[Wang et al. 2012]

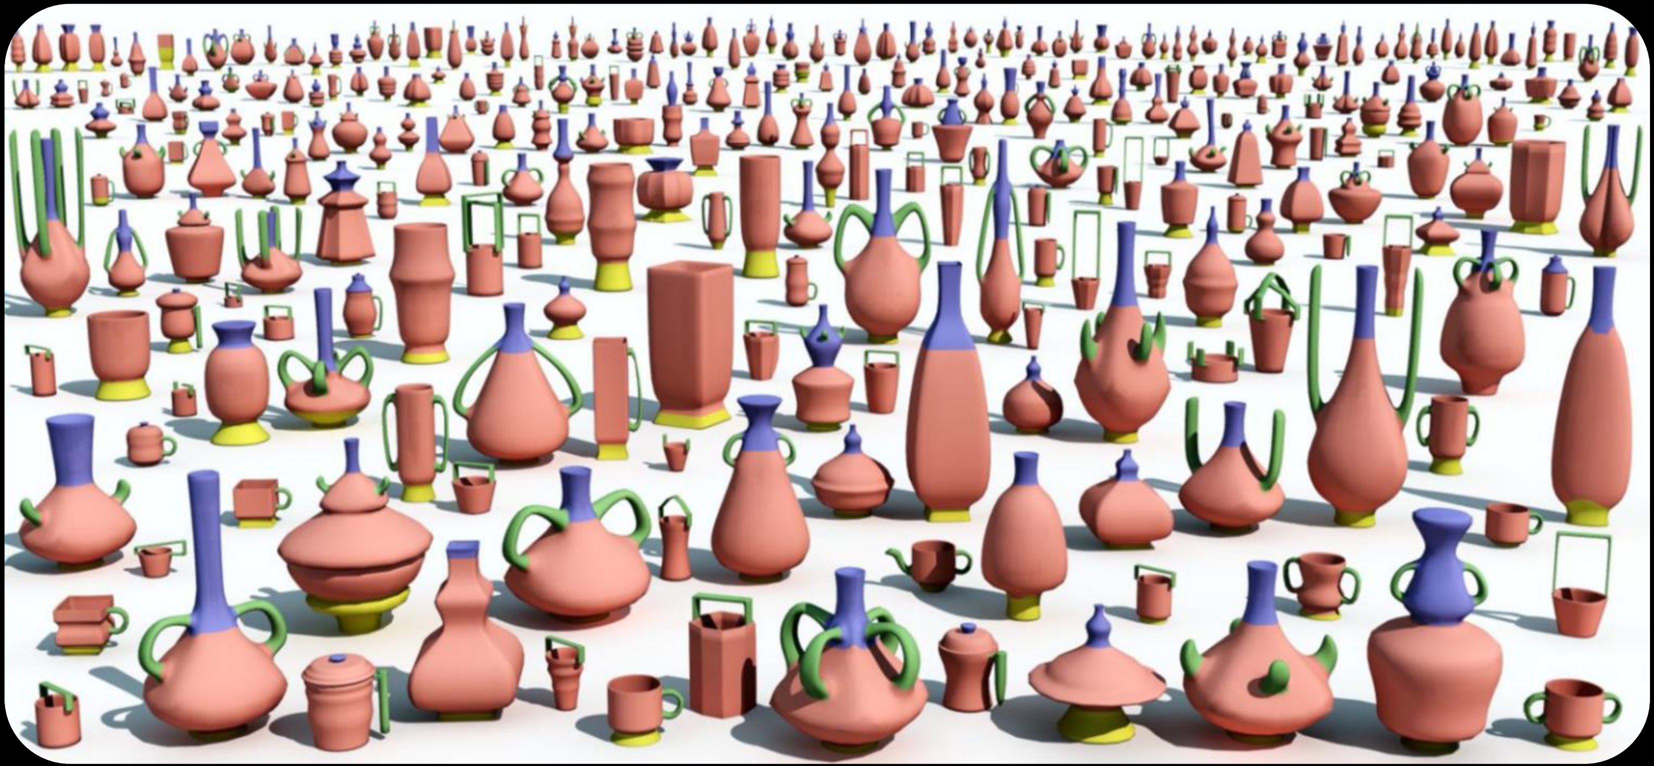

Consistent Correspondence

Justin Solomon

MIT, Spring 2017

Most slides from CS 468, Stanford (Kim & Huang)

Previously

Map between two shapes.

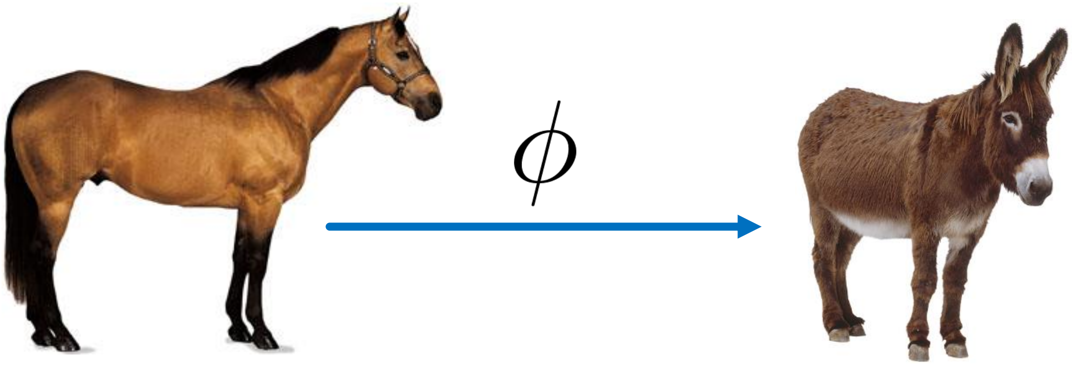

Question

What happens if you **compose** these maps?

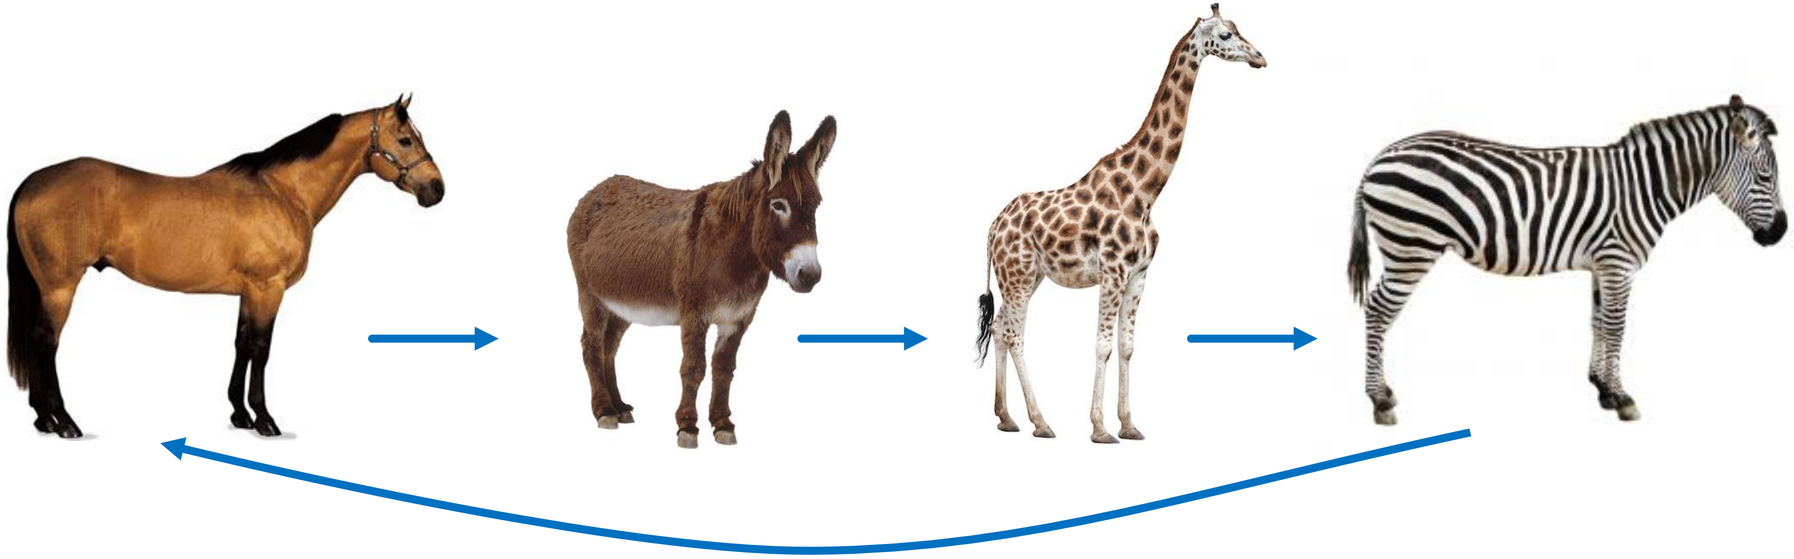

Q:

What do you expect if you compose around a cycle?

Cycle consistency

[**sahy**-k*uh* l k*uh* n-**sis**-t*uh* n-see]:

Composing maps in a cycle yields the identity

An Unpleasant Constraint

$$\phi_1(\phi_2(\phi_3(x))) = \mathrm{Id}$$

Cycle consistency

Contrasting Viewpoint

Many possible pairwise matches!

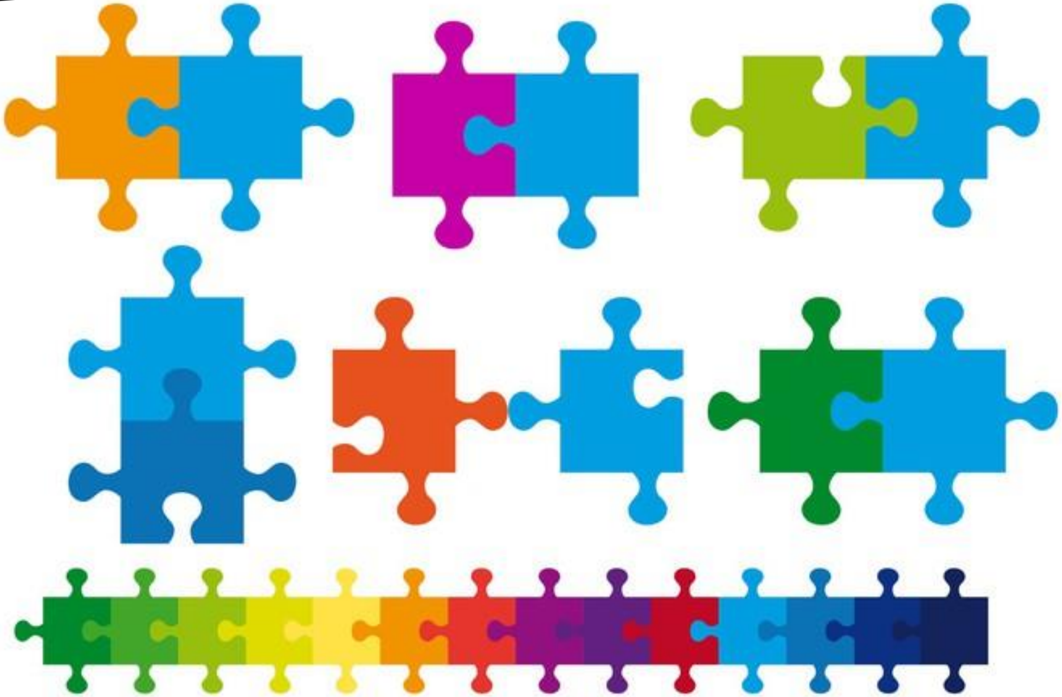

https://s3.pixers.pics/pixers/700/FO/39/51/09/46/700_FO39510946_cd54b90a83d46f5dbd96440271eadfec.jpg

Additional data should <u>help</u>!

Philosophical Point

You should have a good reason if your mapping tool is inconsistent.

