

Registration and Alignment

Justin Solomon

MIT, Spring 2019




Acknowledgements

Many slides from


- **Szymon Rusinkiewicz**, Princeton
ICCV Course, 2005
- **Hao Li**, USC
CSCI 599, 2015

Registration Problem



Align two overlapping objects

3D Reconstruction Pipeline



Rough Plan

- ICP algorithm


A classic!

- ICP variants

- Related problems

Synchronization, non-rigid registration

Starting Point



$$q_i = Rp_i + t$$

Can align given enough matches




How many
correspondences
determine R and t ?




How do you get
correspondences?

Rough Approximation



Closest points correspond

Try a Second Time...



Iterative Closest Point (ICP)

- **Choose** e.g. 1000 random points
- **Match** each to closest point on other scan
- **Reject** pairs with distance $> k$ times median
- **Minimize**

$$E[R, t] := \sum_i \|Rp_i + t - q_i\|^2$$

- **Iterate**

“A method for registration of 3-D shapes.”
Besl and McKay, PAMI 1992.

On the Board

$$\min_{t \in \mathbb{R}^3, R^\top R = I} \sum_i \|Rp_i + t - q_i\|^2$$


Closed-form formulas!

Many (!) Variants of ICP

- **Source points** from one or both meshes
- **Matching** to points in the other mesh
 - **Weighting** correspondences
 - Rejecting **outlier** point pairs
 - Alternative **error metrics**

Point-to-Plane Error Metric

Flat parts can slide along each other



$$E[R, t] := \sum_i ((Rp_i + t - q_i)^\top n_i)^2$$

$$\approx \sum_i [(p_i - q_i)^\top n_i + r^\top (p_i \times n_i) + t^\top n_i]^2 \text{ after linearizing}$$

where $r := (r_x, r_y, r_z)$

Least-squares!

“Object modelling by registration of multiple range images”


Chen and Medioni, Image and Vision Computing 10.3 (1992); image courtesy N. Mitra

Closest Compatible 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

Choose Points to Improve Stability




Uniform Sampling




Stable Sampling

Sample discriminative points


Local Covariance




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



5 DOFs stable




4 DOFs stable




6 DOFs stable

Alternative: Uniform Normals



Random Sampling





Normal-space Sampling





What is the **bottleneck**
of ICP iteration?



BSP Tree




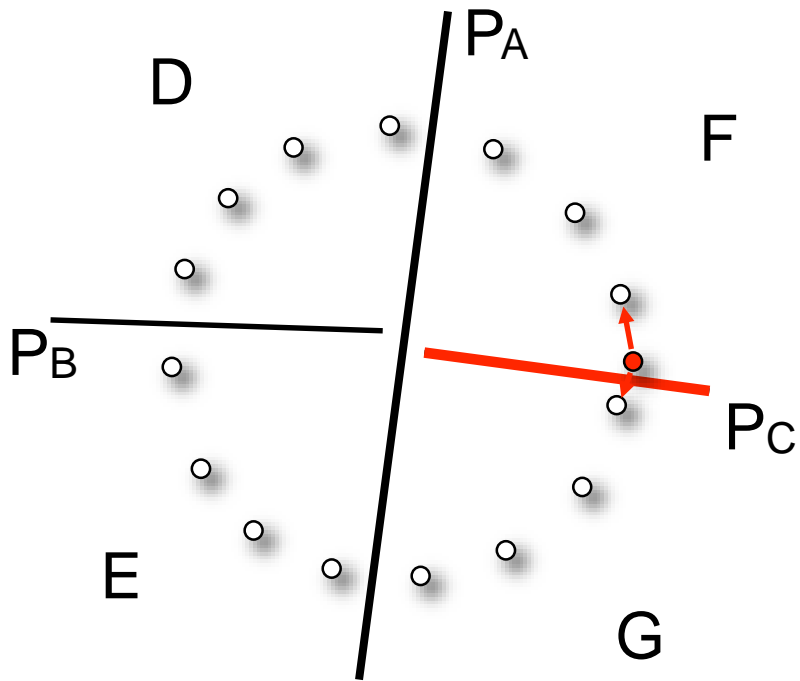
Tree Traversal



Subtlety: Is this right?



