Seminar: Recent Advances in 3D Computer Vision

Erik Bylow, Christiane Sommer
Computer Vision Group
Technical University of Munich

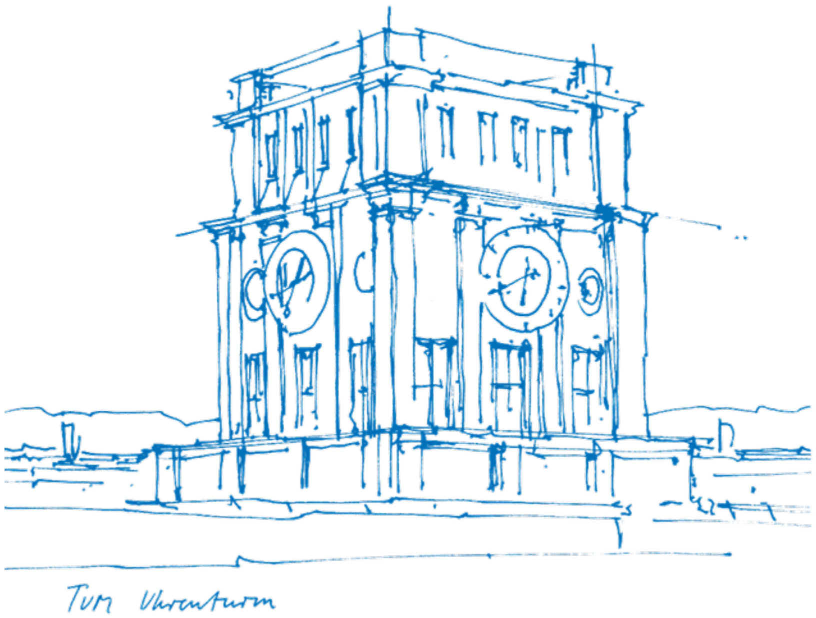

How can I access these slides?

- Option 1 (preferred): seminar web page
  - `https://vision.in.tum.de/teaching/ws2019/seminar_realtime3d` and `https://vision.in.tum.de/teaching/ws2019/seminar_3dgeometry`
  - Passwords for material pages: `realtime3d-ws19` and `3dgeometry-ws19`
  - Material pages will go online after this pre-meeting
- Option 2: contact organizers
  - `realtime3d-ws19@vision.in.tum.de` and `3dgeometry-ws19@vision.in.tum.de`
  - **Only use this option if you forgot the password**

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers: The Evolution of Motion Estimation and Real-time 3D Reconstruction
  - Depth Sensors
  - Monocular Cameras
  - Semantic Tracking and Mapping
  - Event-based Cameras
  - Learning-based Methods
- Possible Papers: An Overview of Methods for Accurate Geometry Reconstruction
  - Multi-view Reconstruction
  - Shading-based Approaches
  - Reconstruction from Point Clouds
  - Learning-based Geometry Estimation
- Questions

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers: The Evolution of Motion Estimation and Real-time 3D Reconstruction
  - Depth Sensors
  - Monocular Cameras
  - Semantic Tracking and Mapping
  - Event-based Cameras
  - Learning-based Methods
- Possible Papers: An Overview of Methods for Accurate Geometry Reconstruction
  - Multi-view Reconstruction
  - Shading-based Approaches
  - Reconstruction from Point Clouds
  - Learning-based Geometry Estimation
- Questions

How is the seminar organized?

- Seminar meetings: talks and discussion
  - Time: Wednesdays, 10:00 - 12:00, every second week
  - Room: MI 02.09.023
  - Starting date: TBA (web page)
  - Two talks per week
  - **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
  - **One week before** talk (mandatory): send slides to your supervisor
  - **Two weeks after** talk (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 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: 3-4 pages
  - Send final report as pdf by email to `realtime3d-ws19@vision.in.tum.de` or `3dgeometry-ws19@vision.in.tum.de`, depending on which seminar you attend
  - Submission deadline: **two weeks after talk**
- Recommended structure (main text only):
  - Introduction
  - Method description
  - Experiments and results
  - Discussion of results
  - Summary

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers: The Evolution of Motion Estimation and Real-time 3D Reconstruction
  - Depth Sensors
  - Monocular Cameras
  - Semantic Tracking and Mapping
  - Event-based Cameras
  - Learning-based Methods
- Possible Papers: An Overview of Methods for Accurate Geometry Reconstruction
  - Multi-view Reconstruction
  - Shading-based Approaches
  - Reconstruction from Point Clouds
  - Learning-based Geometry Estimation
- 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 you want to attend
- 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: “[seminar short title] 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: July 24, 2019

How do you register for the seminar?

Example registration email:

To: realtime3d-ws19@vision.in.tum.de

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

Body Text | Variable Width

Hi Erik and Christiane,

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 (Summer 19)

Best,
John

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 (July 30, 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: The Evolution of Motion Estimation and Real-time 3D Reconstruction
  - Depth Sensors
  - Monocular Cameras
  - Semantic Tracking and Mapping
  - Event-based Cameras
  - Learning-based Methods
- Possible Papers: An Overview of Methods for Accurate Geometry Reconstruction
  - Multi-view Reconstruction
  - Shading-based Approaches
  - Reconstruction from Point Clouds
  - Learning-based Geometry Estimation
- Questions

KinectFusion: Real-Time Dense Surface Mapping and Tracking

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

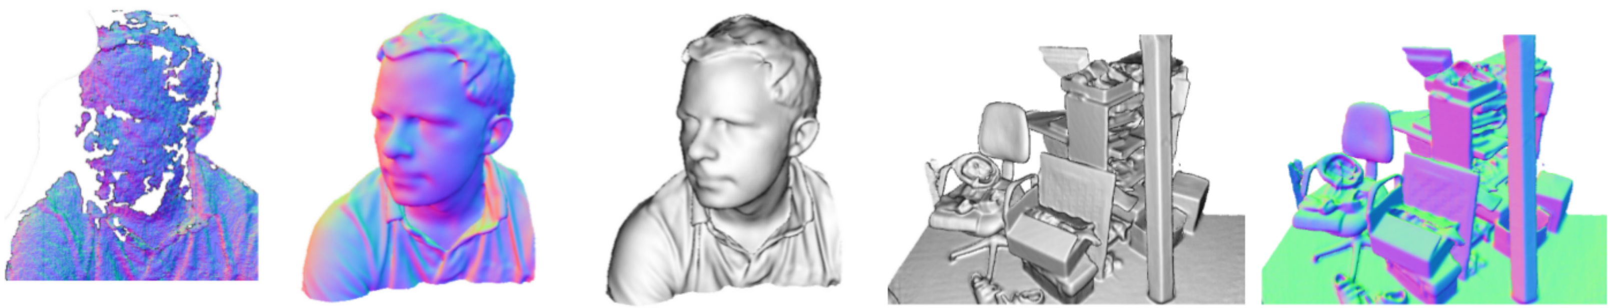

- First paper to generate dense 3D models in real-time using depth sensor and GPU
- Highly cited, impactful, baseline method for 3D reconstruction using RGB-D cameras

Dense Visual SLAM for RGB-D Cameras

Kerl, Sturm, Cremers 2013

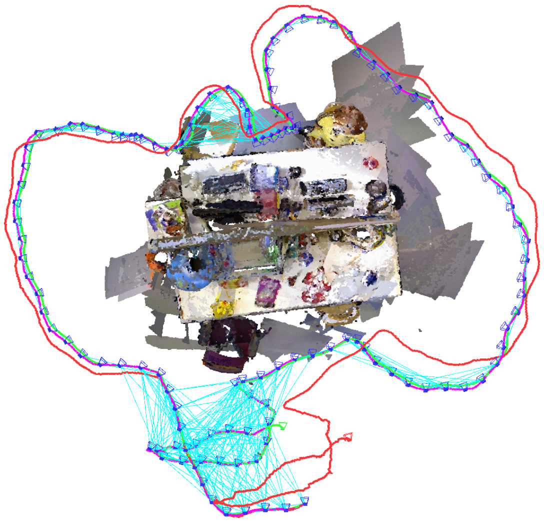

- Odometry method that uses both geometry and color information
- Uses a new distribution of the depth to improve the pose estimation

