Seminar: Recent Advances in 3D Computer Vision

Björn Häfner, Christiane Sommer
Computer Vision Group
Technical University of Munich

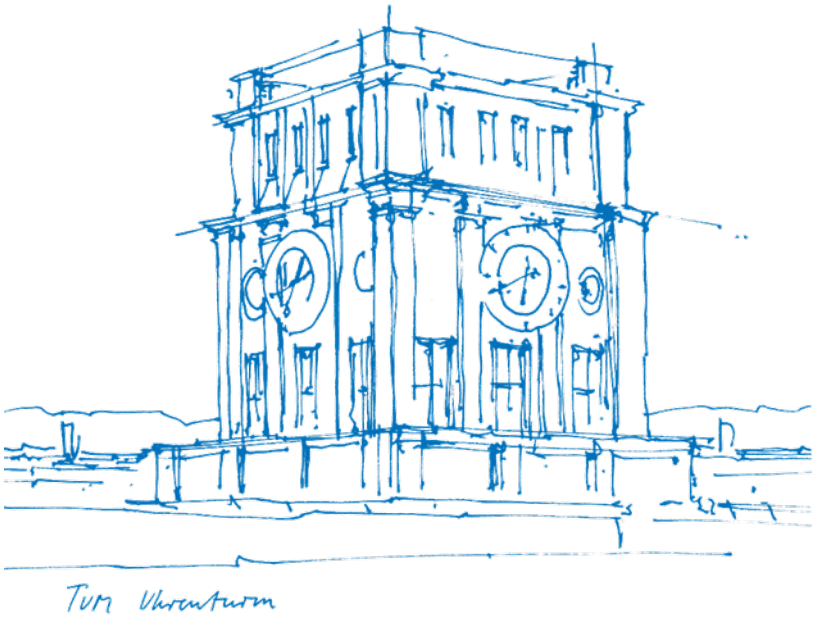

How can I access these slides?

- Option 1 (preferred): seminar web page
  - `https://vision.in.tum.de/teaching/ss2019/seminar_3dcv`
  - Password for material page: `3dcv_ss2019`
  - Material page will soon go online
- Option 2: contact organizers
  - `seminar3dcv@vision.in.tum.de`
  - **Only use this option if you forgot the password**

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - Photometric 3D Reconsruction
  - Static (RGB-)D scanning
  - Dynamic RGB-D scanning
  - Object Detection in Point Clouds
  - Suggest Your Own Paper
- Questions

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - Photometric 3D Reconsruction
  - Static (RGB-)D scanning
  - Dynamic RGB-D scanning
  - Object Detection in Point Clouds
  - Suggest Your Own Paper
- Questions

How is the seminar organized?

- Seminar meetings: talks and discussion
  - Time: Tuesdays, 10:00 - 12:00
  - Room: MI 02.09.023
  - Starting date: TBA (web page)
  - Two talks per week
  - 14 participants $\rightarrow$ 7 weeks
  - **Attendance is mandatory!**
- Talk preparation / contact with supervisor
  - Read through your paper and write down what you don't understand
  - Approx. **one month before** (optional, but recommended) talk: meet supervisor for questions
  - **One week before** (optional, but recommended) talk: meet supervisor to go through slides
  - **One week before** (mandatory) talk: send slides to your supervisor
  - **Two weeks after** talk: submit your report via email

What about the presentation?

- General set-up:
  - Duration: 20-25 minutes talk + 10-15 minutes discussion
  - Make sure to finish on time!
  - Rule of thumb: 1-2 minutes per slide $\rightarrow$ 10-20 slides
  - Do not put too much information on the slides!
- Recommended structure (talk only):
  - Introduction
  - Overview / Outline
  - Method description
  - Experiments and results
  - Personal comments
  - Summary

What about the discussion after each talk?

- Discussion afterwards **will** influence your grade
- Ask questions!
- There are **no** stupid questions!

What about the final report?

- General set-up:
  - Use LaTeX template provided on web page
  - Length: 4-5 pages
  - Send final report as pdf by email to `seminar3dcv@vision.in.tum.de`
  - Submission deadline: two weeks after talk
- Recommended structure (main text only):
  - Introduction
  - Related work
  - Method description
  - Experiments and results
  - Discussion of results
  - Summary

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - Photometric 3D Reconsruction
  - Static (RGB-)D scanning
  - Dynamic RGB-D scanning
  - Object Detection in Point Clouds
  - Suggest Your Own Paper
- Questions

How do you register for the seminar?

