Seminar: The Evolution of Motion Estimation and Real-time 3D Reconstruction

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

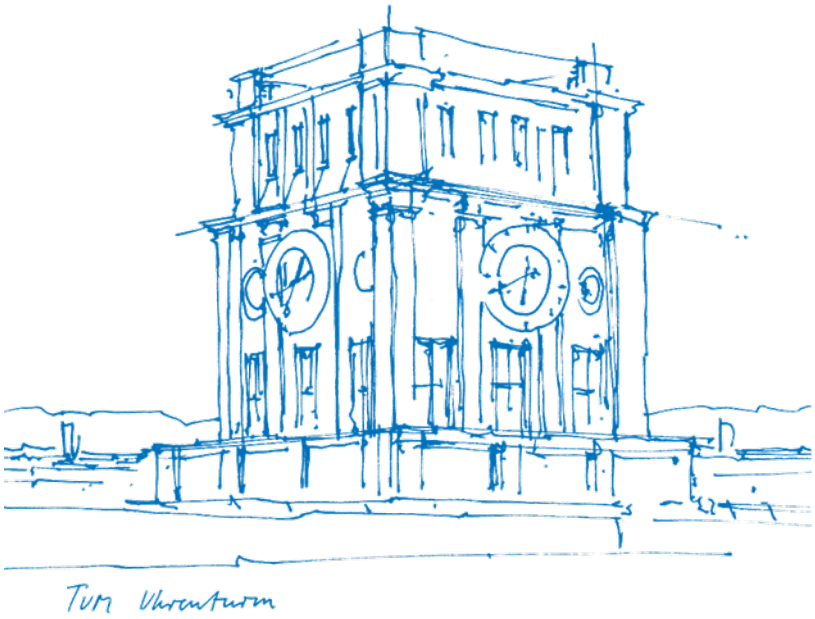

How can I access these slides?

- Option 1 (preferred): seminar web page
  - `https://vision.in.tum.de/teaching/ss2020/seminar_realtime3d`
  - Password for material page: `ss20-realtime3d`
  - Material page will go online after this pre-meeting
- Option 2: contact organizers
  - `realtime3d-ss20@vision.in.tum.de`
  - **Only use this option if you forgot the password**

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - Depth Sensors
  - Monocular Cameras
  - Shading-based Approaches
- Questions

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - Depth Sensors
  - Monocular Cameras
  - Shading-based Approaches
- Questions

How is the seminar organized?

- Seminar meetings: talks and discussion
  - Days: Wednesday, April 15th and Thursday, April 16th
  - Time: 9:00 - 17:00
  - Room: MI 02.09.023
  - 6-7 talks per day
  - **Attendance is mandatory!**
- Talk preparation / contact with supervisor
  - Read through your paper and write down what you don't understand
  - Approx. **one month before** talk (optional, but recommended): meet supervisor for questions
  - **One week before** talk (optional, but recommended): meet supervisor to go through slides
  - **April 9th** (mandatory): send slides to your supervisor
  - **May 3rd** (mandatory): 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 - not too early and not too late!
  - Rule of thumb: 1-2 minutes per slide $\rightarrow$ 10-20 slides
  - Do not put too much information on the slides!
- Recommended structure (talk):
  - 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: 3-4 pages
  - Send final report as pdf by email to `realtime3d-ss20@vision.in.tum.de`
  - Submission deadline: **May 3rd, 2020**
- Recommended structure (main text only):
  - Introduction
  - Method description
  - Experiments and results
  - Discussion of results
  - Summary

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - Depth Sensors
  - Monocular Cameras
  - Shading-based Approaches
- 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 with the title *The Evolution of Motion Estimation and Real-time 3D Reconstruction*
- 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 the seminar email address
  - Email subject: “Real-time 3D 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 12th, 2020

How do you register for the seminar?

Example registration email:

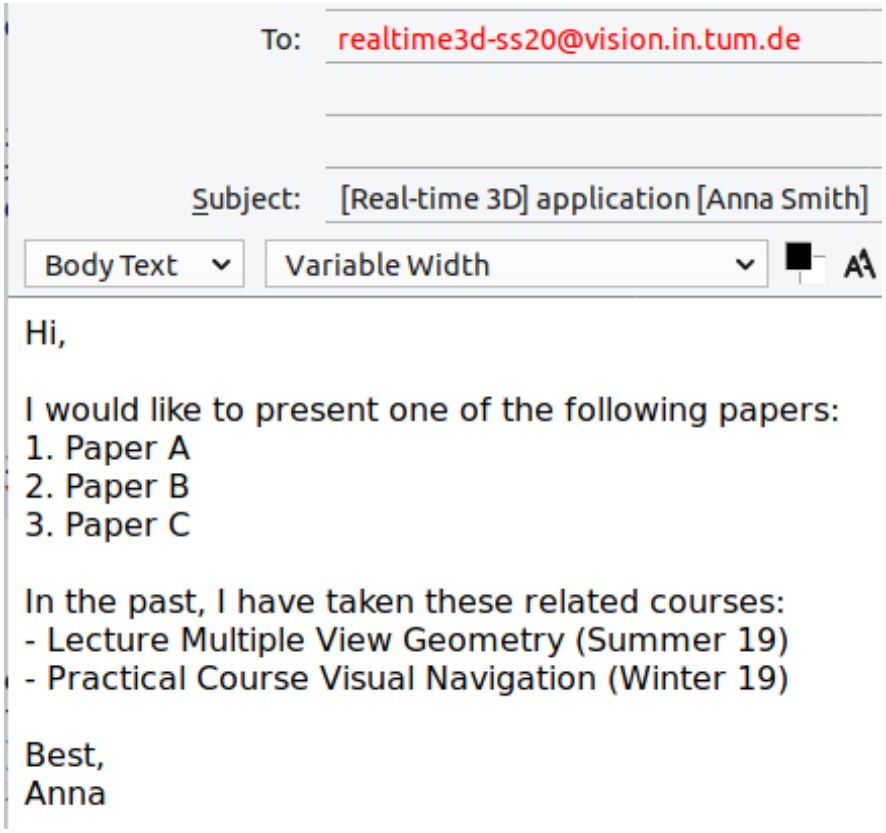
To: realtime3d-ss20@vision.in.tum.de

Subject: [Real-time 3D] application [Anna Smith]

Body Text | Variable Width

Hi,

I would like to present one of the following papers:
1. Paper A
2. Paper B
3. Paper C

In the past, I have taken these related courses:
- Lecture Multiple View Geometry (Summer 19)
- Practical Course Visual Navigation (Winter 19)

Best,
Anna

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 20th, 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
  - Depth Sensors
  - Monocular Cameras
  - Shading-based Approaches
- Questions

KinectFusion: Real-Time Dense Surface Mapping and Tracking

Newcombe, Izadi, Hilliges, Molyneaux, Kim, Davison, Kohli, Shotton, Hodges, Fitzgibbon 2011

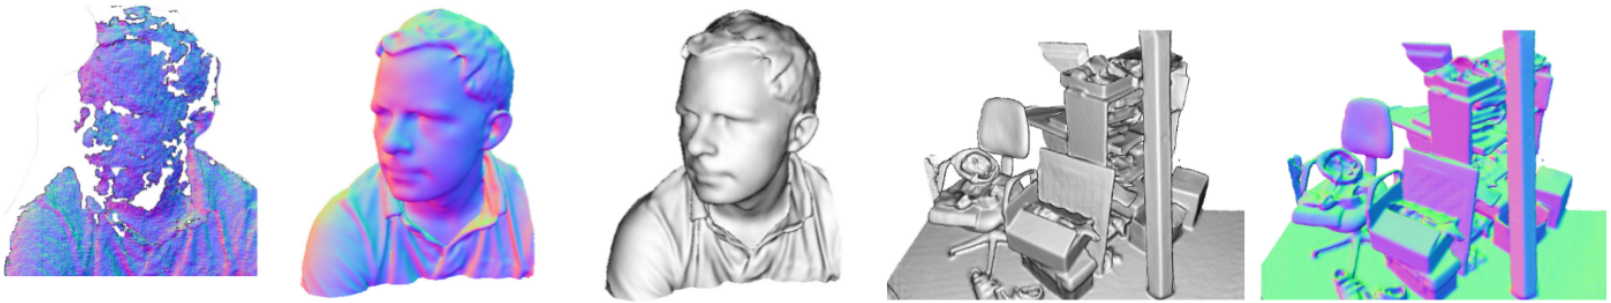

- First paper to generate dense 3D models in real-time using depth sensor and GPU
- Highly cited, impactful, baseline method for 3D reconstruction using RGB-D cameras

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

Bylow et al. 2013, RSS

- Nice introduction to SDFs using RGB-D cameras

Robust Odometry Estimation for RGB-D Cameras

Kerl, Sturm, Cremers 2013

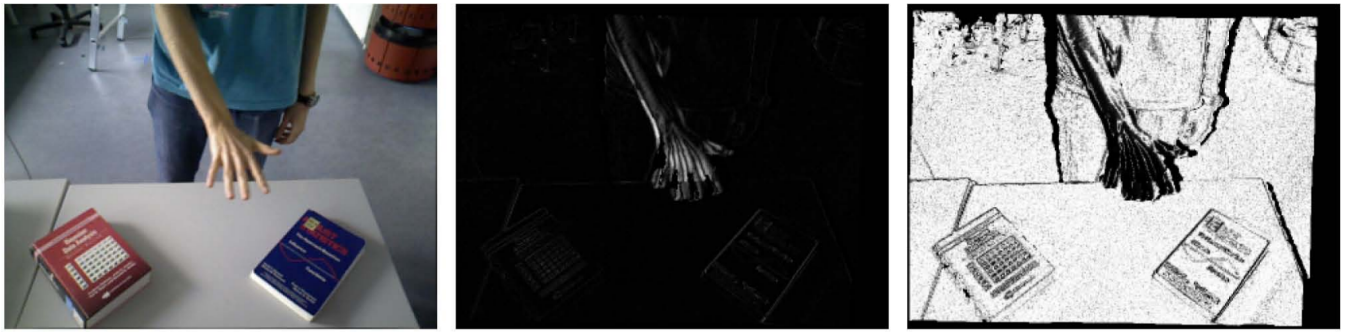

- Odometry method that minimizes photometric cost using depth measurements
- Improved weighting function for robustness in the presence of outliers

