


Computer Vision and Machine Learning for Computer Graphics

Summer Semester 2022

Christian Theobalt, Marc Habermann, **Thomas Leimkühler**

Visual Computing and Artificial Intelligence, MPI Informatik

max planck institut informatik

Overview

- Organization
- Introduction
- Topics
- Summary

Overview

max planck institut
informatik

- Organization
- Introduction
- Topics
- Summary

Organizers

Christian Theobalt
MPI Informatik, Office 228
theobalt@mpi-inf.mpg.de

Marc Habermann
MPI Informatik, Office 218
mhaberma@mpi-inf.mpg.de

Thomas Leimkühler
MPI Informatik, Office 212
thomas.leimkuehler@mpi-inf.mpg.de

Basic Coordinates

- Time: Thursdays, 14:15 – 15:45
- Join in-person: MPI SWS (E1 5), room 029
  - Exception on 12 May: MPI INF (E1 4), room 007
  - Please adhere to the 3G rule (vaccinated, recovered, or tested)
  - Wearing masks is mandatory
- Join remotely: https://zoom.us/j/92103211479?pwd=Y2lRK0JDZFVnbThFNjJFclVreWhtdz09
  - Please turn on your camera
- Website:

  https://vcai.mpi-inf.mpg.de/teaching/vcai_seminar_2022

Formal requirements in a nutshell

- You read all the papers
- Your presence is required
- Submit questions & participate in discussion
- One topic is “Your Topic” (2 papers):
  - Deliver a 40 minute presentation
  - Write a 5–7 page report
- Grade: talk 30%, discussion 30%, report 40%

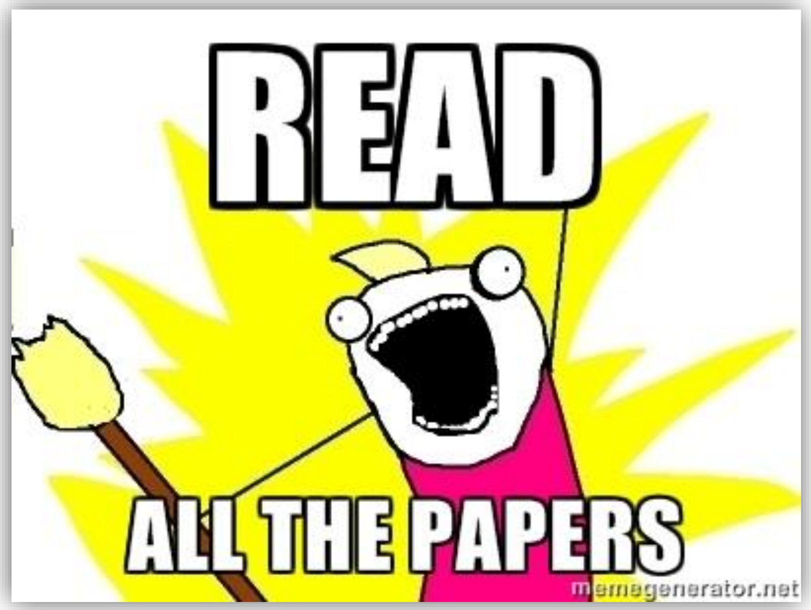

Prior knowledge

- Not for beginners in visual computing
- You need experience in:
  - computer vision
  - computer graphics
  - geometric modeling
  - basic numerical methods
- Examples: you should know how …
  - … a camera is modeled mathematically
  - … 3D transformations are described
  - … a system of equations is solved, etc.

Organization

- 28 topics to choose from
  - listed on seminar website + introduced later today
- 11-12 presentation slots:
  - First presentation: Thursday, 28 April 2022
  - Each week until Thursday, 21 July 2022 (including)
- Each topic comes with a supervisor:
  - You can ask questions by e-mail at any time
    - about your topic, the papers, your presentation and report
  - Up to one office hour per week

Presentations

- Order of presentation will be determined after topic assignment
  - Slots can be swapped if necessary: talk to other participants first
- About 40 minutes long:
  - Introduction (about 5 minutes):
    - summary of previous week
    - finding themes that join the two papers
  - Technical content (about 35 minutes):
    - presentation of the two papers
    - again finding the common links between the papers
- Public feedback from other students after discussion

Suggested presentation preparation

- Schedule two meetings with your supervisor:
  - First meeting: 2–3 weeks before presentation
    - Read the papers for this meeting
    - Ask questions if you have difficulties
    - Discuss your plans for presentation
  - Second meeting: 1 week before presentation
    - Prepare a preliminary presentation (not a full rehearsal)
    - We can provide feedback
  - It is your responsibility to arrange the meetings
  - Do not rely on us providing last-minute feedback

Discussion (45–60 minutes)

- Day before the seminar:
  - Submit 2+ questions for discussion to mhaberma@mpi-inf.mpg.de
  - Important: your contribution will be marked
- At the seminar:
  - One person assigned in advance to lead the discussion
  - Will get the collected questions submitted before the seminar
  - Gives summary of the talk
  - Moderates and guides discussion
  - Raises open questions that remain
  - Discussion of the strengths and weaknesses of the two papers
  - This will also be marked

Report

- 5–7 page summary of the major ideas in your topic:
  - 3–4 pages on the two papers
  - 2–3 pages with your own ideas, for example:
    - Novel ideas based on content described in the papers
    - Limitations not mentioned in the paper + sketch of potential solution
    - Try to suggest improvements
    - Can be the result of the discussion after your presentation
  - 3–4 additional paper references
- The idea is that you get a feeling for your specific topic surpassing the level of simply understanding a paper.

Report

- Due date: **Friday, 18 August 2022**
  (4 weeks after the last seminar)
- Send PDF to thomas.leimkuehler@mpi-inf.mpg.de
- We will provide a LaTeX template on the seminar website
  - If you use other software, make it look like the LaTeX template
    - this is your responsibility
  - Strongly recommended to learn LaTeX
    - used by nearly all research papers in visual computing

Grading scheme

- **Presentation** (overall: 30%)
  - Form (30%): time, speed, structure of slides
  - Content (50%): structure, story line and connections, main points, clarity
  - Questions (20%): answers to questions
  - First presenter (28 April) gets a grade bonus of 0.3
- **Discussion** (overall: 30%)
  - Submitted questions (33%): insight, depth
  - Participation (33%): willingness, debate, ideas
  - Moderation (33%): strengths and weaknesses, integration of questions
- **Report** (overall: 40%)
  - Form (10%): diligence, structure, appropriate length
  - Context (20%): the big picture, topic in context
  - Technical correctness (30%)
  - Discussion (40%): novelty, transfer, own ideas / in own words

Benefits to you

- Practise important skills in research
  - Read and understand technical papers
  - Present scientific results and convince other people
  - Analyse and develop new ideas through discussions
- Discussion is essential
  - If you don’t participate, you miss a big chance
  - Most ideas are developed in discussions about other papers
- Therefore:
  - Prepare for the seminar classes
  - Participate actively in the discussions
  - Benefit from the interaction in the group

Schedule

- **14 April** – Introduction ◄ You are here
- **21 April** – Lectures:
  - "How to read an academic paper"
  - "How to give a good talk"
- **28 April** – First presentation by a student
- … 10 more weekly presentations …
- **21 July** – Last presentation by a student
- **18 August** – Report deadline

Overview

- Organization
- Introduction
- Topics
- Summary

Overview

- Organization
- ► Introduction ◄
- Topics
- Summary

Basics (Image Formation)

Basics (Image Formation)

- **Geometry**

Basics (Image Formation)

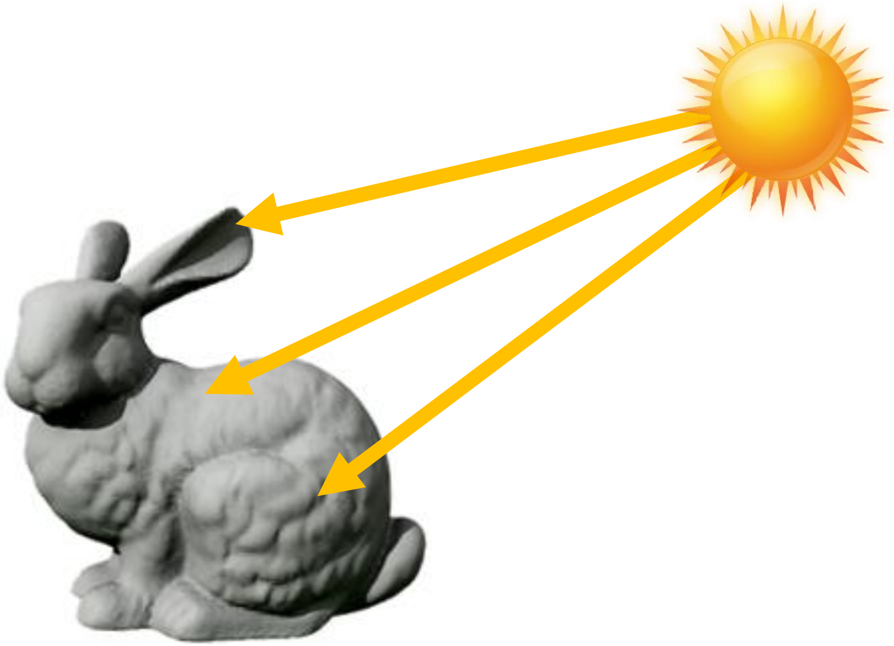

- **Geometry**
- **Illumination**

Basics (Image Formation)

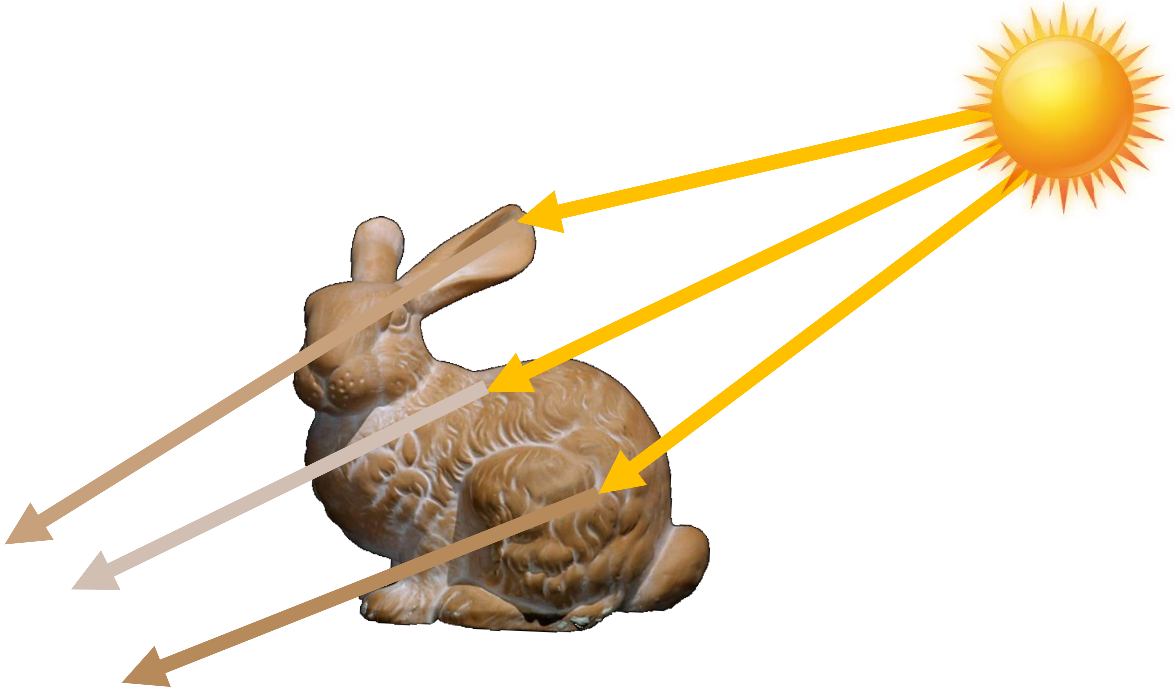

- **Geometry**
- **Illumination**
- **Reflectance**

Basics (Image Formation)

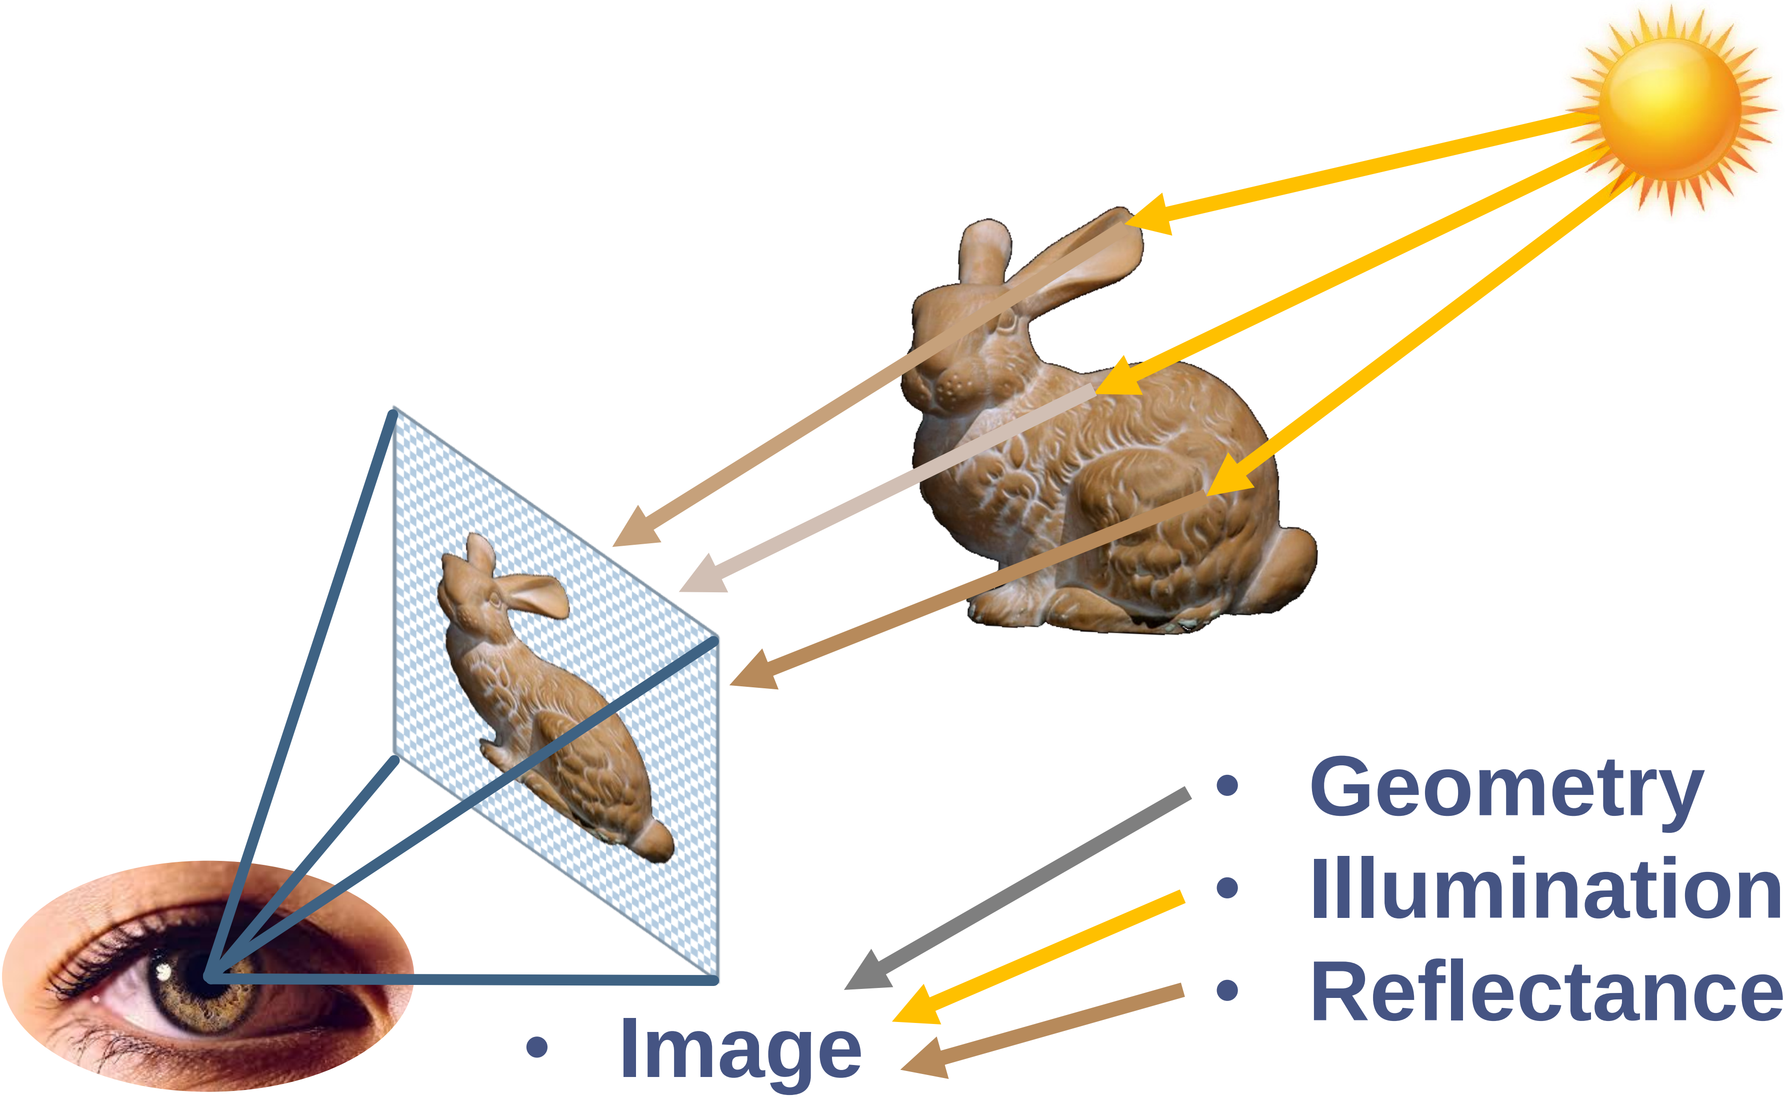

Basics (Image Formation)

