Seminar: Recent Advances in 3D Computer Vision

Mohammed Brahimi, David Schubert
Computer Vision Group
Technical University of Munich

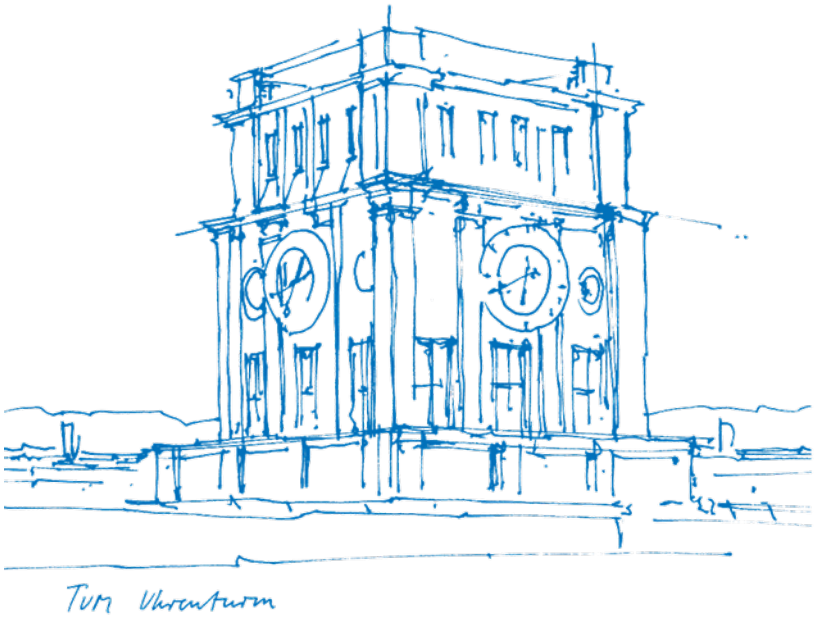

How can I access these slides?

- Option 1 (preferred): seminar web page
  - `https://vision.in.tum.de/teaching/ss2022/seminar_3dcv`
  - Password for material page: `SS22-3dcv`
  - Material page will go online after this pre-meeting
- Option 2: contact organizers
  - `3dcv-ss22@vision.in.tum.de`
  - **Only use this option if you forgot the password**

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - 3D Reconstruction
  - SLAM / SfM
- Questions

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - 3D Reconstruction
  - SLAM / SfM
- Questions

How is the seminar organized?

- Seminar meetings: Talks and discussion
  - Day: Monday, 11th April and Tuesday, 12th April (block seminar)
  - Time: tba
  - Location: tba
  - In case of special circumstances please let us know and we will find a solution
  - Each day will consist of 4 or 5 talks which are held in English
  - **Attendance is mandatory!**

How is the seminar organized?

- Talk preparation / contact with supervisor
  - Read through your paper and write down what you don't understand
  - Three or four weeks before talk: meet supervisor for questions (optional, but recommended)
  - One week before talk: meet supervisor to go through slides (optional, but recommended)
  - Before first presentation on 11th April: submit your slides via submission system (mandatory)
  - Until 1st May, 23:59: 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 one slide!
- 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
  - Upload final report as pdf via submission system
  - Submission deadline: **1st May 2022, 23:59**
- 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 Reconstruction
  - SLAM / SfM
- 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 matching system: 15th February 2022, for email: 19th February 2022

How do you register for the seminar?

