



Eurographics 2012

Cagliari, Italy

May 13 -18



33rd ANNUAL CONFERENCE OF THE EUROPEAN ASSOCIATION FOR COMPUTER GRAPHICS

Dynamic Geometry Processing

EG 2012 Tutorial



Local, Rigid, Pairwise

The ICP algorithm and its extensions

Niloy J. Mitra



University College London

Geometric Matching

 M_1  M_2 

$$M_1 \approx T(M_2)$$



Matching with Translation

 M_1  M_2 

$$M_1 \approx T(M_2)$$

T : translation



Matching with Rigid Transforms

 M_1  M_2 

$$M_1 \approx T(M_2)$$

T : translation + rotation



Partial Matching

 M_1  M_2 


$$M_1 \approx T(M_2)$$

T : translation + rotation

Local vs. Global Matching



global registration
any rigid transform



local registration
nearly aligned

Given M_1, \dots, M_n , find T_2, \dots, T_n such that

$$M_1 \approx T_2(M_2) \cdots \approx T_n(M_n)$$

ICP: Local, partial, rigid transforms

How many point-pairs are needed to *uniquely* define a rigid transform?

$$\mathbf{p}_1 \rightarrow \mathbf{q}_1$$

$$\mathbf{p}_2 \rightarrow \mathbf{q}_2$$


$$\mathbf{p}_3 \rightarrow \mathbf{q}_3$$

$$R\mathbf{p}_i + t \approx \mathbf{q}_i$$

Correspondence problem: $\mathbf{p}_i \overset{?}{\rightarrow} \mathbf{q}_j$

Pairwise Rigid Registration Goal

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



Outline

ICP: Iterative Closest Points

Classification of ICP variants


- Faster alignment
- Better robustness

ICP as function minimization

Thin-plate spline (non-rigid ICP)

Aligning 3D Data

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




Aligning 3D Data

How to find correspondences:

User input?

Feature detection?


Signatures?



Aligning 3D Data

Assume: Closest points as corresponding

$$\mathbf{p}_i \rightarrow \mathcal{C}(\mathbf{p}_i)$$




Aligning 3D Data

... and iterate to find alignment

Iterative Closest Points (ICP) [Besl and McKay 92]

Converges if starting poses are *close enough*



Basic ICP

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

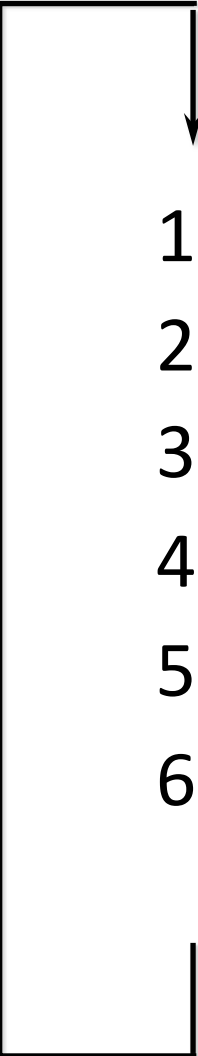
Construct **error function**:

$$E := \sum_i (R\mathbf{p}_i + t - \mathbf{q}_i)^2$$

Minimize (closed form solution in [Horn 87])

ICP Variants

Variants of basic ICP

- 
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

Performance of Variants

Can analyze various aspects of performance:

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

Comparisons of many variants of ICP
[Rusinkiewicz & Levoy, 3DIM 2001]

ICP Variants


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 allows flat regions slide along each other


[Chen and Medioni 91]



Point-to-Plane Error Metric



point-to-point



point-to-plane

Point-to-Plane Error Metric

Error function:

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

where R is a rotation matrix, t is translation vector

$$\mathbf{p}_i \rightarrow R\mathbf{p}_i + t$$

$$\mathbf{p}_i \rightarrow \bar{\mathbf{c}} + \mathbf{p}_i \times \mathbf{c}$$

Point-to-Plane Error Metric

Overconstrained linear system

$$\mathbf{A}x = b,$$

$$\mathbf{A} = \begin{pmatrix} \leftarrow & p_1 \times n_1 & \rightarrow & \leftarrow & n_1 & \rightarrow \\ \leftarrow & p_2 \times n_2 & \rightarrow & \leftarrow & n_2 & \rightarrow \\ \vdots & & & \vdots & & \end{pmatrix}, \quad x = \begin{pmatrix} r_x \\ r_y \\ r_z \\ t_x \\ t_y \\ t_z \end{pmatrix}, \quad b = \begin{pmatrix} -(p_1 - q_1) \cdot n_1 \\ -(p_2 - q_2) \cdot n_2 \\ \vdots \end{pmatrix}$$

