

Kartik Bali

✉ kartikbali1297@gmail.com

☎ +49 1778896492

📍 Geesthacht, Germany

in linkedin.com/in/kartikbali1296

🐙 github.com/kbali1297

🎓 Google Scholar

🌐 kbali1297.github.io

PhD candidate in **3D computer vision, multimodal LLMs (VLMs), and generative models**. First author at **ECCV 2026** and **NeurIPS 2026**; interested in language-driven 3D understanding and VLM post-training.

Available for research internships: **Summer/Fall 2027**.

EDUCATION

- **Ph.D. Candidate, Artificial Intelligence**, TU Hamburg & Helmholtz-Zentrum Hereon *Jan 2025 – expected Dec 2027*
Advisors: Roland Aydin (DFKI) and Robert West (EPFL).
- **M.Sc. Computational Science & Engineering**, Technical University of Munich *Oct 2021 – June 2024*
Grade **1.5** (German scale; 1.0 = best); focus in Machine Learning and Simulation Sciences; Deutschland Stipendium.
- **B.Eng. Mechanical Engineering**, BITS Pilani, Hyderabad *July 2015 – May 2019*
Grade **9.63/10**; Merit Scholarship (top 2% academic performers).

EXPERIENCE

- **Doctoral Researcher**, Helmholtz-Zentrum Hereon, Hamburg *Jan 2025 – Present*
 - **Language-driven 3D grounding**: prompt-conditioned **view selection** over a **VLM** reasoning-segmentation backbone to localize CAD faces and edges on meshes; up to **1.7×** face IoU and **4.5×** edge F1 over CLIP-based baselines [1].
 - **Zero-shot physical reasoning of MLLMs**: multimodal LLMs localize stress-critical regions for **FEM mesh refinement**; without any FEM solve, closes **58%** of the gap to a reference-solve oracle [2].
 - **Active view selection under occlusion**: view-selection policy and CadQuery-generating **VLM** jointly trained with **RL on execution rewards**; **+21/+17** effective-mIoU over best baselines, $\leq 1.5\%$ invalid code; **mechanistic** view-token ablations [3].
 - **Point2BIM** (ongoing, with DFKI): real laser scans $\rightarrow$ executable **BIM programs** via VLMs; **SFT + GRPO** with execution rewards; multi-GPU training (4×H200), **vLLM** inference.
 - Co-author on LLM agents for scientific modeling [6, 7] and pretraining-time LLM alignment [4].
- **AI Engineer**, Project-S and Beyond Technologies, Munich: ML for **radar perception** *Sept 2024 – Dec 2024*
- **Master Thesis Intern**, BMW AG, Munich *Sept 2023 – May 2024*
 - Accelerated design of 2D/3D flow geometries by **2×** using **Physics-Guided Diffusion Models** and **GANs**, integrated with an autograd-based 3D topology optimization tool. Grade **1.0**.
- **Machine Learning Werkstudent**, TUM Graphics & Visualization (Thuerey group) *May 2022 – Aug 2023*
 - Implemented **Graph Neural Networks** and **Convolutional ResNets** for chaotic PDEs; built CI/CD and tests for PhiFlow; co-authored the resulting paper [5].
- **Software Engineer**, L&T Construction, Pune: data-pipeline automation (Praise Award) *July 2019 – Sept 2021*

PUBLICATIONS

- [1] **Bali, K.** and Aydin, R., “MV-GEL: Language-Driven Multi-View Geometric Entity Localization on Meshes”, *ECCV*, 2026 (**accepted**). arXiv:2606.31533. Code.
- [2] **Bali, K.**, Guru, M. K., Cyron, C. J., and Aydin, R., “GRFEM: Multimodal LLMs as Zero-Shot Semantic Assistants for Physics-Guided 3D Mesh Refinement”, *NeurIPS*, 2026 (**accepted, main track**). arXiv:2607.08798. Code.
- [3] **Bali, K.**, Guru, M., Borjigin, Y., Starostina, A., Cyron, C. J., and Aydin, R., “Look Where You Can: Active View Selection for CAD Reconstruction under Occlusion”, under review at *ICLR*, 2027. arXiv:2610.11954.
- [4] Minder, J., ..., **Bali, K.**, ..., Aydin, R., West, R., “Synthetic Persona Pretraining: Alignment from Token Zero”, preprint, 2026. arXiv:2608.13482.
- [5] List, B., Chen, L.-W., **Bali, K.**, Thuerey, N., “Differentiability in Unrolled Training of Neural Physics Simulators on Transient Dynamics”, *Computer Methods in Applied Mechanics and Engineering*, 2025.
- [6] Tacke, M., ..., **Bali, K.**, et al., “Constitutive Scientific Generative Agent (CSGA): Leveraging Large Language Models for Automated Constitutive Model Discovery”, *Mach. Learn. Comput. Sci. Eng.*, 2025.
- [7] Guru, M. K., Gupta, V., **Bali, K.**, et al., “From Finite Element Analysis to Self-Optimization: FRAME and the Era of Agentic AI in Engineering”, under review at *npj Computational Materials*, 2026.

SKILLS

- **ML/AI**: Vision-Language Models, Multimodal LLMs, Agentic VLMs, LLM Fine-tuning (LoRA), RL Post-training (GRPO/RLHF), Capability Evaluation of Foundation Models, Diffusion Models, GANs, VAEs, GNNs
- **Computer Vision**: 3D Scene Understanding, Point Clouds, 3D Reconstruction, Multi-View Geometry, Referring/Reasoning Segmentation, Semantic Segmentation (U-Net, PointNet), Meshes, SDFs
- **ML Systems & Deployment**: PyTorch (FSDP/DDP), HuggingFace (Transformers, PEFT), vLLM, FlashAttention-2, network pruning, Docker, Slurm, W&B, TensorBoard
- **Tools**: Python, C++, CUDA/Triton, JAX, OpenCV, CadQuery, BlenderProc, PyVista, Git, VS Code, Linux
- **Domain**: Scientific ML, Differentiable Simulation, PDEs, Physics-guided Generative Design
- **Languages**: English (C1), German (B1)