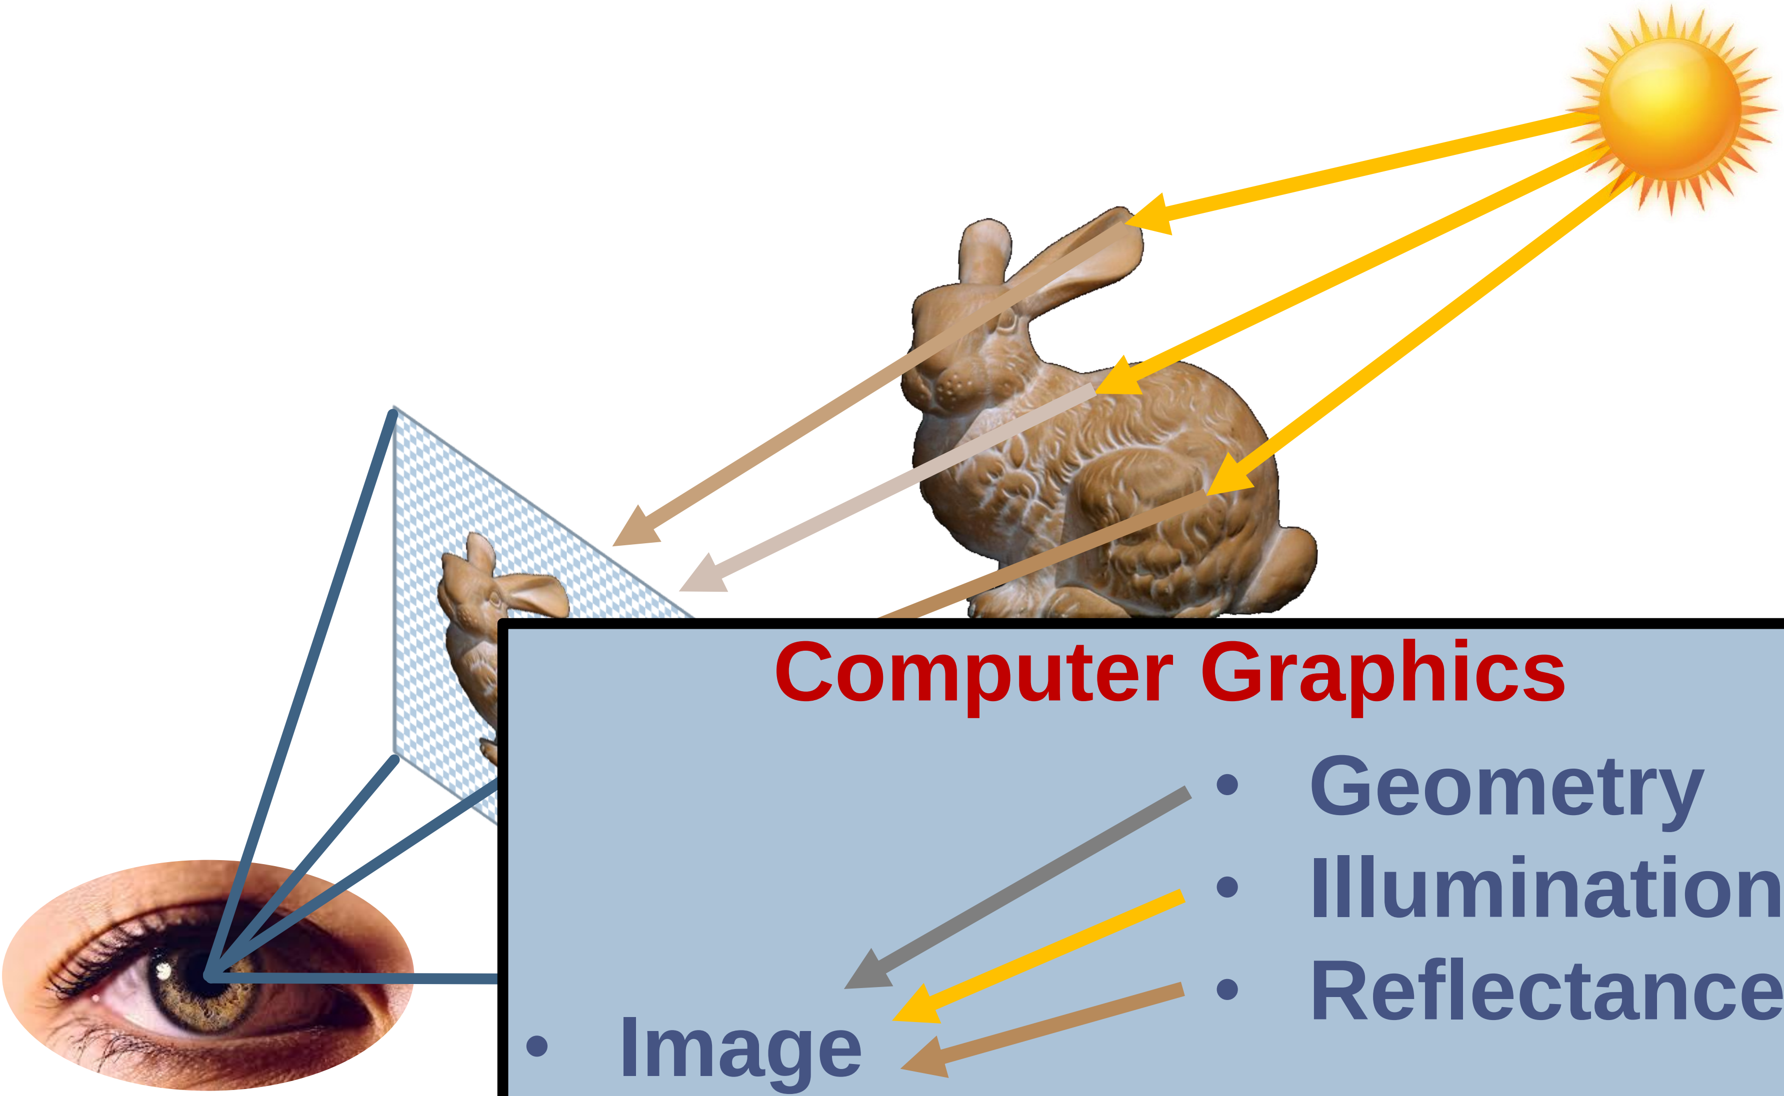

Computer Graphics

- Geometry
- Illumination
- Reflectance

- Image

Photo-real Virtual Humans

The Curious Case of Benjamin Button, 2008

Real or rendered?

Real or rendered?

Basics (Image Formation)

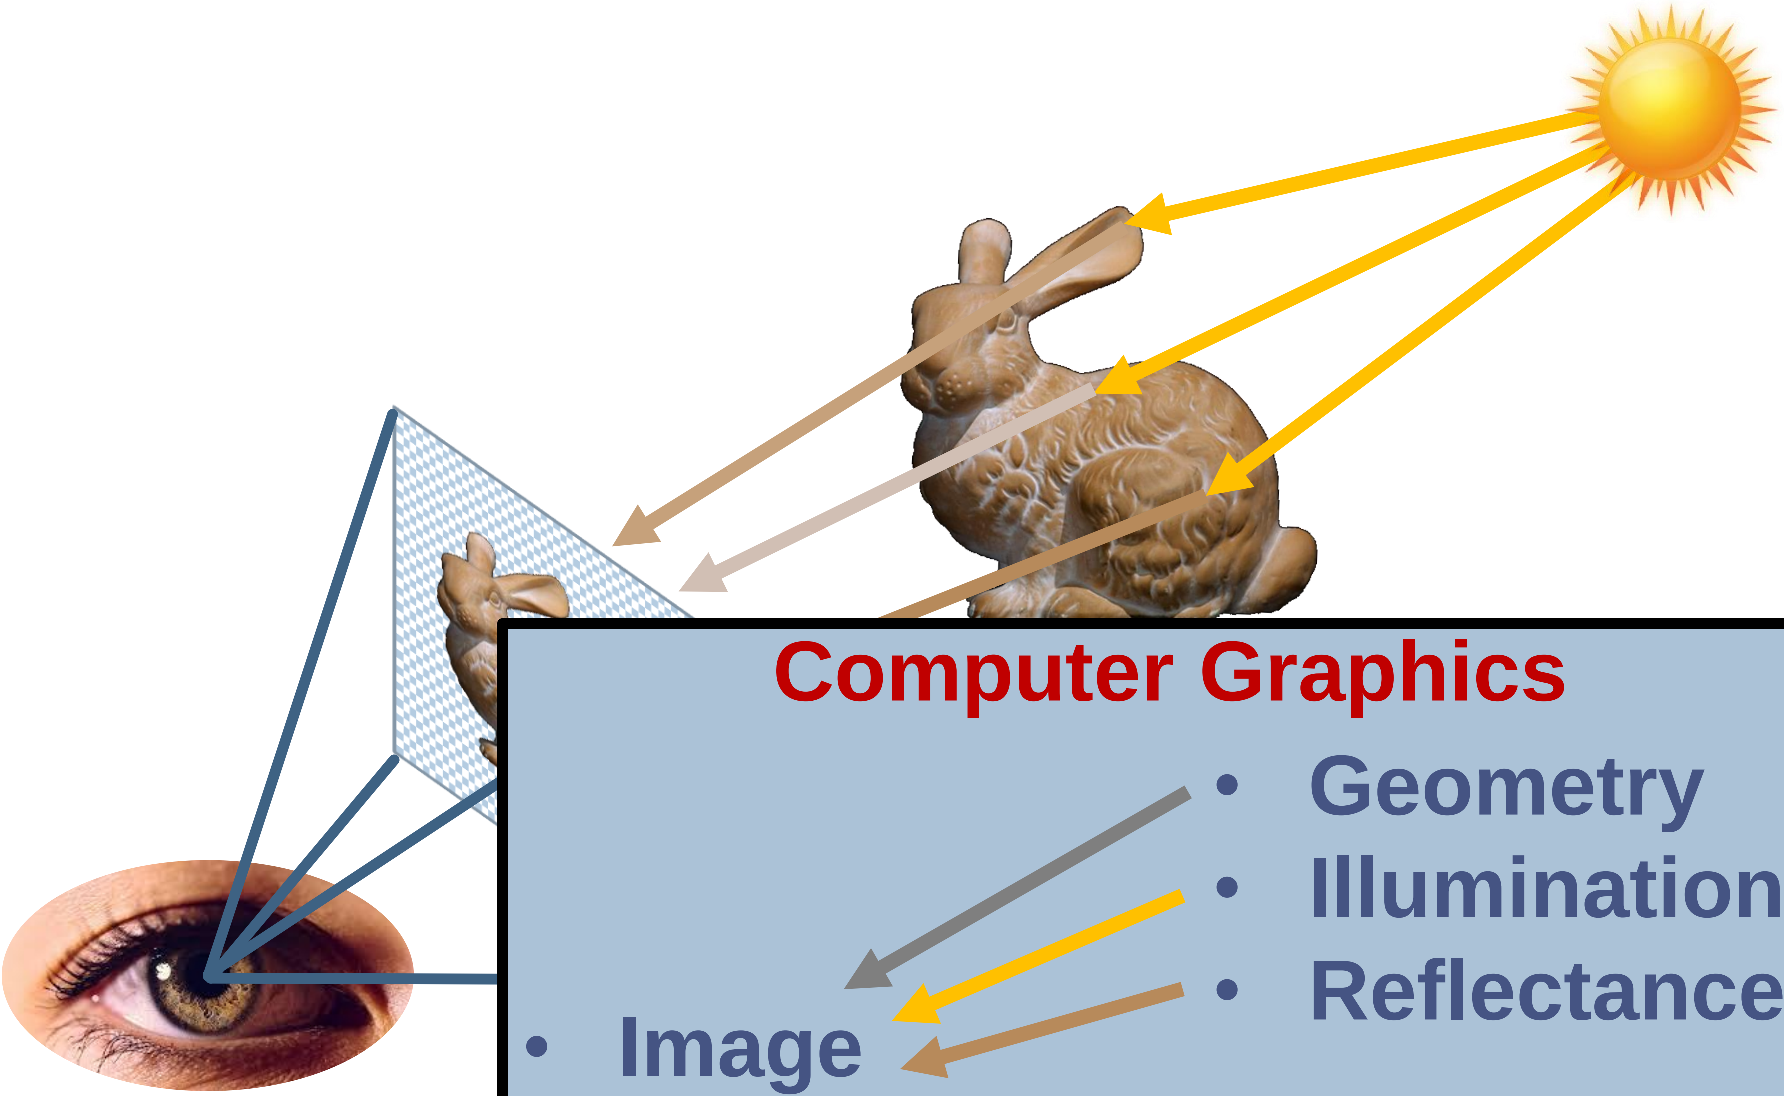

Computer Graphics

- Geometry
- Illumination
- Reflectance

- Image

Basics (Image Formation)