- Step 1: Official registration via TUM matching system
  - Go to `matching.in.tum.de`
  - Register for seminar named “Recent Advances in 3D Computer Vision”
- Step 2: Personal registration via email
  - In the list of papers on the web page, select your three favorites
  - Write an email ranking these three favorites to `seminar3dcv@vision.in.tum.de`
  - Email subject: “3DCV seminar application [your name]”
  - Include information about related lectures / courses you have taken so far.
  - We do not need your CV or a motivation letter!
  - Registrations without email / emails with missing information will be ignored!
- Deadline for both registrations: February 13, 2019

How do you register for the seminar?

Example registration email:

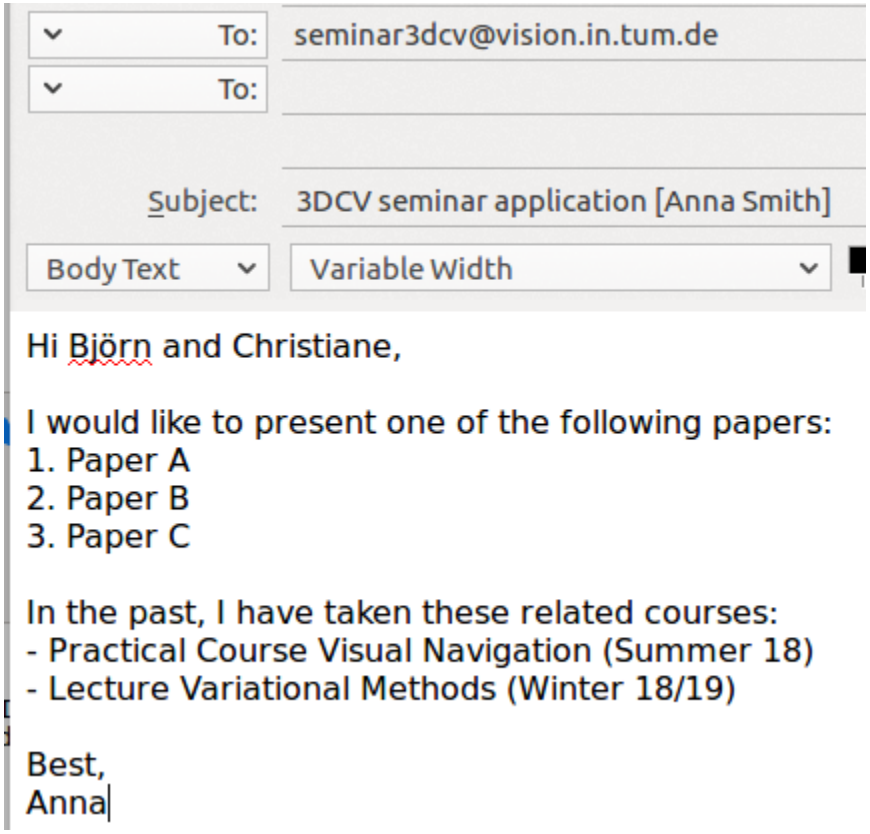

How do we select candidates and assign papers?

- Candidate selection
  - Only students registered in the matching system **AND**

    emails containing all required information will be considered
  - Among students meeting the formal criteria, selection will be random
  - You will get notified by the matching system about the decision (February 20, 2019)
- Paper assignment
  - Papers are assigned after the participant list is finalized
  - We give our best to accommodate your preference list in the assignment

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - Photometric 3D Reconsruction
  - Static (RGB-)D scanning
  - Dynamic RGB-D scanning
  - Object Detection in Point Clouds
  - Suggest Your Own Paper
- Questions

Color Constancy, Intrinsic Images, and Shape Estimation

Barron & Malik 2012, ECCV

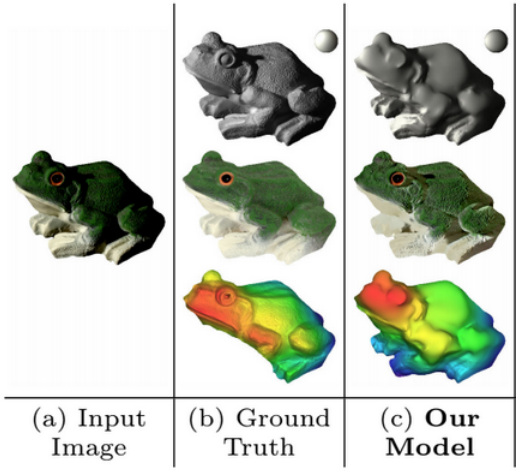

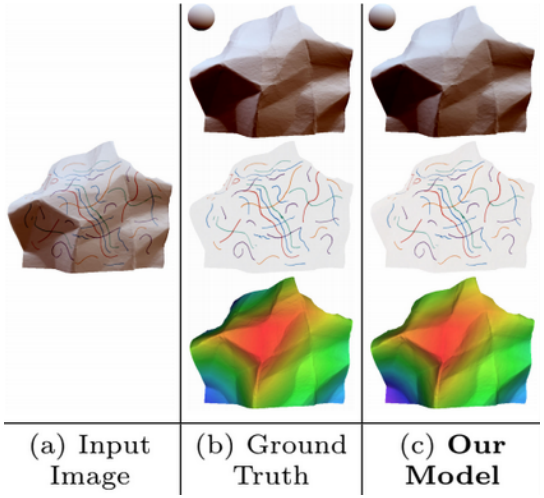

- From a single RGB image estimate the underlying depth map, true color, and illumination
- Makes use of shape-from-shading and machine learning

