



Master Seminar - Recent Advances in 4D Computer Vision

Cecilia Curreli, Mariia Gladkova



Master Seminar - Recent Advances in 4D Computer Vision

Cecilia Curreli, Mariia Gladkova



Mariia Gladkova



Cecilia Curreli

What do you do in a Master's Seminar?

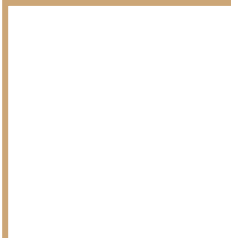
Read a very recent research paper and analyze its contributions within the academic landscape.

Deliverables (5 ECTS):

- Present the paper in front of an audience
- Write a report
- Attend and discuss other students' presentations (mandatory)

Not in the course:

- Practical / coding exercises
- Group project

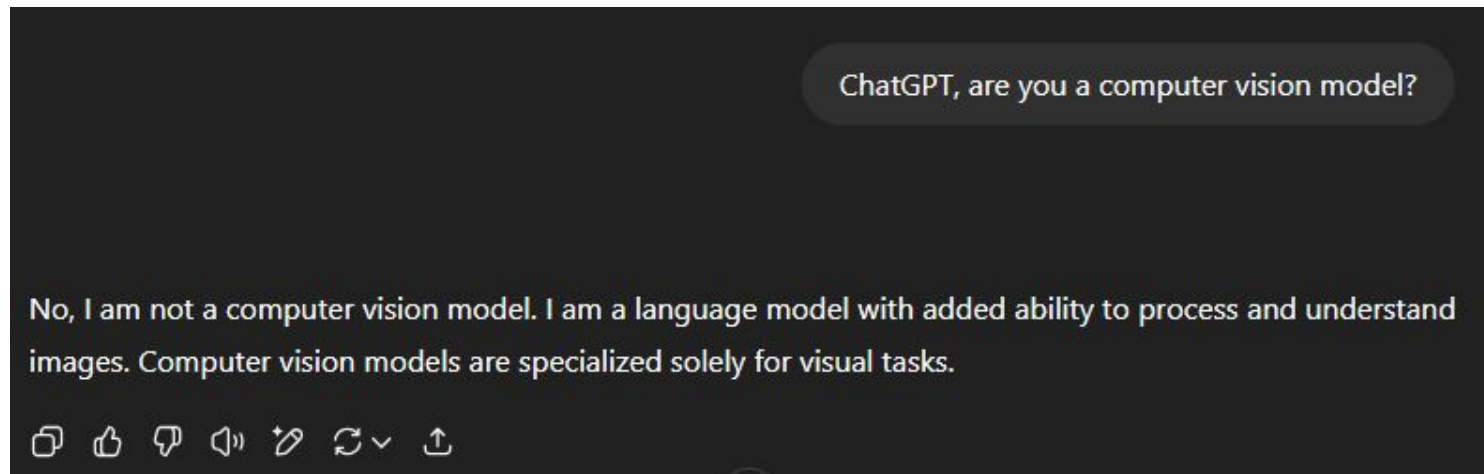


What is 4D Computer Vision?

Cecilia Curreli, Mariia Gladkova




What is 4D Computer Vision?



Vision is not language - but it can leverage language models and information


What does 4D mean?

2D



Image

3D



3D Parametrization

What does 4D mean?



Take a matrix



add the time dimension

What does 4D mean?

2D + Time



Video

Video Generation
From Text or Image

Sora by OpenAI



Video Generation is not 4D Computer Vision ;)

What does 4D mean?

3D + Time




Example of dynamic scene reconstruction from monocular video

4D Computer Vision has the same dimension as our world - 4D!

What does 4D mean?

