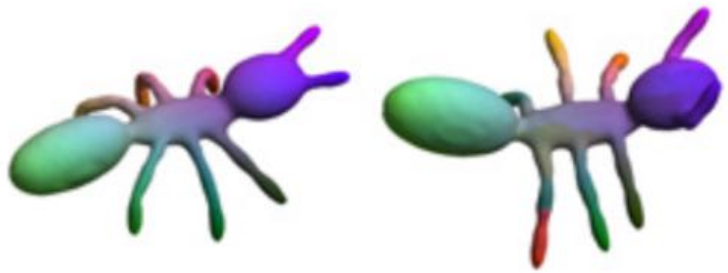

3D Shape Analysis

Summer Semester 2018

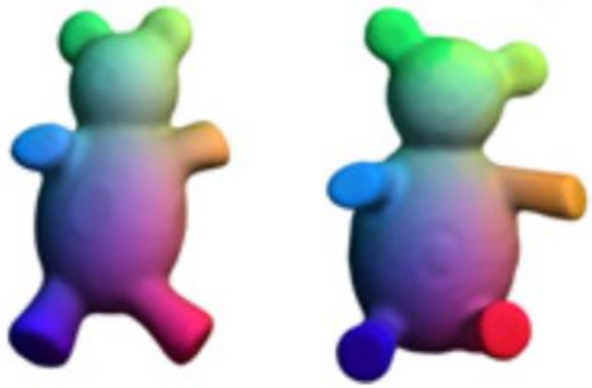

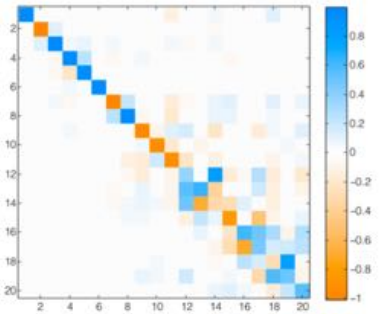

Dr. Florian Bernard

Prof. Dr. Christian Theobalt

Graphics, Vision and Video Group, MPI Informatik

Graphics Vision & Video

max planck institut informatik

Overview

- Organization
- Introduction
- Topics
- Summary

Overview

- Organization
- Introduction
- Topics
- Summary

Organizers

Florian Bernard
MPI Informatik, room 212
fbernard@mpi-inf.mpg.de

Christian Theobalt
MPI Informatik, room 228
theobalt@mpi-inf.mpg.de

Organizational contact: sa-seminar@lists.mpi-inf.mpg.de

Basic Coordinates

- Time: Thursdays, 14:15 – 16:00
- Place: MPI Informatik (E1 4), rooms 019 / 021
- Website:
  - http://gvv.mpi-inf.mpg.de/teaching/gvv_seminar_2018/

- Slides will be made available on the website
  - user: **shape**
  - password: **procrustes**

Formal requirements in a nutshell

- You read all the papers
- Your presence is compulsory
  - We will keep track of attendance.
- Submit questions & participate in discussion
- One topic is “Your Topic” (2 papers):
  - Deliver a **45** minute presentation
  - Write a 5–7 page report
- Grade: talk 30%, discussion 30%, report 40%

Prior knowledge

- Not for novices in computer science / visual computing
- You need to have an understanding of:
  - mathematics (linear algebra & calculus)
  - at least one of the following
    - convex/non-convex optimization
    - 3D geometry
    - computer vision/image processing
    - computer graphics
- Examples: you should know about…
  - eigendecomposition
  - gradients
  - more about this later

Registration

- Register your interest by email – sa-seminar@lists.mpi-inf.mpg.de
  - Name, matriculation number, degree program, semester, previous courses and/or experience (if you haven't done this yet)
  - So far, there are 11 interested people (10 "reserved slots", 1 on waiting list)
- Email us **AFTER** this lecture to get a topic assignment (=final registration)
  - first-come first-served basis
  - "reserved slots" will be filled before considering the waiting list
  - if you have a "reserved slot" and do not use it, let us know immediately
- Topic assignment:
  - Send a list of your 3 preferred topics (in order of preference) **until Sunday, 15 April 2018**
  - We will try to accommodate wishes as much as possible
- Register in HISPOS within 3 weeks (email to come)

Organization

- 13 topics to choose from
  - listed on seminar website + introduced later today
- Maximum of 10 presentation slots:
  - First presentation: Thursday, 03 May 2018
  - Each week until Thursday, 19 July 2018 (including)
    - Except: 10 May and 31 May (no seminar due to public holidays)
- Each topic comes with a supervisor:
  - You can ask questions by e-mail at any time
    - about your topic, the papers, your presentation and report
  - Up to one office hour per week

Presentations

- Order of presentation will be determined after topic assignment
  - Check with us if you need to swap slots, but find someone who is willing to swap **first**
- About 45 minutes long:
  - Introduction (about 10 minutes):
    - summary of previous week
    - finding themes that join the two weeks
  - Technical content (about 35 minutes):
    - detailled and comprehensive problem statement
    - presentation of the two papers
    - describe the common links between the papers
- Public feedback from other students after discussion

Suggested presentation preparation

- Schedule two meetings with your supervisor:
  - First meeting: 2–3 weeks before presentation
    - Read the papers for this meeting
    - Ask questions if you have difficulties
    - Discuss your plans for presentation
  - Second meeting: 1 week before presentation
    - Prepare a preliminary presentation (not a full rehearsal)
    - We can provide feedback
  - It is your responsibility to arrange the meetings
  - Do not rely on us providing last-minute feedback

Discussion (30-45 minutes)

- Day before the seminar:
  - **Everyone** submits 2+ questions for discussion to sa-seminar@lists.mpi-inf.mpg.de Important: your contribution will be **marked**
- At the seminar:
  - One person (“moderator”) is chosen at random (at beginning of each seminar session) and leads the discussion
  - We provide the submitted questions just before the seminar the moderator
  - Responsibilities of moderator
    - give summary of the talk
    - moderate and guide discussion
    - raise open questions that remain
    - stimulate discussion of the strengths and weaknesses of the two papers
    - this will also be **marked**

Report

- 5–7 page summary of the major ideas in your topic:
  - 3–4 pages on the two papers
  - 3–4 additional paper references
  - 2–3 pages with your own ideas, for example:
    - Limitations not mentioned in the paper + sketch of potential solution
    - Try to suggest improvements
    - Novel ideas based on content described in the papers
    - Can be the result of the discussion after your presentation
- The idea is that you get a feeling for your specific topic surpassing the level of simply understanding a paper.

Report

- Due date: **Friday, 16 August 2017**

  (4 weeks after the last seminar)
- Send PDF by e-mail
- We will provide a LaTeX template on the seminar website
  - If you use other software, make it look like the LaTeX template
    - this is your responsibility
  - Strongly recommended to learn LaTeX
    - used by nearly all research papers in visual computing

Grading scheme

- Presentation (overall: 30%)
  - Form (30%): time, speed, structure of slides
  - Content (50%): structure, story line and connections, main points, clarity
  - Questions (20%): answers to questions
- Discussion (overall: 30%)
  - Submitted questions (33%): insight, depth
  - Participation (33%): willingness, debate, ideas
  - Moderation (33%): strengths and weaknesses, integration of questions
- Report (overall: 40%)
  - Form (10%): diligence, structure, appropriate length
  - Context (20%): the big picture, topic in context
  - Technical correctness (30%)
  - Discussion (40%): novelty, transfer, own ideas / in own words

Benefits to you

- Practise important skills in research
  - Read and understand technical papers
  - Present scientific results and convince other people
  - Analyse and develop new ideas through discussions
- Discussion is essential:
  - If you don’t participate, you miss a big chance
  - Most ideas are developed in discussions about other papers
- Therefore:
  - Prepare for the seminar classes
  - Participate actively in the discussions
  - Benefit from the interaction in the group

What this seminar is NOT …

- A course to just sit and listen
  - Come prepared
  - Read all papers before class, think about problems, submit questions and discuss them in class
  - Your participation is beneficial for everyone
    - the group makes the seminar
- "Cheap" credit points
  - Don't underestimate the time it takes to understand a paper, prepare a talk, and write a report
  - So please do take it seriously!

Schedule

- 10 April – Introduction ◀ You are here
- 19 April – Lectures:
  - “How to read an academic paper”
  - “How to give a good talk”
- 03 May – First presentation by a student
- … 8 more weekly presentations …
- 19 July – Last presentation by a student
- 16 August – Report deadline

Overview

- Organization
- Introduction
- Topics
- Summary

What is ”shape”?

- an attribute relevant for image understanding (besides colour, texture, position, size ...)

Copyright by Seth Casteel

What is ”shape”?

- **Shape representation**: labelled point-cloud
- **Definition of shape:** “what is left when differences which can be attributed to translations, rotations, and dilations have been quotiened out” (Kendall, 1984)
- Examples:

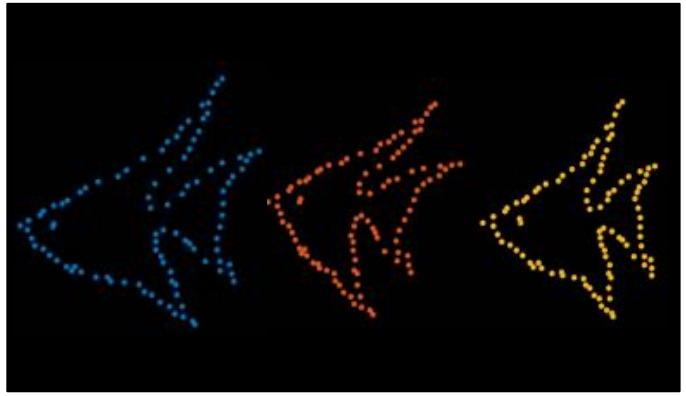

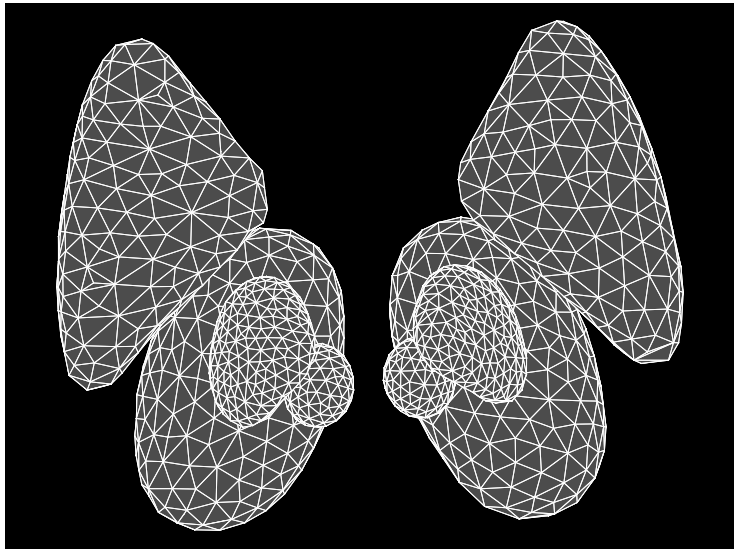

Shape Analysis

- **Objective:** recognise, classify or quantify morphological patterns or regularities of some objects of interest
- **Motivation:**
  - develop theories of underlying concepts (understand)
  - build and use models of shape (e.g. for recognition, reconstruction, tracking, ...)
- **Challenges:**
  - generally, shapes are continuous surfaces
    - manifolds, non-Euclidean
    - usually they are discretised to perform computations
  - the correspondence problem is extremely challenging
    - e.g. the related quadratic assignment problem is NP-hard

Shape Analysis

- **Fields related to shape analysis**
  - computer science & applied mathematics
  - (medical) image analysis & image processing
  - computer vision & computer graphics
  - augmented & virtual reality
  - geometry processing
  - medicine, biology, archaeology, ...