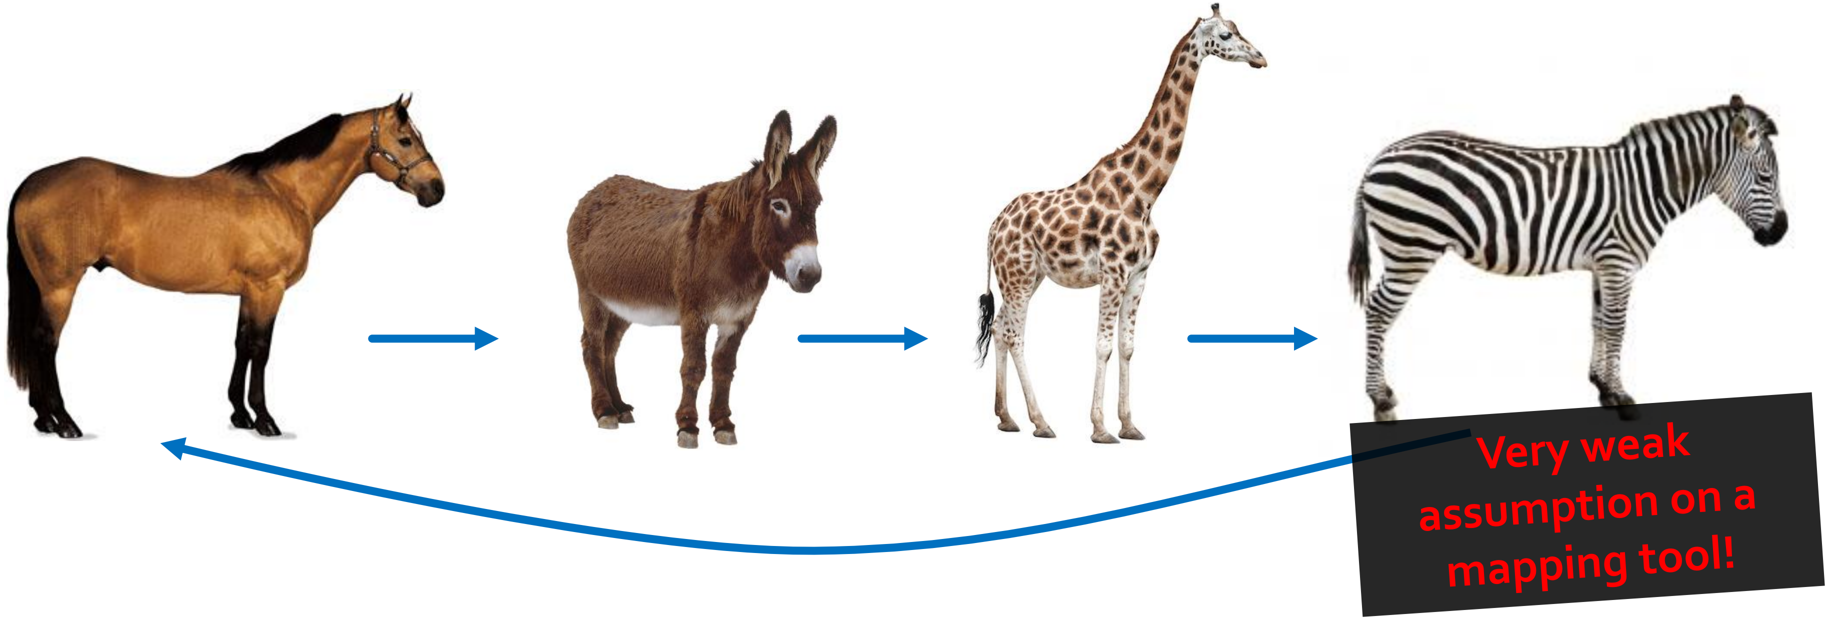

Joint Matching: Simplest Formulation

- **Input**
  - ***N*** **shapes**
  - **N^2 maps (see last lecture)**
- **Output**
  - **Cycle-consistent approximation**

Holy Grail

Simultaneously optimize all maps in a collection.

Open problem!

Unsurprisingly…

Given: Model graph $G = (S, E)$

Find: Largest consistent spanning tree

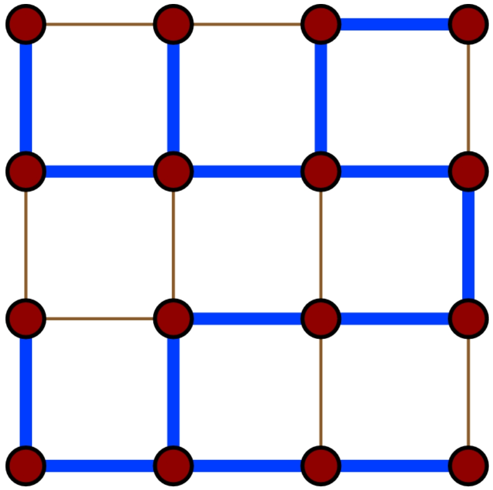

"Automatic Three-Dimensional Modeling from Reality" (Huber, 2002)

NP-hard

Today

Sampling of methods for consistent correspondence.

- Spanning tree
- Inconsistent cycle detection
- Convex optimization

Spanning Tree: Original Context

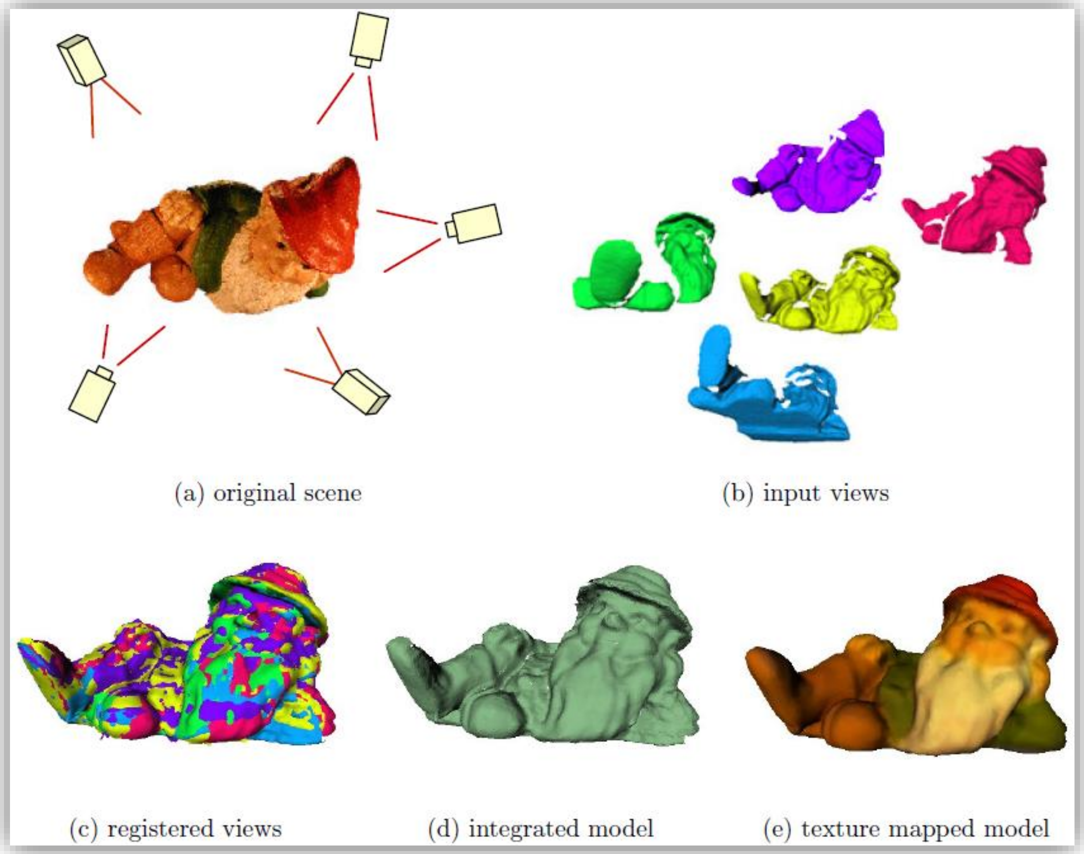

"Automatic Three-Dimensional Modeling from Reality" (Huber, 2002)

