Seminar: Recent Advances in 3D Computer Vision

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

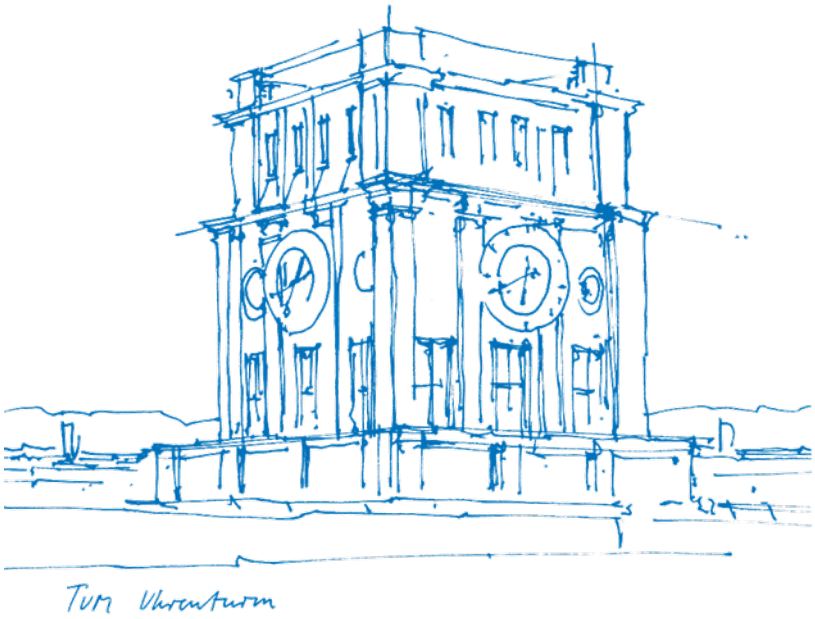

How can I access these slides?

- Option 1 (preferred): seminar web page
  - `https://vision.in.tum.de/teaching/ss2021/seminar_3dcv`
  - Password for material page: `ss21-3dcv`
  - Material page will go online after this pre-meeting
- Option 2: contact organizers
  - `3dcv-ss21@vision.in.tum.de`
  - **Only use this option if you forgot the password**
  - Do not write to our personal email address

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - 3D Scene Representation using Neural Networks
  - Photometric reconstruction
  - RGB-D scanning and refinement
- Questions

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - 3D Scene Representation using Neural Networks
  - Photometric reconstruction
  - RGB-D scanning and refinement
- Questions

How is the seminar organized?

- Seminar meetings: talks and discussion
  - Days: Wednesday, April 7th and Thursday, April 8th
  - Time: 9:00 - 17:00
  - Location: virtual (on BBB)
  - 6 talks per day
  - **Attendance is mandatory!** If you do not have time on one of the two days, do not register!

    (Exception: exams)
- Talk preparation / contact with supervisor
  - Read through your paper yourself and write down what you do not understand
  - Approx. **one month before** talk: meet supervisor for questions (optional, but recommended)
  - **One week before** talk: meet supervisor to go through slides (optional, but recommended)
  - **April 2nd**: submit your slides via submission system (mandatory)
  - **April 23rd**: submit your report via submission system (mandatory)

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 (without bibliography)
  - Upload final report as pdf via submission system
  - Submission deadline: **April 23rd**
- Recommended structure (main text only):
  - Introduction
  - Method description
  - Experiments and results
  - Discussion of results
  - Summary

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - 3D Scene Representation using Neural Networks
  - Photometric reconstruction
  - RGB-D scanning and refinement
- Questions

How do you register for the seminar?

- Step 1: Official registration via TUM matching system
  - Go to `https://matching.in.tum.de`
  - Register for seminar with the title *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 the seminar email address
  - Email subject: "[3DCV] 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 16th, 2021

How do you register for the seminar?

