Shape Reconstruction

Pierre Alliez

Inria

Outline

- Sensors
- Problem statement
- Computational Geometry
  - Convex hull, Voronoi/Delaunay, alpha-shapes
- Variational formulations
  - Poisson / spectral

SENSORS

Laser scanning

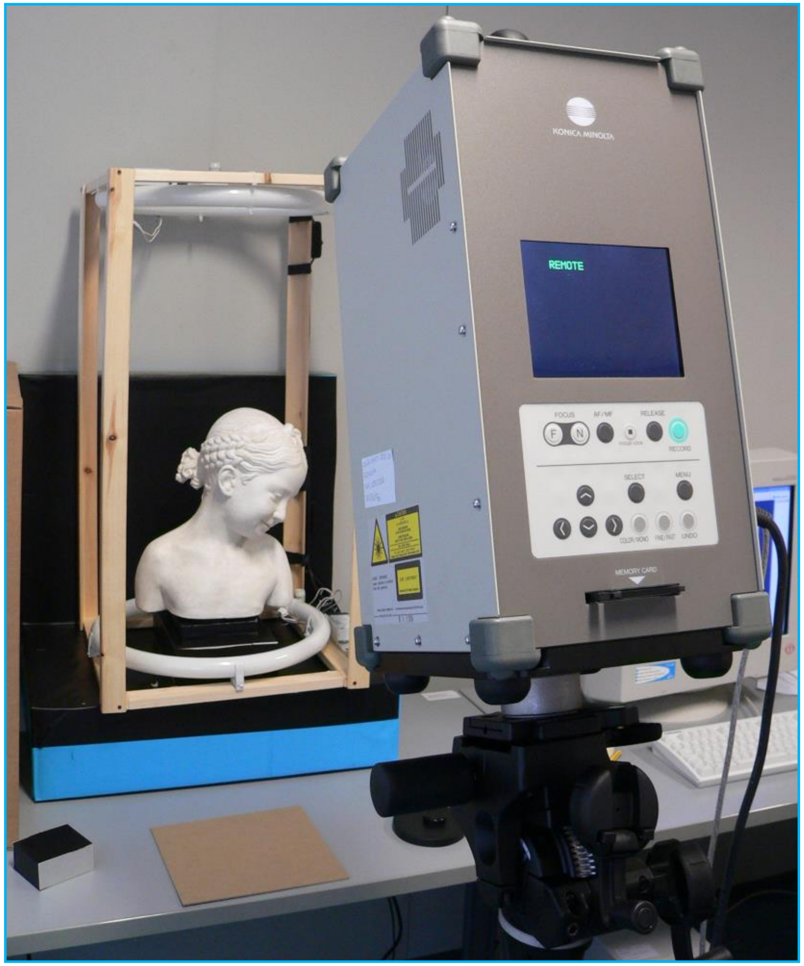

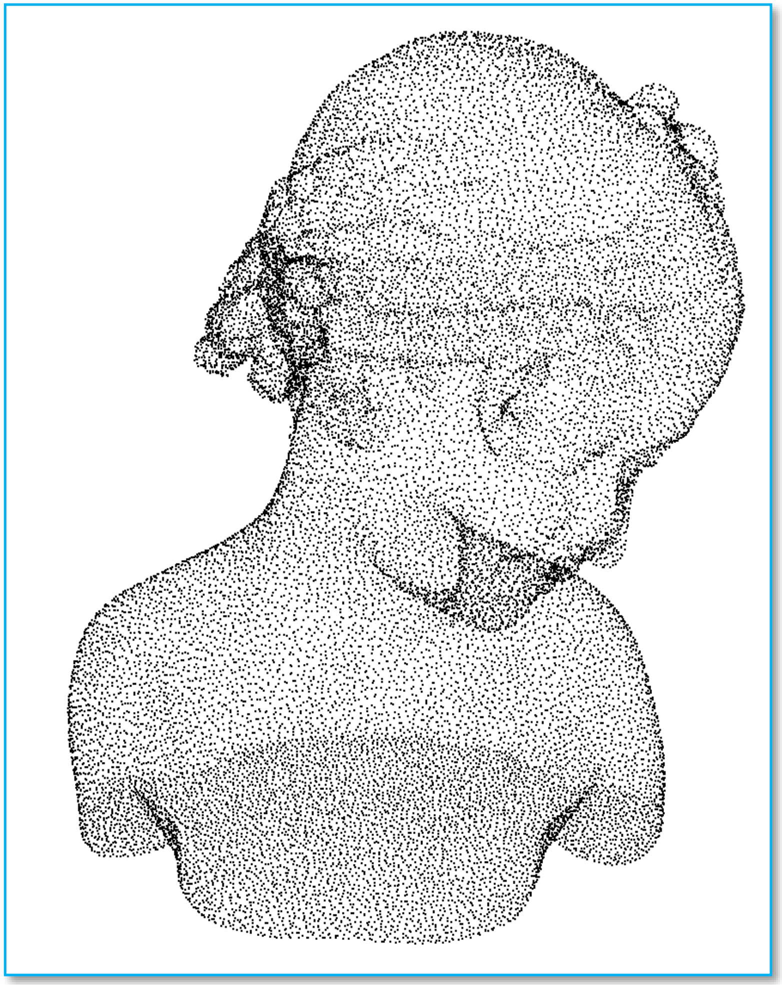

Car-based Laser

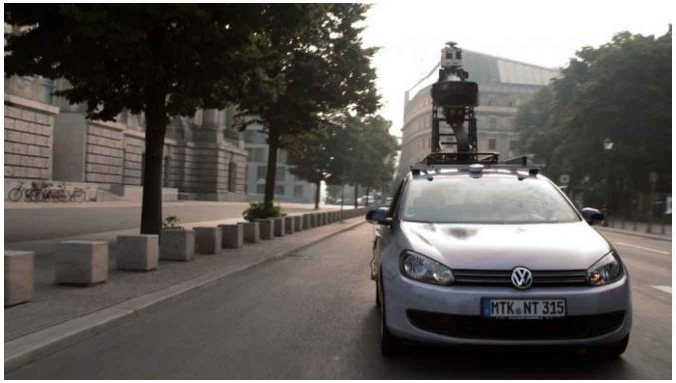

Airborne Lidar

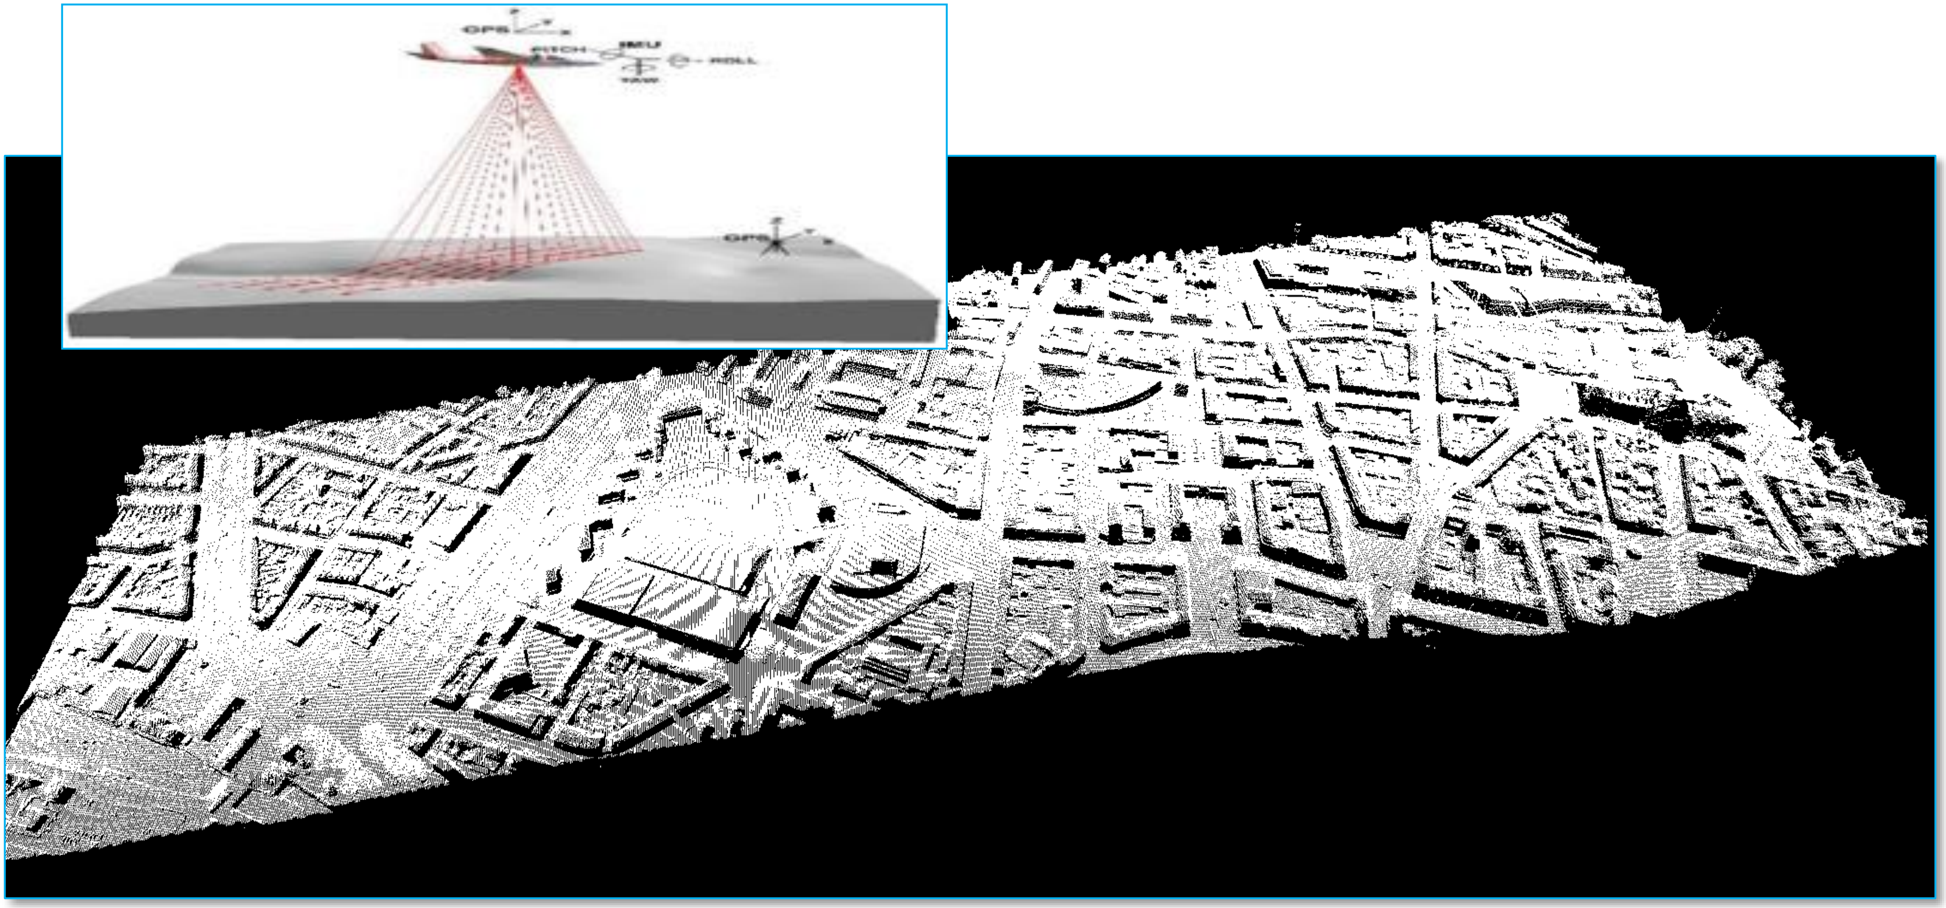

Multi-View Stereo (MVS)

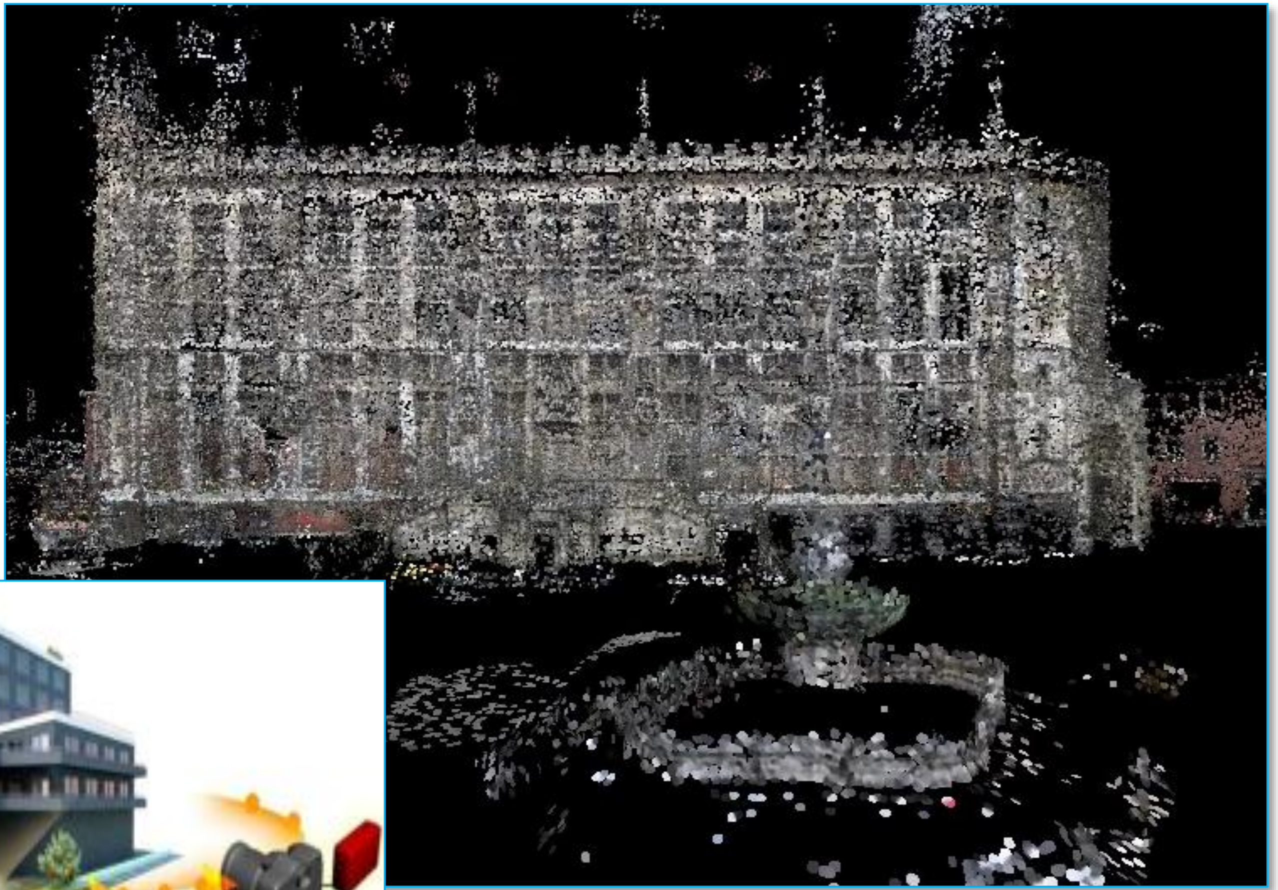

Depth Sensors

PROBLEM STATEMENT

Reconstruction Problem

Input: point set *P* sampled over a surface *S*:

- Non-uniform sampling
- With holes
- With uncertainty (noise)

Output: surface

- Approximation of *S* in terms of topology and geometry
- Desired:
  - Watertight
  - Intersection free

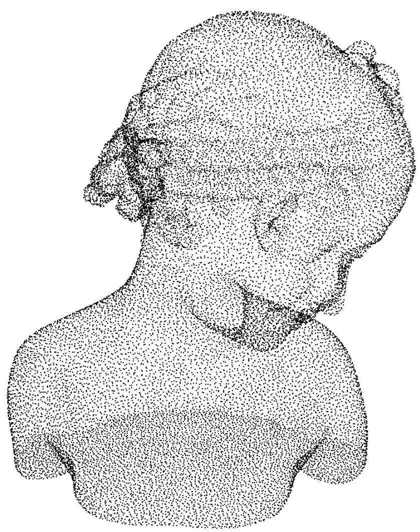

point set

reconstruction

surface

Ill-posed Problem

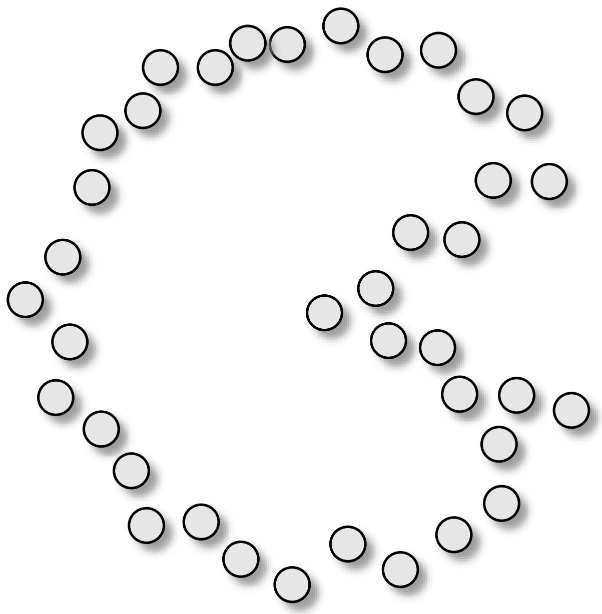

Many candidate surfaces for the reconstruction problem!

Ill-posed Problem

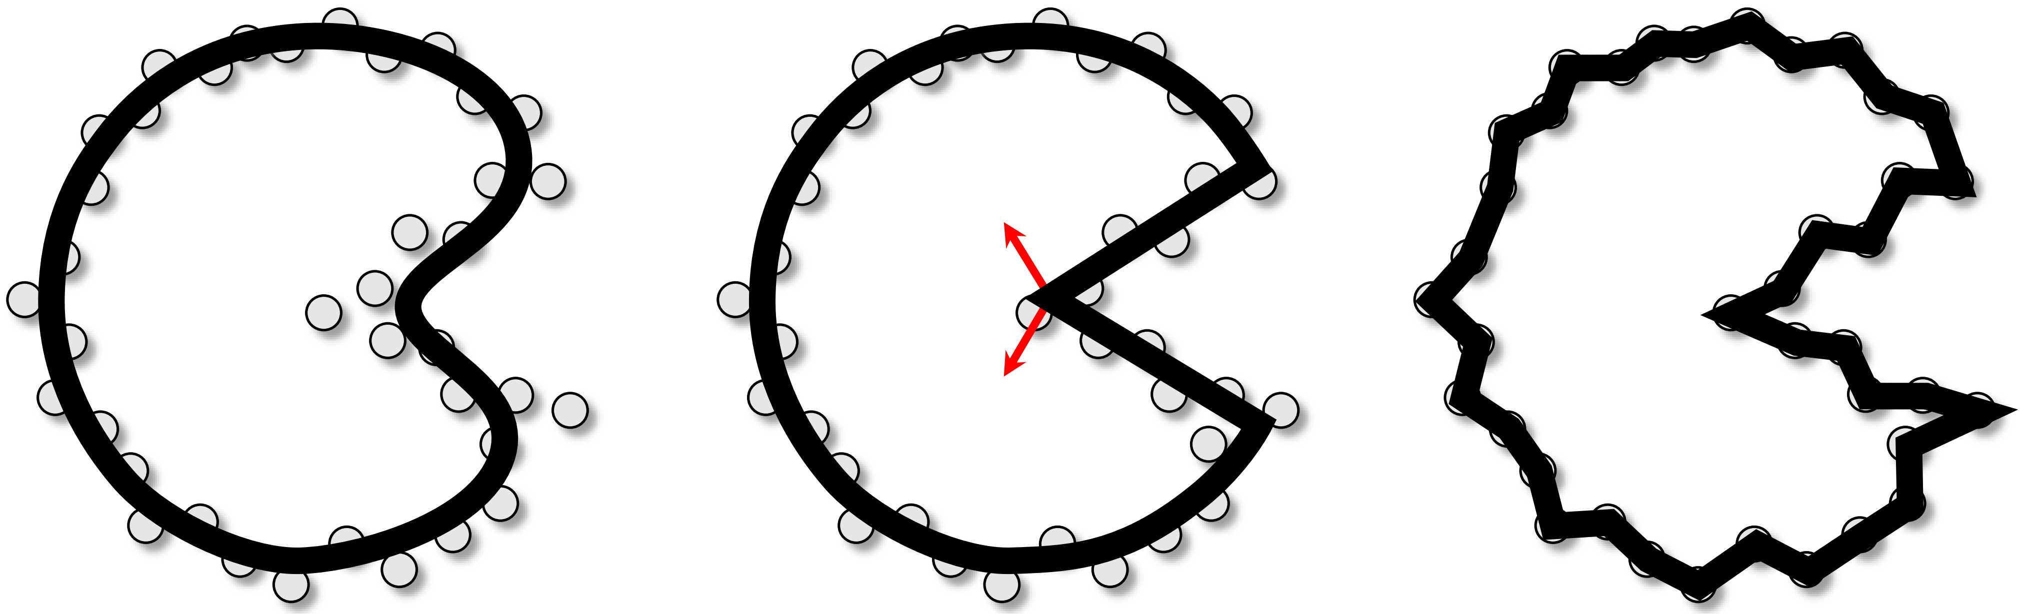

Many candidate surfaces for the reconstruction problem! How to pick?

Priors

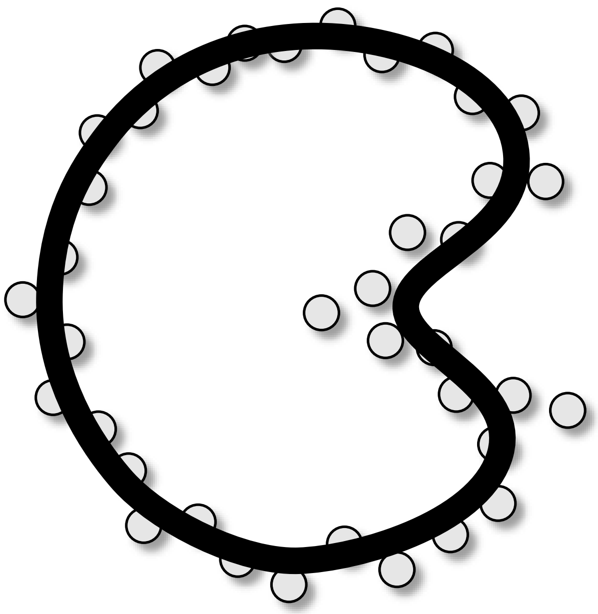

Smooth

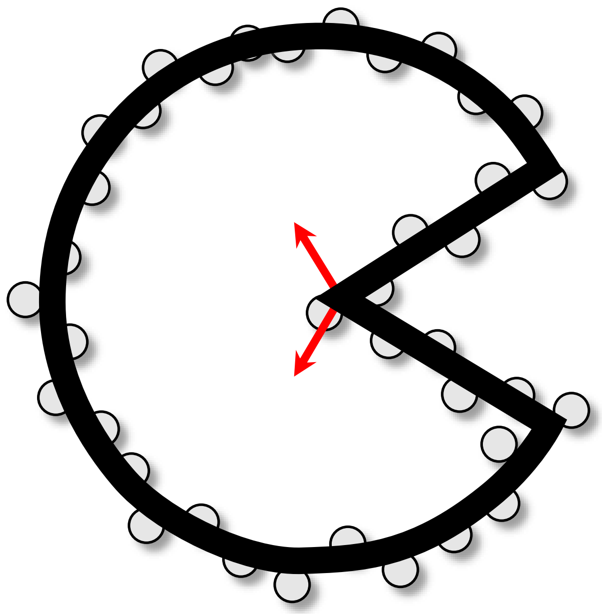

Piecewise Smooth

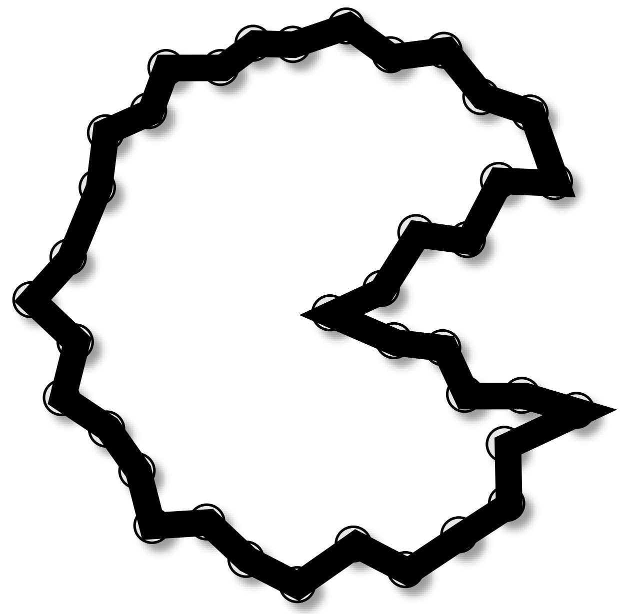

“Simple”

Surface Smoothness Priors

Local Smoothness

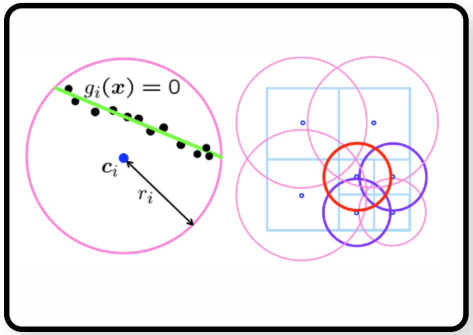

Local fitting
No control away from data
Solution by interpolation

Global Smoothness

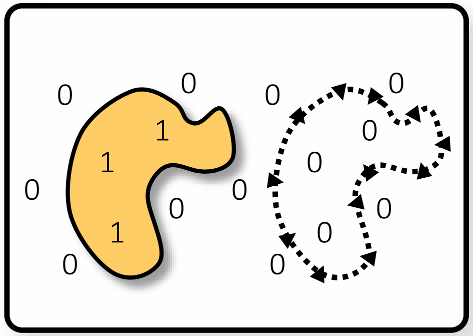

Global: linear, eigen, graph cut, …
Robustness to missing data

Piecewise Smoothness

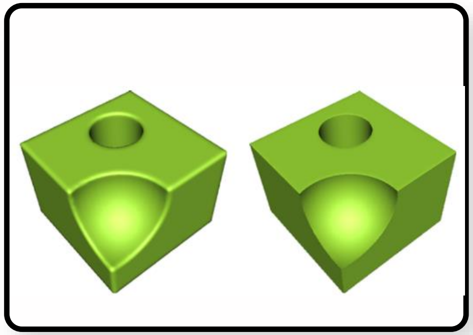

Sharp near features
Smooth away from features

Domain-Specific Priors

Surface Reconstruction
by Point Set Structuring

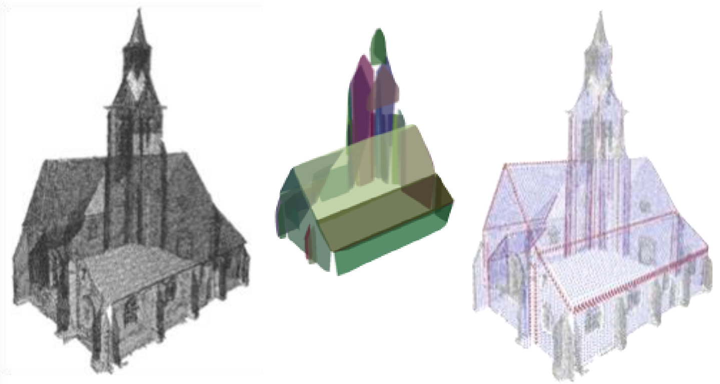

[Lafarge - A. EUROGRAPHICS 2013]

LOD Reconstruction
for Urban Scenes

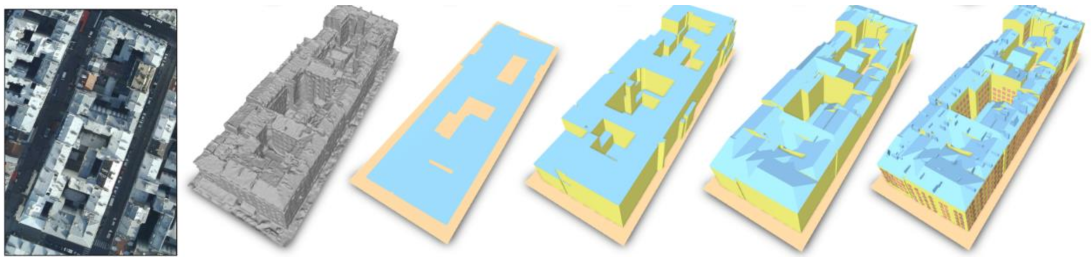

[Verdie, Lafarge - A. ACM Transactions on Graphics 2015]