Multi-view registration

Basic Algorithm

Extract consistent spanning tree in model graph

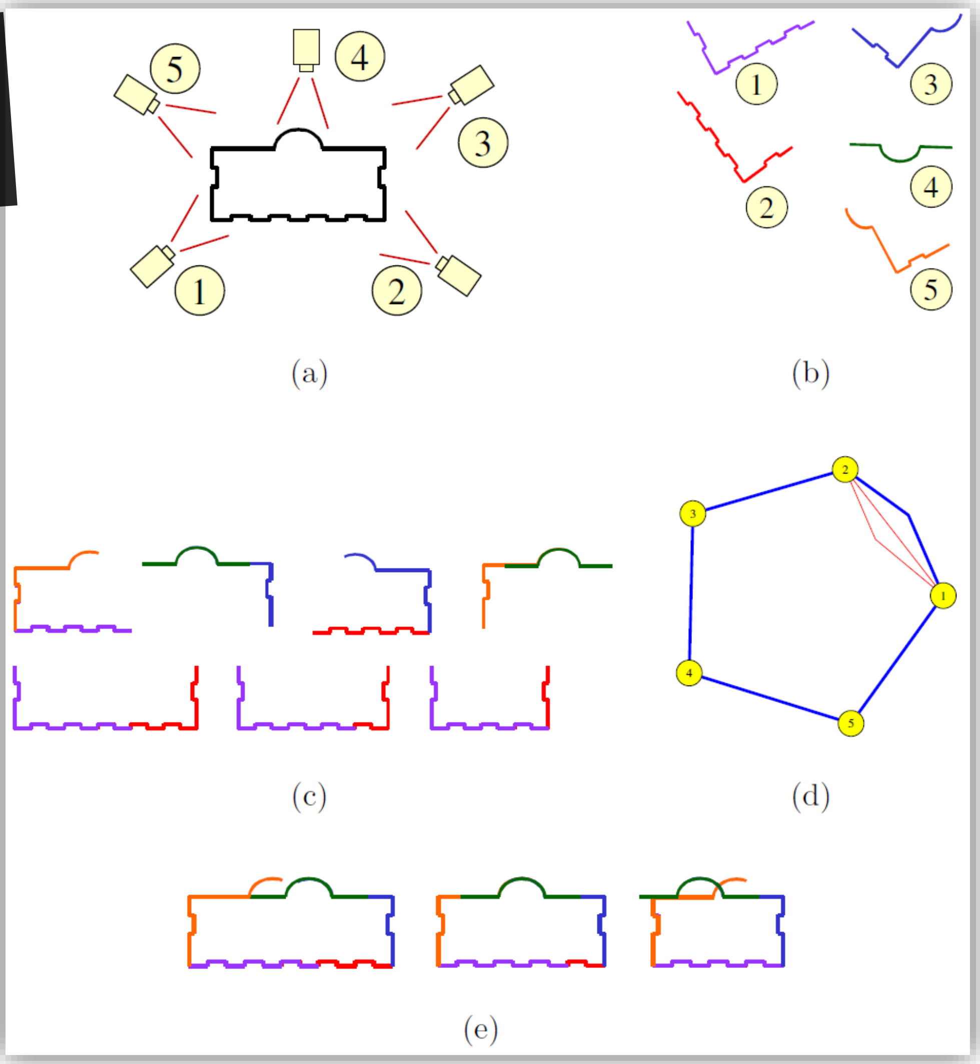

Issues

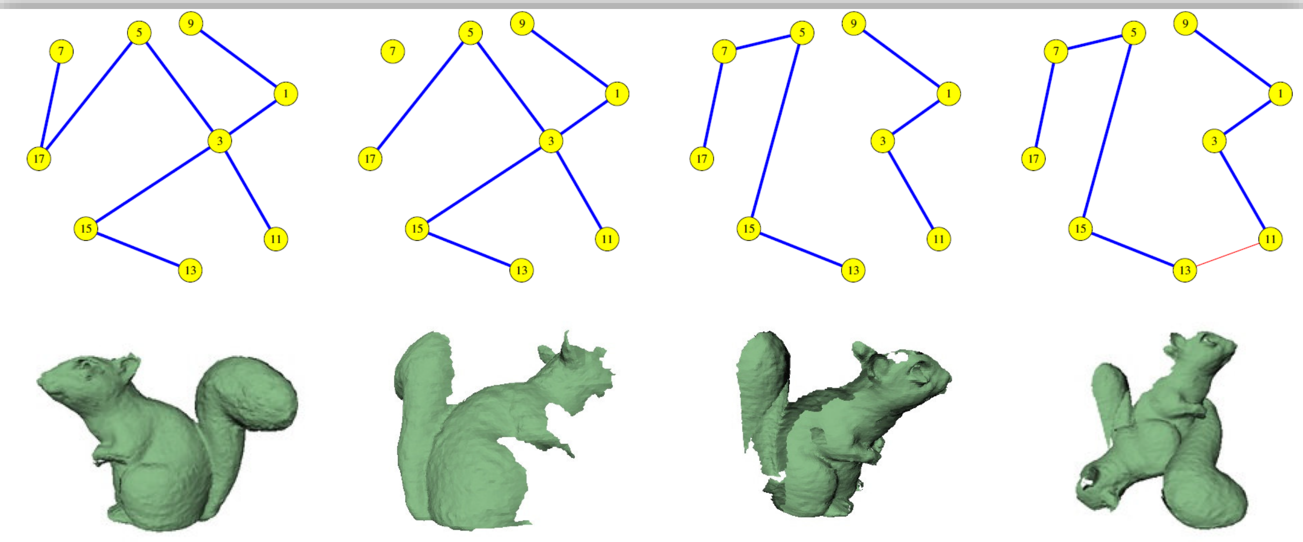

(a) $Q_G = 115$ (b) $Q_G = 97.3$ (c) $Q_G = 69.9$ (d) $Q_G = -3183$

Figure 3.13: Global quality values for several versions of the squirrel model. The model hypothesis is shown in the top row with the corresponding 3D visualization in the bottom row. a) Correct model. b) Correct model with a single view detached; c) Correct model split into equally sized two parts (only one part shown in 3D). d) Model with one error.

- Many spanning trees
- Single incorrect match can destroy the maps

Inconsistent Loop Detection

Used to deal with repeating structures like windows!

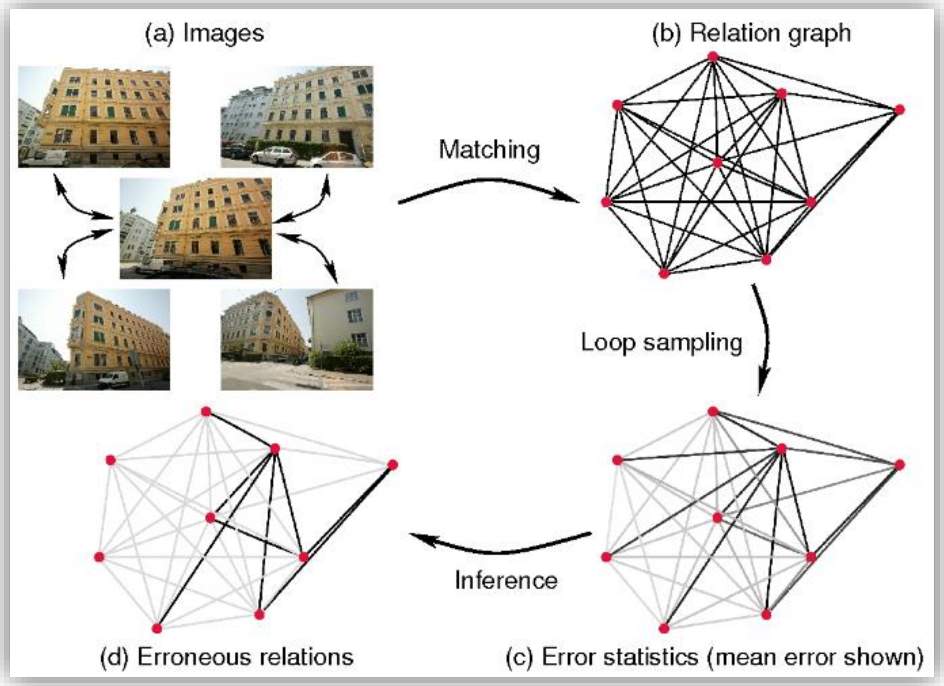

Large for inconsistent cycles

$$\begin{aligned} \max \quad & \textstyle\sum_L \rho_L x_L \\ \text{s.t.} \quad & x_L \geq x_e \;\forall e \in L \\ & x_L \leq \textstyle\sum_{e \in L} x_e \\ & x_L, x_e \in [0, 1] \end{aligned}$$

x_e = 1 for false positive edge

x_L = max of x_e over loop

"Disambiguating Visual Relations Using Loop Constraints" (Zach et al., CVPR 2010)

Relationship: Consistency vs. Accuracy

DOI: 10.1111/j.1467-8659.2011.02022.x

Eurographics Symposium on Geometry Processing 2011
Mario Botsch and Scott Schaefer
(Guest Editors)

Volume 30 (2011), Number 5

An Optimization Approach to Improving Collections of Shape Maps

Andy Nguyen[1] Mirela Ben-Chen[1] Katarzyna Welnicka[2] Yinyu Ye[1] Leonidas Guibas[1]

[1]Stanford University
[2]Technical University of Denmark

Abstract

Finding an informative, structure-preserving map between two shapes has been a long-standing problem ... ometry processing, involving a variety of solution approaches and applications. However, in many cases, ... given not only two related shapes, but a collection of them, and considering each pairwise map independently ... not take full advantage of all existing information. For example, a notorious problem with computing shape ... is the ambiguity i... there exist two m... them based on the... sensitivity to how... Given the context... map consistency, ... chosen in the net... help us replace a... sense interpolate... mization problem... and individually a... shapes, as long as... for improving map...

Categories and Su...

Iteratively fix triplets and reweight

Definition 3 Given a collection of maps $\mathcal{M}$, let $\mathcal{B}(\mathcal{M}) = \{m_{i,j} \in \mathcal{M} \mid E_{acc}(m_{i,j}) > 0\}$ — the collection of inaccurate maps. Then we say that $\mathcal{M}$ is *almost accurate*, if there do not exist two maps $m_1, m_2 \in \mathcal{B}(\mathcal{M})$, which both belong to the same 3-cycle in $G_{\mathcal{M}}$. We call such maps *isolated*.

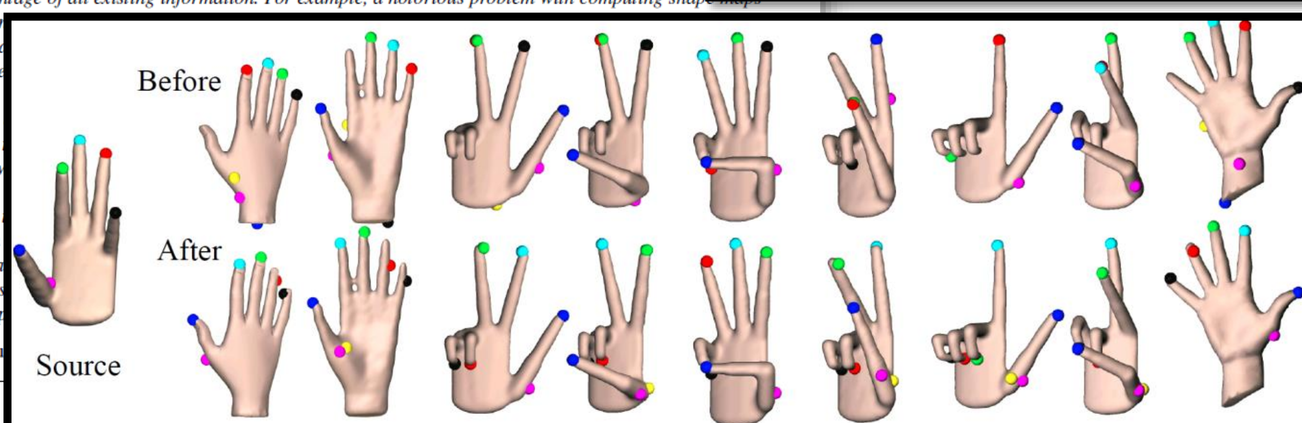

1. Introduction

a common measure of distortion is the difference between

Fuzzy Correspondences

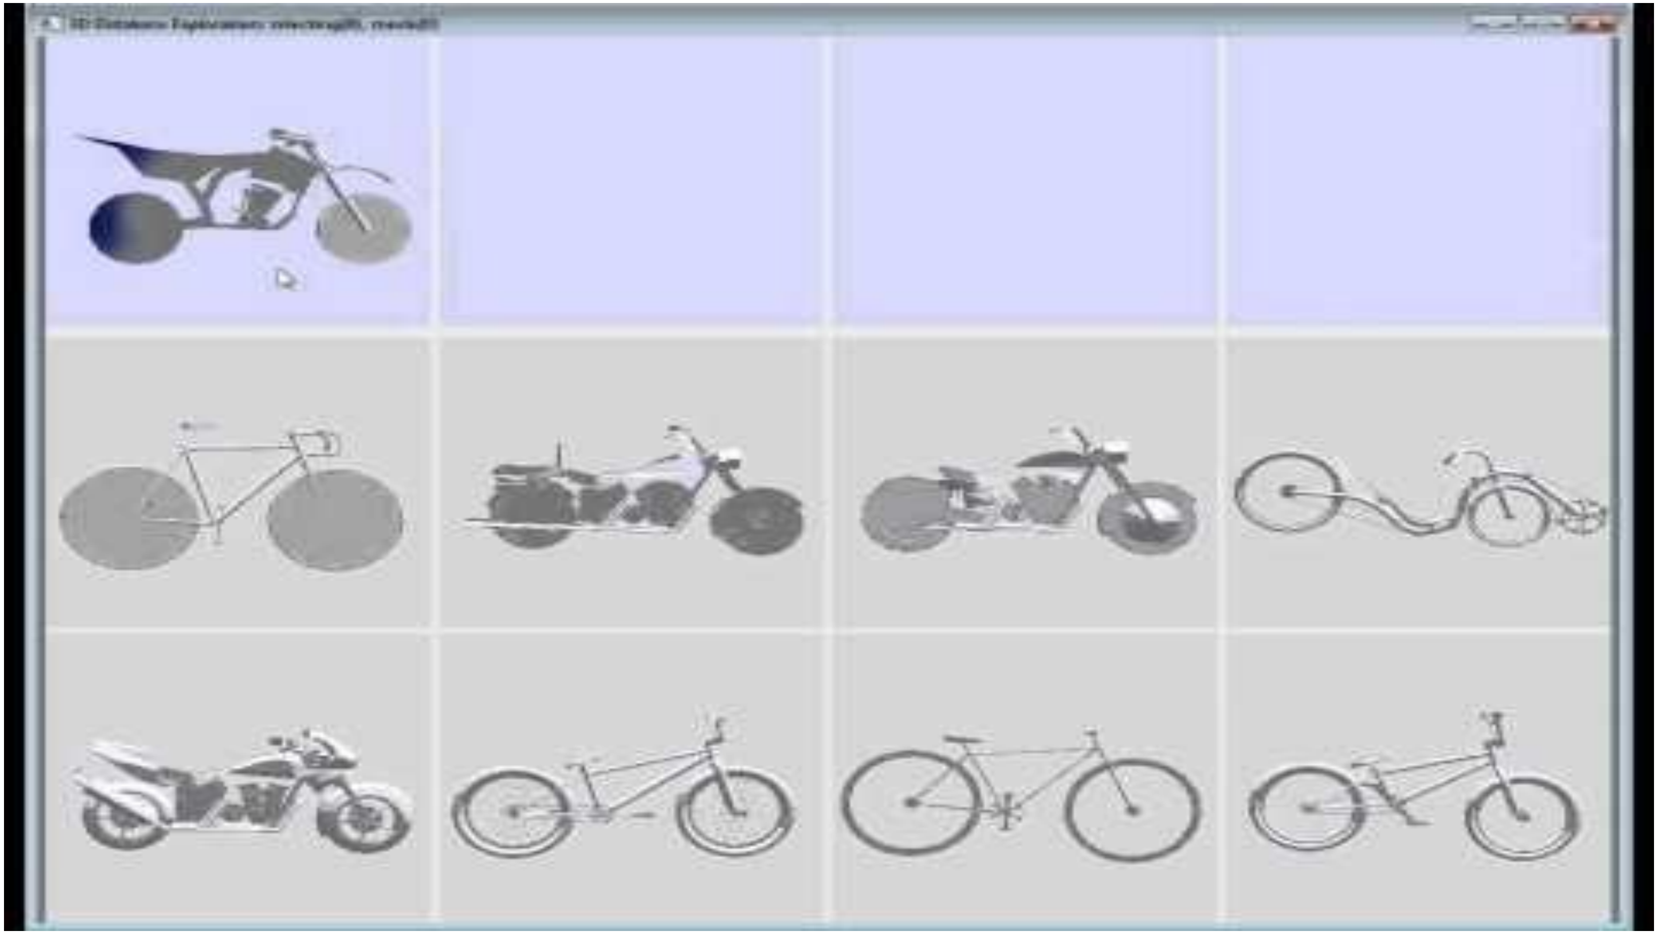

Exploring Collections of 3D Models using Fuzzy Correspondences (Kim et al., SIGGRAPH 2012)