Example registration email:

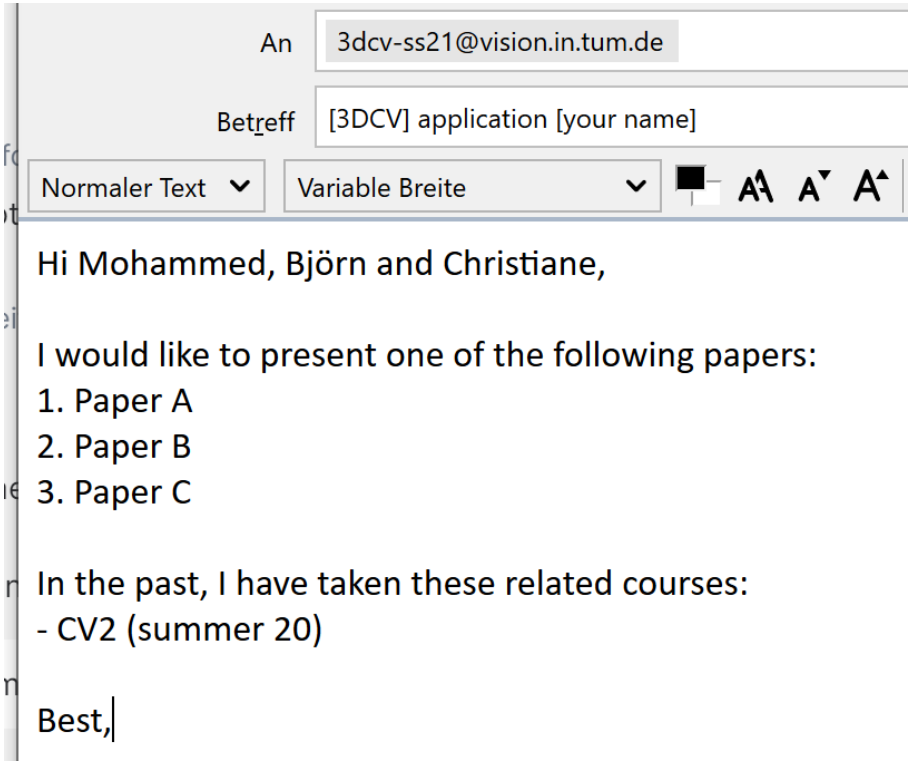

How do we select candidates and assign papers?

- 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
  - Note that if you have not taken any related course, you must be willing to invest a lot of work to learn the required basics
  - You will get notified by the matching system about the decision (February 25th, 2021)
- 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
  - 3D Scene Representation using Neural Networks
  - Photometric reconstruction
  - RGB-D scanning and refinement
- Questions

Occupancy Networks: Learning 3D Reconstruction in Function Space

Mescheder, Oechsle, Niemeyer, Nowozin, Geiger 2019

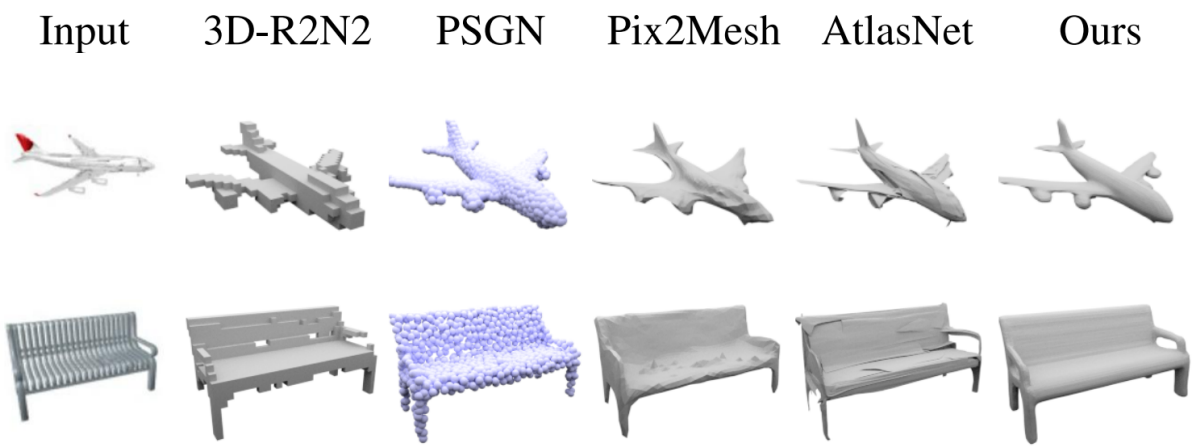

- Train a NN to predict the occupancy probability for a point in space and extract the isosurface
- The above representation can be condition on an additional input which allows for: Single Image 3D Reconstruction, Super Resolution, etc.