Aspects of shape analysis

- modelling shape deformation

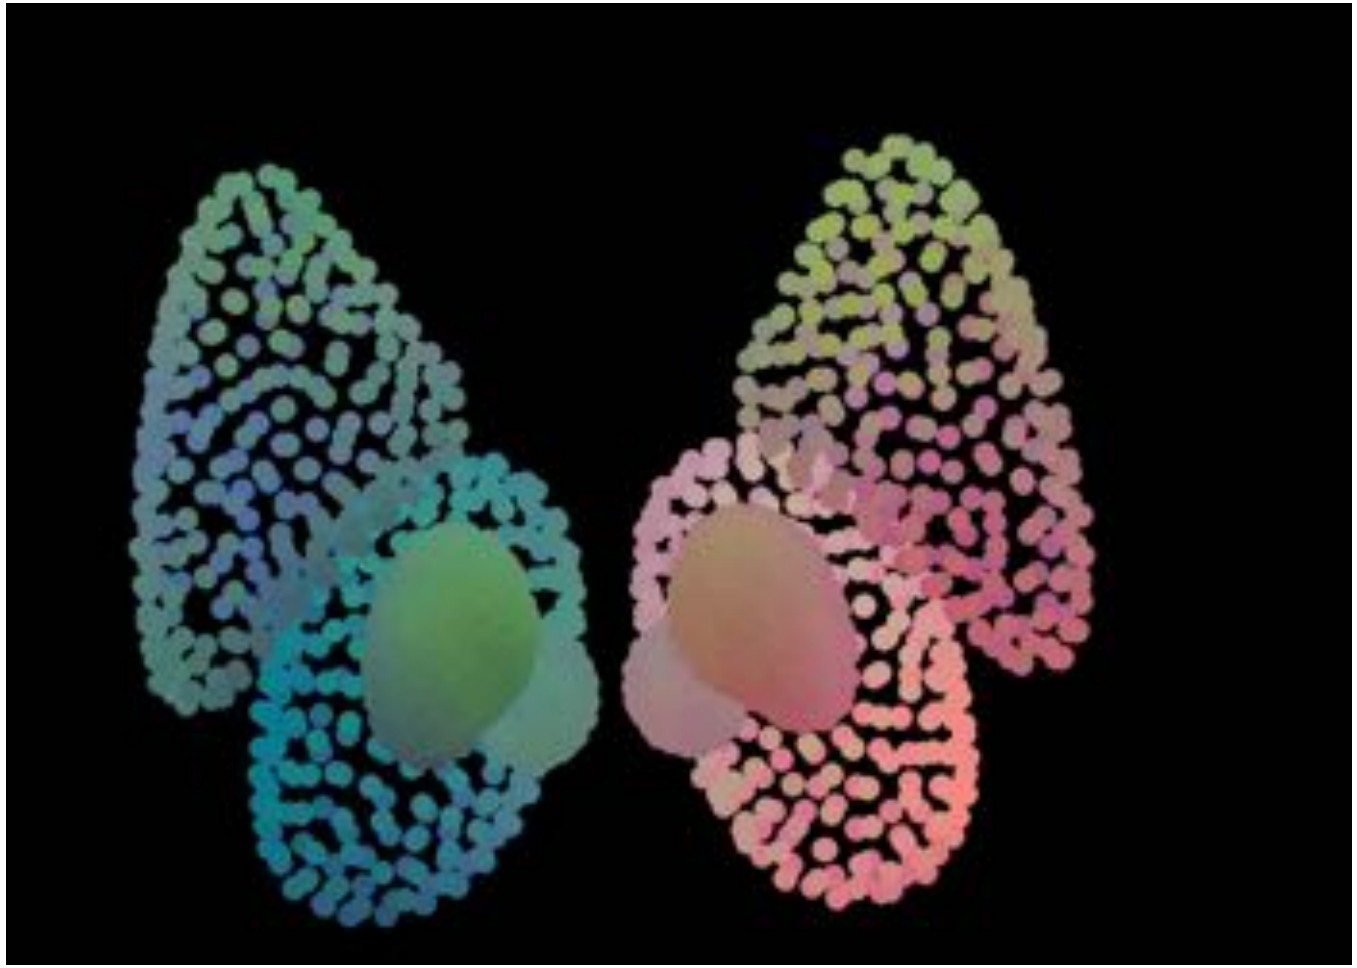

Aspects of shape analysis

- correspondence problem

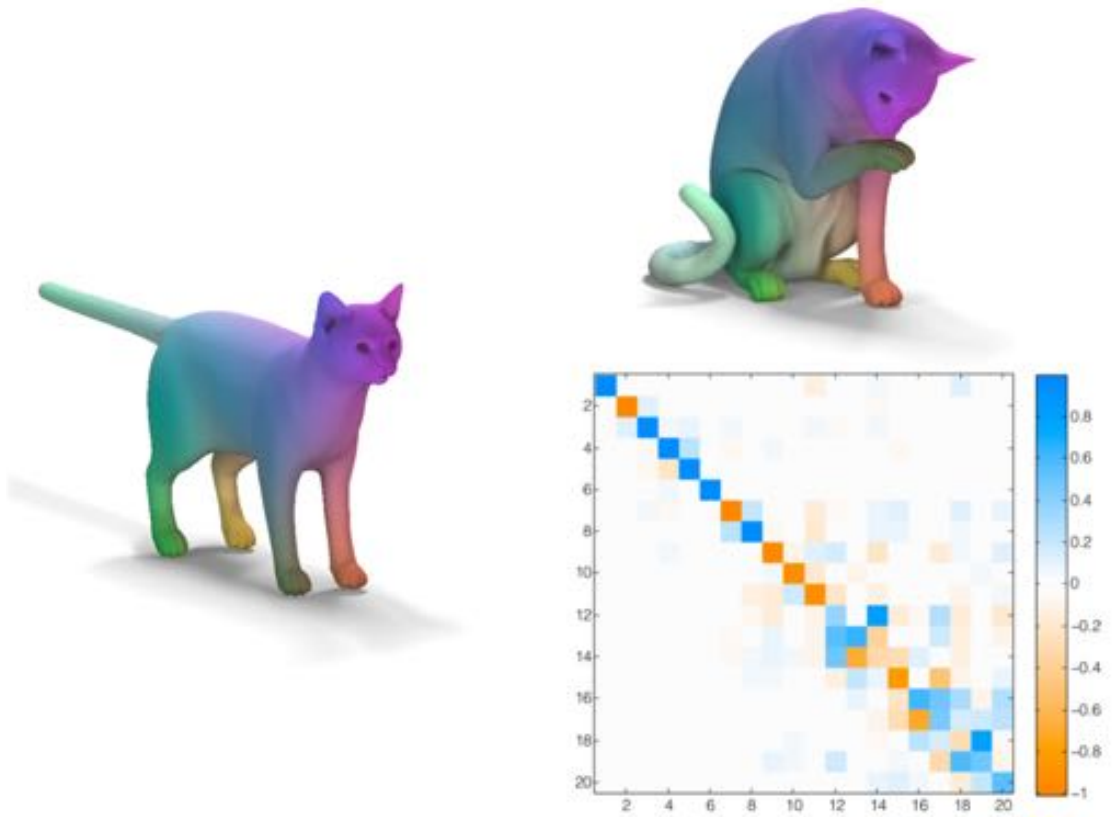

Ovsjanikov et al., SIGGRAPH 2012

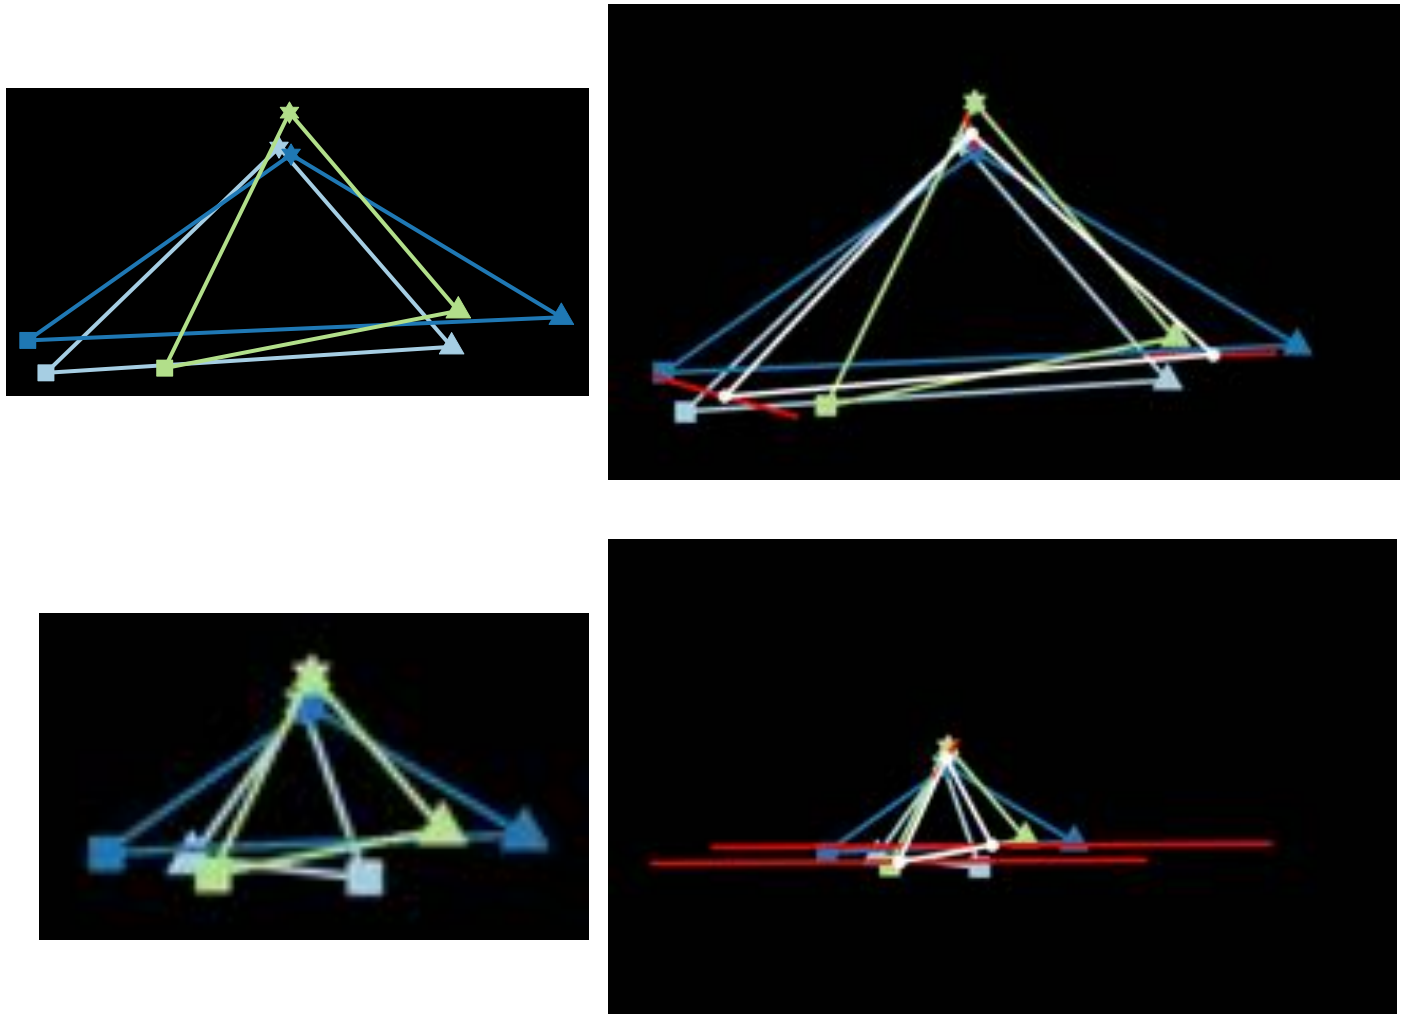

Aspects of shape analysis

- shape matching

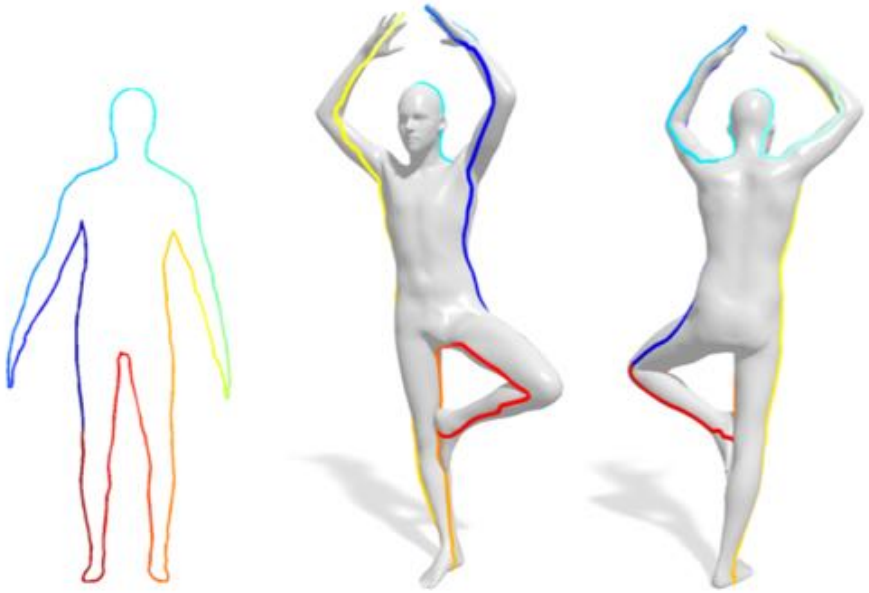

2D shape to 3D shape
(Laehner et al., CVPR 2016)

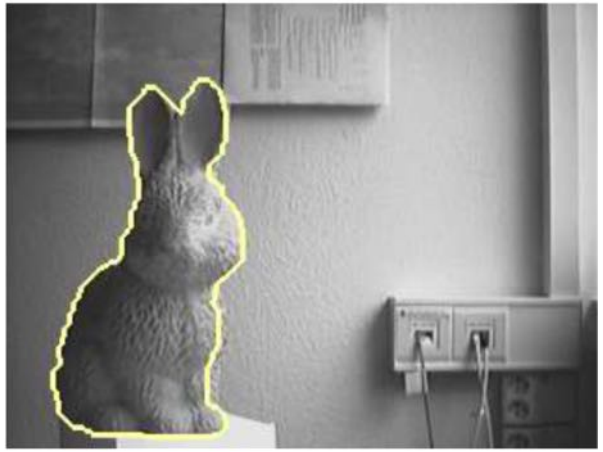

2D shape to 2D image
(Schoenemann & Cremers,TPAMI, 2010)

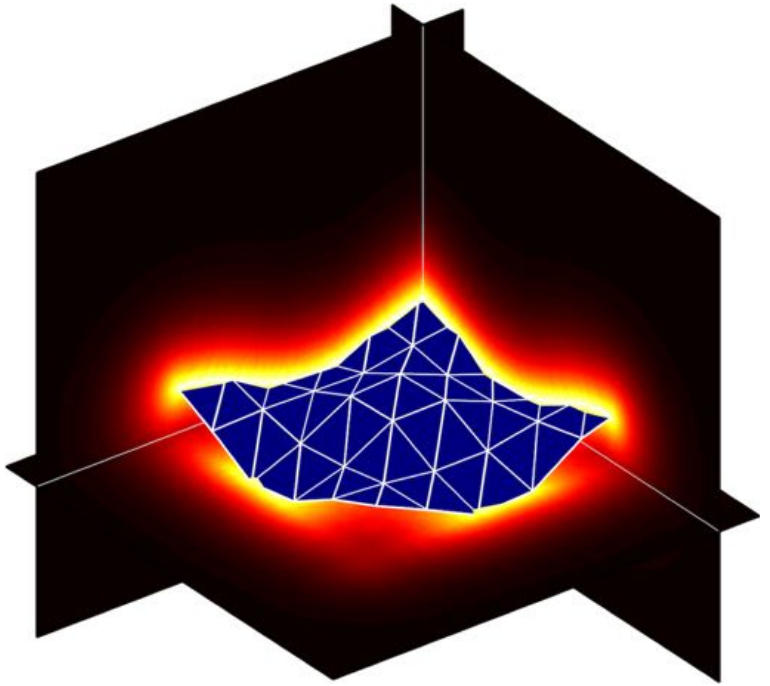

3D shape to 3D image
(Bernard et al., CVPR 2017)

Aspects of shape analysis

- shape fitting

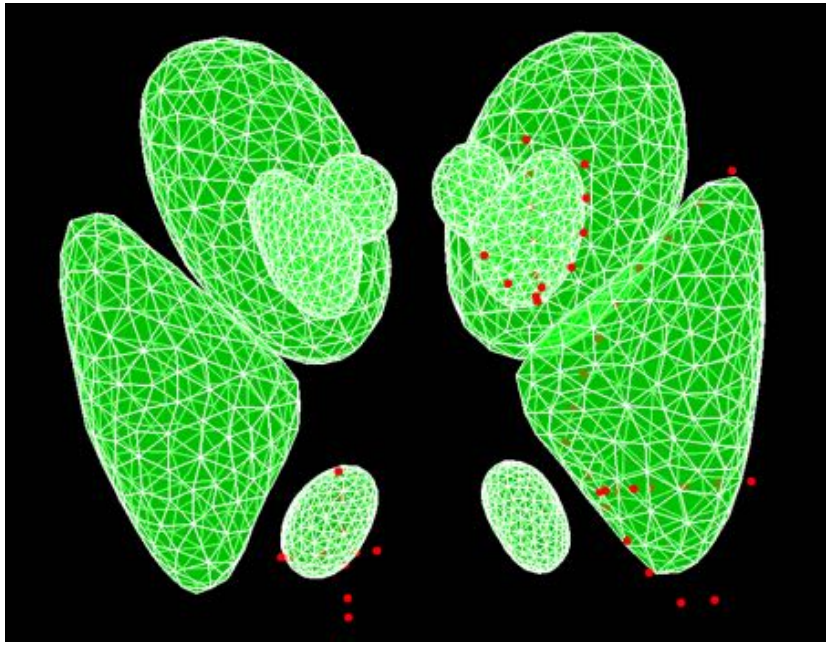

Fit shape to points
(Bernard et al., MedIA, 2017)

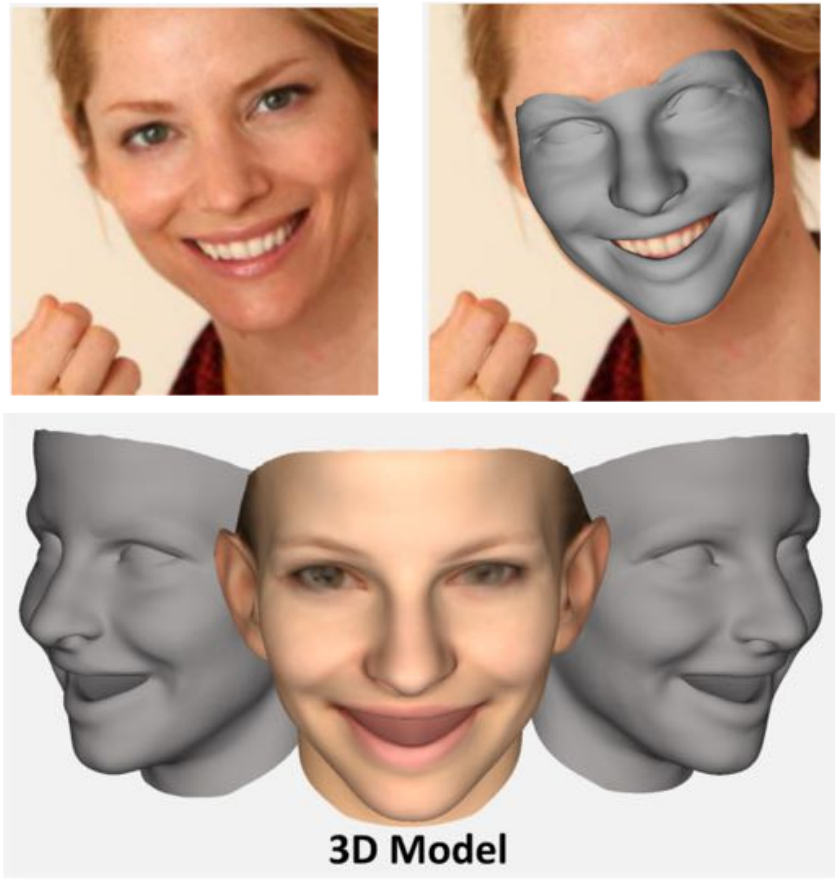

Fit shape model to image
(Tewari et al., ICCV, 2017)