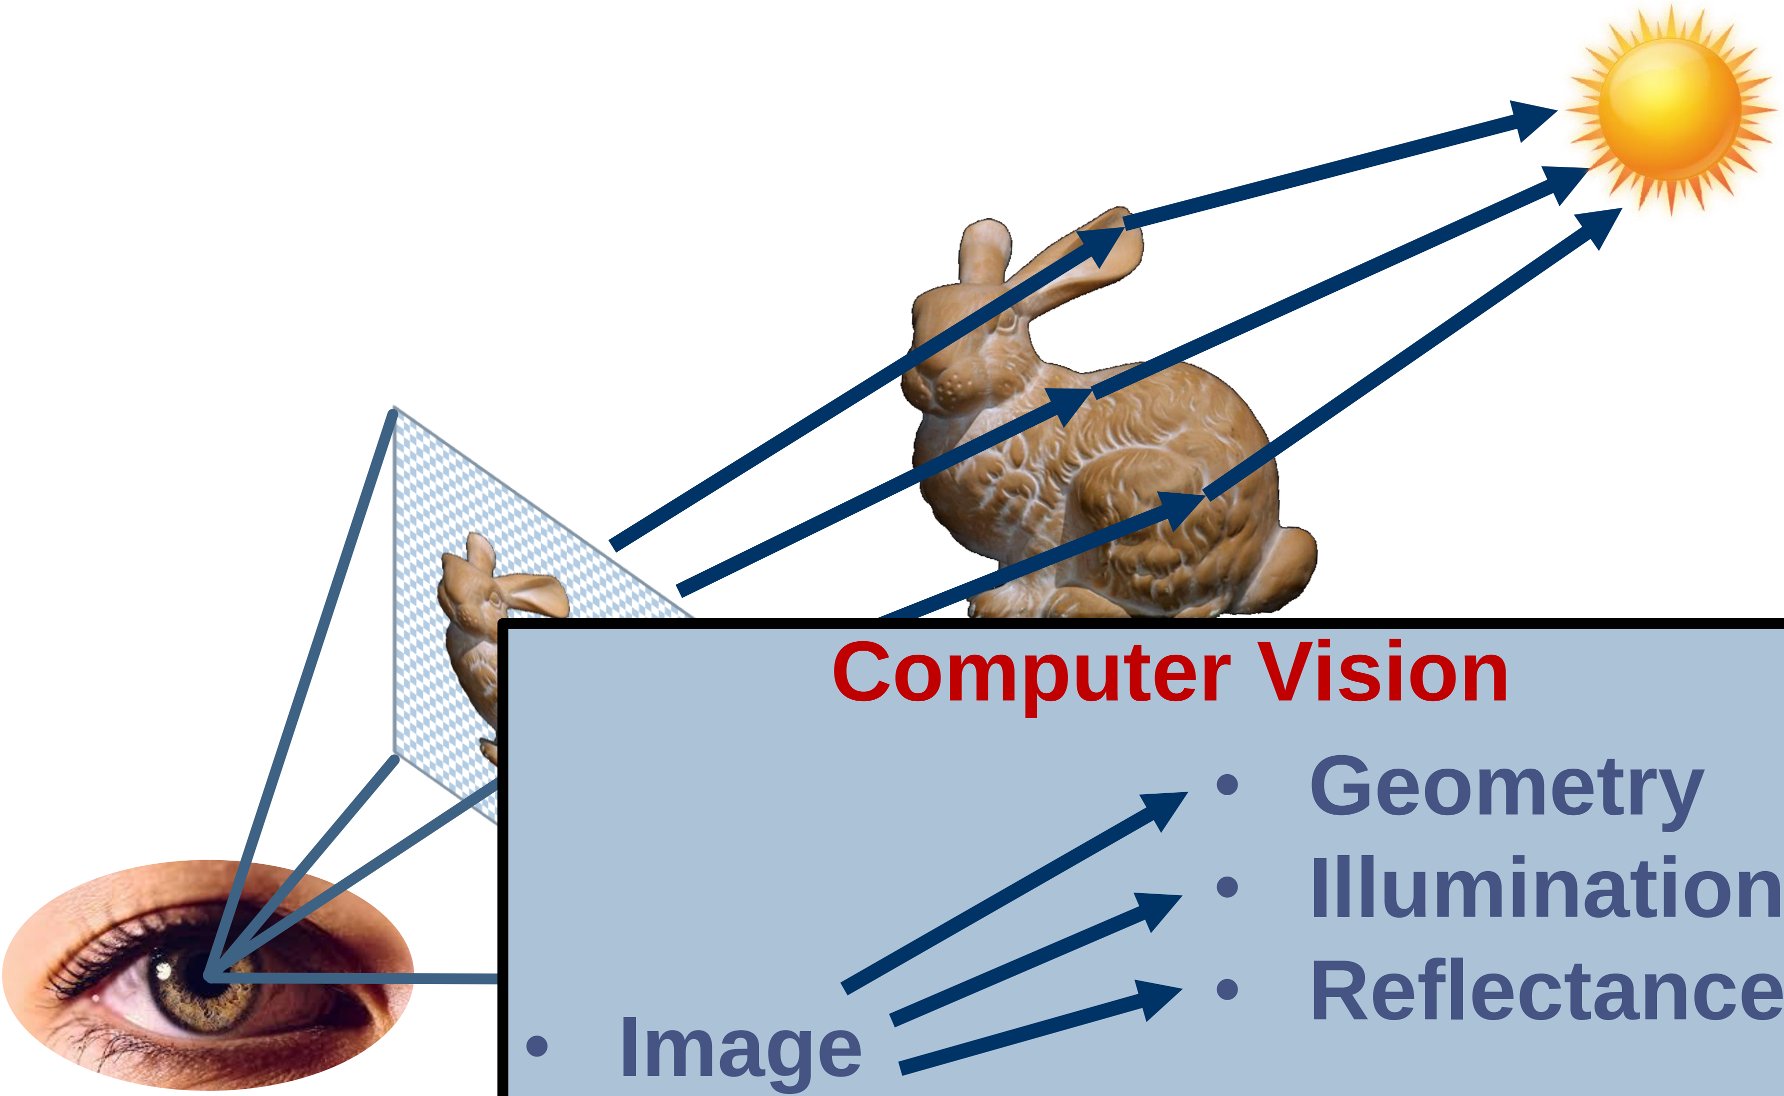

Computer Vision

- Image
- Geometry
- Illumination
- Reflectance

Computer Graphics / Computer Vision

Real world

- Images
- Videos
- Sensor data
- …

← **Computer Graphics**

Machine Learning

Computer Vision →

Scene model

- Geometry
- Illumination
- Reflectance
- Motion

Overview

- Organization
- ▶ Introduction ◀
- Topics
- Summary

Overview

- Organization
- Introduction
- Topics
- Summary

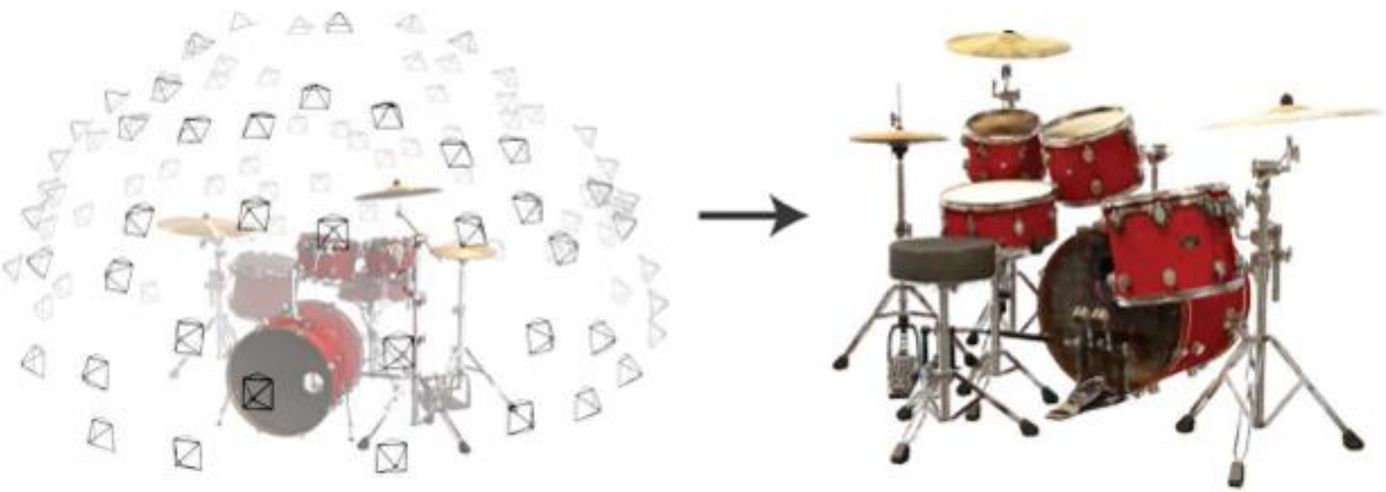

NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis

(Mildenhall et al. ECCV 2020)

Supervisor: Mohit Mendiratta

Multiview Synthesis and Reconstruction

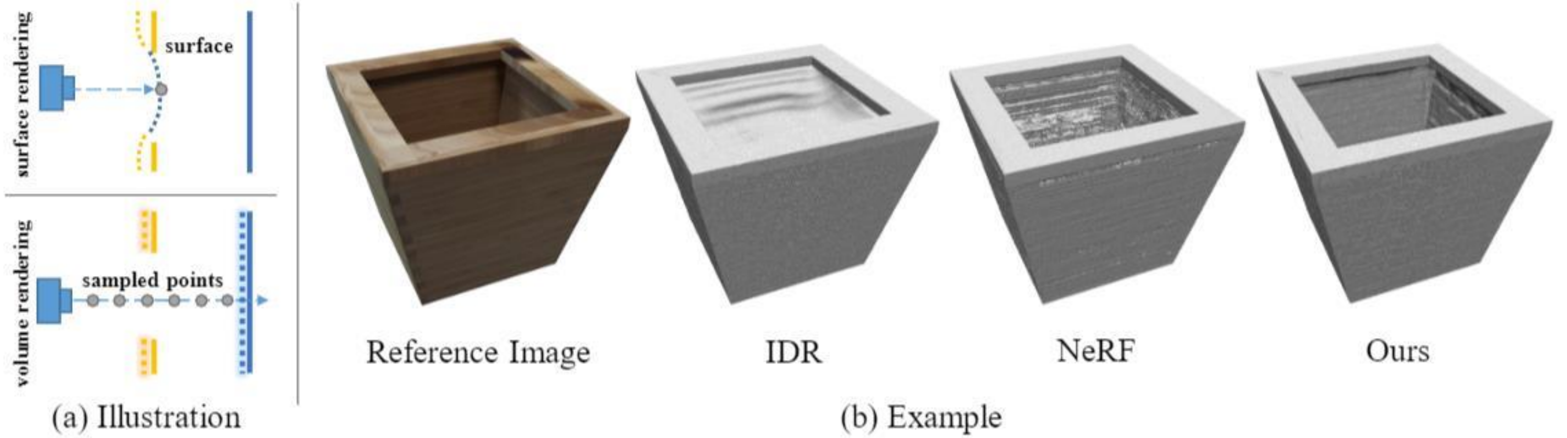

Figure 1: (a) Illustration of the surface rendering and volume rendering. (b) A toy example of bamboo planter, where there are occlusions on the top of the planter. Compared to the state-of-the-art methods, our approach can handle the occlusions and achieve better reconstruction quality.

NeuS: Learning Neural Implicit Surfaces by Volume Rendering for Multi-view Reconstruction

(Wang et al. NeurIPS 2021)

Supervisor: Mohit Mendiratta

Physics based Character Animation

SuperTrack – Motion Tracking for Physically Simulated Characters using Supervised Learning

(Fussell et al., Siggraph 2021)

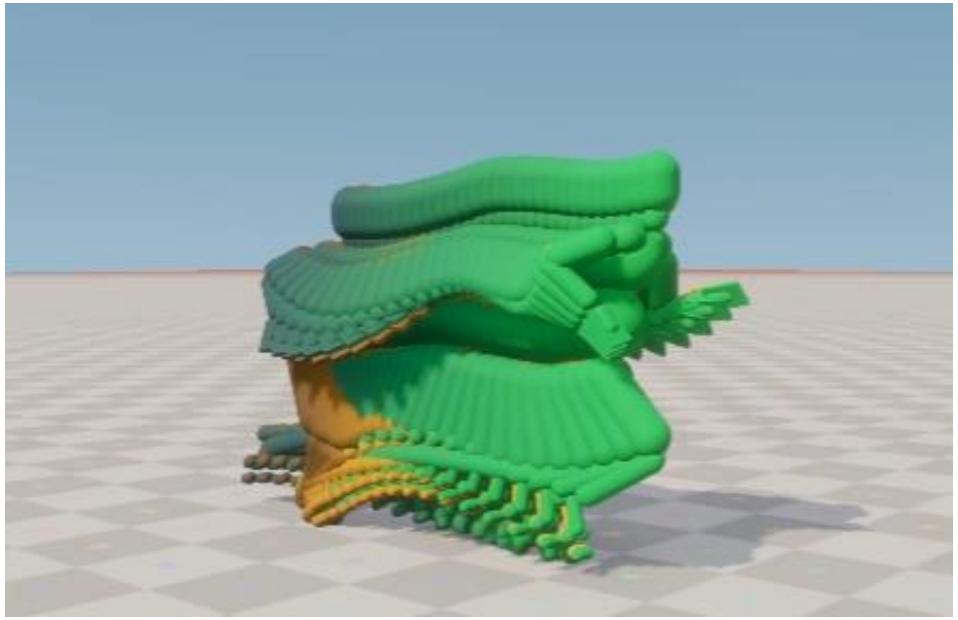

Data Driven Physics for Human Soft Tissue Animation

(Kim et al., Siggraph 2017)

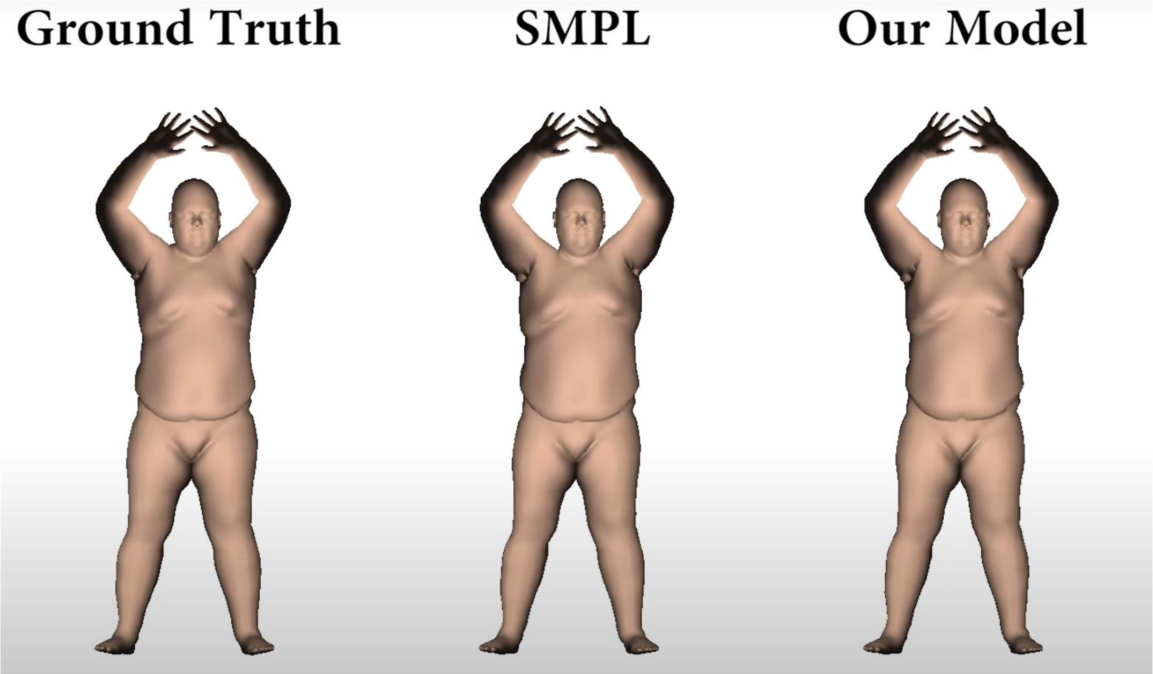

Supervisor: Soshi Shimada

Deeper Human Modeling

Multi-Segment Foot for Human Modeling and Simulation

(Park et al., CGF 2020)

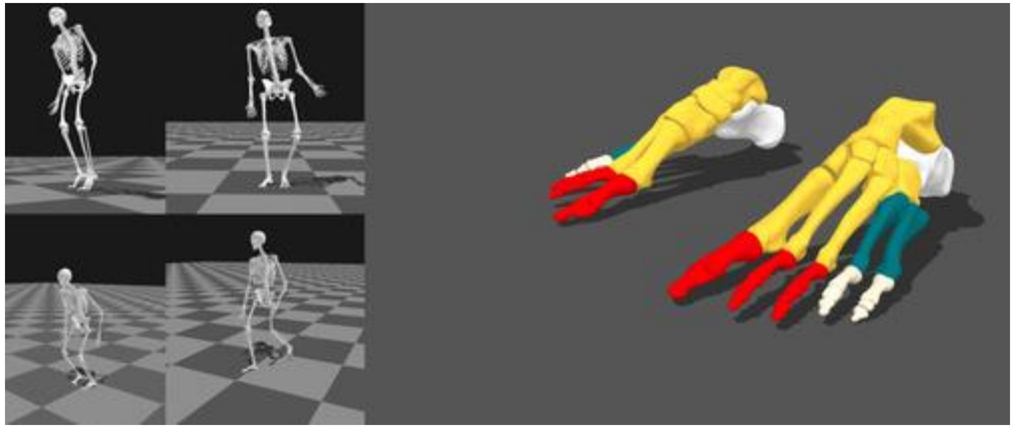

A Multilinear Model for Bidirectional Craniofacial Reconstruction

(Achenbach et al., Eurographics 2018)

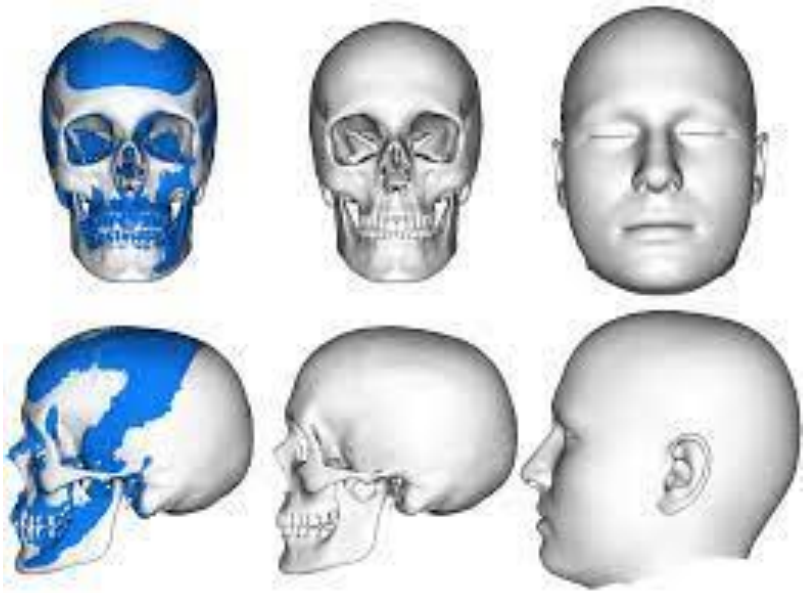

Supervisor: Soshi Shimada

Nerfies: Deformable Neural Radiance Fields

(Park et al., ICCV 2021)

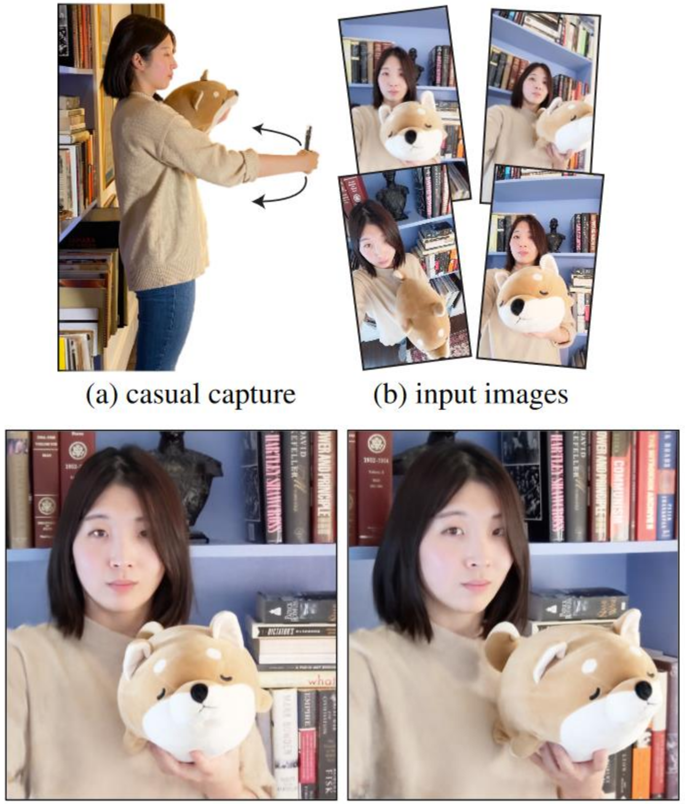

(a) casual capture (b) input images

(c) nerfie novel views

Neural Scene Flow Fields for Space-Time View Synthesis of Dynamic Scenes

(Li et al., CVPR 2021)

