Classical Concepts of Computer Vision and Computer Graphics in the Neural Age

Seminar - Summer Semester 2024

Thomas Leimkühler, Marc Habermann, Rishabh Dabral, Christian Theobalt

MAX PLANCK INSTITUTE FOR INFORMATICS

UNIVERSITÄT DES SAARLANDES

Organizers

Marc Habermann
MPI Informatik, Office 216
mhaberma@mpi-inf.mpg.de

Christian Theobalt
MPI Informatik, Office 224
theobalt@mpi-inf.mpg.de

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

Rishabh Dabral
MPI Informatik, Office 221
rdabral@mpi-inf.mpg.de

Basic Coordinates

Time: Thursdays, 14:15 – 15:45

Location: E 1 5 Room 630 (6th floor)

In-person only (no virtual version)

Website: https://vcai.mpi-inf.mpg.de/teaching/vcai_seminar_2024/index.html

Content of this seminar

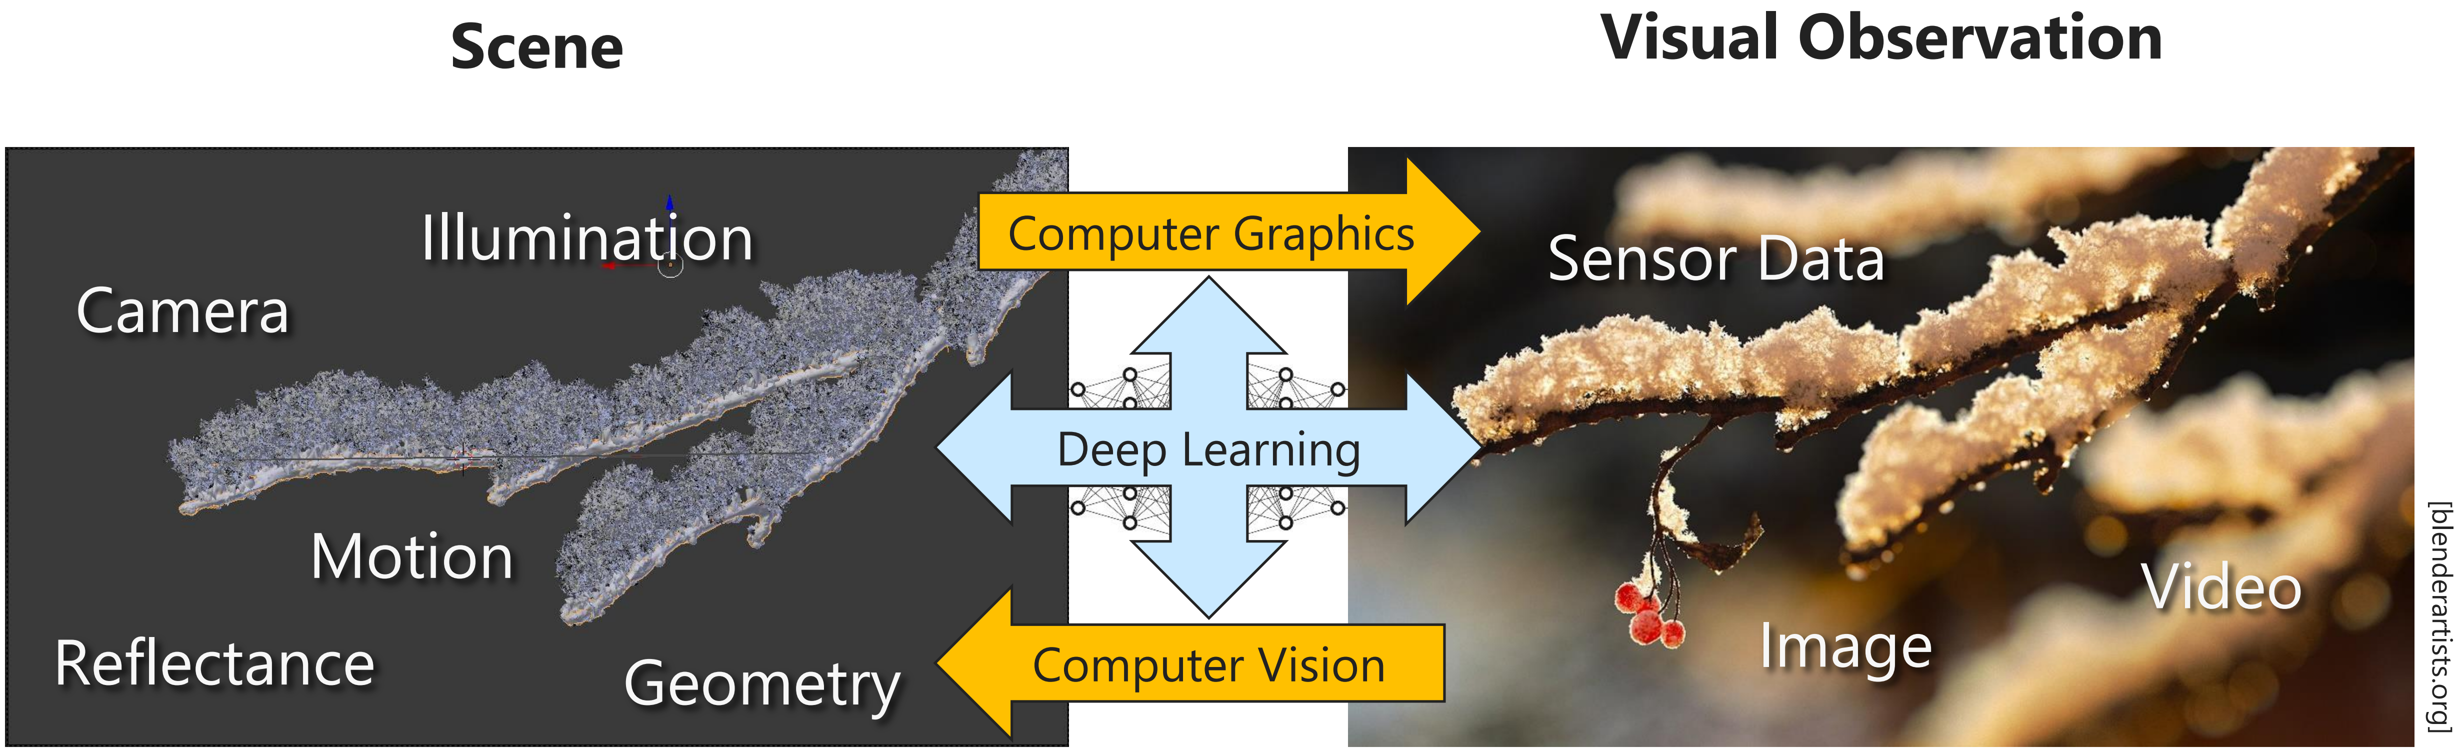

Content of this seminar

- **Computer Graphics** and **Computer Vision** have witnessed the transformative effects of Deep Learning → the **Neural Age**
- Neural methods have…

  … changed how we think about problems in visual computing

  … significantly propelled the field forward
- We discuss seminal papers that shaped the field, always paired with a recent paper that has further developed the idea in a modern context
- Previous editions of this seminar were called “Computer Vision and Machine Learning for Computer Graphics” → slightly different focus this time

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%

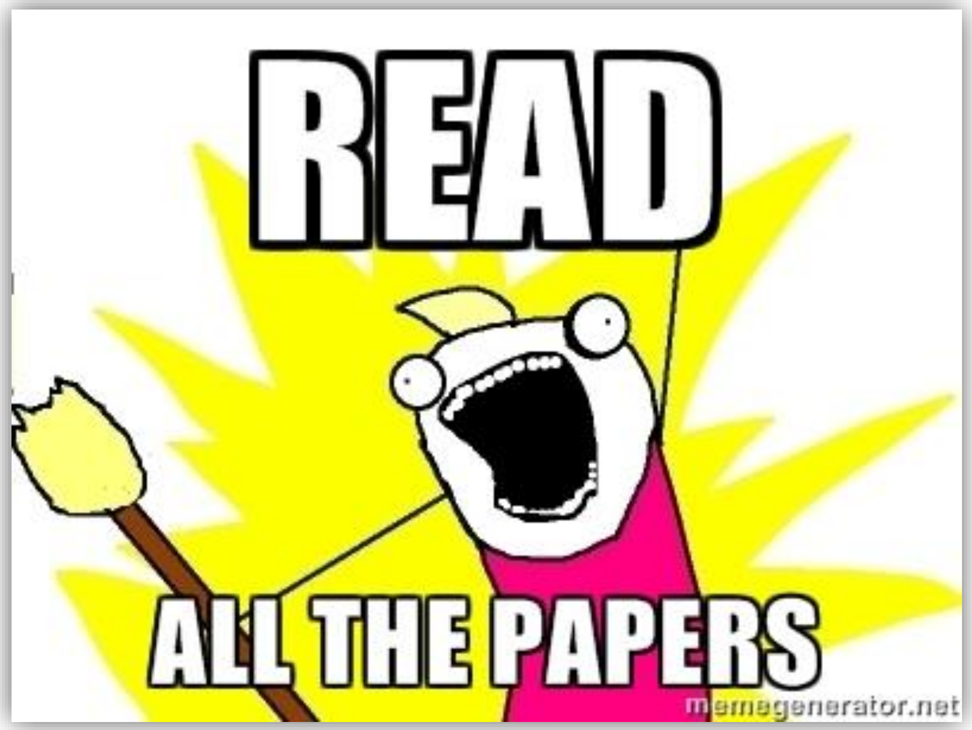

Requirements

- **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.