Aspects of shape analysis

- reconstruction by fitting

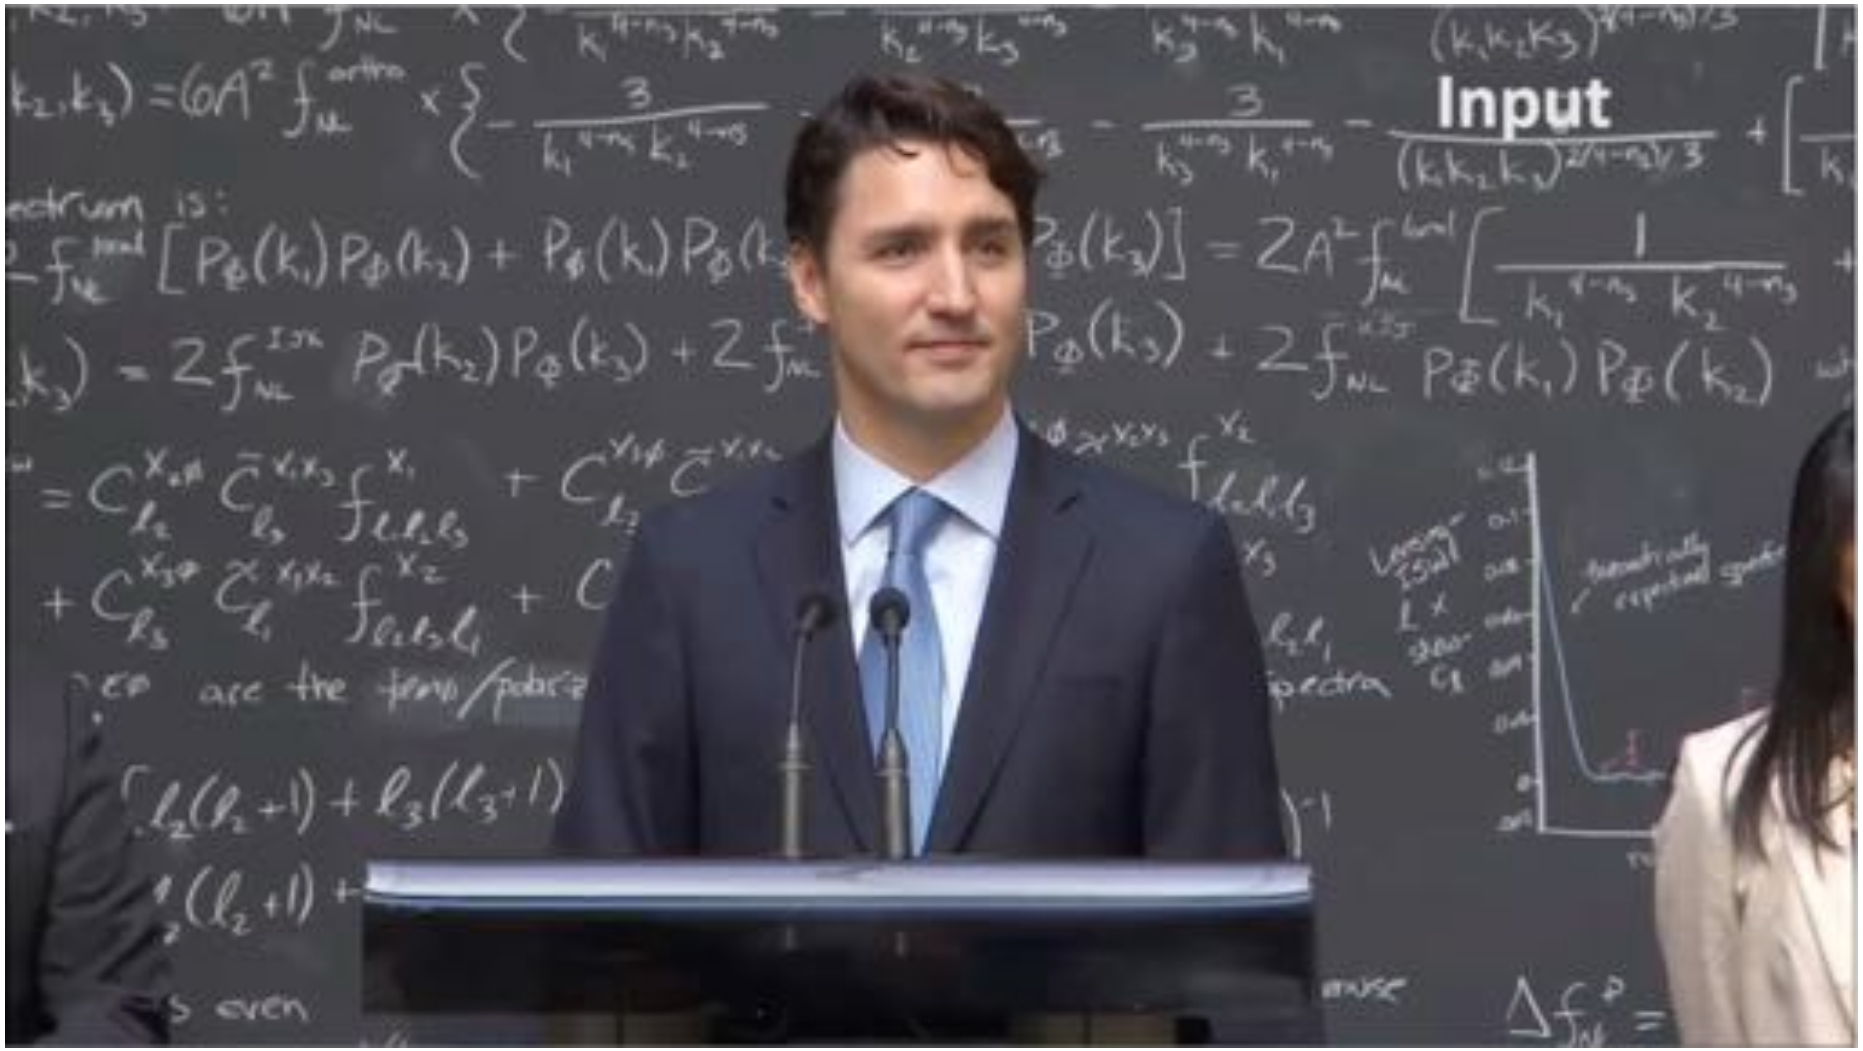

Fit shape model to image (Tewari et al., ICCV, 2017)

Aspects of shape analysis

- tracking by fitting

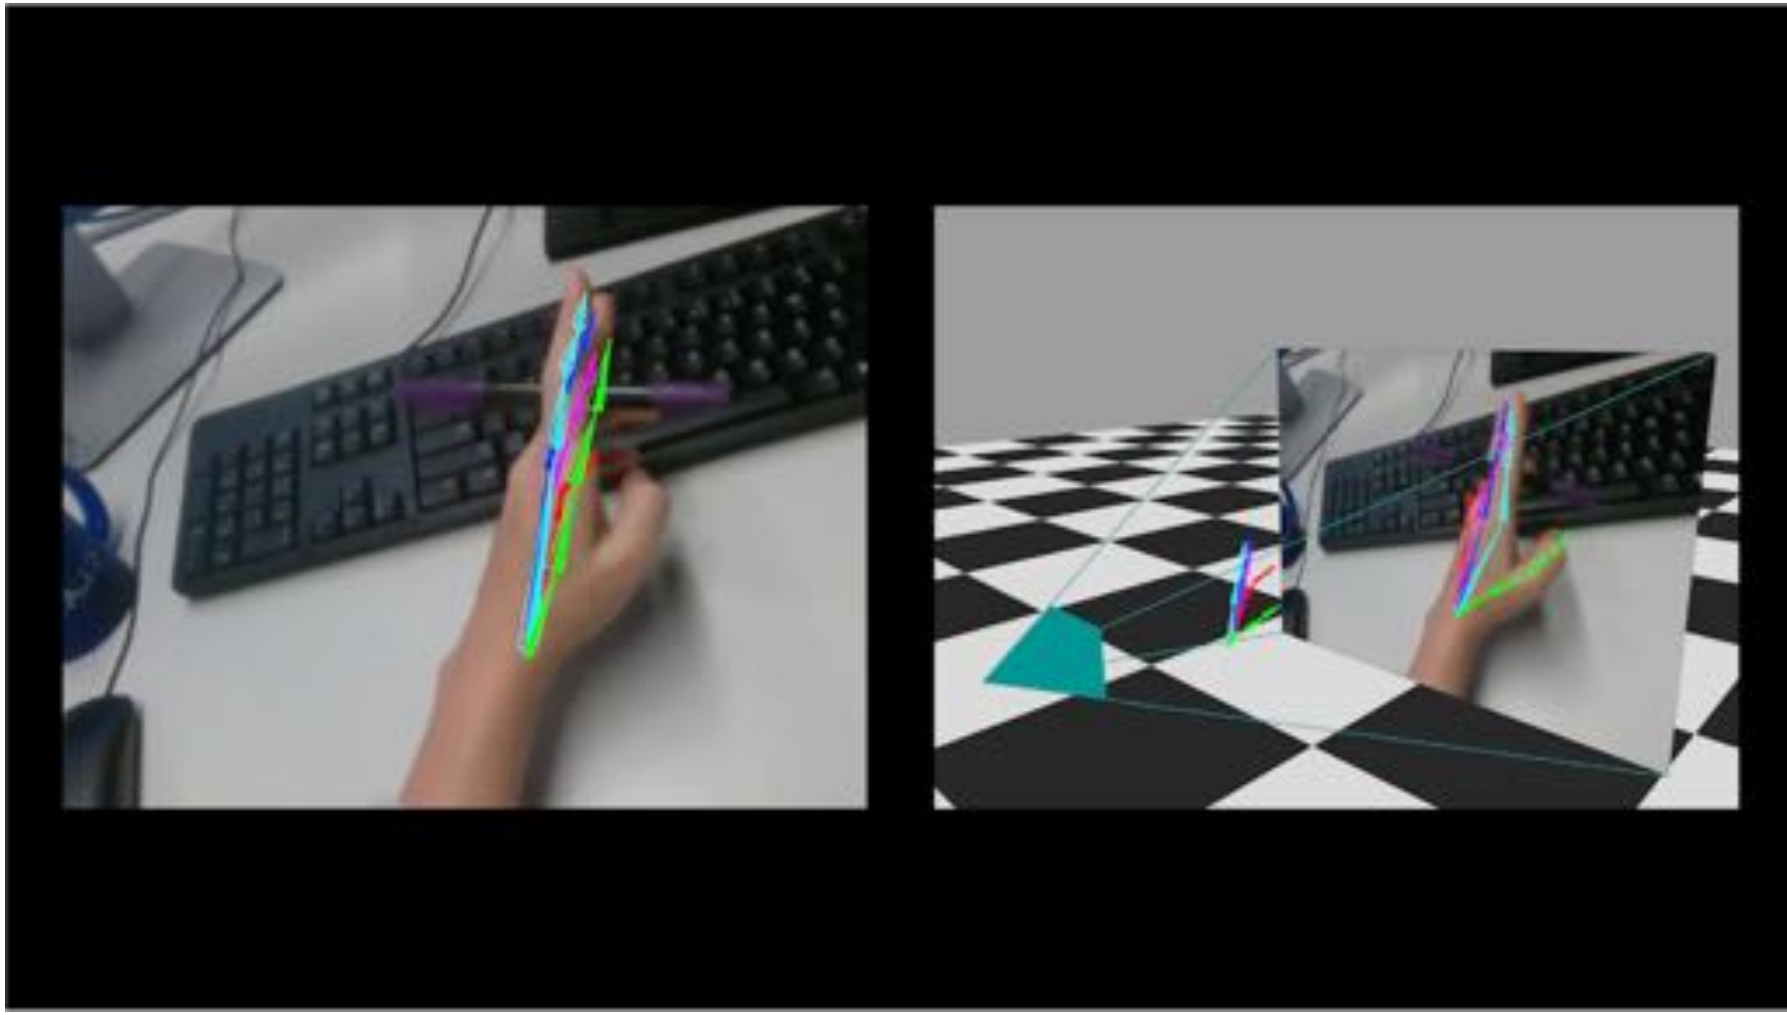

Fit hand skeleton model to image (Mueller et al., CVPR, 2018)