Supervisor: Edith Tretschk

Taking Transformers from Language to 3D Reconstruction

Perceiver IO: A General Architecture for Structured Inputs & Outputs

(Jaegle et al., arXiv 2021)

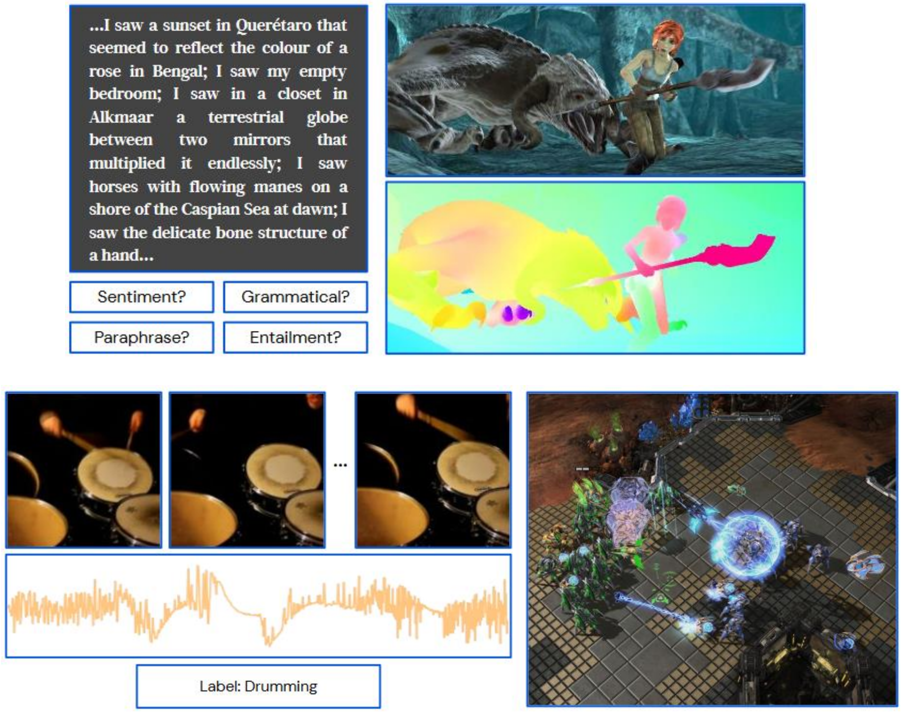

Input-level Inductive Biases for 3D Reconstruction

(Yifan et al., arXiv 2021)

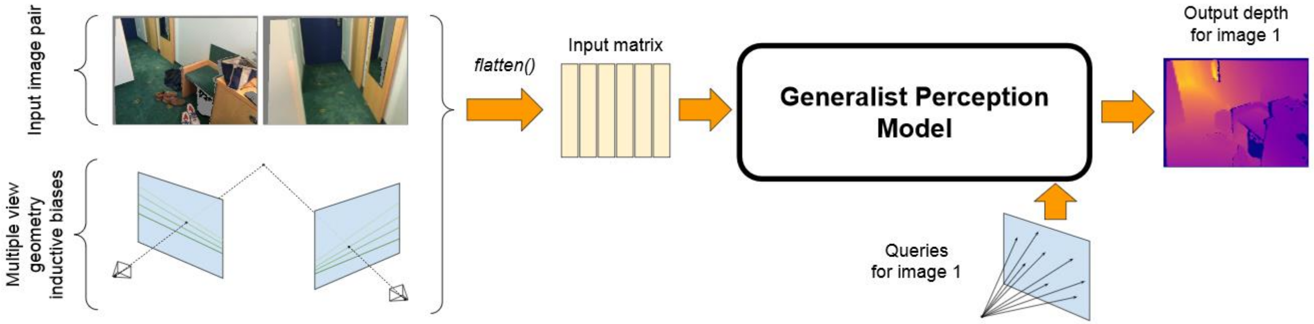

Supervisor: Edith Tretschk

pi-GAN: Periodic Implicit Generative Adversarial Networks for 3D-Aware Image Synthesis

(Chan et al., CVPR 2021)

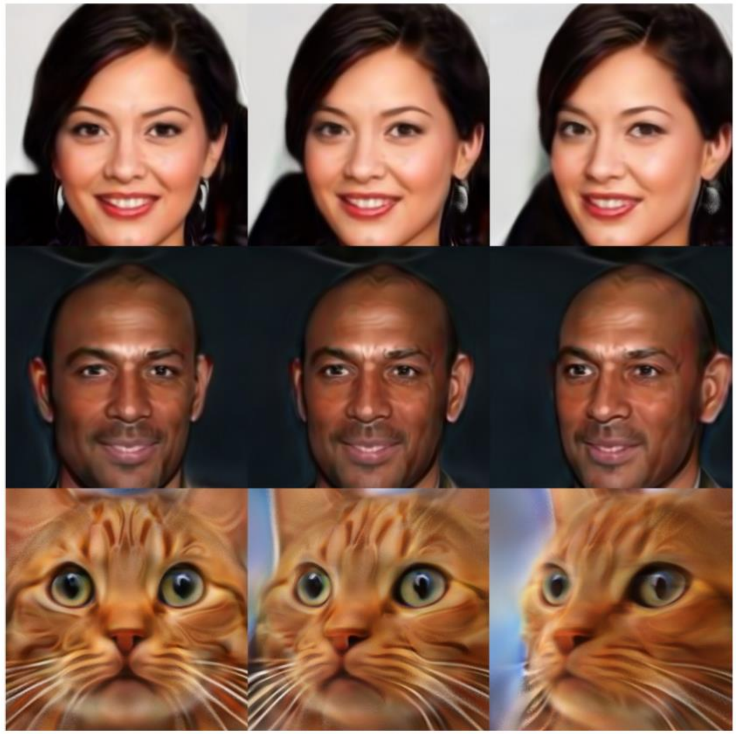

Efficient Geometry-aware 3D Generative Adversarial Networks

(Chan et al., Arxiv 2021)

Supervisor: Xingang Pan

Multiscale Neural Scene Representation

Mip-NeRF: A Multiscale Representation for Anti-Aliasing Neural Radiance Fields

(Barron et al., ICCV 2021)

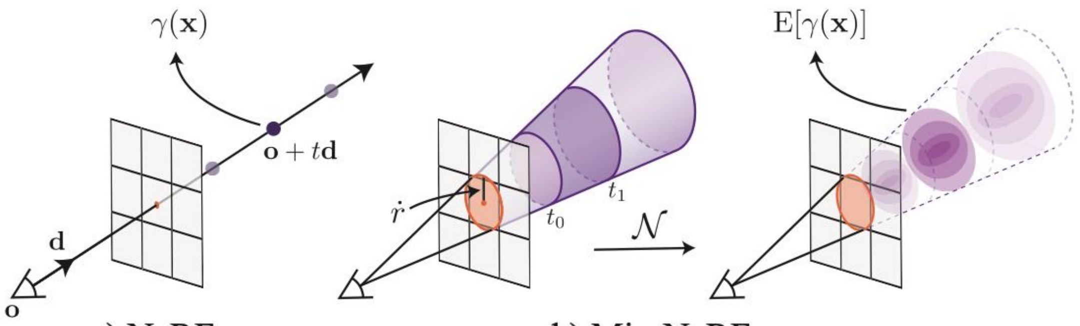

a) NeRF b) Mip-NeRF

BACON: Band-limited Coordinate Networks for Multiscale Scene Representation

(Lindell et al., Arxiv 2021)

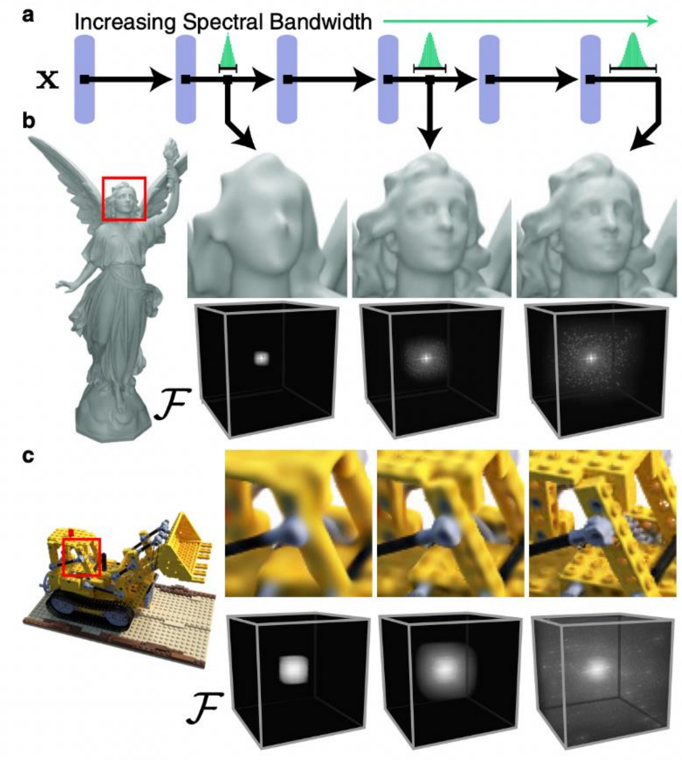

Supervisor: Xingang Pan

Shape Matching with Quantum Algorithms

Matching Point Sets with Quantum Circuit Learning

(Noormandipour and Wang, arXiv, 2021)

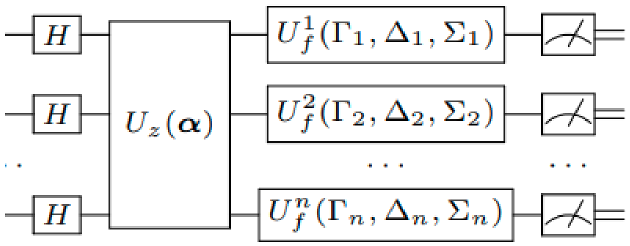

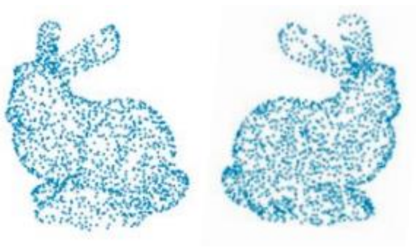

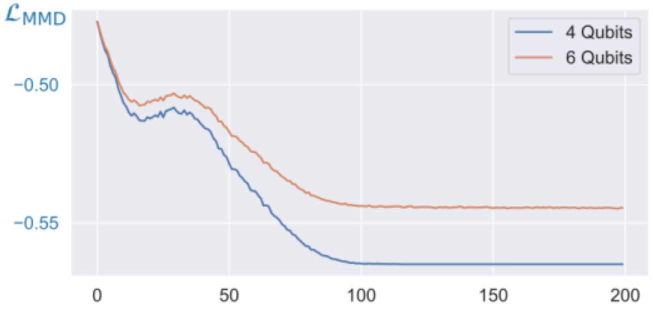

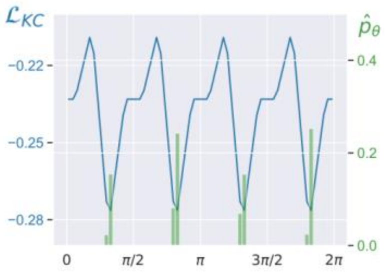

Q-Match: Iterative Shape Matching via Quantum Annealing

(Benkner et al., ICCV 2021)

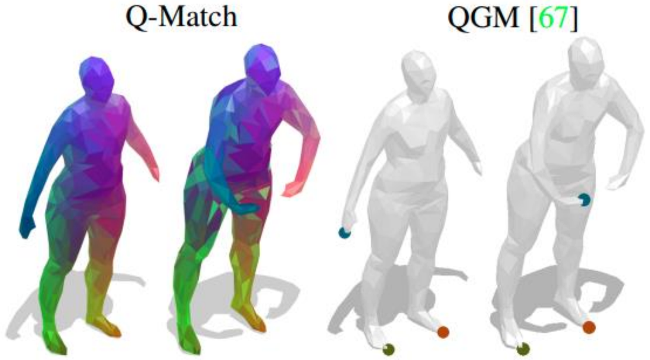

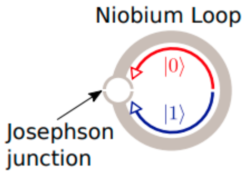

Figure 2. Superconducting flux qubit.

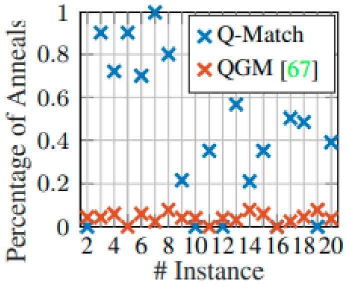

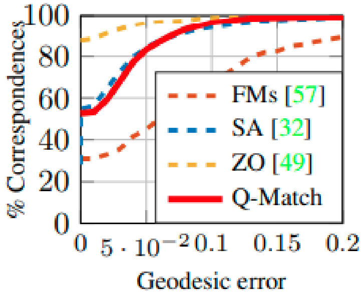

Supervisor: Dr. Vladislav Golyanik

Monocular Non-Rigid 3D Reconstruction

BANMo: Building Animatable 3D Neural Models from Many Casual Videos

(Yang et al., arXiv, 2021)

Physics-Based Monocular Non-Rigid 3D Reconstruction

(currently under review; more details coming)

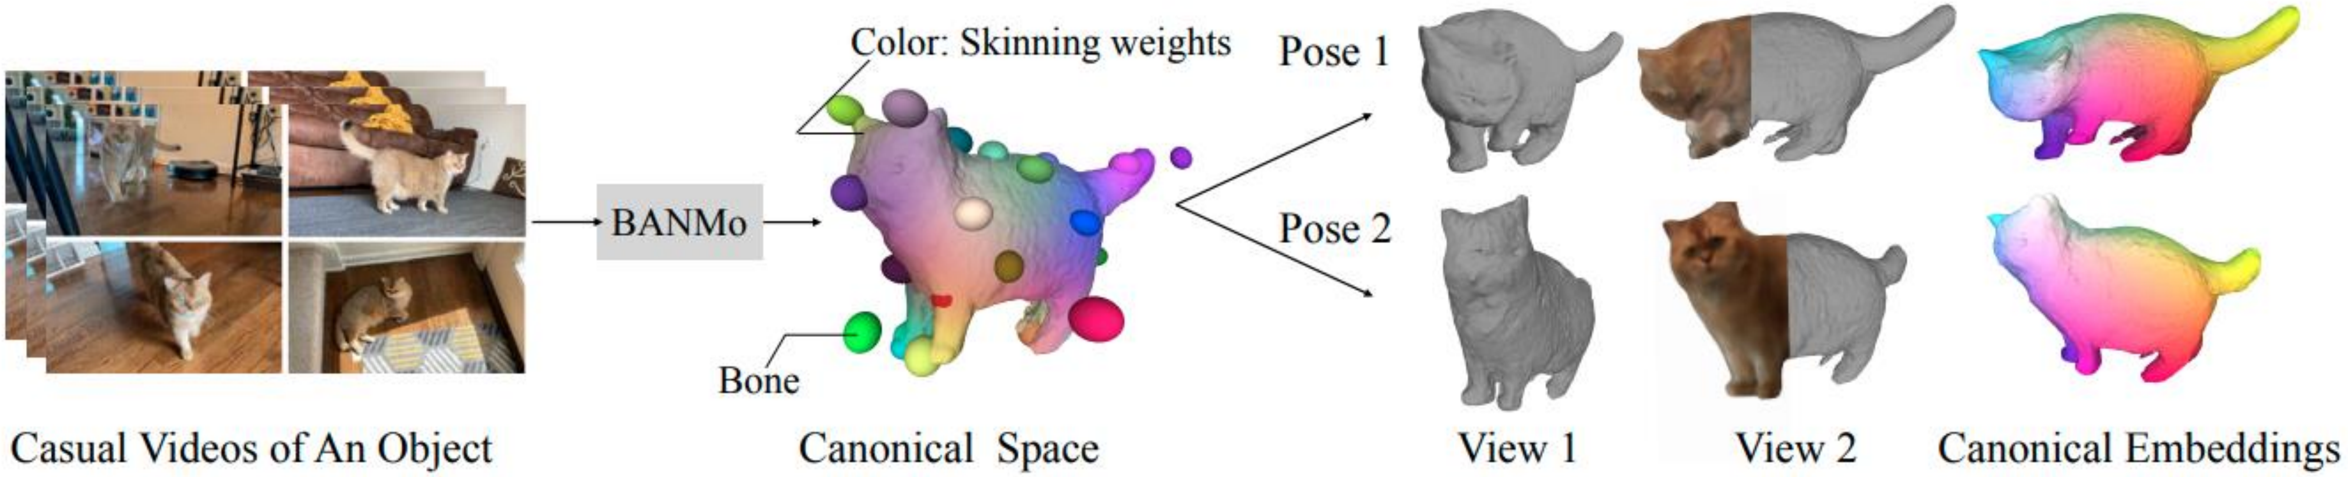

Supervisor: Dr. Vladislav Golyanik

GANs through space and time

Supervisor: Gereon Fox

GANs through space and time

A Good Image Generator is What You Need for High-Resolution Video Synthesis

(Tian et al., ICLR 2021)

