Well-spaced samples of generic surfaces have sparse Delaunay triangulations

Jeff Erickson
University of Illinois

Independent/joint work with
Dominique Attali, Jean-Daniel Biossonnat, and André Lieuiter
(to appear at SoCG 2003)

Nice samples of nice surfaces have nice Delaunay triangulations

Jeff Erickson
University of Illinois

Delaunay triangulations are neat!

Jeff E.

Surface reconstruction

Input: set P of *sample points* from an unknown smooth surface ∑

Output: an geometric approximation of ∑ with the correct topology

Several provable algorithms
[Amenta, Bern, Boissonnat, Cazals, Dey, Edelsbrunner, Eppstein, Funke, Giesen, Hiyoshi, ...]

Lots of practical heuristics and improvements!

Delaunay triangulation

- ≤ 4 points form a *Delaunay simplex* if they lie on the boundary of an empty *Delaunay ball*
- n points in space can have $\Omega(n^2)$ Delaunay simplices

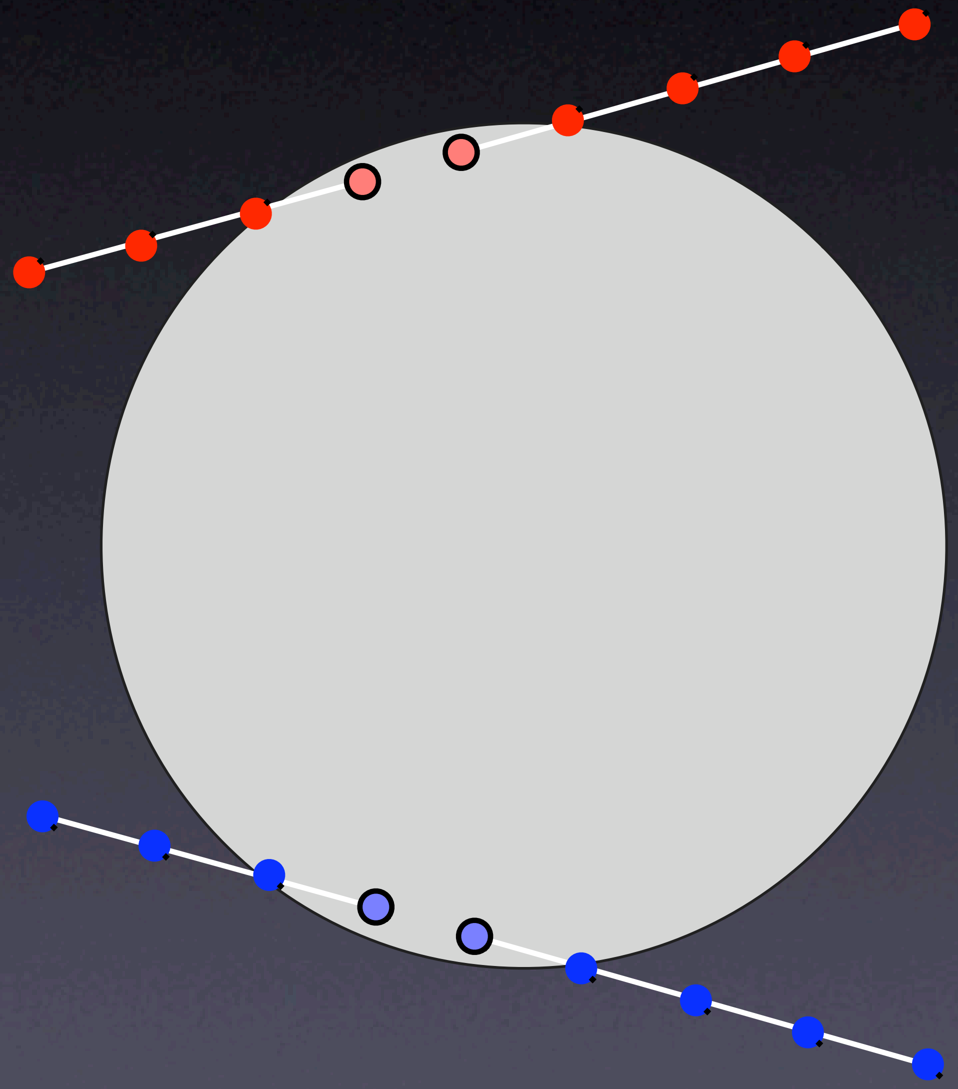

Theory ≠ Practice

Theory:
Delaunay triangulations have *quadratic* complexity (in the worst case).

Practice:
Delaunay triangulations have *linear* complexity.

Well, then it’s not a very good “theory”, is it?

Practical Delaunay complexity

Random points:

- in space: $O(n)$ [Meijering '53, Miles '72; Dwyer '91]
- on fixed convex polyhedron: $O(n)$
  [Golin and Na '00]
- on fixed nonconvex polyhedron: $O(n \log^4 n)$
  [Golin and Na '02]

Practical Delaunay complexity

(ϵ,k)-sample of Σ: Any ball of radius ϵ centered on Σ contains at least 1 and at most k samples

- fixed polyhedron: $O(k^2 n)$
  [Attali and Boissonnat '01]⇐!!!
- arbitrary fixed surface: $\Theta(k^2 n^{3/2})$
  Lower bound: [E'01], Upper bound: [E'02]
- New: fixed *generic* surface: $O(k^2 n \log n)$

Warning: fixed surfaces

Any finite quantity that depends on the fixed surface Σ, but not on ϵ or k or any particular point, is considered a *constant*.

- surface area
- number of facets
- aspect ratios of facets
- angles between facets
- angles between edges
- diameter
- minimum local feature size
- min and max principal curvatures
- bounds on partial derivatives
- "genericity"

These constants are hidden in the O() notation.

"Delaunay condition number"

Why fix the surface?