Example registration email:

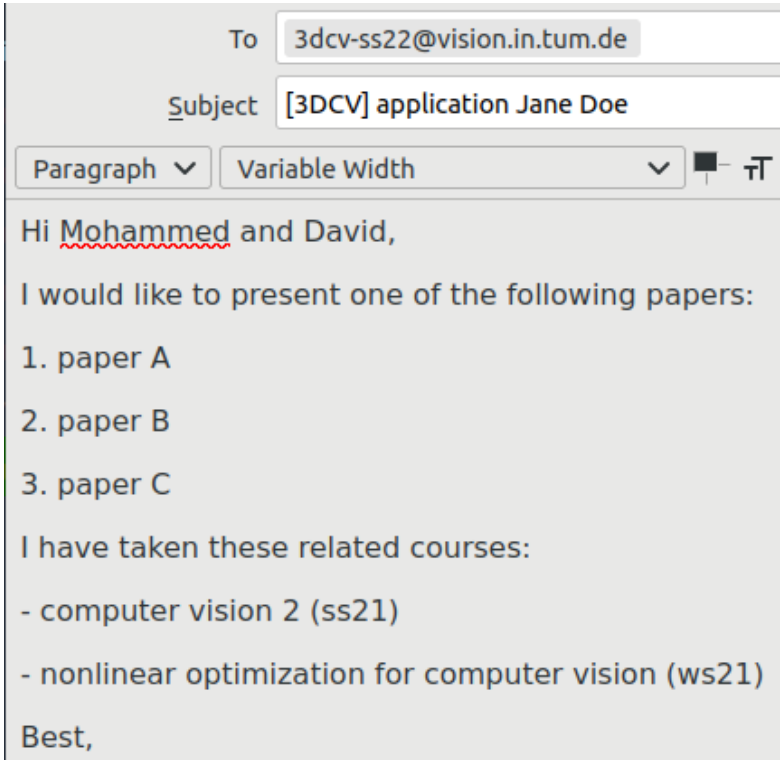

How do we select candidates and assign papers?

- 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.
- 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 Reconstruction
  - SLAM / SfM
- Questions

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

Erik Bylow et al. SCIA 2019

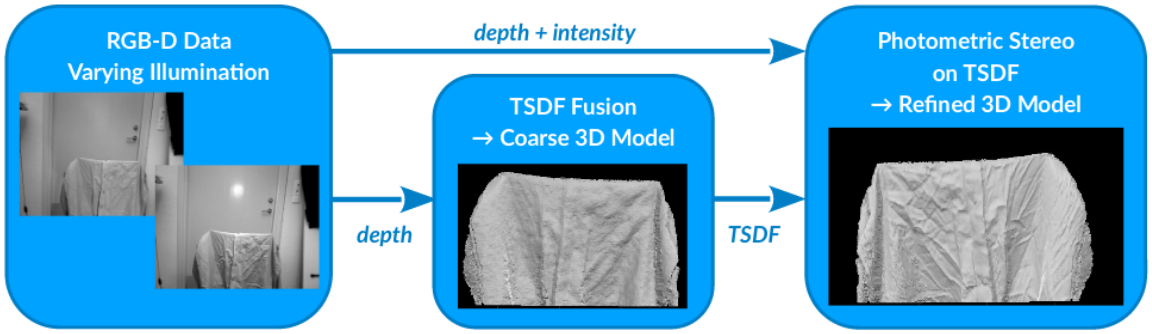

- For each view, many RGB images with different lighting are captured
- Refine a coarse reconstruction obtained with RGB-D fusion
- Geometry is represented using a discrete SDF (voxel grid)