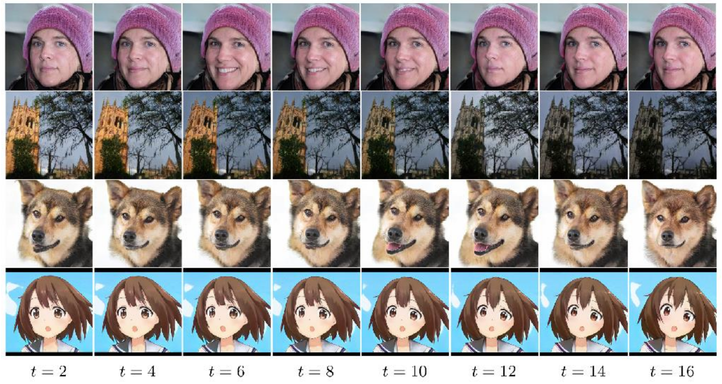

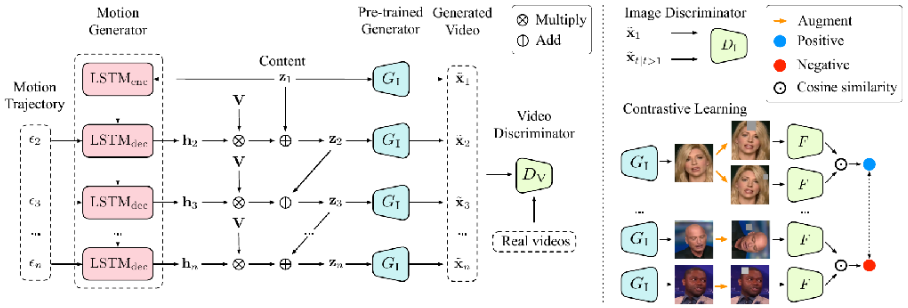

StyleVideoGAN: A Temporal Generative Model using a Pretrained StyleGAN

(Fox et al, BMVC 2021)

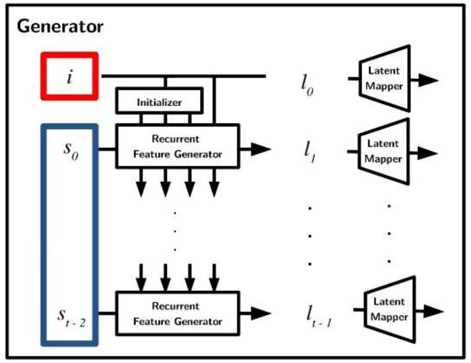

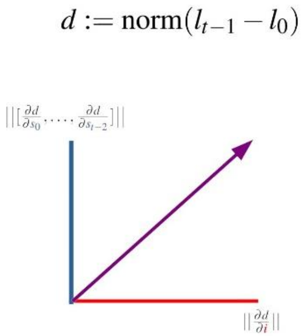

Supervisor: Gereon Fox

Coherent 3D Human Pose Estimation

Body Size and Depth Disambiguation in Multi-Person Reconstruction from Single Images

(Ugrinovic et al., 3DV 2021)

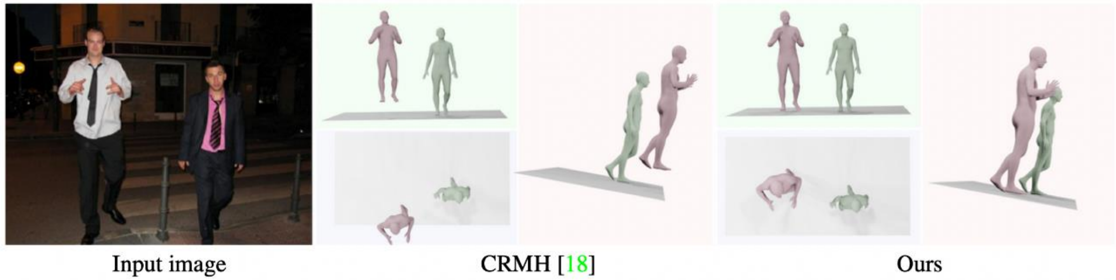

Putting People in their Place: Monocular Regression of 3D People in Depth

(Sun et al., arxiv 2021)

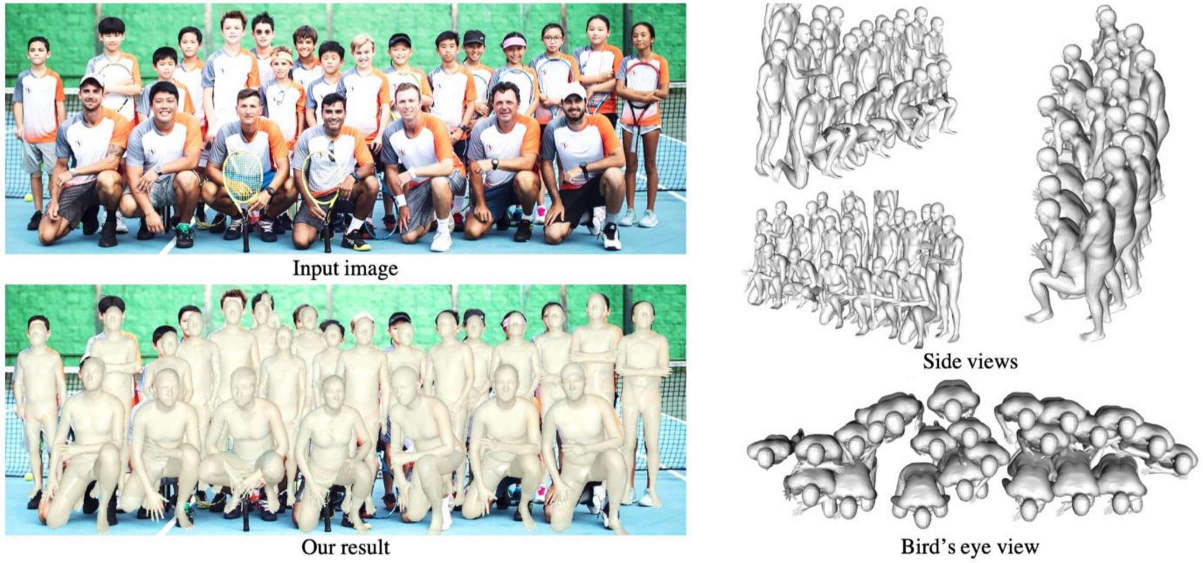

Supervisor: Diogo Luvizon

Tracking People with 3D Representations

(Rajasegaran et al., NeurIPS 2021)

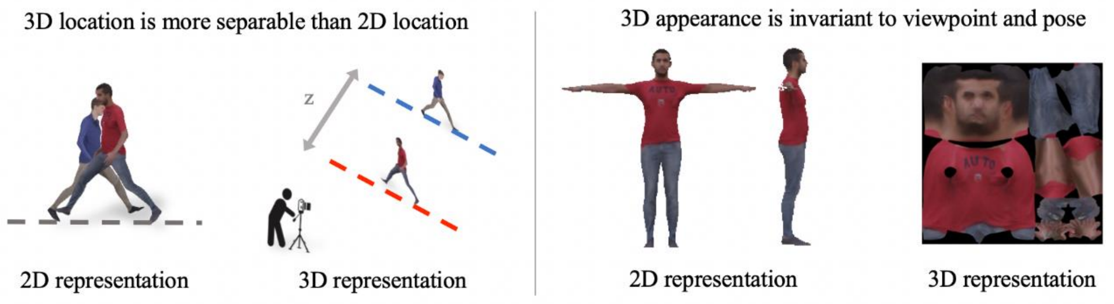

GLAMR: Global Occlusion-Aware Human Mesh Recovery **with Dynamic Cameras**

(Yuan et al., arxiv 2021)

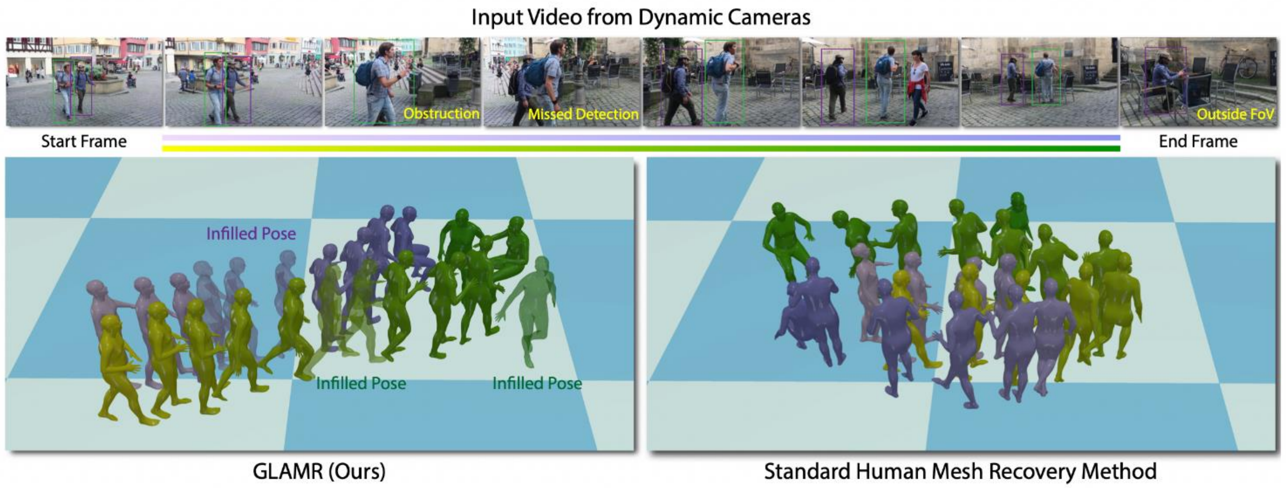

Supervisor: Diogo Luvizon

Hand Tracking

Probabilistic Modeling for Human Mesh Recovery

(Kolotouros et al., ICCV2021)

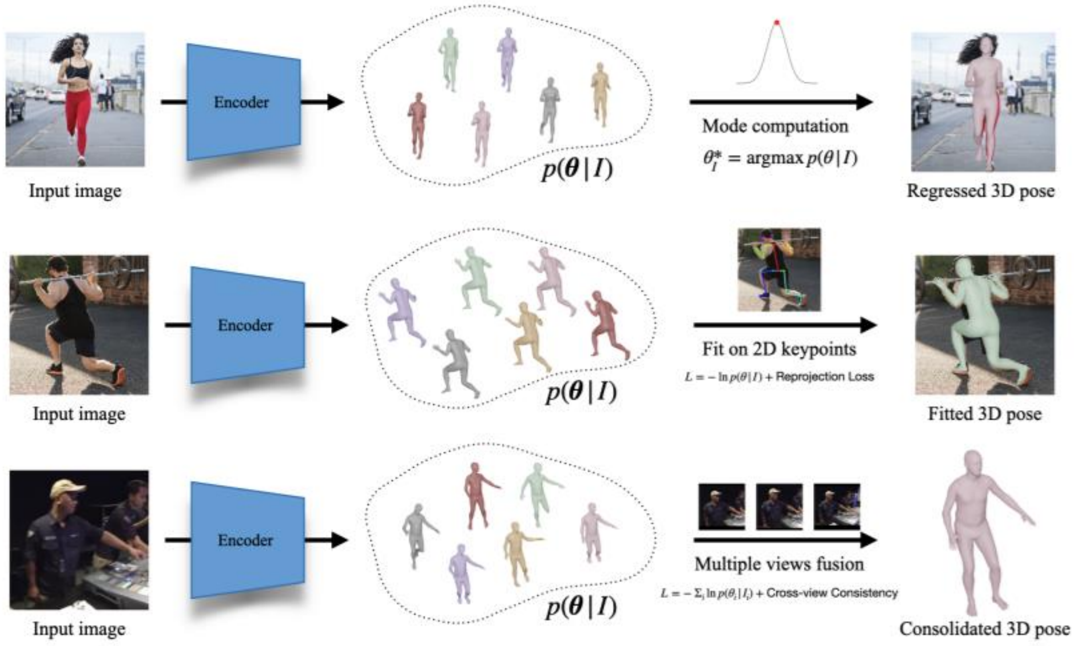

Probabilistic Monocular 3D Human Pose Estimation with Normalizing Flows

(Wehrbein et al., ICCV2021)

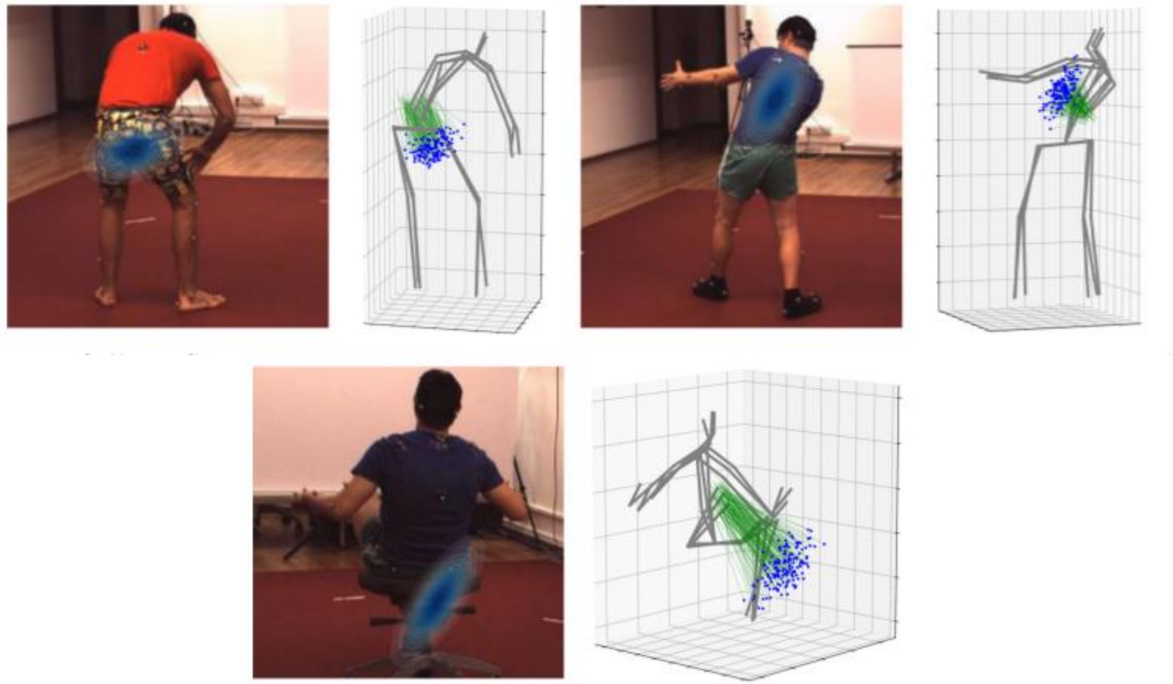

Supervisor: Jiayi Wang

NeRFactor: Neural Factorization of Shape and Reflectance

(Xiuming et al., TOG 2021)

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

(Kai et al., CVPR 2021)

