


Master Lab, Master Seminar and BACHELOR PROJEKTGRUPPE


OFFERED BY THE
LOVC, GML & AISD Groups

INSTITUTE OF COMPUTER SCIENCE II VISUAL COMPUTING DEPARTMENT

LOVC Group (Prof. Dr. Florian Bernard):



Geometry in ML Group (Prof. Dr. Zorah Lähner):




AI aided Scientific Discovery Group (Prof. Dr. Vahid Babaei):



Schedule

1. General rules for Bachelor Projektgruppe/Master Lab/Master Seminar
2. **13** Topics for Bachelor Projektgruppe/Master Lab/Master Seminar
offered by Learning and Optimisation for Visual Computing Group (LOVC)
3. **3** Topics for Master Lab
offered by AI aided Scientific Discovery Group (AISDG)
4. **9** Topics for Bachelor Projektgruppe/Master Lab/Master Seminar
offered by Geometry in Machine Learning Group (GML)

Timeline




Overview

Workload:

- Literature reading (Seminar, MLab, BPG)
- Implementation (MLab, BPG)

Grading:

- Presentation
- Report




**MILESTONE
SETTING**

Milestones setting



Meeting (usually weekly)



Answering questions



**GIVING ADVICE ON A
PRESENTATION
REHEARSAL**

Presentation practice

Presentation

- Present your work at the end of the semester:
20mins talk + 5mins Q&A/person
- Slides for the talk:
Discuss the structure with your supervisor before you start.
Trial talk is always a nice option to see if the talk matches the expectations
- Basic requirements:
Slide numbers
References for figures or text that you directly took from somewhere

Report

- Hand in until **1 weeks after the presentation!**

- Written in LaTeX:

Template available at <https://cg.cs.uni-bonn.de/courses>.

BPG & M Lab (around **10** pages) and **M Sem** (around **5** pages) in length. (excluding the reference)

- Sound and comprehensible summary of provided literature and the experiments in your own words.

- Adhere to scientific standards:

Provide proof for your statements(if any).

Use references to make your sources transparent.

- State the work division if working in group

Report (NEW)

- **Only for Master Seminar.**
- It is possible to write the report as a blog post.
- The blog can:
 - Be published online (voluntarily).
 - Should be roughly equivalent to around 5 PDF pages.

Embargo restrictions (MLab Only)

1. For those affected by embargo restrictions, please also apply to the examination office
2. Application should be done together with registration

Topic selection

- Interested in a topic?

Contact Moritz Heep (mheep@uni-bonn.de) (≥ 3 topics - 20.10.2026)

- Slides for organization and topics are available on our homepage:

<https://lovc.cs.uni-bonn.de/index.php/teaching/>

<https://geometryinml.cs.uni-bonn.de/teaching/>

( slides only accessible while in Uni Bonn network)



Master Lab, Master Seminar and BACHELOR PROJEKTGRUPPE



OFFERED BY THE

**Learning and Optimisation
for Visual Computing (LOVC) GROUP**

INSTITUTE OF COMPUTER SCIENCE II VISUAL COMPUTING DEPARTMENT

LOVC1: Extending BeCoS benchmark


- BeCoS is a benchmark that provides ground truth correspondences between thousands of 3D deformable shapes from multiple categories.




- Reproduce (part of) the results. (B PG, M Lab)
- The results of current shape matching methods on BeCoS are suboptimal. Investigate the reason behind this lack of performance and propose new ways to deal with matching of a highly non-isometric dataset. (B PG, M Lab)

B PG	✓
M Sem	✗
M Lab	✓

LOVC2: Investigating Neural Functional Maps for Statistical Shape Models



- FUSS combines deep functional maps and neural deformation models to build a statistical shape model.
- HyperFM learns a neural functional map to overcome the linear assumptions of standard functional maps.