Multiview Neural Surface Reconstruction by Disentangling Geometry and Appearance

Lior Yariv et al. NeurIPS 2020

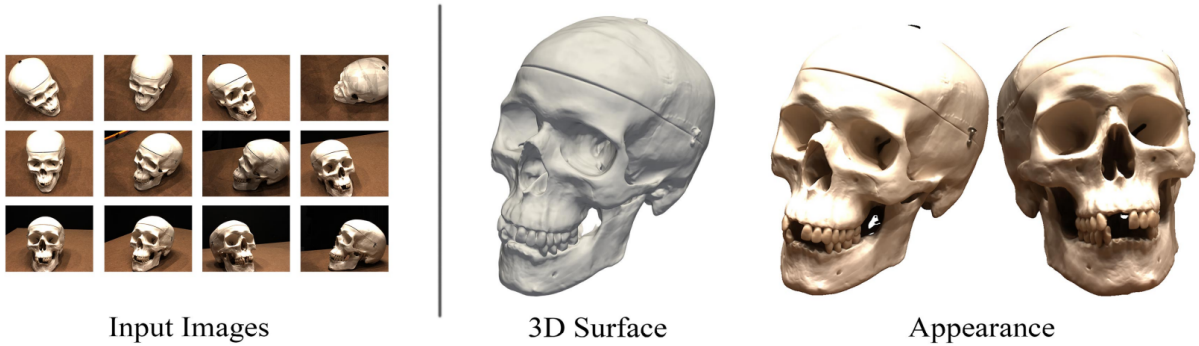

- One RGB image is captured for many known camera positions with static illumination.
- Two MLPs representing respectively the SDF and the appearance are optimized jointly.

PhySG: Inverse Rendering with Spherical Gaussians for Physics-based Material Editing and Relighting

Kai Zhang et al. CVPR 2021

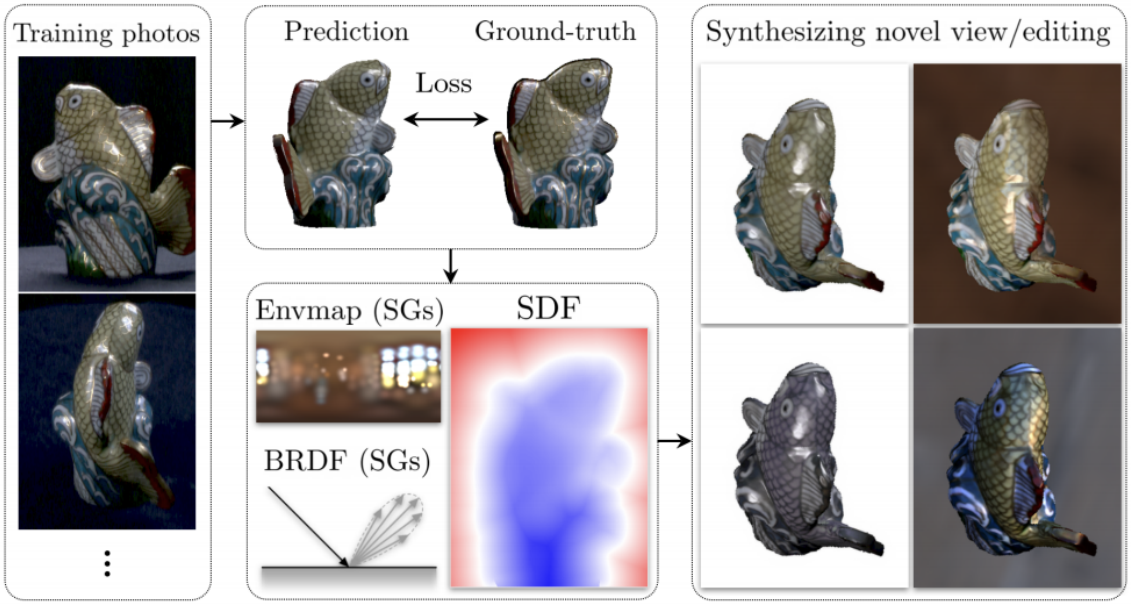

- The appearance is decoupled into lighting and material, which are optimized jointly with the geometry.