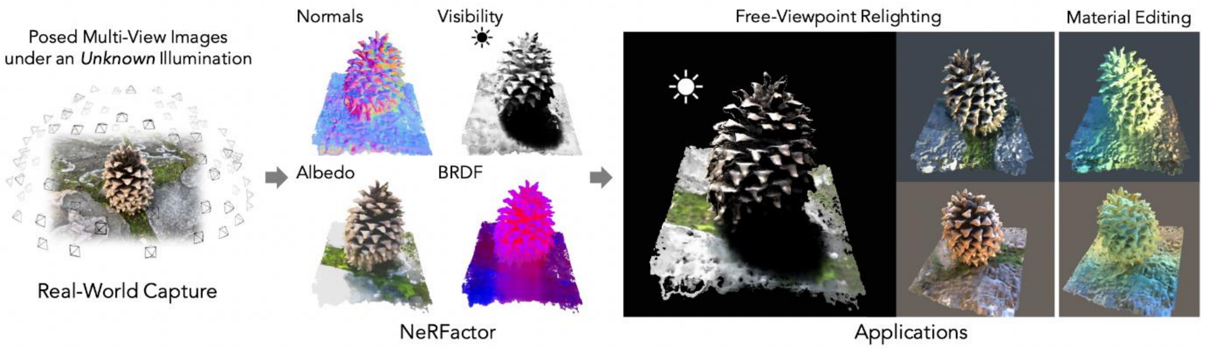

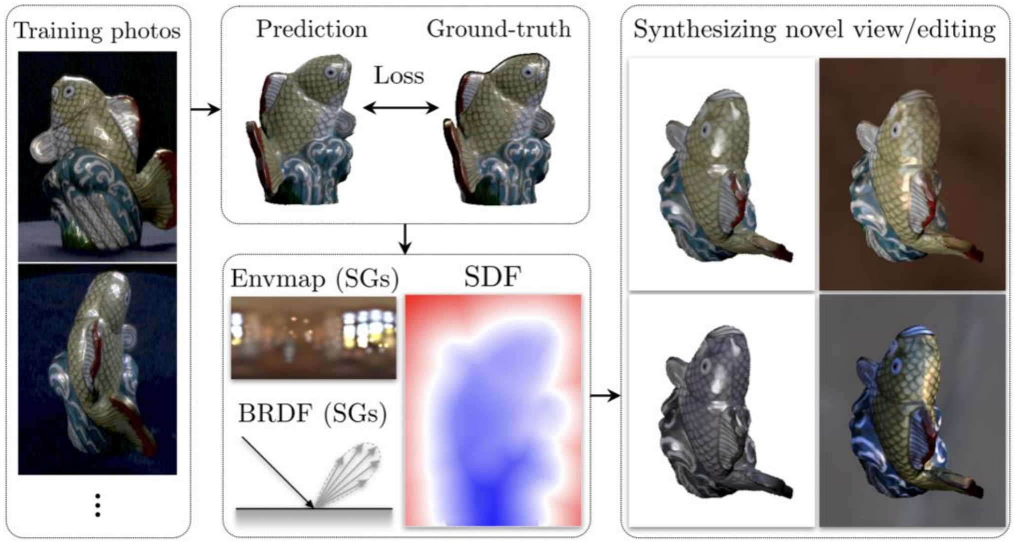

Supervisor: Linjie Lyu

Differentiable Physics Simulation for Inverse Problems

∇Sim: Differentiable Simulation for System Identification and Visuomotor Control (Jatavallabhula et al., ICLR 2021)

Inverse Simulation: Reconstructing Dynamic Geometry of Clothed Humans via Optimal Control (Guo et al., CVPR 2021)

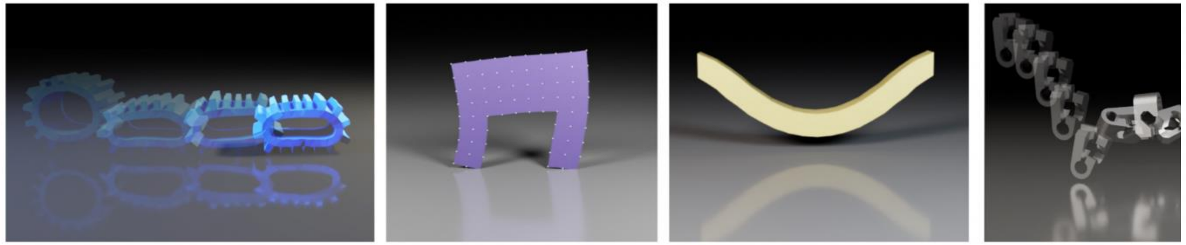

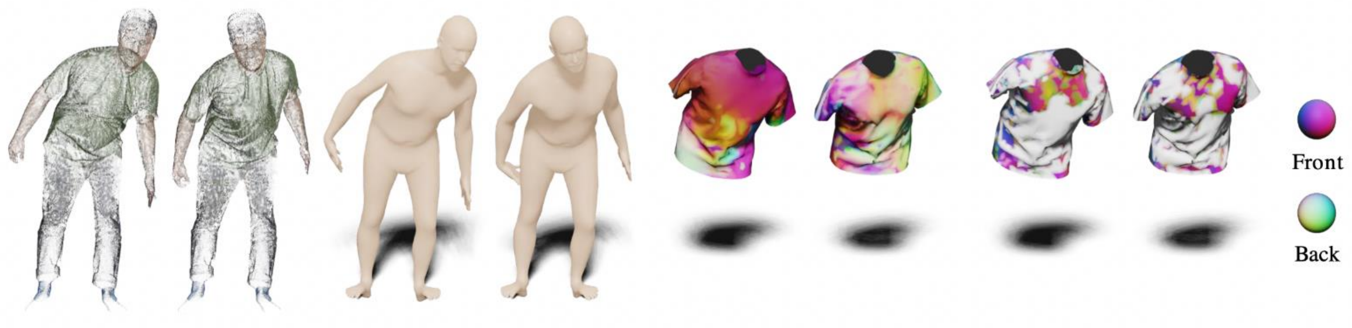

(a) Input point cloud (b) Estimated body (c) Velocity direction (d) Force direction

Supervisor: Navami Kairanda

Phase-Functioned Neural Networks for Character Control

(Holden et al., SIGGRAPH 2017)

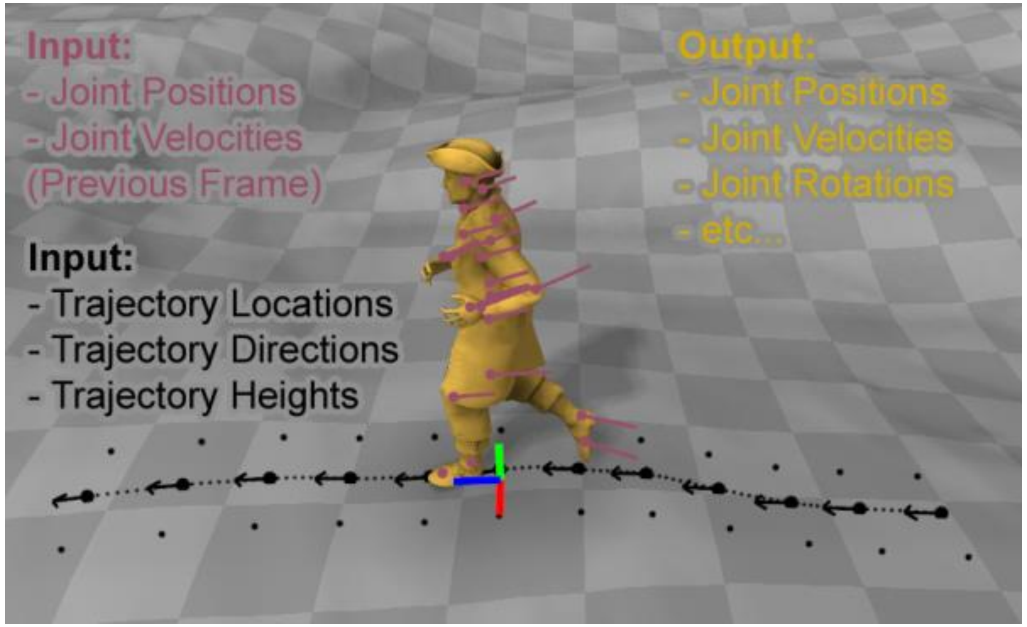

Character Controllers using Motion VAEs

(Ling et al., SIGGRAPH 2020)

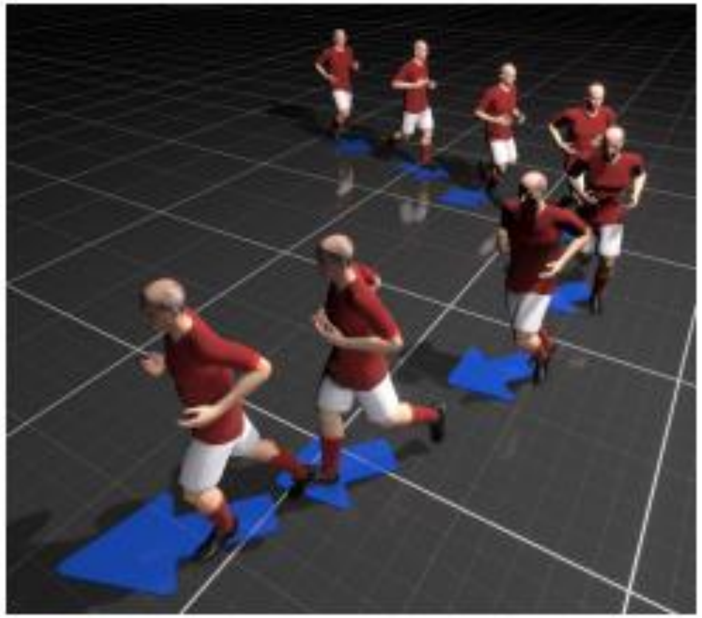

Supervisor: Ikhsanul Habibie

Discrete Neural Representation

Zero-Shot Text-to-Image Generation (Ramesh et al., PLMR 2021)

(a) a tapir made of accordion. a tapir with the texture of an accordion.

(b) an illustration of a baby hedgehog in a christmas sweater walking a dog

Taming Transformers for High-Resolution Image Synthesis

(Esser et al., SIGGRAPH 2021)

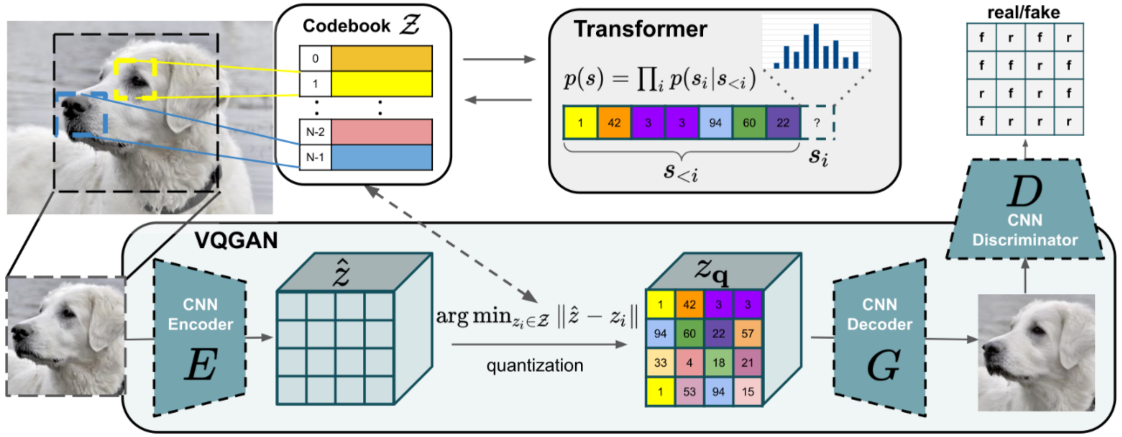

Supervisor: Ikhsanul Habibie

Event-based Pose Estimation

Visual Computing
and
Artificial Intelligence

max planck institut
informatik

EventCap: Monocular 3D Capture of High-Speed Human Motions using an Event Camera

(Xu et al., CVPR 2020)

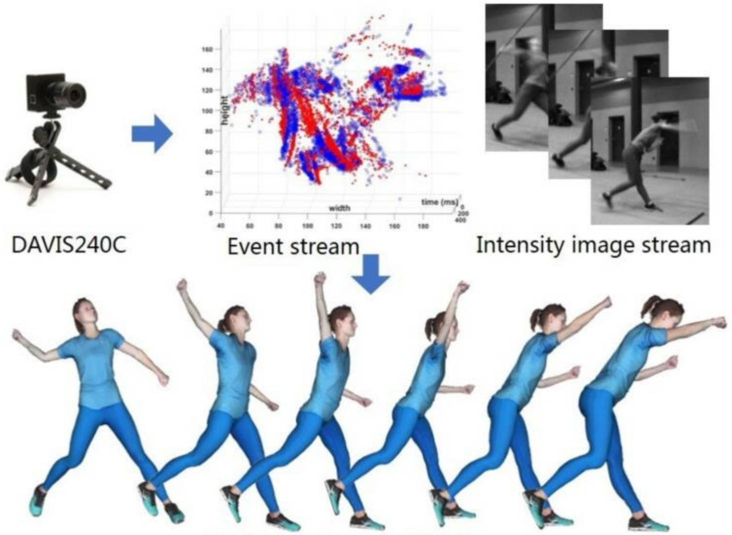

EventHands: Real-Time Neural 3D Hand Pose Estimation from an Event Stream

(Rudnev et al., ICCV 2021)

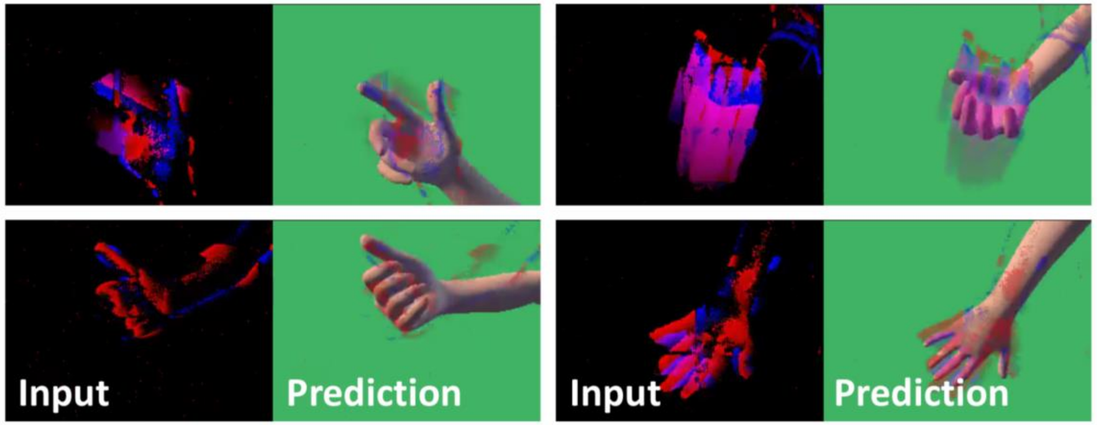

Supervisor: Viktor Rudnev

Learning NeRF from Few Views

pixelNeRF: Neural Radiance Fields from One or Few Images
(Yu et al., CVPR 2021)

RegNeRF: Regularizing Neural Radiance Fields for View Synthesis from Sparse Inputs
(Niemeyer et al., arXiv 2021)

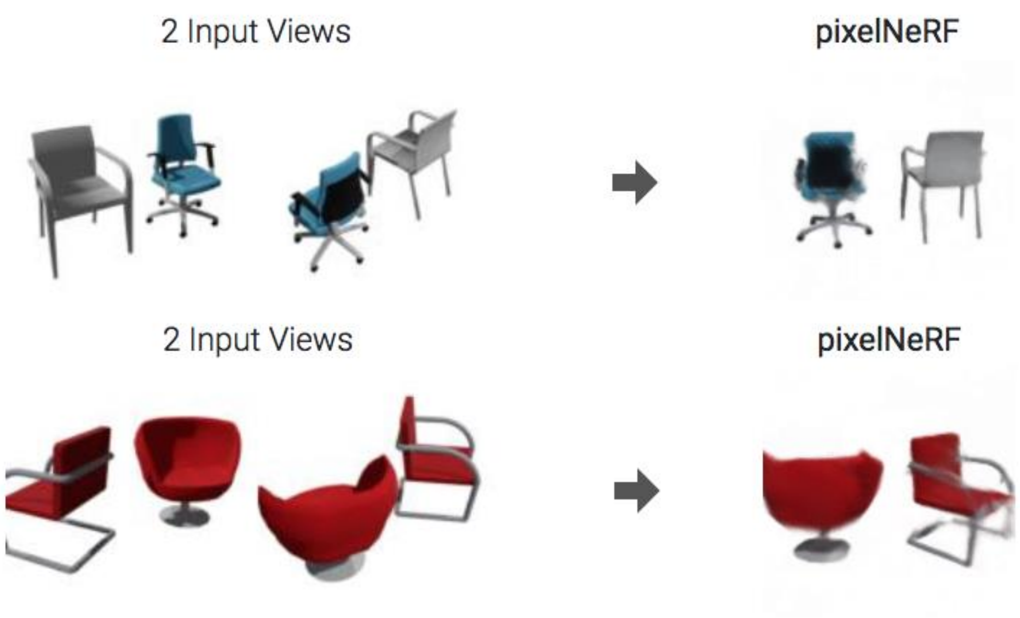

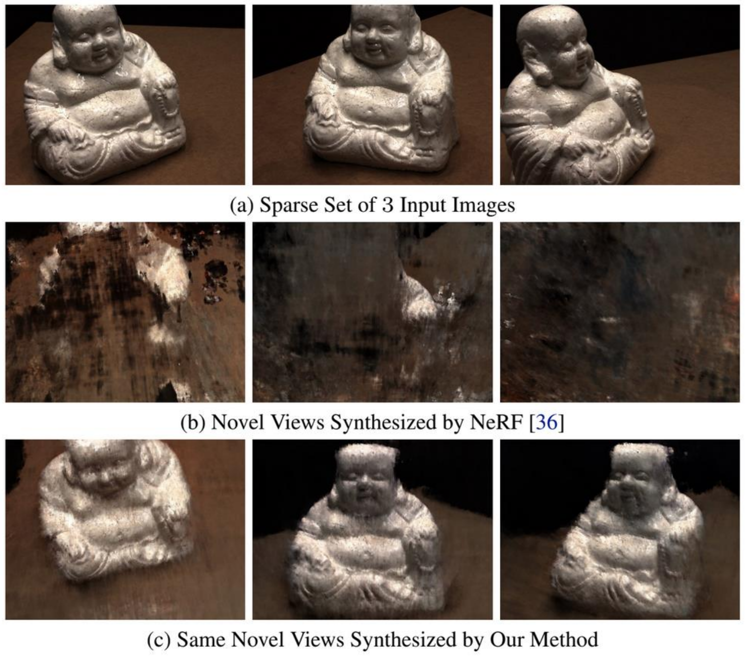

Supervisor: Viktor Rudnev

Inverse Graphics with Neuro-Explicit Models

Neural Body Fitting: Unifying Deep Learning and Model Based Human Pose and Shape Estimation
(Omran et al., 3DV 2018)