Topics

- 20 topics to choose from (listed on seminar website + introduced later today)
- 12-13 presentation slots
- First presentation: Thursday, 2 May 2024
- Each week until Thursday, 25 July 2024 (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

Seminar Session

- Order of presentation will be determined after topic assignment

- Slots can be swapped if necessary: talk to other participants first

- **Talk (40 minutes):**
  - 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

- **Moderator summarizes the content and asks for feedback about the presentation (10 minutes)**

- **Paper discussion (40 minutes)**

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

Moderator and Discussion

- **For each week (till Wednesday 8am):**

  Everyone submits 2+ questions for discussion to mhaberma@mpi-inf.mpg.de

  Important: your contribution will be marked

  One person assigned to be the moderator, who will lead the discussion

  Will receive the collected questions submitted before the seminar

- **At the seminar:**

  Moderator gives a summary of the talk (5 minutes)

  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, 22 August 2024** (4 weeks after the last seminar)
- Send PDF to rdabral@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 (2 May) 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

- 11 April – Introduction ◄ You are here
- 18 April – Lecture: “How to read an academic paper”
- 25 April – Lecture: “How to give a good talk”
- 2 May – First presentation by a student

  … 12-13 more weekly presentations …

- 25 July – Last presentation by a student

  Possibly one more session in between

- 22 August – Report deadline

Spherical Harmonics Lighting

Sloan et al. 2002.
Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments.
ToG, 2002.

Chen et al. 2020.
A Neural Rendering Framework for Free-Viewpoint Relighting .
CVPR 2020.

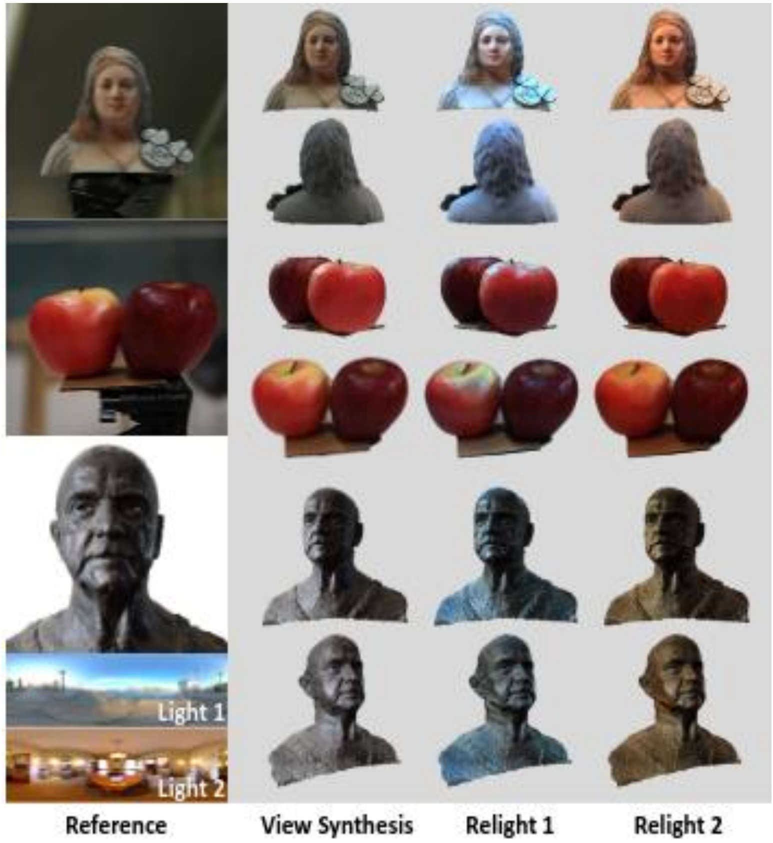

Supervisor: Pulkit Gera

Spatial Regularization

Sorkine et al. 2004.
Laplacian Surface Editing. SGP 2004.

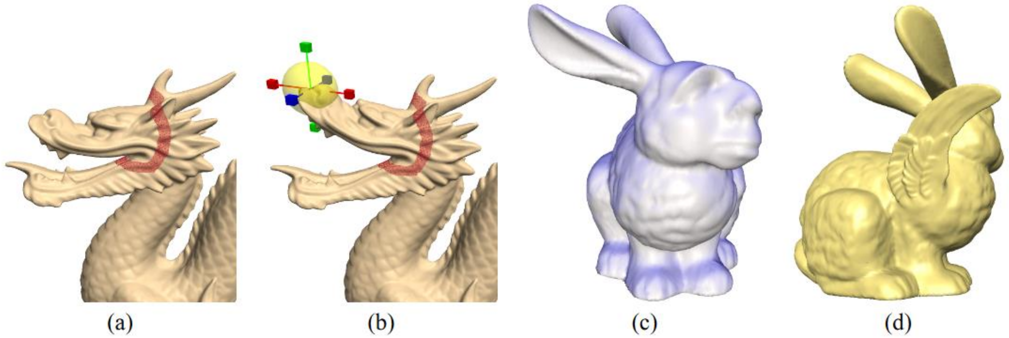

Figure 1: Advanced mesh editing operations using Laplacian coordinates: free-form deformations (a-b), coating transfer (c) and mesh transplanting (d). Representing the geometry using the Laplacian coordinates enables preservation of detail.

Zhang et al. 2022.
Critical Regularizations for Neural Surface Reconstruction in the Wild. CVPR 2022.

