Non-Rigid Range-Scan Alignment Using Thin-Plate Splines

Benedict J. Brown   Szymon Rusinkiewicz

Princeton University
{bjbrown,smr}@cs.princeton.edu

9 September 2004

Some Large Scanning Projects

Digital Michelangelo

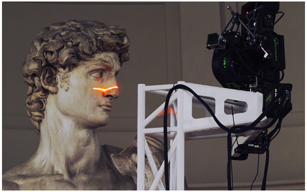

Stanfod Graphics Group

Forma Urbis Romae

Stanford Graphics Group

Great Buddha

CVL, University of Tokyo

Florentine Pietà

IBM Research

Range Scanning Pipeline

Acquisition

- Scanners acquire data from a single viewpoint

Range Scanning Pipeline

Acquisition

Alignment

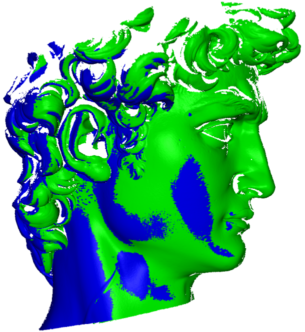

Range Scanning Pipeline

Acquisition

Alignment

Merging

Range Scanning Pipeline

Acquisition

↓

Alignment

↓

Merging

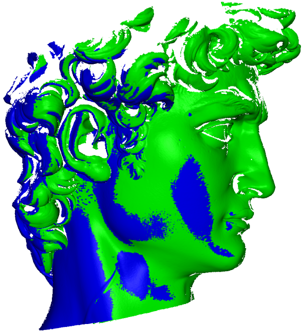

Alignment Problems

- Existing alignment algorithms fail
- Non-rigid calibration error

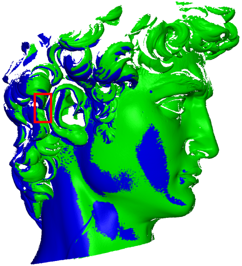

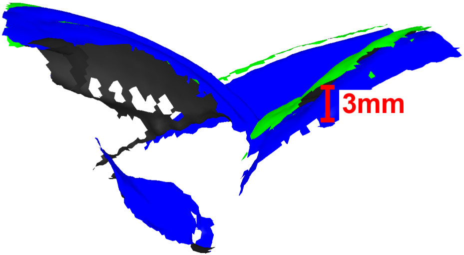

Digital Michelangelo Project [Levoy00]

- Scan Michelangelo's statues in Florence using a laser range finder
- .25mm precision over entire statue
- David is 5m high
- Gantry is 7.6m high
- Extreme requirements cause unavoidable calibration errors

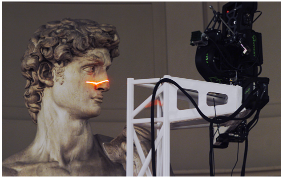

Paul Debevec

Stanford Graphics Group

Overview

- Previous Work
  - Iterative Closest Points (Rigid-Body Alignment)
  - Hierarchical ICP
- Non-Rigid Alignment
  - Thin-Plate Splines
- Feature Correspondences
  - Piecewise ICP
- Pairwise Non-Rigid Alignment
- Results & Future Work

Rigid-Body Alignment

- Range scans are of a single, rigid object
  - Rigid-body transormation should yield exact aligment

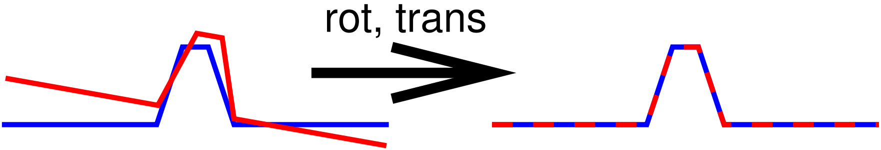

Rigid-Body Alignment

- Range scans are of a single, rigid object
  - Rigid-body transormation should yield exact aligment
- Not possible even with a perfectly calibrated camera
  - Noise and irregular sampling