Fuzzy Correspondences: Idea

- Compute Nk x Nk **similarity matrix**
  - Same number of samples per surface
  - Align similar shapes
- Compute **spectral** embedding
- Use as **descriptor**: Display $e^{-\left|d_i - d_j\right|^2}$

Consistent Segmentation

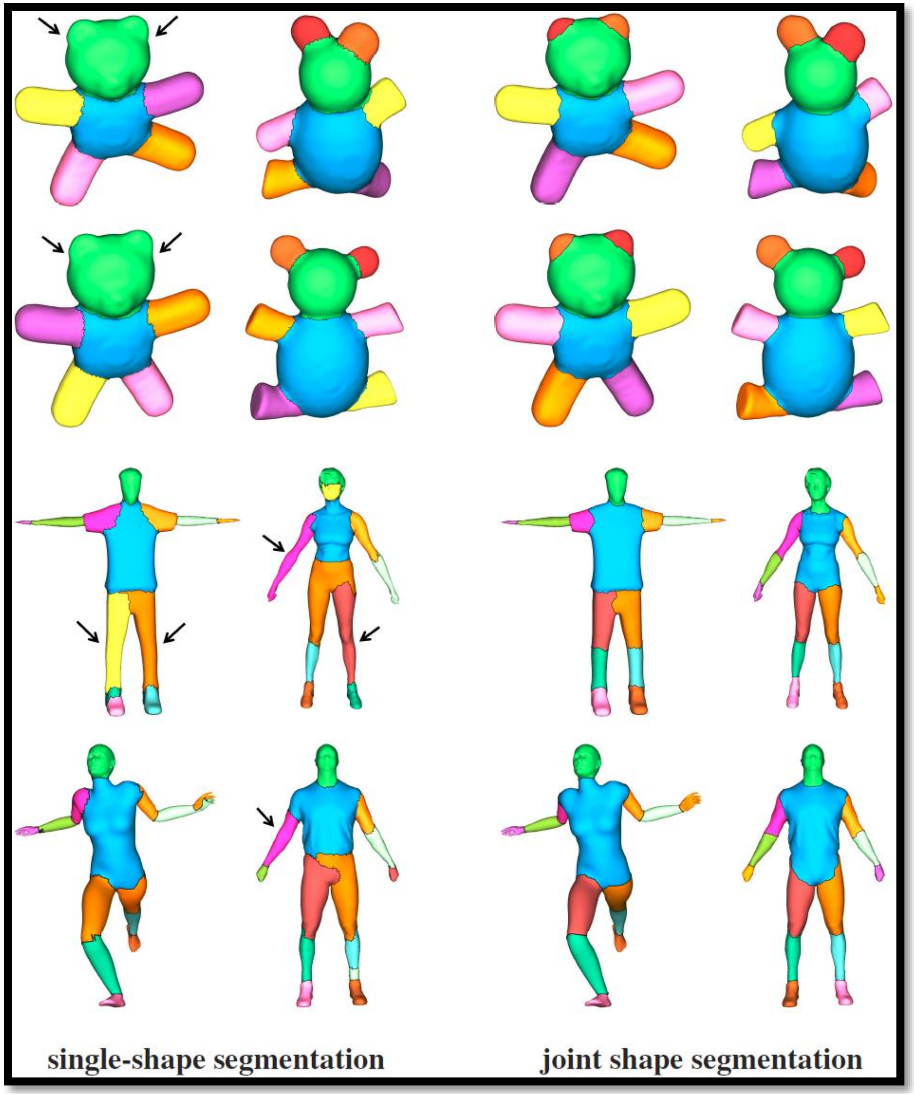

Global optimization to choose among many possible segmentations

"Joint Shape Segmentation with Linear Programming"
(Huang, Koltun, Guibas; SIGGRAPH Asia 2011)

Joint Segmentation: Motivation

Structural similarity of segmentations

- Extraneous geometric clues

Single shape segmentation
[Chen et al. 09]

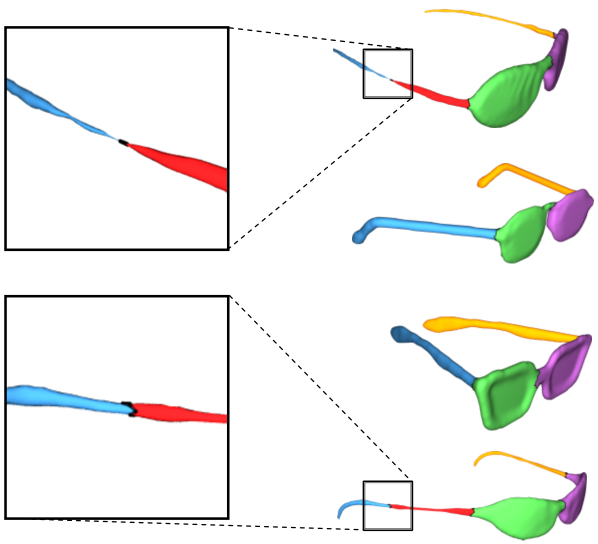

Joint shape segmentation
[Huang et al. 11]

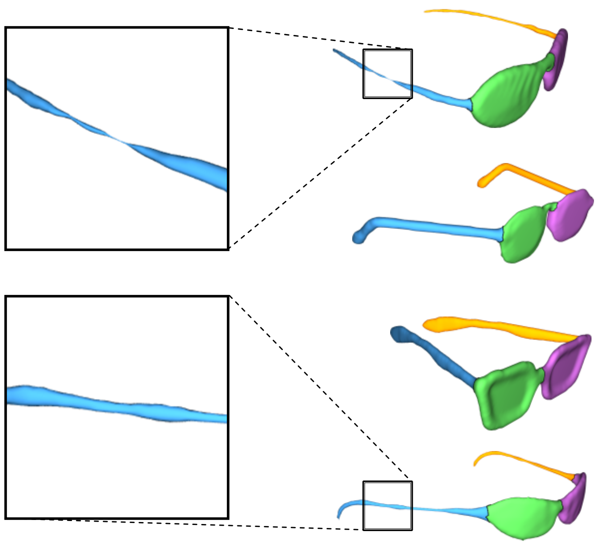

"Joint Shape Segmentation with Linear Programming"
(Huang, Koltun, Guibas; SIGGRAPH Asia 2011; slides provided by authors)

Joint Segmentation: Motivation

Structural similarity of segmentations

- Low saliency

Single shape segmentation
[Chen et al. 09]

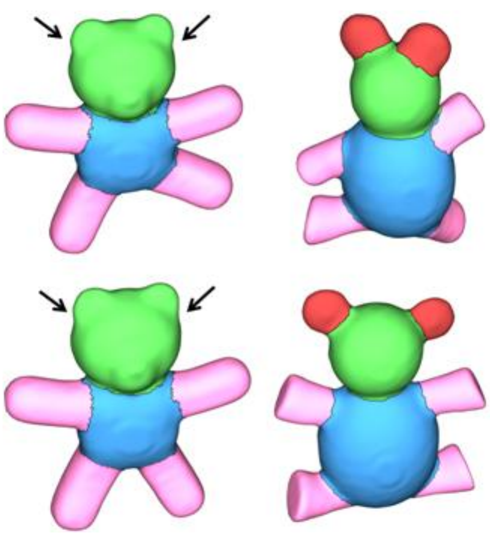

Joint shape segmentation
[Huang et al. 11]

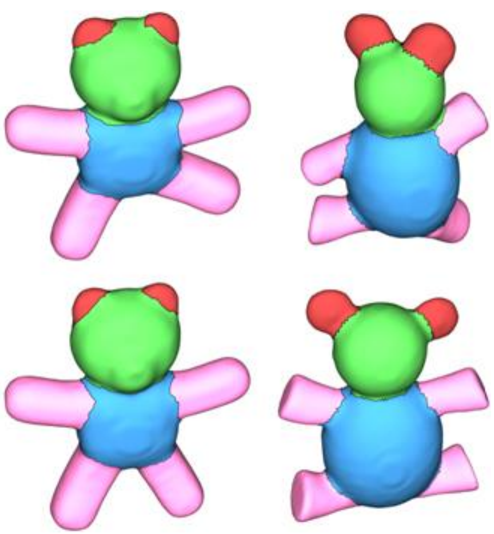

"Joint Shape Segmentation with Linear Programming"
(Huang, Koltun, Guibas; SIGGRAPH Asia 2011; slides provided by authors)

Joint Segmentation: Motivation

(Rigid) invariance of segments

- Articulated structures

Single shape segmentation
[Chen et al. 09]

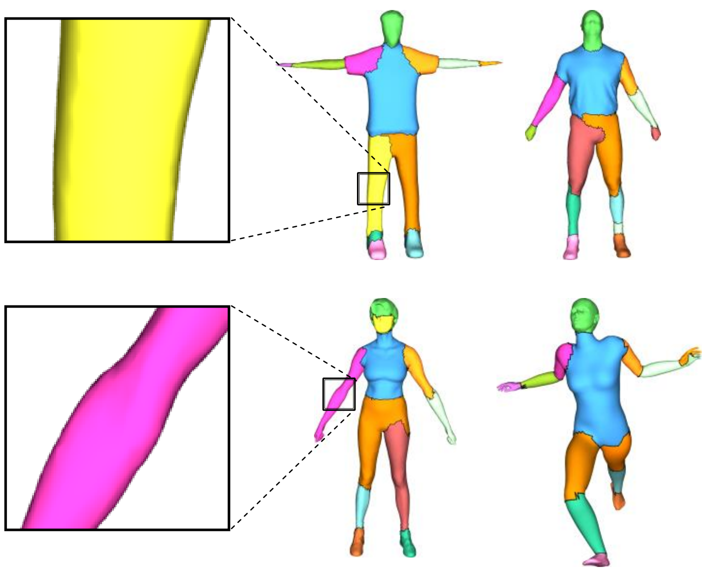

Joint shape segmentation
[Huang et al. 11]