Aspects of shape analysis

- shape representation

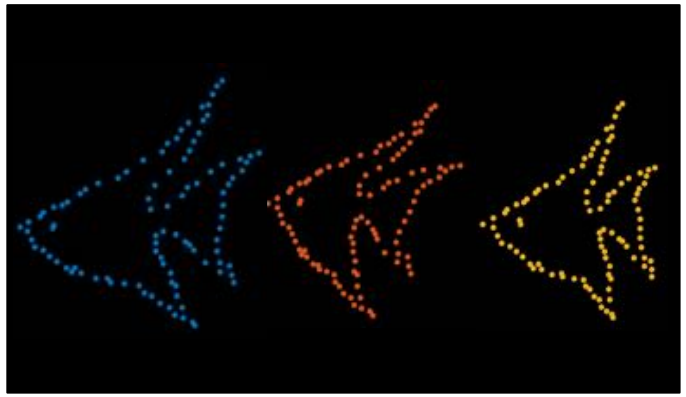

Point-cloud

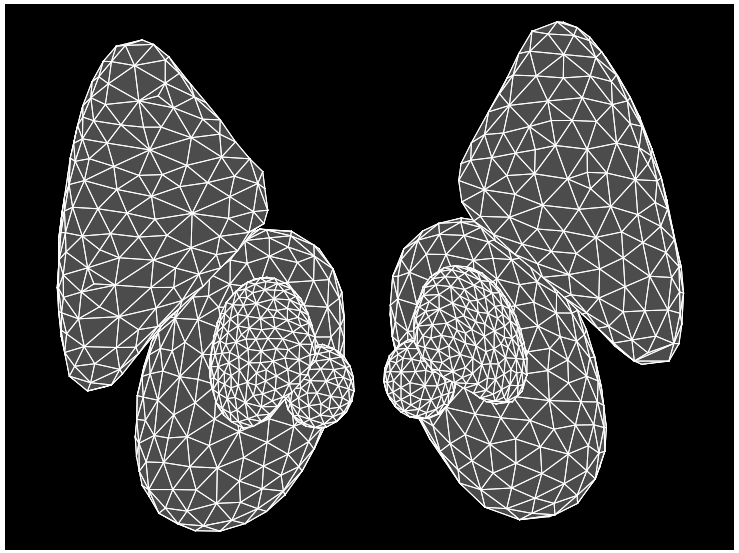

Triangular surface mesh

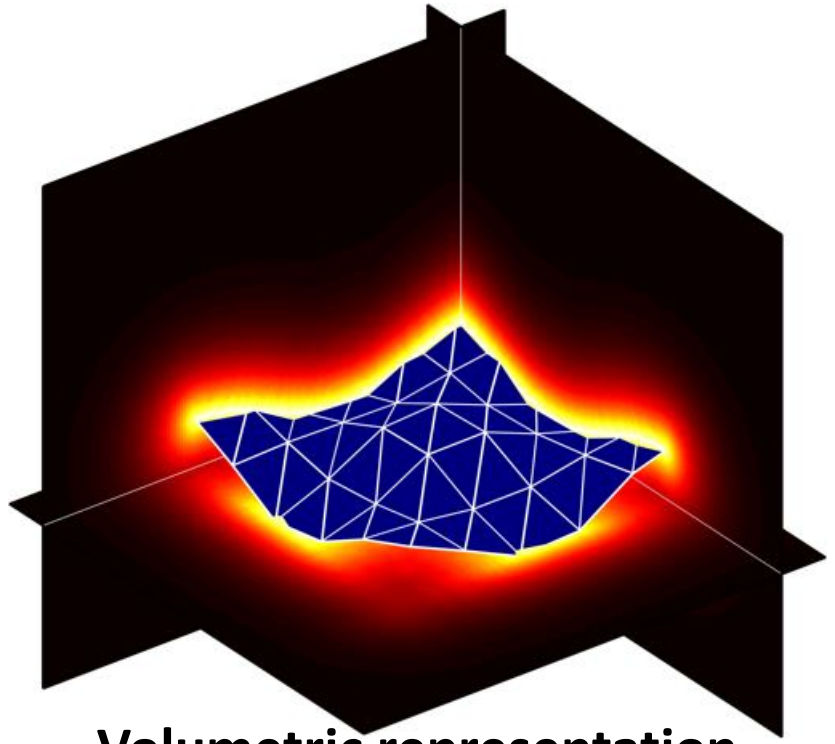

Volumetric representation

Aspects of shape analysis

- multi-shape analysis

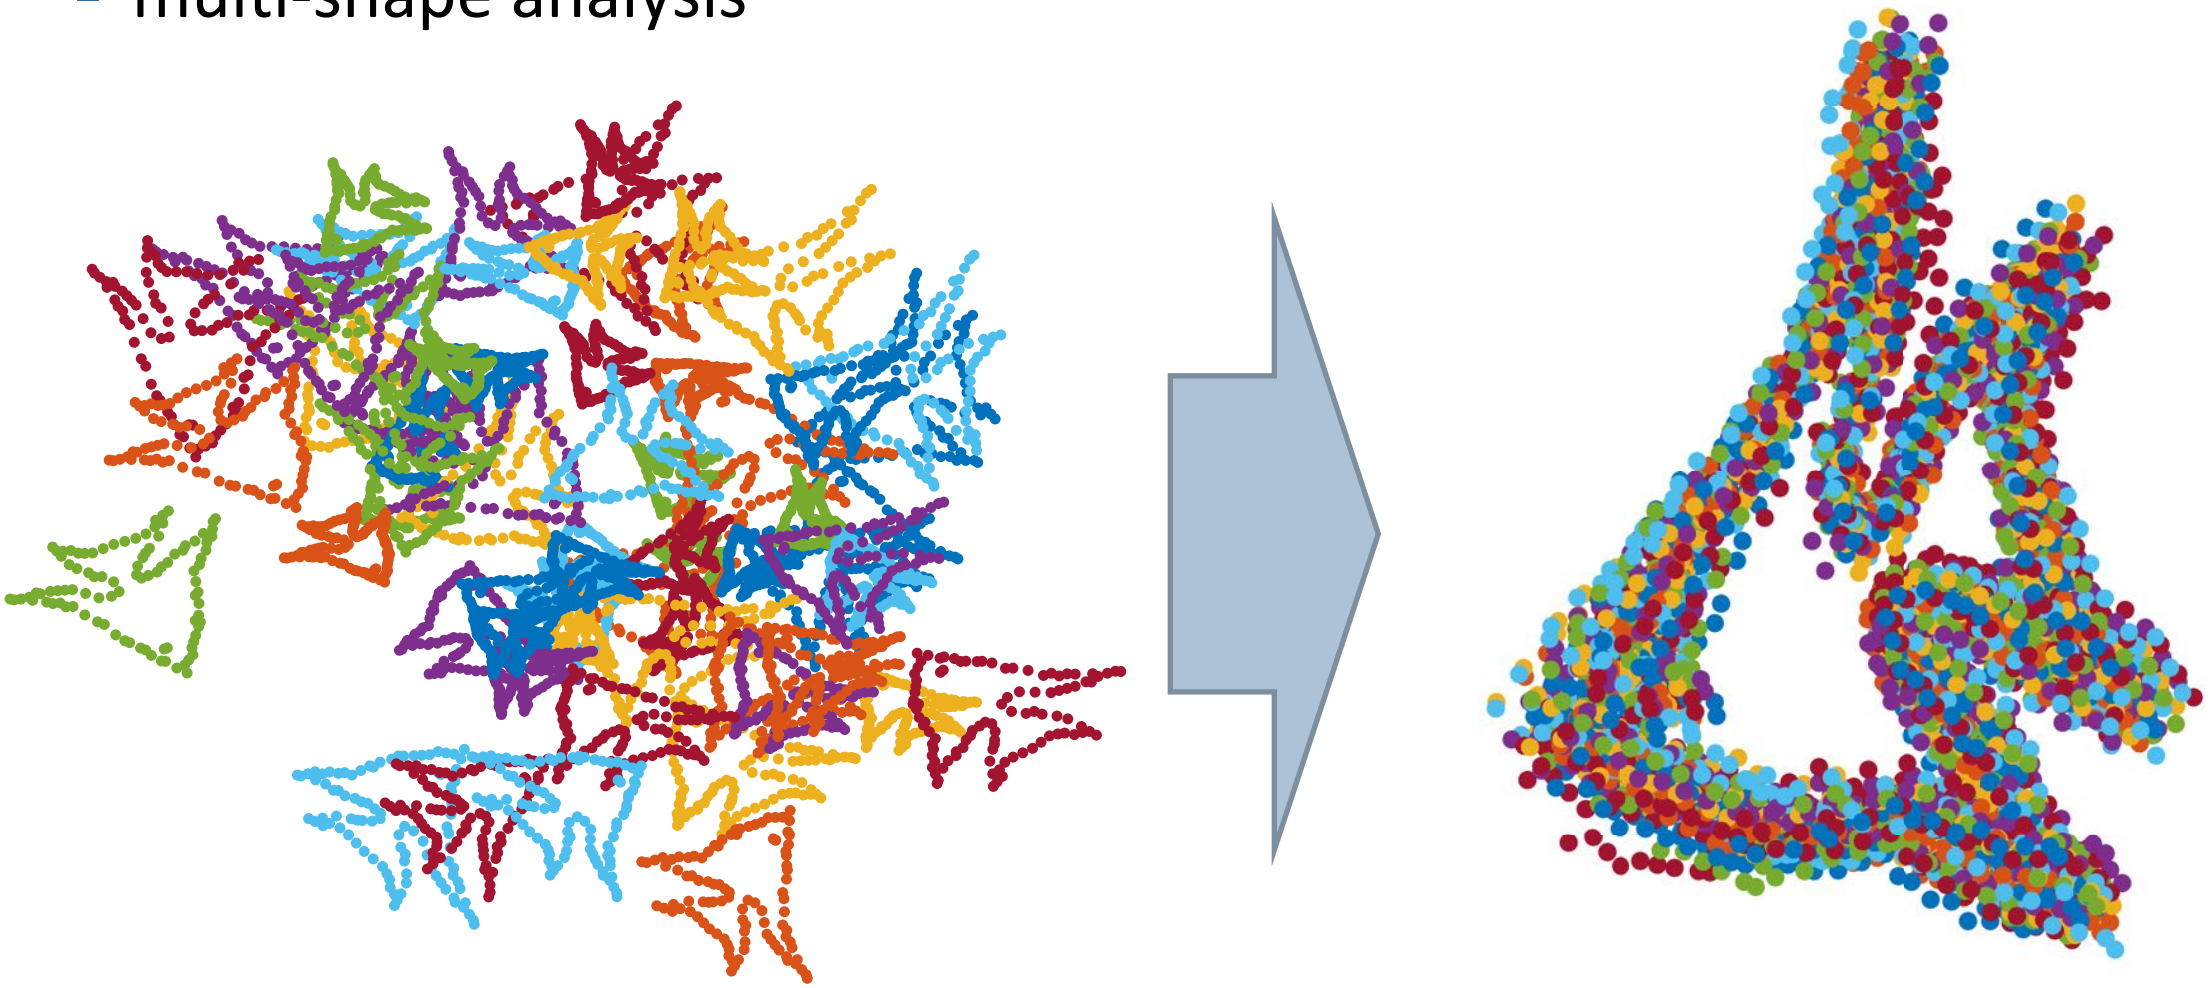

Toolbox for shape analysis

Linear algebra:

- the very basics (matrix multiplication, inverse, vector space, (linear) subspace, basis, orthogonality, trace, ...)
- eigenvalues/eigenvectors: $Ax = sx$
  - relation to Rayleigh quotient $R(A,x) = x'Ax/x'x$
  - spectral decomposition: $A = QSQ^{-1}$
- singular value decomposition: $A = USV'$
  - principal components analysis (eigendecomposition of covariance matrix)
- solving linear systems of equations: $Ax = b$
- orthogonal matrices $O(d)$, the groups $SO(d)$, $SE(d)$

Toolbox for shape analysis

Discrete/continuous optimisation

- convexity of functions/sets
- understand why convexity is ”nice” in optimisation
- positive definiteness of matrices
- permutation matrices
- linear assignment problem

Basic concepts of geometry processing

- discrete representation of surfaces as triangular surface meshes
- normals

Differential geometry (in this seminar not crucial)

Toolbox for shape analysis

Homework for next week: Fill your knowledge gaps in these areas (if any)

Overview

- Organization
- Introduction
- Topics
- Summary

Offered Topics

- Covering a combination of well-established and state-of-the-art research papers
- Strong focus on top conferences and journals in vision, graphics and machine learning:
  - SIGGRAPH & SIGGRAPH Asia (Transactions on Graphics)
  - Transactions on Pattern Analysis and Machine Intelligence (PAMI)
  - IEEE Computer Vision and Pattern Recognition (CVPR)
  - International Conference on Computer Vision (ICCV)
  - European Conference on Computer Vision (ECCV)
  - Neural Information Processing Systems (NIPS)
  - Symposium on Geometry Processing (SGP)

Topic Overview

- Shape Models
- Shape Deformation
- Spectral Graph Matching
- Graph Matching (Path-following)
- Semidefinite Relaxations for Matching Problems
- Lifting-free Convex Relaxations of Matching Problems
- Convex Point Registration
- Spectral Multi-Shape Matching and Alignment
- Combinatorial Shape-to-Image Matching
- Geometric Deep Learning
- Deep Learning on Point Clouds
- Functional Maps
- Volumetric Shape Representation

TOPIC: Shape Models
(Marc Habermann)

TOPIC: Shape Models

Active Shape Models - 'Smart Snakes'

Timothy F. Cootes, Christopher J. Taylor
BMVC 1992

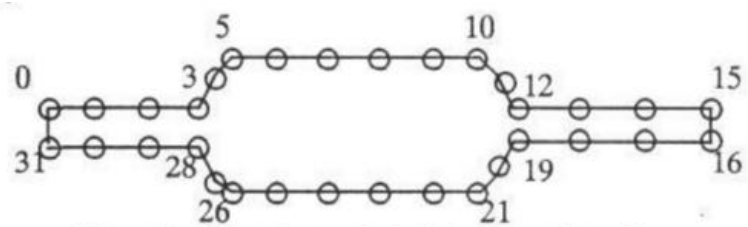

Figure 1 : 32 point model of the boundary of a resistor.

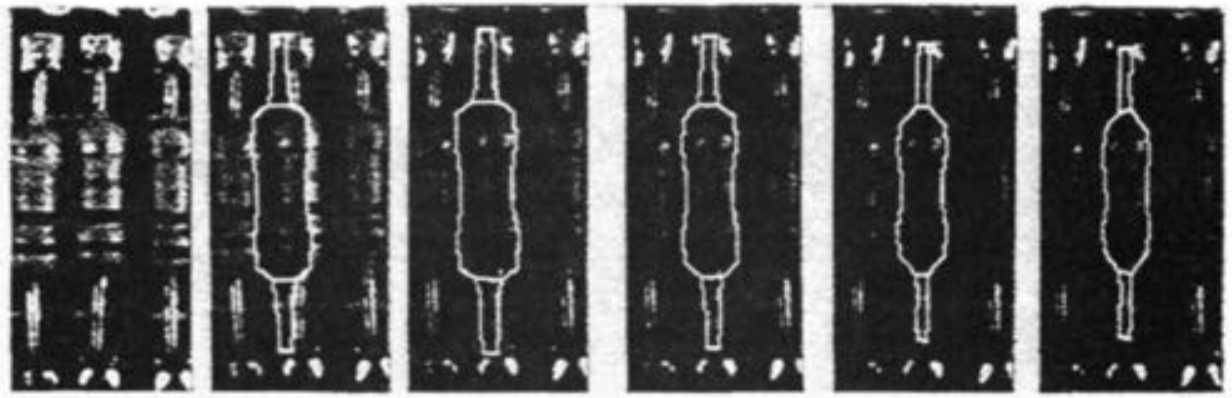

(a) Original Image *(b) Initial Position* *(c) After 30 iterations* *(d) After 60 iterations* *(e) After 90 iterations* *(f) After 120 iterations*

TOPIC: Shape Models

Lie Bodies: A Manifold Representation of 3D Human Shape

Oren Freifeld, Michael J. Black

ECCV 2012

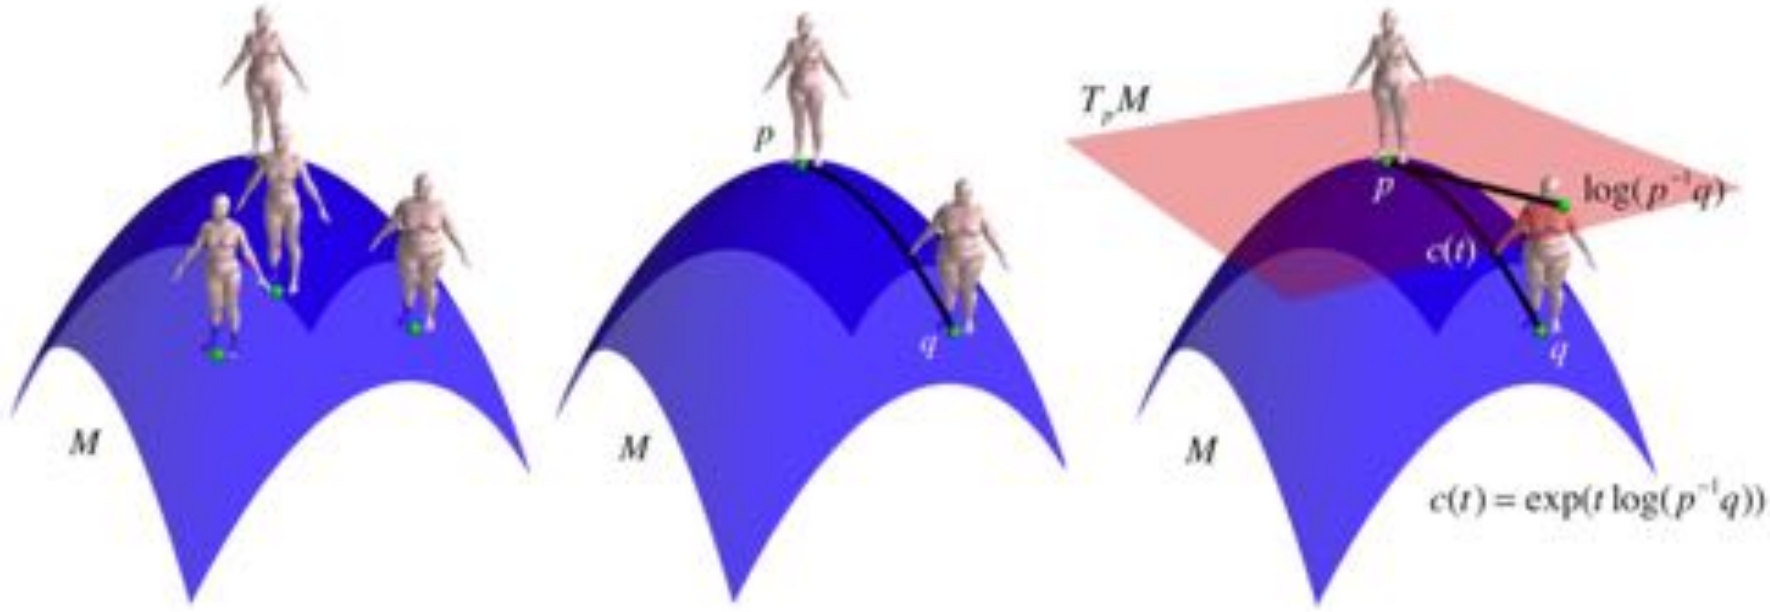

TOPIC: Shape Deformation
(Ayush Tewari)