ElasticFusion: Dense SLAM Without A Pose Graph

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

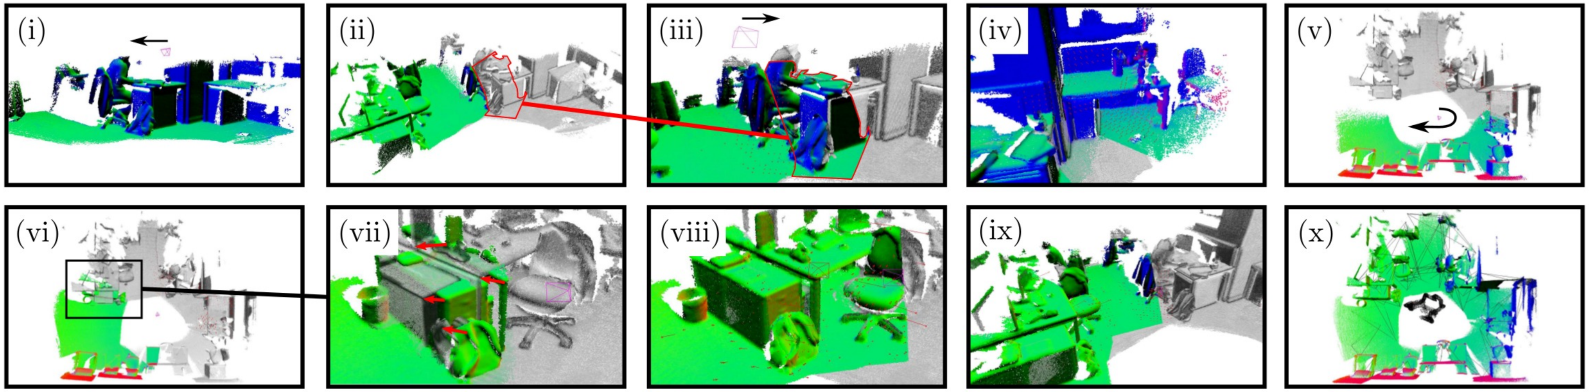

- Uses surfels instead of a TSDF to represent the 3D model
- First method to update the surface in online manner

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

Dai, Nießner, Zollhöfer, Izadi, Theobalt 2017

- Use all depth and color data to obtain consistent mapping
- Uses bundle adjustment to avoid drift and can recompute the surface during scan.

BAD SLAM: Bundle Adjusted Direct RGB-D SLAM

Schöps, Sattler, Pollefeys 2019

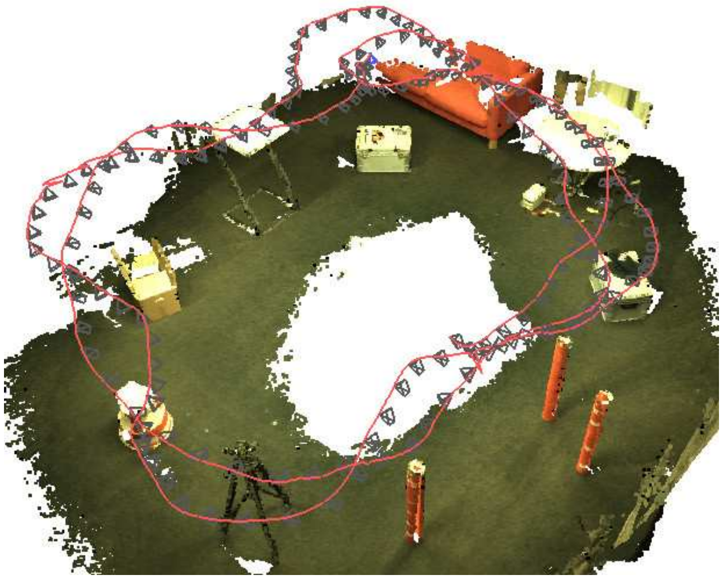

- Published on CVPR 2019 this year
- Perform bundle adjustment on surfels to get a high quality pose

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers: The Evolution of Motion Estimation and Real-time 3D Reconstruction
  - Depth Sensors
  - Monocular Cameras
  - Semantic Tracking and Mapping
  - Event-based Cameras
  - Learning-based Methods
- Possible Papers: An Overview of Methods for Accurate Geometry Reconstruction
  - Multi-view Reconstruction
  - Shading-based Approaches
  - Reconstruction from Point Clouds
  - Learning-based Geometry Estimation
- Questions

PTAM: Parallel Tracking and Mapping

Klein, Murray 2007

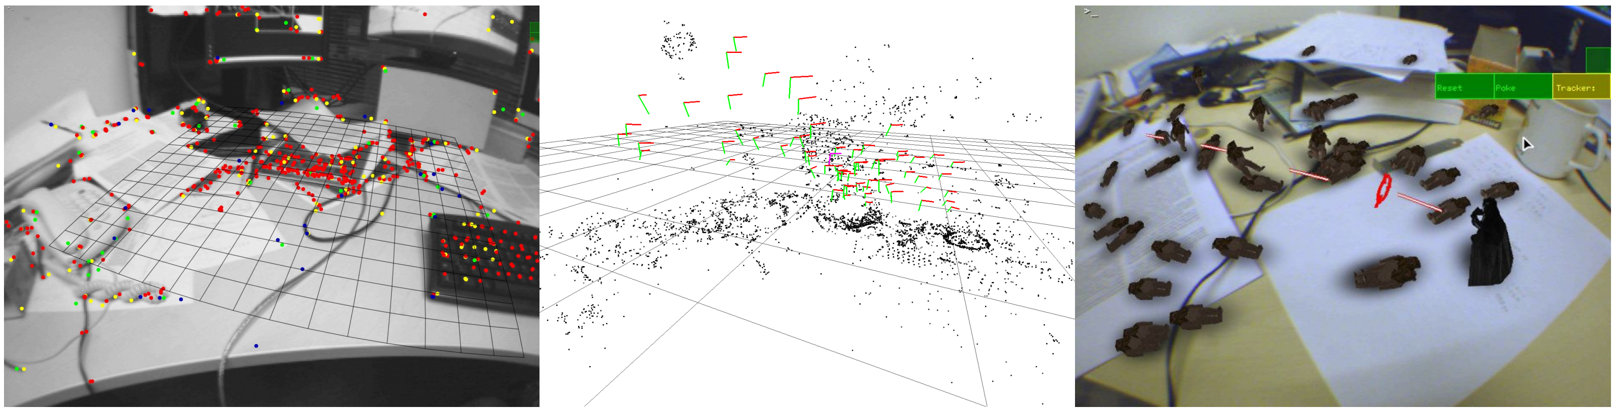

- One of the first systems capable of estimating both pose and geometry in real-time for handheld cameras
- Simple AR applications

