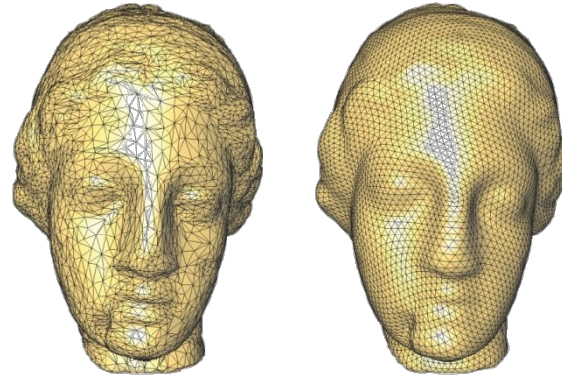
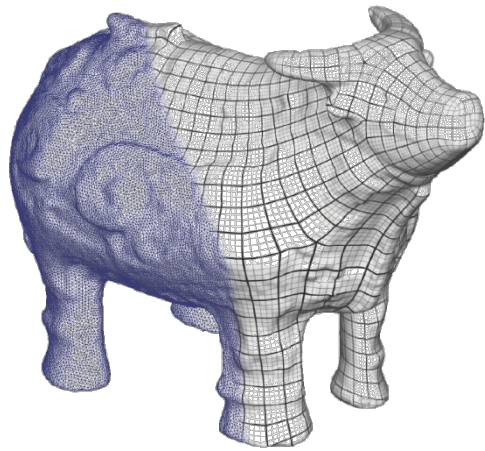
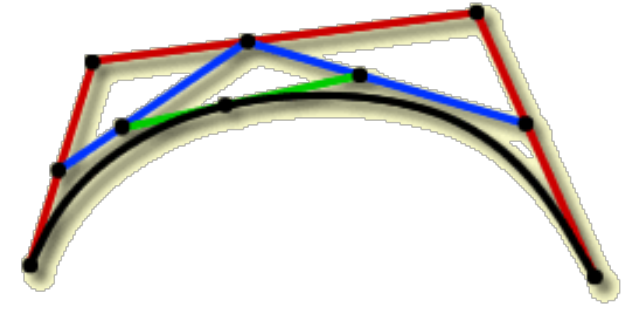


CS348a: Geometric Modeling and Processing



Leonidas Guibas
Computer Science Department
Stanford University



**Last Time:
Overview of Acquired Shapes
& Geometry Processing**

From Point Clouds to Surfaces



physical
model



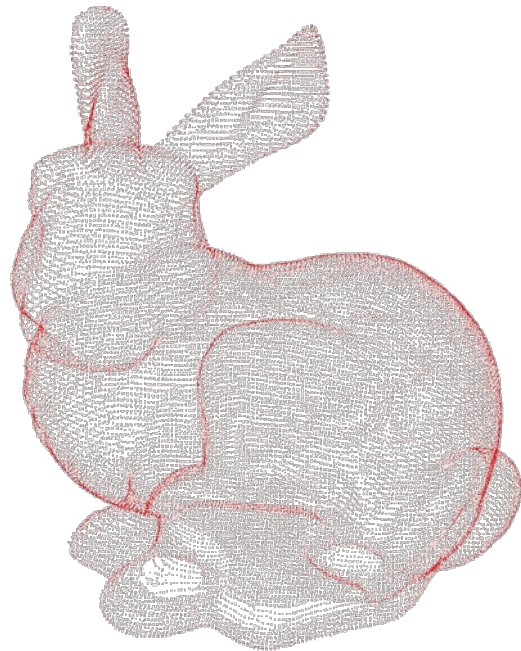
acquired
point cloud



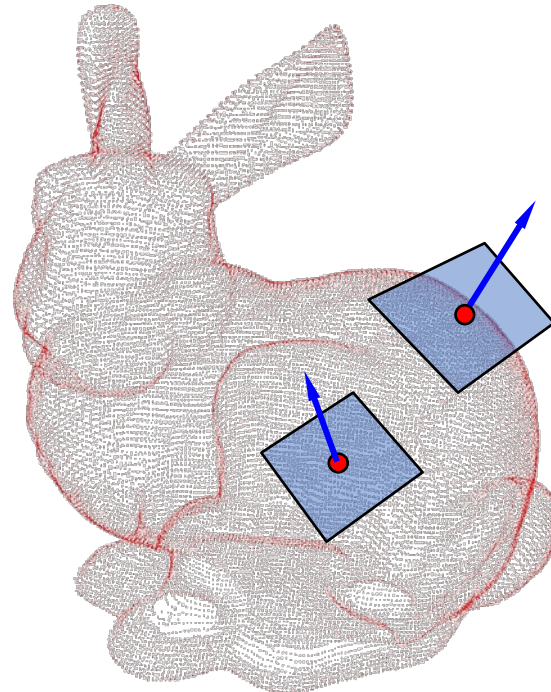
Reconstructed mesh or CAD
3D model

Point Clouds

- Simplest representation: **only points**, **no** connectivity
- Collection of (x,y,z) coordinates, possibly with normals

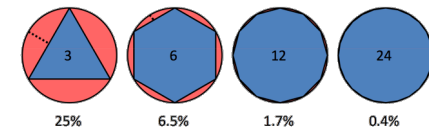
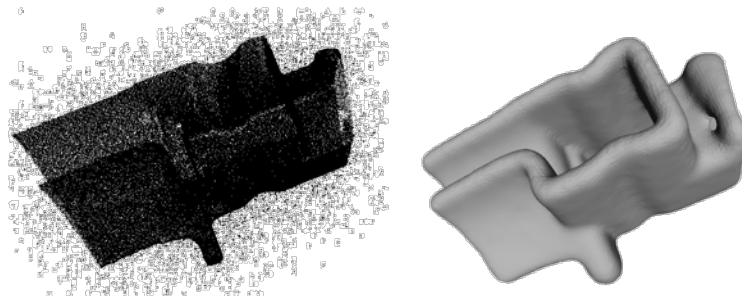


Stanford bunny



Point Clouds

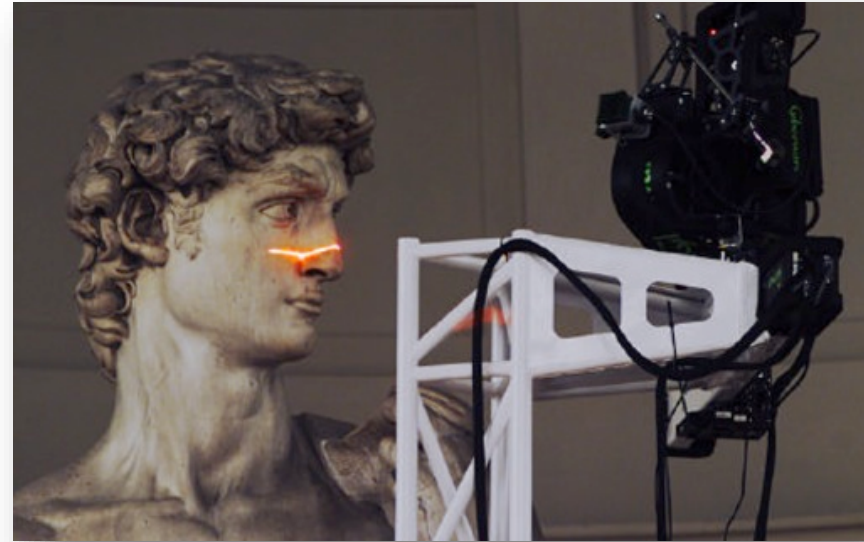
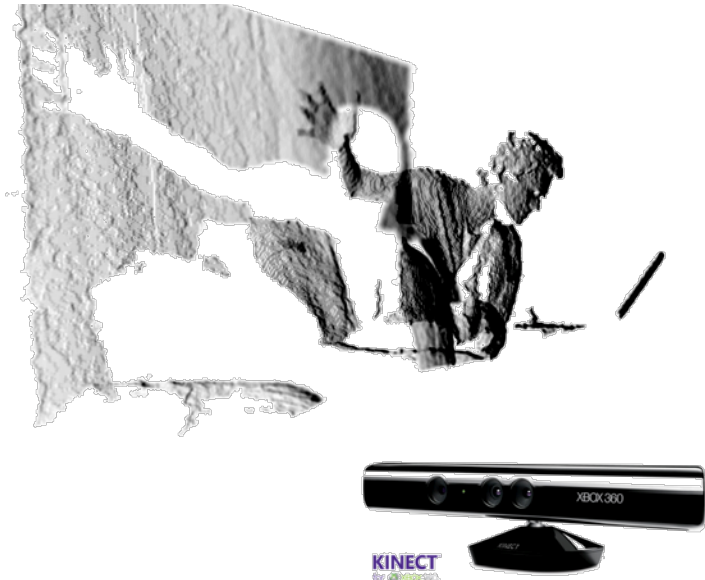
- Simplest representation: **only points**, no connectivity
- Collection of (x,y,z) coordinates, possibly with normals
- Points with orientation are called **surfels**
- Several limitations:
 - **no** simplification or subdivision
 - **no** direct smooth rendering
 - **no** topological information
 - weak approximation power: $O(h)$ for point clouds
 - noise and outliers



vs. $O(h^2)$ for meshes

3D Acquisition: Surface Scanning

- Typically, that's the only thing that's available
 - Nearly all 3D scanning devices produce point clouds



Surface Scanning Types

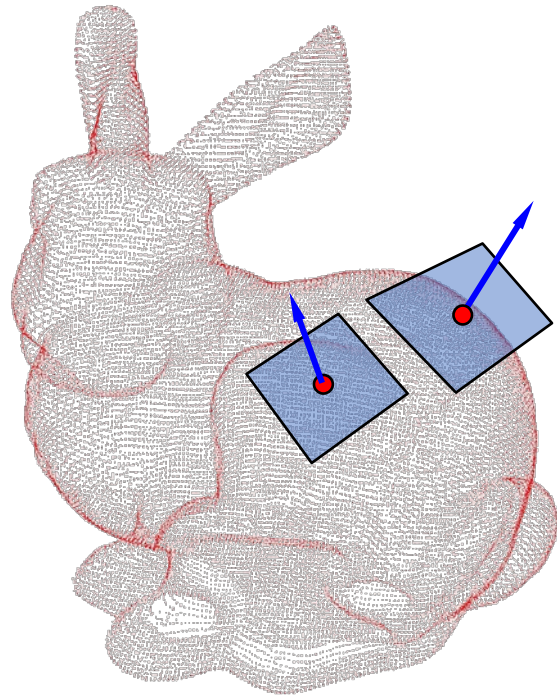
Major types of 3D scanners

- **Range (emission-based) scanners**
 - Time-of-flight laser scanner
 - Phase-based laser scanner
- **Triangulation**
 - Laser line sweep
 - Structured light
- **Stereo / computer vision**
 - Passive stereo
 - Active stereo / space-time stereo

Normal Estimation and Outlier Removal

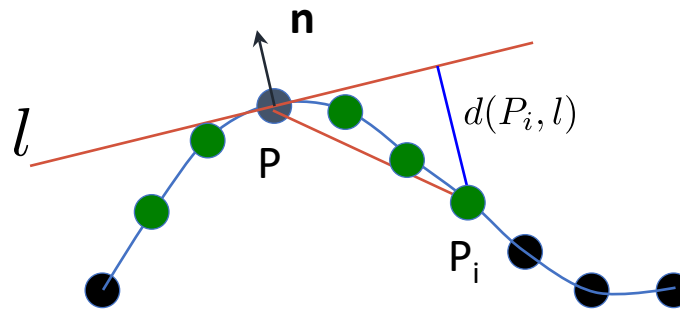
Fundamental problems in point cloud processing.

Although seemingly very different, can be solved with the same general approach – look at the “shape of neighborhoods” ...



Normal Estimation via Covariance

Assume we have a clean sampling of the surface. OK, start with a curve.



The normal \mathbf{n} must be an eigenvector of the matrix:

$$C\mathbf{n} = \lambda\mathbf{n} \quad C = \sum_{i=1}^k (p_i - P)(p_i - P)^T$$

So, \mathbf{n}_{opt} must be the eigenvector corresponding to the **smallest eigenvalue** of C .

Normal Estimation – Neighborhood Size



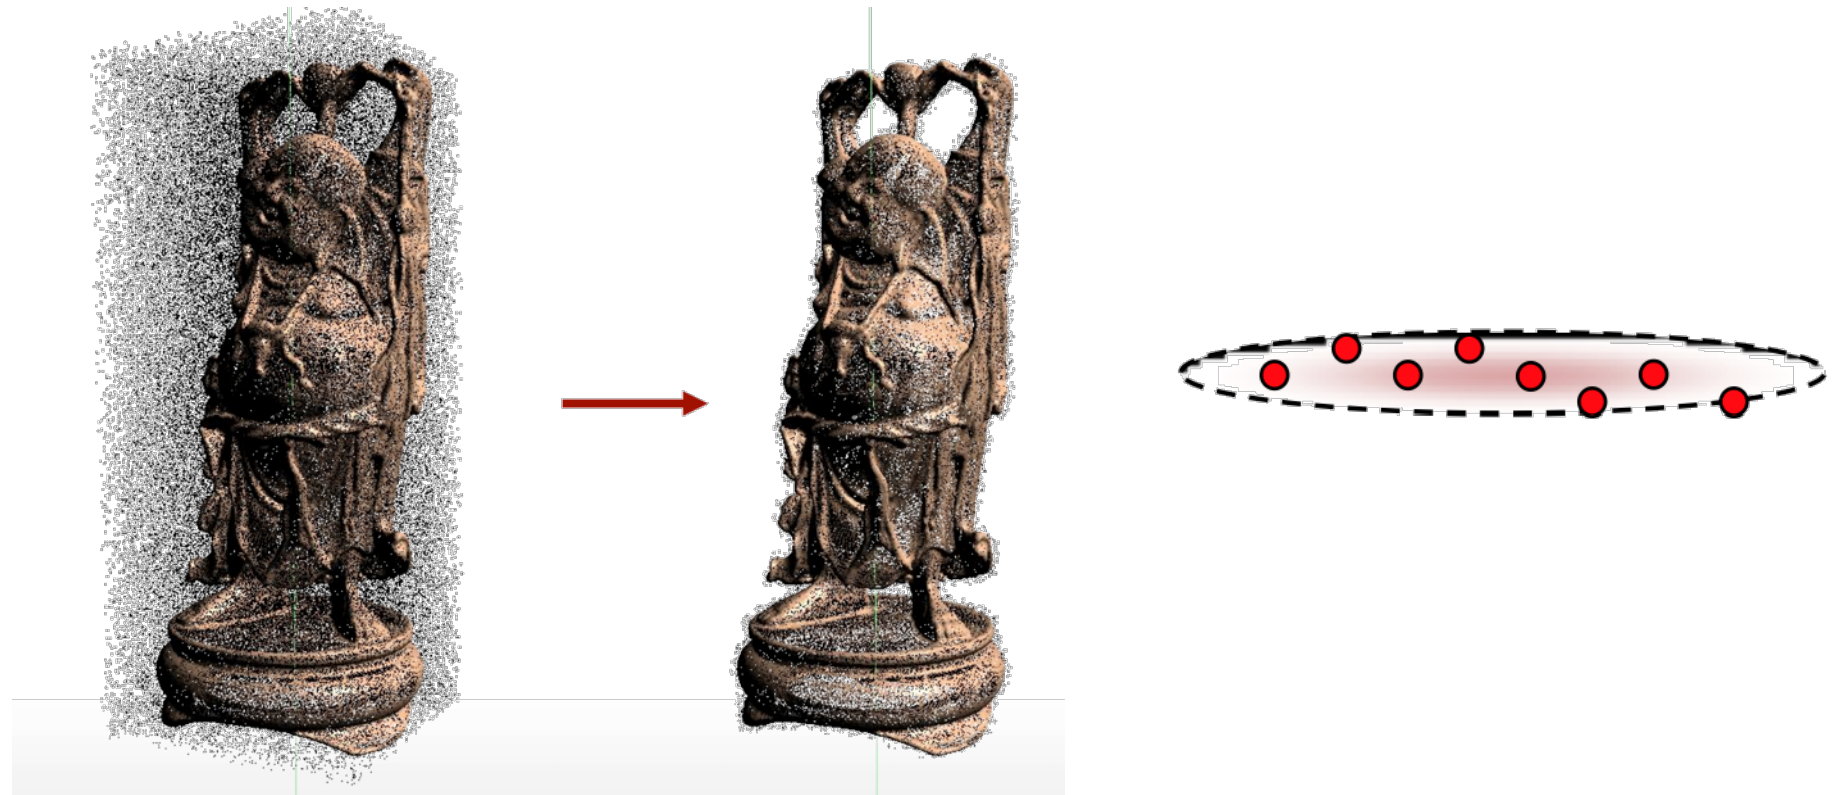
1x noise



2x noise

source: Mitra et al. '04

Outlier Removal



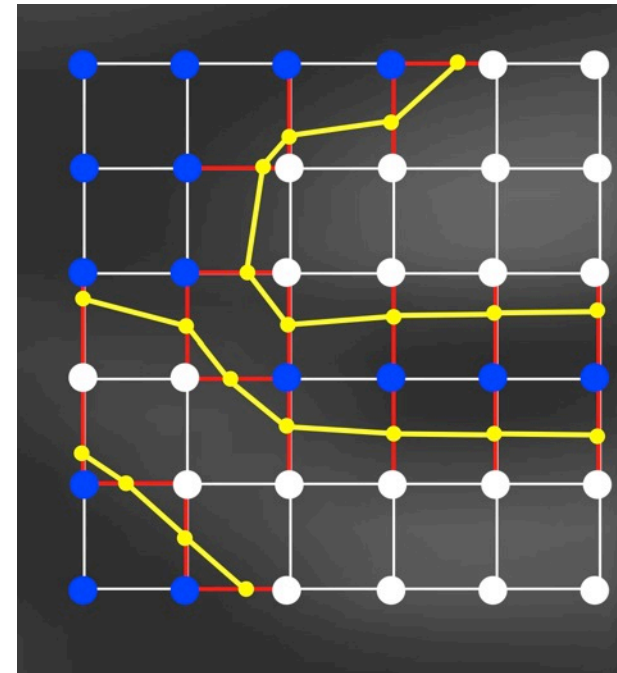
Goal: remove points that do not lie close to a surface.

Marching Squares (2D) : Curve Extraction

Given a function: $f(x)$

- $f(\mathbf{x}) < 0$ inside
- $f(\mathbf{x}) > 0$ outside

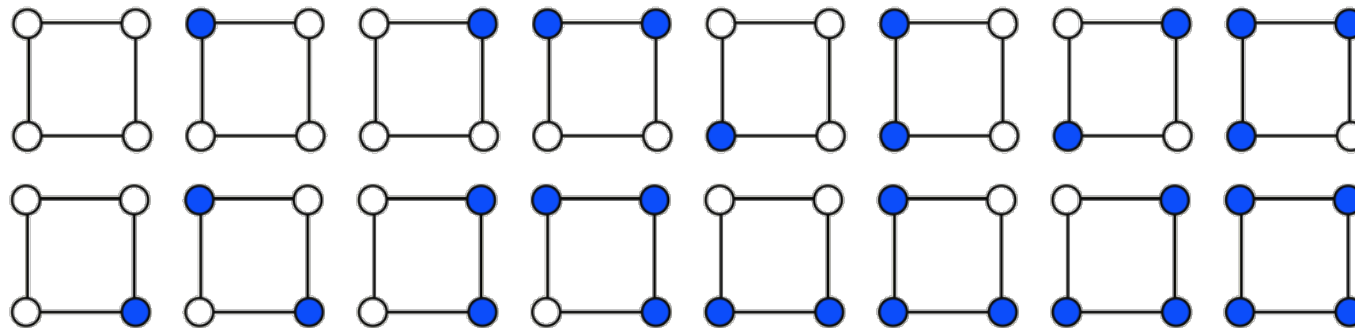
1. Discretize space.
2. Evaluate $f(x)$ on a grid.
3. Classify grid points (+/-)
4. Classify grid edges
5. Compute intersections
6. Connect intersections



Marching Squares (2D)

Connecting the intersections:

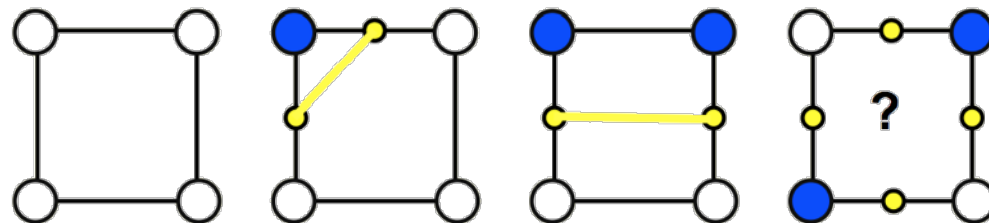
- Grand principle: treat each cell separately!
- Enumerate all possible inside/outside combinations.
- Group those leading to the same intersections



Marching Squares (2D)

Connecting the intersections:

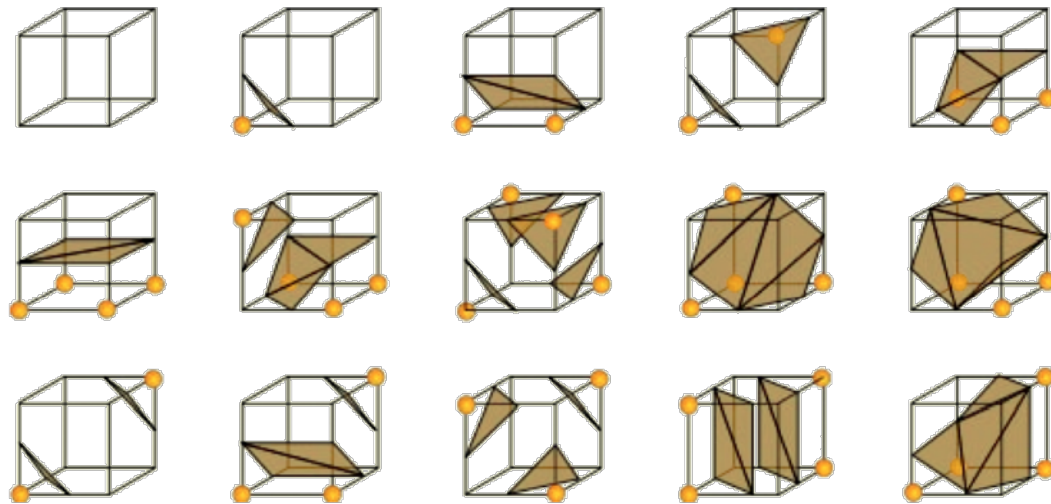
- Grand principle: treat each cell separately!
- Enumerate all possible inside/outside combinations.
- Group those leading to the same intersections.
- Group equivalent after rotation.
- Connect intersections