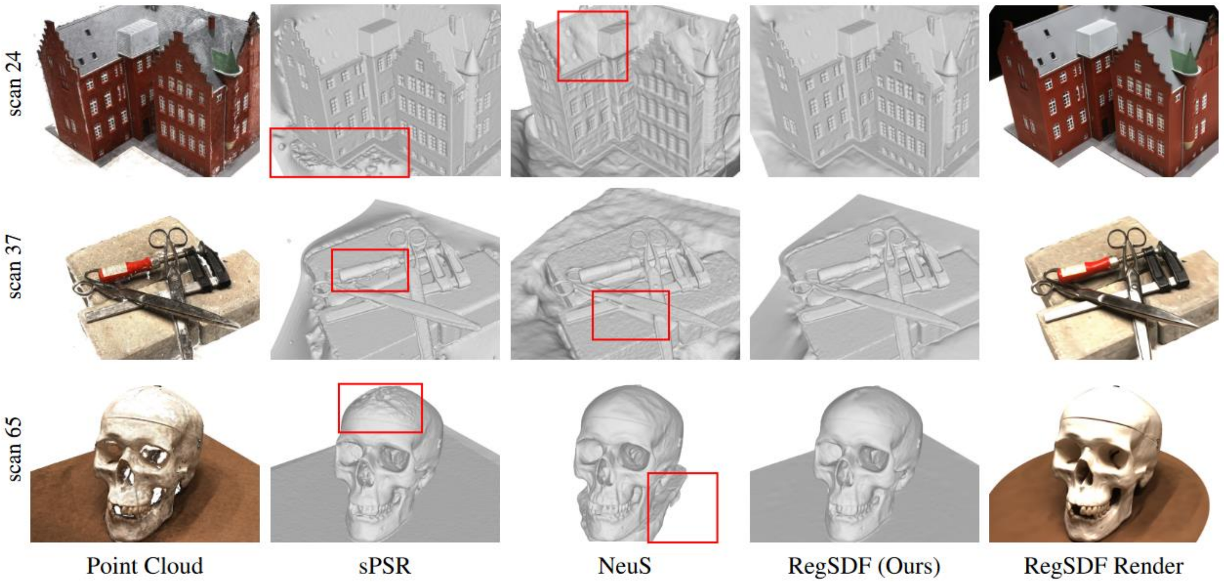

Supervisor: Mohit Mendiratta

Deformation Representations and Parameterizations

Sumner et al., Embedded deformation for shape manipulation, ToG 2007

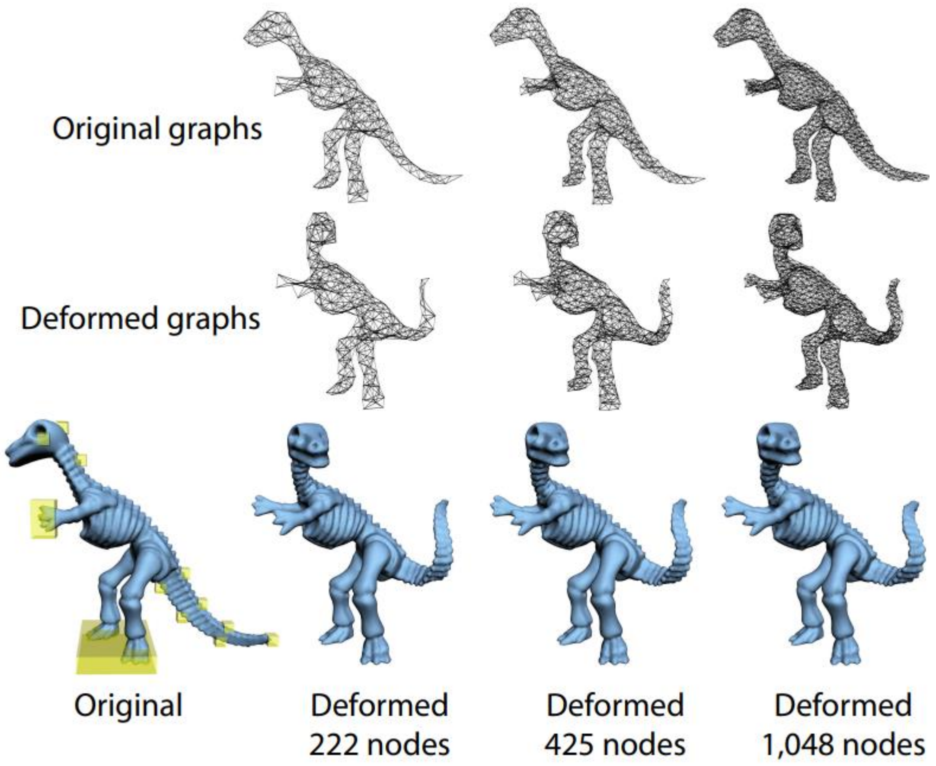

Xu et al., Deforming Radiance Fields with Cages, ECCV 2022

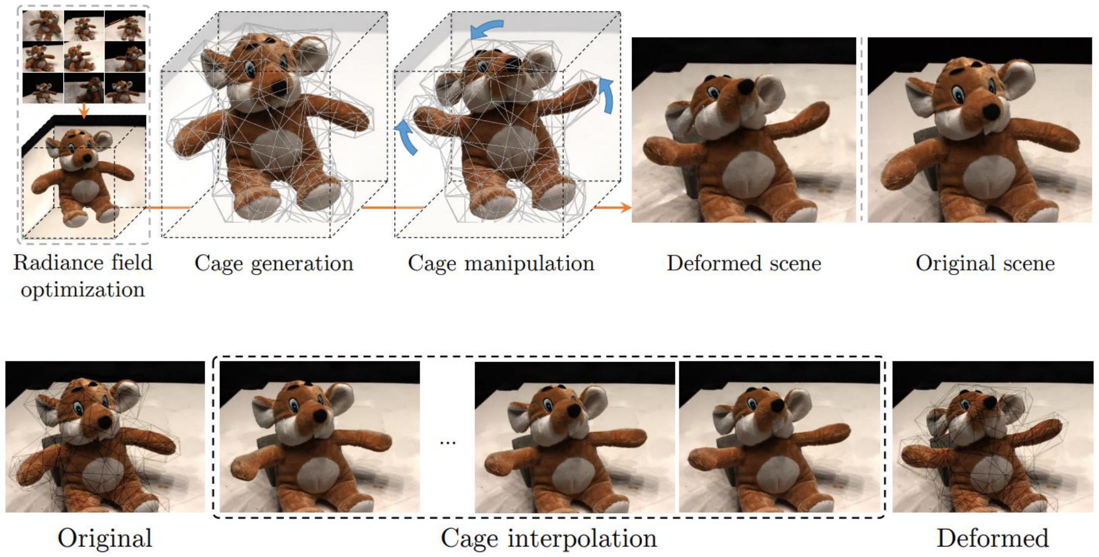

Supervisor: Guoxing Sun

Light Fields

Levoy et al.
Light Field Rendering. Siggraph 1996.

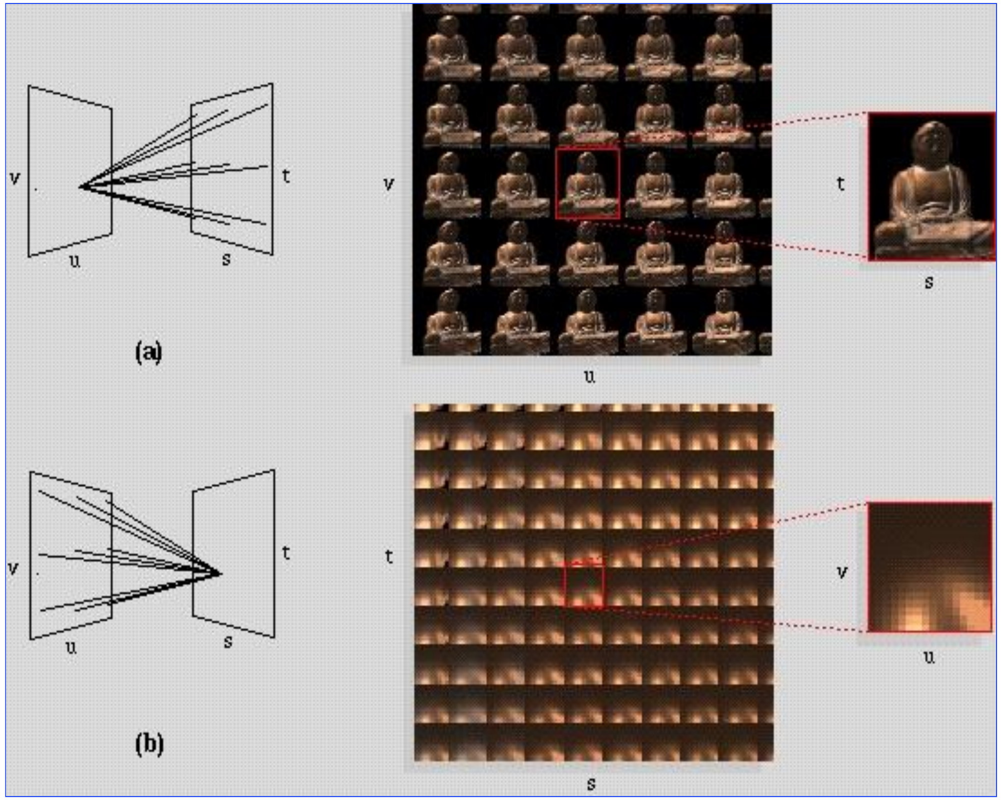

Sitzsmann et al.
Light Field Neural Network. NeurIPS 2021.

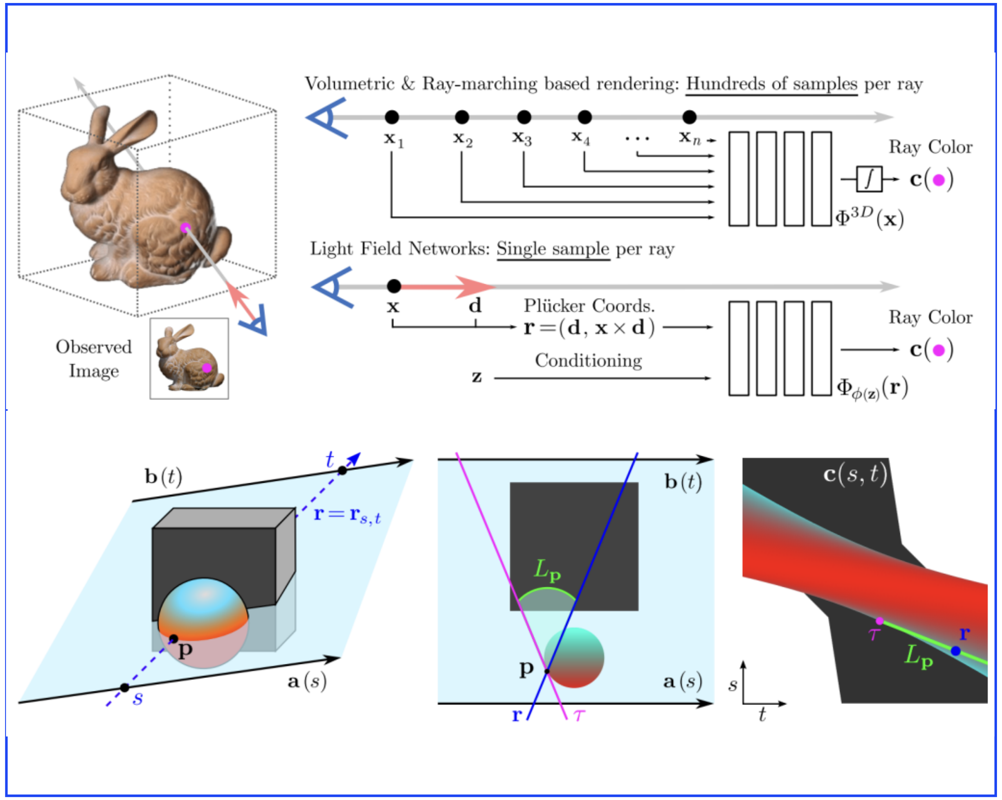

Supervisor: Jianchun Chen

Gaussians in Graphics

Stoll et al. 2011.
Fast Articulated Motion Tracking using a Sums of Gaussians Body Model. ICCV 2011.

Kerbl and Kopanas et al. 2023.
3D Gaussian Splatting for Real-Time Radiance Field Rendering. ACM ToG 2023.

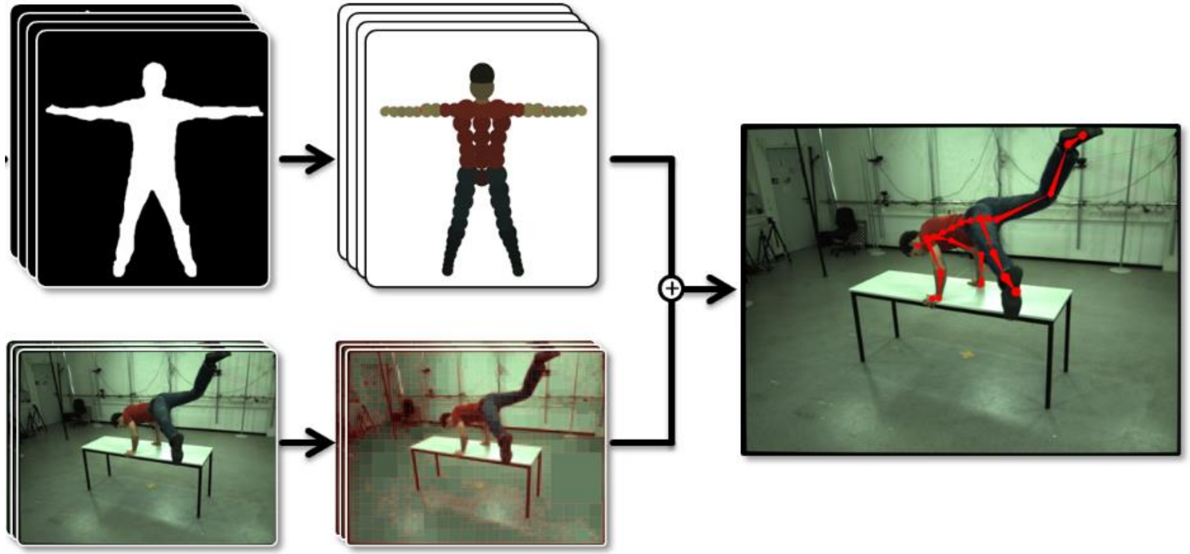

Supervisor: Navami Kairanda

Performance Capture

de Aguiar et al.,
Performance Capture from Sparse Multi-view Video,
SIGGRAPH 2008.