RGBD-Fusion Real-Time High Precision Depth Recovery

Or-El et al. 2015, CVPR

- From a single RGB-D shot improve the underlying depth map
- Use shape-from-shading and smart optimization

A Non-Convex Variational Approach to Photometric Stereo under Inaccurate Lighting

Quéau et al. 2017, CVPR

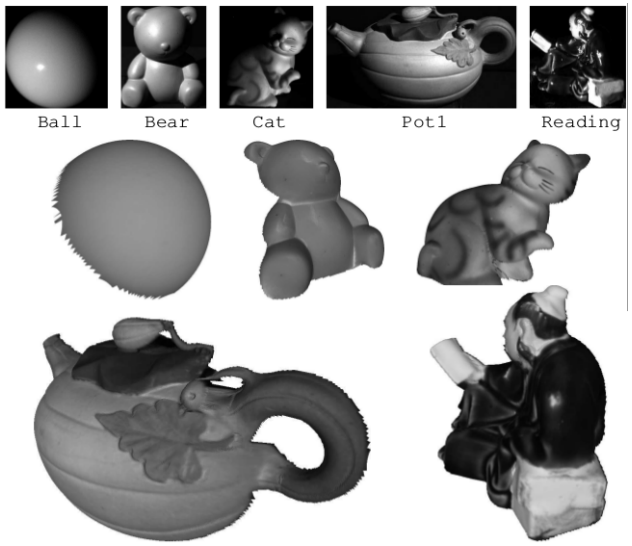

- From multiple (differently illuminated) RGB images estimate depth, true color and lighting
- Make use of Photometric Stereo

Depth Super-Resolution Meets Uncalibrated Photometric Stereo

Peng et al. 2017, ICCVW

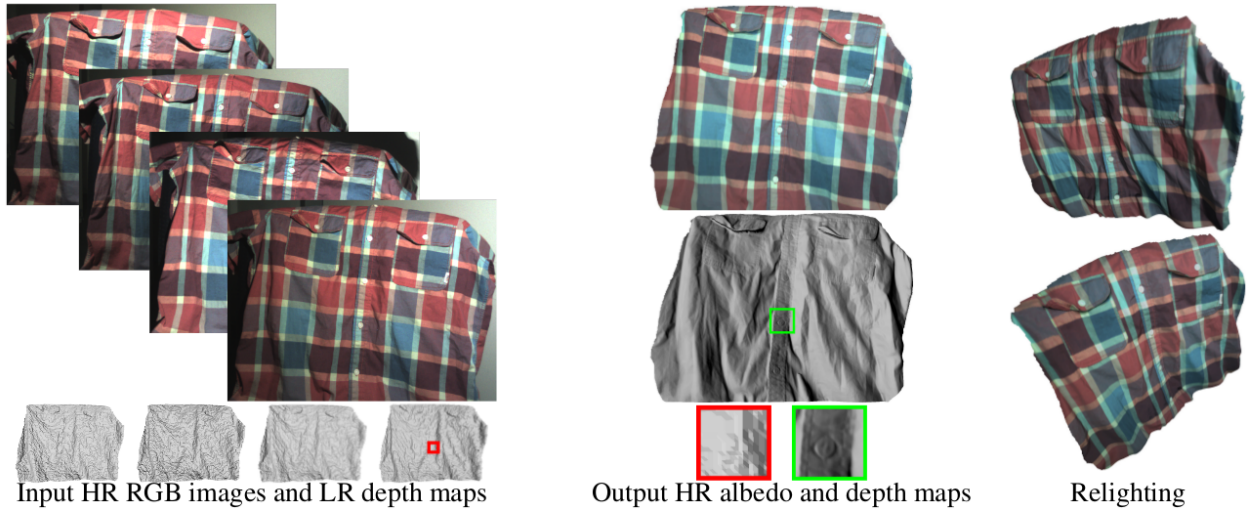

- From multiple (differently illuminated) RGB-D images estimate depth, true color and lighting
- Make use of Photometric Stereo