Marching Cubes (3D)

Same basic machinery applies to 3D.
cells become **cubes** (voxels)
lines become **triangles**

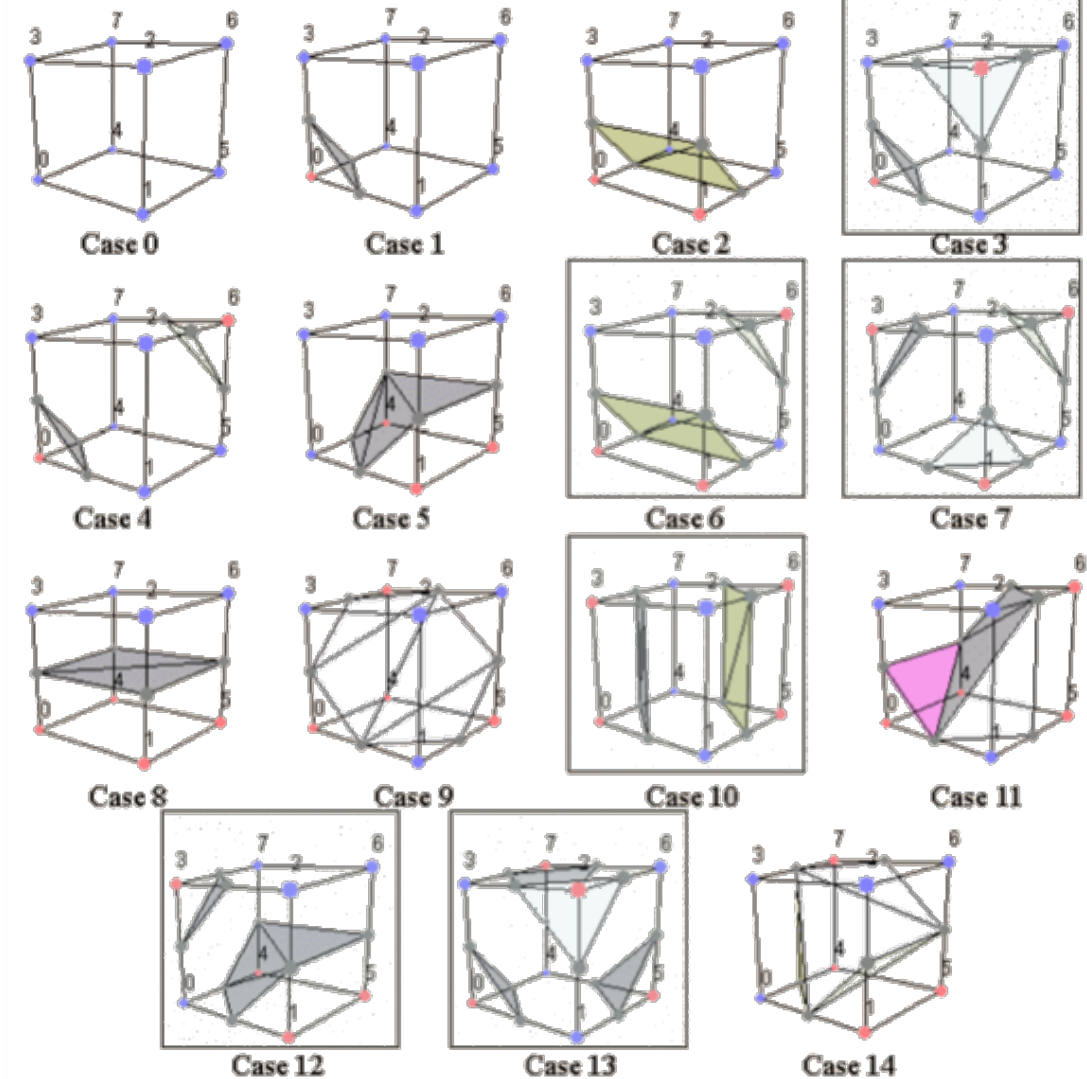
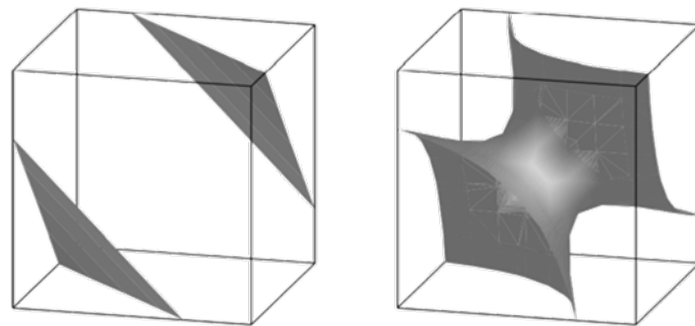
- 256 different cases
- 15 after symmetries



Marching Cubes (3D): Surface Extraction

Same basic machinery applies to 3D.
cells become **cubes** (voxels)
lines become **triangles**

- 256 different cases
- 15 after symmetries
- 6 ambiguous cases (in boxes)



Marching Cubes (3D)

Main Strengths:

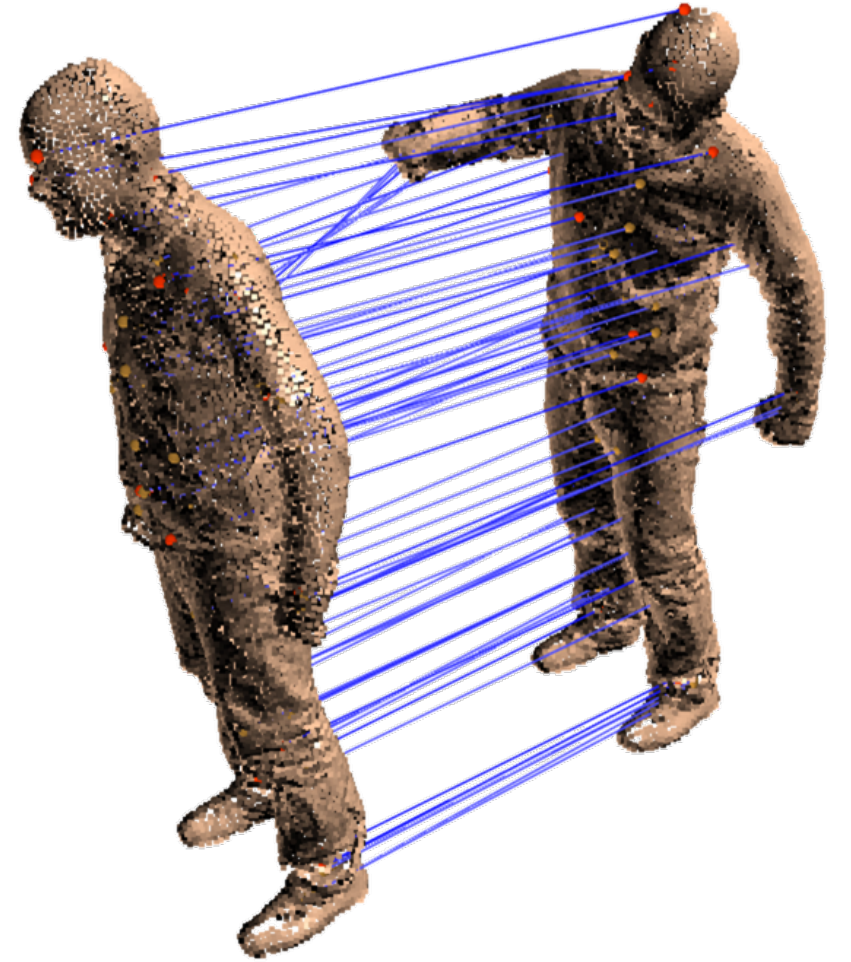
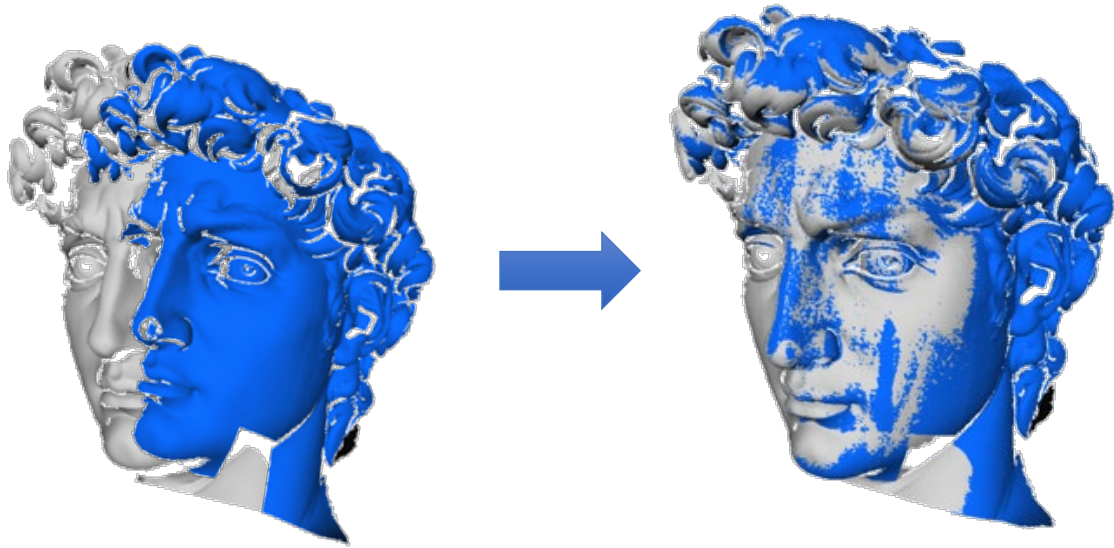
- Very multi-purpose.
- Extremely fast and parallelizable.
- Relatively simple to implement.
- Virtually parameter-free.

Main Weaknesses:

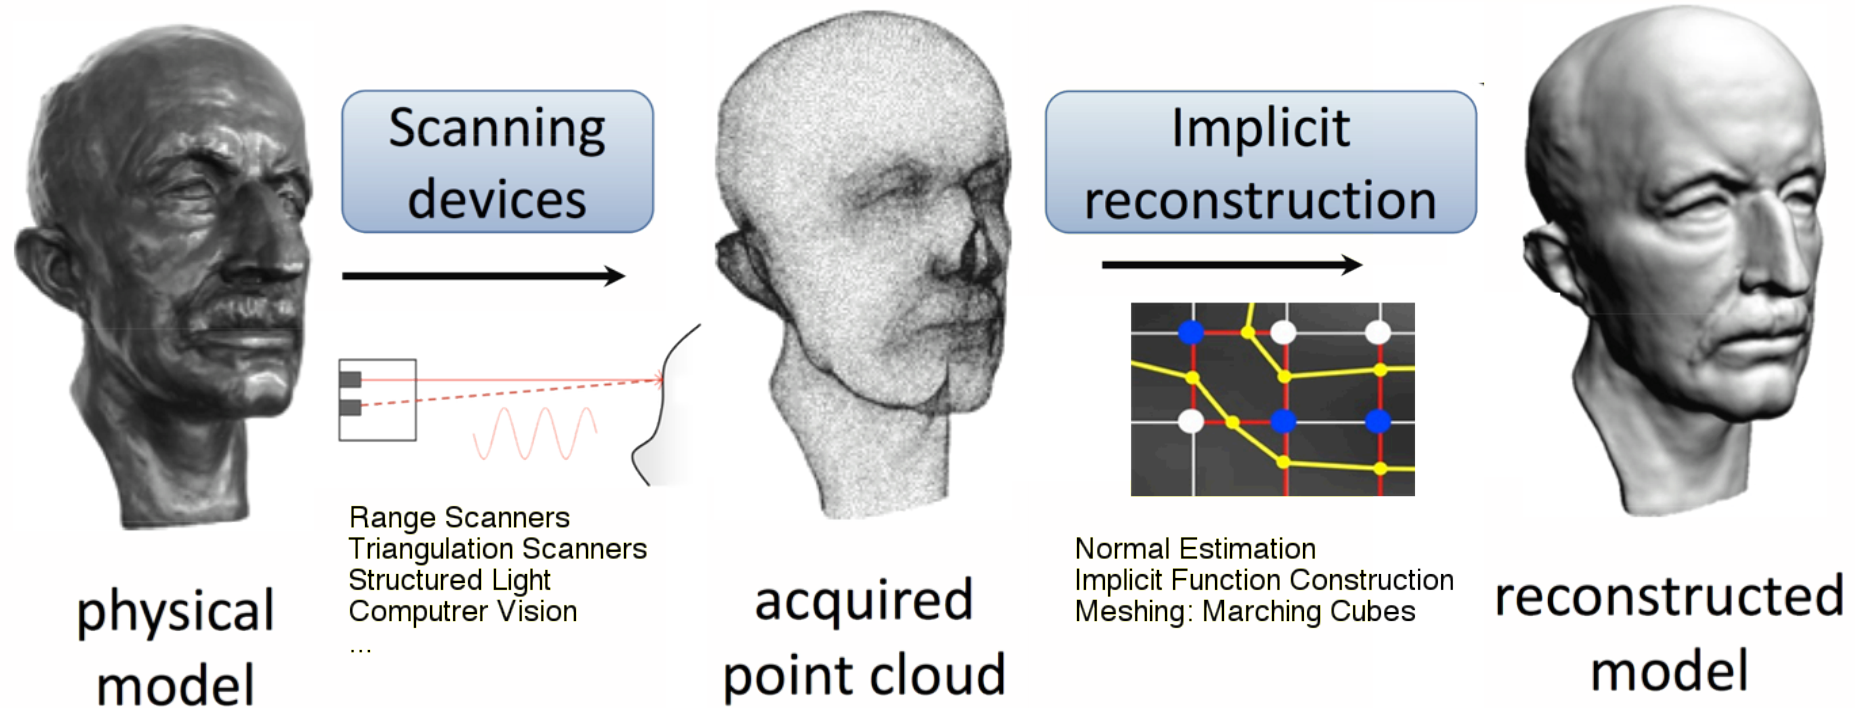
- Can create badly shaped (skinny) triangles.
- Basic versions do not provide topological guarantees.
- Many special cases (implemented as big lookup tables).
- No sharp features.

Today:
Shape Registration, Matching,
Correspondences

Alignment, Registration and Matching



3D Point Cloud Processing



3D Acquisition Pipeline

3D Point Cloud Processing



This lecture

Registration Pipeline

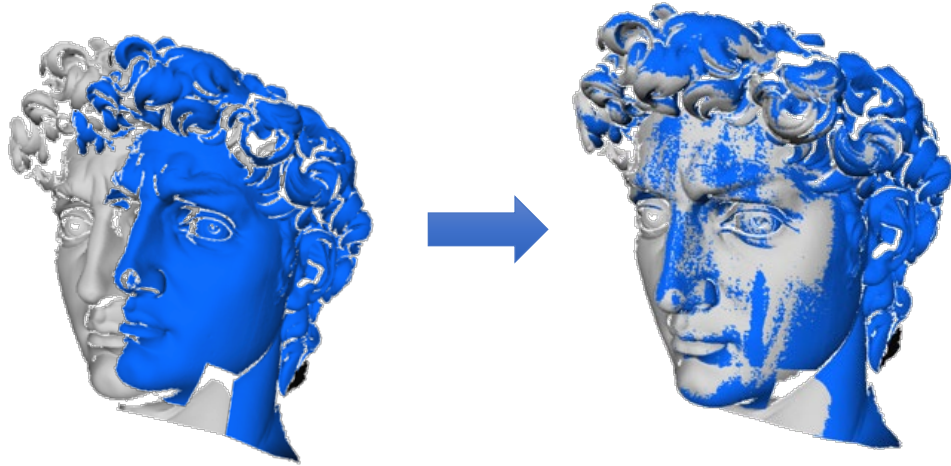


Source: Rusinkiewicz et al.

Steps:

1. Initial registration
2. Pairwise refinement
3. Global relaxation to distribute error
4. Generation of surface

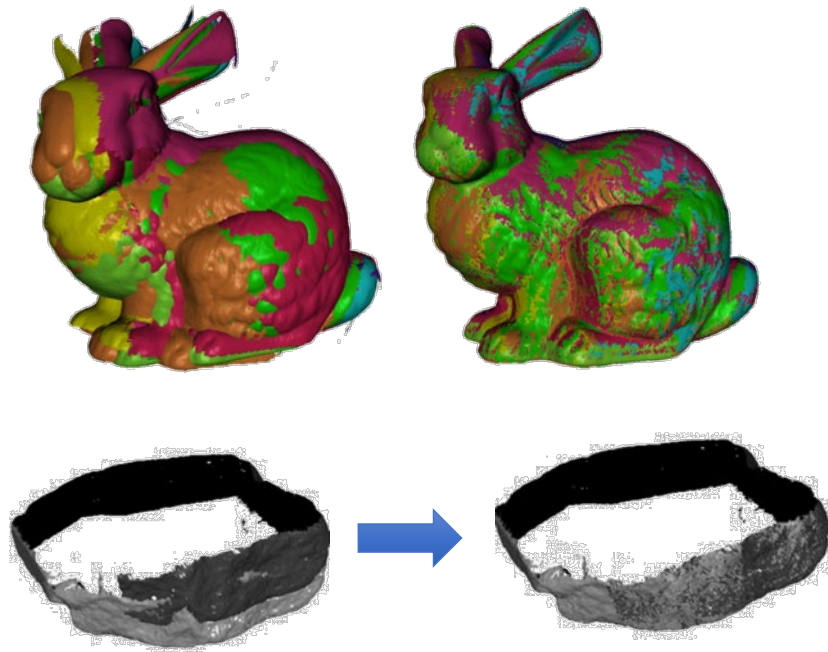
Registration Pipeline



Steps:

1. Initial registration
2. **Pairwise registration**
3. Global relaxation to distribute error
4. Generation of surface

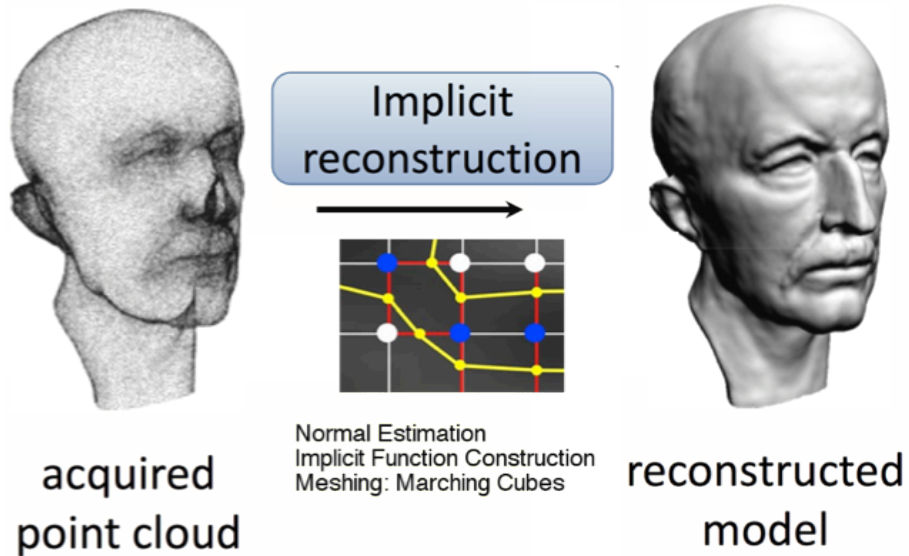
Registration Pipeline



Steps:

1. Initial registration
2. Pairwise registration
3. Global relaxation to distribute error
4. Generation of surface

Registration Pipeline



Steps:

1. Initial registration
2. Pairwise registration
3. Global relaxation to distribute error
4. **Generation of surface**

Fundamental Registration Problem



Given two shapes with partially overlapping geometry,
find an alignment between them

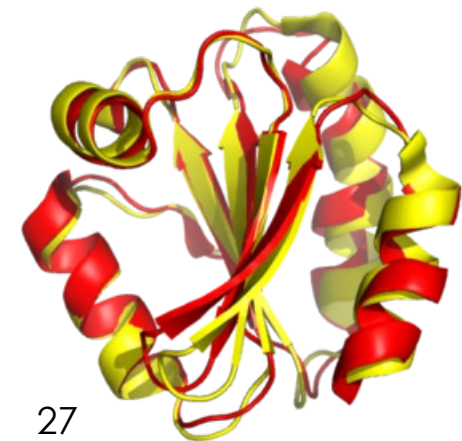
Measuring Success: Shape Distances

Given two shapes A and B , we are interested in defining a distance or (dis-)similarity measure under a transform T

$$\min_T \delta(A, T(B))$$

Such measures are crucial in shape similarity search, shape classification, etc.

As another example, shape registration and matching is very important in modern structural biology



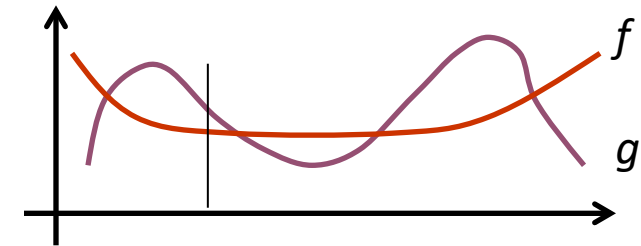
27

Issues about Distance Metrics Between Shapes

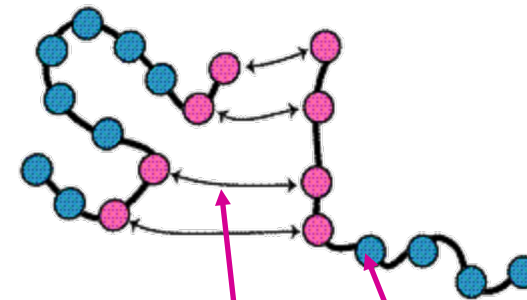
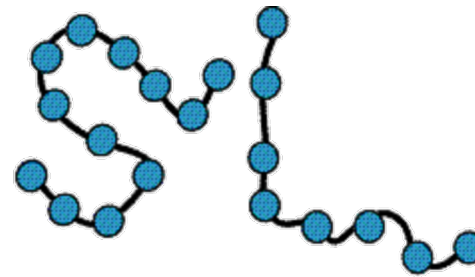
- We are all familiar with function norms (L_2 , etc.). The **common parametrization** establishes **correspondences**. We don't have that for structures or shapes.
- Partial matches need to be considered -- notion of **support** σ for the match.
- What group of **aligning transforms** is to be considered?
- Is the resulting distance a **metric**?