Saito et al.,
PIFuHD: Multi-Level Pixel-Aligned Implicit Function for
High-Resolution 3D Human Digitization, CVPR 2020.

Supervisor: Hiroyasu Akada

Warp Fields and Fusion

Newcombe et al. 2015.
DynamicFusion: Reconstruction and Tracking of Non-rigid Scenes in Real-Time. CVPR 2015.

Park, Keunhong et al. 2021.
Nerfies: Deformable neural radiance fields. ICCV 2021.

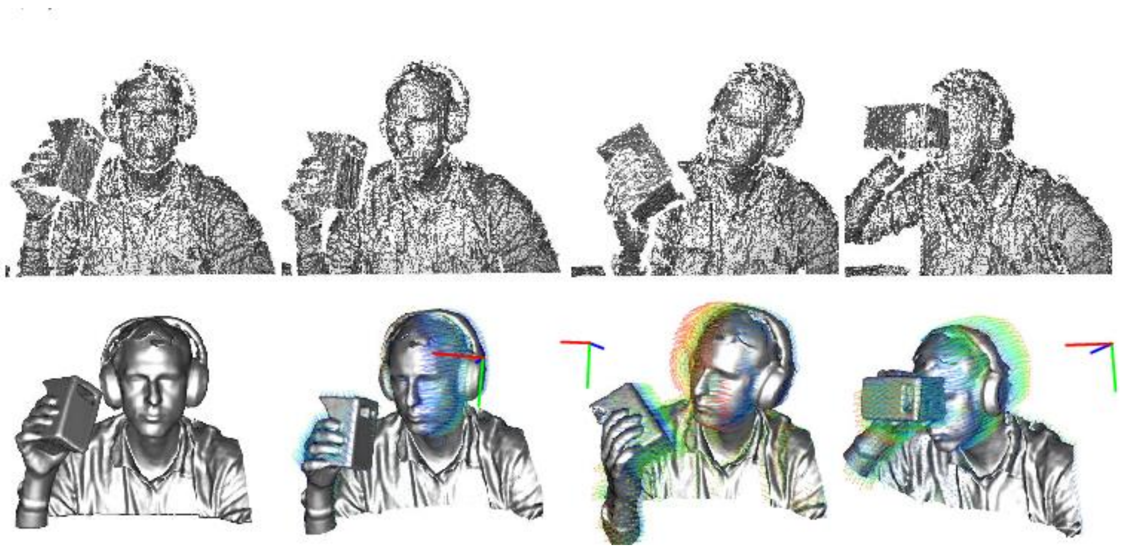

(a) Capture Process (b) Input (c) Nerfie (d) Nerfie Depth

Supervisor: Jian Wang

Volume Rendering

Kajiya et al.,
Ray tracing volume densities,
Computer Graphics 1984.

Mildenhall et al.,
Nerf: Representing scenes as neural radiance fields for view synthesis, ECCV 2020.

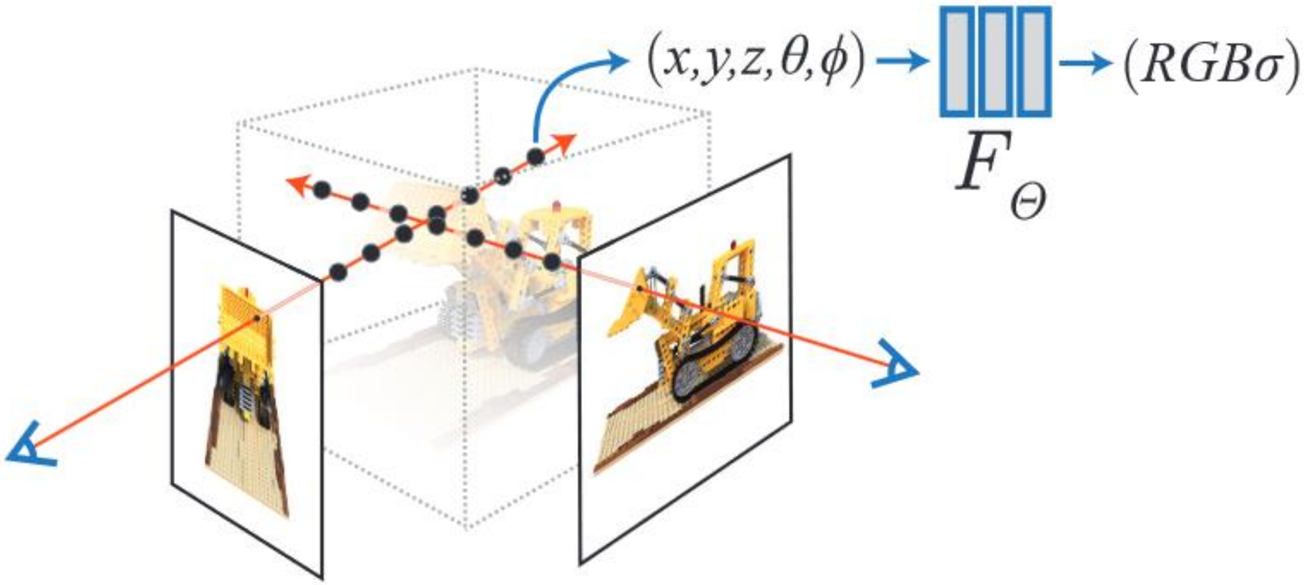

Supervisor: Hamza Pehlivan

Generative Adversarial Networks

Goodfellow et al. 2014.
Generative Adversarial Nets. NeurIPS 2014.

$$\min_G \max_D V(D, G) = \mathbb{E}_{\boldsymbol{x} \sim p_{data}(\boldsymbol{x})}[\log D(\boldsymbol{x})] + \mathbb{E}_{\boldsymbol{z} \sim p_{\boldsymbol{z}}(\boldsymbol{z})}[\log(1 - D(G(\boldsymbol{z})))]$$

Karras et al. 2019.
A Style-Based Generator Architecture
for Generative Adversarial Networks. CVPR 2019.

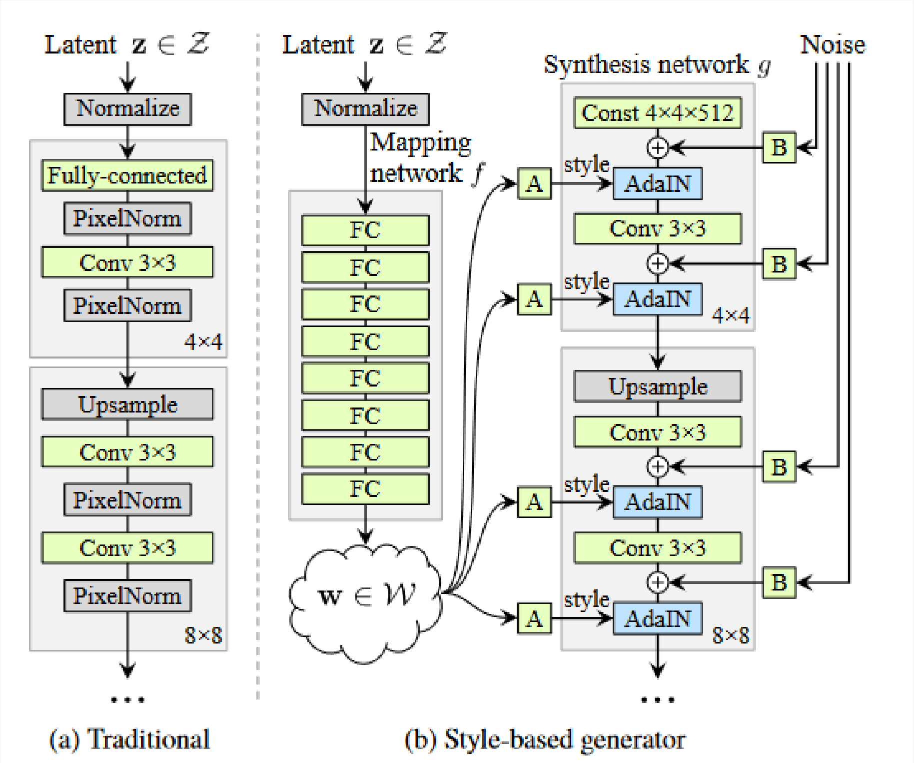

(a) Traditional (b) Style-based generator

Supervisor: Anton Zubekhin

From Convolution Pyramids to UNet

Farbman et al. 2011.
Convolution Pyramids. SIGGRAPH 2011.

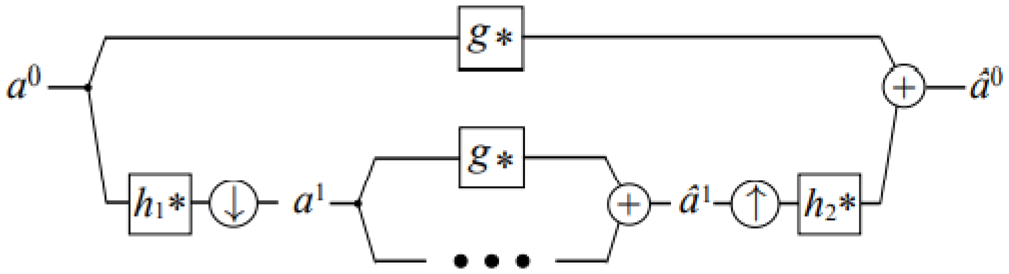

Ronnerberger et al. 2015.
UNet. MICCAI 2015.