Project goals:

- Understand the background (both papers) in shape matching and deformation.
- Investigate the possibility of using neural functional maps in a joint matching and deformation framework.
- Implement and evaluate the model using both publicly available repositories.
- Report on the outcomes of the project in a report.


El Amrani et al. A Universal and Flexible Framework for Unsupervised Statistical Shape Model Learning (MICCAI 2024)

Cao et al. Hyper-Network Neural Functional Maps for Unsupervised Robust 3D Shape Matching (ECCV 2026)

B PG ✓
M Sem ✓
M Lab ✓

LOVC3: Partial Shape Matching using Coarse-to-Fine Language Guidance

- VLMs are getting more and more powerful and can be used for shape matching tasks.
- So far, only full shape matching has benefited from VLMs.
- In this project, we look into the question:
 - Can we use **Language Guidance** to perform better on **partial shape** matching benchmarks?



Project goals:

- Understand the background literature in full-to-full and partial-to-partial shape matching.
- Implement and evaluate the possibility of using VLMs for guiding matching.
- Report on the outcomes of the project in a report.

LOVC4: Multigrid Solvers for Irregular Image Domains

Domains


- Coarse-to-Fine approaches (for image processing tasks) and multigrid solvers (for solving PDEs) are well-established in their respective fields
- Combining both allows for example the reconstruction of a surface from its normals
- Performant solutions to PDEs in the image domain are underexplored
- **Goal:** Investigate multigrid strategies and apply them to irregular image domains



B PG	✓
M Sem	✗
M Lab	✓

LOVC5: One Target to Calibrate them All



- Modern lightstages combine a multitude of components, e.g. cameras and lights but potentially also laser scanners.
- All these components need to be calibrated (positions, orientations, etc.)
- **Goal:** Combine all these modalities in a single calibration target



Light dome



Camera and laser scanner calibration [2]



Camera Calibration [1]


[1] H. Santo, M. Waechter, W.-Y. Lin, Y. Sugano and Y. Matsushita: Light Structure from Pin Motion: Geometric Point Light Source Calibration. (IJCV, 2020)

[2] L. Makabe, H. Santo, F. Okura, and Y. Matsushita: Shape-coded aruco: Fiducial marker for bridging 2d and 3d modalities. (WACV, 2022)

LOVC6: Mesh-based Attention Schemes for Shape Matching

- Deep learning on triangle meshes is driven by the requirement that different triangulations of the same shape must produce the same results
- This invariance property has been used to develop convolutions [1, 2] and more recently an attention mechanism [3] for triangle meshes

Goal: Use latest developments in deep learning techniques for shape matching




[1] N. Sharp, S. Attaki, K. Crane and M. Ovsjanikov: *Diffusionnet: Discretization agnostic learning on surfaces.* (ToG, 2022)

[2] Maesumi et al.: *PoissonNet: A Local-Global Approach for Learning on Surfaces.* (ToG, 2025)

[3] A. Shetty, Z. Zhu, S. Pirk and N. Aigerman: *Intrinsic and Triangulation-Agnostic Attention: A Simple and Powerful Approach for Learning on Meshes.* (arXiv, 2026)

LOVC7: Low-Rank Adaptation (LoRA) for Weight-level Merging

- LoRA enables adaptation of pre-trained models to specific tasks. NP-LoRA adopts a geometric perspective on LoRA fusion, interpreting content and style LoRAs as occupying overlapping, non-orthogonal low-rank subspaces.



- Understand the core idea of the NP-LoRA paper. (B PG, M Sem, M Lab)
- Run (official) code implementation on in-the-wild images and find possible drawbacks of this method. (B PG, M Lab)
- Compare NP-LoRA with other approaches like K-LoRA and LoRA.rar . (M Lab).


