Three-Dimensional α Shapes

Herbert Edelsbrunner and Ernst P. Mücke

ACM Tran. Graph. 13(1), 1994

Presented by Matthew Bolitho

20 September 2005

Shape

Question

Given a set of points, how can we determine its “shape”?

Shape

Question

Given a set of points, how can we determine its "shape"?

- The convex hull

Shape

Question

Given a set of points, how can we determine its "shape"?

- The convex hull
- Something else?

Shape

Question

Given a set of points, how can we determine its "shape"?

- The convex hull
- Something else?
- The α-shape

α-Shapes

- A generalisation of the convex hull

α-Shapes

- A generalisation of the convex hull
- The α-shape is actually a family of shapes, parameterised by the α value

α-Shapes

- A generalisation of the convex hull
- The α-shape is actually a family of shapes, parameterised by the α value
  - As $\alpha \to \infty$, the α-shape is the convex hull
  - As $\alpha \to 0$, the α-shape is the point set S

α-Shapes

- A generalisation of the convex hull
- The α-shape is actually a family of shapes, parameterised by the α value
  - As $\alpha \rightarrow \infty$, the α-shape is the convex hull
  - As $\alpha \rightarrow 0$, the α-shape is the point set S
  - Other values for α give shapes somewhere in-between

α-Shapes

- A generalisation of the convex hull
- The α-shape is actually a family of shapes, parameterised by the α value
  - As $\alpha \rightarrow \infty$, the α-shape is the convex hull
  - As $\alpha \rightarrow 0$, the α-shape is the point set S
  - Other values for α give shapes somewhere in-between
- α-shape may be concave or disjoint

α-Shapes

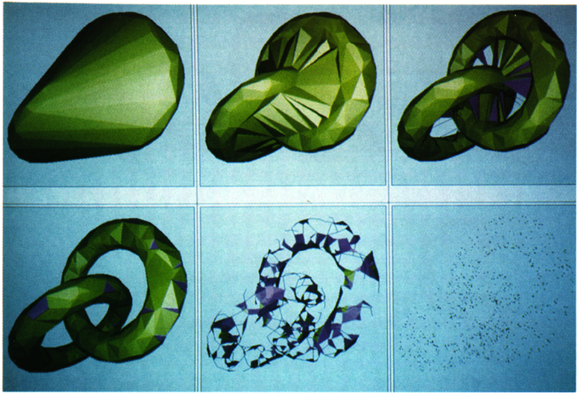

Figure from Edelsbrunner94

Algorithm

- Imagine a spherical scoop, with size α.

Algorithm

- Imagine a spherical scoop, with size α.
- P defines a set of points in space that the scoop cannot pass through.

Algorithm

- Imagine a spherical scoop, with size α.
- P defines a set of points in space that the scoop cannot pass through.
- The general idea is to carve out as much of space as possible:

Algorithm

- Imagine a spherical scoop, with size α.
- P defines a set of points in space that the scoop cannot pass through.
- The general idea is to carve out as much of space as possible:
    - There are places that the scoop is blocked
      i.e. $|p_i - p_j| < \alpha, i \neq j$
      At these positions there are at least d points in P touching the scoop.
    - An edge of the α-shape is defined by connecting those points.

Algorithm

- Imagine a spherical scoop, with size α.
- P defines a set of points in space that the scoop cannot pass through.
- The general idea is to carve out as much of space as possible:
  - There are places that the scoop is blocked
    i.e. $|p_i - p_j| < \alpha, i \neq j$
    At these positions there are at least d points in P touching the scoop.
  - An edge of the α-shape is defined by connecting those points.
  - Then boundary of the α-shape is a collection of these edges

Basic Definitions

Let $P \subset \mathbb{R}^d$ be a set of n points

Basic Definitions

Let $P \subset \mathbb{R}^d$ be a set of n points

Let $S_\alpha(P)$ be the α-shape of P for a given α value

General Position

For simplicity, we assume that P is in *general position* form:

General Position

For simplicity, we assume that P is in *general position* form:

- No 4 points in P may lie in the same line

General Position

For simplicity, we assume that P is in *general position* form:

- No 4 points in P may lie in the same line
- No 5 points in P may lie on the same sphere

General Position

For simplicity, we assume that P is in *general position* form:

- No 4 points in P may lie in the same line
- No 5 points in P may lie on the same sphere
- For a fixed α, the smallest sphere through any 2,3 or 4 points in P has a radius different to α

General Position

For simplicity, we assume that P is in *general position* form:

- No 4 points in P may lie in the same line
- No 5 points in P may lie on the same sphere
- For a fixed α, the smallest sphere through any 2,3 or 4 points in P has a radius different to α

This allows us to ignore special cases

α-balls

Let $B_\alpha(p)$ be an open half ball with radius α covering a space p for $0 \leq \alpha \leq \infty$, $p \subset \mathbb{R}^d$

α-balls

Let $B_\alpha(p)$ be an open half ball with radius α covering a space p for $0 \leq \alpha \leq \infty$, $p \subset \mathbb{R}^d$

- $B_0(p)$ is the point p
- $B_\infty(p)$ is the half-space p

α-balls

Let $B_\alpha(p)$ be an open half ball with radius α covering a space p for $0 \leq \alpha \leq \infty$, $p \subset \mathbb{R}^d$

- $B_0(p)$ is the point p
- $B_\infty(p)$ is the half-space p

$B_\alpha(p)$ is *empty* if $p \cap P = 0$

Simplices I

- An *n-simplex* is an n-dimensional analogue of a triangle:
  - The 0-simplex is a point
  - The 1-simplex is a line
  - The 2-simplex is a triangle
  - The 3-simplex is a tetrahedron
- A n-simplex has $n + 1$ vertices

Simplices II

Let $T \subset P$, and $|T| = k + 1 \leq d + 1$

The polytope $\triangle_T = conv(T)$ has dimension k and is therefore a k-simplex

Exposed Simplices

Let δp be the surface of $B_\alpha(p)$

A k-simplex $\triangle_T$ is said to be *exposed* if there is an empty $B_\alpha(p)$ where $T \in \delta p$

Exposed Simplices

Let δp be the surface of $B_\alpha(p)$

A k-simplex $\triangle_T$ is said to be *exposed* if there is an empty $B_\alpha(p)$ where $T \in \delta p$

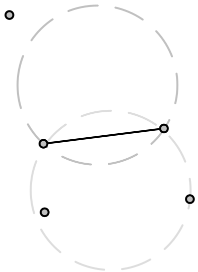

Building the α-shape

- $S_\alpha(P)$ is constructed from all exposed simplices:

$$\delta S_\alpha(P) = \{\triangle_T | T \subset P, |T| \leq d \text{ and } \triangle_T \text{ is exposed } \}$$

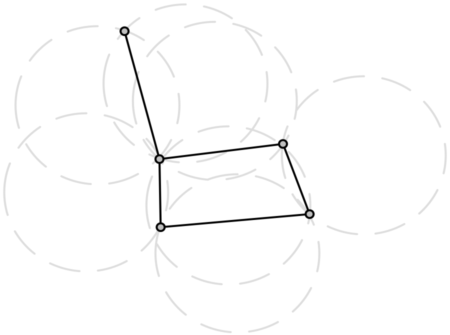

Properties

Observations

It is easy to show:

- $\lim_{\alpha \to \infty} S_\alpha(P)$ is the convex hull of P
- $\lim_{\alpha \to 0} S_\alpha(P)$ is the original set P

Properties

Observations

It is easy to show:

- $\lim_{\alpha \to \infty} S_\alpha(P)$ is the convex hull of P
- $\lim_{\alpha \to 0} S_\alpha(P)$ is the original set P

Claim

$S(\alpha)$ is a subset of the Delaunay triangulation of P

Delaunay Triangulation Equivelence

Let $DT(P)$ be a set of k-simplices $0 \leq k \leq d$ such that $\triangle_T = conv(T)$, $T \subset P$, $|T| = k + 1$

Delaunay Triangulation Equivelence

Let $DT(P)$ be a set of k-simplices $0 \leq k \leq d$ such that
$\triangle_T = conv(T),\ T \subset P,\ |T| = k + 1$

- For $k = d$, an α-ball $B_\alpha(p)$ coincides with the circumsphere of $\triangle_T$

Delaunay Triangulation Equivelence

Let $DT(P)$ be a set of k-simplices $0 \leq k \leq d$ such that $\triangle_T = conv(T)$, $T \subset P$, $|T| = k + 1$

- For $k = d$, an α-ball $B_\alpha(p)$ coincides with the circumsphere of $\triangle_T$
- By definition, this does not contain any other points from P, therefore $B_\alpha(p)$ is empty

Delaunay Triangulation Equivelence

Let $DT(P)$ be a set of k-simplices $0 \leq k \leq d$ such that $\triangle_T = conv(T)$, $T \subset P$, $|T| = k+1$

- For $k = d$, an α-ball $B_\alpha(p)$ coincides with the circumsphere of $\triangle_T$
- By definition, this does not contain any other points from P, therefore $B_\alpha(p)$ is empty
- Thus the simplices that form the edges of the d-simplex are *exposed*, and form the boundary for some α-shape.