Unified Shape and SVBRDF Recovery using Differentiable Monte Carlo Rendering

Luan et al., EGSR 2021

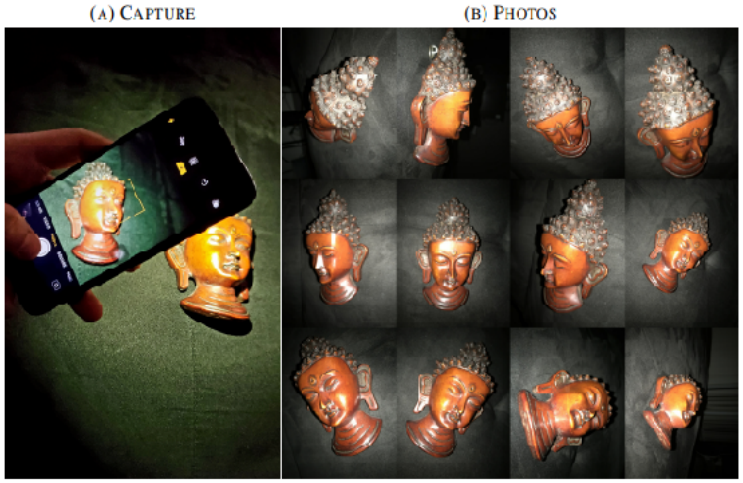

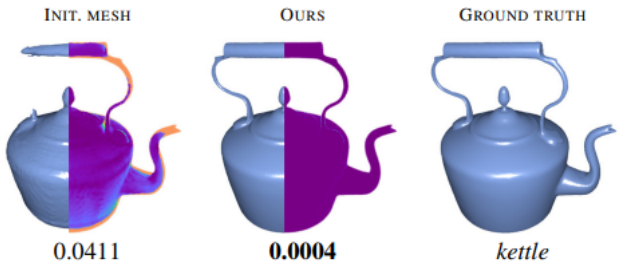

- The light source is collocated with the camera.
- Geometry is represented using a Mesh, whose topology can evolve!

NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis

Ben Mildenhall et al. ECCV 2020

- One RGB image is captured for many known camera positions with static illumination.
- Allow to render the scene from arbitrary viewpoint.