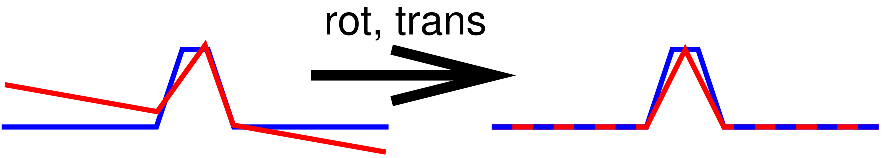

Rigid-Body Alignment

- Range scans are of a single, rigid object
  - Rigid-body transormation should yield exact aligment
- Not possible even with a perfectly calibrated camera
  - Noise and irregular sampling
- Instead find least squares fit

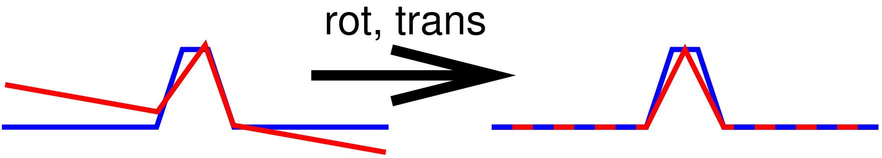

Iterative Closest Points [Besl92]

- To fit two meshes, need correspondence between points
  - Assume points correspond to closest points on other mesh
  - Compute best fit on a subset of all points
- If starting point was good, result should be better
  - Iterate until fit converges to minimum error

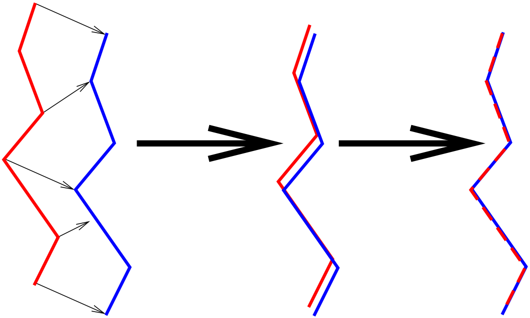

Advantages of ICP

- Efficient (seconds per pair of meshes)

Advantages of ICP

- Efficient (seconds per pair of meshes)
- Robust and stable, especially with advanced sampling
  - Select samples which constrain ICP [Gelfand03]
  - Sample features heavily

Advantages of ICP

- Efficient (seconds per pair of meshes)
- Robust and stable, especially with advanced sampling
  - Select samples which constrain ICP [Gelfand03]
  - Sample features heavily
- Instead of point-to-point distance, use point-to-plane [Chen92]
  - Features lock onto each other, while flat areas can slide freely
  - Convergence is both more stable and faster

Hierarchical ICP [Ikemoto03]

- Dice meshes into small pieces
  - Do global alignment on all pieces
- Neighboring pieces must contain some overlap
  - Too little overlap leads to discontinuities
  - Too much overlap prevents freedom in warping
- Not smooth, slow (hours per alignment)

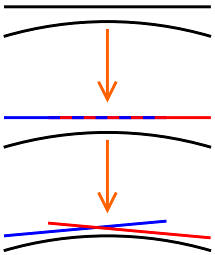

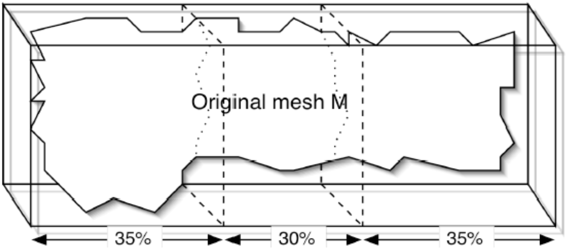

Ikemoto *et al.*

Overview

- Previous Work
  - Iterative Closest Points (Rigid-Body Alignment)
  - Hierarchical ICP
- Non-Rigid Alignment
  - Thin-Plate Splines
- Feature Correspondences
  - Piecewise ICP
- Pairwise Non-Rigid Alignment
- Results & Future Work

Non-Rigid Alignment

- Calibration error is
  - low-frequency $\Rightarrow$ smooth, slowly-varying function
  - hard to characterize, so need flexible function
- Use non-rigid warp to compensate for calibration error
- Represent with thin-plate splines [Bookstein89] [Wahba90]