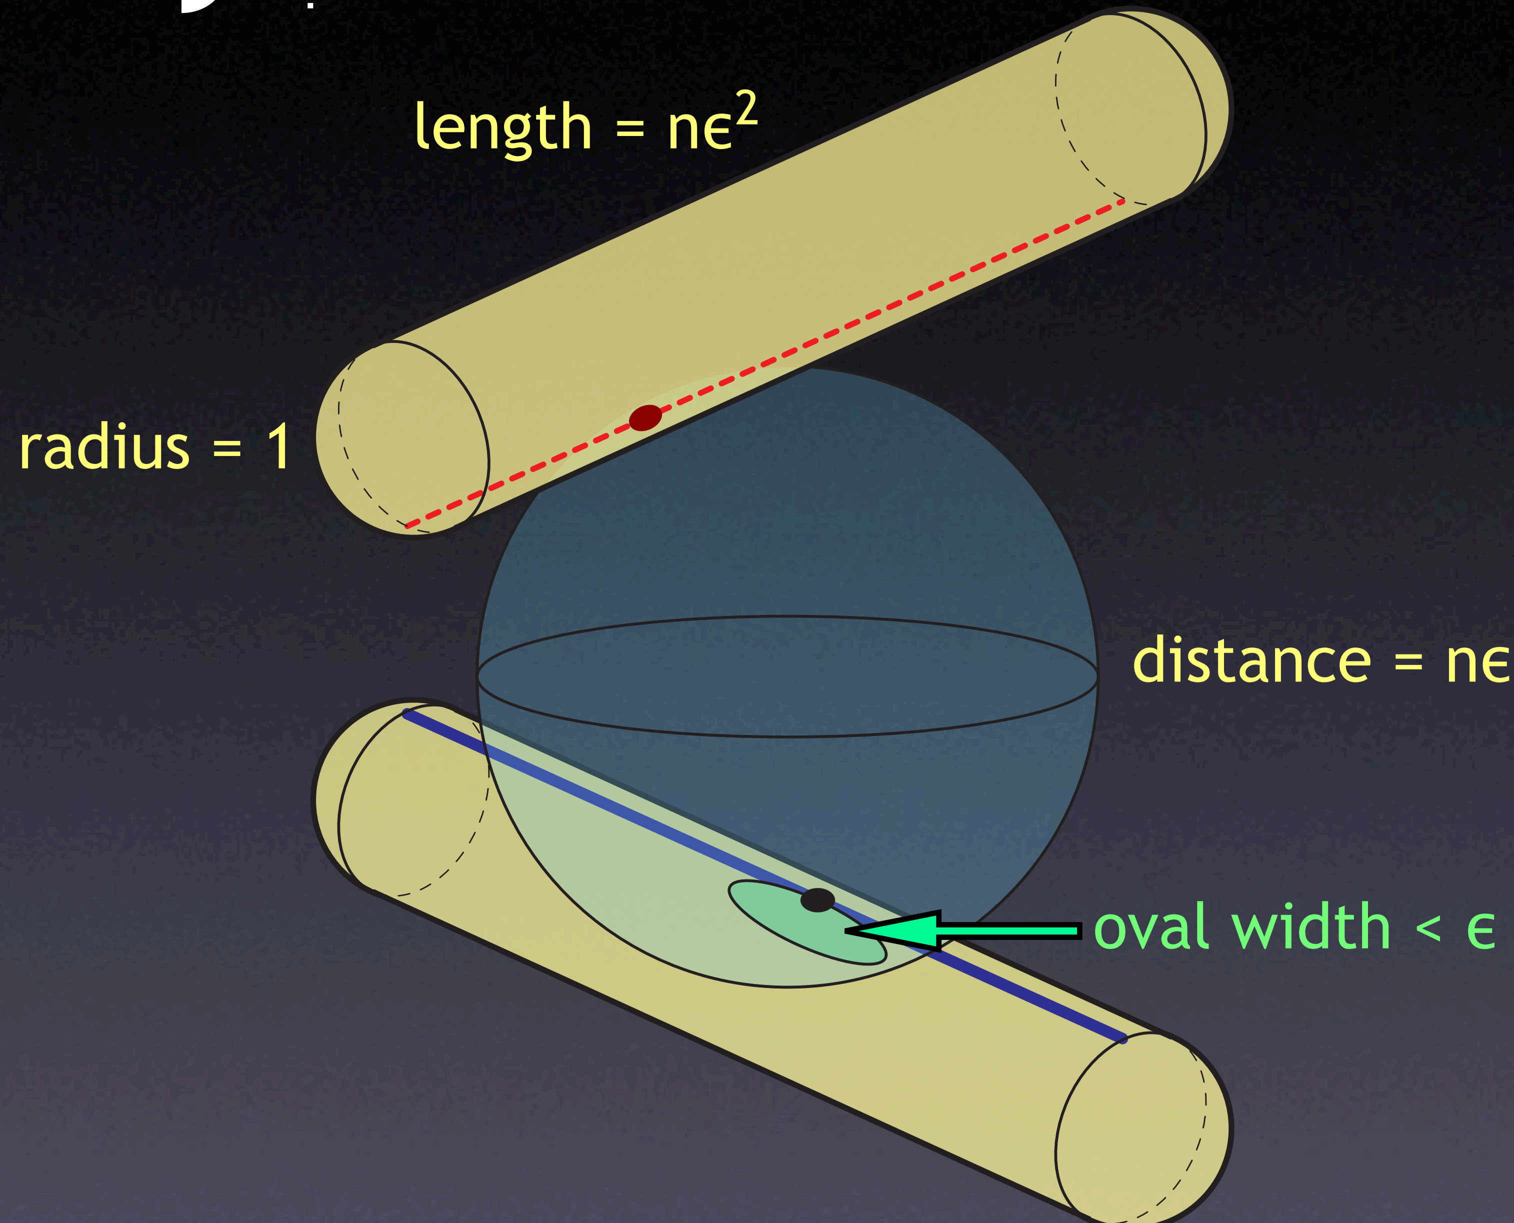

If the surface varies with n and ϵ, we can get $\Omega(n^2\epsilon^2)$ Delaunay simplices.

[E'01]

Why limit sample density?

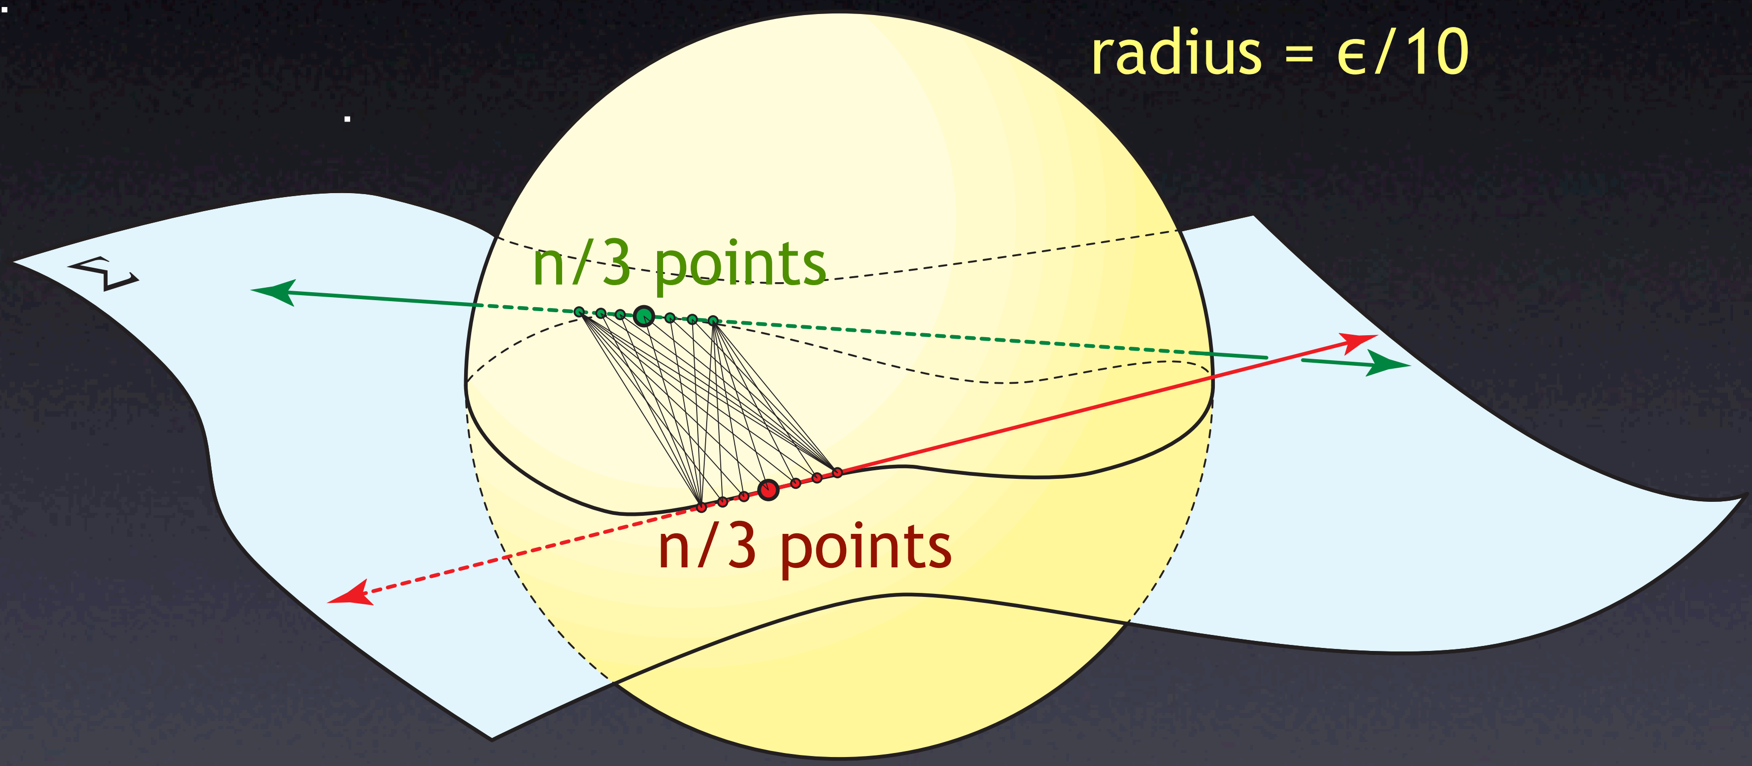

For **any** surface except the sphere, we can get $\Omega(n^2)$ Delaunay simplices by locally oversampling.