NeMo: Neural Mesh Models of Contrastive Features for Robust 3D Pose Estimation
(Wang et al., ICLR 2021)

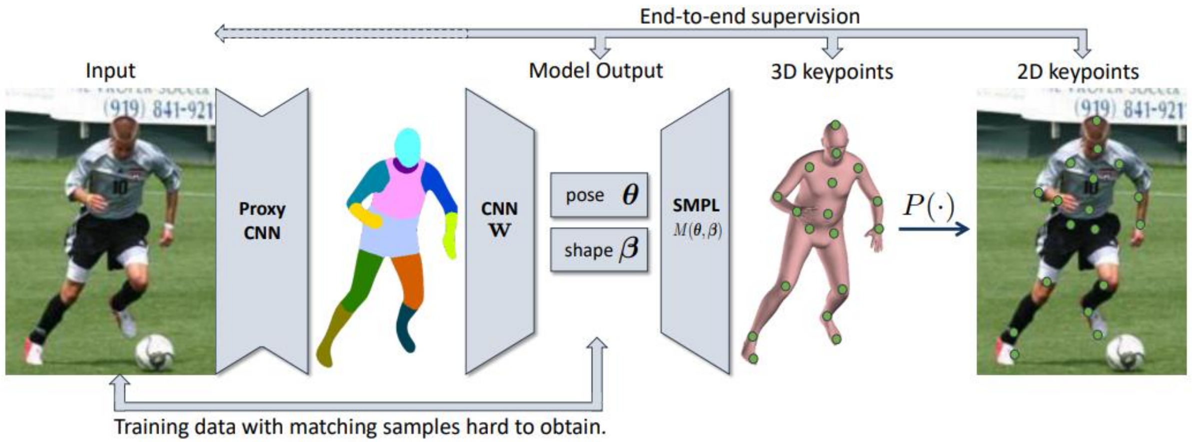

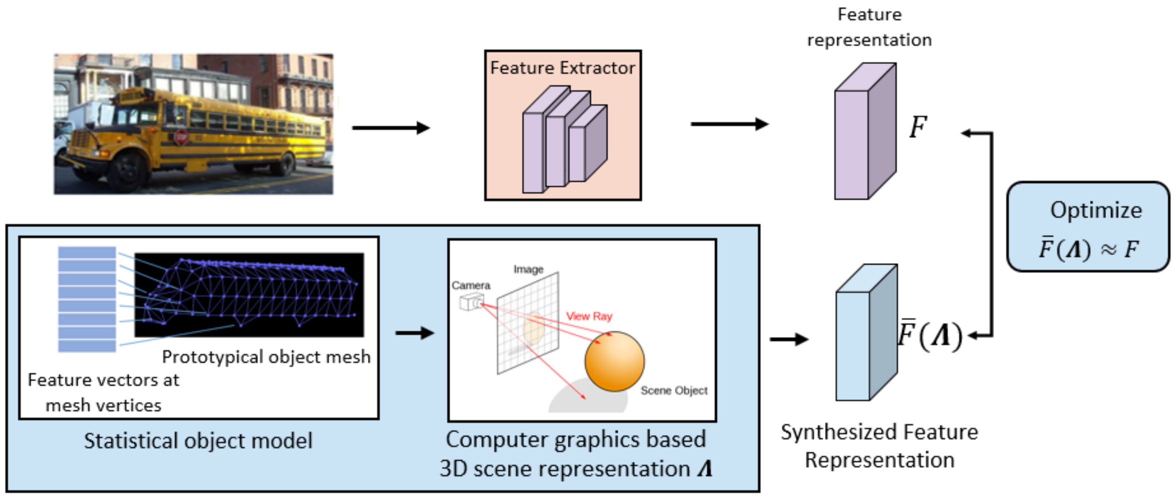

Supervisor: Dr. Adam Kortylewski

Model-based 3D Reconstruction of Faces

MoFA: Model-Based Deep Convolutional Face Autoencoder for Unsupervised Monocular Reconstruction
(Tewari et al., CVPR 2017)

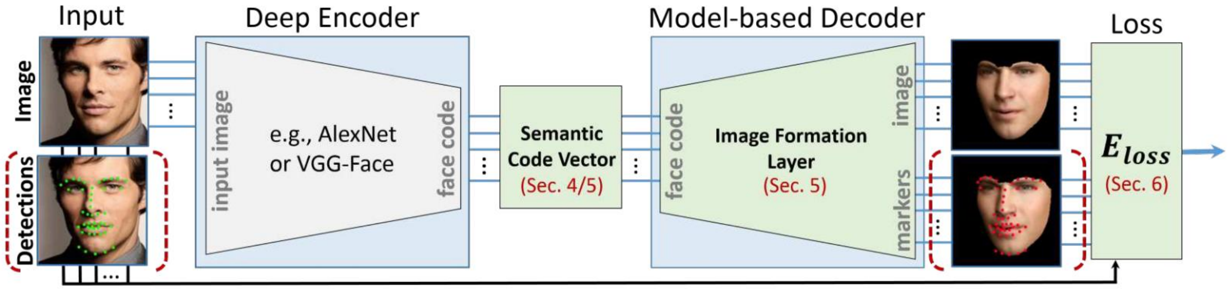

Markov Chain Monte Carlo for Automated Face Image Analysis
(Schoenborn et al., IJCV 2017)

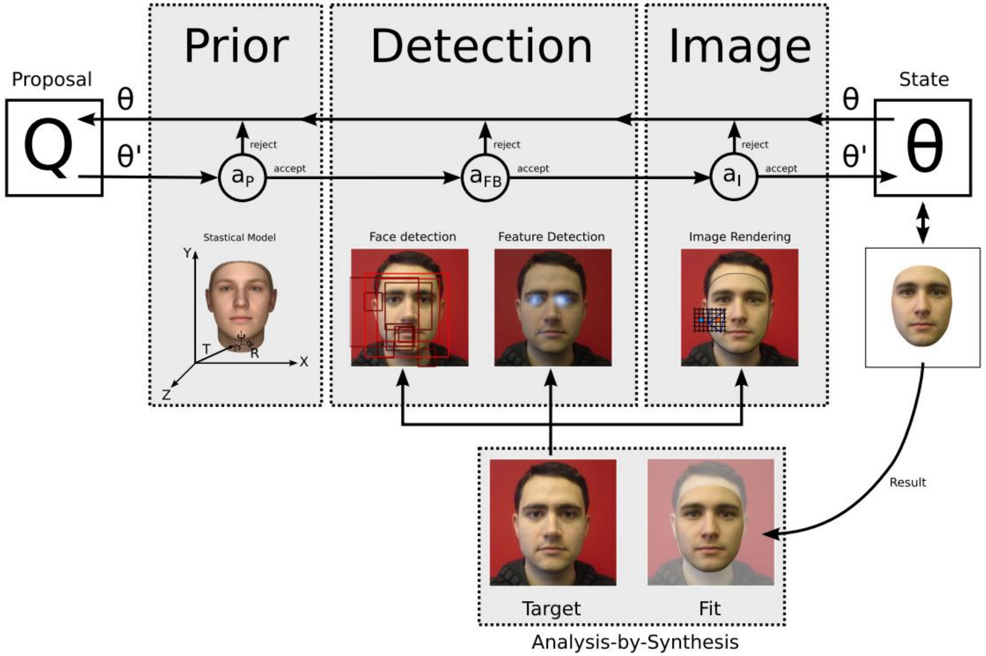

Supervisor: Dr. Adam Kortylewski

Photorealistic Face Relighting

Monocular Reconstruction of Neural Face Reflectance Fields

(B R et al., CVPR 2021)

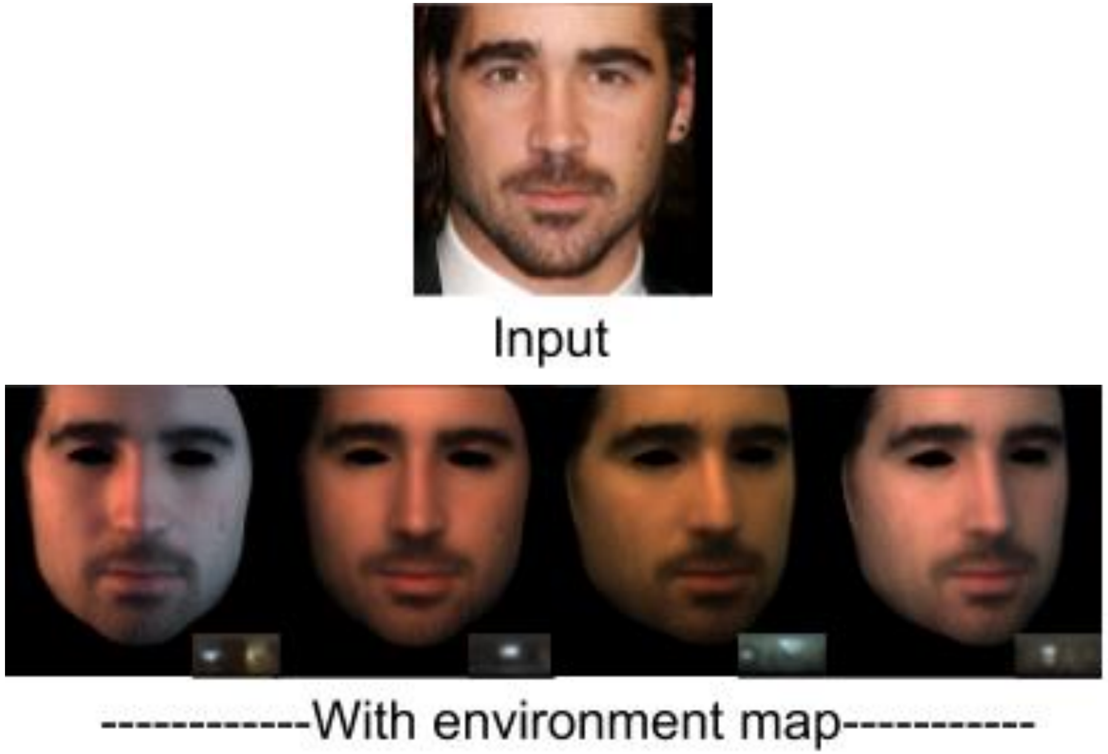

PhotoApp: Photorealistic Appearance Editing of Head Portraits

(B R et al., SIGGRAPH 2021)

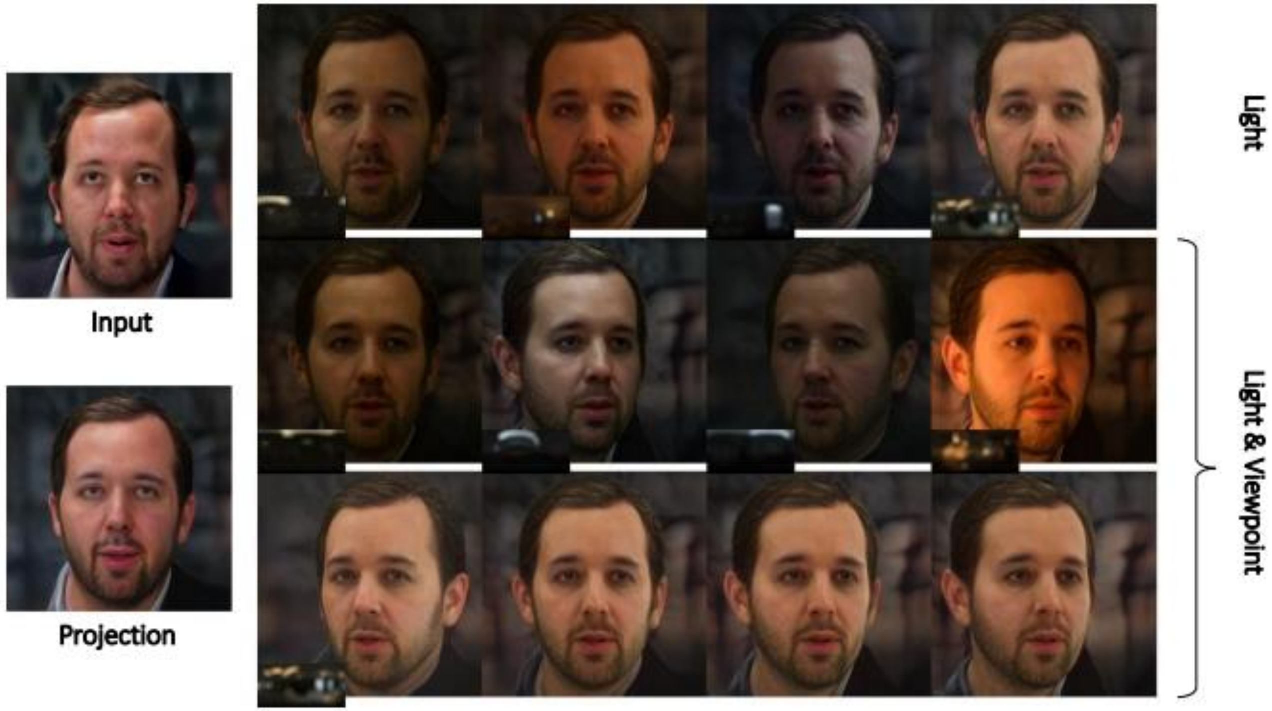

Supervisor: Dr. Mohamed Elgharib

Photorealistic Face Editing

PIE: Portrait Image Embedding for Semantic Control
(Tewari et al., SIGGRAPH (Asia) 2020)

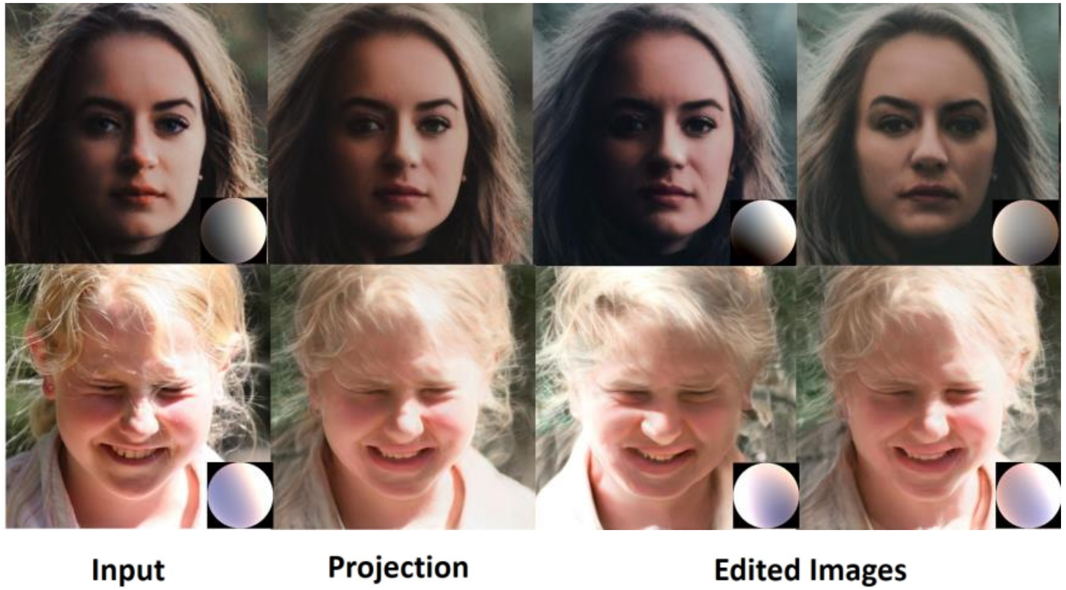

NerFACE: Dynamic Neural Radiance Fields for Monocular 4D Facial Avatar Reconstruction
(Gafni et al., CVPR 2021)

