

Pairwise, Rigid Registration

The ICP Algorithm and Its Variants

Correspondence Problem Classification

How many meshes?

- **Two:** Pairwise registration
- More than two: multi-view registration

Initial registration available?


- **Yes:** Local optimization methods
- No: Global methods

Class of transformations?

- **Rotation and translation:** Rigid-body
- Non-rigid deformations

Pairwise Rigid Registration Goal

Align two partially-overlapping meshes given initial guess for relative transform



ICP: Iterative Closest Points

- Find correspondences
- Minimize surface-to-surface distance


Classification of ICP variants

- Faster alignment
- Better robustness

ICP as function minimization

Aligning 3D Data

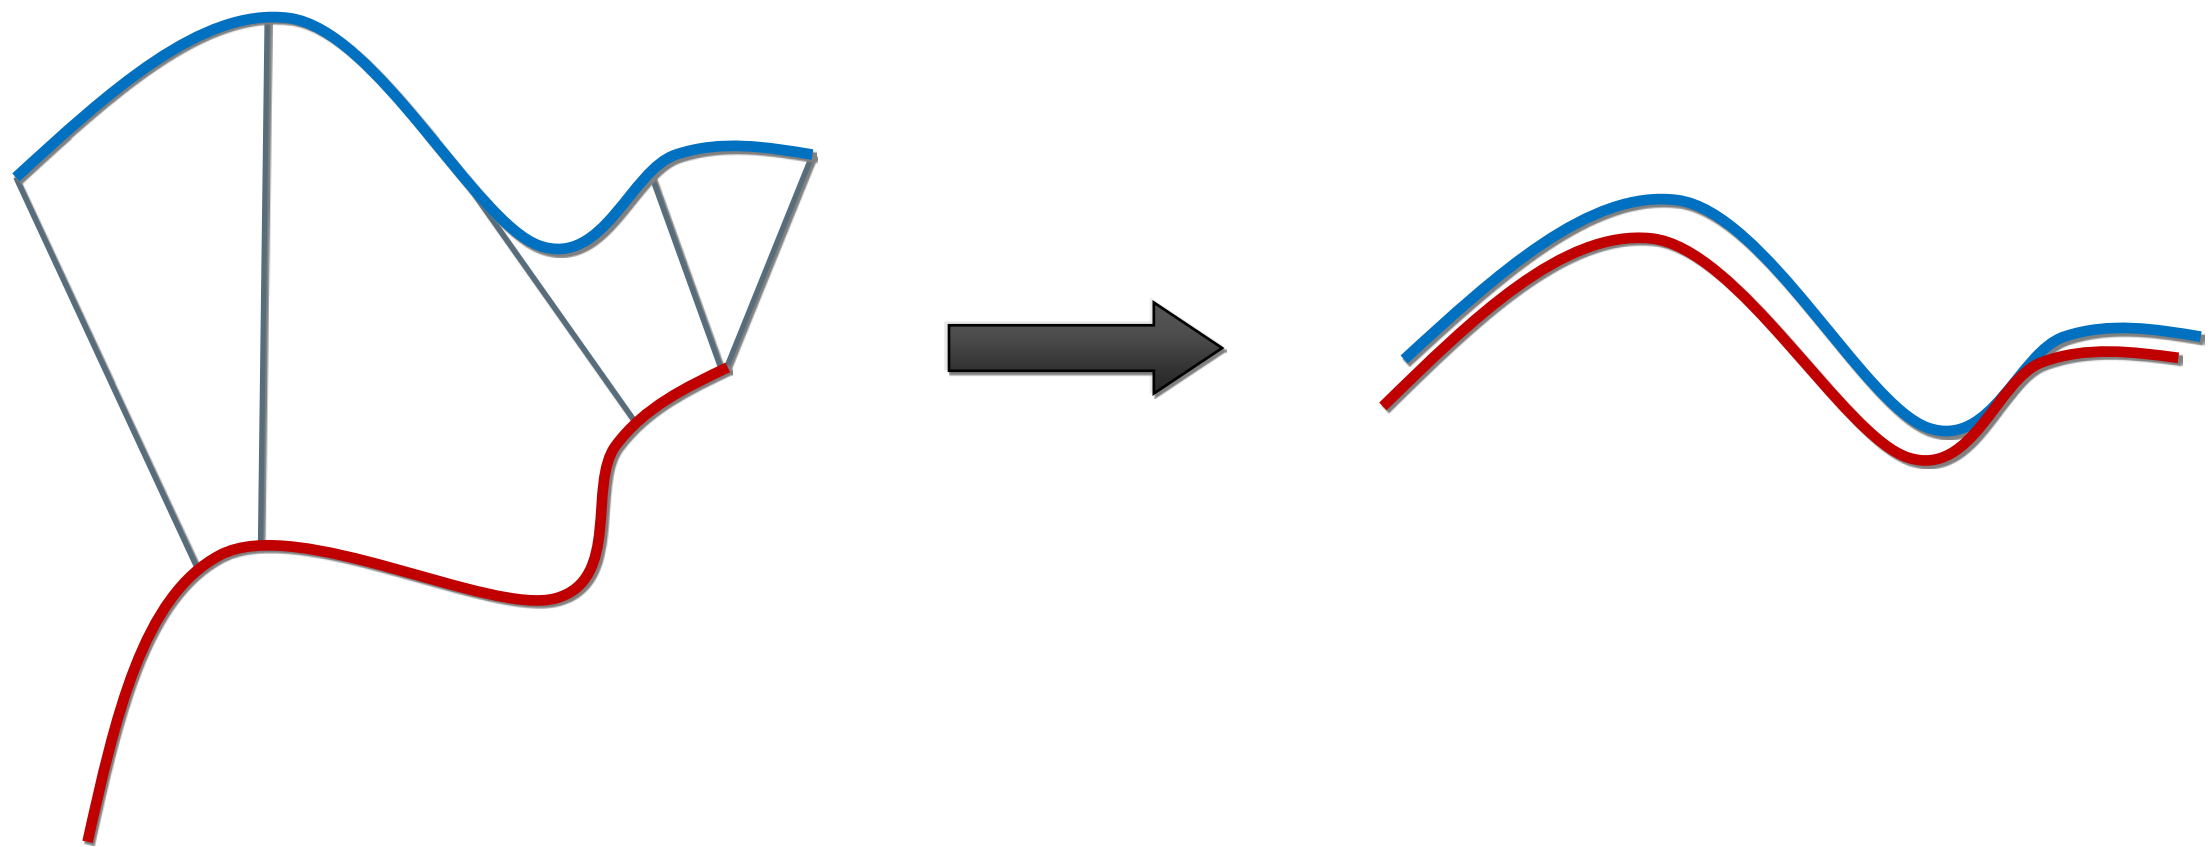
If correct correspondences are known, can find correct relative rotation/translation