α-Complexes

- In order to compute the α-shape, we use the α-complex

α-Complexes

- In order to compute the α-shape, we use the α-complex
- The α-complex is a simplicial complex that is a subset of $DT(P)$

α-Complexes

- In order to compute the α-shape, we use the α-complex
- The α-complex is a simplicial complex that is a subset of $DT(P)$
- The subset of $DT(P)$ is determined by α

Definitions

Let σ_T be the radius of the circumsphere of a simplex $\triangle_T$

Definitions

Let σ_T be the radius of the circumsphere of a simplex $\triangle_T$

Let μ_T be the center of the circumsphere of a simplex $\triangle_T$

Definitions

Let σ_T be the radius of the circumsphere of a simplex $\triangle_T$

Let μ_T be the center of the circumsphere of a simplex $\triangle_T$

- A simplex $\triangle_T$ from $DT(P)$ is in $C_\alpha(P)$ if either:
  - $\sigma_T < \alpha$ and the α-ball at μ_T is empty
  - $\triangle_T$ is the face of another $\triangle_T$ in $C_\alpha(P)$

Description

- Compute the Delaunay Triangulation

Description

- Compute the Delaunay Triangulation
- Generate the α-complex
  - The algorithm actually computes the α-complex for all α values.

Description

- Compute the Delaunay Triangulation
- Generate the α-complex
    - The algorithm actually computes the α-complex for all α values.

Description

- Compute the Delaunay Triangulation
- Generate the α-complex
  - The algorithm actually computes the α-complex for all α values.
- Extract the boundary of the α-shape from the α-complex

Description

- Compute the Delaunay Triangulation
- Generate the α-complex
  - The algorithm actually computes the α-complex for all α values.
- Extract the boundary of the α-shape from the α-complex
- Steps 1 and 2 can be precomputed for a given P

Complexity

- Delaunay Triangulation: $O(n \log n)$
- Generate α-complex: $O(m \log m)$
- Extract boundary of α-shape: $O(m)$ (could be better?)

General Position Assumption

- We assumed that P was in *general position* form

General Position Assumption

- We assumed that P was in *general position* form
- In reality, point sets often break these contraints

General Position Assumption

- We assumed that P was in *general position* form
- In reality, point sets often break these contraints
- Solution: *Simulation of simplicity* H. Edelsbrunner and E. P. Mcke, ACM Trans. Graph. 9(1) 1990

Properties of α-Shapes

- $S_\infty(P) = conv(P)$
- $S_0(P) = P$

Properties of α-Shapes

- $S_\infty(P) = conv(P)$
- $S_0(P) = P$
- $S_{\alpha_1}(P) \subset S_{\alpha_2}(P)$ if $\alpha_1 < \alpha_2$

Properties of α-Shapes

- $S_\infty(P) = conv(P)$
- $S_0(P) = P$
- $S_{\alpha_1}(P) \subset S_{\alpha_2}(P)$ if $\alpha_1 < \alpha_2$
- $S_\alpha(P) \subset DT(P)$

Advantages

- α-shape reconstructions can have arbitrary topology
- The α-shape interpolates the set P

Disadvantages

- The choice of α-value is non-intuative
- The reconstruction may not be water tight
- The reconstruction may be disjoint

Sampling Density

- Given a point set P, one can only choose a single α value.

Sampling Density

- Given a point set P, one can only choose a single α value.
  - The α-value is determined by the smallest feature in the reconstruction

Sampling Density

- Given a point set P, one can only choose a single α value.
  - The α-value is determined by the smallest feature in the reconstruction
  - Thus, the α-value determines the sampling density everywhere

Sampling Density

- Given a point set P, one can only choose a single α value.
  - The α-value is determined by the smallest feature in the reconstruction
  - Thus, the α-value determines the sampling density everywhere

P must be uniformly sampled, at a resolution constrained by the locally highest resoultion desired.

Sampling Density

- Given a point set P, one can only choose a single α value.
  - The α-value is determined by the smallest feature in the reconstruction
  - Thus, the α-value determines the sampling density everywhere

P must be uniformly sampled, at a resolution constrained by the locally highest resoultion desired.

- What about blending several α-shapes together?
- What about defining $\alpha(q)$ for $q \in \mathbb{R}^d$ such that $\alpha(q)$ is proportional to sampling density near q