TOPIC: Shape Deformation

Deformation Transfer for Triangle Meshes
Robert W. Sumner, Jovan Popovic
SIGGRAPH 2004
As-Rigid-As-Possible Surface Modeling
Olga Sorkine, Marc Alexa
SGP 2007

TOPIC: Shape Deformation

As-Rigid-As-Possible Surface Modeling:

TOPIC: Shape Deformation

Deformation Transfer for Triangle Meshes:

TOPIC: Shape Deformation

Deformation Transfer for Triangle Meshes:

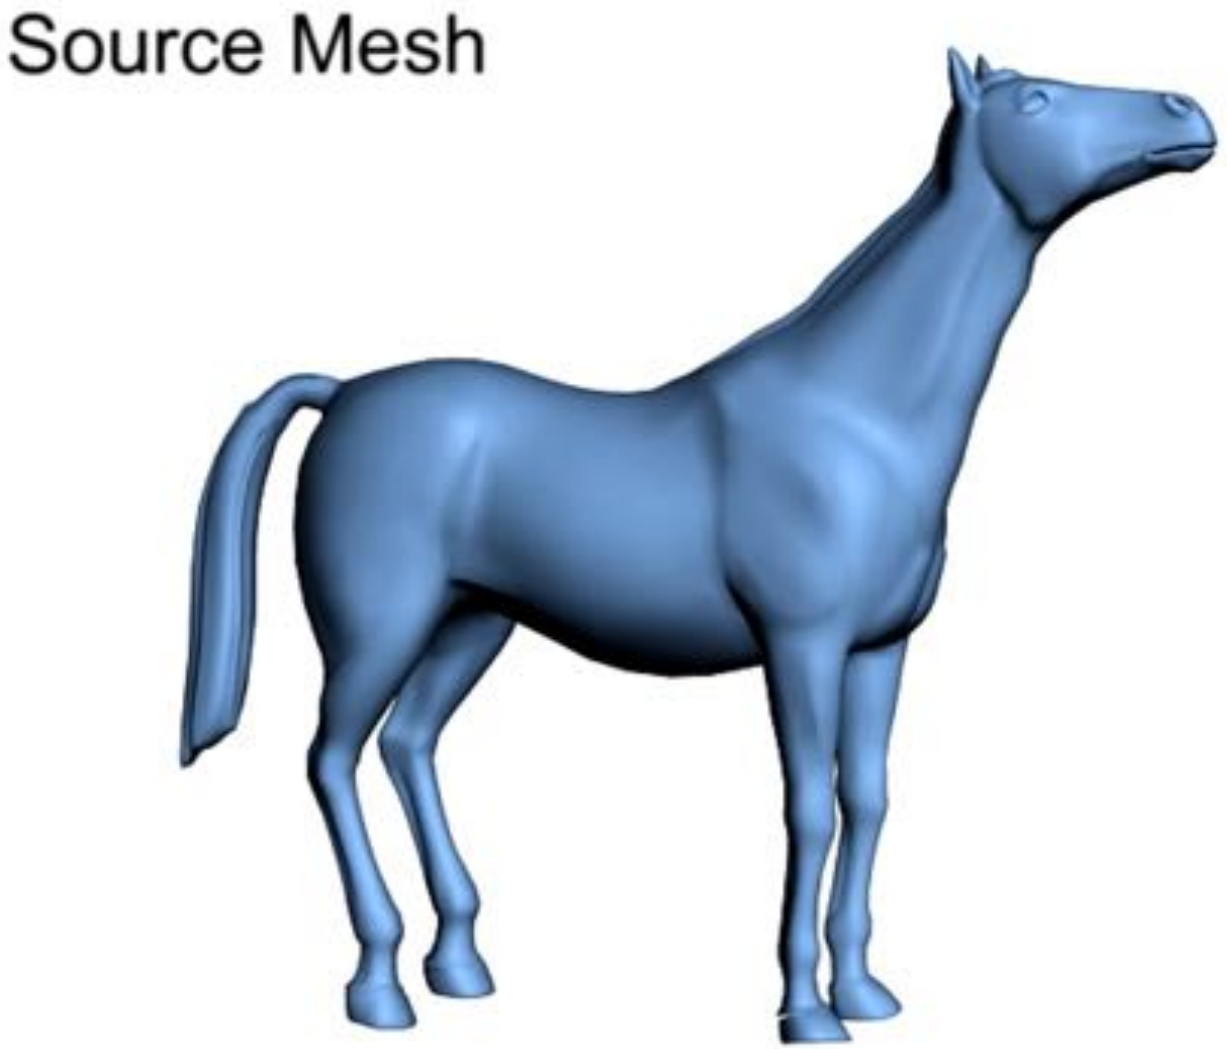

TOPIC: Spectral Graph Matching
(Hyeongwoo Kim)

TOPIC: Spectral Graph Matching

A Spectral Technique for Correspondence Problems Using Pairwise Constraints

Marius Leordeanu, Martial Hebert

ICCV 2005

Balanced Graph Matching

Timothee Cour, Praveen Srinivasan Jianbo Shi

NIPS 2006

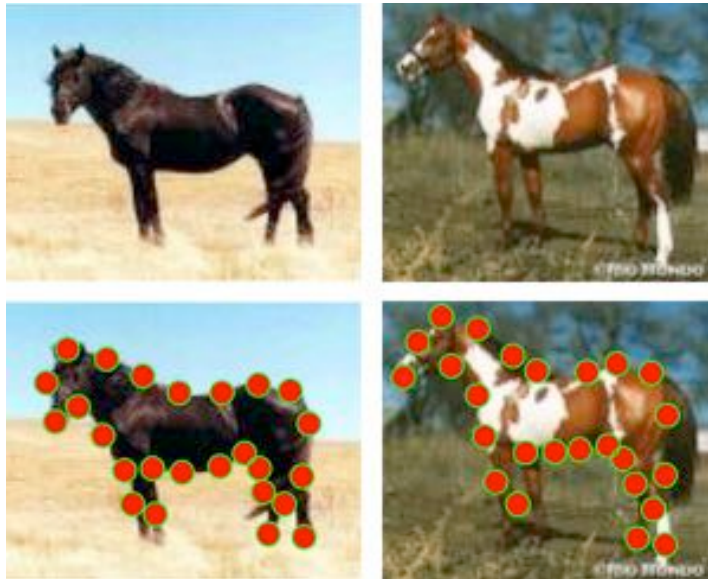

$$S = \sum_{a,b \in C} M(a,b) = x^T M x$$

$$\max \quad \frac{x^\mathsf{T} W x}{x^\mathsf{T} x} \quad \text{s.t.} \quad Cx = b$$

TOPIC: Graph Matching (Path-following)
(Florian Bernard)

TOPIC: Graph Matching (Path-following)

Factorized Graph Matching
Feng Zhou, Fernando De la Torre
TPAMI 2016

Binary Constraint Preserving Graph Matching
Bo Jiang, Jin Tang, Chris Ding, Bin Luo
CVPR 2017

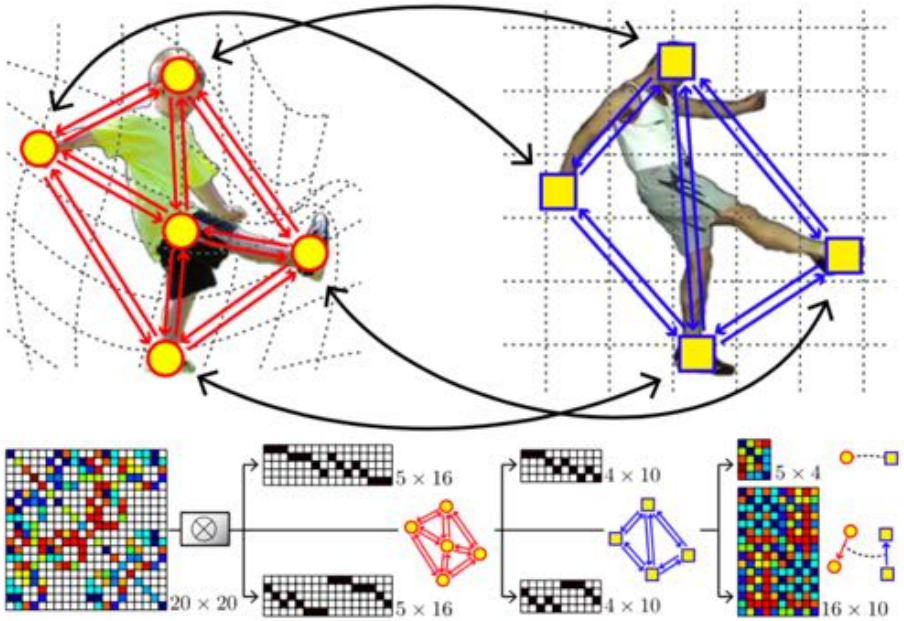

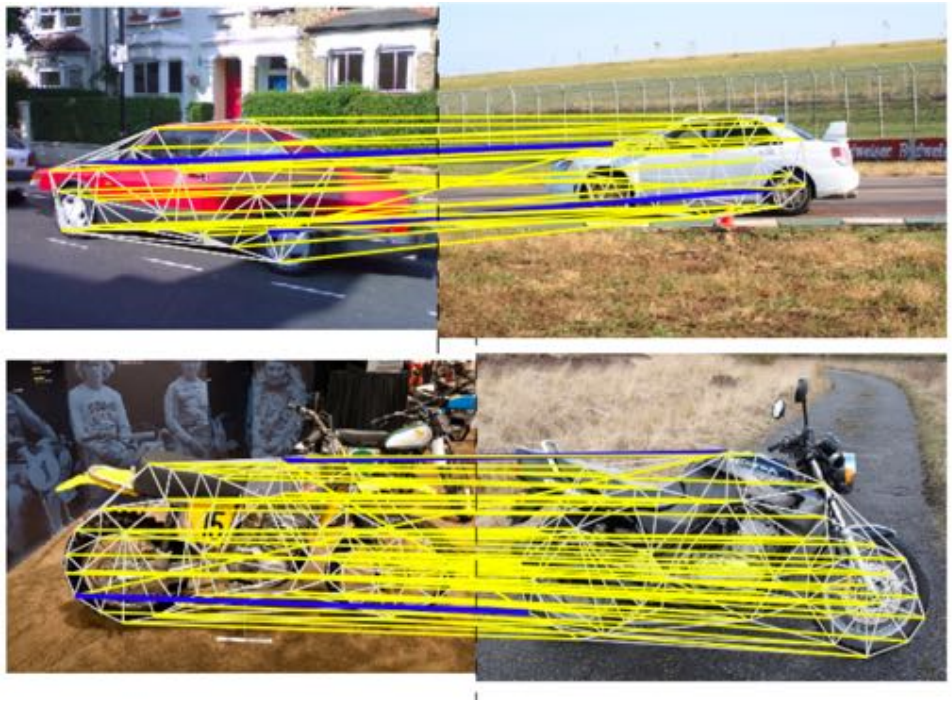

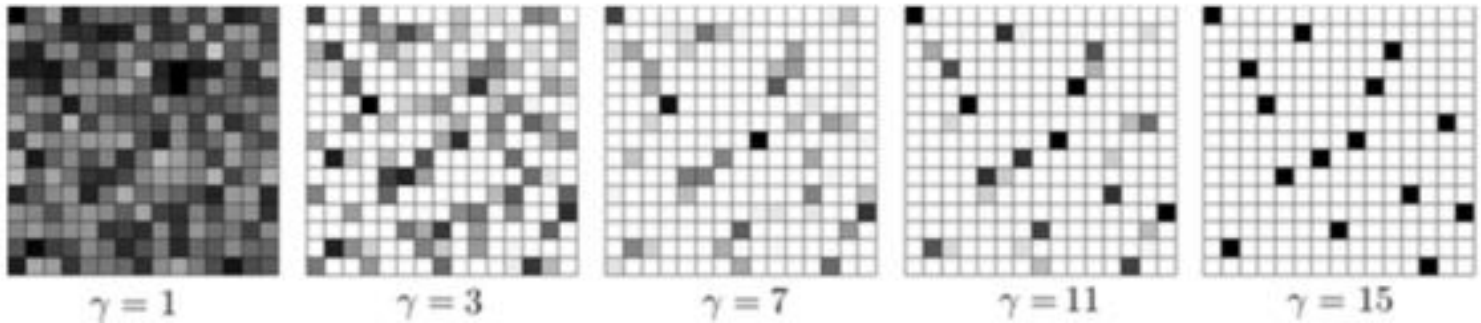

TOPIC: Semidefinite Relaxations for Matching Problems
(Florian Bernard)

TOPIC: Semidefinite Relaxations for Matching Problems

A Fast Semidefinite Approach to Solving Binary Quadratic Problems

Peng Wang, Chunhua Shen, Anton van den Hengel

CVPR 2013

Biconvex Relaxation for Semidefinite Programming in Computer Vision

Sohil Shah, Abhay Kumar, Carlos Castillo, David Jacobs, Christoph Studer, Tom Goldstein

ECCV 2016

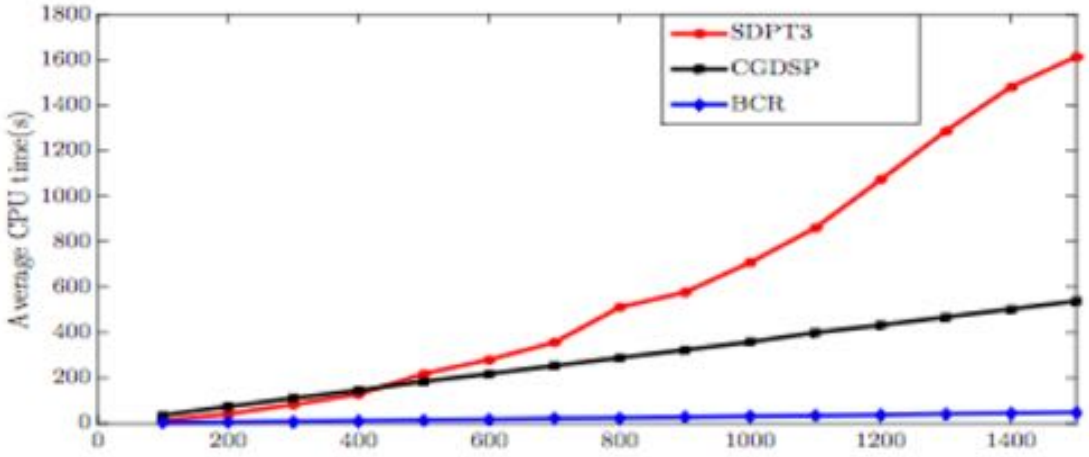

TOPIC: Lifting-free Convex Relaxations of Matching Problems
(Florian Bernard)

TOPIC: Lifting-free Convex Relaxations of Matching Problems

DS++: a flexible, scalable and provably tight relaxation for matching problems
Nadav Dym, Haggai Maron, Yaron Lipman
SIGGRAPH Asia 2017

DS*: Tighter Lifting-Free Convex Relaxations for Quadratic Matching Problems
Florian Bernard, Christian Theobalt, Michael Moeller
CVPR 2018