[1] NP-LoRA: Null Space Projection for Subject-Style LoRA Fusion, Chen et al. 2026

[2] K-LoRA: Unlocking training-free fusion of any subject and style LoRAs, Ouyang et al. 2025

[3] LoRA.rar: Learning to merge lorAs via hypernetworks for subject-style conditioned image generation, Shenaj et al. 2025

LOVC8: Personality Modulation of Large Language Models by Model Merging

- Personality vectors are obtained by subtracting the parameters of a base LLM model from those of a fine-tuned personality model (e.g., openness, extraversion, agreeableness, etc.) LLMs can be made to exhibit desired personality traits (without additional training) by merging personality vectors.
- Understand the core idea of the paper. (B PG, M Sem, M Lab)
- Use official code implementation to reproduce results on newer base models and find possible drawbacks of this method. (B PG, M Lab)
- Adapt the method to use Low-Rank Adaptation (LoRA) to constrain updates to a low-rank subspace instead of the full-parameter space. (M Lab).




[1] Personality Vector: Modulating Personality of Large Language Models by Model Merging. Sun et al. 2025

[2] LoRA: Low-Rank Adaptation of Large Language Models. Hu et al. 2021

LOVC9: Model Merging-inspired Novel Eigenvoice Synthesis


- Murata et al. (2025) propose a speech synthesis method that defines a low-dimensional speaker space directly within the model parameters of a speech synthesis model. This enables controllable speech generation across attributes like gender and voice traits
- Understand the core idea of the paper. (B PG, M Sem, M Lab)
- Implement the method described in the paper from scratch and reproduce the results (M Lab)



B PG ✓
M Sem ✓
M Lab ✓

[1] Eigenvoice Synthesis based on Model Editing for Speaker Generation. Murata et al. 2025

LOVC10: FlowFeat: Pixel-Dense Embedding of Motion Profiles




- Understand the core idea of the paper. (B PG, M Sem, M Lab)
- Run official code implementation on other datasets and find possible drawbacks of this method. (B PG, M Lab)
- Improve possible drawbacks within the pipeline: longer time-dependent information during training. (M Lab)

[1] Araslanov, Nikita, Anna Sonnweber, and Daniel Cremers. "FlowFeat: Pixel-dense embedding of motion profiles." *NeurIPS* (2025)

B PG ✓
M Sem ✓
M Lab ✓

LOVC11: Multi-View Foundation Models




- Understand the core idea of the paper. (B PG, M Sem, M Lab)
- Run official code implementation on other datasets and find possible application with applying this method. (B PG, M Lab)
- Improve possible drawbacks within the pipeline including (1) better semantics, or (2) using video as input. (M Lab)

[1] Segre, Leo, Or Hirschorn, and Shai Avidan. "Multi-view foundation models." ECCV (2026)

B PG ✓
M Sem ✓
M Lab ✓

LOVC12: Neural Representations for Physics-based Inverse Geometry Design / Metamaterials




Metamaterials:

- “properties arise not from the chemical composition of its base substances, but from their deliberately designed internal structure”

Goals:

- Research Topology Optimization using Neural Methods (B PG, M Sem, M Lab).
- Develop a general approach for neural, physics-based inverse design of structures with desired mechanical behaviour. (B PG, M Lab)



[1] <https://www.youtube.com/watch?v=IsTiWYSfPck>



[2] Chandrasekhar, A., Suresh, K. TOuNN: Topology Optimization using Neural Networks. *Struct Multidisc Optim* 63, 1135–1149 (2021)

LOVC13: Speeding up Stellarator Design by learning Simulation Initialisation

Stellarators are a fusion device that requires careful geometric design which needs a lot of simulation.