ElasticFusion: Dense SLAM Without A Pose Graph

Whelan, Leutenegger, Salas-Moreno, Glocker, Davison 2015

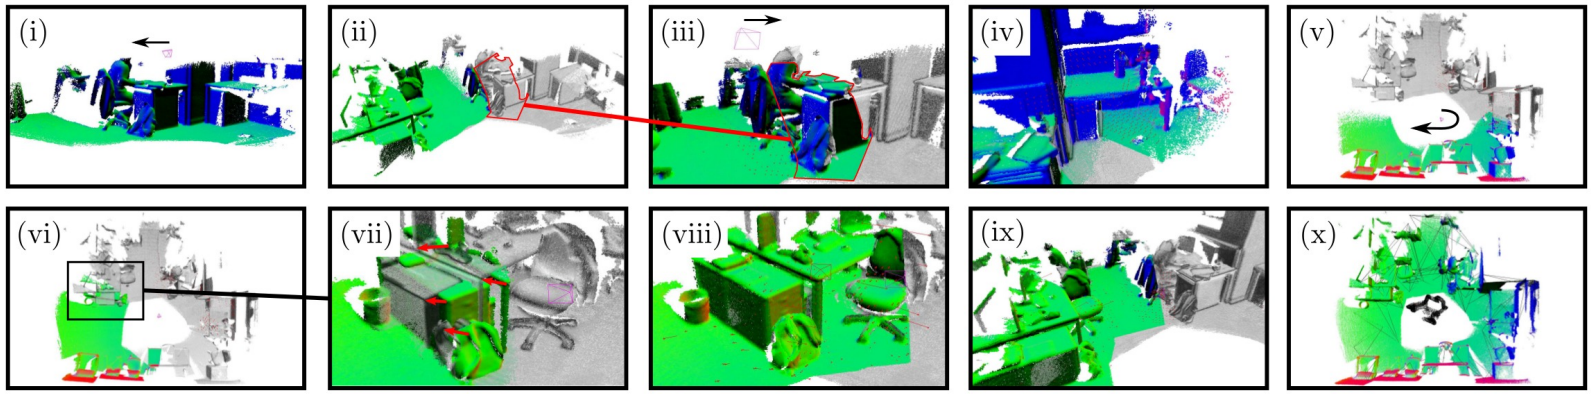

- Uses surfels instead of a TSDF to represent the 3D model
- First method to update the surface in online manner

BAD SLAM: Bundle Adjusted Direct RGB-D SLAM

Schöps, Sattler, Pollefeys 2019

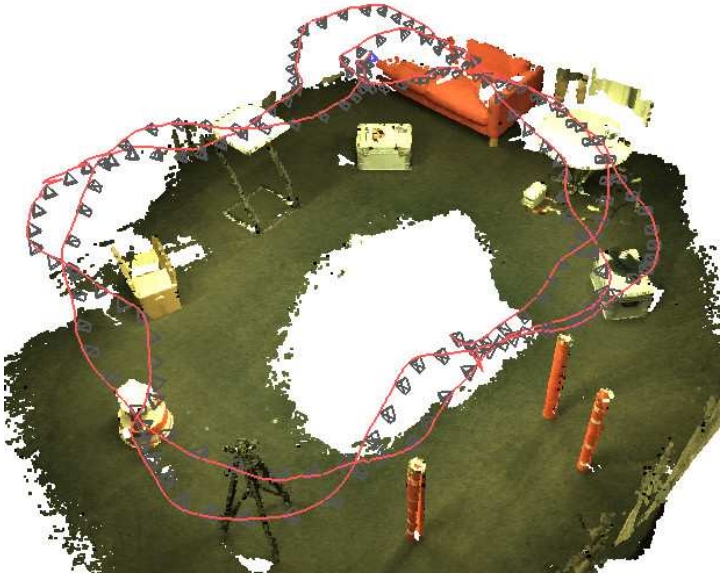

- Published on CVPR 2019
- Perform bundle adjustment on surfels to get a high quality pose

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - Depth Sensors
  - Monocular Cameras
  - Shading-based Approaches
- Questions