Practical SVBRDF Acquisition of 3D Objects with Unstructured Flash Photography

NAM et al. 2018, TOG

- Reconstruction of non-lambertian objects using mobile devices

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - Photometric 3D Reconsruction
  - Static (RGB-)D scanning
  - Dynamic RGB-D scanning
  - Object Detection in Point Clouds
  - Suggest Your Own Paper
- Questions

Real-Time Camera Tracking and 3D Reconstruction Using Signed Distance Functions

Bylow et al. 2013, RSS

- Nice introduction to SDFs using RGB-D cameras

KinectFusion: Real-Time Dense Surface Mapping and Tracking

Newcombe et al. 2011, ISMAR

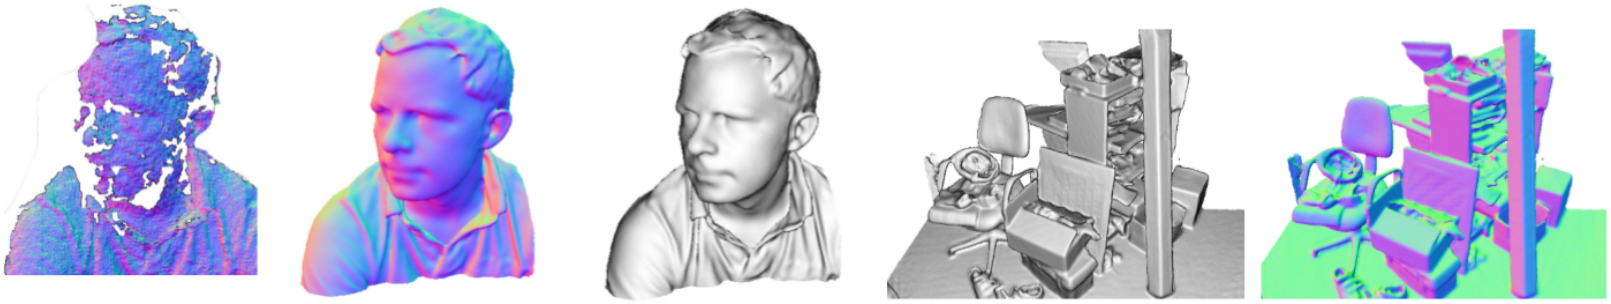

- Highly cited, impactful, baseline method for 3D reconstruction using RGB-D cameras

BundleFusion: Real-Time Globally Consistent 3D Reconstruction Using On-the-Fly Surface Reintegration

Dai et al. 2016, TOG

- Use all depth and color data to obtain consistent mapping

Shading-based Refinement on Volumetric Signed Distance Functions

Zollhöfer et al. 2014, TOG

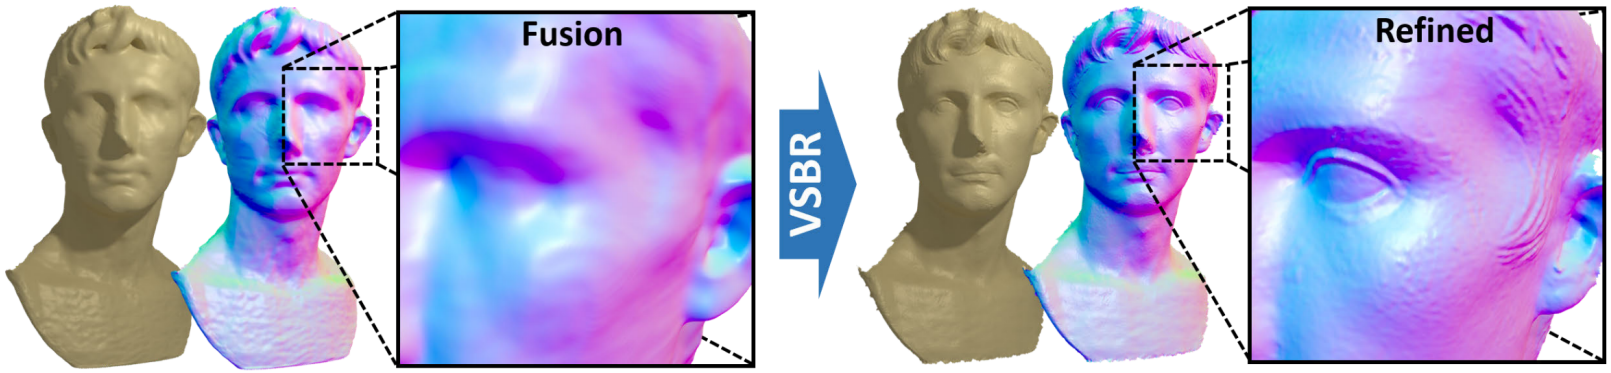

- Optimize geometry using shading information in a shape-from-shading manner
- Use signed distance functions to represent geometry