Warm-up

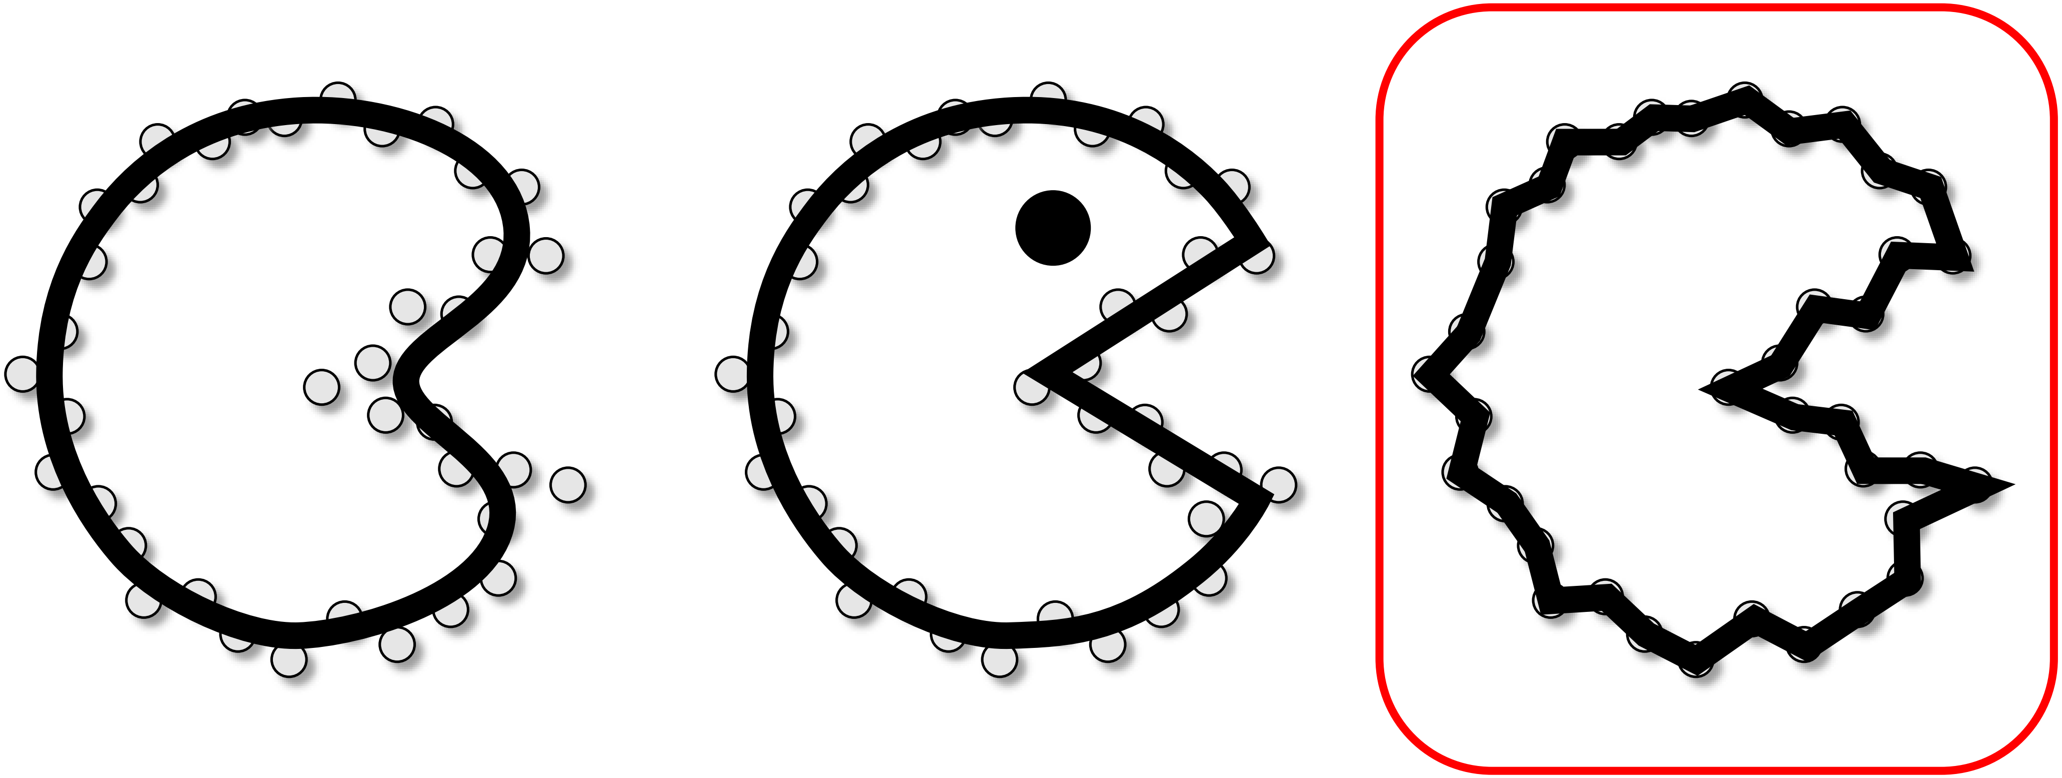

Smooth | Piecewise Smooth | “Simple”

CONVEX HULL

Convex Hull

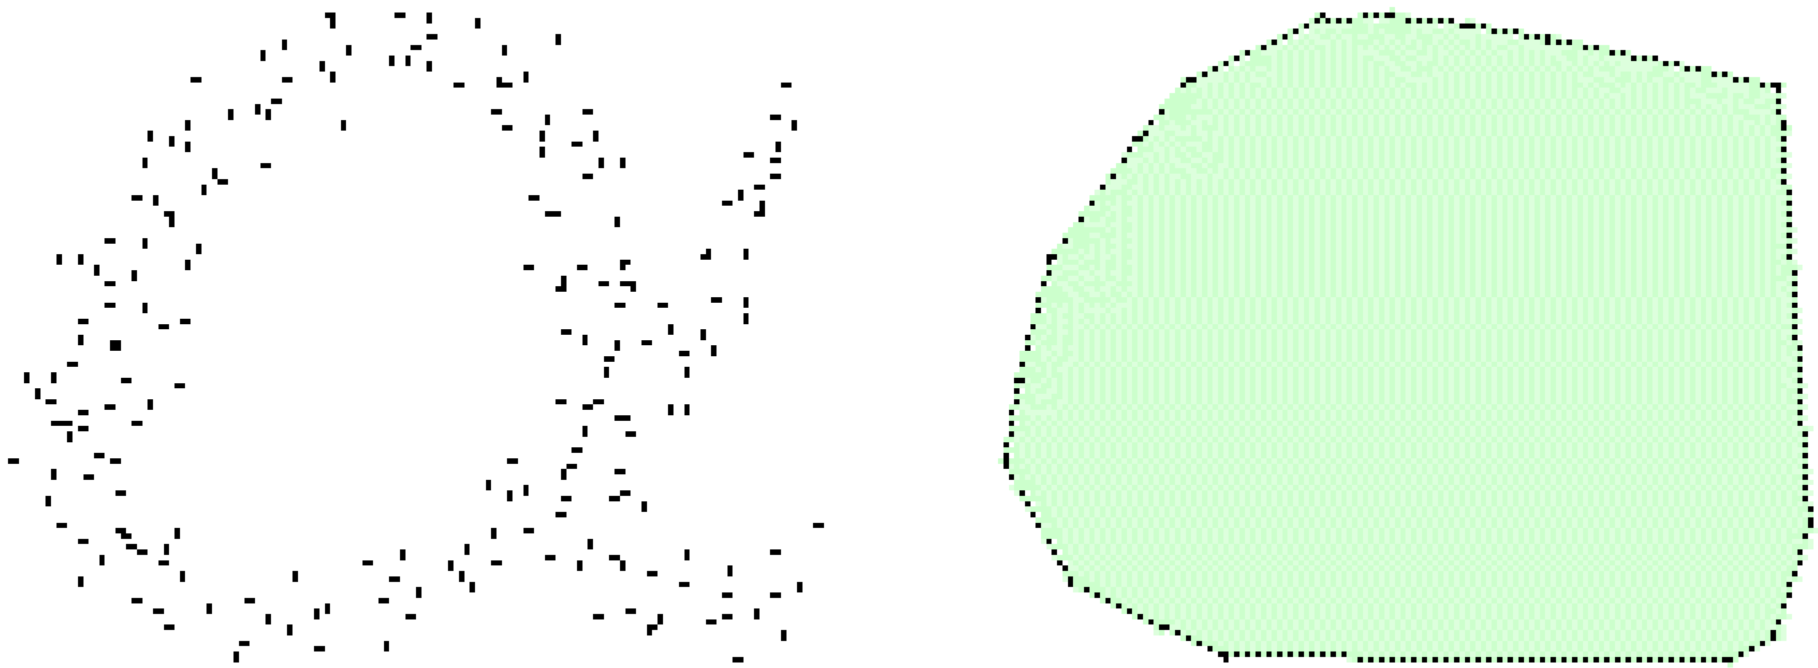

VORONOI / DELAUNAY

Voronoi Diagram

Let $\mathcal{E} = \{\mathbf{p_1}, \ldots, \mathbf{p_n}\}$ be a set of points (so-called sites) in $\mathbb{R}^d$. We associate to each site $\mathbf{p_i}$ its Voronoi region $V(\mathbf{p_i})$ such that:

$$V(\mathbf{p_i}) = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x} - \mathbf{p_i}\| \leq \|\mathbf{x} - \mathbf{p_j}\|, \forall j \leq n\}.$$

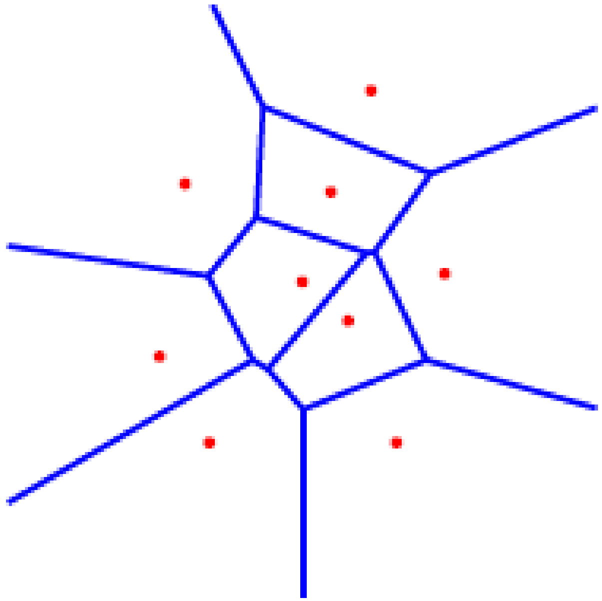

Delaunay Triangulation

Dual structure of the Voronoi diagram.

The Delaunay triangulation of a set of sites E is a simplicial complex such that k+1 points in E form a Delaunay simplex if their Voronoi cells have nonempty intersection

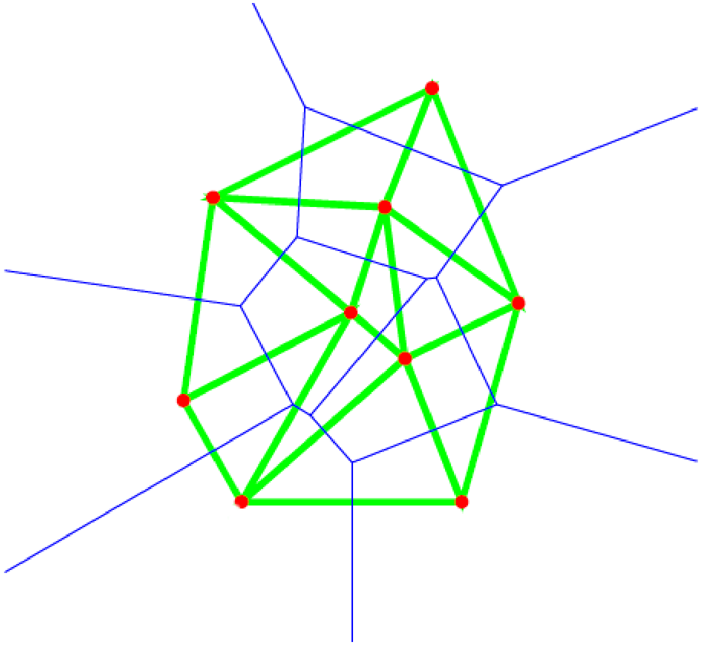

Delaunay-based

Key idea: assuming dense enough sampling, reconstructed triangles are Delaunay triangles.

First define

Medial axis

Local feature size

Epsilon-sampling

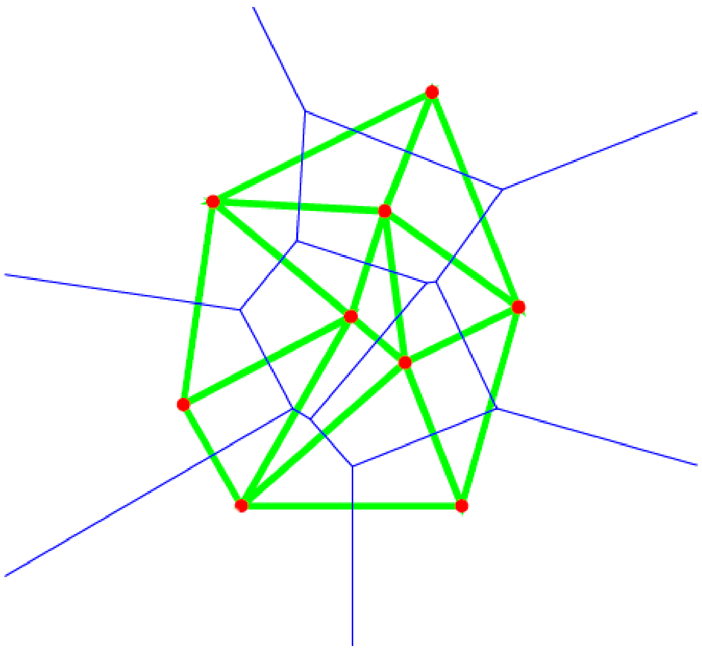

Alpha-Shapes [Edelsbrunner, Kirkpatrick, Seidel]

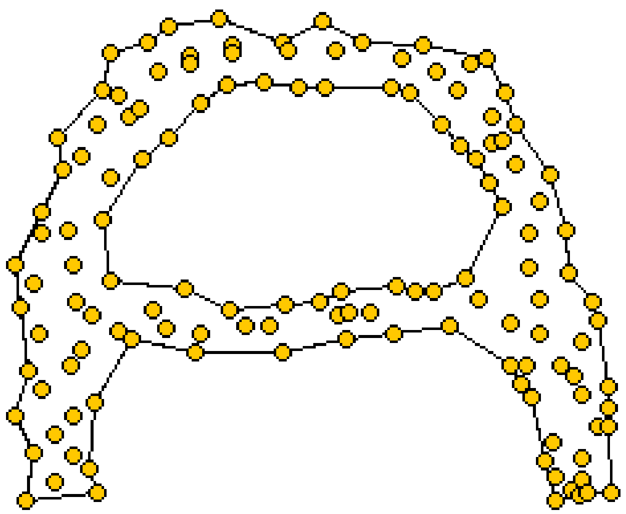

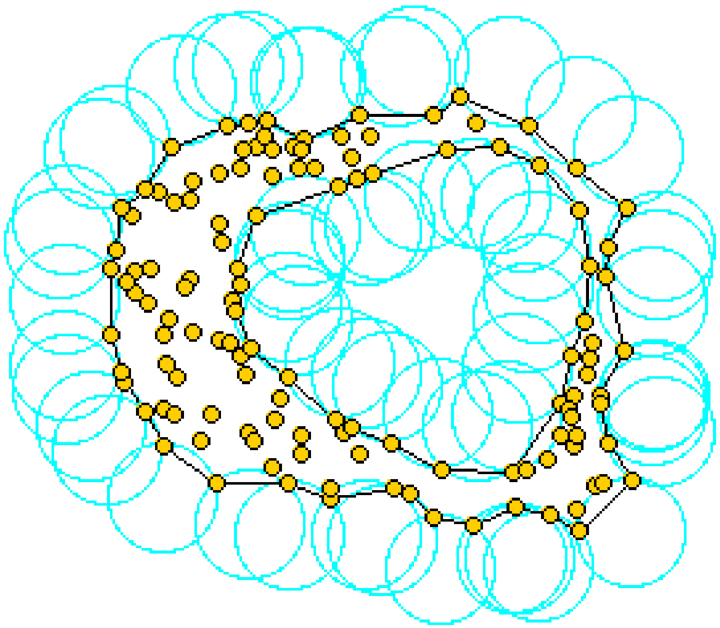

Segments: point pairs that can be touched by an empty disc of radius alpha.

Alpha-Shapes

In 2D: family of piecewise linear simple curves constructed from a point set P.

Subcomplex of the Delaunay triangulation of P.

Generalization of the concept of the convex hull.

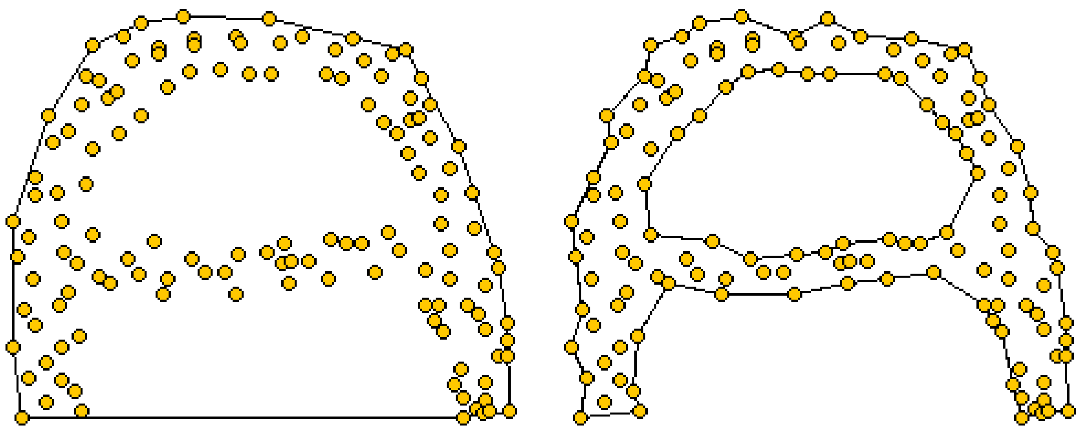

Alpha-Shapes

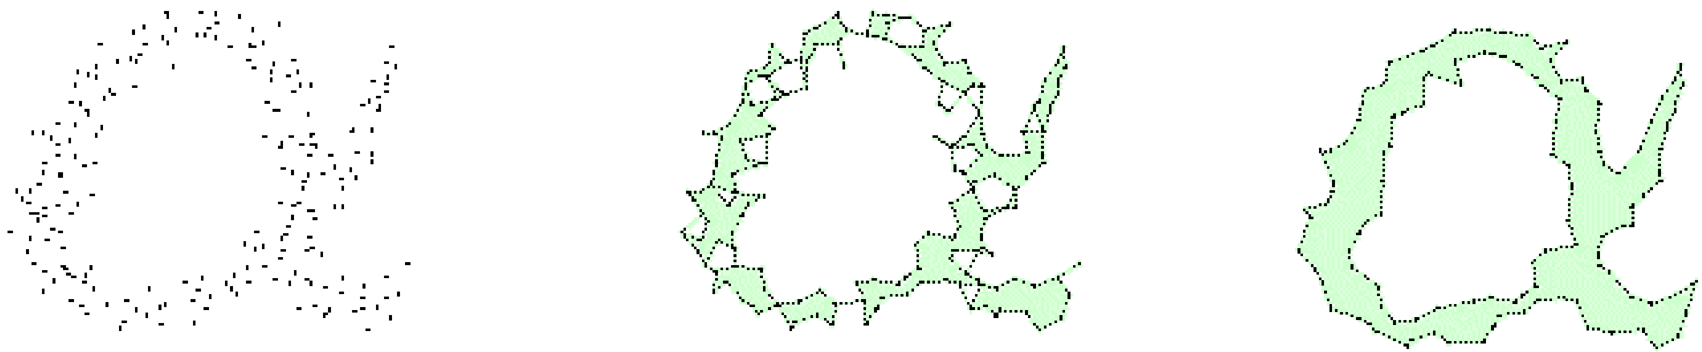

$\alpha = 0$

Alpha controls the desired level of detail.

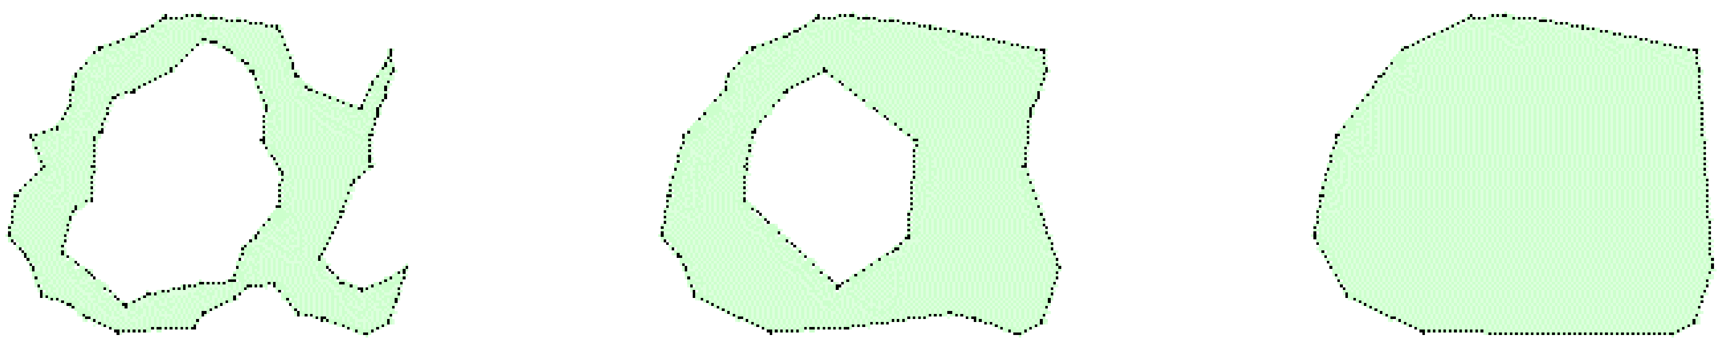

$\alpha = \infty$

Computing Alpha-Shapes

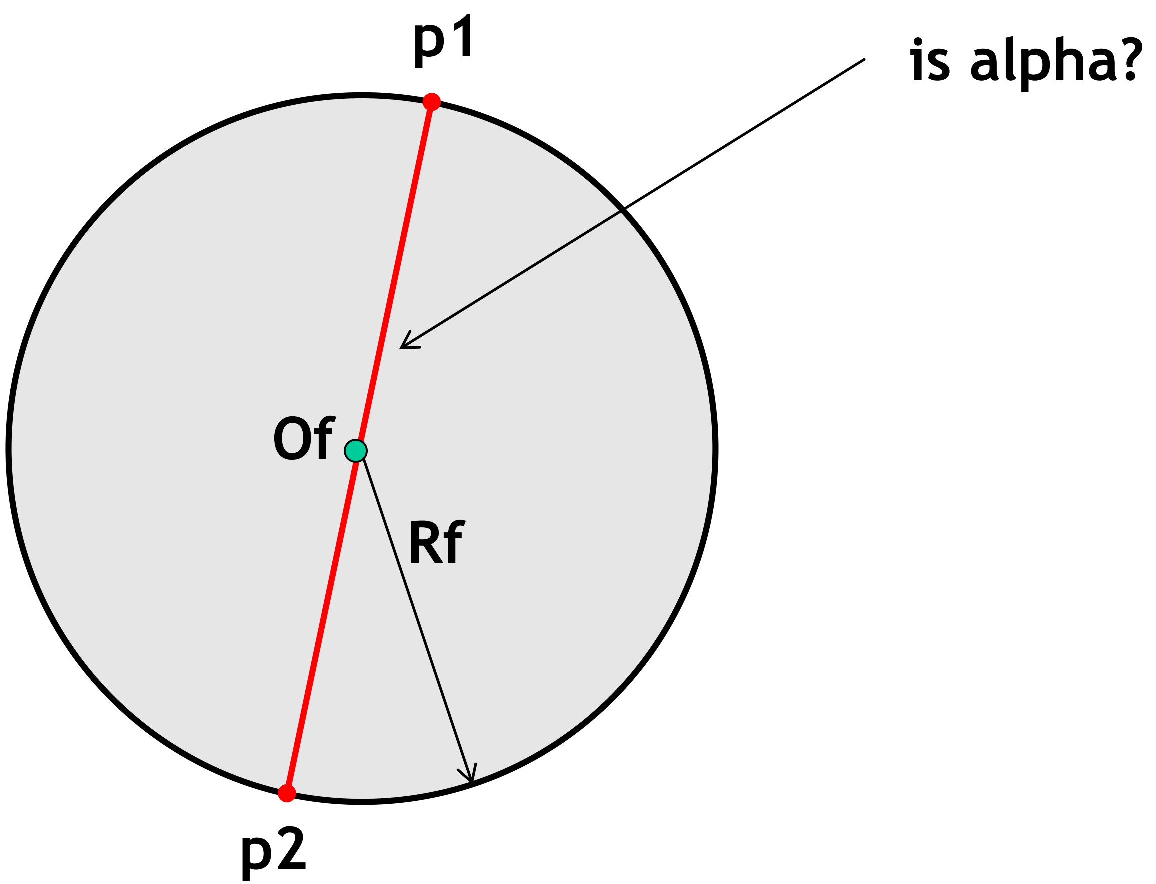

Computing Alpha-Shapes

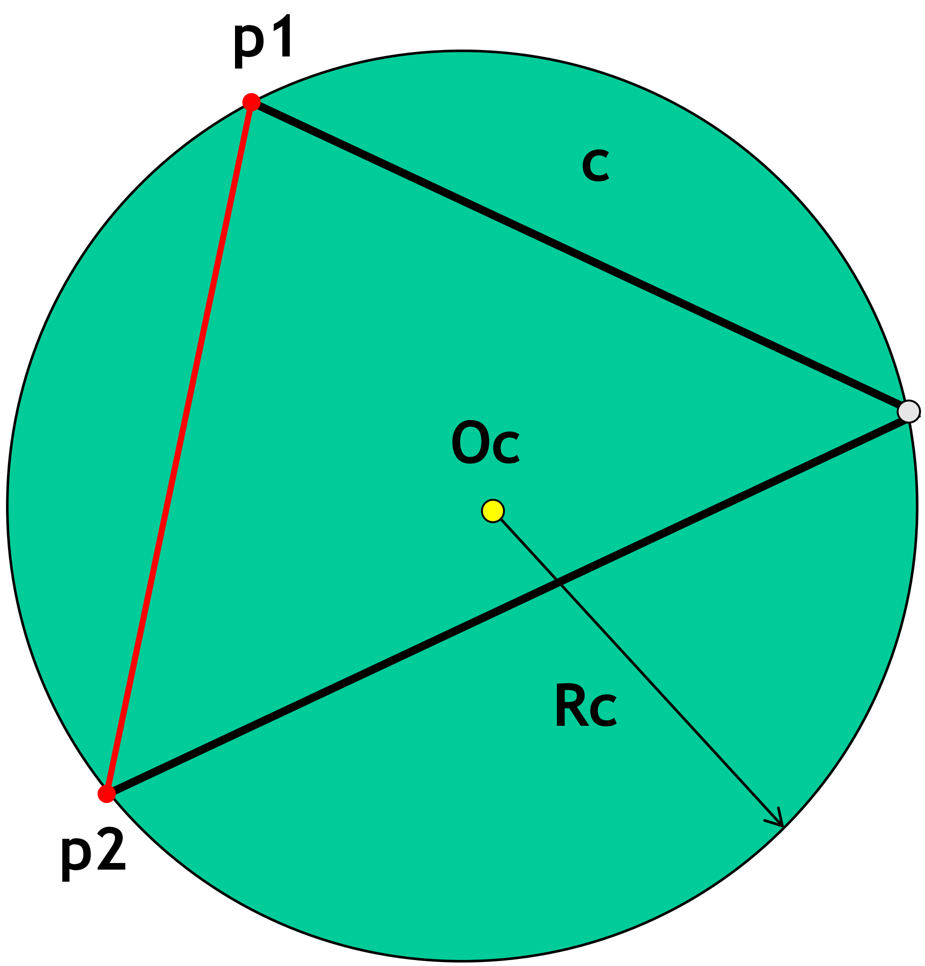

Computing Alpha-Shapes

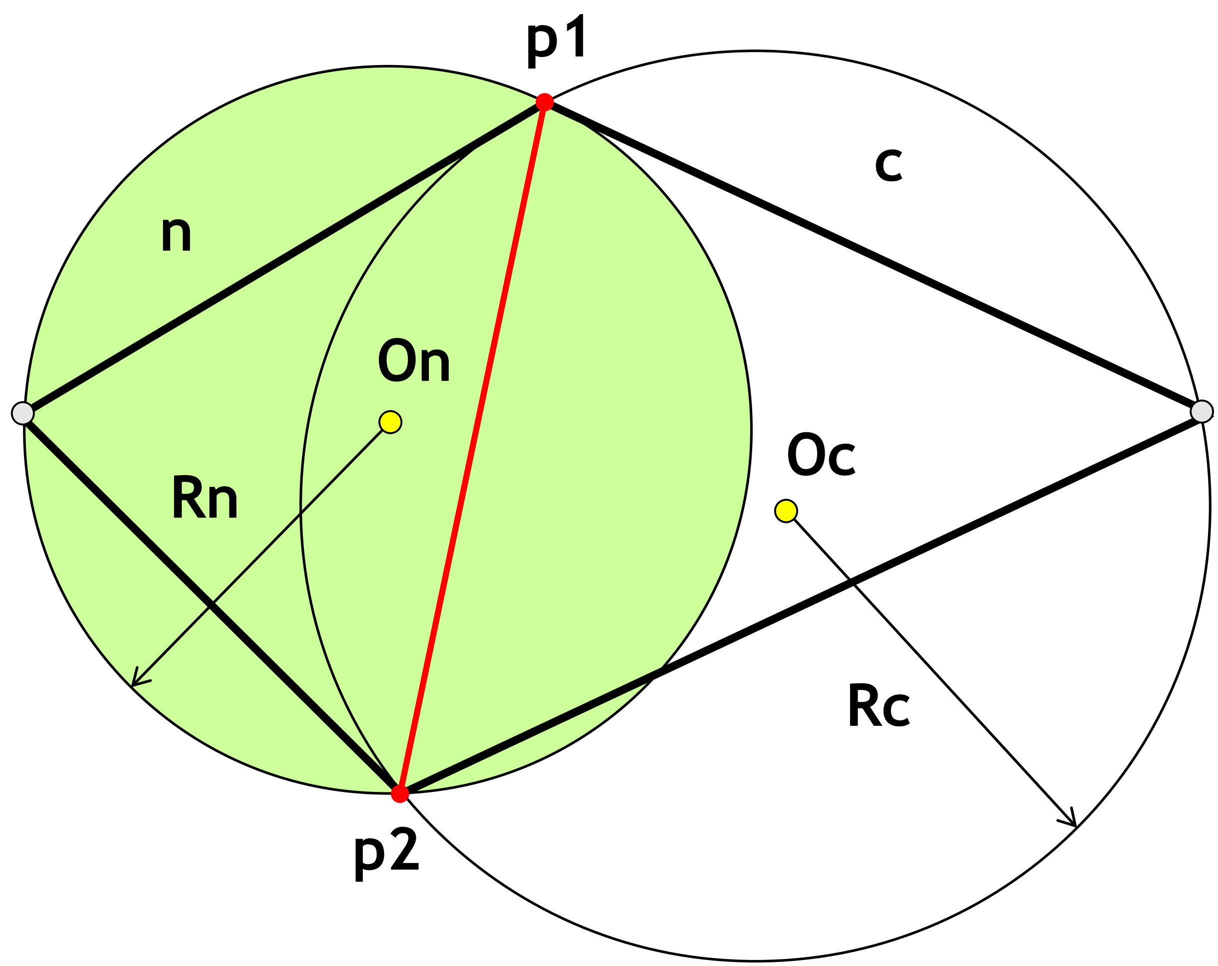

Computing Alpha-Shapes

is_alpha if
alpha > Rf
AND
alpha <
max(Rc,Rn)