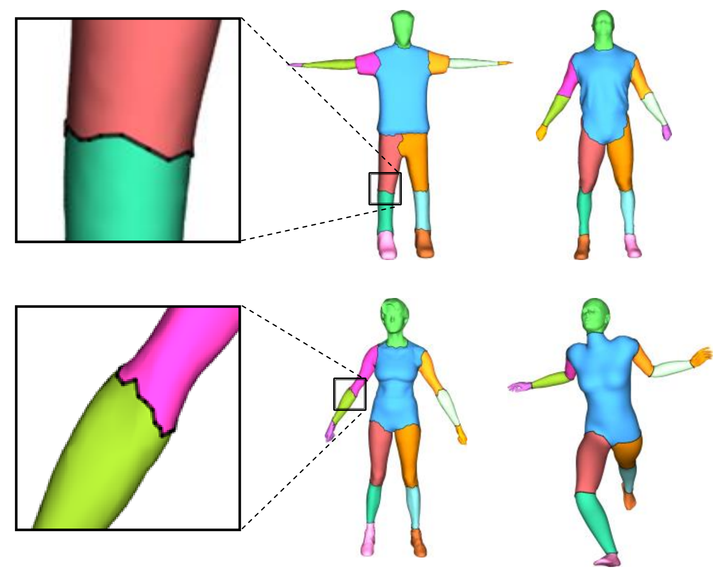

"Joint Shape Segmentation with Linear Programming"
(Huang, Koltun, Guibas; SIGGRAPH Asia 2011; slides provided by authors)

Parameterization

- Subsets of initial randomized segmentations

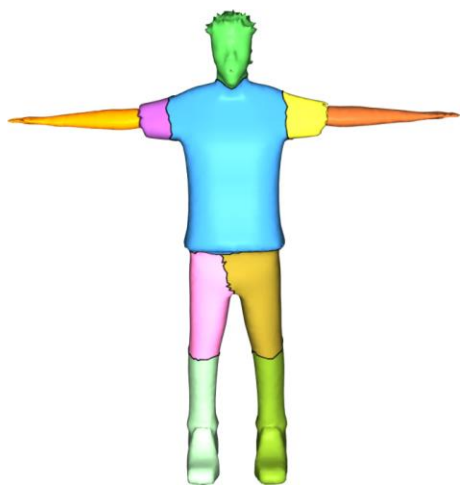

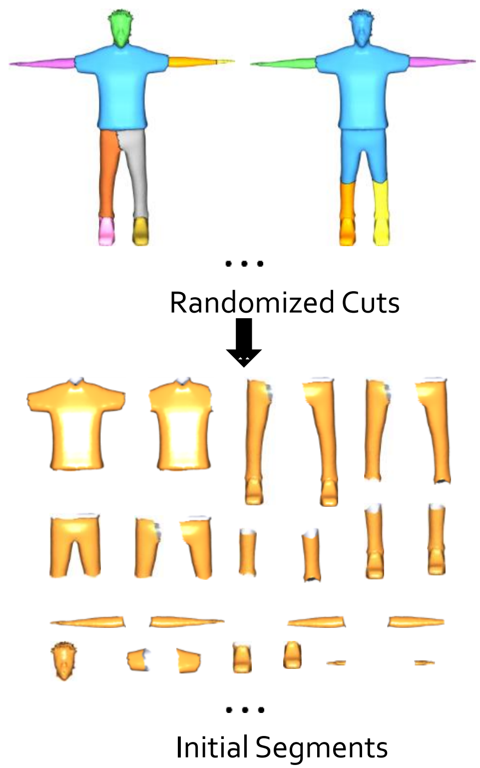

"Joint Shape Segmentation with Linear Programming"
(Huang, Koltun, Guibas; SIGGRAPH Asia 2011; slides provided by authors)

Segmentation Constraint/Score

- Each point covered by one segment

$$|\mathrm{cover}(p)| = 1 \ \forall p \in W$$

- Avoid tiny segments

$$\mathrm{score}(S) = \sum_{s \in S} \mathrm{area}(s) \cdot \mathrm{repetitions}_s$$

"Joint Shape Segmentation with Linear Programming"
(Huang, Koltun, Guibas; SIGGRAPH Asia 2011; slides provided by authors)

Consistency Term

- Defined in terms of mappings
  - Oriented
  - Partial

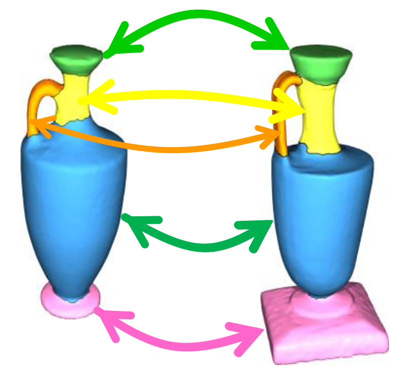

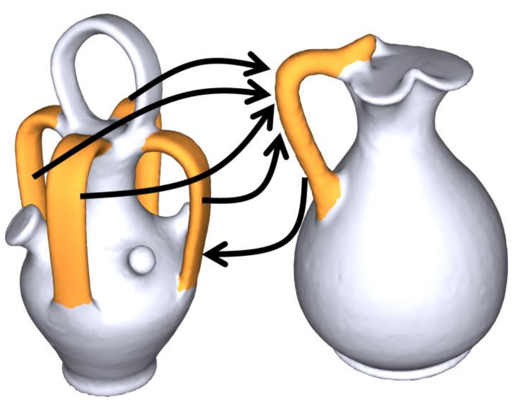

Many-to-one correspondences

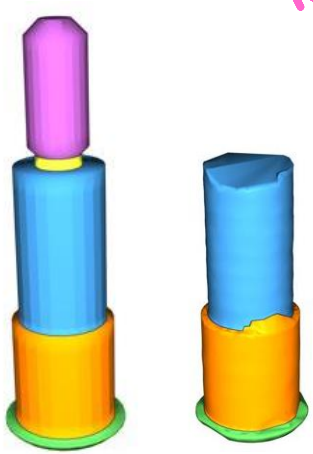

Partial similarity

"Joint Shape Segmentation with Linear Programming"
(Huang, Koltun, Guibas; SIGGRAPH Asia 2011; slides provided by authors)

Multi-Way Joint Segmentation

- Objective function

$$\sum_{i=1}^{n} \text{score}(S_i) + \sum_{(S_i,S_j)\in\mathcal{E}} \text{consistency}(S_i, S_j)$$

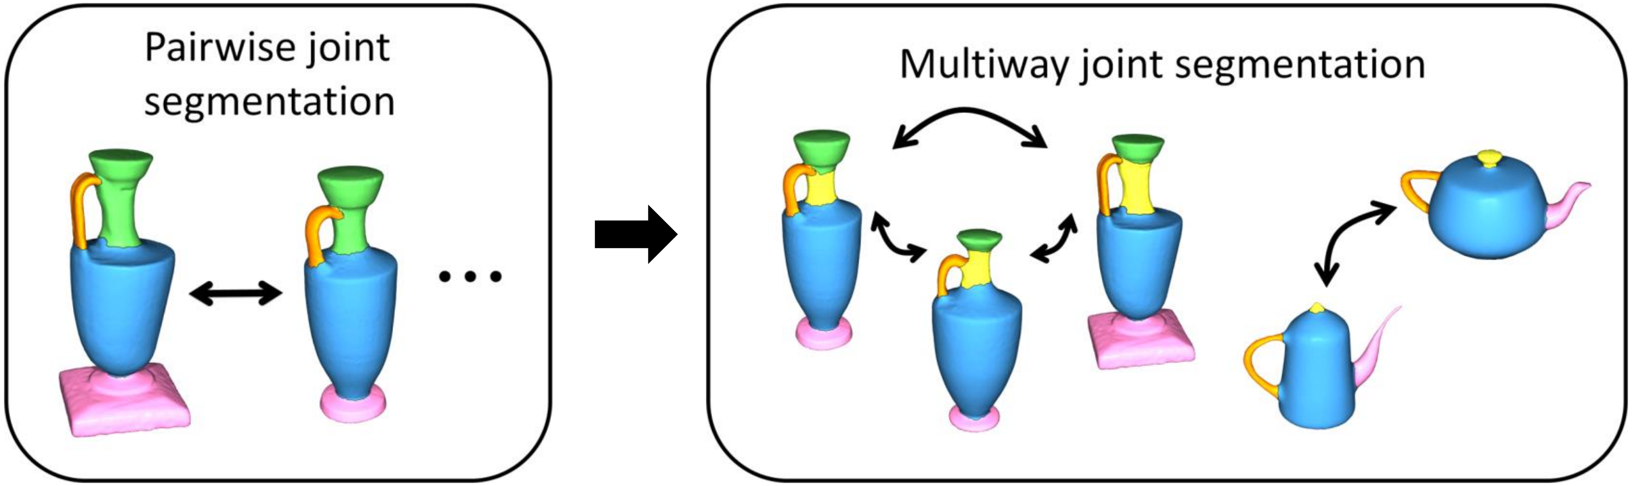

See paper: Linear program relaxation

Q:

Can you extract consistent maps in an optimal way?

Basic Setup

$$\begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 1 \\ 0 & 1 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 1 & \cdots & 0 \end{pmatrix}$$

Map as a permutation matrix

Q:

What is the inverse of a permutation matrix?

Discrete Relaxation

$$\begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 1 \\ 0 & 1 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 1 & \cdots & 0 \end{pmatrix}$$

Sums to 1

Sums to 1

Map as a doubly-stochastic matrix

Basic Setting

- Given n objects
- Each object sampled with m points

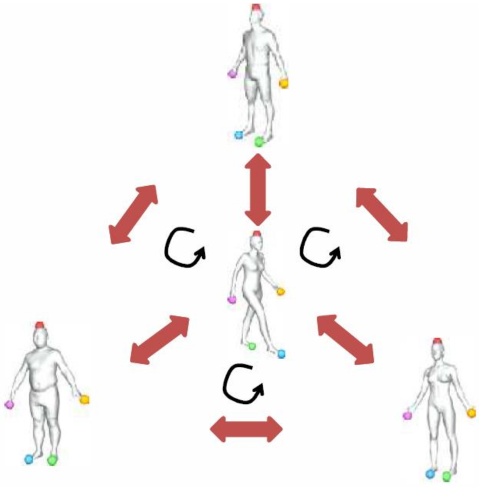

"Consistent Shape Maps via Semidefinite Programming" (Huang & Guibas, SGP 2013)

Map Collection: Matrix Representation

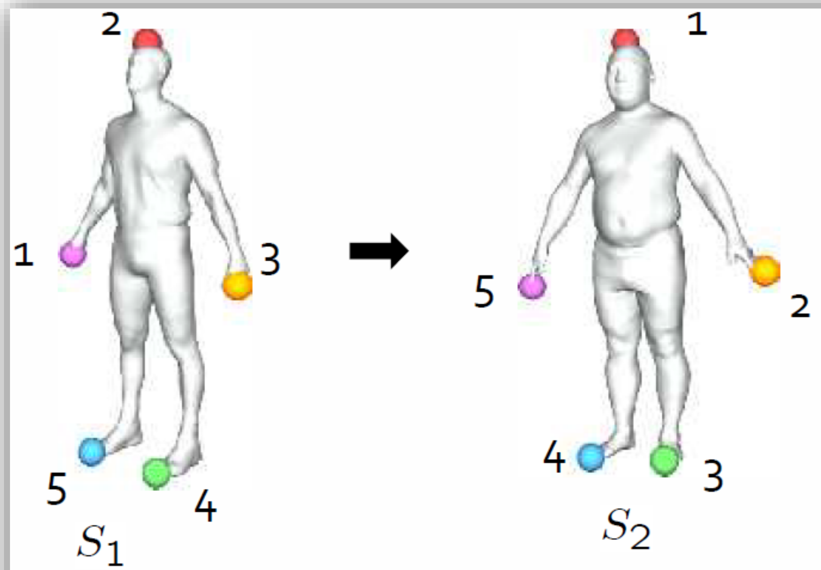

$$X_{12} = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

$$X = \begin{bmatrix} I_m & X_{12} & \cdots & X_{1n} \\ X_{12}^T & I_m & \cdots & \vdots \\ \vdots & \vdots & \ddots & X_{(n-1),n} \\ X_{1n}^T & \vdots & X_{(n-1),n}^T & I_m \end{bmatrix}$$

- Diagonal blocks are identity matrices
- Off diagonal blocks are permutation matrices
- Symmetric

Q:

What is the rank of a consistent map collection matrix?