Occupancy Networks: Learning 3D Reconstruction in Function Space

Mescheder et al., CVPR 2019

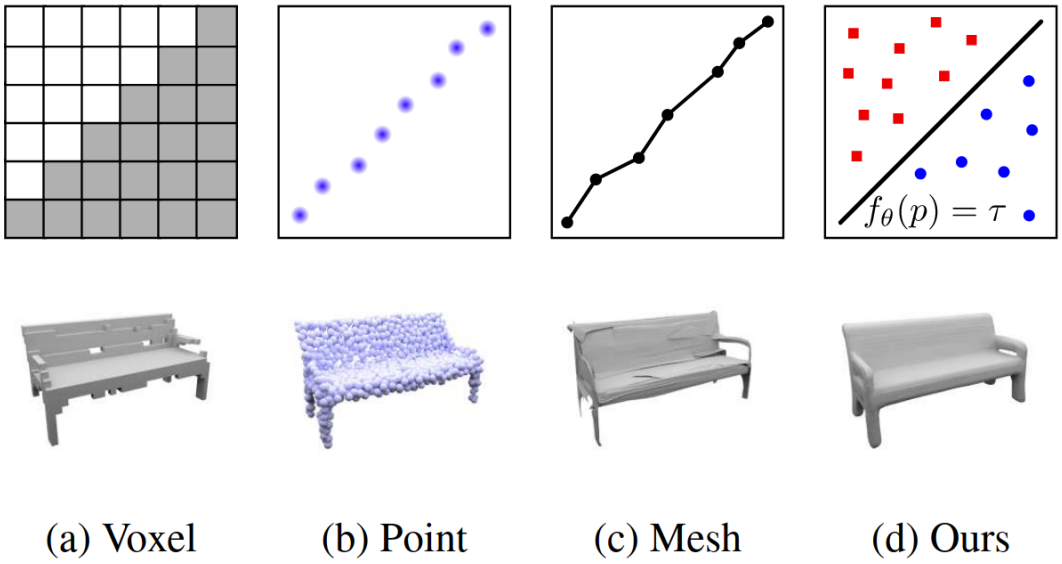

Represent object surface as the decision boundary of a neural network

PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation

Qi et al., CVPR 2017

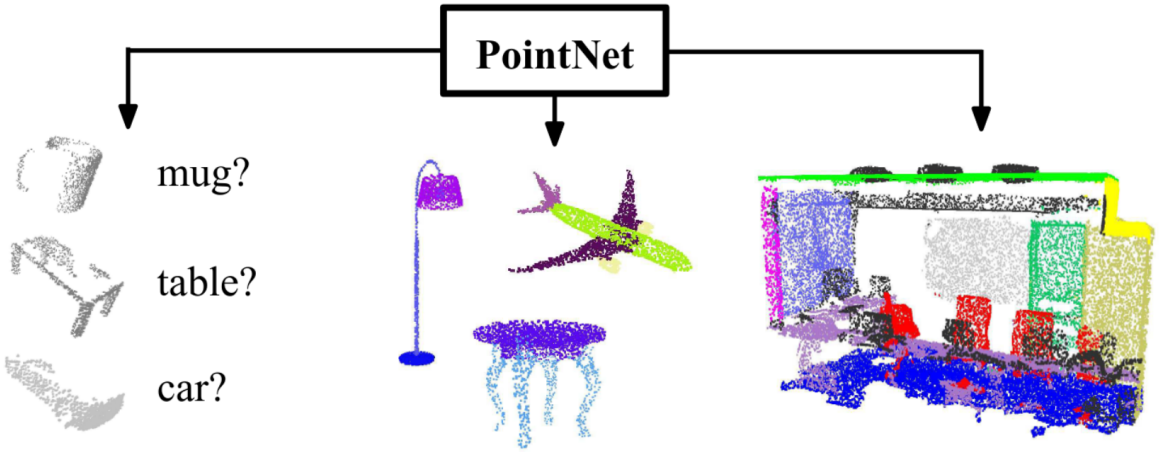

A neural network that can consume point clouds without intermediate voxel representation

Outline

- General Information
  - About the Seminar
  - Registration
- Possible Papers
  - 3D Reconstruction
  - SLAM / SfM
- Questions

