Seminar: Recent Advances in 3D Computer Vision

Christiane Sommer
Computer Vision Group
Technical University of Munich

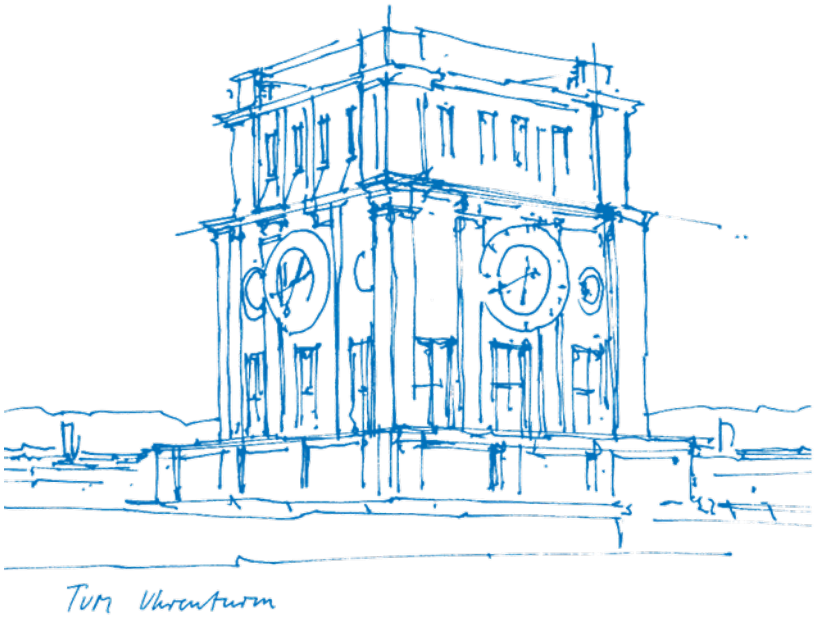

How can I access these slides?

- Option 1 (preferred): seminar web page
  - `vision.in.tum.de/teaching/ws2018/seminar_3dcv`
  - Password for material page: `3dcv_ws2018`
  - Material page will go online June 26, 2018
- Option 2: contact organizer
  - Christiane Sommer, `sommerc@in.tum.de`
  - Only use this option if you forgot password

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - Static (RGB-)D scanning
  - RGB-D scan refinement (color/geometry)
  - Dynamic RGB-D scanning
  - Semantic monocular scanning
  - Scan completion
  - Learning 3D descriptors
- Questions

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - Static (RGB-)D scanning
  - RGB-D scan refinement (color/geometry)
  - Dynamic RGB-D scanning
  - Semantic monocular scanning
  - Scan completion
  - Learning 3D descriptors
- Questions

How is the seminar organized?

- Seminar meetings: talks and discussion
  - Time: Thursdays, 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** talk: meet supervisor to clarify questions
  - **One week before** talk: meet supervisor to go through slides
  - **Two weeks after** talk: submit your report via email

What are the requirements for the talk?

- General set-up:
  - Duration: 25-30 minutes talk + 10-15 minutes discussion
  - Make sure to finish on time!
  - Rule of thumb: 1-2 minutes per slide $\rightarrow$ 15-30 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 final report?

- General set-up:
  - Use LaTeX template provided on web page
  - Length: 6-10 pages
  - Send final report as pdf to supervisor by email
  - 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
  - Static (RGB-)D scanning
  - RGB-D scan refinement (color/geometry)
  - Dynamic RGB-D scanning
  - Semantic monocular scanning
  - Scan completion
  - Learning 3D descriptors
- 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 containing these three favorites to `sommerc@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: July 4, 2018

How do you register for the seminar?