A Differential Volumetric Approach to Multi-View Photometric Stereo

Logothetis et al. 2018, arxiv

- Optimize geometry using shading information in a photometric stereo manner
- Use signed distance functions to represent geometry

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - Photometric 3D Reconsruction
  - Static (RGB-)D scanning
  - Dynamic RGB-D scanning
  - Object Detection in Point Clouds
  - Suggest Your Own Paper
- Questions

DynamicFusion: Reconstruction and tracking of non-rigid scenes in real time

Newcombe et al. 2015, CVPR

(d) Canonical Model (e) Canonical model warped into its live frame (f) Model Normals

- dynamic (non-rigid) scanning using one Kinect
- generalization of KinectFusion to non-rigid scenes

Fusion4D: real-time performance capture of challenging scenes

Dou et al. 2016, TOG

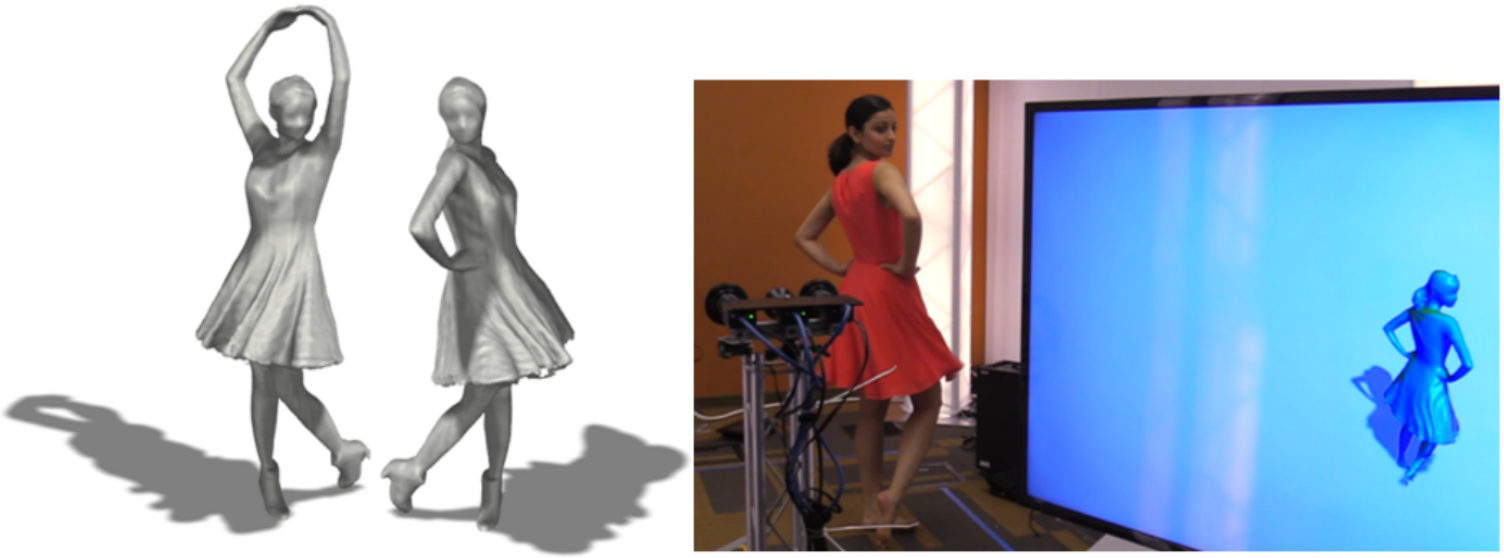

- dynamic scanning in multi-view set-up

Real-time geometry, albedo, and motion reconstruction using a single RGB-D camera