Hint: "Urshape" Factorization

$$X = \begin{bmatrix} I_m & X_{12} & \cdots & X_{1n} \\ X_{12}^T & I_m & \cdots & \vdots \\ \vdots & \vdots & \ddots & X_{(n-1),n} \\ X_{1n}^T & \vdots & X_{(n-1),n}^T & I_m \end{bmatrix}$$

- Diagonal blocks are identity matrices
- Off diagonal blocks are permutation matrices
- Symmetric

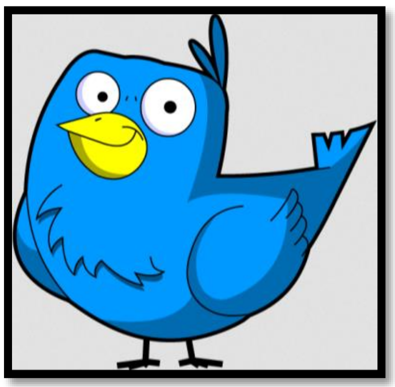

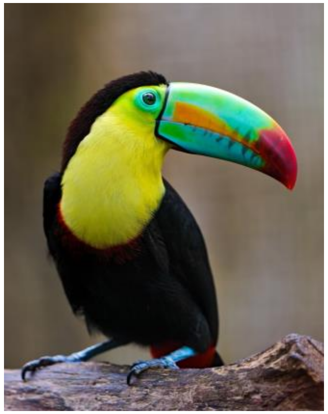

Rank *m*, Number of Samples

$$X_{ij} = X_{j1}^\top X_{i1} \iff X = \begin{pmatrix} I_m \\ \vdots \\ X_{n1}^\top \end{pmatrix} \begin{pmatrix} I_m & \cdots & X_{n1} \end{pmatrix}$$

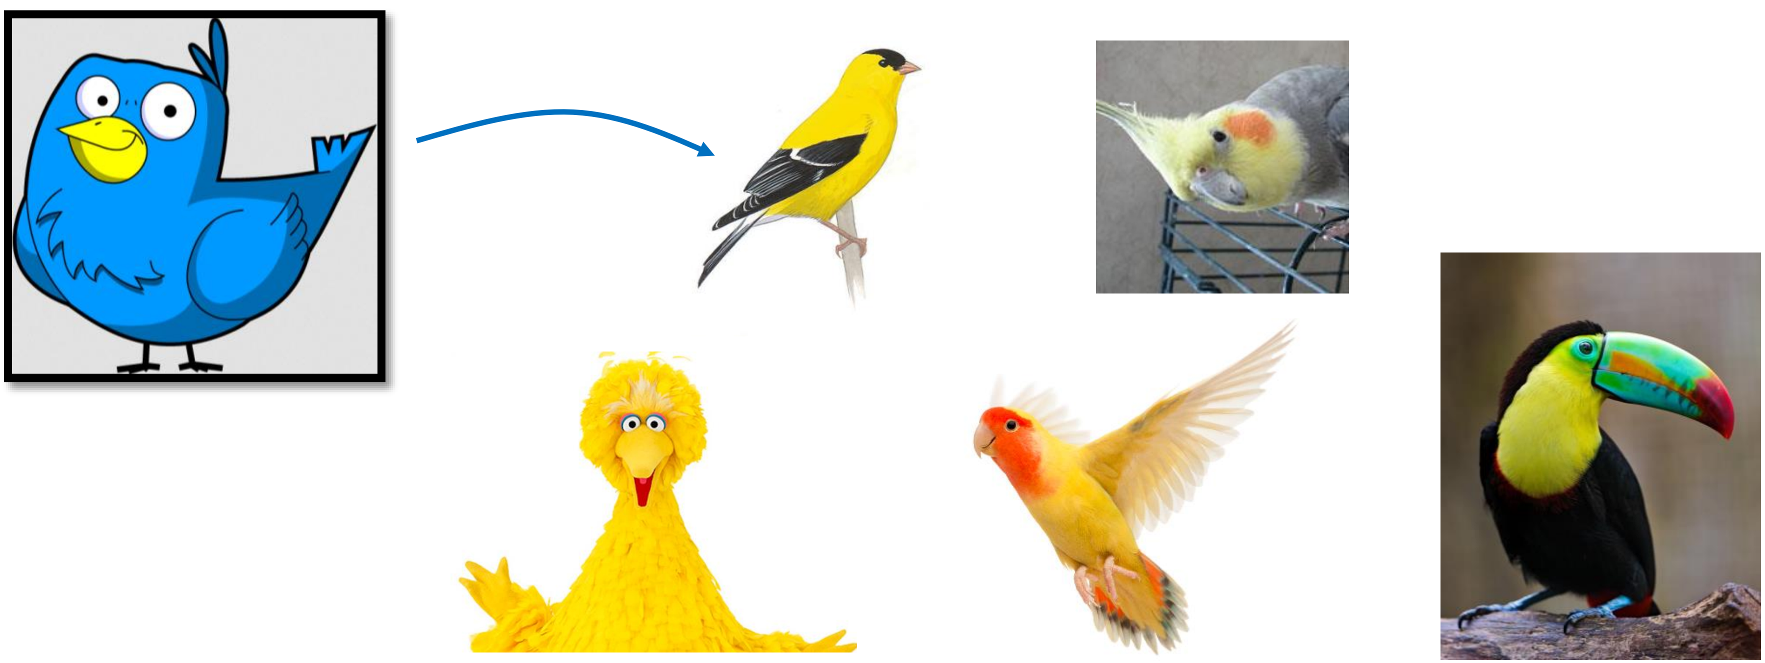

On the Board ...

Definition 2.1 Given a shape collection $\mathcal{S} = \{S_1, \cdots, S_n\}$ of n shapes where each shape consists of the same number of samples, we say a map collection $\Phi = \{\phi_{ij} : S_i \to S_j | 1 \leq i, j \leq n\}$ of maps between all pairs of shapes is cycle consistent if and only if the following equalities are satisfied:

$$\phi_{ii} = id_{S_i}, \quad 1 \leq i \leq n, \qquad \text{(1-cycle)}$$
$$\phi_{ji} \circ \phi_{ij} = id_{S_i}, \quad 1 \leq i < j \leq n, \qquad \text{(2-cycle)}$$
$$\phi_{ki} \circ \phi_{jk} \circ \phi_{ij} = id_{S_i}, \quad 1 \leq i < j < k \leq n, \qquad \text{(3-cycle)} \quad (1)$$

where id_{S_i} denotes the identity self-map on S_i.

Equivalence for binary map matrix Φ:

1. Φ is cycle-consistent
2. $X = Y_i^\top Y_i$, where $Y_i = (X_{i1}, \ldots, X_{in})$
3. $X \succeq 0$

Approximation by Consistent Maps

$$
\begin{aligned}
\max_X \quad & \sum_{ij \in E} \langle X_{ij}^{\text{in}}, X_{ij} \rangle \\
\text{s.t.} \quad & X \in \{0,1\}^{nm \times nm} \\
& X \succeq 0 \\
& X_{ii} = I_m \\
& X_{ij}\mathbf{1} = \mathbf{1} \\
& X_{ij}^\top \mathbf{1} = \mathbf{1}
\end{aligned}
$$

Approximation by Consistent Maps

$$\begin{aligned} \max_X \quad & \sum_{ij \in E} \langle X_{ij}^{\text{in}}, X_{ij} \rangle \\ \text{s.t.} \quad & X \in \{0,1\}^{nm \times nm} \\ & X \succeq 0 \\ & X_{ii} = I_m \\ & X_{ij}\mathbf{1} = \mathbf{1} \\ & X_{ij}^{\top}\mathbf{1} = \mathbf{1} \end{aligned}$$

Maximize number of preserved matches

Approximation by Consistent Maps

$$
\begin{aligned}
\max_X \quad & \sum_{ij \in E} \langle X_{ij}^{\text{in}}, X_{ij} \rangle \\
\text{s.t.} \quad & X \in \{0,1\}^{nm \times nm} \\
& X \succeq 0 \\
& X_{ii} = I_m \\
& X_{ij} \mathbf{1} = \mathbf{1} \\
& X_{ij}^\top \mathbf{1} = \mathbf{1}
\end{aligned}
$$

Binary matrix

Approximation by Consistent Maps

$$
\begin{aligned}
\max_X \quad & \textstyle\sum_{ij \in E} \langle X_{ij}^{\text{in}}, X_{ij} \rangle \\
\text{s.t.} \quad & X \in \{0,1\}^{nm \times nm} \\
& X \succeq 0 \\
& X_{ii} = I_m \\
& X_{ij} 1 = 1 \\
& X_{ij}^{\top} 1 = 1
\end{aligned}
$$

Every block is a permutation

Approximation by Consistent Maps

$$
\begin{aligned}
\max_X \quad & \textstyle\sum_{ij \in E} \langle X_{ij}^{\text{in}}, X_{ij} \rangle \\
\text{s.t.} \quad & X \in \{0,1\}^{nm \times nm} \\
& X \succeq 0 \\
& X_{ii} = I_m \\
& X_{ij} \mathbf{1} = \mathbf{1} \\
& X_{ij}^\top \mathbf{1} = \mathbf{1}
\end{aligned}
$$

Self maps are identity

Approximation by Consistent Maps

$$\begin{aligned} \max_X \quad & \textstyle\sum_{ij \in E} \langle X_{ij}^{\text{in}}, X_{ij} \rangle \\ \text{s.t.} \quad & X \in \{0,1\}^{nm \times nm} \\ & X \succeq 0 \\ & X_{ii} = I_m \\ & X_{ij}\mathbf{1} = \mathbf{1} \\ & X_{ij}^\top \mathbf{1} = \mathbf{1} \end{aligned}$$

Already showed:
Equivalent to low-rank

Approximation by Consistent Maps

$$
\begin{aligned}
\max_X \quad & \sum_{ij \in E} \langle X_{ij}^{\text{in}}, X_{ij} \rangle \\
\text{s.t.} \quad & X \in \{0,1\}^{nm \times nm} \\
& X \succeq 0 \\
& X_{ii} = I_m \\
& X_{ij} \mathbf{1} = \mathbf{1} \\
& X_{ij}^\top \mathbf{1} = \mathbf{1}
\end{aligned}
$$

Nonconvex!

Convex Relaxation

$$
\begin{aligned}
\max_X \quad & \sum_{ij \in E} \langle X_{ij}^{\text{in}}, X_{ij} \rangle \\
\text{s.t.} \quad & X \geq 0 \\
& X \succeq 0 \\
& X_{ii} = I_m \\
& X_{ij}\mathbf{1} = \mathbf{1} \\
& X_{ij}^\top \mathbf{1} = \mathbf{1}
\end{aligned}
$$

Rounding Procedure

Guaranteed to give permutation

$$\begin{aligned} \max_X \quad & \langle X, X_0 \rangle \\ \text{s.t.} \quad & X \geq 0 \\ & X\mathbf{1} = \mathbf{1} \\ & X^\top \mathbf{1} = \mathbf{1} \end{aligned}$$

Linear assignment problem

Recovery Theorem

Can tolerate $\lambda_2/4(n-1)$ incorrect correspondences from each sample on one shape.

λ_2 is algebraic connectivity; bounded above by two times maximum degree

`\omit{proof}`

Recovery Theorem: Complete Graph

Can tolerate 25% incorrect correspondences from each sample on one shape.