Example registration email:

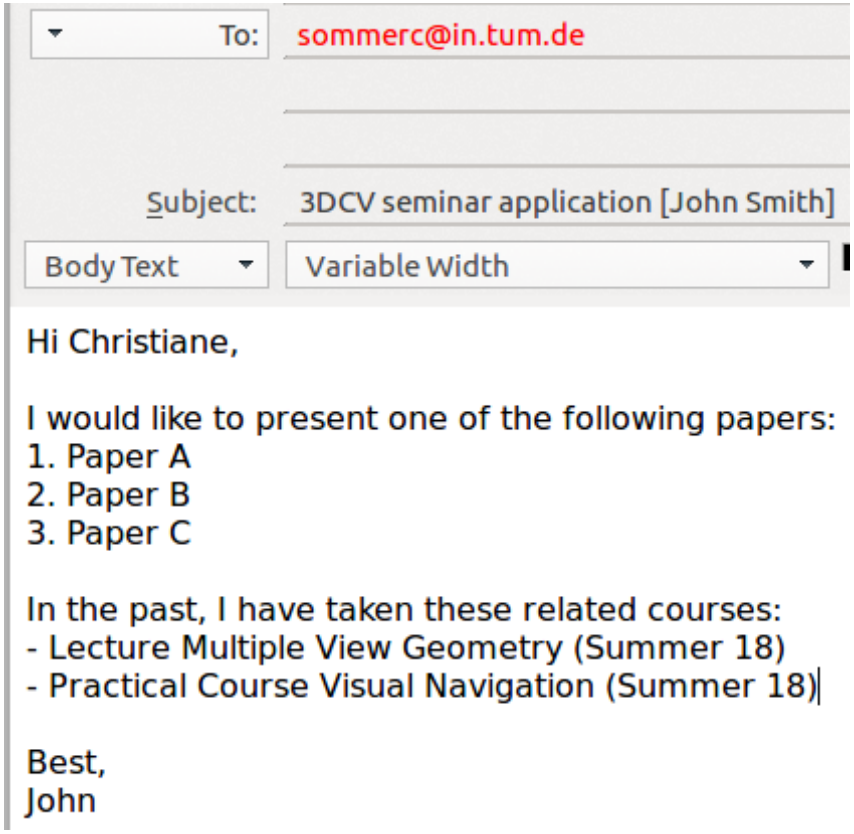

How do we select candidates and assign papers?

- 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 (July 12, 2018)
- 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
  - Static (RGB-)D scanning
  - RGB-D scan refinement (color/geometry)
  - Dynamic RGB-D scanning
  - Semantic monocular scanning
  - Scan completion
  - Learning 3D descriptors
- Questions

Real-time 3D reconstruction at scale using voxel hashing

Nießner et al. 2013

- efficient storage of volumetric SDF grid using hash table
- track camera using depth values only

Large-Scale Multi-Resolution Surface Reconstruction from RGB-D Sequences

Steinbrücker et al. 2013

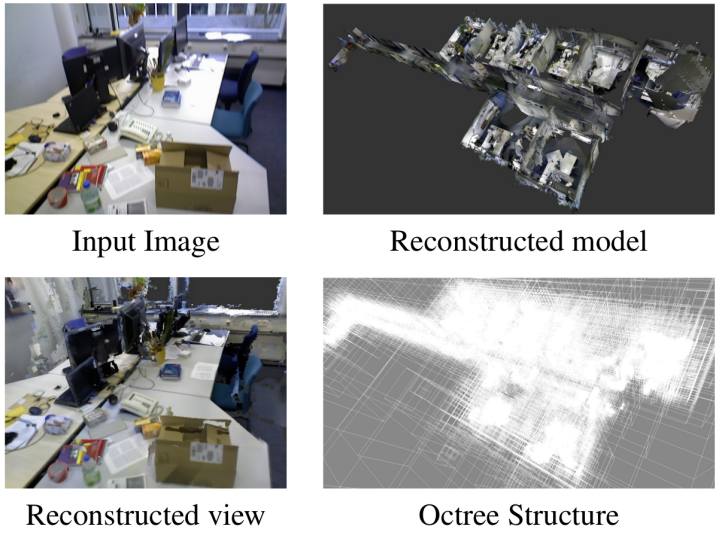

- efficient storage of volumetric SDF grid using octree
- track camera using direct RGB-D image alignment