[E'01]

The helix

ϵ-sample of a cylinder:
√n turns of a helix,
√n points on each turn

Two points are Delaunay neighbors iff they are less than a full turn apart.

$\Omega(n^{3/2})$ Delaunay simplices!

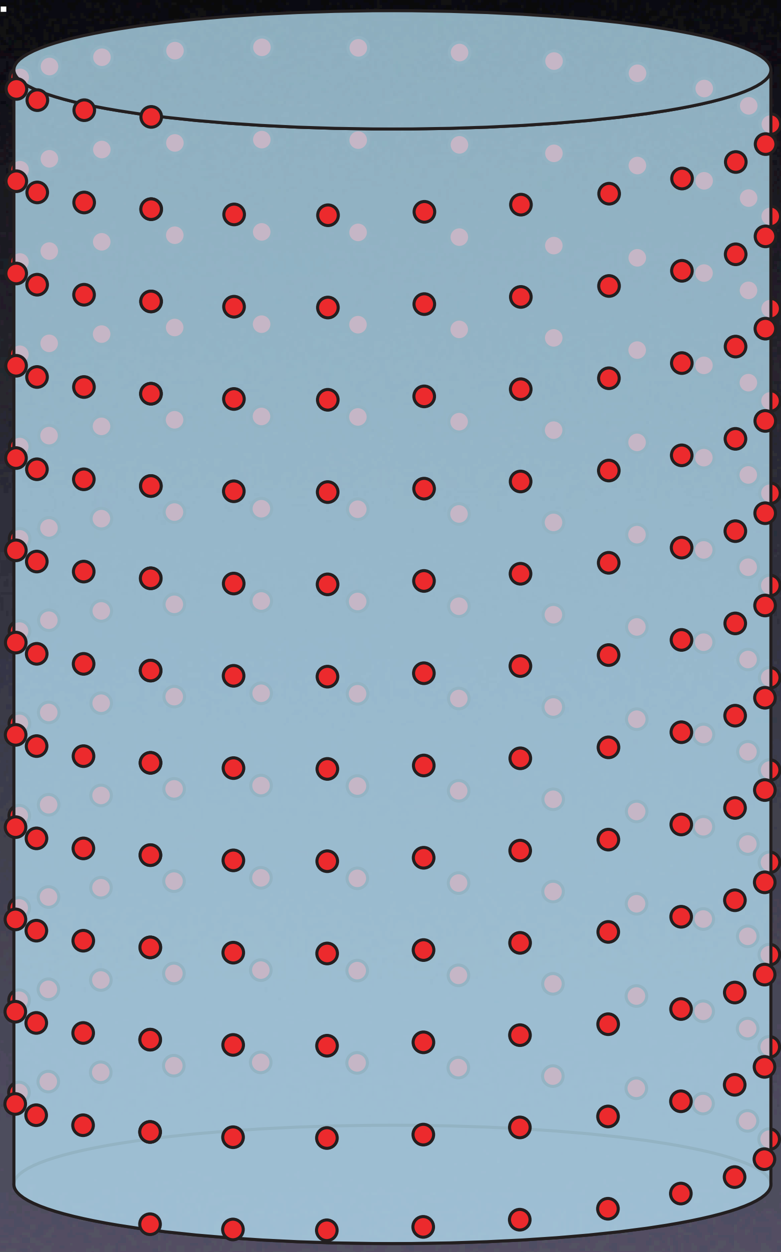

[E'01]

Spread Δ

Spread = diameter/closest pair distance
[Goodman, Pollack, and Sturmfels '89; Valtr *et al*. '93-'97]

Roughly related to dimensionality:

- nicely distributed in a volume ⇔ $\Delta \approx n^{1/3}$
- nicely distributed on a surface ⇒ $\Delta \approx n^{1/2}$
- nicely distributed on a curve ⇒ $\Delta \approx n$

Spread upper bound

- The Delaunay triangulation of any set of points with spread Δ has complexity $O(\Delta^3)$.
- The Delaunay triangulation of the union of k sets, each with spread Δ, has complexity $O(k^2\Delta^3)$.
- An (ϵ,k)-sample of a fixed surface is the union of k sets with spread $\Theta(\sqrt{n})$, so its Delaunay complexity is $O(k^2n^{3/2})$.

[E'02]

Theory ≠ Practice

Theory:
Delaunay triangulations of nice surface samples have complexity $\Theta(n^{3/2})$ in the worst case, and the worst case example is simple!

Practice:
Okay, sure, but Delaunay triangulations of *real* surface samples always have *linear* complexity!

So it's still not a very good "theory", is it?