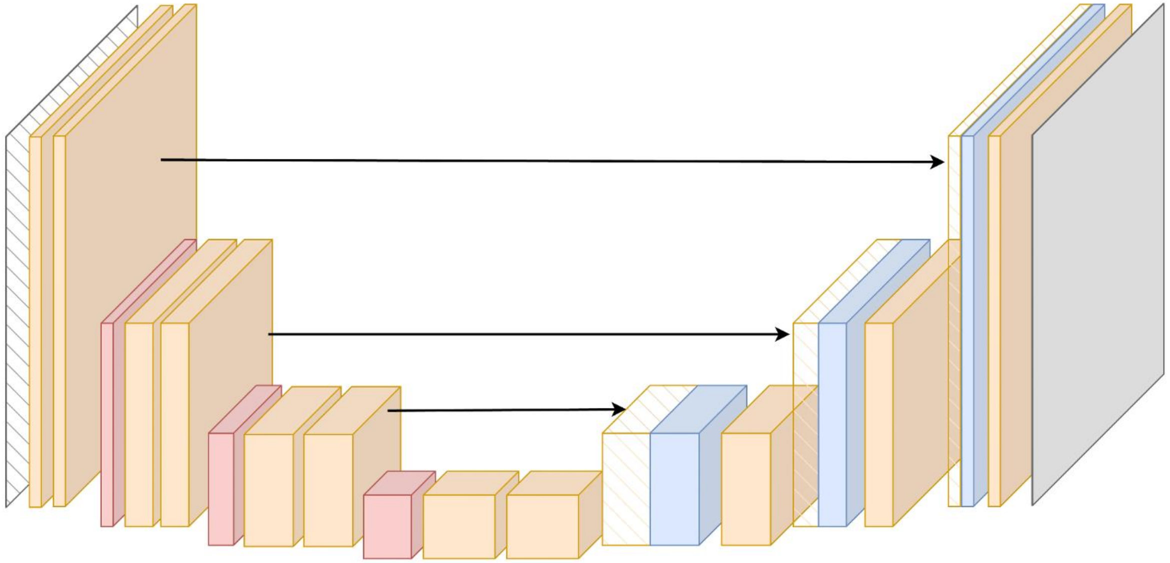

source: Conor O’Sullivan

Supervisor: Andrea Boscolo Camiletto

Residuals and ODEs

He et al. 2016.
Deep Residual Learning for Image Recognition. CVPR 2016.

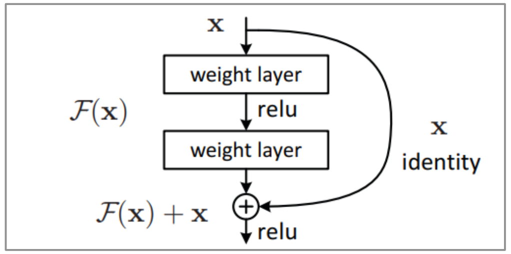

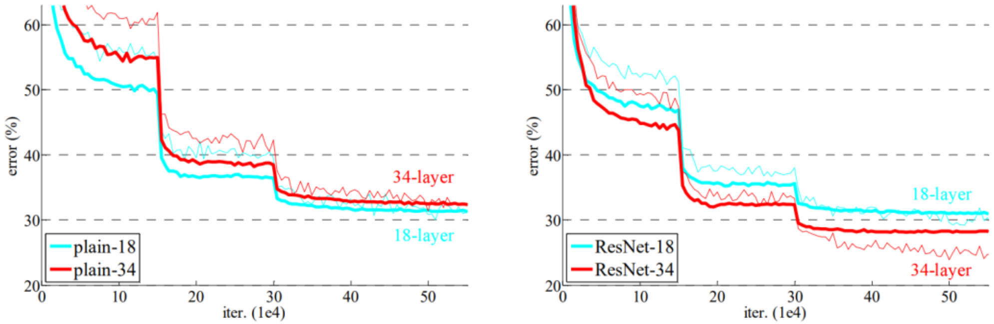

Chen et al. 2018.
Neural Ordinary Differential Equations. NeurIPS 2018.

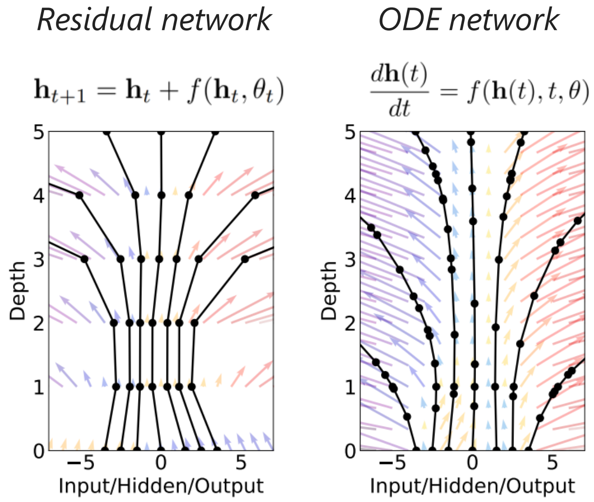

	Test Error	# Params	Memory	Time
1-Layer MLP†	1.60%	0.24 M	-	-
ResNet	0.41%	0.60 M	$\mathcal{O}(L)$	$\mathcal{O}(L)$
ODE-Net	0.42%	0.22 M	$\mathbf{\mathcal{O}(1)}$	$\mathcal{O}(\tilde{L})$

Supervisor: Olaf Dünkel

Filtering By repeated Integration

Paul Heckbert. 1986.
Filtering by Repeated Integration. Siggraph 1986.

Nsampi et al. 2023.
Neural Field Convolutions by Repeated Differentiation
Siggraph Asia 2023.

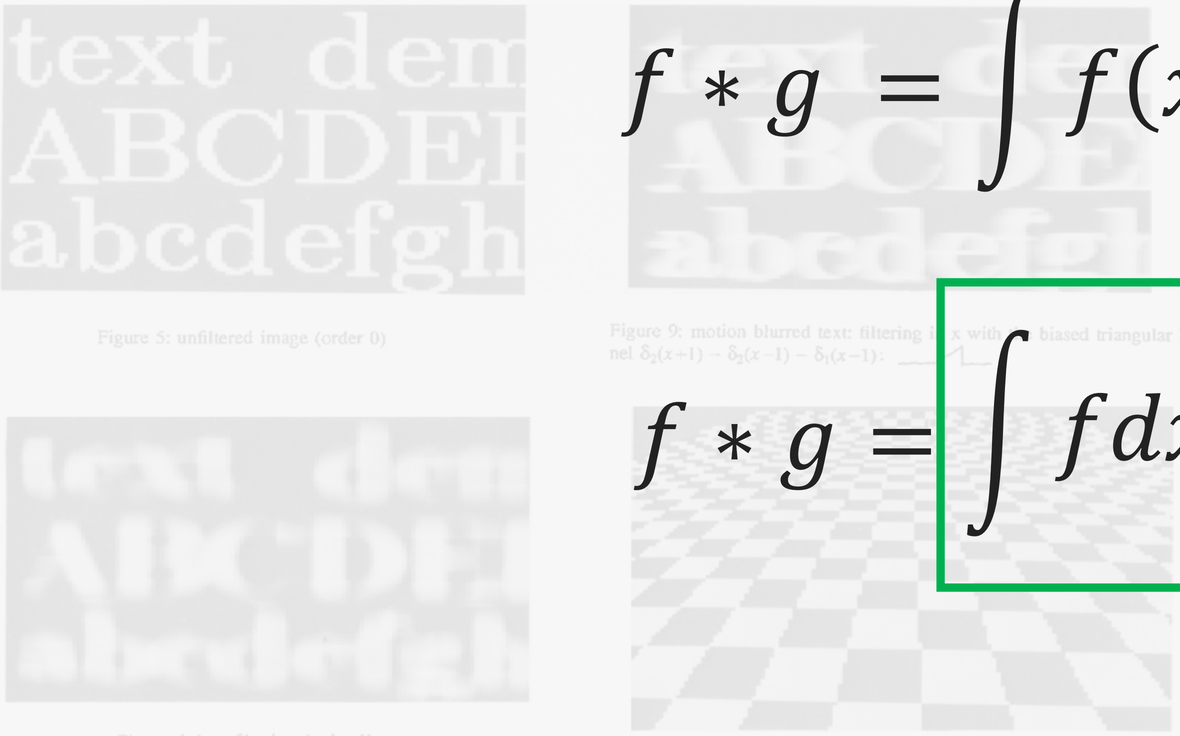

Figure 5: unfiltered image (order 0)

Figure 9: motion blurred text: filtering in x with biased triangular kernel $\delta_2(x+1) - \delta_2(x-1) - \delta_1(x-1)$

Figure 6: box filtering (order 1)

Figure 10: point sampled (order 0)

