

Introduction to Neural Scene Representation and Neural Rendering

Lingjie Liu




We Live in a World that is 3D and Contains Dynamics





We Digitize Our World in 3D




Future AI: Towards 3D Aware



3D Reconstruction of Real-world Scenes



Geometry
+ Appearance



Motion
+ Deformation

Photo-realistic Rendering

- **Image Synthesis** of Real-world Scenes with 3D Control.



Applications



AR / VR



Gaming / Movie



Healthcare



Autonomous Driving



Robot Grasping



Human-robot Interaction

Why are they challenging?

Problem formulation






Captured images



Processing





Rendering of real-world place




[Mildenhall et al., Neural Radiance Fields (NeRF), ECCV 2020]
[Wu et al., Scalable Neural Indoor Scene Rendering, SIGGRAPH 2022]

Classical Computer Graphics Pipeline




Computer Graphics Rendering

Image-based 3D Reconstruction



**COLMAP [Johannes et al. 2016, Schoenberger et al. 2016]
(Input: 100 images)**

Challenges in Image-based Reconstruction




Hard to extract reliable correspondences!

Computer Graphics Rendering

Rendering requires very high-quality 3D models







VS




Neural Scene Representation and Neural Rendering

To the rescue


Neural Scene Representation and Neural Rendering



Neural Scene Representation and Neural Rendering



[Tewari et al. 2020]



[Tewari et al. 2021]

Neural Rendering - Definition

- Definition:

*"Deep neural networks for **image or video generation** that enable **explicit or implicit control** of **scene properties**"*

1)

Generative networks that
synthesize raw pixel output


2)

Controllable by
interpretable parameters
or by video/audio input.

3)

Illumination, camera, pose,
geometry, appearance, or
semantic structure controllable


Generative Query Network (GQN)




Neural scene representation and rendering, Eslami et al. 2018

Generative Query Network (GQN)

observation




neural rendering




Neural scene representation and rendering, Eslami et al. 2018

Generative Query Network (GQN)

observation




neural rendering



Neural scene representation and rendering, Eslami et al. 2018

Generative Query Network (GQN)



Neural scene representation and rendering, Eslami et al. 2018


Neural Rendering Zoo

“Regress it”




e.g., GQN

“Make it more real”




e.g., DVP or DNR

Deep Video Portraits (DVP)



Deep Video Portraits, Kim et al. 2018

Deep Video Portraits (DVP)



Deep Video Portraits, Kim et al. 2018


Neural Rendering Zoo

“Regress it”




e.g., GQN

“Make it more real”




e.g., DVP or DNR

“Regress & render”



e.g., Neural Volumes

Neural Volumes



Neural Volumes: Learning Dynamic Renderable Volumes from Images, Lombardi et al. 2019