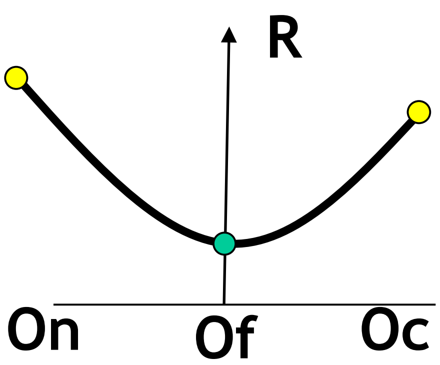

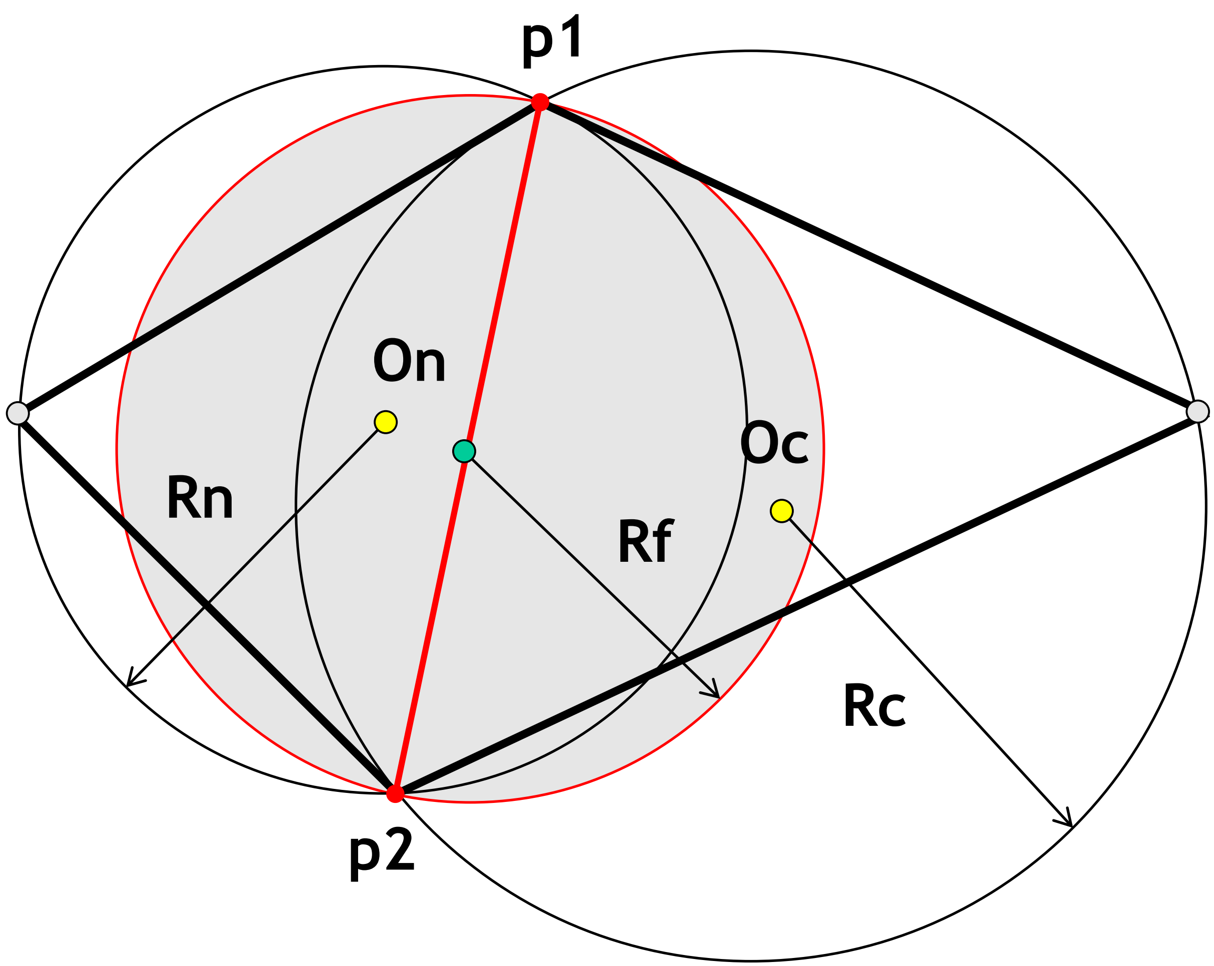

Computing Alpha-Shapes

is_alpha if
alpha > min(Rc,Rn)
AND
alpha < max(Rc,Rn)

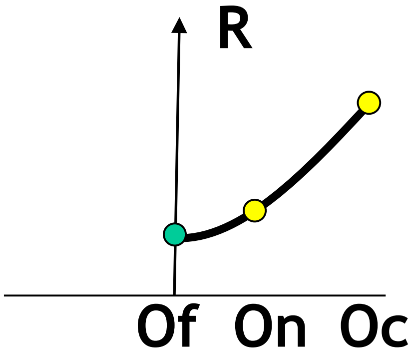

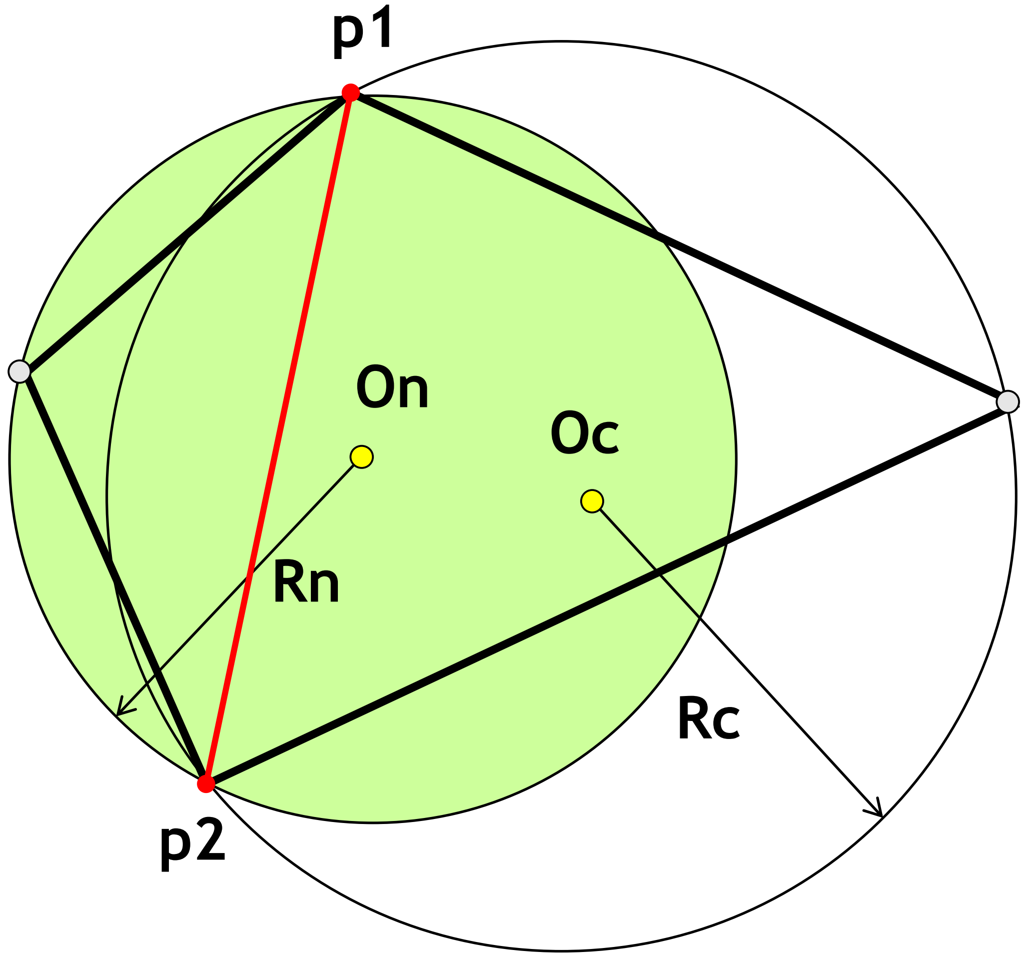

Computing Alpha-Shapes

is_alpha if
alpha > Rf

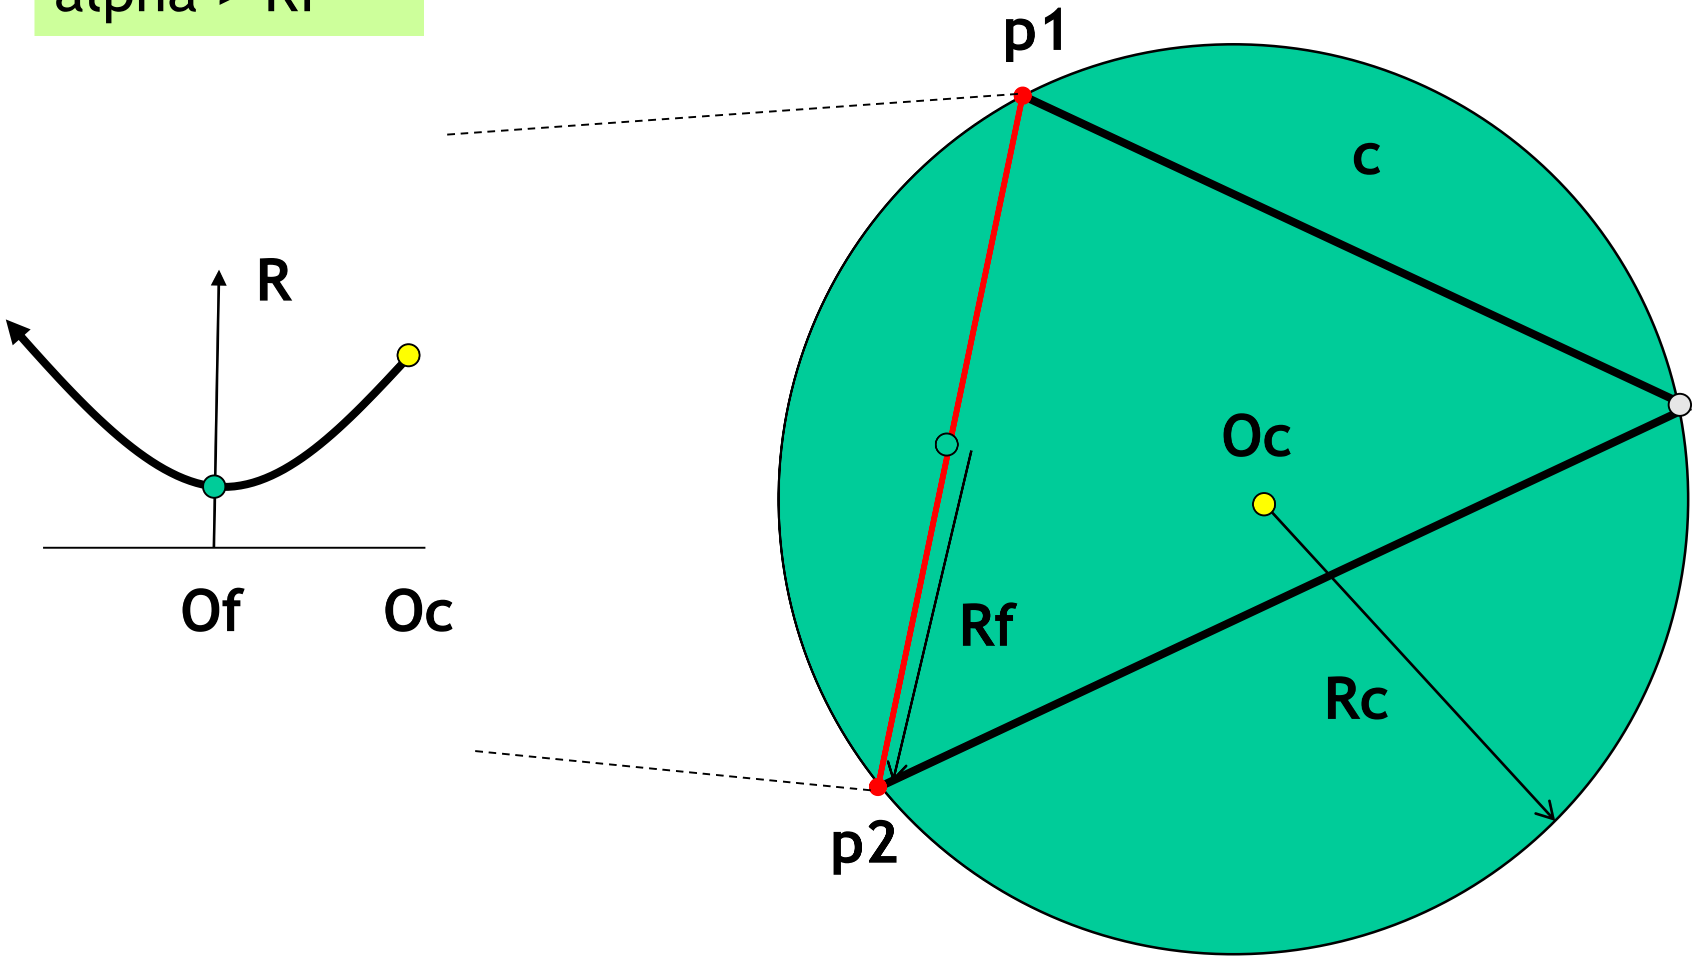

MEDIAL AXIS

Medial Axis

For a shape (curve/surface) a *Medial Ball* is a circle/sphere that only meets the shape tangentially, in at least two points.

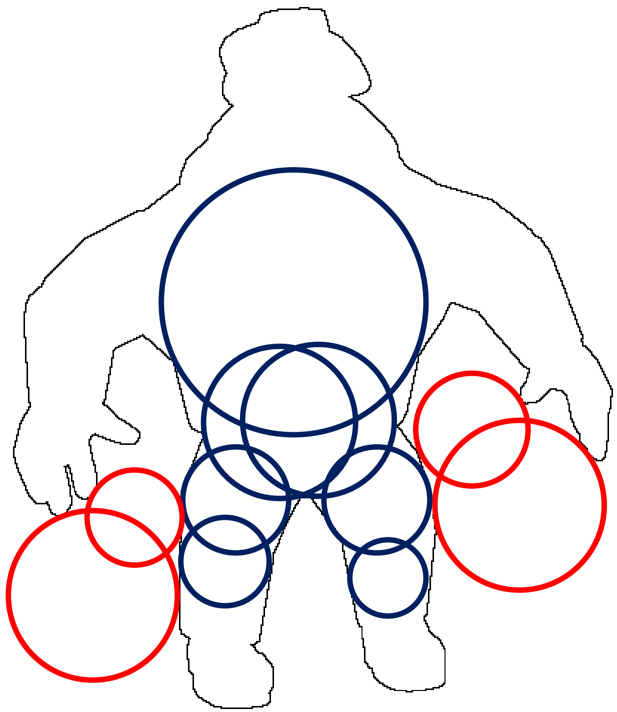

Medial Axis

For a shape (curve/surface) a *Medial Ball* is a circle/sphere that only meets the shape tangentially, in at least two points.

The centers of all such balls make up the *medial axis/skeleton*.

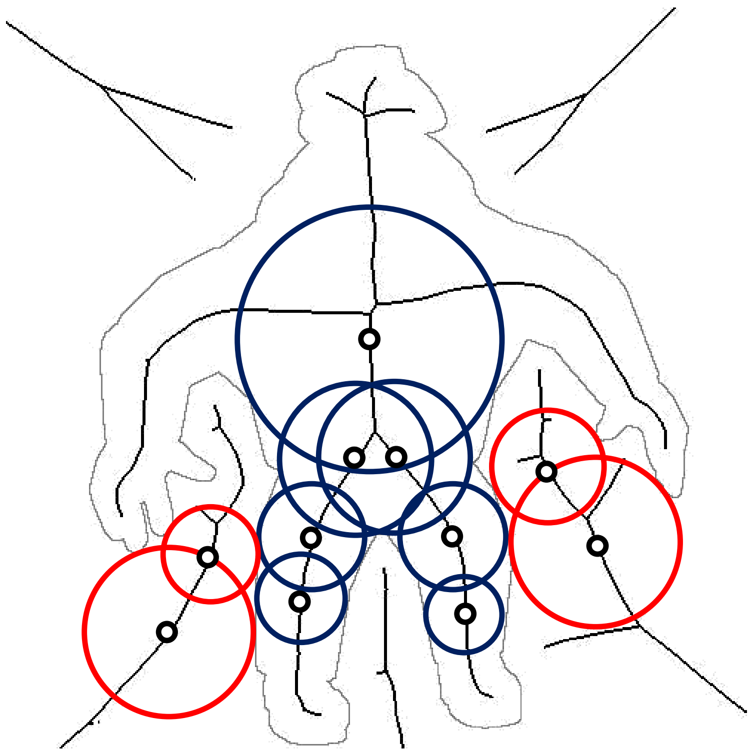

Medial Axis

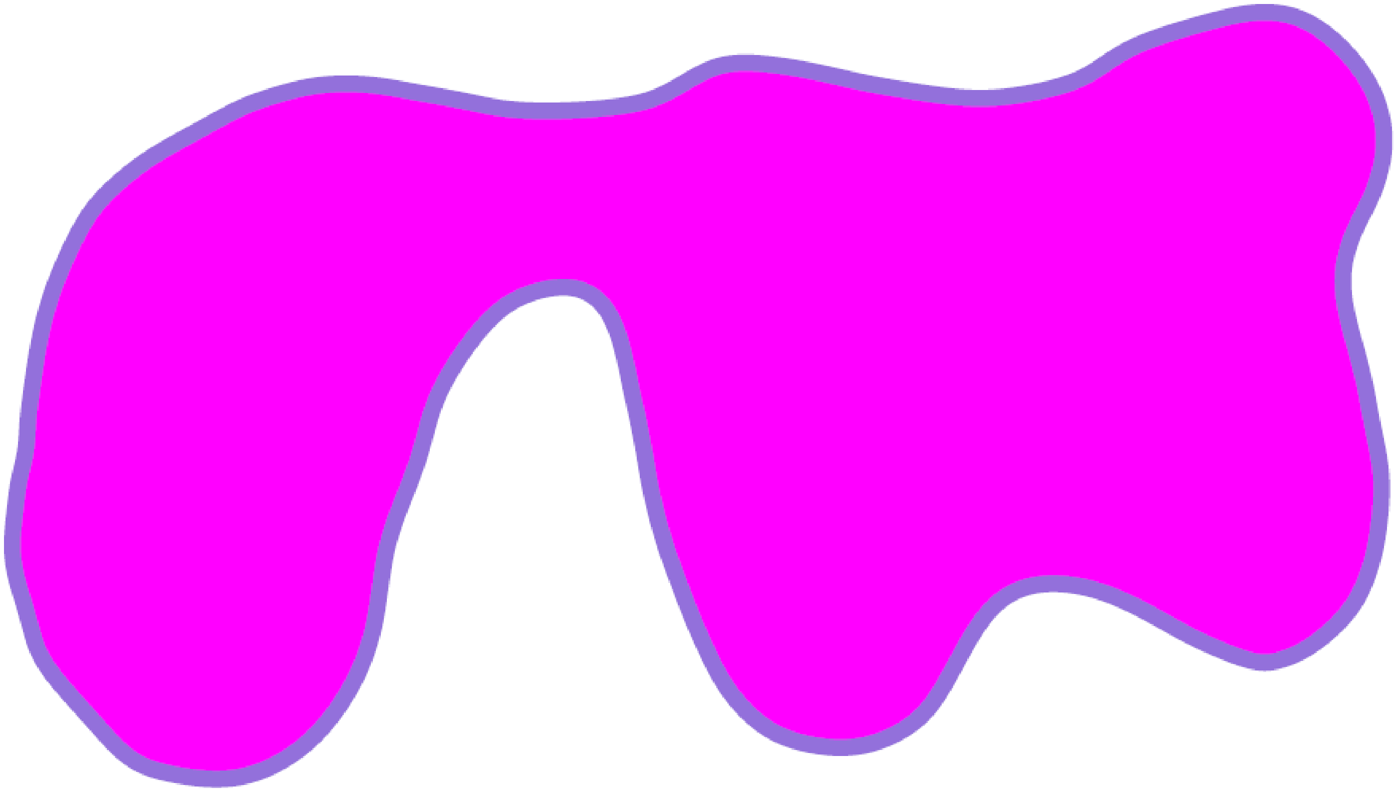

Medial Axis

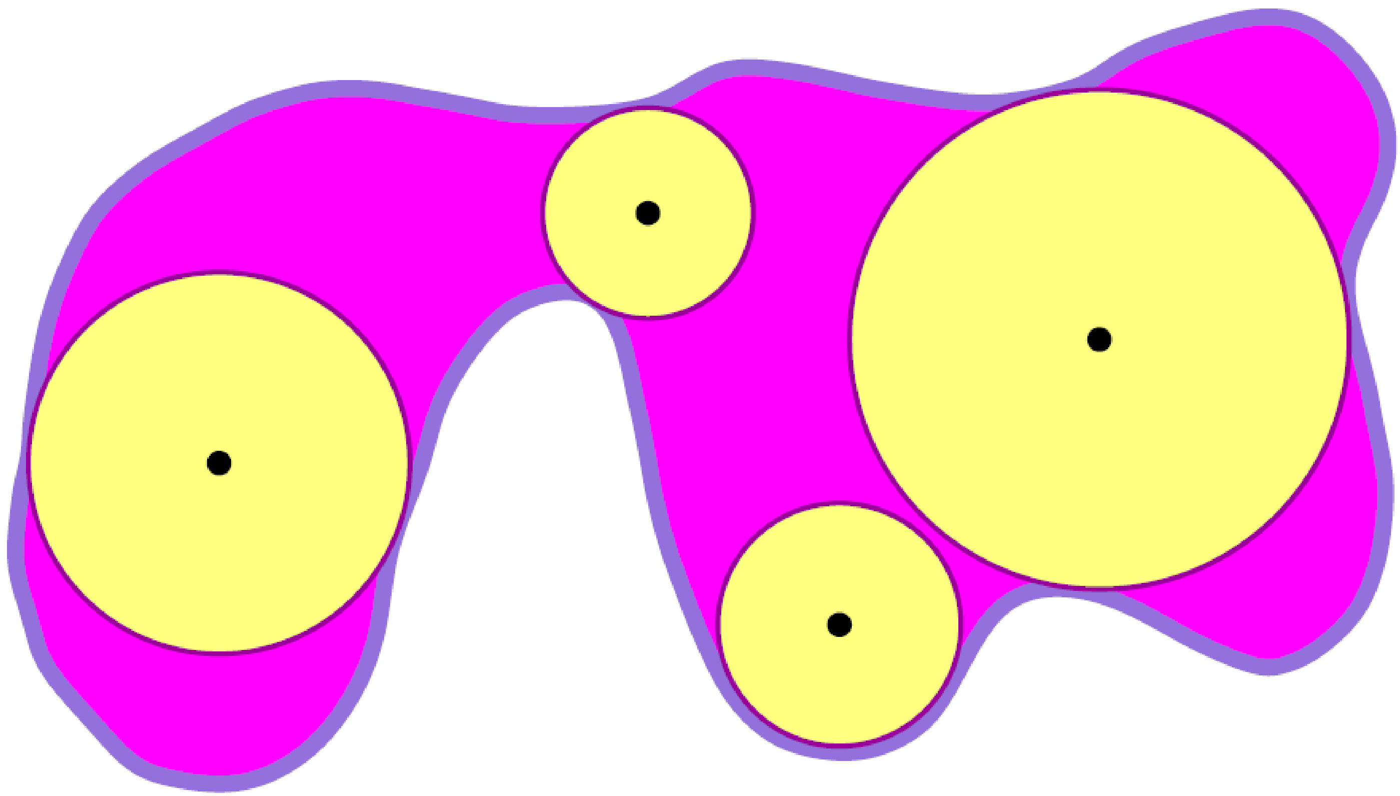

Medial Axis

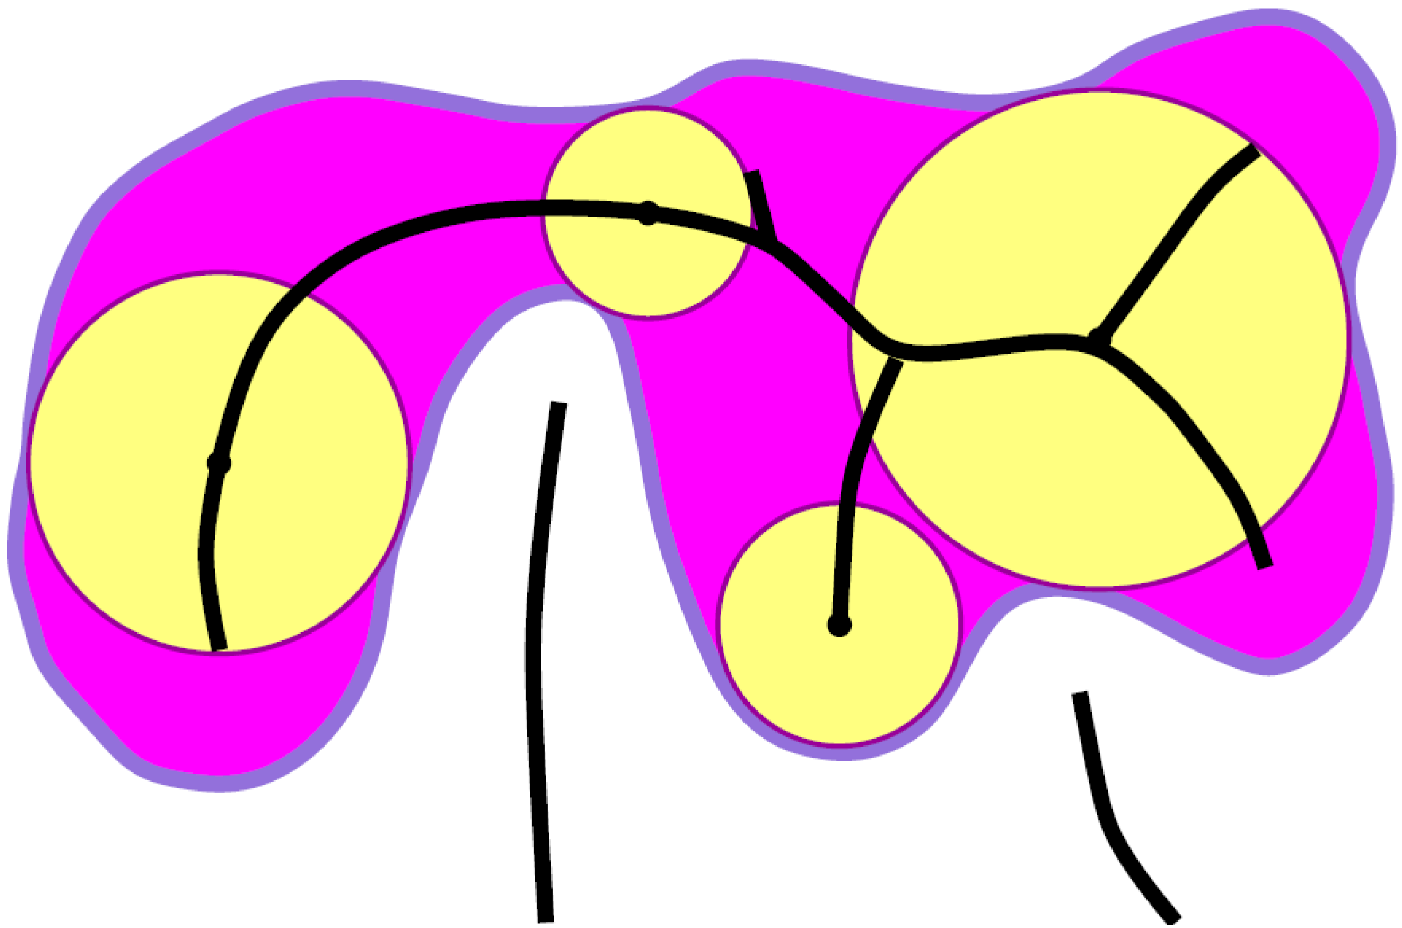

Medial Axis

<u>Observation*</u>:

For a reasonable point sample, the medial axis is well-sampled by the Voronoi vertices.

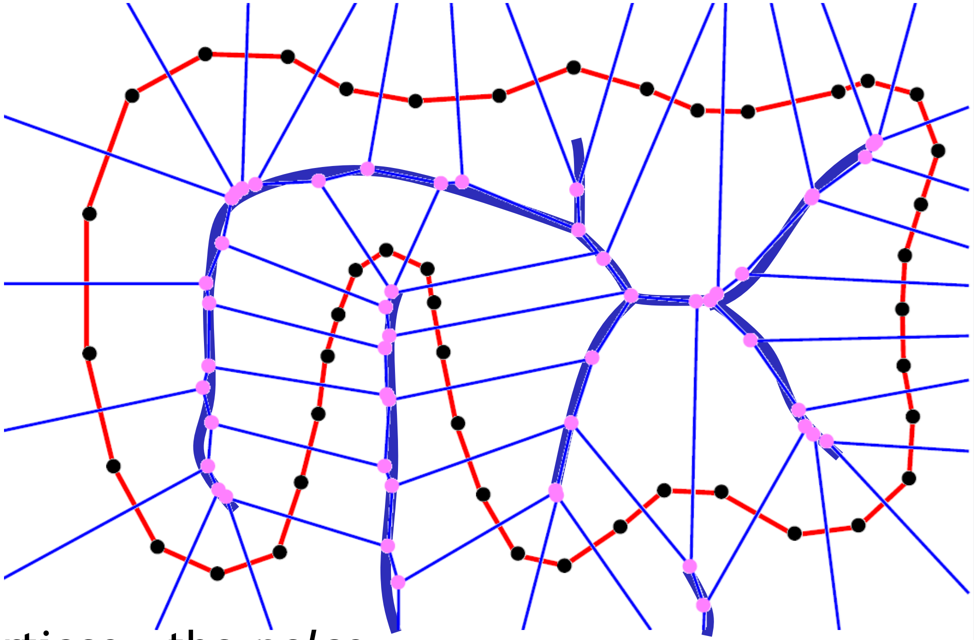

*In 3D, this is only true for a subset of the Voronoi vertices – the *poles*.