Guo et al. 2017, TOG

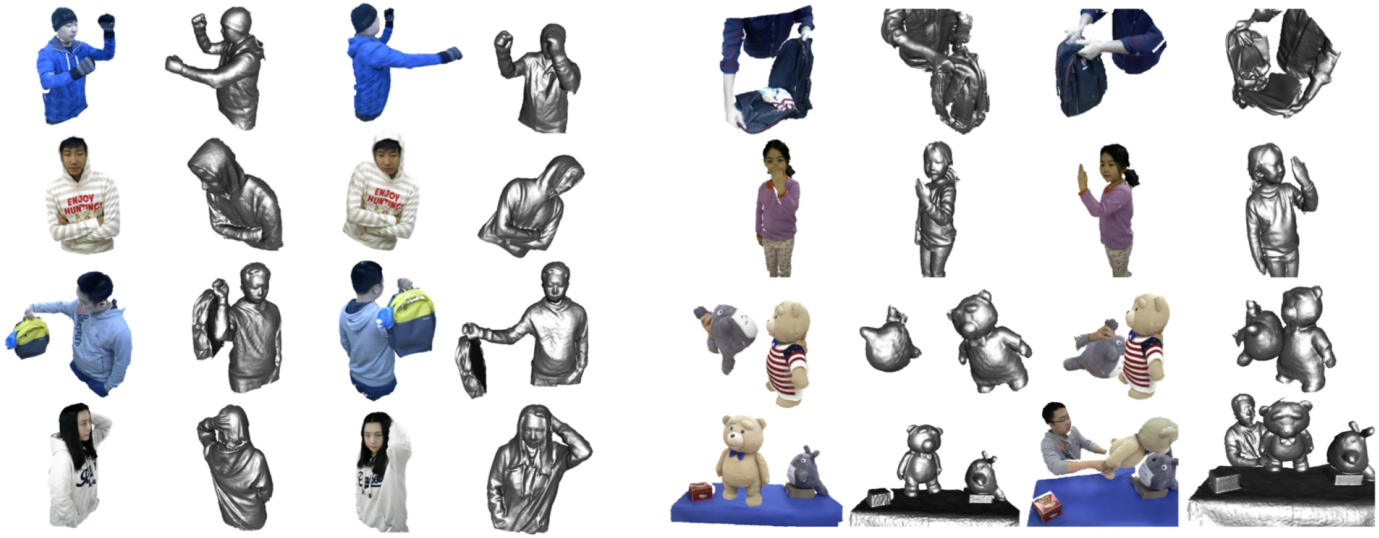

- dynamic scanning with albedo reconstruction

SobolevFusion: 3D Reconstruction of Scenes Undergoing Free Non-rigid Motion

Slavcheva et al. 2018, CVPR

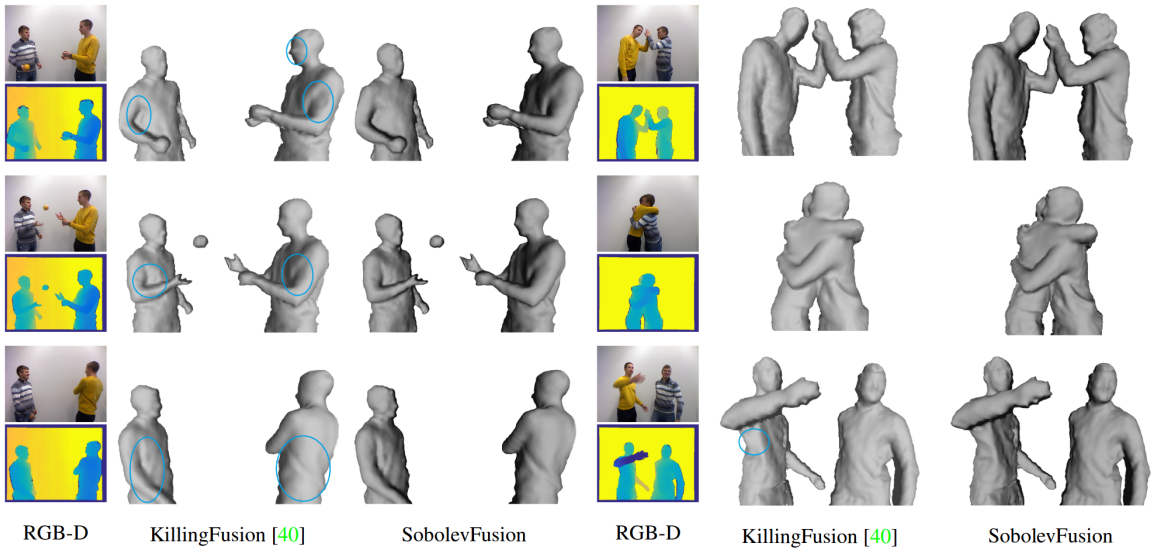

- dynamic scanning using Sobolev gradients
- quite mathematical, but very nice theory

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - Photometric 3D Reconsruction
  - Static (RGB-)D scanning
  - Dynamic RGB-D scanning
  - Object Detection in Point Clouds
  - Suggest Your Own Paper
- Questions

