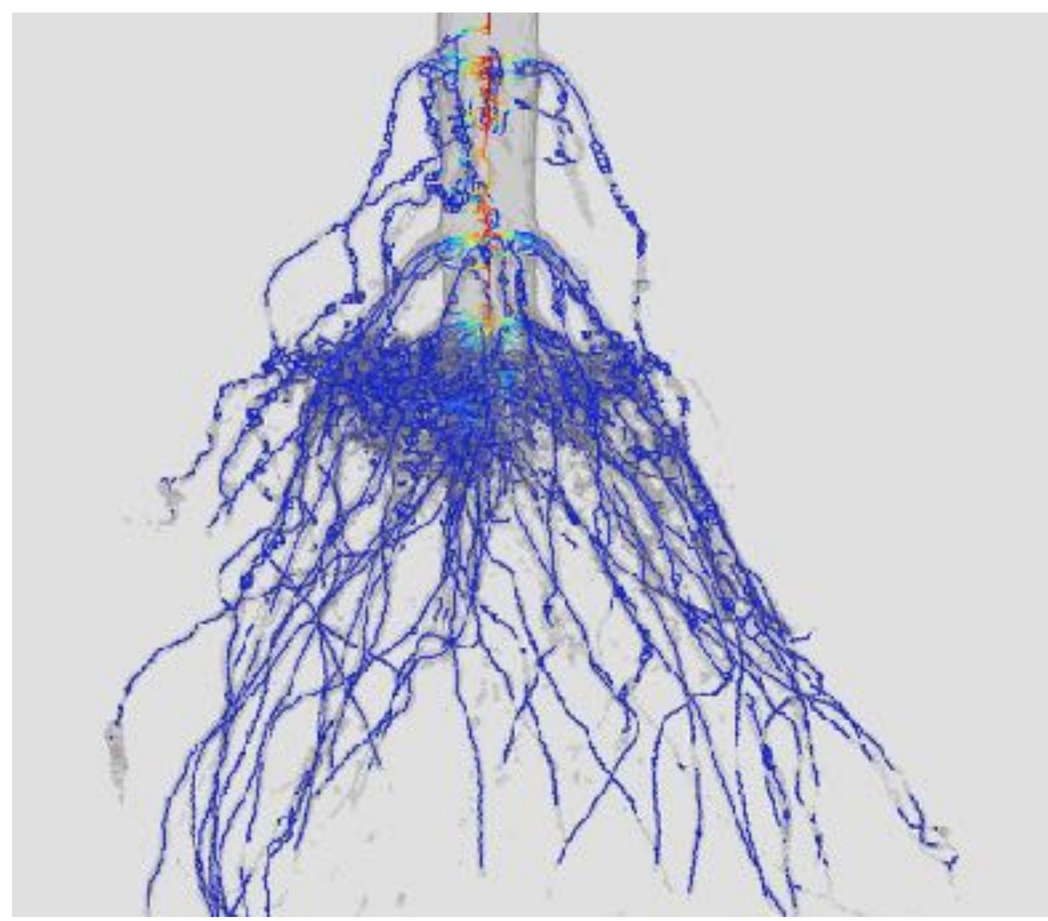

Messy and complicated root

Thanks!

- Questions?