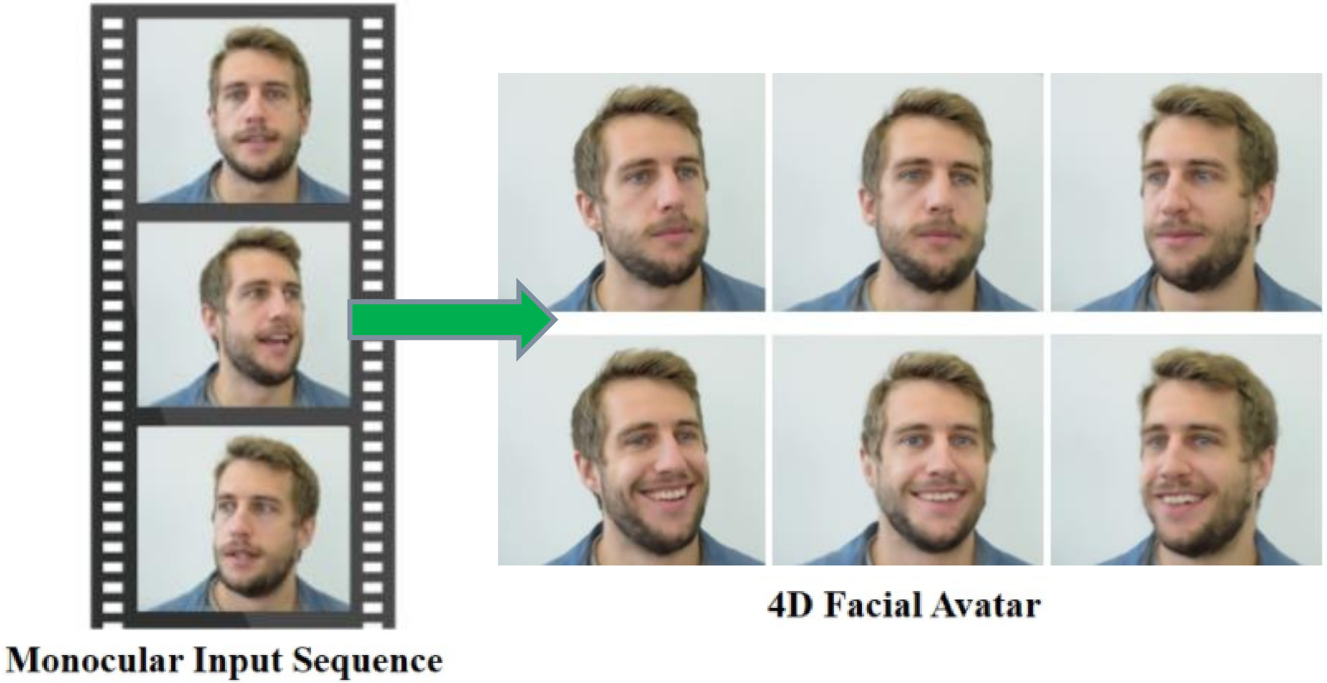

Supervisor: Dr. Mohamed Elgharib

Placing 3D People in 3D Scenes: Analyzing the Human-Scene Interaction in 3D Space

Resolving 3D Human Pose Ambiguities with 3D Scene Constraints

(Hassan et al. ICCV 2019)

Generating 3D People in Scenes without People

(Zhang et al. CVPR, 2020)

Supervisor: Jian Wang

Egocentric Human Body Pose Estimation

Mo2Cap2: Real-time Mobile 3D Motion Capture with a Cap-mounted Fisheye Camera

(Xu et al., IEEE VR 2019)

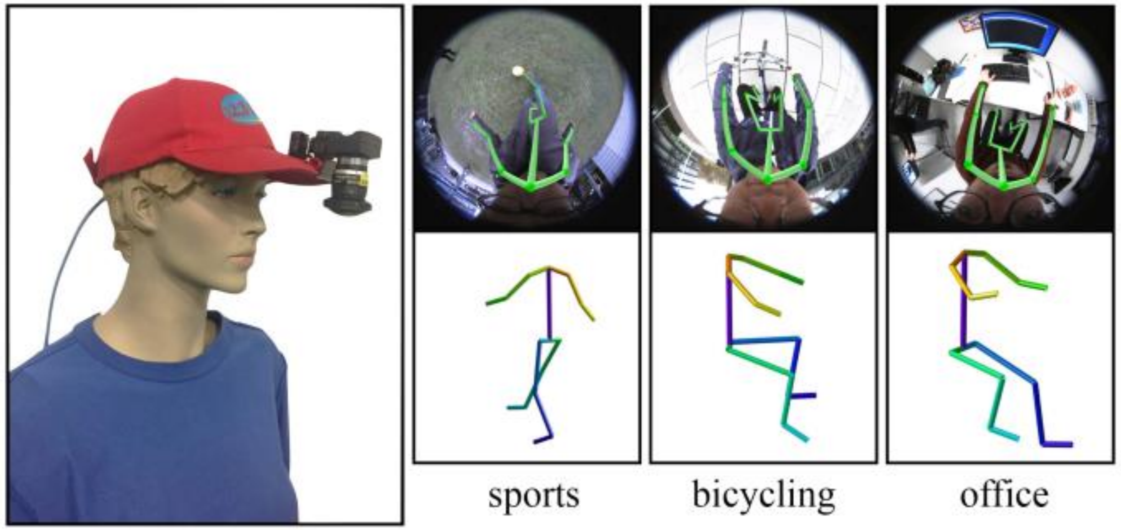

Estimating Egocentric 3D Human Pose in Global Space.

(Wang et al., ICCV 2021)

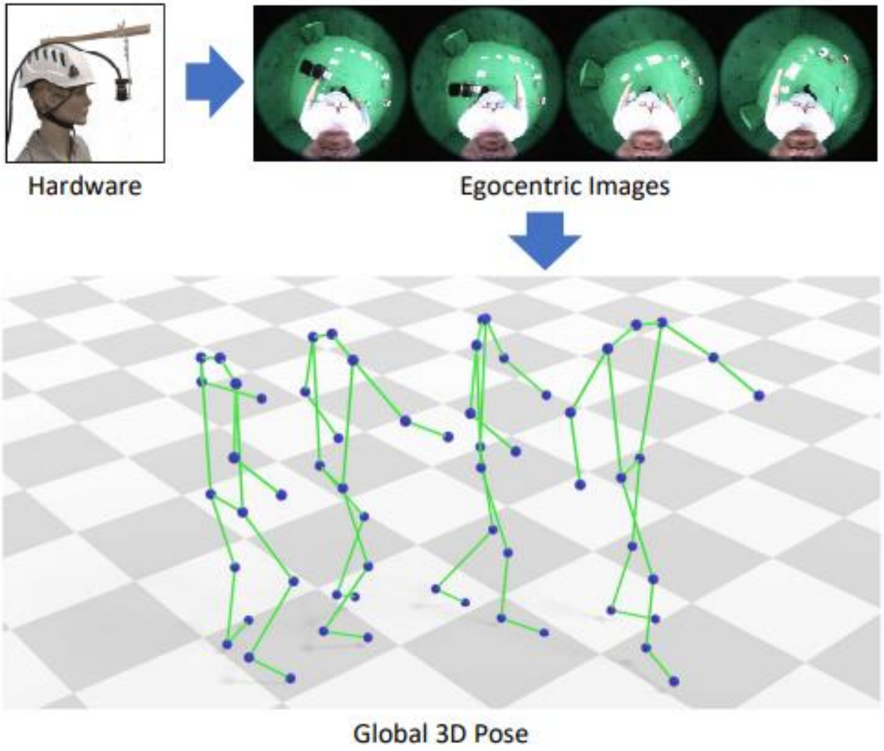

Supervisor: Jian Wang

Human Motion and Language

Synthesis of Compositional Animations from Textual Descriptions
(Ghosh et al. ICCV2021)

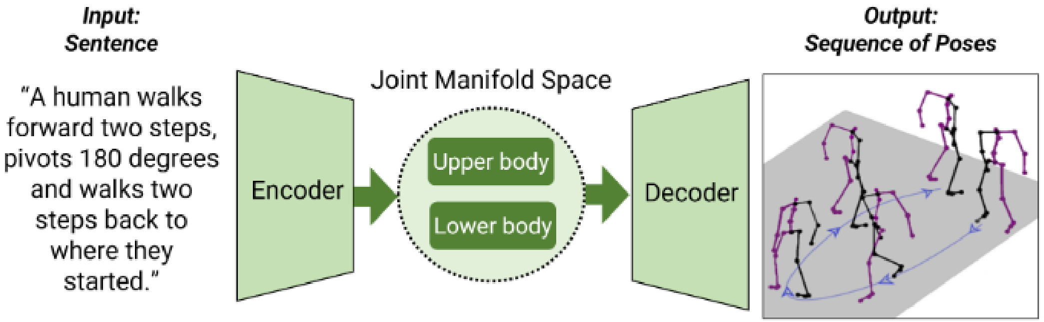

Action-Conditioned 3D Human Motion Synthesis with Transformer VAE
(Petrovich et al. ICCV2021)

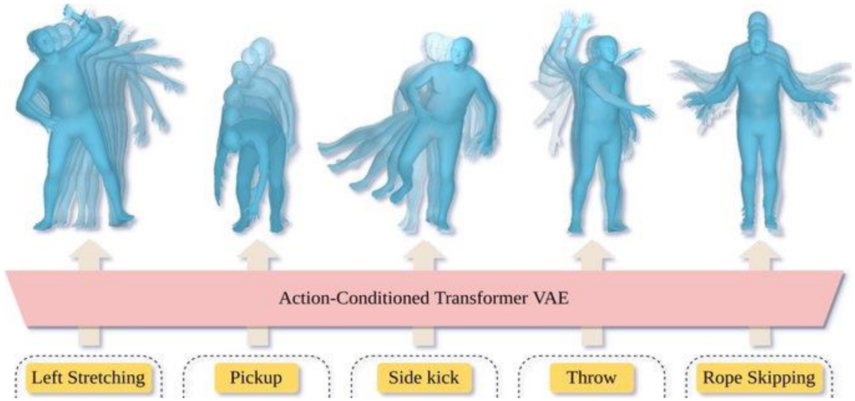

Supervisor: Rishabh Dabral

3D Body Models

Visual Computing
and
Artificial Intelligence

max planck institut
informatik

GHUM & GHUML: Generative 3D Human Shape and Articulated Pose Models
(Xu et al. CVPR2020)

imGHUM: Implicit Generative Models of 3D Human Shape and Articulated Pose
(Alldieck et al. ICCV2021)

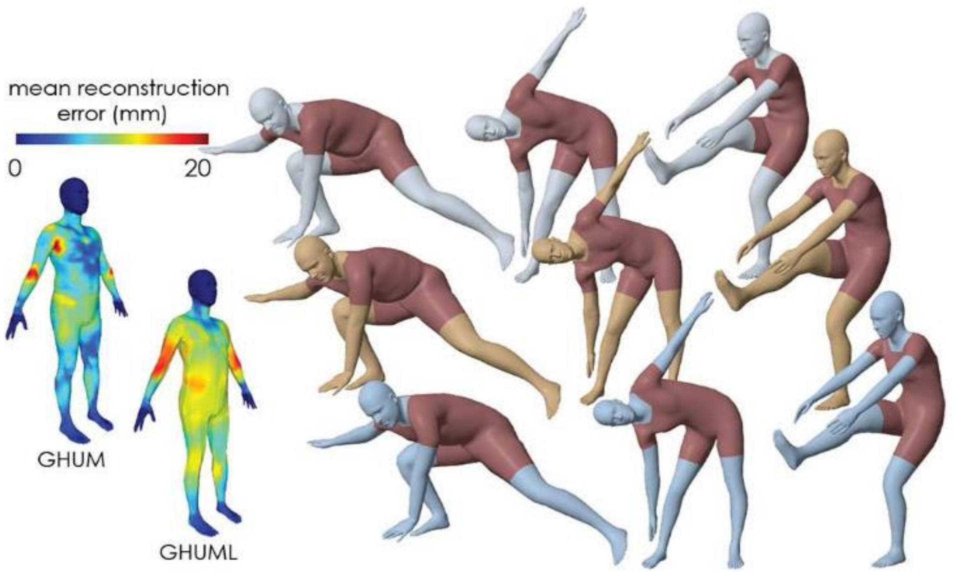

Figure 1. **imGHUM** is the first parametric full human body model represented as an implicit signed distance function. **imGHUM** successfully models broad variations in pose, shape, and facial expressions. The level sets of **imGHUM** are shown in blue-scale.

Supervisor: Rishabh Dabral

Compositional Models

Learning Object-Compositional Neural Radiance Field for Editable Scene Rendering
(Yang et al., ICCV 2021)

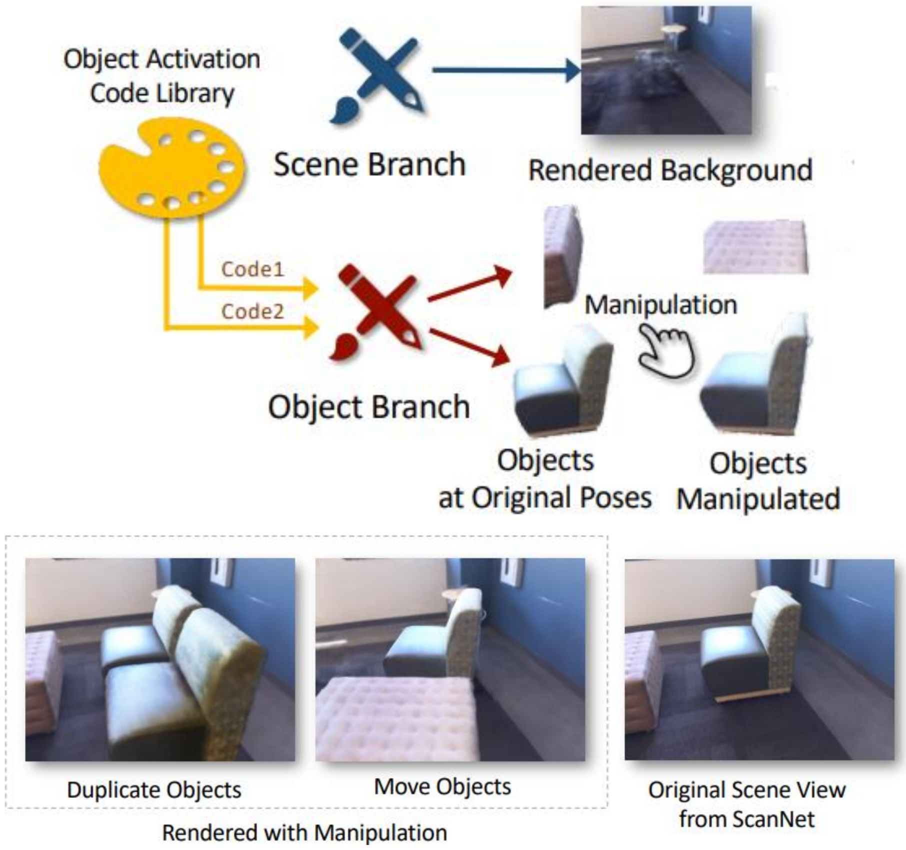

GIRAFFE: Representing Scenes as Compositional Generative Neural Feature Fields

(Niemeyer et al., CVPR 2021)

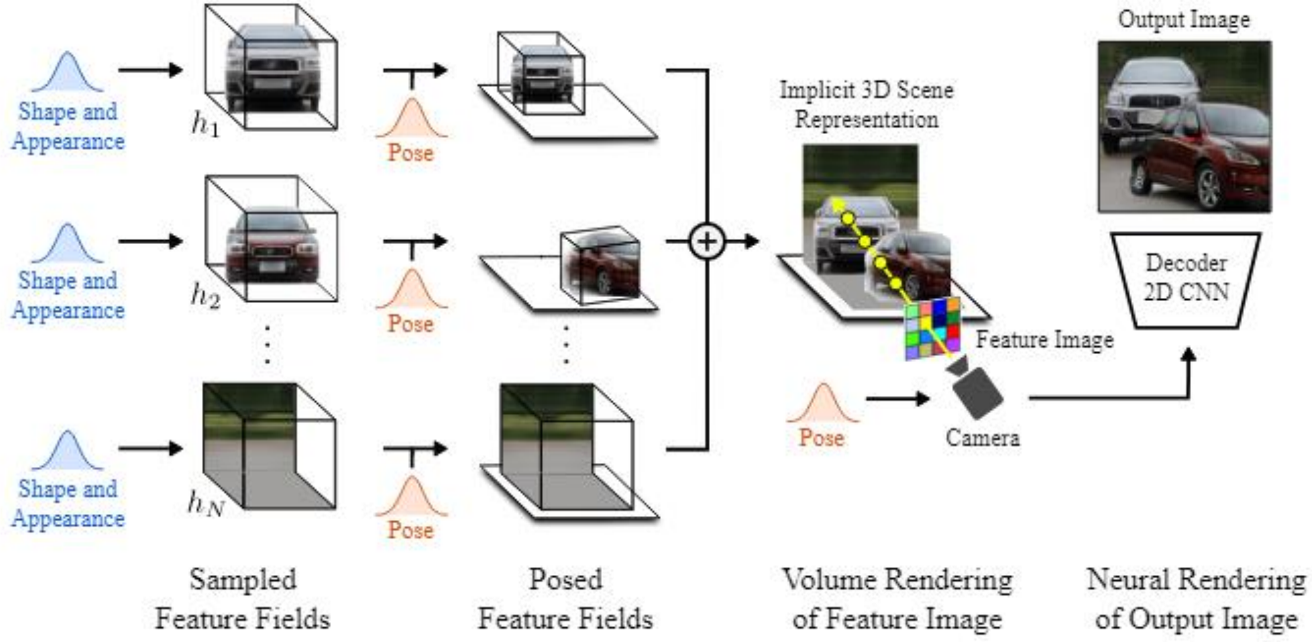

Supervisor: Mallikarjun B R

Overview

- Organization
- Introduction
- Topics
- Summary

Overview

- Organization
- Introduction
- Topics
- Summary

Summary: What's Next

- Topic assignment:
  - Wait for central admission to this seminar (**Friday, 15 April**)
  - Send me a list of 3 topics (in order of preference) until **Tuesday, 19 April**: thomas.leimkuehler@mpi-inf.mpg.de
  - We will try to accommodate wishes as much as possible
- Next week (**Thursday, 21 April**):
  - Topics will be assigned
  - Lecture 1: "How to read an academic paper"
  - Lecture 2: "How to give a good scientific talk"
- First topic presentation: **Thursday, 28 April**

Thank you!

Any questions?