Implicit geometric regularization for learning shapes

Gropp, Yariv, Haim, Atzmon, Lipman 2020

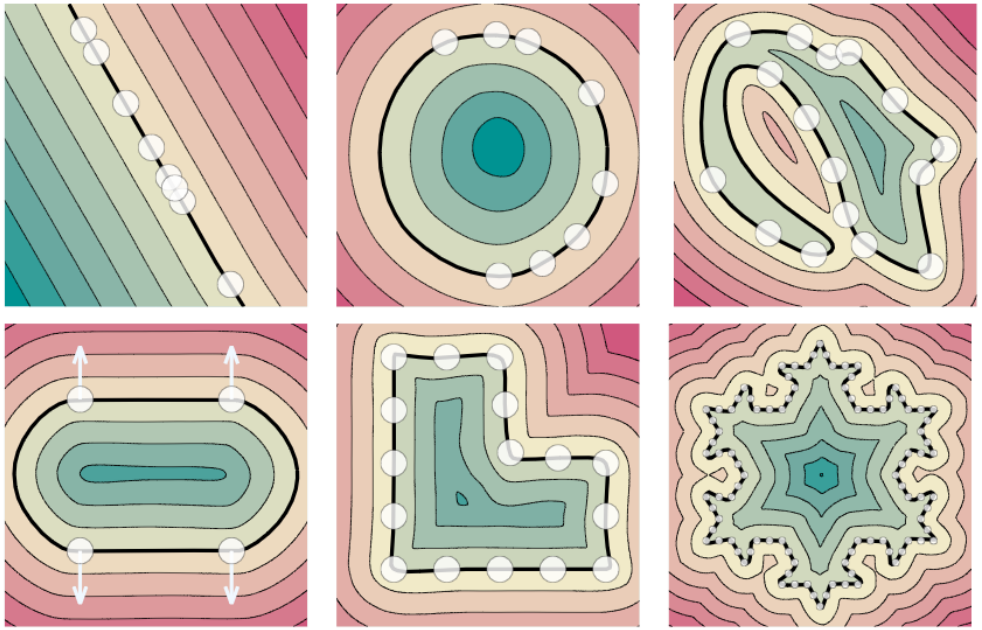

- Train a NN to predict the SDF for a point in space and extract the isosurface
- Use norm of gradient as a regularizer