Not for partial matches

$$\delta(A, C) \leq \delta(A, B) + \delta(B, C)$$



$$\|f(x) - g(x)\|_2$$

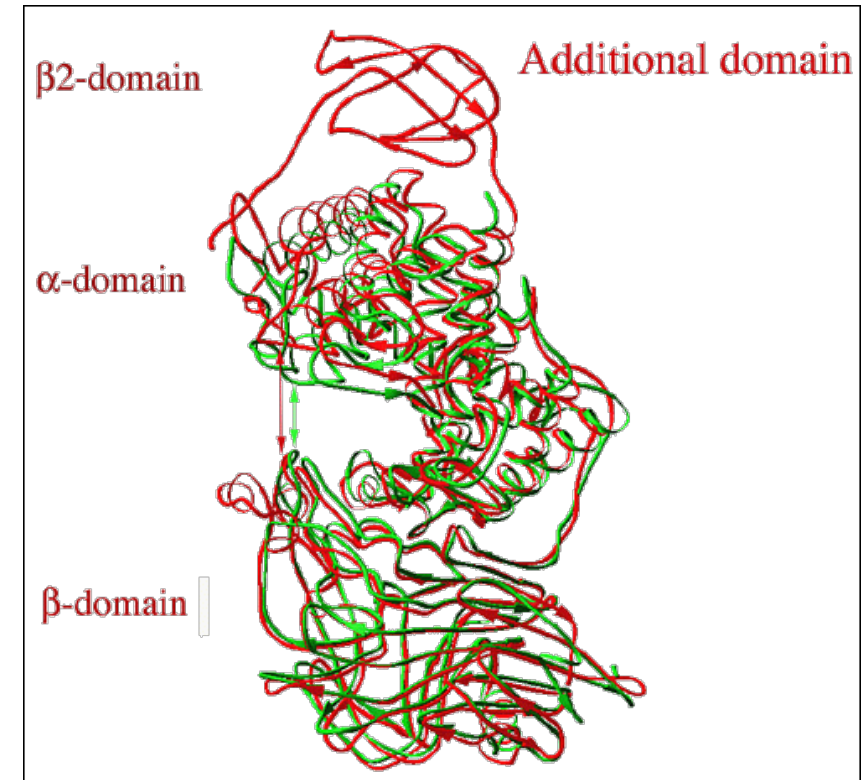
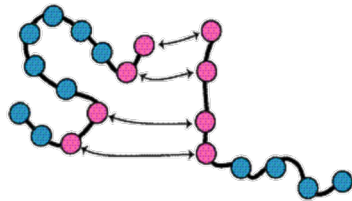


$$\delta(\sigma(A), \sigma(B))$$



Simultaneous Estimation

- We are given two shapes A and B , each in its own coordinate system
- We must establish **correspondences** between certain parts (the **alignment supports**) of A and B
- We must find an optimal **transform** that best **aligns** the supports of A and B
- We must **score** this choice of supports and transform to produce a **distance measure** δ



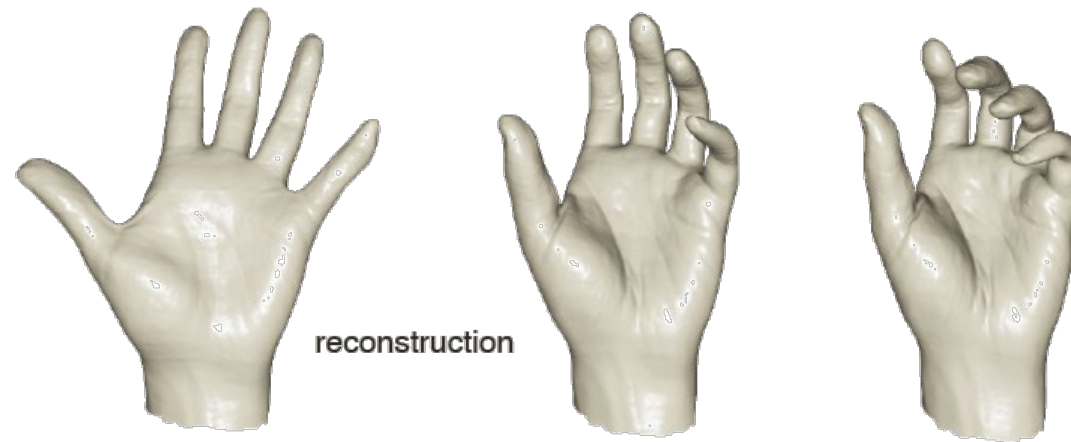
In computing the score, how do we

1. aggregate distances?
2. trade-off larger supports for larger aggregate distance?

Transform Degrees of Freedom

- Transform estimation

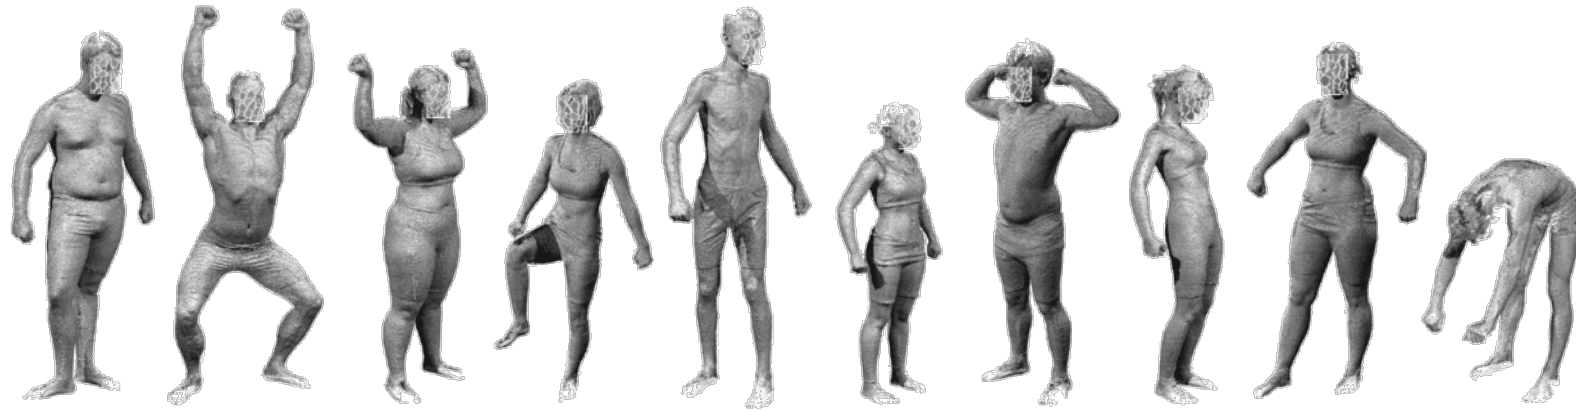
- A rigid motion has 6 degrees of freedom (3 for translation and 3 for rotation)
- We typically estimate the motion using many more pairs of corresponding points, so the problem is **overdetermined** (which is good, given noise, outliers, etc – use least squares approaches)
- More general transforms require more degrees of freedom. When shape deformations are allowed, the degrees of freedom can grow very rapidly



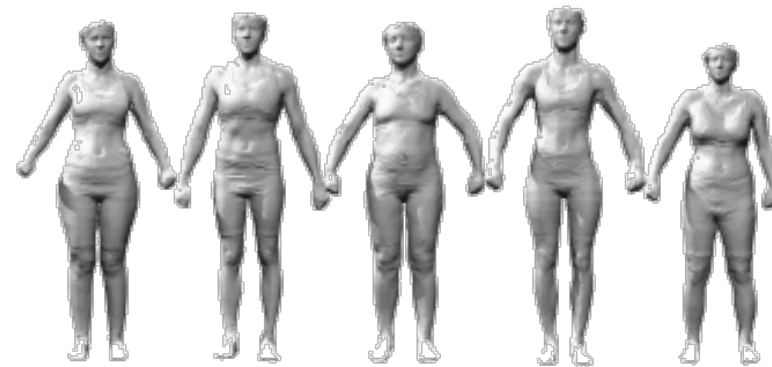
Other Applications of Alignments

- Manufacturing / Quality Control:
One shape is a **model** and the other is a **scan** of a product. Useful for finding defects.
- Medicine:
Finding correspondences between 3D MRI scans of the same person to diagnose or monitor disease.
- Animation Reconstruction & 3D Video.
- Statistical Shape Analysis:
Building models for a collection of shapes.

Applications – Statistical Shape Analysis



- Scan many people. Learn a deformation model (e.g. PCA).
- Find the principal variation modes; create new random instances.
- Requires alignment.



female, 1.6 m, 65kg

Registration Method Taxonomy

Local vs. Global

refinement (e.g. ICP) | alignment (search)

Rigid vs. Deformable

rotation, translation | general deformation

Pair vs. Collection

two shapes | multiple shapes

Registration Method Taxonomy

Today

Local vs. Global

refinement (e.g. ICP) alignment (search)

Rigid vs. Deformable

rotation, translation general deformation

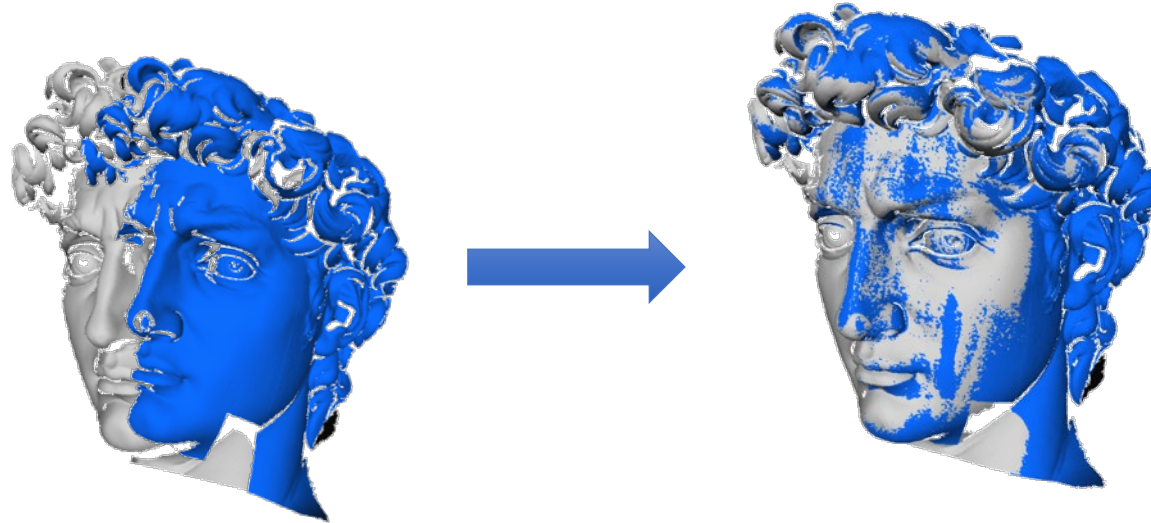
Pair vs. Collection

two shapes multiple shapes

Local Rigid Alignment

Local Rigid Alignment

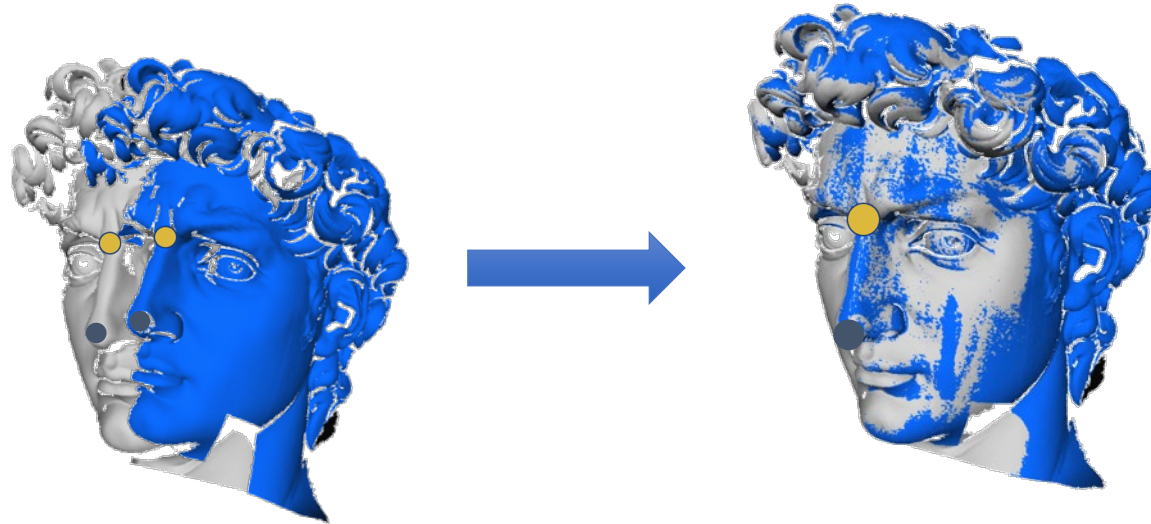
- Simplest instance of the registration problem



Given two shapes that are **approximately aligned** (e.g. by a human, or via prior knowledge) we want to find the **optimal rigid transformation** that brings them into correspondence.

Local Rigid Alignment

- What does it mean for an alignment to be **good**?



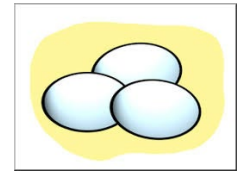
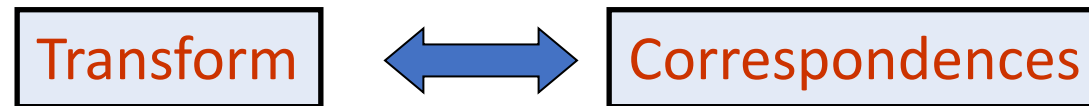
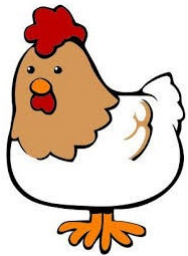
Intuition: we want “corresponding points” to be close after transformation.

Problems

1. We don't know what points correspond.
2. We don't know the optimal alignment.

How to Get Correspondences?

A *chicken-and-egg problem*: if we knew the optimal aligning transform, then we could get correspondences by **proximity** (possibly with the aid of some global adjustment, e.g., dynamic programming)



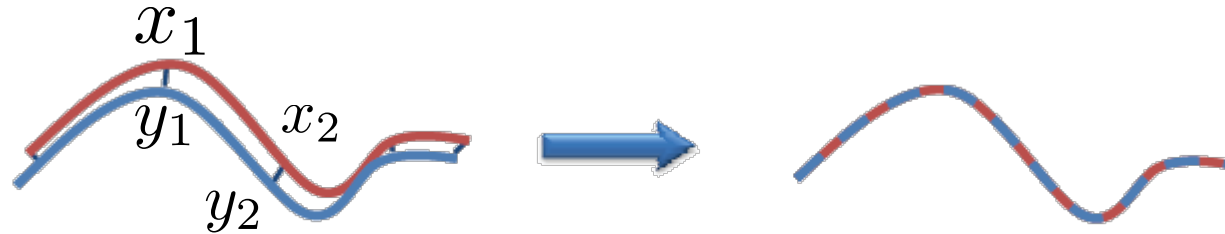
Guess one, estimate the other, and *iterate!*

EM like

- Correspondences from proximity (**Iterated Closest Pair**)
- Correspondences from local shape descriptors (**Shape Features**)
- Transform from voting schemes (**Geometric Hashing**)
- Combinations

Iterative Closest Point (ICP)

- Approach: iterate between finding correspondences and finding the transformation:



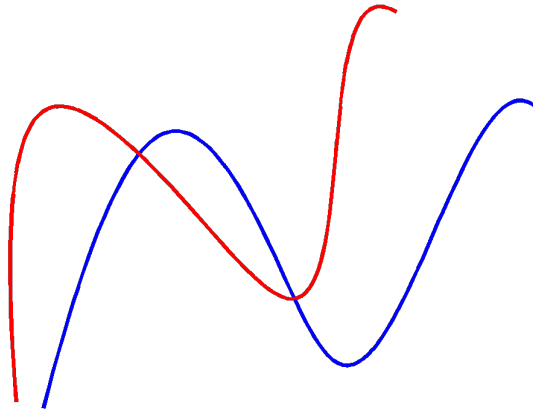
Given a pair of shapes, X and Y , iterate:

1. For each $x_i \in X$ find **nearest** neighbor $y_i \in Y$.
2. Find rigid motion \mathbf{R}, t minimizing:

$$\sum_{i=1}^N \|\mathbf{R}x_i + t - y_i\|_2^2$$

Iterative Closest Point (ICP)

- Approach: iterate between finding correspondences and finding the transformation:



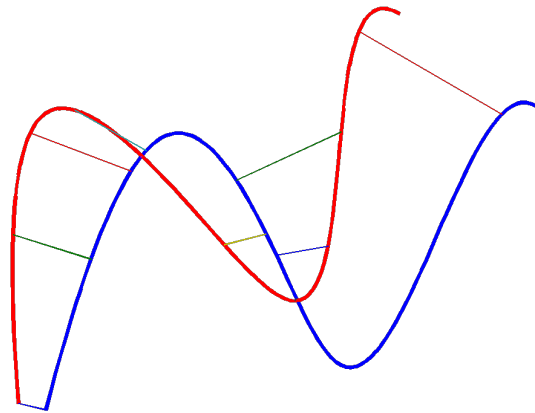
Given a pair of shapes, X and Y , iterate:

1. For each $x_i \in X$ find **nearest** neighbor $y_i \in Y$.

2. Find rigid motion \mathbf{R} , t minimizing:
$$\sum_{i=1}^N \|\mathbf{R}x_i + t - y_i\|_2^2$$

Iterative Closest Point

- Approach: iterate between finding correspondences and finding the transformation:



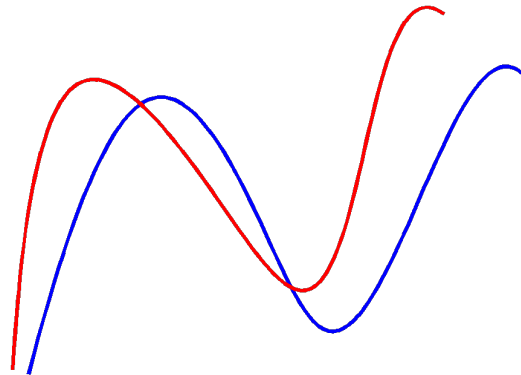
Given a pair of shapes, X and Y , iterate:

1. For each $x_i \in X$ find **nearest** neighbor $y_i \in Y$.

2. Find rigid motion \mathbf{R} , t minimizing:
$$\sum_{i=1}^N \|\mathbf{R}x_i + t - y_i\|_2^2$$

Iterative Closest Point

- Approach: iterate between finding correspondences and finding the transformation:



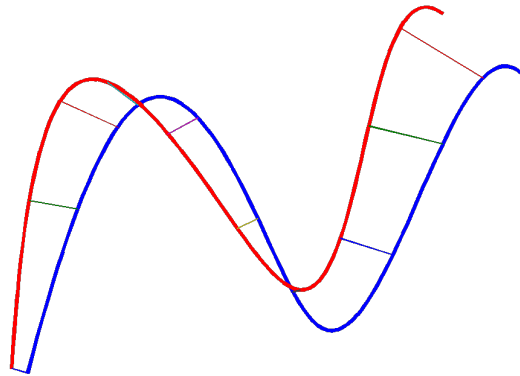
Given a pair of shapes, X and Y , iterate:

1. For each $x_i \in X$ find **nearest** neighbor $y_i \in Y$.

2. Find rigid motion \mathbf{R} , t minimizing:
$$\sum_{i=1}^N \|\mathbf{R}x_i + t - y_i\|_2^2$$

Iterative Closest Point

- Approach: iterate between finding correspondences and finding the transformation:



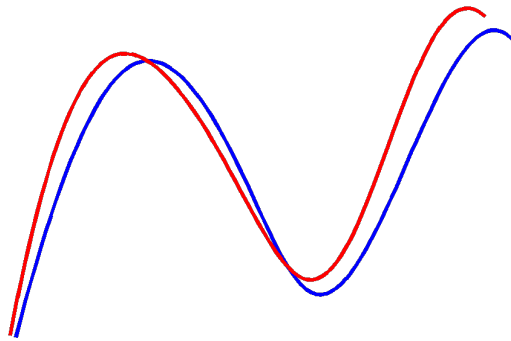
Given a pair of shapes, X and Y , iterate:

1. For each $x_i \in X$ find **nearest** neighbor $y_i \in Y$.

2. Find rigid motion \mathbf{R} , t minimizing:
$$\sum_{i=1}^N \|\mathbf{R}x_i + t - y_i\|_2^2$$

Iterative Closest Point

- Approach: iterate between finding correspondences and finding the transformation:



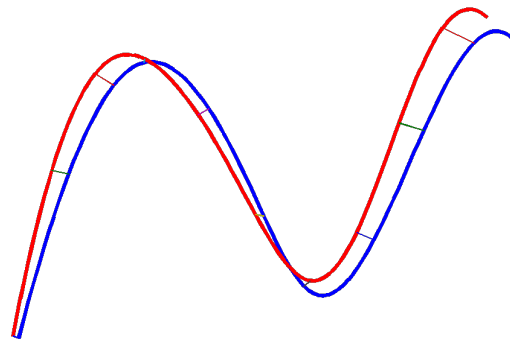
Given a pair of shapes, X and Y , iterate:

1. For each $x_i \in X$ find **nearest** neighbor $y_i \in Y$.

2. Find rigid motion \mathbf{R} , t minimizing:
$$\sum_{i=1}^N \|\mathbf{R}x_i + t - y_i\|_2^2$$

Iterative Closest Point

- Approach: iterate between finding correspondences and finding the transformation:



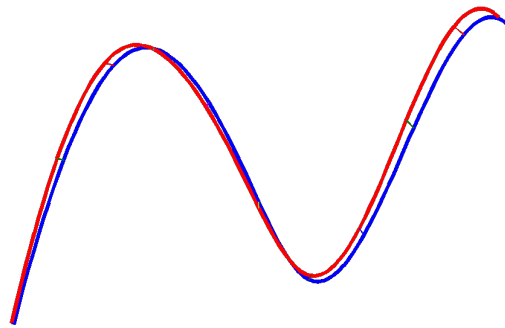
Given a pair of shapes, X and Y , iterate:

1. For each $x_i \in X$ find **nearest** neighbor $y_i \in Y$.

2. Find rigid motion \mathbf{R} , t minimizing:
$$\sum_{i=1}^N \|\mathbf{R}x_i + t - y_i\|_2^2$$

Iterative Closest Point

- Approach: iterate between finding correspondences and finding the transformation:



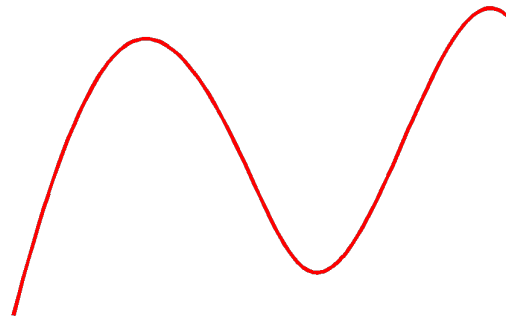
Given a pair of shapes, X and Y , iterate:

1. For each $x_i \in X$ find **nearest** neighbor $y_i \in Y$.

2. Find rigid motion \mathbf{R} , t minimizing:
$$\sum_{i=1}^N \|\mathbf{R}x_i + t - y_i\|_2^2$$

Iterative Closest Point

- Approach: iterate between finding correspondences and finding the transformation:



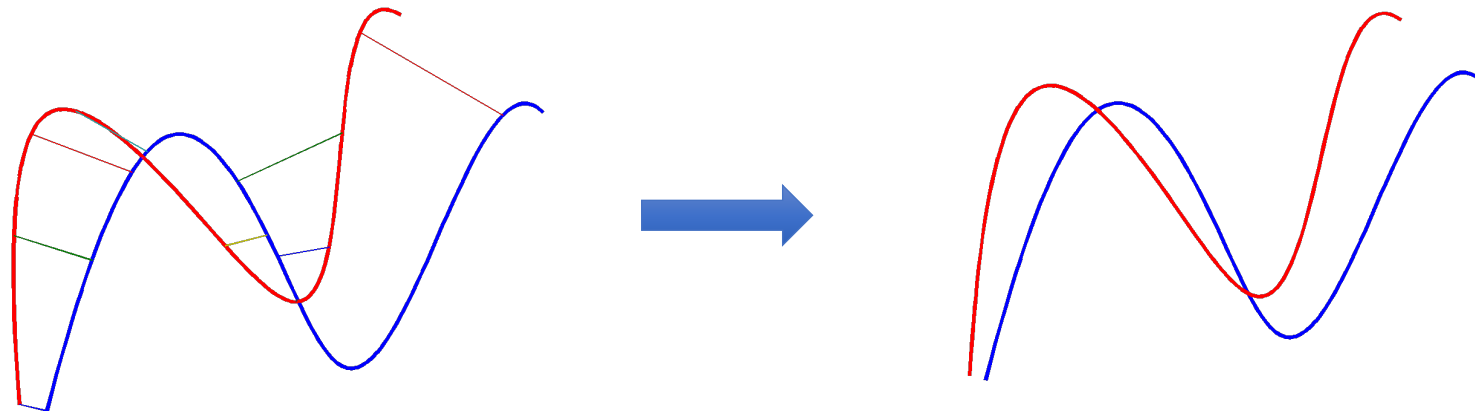
Given a pair of shapes, X and Y , iterate:

1. For each $x_i \in X$ find **nearest** neighbor $y_i \in Y$.

2. Find rigid motion \mathbf{R} , t minimizing:
$$\sum_{i=1}^N \|\mathbf{R}x_i + t - y_i\|_2^2$$

Iterative Closest Point

- Requires two key computations:
 1. Computing nearest neighbors
 2. Computing the optimal transformation



ICP: Nearest Neighbor Computation

Closest points

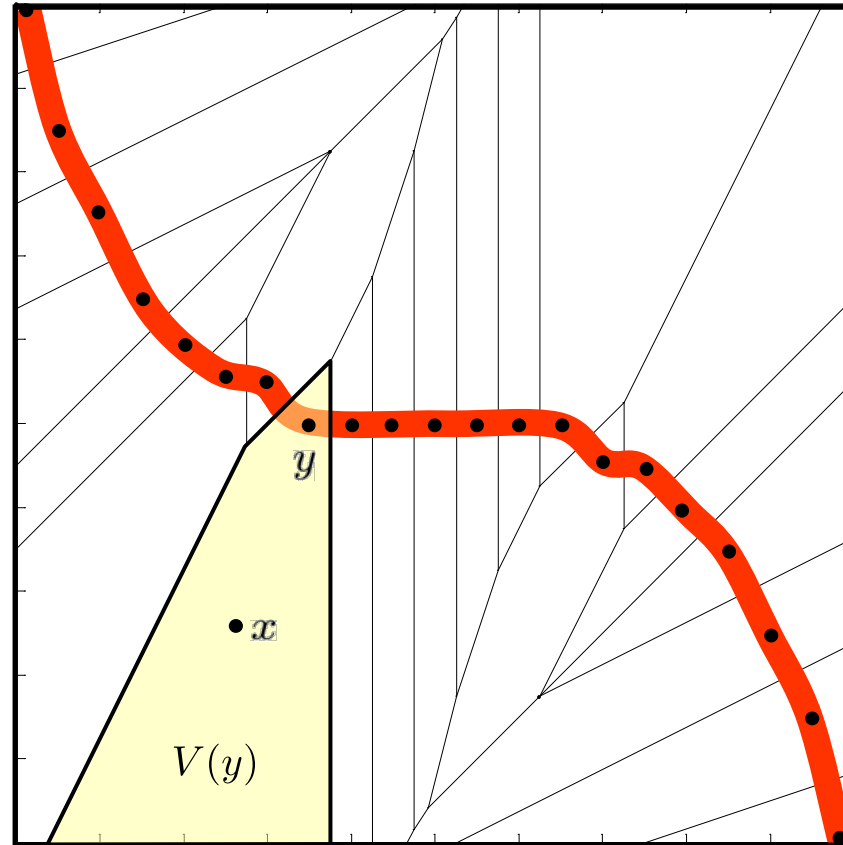
$$y_i = \arg \min_{y \in Y} \|y - x_i\|$$

- How to find closest points **efficiently**?
- Straightforward complexity: $\mathcal{O}(MN)$
- M number of points on X , N number of points on Y .
- More sophisticated: X divides the space into **Voronoi cells**

$$V(y \in Y) = \{z \in \mathbb{R}^3 : \|y - z\| < \|y' - z\| \ \forall y' \in Y \neq y\}$$

- Given a query point y , **determine to which cell it belongs.**

Closest Points: Voronoi



Computational
Geometry

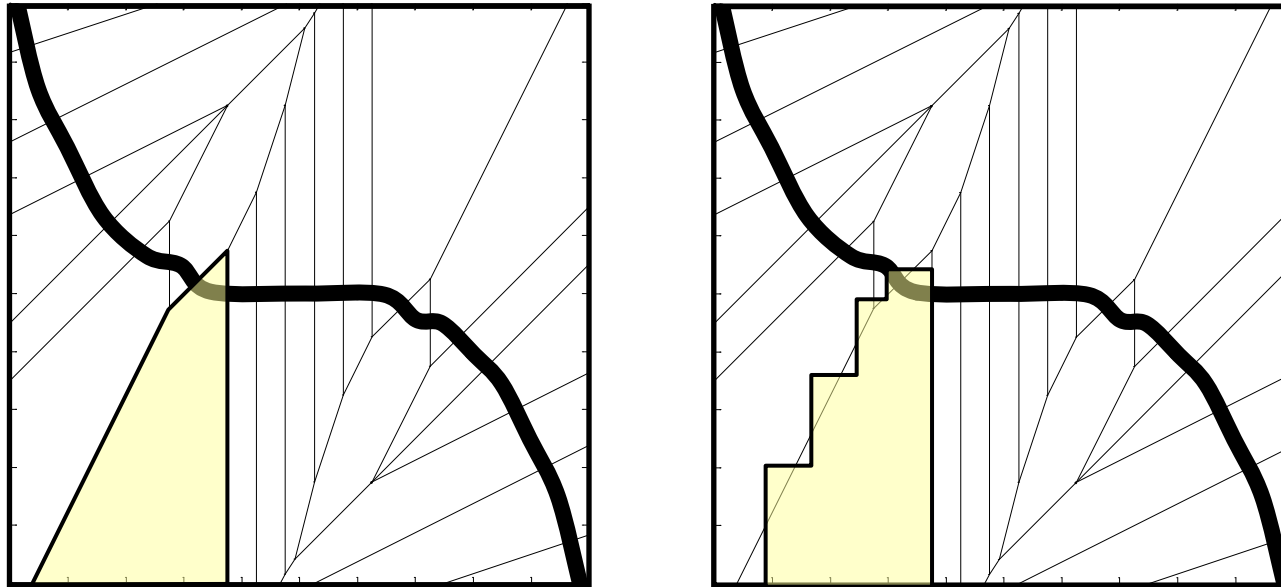
(→ CS268)

$$V(y \in Y) = \{z \in \mathbb{R}^3 : \|y - z\| < \|y' - z\| \forall y' \in Y \neq y\}$$

Source:
M. Bronstein

Closest Points: Voronoi Cells

Approximate nearest neighbors



M. Bronstein

- To reduce search complexity, **approximate** Voronoi cells.
- Use **binary space partition trees** (e.g. **kd-trees** or **octrees**).
- Approximate nearest neighbor search complexity: $\mathcal{O}(N \log M)$.

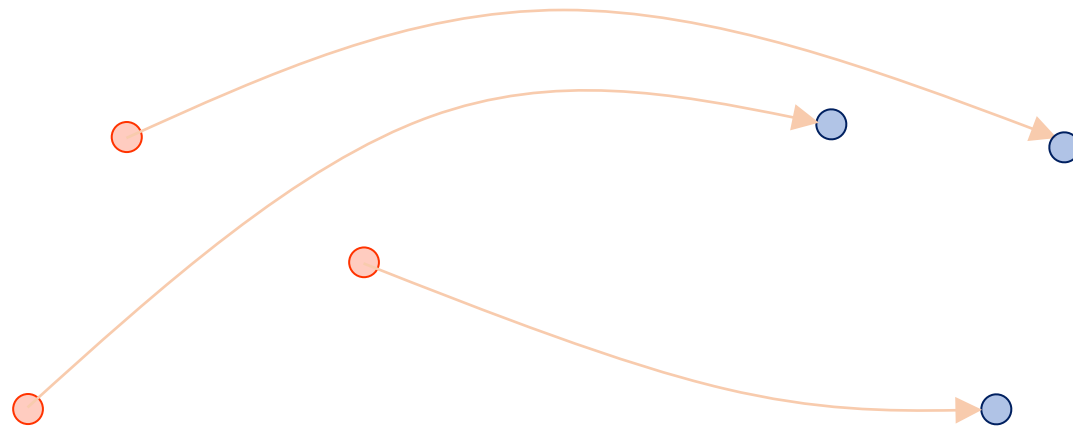
ICP: Optimal Transformation

Problem Formulation:

1. Given two sets points: $\{x_i\}, \{y_i\}, i = 1..n$ in \mathbb{R}^3 Find the rigid transform:

\mathbf{R}, t that minimizes:

$$\sum_{i=1}^N \|\mathbf{R}x_i + t - y_i\|_2^2$$



Simplest Case: Rigid Alignment, Given Correspondences

- We are given two sets of **corresponding** points x_1, x_2, \dots, x_n and y_1, y_2, \dots, y_n in \mathbb{R}^3 . We wish to compute the rigid transform T that best aligns x_1 to y_1, x_2 to y_2, \dots , and x_n to y_n .
- We define the error to be minimized by

$$\min_T \sum_{i=1}^n \|T(x_i) - y_i\|^2$$

MSE error, RMS distance, ...



- Old Problem:
 - Known and solved as the *orthogonal Procrustes problem* in Factor Analysis (Statistics) [Shönemann, 1966]
 - Known and solved as the *absolute orientation problem* in Photogrammetry [Horn, 1986]
 - Also in robotics, graphics, medical image analysis, statistical theories of shape, etc ...

SVD-Based Solution

- A rigid motion T is a combination of a translation a and a rotation R , so that $T(x) = R(x) + a$.
- The quantity to be minimized is:

$$\min_{a, R} \left(\sum_{i=1}^n |R(x_i) + a - y_i|^2 \right)$$

| The unknowns |

SVD-Based Solution

- A rigid motion T is a combination of a translation a and a rotation R , so that $T(x) = R(x) + a$.
- If we place the origin of our coordinate system at the mean of the x_i 's, then the quantity to be minimized simplifies to (up to some constants):

$$\min_{a, R} \left(\sum_{i=1}^n |y_i - a|^2 - 2 \sum_{i=1}^n \langle R(x_i), y_i \rangle \right)$$

- Note that the translational and rotational parts separate. The translational part a can easily be seen to be optimized by

$$a = \frac{1}{n} \sum_{i=1}^n y_i$$

The centroids of the two point sets have to be aligned!

The Rotation Part via SVD

- Define

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i$$

$$\bar{y} = \frac{1}{n} \sum_{i=1}^n y_i$$

$$X = [x_1 - \bar{x}, \dots, x_n - \bar{x}]^T$$

$$Y = [y_1 - \bar{y}, \dots, y_n - \bar{y}]^T$$

- Here X and Y are 3 by n matrices.

A diagram illustrating matrix multiplication. On the left, a yellow horizontal rectangle labeled 'X' is multiplied by a yellow vertical rectangle labeled 'Y^T'. The result is a yellow square labeled '3x3'.

*SVD = singular value decomposition

- Now compute the SVD*

$$XY^T = UDV^T \quad (3 \times 3)$$

- U and V are 3 by 3 orthogonal matrices, and D is a diagonal matrix with decreasing non-negative entries along the diagonal (the singular values).

- Define S by

$$S = \begin{cases} I, & \text{if } \det U \det V = 1 \\ \text{diag}(1, \dots, 1, -1), & \text{otherwise} \end{cases}$$

- Then

$$R = USV^T$$

$O(n)$ algorithm!

ICP: Optimal Transformation

Problem Formulation:

1. Given two sets points: $\{x_i\}, \{y_i\}, i = 1..n$ in \mathbb{R}^3 . Find the rigid transform:

\mathbf{R}, t that minimizes:
$$\sum_{i=1}^N \|\mathbf{R}x_i + t - y_i\|_2^2$$

2. Closed form solution:

1. Construct: $C = \sum_{i=1}^N (y_i - \mu^Y)(x_i - \mu^X)^T$, where $\mu^X = \frac{1}{N} \sum_i x_i$,

2. Compute the SVD of C: $C = U\Sigma V^T$ $\mu^Y = \frac{1}{N} \sum_i y_i$

1. If $\det(UV^T) = 1, R_{\text{opt}} = UV^T$

2. Else $R_{\text{opt}} = U\tilde{\Sigma}V^T, \tilde{\Sigma} = \text{diag}(1, 1, \dots, -1)$

3. Set $t_{\text{opt}} = \mu^Y - R_{\text{opt}}\mu^X$

Note that C is a 3x3 matrix. SVD is very fast.

Arun et al., Least-Squares Fitting of Two 3-D Point Sets

Iterative Closest Point

Given a pair of shapes, X and Y , iterate:

1. For each $x_i \in X$ find **nearest** neighbor $y_i \in Y$.
2. Find rigid motion \mathbf{R}, t minimizing:
$$\sum_{i=1}^N \|\mathbf{R}x_i + t - y_i\|_2^2$$

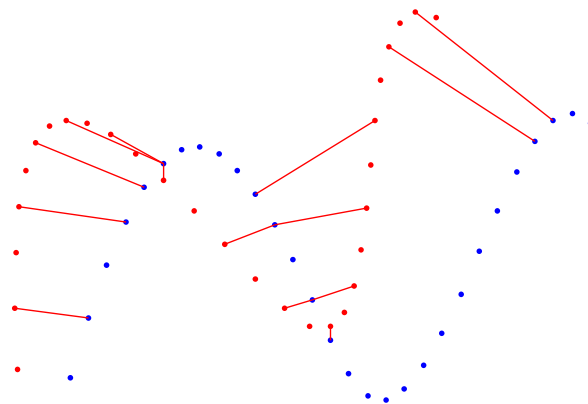
Convergence:

- at each iteration $\sum_{i=1}^N d^2(x_i, Y)$ decreases.
- Converges to local minimum
- Good initial guess: hopefully global minimum.

[Besl&McKay92]

Variations of ICP

1. **Selecting** source points (from one or both scans): sampling
2. **Matching** to points in the other mesh
3. **Weighting** the correspondences
4. **Rejecting** certain (outlier) point pairs
5. **Assigning** an error metric to the current transform
6. **Minimizing** the error metric w.r.t. the transformation

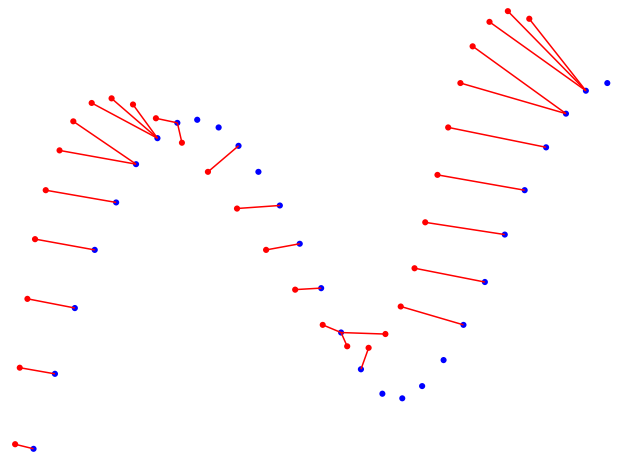


Iterative Closest Point

Given a pair of shapes, X and Y , iterate:

1. For each $x_i \in X$ find **nearest** neighbor $y_i \in Y$.
2. Find rigid motion \mathbf{R}, t minimizing:

$$\sum_{i=1}^N \|\mathbf{R}x_i + t - y_i\|_2^2$$



Iterative Closest Point

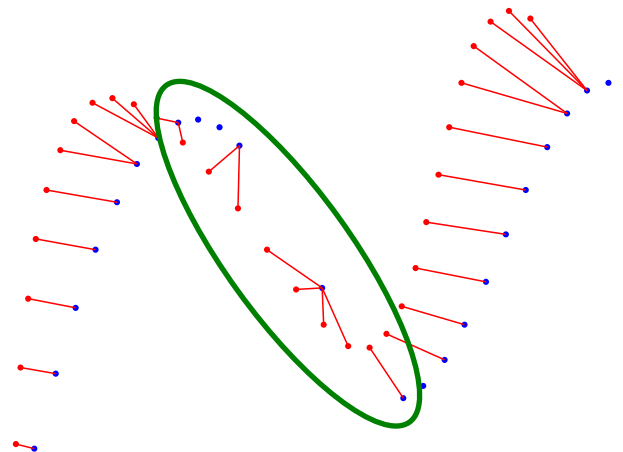
Given a pair of shapes, X and Y , iterate:

1. For each $x_i \in X$ find **nearest** neighbor $y_i \in Y$.
2. Find rigid motion \mathbf{R}, t minimizing:

$$\sum_{i=1}^N \|\mathbf{R}x_i + t - y_i\|_2^2$$

Problem:

uneven sampling



Iterative Closest Point

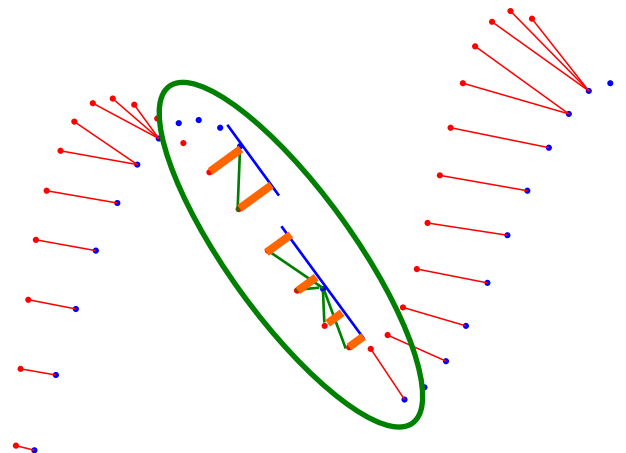
Given a pair of shapes, X and Y , iterate:

1. For each $x_i \in X$ find **nearest** neighbor $y_i \in Y$.
2. Find rigid motion \mathbf{R}, t minimizing:

$$\sum_{i=1}^N d(\mathbf{R}x_i + t, P(y_i))^2 = \sum_{i=1}^N ((\mathbf{R}x_i + t - y_i)^T \mathbf{n}_{y_i})^2$$

Solution:

Minimize distance to
the tangent plane



Iterative Closest Point

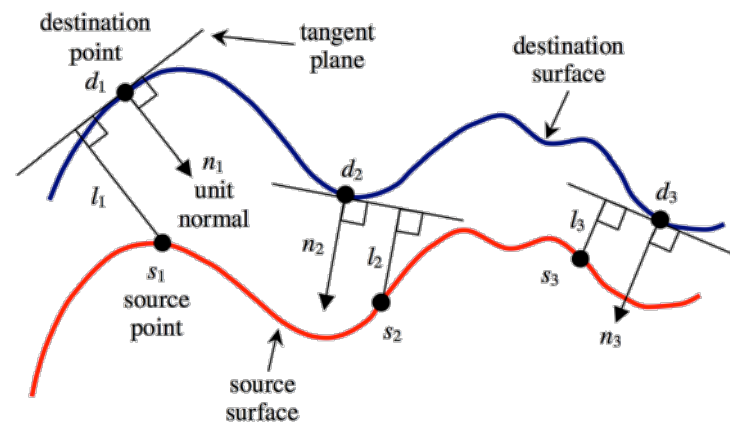
Given a pair of shapes, X and Y , iterate:

1. For each $x_i \in X$ find **nearest** neighbor $y_i \in Y$.
2. Find rigid motion \mathbf{R}, t minimizing:

$$\sum_{i=1}^N d(\mathbf{R}x_i + t, P(y_i))^2 = \sum_{i=1}^N ((\mathbf{R}x_i + t - y_i)^T \mathbf{n}_{y_i})^2$$

Solution:

Minimize distance to the tangent plane



Iterative Closest Point

Given a pair of shapes, X and Y , iterate:

1. For each $x_i \in X$ find **nearest** neighbor $y_i \in Y$.
2. Find rigid motion \mathbf{R}, t minimizing:

$$\mathbf{R}_{\text{opt}}, t_{\text{opt}} = \underset{\substack{\mathbf{R}^T \mathbf{R} = \text{Id}, t}}{\text{arg min}} \sum_{i=1}^N ((\mathbf{R}x_i + t - y_i)^T \mathbf{n}_{y_i})$$

Question:

How to minimize the error?