Robust Reconstruction of Indoor Scenes

Choi, Zhou, Koltun 2015

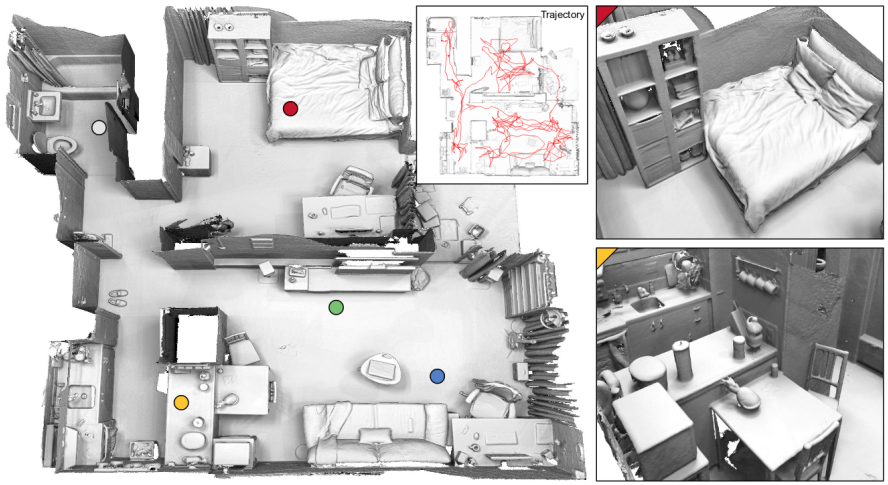

- register scans based on geometric information
- globally optimize using line processes

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

Dai et al. 2016

- use all depth and color data to obtain consistent mapping

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - Static (RGB-)D scanning
  - RGB-D scan refinement (color/geometry)
  - Dynamic RGB-D scanning
  - Semantic monocular scanning
  - Scan completion
  - Learning 3D descriptors
- Questions

Color map optimization for 3D reconstruction with consumer depth cameras

Zhou, Koltun 2014

Input

Optimized reconstruction

- optimize color using photoconsistency assumption
- non-rigid correction

Shading-based Refinement on Volumetric Signed Distance Functions

Zollhöfer et al. 2014

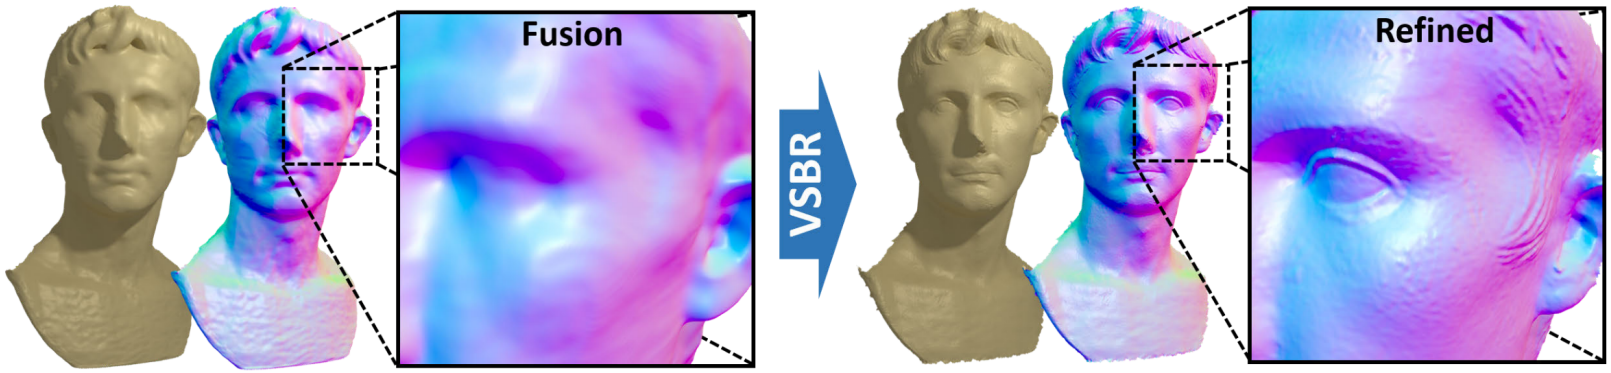

- optimize geometry using shading (color) information
- use signed distance functions to represent geometry

Intrinsic3D: High-Quality 3D Reconstruction by Joint Appearance and Geometry Optimization with Spatially-Varying Lighting