Implicit neural representations with periodic activation functions

Sitzmann, Martel, Bergman, Lindell, Wetzstein 2020

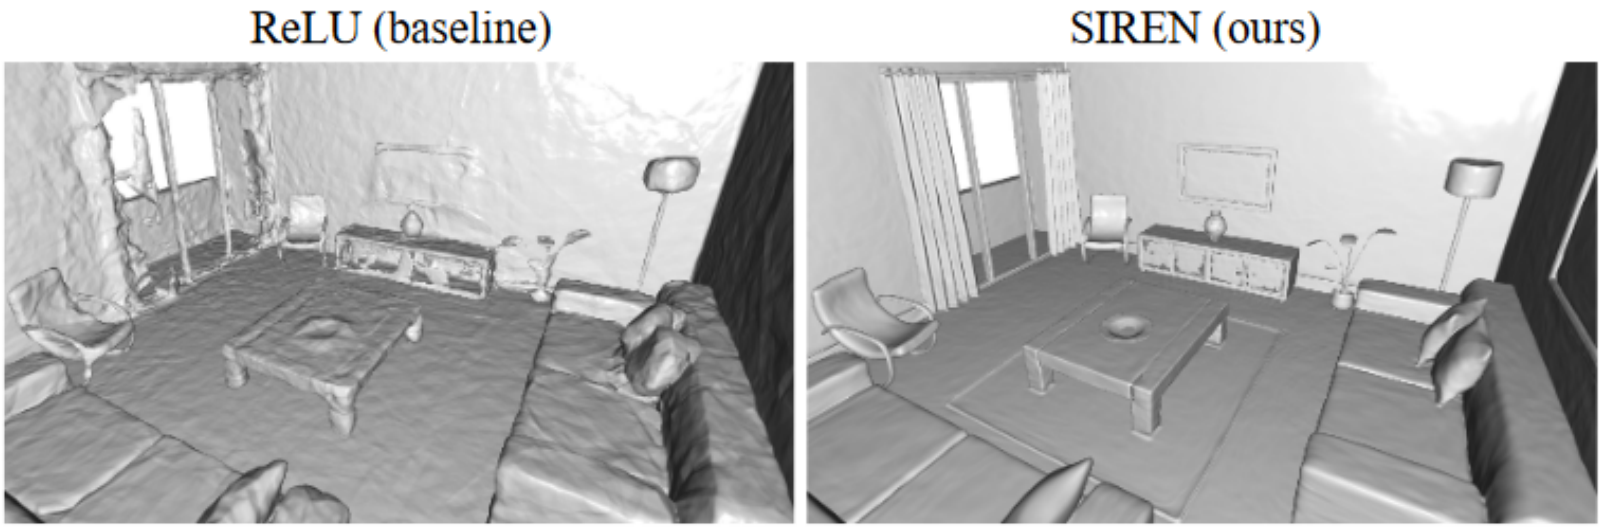

- Focus on activation functions that enable high resolution

NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis

Mildenhall et al. 2020

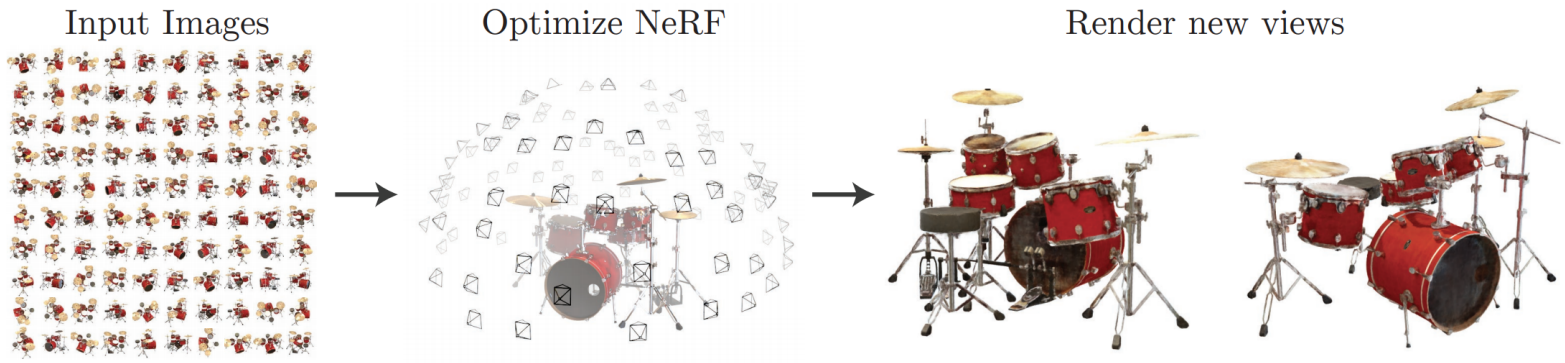

- For a set of images of one scene train one NN that maps $(x, y, z, \theta, \phi) \mapsto (RGB\sigma)$ using volume rendering and image based losses
- The NN can be used to generate novel views and gives very impressive results