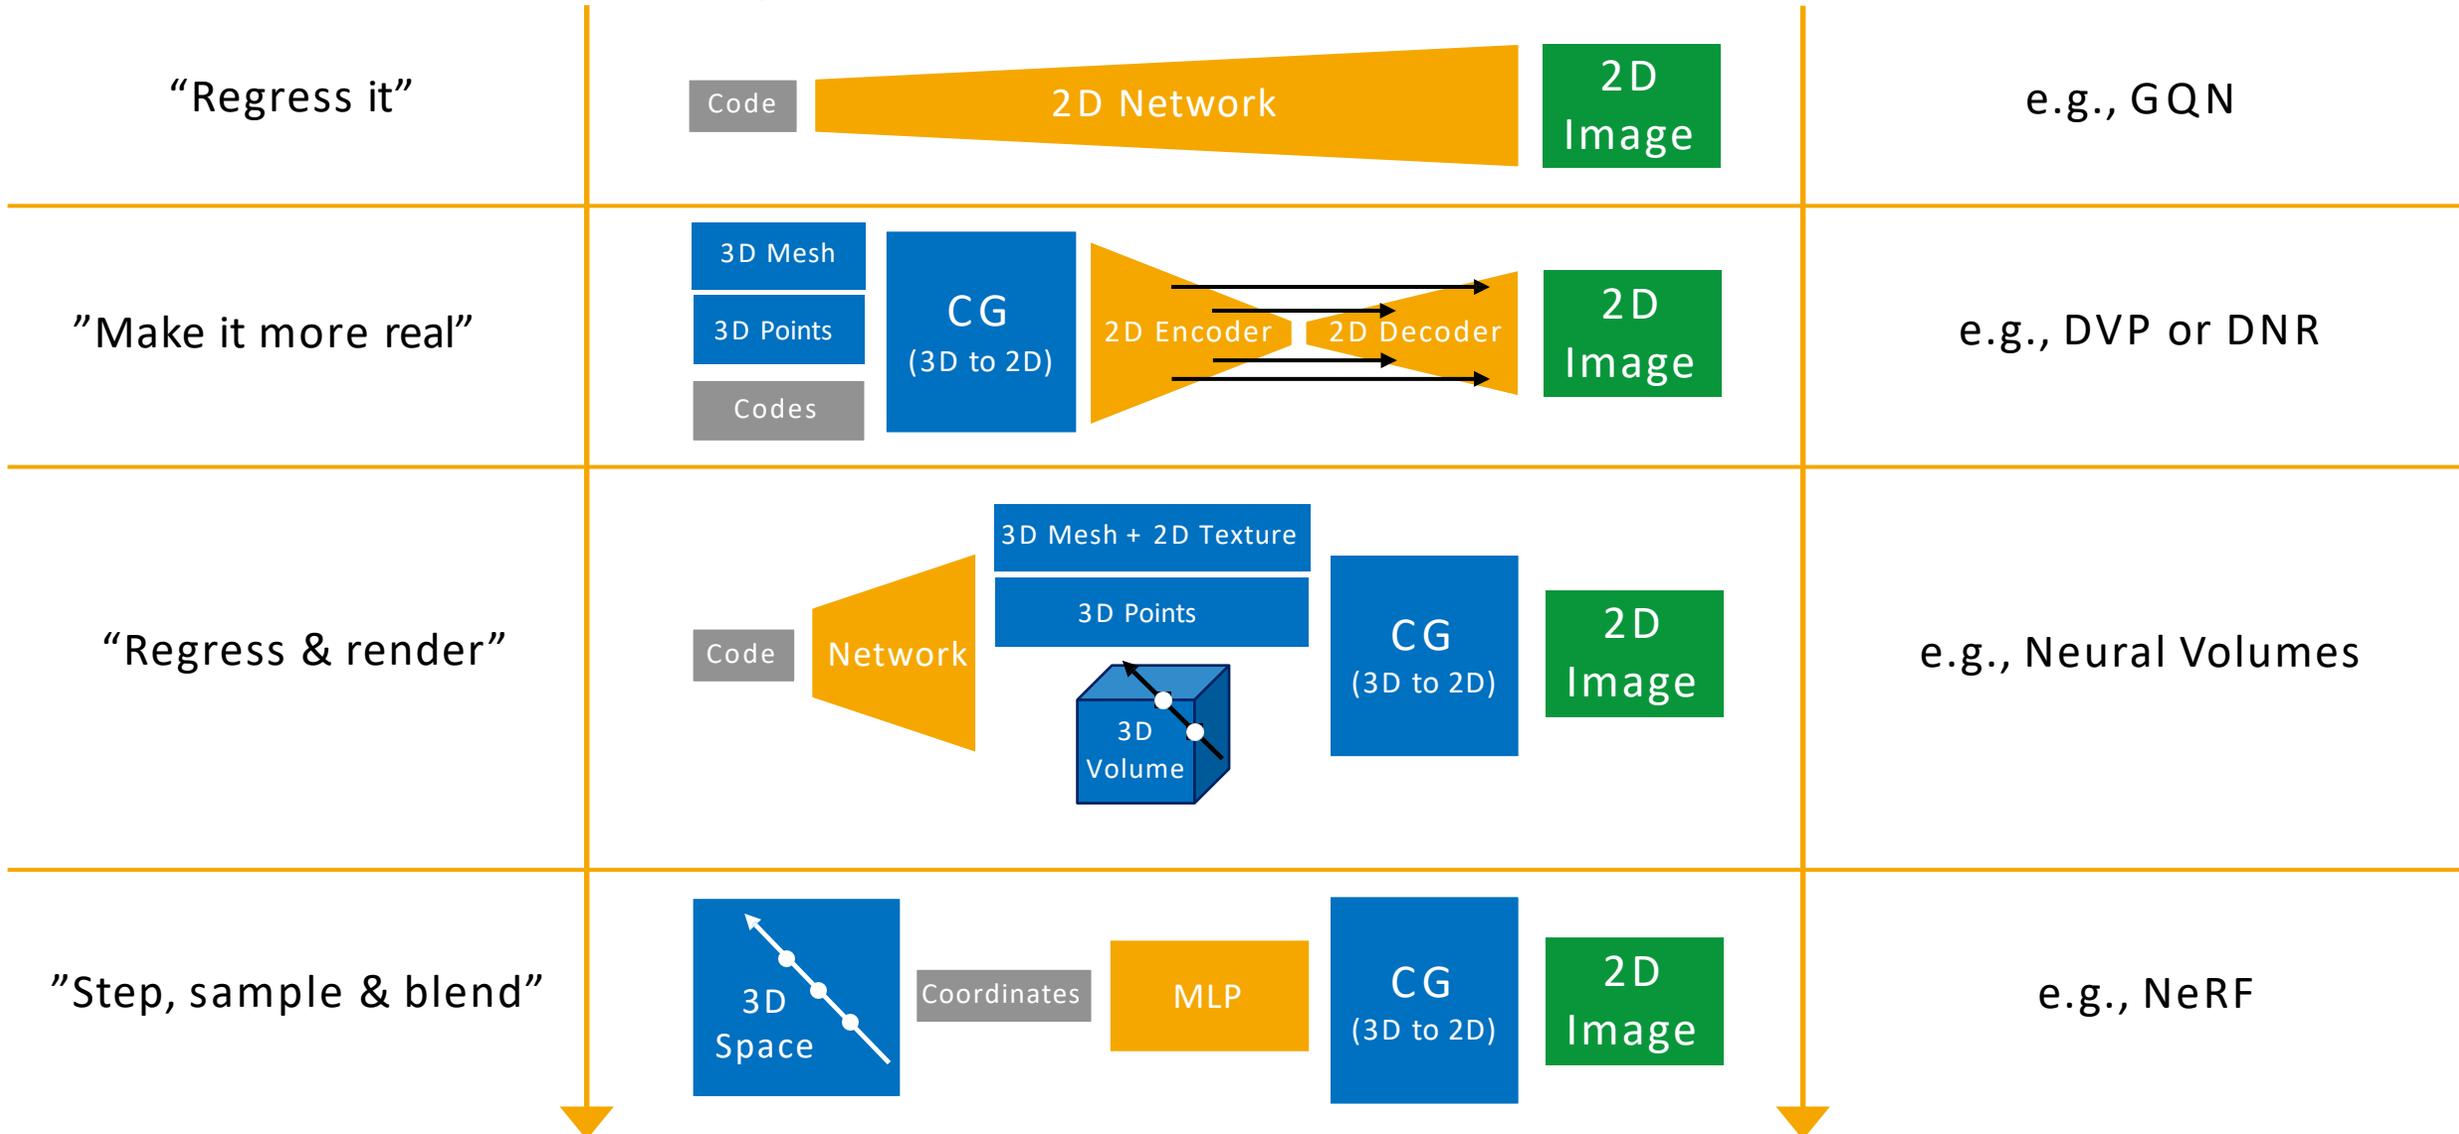
Neural Volumes




Neural Volumes



Neural Rendering Zoo



Neural Radiance Fields (NeRF)




[Mildenhall et al. 2020]

Neural Radiance Fields (NeRF)




[Mildenhall et al. 2020]

Neural Rendering Zoo




Overview


Scene Representation



Voxelgrids



Implicit Function



Hybrid
Implicit/Explicit



(We'll talk more about other representations in our next class)

Renderer

Volumetric


Sphere-Tracing
Volumetric

Volumetric


Both Scene Representation and Differentiable Renderer often adapted from traditional computer graphics.