PTAM: Parallel Tracking and Mapping

Klein, Murray 2007

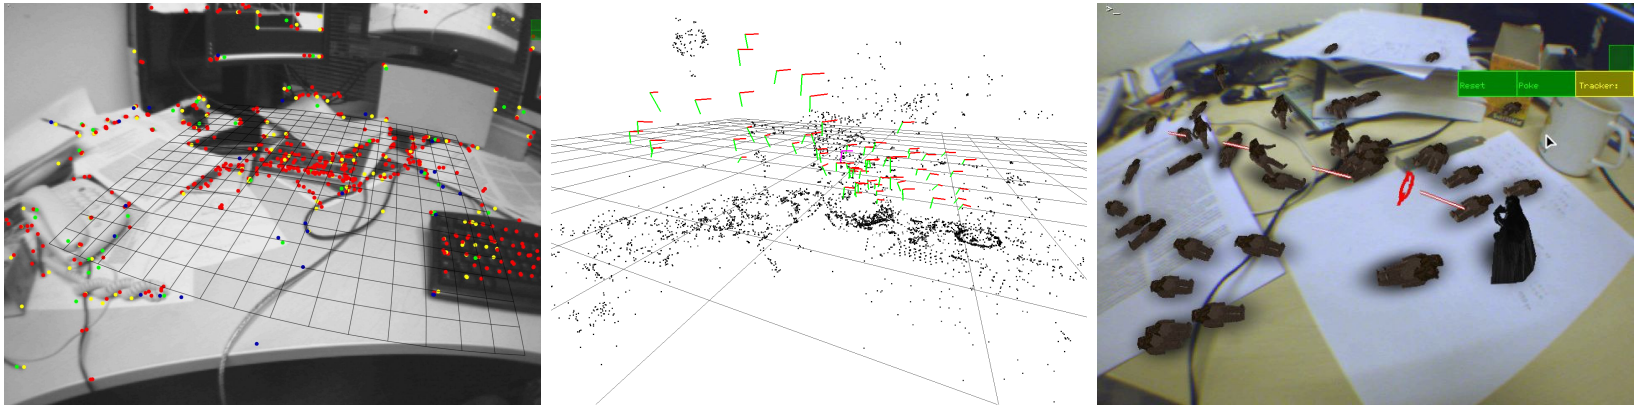

- One of the first systems capable of estimating both pose and geometry in real-time for handheld cameras
- Simple AR applications