MeshSDF: Differentiable Iso-Surface Extraction

Remelli et al. 2020, NeurIPS

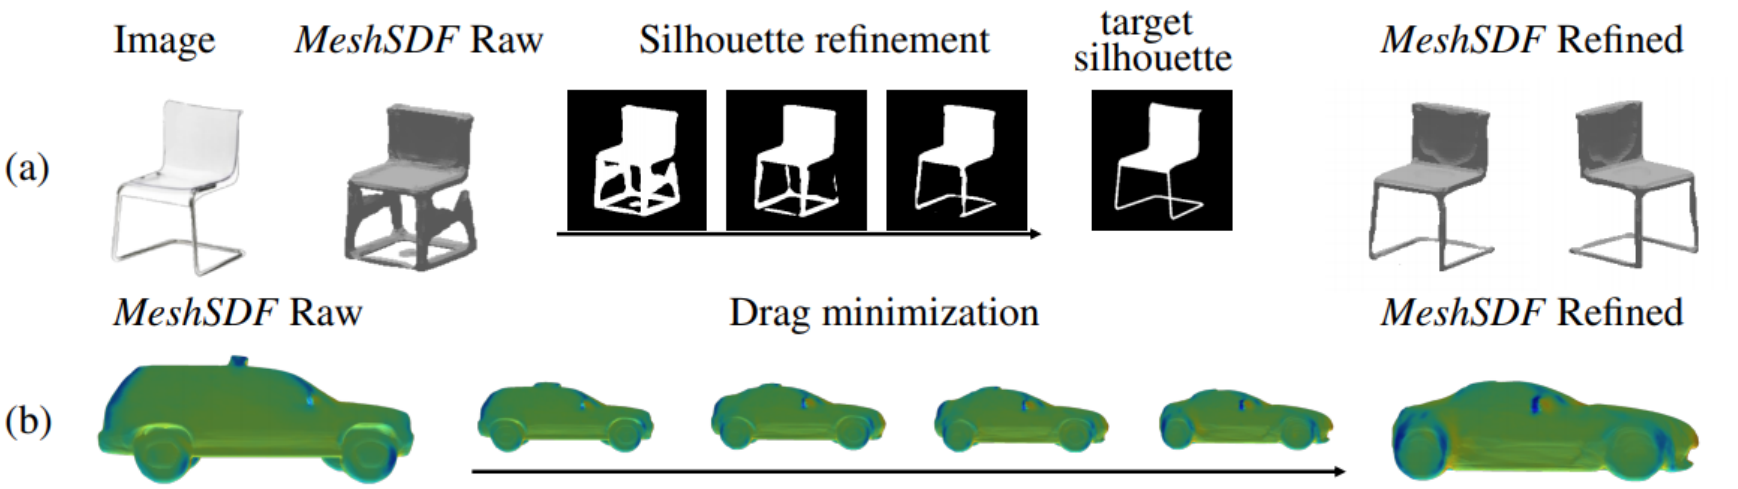

- Differentiable mapping from deep signed distance functions to mesh-based representation
- Makes “marching cubes” differentiable

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - 3D Scene Representation using Neural Networks
  - Photometric reconstruction
  - RGB-D scanning and refinement
- Questions

Sparse Photometric 3D Face Reconstruction Guided by Morphable Models

Xuan Cao et al. 2018

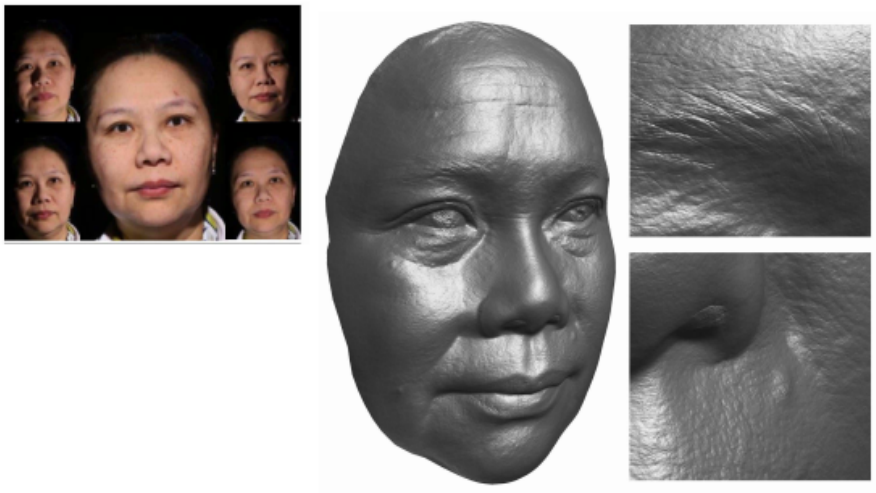

- A highly accurate 3D face is generated by solving an inverse rendering problem with the help of a 3D morphable model, using few input images taken by a fix camera, under unknown changing lighting.