DTAM: Dense Tracking and Mapping in Real-Time

Newcombe, Lovegrove, Davison 2011

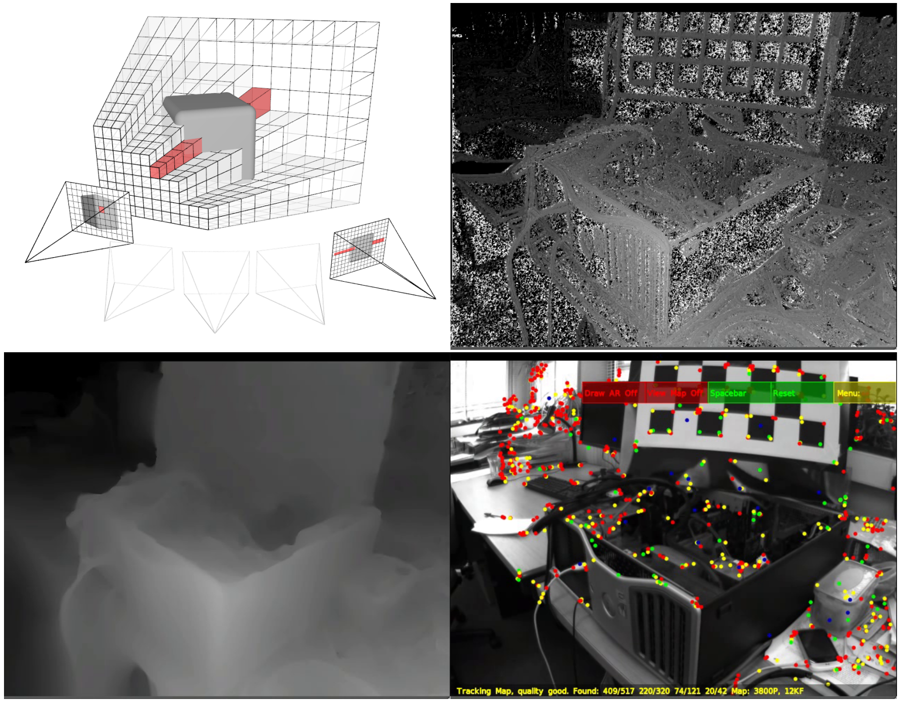

- One of the first monocular systems to create dense 3D models

ORB-SLAM: a Versatile and Accurate Monocular SLAM System

Mur-Artal, Montiel, Tardós 2015

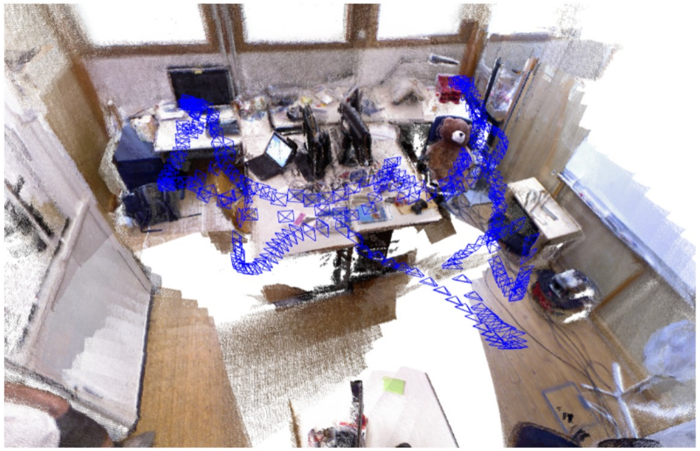

- Use all depth and color data to obtain consistent mapping

Direct Sparse Odometry

Engel, Koltun, Cremers 2016

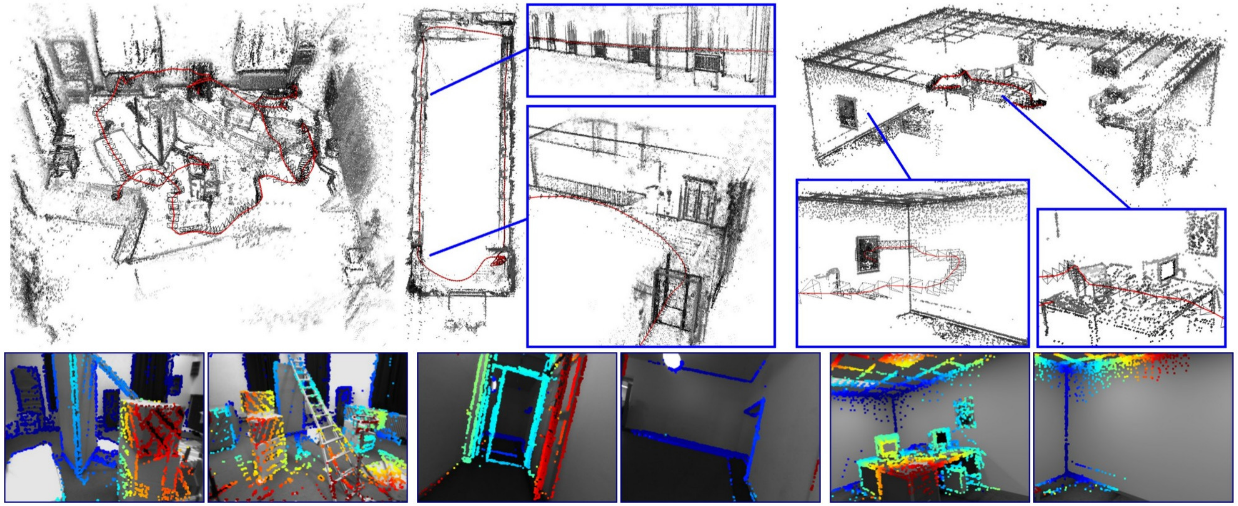

- Large-scale odometry
- Does not rely on keypoint detections

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers: The Evolution of Motion Estimation and Real-time 3D Reconstruction
  - Depth Sensors
  - Monocular Cameras
  - Semantic Tracking and Mapping
  - Event-based Cameras
  - Learning-based Methods
- Possible Papers: An Overview of Methods for Accurate Geometry Reconstruction
  - Multi-view Reconstruction
  - Shading-based Approaches
  - Reconstruction from Point Clouds
  - Learning-based Geometry Estimation
- Questions