Voronoi & Medial Axis

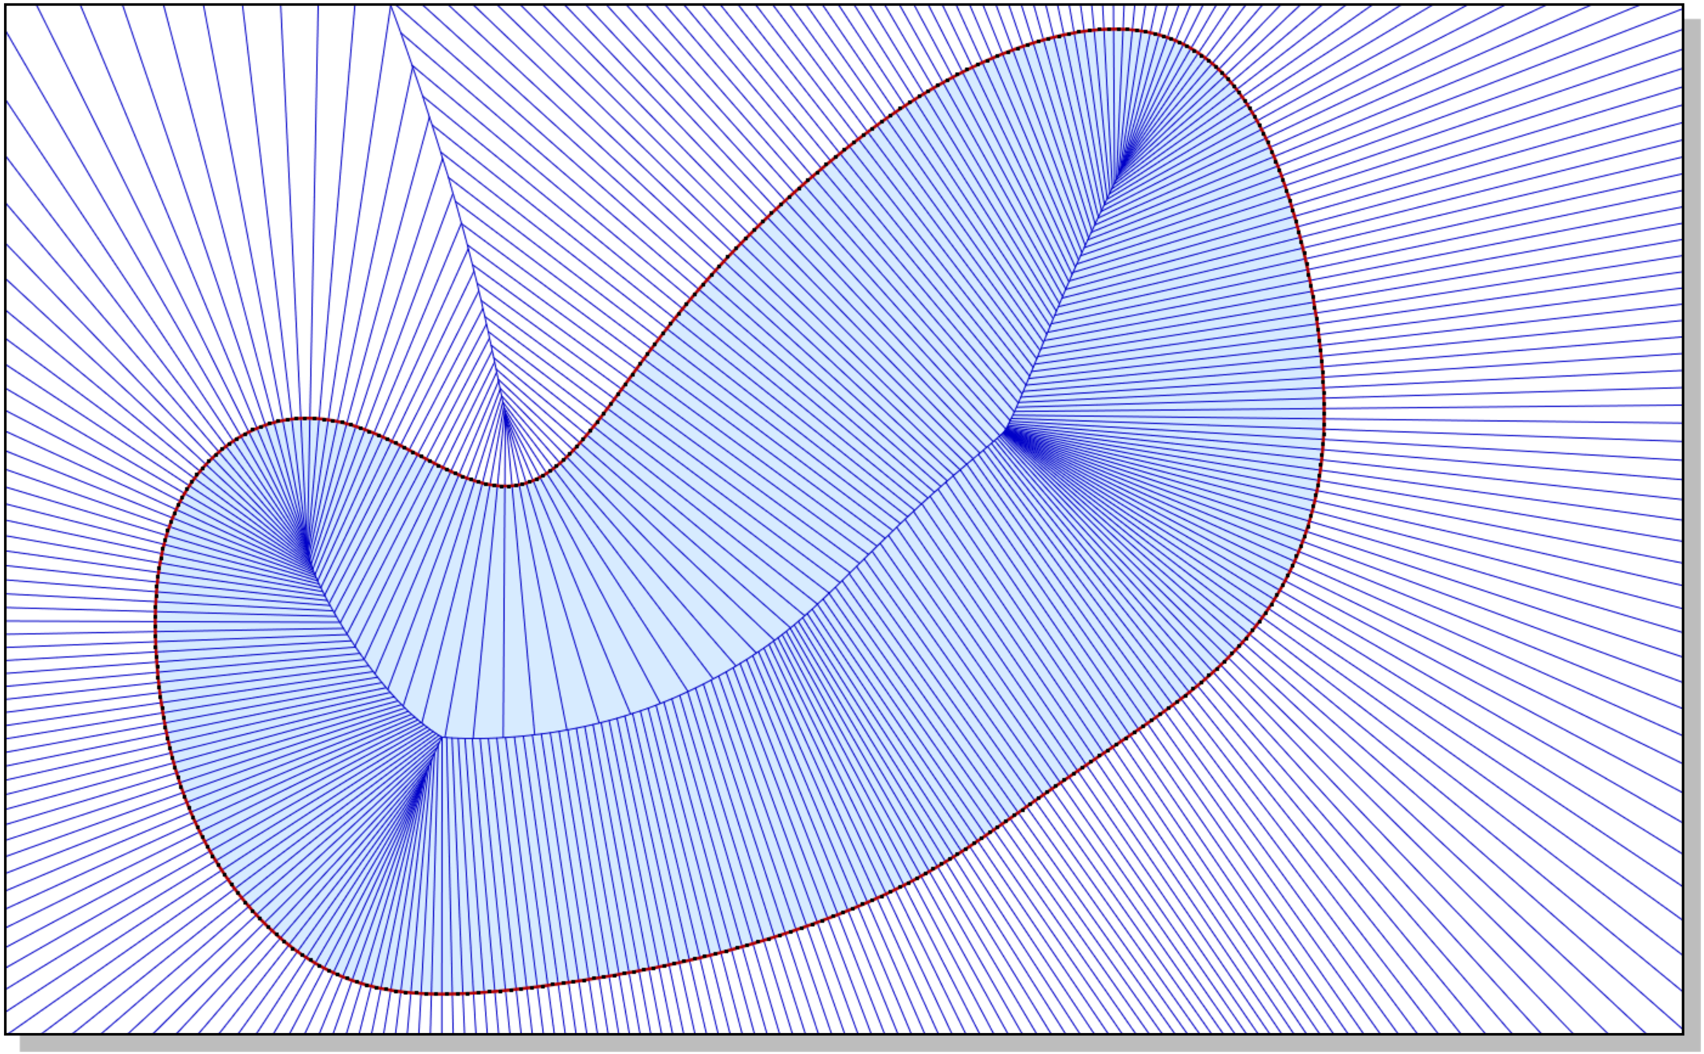

Local Feature Size

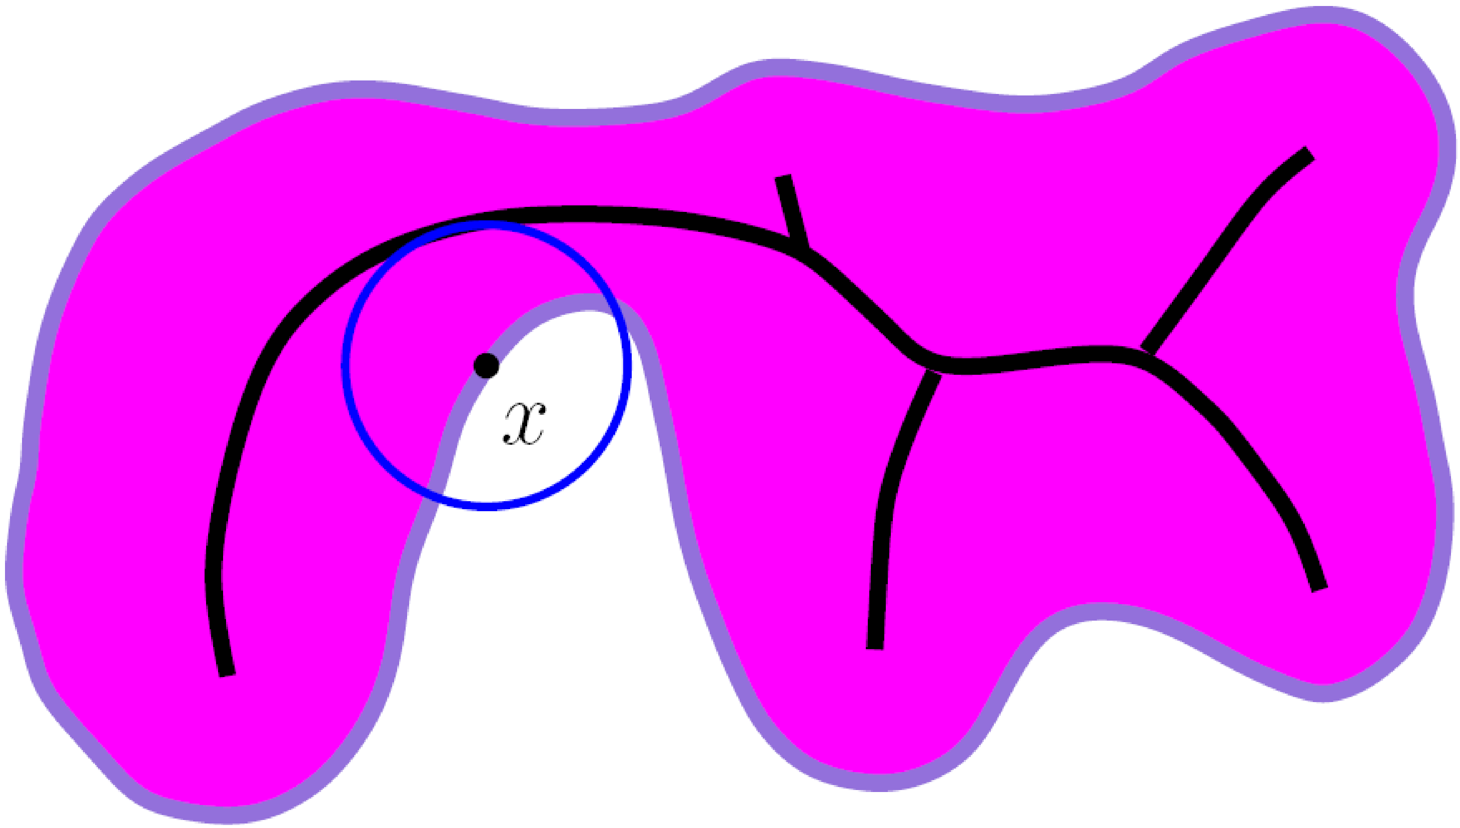

Epsilon-Sampling

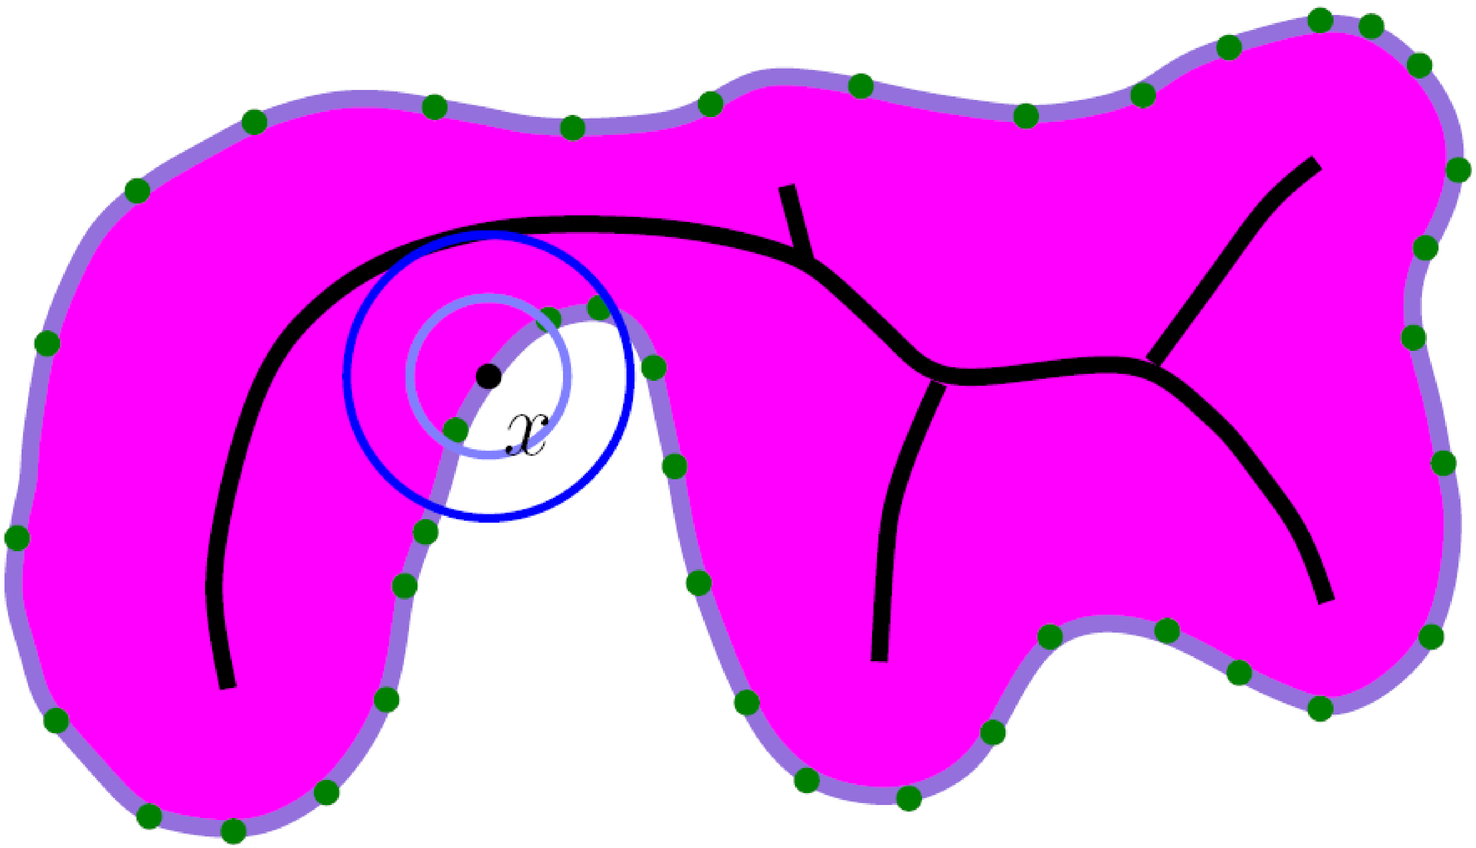

Crust [Amenta et al. 1998]

If we consider the Delaunay Triangulation of a point set, the shape boundary can be described as a subset of the Delaunay edges.

Q: How do we determine which edges to keep?

A: Two types of edges:

1. Those connecting adjacent points on the boundary
2. Those traversing the shape.

Discard those that traverse.

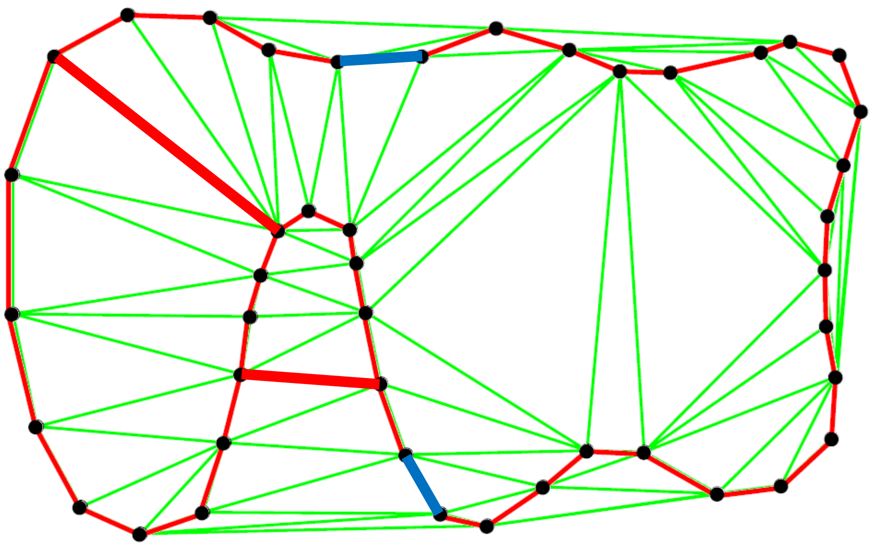

Crust [Amenta et al. 1998]

<u>Observation</u>:

Edges that traverse cross the medial axis.

Although we don't know the axis, we can sample it with the Voronoi vertices.

Edges that traverse must

be near the Voronoi vertices.

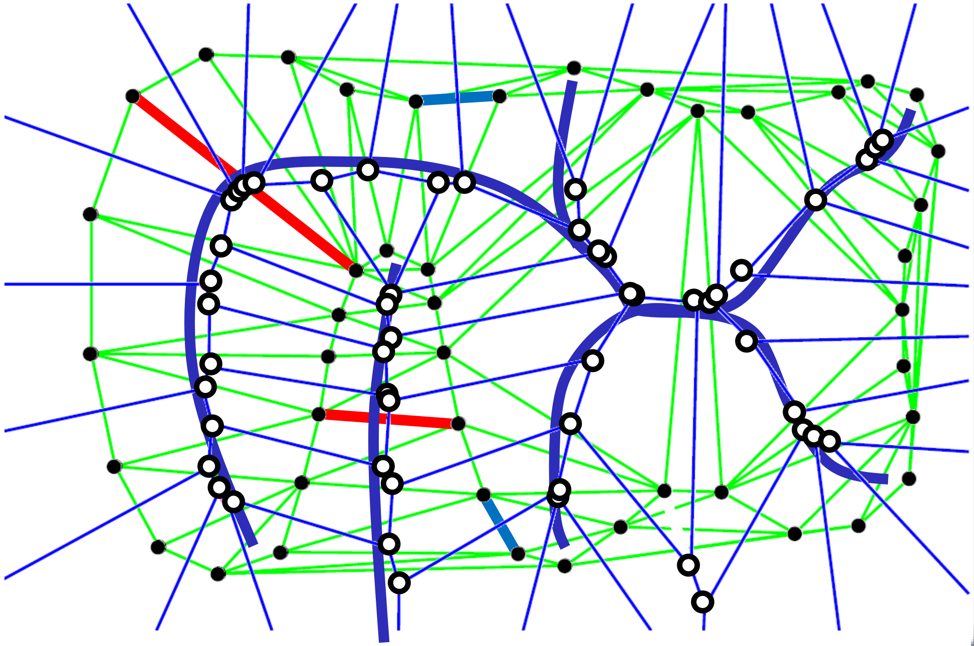

Crust [Amenta et al.]