Medial axis

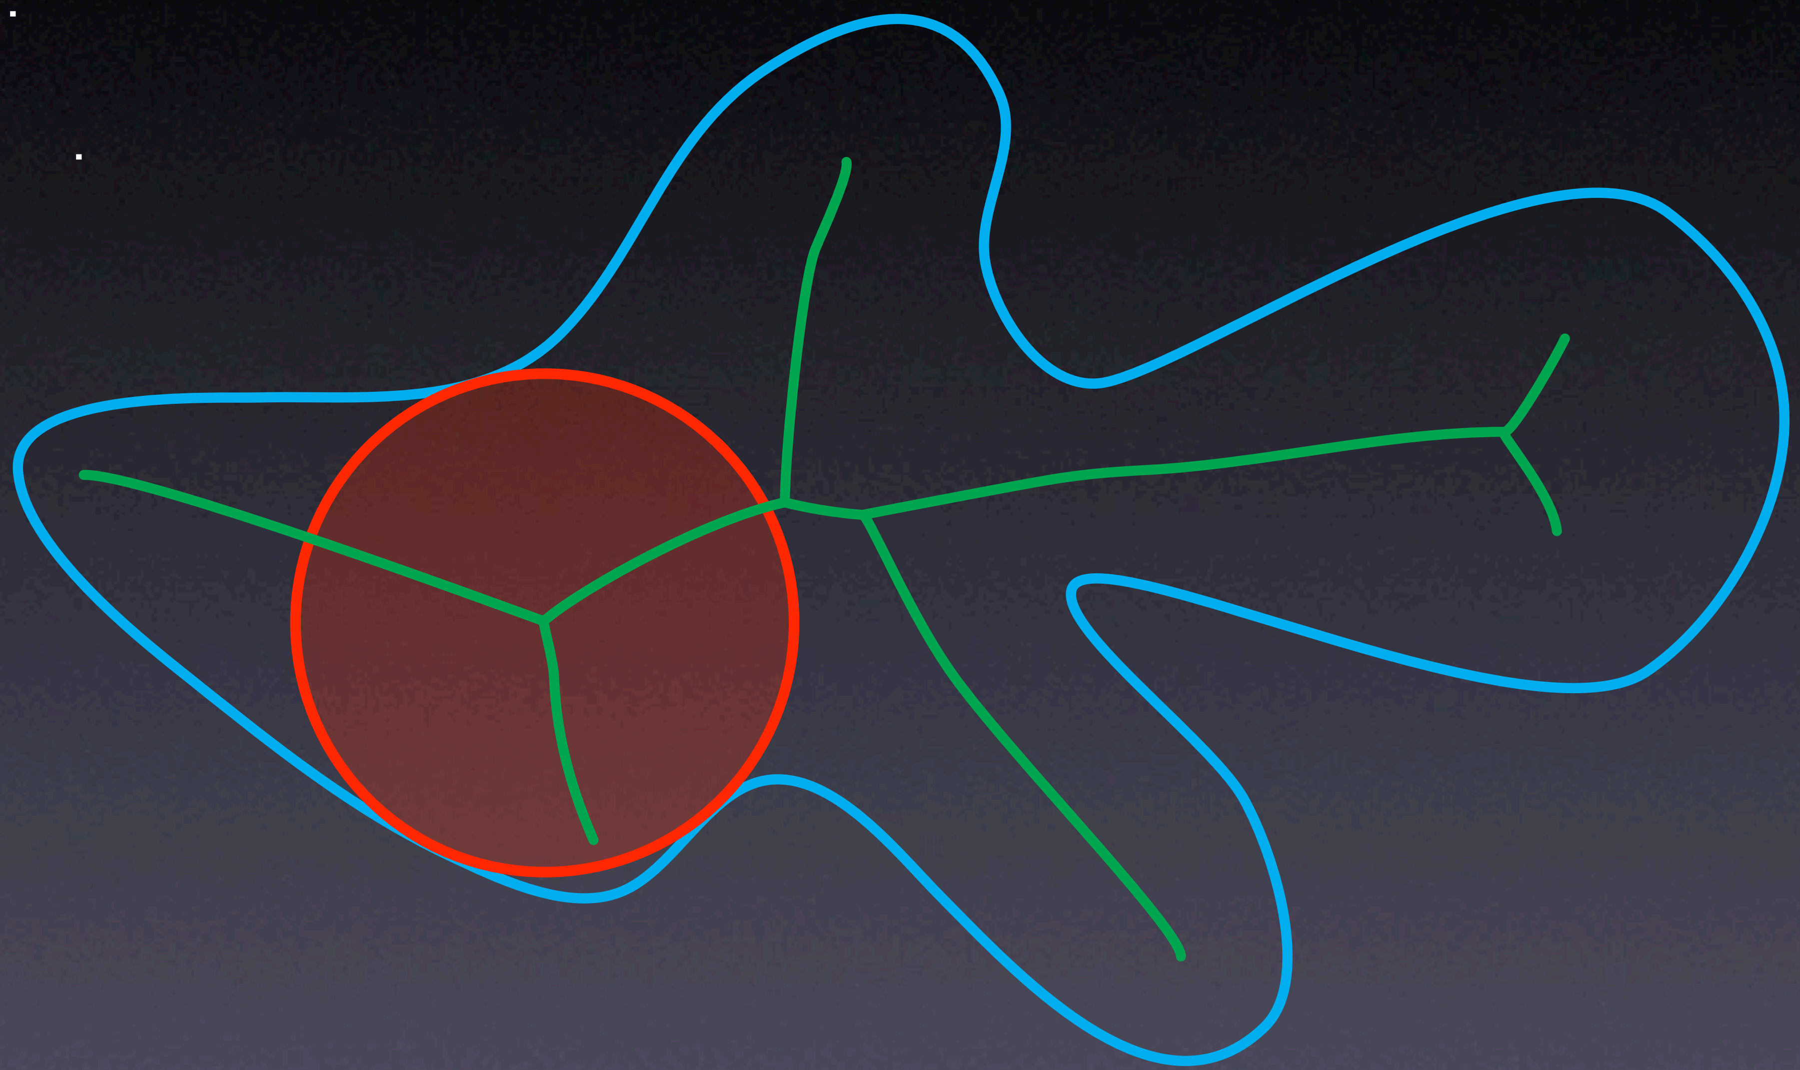

medial ball: empty interior, touches Σ more than once

medial axis: centers of medial balls

Generic contact types

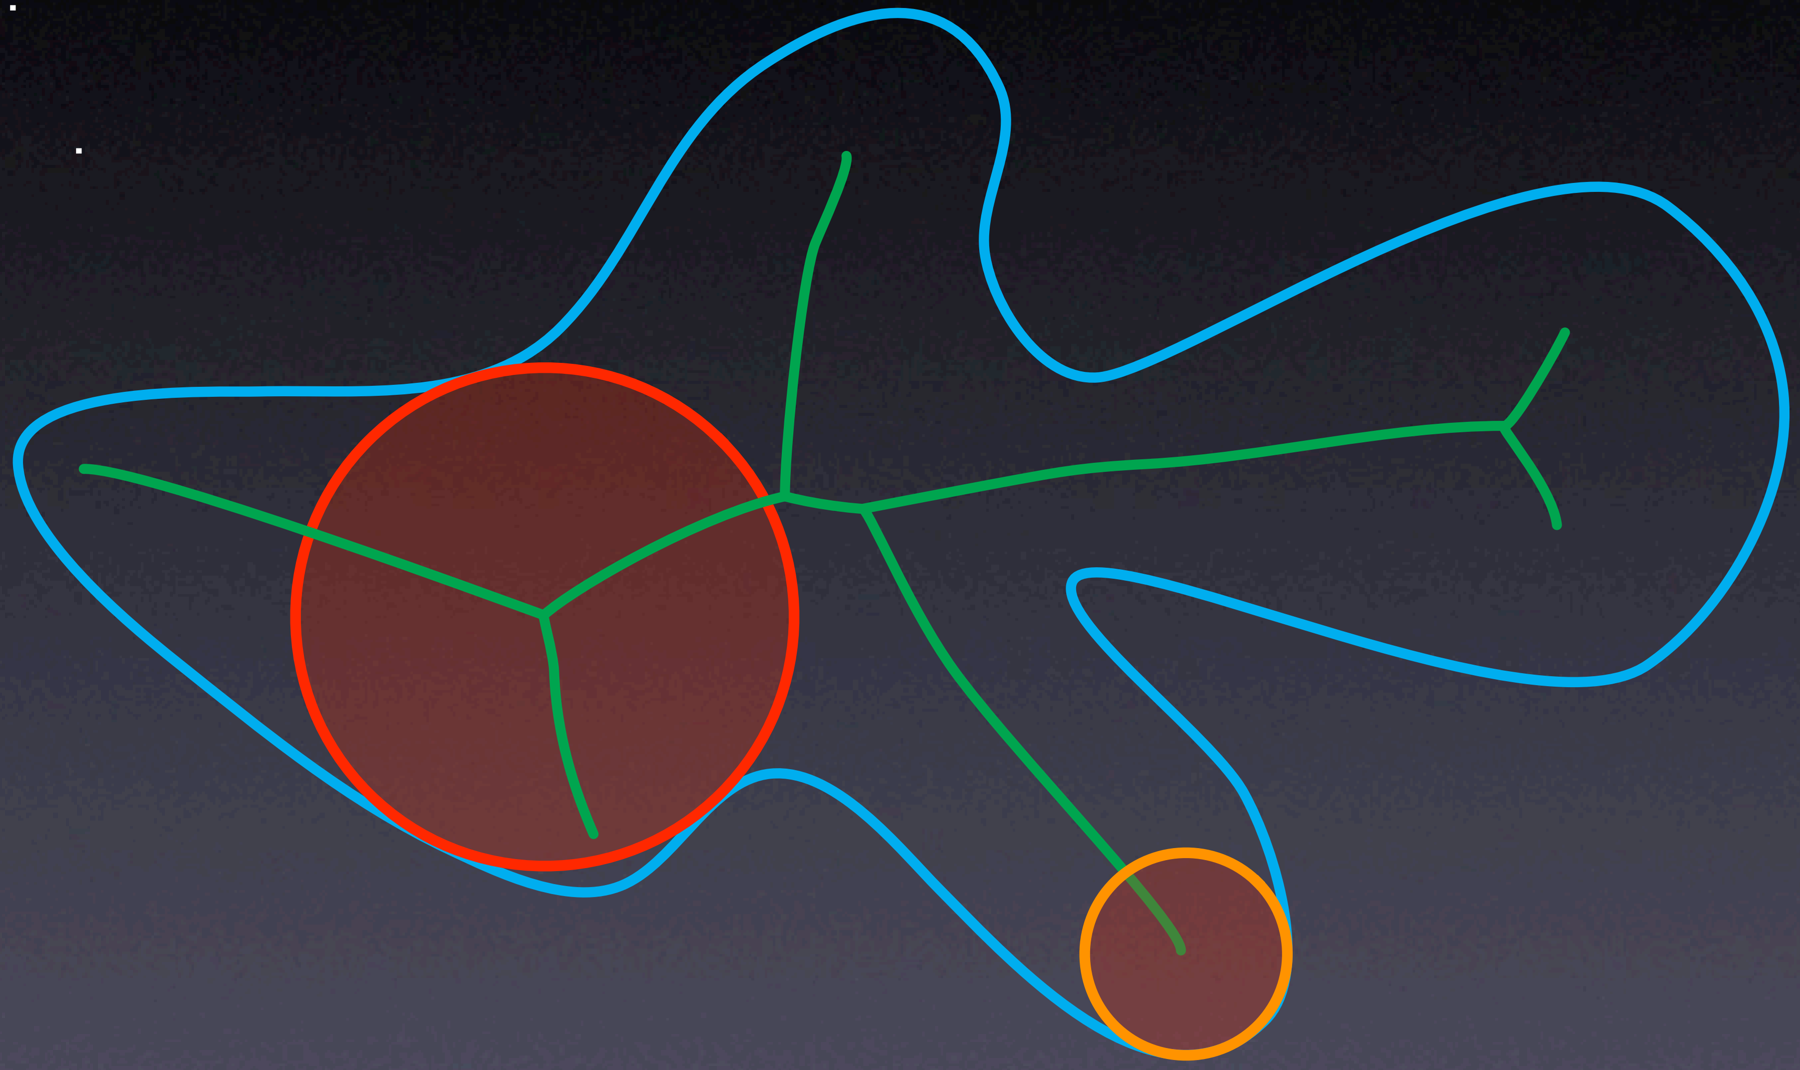

A_1: simple tangency

A_3: osculation ⇒ local maximum of curvature

Generic surfaces

- Σ is *generic* if no medial ball touches Σ more than four times, counting with multiplicity
- Generically, only A_1 and A_3 contacts.