Goals:

- Research Acceleration of Iterative Numerical Methods using ML. (M Sem)
- Develop a method to speedup simulation by finding initialisations from which the simulation converges quickly. (B PG, M Lab)



Arisaka, S. & Li, Q.. (2023). Principled Acceleration of Iterative Numerical Methods Using Machine Learning. ICML

B PG ✓
M Sem ✓
M Lab ✓

Master Lab, Master Seminar and BACHELOR PROJEKTGRUPPE

OFFERED BY THE




Geometry in Machine Learning (GML) GROUP

geometryinml.cs.uni-bonn.de

INSTITUTE OF COMPUTER SCIENCE II VISUAL COMPUTING DEPARTMENT



GML1-intro: Segmentation on NeRF / 3DGS

Segmentation on NeRF / 3DGS via contrastive learning



OmniSeg3D: Omniversal 3D Segmentation via Hierarchical Contrastive Learning (CVPR2024)


GML1-task: Segmentation on NeRF / 3DGS



1. Understand the original paper in NeRF and Present the core idea (M Sem)
2. Reproduce the paper in 3DGS (code available) (BPG/M Lab)
3. Replace hierarchical contrastive learning with Hyperbolic learning and retrain (BPG/M Lab)

B PG ✓
M Sem ✓
M Lab ✓

GML2: Multi-view consistent image generation with 2D diffusion model





1. Fully understand and able to present the core idea of the paper (M Sem)
2. Reproduce the multiview-consistent image generation part (B PG & M Lab)
3. Explore the multiview-consistency under different circumstances (eg. relighting) or try to improve the mesh reconstruction via 3DGS (B PG & M Lab)

B PG ✓
M Sem ✓
M Lab ✓


SyncDreamer 2024 ICLR.

GML3: Neural Field Editing (M Sem)


Natural deformations of 3D neural fields via user constraints:



Compare approaches and properties of two recent state of the art methods [1,2]




[1] Baieri et al, "Implicit-ARAP: Efficient Handle-Guided Neural Field Deformation via Local Patch Meshing", NeurIPS 2025.



[2] Djuren et al., "As-Rigid-As-Possible Regularization for Implicit Surfaces", SGP 2026.

B PG	✗
M Sem	✓
M Lab	✗

GML4: Interactive Neural Field Editing (M Lab/BPG)



- Developing solutions for known bottlenecks in Implicit-ARAP implementation
- Integration with Instant-NGP [3] positional encoding
- **Goal:** interactive viewer showing the neural field deformation as it is being optimized

[3] Müller et al., “Instant Neural Graphics Primitives with a Multiresolution Hash Encoding”. ACM Transactions on Graphics, 2022.

GML5: Build your own Inverse Renderer

Goal: Recover an object's materials and lighting from photos.

Step 1: Recover the color of a synthetic object.

Step 2: Add shiny materials (roughness, specular reflections).


Step 3: Recover the lighting with environment maps.

Step 4: Apply it to a real object you photograph yourself.

You'll learn: Rendering equation, BRDFs, differentiable rendering

Tools: Python, PyTorch, Mitsuba 3 or Nvdiffast

Prerequisites: Python and basic linear algebra; graphics experience is a plus



Photos → shape, material and light

GML6: Inverse Rendering on 3DGS

Goal: Turn photos of real objects into 3D assets that can be relit and edited

Capture: Photograph real objects and reconstruct them with 3D Gaussian Splatting

Geometry: Estimate reliable surface normals

Materials and lighting: Recover albedo and roughness, and learn an environment map


Shadows: Model which parts are occluded from light

Relighting: Render objects under new lighting and evaluate on real-world benchmarks

Team: 2 students, one on capture and geometry, one on materials and lighting

You'll learn: Gaussian Splatting, physically based shading, working with research code

Prerequisites: Python, PyTorch, a course in vision, graphics or deep learning




Liang, Zhihao et al. "GS-IR: 3D Gaussian Splatting for Inverse Rendering." 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023): 21644-21653.

B PG 
M Sem 
M Lab 

GML7-intro: Fusion Energy Generation

- Goal: safe and sustainable energy with fusion energy generation
- Requires plasma state (hot and high pressure)
- Current prototypes are not yet efficient enough
- Paper with dataset and background information:
<https://arxiv.org/pdf/2506.19583>







MPI for Plasma Physics

B PG	✓
M Sem	✗
M Lab	✓


GML7-intro: Stellarator - Plasma Shape Optimization

- Stellarator: fusion generator type for continuous operation (But more complex shape)
- Based on magnetic confinement
- Challenge: Many parameters to optimize (stability, magnetic field properties, coil shape, ...)