2D




Image



2D pixel array

3D



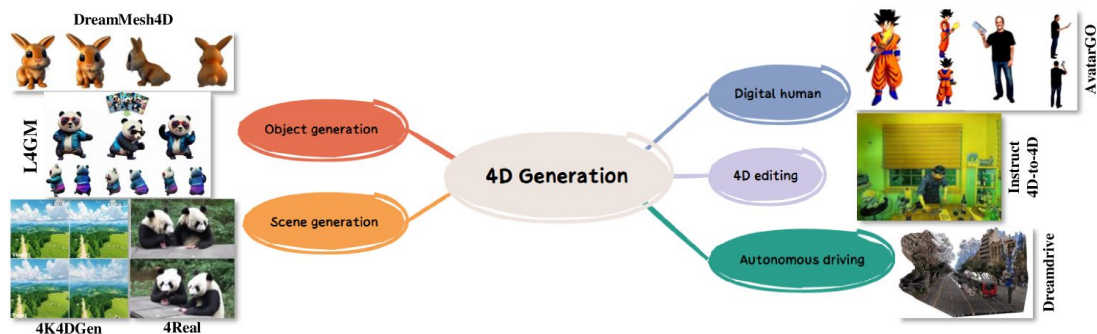
Understanding 4D reality is an exciting direction of great research interest!

[1] Ramesh, et al. DALL-E 2 "Hierarchical Text-Conditional Image Generation with CLIP Latents." (2022).

[2] Poole, Ben, et al. "Dreamfusion: Text-to-3d using 2d diffusion." (2022).

What are the topics discussed?

- Generation of dynamic scenes
 - Text-to-4D
 - Image-to-4D
- Reconstruction
 - 4D reconstruction
 - Novel view synthesis
 - Motion estimation
- Various application domains
 - Autonomous driving
 - Robotics & XR
 - Humans
- ... and many more!



Text-to-4D Generation



A turtle swimming.



Waves crashing against a lighthouse.



A bee fluttering its wings fast.



A cat singing, best quality, 4K, HD



From text snippet to a video or a 4D model

4D Novel View Synthesis




From posed images to a 4D model with high fidelity and photorealism

[1] Wu et al. "4D Gaussian Splatting for Real-Time Dynamic Scene Rendering" CVPR (2024).

[2] Li et al. "Spacetime Gaussian Feature Splatting for Real-Time Dynamic View Synthesis" CVPR (2024).

4D Autonomous Driving



Reconstruction

From images to scene analysis

4D Autonomous Driving



left



forward



right

Generation

From image + controls to 4D prediction

4D Humans



Source Actor



Animated Avatar



Reconstruction

Multi-billion net-worth companies invest heavily in human avatars ...


[1] Quian et al. "GaussianAvatars: Photorealistic Head Avatars with Rigged 3D Gaussians." CVPR (2024).

[2] Shin, Kim et al. "WHAM: Reconstructing World-grounded Humans with Accurate 3D Motion." CVPR (2024).

4D Humans

Human Motion **Composition**
with FlowMDM

from the **past** ... predict the **future**




Generation

..and look closer into their realistic motion and interactions!

[1] Barquero et al. "FlowMDM: Seamless Human Motion Composition with Blended Positional Encodings." CVPR (2024).

[2] Curreli et al. "Nonisotropic Gaussian Diffusion for Realistic 3D human Motion Prediction." CVPR (2025).



Seminar Layout & Organization

Cecilia Curreli, Mariia Gladkova



What do you do in a Master's Seminar?

Read a very recent research paper and analyze its contributions within the academic landscape.