SLAM++: Simultaneous Localisation and Mapping ad the Level of Objects

Salas-Moreno, Newcombe, Strasdat, Kelly, Davison 2013

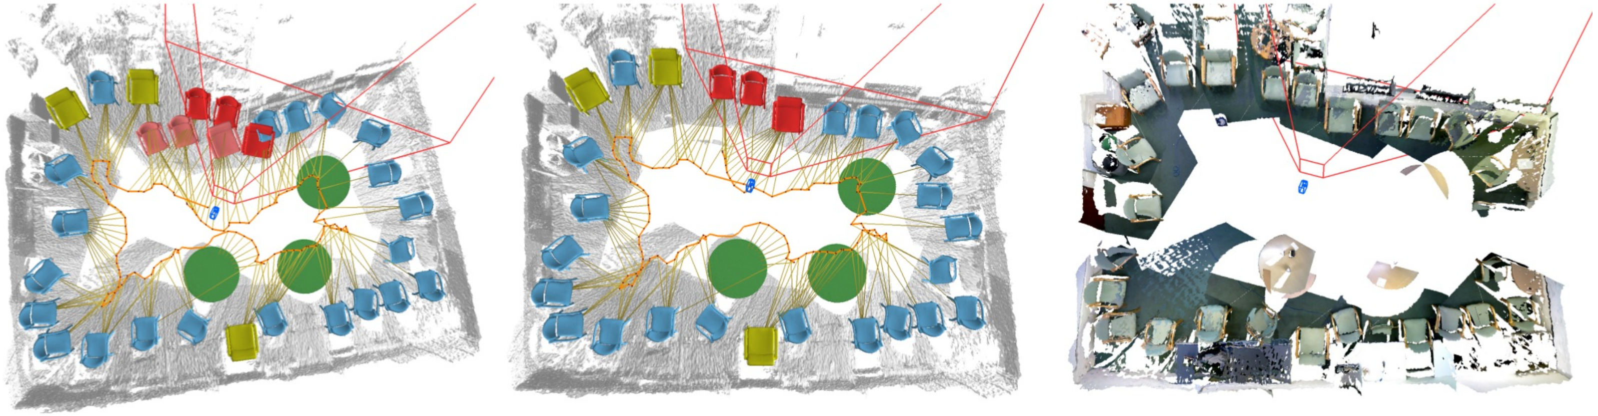

- Semantic 3D reconstruction
- Classifies different objects on the fly

Co-Fusion: Real-time Segmentation, Tracking and Fusion of Multiple Objects

Rünz, Agapito 2017

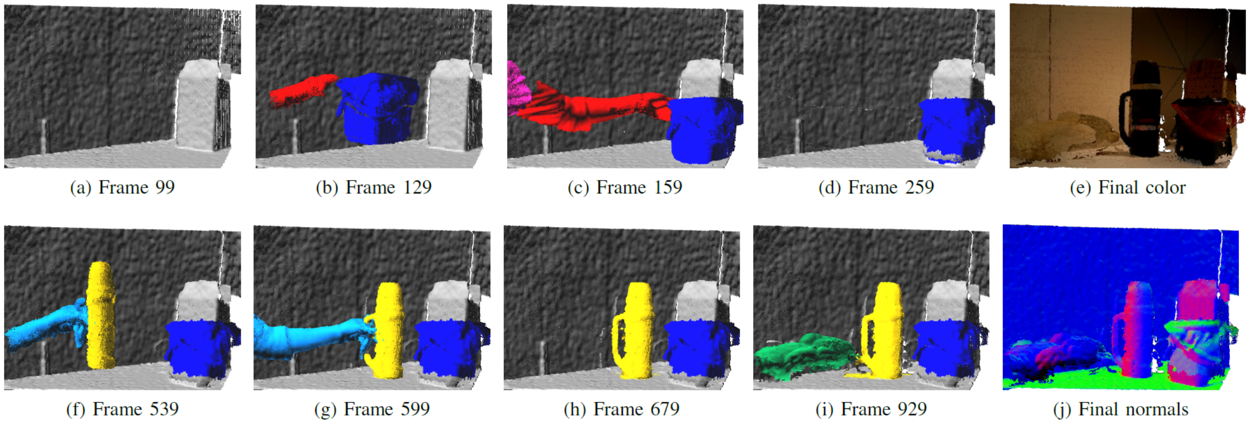

- Semantic 3D Reconstruction
- Allows objects to move

Fusion++: Volumetric Object-Level SLAM

McCormac, Clark, Bloesch, Davison, Leutenegger 2018

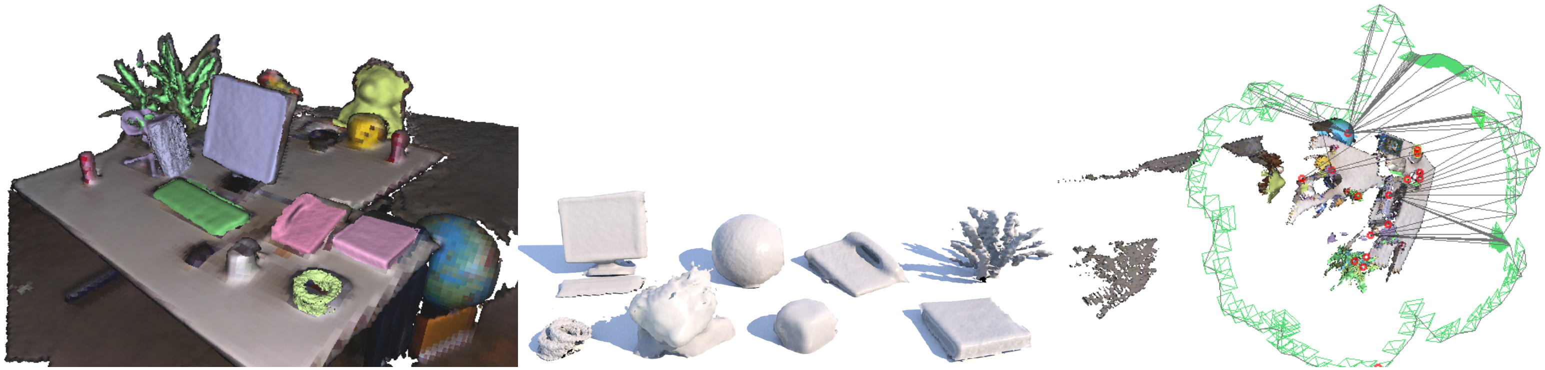

- Volumetric Object-Level Reconstructions
- Can handle dynamic motions

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers: The Evolution of Motion Estimation and Real-time 3D Reconstruction
  - Depth Sensors
  - Monocular Cameras
  - Semantic Tracking and Mapping
  - Event-based Cameras
  - Learning-based Methods
- Possible Papers: An Overview of Methods for Accurate Geometry Reconstruction
  - Multi-view Reconstruction
  - Shading-based Approaches
  - Reconstruction from Point Clouds
  - Learning-based Geometry Estimation
- Questions

Event-Based Visual Inertial Odometry

Zhu, Atanasov, Daniilidis 2017

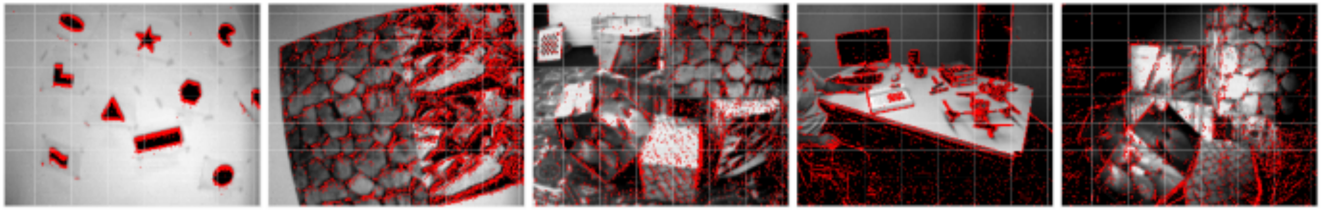

- Each pixel detects a change in the scene
- Much higher frame rate and no motion-blur
- Golf images obtained from https://www.prophesee.ai/2018/06/25/event-based-vision-2/