- A_3 contacts lie on *ridge curves* on Σ: local max of principal curvature
- Almost every surface is generic, but *not* surfaces of revolution or Herbert's skin.

Generic medial axes

- Exactly 5 generic medial axis features
  [Bryzgalova '77; Mazov '82; Bruce, Giblin, and Gibson '85; Bogaevsky '89; Giblin and Kimia '00; Leymarie and Kimia '03]
-
-

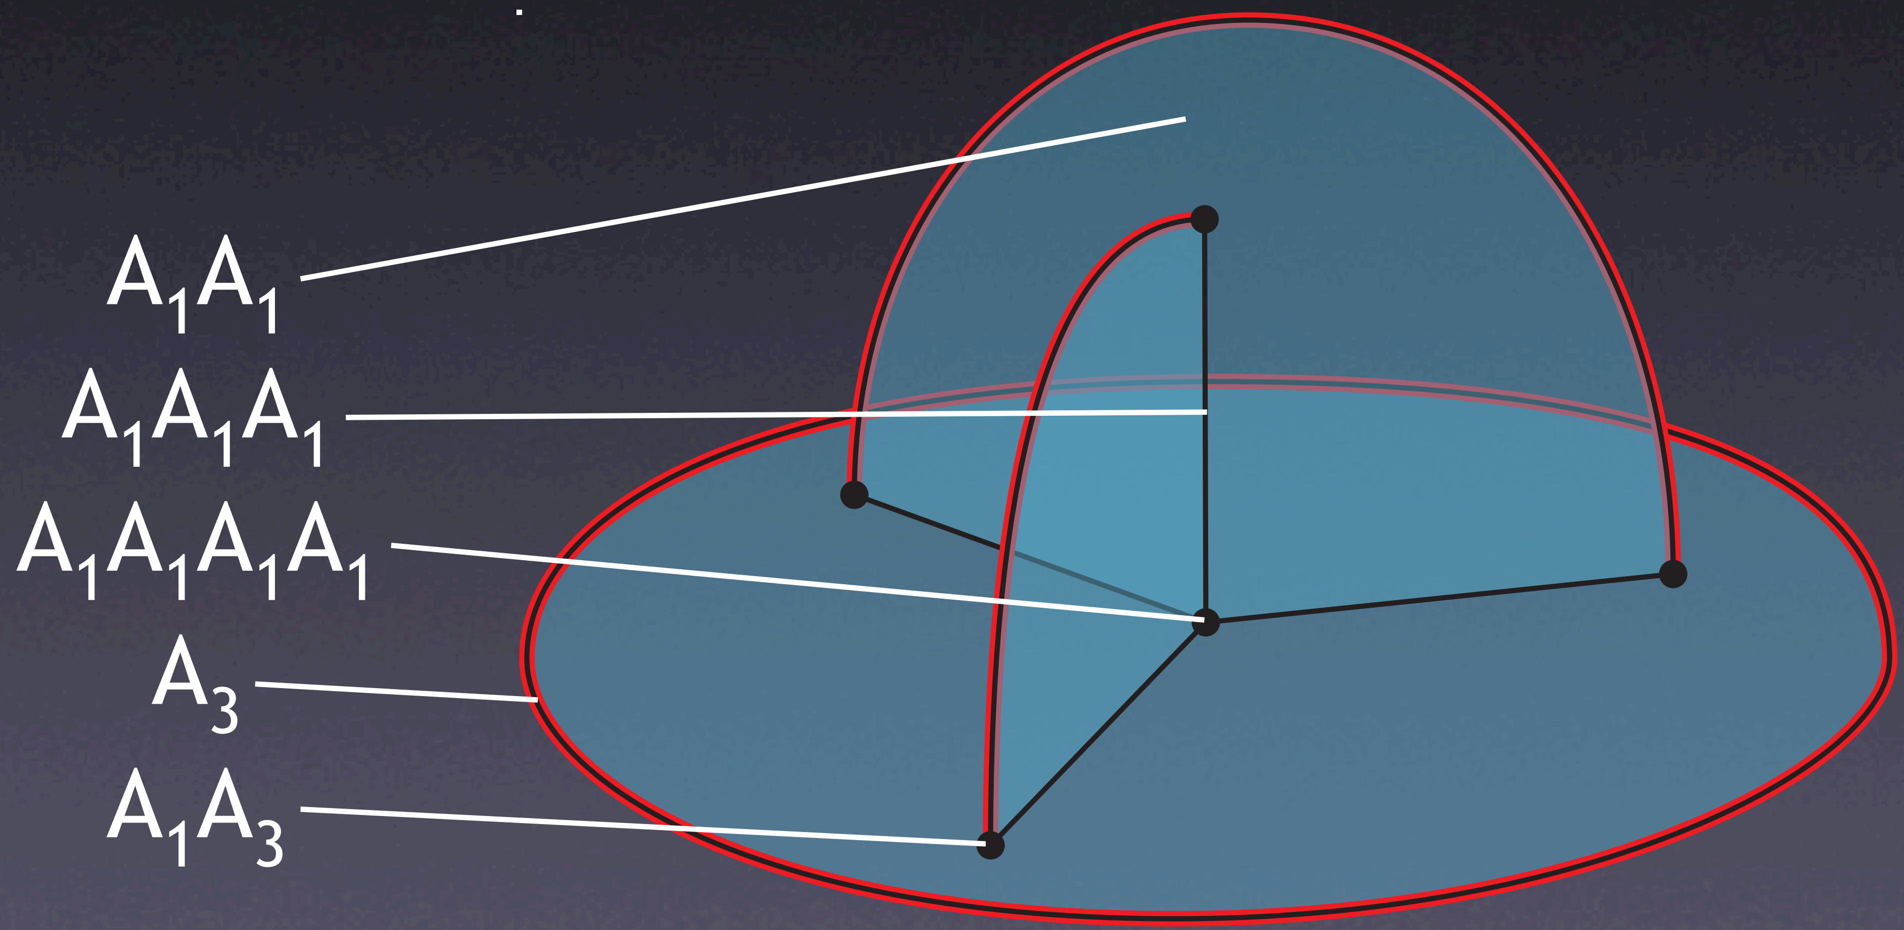

Intuition

As the sampling density increases, Delaunay balls behave more and more like medial balls.