Challenge:

Although the error is **quadratic** (linear derivative), the space of rotation matrices is **not linear**.

Problem:

No closed form solution.

Iterative Closest Point

Given a pair of shapes, X and Y, iterate:

1. For each $x_i \in X$ find **nearest** neighbor $y_i \in Y$.
2. Find rigid motion \mathbf{R}, t minimizing:

$$\mathbf{R}_{\text{opt}}, t_{\text{opt}} = \underset{\substack{\mathbf{R}^T \mathbf{R} = \text{Id}, t}}{\text{arg min}} \sum_{i=1}^N ((\mathbf{R}x_i + t - y_i)^T \mathbf{n}_{y_i})$$

Common Approach:

Linearize rotation. Assume rotation angle is small.

$$\mathbf{R}x_i \approx x_i + r \times x_i \quad \begin{array}{l} r : \text{axis,} \\ \|r\|_2 : \text{angle of rotation.} \end{array}$$

Note: follows from Rodrigues's formula

$$R(r, \alpha)x_i = x_i \cos(\alpha) + (r \times x_i) \sin(\alpha) + r(r^T x_i)(1 - \cos(\alpha))$$

and first order approximations: $\sin(\alpha) \approx \alpha, \cos(\alpha) \approx 1$

Iterative Closest Point

Given a pair of shapes, X and Y , iterate:

1. For each $x_i \in X$ find **nearest** neighbor $y_i \in Y$.
2. Find rigid motion r, t minimizing:

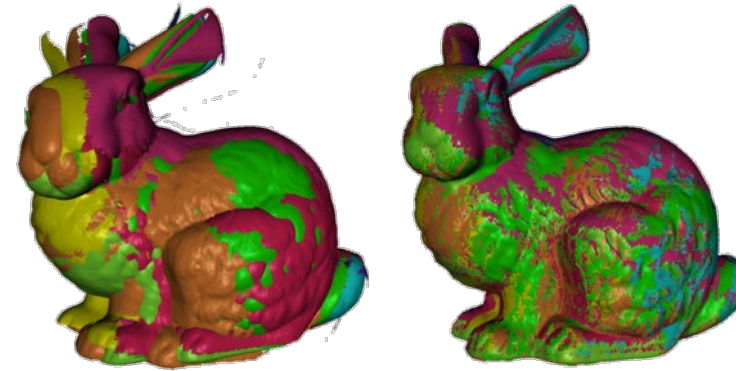
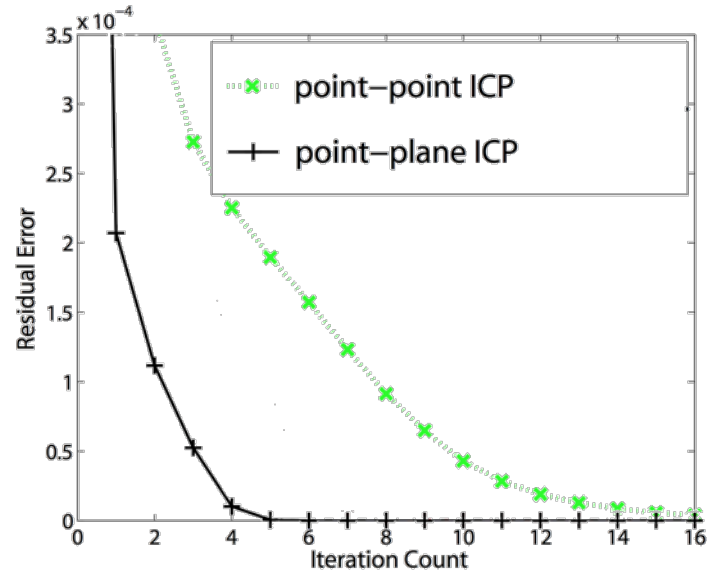
$$\begin{aligned} E(r, t) &= \sum_{i=1}^N ((x_i + r \times x_i + t - y_i)^T \mathbf{n}_{y_i}) \\ &= \sum_{i=1}^N ((x_i - y_i)^T \mathbf{n}_{y_i} + r^T (x_i \times \mathbf{n}_{y_i}) + t^T \mathbf{n}_{y_i})^2 \end{aligned}$$

Setting: $\frac{\partial}{\partial r} E(r, t) = 0$ and $\frac{\partial}{\partial t} E(r, t) = 0$ leads to a 6x6 linear system

$$Ax = b$$

$$x = \begin{pmatrix} r \\ t \end{pmatrix} \quad A = \sum \begin{pmatrix} x_i \times \mathbf{n}_{y_i} \\ \mathbf{n}_{y_i} \end{pmatrix} \begin{pmatrix} x_i \times \mathbf{n}_{y_i} \\ \mathbf{n}_{y_i} \end{pmatrix}^T \quad b = \sum (y_i - x_i)^T \mathbf{n}_{y_i} \begin{pmatrix} x_i \times \mathbf{n}_{y_i} \\ \mathbf{n}_{y_i} \end{pmatrix}$$

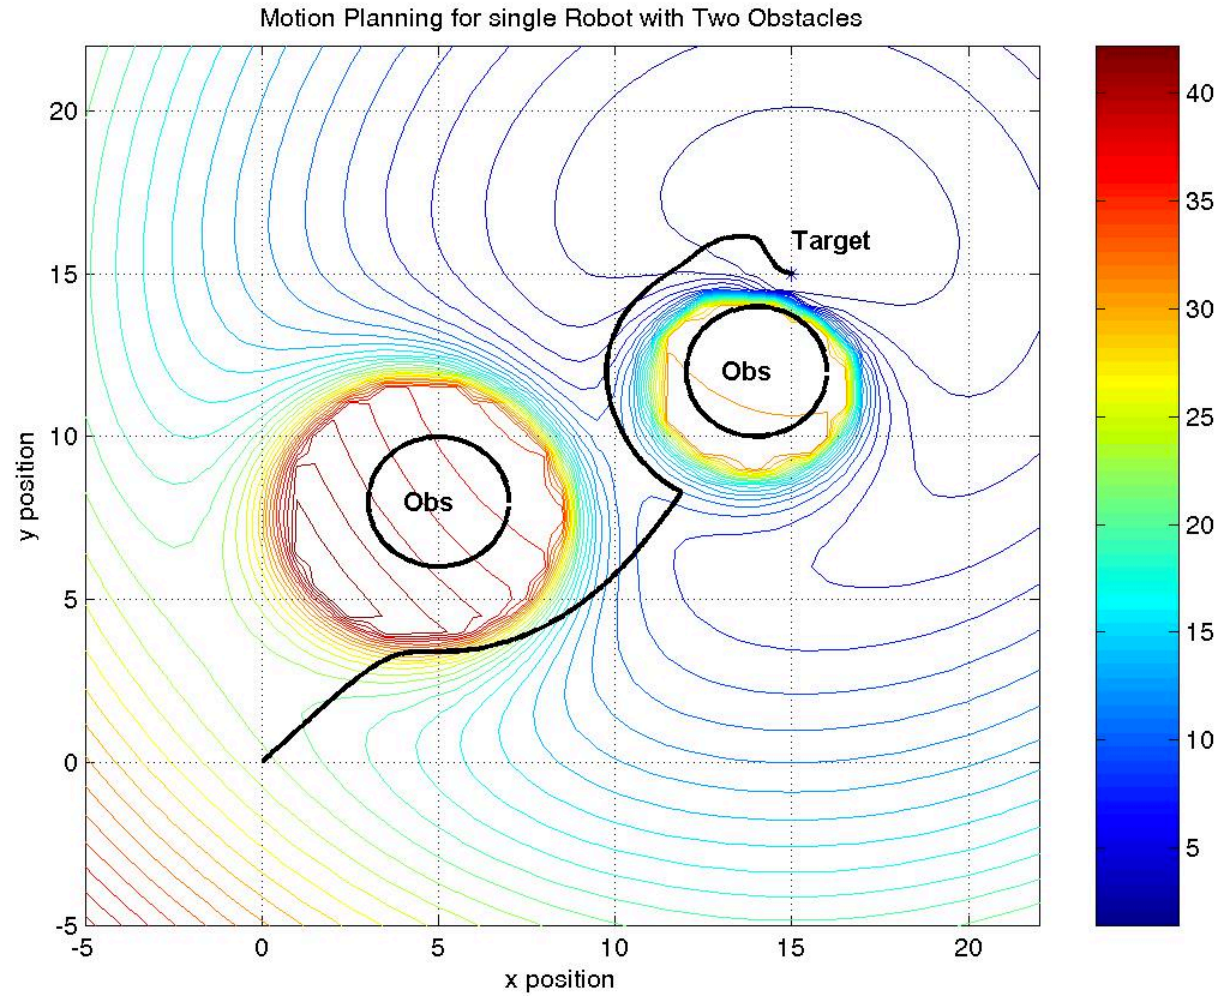
Iterative Closest Point



Aligning the bunny to itself:
Point-to-plane always wins in the end-game.

Distance Fields for Registration

“Gravitational” Potential

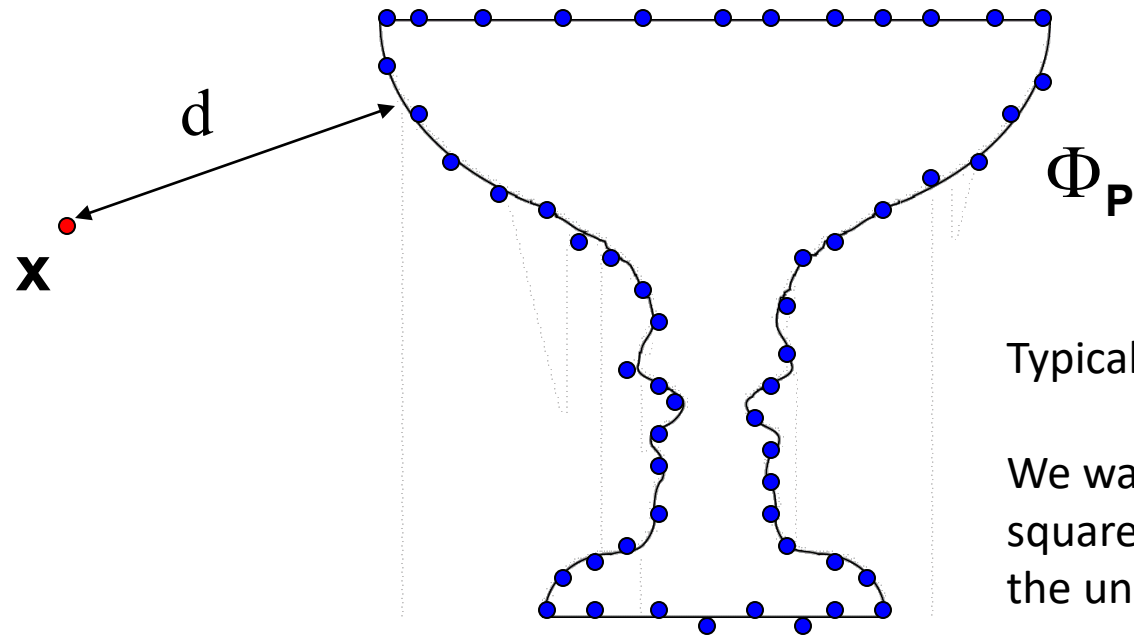


Robot motion planning via potential fields

“Gravitational” Potential

- Given two related shapes, the “data” A and the “model” B, create a potential field that **pulls** B to the correct alignment with A
- Key tasks
 - Define the potential field
 - Formulate the optimization problem
 - Do gradient descent using approximate linearization
 - Iterative approach

Squared Distance Function (F)

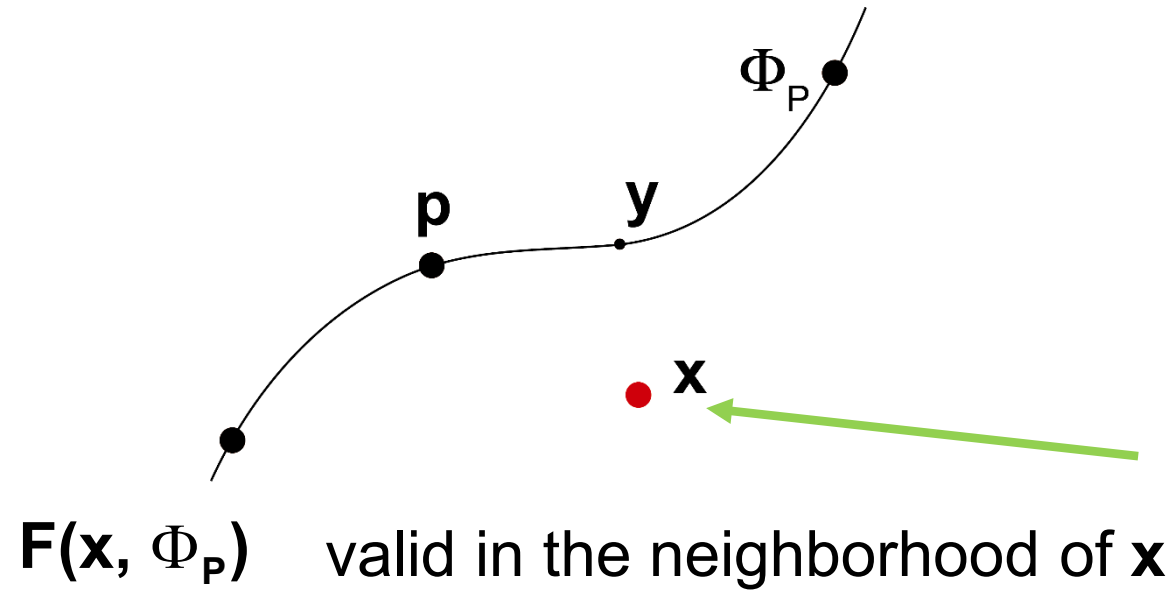


Typically A is a point cloud

We want to approximate the squared distance function to the underlying object

$$\mathbf{F}(\mathbf{x}, \Phi_P) = d^2$$

Approximate Squared Distance

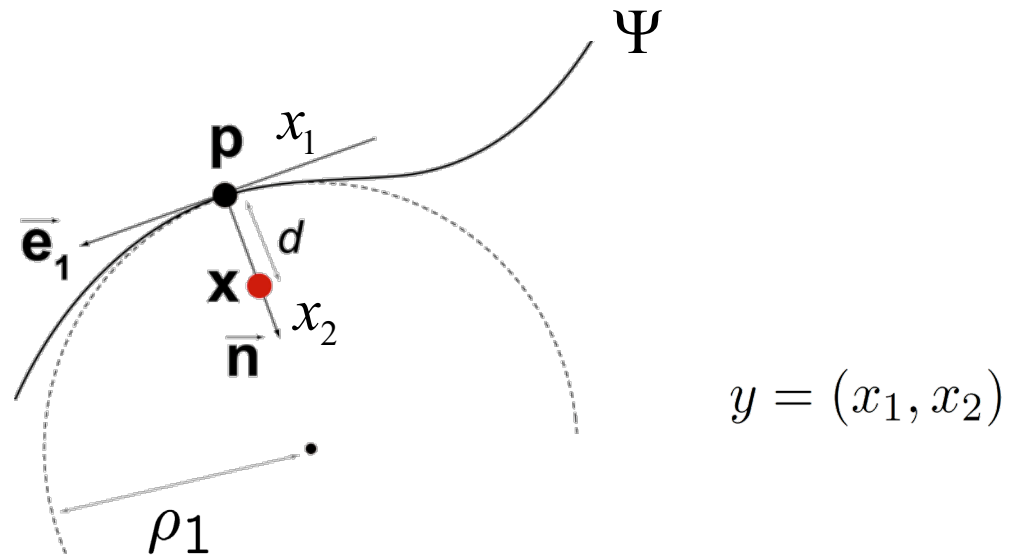


Aim for 2nd order approximation, because we want to take derivatives.

Pairwise Rigid Correspondence

Geometry of the square distance function

For a curve Ψ ,
around point x



To second order:

$$d^2(y, \Psi) \approx \frac{d}{d - \rho_1} x_1^2 + x_2^2 \quad \text{in the Frenet frame at } p$$

[Pottmann and Hofer 2003]

Approximate Squared Distance

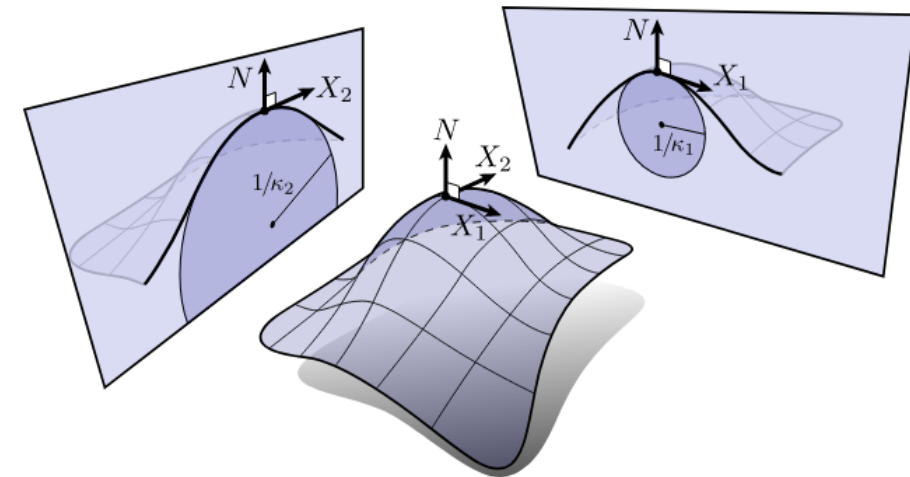
For a curve Ψ , to second order:

$$d^2(y, \Psi) \approx \frac{d}{d - \rho_1} x_1^2 + x_2^2$$

For a surface Φ , to second order:

$$d^2(y, \Phi) \approx \frac{d}{d - \rho_1} x_1^2 + \frac{d}{d - \rho_2} x_2^2 + x_3^2$$

$\rho_1 = 1/\kappa_1$ and $\rho_2 = 1/\kappa_2$ are inverse principal curvatures



[Pottmann and Hofer 2003]

Approximate Squared Distance

For a surface Φ , to second order:

$$d^2(y, \Phi) \approx \frac{d}{d - \rho_1} x_1^2 + \frac{d}{d - \rho_2} x_2^2 + x_3^2$$

$\rho_1 = 1/\kappa_1$ and $\rho_2 = 1/\kappa_2$ are inverse principal curvatures

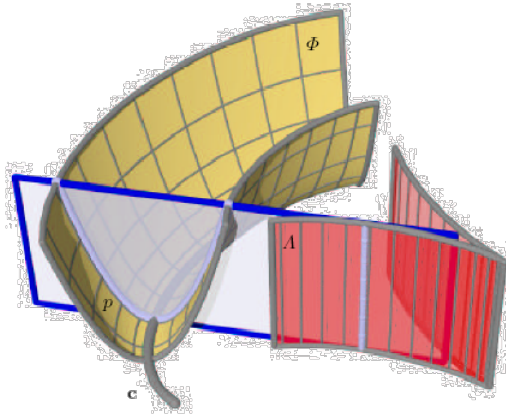
Note that as $d \rightarrow 0$, $d^2(y, \Phi) \rightarrow x_3^2$ point-to-plane

$d \rightarrow \infty$, $d^2(y, \Phi) \rightarrow x_1^2 + x_2^2 + x_3^2$ point-to-point

ICP Without Correspondences

- ICP without correspondences
 - define a **quadratic approximant to the square distance function**

[Pottman & Hofer, 02]

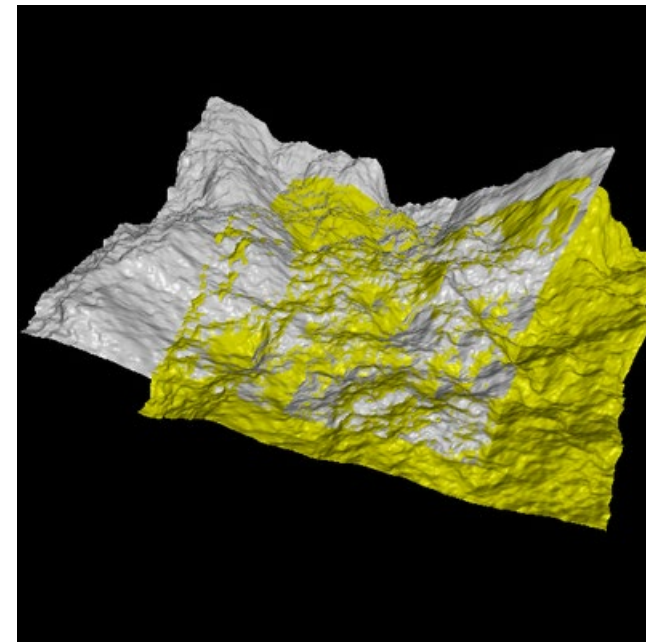


$$F(x_1, x_2, x_3) = \frac{d}{d - \rho_1} x_1^2 + \frac{d}{d - \rho_2} x_2^2 + x_3^2$$

$$F(x_1, x_2) = \frac{d}{d - \rho} x_1^2 + x_2^2$$

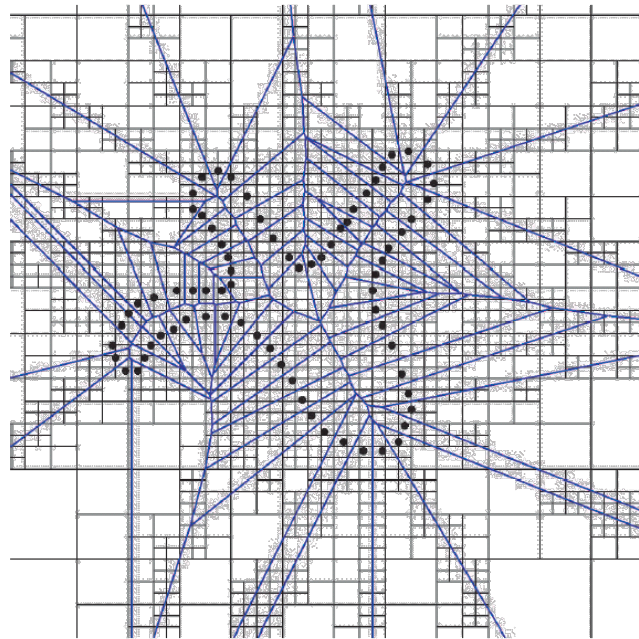
← curvature

- perform iterative gradient-descent in this field
- point to foot-point distance
 - case d is large: classical ICP
 - case d is small: point-to-plane ICP

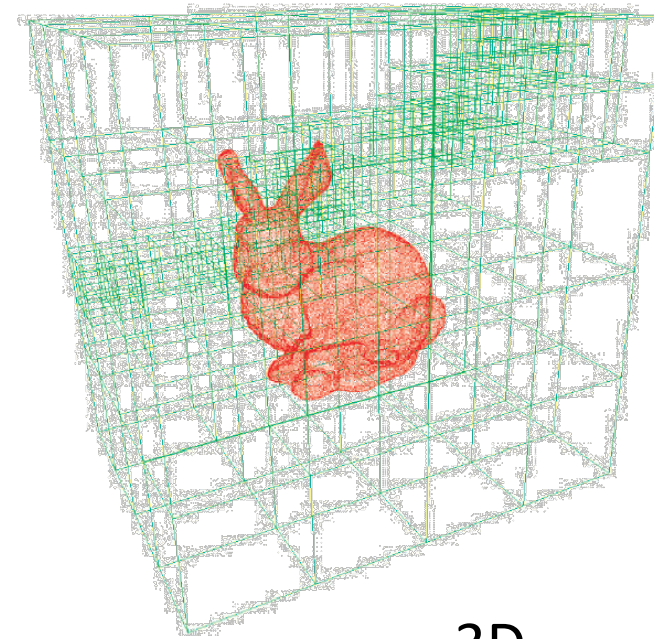


$d^2(y, \Phi_p)$ Using d2 Tree

Partition the space into cells where each cell stores a quadratic approximant of the squared distance function.