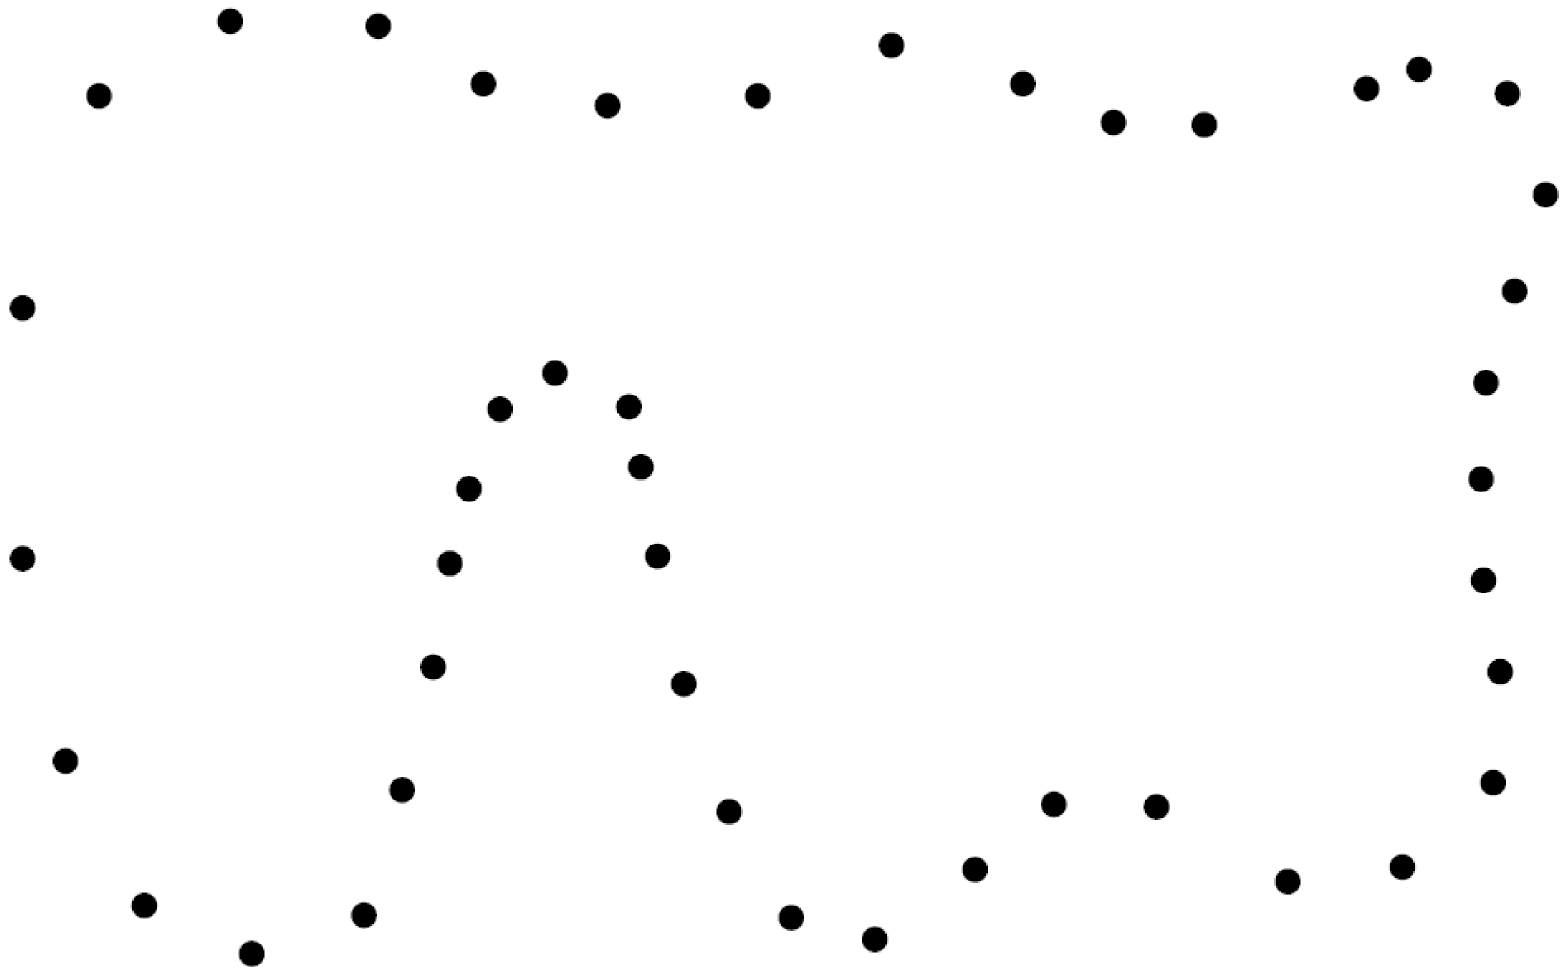

Delaunay Triangulation

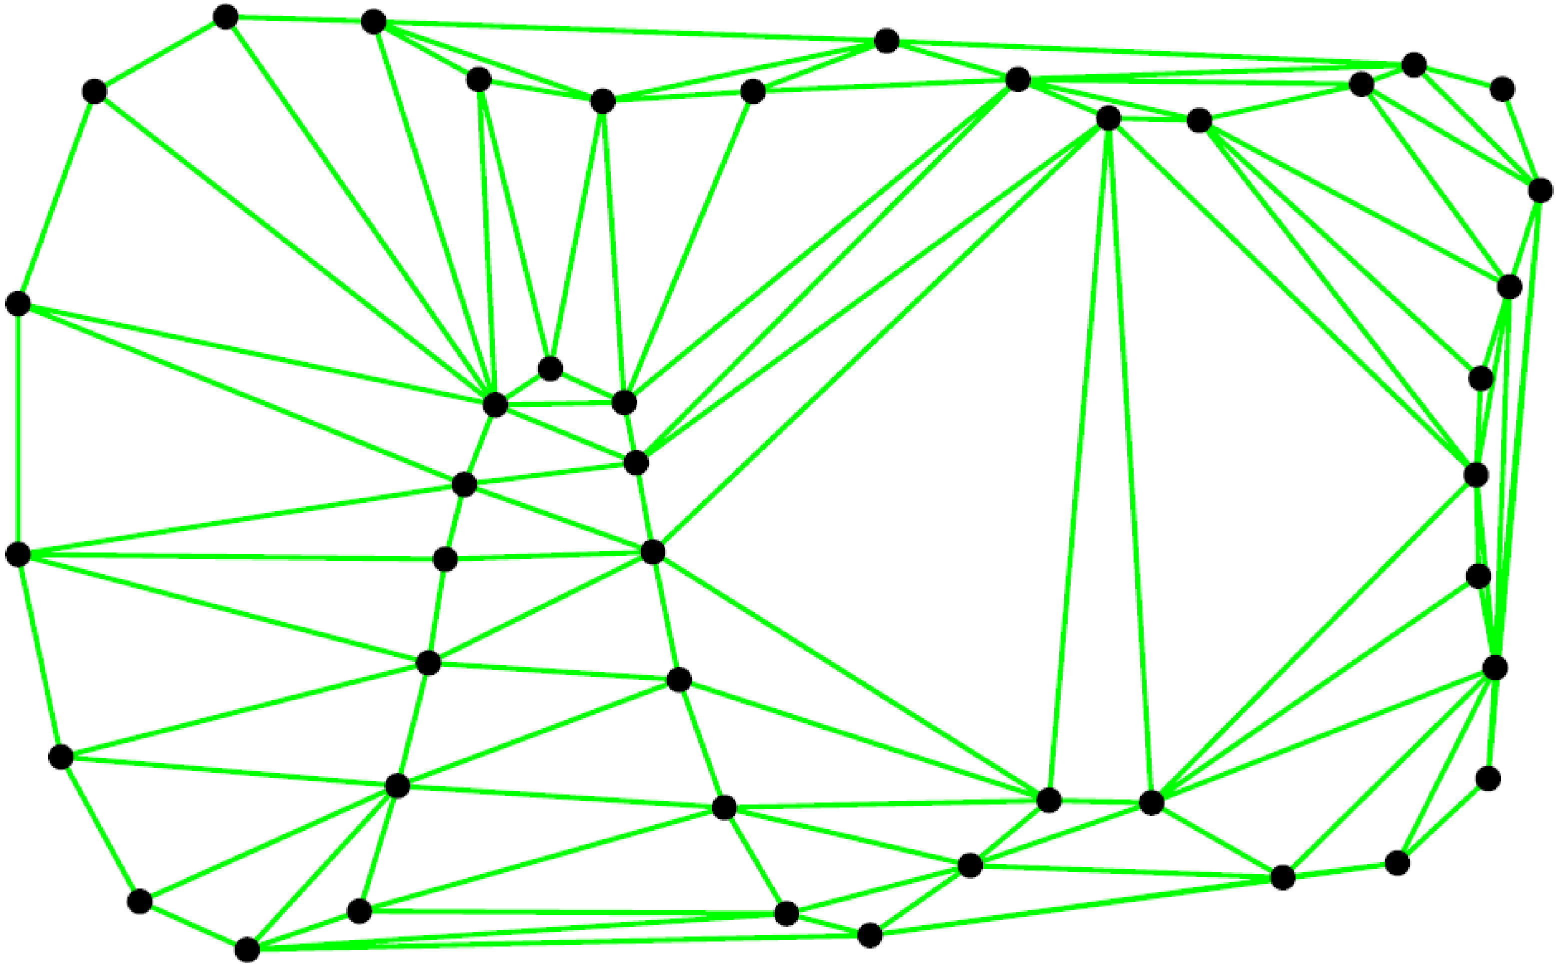

Delaunay Triangulation & Voronoi Diagram

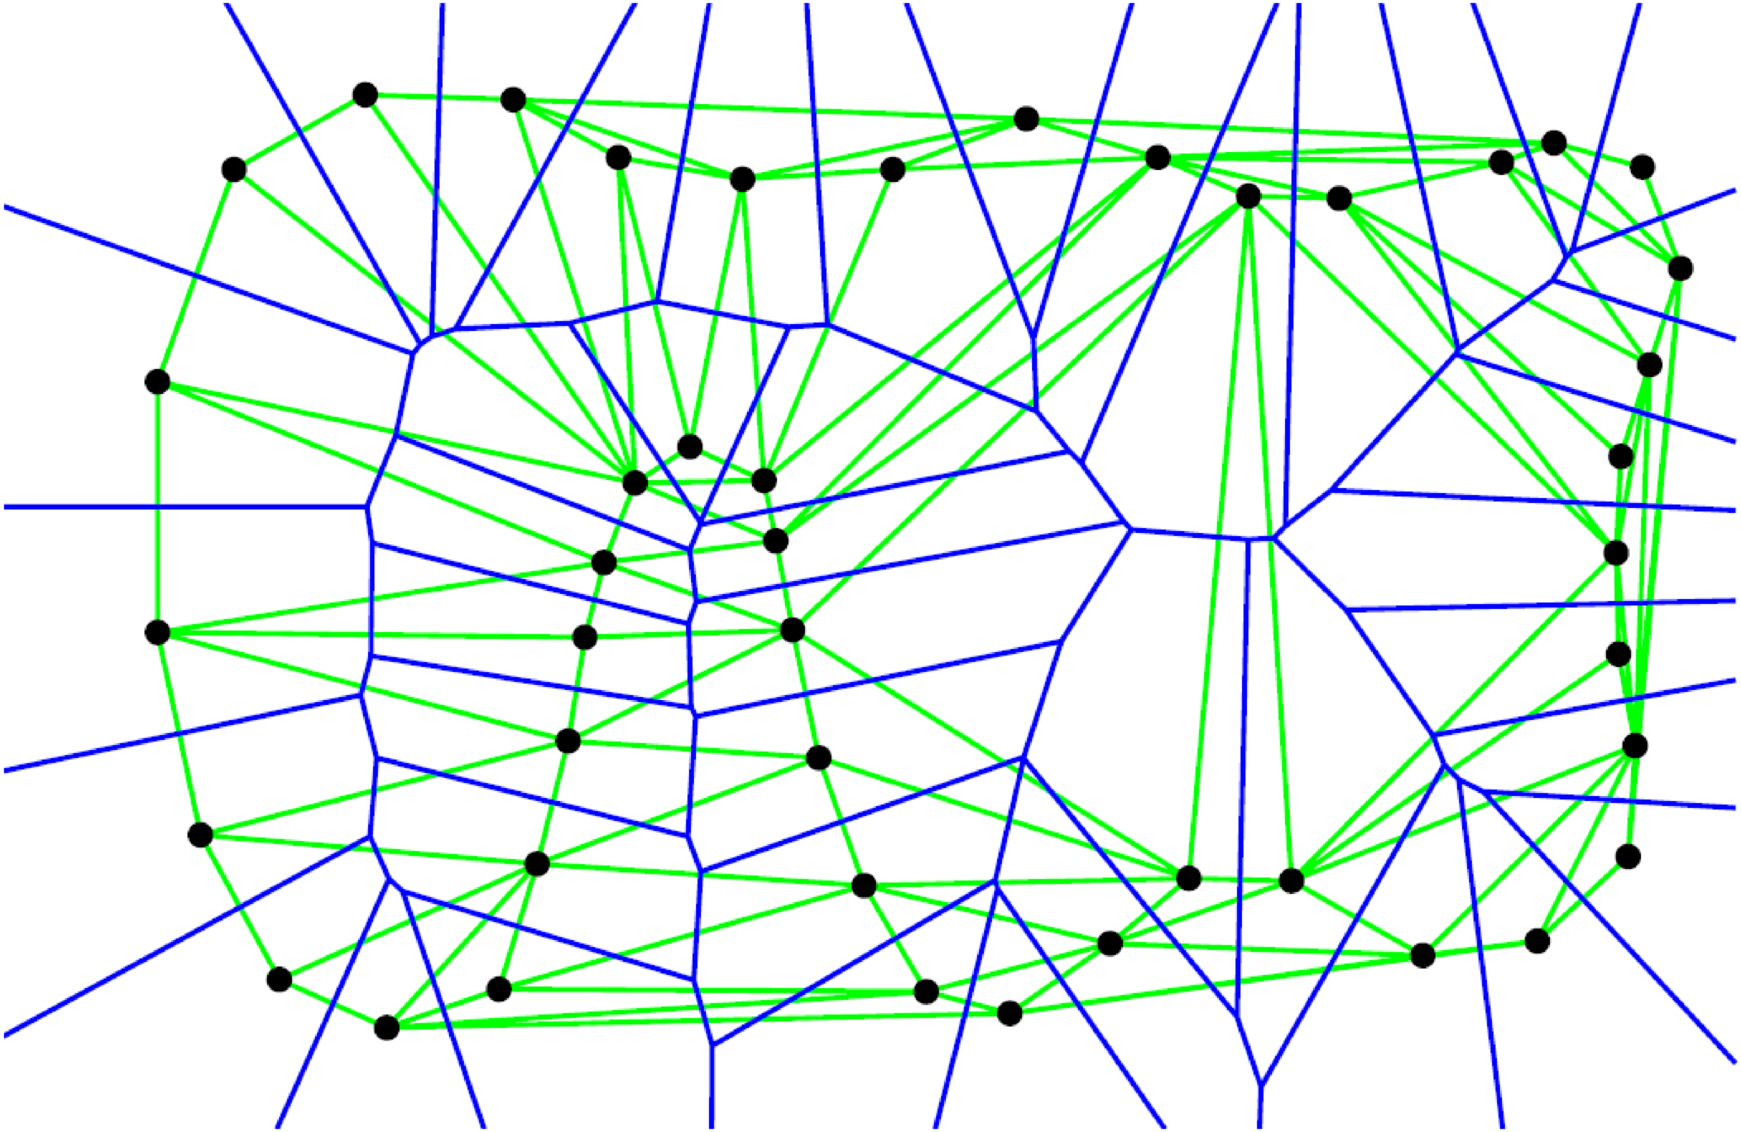

Voronoi Vertices

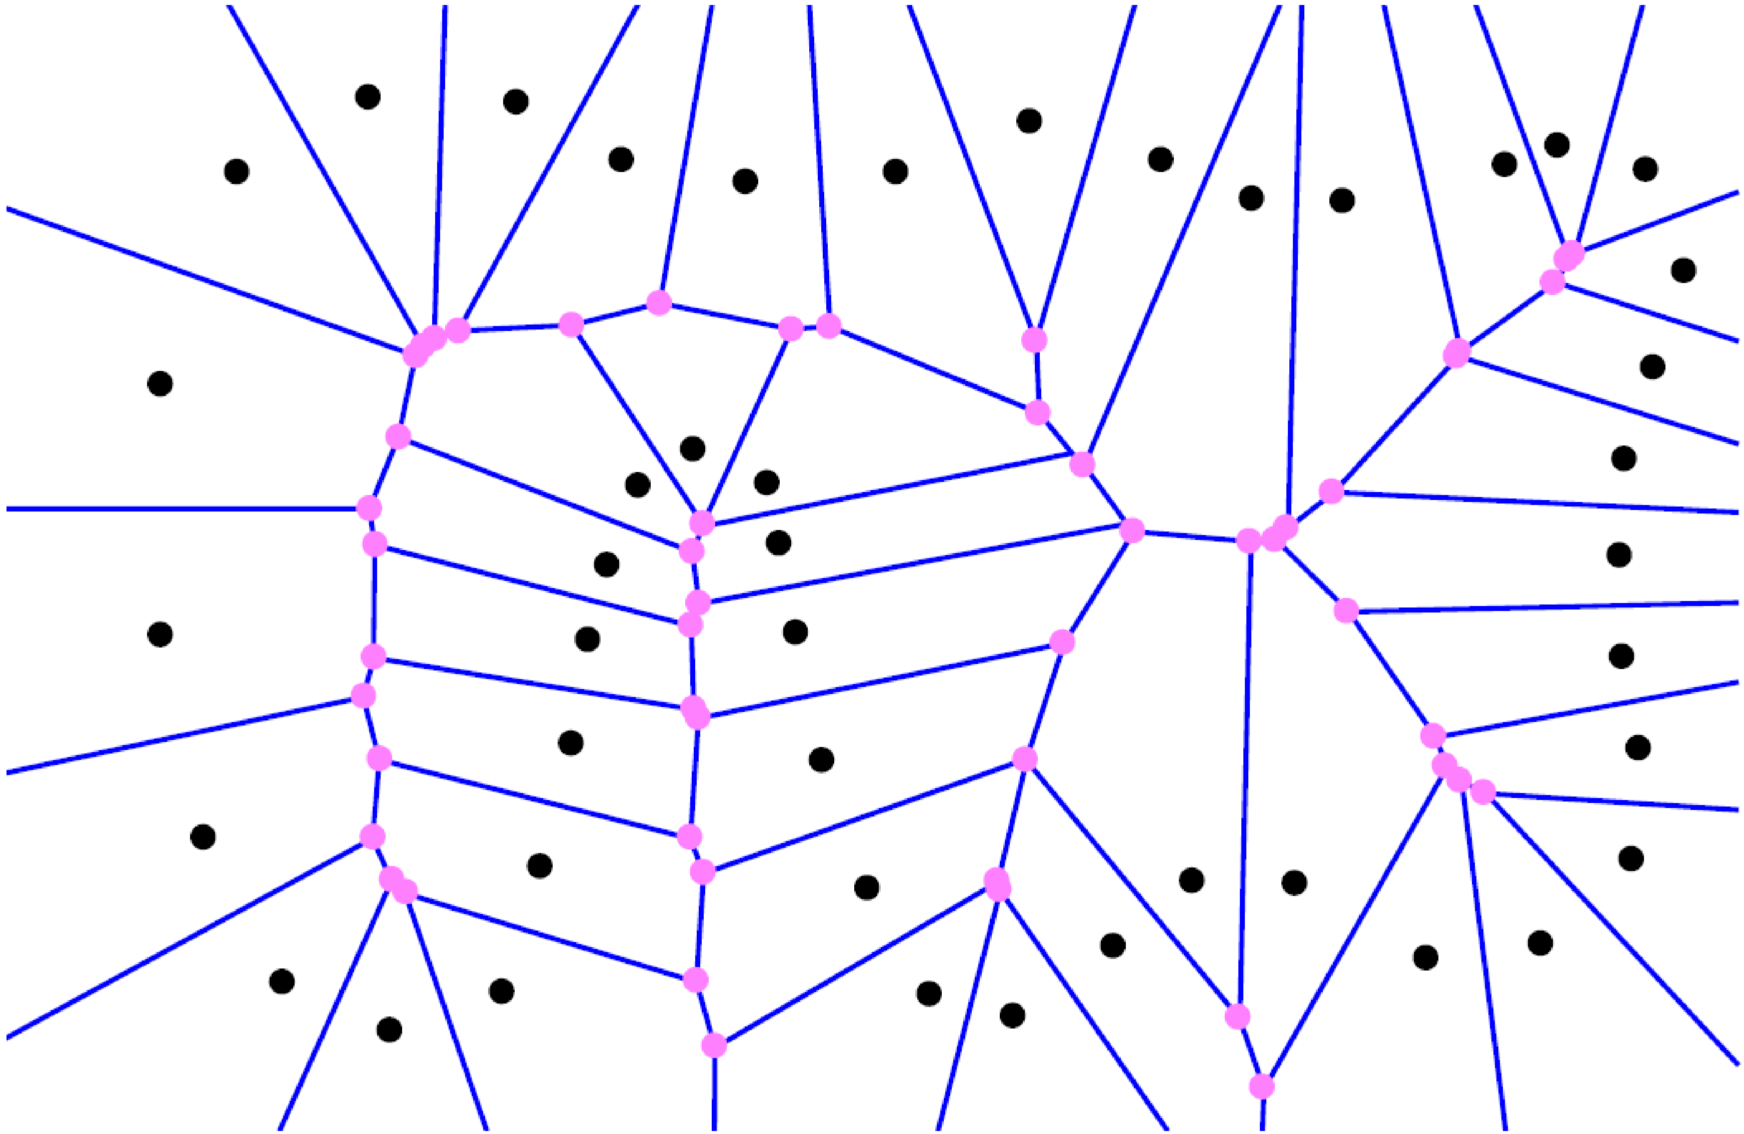

Refined Delaunay Triangulation

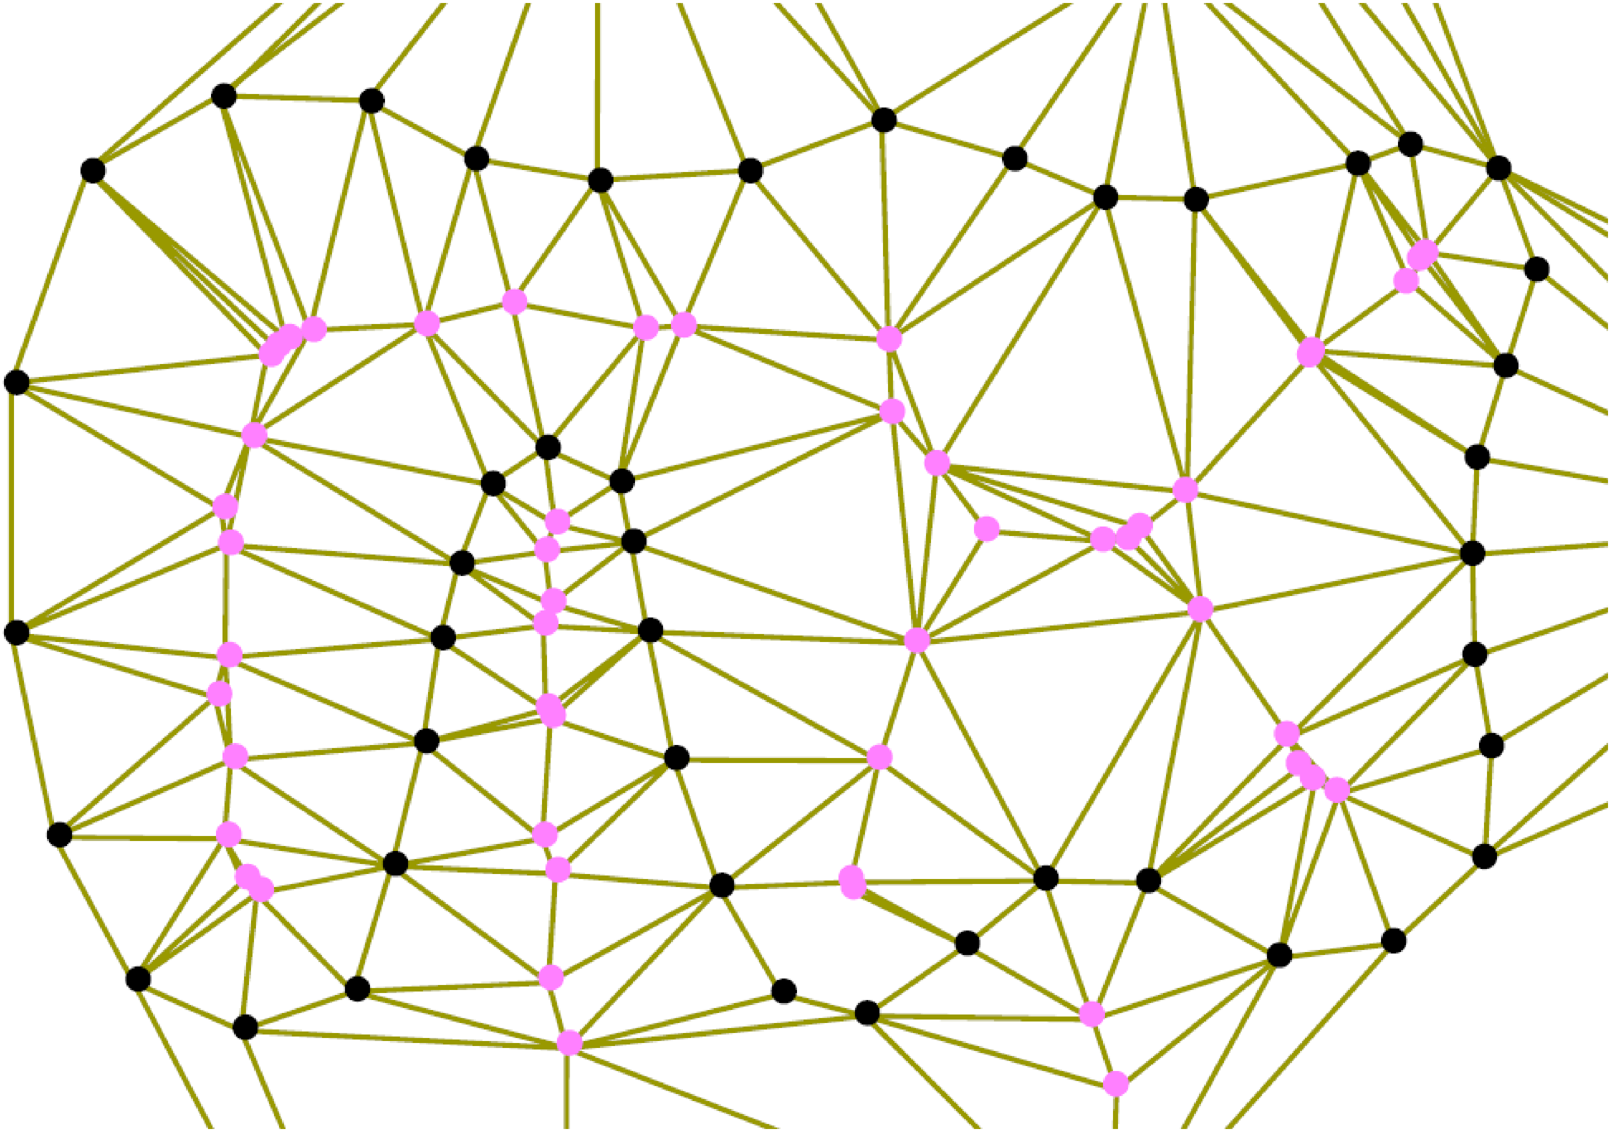

Crust

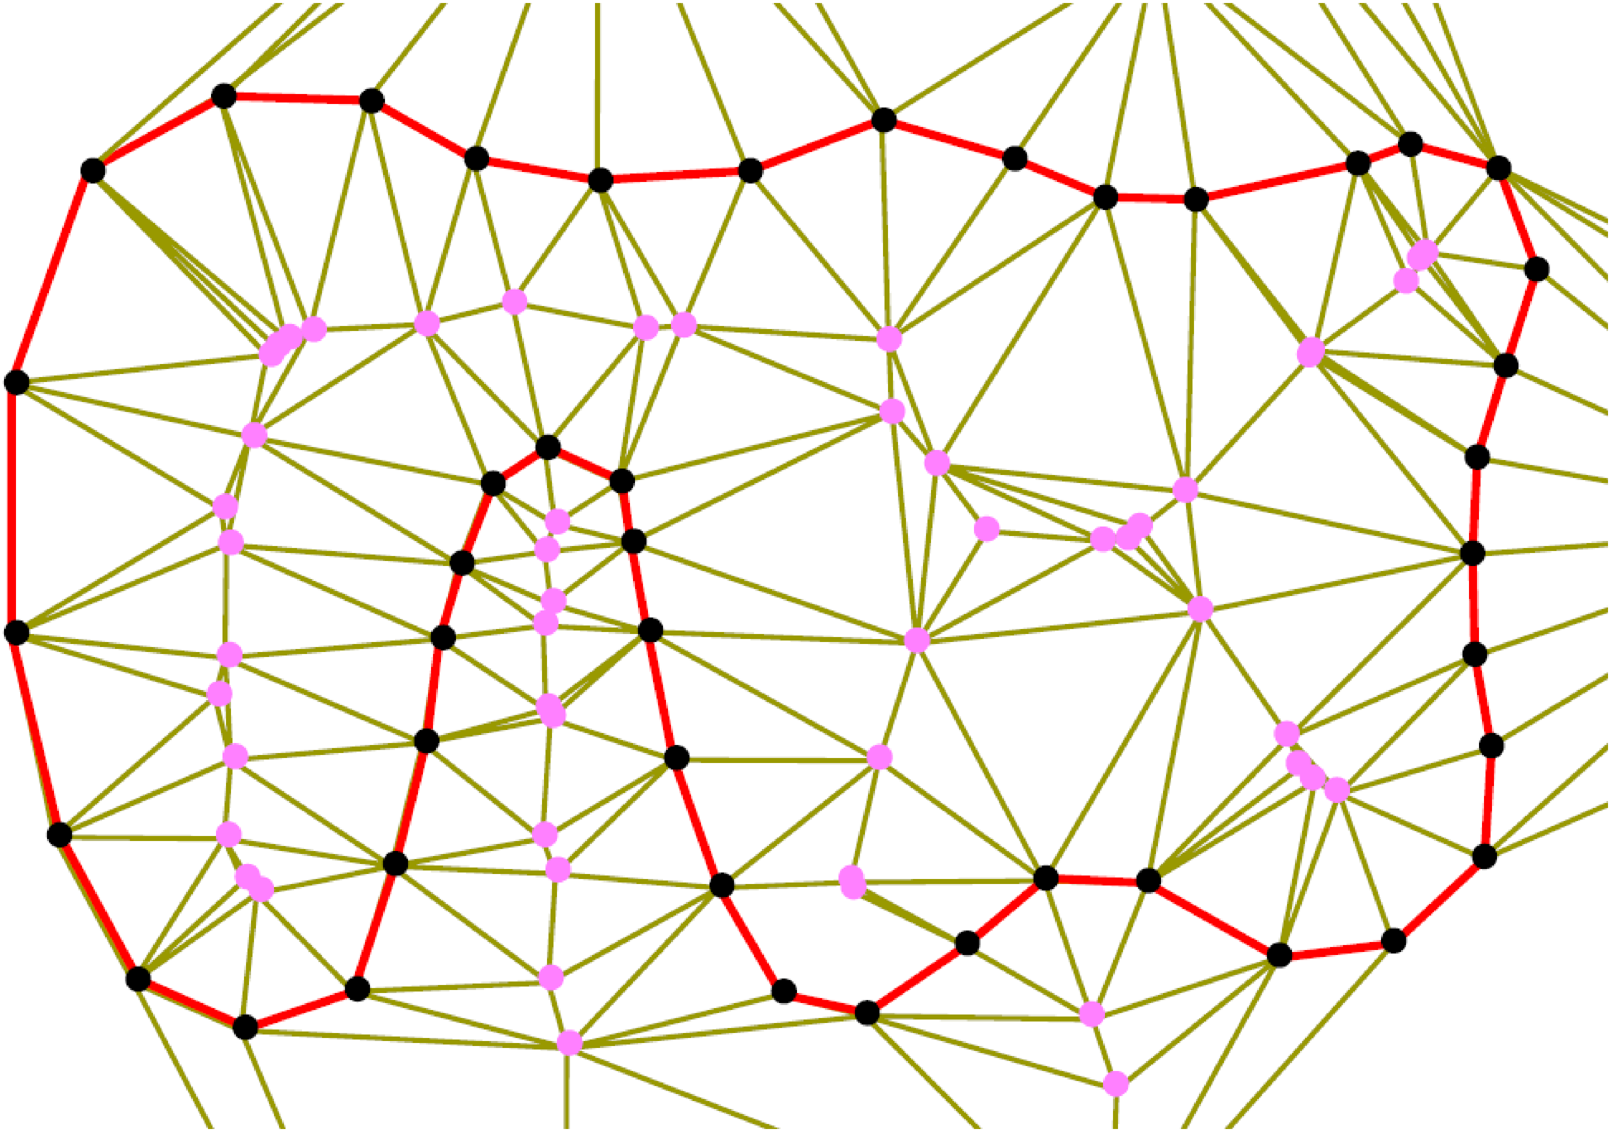

Crust

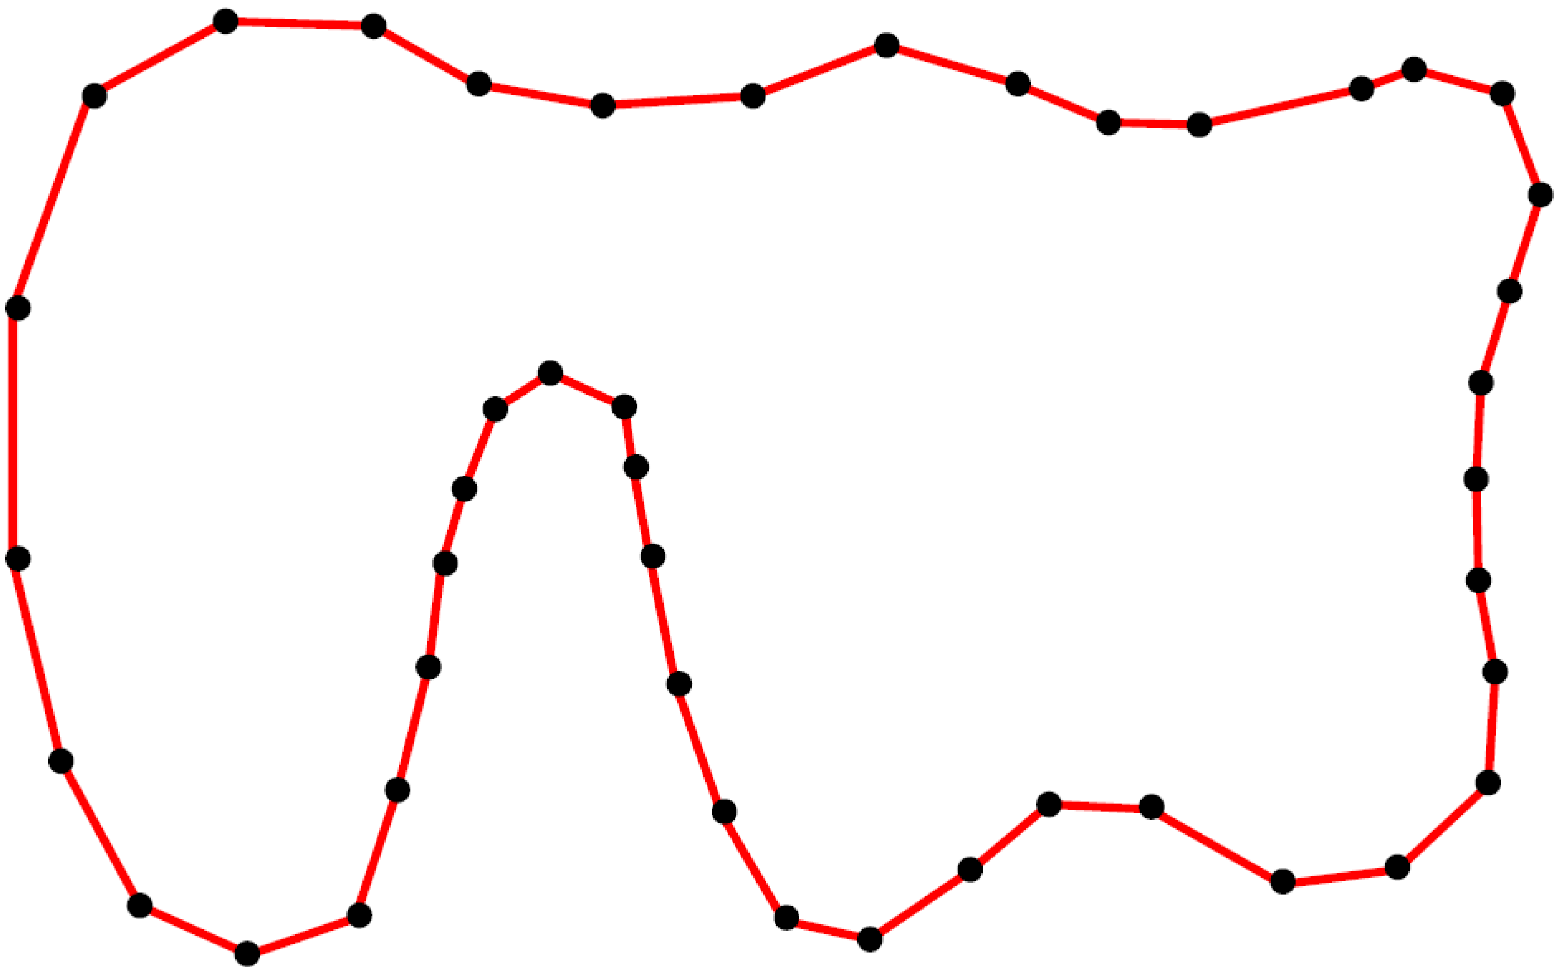

Crust (variant)

<u>Algorithm</u>:

1. Compute the Delaunay triangulation.
2. Compute the Voronoi vertices
3. Keep all edges for which there is a circle that contains the edge but no Voronoi vertices.

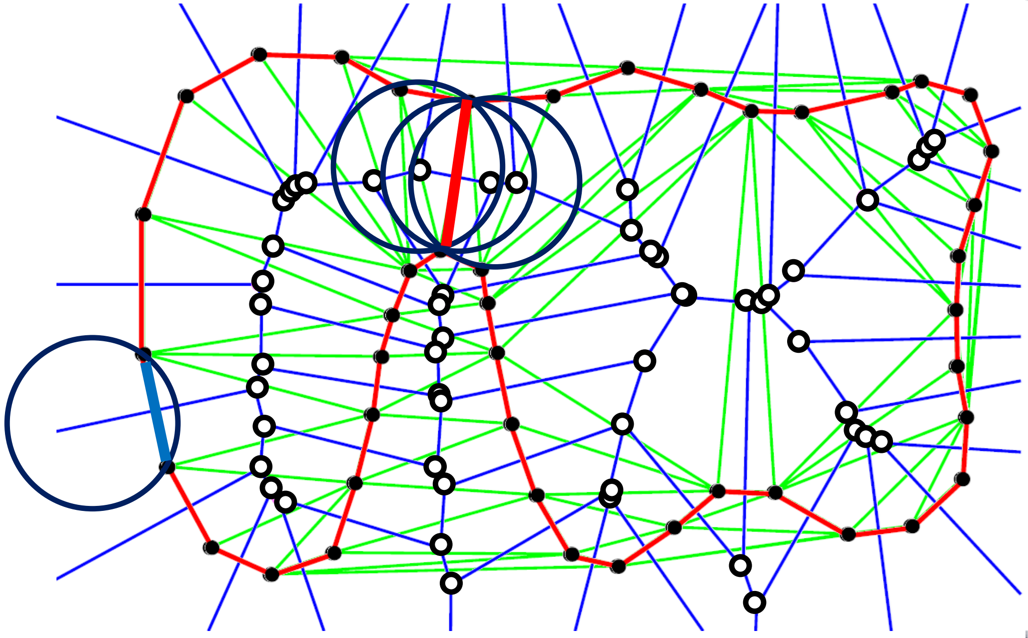

SPECTRAL « CRUST »

Space Partitioning

Given a set of points, construct the Delaunay triangulation.

If we label each triangle as inside/outside, then the surface of interest is the set of edges that lie between inside and outside triangles.

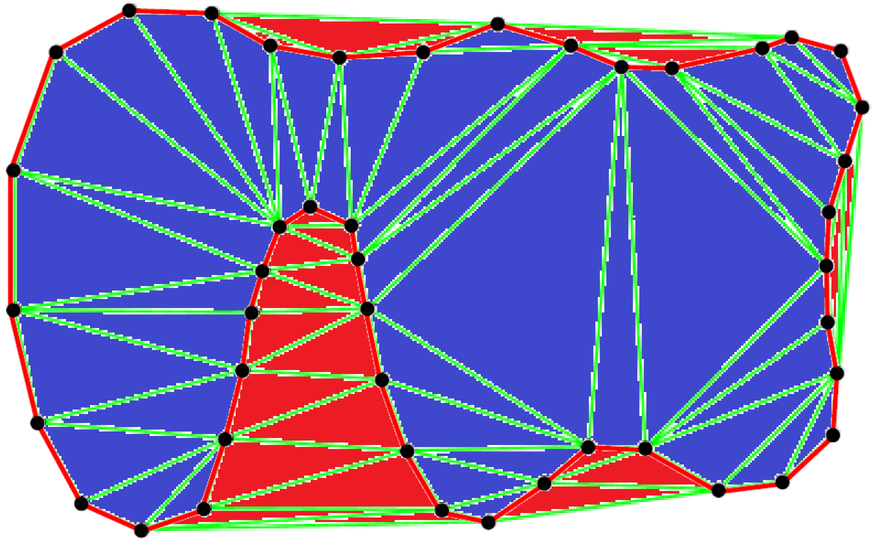

Space Partitioning

Q: How to assign labels?