So to understand the *combinatorial* behavior of Delaunay balls in the limit as $\epsilon \to 0$, we need to study the *differential* behavior of medial balls!

Curvature measures

- $r(p)$ = radius of medial ball tangent at p
- $\kappa_1(p)$ = maximum principal curvature at p
- $\kappa_2(p)$ = minimum principal curvature at p

- $\kappa_2 < \kappa_1 < 1/r$ except
  - $\kappa_2 = \kappa_1 < 1/r$ only at umbilic points
  - $\kappa_2 < \kappa_1 = 1/r$ only at A_3 contact curves

Delaunay ball intersecting surface

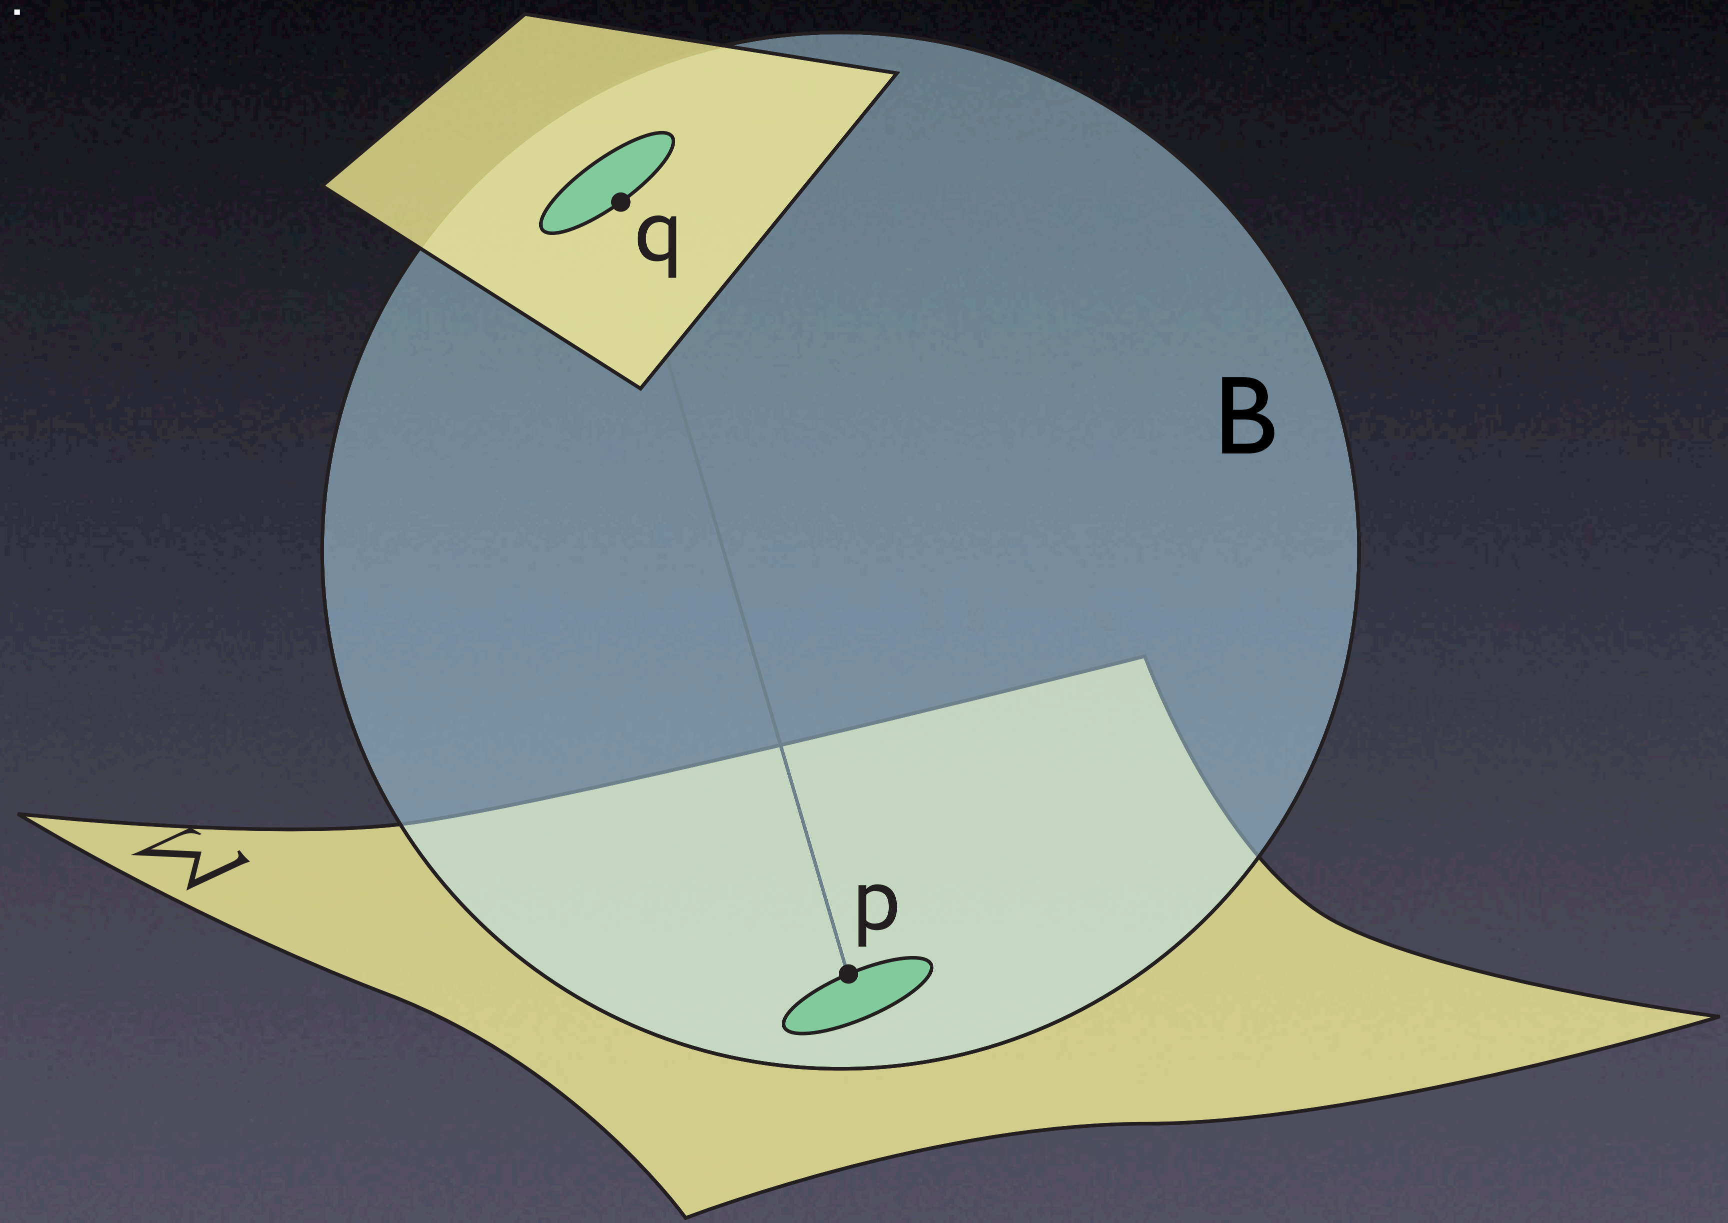

Expand a medial ball tangent at some point p far from A_3 contact curves, so $\kappa_1 < 1/r$.