How to find correspondences: User input?

Feature detection? Signatures?


Alternative: assume closest points correspond



... and iterate to find alignment

- Iterative Closest Points (ICP) [Besl & McKay 92]

Converges if starting position “close enough”



Select e.g. 1000 random points

Match each to closest point on other scan,
using data structure such as *k*-d tree

Reject pairs with distance $> k$ times median

Construct **error function**:

Minimize (closed form solution in [Horn 87])

$$E = \sum \|\mathbf{R}\mathbf{p}_i + \mathbf{t} - \mathbf{q}_i\|^2$$

Variants on the following stages of ICP have been proposed:

1. **Selecting** source points (from one or both meshes)
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. transformation

Can analyze various aspects of performance:

- Speed
- Stability
- Tolerance of noise and/or outliers
- Maximum initial misalignment



Comparisons of many variants in

[Rusinkiewicz & Levoy, 3DIM 2001]

1. Selecting source points (from one or both meshes)
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. transformation

Point-to-Plane Error Metric

Using point-to-plane distance instead of point-to-point lets flat regions slide along each other [Chen & Medioni 91]



Error function:

$$E = \sum [(\mathbf{R}\mathbf{p}_i + \mathbf{t} - \mathbf{q}_i) \cdot \mathbf{n}_i]^2$$

where \mathbf{R} is a rotation matrix, \mathbf{t} is translation vector

Linearize (i.e. assume that $\sin \theta \approx \theta$, $\cos \theta \approx 1$):

$$E \approx \sum ((\mathbf{p}_i - \mathbf{q}_i) \cdot \mathbf{n}_i + \mathbf{r} \cdot (\mathbf{p}_i \times \mathbf{n}_i) + \mathbf{t} \cdot \mathbf{n}_i)^2, \quad \text{where } \mathbf{r} = \begin{pmatrix} \mathbf{r}_x \\ \mathbf{r}_y \\ \mathbf{r}_z \end{pmatrix}$$