2D



3D

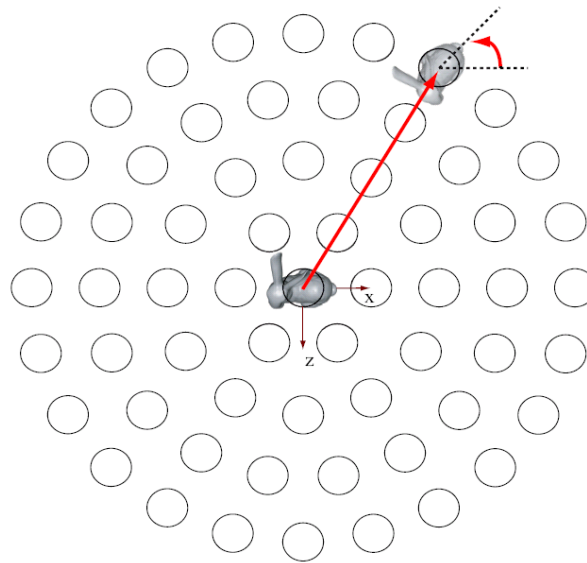
Leopoldseder S et al. *d2-tree: A hierarchical representation of the squared distance function*

Registration Using d2 Tree

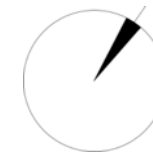
Build using bottom-up approach: fit a quadratic approximation to a fine grid.

Merge cells if they have similar approximations.

Funnel of convergence:



Translation in x-z plane. Rotation about y-axis.

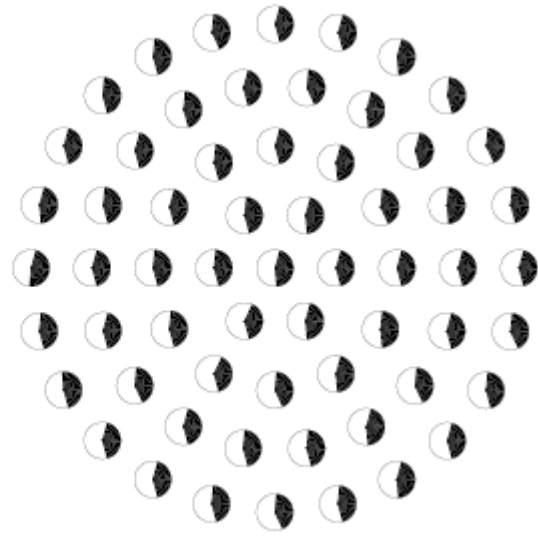
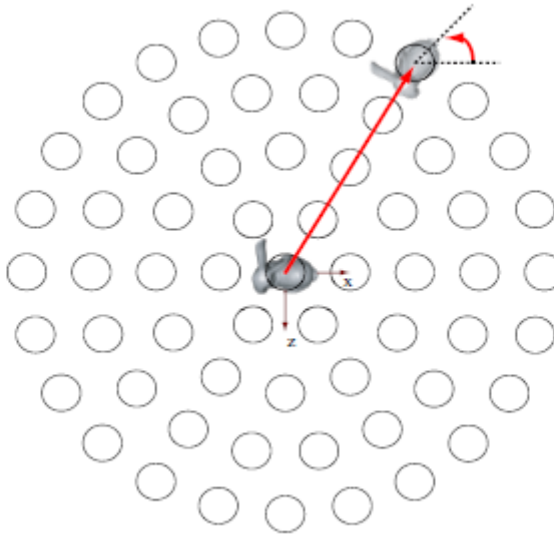


- Converges
- Does not converge

Matching the Bunny to Itself



point-to-plane

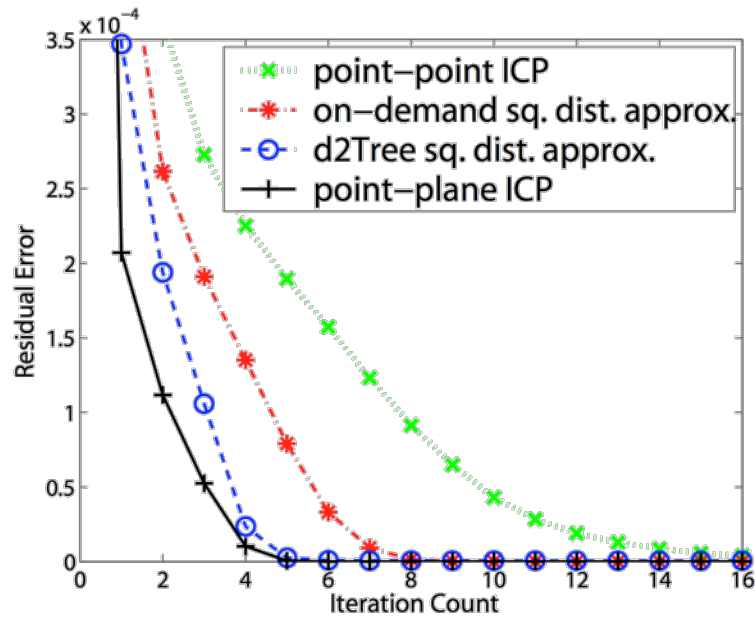


d2Tree

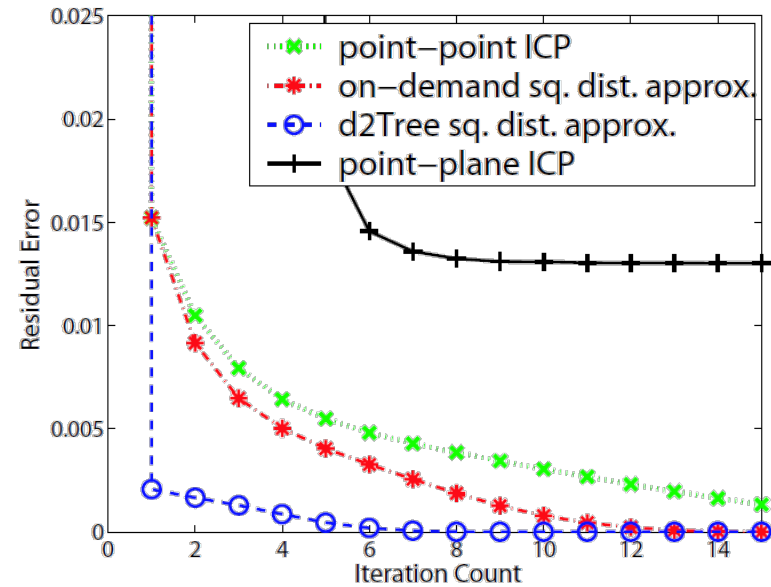
*Registration of Point Cloud Data from a
Geometric Optimization Perspective*

Mitra et al., SGP 2004

Matching the Bunny to Itself



Well-aligned



Noisy, far away

Registration of Point Cloud Data from a Geometric Optimization Perspective

Mitra et al., SGP 2004

Local Rigid Matching – ICP

The upshot is that

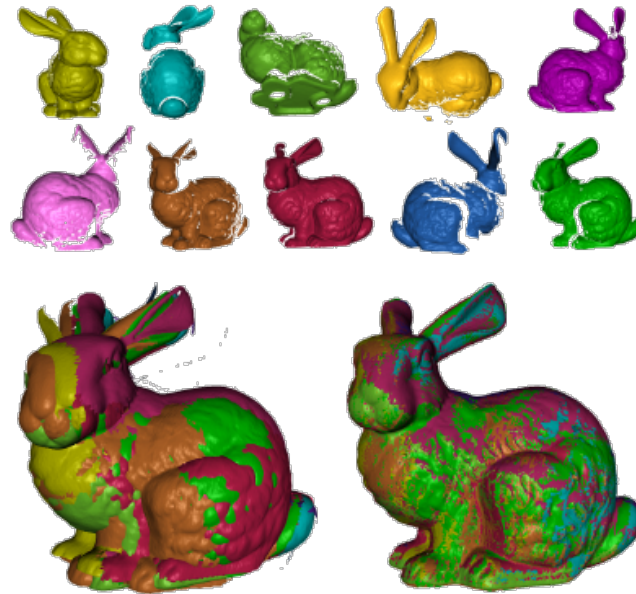
- Locally, the point-to-plane metric provides a second order approximation to the squared distance function.
- Optimization based on point-to-plane will converge quadratically to a local minimum.
- Convergence funnel can be narrow, but can improve it with either d2tree or point-to-point.

What if we are outside the convergence funnel?

Global Matching

Global Matching

Given shapes in *arbitrary* positions, find their alignment:



Robust Global Registration
Gelfand et al. SGP 2005

Can be approximate, since will refine later using e.g. ICP

Global Matching – Approaches

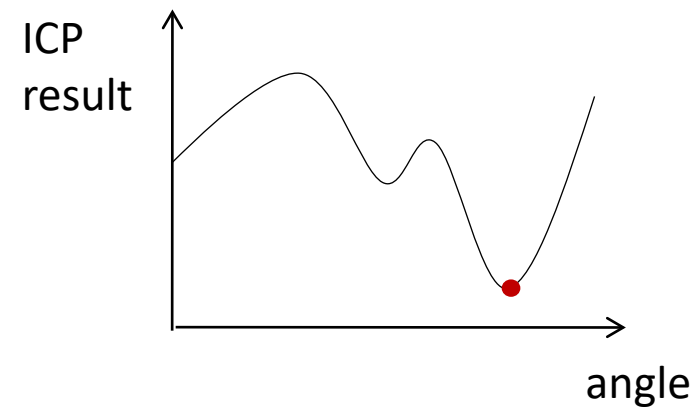
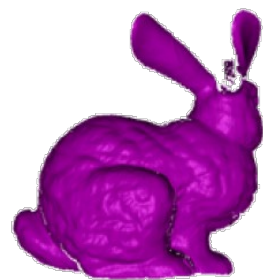
Several classes of approaches:

1. Exhaustive Search
2. Normalization (PCA)
3. Random Sampling (RANSAC)
4. Invariant Features

Exhaustive Search:

Compare (ideally) all alignments

- Sample the space of possible initial alignments.
- Correspondence is determined by the alignment at which models are closest.



Very common in biology: e.g., protein docking

Exhaustive Search:

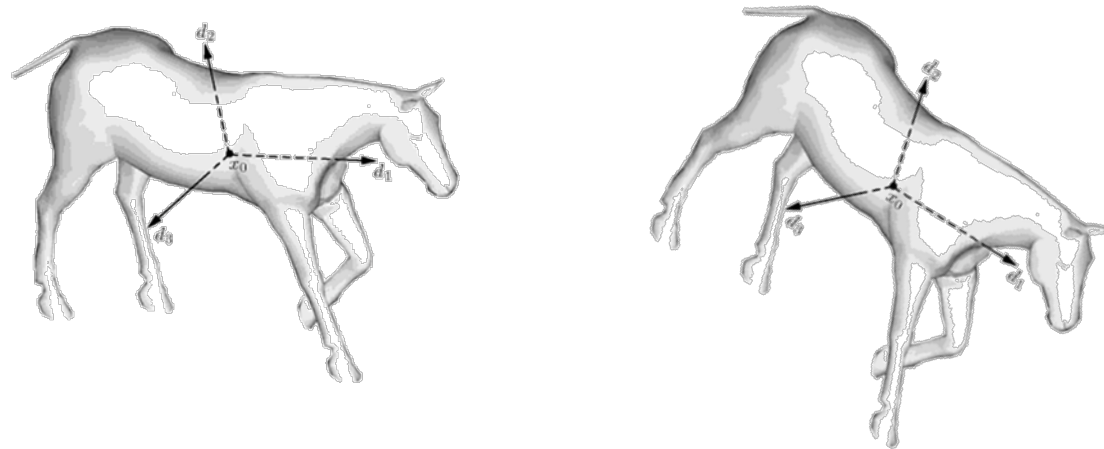
Compare at all alignments

- Sample the space of possible initial alignments
- Correspondence is determined by the alignment at which models are closest
- Provides optimal result
- Can be unnecessarily slow
- Does not generalize to non-rigid deformations

Normalization – Canonical Poses / Frames

There are only a handful of initial configurations that are important.

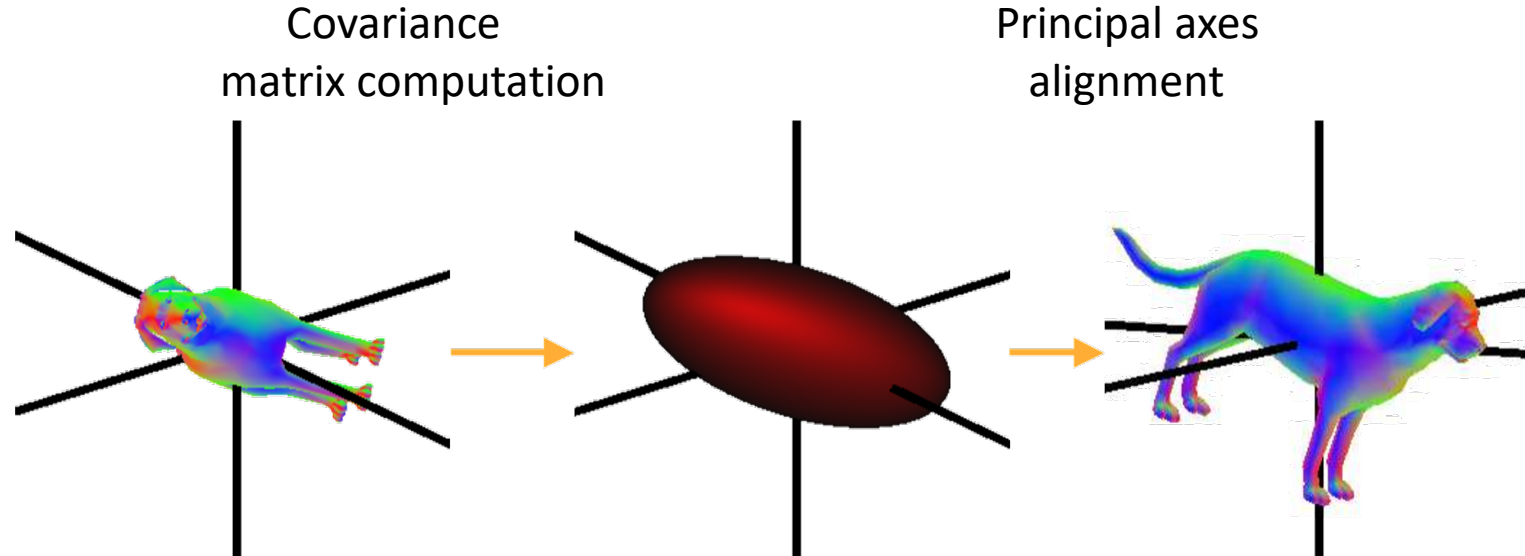
Can center all shapes at the origin and use PCA to find the principal directions of the shape.



In addition sometimes try all permutations of x-y-z.

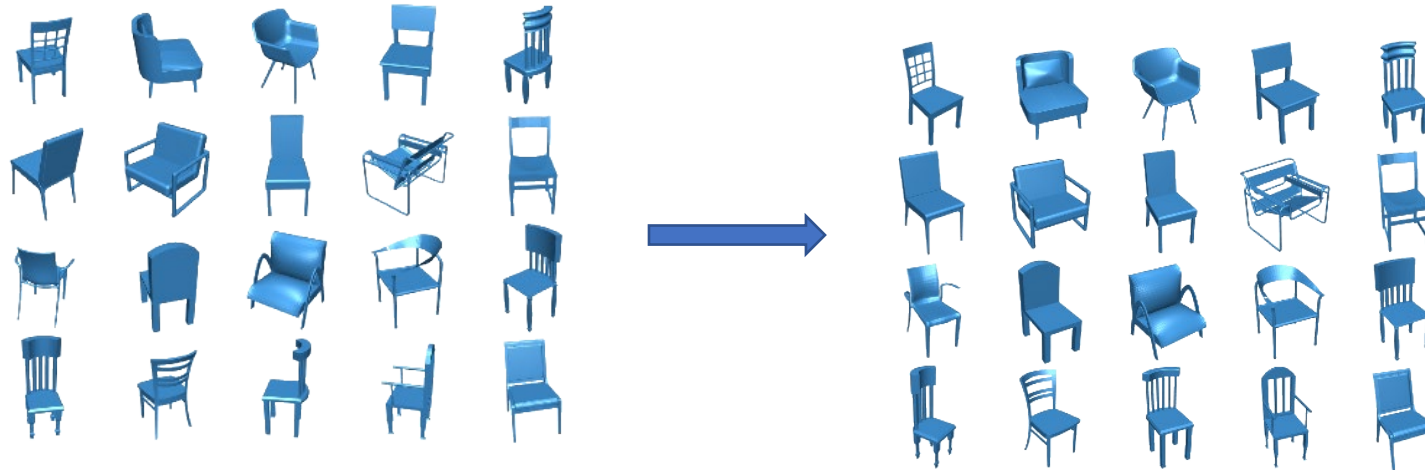
PCA-Based Alignment

- ◆ Use PCA to place models into canonical coordinate frames
- ◆ Then align those frames



Normalization – Canonical Poses

There are only a handful of initial configurations that are important.

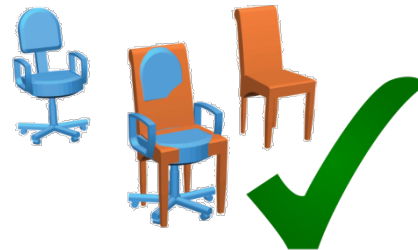
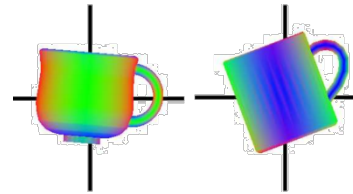
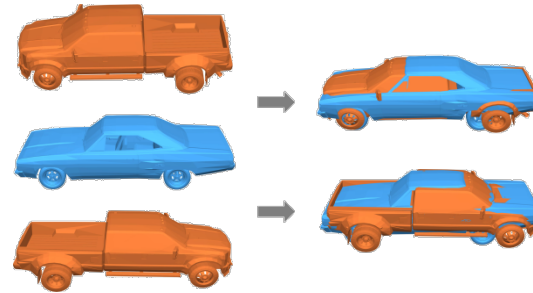


Works well if we have complete shapes and no noise.

Fails for partial scans, outliers, high noise, etc.