Requirements


Scene Representation



Voxelgrids



Implicit Function



Hybrid
Implicit/Explicit



Renderer

Volumetric

Sphere-Tracing
Volumetric


Volumetric

Pros

Cons


Voxel-based methods

DeepVoxels




Sitzmann et al., CVPR 2018

Neural Volumes




Lombardi et al., SIGGRAPH 2019

HoloGAN



Phuoc et al., ICCV 2019


Voxel-based methods




Trilinear Interpolation

Requirements


Scene Representation



Voxelgrids



Implicit Function



Hybrid
Implicit/Explicit



Renderer

Volumetric

Sphere-Tracing
Volumetric

Volumetric

Pros


Fast rendering

Cons


Memory $O(n^3)$
Limited spatial
resolution

Requirements


Scene Representation



Voxelgrids



Implicit Function



Hybrid
Implicit/Explicit



Renderer

Volumetric

Sphere-Tracing
Volumetric

Volumetric


Pros

Fast rendering

Cons

Memory $O(n^3)$
Limited spatial
resolution

Neural Implicit Approaches



Scene Representation Networks
Generalizes across scenes
Sitzmann et al., NeurIPS 2019




Differentiable Volumetric Rendering
Generalizes across scenes
Niemeyer et al., CVPR 2020




NeRF
Single-scene
Mildenhall et al., ECCV 2020



Implicit Differentiable Renderer
Single-scene
Yariv et al., NeurIPS 2020




Sphere tracing



Volumetric

Sphere Tracing


$$f(x, y, z) = d$$



[Source: Takikawa et al]

Sphere Tracing


$$f(x, y, z) = d$$



[Source: Takikawa et al]

Sphere Tracing


$$f(x, y, z) = d$$



[Source: Takikawa et al]


Sphere Tracing

$$f(x, y, z) = d$$




[Source: Takikawa et al]

Sphere Tracing




[Source: Takikawa et al]

Sphere Tracing




[Source: Takikawa et al]

Sphere Tracing




[Source: Takikawa et al]

Sphere Tracing




[Source: Takikawa et al]

Sphere Tracing




[Source: Takikawa et al]

Sphere Tracing




[Source: Takikawa et al]

Sphere Tracing



[Source: Takikawa et al]


Sphere Tracing



[Source: Takikawa et al]


Sphere Tracing

$$f(x, y, z) = d$$



[Source: Takikawa et al]

Neural Implicit Approaches



Scene Representation Networks
Generalizes across scenes
Sitzmann et al., NeurIPS 2019




Differentiable Volumetric Rendering
Generalizes across scenes
Niemeyer et al., CVPR 2020



NeRF
Single-scene
Mildenhall et al., ECCV 2020




Implicit Differentiable Renderer
Single-scene
Yariv et al., NeurIPS 2020



Sphere tracing

- Faster
- Fewer network evaluations
- Convergence more difficult




Volumetric


- Higher Quality
- Easy convergence
- Very expensive

Requirements


Scene Representation



Voxelgrids



Implicit Function



Hybrid
Implicit/Explicit

...

Renderer

Volumetric

Sphere-Tracing
Volumetric

Volumetric

Pros

Fast rendering

High quality
Compact
Admits *global* priors


Cons

Memory $O(n^3)$
Limited spatial
resolution


Extremely expensive,
slow rendering

Requirements


Scene Representation



Voxelgrids



Implicit Function



Hybrid
Implicit/Explicit



Renderer

Volumetric

Sphere-Tracing
Volumetric

Volumetric

Pros

Fast rendering

High quality
Compact
Admits *global* priors

Cons

Memory $O(n^3)$
Limited spatial
resolution

Extremely expensive,
slow rendering

Hybrid Implicit / Explicit

Wine Holder



SRN (Sitzmann et al. 2019)
(Rendering speed: 1.10 s/frame)




NSVF
(Rendering speed: 1.68 s/frame)

Neural Sparse Voxel Fields,
Liu et. al., NeurIPS 2020



PiFU, Saito et al., ICCV 2019
GRF, Trevithick et al., arXiv 2020
pixelNeRF, Yu et. al., CVPR 2021
MVSNeRF, Chen et al., arXiv 2021
[Learn *local* \(image patch-based\) priors](#)



Unconstrained Scene Generation with
Locally Conditioned Radiance Fields,
DeVries et al., arXiv 2021

Neural Sparse Voxel Fields (NSVF)

- Avoid sampling points in empty space as much as possible.