Semi-Dense 3D Reconstruction with a Stereo Event Camera

Zhou, Gallego, Rebecq, Kneip, Li, Scaramuzza 2018

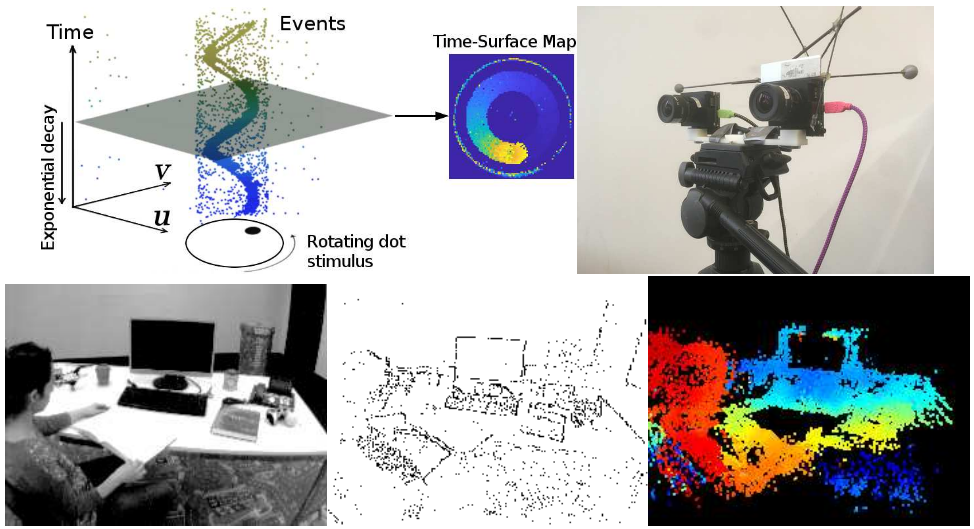

- Creates semi-dense depth images using an event-based stereo rig

Outline

DeepTAM: Deep Tracking and Mapping

Zhou, Ummenhofer, Brox 2018

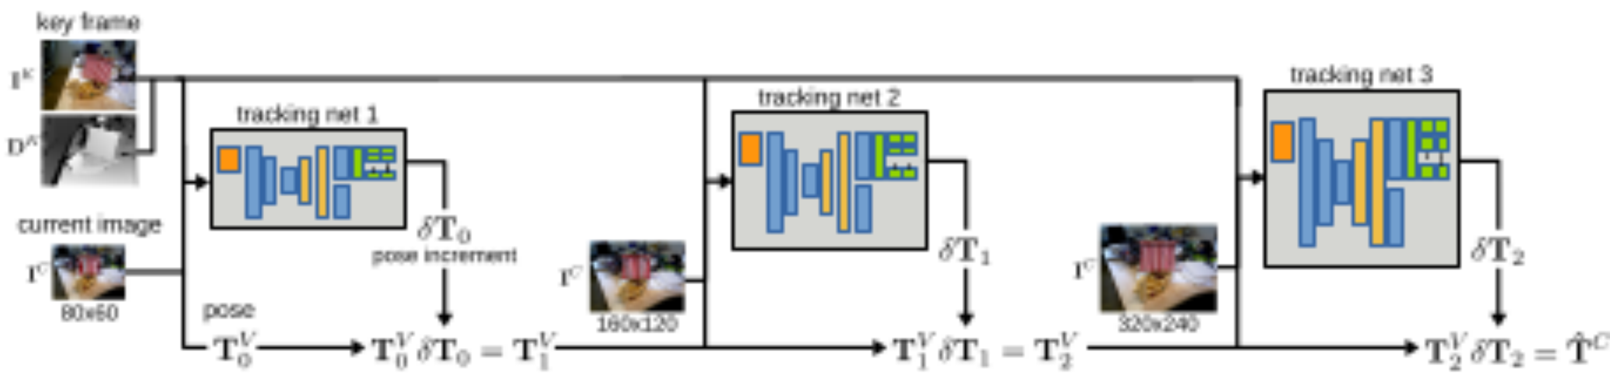

- Learn a network to predict the pose and generate depth images

Outline

Non-Sequential Structure from Motion

Enqvist, Kahl, Olsson 2011

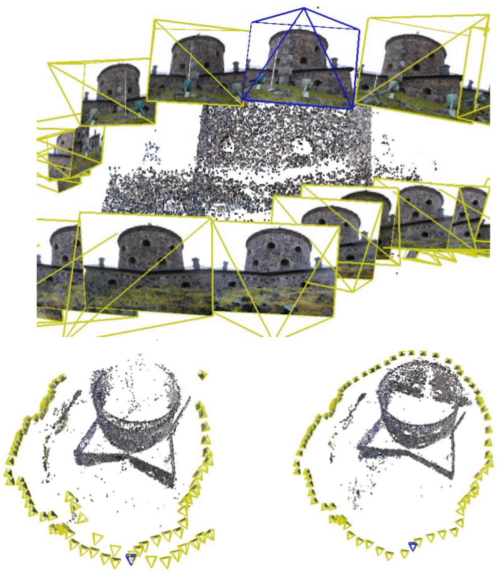

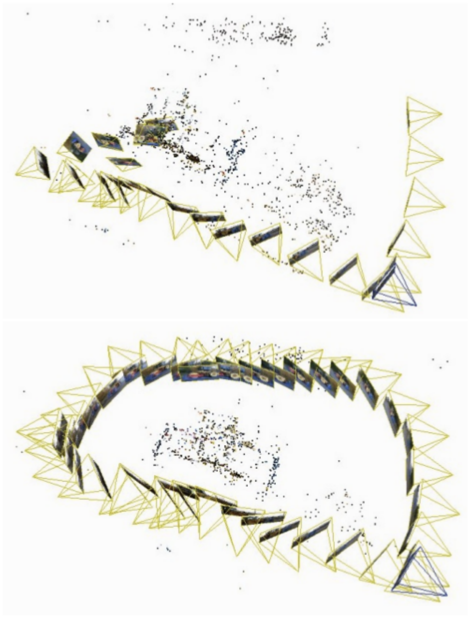

- Simultaneous consistent alignment of multiple RGB images
- Uses concepts from epipolar geometry, such as essential matrices

Color map optimization for 3D reconstruction with consumer depth cameras

Zhou, Koltun 2014

Input

Optimized reconstruction

- Optimize color using photoconsistency assumption
- Non-rigid correction

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers: The Evolution of Motion Estimation and Real-time 3D Reconstruction
  - Depth Sensors
  - Monocular Cameras
  - Semantic Tracking and Mapping
  - Event-based Cameras
  - Learning-based Methods
- Possible Papers: An Overview of Methods for Accurate Geometry Reconstruction
  - Multi-view Reconstruction
  - Shading-based Approaches
  - Reconstruction from Point Clouds
  - Learning-based Geometry Estimation
- Questions