A: Spectral Partitioning

Assign a weight to each edge indicating if the two triangles are likely to have the same label.

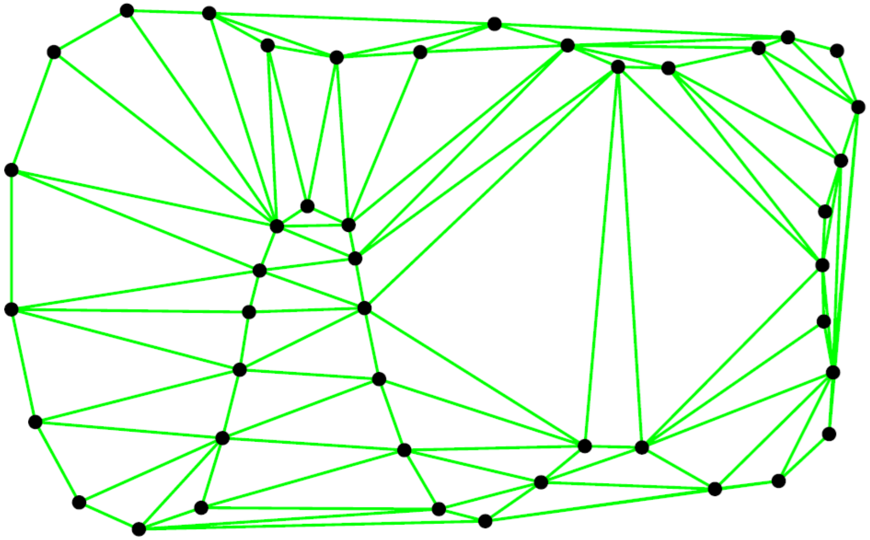

[Kolluri et al., 2004]

Space Partitioning

Assigning edge weights

Q: When are triangles on opposite sides of an edge likely to have the same label?

A: If the triangles are on the same side, their circumscribing circles intersect deeply.

Use the angle of intersection to set the weight.

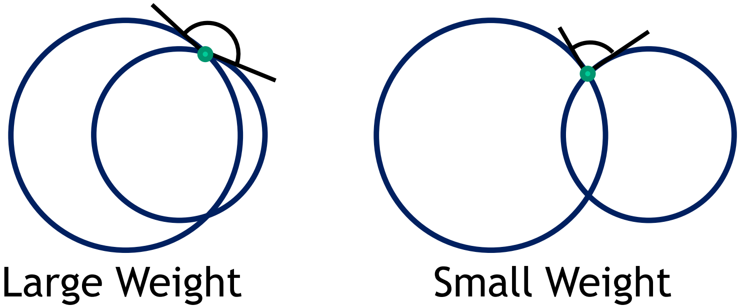

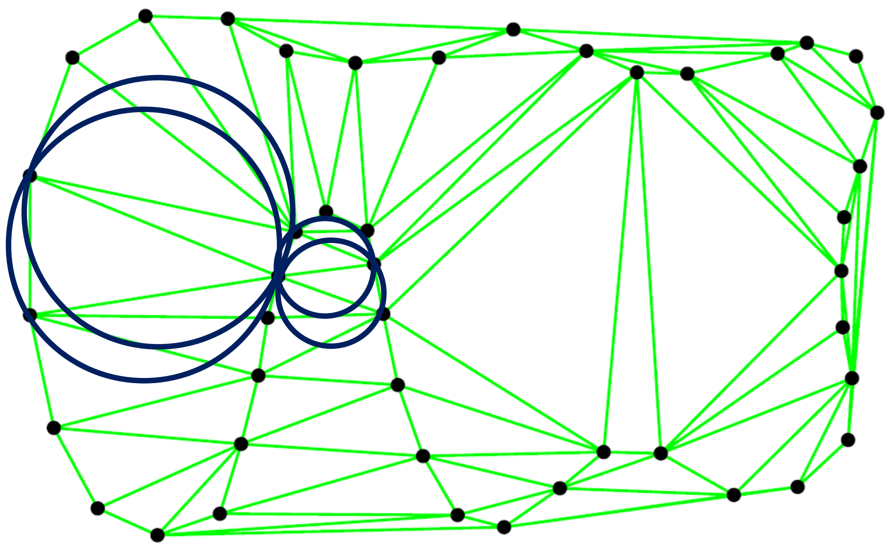

Crust

Several Delaunay algorithms <u>provably correct</u>

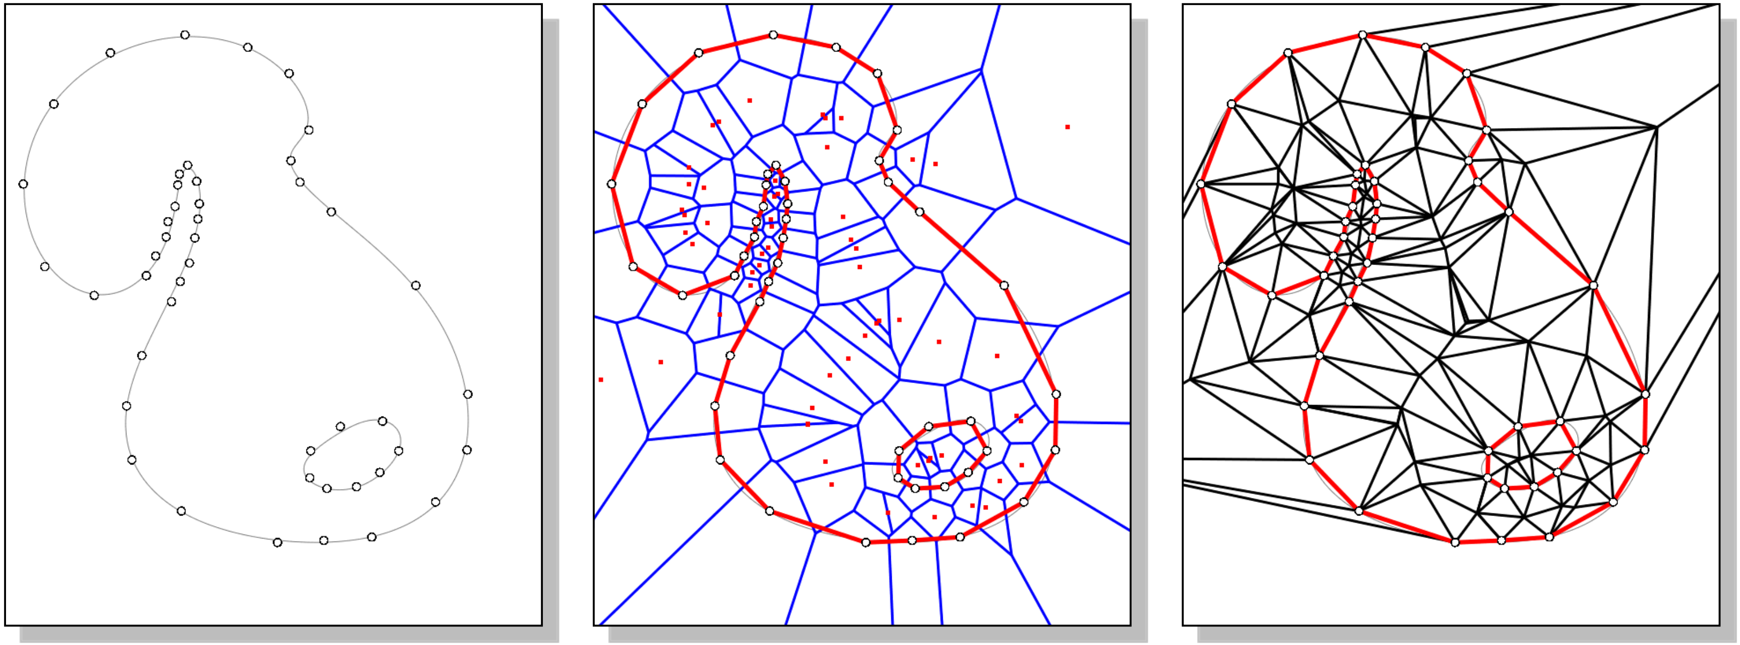

Delaunay-based

Several Delaunay algorithms are **provably correct**… in the absence of noise and undersampling.

perfect data ?

Noise & Undersampling

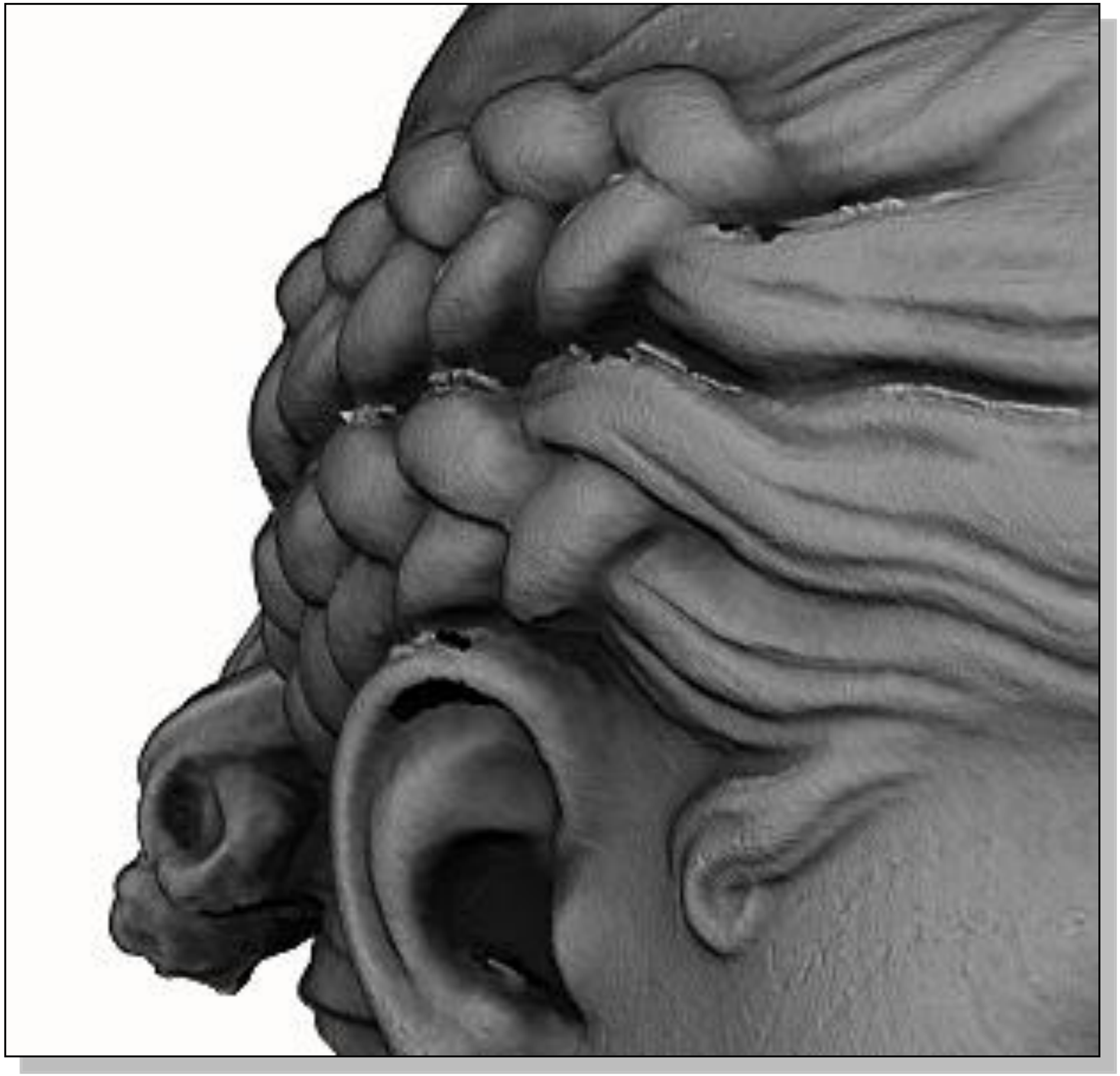

Delaunay-based

Several Delaunay algorithms are **provably correct**… in the absence of noise and undersampling.

Motivates reconstruction by fitting **approximating** implicit surfaces

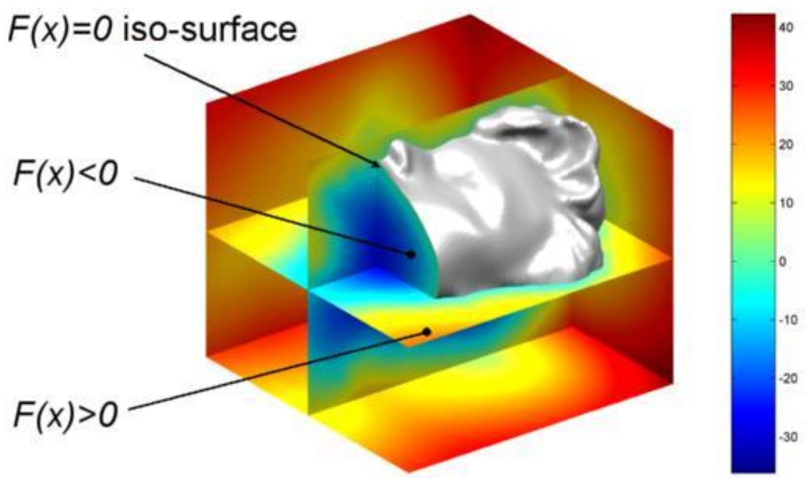

VARIATIONAL FORMULATIONS

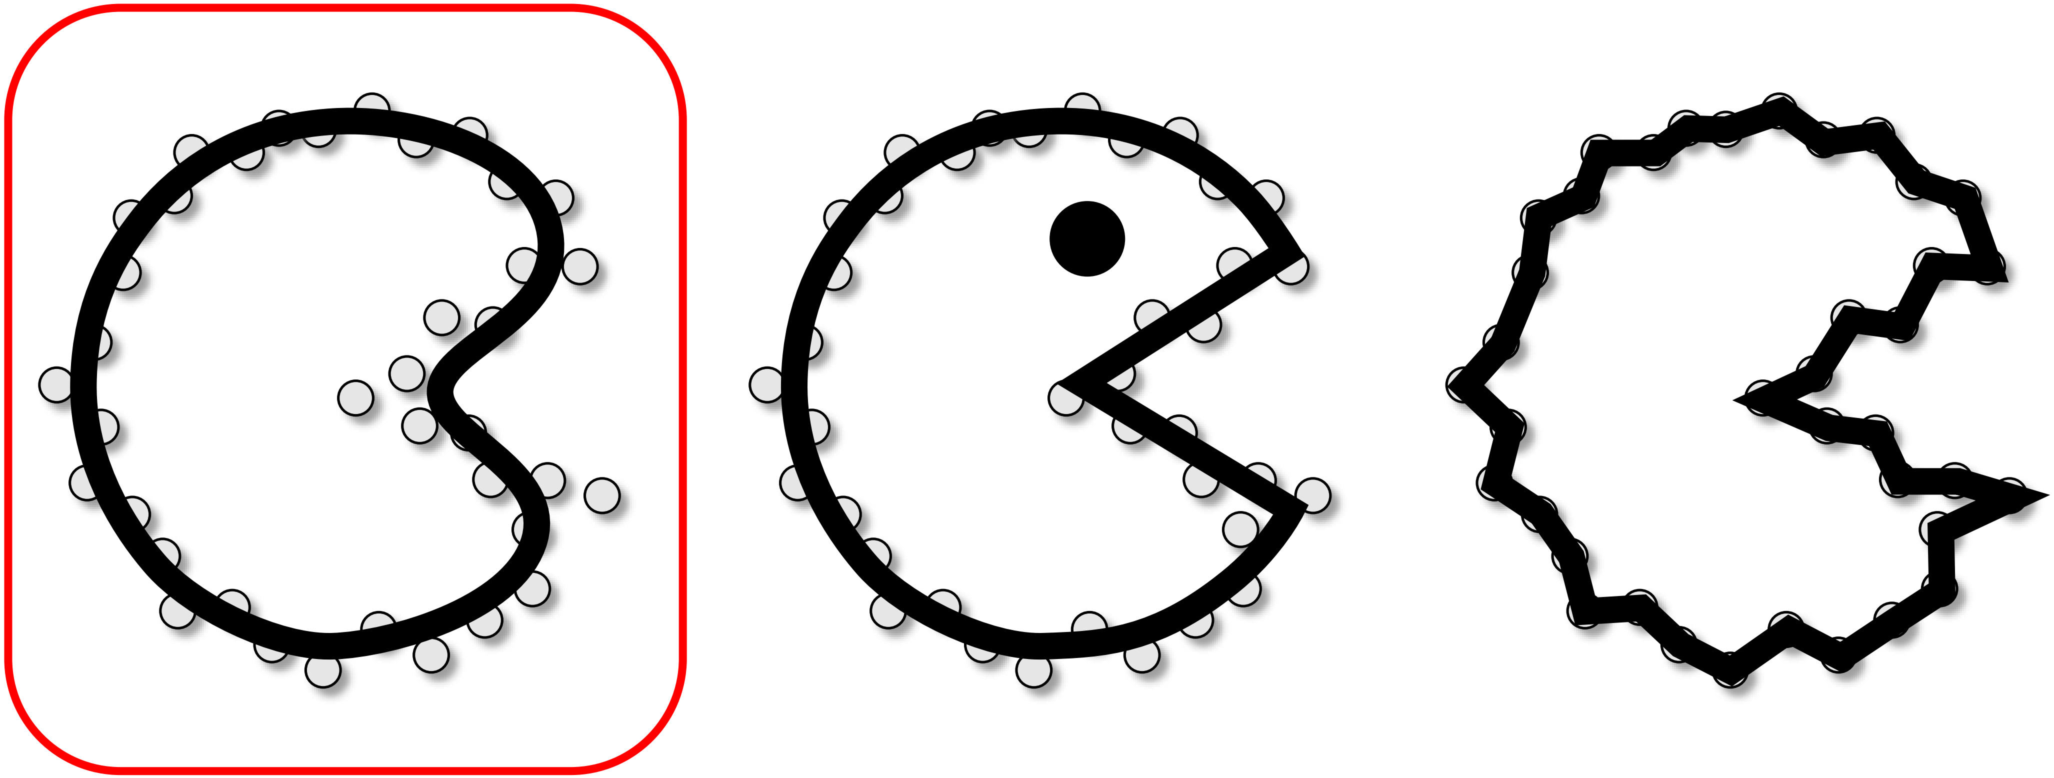

Smooth

Piecewise Smooth

“Simple”

Poisson Surface Reconstruction

[**Kazhdan et al.** SGP'06]

Indicator Function

Construct indicator function from point samples

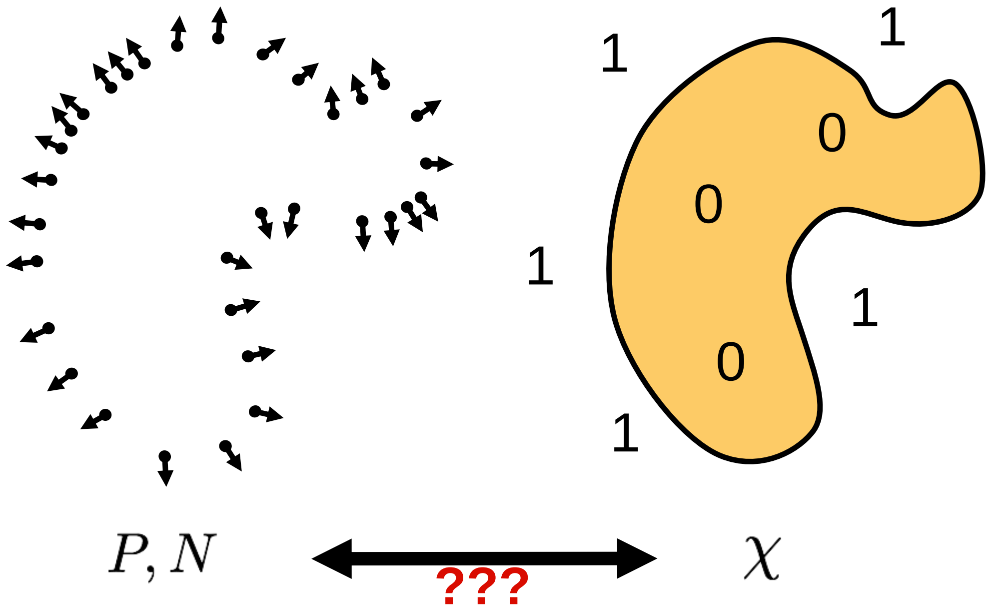

Indicator Function

Construct indicator function from point samples

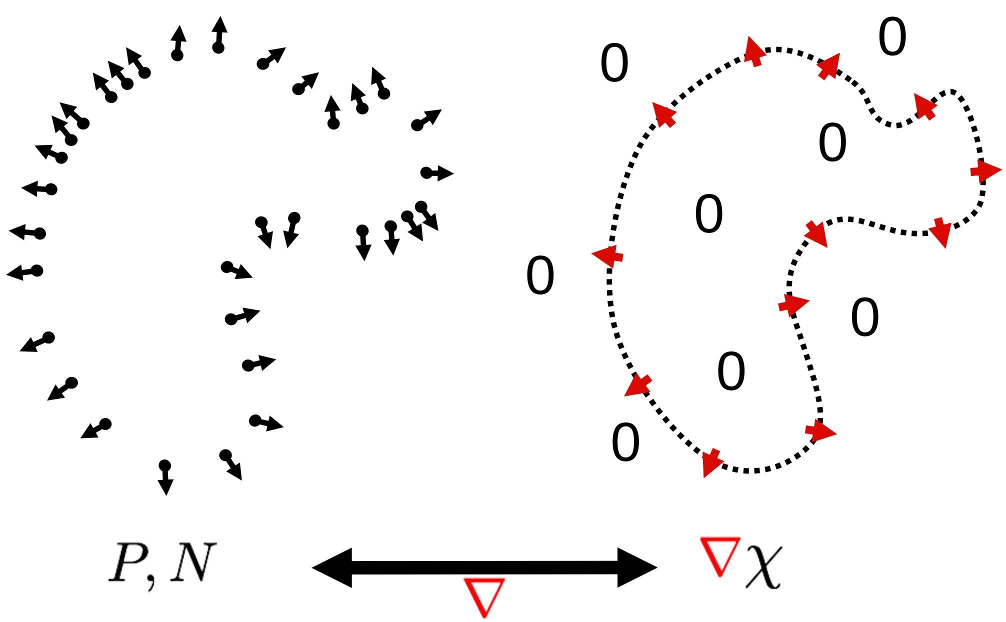

splatted normals

$$\min_{\chi} \int \|\nabla \chi(\mathbf{x}) - \boxed{\mathcal{N}(\mathbf{x})}\|_2^2 \, dx$$

variational calculus

$$\boxed{\Delta \chi} = \nabla \cdot \mathcal{N}$$

sparse linear system

2D Poisson Reconstruction

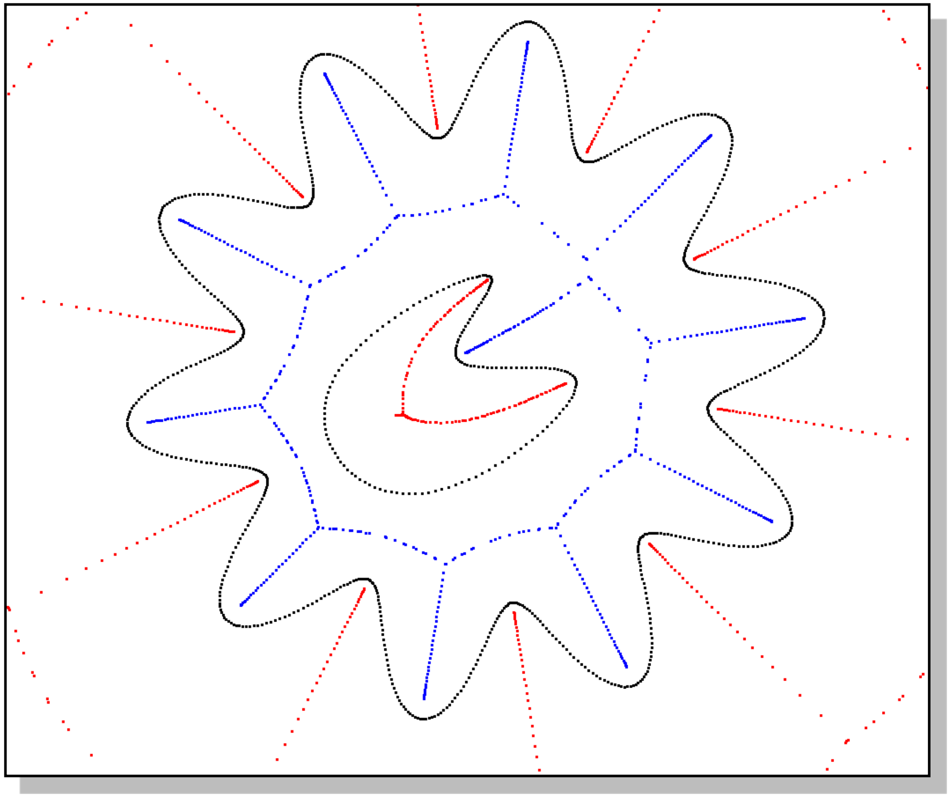

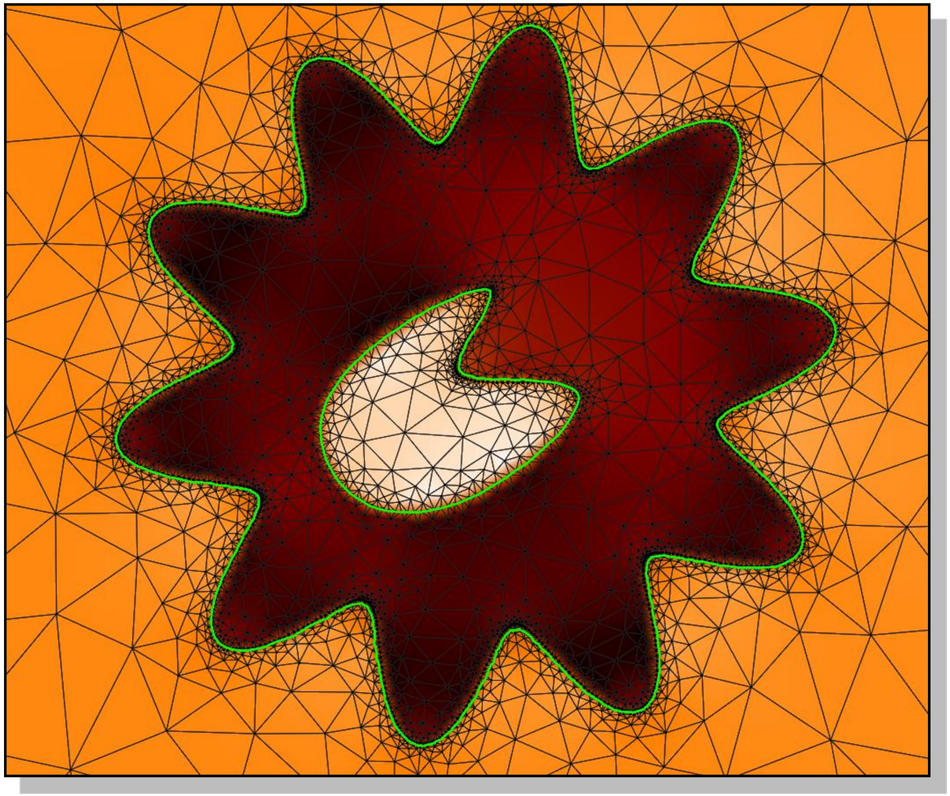

Poisson Reconstruction

Requires oriented normals, as many other implicit approaches.

Poisson Reconstruction

Requires oriented normals, as many other implicit approaches.

Normal estimation

Normal orientation

$\Big\}$ ill-posed problems

Poisson Reconstruction

Can we deal with unoriented normals?

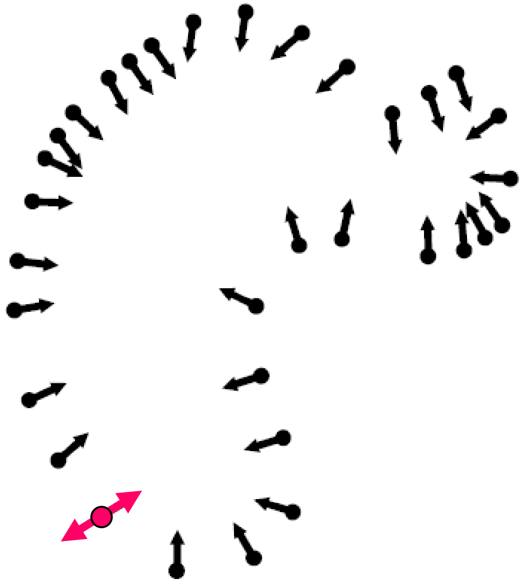

Spectral Surface Reconstruction

[**A., Cohen-Steiner, Tong & Desbrun**. SGP'07]

Unoriented Normals?