Illustration of Sparse Voxels





Illustration of a voxel-bounded neural field


Neural Sparse Voxel Fields, Liu et al. 2020

Neural Sparse Voxel Fields (NSVF)

- Avoid sampling points in empty space as much as possible.



Sample in the whole space



Only sample inside the sparse-voxels

Comparison




NeRF (Mildenhall et al. 2020)
(Rendering speed: 100 s/frame)




Ours (NSVF)
(Rendering speed: 2.62 s/frame)

Requirements


Scene Representation



Voxelgrids



Implicit Function



Hybrid
Implicit/Explicit



Renderer

Volumetric

Sphere-Tracing
Volumetric

Volumetric

Pros

Fast rendering

High quality
Compact
Admits *global* priors

Significant Speedup
Admits *local* priors

Cons

Memory $O(n^3)$
Limited spatial
resolution


Extremely expensive,
slow rendering

No compact
representation
No *global* priors


Neural Scene Representation and Neural Rendering

Neural Fields


Scene Representation



Voxelgrids



Implicit Function



Hybrid Implicit/Explicit

Renderer

Volumetric

Sphere-Tracing
Volumetric

Volumetric

Pros

Fast rendering

High quality
Compact
Admits *global* priors

Significant Speedup
Admits *local* priors

Cons

Memory $O(n^3)$
Limited spatial resolution

Extremely expensive,
slow rendering

No compact representation
No *global* priors

EUROGRAPHICS 2022
D. Meneveaux and G. Patané (Guest Editors)

Volume 41 (2022), Number 2
STAR – State of The Art Report

Neural Fields in Visual Computing and Beyond

Yibeng Xie^{1,2} · Towaki Takikawa^{3,4} · Shunsuke Saito⁵ · Or Litany⁶ · Shiqin Yan¹ · Numair Khan¹ · Federico Tombari^{6,7} · James Tompkin¹ · Vincent Sitzmann⁸ · Srinath Sridhar¹ · G. Patané

¹Brown University ²Unity Technologies ³University of Toronto ⁴NVIDIA ⁵Meta Reality Labs Research ⁶Google ⁷Technical University of Munich ⁸Massachusetts Institute of Technology ⁹Equal advising

<https://neuralfields.cs.brown.edu/>

Part I: Techniques
 $f_\theta: \mathbb{R}^m \rightarrow \mathbb{R}^n$
 Conditioning, Hybrid Representations, Forward Maps, Architectures, Manipulation

Part II: Applications
 2D and 3D Reconstruction, Generative Models, Digital Humans, Compression, Robotics, ...and Beyond!


Figure 1: Contribution of this report. Following a survey of over 250 papers, we provide a review of (Part I) techniques in neural fields such as prior learning and conditioning, representations, forward maps, architectures, and manipulation, and of (Part II) applications in visual computing including 2D image processing, 3D scene reconstruction, generative modeling, digital humans, compression, robotics, and beyond. This report is complemented by a community-driven website with search, filtering, bibliographic, and visualization features.

Abstract
 Recent advances in machine learning have led to increased interest in solving visual computing problems using methods that employ coordinate-based neural networks. These methods, which we call *neural fields*, parameterize physical properties of scenes or objects across space and time. They have seen widespread success in problems such as 3D shape and image synthesis, animation of human bodies, 3D reconstruction, and pose estimation. Rapid progress has led to numerous papers, but a consolidation of the discovered knowledge has not yet emerged. We provide context, mathematical grounding, and a review of over 250 papers in the literature on neural fields. In Part I, we focus on neural field techniques by identifying common components of neural field methods, including different conditioning, representation, forward map, architecture, and manipulation methods. In Part II, we focus on applications of neural fields to different problems in visual computing, and beyond (e.g., robotics, audio). Our review shows the breadth of topics already covered in visual computing, both historically and in current incarnations, and highlights the improved quality, flexibility, and capability brought by neural field methods. Finally, we present a companion website that acts as a living database that can be continually updated by the community.


CCS Concepts
 • Computing methodologies → Machine Learning; Artificial Intelligence;

A field is a quantity defined for all spatial and / or temporal coordinates.