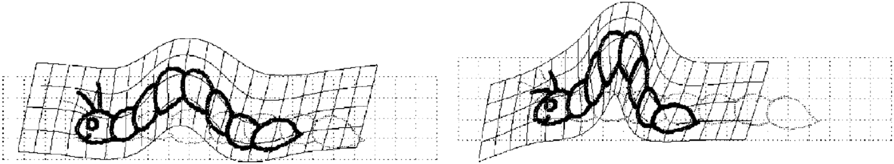

Chui & Rangarajan

Thin-Plate Splines

- Maps source point set X to target set Y: $f(x_i) = y_i$

Thin-Plate Splines

- Maps source point set X to target set Y: $f(x_i) = y_i$
- Minimizes *bending energy*:

$$J = \iiint f_{xx}^2 + f_{yy}^2 + f_{zz}^2 + 2(f_{xy}^2 + f_{yz}^2 + f_{zx}^2)\,\mathrm{d}x\,\mathrm{d}y\,\mathrm{d}z$$

Thin-Plate Splines

- Maps source point set X to target set Y: $f(x_i) = y_i$
- Minimizes *bending energy*:

$$J = \iiint f_{xx}^2 + f_{yy}^2 + f_{zz}^2 + 2(f_{xy}^2 + f_{yz}^2 + f_{zx}^2)\,\mathrm{d}x\,\mathrm{d}y\,\mathrm{d}z$$

- Gives the “minimal” deformation from an affine transformation necessary to map X onto Y

Thin-Plate Splines

- Maps source point set X to target set Y: $f(x_i) = y_i$
- Minimizes *bending energy*:

$$J = \iiint f_{xx}^2 + f_{yy}^2 + f_{zz}^2 + 2(f_{xy}^2 + f_{yz}^2 + f_{zx}^2)\,\mathrm{d}x\,\mathrm{d}y\,\mathrm{d}z$$

- Gives the "minimal" deformation from an affine transformation necessary to map X onto Y
- Calculate by minimizing energy functional

$$E_{TPS} = \sum_{i=1}^{n} |y_i - f(x_i)|^2 + n\lambda J$$

for a fixed λ.

Thin-Plate Splines

$$E_{TPS} = \sum_{i=1}^{n} |y_i - f(x_i)|^2 + n\lambda J$$

- Interpolates control points while minimizing warp
  - Reduces to affine transformation when that is sufficient
- E_{TPS} is minimized by a linear system of equations
- λ provides tradeoff of warp smoothness and interpolation
  - λ corresponds to the measurement variance
  - To achieve good alignment, we must have low variance
    - $\Rightarrow \lambda \approx 0$

Overview

- Previous Work
  - Iterative Closest Points (Rigid-Body Alignment)
  - Hierarchical ICP
- Non-Rigid Alignment
  - Thin-Plate Splines
- Feature Correspondences
  - Piecewise ICP
- Pairwise Non-Rigid Alignment
- Results & Future Work

Finding Correspondences for TPS

Thin-plate splines can interpolate their control points exactly, regardless of the amount of warp, so iterative correspondences don't work.

- Progressively anneal λ

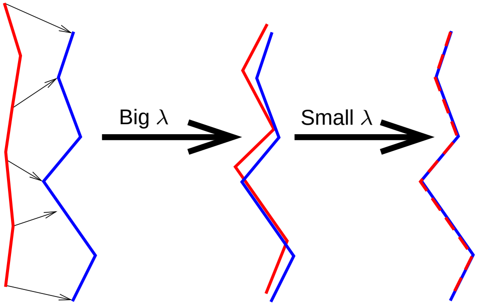

Finding Correspondences for TPS

Thin-plate splines can interpolate their control points exactly, regardless of the amount of warp, so iterative correspondences don't work.

- Progressively anneal λ
- Fuzzy correspondences [Chui03]

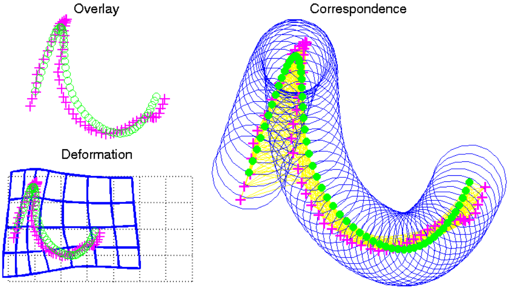

Chui & Rangarajan

Piecewise ICP

- Record ICP alignment errors for each piece
- Dice piece with highest error
- Stop when alignment becomes unstable