Maier et al. 2017

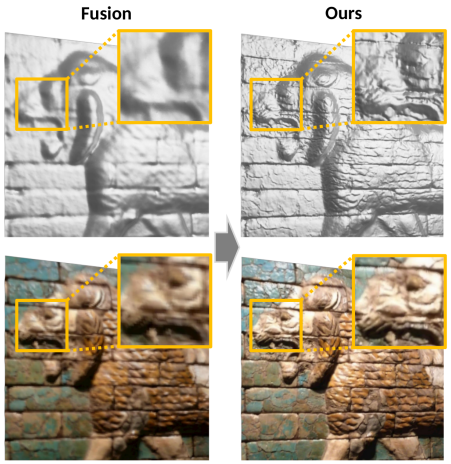

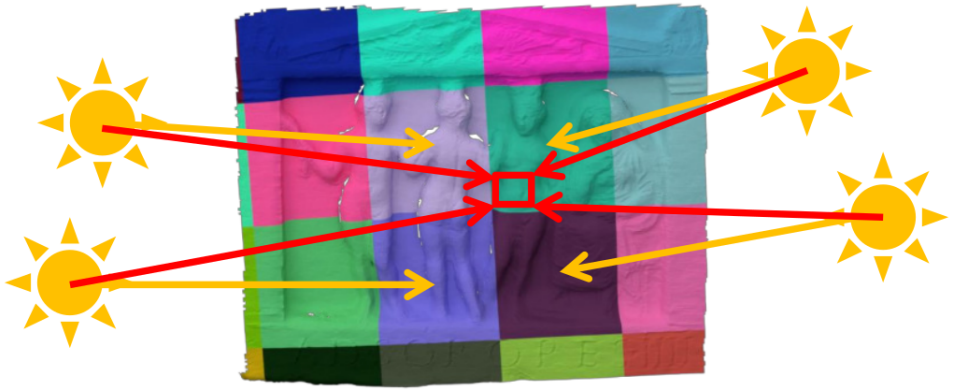

- optimize geometry and color using shading information
- allow for spatially-varying lighting

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - Static (RGB-)D scanning
  - RGB-D scan refinement (color/geometry)
  - Dynamic RGB-D scanning
  - Semantic monocular scanning
  - Scan completion
  - Learning 3D descriptors
- Questions

Fusion4D: real-time performance capture of challenging scenes

Dou et al. 2016

- dynamic (non-rigid) scanning in multi-view set-up

VolumeDeform: Real-time Volumetric Non-rigid Reconstruction

Innmann et al. 2016

- dynamic (non-rigid) scanning with a single camera

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - Static (RGB-)D scanning
  - RGB-D scan refinement (color/geometry)
  - Dynamic RGB-D scanning
  - Semantic monocular scanning
  - Scan completion
  - Learning 3D descriptors
- Questions