Problems with PCA

- ◆ Principal axes are not consistently oriented
- ◆ Axes are unstable when principal values are similar
- ◆ Partial similarity

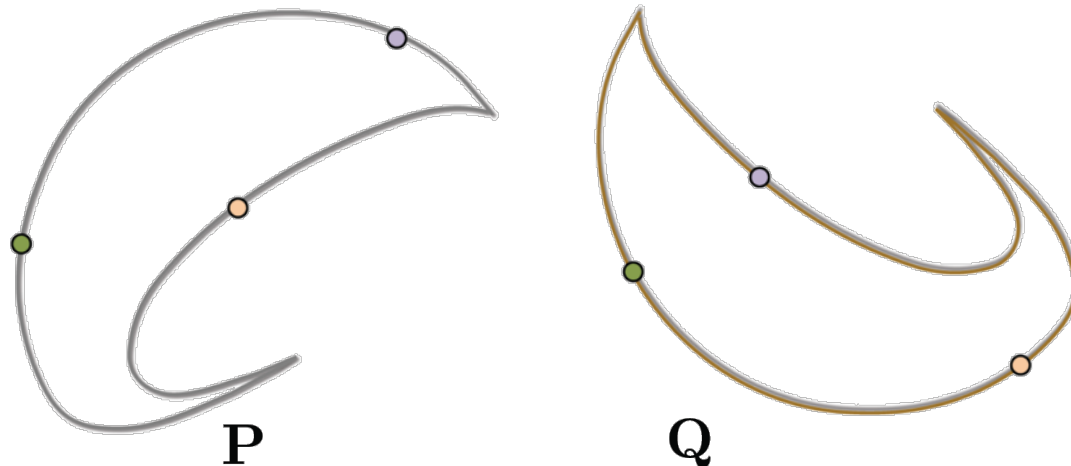


Random Sampling (RANSAC)

ICP only needs 3 point pairs! – rigid motions are low-dimensional (6 params)

Robust and simple approach. Iterate between:

1. Pick a random pair of 3 points on model & scan
2. Estimate alignment, and check for error.



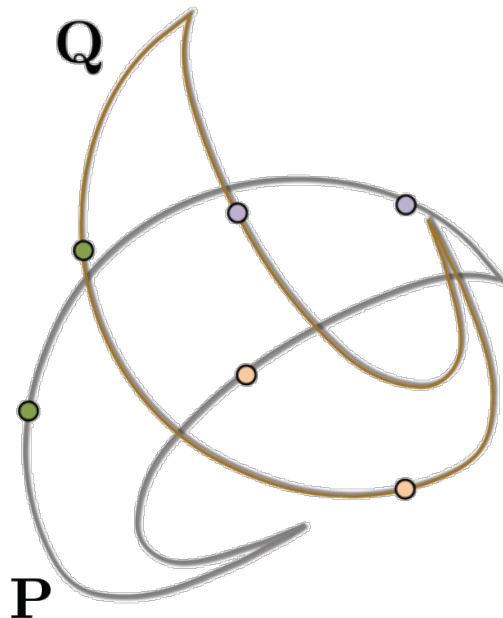
Guess and
verify

Random Sampling (RANSAC)

ICP only needs 3 point pairs!

Robust and simple approach. Iterate between:

1. Pick a random pair of 3 points on model & scan
2. Estimate alignment, and check for error.



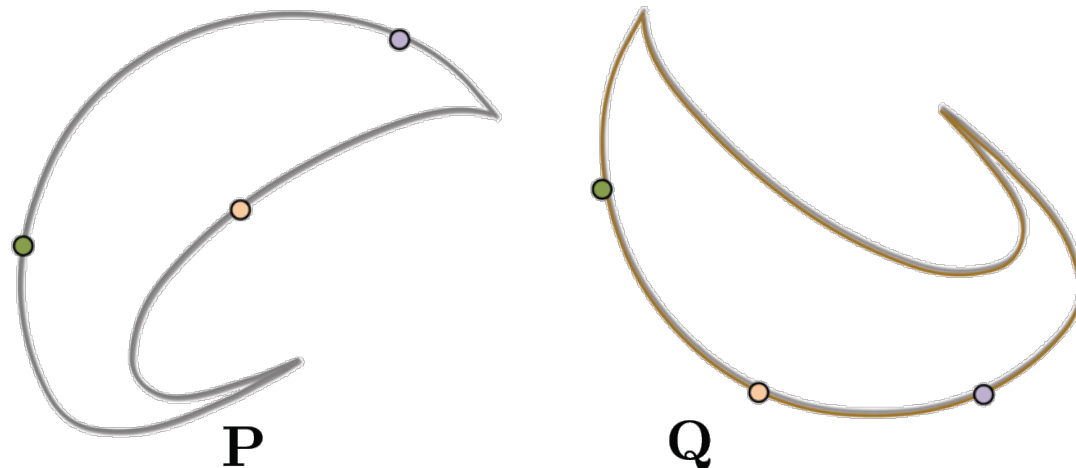
Guess and
verify

Random Sampling (RANSAC)

ICP only needs 3 point pairs!

Robust and Simple approach. Iterate between:

1. Pick a random pair of 3 points on model & scan
2. Estimate alignment, and check for error.



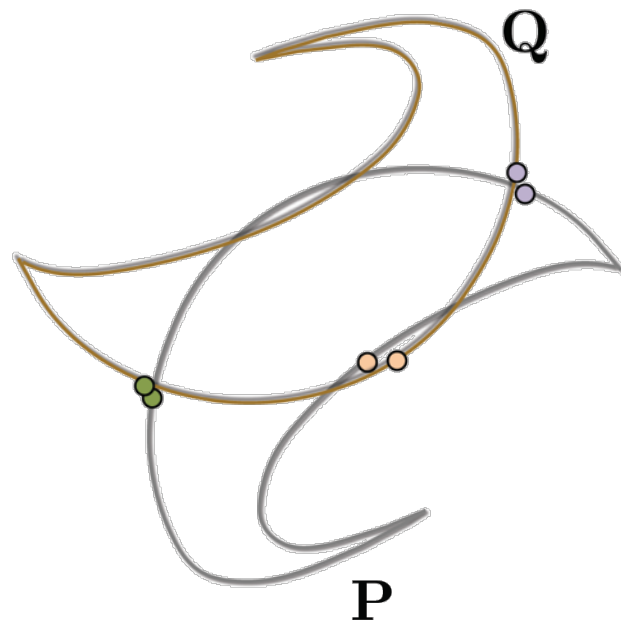
Guess and
verify

Random Sampling (RANSAC)

ICP only needs 3 point pairs!

Robust and simple approach. Iterate between:

1. Pick a random pair of 3 points on model & scan
2. Estimate alignment, and check for error.



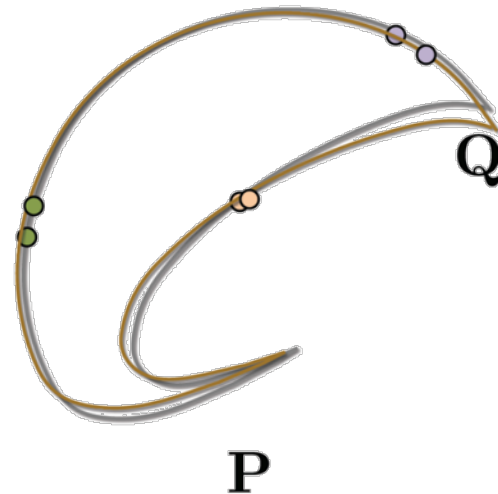
Guess and
verify

Random Sampling (RANSAC)

ICP only needs 3 point pairs!

Robust and simple approach. Iterate between:

1. Pick a random pair of 3 points on model & scan
2. Estimate alignment, and check for error.



Can also refine the final result. Picks don't have to be exact.

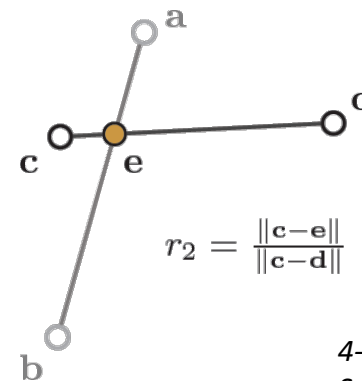
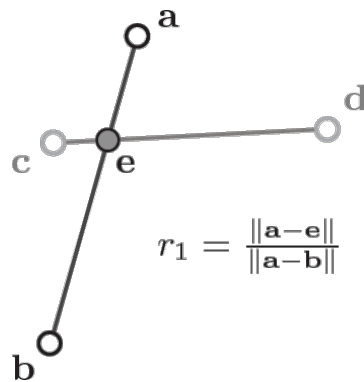
Guess and
verify

Random Sampling (RANSAC)

A pair of triples (from **P** and **Q**) are enough to determine a **rigid transform**,
resulting in $O(n^3)$ RANSAC.

Surprisingly, a special set of 4 points, **congruent sets**, makes the problem
simpler leading to $O(n^2)$!

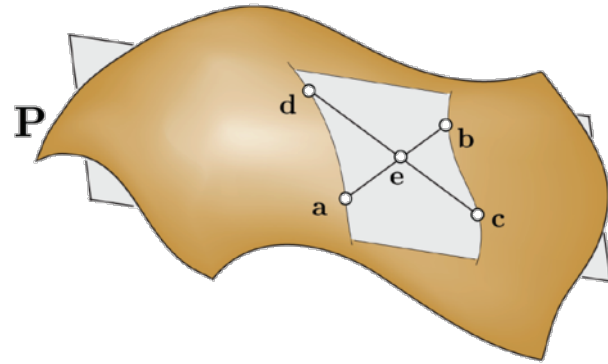
Co-planar points
remain coplanar



*4-points Congruent Sets for Robust
Surface Registration,
Aiger et al., SIGGRAPH 2008*

Method Overview

On the source shape, pick 4 (approx.) coplanar points.



Compute

$$r_1 = \frac{\|a - e\|}{\|a - b\|} \quad r_2 = \frac{\|d - e\|}{\|d - c\|}$$

For every pair (q_1, q_2) of points on the destination compute

$$p_1 = q_1 + r_1(q_2 - q_1)$$

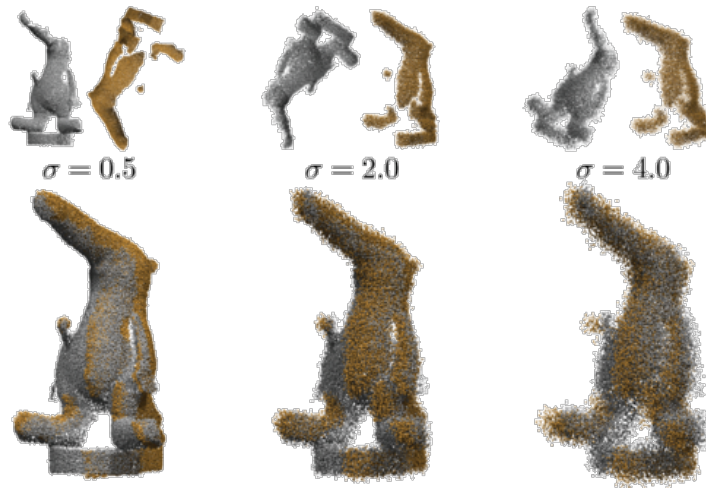
$$p_2 = q_1 + r_2(q_2 - q_1)$$

Those pairs (q_1, q_2) , (q_3, q_4) for which $p_1(q_1, q_2) = p_2(q_3, q_4)$ are a good candidate correspondence for (a, b, c, d) .

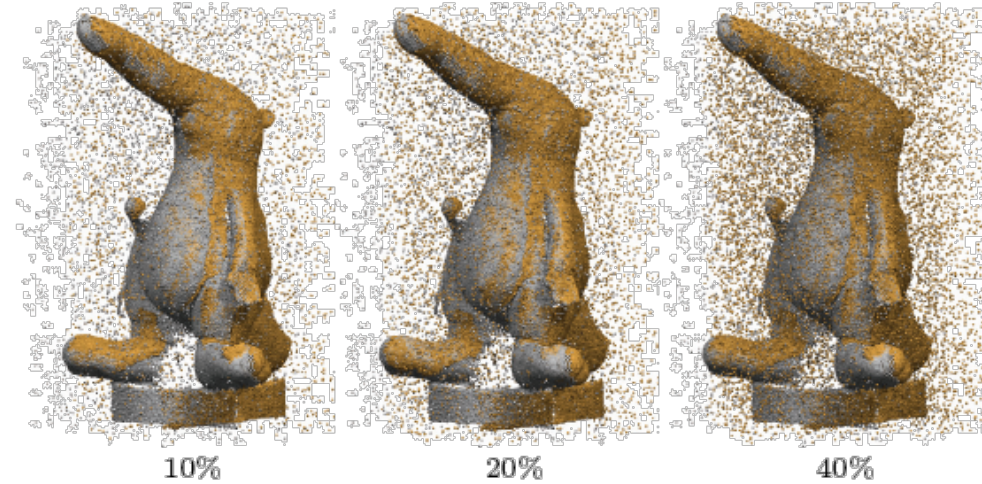
Under mild assumptions the procedure runs in $O(n^2)$ time.

Method Overview

Can pick a few base points for partial matching.



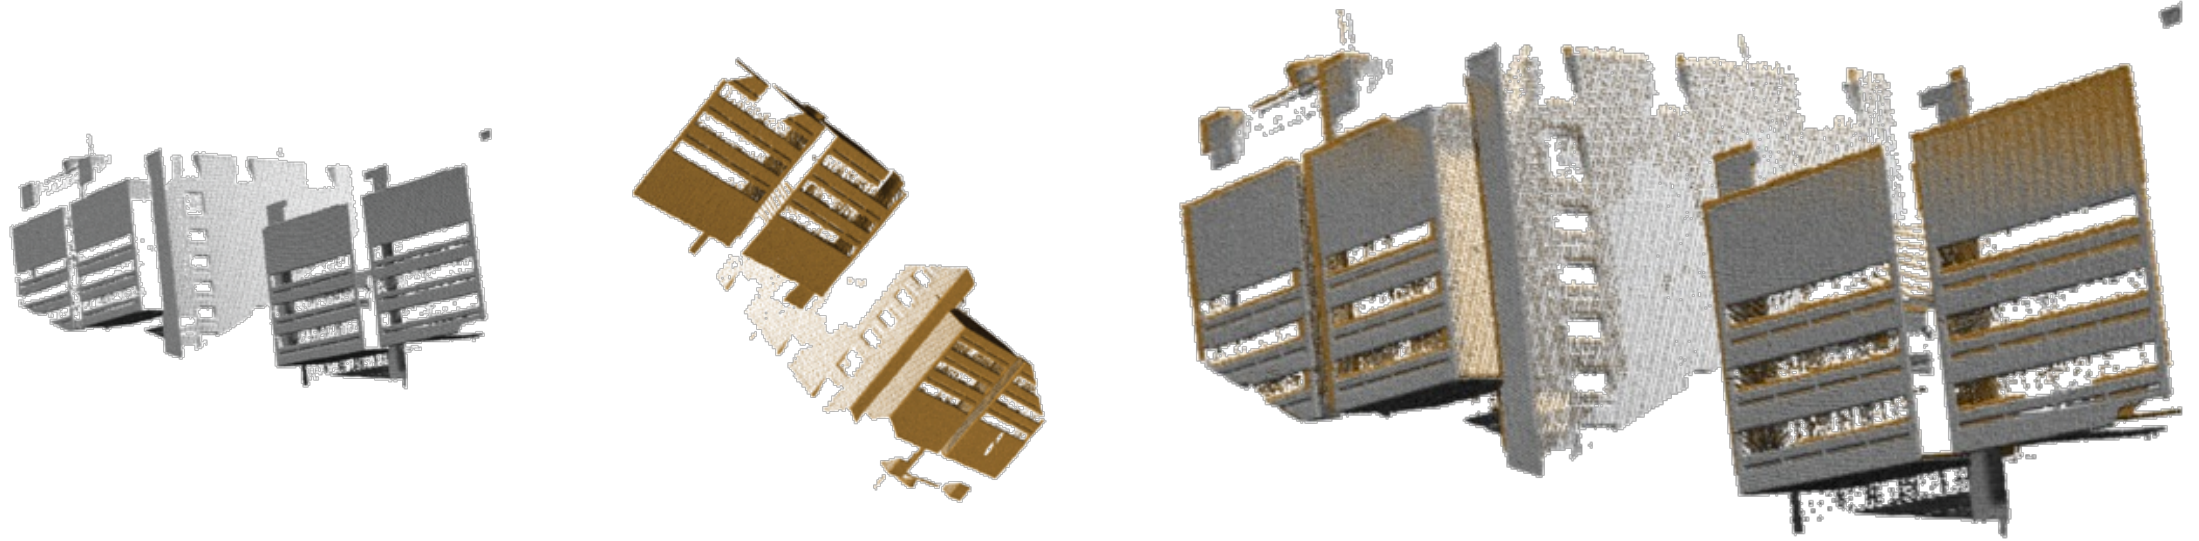
Random sampling
handles noise



and outliers

Method Overview

Can pick a few base points for partial matching.



Partial matches

Global Matching – Approaches

Several classes of approaches:

1. Exhaustive Search
2. Normalization
3. Random Sampling
4. Invariant Features

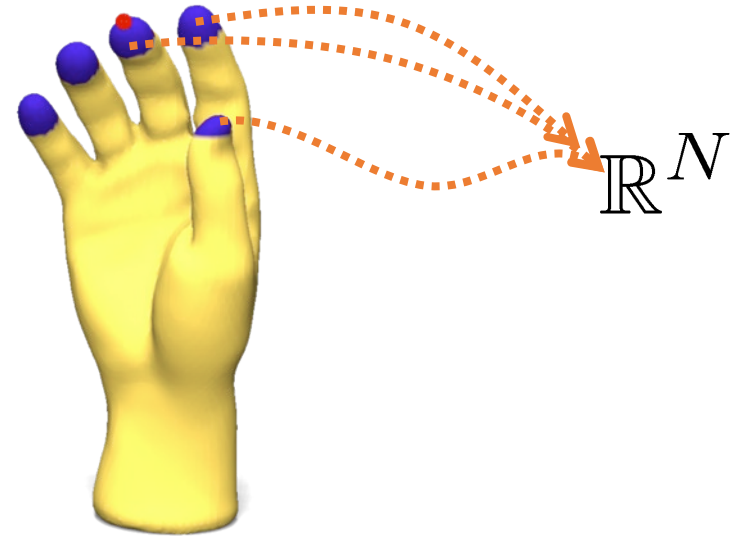
Global Matching – Invariant Features

Try to characterize the shape using properties that are invariant under the desired set of transformations.

Conflicting interests – invariance vs. informativeness.

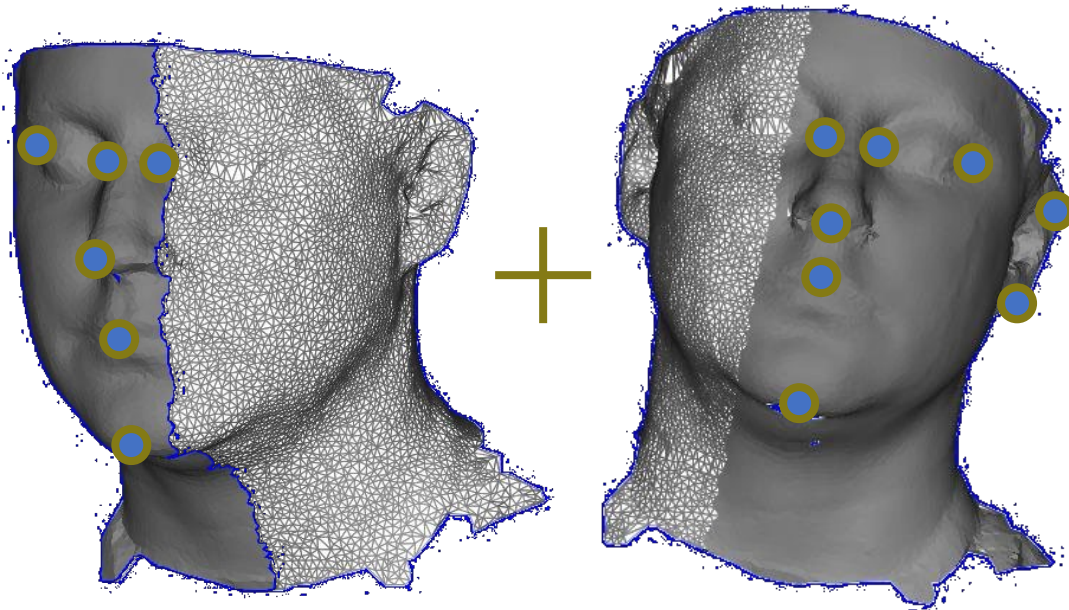
The most common pipeline:

1. identify salient feature points
2. compute informative and commensurable descriptors.