Photometric Stereo with General, Unknown Lighting

Basri, Jacobs, Kemelmacher 2007

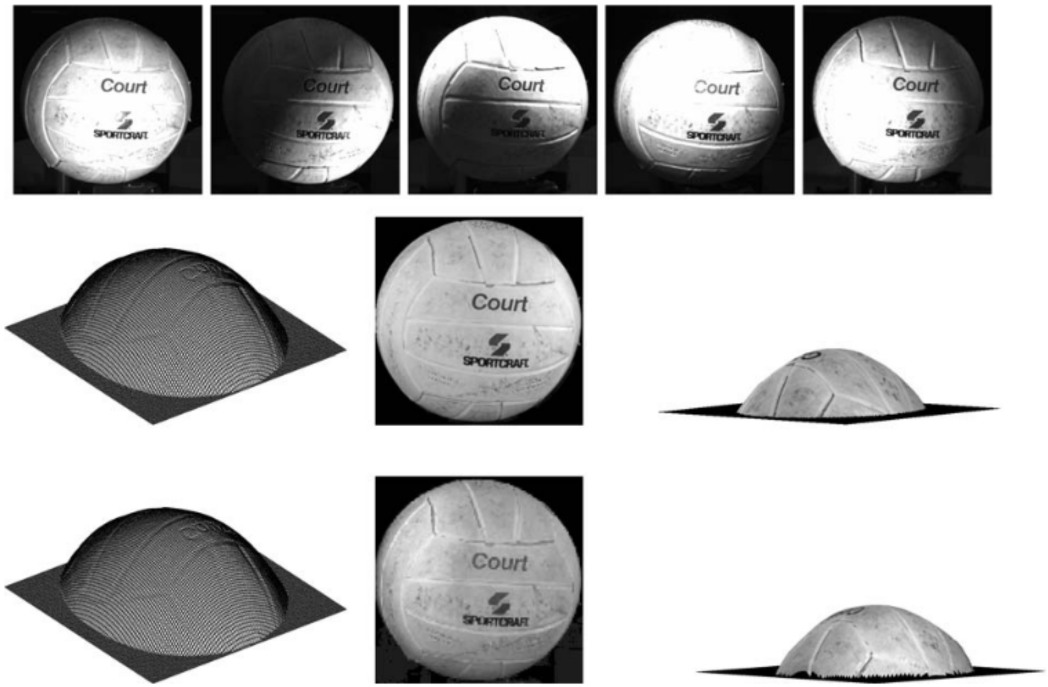

- Reconstruct shape from multiple images with different lighting

Shading-based Refinement on Volumetric Signed Distance Functions

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

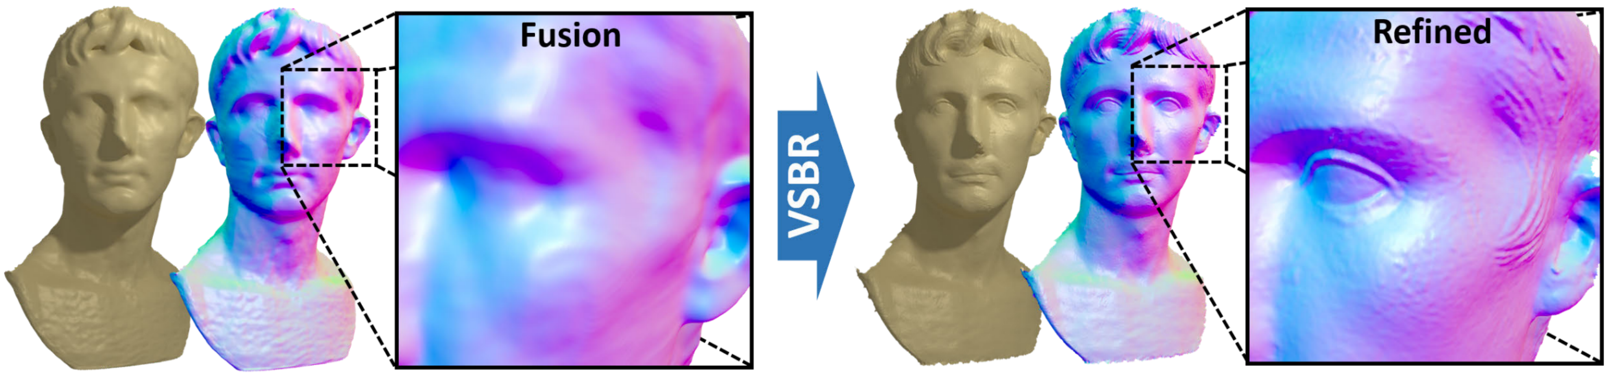

- Optimize geometry using shading (color) information
- Use signed distance functions to represent geometry

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

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

Haefner, Quéau, Möllenhoff, Cremers 2018

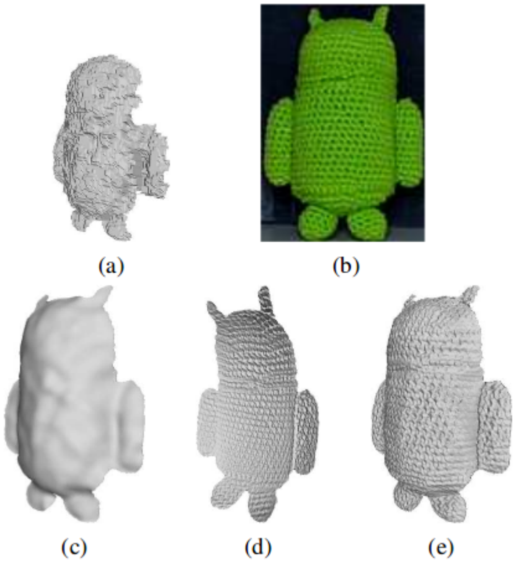

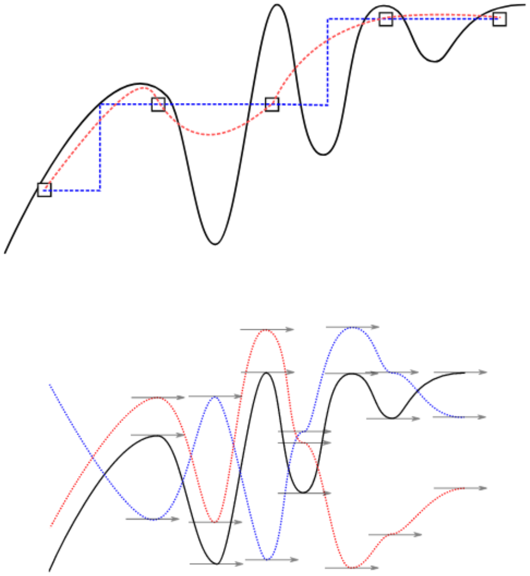

- Generate super-resolved depth image using shading clues
- Overcome shading ambiguity using depth clues

Combining Depth Fusion and Photometric Stereo for Fine-Detailed 3D Models

Bylow, Maier, Kahl, Olsson 2019

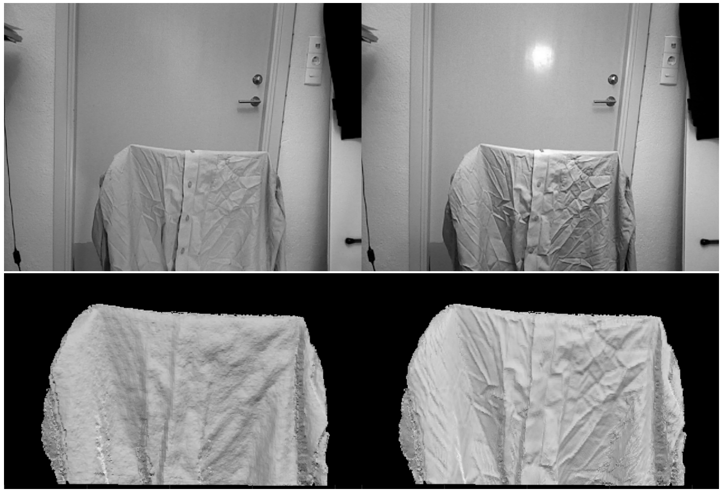

- Vary light source and move the depth sensor
- Refine the normals to get a higher accuracy

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers: The Evolution of Motion Estimation and Real-time 3D Reconstruction
  - Depth Sensors
  - Monocular Cameras
  - Semantic Tracking and Mapping
  - Event-based Cameras
  - Learning-based Methods
- Possible Papers: An Overview of Methods for Accurate Geometry Reconstruction
  - Multi-view Reconstruction
  - Shading-based Approaches
  - Reconstruction from Point Clouds
  - Learning-based Geometry Estimation
- Questions

Poisson surface reconstruction

Kazhdan, Bolitho, Hoppe 2006

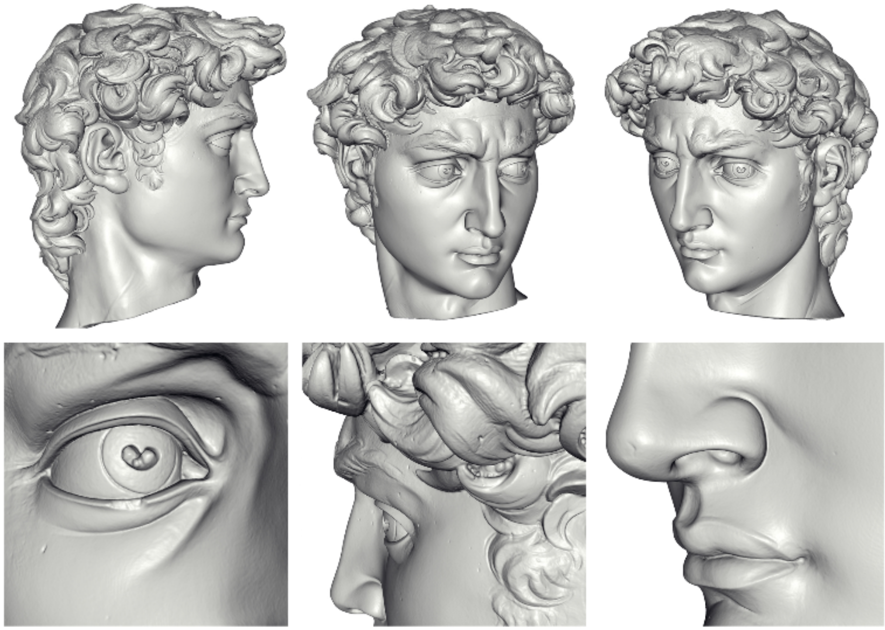

- Solve differential equation to obtain volumetric indicator function from point clouds