DTAM: Dense Tracking and Mapping in Real-Time

Newcombe, Lovegrove, Davison 2011

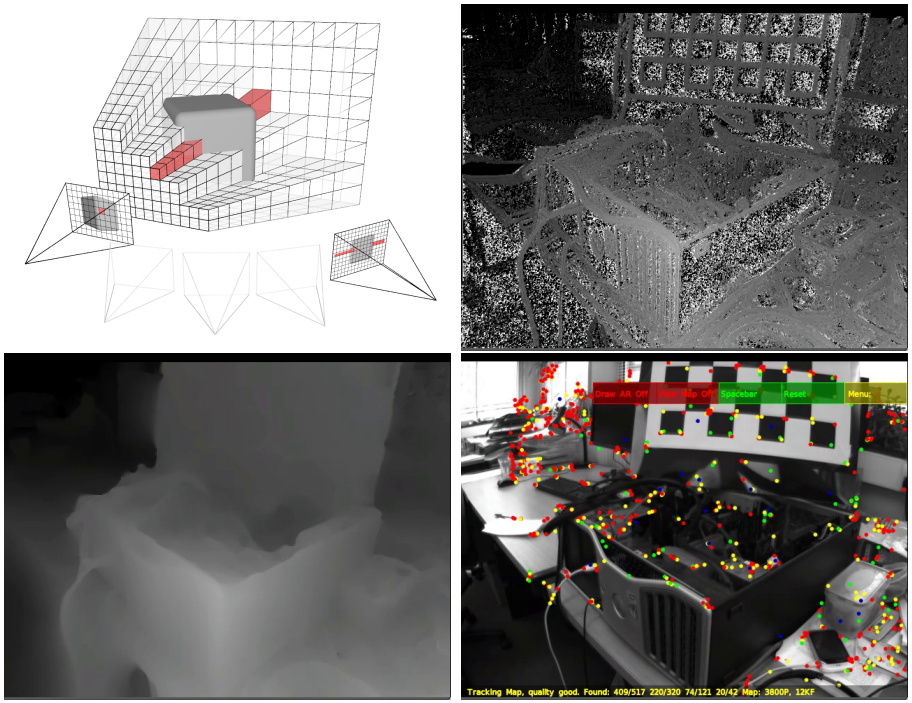

- One of the first monocular systems to create dense 3D models

ORB-SLAM: a Versatile and Accurate Monocular SLAM System

Mur-Artal, Montiel, Tardós 2015

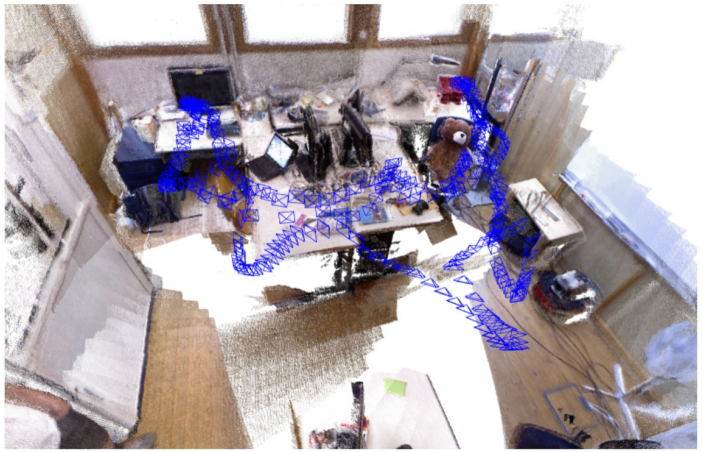

- Use all depth and color data to obtain consistent mapping

Direct Sparse Odometry

Engel, Koltun, Cremers 2016

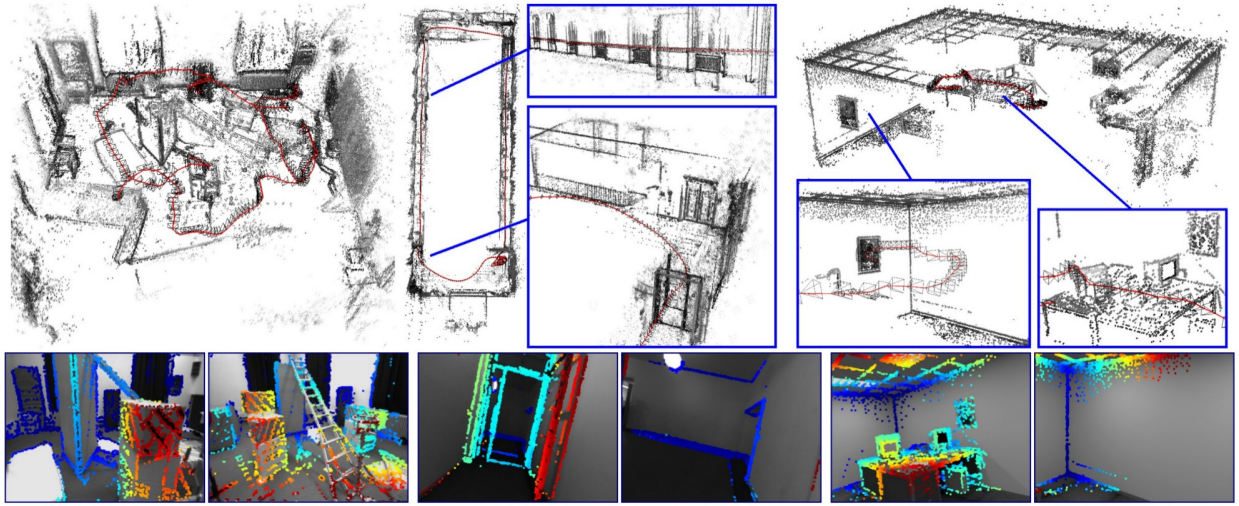

- Large-scale odometry
- Does not rely on keypoint detections

CNN-SLAM

Tateno et al. 2017

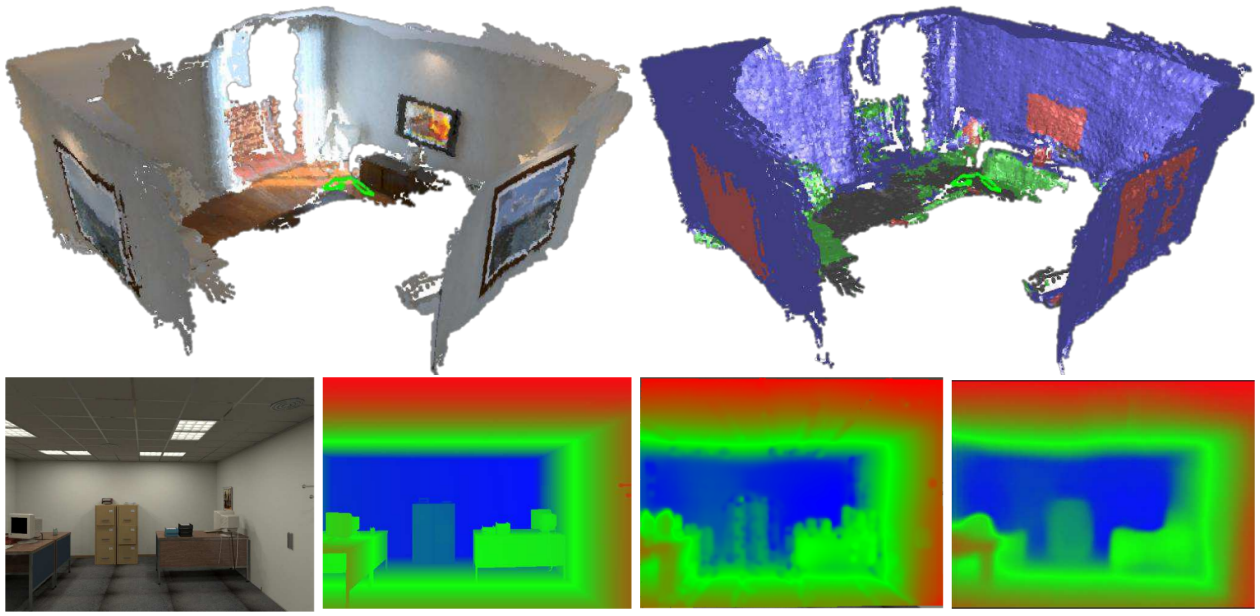

- Dense monocular SLAM