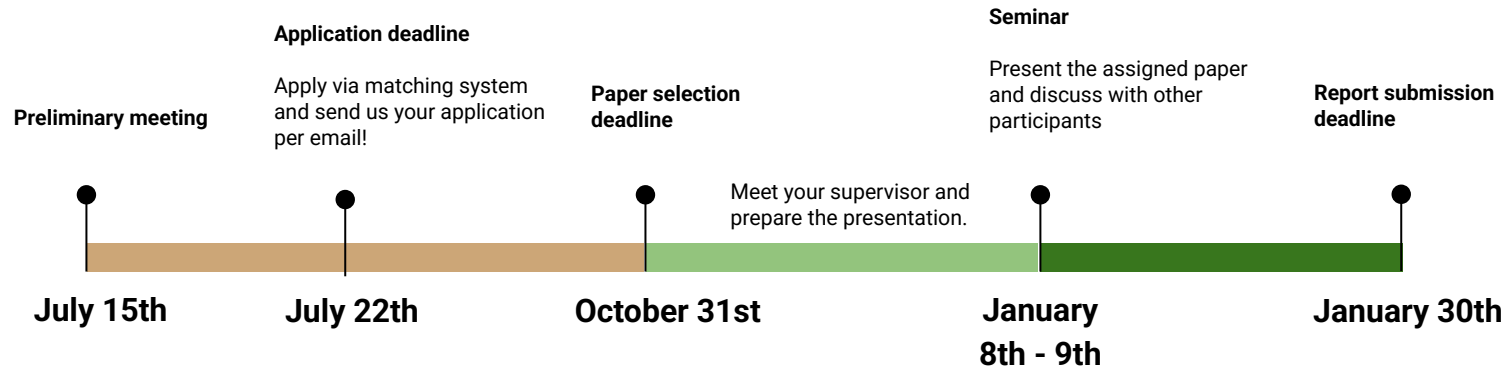
Deliverables (5 ECTS):

- Present the paper in front of an audience (ca. 20-25 mins + 5-10 mins questions)
- Write a report (max 5 pages)
- Attend and discuss other students' presentations (mandatory)

Not in the course:

- Practical / coding exercises
- Group project

How does the schedule look like?



January 2026, mark your calendars ;)

Further orga stuff

- Meet your supervisor
 - 1 week before presentation: paper and slides submission
- Write a report
 - LaTeX template would be provided
 - 5 pages summarizing the paper and providing your high level insights
 - Report deadline is due in 3 weeks after your presentation
- All meetings and seminar sessions are mandatory to attend

How does the grading look like?

- Presentation: 50%
- Report: 40%
- Active participation and questions during seminar: 10%

What do I need to know beforehand?

- Previously attended DL lectures, such as I2DL, CV III, 3D Scanning & Motion Capture, and other
- Knowledge of linear algebra, probabilities, non-linear optimization is highly beneficial

We focus on learning methods, so brushing off the dust from DL notes will help!

I'm interested, how do I apply?

- Apply via [matching system](#) and assign our seminar high priority
- Send an email to 4dvision-ws25@vision.in.tum.de (example on the next slide) with
 - Short info about your background (see template in the next slide)
 - Attachment 1: Transcript of records (Bachelor's + Master's)
 - Attachment 2: CV/Resume

**Deadline for both
matching system and
email: 22th July!**

Two-step verification: matching system + email

Any tips for the email?

Do not write to our
private emails!
Your email may get lost

To  4dvision-ws25@vision.in.tum.de ✕

Cc

Application 01234567

Hi Cecilia and Mariia,

My name is John Doe, and here are my application details:

- Matriculation number: 01234567
- Bachelor's grade: 1.7
- Master's grade: N/A (it is my first semester)
- Relevant courses: I2DL (grade pending), Deep Learning (Bachelor course, 1.0)
- Relevant experience: In my bachelor thesis, I collected image data and implemented a network to count the number of bees on the farm.

I have applied for the seminar via a matching system with a high-priority ranking. Looking forward to participation!

Best regards,
John

Make it concise and relevant to the seminar scope.

Send

Discard



Any remaining questions?

Master Seminar -
Recent Advances in
4D Computer
Vision

Cecilia Curreli, Mariia Gladkova



Mariia Gladkova



Cecilia Curreli

4dvision-ws25@vision.in.tum.de

Reach out via email for lost+found questions!