Two Possibilities



Pseudocode: Be Conservative!

```
BSPNode::dist(Point x, Scalar& dmin) {
    if (leaf_node())
        for each sample point p[i]
            dmin = min(dmin, dist(x, p[i]));
    else {
        d = dist_to_plane(x);
        if (d < 0) {
            left_child->dist(x, dmin);
            if (|d| < dmin) right_child->dist(x, dmin);
        } else {
            right_child->dist(x, dmin);
            if (|d| < dmin) left_child->dist(x, dmin);
        }
    }
}
```

k-d Tree





Image courtesy R. Gvili

Axis-aligned tree




Speed aside, is ICP
always successful?

Convergence Funnel Visualization



Translation in xz plane
Rotation about y




Converges




Does not converge

Distance Field Method




Translation in xz plane
Rotation about y




- Converges
- Does not converge

Point-to-Plane


Translation in xz plane
Rotation about y



- Converges
- Does not converge




Issue: ICP Three Times



Usually have ≥ 2 scans

Improve Sequential Alignment?




Prevent "drift"

Simple Methods

- Align everything to **anchor** scan
Which to choose? Dependence on anchor?
- Align to **union** of previous scans
Order dependence? Speed?
- **Simultaneously** align everything using ICP
Local optima? Computational expense?

Graph Approach



Align similar scans, then assemble

Lu and Milios

- **Pairwise phase**

Compute pairwise ICP on graph

- **Global alignment**


Least-squares rotation/translation

**Linearize for
global alignment**

Failed ICP in Global Registration



Correct global registration



Possible to
make robust?

Global registration including bad ICP

Digression: Angular Synchronization

Given: $\delta_{ij} \approx \theta_i - \theta_j \pmod{2\pi}$, $(i, j) \in E$

Find: $\{\theta_i\}$ up to constant shift



2D version without translation

*On the board:
Eigenvalue and/or SDP relaxations*

Open problem:
Synchronization on non-compact groups (e.g. SE(3)!)

“Angular synchronization by eigenvectors and semidefinite programming.”
Singer, ACHA 2010.

Non-Rigid Registration




data provided by Paramount Pictures and Aguru Images

Problems


- Noisy data
- Acquisition holes (incomplete)
 - No correspondence
 - Deformation



missing correspondences



noise



holes

Example Paper

Eurographics Symposium on Geometry Processing 2008
Pierre Alliez and Szymon Rusinkiewicz
(Guest Editors)

Volume 27 (2008), Number 5

Global Correspondence Optimization for Non-Rigid Registration of Depth Scans

Hao Li Robert W. Sumner Mark Pauly
Applied Geometry Group
ETH Zurich

Abstract

We present a registration algorithm for pairs of deforming and partial range scans that addresses the challenges of non-rigid registration within a single non-linear optimization. Our algorithm simultaneously solves for correspondences between points on source and target scans, confidence weights that measure the reliability of each correspondence and identify non-overlapping areas, and a warping field that brings the source scan into alignment with the target geometry. The optimization maximizes the region of overlap and the spatial coherence of the deformation while minimizing registration error. All optimization parameters are chosen automatically; hand-tuning is not necessary. Our method is not restricted to part-in-whole matching, but addresses the general problem of partial matching, and requires no explicit prior correspondences or feature points. We evaluate the performance and robustness of our method using scan data acquired by a structured light scanner and compare our method with existing non-rigid registration algorithms.


Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling

1. Introduction

Surface registration is a fundamental problem in geometric modeling and 3-D shape acquisition. Most scanning systems provide partial surface data that must be aligned and merged