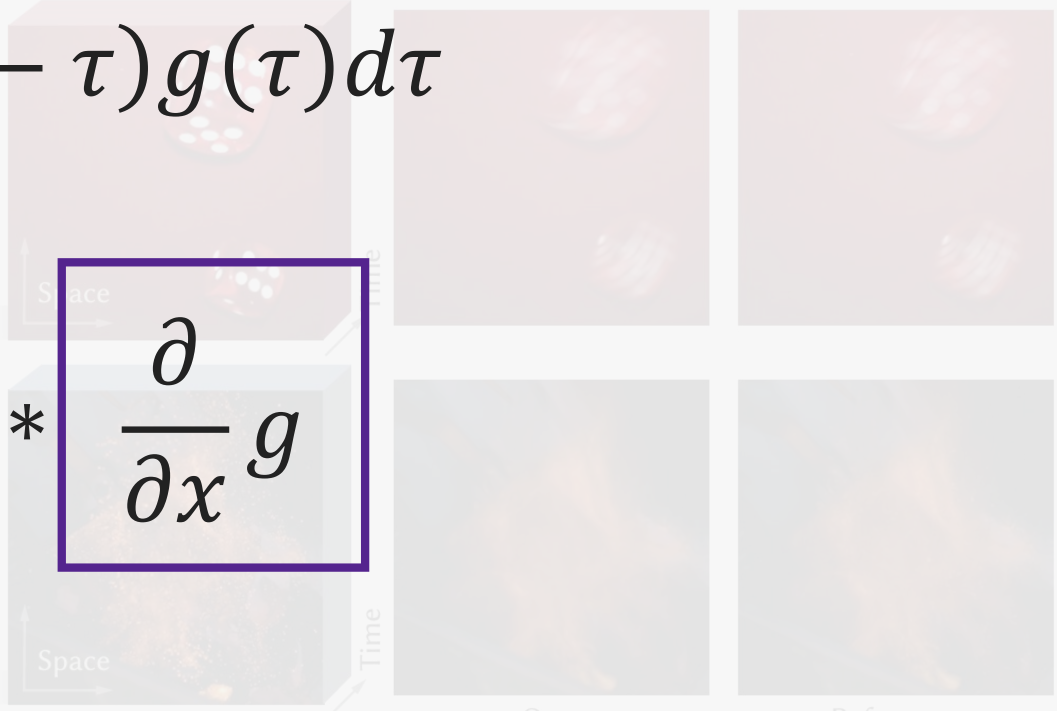

Input signal

Motion blur

$$f * g = \int f(x-\tau)g(\tau)d\tau$$

$$f * g = \int f dx * \frac{\partial}{\partial x} g$$

Supervisor: Ntumba Elie Nsampi

Animation

Kovar et al. 2002.
Motion Graphs. SIGGRAPH.

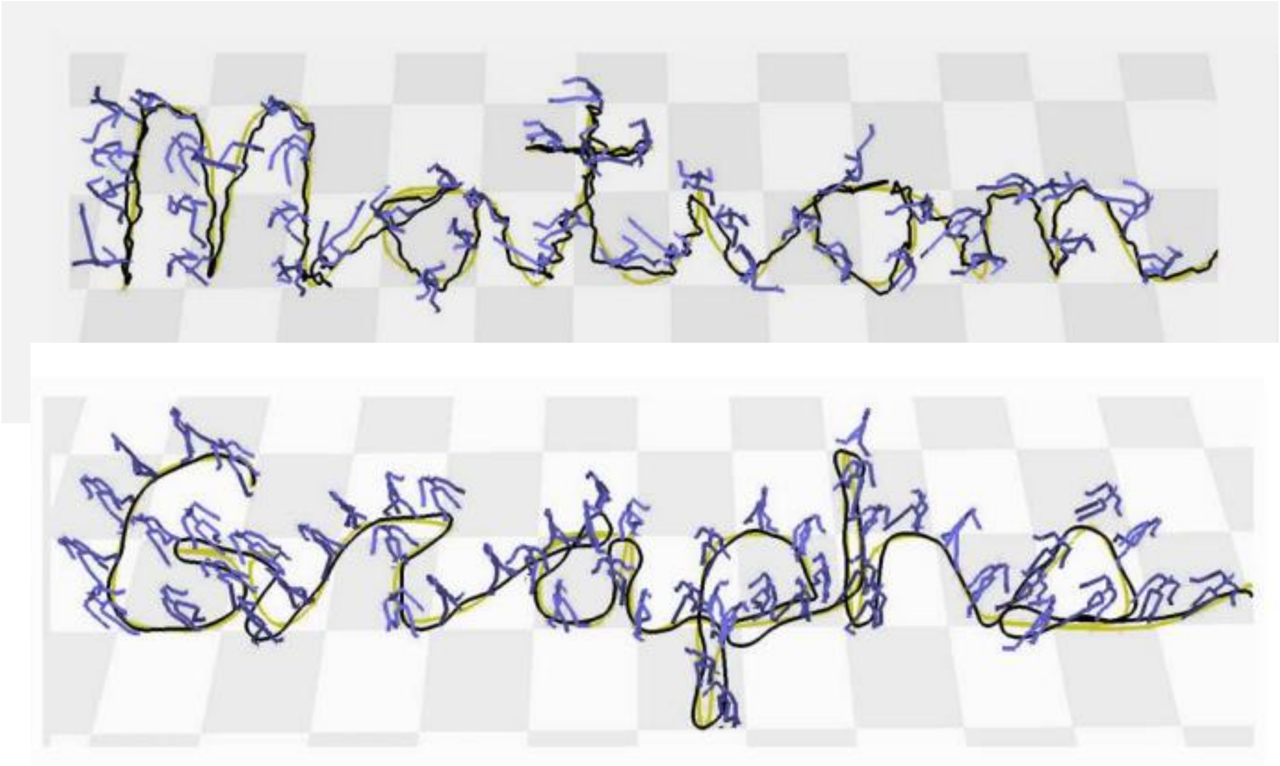

Holden et al. 2020.
Learned Motion Matching. TOG 2020

Supervisor: Wanyue Zhang

Uncertainty Modeling

Kendall and Gal.

What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision?

NeurIPS 2017

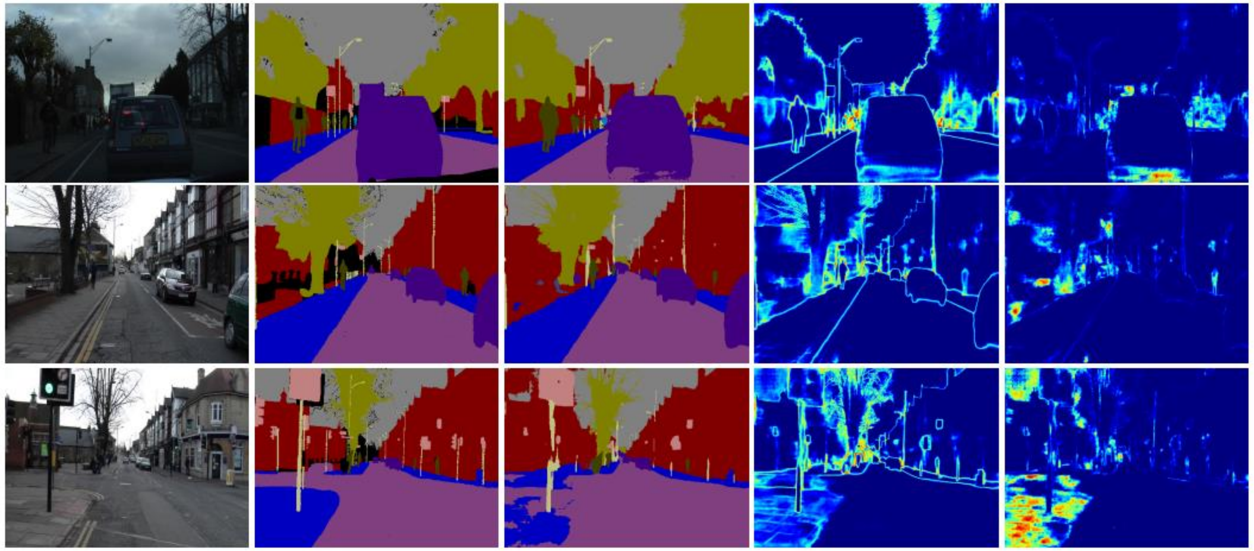

(a) Input Image (b) Ground Truth (c) Semantic Segmentation (d) Aleatoric Uncertainty (e) Epistemic Uncertainty

Goli et al.

Bayes' Rays: Uncertainty Quantification for Neural Radiance Fields

CVPR 2024

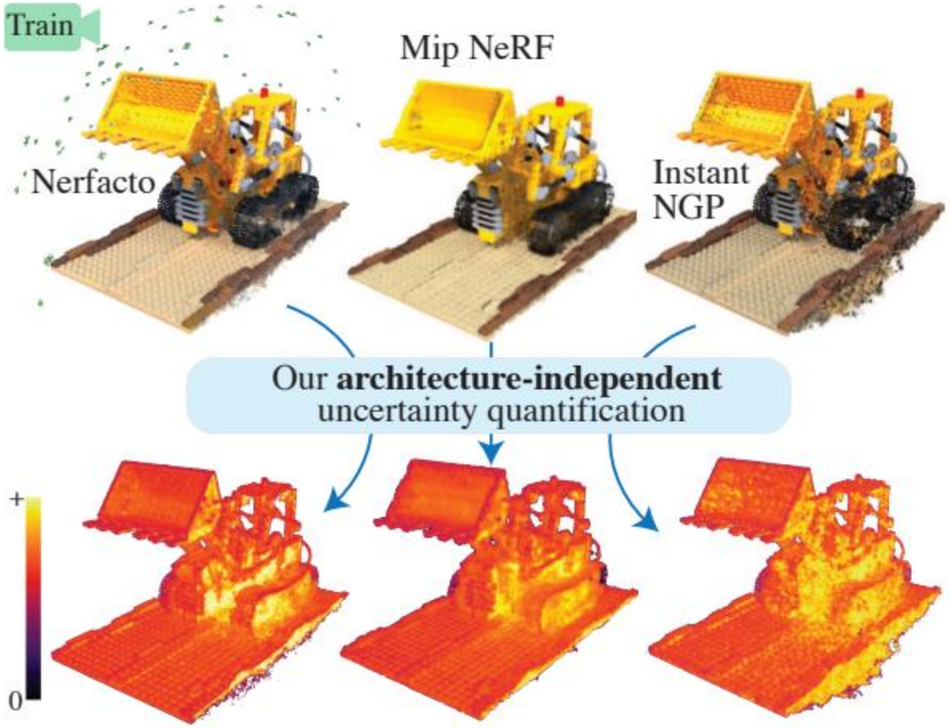

Supervisor: Linjie Lyu

Interval Arithmetic Meets Implicit Surfaces

Interval Arithmetic and Recursive Subdivision
for Implicit Functions and Constructive Solid Geometry
Duff, SIGGRAPH 1992