SSD: Smooth signed distance surface reconstruction

Calakli, Taubin 2011

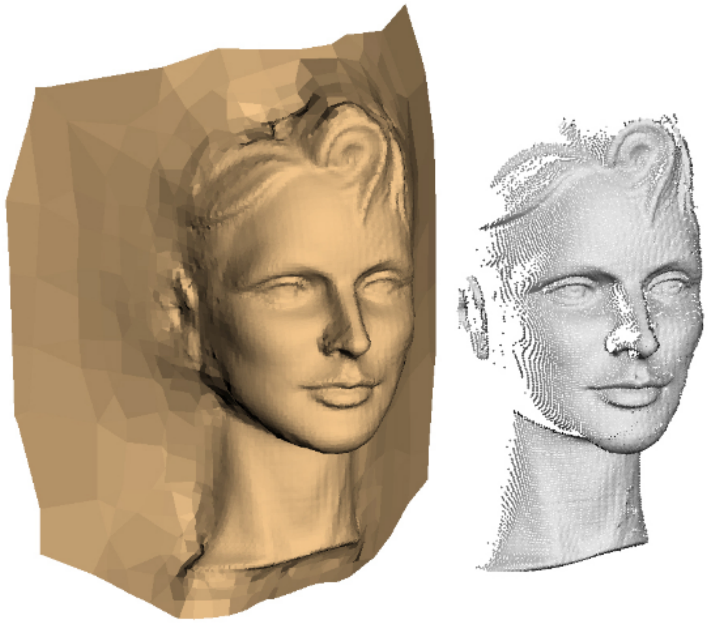

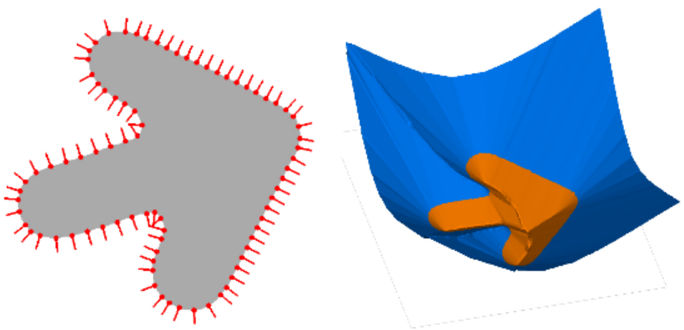

- Solve differential equation to obtain volumetric signed distance field from point clouds

Globfit: Consistently fitting primitives by discovering global relations

Li, Wu, Chrysathou, Sharf, Cohen-Or, Mitra 2011

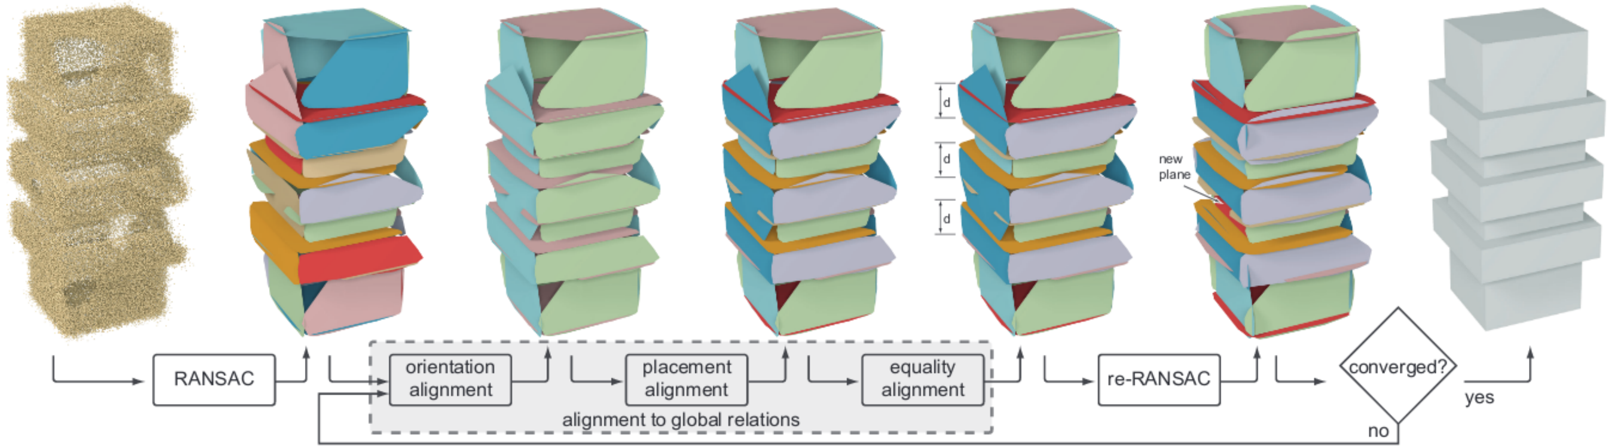

- Detect geometric primitives using RANSAC
- Align primitive orientation and location to get more consistent reconstruction

Robust Fitting of Subdivision Surfaces for Smooth Shape Analysis

Estellers, Schmidt, Cremers 2018

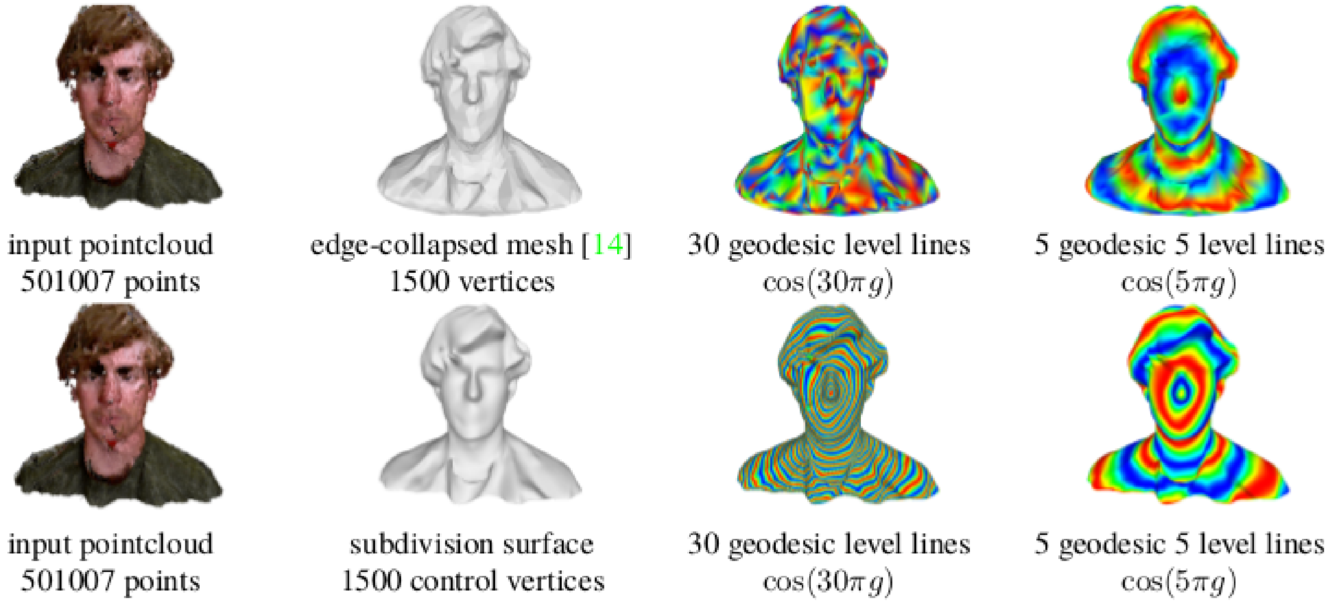

- Reconstruct surfaces using subdivision surface representation
- Can compute differential properties on the surface