Robust Multiview Photometric Stereo using Planar Mesh Parameterization

Jaesik Park et al. 2016

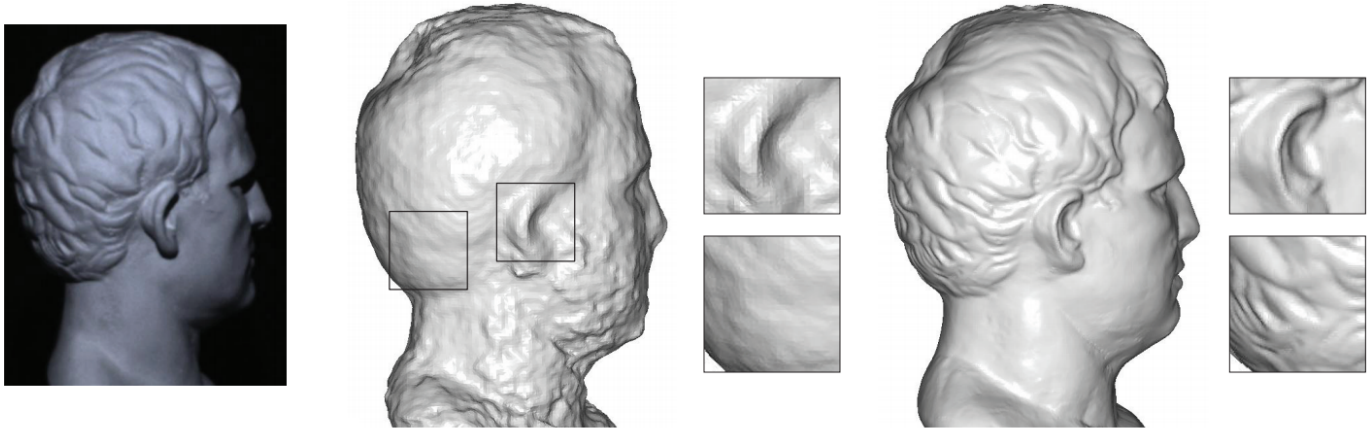

Initial 3D reconstruction

Refined result

- Many images of an object are taken by several fix cameras, under unknown changing lighting.
- By using the input images and an initial coarse mesh obtained with an existing multiview technique, a displacement map is optimized in order to produce a detailed mesh.

Inverse Path Tracing for Joint Material and Lighting Estimation

Azinovic et al. 2019, CVPR

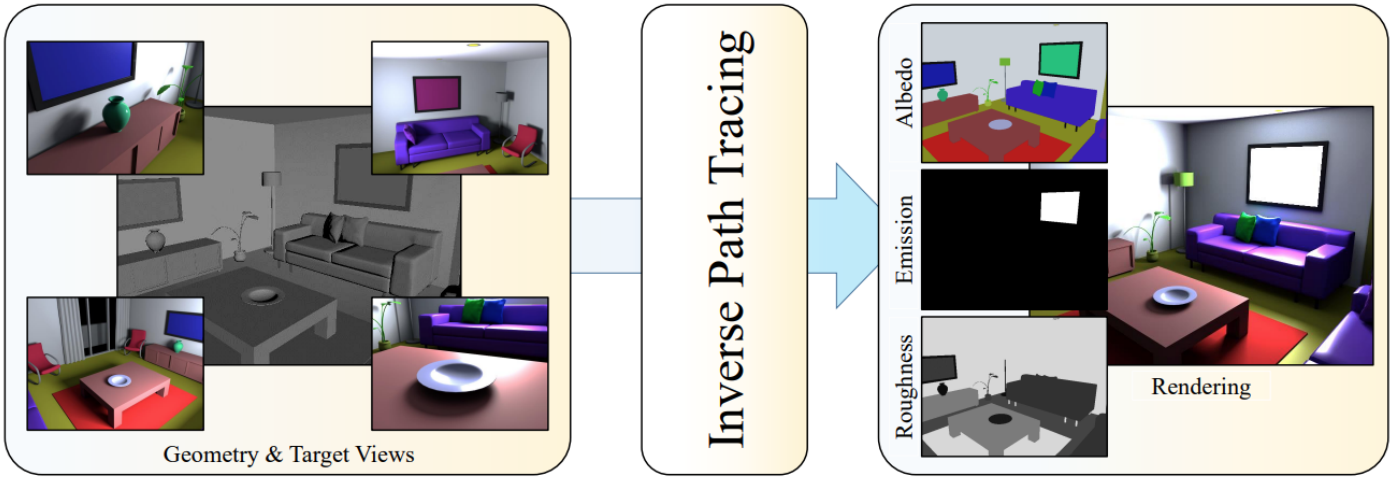

- Inverse Rendering (hot topic in graphics and vision)
- Recover scene properties (material and light sources) in room-scale environments
- Formulates inverse differentiable path tracer