λ_2 is algebraic connectivity; bounded above by two times maximum degree

\omit{proof}

Phase Transition

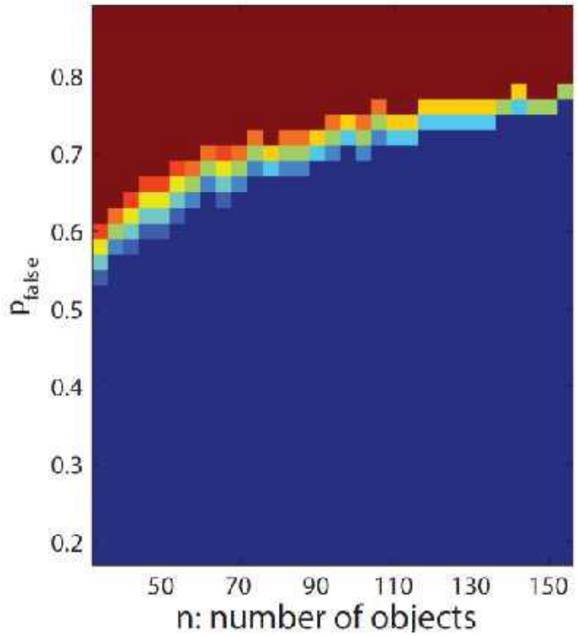

Always recovers / Never recovers

Example Result

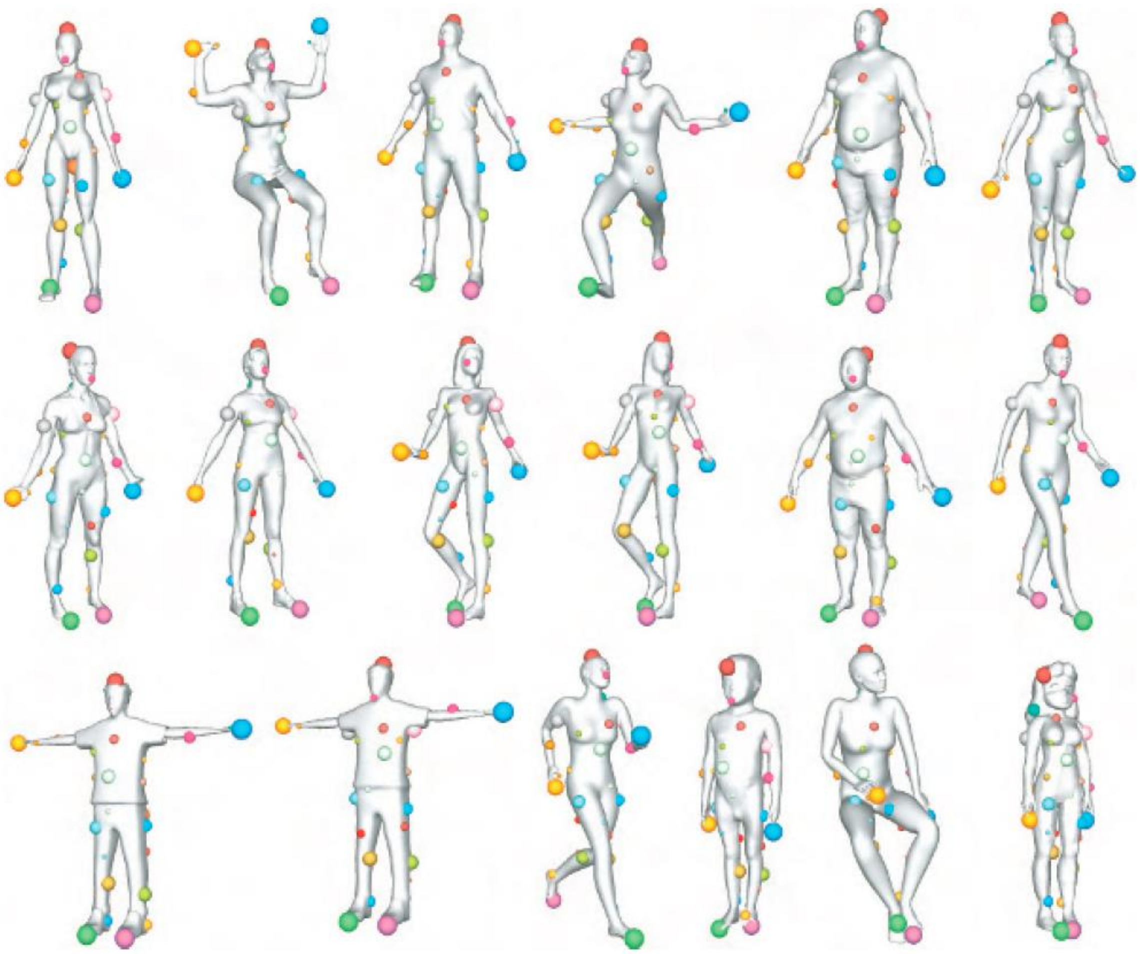

Q:

Where do the pairwise **input maps** come from?

Possible Extension with Guarantees

Eurographics Symposium on Geometry Processing 2015
Mirela Ben-Chen and Ligang Liu
(Guest Editors)

Volume 34 (2015), Number 5

Tight Relaxation of Quadratic Matching

Itay Kezurer† Shahar Z. Kovalsky† Ronen Basri Yaron Lipman

Weizmann Institute of Science

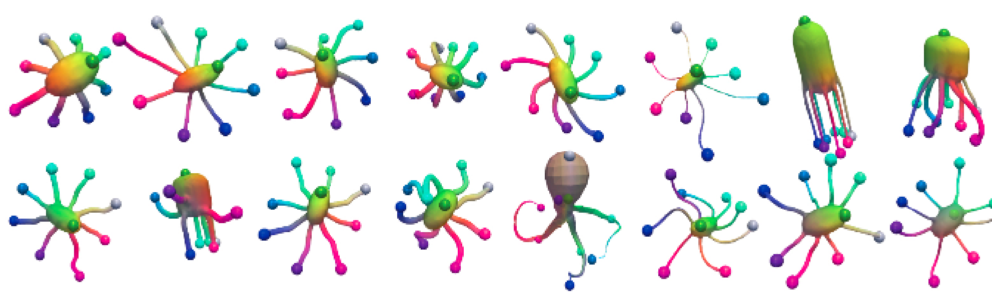

Figure 1: Consistent Collection Matching. Results of the proposed one-stage procedure for finding consistent correspondence between shapes in a collection showing strong variability and non-rigid deformations.

Abstract

Establishing point correspondences between shapes is extremely challenging as it involves both finding sets of semantically persistent feature points, as well as their combinatorial matching. We focus on the latter and consider the Quadratic Assignment Matching (QAM) model. We suggest a novel convex relaxation for this NP-hard problem that builds upon a rank-one reformulation of the problem in a higher dimension, followed by relaxation into a semidefinite program (SDP). Our method is shown to be a certain hybrid of the popular spectral and doubly-stochastic relaxations of QAM and in particular we prove that it is tighter than both.

Experimental evaluation shows that the proposed relaxation is extremely tight: in the majority of our experiments it achieved the certified global optimum solution for the problem, while other relaxations tend to produce sub-optimal solutions. This, however, comes at the price of solving an SDP in a higher dimension.

Our approach is further generalized to the problem of Consistent Collection Matching (CCM), where we solve

Heavy optimization problem!

$$
\begin{aligned}
\max_{Y} \quad & \mathrm{tr}\,(WY) && \text{(7a)}\\
\text{s.t.} \quad & Y \succeq [X][X]^T && \text{(7b)}\\
& X \in \mathrm{conv}\,\Pi_n^k && \text{(7c)}\\
& \mathrm{tr}Y = k && \text{(7d)}\\
& Y \geq 0 && \text{(7e)}\\
& \sum_{qrst} Y_{qrst} = k^2 && \text{(7f)}\\
& Y_{qrst} \leq \begin{cases} 0, & \text{if } q = s,\ r \neq t \\ 0, & \text{if } r = t,\ q \neq s \\ \min\{X_{qr}, X_{st}\}, & \text{otherwise} \end{cases} && \text{(7g)}
\end{aligned}
$$

$$
\begin{aligned}
\max_{\mathbf{X},\mathbf{Y}} \quad & \sum_{i,j} \mathrm{tr}\left(W^{ij} Y^{ij}\right) && && \text{(10a)}\\
\text{s.t.} \quad & \left(X^{ij}, Y^{ij}\right) \in \mathcal{C}^k && \forall i < j && \text{(10b)}\\
& X^{ii} \in \mathcal{D} \cap \mathrm{conv}\,\Pi_n^k && \forall i && \text{(10c)}\\
& \mathbf{X} \succeq 0 && && \text{(10d)}
\end{aligned}
$$

Approximate Methods

Consistent Partial Matching of Shape Collections via Sparse Modeling

L. Cosmo[1], E. Rodolà[2], A. Albarelli[1], F. Mémoli[3], D. Cremers[2]

[1]University of Venice, Italy [2]TU Munich, Germany [3]Ohio State University, U.S.

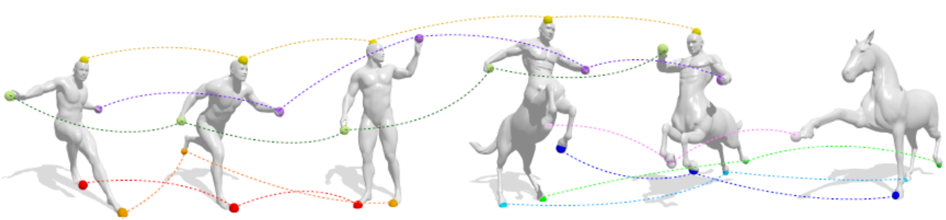

Figure 1: *A partial multi-way correspondence obtained with our approach on a heterogeneous collection of shapes. Our method does not require initial pairwise maps as input, as it actively seeks a reliable corresponde… space of joint, cycle-consistent matches. Partially-similar as well as outlier shapes are aut… for by adopting a sparse model for the joint correspondence. A subset of all matches is show…*

Abstract

Recent efforts in the area of joint object matching approach the problem by taking as … which are then jointly optimized across the whole collection so that certain accuracy … satisfied. One natural requirement is cycle-consistency – namely the fact that map … same result regardless of the path taken in the shape collection. In this paper, we in… obtain consistent matches without requiring initial pairwise solutions to be given as i… a joint measure of metric distortion directly over the space of cycle-consistent maps; in… similar and extra-class shapes, we formulate the problem as a series of quadratic prog… constraints, making our technique a natural candidate for analyzing collections with … The particular form of the problem allows us to leverage results and tools from the … theory. This enables a highly efficient optimization procedure which assures accur… solutions in a matter of minutes in collections with hundreds of shapes.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphic… and Object Modeling—Shape Analysis

1. Introduction

Finding matches among multiple objects is a research topic …

this end, a natural and widely accepted criterion is *cycle-consistency* [ZKP10], namely that composition of maps …

Sequence of quadratic programs; based on metric distortion and WKS descriptor match

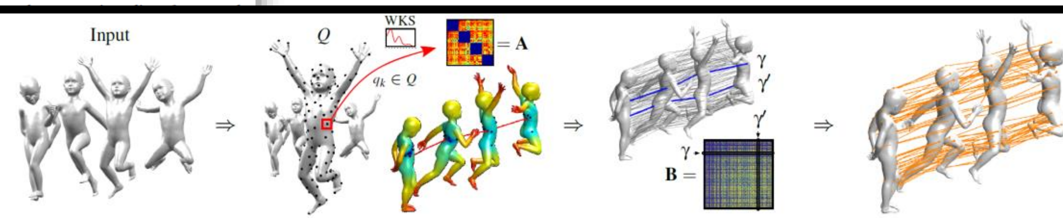

Figure 6: *Our matching pipeline. First sub-problem (from left): Given a collection of shapes as input, a set Q of queries are generated (e.g., by farthest point sampling in the joint WKS space); we then compute distance maps (shown here as heat maps over the shapes) in descriptor space from each shape point to each query $q_k \in Q$, and keep the vertices having distance smaller than a threshold; finally, a single multi-way match is extracted by solving problem (11). Second sub-problem: The multi-way matches extracted by iterating the previous step are compared using a measure of metric distortion; the final solution (in orange) is obtained by solving problem (13) over the reduced feasible set.*