Model globally, match locally: Efficient and robust 3D object recognition

Drost et al. 2010, CVPR

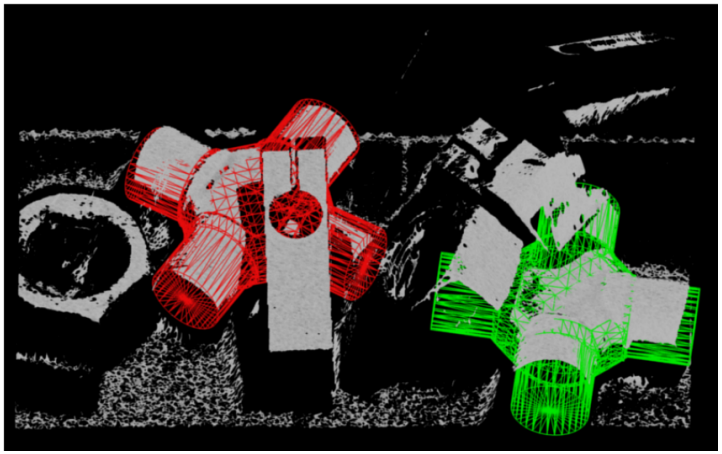

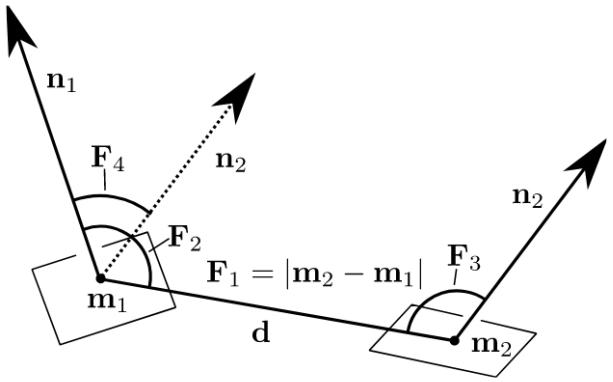

- point pair features to detect rigid body objects in point clouds

Local Hough transform for 3d primitive detection

Drost and Ilic 2015, 3DV

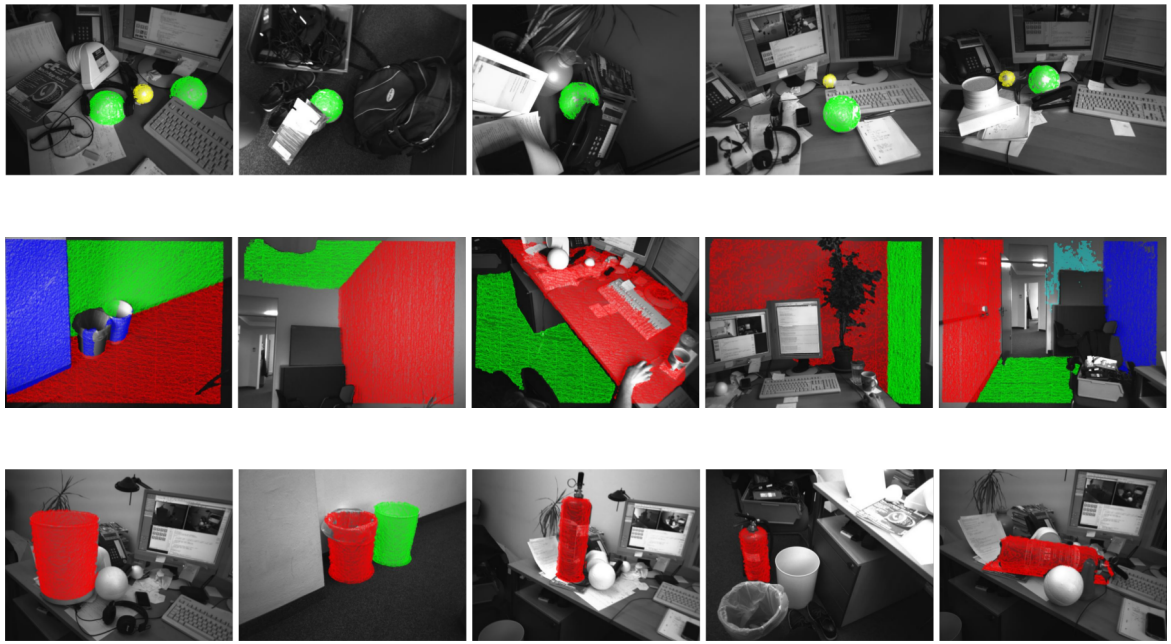

- point pair features to detect geometric primitives in point clouds

PPF-FoldNet: Unsupervised Learning of Rotation Invariant 3D Local Descriptors

Deng et al. 2018, ECCV

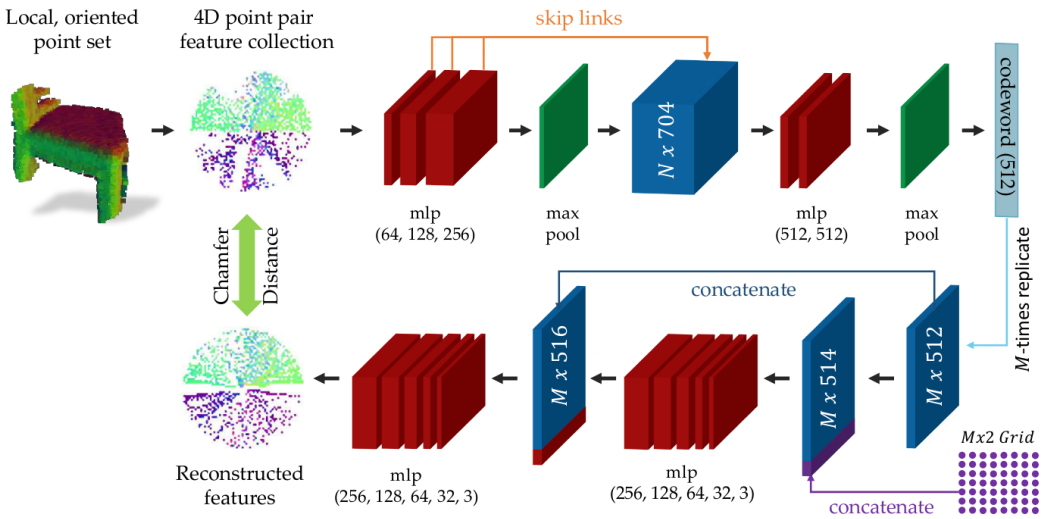

- point pair features as input to deep neural network
- autoencoder architecture for unsupervised learning

A Minimalist Approach to Type-Agnostic Detection of Quadrics in Point Clouds

Birdal et al. 2018, CVPR

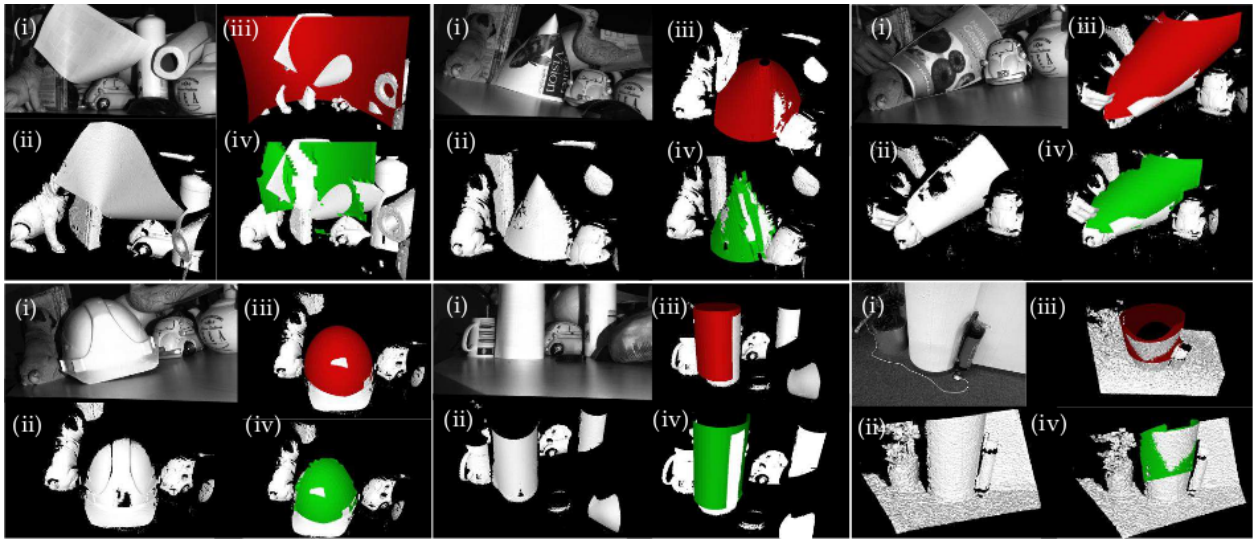

- write quadric detection as linear system
- rather mathy paper

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - Photometric 3D Reconsruction
  - Static (RGB-)D scanning
  - Dynamic RGB-D scanning
  - Object Detection in Point Clouds
  - Suggest Your Own Paper
- Questions

Suggest Your Own Paper

You have **a paper** that

- is not one of the suggested papers above and
- does fit in one of the four topics (Photometric 3D Reconstruction, Static RGB-D Scanning, Dynamic RGB-D Scanning, Object Detection in Point Clouds)
- that you would like to discuss in this seminar?

Then

- state the paper suggested by you as one preference in your personal registration via email
- state in which of the four topics this paper would fit in.

Questions?

Reminder:

- Web page: `vision.in.tum.de/teaching/ss2019/seminar_3dcv`
- Password: `3dcv_ss2019`
- Contact: `seminar3dcv@vision.in.tum.de`