Result: overconstrained linear system

Overconstrained linear system

$$\mathbf{Ax} = \mathbf{b},$$

$$\mathbf{A} \equiv \begin{pmatrix} \leftarrow \mathbf{p}_1 \times \mathbf{n}_1 \Rightarrow \leftarrow \mathbf{n}_1 \Rightarrow \\ \leftarrow \mathbf{p}_2 \times \mathbf{n}_2 \Rightarrow \leftarrow \mathbf{n}_2 \Rightarrow \\ \vdots \end{pmatrix}, \quad \mathbf{x} \equiv \begin{pmatrix} r_x \\ r_y \\ r_z \\ t_x \\ t_y \\ t_z \end{pmatrix}, \quad \mathbf{b} \equiv \begin{pmatrix} -(\mathbf{p}_1 - \mathbf{q}_1) \cdot \mathbf{n}_1 \\ -(\mathbf{p}_2 - \mathbf{q}_2) \cdot \mathbf{n}_2 \\ \vdots \end{pmatrix}$$

Solve using least squares

$$\mathbf{A}^T \mathbf{Ax} = \mathbf{A}^T \mathbf{b}$$

$$\mathbf{x} \equiv (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b}$$

Can select variants to improve likelihood of reaching correct local optimum

1. Selecting source points (from one or both meshes)
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. transformation

Closest point often a bad approximation to corresponding point

Can improve matching effectiveness by restricting match to compatible points

- Compatibility of colors [Godin et al. 94]
- Compatibility of normals [Pulli 99]
- Other possibilities: curvatures, higher-order derivatives, and other local features

1. **Selecting** source points (from one or both meshes)
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. transformation

Use all points


Uniform subsampling

Random sampling


Stable sampling [Gelfand et al. 2003]

- Select samples that constrain all degrees of freedom of the rigid-body transformation

Stable Sampling



Uniform Sampling



Stable Sampling

Covariance Matrix


Aligning transform is given by $\mathbf{A}^T \mathbf{A} \mathbf{x} = \mathbf{A}^T \mathbf{b}$, where

$$\mathbf{A} \equiv \begin{pmatrix} \leftarrow \mathbf{p}_1 \times \mathbf{n}_1 \rightarrow & \leftarrow \mathbf{n}_1 \rightarrow \\ \leftarrow \mathbf{p}_2 \times \mathbf{n}_2 \rightarrow & \leftarrow \mathbf{n}_2 \rightarrow \\ \vdots & \vdots \end{pmatrix}, \quad \mathbf{x} \equiv \begin{pmatrix} r_x \\ r_y \\ r_z \\ t_x \\ t_y \\ t_z \end{pmatrix}, \quad \mathbf{b} \equiv \begin{pmatrix} -(\mathbf{p}_1 - \mathbf{q}_1) \cdot \mathbf{n}_1 \\ -(\mathbf{p}_2 - \mathbf{q}_2) \cdot \mathbf{n}_2 \\ \vdots \end{pmatrix}$$

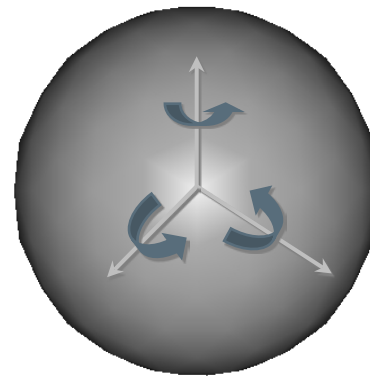
Covariance matrix $\mathbf{C} = \mathbf{A}^T \mathbf{A}$ determines the change in error when surfaces are moved from optimal alignment

Sliding Directions


Eigenvectors of \mathbf{C} with small eigenvalues correspond to sliding transformations