Joint Detection, Tracking and Mapping by Semantic Bundle Adjustment

Fioraio, Di Stefano 2013

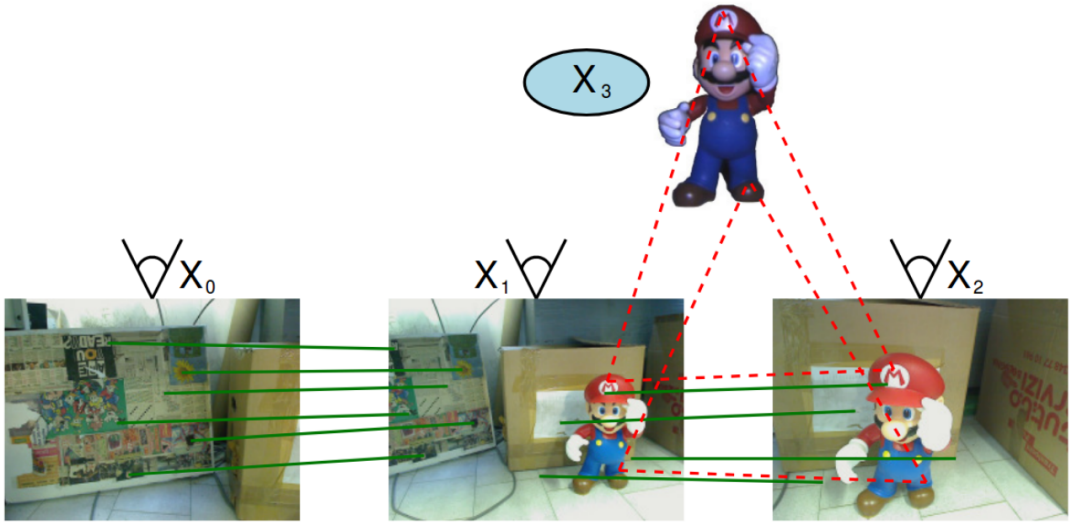

- use rigid bodies as features for camera tracking

Real-time monocular object SLAM

Gálvez-López et al. 2016

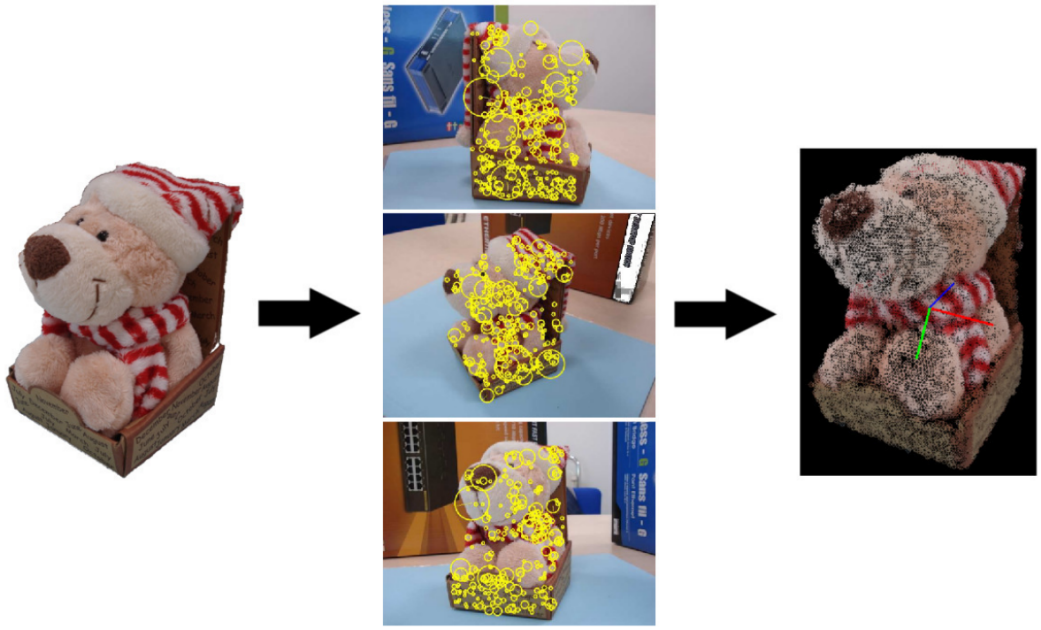

- large database of objects for semantic SLAM

Probabilistic data association for semantic SLAM

Bowman et al. 2017

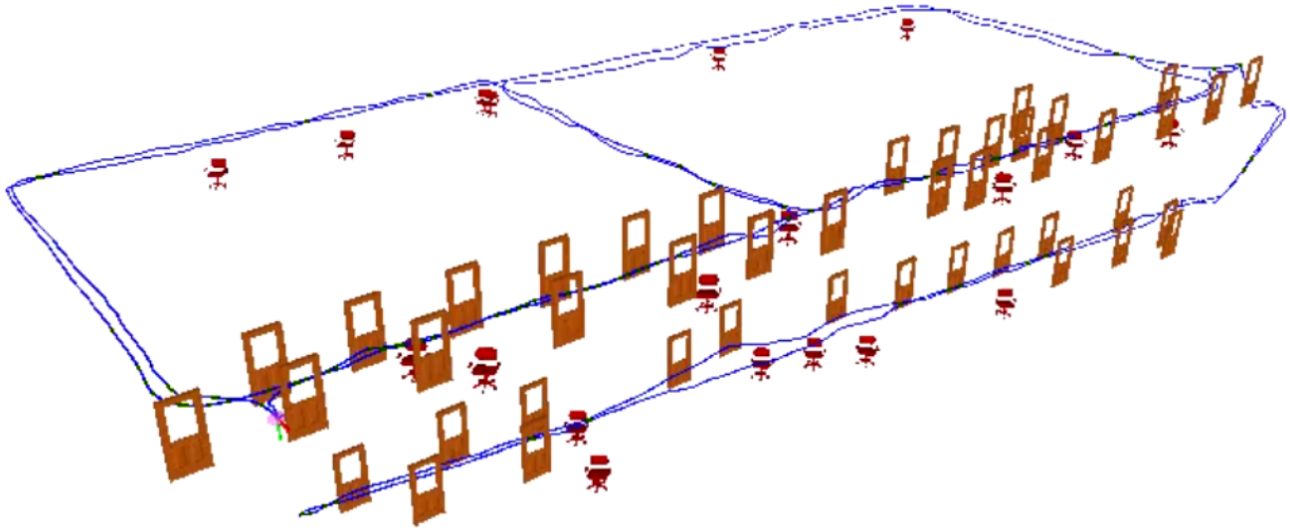

- focus on data association (which object observations correspond to same object)