Matching Using Feature Points

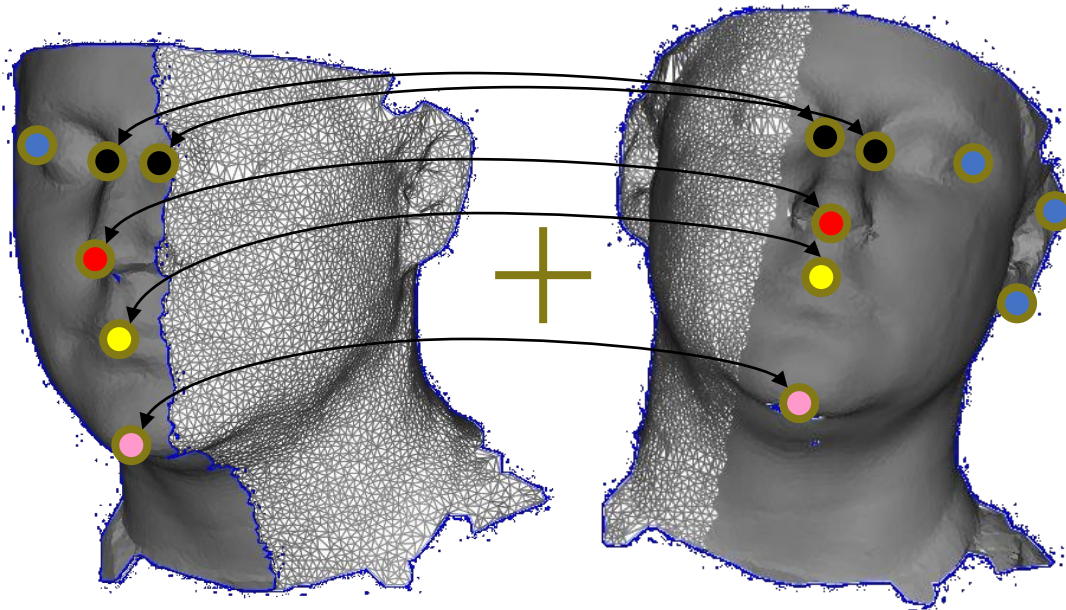
1. Find **feature points** on the two scans (we'll come back to that issue)



Partially Overlapping Scans

Approach

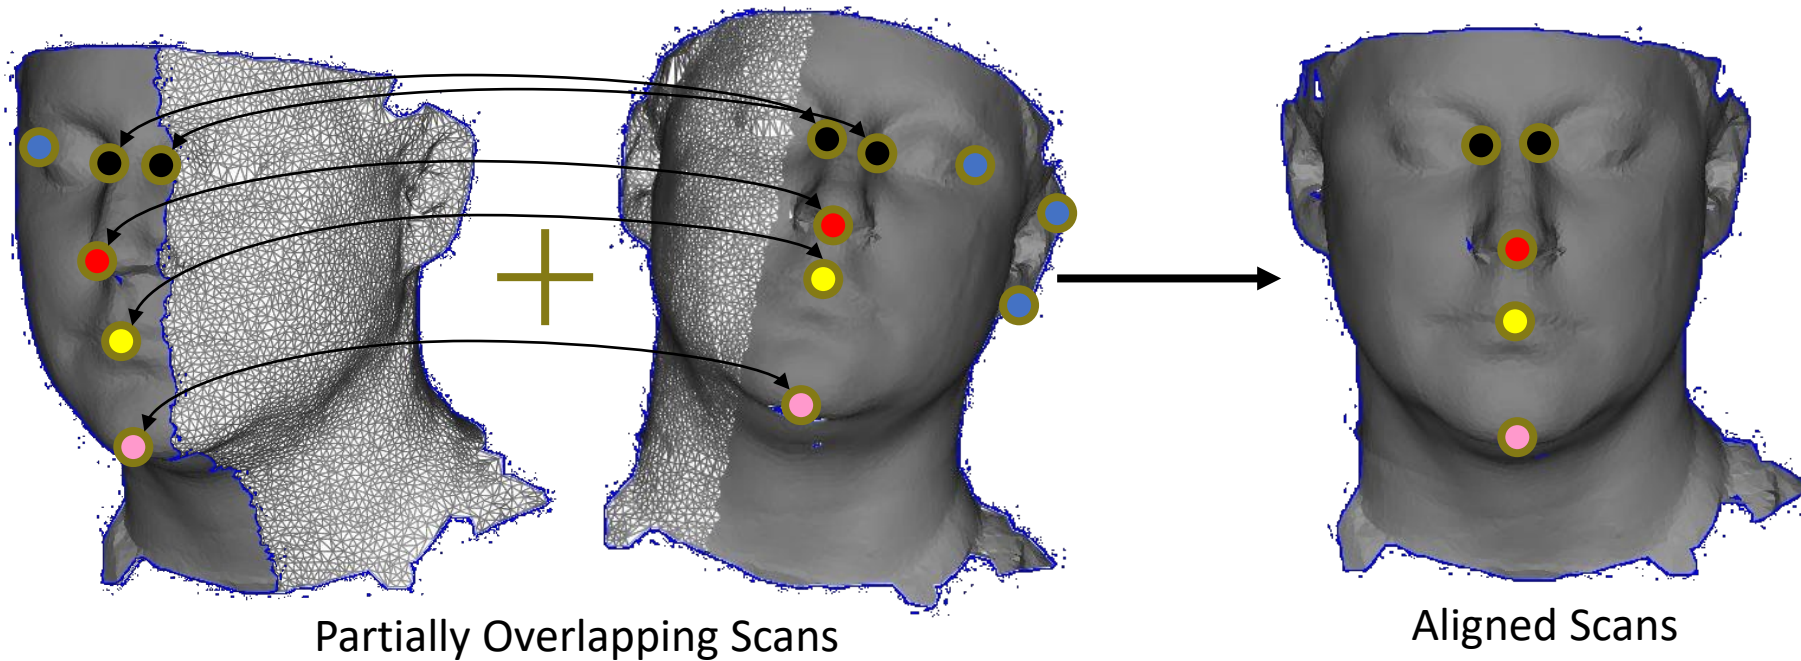
1. (Find feature points on the two scans)
2. Establish **correspondences**



Partially Overlapping Scans

Approach

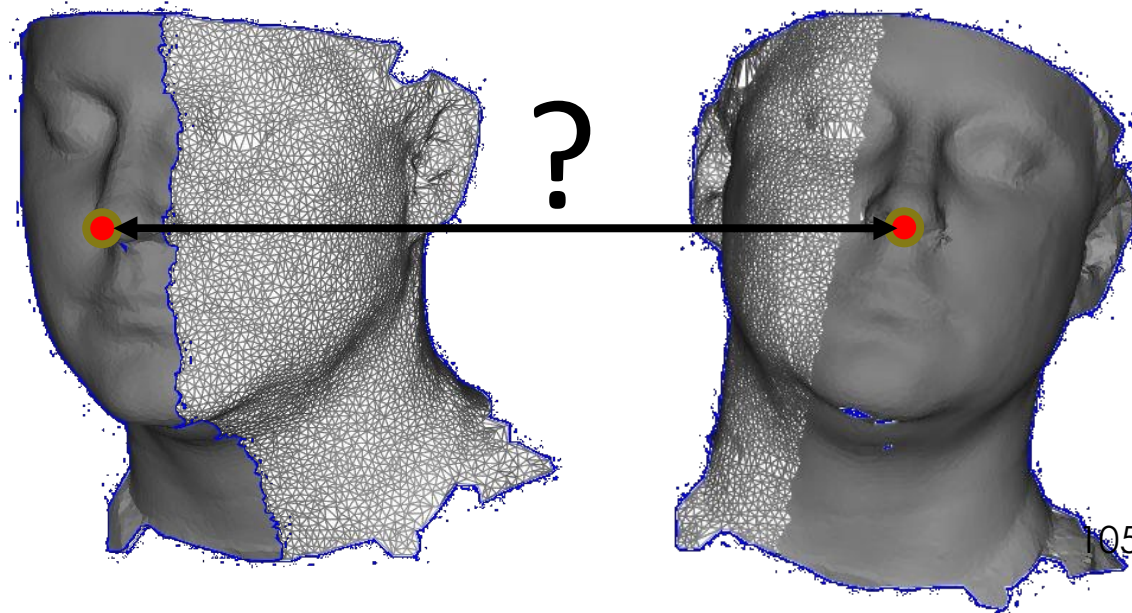
1. (Find feature points on the two scans)
2. Establish correspondences
3. Compute the aligning transformation



Correspondence

Goal:

Identify when two points on different scans represent the same feature

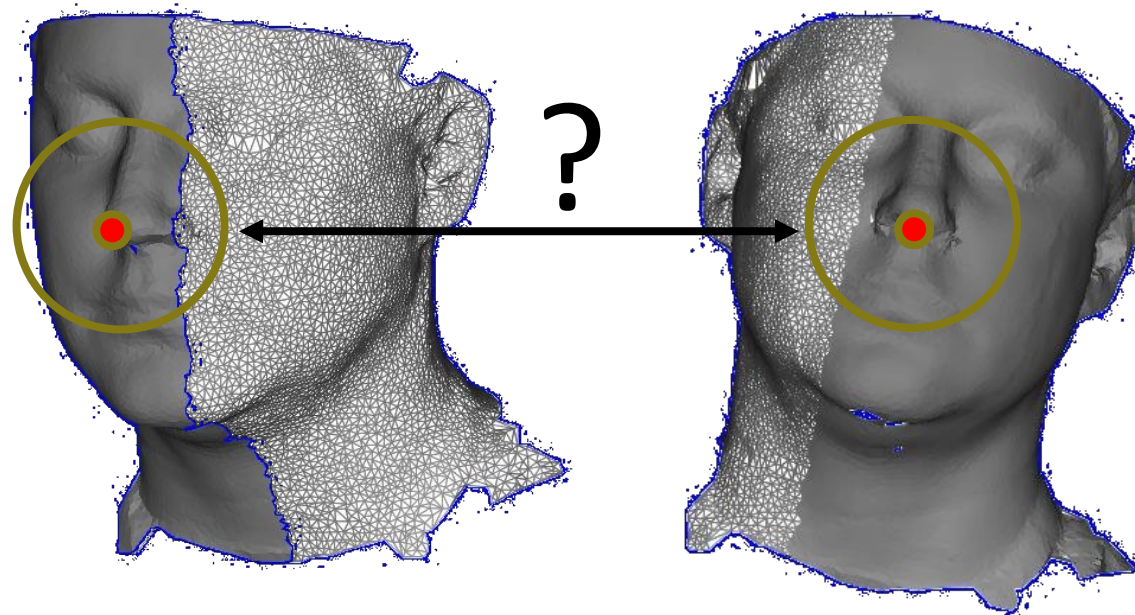


Correspondence

Goal:

Identify when two points on different scans represent the same feature:

Are the surrounding regions similar?

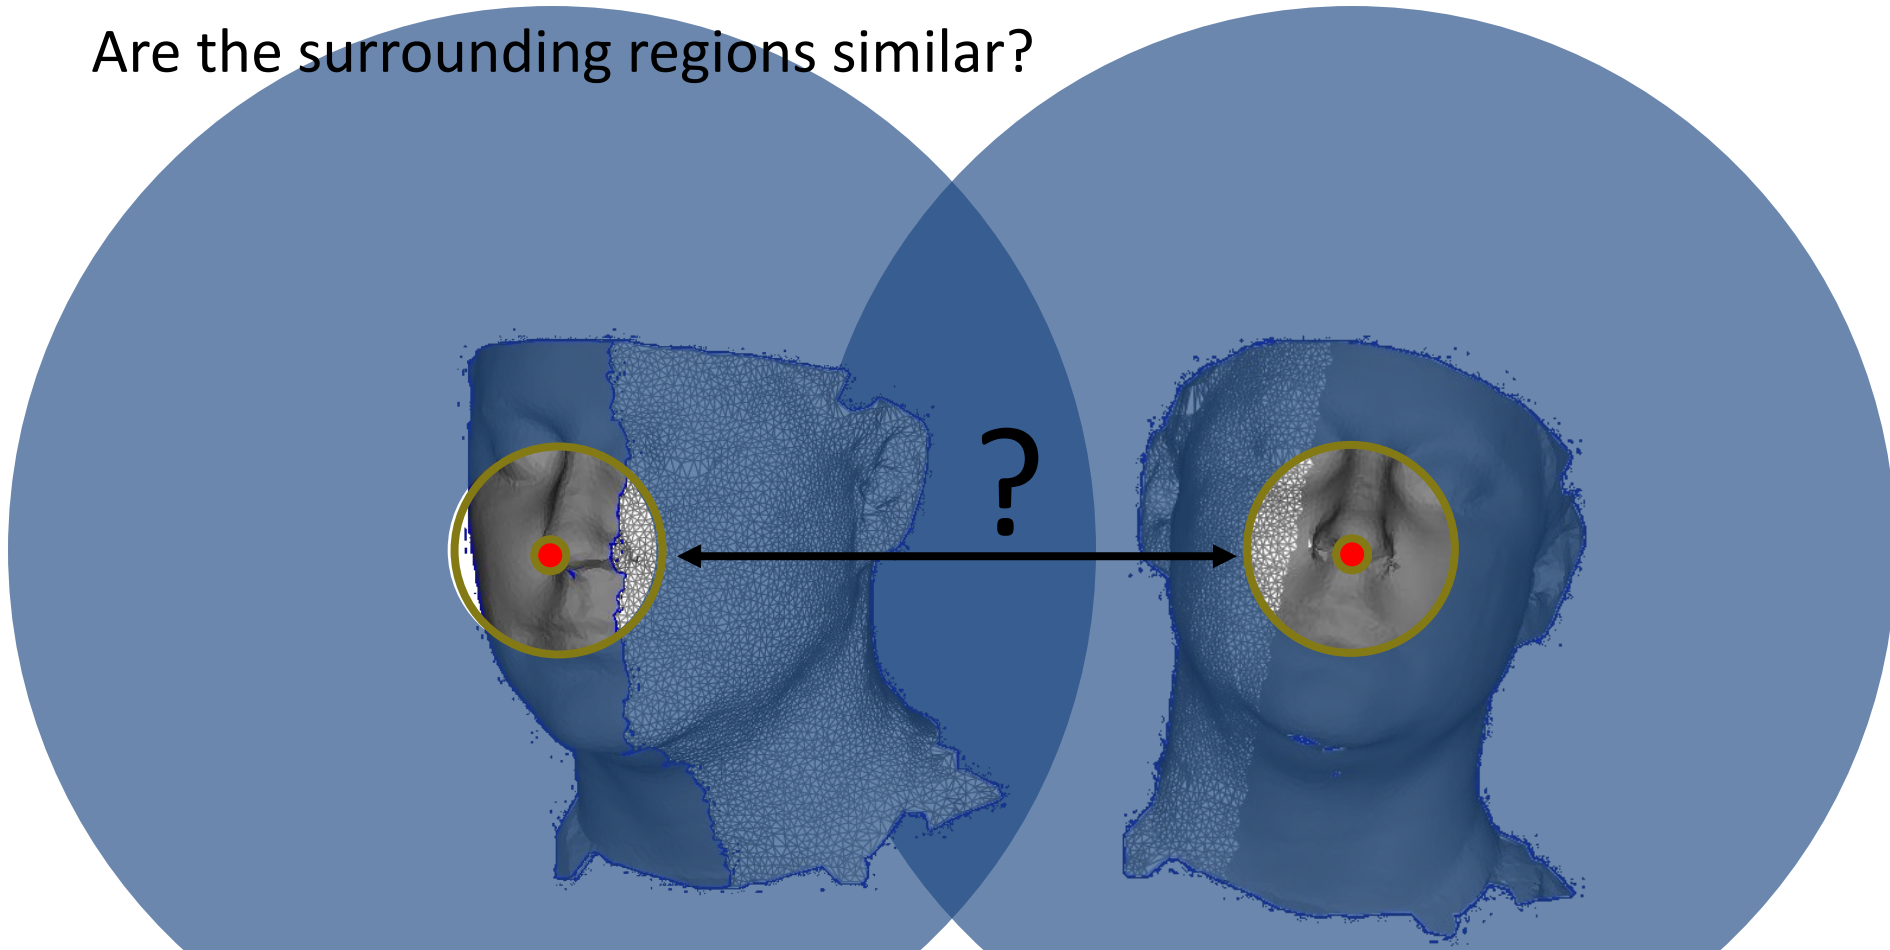


Correspondence

Goal:

Identify when two points on different scans represent the same feature:

Are the surrounding regions similar?



Main Question

How to compare regions on the shape in an invariant manner?

A large variety of *descriptors* have been suggested.

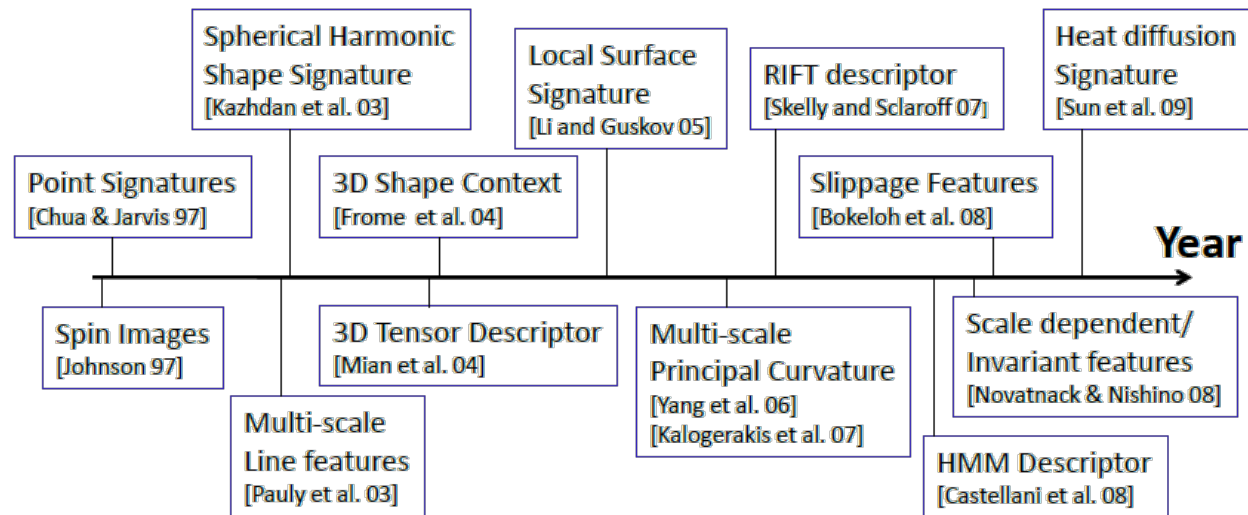


table by Will Chang

To give an example, we describe two.

Spin Images

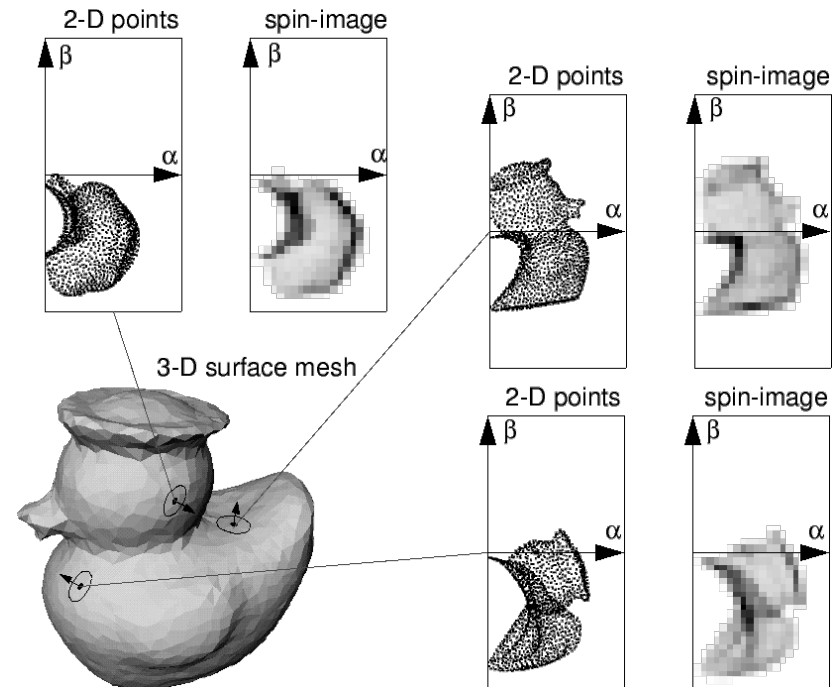
Creates an image associated with a neighborhood of a point.

Compare points by comparing their *spin images* (2D).

Given a point and a normal, every other point is indexed by two parameters:

β distance to tangent plane

α distance to normal line



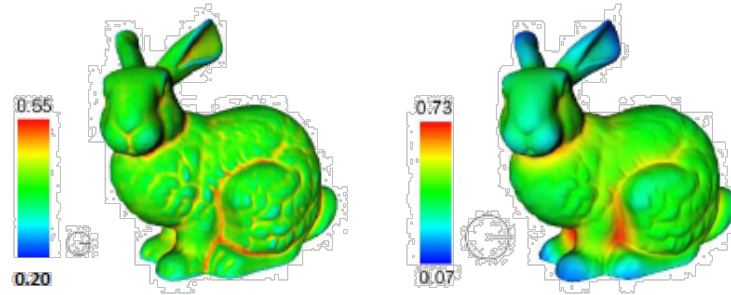
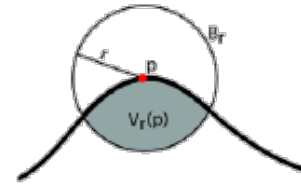
Using Spin Images for Efficient Object Recognition in Cluttered 3D Scenes
Johnson et al, PAMI 99

Integral Volume Descriptor

Integral invariant signatures, Manay et al. ECCV 2004

Integral Invariants for Robust Geometry Processing, Pottmann et al. 2007-2009

$$V_r(p) = \int_{B_r(p) \cap S} dx$$



Relation to mean curvature

$$V_r(p) = \frac{2\pi}{3}r^3 - \frac{\pi H}{4}r^4 + O(r^5)$$

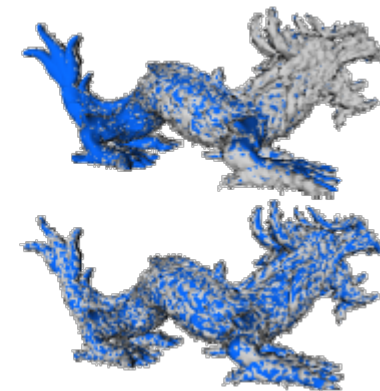
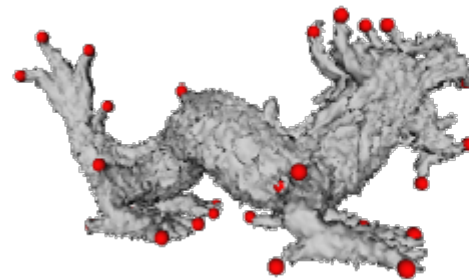
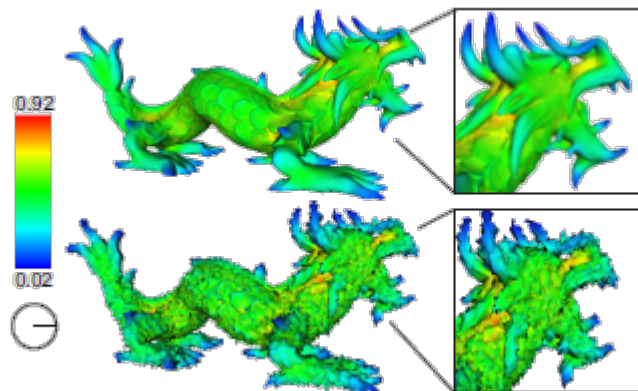
Robust Global Registration,
Gelfand et al. 2005

Feature Based Methods

Once we have a feature descriptor, we can find the most *unusual* one: feature detection.

Establish correspondences by first finding *reliable* ones. Propagate the matches everywhere.

To backtrack use branch-and bound.



Method Taxonomy

Local vs. Global

refinement (e.g. ICP) | alignment (search)

Rigid vs. Deformable

rotation, translation | general deformation

Pair vs. Collection

two shapes | multiple shapes

Conclusion

- Shape matching is an active area of research.
- Local rigid matching works well. Many approaches to global matching. Works well, depending on the domain.
- Non-rigid matching is much harder. Isometric deformation model is common and useful, but limiting.
- Research problems: other deformation models, consistent matching with many shapes, robust deformable matching.

That's All