3 small eigenvalues
2 translation
1 rotation




3 small eigenvalues
3 rotation



2 small eigenvalues
1 translation
1 rotation

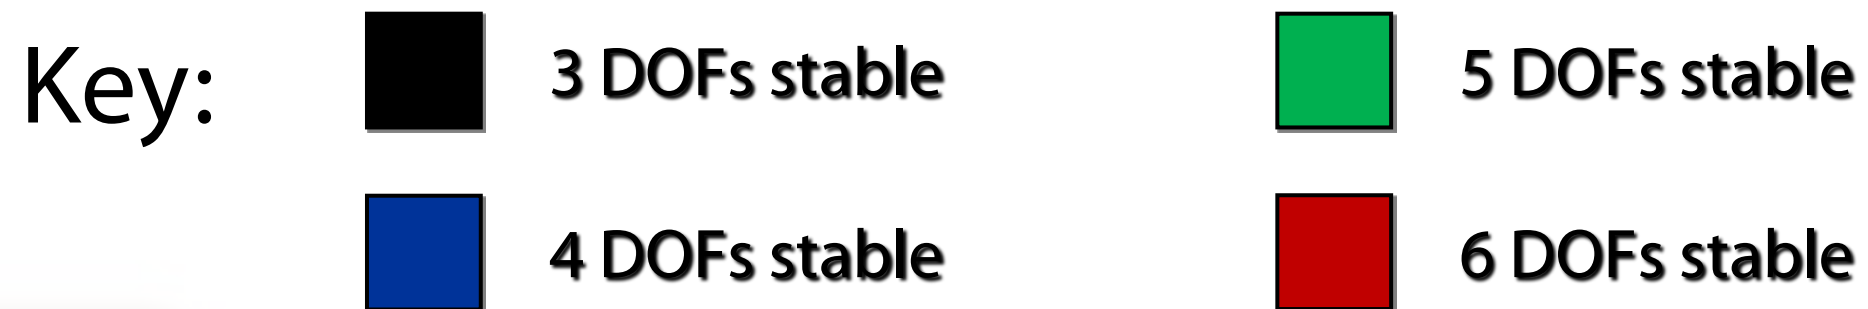



1 small eigenvalue
1 rotation



1 small eigenvalue
1 translation

Stability Analysis




Select points to prevent small eigenvalues

- Based on \mathbf{C} obtained from sparse sampling


Simpler variant: normal-space sampling

- Select points with uniform distribution of normals
- **Pro:** faster, does not require eigenanalysis
- **Con:** only constrains translation

Stability-based or normal-space sampling
important for smooth areas with small features



Random sampling



Normal-space sampling

Selection vs. Weighting

Could achieve same effect with weighting

Hard to ensure enough samples in features except at high sampling rates

However, have to build special data structure

Preprocessing / run-time cost tradeoff


Projection-based matching

1. Selecting source points (from one or both meshes)
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. transformation

Finding Corresponding Points

Finding closest point is most expensive stage of the ICP algorithm


- Brute force search – $O(n)$
- Spatial data structure (e.g., k-d tree) – $O(\log n)$



Projection to Find Correspondences

Idea: use a simpler algorithm to find correspondences
For range images, can simply project point [Blais 95]


- Constant-time
- Does not require precomputing a spatial data structure



Slightly worse performance per iteration

Each iteration is one to two orders of magnitude faster than closest-point

Result: can align two range images in a few milliseconds, vs. a few seconds




Given:

- A scanner that returns range images in real time
- Fast ICP
- Real-time merging and rendering

Result: 3D model acquisition

- Tight feedback loop with user
- Can see and fill holes while scanning

Scanner Layout




Photograph






Real-Time Result



Real-Time Result



One way of studying performance is via empirical tests on various scenes

How to analyze performance analytically?

For example, when does point-to-plane help?


Under what conditions does projection-based matching work?

What Does ICP Do?

Two ways of thinking about ICP:

- Solving the correspondence problem (expectation maximization)
- **Minimizing point-to-surface squared distance**

ICP is like (Gauss-) Newton method on an approximation of the distance function




What Does ICP Do?

Two ways of thinking about ICP:

- Solving the correspondence problem (expectation maximization)
- **Minimizing point-to-surface squared distance**

ICP is like (Gauss-) Newton method on an approximation of the distance function




What Does ICP Do?