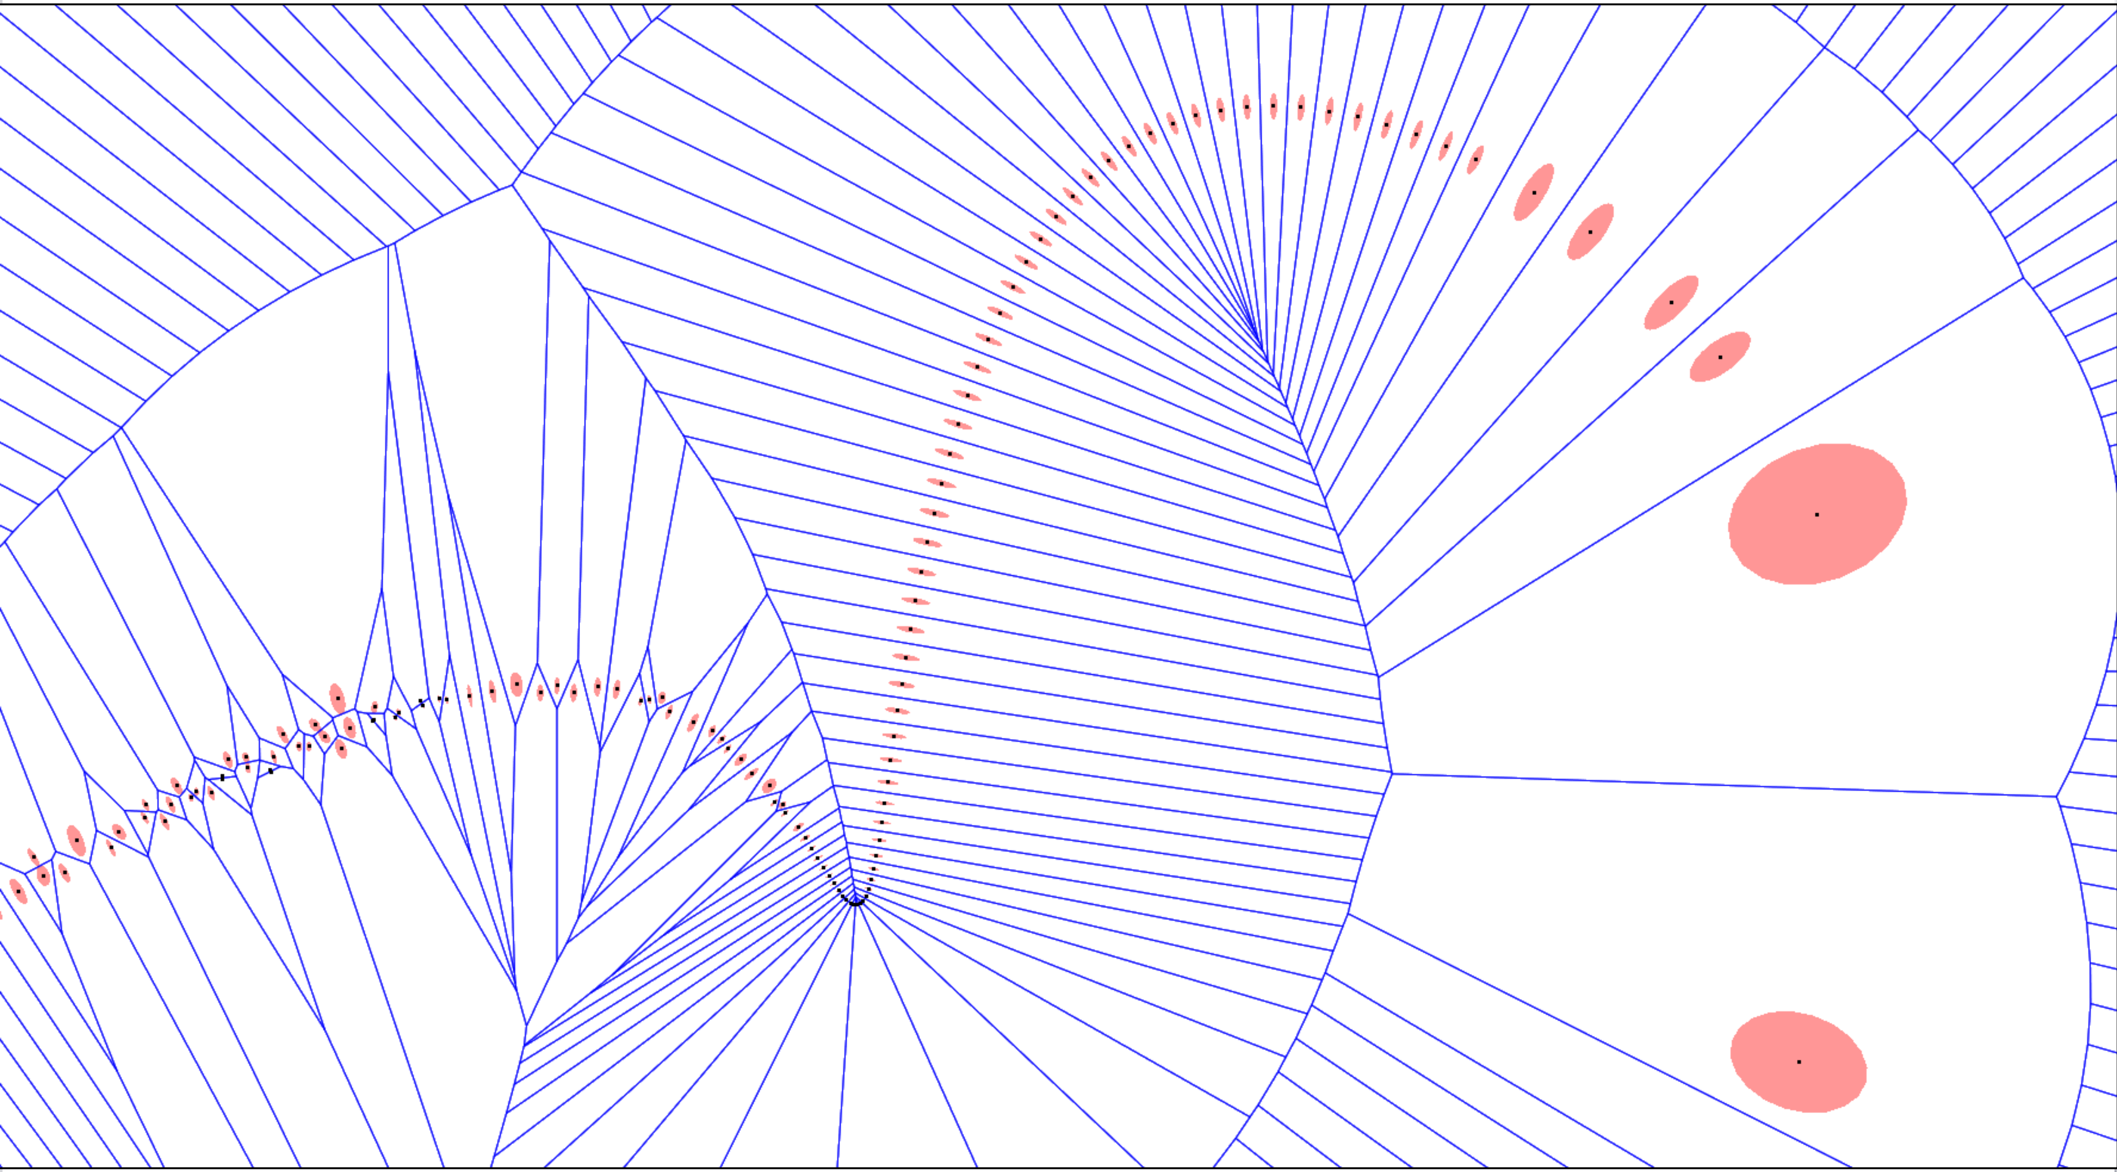

Algorithm Overview

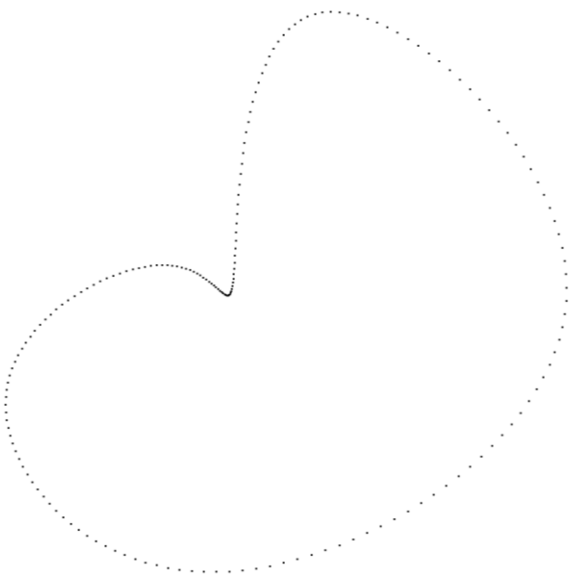

point set

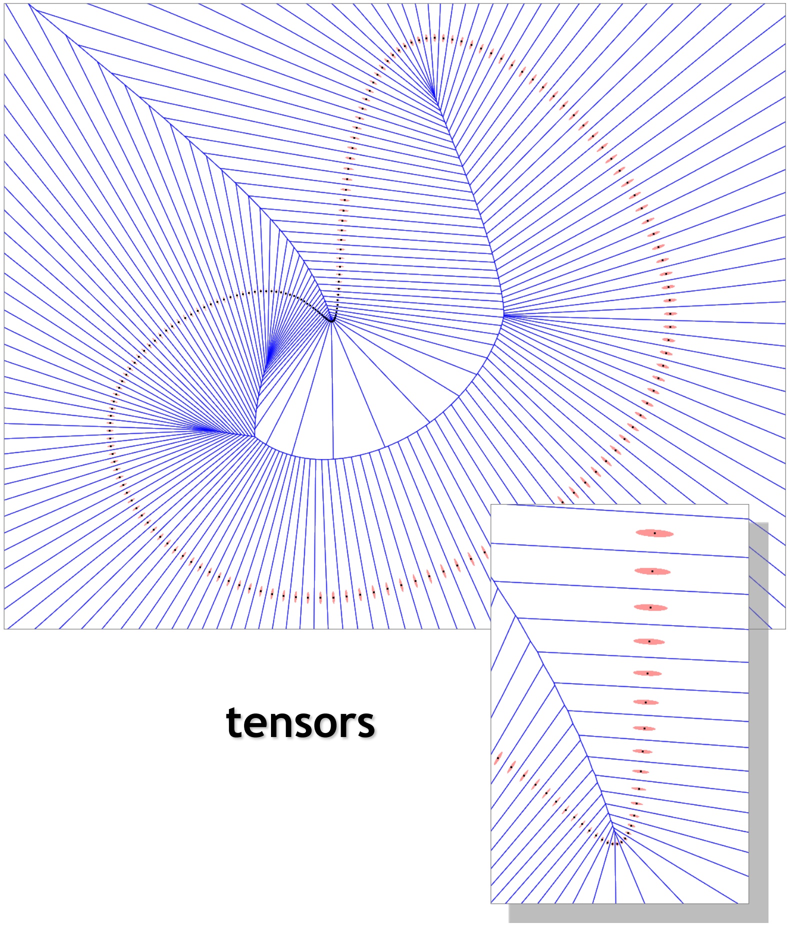

tensors

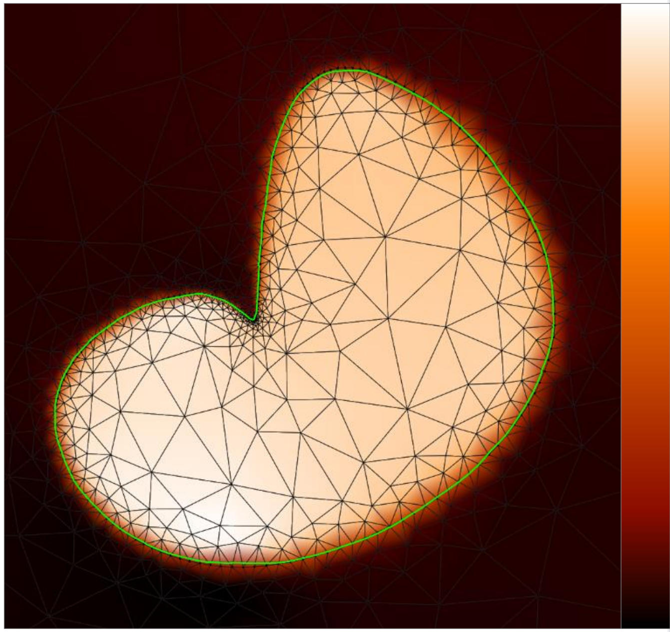

implicit function
+ contouring

Tensor Estimation

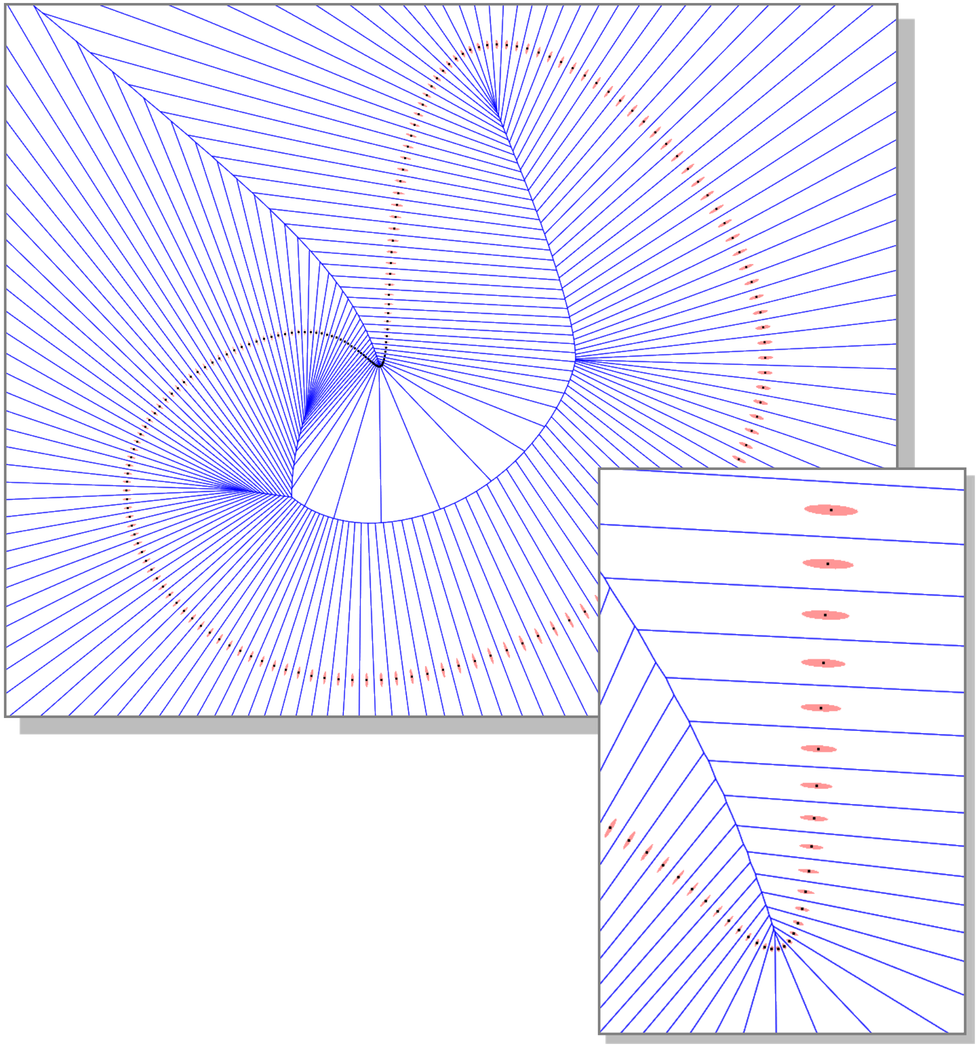

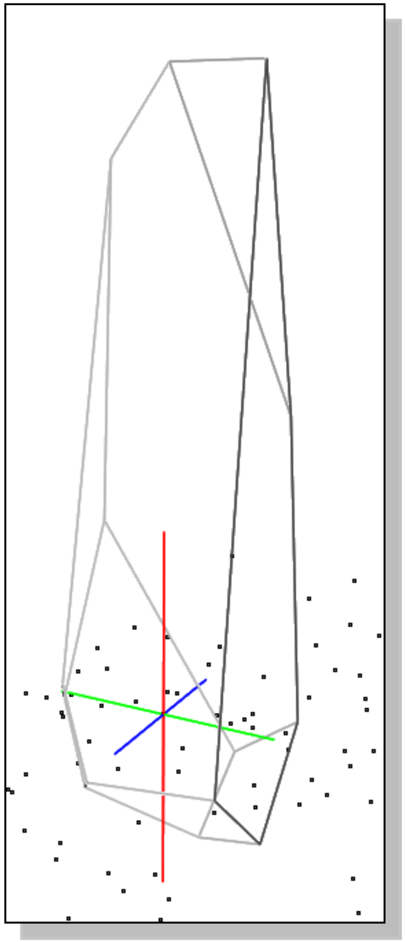

$$\int_{\Omega} (X - p)(X - p)^T dV$$

Noise-free vs Noisy

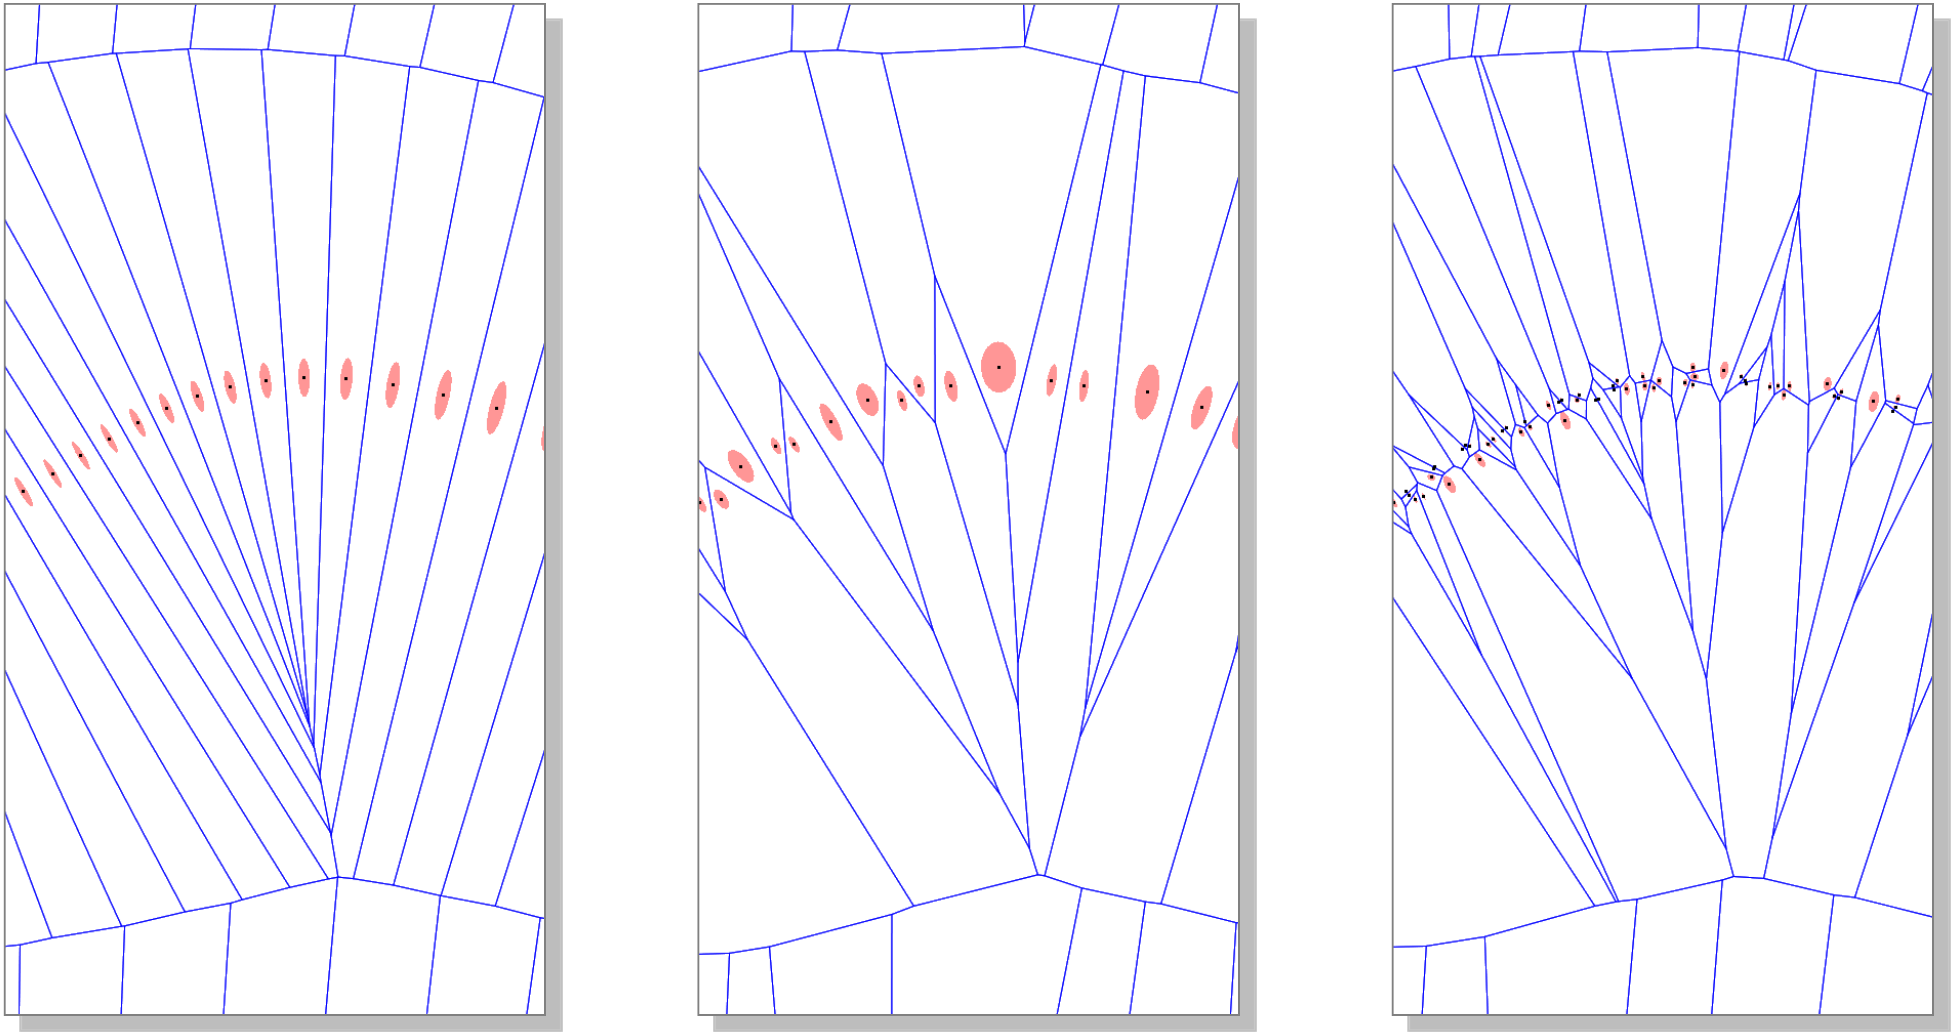

Dealing with Noise

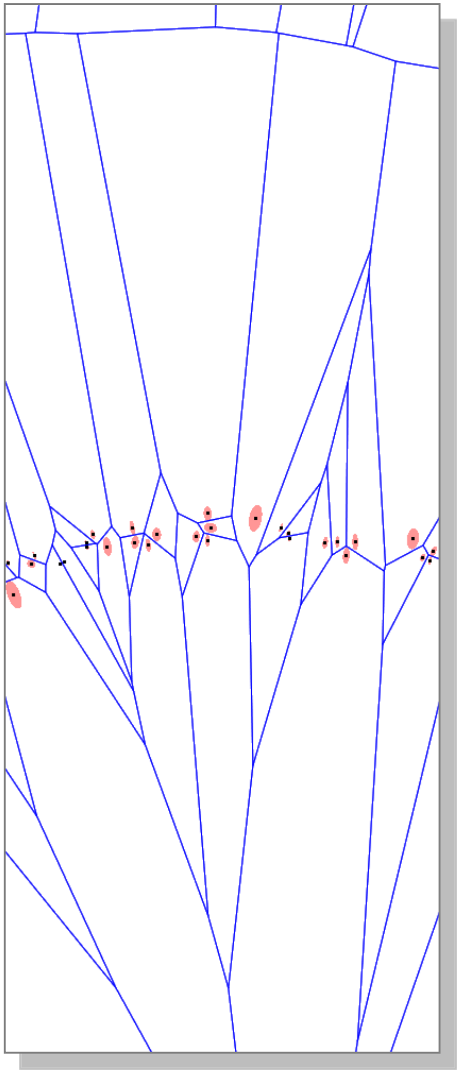

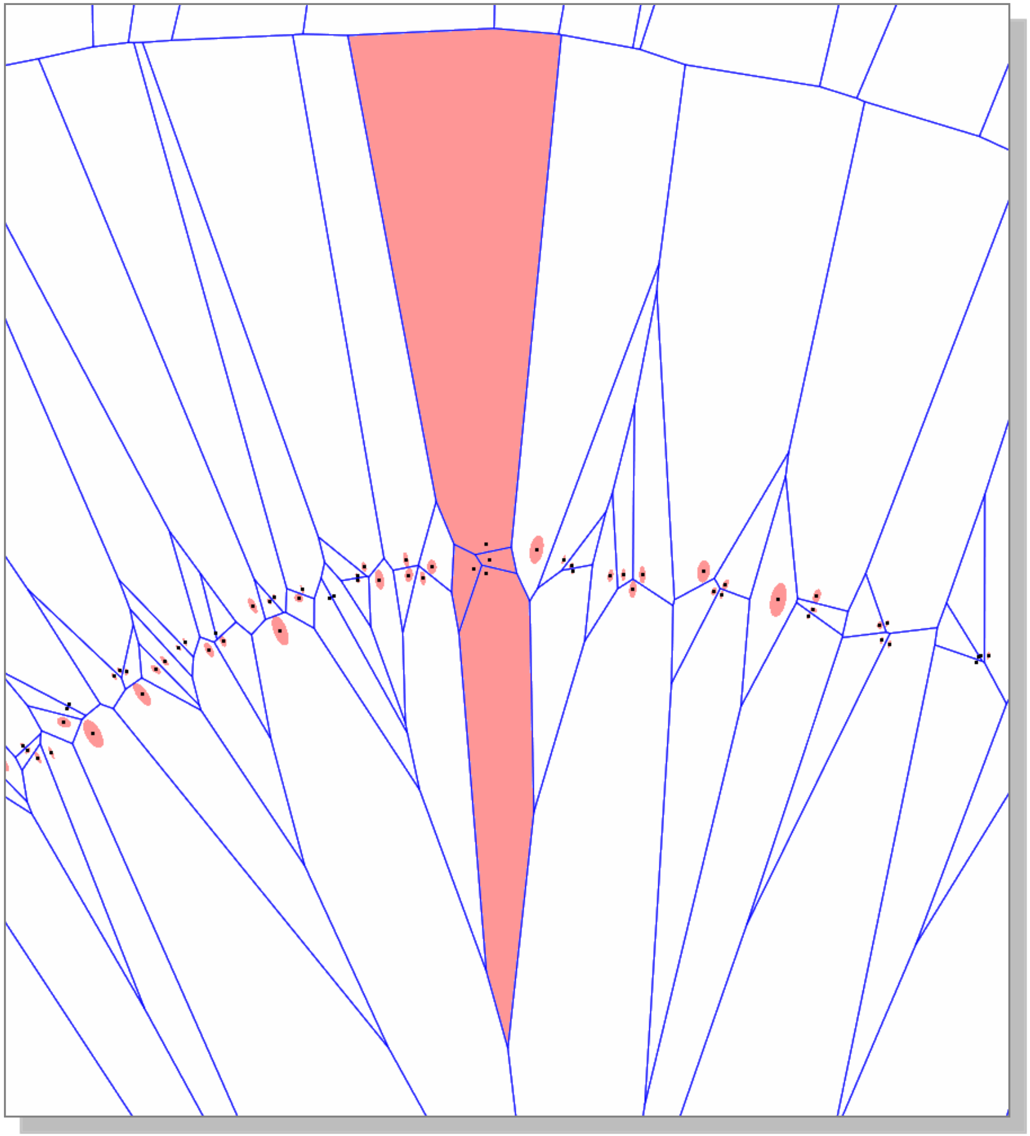

Implicit Function

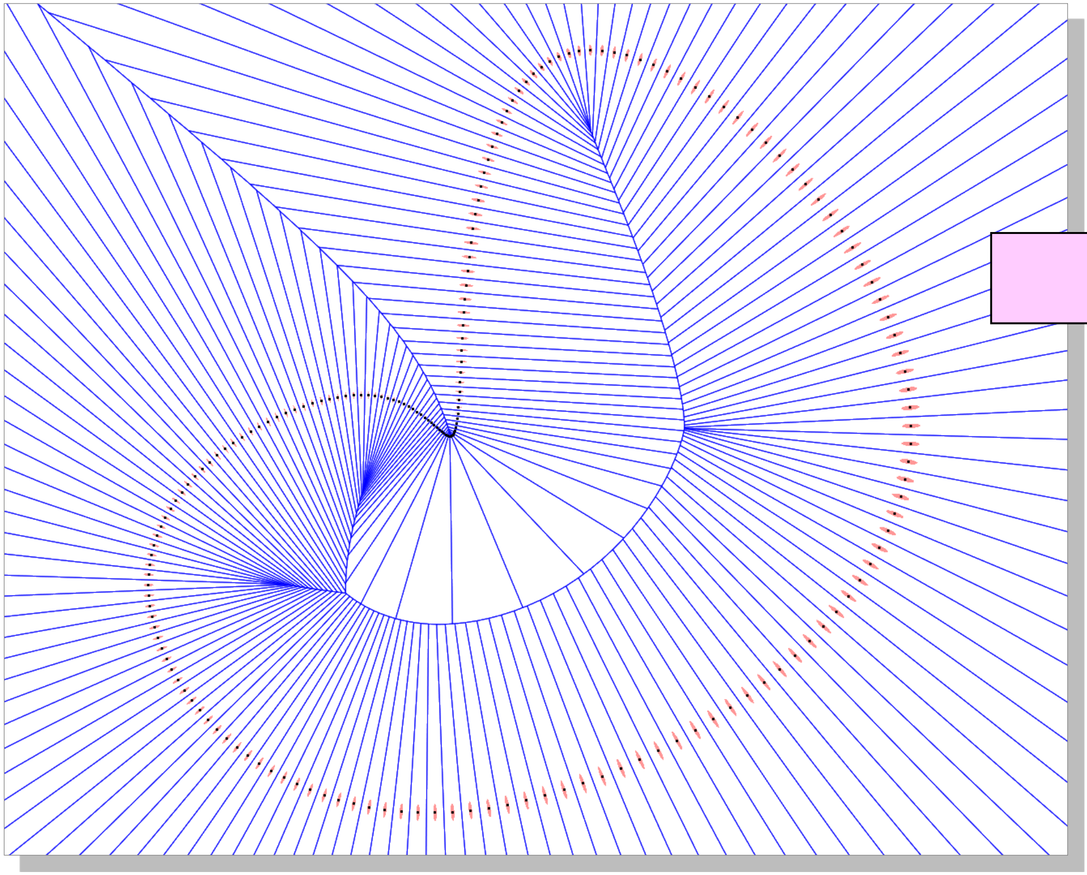

tensors

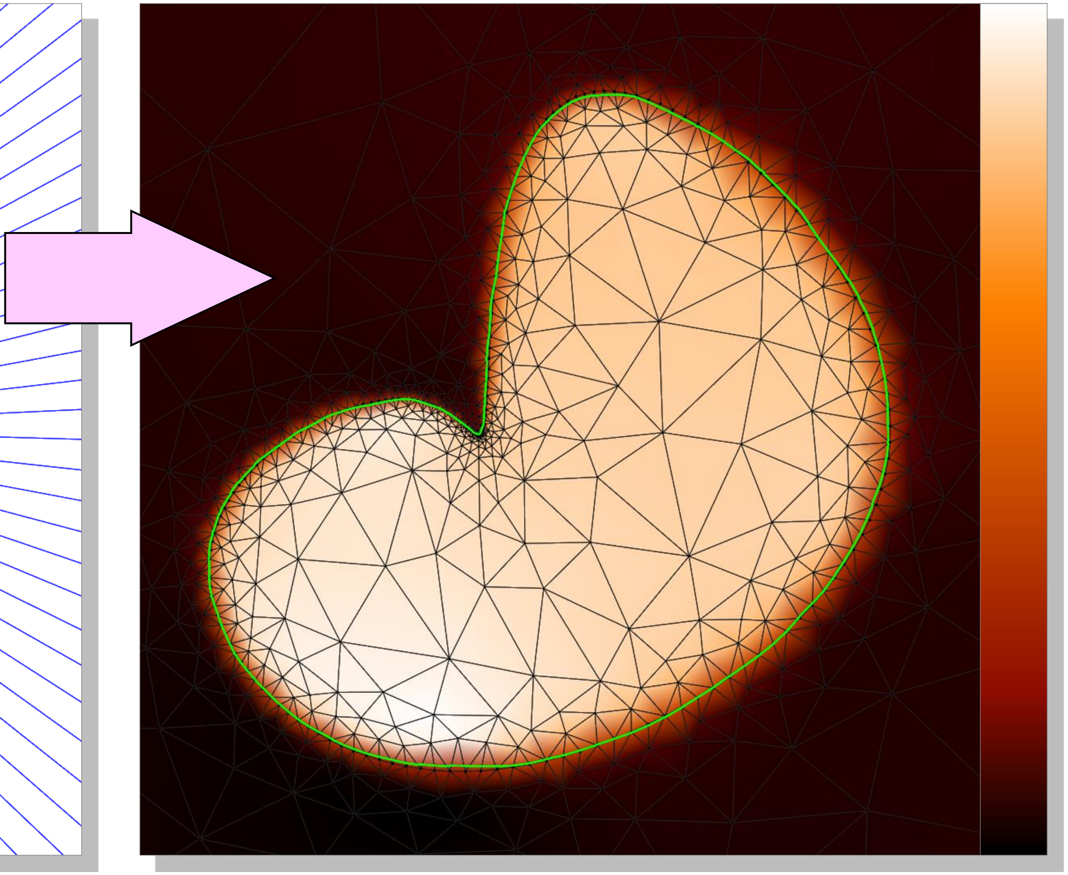

implicit function

Formulation

Find implicit function f such that its gradient ∇f best aligns to the principal component of the tensors.