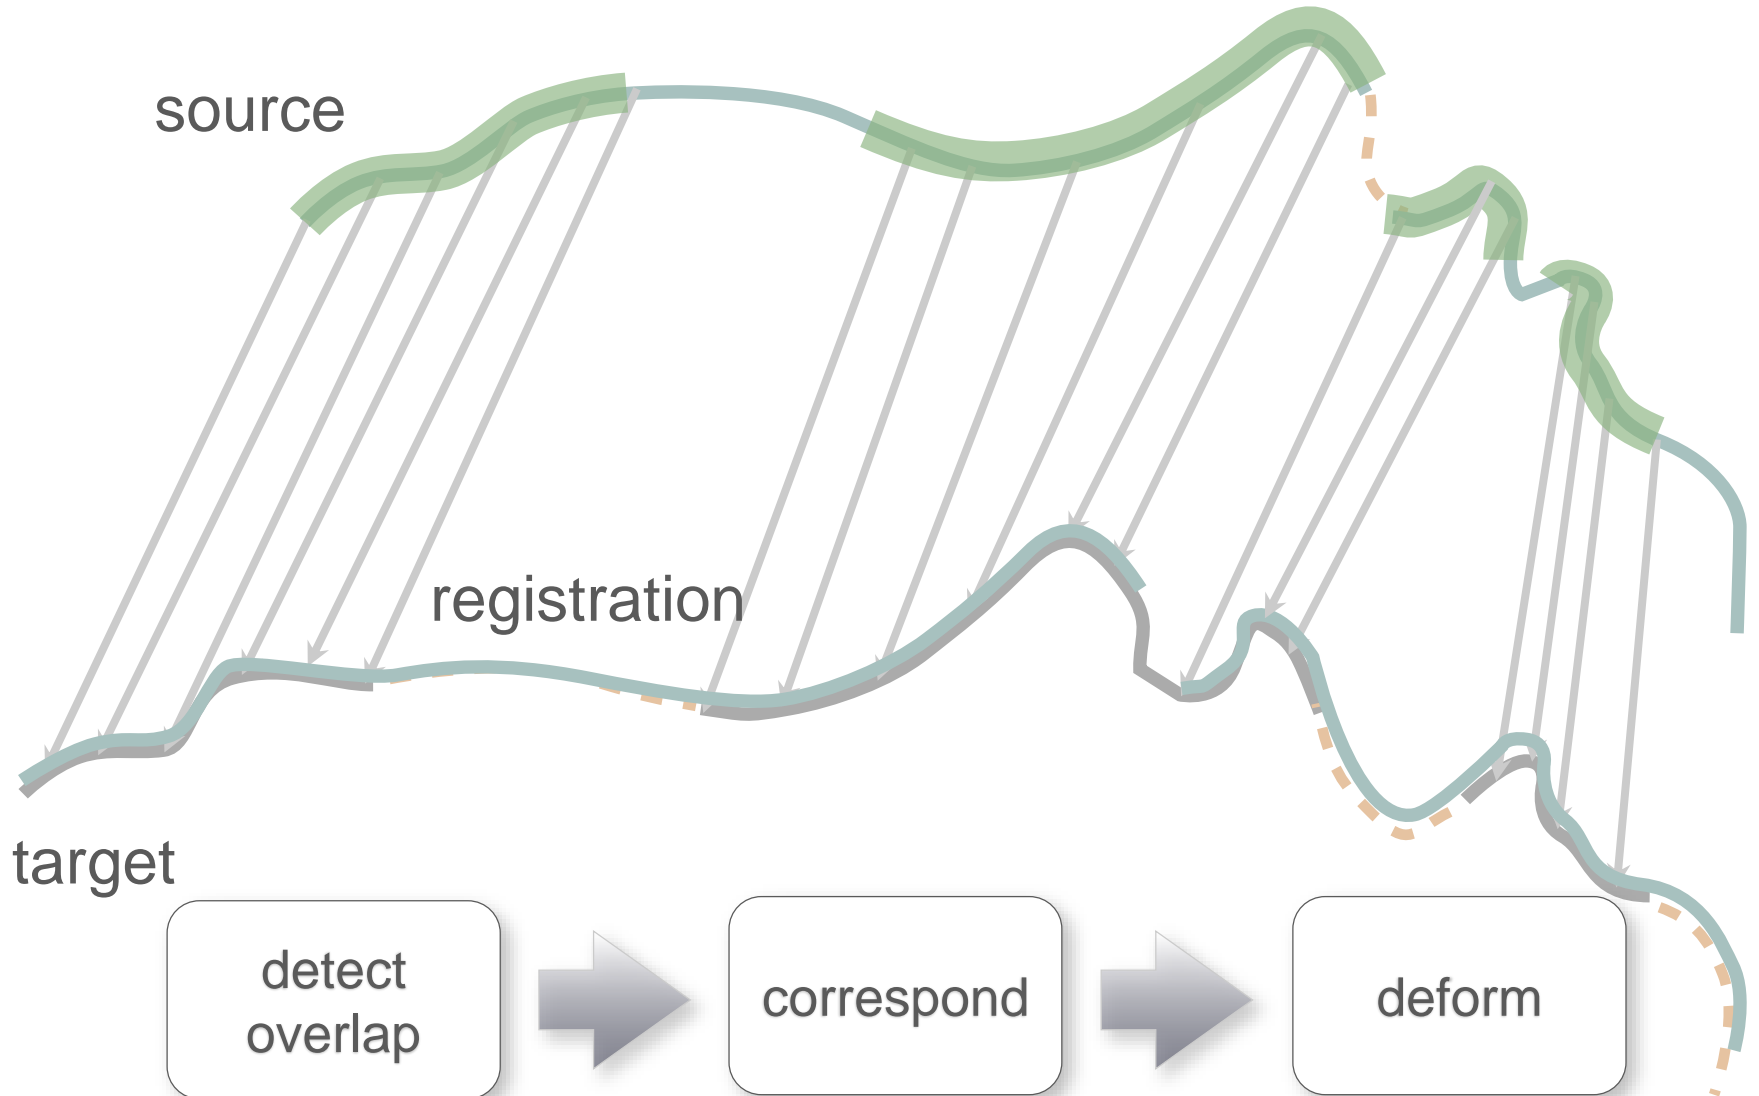


Concrete Example

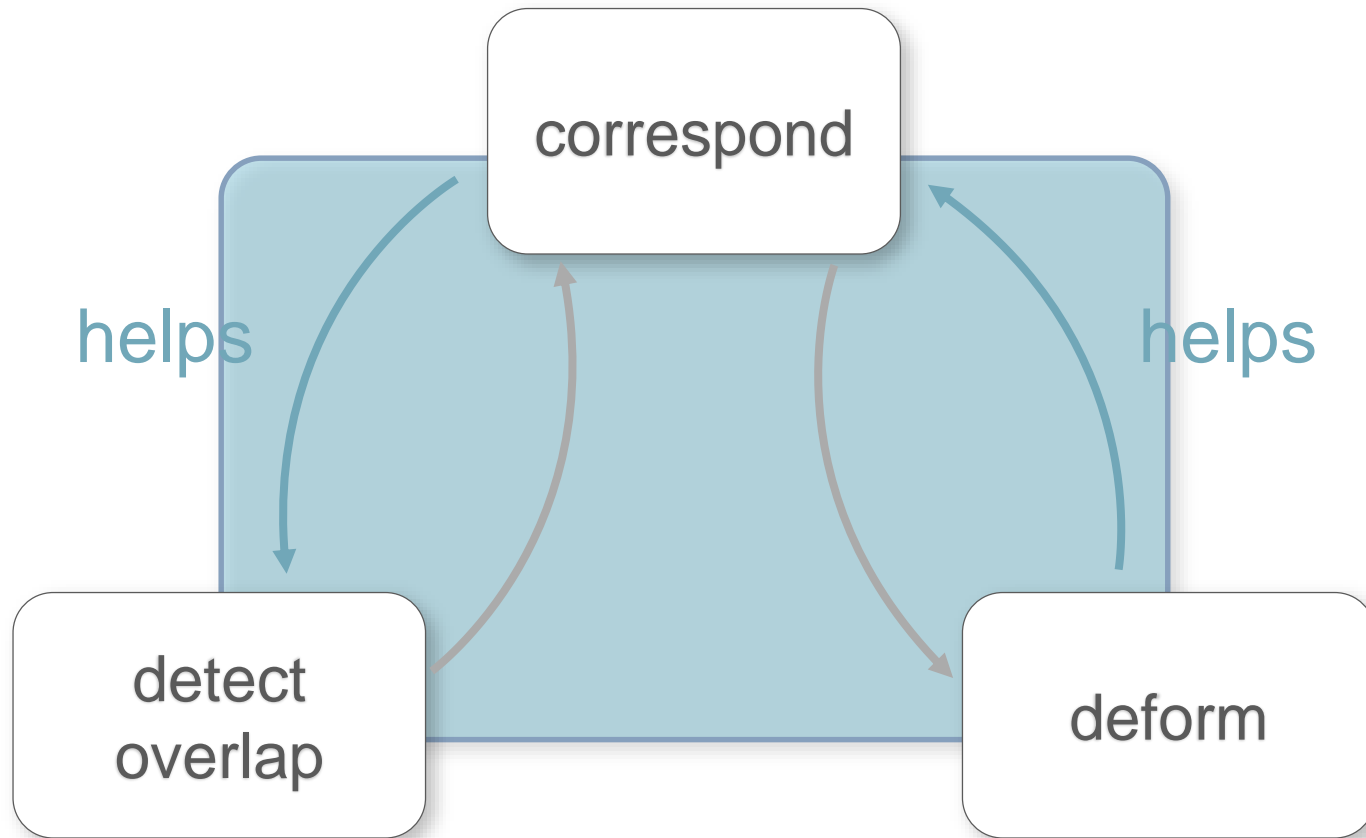


In addition to deformation