B PG	✓
M Sem	✗
M Lab	✓

GML7-task: Improve B-Spline Surface Representation for Plasma Shapes




$$R(\theta, \varphi) = \sum_{m=0}^M \sum_{n=-N}^N R_{m,n} \cos(m\theta - n\varphi),$$

$$Z(\theta, \varphi) = \sum_{m=0}^M \sum_{n=-N}^N Z_{m,n} \sin(m\theta - n\varphi),$$



$$S(x) = \sum_{j=0}^{n-1} c_j B_{j,k;t}(x)$$




- Improve conversion between angle-based Fourier Representation and B-Spline Surfaces
- Apply local deformations and analyse impact on physics metrics

This project is affected by embargo restrictions.

B PG	✓
M Sem	✗
M Lab	✓

GML8-intro: GLOT across Languages (B PG, M Sem, M Thesis)




Towards Improved Sentence Representations using Token Graphs, Mantri et al. ICLR 2026

- Sentence-level representations from pretrained LLMs
- Treat token embeddings as a latent graph where edges represent complex semantic dependencies

B PG ✓ M Sem ✓ M Lab ✓

GML8-task: Extending GLOT to Low-Resource Languages



Does the Token-GNN transfer to other languages?



- Same meaning, **different token dependencies** in every language
- Goal: **one shared Token-GNN**
- How to build tokenizer ? How to build the graph for each language?

B PG ✓ M Sem ✓ M Lab ✓


GML9-intro: Foundation Models for Astrophysics



- Both models are pretrained without labels on millions of optical survey galaxies, giving general-purpose embeddings that can be reused for new tasks without retraining.
- Do astro-specific embeddings beat general vision models (ImageNet, DINOv2), and how far beyond their training data do they still work?

B PJ ✓ **M Sem** ✓ **M Lab**

GML9: How Far Do Astro Foundation Models Transfer?



On data none of the models has seen, does astro pretraining still beat general vision pretraining, and where does that advantage break down?

B PJ ✓ M Sem ✓ M Lab


Master Lab

OFFERED BY THE

AI aided Scientific Discovery Group (AISD) GROUP


INSTITUTE OF COMPUTER SCIENCE II VISUAL COMPUTING DEPARTMENT

AISDG1: Label-scarce Regression from Pairwise Rankers



- Can we improve a data-driven regression model without collecting new labels?
- Explore how pretrained models (e.g., LLMs, Prior-fitted Networks) can provide pairwise rankings that refine numerical predictions when labeled data is scarce.


AISDG2: Inverse Design of Synthesizable Colorants



- Can we design a molecule with a desired color while ensuring they are synthesizable?
- Explore inverse molecular design for colorant using predefined chemical building blocks.

Wijaya et al. Trustworthy Inverse Molecular Design via Alignment with Molecular Dynamics. *Advanced Science* 2025.
 Lo et al. A Genetic Algorithm for Navigating Synthesizable Molecular Spaces. *ICLR* 2026.

AISDG3: Evolving Virtual Creatures in the Era of Scalable Computing



- Karl Sims showed that fairly complex virtual creatures can emerge from simple geometry and behavior primitives via genetic algorithm.
- Can we evolve more complex and surprising virtual creatures using modern scalable computing power?

Sims. Evolving Virtual Creatures. SIGGRAPH 1994.

Any questions?

- Interested in a topic?

Contact Moritz Heep (mheep@uni-bonn.de) (≥ 3 topics - 20.10.2026)

- Slides for organization and topics are available on our homepage:

<https://lovc.cs.uni-bonn.de/index.php/teaching/>

<https://geometryinml.cs.uni-bonn.de/teaching/>

( slides only accessible while in Uni Bonn network)