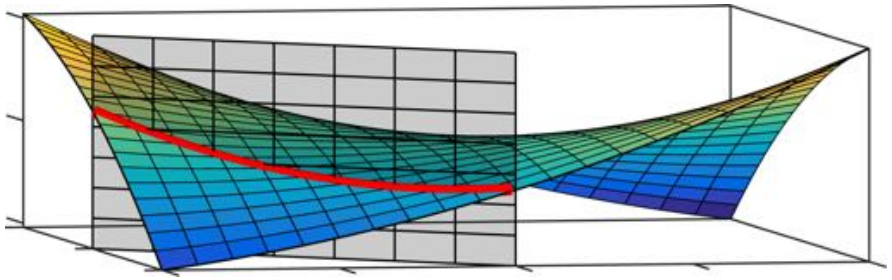

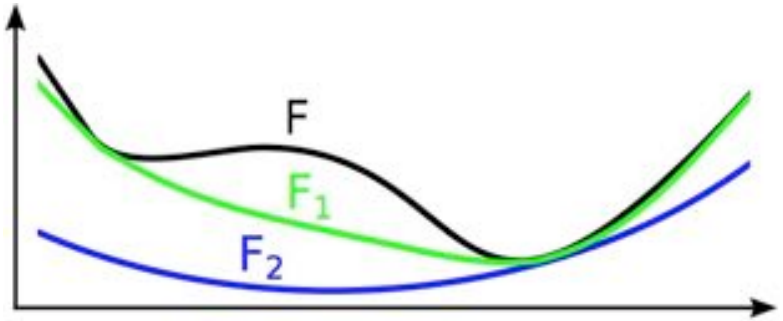

TOPIC: Convex Point Registration
(Jiayi Wang)

TOPIC: Convex Point Registration

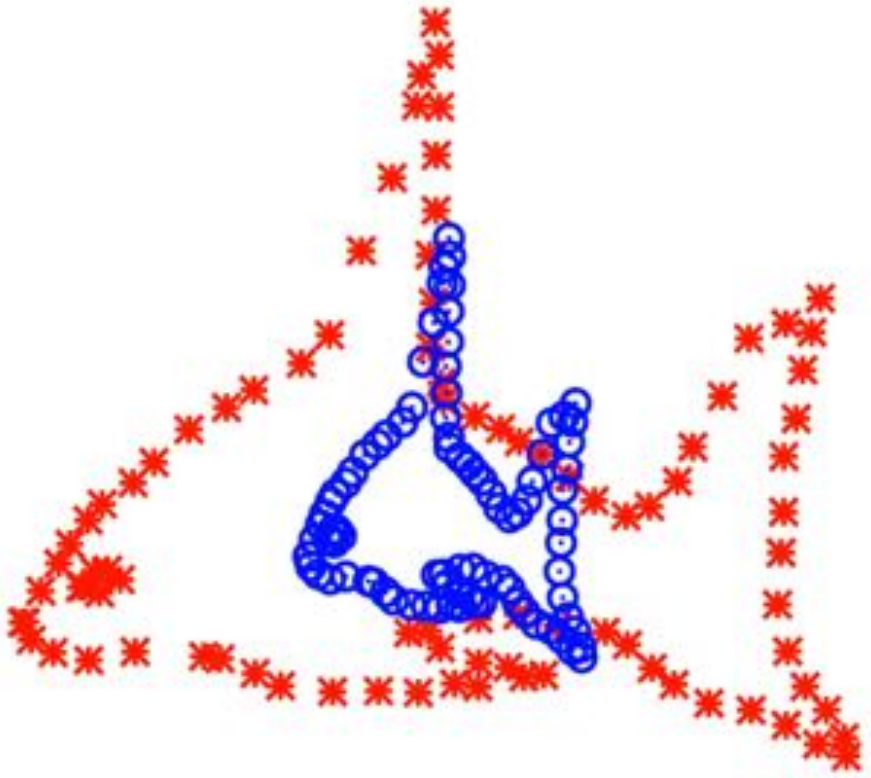

TOPIC: Convex Point Registration

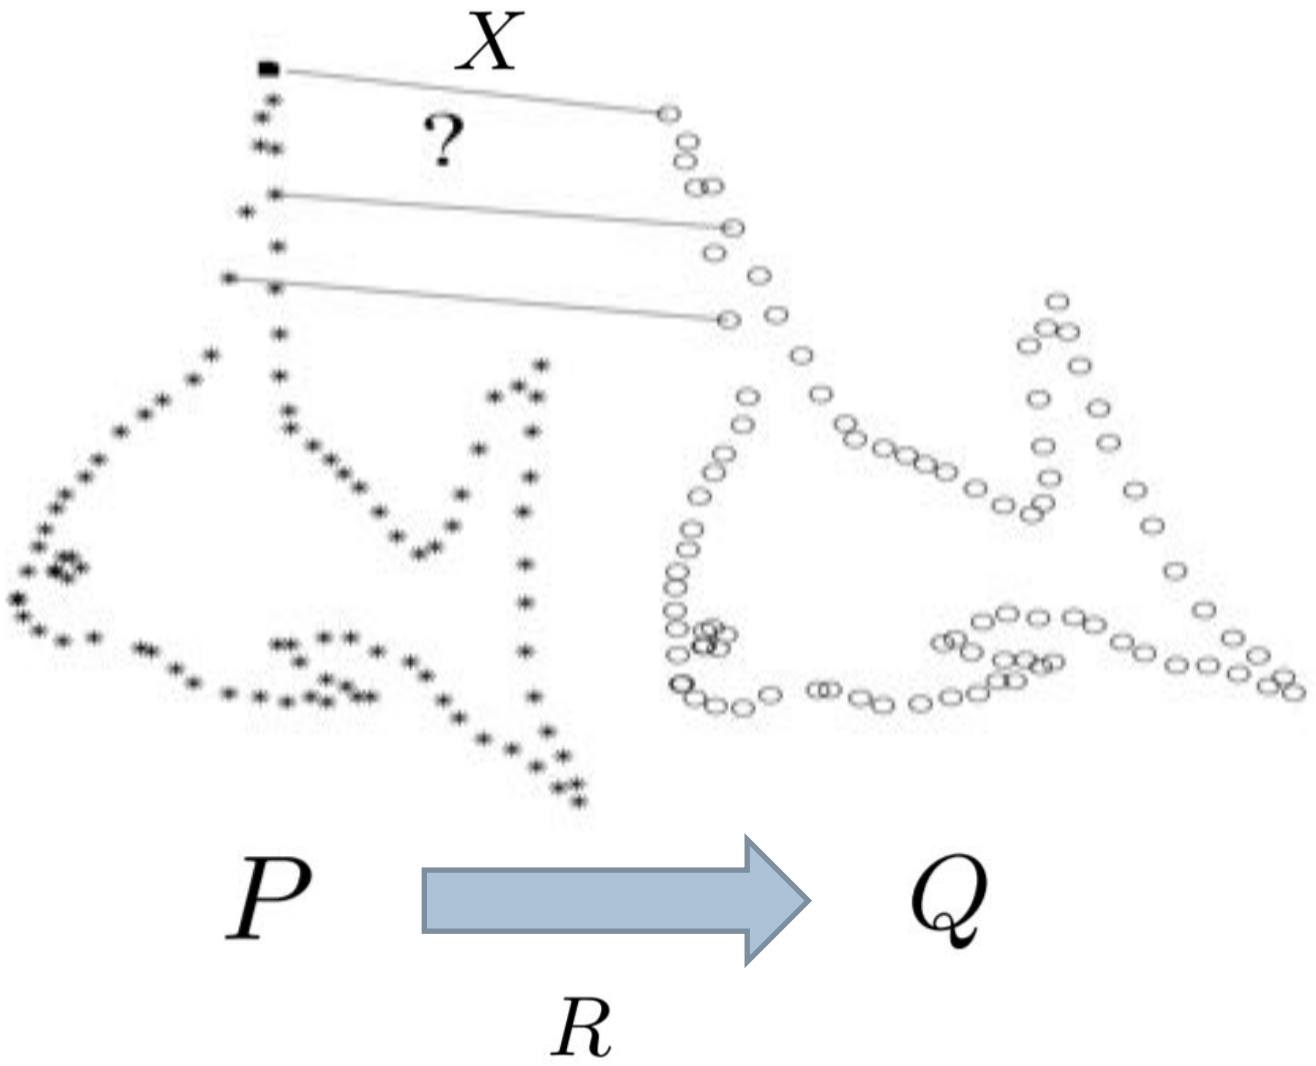

P, Q - Two Point Sets

X - Correspondence

R - Transformation

$$d(P,Q) = \min_{X,R} \|RP - QX\|_F^2$$

$$\text{s.t.} \quad X \in \Pi_n$$

$$R \in \mathcal{O}(d)$$

TOPIC: Convex Point Registration

$$d(P,Q) = \min_{X,R} \|RP - QX\|_F^2$$

$$\text{s.t.} \quad X \in \Pi_n$$
$$R \in \mathcal{O}(d)$$

Not Usually Convex! ⇒ Stuck In Local Optima

TOPIC: Convex Point Registration

Main Ideas

Solve easier but related problems (Semi-Definite Programming Relaxation)

Convex global 3D registration with Lagrangian duality

Jesus Briales, Javier Gonzalez-Jimenez

CVPR 2017

Given the correspondences, adding additional constraints to a SDP relaxation gives experimentally globally optimal results

Point Registration via Efficient Convex Relaxation

Haggai Maron, Nadav Dym, Itay Kezurer, Shahar Kovalsky, Yaron Lipman

SIGGRAPH 2016

The size of a SDP relaxation can be reduced by decomposing the original constraints. The resulting problem can be solved faster and gives the same solution as the original relaxation.

TOPIC: Spectral Multi-Shape Matching and Alignment
(Florian Bernard)

TOPIC: Spectral Multi-Shape Matching and Alignment

Solving the multi-way matching problem by permutation synchronization
Deepti Pachauri, Risi Kondor, Vikas Singh
NIPS 2013

A Solution for Multi-Alignment by Transformation Synchronisation
Florian Bernard, Johan Thunberg, Peter Gemmar, Frank Hertel, Andreas Husch, Jorge Goncalves
CVPR 2015

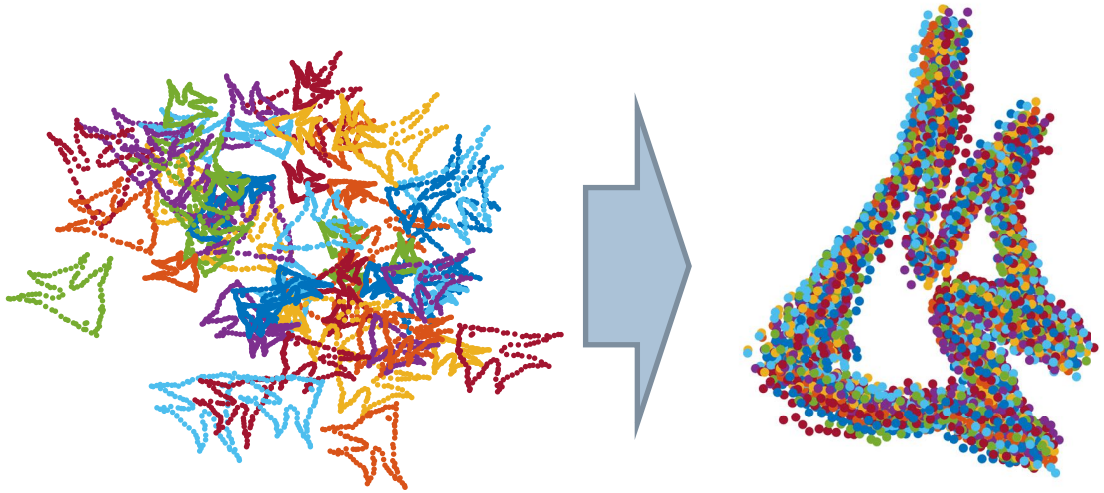

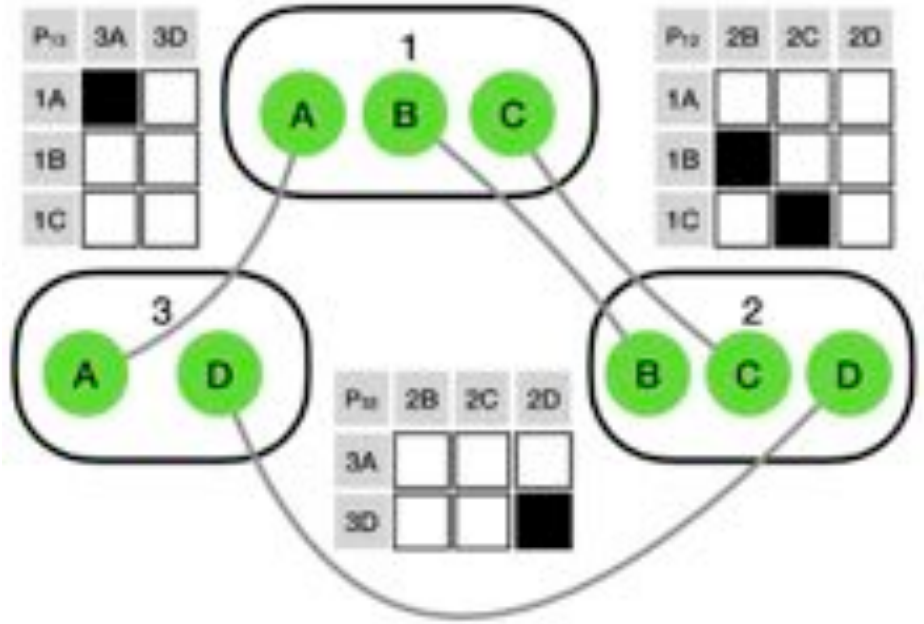

TOPIC: Combinatorial Shape-to-Image Matching

(Florian Bernard)

TOPIC: Combinatorial Shape-to-Image Matching

A Combinatorial Solution for Model-based Image Segmentation and Real-time Tracking
Thomas Schoenemann, Daniel Cremers
TPAMI 2010

A Combinatorial Solution to Non-Rigid 3D Shape-to-Image Matching
Florian Bernard, Frank R. Schmidt, Johan Thunberg, Daniel Cremers
CVPR 2017