- Use depth map predicted from CNN

CodeSLAM

Michael Bloesch et al. 2018

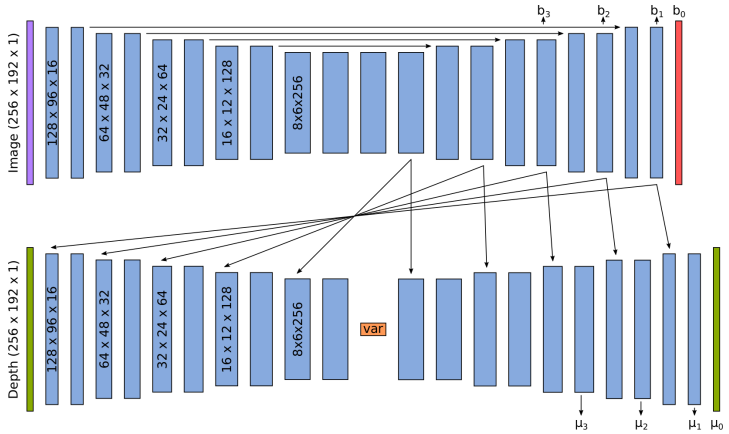

- Learning a compact, optimisable representation of the scene geometry

DeepTAM: Deep Tracking and Mapping

Zhou, Ummenhofer, Brox 2018

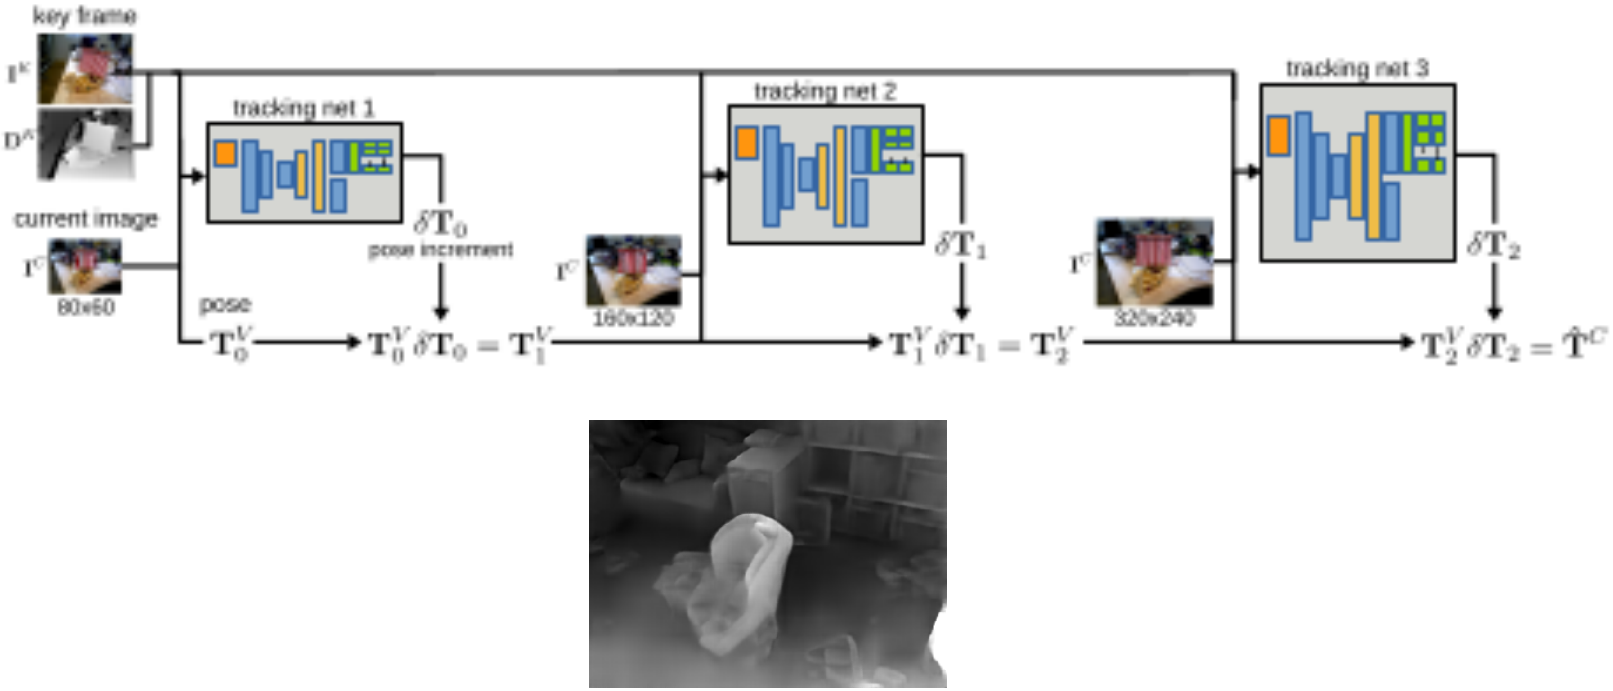

- Learn a network to predict the pose and generate depth images

Deep Virtual Stereo Odometry

Yang et al. 2018

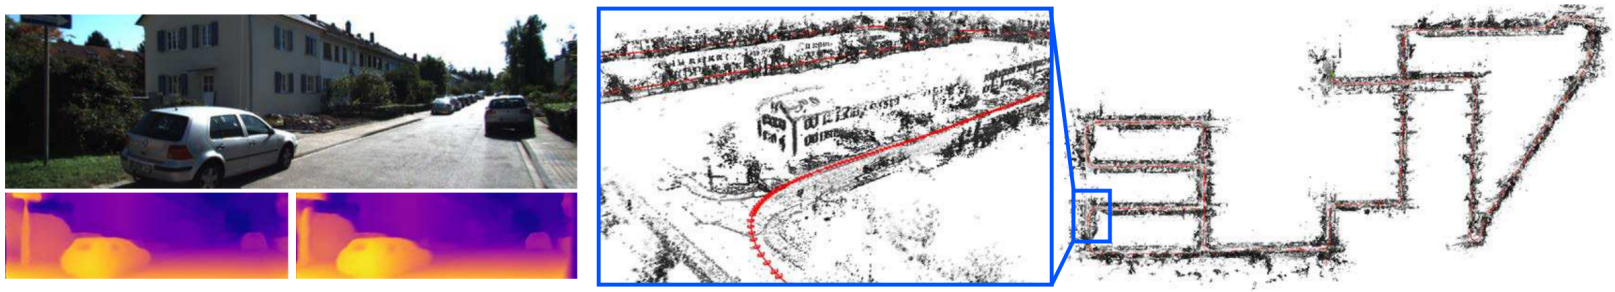

- Use virtual stereo constraints
- Obtain stereo performance on monocular data

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - Depth Sensors
  - Monocular Cameras
  - Shading-based Approaches