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - Static (RGB-)D scanning
  - RGB-D scan refinement (color/geometry)
  - Dynamic RGB-D scanning
  - Semantic monocular scanning
  - Scan completion
  - Learning 3D descriptors
- Questions

Shape Completion using 3D-Encoder-Predictor CNNs and Shape Synthesis

Dai, Qi, Nießner 2017

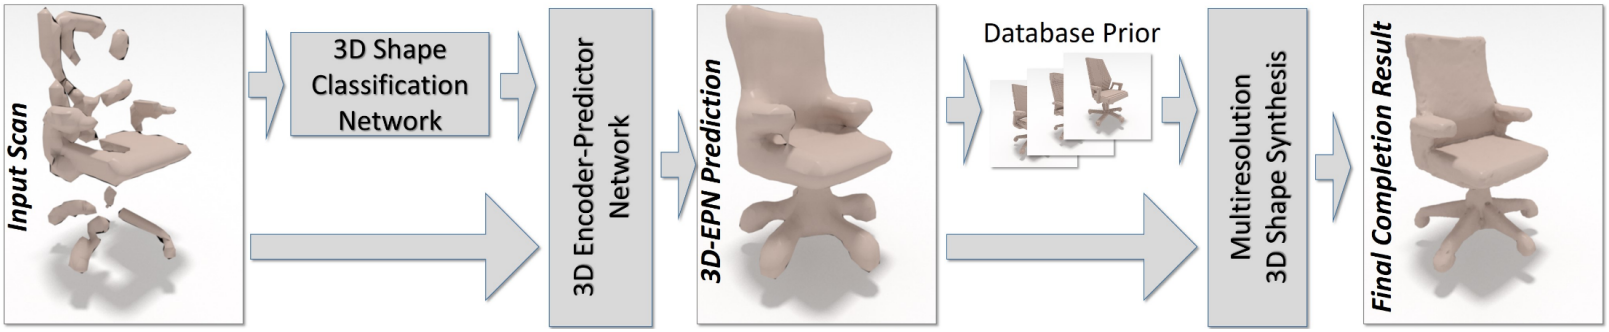

- use partial scans as input
- complete shape using encoder-predictor network

ScanComplete: Large-Scale Scene Completion and Semantic Segmentation for 3D Scans