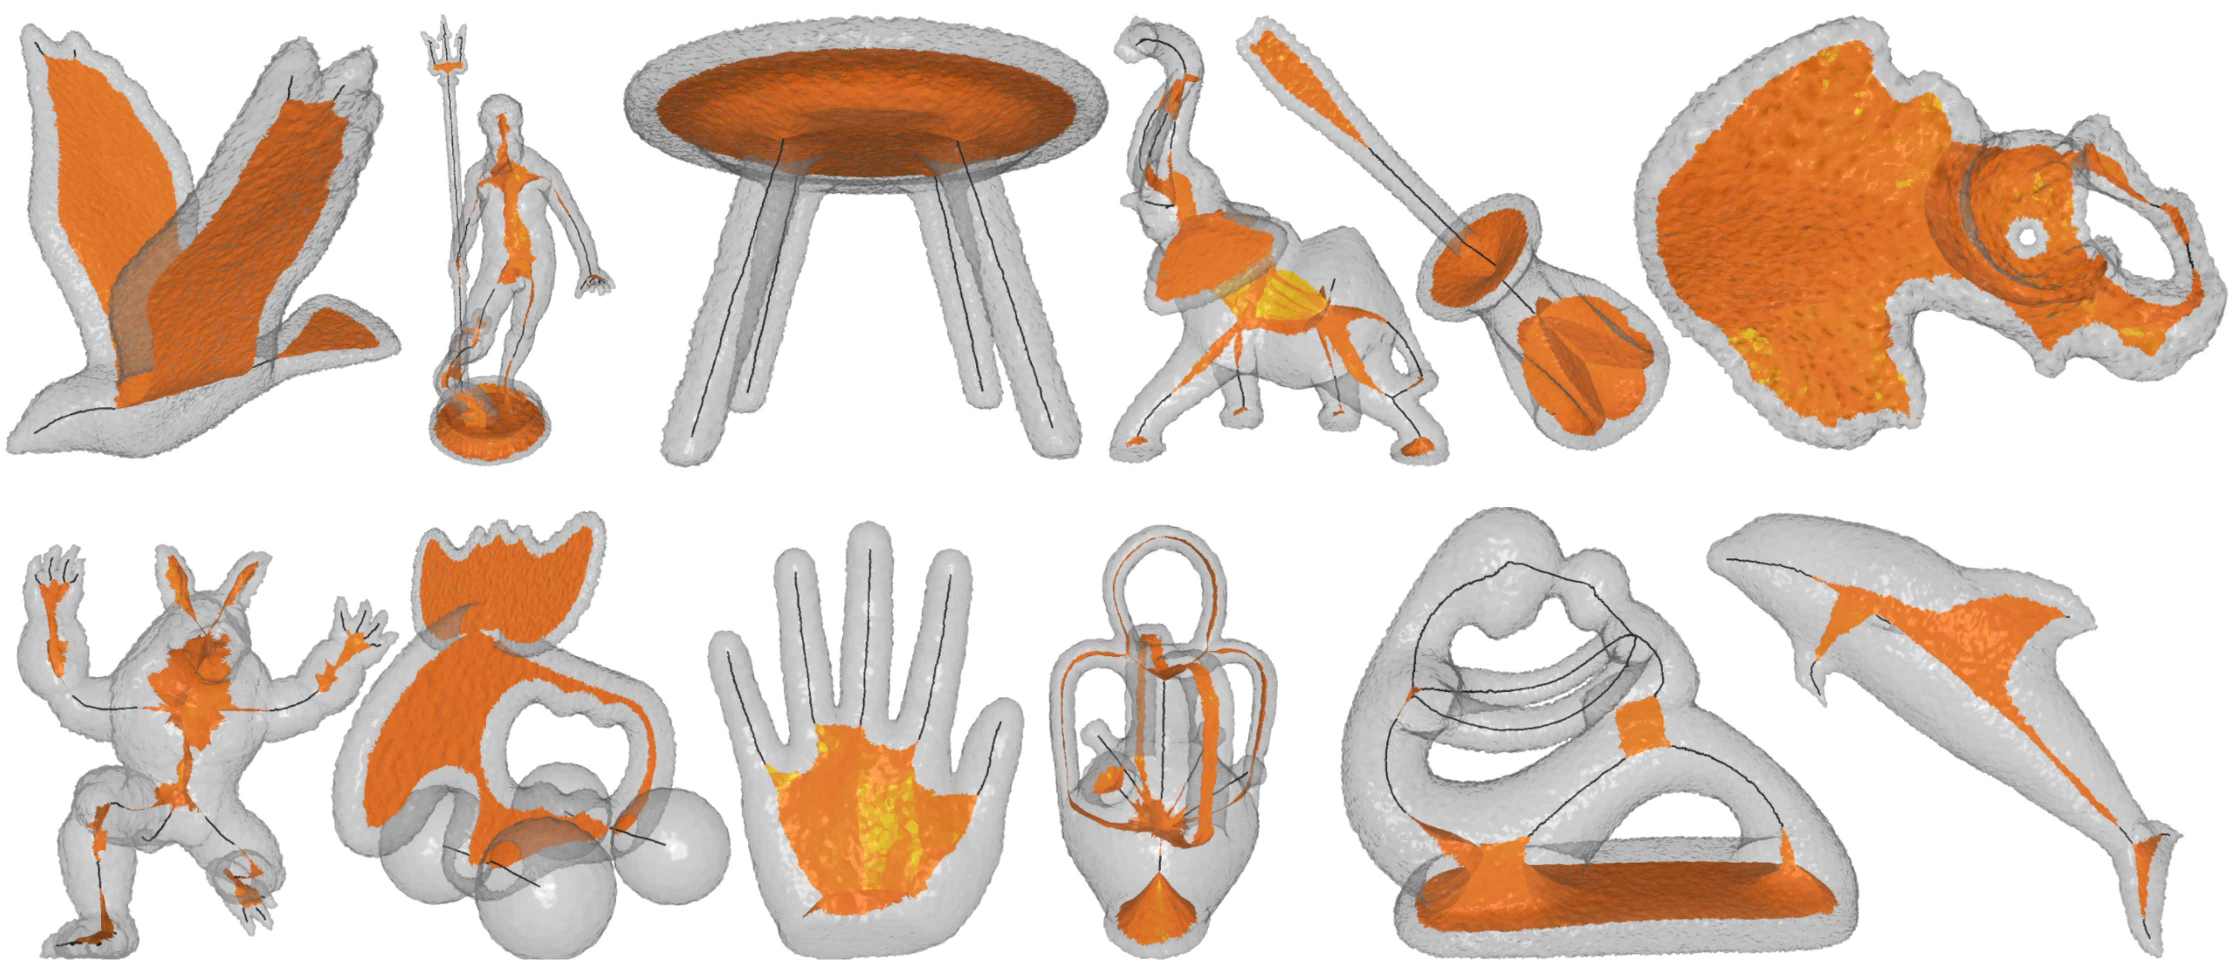

Extra slides