- Questions

Shading-based Refinement on Volumetric Signed Distance Functions

Zollhöfer, Dai, Innmann, Wu, Stamminger, Theobalt, Nießner 2015

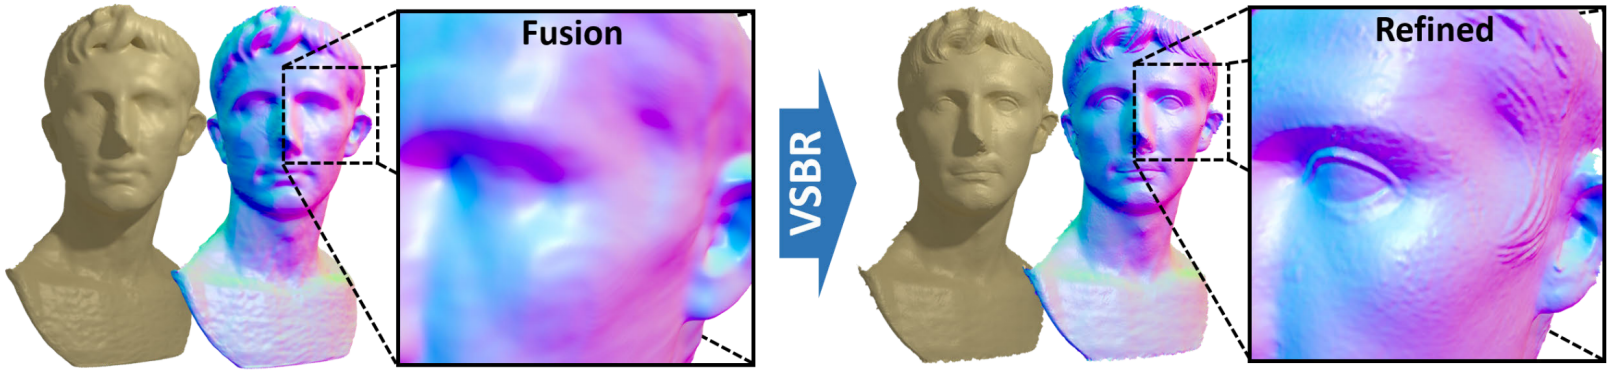

- Optimize geometry using shading (color) information
- Use signed distance functions to represent geometry

RGBD-Fusion: Real-Time High Precision Depth Recovery

Or-El, Rosman, Wetzler, Kimmel, Bruckstein 2015

- 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

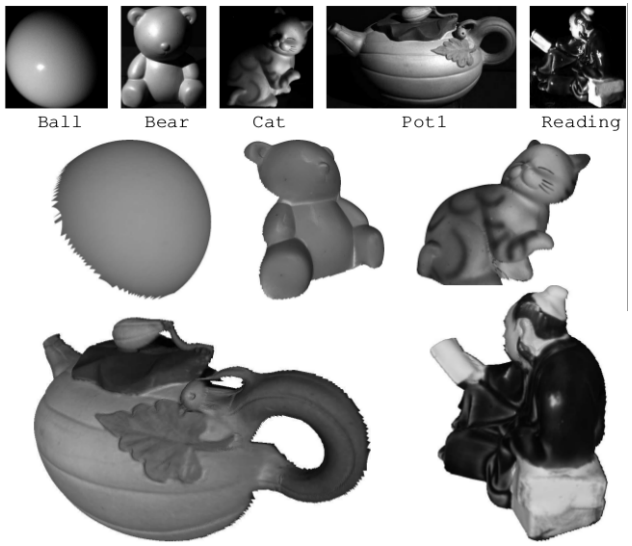

- From multiple (differently illuminated) RGB images estimate depth, true color and lighting
- Make use of Photometric Stereo