Target scan

Diced source scan

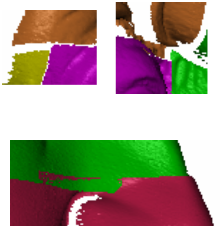

Pairwise Alignment Pipeline

- Piecewise ICP
  - Diced 15 times, 42.47 sec (first 5 shown)

Pairwise Alignment Pipeline

- Piecewise ICP
  - Diced 15 times, 42.47 sec (first 5 shown)
- TPS Warp
  - 250 points, .06 sec to compute warp, 6.92 sec to apply it to 323,098 vertices

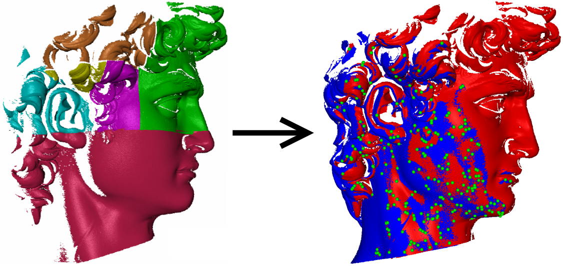

Pairwise Alignment Pipeline

- Piecewise ICP
  - Diced 15 times, 42.47 sec (first 5 shown)
- TPS Warp
  - 250 points, .06 sec to compute warp, 6.92 sec to apply it to 323,098 vertices
- VRIP volumetric merging step [Curless96]

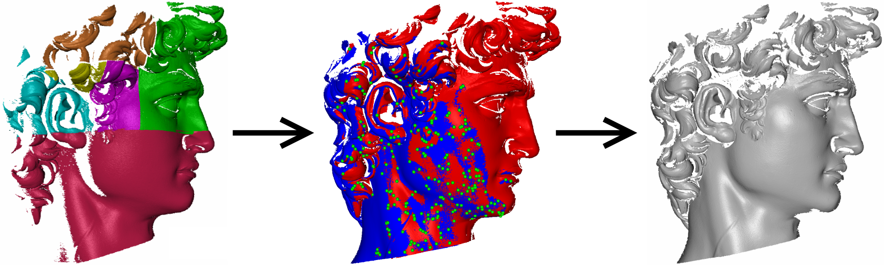

Overview

- Previous Work
  - Iterative Closest Points (Rigid-Body Alignment)
  - Hierarchical ICP
- Non-Rigid Alignment
  - Thin-Plate Splines
- Feature Correspondences
  - Piecewise ICP
- Pairwise Non-Rigid Alignment
- Results & Future Work

Pairwise Alignment Results

right_shoulder_chest_a
295,012 vertices

face4_e
323,098 vertices

Alignment Quality

- No overlap
- Good ICP alignment
- Good TPS alignment
- Good ICP & TPS alignments

Pairwise Alignment Results

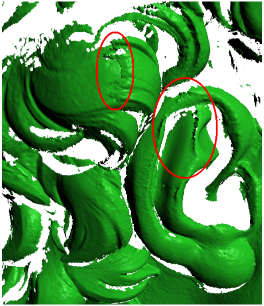

ICP Alignment

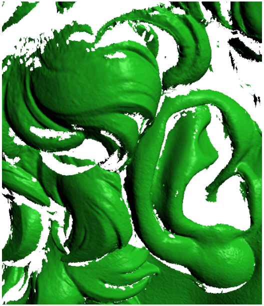

TPS Alignment

Merged Mesh

Current Work: Global Alignment Results

16 meshes after global alignment

VRIP merged output

Future Work and Acknowledgments

- **Global Registration**
- Stability and Performance Improvements
  - Better ICP sampling heuristics
  - Improve dicing heuristic to reduce number of iterations
- TPS feature point selection improvements

Thanks to

- National Science Foundation
- Princeton Graphics Group
- Prof. Ken Steiglitz
- Natasha Gelfand and Leslie Ikemoto (Hierarchical ICP)