CodeSLAM — Learning a Compact, Optimisable Representation for Dense Visual SLAM

Bloesch et al., CVPR 2018

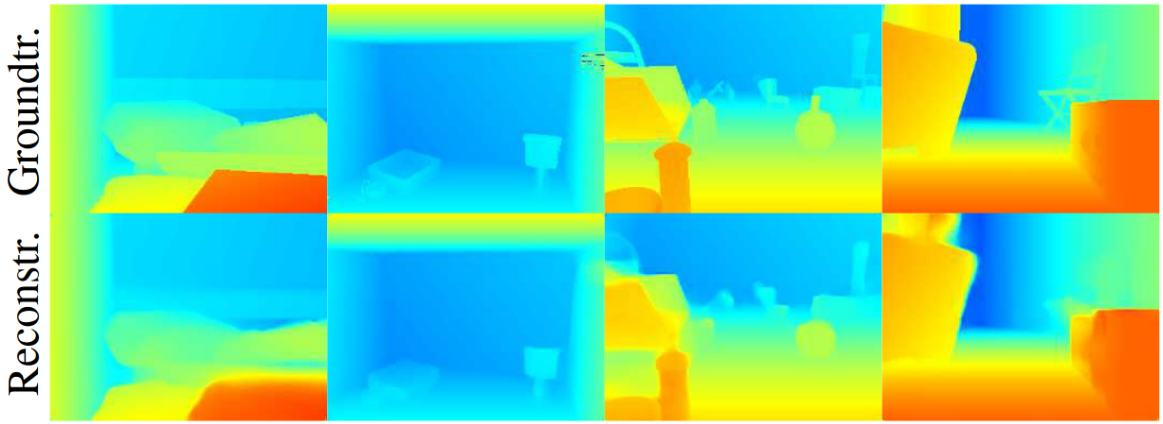

Learn a low-dimensional depth encoding that can be used to optimize dense depth

BAD SLAM: Bundle Adjusted Direct RGB-D SLAM

Schöps et al., CVPR 2019

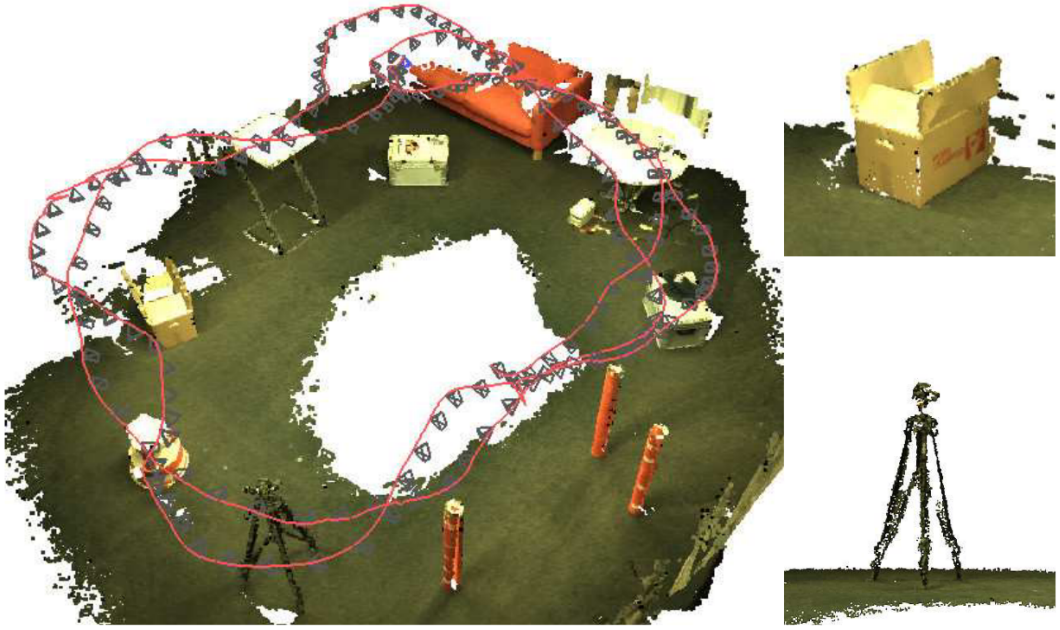

Perform bundle adjustment on surfels to get a high-quality pose

Pixel-Perfect Structure-from-Motion with Featuremetric Refinement

Lindenberger et el., ICCV 2021

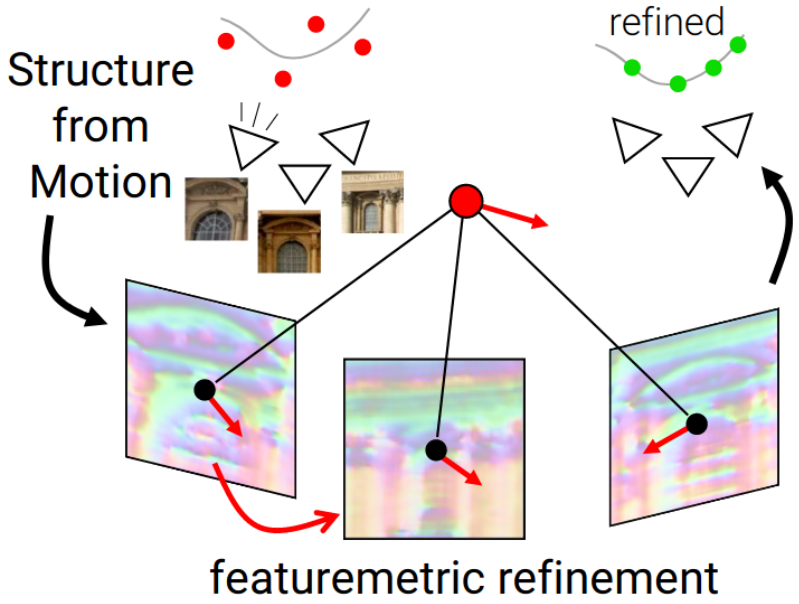

Refine feature locations using dense features predicted by a neural network

Questions?

Reminder:

- Web page: `https://vision.in.tum.de/teaching/ss2022/seminar_3dcv`
- Password: `SS22-3dcv`
- Contact: `3dcv-ss22@vision.in.tum.de`