CGF 2017

Approximate Methods

Multiplicative updates for nonconvex nonnegative matrix factorization

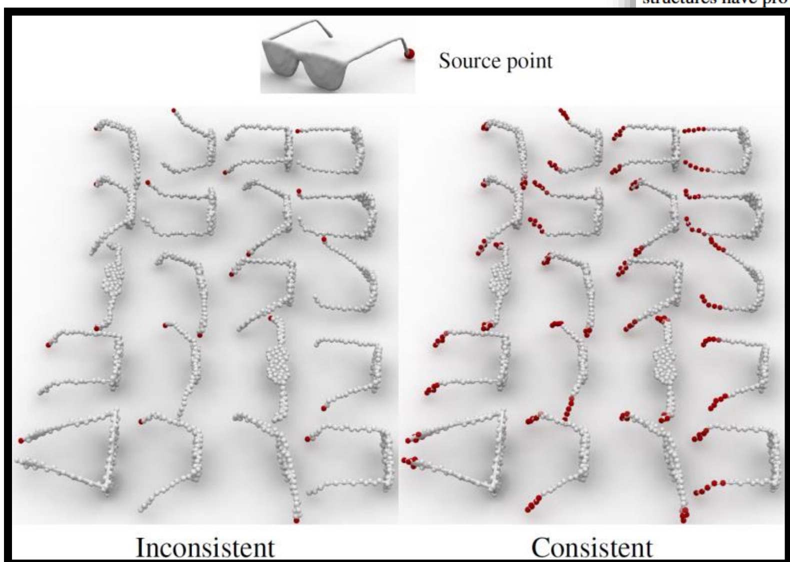

$$\min_{A} \mathrm{KL}(G|AA^{\top})$$

Entropic Metric Alignment for Correspondence Problems

Justin Solomon* MIT
Gabriel Peyré CNRS & Univ. Paris-Dauphine
Vladimir G. Kim Adobe Research
Suvrit Sra MIT

Abstract

Many shape and image processing tools rely on computation of correspondences between geometric domains. Efficient methods that stably extract "soft" matches in the presence of diverse geometric structures have proven to be valuable for shape retrieval and transfer
ic information. With these applications in mind,
ithm for probabilistic correspondence that opti-
egularized Gromov-Wasserstein (GW) objective.
evelopments in numerical optimal transportation,
mpact, provably convergent, and applicable to
ain expressible as a metric measure matrix. We
nsive experiments illustrating the convergence
f our algorithm to a variety of graphics tasks.
xpand entropic GW correspondence to a frame-
ching problems, incorporating partial distance
nce, shape exploration, symmetry detection, and
re than two domains. These applications expand
ic GW correspondence to major shape analysis
able to distortion and noise.

ov-Wasserstein, matching, entropy
uting methodologies → Shape analysis;

on

of the geometry processing toolbox is a tool for
ondence, the problem of finding which points on
respond to points on a source. Many variations
e been considered in the graphics literature, e.g.
with some sparse correspondences provided by the user. Regardless, the basic task of geometric correspondence facilitates the transfer of properties and edits from one shape to another.

The primary factor that distinguishes correspondence algorithms is the choice of objective functions. Different choices of objective functions express contrasting notions of which correspondences are "desirable." Classical theorems from differential geometry and most modern algorithms consider *local* distortion, producing maps that take tangent planes to tangent planes with as little stretch as possible; slightly larger neighborhoods might be taken into account by e.g.

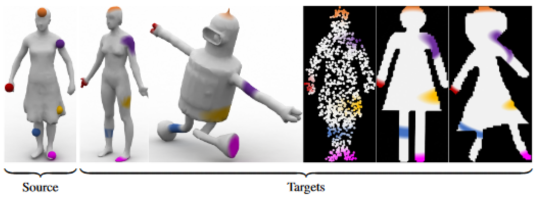

Figure 1: *Entropic GW can find correspondences between a source surface (left) and a surface with similar structure, a surface with shared semantic structure, a noisy 3D point cloud, an icon, and a hand drawing. Each fuzzy map was computed using the same code.*

are violated these algorithms suffer from having to patch together local elastic terms into a single global map.

In this paper, we propose a new correspondence algorithm that minimizes distortion of long- and short-range distances alike. We study an entropically-regularized version of the *Gromov-Wasserstein* (GW) mapping objective function from [Mémoli 2011] measuring the distortion of geodesic distances. The optimizer is a probabilistic matching expressed as a "fuzzy" correspondence matrix in the style of [Kim et al. 2012; Solomon et al. 2012]; we control sharpness of the correspondence via the weight of an entropic regularizer.

Although [Mémoli 2011] and subsequent work identified the possibility of using GW distances for geometric correspondence, computational challenges hampered their practical application. To overcome these challenges, we build upon recent methods for regularized optimal transportation introduced in [Benamou et al. 2015; Solomon et al. 2015]. While optimal transportation is a fundamentally different optimization problem from regularized GW computation (linear versus quadratic matching), the core of our method relies upon solving a sequence of regularized optimal transport problems.

Our remarkably compact algorithm (see Algorithm 1) exhibits global convergence, i.e., it *provably* reaches a local minimum of the regularized GW objective function regardless of the initial guess. Our algorithm can be applied to any domain expressible as a metric measure space (see §2). Concretely, only distance matrices are required as input, and hence the method can be applied to many classes of domains including meshes, point clouds, graphs, and even more

Computer Vision Application

Learning Dense Correspondence via 3D-guided Cycle Consistency

Tinghui Zhou, UC Berkeley
Philipp Krähenbühl, UC Berkeley
Mathieu Aubry, ENPC ParisTech
Qixing Huang, TTI-Chicago
Alexei A. Efros, UC Berkeley

Abstract

Discriminative deep learning approaches have shown impressive results for problems where human-labeled ground truth is plentiful, but what about tasks where labels are difficult or impossible to obtain? This paper tackles one such problem: establishing dense visual correspondence across different object instances. For this task, although we do not know what the ground-truth is, we know it should be ***consistent*** *across instances of that category. We exploit this consistency as a supervisory signal to train a convolutional neural network to predict cross-instance correspondences between pairs of images depicting objects of the same category. For each pair of training images we find an appropriate 3D CAD model and render two synthetic views to link in with the pair, establishing a correspondence flow 4-cycle. We use ground-truth synthetic-to-synthetic correspondences, provided by the rendering engine, to train a ConvNet to predict synthetic-to-real, real-to-real and real-to-synthetic correspondences that are cycle-consistent with the ground-truth. At test time, no CAD models are required. We demonstrate that our end-to-end trained ConvNet supervised by cycle-consistency outperforms state-of-the-art pairwise matching methods in correspondence-related tasks.*

1. Introduction

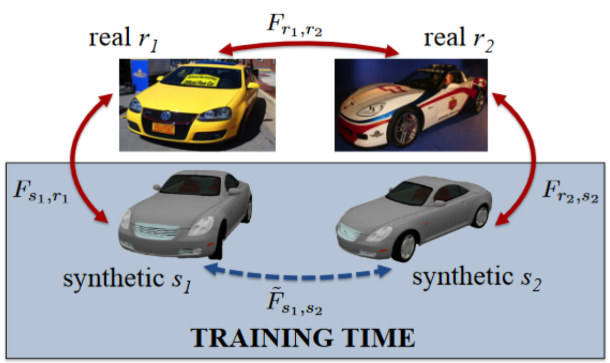

Figure 1. Estimating a dense correspondence flow field F_{r_1,r_2} between two images r_1 and r_2 — essentially, where do pixels of r_1 need to go to bring them into correspondence with r_2 — is very difficult. There is a large viewpoint change, and the physical differences between the cars are substantial. We propose to *learn* to do this task by training a ConvNet using the concept of cycle consistency in lieu of ground truth. At training time, we find an appropriate 3D CAD model to establish a correspondence 4-cycle, and train the ConvNet to minimize the discrepancy between $\tilde{F}_{s_1,s_2}$ and $F_{s_1,r_1} \circ F_{r_1,r_2} \circ F_{r_2,s_2}$, where $\tilde{F}_{s_1,s_2}$ is known by construction. At test time, no CAD models are used.

matching, but many other computer vision tasks, including recognition, segmentation, depth estimation, etc. could be posed as finding correspondences in a large visual database followed by label transfer.

[Wang et al. 2012]

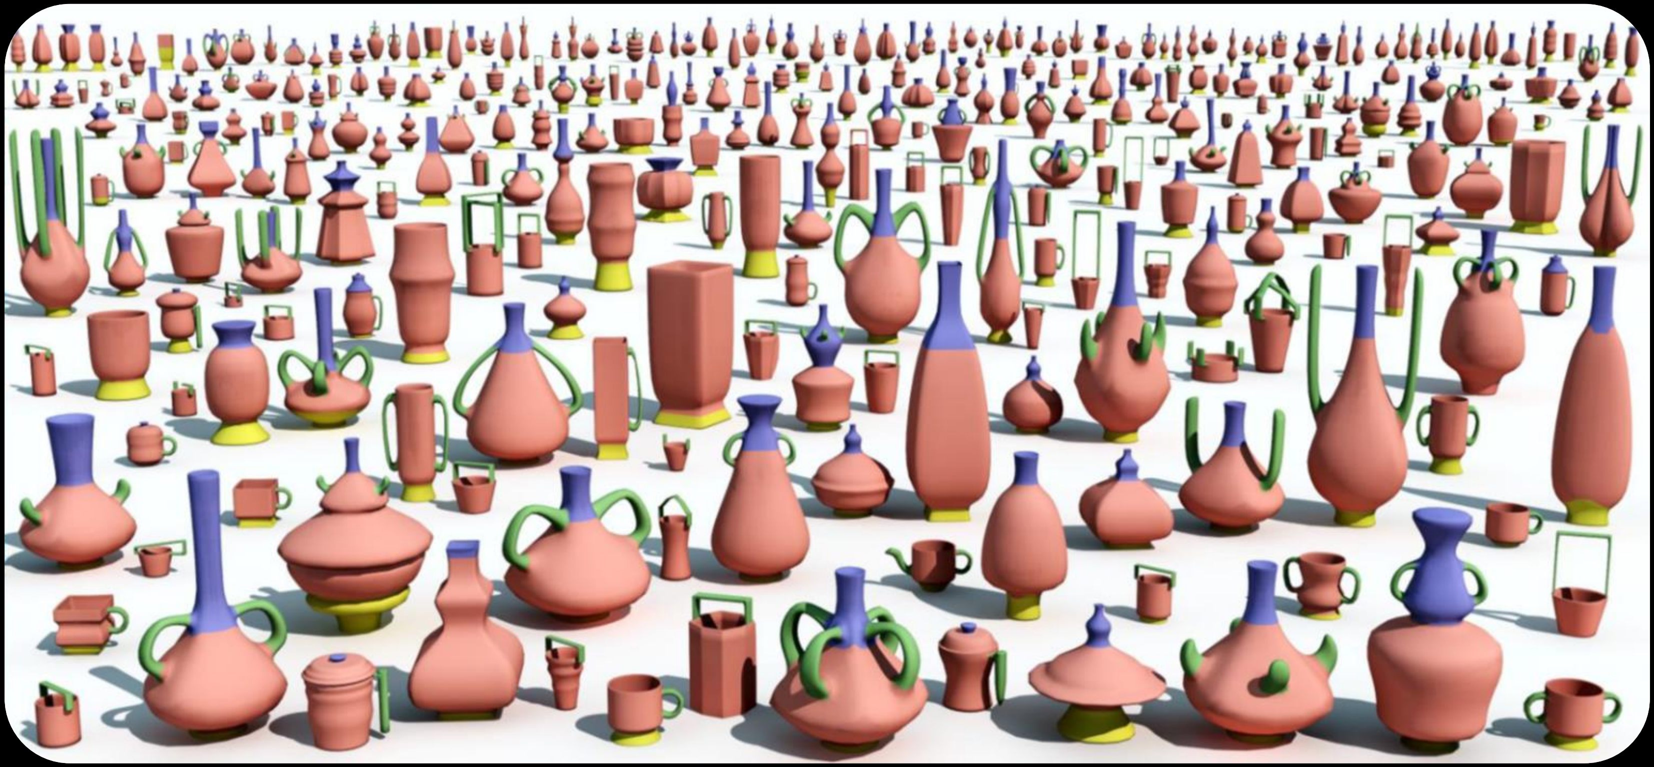

Consistent Correspondence

Justin Solomon

MIT, Spring 2017

Many slides from CS 468, Stanford (Kim & Huang)