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers: The Evolution of Motion Estimation and Real-time 3D Reconstruction
  - Depth Sensors
  - Monocular Cameras
  - Semantic Tracking and Mapping
  - Event-based Cameras
  - Learning-based Methods
- Possible Papers: An Overview of Methods for Accurate Geometry Reconstruction
  - Multi-view Reconstruction
  - Shading-based Approaches
  - Reconstruction from Point Clouds
  - Learning-based Geometry Estimation
- Questions

Shape from Shading through Shape Evolution

Yang, Deng 2018

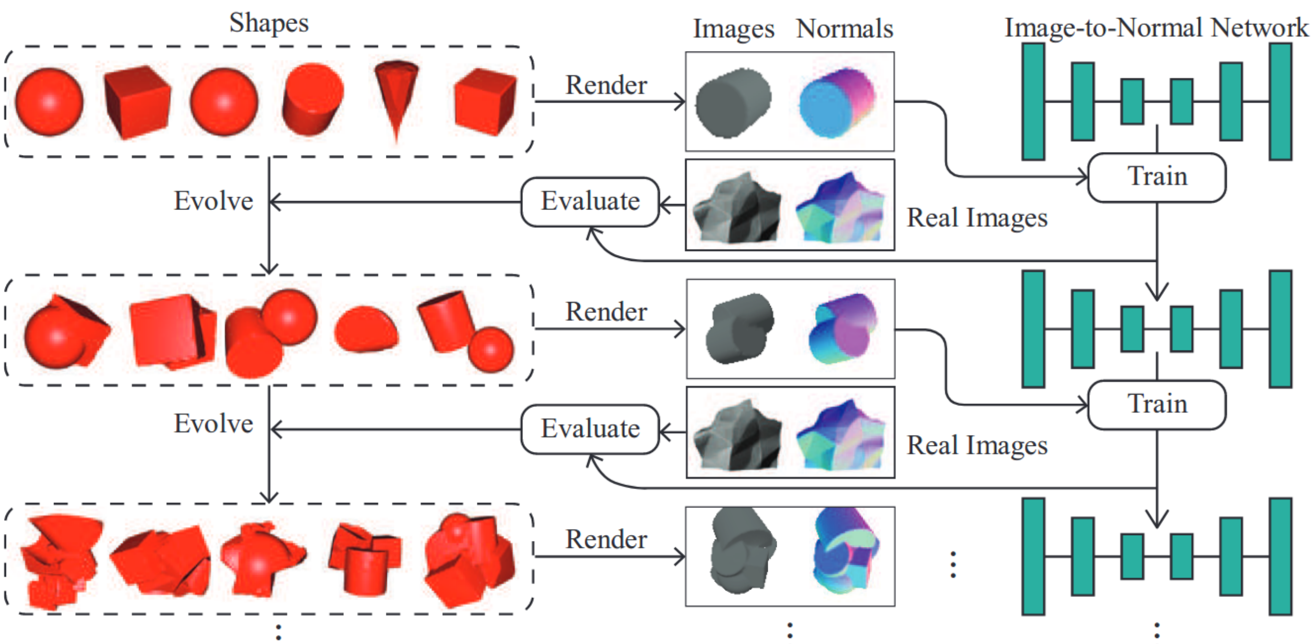

- Generate synthetic shapes by primitive evolution
- Use evolved shape to train network for shape-from-shading

Learning to Reconstruct Shape and Spatially-Varying Reflectance from a Single Image

Li, Xu, Ramamoorthi, Sunkavalli, Chandraker 2018

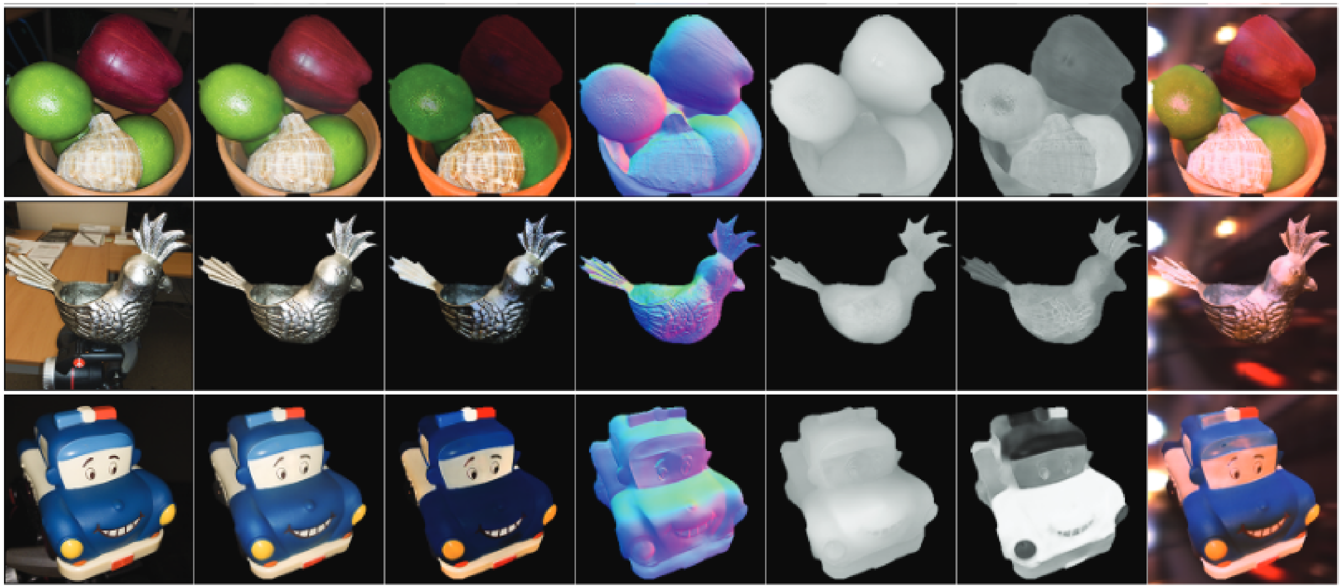

- Inverse rendering using neural network architecture
- Only use a single image to infer object properties

Questions?

Reminder:

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

- Web page: `vision.in.tum.de/teaching/ws2019/seminar_realtime3d`
- Password: `realtime3d-ws19`
- Contact: `realtime3d-ws19@vision.in.tum.de`

An Overview of Methods for Accurate Geometry Reconstruction

- Web page: `vision.in.tum.de/teaching/ws2019/seminar_3dgeometry`
- Password: `3dgeometry-ws19`
- Contact: `3dgeometry-ws19@vision.in.tum.de`