Dai et al. 2018

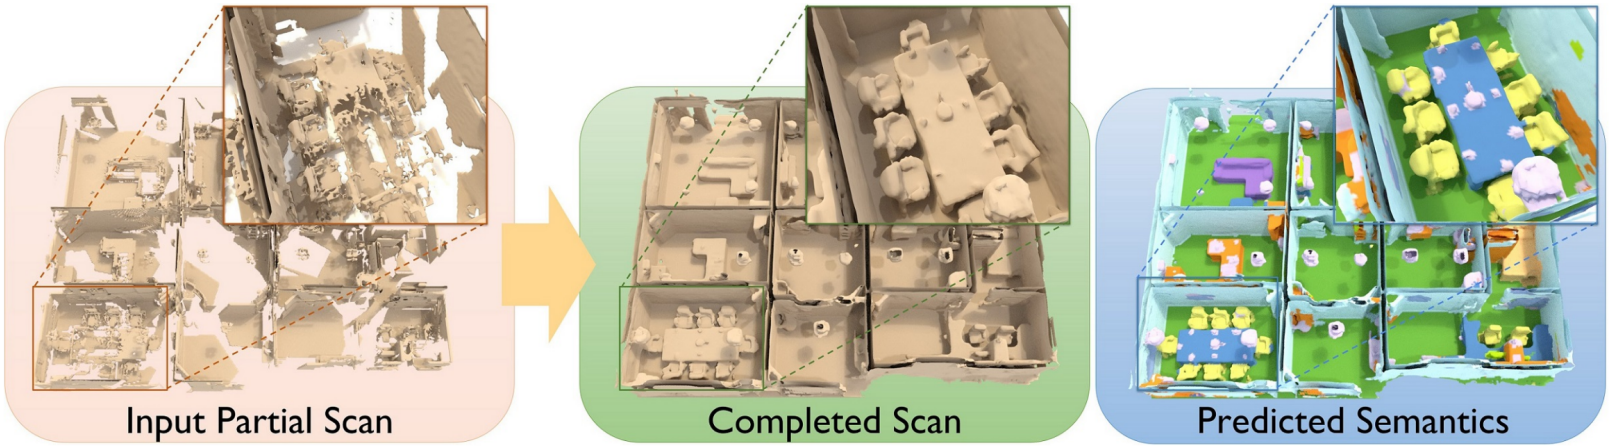

- scan completion for whole scenes

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - Static (RGB-)D scanning
  - RGB-D scan refinement (color/geometry)
  - Dynamic RGB-D scanning
  - Semantic monocular scanning
  - Scan completion
  - Learning 3D descriptors
- Questions

3DMatch: Learning Local Geometric Descriptors from RGB-D Reconstructions

Zeng et al. 2017

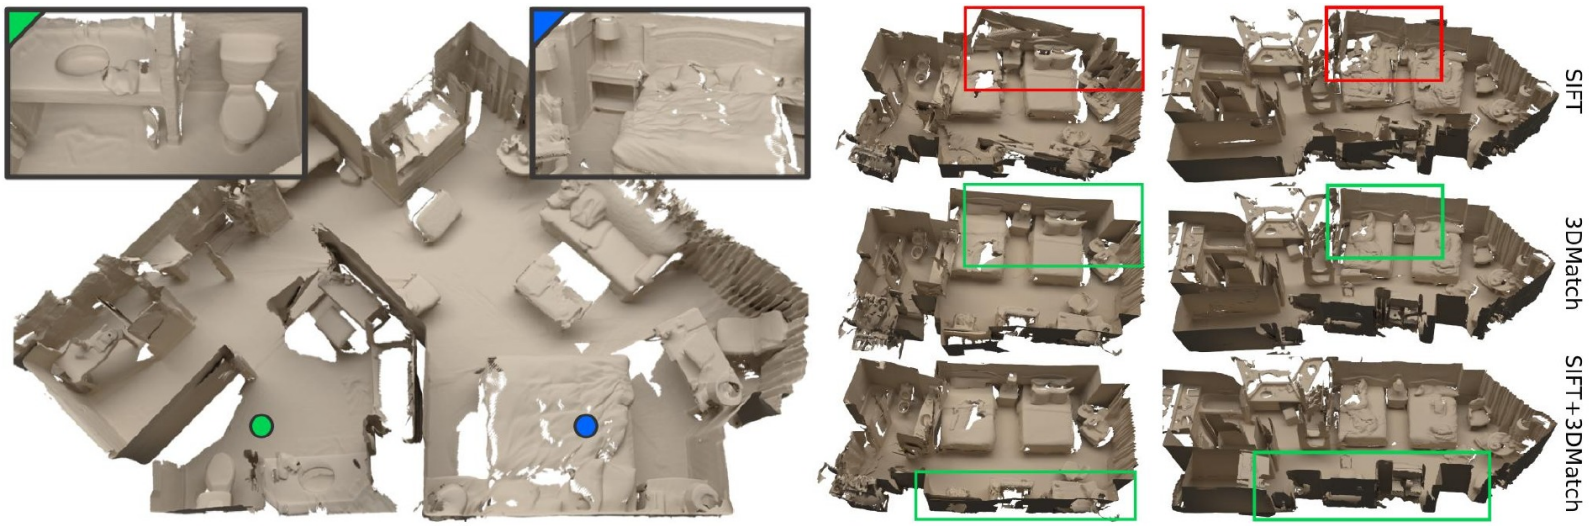

- represent volume as truncated distance function
- extract descriptors for local patches