Given a tensor field C, find the *maximizer* f of:

$$E_C^D(f) = \int_\Omega \nabla f^t \, C \, \nabla f \text{ subject to: } \int_\Omega \left[|\Delta f|^2 + \varepsilon |f|^2 \right] = 1$$

Anisotropic Dirichlet energy
Measures alignment with tensors

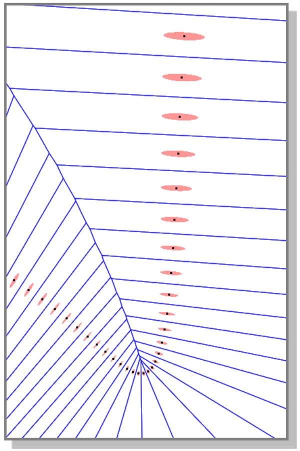

Biharmonic energy
Measures smoothness of ∇f

Formulation

Find implicit function *f* such that its gradient ∇f <u>best aligns to the principal component</u> of the tensors.

Given a tensor field C, find the *maximizer* f of:

$$E_C^D(f) = \int_\Omega \nabla f^t \, C \, \nabla f \text{ subject to: } \int_\Omega \left[|\Delta f|^2 + \varepsilon |f|^2 \right] = 1$$

Anisotropic Dirichlet energy
Rewards <u>alignment</u> with tensors

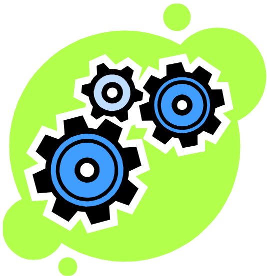

Biharmonic energy
Favors <u>smoothness</u> of ∇f

fitting vs smoothness

Rationale

On areas with:

- anisotropic tensors: favors alignment
- isotropic tensors: favors smoothness

Large aligned gradients + smoothness

leads to consistent orientation of ∇f

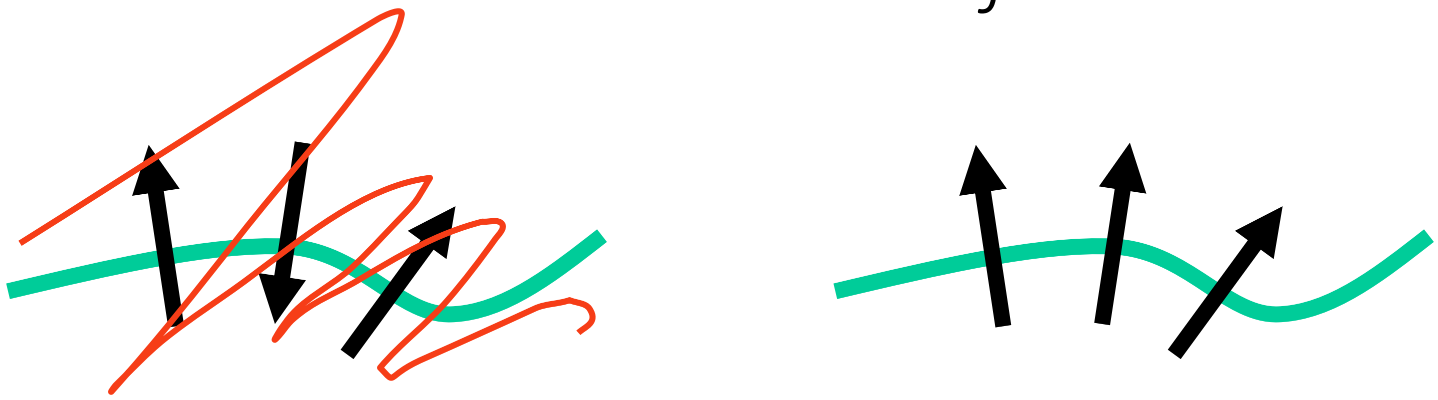

Generalized Eigenvalue Problem

Given a tensor field C, find the *maximizer* f of:

$$E_C^D(f) = \int_\Omega \nabla f^t \, C \, \nabla f \text{ subject to: } \int_\Omega \left[|\Delta f|^2 + \varepsilon |f|^2 \right] = 1$$

A: anisotropic Laplacian operator

$$E_C^D(F) \approx F^t \, A \, F$$

B: isotropic Bilaplacian operator

$$E^B(f) \approx F^t \, B \, F$$

$$\boxed{AF = \lambda BF}$$

max

Eigenvector (PWL function)

Generalized Eigenvalue Problem

$$AF = \lambda BF$$

eigenvector
(PWL function)

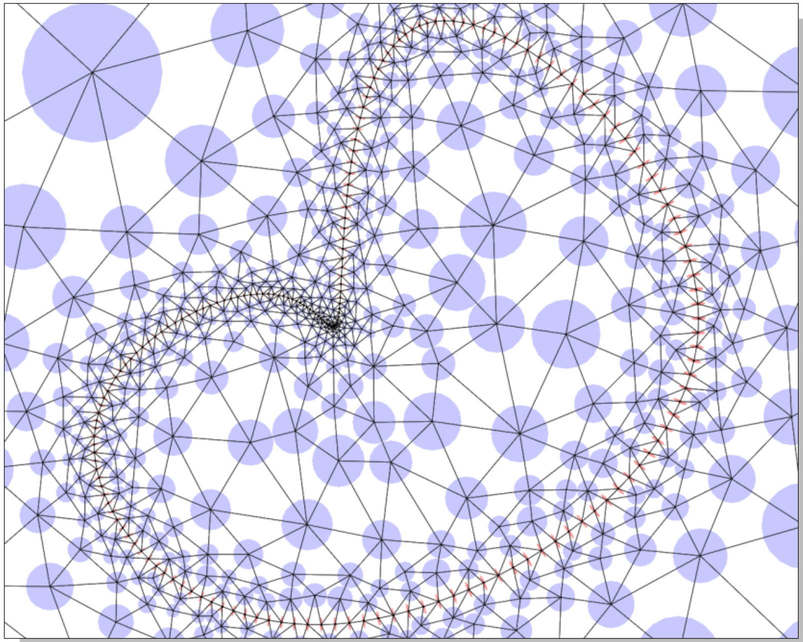

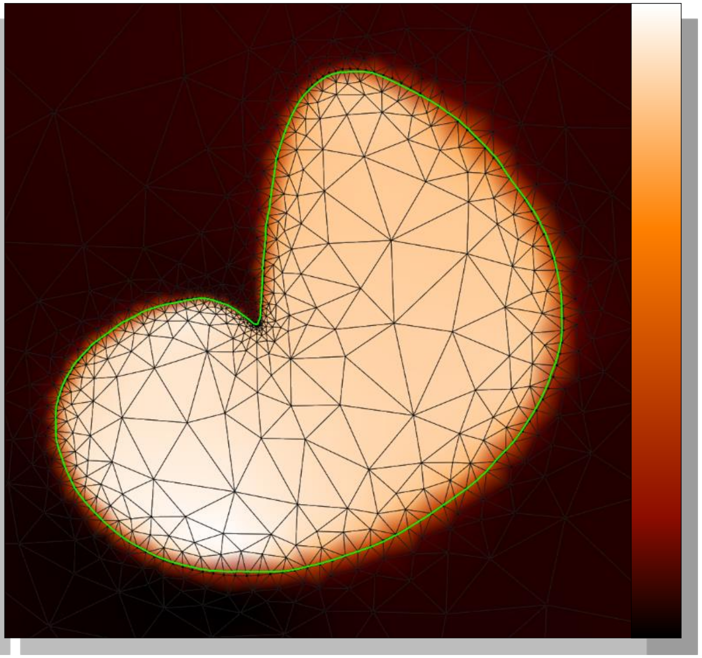

Sparse Sampling

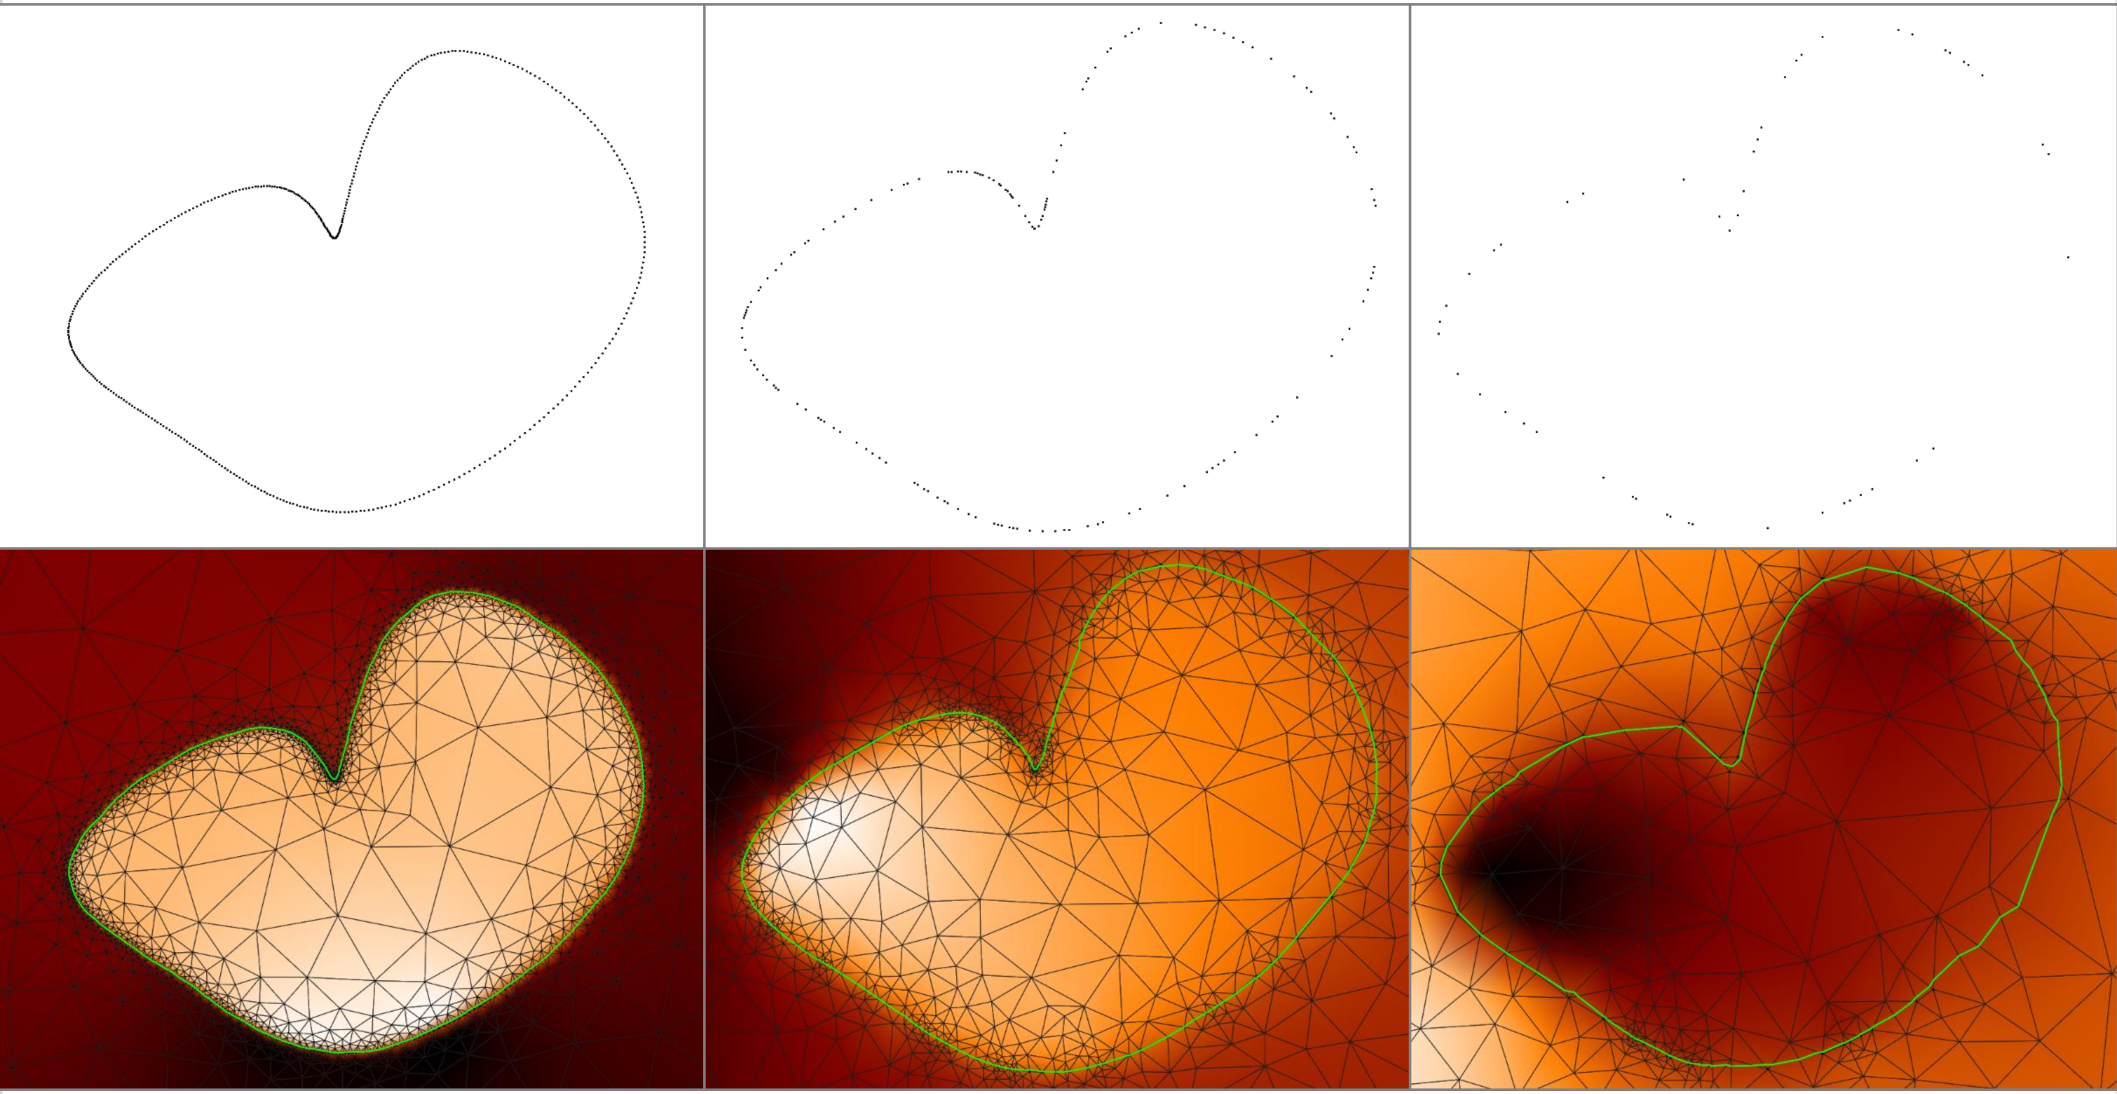

Noise

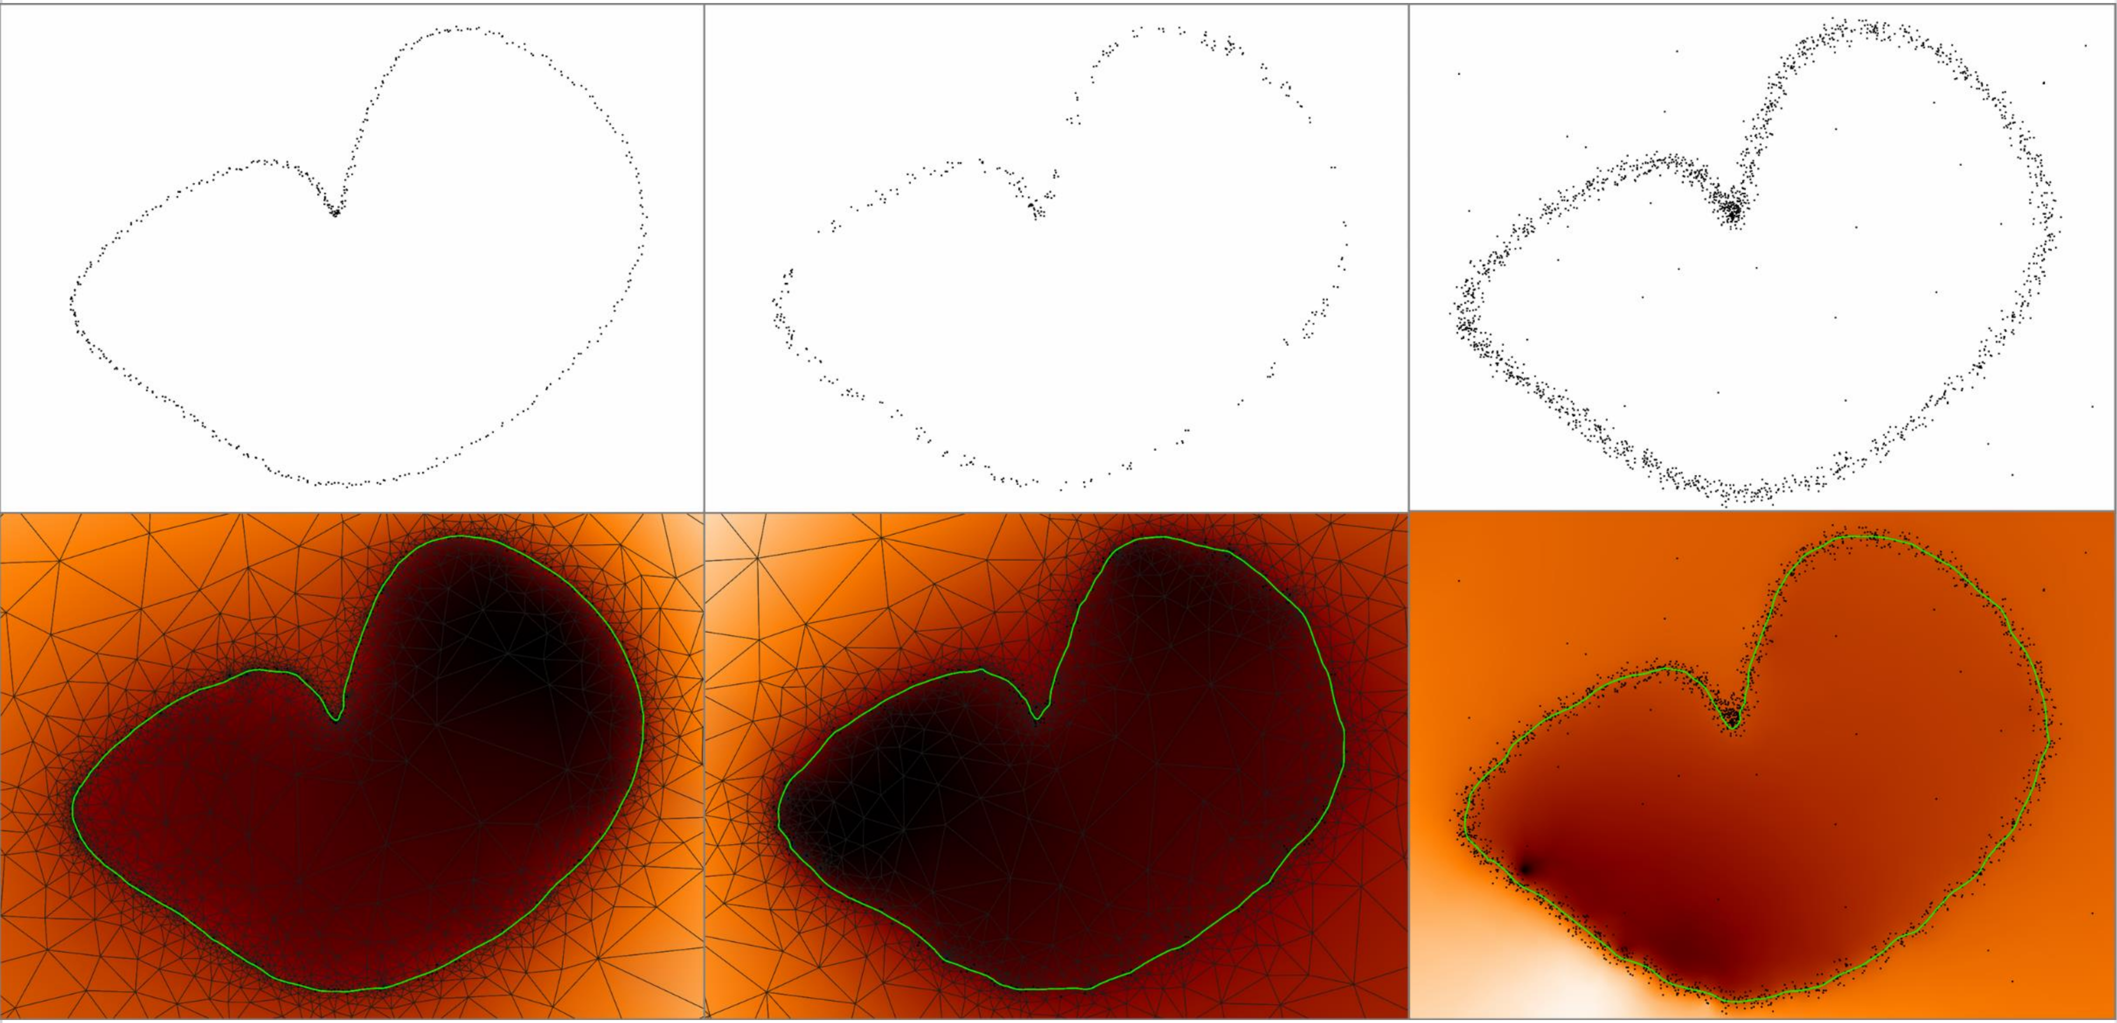

vs Poisson Reconstruction

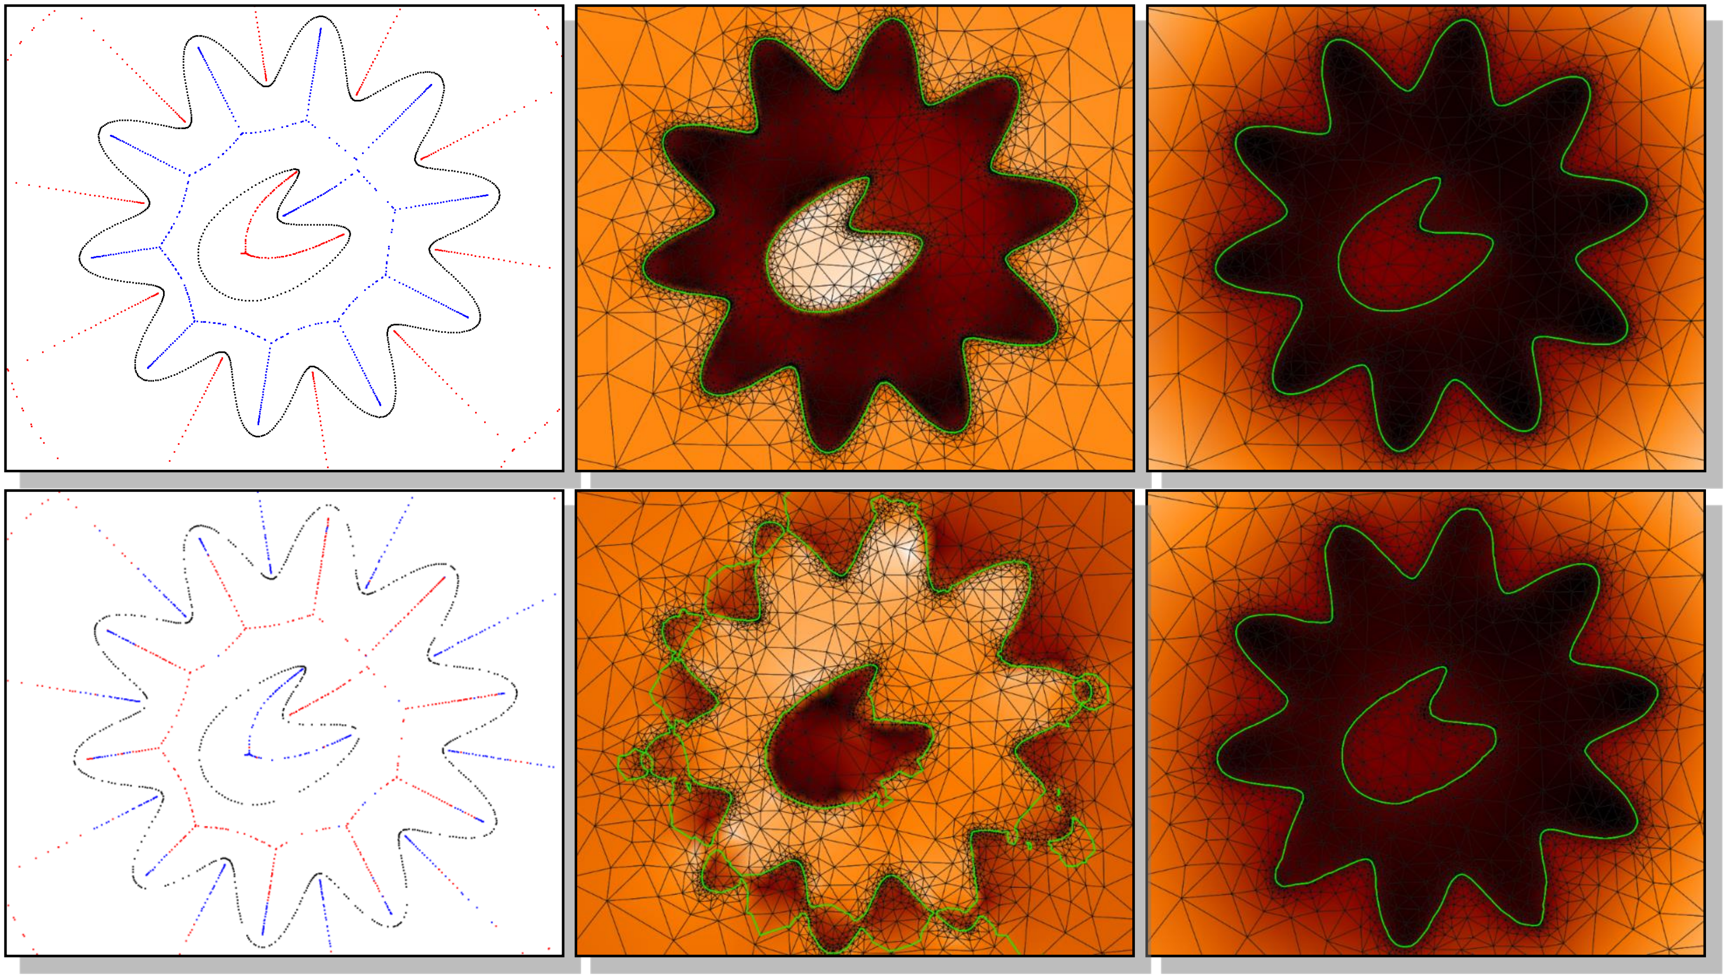

Oriented points | Poisson | Spectral

Sforza