Examples of Fields




Vector Field



Fields



Examples of Fields




Image

Vector Field

Fields




Examples of Fields



3D Parabola
(Explicit Surface)




Image




Vector Field


Fields



Examples of Fields




3D Signed Distance Fields
(Implicit Surface)



3D Parabola
(Explicit Surface)




Image




Vector Field


Fields



Examples of Fields




3D Signed Distance Fields
(Implicit Surface)



3D Parabola
(Explicit Surface)




Image




Vector Field

Fields




Audio




What are neural fields?


Fields / signals can be represented in many ways.



Continuous





Discrete





Neural

What are neural fields?





Magnetic Field




Eulerian Flow Field of a Fluid

[Koldora CC]


What are neural fields?



Neural Fields General Framework




Differentiable Rendering



Figures adapted from:
Mildenhall et al. 2020 (NeRF)
Sitzmann et al. 2019 (SRN)
Lingjie Liu

BRDF Shading



$$L(\mathbf{x}, \vec{\omega}_o) = L_e(\mathbf{x}, \vec{\omega}_o) + \int_{\mathcal{C}} f_r(\mathbf{x}, \vec{\omega}_i \rightarrow \vec{\omega}_o) L(\mathbf{x}', \vec{\omega}_i) G(\mathbf{x}, \mathbf{x}') V(\mathbf{x}, \mathbf{x}') d\omega_i$$

Course Link:

<https://neural-representation-2025.github.io/index.html>



TAs



Abhimanyu Suthar
Email: absutharus@gmail.com



Chuhao Chen
Email: chuhaoc@seas.upenn.edu

Tentative Syllabus


	Presentation	Paper Presenting Schedule
1	Sep 2 (Tue)	Introduction
2	Sep 4	Introduction 2
3	Sep 9 (Tue)	Presentation Round 1 Start ->
4	Sep 11	
5	Sep 16 (Tue)	
6	Sep 18	
7	Sep 23 (Tue)	
8	Sep 25	
9	Sep 30 (Tue)	
10	Oct 2	Presentation Round 1 End <-
11	Oct 7 (Tue)	Guest Talk 1
	Oct 9	No Class
12	Oct 14 (Tue)	Guest Talk 2

13	Oct 16	Presentation Round 2 Start ->
	Oct 21 - Oct 23	No Class
14	Oct 28 (Tue)	
15	Oct 30	
16	Nov 4 (Tue)	
17	Nov 6	
18	Nov 11 (Tue)	
19	Nov 13	
20	Nov 18 (Tue)	Presentation Round 2 End <-
21	Nov 20	Favorite Paper + New Ideas
22	Nov 25 (Tue)	Practice Lecture + Summary
	Nov 27	No Class
23	Dec 2	No Class (Tentative)
24	Dec 4	No Class (Tentative)

Next Class

1. Present some pioneering works in this field, e.g., NeRF, SRN, Neural Volumes, Gaussian Splatting, NeuS, ...
2. Fundamentals of Classical 3D Representations and Rendering in Computer Graphics