Learning Compact Geometric Features

Khoury, Zhou, Koltun 2017

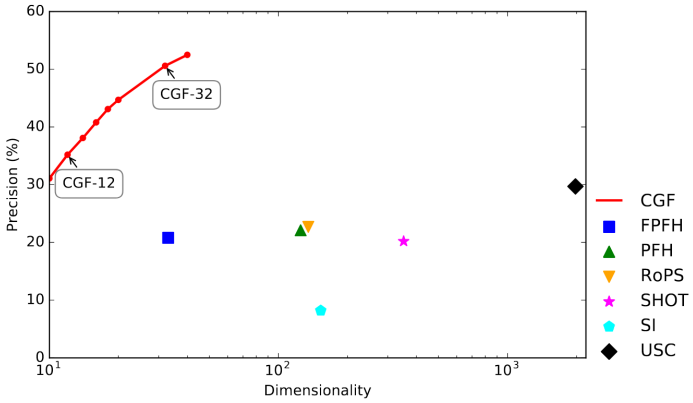

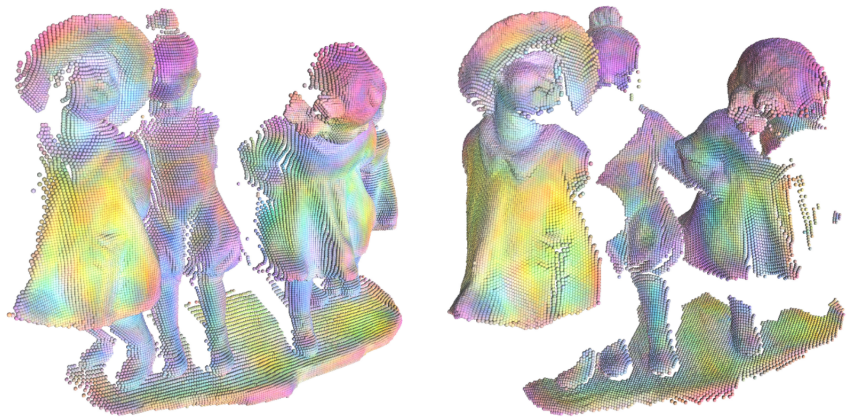

- extract descriptors for points in point cloud
- use histogram of points as network input

PPFNet: Global Context Aware Local Features for Robust 3D Point Matching

Deng, Birdal, Ilic 2018

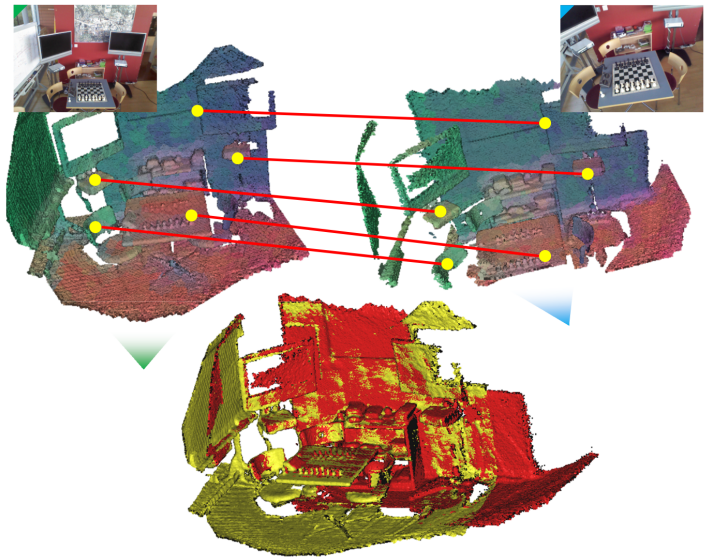

- use point pair features (PPFs) as network input

Questions?

Reminder:

- Web page: `vision.in.tum.de/teaching/ws2018/seminar_3dcv`
- Password: `3dcv_ws2018`
- Contact: Christiane Sommer, `sommerc@in.tum.de`