Depth Super-Resolution Meets Uncalibrated Photometric Stereo

Peng et al. 2017, ICCVW

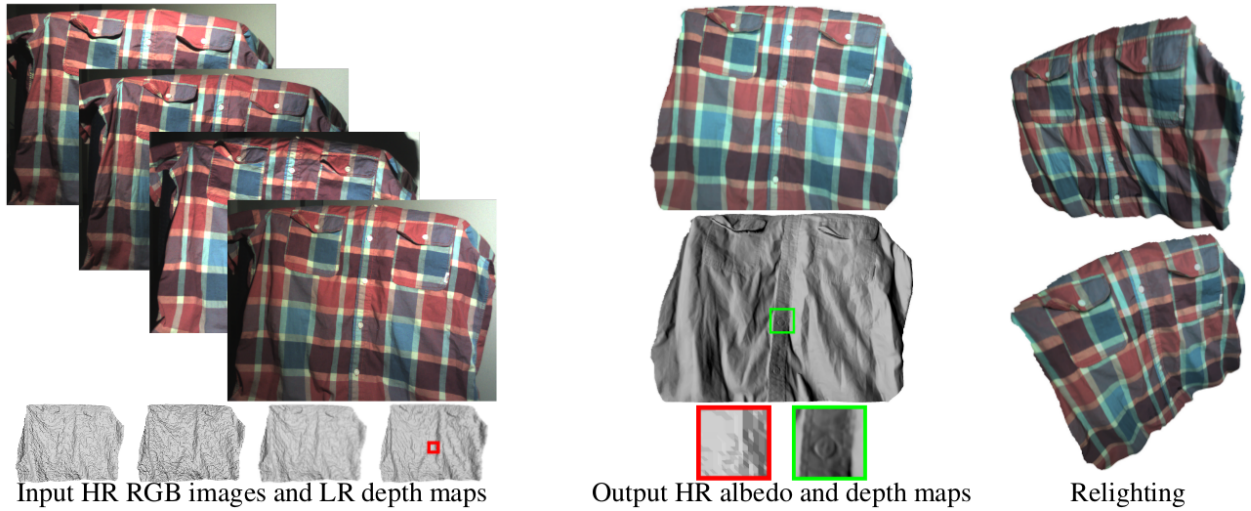

- From multiple (differently illuminated) RGB-D images estimate depth, true color and lighting
- Make use of Photometric Stereo

Fight ill-posedness with ill-posedness: Single-shot variational depth super-resolution from shading

Haefner, Quéau, Möllenhoff, Cremers 2018

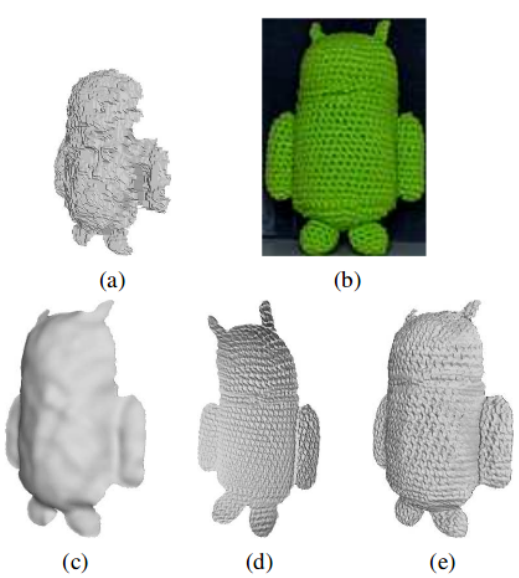

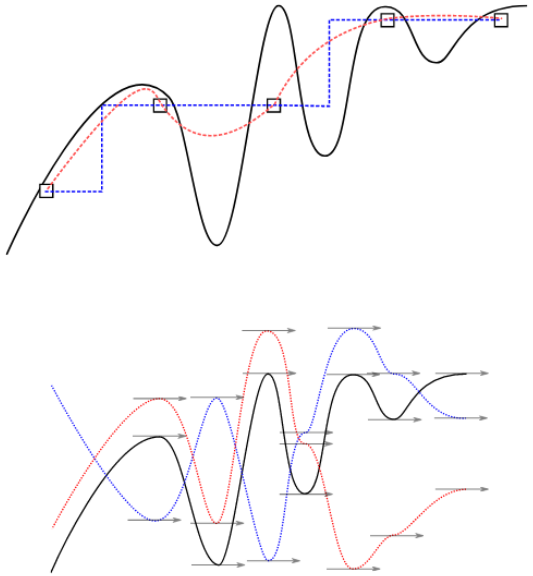

- Generate super-resolved depth image using shading clues
- Overcome shading ambiguity using depth clues

Practical SVBRDF Acquisition of 3D Objects with Unstructured Flash Photography

NAM et al. 2018, TOG

- Reconstruction of non-lambertian objects using mobile devices

A Differential Volumetric Approach to Multi-View Photometric Stereo

Logothetis et al. 2019, ICCV

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

Questions?

Reminder:

- Web page: `vision.in.tum.de/teaching/ss2020/seminar_realtime3d`
- Password: `ss20-realtime3d`
- Contact: `realtime3d-ss20@vision.in.tum.de`