2	(Remarks: the papers within each blue block are similar, choosing one of them to share would be enough, doing a brief summary of others would be greatly welcome)	
3	Fast Training	<u>Taming 3DGS: High-Quality Radiance Fields with Limited Resources</u>
4		<u>DashGaussian: Optimizing 3D Gaussian Splatting in 200 Seconds</u>
5	High-quality Rendering	<u>Volumetrically Consistent 3D Gaussian Rasterization</u>
6	Generalisable 3D Reconstruction	<u>MVSplat: Efficient 3D Gaussian Splatting from Sparse Multi-View Images</u>
7		<u>Flash3D: Feed-Forward Generalisable 3D Scene Reconstruction from a Single Image</u>
8	Fast Visual Geometry	<u>DUST3R: Geometric 3D Vision Made Easy</u>
9		<u>Spann3R: 3D Reconstruction with Spatial Memory</u>
10		<u>Fast3R: Towards 3D Reconstruction of 1000+ Images in One Forward Pass</u>
11		<u>VGGSfM: Visual Geometry Grounded Deep Structure From Motion</u>
12		<u>VGGT: Visual Geometry Grounded Transformer</u>
13		<u>Pi³: Scalable Permutation-Equivariant Visual Geometry Learning</u>
14	Transformer-based Neural Rendering	<u>LVSM: A Large View Synthesis Model with Minimal 3D Inductive Bias</u>
15		<u>RenderFormer: Transformer-based Neural Rendering of Triangle Meshes with Global Illumination</u>
16	3D Generation	<u>MeshFormer: High-Quality Mesh Generation with 3D-Guided Reconstruction Model</u>
17		<u>Structured 3D Latents for Scalable and Versatile 3D Generation (TRELLIS)</u>
18	4D Reconstruction & Generation	<u>SV4D: Dynamic 3D Content Generation with Multi-Frame and Multi-View Consistency</u>
19		<u>SV4D 2.0: Enhancing Spatio-Temporal Consistency in Multi-View Video Diffusion for High-Quality 4D Generation</u>
20		<u>CAT4D: Create Anything in 4D with Multi-View Video Diffusion Models</u>
21		<u>MoSca: Dynamic Gaussian Fusion from Casual Videos via 4D Motion Scaffolds</u>
22		<u>Shape of Motion: 4D Reconstruction from a Single Video</u>
23		<u>DIMO: Diverse 3D Motion Generation for Arbitrary Objects</u>



Note: For a paper bundle, you only need to present one of the papers in the bundle according to their preference, but you are encouraged to discuss the connections between the papers in the bundle.

Physics-grounded Reconstruction & Generation	<u>Reconstruction and Simulation of Elastic Objects with Spring-Mass 3D Gaussians</u>	Oct 16 (Tue)
	<u>PhysTwin: Physics-Informed Reconstruction and Simulation of Deformable Objects from Videos</u>	
	<u>Vid2Sim: Generalizable, Video-based Reconstruction of Appearance, Geometry and Physics for Mesh-free Simulation</u>	Oct 16 (Tue)
	<u>Gaussian Splashing: Unified Particles for Versatile Motion Synthesis and Rendering</u>	Oct 28 (Tue)
	<u>PhysMotion: Physics-Grounded Dynamics From a Single Image</u>	Oct 28 (Tue)
<u>WonderPlay: Dynamic 3D Scene Generation from a Single Image and Actions</u>		
Multi-modality	<u>Dr. Splat: Directly Referring 3D Gaussian Splatting via Direct Language Embedding Registration</u>	Oct 30 (Thu)
3D Editing	<u>EditSplat: Multi-View Fusion and Attention-Guided Optimization for View-Consistent 3D Scene Editing with 3D Gaussian Splatting</u>	Oct 30 (Thu)
Robotics	<u>MASt3R-SLAM: Real-Time Dense SLAM with 3D Reconstruction Priors.</u>	Nov 4 (Tue)
	<u>GWM: Towards Scalable Gaussian World Models for Robotic Manipulation</u>	Nov 4 (Tue)
	<u>Splat-Nav: Safe Real-Time Robot Navigation in Gaussian Splatting Maps</u>	Nov 6 (Thu)
	<u>Pre-training Auto-regressive Robotic Models with 4D Representations</u>	Nov 6 (Thu)
Surface Reconstruction	<u>Objects as volumes: A stochastic geometry view of opaque solids</u>	Nov 11 (Tue)
	<u>Geometry Field Splatting with Gaussian Surfels</u>	Nov 11 (Tue)
3D Segmentation & Part-aware Generation	<u>Part123: Part-aware 3D Reconstruction from a Single-view Image</u>	Nov 13 (Thu)
	<u>HoloPart: Generative 3D Part Amodal Segmentation</u>	Nov 13 (Thu)
	<u>PartField: Learning 3D Feature Fields for Part Segmentation and Beyond</u>	
Video Generation	<u>GEN3C: 3D-Informed World-Consistent Video Generation with Precise Camera Control</u>	Nov 18 (Tue)
	<u>Voyaging into Perpetual Dynamic Scenes from a Single View</u>	Nov 18 (Tue)

Before the seminar

- Read the papers of the week.
- Submit at least two questions for discussion before the seminar to a Google form (<https://docs.google.com/forms/d/e/1FAIpQLSfSxryv JO9Ffbd7iKClqnczqPWJUqv3O GFI6K-2sAKOJmBYQ/viewform>). This is important – your contribution will be marked. The deadline for submitting questions is one hour before each class session (so Tuesday 2:30 PM and Thursday 2:30 PM).