Solve using least squares

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

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

Improving ICP Stability

Closest *compatible* point

Stable sampling

ICP Variants

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 Compatible Point

Closest points are often bad as corresponding points

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

ICP Variants



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

Selecting Source Points

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 $A^T A x = A^T b$, where

$$\mathbf{A} = \begin{pmatrix} \leftarrow & p_1 \times n_1 & \rightarrow & \leftarrow & n_1 & \rightarrow \\ \leftarrow & p_2 \times n_2 & \rightarrow & \leftarrow & n_2 & \rightarrow \\ & \vdots & & & \vdots & \end{pmatrix}, \quad \mathbf{x} = \begin{pmatrix} r_x \\ r_y \\ r_z \\ t_x \\ t_y \\ t_z \end{pmatrix}, \quad \mathbf{b} = \begin{pmatrix} -(p_1 - q_1) \cdot n_1 \\ -(p_2 - q_2) \cdot n_2 \\ \vdots \end{pmatrix}$$


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 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



Key:



3 DOFs stable



4 DOFs stable



5 DOFs stable



6 DOFs stable

Sample Selection

Select points to prevent small eigenvalues


- Based on 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

Result

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

Improving ICP Speed

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




Projection-Based Matching

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*



Application


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



Theoretical Analysis of ICP Variants

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 Work?

Two ways of thinking about ICP:

- Solving the correspondence problem
- Minimizing point-to-surface squared distance

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




What Does ICP Work?

Two ways of thinking about ICP:

- Solving the correspondence problem
- 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
- **Minimizing point-to-surface squared distance**

ICP is like Newton's 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])


Distance-field Based Optimization

Precompute piecewise-quadratic approximation to distance field throughout space

Store in "d2tree" data structure



2D



3D

Distance-field Based 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

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

Convergence Funnel



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

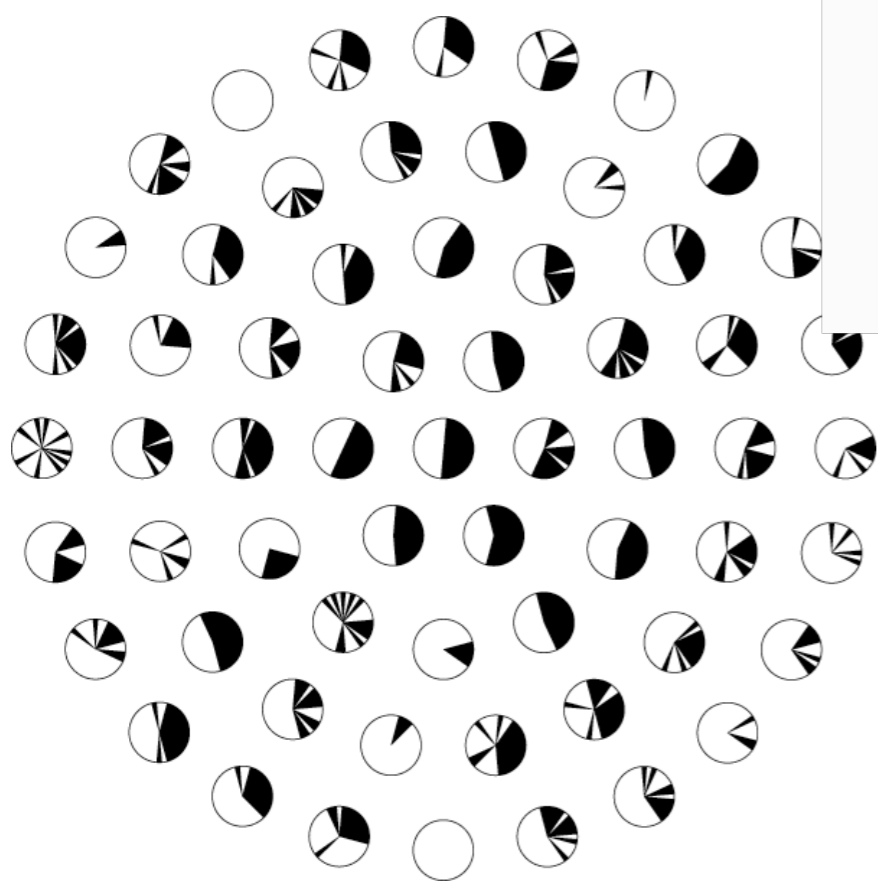


Converges





Does not converge

Convergence Funnel



Plane-to-plane ICP



distance-field
formulation


Non-rigid ICP

Thin plate spline [Bookstein '89]