Two ways of thinking about ICP:


- Solving the correspondence problem (expectation maximization)
- **Minimizing point-to-surface squared distance**

ICP is like (Gauss-) Newton method on an approximation of the distance function


- ICP variants affect shape of global error function or local approximation




Point-to-Surface Distance




Point-to-Point Distance



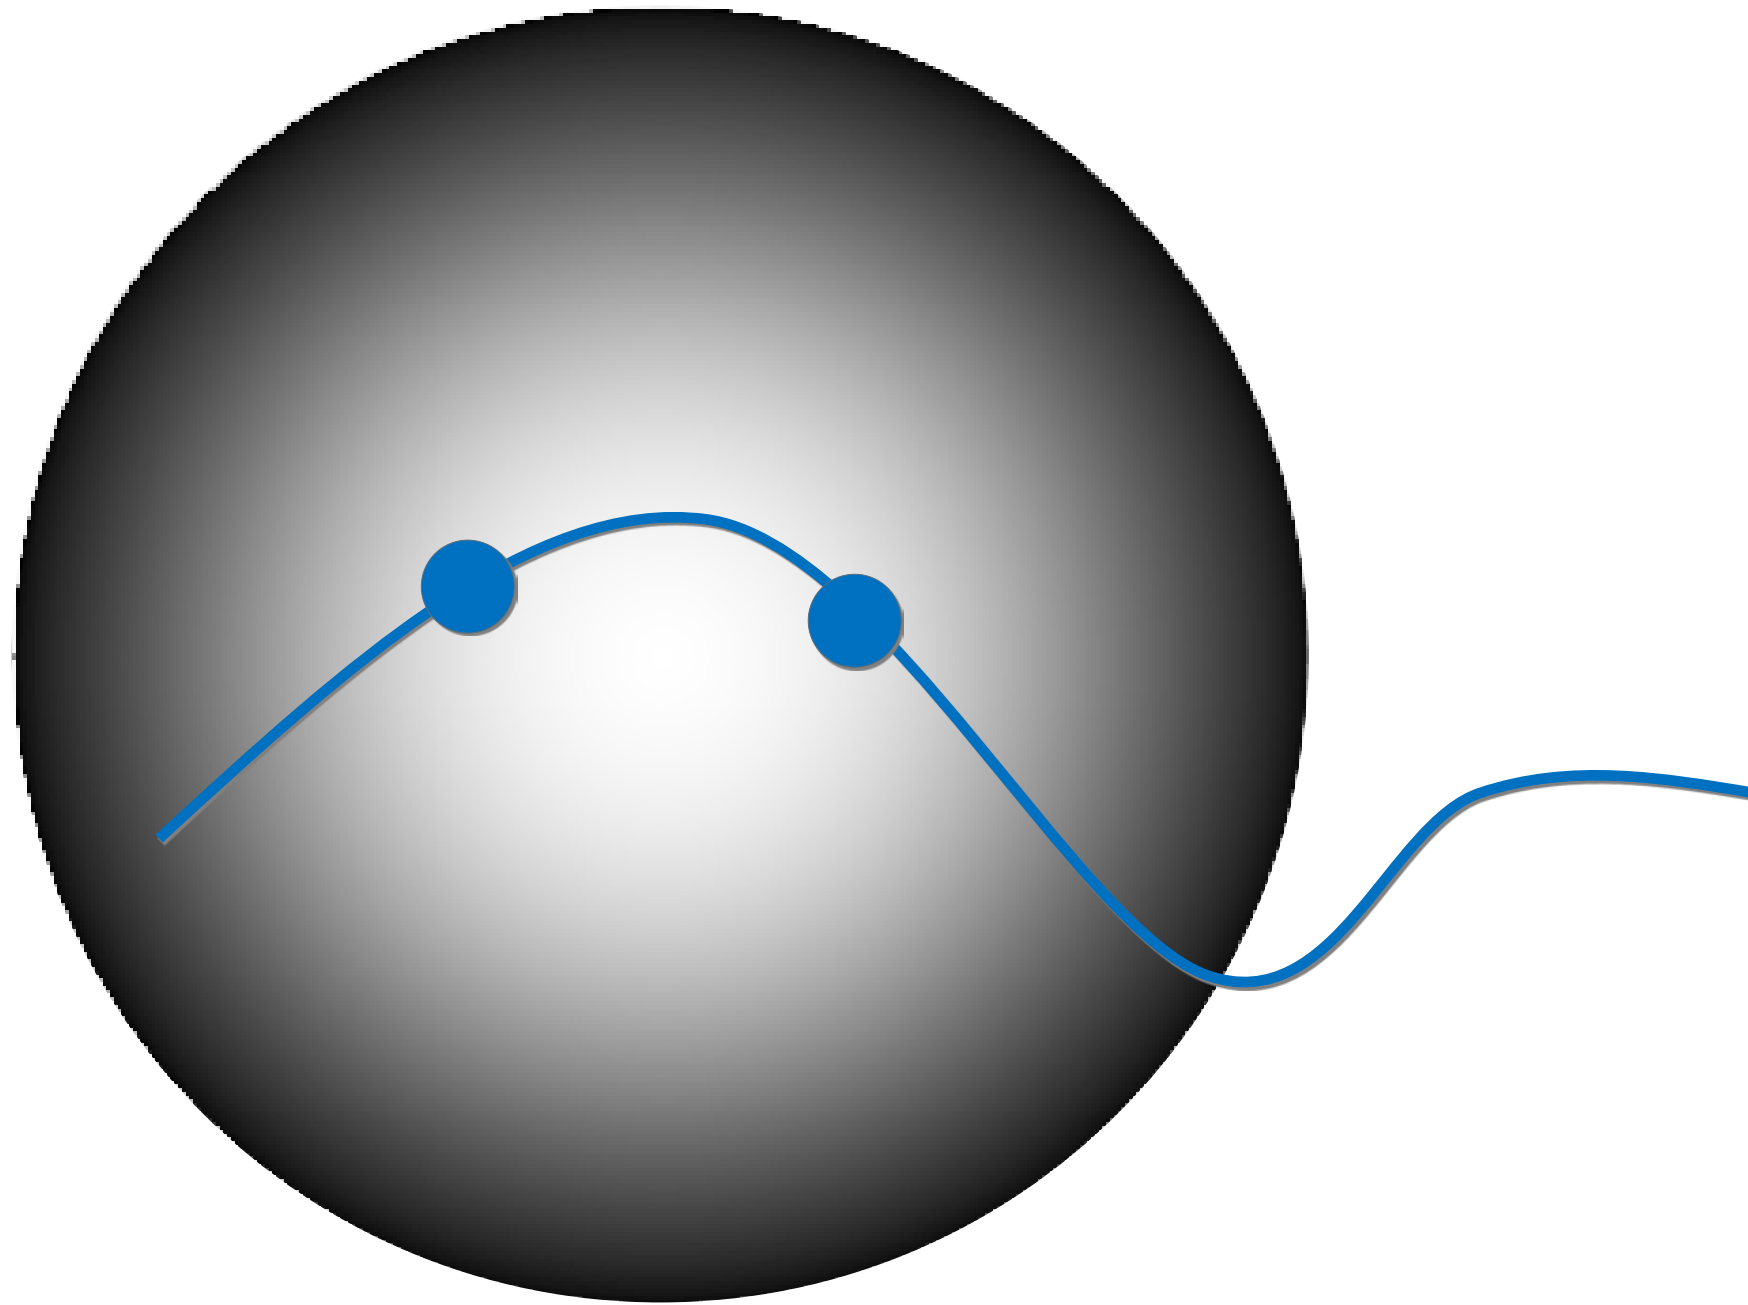
Point-to-Plane Distance



Point-to-Multiple-Point Distance



Point-to-Multiple-Point Distance



Soft Matching and Distance Functions

Soft matching equivalent to standard ICP on (some) filtered surface

Produces filtered version of distance function
⇒ fewer local minima


Multiresolution minimization [Turk & Levoy 94]
or softassign with simulated annealing
(good description in [Chui 03])

Mitra et al.'s Optimization


[Mitra et al. 2004]

Precompute piecewise-quadratic approximation
to distance field throughout space

Store in "d2tree" data structure



2D



3D

Mitra et al.'s Optimization


Precompute piecewise-quadratic approximation to distance field throughout space

Store in “d2tree” data structure


At run time, look up quadratic approximants and optimize using Newton's method

- Often fewer iterations, but more precomputation
- More robust, wider basin of convergence


Convergence Funnel



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




Converges





Does not converge

Convergence Funnel



Plane-to-plane ICP



distance-field
formulation

ICP: Prototypical local alignment algorithm

Either:

- Find unknown correspondences using expectation maximization

or

- Find unknown transformation by iteratively minimizing surface-to-surface distance