Inverse Rendering for Complex Indoor Scenes: Shape, Spatially-Varying Lighting and SVBRDF from a Single Image

Li et al. 2020, CVPR

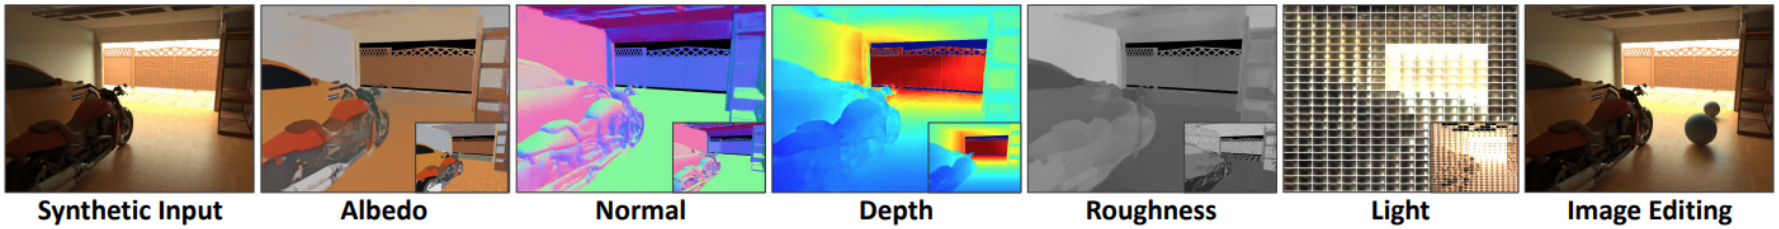

- Inverse Rendering (hot topic in graphics and vision)
- Train Neural Net to recover scene properties from a single image
- Photorealistic object insertion and material editing

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - 3D Scene Representation using Neural Networks
  - Photometric reconstruction
  - RGB-D scanning and refinement
- Questions

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

Erik Bylow et al. 2019

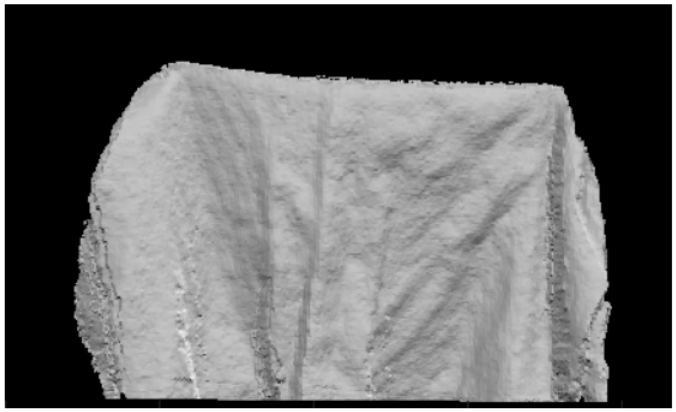

Volumetric fusion

Final result

- Several RGB images of an object are captured with fix camera positions under varying unknown natural illuminations, and they are used to optimize the SDF intialized with a Volumetric Fusion of an RGB-D sequence.

