PRANEET CHANDRAKANT BALA

+1510-507-0333 | balapraneet@gmail.com | https://www.linkedin.com/in/praneet-bala | https://www.balapraneet.com | Saint Paul, MN

Ph.D. candidate in Computer Science at the University of Minnesota, researching self-supervised learning, 3D computer vision, and generative modeling, with 4 peer-reviewed publications in Nature Communications, IJCV, and eLife, and a paper under review at ECCV 2026. Hands-on work spans diffusion-based motion forecasting, multi-view 3D reconstruction, large-scale dataset creation, and neural rendering. Reviewer for CVPR, IJCV, and TPAMI.

EDUCATION

Ph.D. Candidate in *Computer Science,* *University of Minnesota - Twin Cities* — Sep 2019 - May 2026 (Expected)
Advisor: Dr. Hyun Soo Park and Dr. Jan Zimmermann

M.S. in *Computer Science,* *University of Minnesota - Twin Cities* — Sep 2017 - May 2025

B.E. in *Electronics Engineering,* *University of Mumbai, India* — Jul 2013 - May 2017

TECHNICAL SKILLS

- **3D Vision & Rendering:** Gaussian Splatting, NeRF-style neural rendering, COLMAP (SfM), 3D Reconstruction, Multi-view Geometry, Pose Estimation, Visual Odometry, Depth Estimation, Scene Understanding
- **Generative & Foundation Models:** Self-supervised Learning, Diffusion Models, Spatiotemporal Transformers, ControlNet, CLIP-guided generation, Contrastive Learning, Domain Adaptation, Multimodal Learning, Computer Graphics
- **System & Frameworks:** Python (NumPy, Pandas, Matplotlib, Scikit-learn), C/C++, PyTorch, TensorFlow, Keras, Cuda C, JavaScript, OpenCV, Blender, Unreal engine, Git, MATLAB

RESEARCH & PROJECTS

Learning to Forecast Domain Aware 3D Body Motion from Videos [*Python, PyTorch, MATLAB*] — Jul 2024 - Present
- Developed a self-supervised domain adaptation framework to predict 3D motion from videos without 3D annotations.
- Integrated a spatiotemporal transformer with latent diffusion to enhance short-term accuracy and long-term motion realism while reducing error accumulation. *(Under review at ECCV'26)*

Self-supervised Secondary Landmark Detection via 3D Representation Learning [*Python, TensorFlow*] — Sep 2020 - Mar 2023
- Developed a self-supervised framework using 3D representation learning and contrastive learning to detect secondary landmarks without manual annotations, scalable to large unlabeled video datasets.
- Validated generalization across multi-view settings and diverse species including macaques, humans, and flies.

CLIP-Guided Pose-Conditioned Image Generation with ControlNet [*Python, PyTorch*] — Mar 2024 - Oct 2024
- Fine-tuned Stable Diffusion via ControlNet on 80,000 macaque images using CLIP text conditioning and 2D pose control signals, demonstrating hands-on experience with conditional generative model training and post-training adaptation.
- Evaluated generation quality perceptually across pose transfer and facial fidelity metrics, identifying failure modes and informing future fine-tuning strategies.

3D Scene Reconstruction via Gaussian Splatting [*Independent Project*] — Oct 2025 - Present
- Built a video-to-3D pipeline for environment reconstruction from turntable videos using COLMAP and Gaussian Splatting, enabling interactive neural rendering and scene-level optimization.
- Evaluated reconstruction quality against NeRF-style baselines; optimized scene representations for real-time neural rendering.

OpenMonkeyStudio: Markerless 3D Pose Estimation [*Python, TensorFlow, MATLAB*] — Sep 2018 - Jun 2020
- Engineered a 62-camera system, mounted around an enclosure, to track primate motion in 3D space and study the behavior.
- Built a public dataset from multi-view image streams, with 195,228 pose annotations (with 13 landmarks) across diverse activities.

PUBLICATIONS

[4] **Bala** et al., 'Self-supervised Secondary Landmark Detection via 3D Representation Learning'. *IJCV*'23.

[3] Yao, **Bala** et al., 'OpenMonkeyChallenge: Dataset and Benchmark Challenges for Pose Estimation of Non-human Primates'. *IJCV*'23.

[2] Desai, **Bala** et al., 'OpenApePose, A Database of Annotated Ape Photographs for Pose Estimation'. *eLife*'23.

[1] **Bala** et al., 'Automated Markerless Pose Estimation in Freely Moving Macaques with OpenMonkeyStudio'. *Nature Comm*'20.

UNDER REVIEW: **Bala** et al., 'Learning to Forecast Domain Aware 3D Body Motion from Videos'. *ECCV'26*

WORK EXPERIENCE

Graduate Research Assistant, *University of Minnesota, Twin Cities* — May 2019 - Present
- Led R&D across 6 research projects spanning self-supervised learning, 3D reconstruction, generative modeling, multi-camera system design, and large-scale dataset creation, resulting in 4 published papers.
- Released 3 public datasets totaling ~400,000 labeled images; built a real-time 4-view camera system for live 3D reconstruction.

Graduate Teaching Assistant, *University of Minnesota, Twin Cities* — Jan 2019 - May 2024
- *Courses*: Computer Vision, Machine Architecture and Organization, Programming Interactive Computer Graphics and Games.

PROFESSIONAL SERVICE

- **Reviewer:** CVPR '22, '23, '24, IJCV '22, '24, '25, TPAMI '25 (recognized with certification)
- **Invited talks:** UMN Center for Neuroengineering seminar series; UMN Visual Computing & AI seminar series.