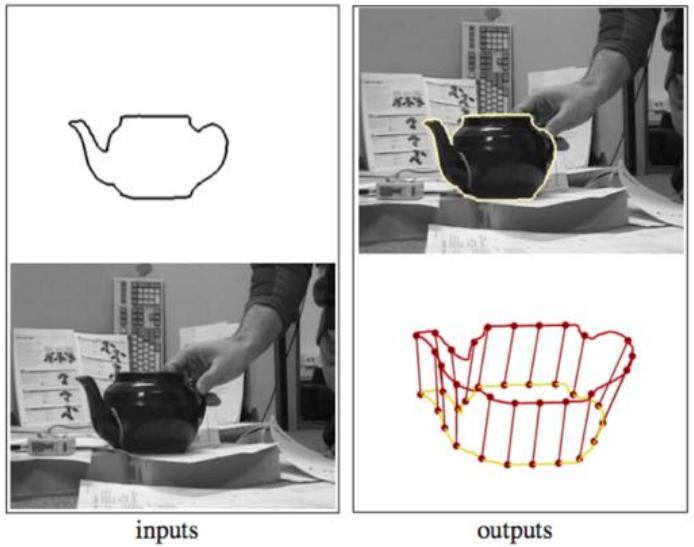

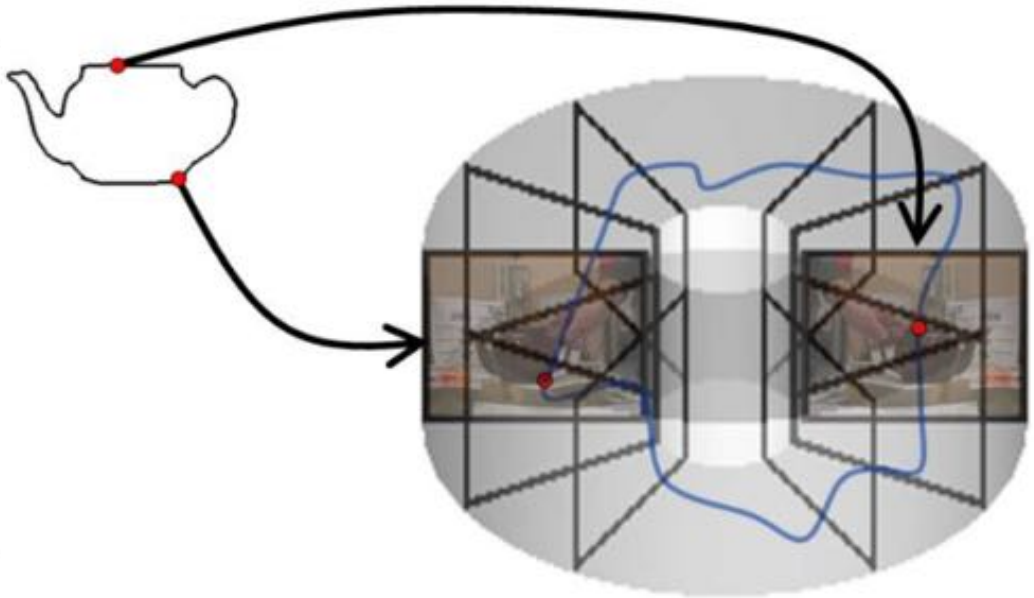

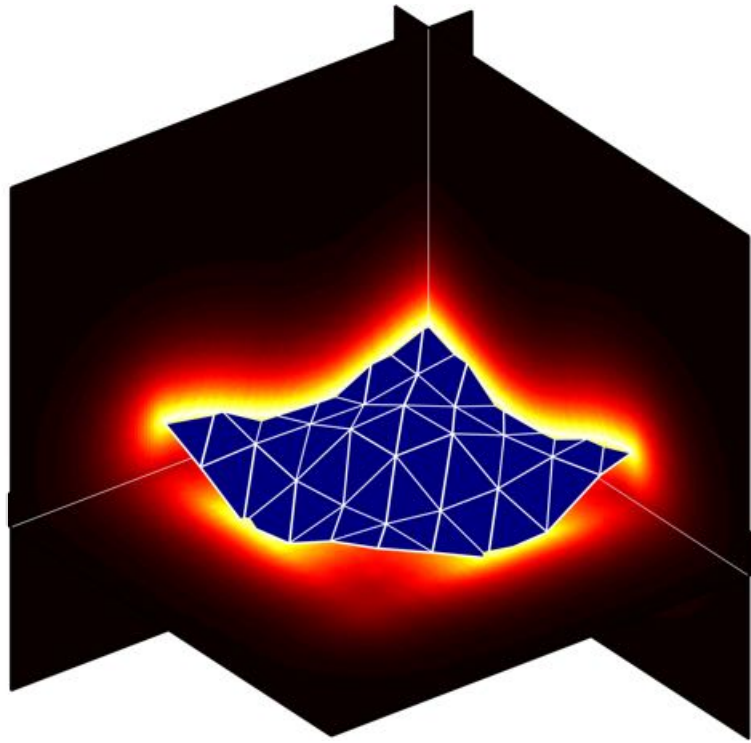

TOPIC: Geometric Deep Learning
(Dushyant Mehta)

TOPIC: Geometric Deep Learning

Deep Learning For 3D Shape Analysis:

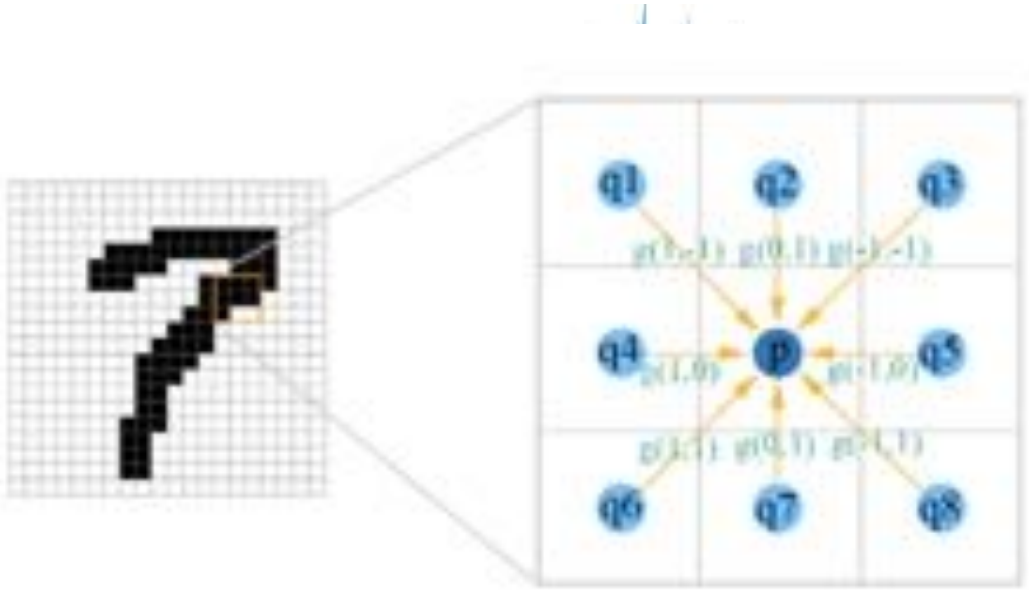

SpiderCNN, arxiv:1803.11527

Data on regular grids
(1D/2D/3D)

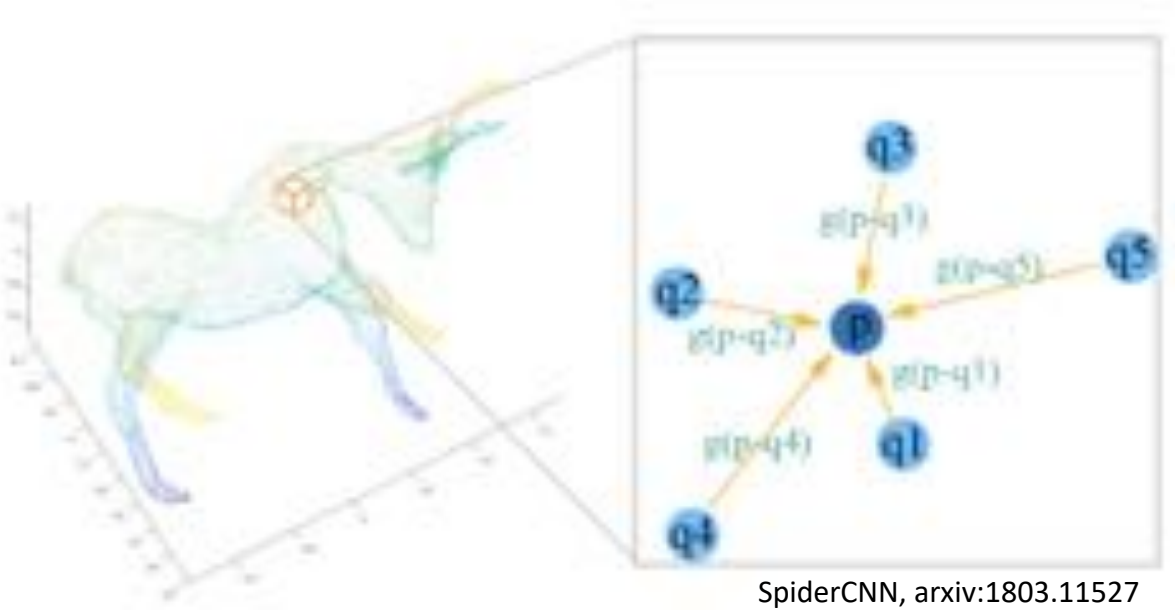

SpiderCNN, arxiv:1803.11527

Unstructured data embedded in
2D or 3D

TOPIC: Geometric Deep Learning

Learning shape correspondence with anisotropic convolutional neural networks

Davide Boscaini, Jonathan Masci, Emanuele Rodolà, Michael Bronstein

NIPS 2016

Anisotropic Kernels

Anisotropic Diffusion Descriptors, Boscaini et al.

TOPIC: Geometric Deep Learning

Geometric deep learning on graphs and manifolds using mixture model CNNs

Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodolà, Jan Svoboda, Michael M. Bronstein

CVPR 2017

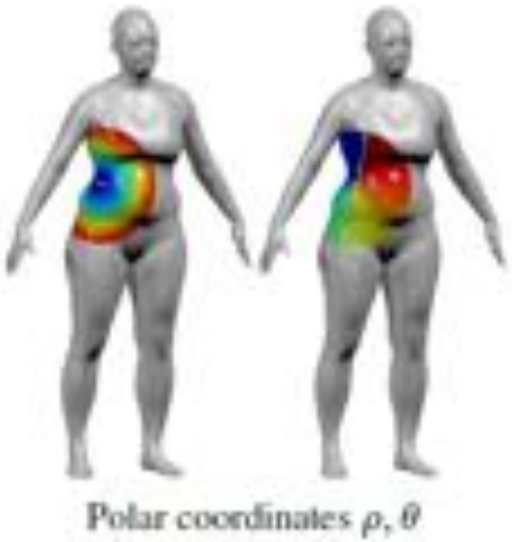

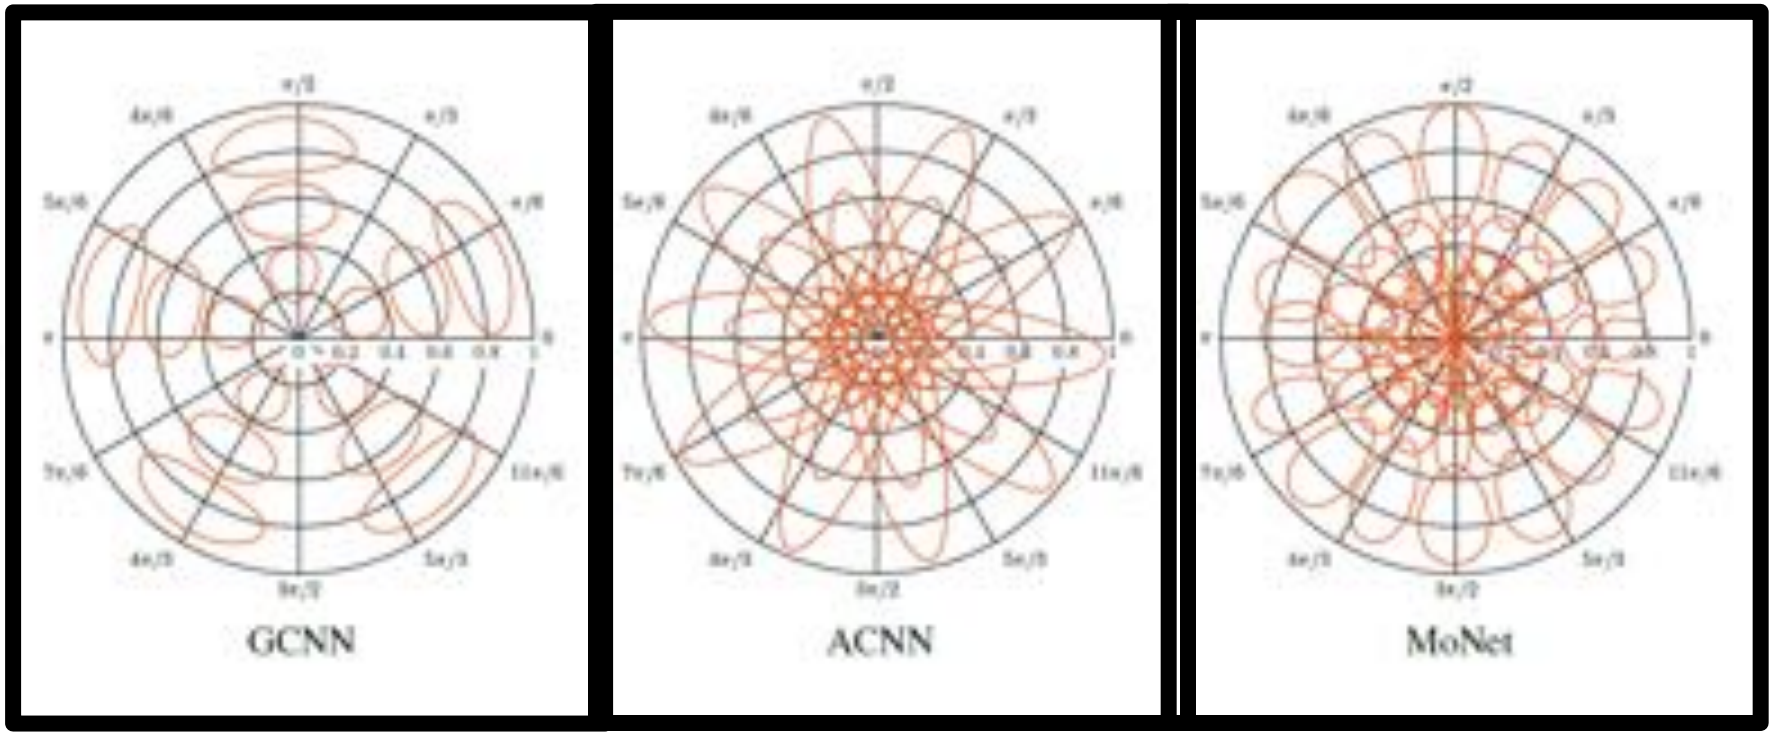

TOPIC: Deep Learning on Point-Clouds
(Ikhsanul Habibie)

TOPIC: Deep Learning on Point Clouds

PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation
Charles R. Qi, Hao Su, Kaichun Mo, Leonidas J. Guibas
CVPR 2017
Point Convolutional Neural Networks by Extension Operators
M. Atzmon, H. Maron, Y. Lipman
SIGGRAPH 2018 (conditionally accepted)

Motivation

- object classification and recognition from point clouds data using deep neural network

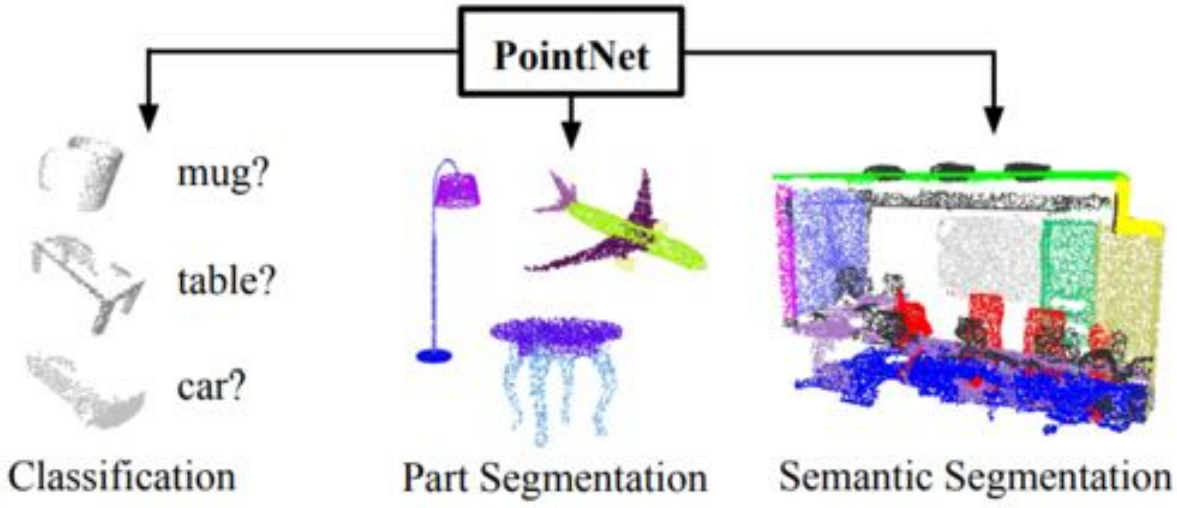

TOPIC: Deep Learning on Point Clouds

Challenges

- Compared to images, point cloud data is unordered
- Data is invariant under transformation
- The network has to capture local interactions across points

How can we use deep learning to solve these problems?

Proposed idea:
transform the space of the set and apply symmetric function

$$f(\{x_1, \ldots, x_n\}) \approx g(h(x_1), \ldots, h(x_n)), \qquad (1)$$

where $f : 2^{\mathbb{R}^N} \to \mathbb{R}$, $h : \mathbb{R}^N \to \mathbb{R}^K$ and $g : \underbrace{\mathbb{R}^K \times \cdots \times \mathbb{R}^K}_{n} \to \mathbb{R}$ is a symmetric function.