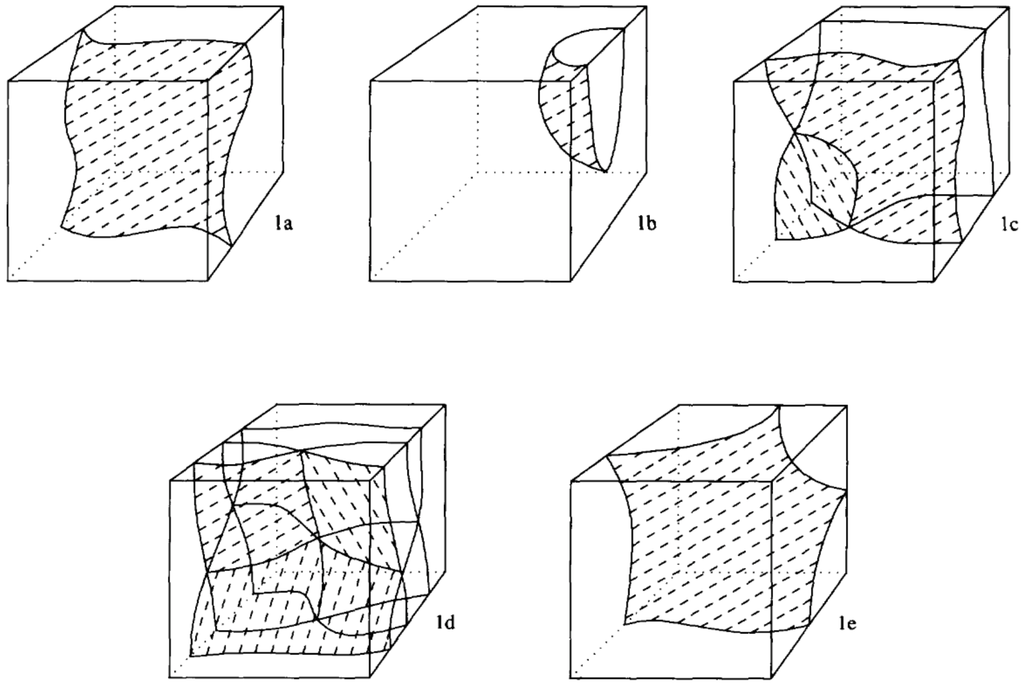

Spelunking the Deep: Guaranteed Queries
on General Neural Implicit Surfaces via Range Analysis
Sharp & Jacobson, SIGGRAPH 2022

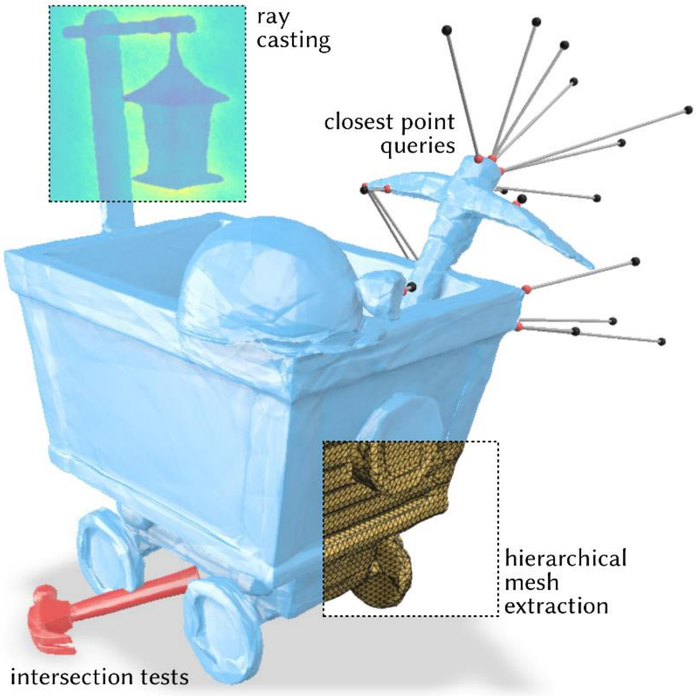

Supervisor: Gereon Fox

Clothing Simulation

Baraff et al.
Large Steps in Cloth Simulation,
Siggraph1998

Artur Grigorev et. al.,
HOOD: Hierarchical Graphs for Generalized
Modelling of Clothing Dynamics, CVPR2023

Supervisor: Heming Zhu

Face Modelling

A Morphable model for
the synthesis of 3D Faces.
Blanz and Vetter, SIGGRAPH 1999

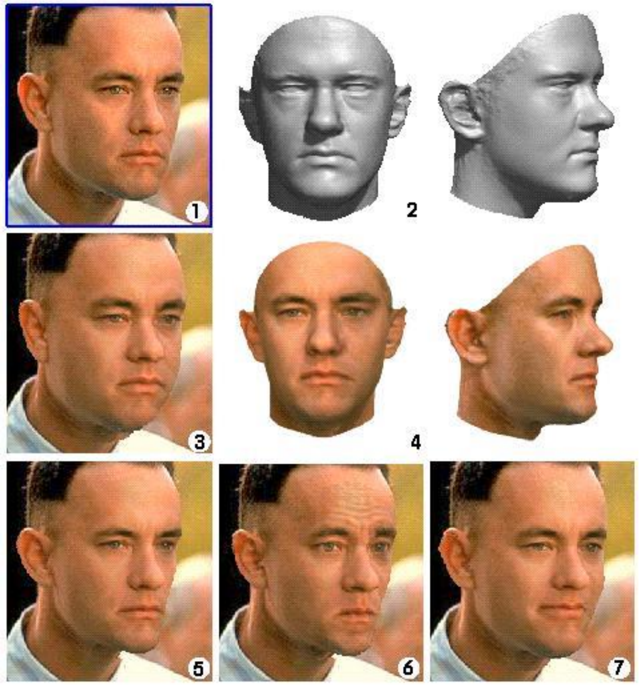

Learning Neural Parametric Head Models,
Giebenhain et. al. , CVPR 2023

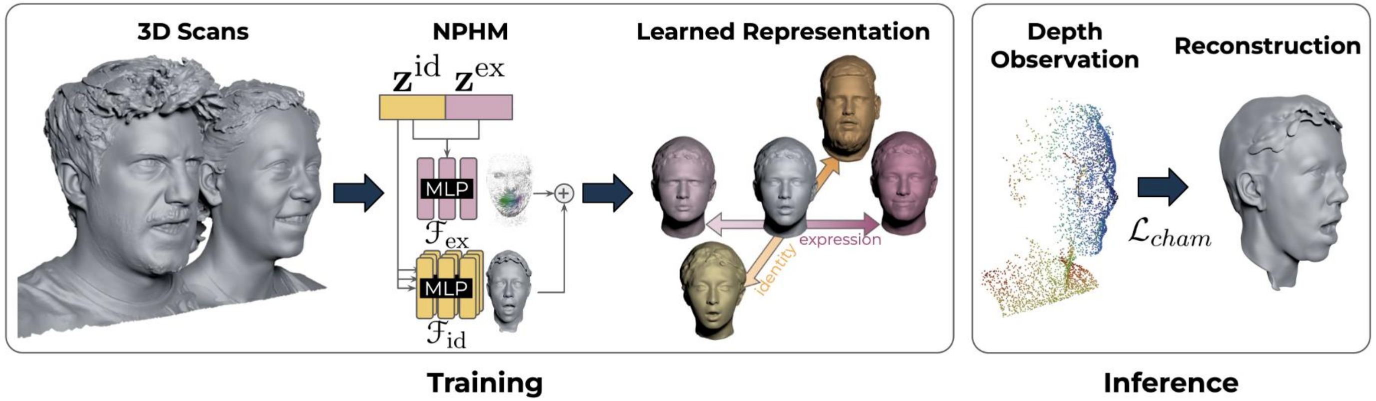

Supervisor: Kartik Teotia

Surface Reconstruction

Kazhdan et al.
Poisson Surface Reconstruction.
SGP 2006.

Sellán et al.
Neural Stochastic Screened Poisson Reconstruction. SIGGRAPH Asia 2023.