Reasonable Approach

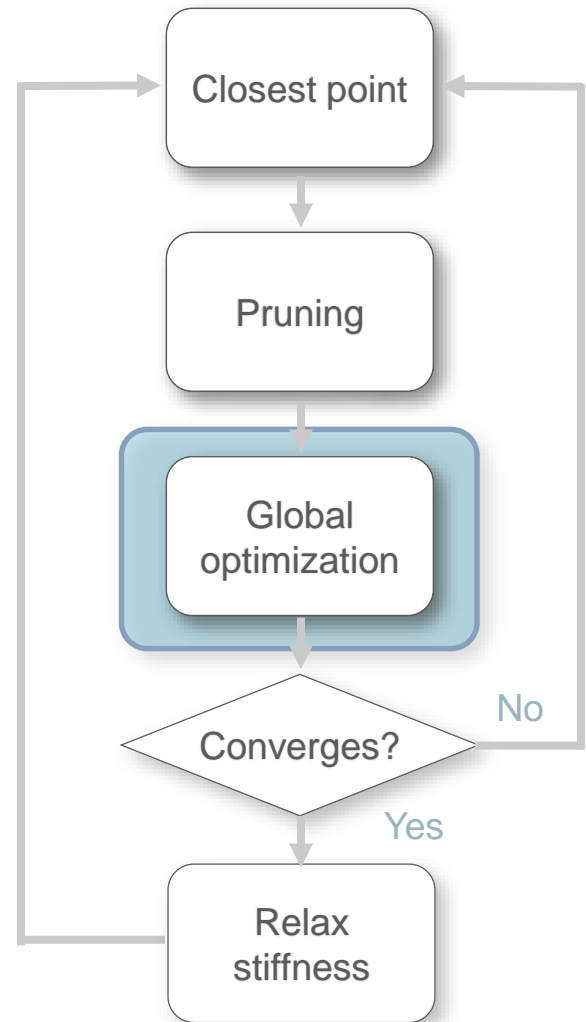



Global Optimization

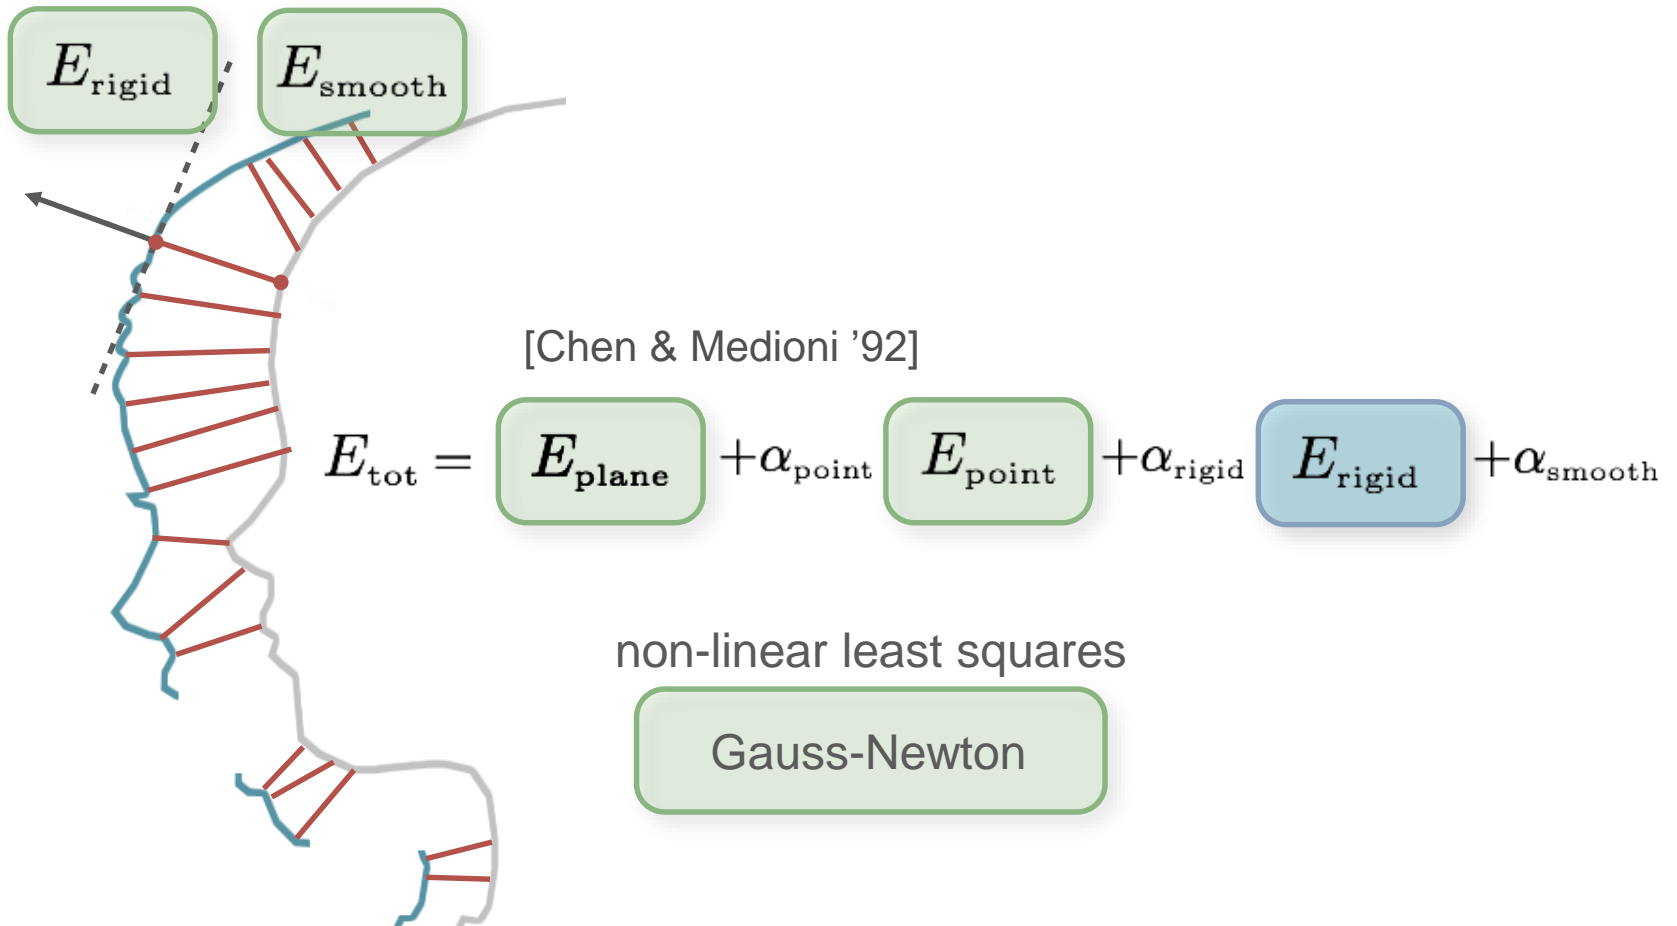