Sneak peek into the architecture

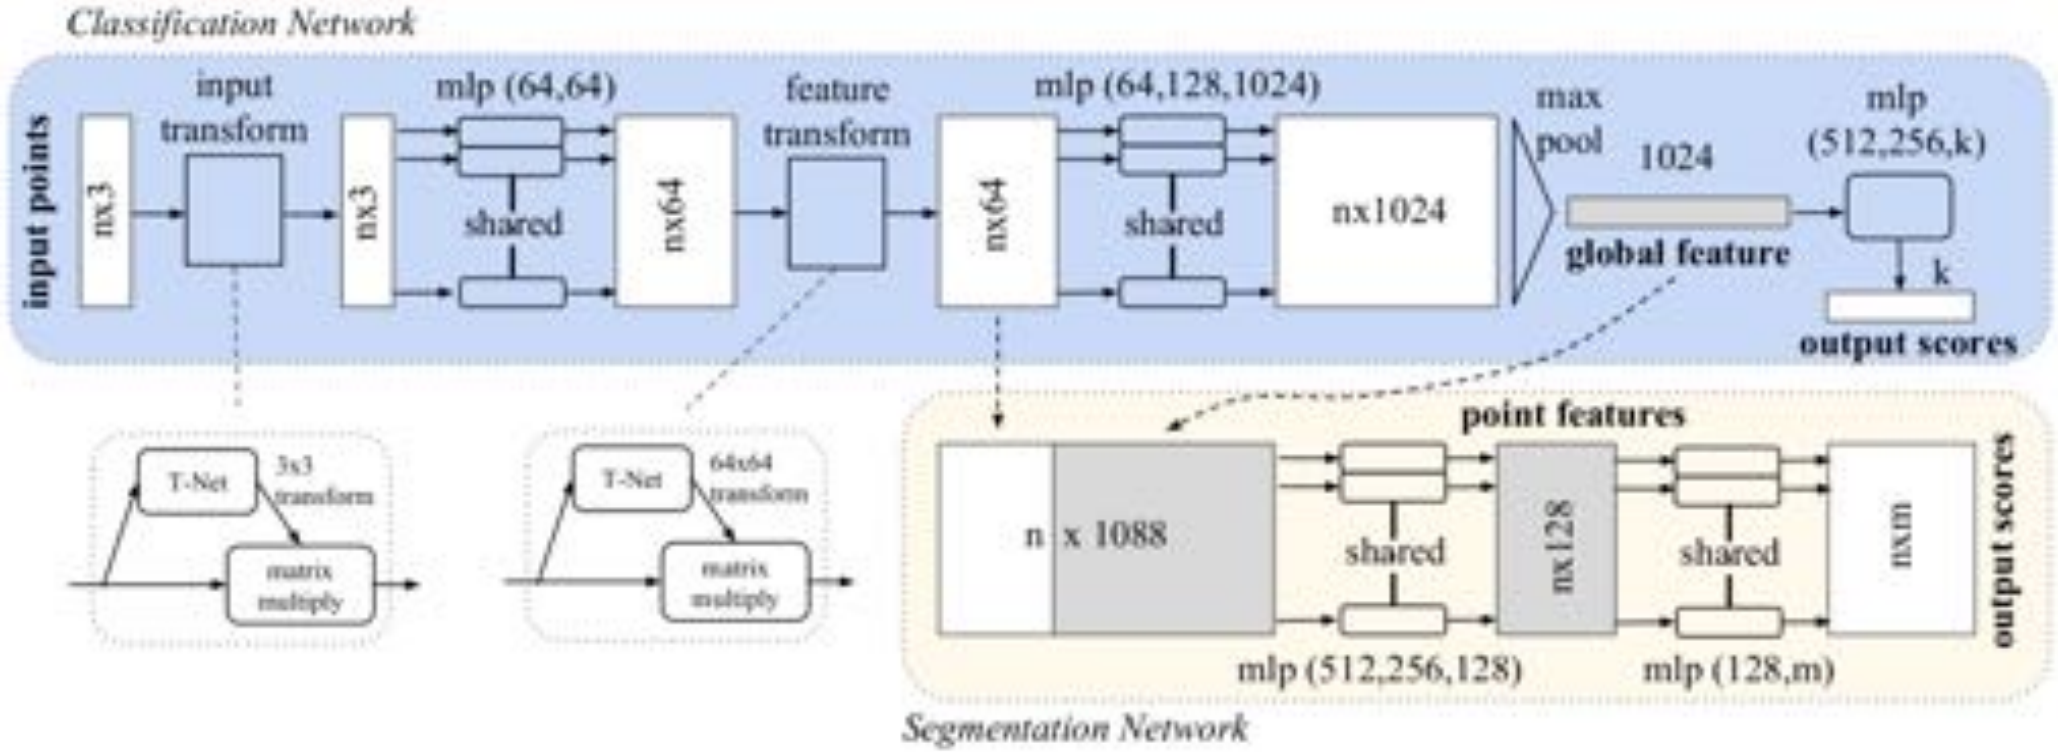

Example results

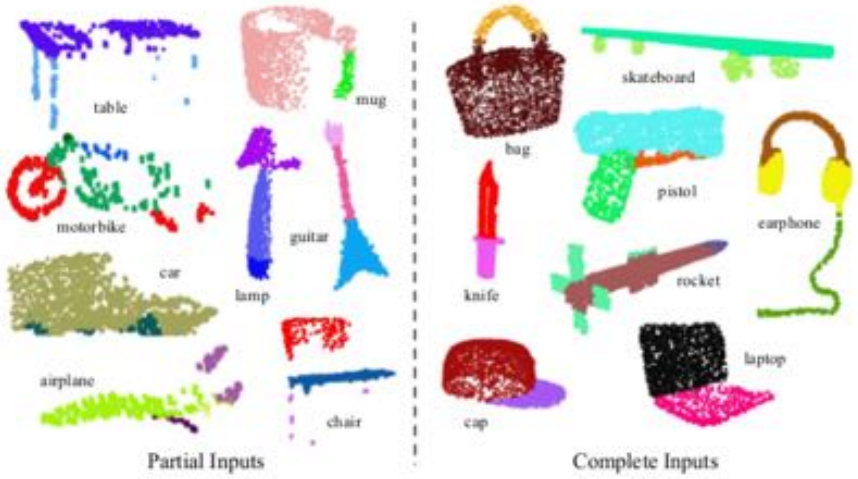

Figure 3. **Qualitative results for part segmentation.** We visualize the CAD part segmentation results across all 16 object categories. We show both results for partial simulated Kinect scans (left block) and complete ShapeNet CAD models (right block).

Figure 4. **Qualitative results for semantic segmentation.** Top row is input point cloud with color. Bottom row is output semantic segmentation result (on points) displayed in the same camera viewpoint as input.

TOPIC: Functional Maps
(Florian Bernard)

TOPIC: Functional Maps

Functional Maps: A Flexible Representation of Maps Between Shapes

Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, Leonidas Guibas

SIGGRAPH 2012

Deep Functional Maps: Structured Prediction for Dense Shape Correspondence

Or Litany, Tal Remez, Emanuele Rodolà, Alex M Bronstein, Michael M Bronstein

ICCV 2017

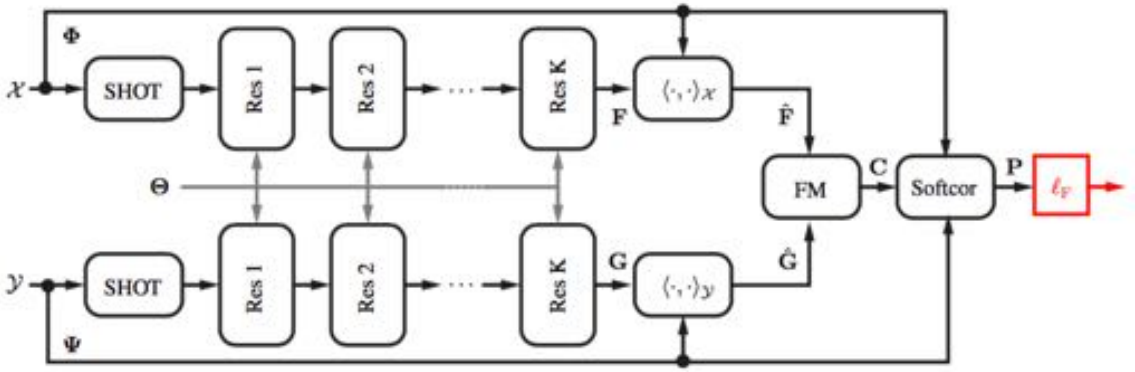

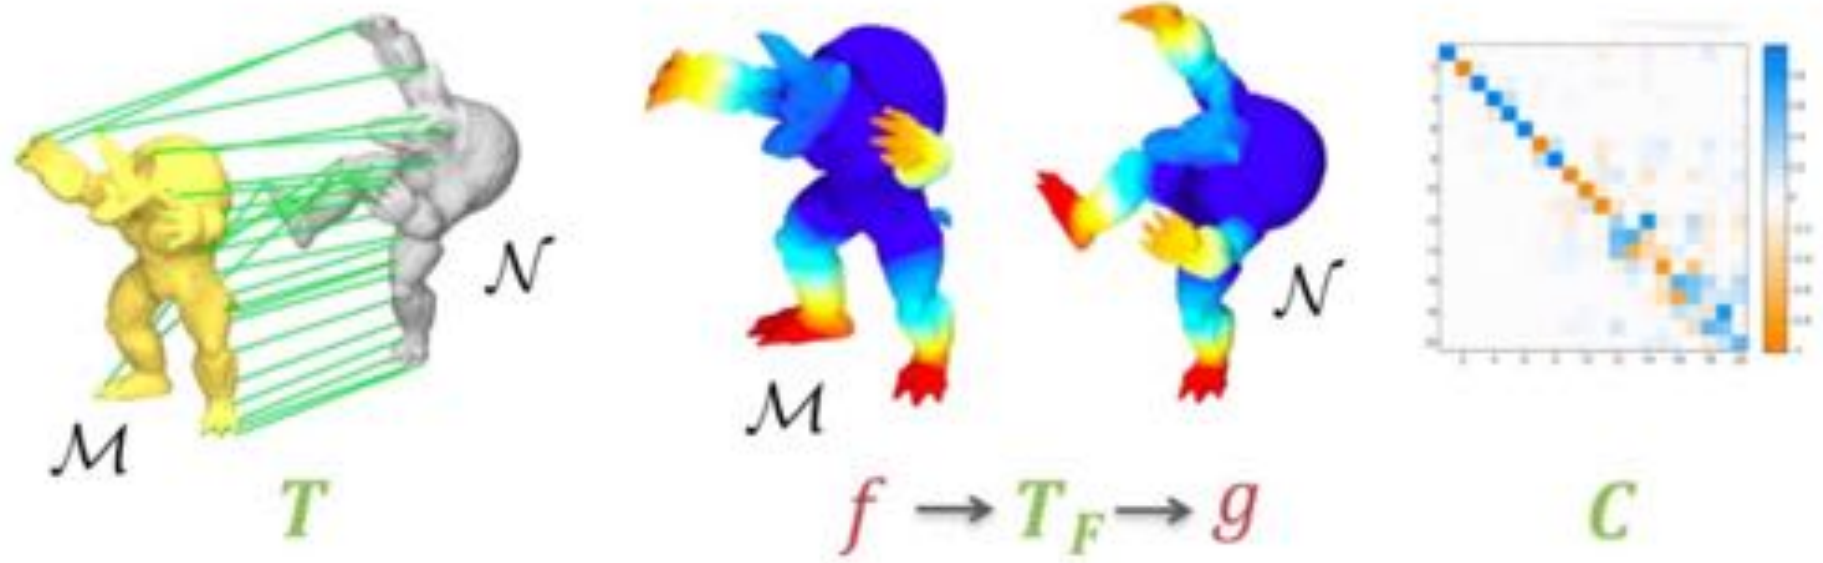

TOPIC: Volumetric Shape Representation
(Franziska Müller)

TOPIC: Volumetric Shape Representation

Shape Transformation Using Variational Implicit Functions

Greg Turk, James F. O'Brien

SIGGRAPH 1999

Reconstruction and Representation of 3D Objects with Radial Basis Functions

J. C. Carr, R. K. Beatson, J. B. Cherrie, T. J. Mitchell, W. R. Fright, B. C. McCallum, T. R. Evans

SIGGRAPH 2001

$f(x) = 0 \Leftrightarrow$

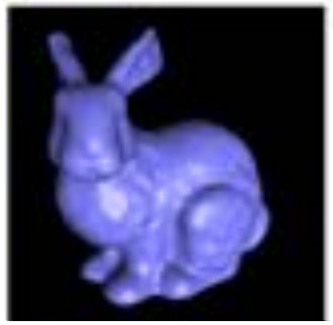

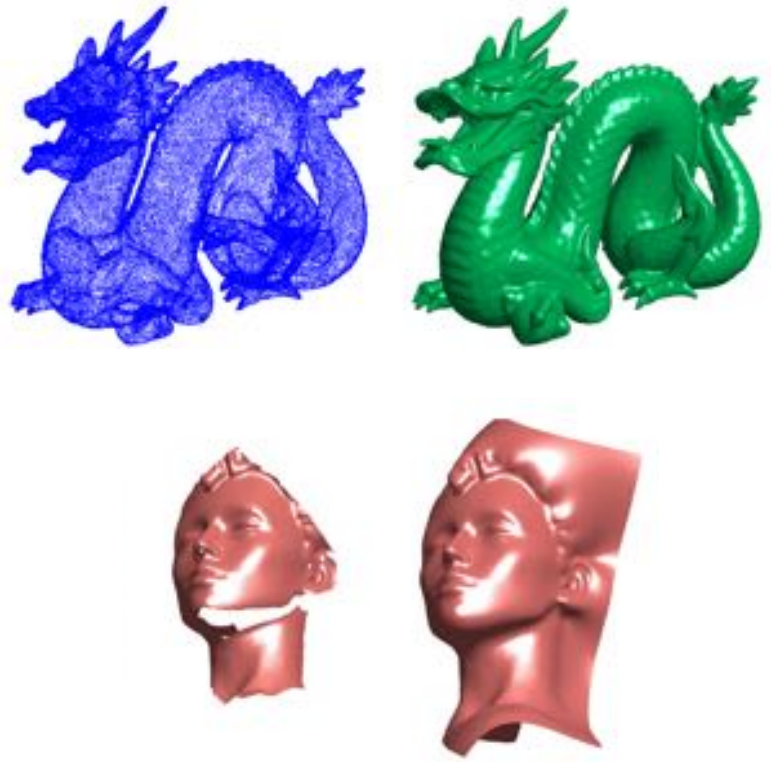

General Remarks on the Topics

- **Disclaimer**: Not all covered papers may still be the state-of-the-art
  - **But:** they contain important ideas/concepts required for pushing the current state-of-the-art forward
  - Your responsibility: cross-read and bring the papers into context (more about this next week)
- **Correspondence problem**: plays a role in (almost) all topics
  - Almost always it will be a good idea to think about assumptions on correspondences (known, unknown, ...)
  - In your presentation, this should be part of the problem statement
- **Problem statement**: Most problems require some time to grasp
  - Modelling (i.e. finding a mathematical description for a given problem) plays a key role in shape analysis
  - In your talk, spend enough time (>10 minutes) in explaining the problem and the modelling choices
    - Often it will be helpful to create additional illustrations and make up toy examples
    - The covered papers have mostly simple and straightforward concepts, even though some technical details may look rather involved
    - We want you to learn how to extract and explain the simple concepts
    - There is no way of explaining all technical details in 30 minutes

Overview

- Organization
- Introduction
- Topics
- Summary

Summary

- Topic assignment:
  - Email us to get a slot (first-come, first-served)
  - Send a list of your 3 preferred topics (in order of preference) until **Sunday, 15 April 2018**
  - We will try to accommodate wishes as much as possible
  - Topics will be assigned on **Monday, 16 April 2018**
- First topic presentation: **Thursday, 03 May 2018**
- Two lectures next week:
  - “How to read an academic paper”
  - “How to give a good scientific talk”

Thanks!

Any questions?