$$J = \int \left(\sum_{i,j} \mathcal{S}_{\mathbf{f}_i \mathbf{f}_j}^2 \right) d\mathbf{f}_1 \dots d\mathbf{f}_n$$

Minimize bending energy (second order partial derivatives)


Affine transforms are linear and hence do not contribute to J.




(a) Rigid Pupil



(b) Non-Rigid




(c) Rigid Hair



(d) Non-Rigid

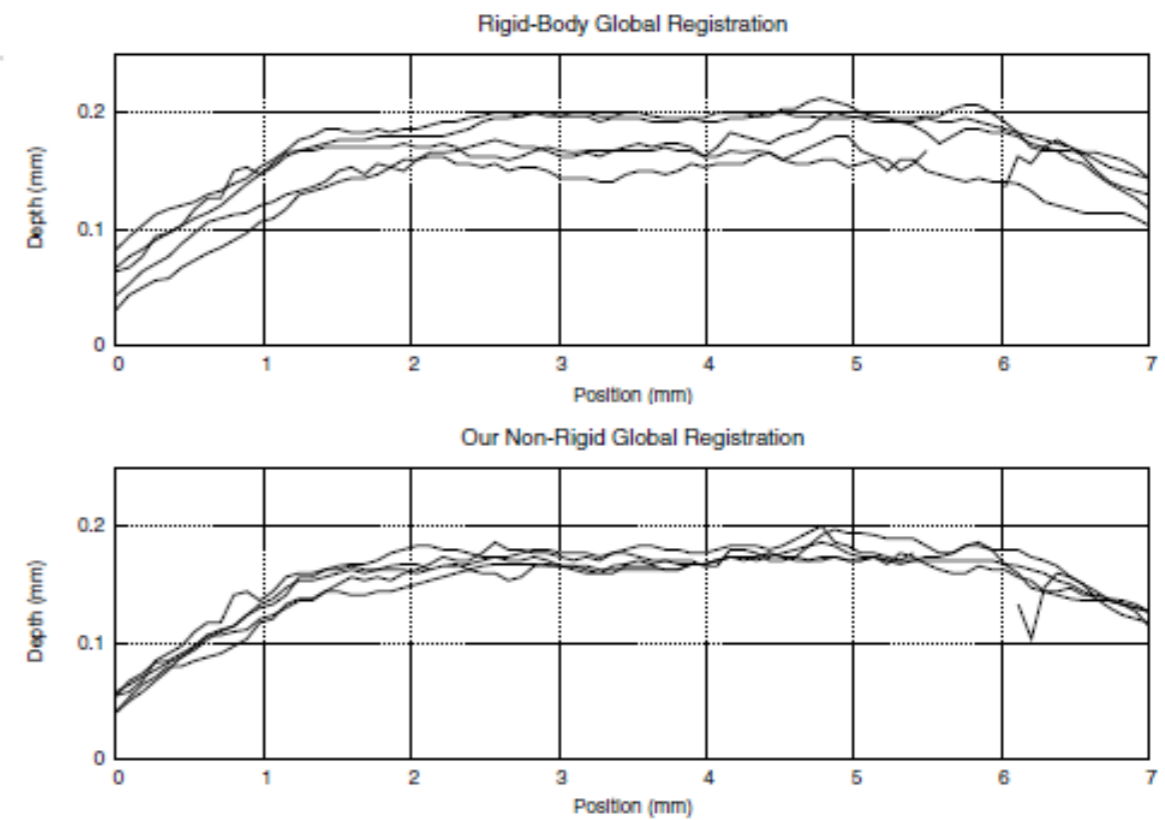
Non-rigid Registration



Rigid



Non-Rigid



[Brown and Rusinkiewicz, 2007]