

## Employment

---



**Senior Research Scientist** **Tavus** **Jan 2025 – Present**

- Working on digital twins.

**Founding Research Scientist** **Hedra** **March 2024 – Jan 2025**

- Optimized the inference code to achieve top performance using Triton, CUDA, torch.compile, and other tools.
- Trained/fine-tuned large-scale video diffusion models for talking digital avatars.
- Built a video scraper, an evaluation framework and parts of the data processing pipeline for video models.
- Worked on 3D Gaussian avatars from monocular videos.

**Senior Research Engineer** **Synthesia** **March 2022 – March 2024**

- Worked on diffusion models for avatars and environments.
- Wrote highly optimized CUDA kernels and used Nvidia OptiX to achieve real-time rates for 3D avatars.
- Researched and implemented a volumetric neural rendering framework for 3D digital humans.
- Devised a compact and efficient method, HumanRF (published at SIGGRAPH 2023), which enables novel-view synthesis for arbitrarily long videos of dynamic 3D humans.
- Worked on various tasks to prepare a processing pipeline for multi-view human avatar datasets including a memory-efficient, on-the-fly dataloading mechanism for terabytes of data.
- Worked on incorporating hand gestures into digital avatars using machine learning techniques.

**Ph.D. Student (incomplete)** **ETH Zurich** **June 2021 – March 2022**

- Devised a method for novel view synthesis of human heads from sparse set of input views.

**Graphics Research Intern** **Adobe Research** **June 2020 – November 2020**

- Invented a deep learning-based image-space denoising algorithm for interactive Monte Carlo rendering.
- The work is patented, published at SIGGRAPH 2021 and shipped in Adobe products.

**Graduate Student Researcher** **Technical University of Munich** **February 2019 – May 2020**

- Researched fast volume rendering and visualization techniques.

## Languages and Technologies

---

- Python, PyTorch, C++, CUDA, Triton, Docker, OptiX, cuDNN

## Publications

---

- **HumanRF: High-Fidelity Neural Radiance Fields for Humans in Motion**  
Mustafa Işık, Martin Rünz, Markos Georgopoulos, Taras Khakhulin, Jonathan Starck, Lourdes Agapito, Matthias Nießner  
*SIGGRAPH 2023 (ACM Transactions on Graphics)*
- **Interactive Monte Carlo Denoising using Affinity of Neural Features**  
Mustafa Işık, Krishna Mullia, Matthew Fisher, Jonathan Eisenmann, Michaël Gharbi  
*SIGGRAPH 2021 (ACM Transactions on Graphics)*
- **Learning Adaptive Sampling and Reconstruction for Volume Visualization**  
Sebastian Weiß, Mustafa Işık, Justus Thies, Rüdiger Westermann  
*IEEE Transactions on Visualization and Computer Graphics 2020*

## Selected Projects

---

- 🗓️ **RGB Face Tracking&Reconstruction (2020)**. This project aims to reconstruct faces from monocular RGB videos (e.g., a video on YouTube) via differentiable rendering. Implemented in **C++**, **CUDA** and **OpenGL**.
- 🗓️ **Kinect Fusion (2019)**. This project aims to reconstruct 3D geometry in real-time using a Kinect camera. Implemented in **C++** and **CUDA**.
- 🗓️ **Glue (2018)**. Glue is a hobby renderer that taught a great deal of theory and practice behind 3D graphics, rendering, and ray tracing. Implemented in **C++17** and **CMake**.

## Education

---

**Munich, Germany**                      **Technical University of Munich**                      **October 2018 – April 2021**

- M.Sc. in Computer Science. GPA: 1.1 (graduated with high distinction)

**Ankara, Turkey**                      **Middle East Technical University**                      **September 2013 – January 2018**

- B.Sc. in Computer Engineering. GPA: 3.75 (placed 2<sup>nd</sup> among graduates of 2017)

## Honors and Awards

---

- **Member of best.in.tum (2019 – 2021)**. best.in.tum comprises the top two percent of the students in the Department of Computer Science at Technical University of Munich.
- **High Honor Student (2013 – 2018)**.