End

Pairwise Alignment Results

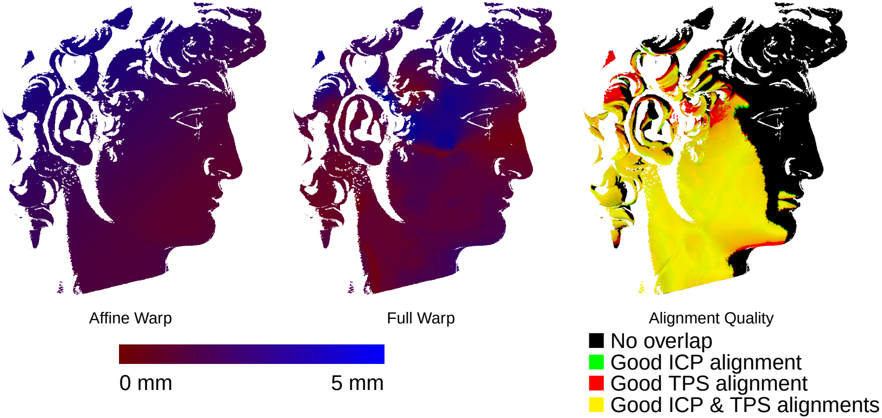

Thin-Plate Splines

- Thin-plate splines always take the form

$4 \times n$ non-affine warping parameters ($WX^t = 0$)

$$y = Ax + W\underbrace{K(x)}$$

4×4 affine transformation **$n \times 1$ control point influence vector**

where $K(x) = (|x - x_1|, \ldots, |x - x_n|)^t$ in 3-D

Thin-Plate Splines

- Thin-plate splines always take the form

$4 \times n$ non-affine warping parameters ($WX^t = 0$)

$$y = Ax + W\underbrace{K(x)}_{n \times 1 \text{ control point influence vector}}$$

4×4 affine transformation

where $K(x) = (|x - x_1|, \ldots, |x - x_n|)^t$ in 3-D

- The warping coefficients A and W are computed by the equation

$$\left(\begin{array}{c|c} A & W \end{array} \right) \left(\begin{array}{c|c} X & 0 \\ \hline K + n\lambda I & X^t \end{array} \right) = \left(\begin{array}{c|c} Y & 0 \end{array} \right)$$

where $K_{ij} = |x_i - x_j|$.

Previous Work: Softassign and Deterministic Annealing

- Maintain probability that each point in X maps to each other point in Y (Softassign)
- Probability of two points corresponding has Gaussian fall-off with respect to distance
- Fall-off narrows at each iteration until we reach exact correspondence (Deterministic annealing)

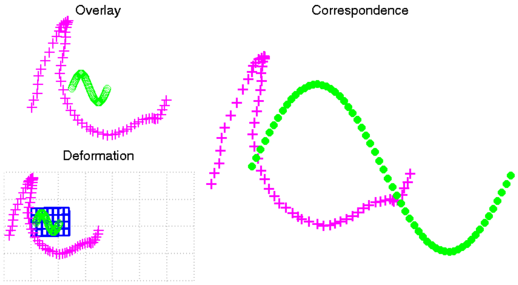

Chui & Rangarajan

Previous Work: Softassign and Deterministic Annealing

- Maintain probability that each point in X maps to each other point in Y (Softassign)
- Probability of two points corresponding has Gaussian fall-off with respect to distance
- Fall-off narrows at each iteration until we reach exact correspondence (Deterministic annealing)

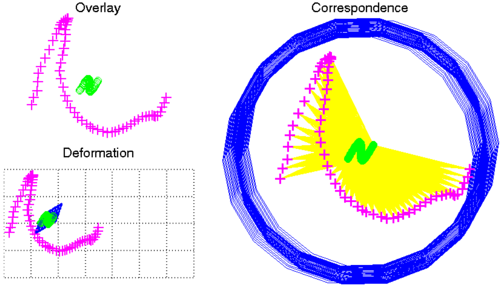

Chui & Rangarajan

Previous Work: Softassign and Deterministic Annealing

- Maintain probability that each point in X maps to each other point in Y (Softassign)
- Probability of two points corresponding has Gaussian fall-off with respect to distance
- Fall-off narrows at each iteration until we reach exact correspondence (Deterministic annealing)

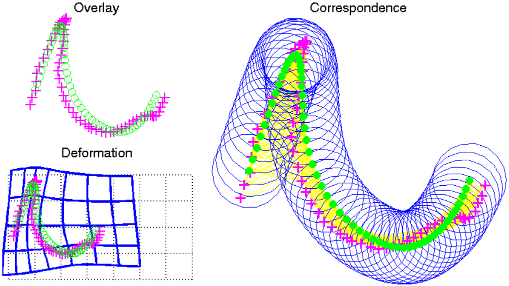

Chui & Rangarajan

Previous Work: Softassign and Deterministic Annealing

- Maintain probability that each point in X maps to each other point in Y (Softassign)
- Probability of two points corresponding has Gaussian fall-off with respect to distance
- Fall-off narrows at each iteration until we reach exact correspondence (Deterministic annealing)