Taylor series approximation ⇒ $B \cap \Sigma$ fits between two ellipses with aspect ratio

$$\sqrt{\frac{1 - r\kappa_2}{1 - r\kappa_1}}$$

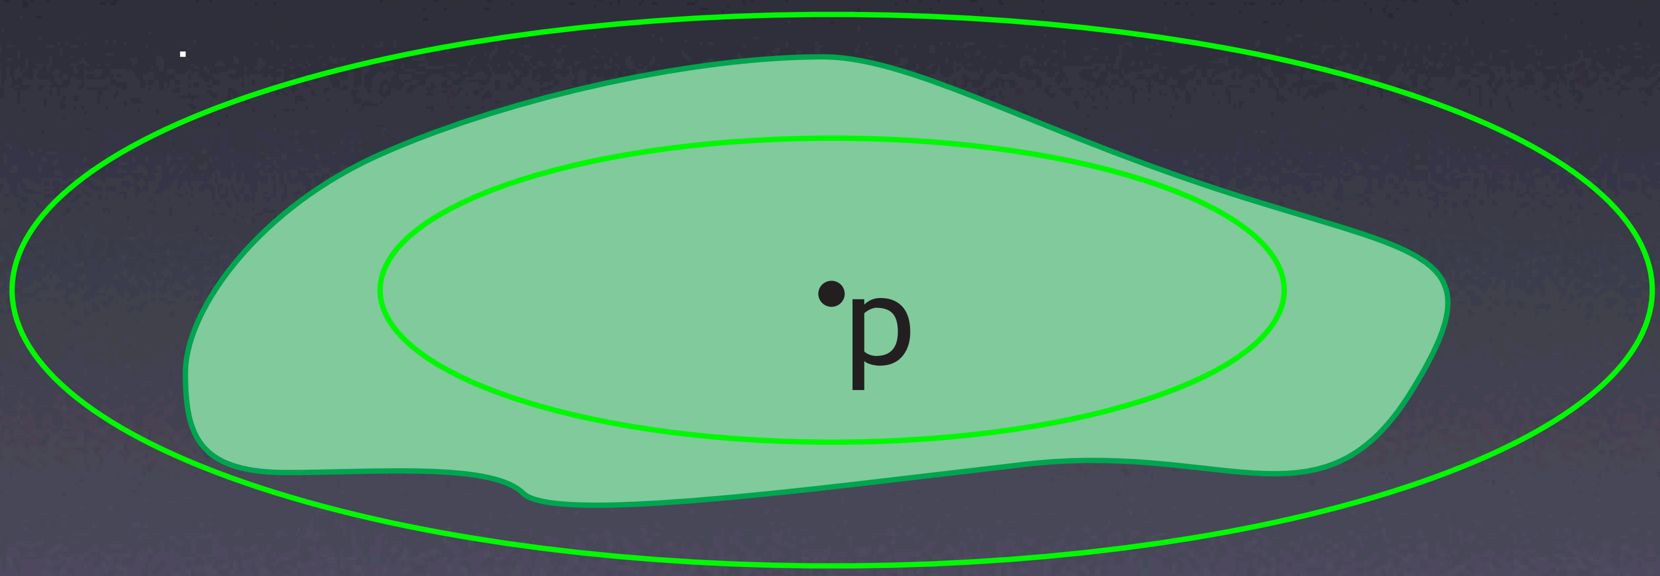

B is a Delaunay ball
⇒ small ellipse can't contain ϵ-disk
⇒ larger ellipse can't contain 2ϵ-disk

Area of blob < Area of larger ellipse < $4\pi\varepsilon^2\sqrt{\frac{1 - r\kappa_2}{1 - r\kappa_1}}$

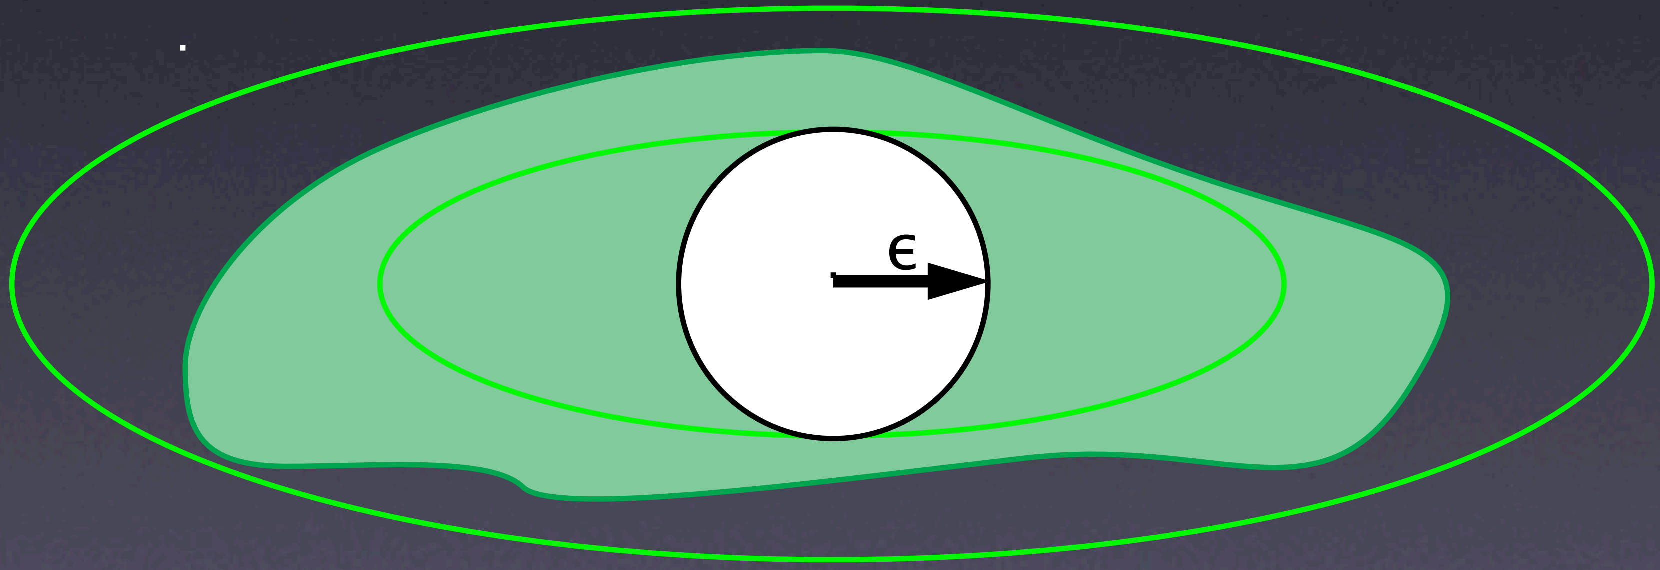

- Let p be a sample far from A_3 contact curves.
- Every local Delaunay neighbor of p carves out an empty blob on Σ.
- Since local neighbors are close, these blobs all look about the same.