During the seminar (Starting from Sept 9, two rounds)

- Overview (10 minutes)
 - The instructor or TAs give a brief introduction on the topic.
- 2x Presentations (each 25 minutes, 25 % of grade):
 - Two pre-assigned participants present the paper of their choice.
 - 5 minutes on motivation, background and related work.
 - 20 minutes of presentation of the paper.
- Discussion and Feedback (30 minutes, 25% of grade across weeks):
 - One participant is assigned at random at the beginning of the seminar to lead the discussion. Everyone leads the discussion at least once in the seminar series.
 - The discussion leader receives a digest of the submitted questions just before the seminar.
 - The discussion leader raises questions appropriately throughout the discussion, covers future work aspects, and finally provides a summary of the strengths and weaknesses of the techniques and of the discipline.
 - The students provide feedback to the presenting student on their presentation with respect to what has worked well, and what could be improved and how.

Grading Criteria

<p>Form (30%) <i>To time? Verbal speed & clarity? Body posture? Engagement with audience?</i></p>	<p>Moderation (30%) <i>Integrates questions well? Pushes forward discussion? Good summary? Strengths and weaknesses of paper?</i></p>	<p>Practice Lecture (30%) <i>Listen attentively to the lecture? Engage in small coding exercises?</i></p>
<p>Content (50%) <i>Structure/storyline? Main points? Paper connections? Valid conclusions?</i></p>	<p>Questions (70%) <i>One question per paper (two questions per class) should be submitted at least one hour before the class during which the paper will be presented. (However, students are permitted to submit questions late (but before the discussion), up to two occurrences, without facing any penalties)</i></p>	<p>Discussion on your favorite papers (35%) <i>Each person has 5 minutes to present their favorite paper. Is your presentation clear? Explain clearly why you chose it and what you like about it?</i></p>
<p>Answers (20%) <i>Good answers to questions? Knowledgeable?</i></p>		<p>Brainstorming (35%) <i>Actively participate in the discussion? Contribute your own ideas or opinions?</i></p>

2x Presentations
(50% of grade)

Discussion
(25% of grade)

Other Activity Participation
(25% of grade)

TODOs After this Class


1. **Paper Selection and Registration:** [Important! Deadline: Sept 9]

Please select and register for the two papers you would like to present using the following Excel link:

https://docs.google.com/spreadsheets/d/1_FJueXqWnKWoyOGZTNiwp2qRmSP0u1H6ayEYE5j3Ib0/edit?gid=0#gid=0

2. **Presentation Preparation:**

- Ensure you are fully prepared **one class before your scheduled class for presentation.**
- Upload your slides to the Google folder (<https://drive.google.com/drive/folders/1OIKCn562KA3sPUGh-9x8mFuqazU92Cyk>) **at least one hour before the class prior to your assigned class for presentation.** This is important in case of an emergency requiring us to reschedule your talk.
- For example, if you're presenting on Tuesday, upload your slides by the previous Thursday at 2:30 PM. If presenting on Thursday, upload by Tuesday at 2:30 PM.



3. **Class Participation:**

- Before each class, please read the papers that will be discussed and submit two questions **at least one hour before the class** using the following link:

<https://docs.google.com/forms/d/e/1FAIpQLSeVtMP-PzgtJ9ZthAisBP5MZZ5JqBDrUQ4qqNkusmqVrnquQ/viewform?usp=sharing&oid=101671697877703922510>

TODOs NOW

Find the first four speakers for the first two classes next week (Sept 9 and Sept 11).

Bonus: You can earn 3 bonus point for Tuesday's class and 2 bonus points for Thursday's class toward your final score!

Acknowledgments

- Advances in Neural Rendering
- Neural Fields in Visual Computing and Beyond
- awesome-NeRF: a curated list of awesome neural radiance fields papers
- MPII Summer Semester 2023: Computer Vision and Machine Learning for Computer Graphics

Any Questions?