Tasks support each other

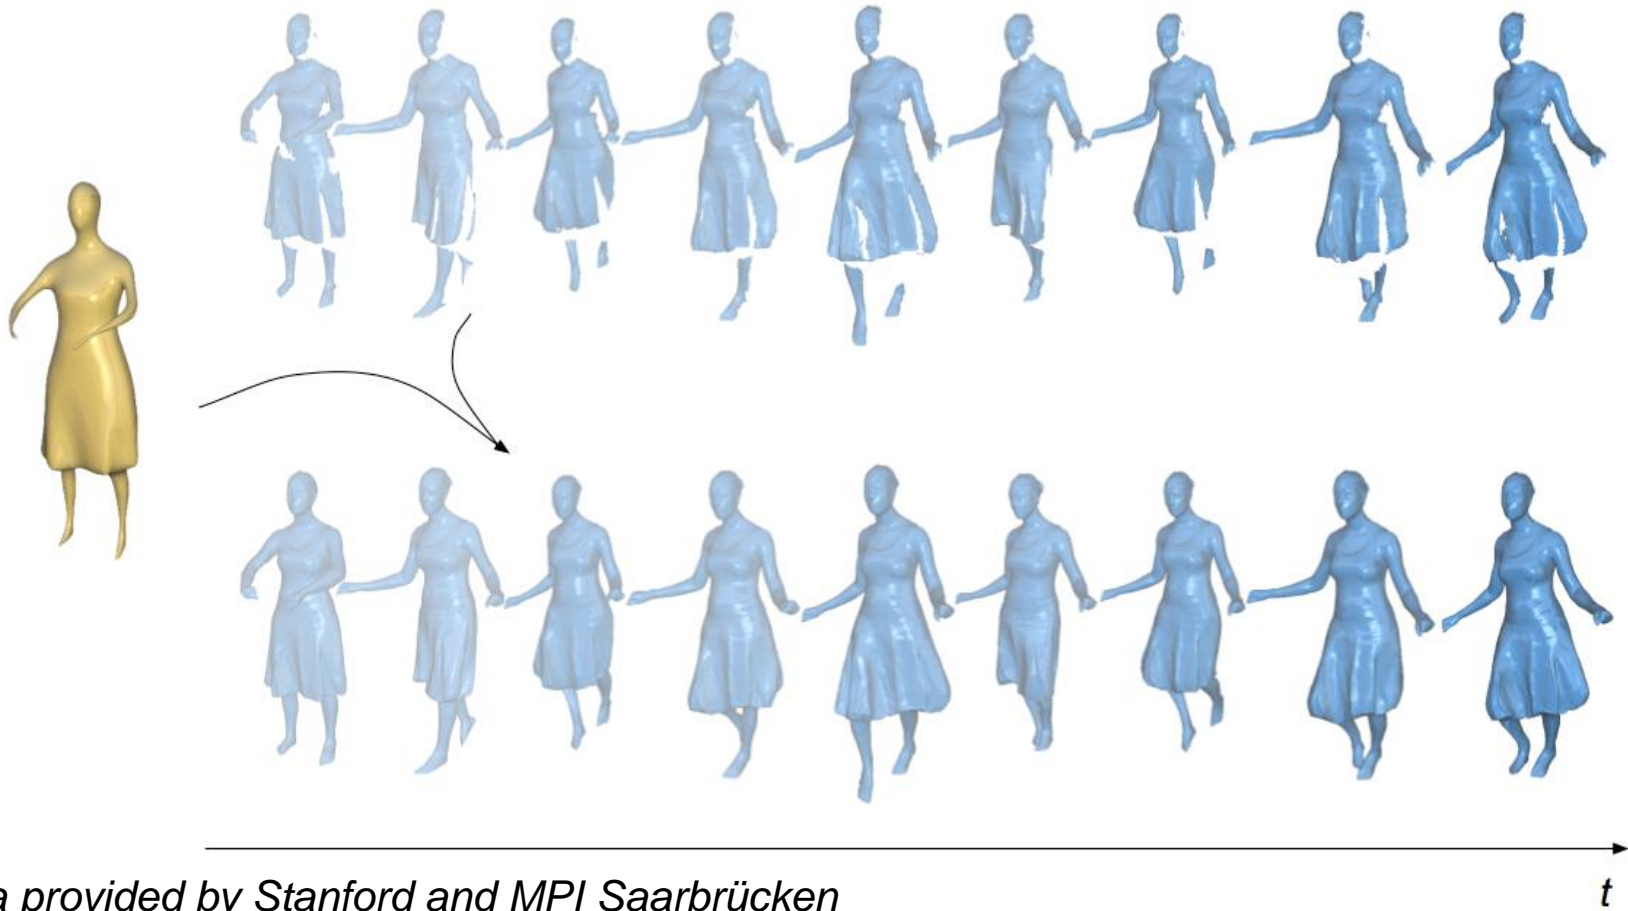
Pipeline



Rough Summary




Alternative Approaches



Template-based matching

Alternative Approaches




Template alignment, blendshapes


Outstanding Challenges




Deformation, clothing
& props



Environment



Low-cost scanners



Registration and Alignment

Justin Solomon
MIT, Spring 2019