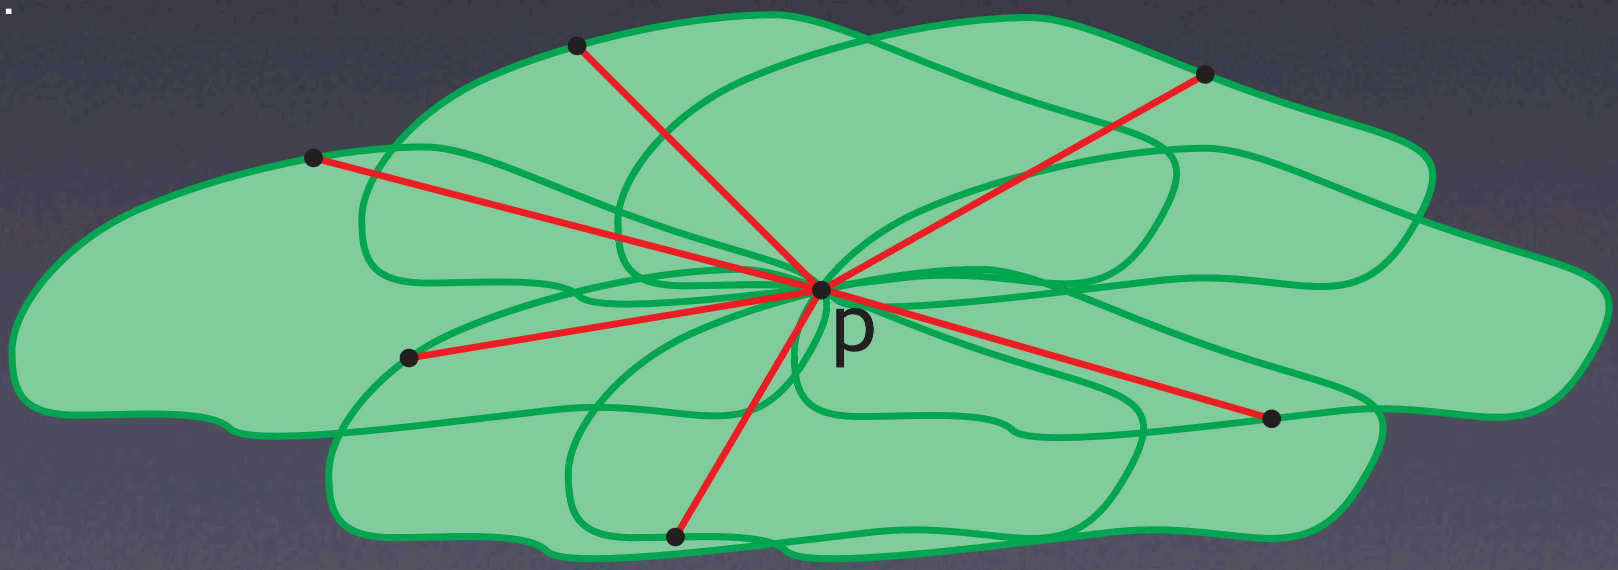

- Union of all blobs fits in ellipse with area

$$64\pi\varepsilon^2\sqrt{\frac{1-r\kappa_2}{1-r\kappa_1}} = O(\varepsilon^2)$$

- Large ellipse covered by O(1) ϵ-balls, each containing O(1) sample points

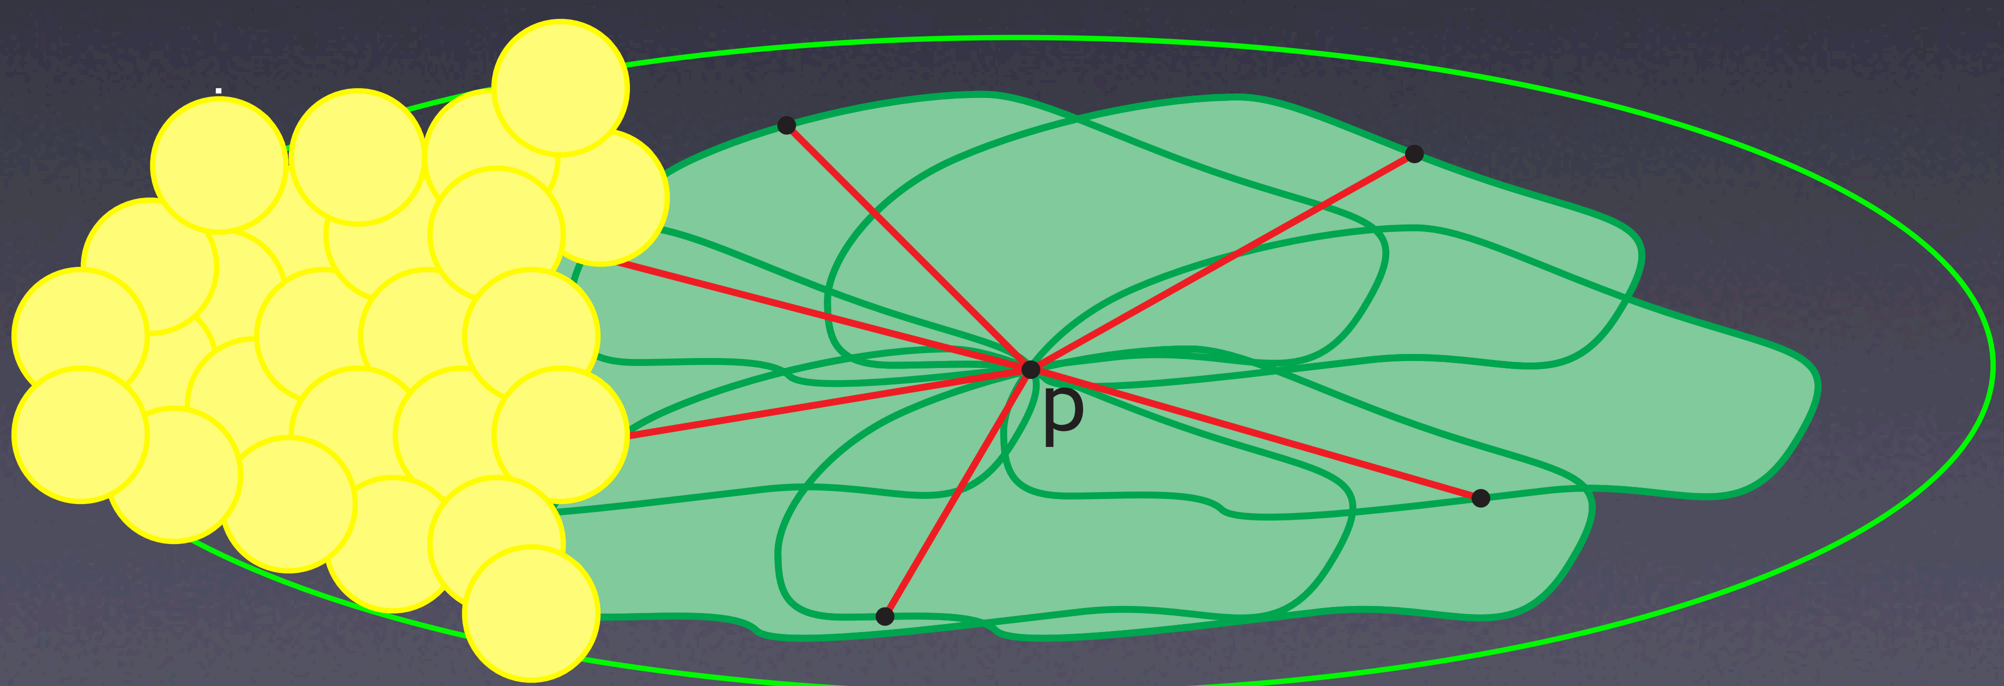

Lemma:

Any point far from A_3 contact curves has O(1) local Delaunay neighbors.

Near A_3 contact curves

Suppose p is distance x from an A_3 curve

- $x > \sqrt{\epsilon}$:
  $1-\kappa_1 r = \Theta(x^2) \Rightarrow O(1/x)$ local neighbors
- $x \leq \sqrt{\epsilon}$
  higher-order Taylor approximation $\Rightarrow$
  $O(1/\sqrt{\epsilon}) = O(n^{1/4})$ local neighbors

Now integrate over x...

Danger!

Theorem:

O(n log n) local Delaunay edges

However, points near A_3 curves might still have high degree!

What's left to do?

- Count *remote* Delaunay edges, which cross from one side of the surface to the other
  - Remote neighborhood is only $O(1)$ bigger than local neighborhood.
- Count *external* Delaunay edges, whose Delaunay balls are centered outside Σ.
  - Apply a conformal transformation to turn the surface inside out!

Main Result

Fix a generic surface Σ.

The Delaunay triangulation of any $(\epsilon, O(1))$-sample of Σ has complexity $O(n \log n)$.

Future work

What is the expected complexity of the Delaunay triangulation of a set of *random* points from a surface?

- Theorem: For cylinder, $\Theta(n \log n)$!
- Conjecture: For any generic surface, $\Theta(n)$
- Conjecture: For any ~~smooth~~ surface, ~~$\Theta(n \log n)$~~
  $\alpha\, n \log n + \Theta(n)$ for some small absolute constant α (independent of Σ)

Thanks for listening, and thanks to the organizers for a great workshop!