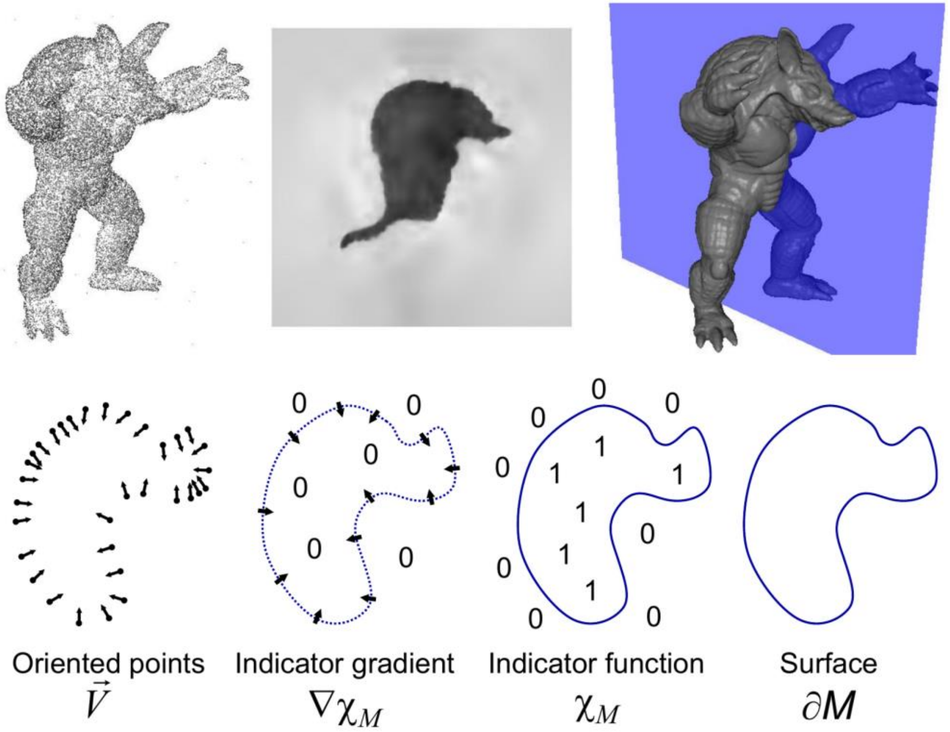

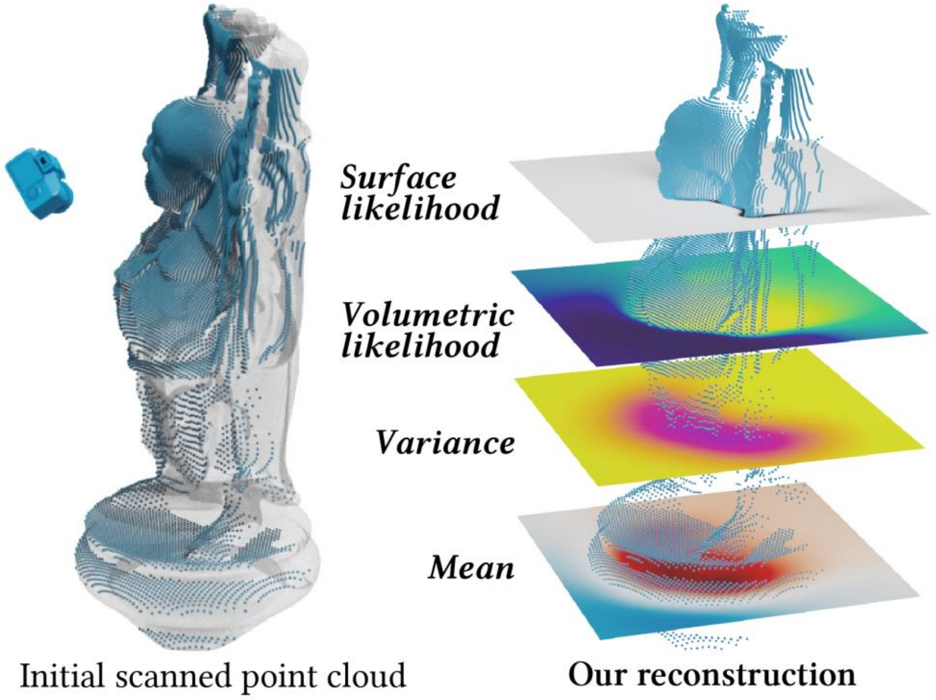

Supervisor: Viktor Rudnev

Physics-based Motion Modeling

Witkin and Kass, 1988
Spacetime Constraints. SIGGRAPH 1988.

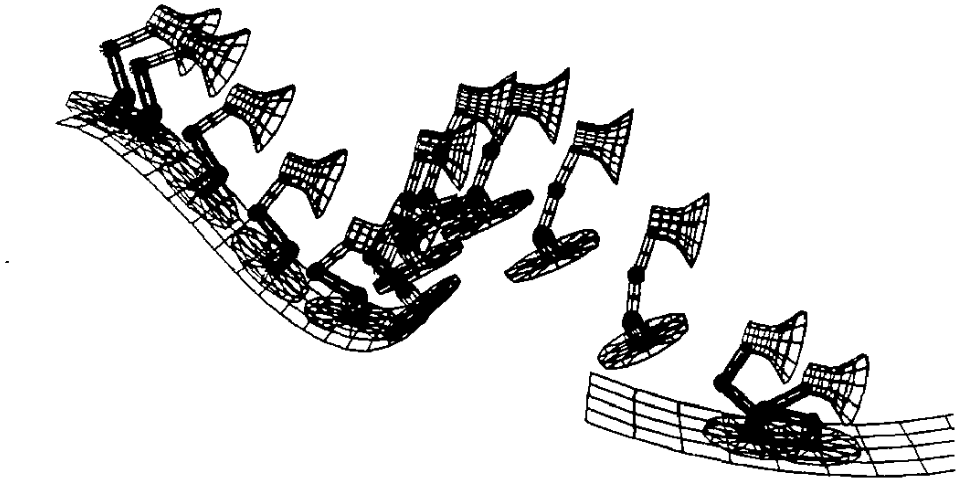

Peng et al. 2021.
AMP: Adversarial Motion Priors. SIGGRAPH 2021.

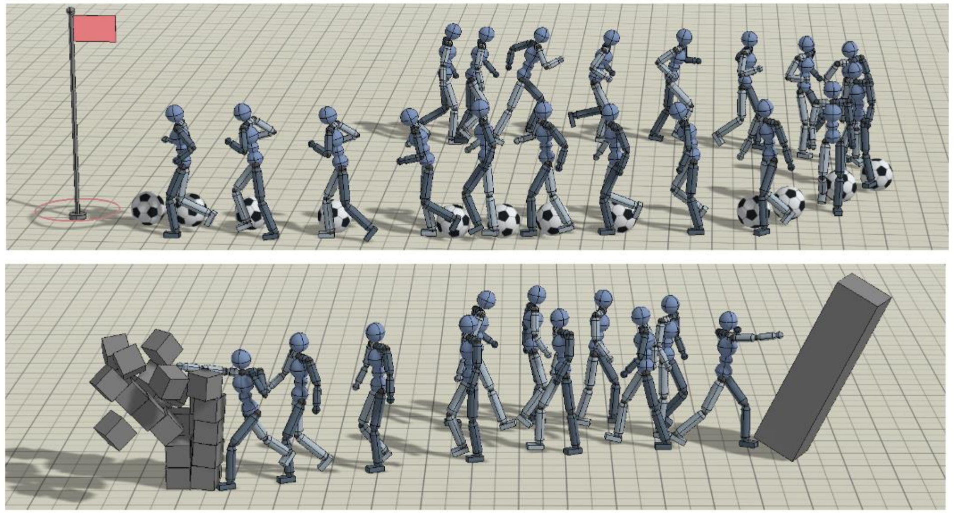

Supervisor: Noshaba Cheema

Reflectance Field Recovery

Devebec et al. 2000
Acquiring the Reflectance Field of a Human Face 2000.

Gotardo et al. 2018
Practical Dynamic Facial Appearance Modeling and Acquisition. SIGGRAPH 2018.

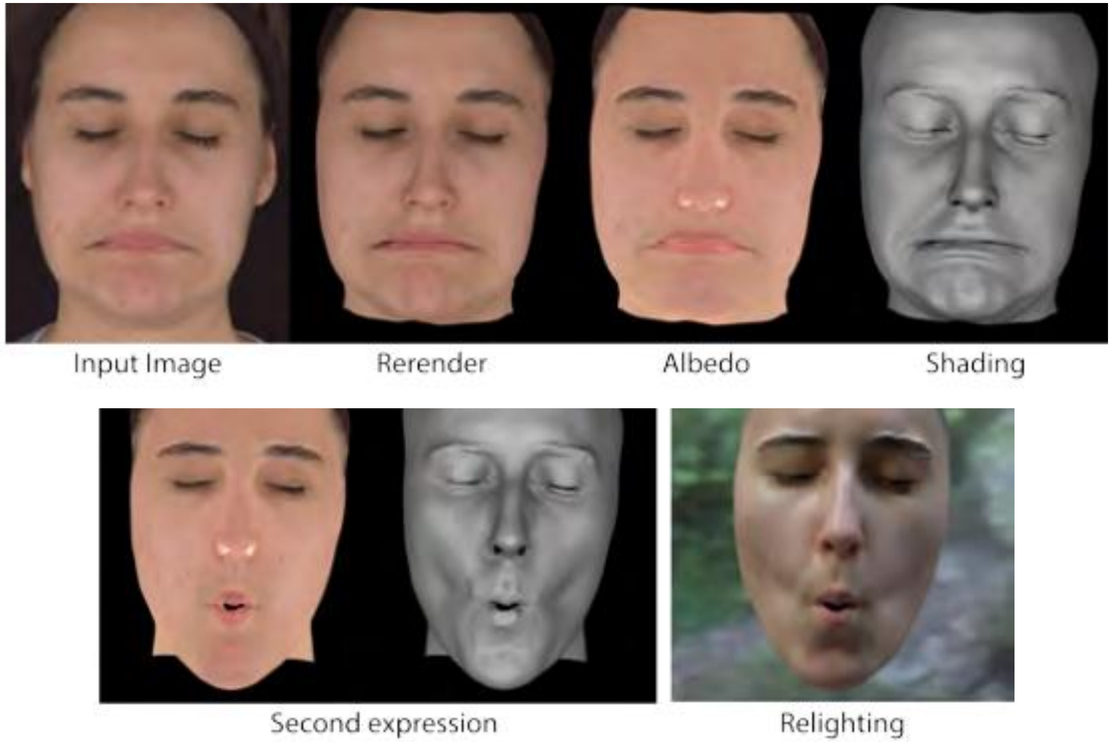

Supervisor: Pramod Rao

Next steps

- **If you have registered -> Wait!**
- **Matchmaking will be completed and you will be notified whether you are assigned to this seminar!**
- **If you are assigned to this seminar:**

Send your 3 preferred topics to mhaberma@mpi-inf.mpg.de with name and immatriculation ID

Example:

Marc Habermann, 2385768

1) Physics-based Motion Modeling — 1st choice

2) Interval Arithmetic Meets Implicit Surfaces — 2nd choice

3) Light Fields — 3rd choice

Thank you!

Questions?