TextureFusion: High-Quality Texture Acquisition for Real-Time RGB-D Scanning

Joo Ho Lee et al. 2020

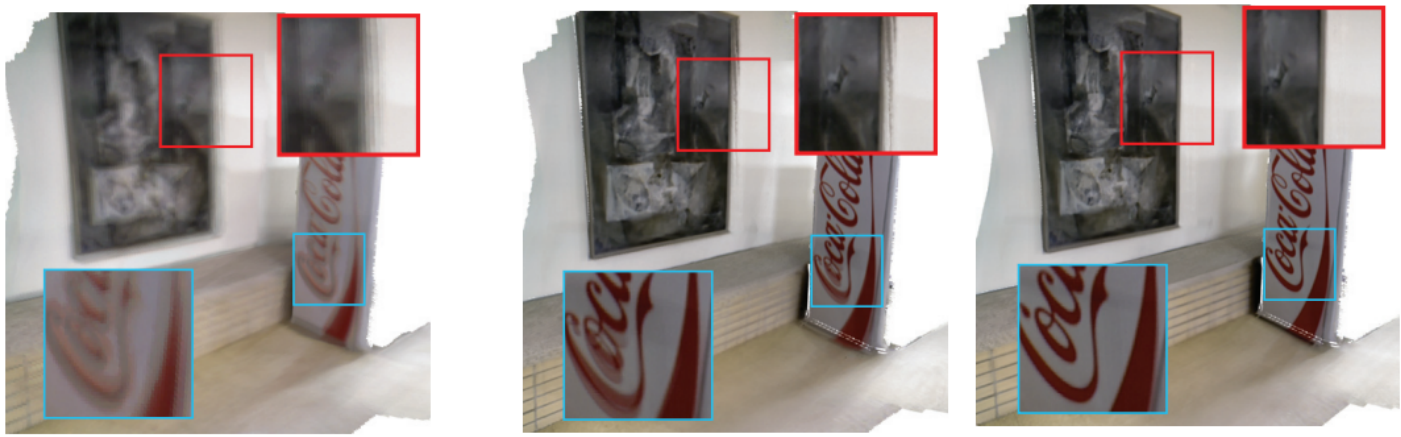

No refinement | Offline optimization | TextureFusion

- While scanning an indoor environment with an RGB-D camera, perform a real-time Texture fusion and refinement with a higher resolution than the Geometric fusion.
- The final texture do not require any additional adjustement and can be even better than those obtained with an offline optimization.

KillingFusion: Non-rigid 3D Reconstruction without Correspondences

Miroslava Slavcheva et al. 2017

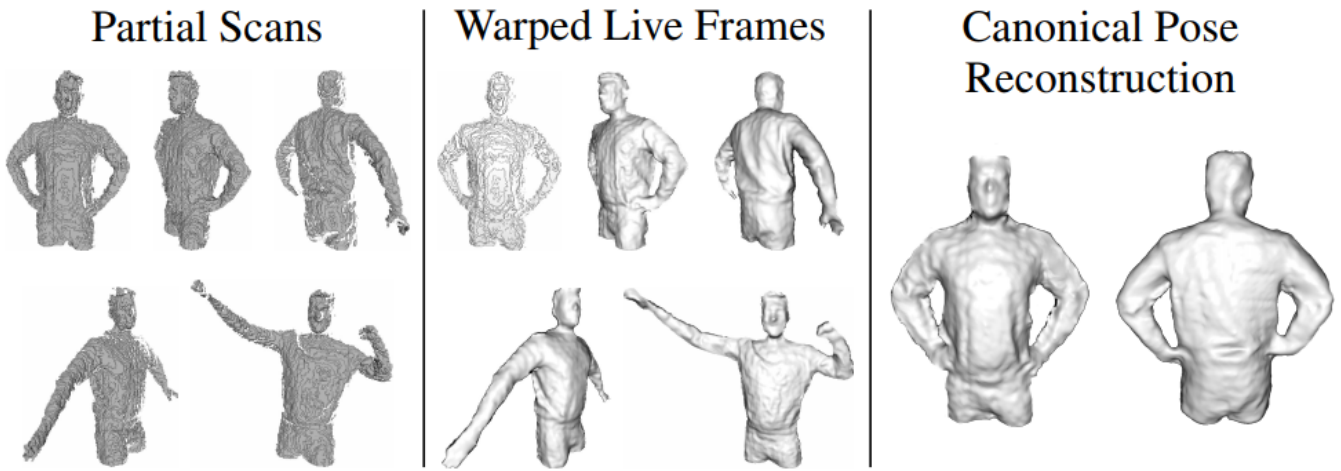

- Estimate a 3D reconstruction of a deforming object from an RGB-D sequence, without relying on additional priors.
- The use of SDF to represent geometry and a dense deformation field allow to handle complicated topological changes.

Questions?

Reminder:

- Web page: `https://vision.in.tum.de/teaching/ss2021/seminar_3dcv`
- Password: `ss21-3dcv`
- Contact: `3dcv-ss21@vision.in.tum.de`