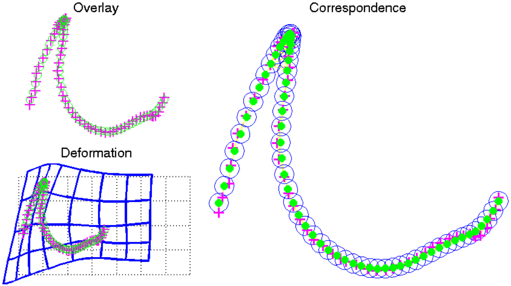

Chui & Rangarajan

Current Work: Global Alignment

- Error accumulates over successive alignments
- Distribute error
  - across warping function (ICP)
  - across positions of global feature points in space

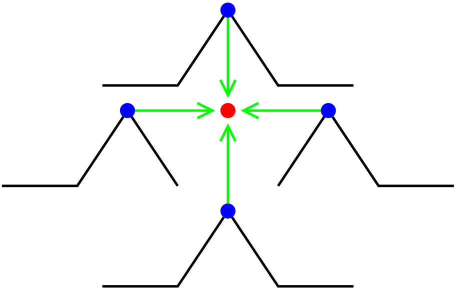

Current Work: Global Alignment

- Perform best global ICP alignment

Current Work: Global Alignment

- Perform best global ICP alignment
- For each range scan S

Current Work: Global Alignment

- Perform best global ICP alignment
- For each range scan S
  - Randomly select features at constant sampling rate

Current Work: Global Alignment

- Perform best global ICP alignment
- For each range scan S
  - Randomly select features at constant sampling rate
  - For each overlapping range scan S'

Current Work: Global Alignment

- Perform best global ICP alignment
- For each range scan S
  - Randomly select features at constant sampling rate
  - For each overlapping range scan S'
    - Align S' to S and find feature correspondences

Current Work: Global Alignment

- Perform best global ICP alignment
- For each range scan S
  - Randomly select features at constant sampling rate
  - For each overlapping range scan S'
    - Align S' to S and find feature correspondences
    - Record position of each feature point on S'

Current Work: Global Alignment

- Perform best global ICP alignment
- For each range scan S
  - Randomly select features at constant sampling rate
  - For each overlapping range scan S'
    - Align S' to S and find feature correspondences
    - Record position of each feature point on S'
  - Record mean position of each feature point across all scans on which it falls as its canonical position

Current Work: Global Alignment

- Perform best global ICP alignment
- For each range scan S
  - Randomly select features at constant sampling rate
  - For each overlapping range scan S'
    - Align S' to S and find feature correspondences
    - Record position of each feature point on S'
  - Record mean position of each feature point across all scans on which it falls as its canonical position
- For each range scan S

Current Work: Global Alignment

- Perform best global ICP alignment
- For each range scan S
  - Randomly select features at constant sampling rate
  - For each overlapping range scan S'
    - Align S' to S and find feature correspondences
    - Record position of each feature point on S'
  - Record mean position of each feature point across all scans on which it falls as its canonical position
- For each range scan S
  - Warp all features on S to their canonical positions

Current Work: Global Alignment

- Perform best global ICP alignment
- For each range scan S
  - Randomly select features at constant sampling rate
  - For each overlapping range scan S'
    - Align S' to S and find feature correspondences
    - Record position of each feature point on S'
  - Record mean position of each feature point across all scans on which it falls as its canonical position
- For each range scan S
  - Warp all features on S to their canonical positions
- $O(nm)$ for m overlapping scans

Current Work: Global Alignment

- Perform best global ICP alignment
- For each range scan S
  - Randomly select features at constant sampling rate
  - For each overlapping range scan S'
    - Align S' to S and find feature correspondences
    - Record position of each feature point on S'
  - Record mean position of each feature point across all scans on which it falls as its canonical position
- For each range scan S
  - Warp all features on S to their canonical positions
- $O(nm)$ for m overlapping scans
- At most two scans in memory at a time