TARUN KALLURI

Foundation models, multimodal post-training, AI safety and evaluation

Website | Google Scholar | LinkedIn | GitHub | tarun.05.kalluri@gmail.com

EXPERIENCE

Meta Superintelligence Labs *Sep 2024 – Present*
Senior Research Scientist, Multimodal Post-Training (Llama-4, Muse Spark)

- Drove multimodal post-training data strategy, including data quality, composition, ablations, and mix decisions adopted in Llama-4 public release and subsequent product iterations.
- Contributed visual reasoning and perception post-training data improvements to Muse Spark.
- Designed taxonomy-driven prompt categorization and difficulty-aware ML classifiers to curate high-signal post-training data from large-scale corpora, yielding 7× training speedup and double-digit gains on human preference evals and public multimodal benchmarks; shipped to tens of millions of DAU on WhatsApp, Instagram, and Facebook through Meta AI .

Google AI Research *Summer 2023*
Research Intern | Mentors: Jeremiah Liu, Kihyuk Sohn

- Leveraged text-to-image generative models to synthesize domain-specific aerial training data, achieving >10% improvement in **cross-domain robustness** for real-time disaster assessment across low-resource geographies.

Facebook AI Research (FAIR) *Summer 2021*
Research Intern | Mentors: Du Tran, Lorenzo Torresani, Heng Wang

- Proposed a unified framework for **open-world instance segmentation** combining bottom-up grouping with top-down recognition, achieving 5–10% mAR improvement on unseen categories and setting SOTA on UVO and MS-COCO.

Facebook AI Research (FAIR) *Summer 2020*
Research Intern | Mentors: Du Tran, Deepak Pathak

- Developed FLAVR, a 3D space-time convolutional architecture for **video frame interpolation** achieving state-of-the-art quality with 6× faster inference. **Best Paper Finalist** at WACV 2023.

IIIT Hyderabad *Sep 2017 – Aug 2019*
Student Researcher | Mentor: C.V. Jawahar

- Proposed a universal feature alignment framework for **semi-supervised semantic segmentation**, achieving SOTA using only 50 labeled images across diverse road scenes. Published at ICCV 2019.

EDUCATION

University of California San Diego (UCSD) *2019 – 2024*
Ph.D., Computer Vision & Machine Learning | Advisor: Manmohan Chandraker
Developed methods to improve robustness and fairness of frontier AI models across low-resource domains via domain adaptation, with multiple **first-authored publications** at CVPR, ICCV, ICML, and ECCV.

Indian Institute of Technology (IIT) Guwahati *2012 – 2016*
B.Tech, Electronics & Communication Engineering | Minor: Computer Science

SELECTED PUBLICATIONS [FULL LIST ON GOOGLE SCHOLAR]

[1] **Scaling Agent Learning via Experience Synthesis.** Z. Chen, Z. Zhao, . . . , **T. Kalluri**, . . . , J. Weston, D. Huynh. *ICLR, 2026.*

[2] **Tell, Don't Show! Language Guidance Eases Transfer Across Domains in Images and Videos.** **T. Kalluri**, B. Majumder, M. Chandraker. *ICML, 2024.*

[3] **UDA-Bench: Revisiting Common Assumptions in Unsupervised Domain Adaptation Using a Standardized Framework.** **T. Kalluri**, S. Ravichandran, M. Chandraker. *ECCV, 2024.*

[4] **Robust Disaster Assessment from Aerial Imagery Using Text-to-Image Synthetic Data.** **T. Kalluri**, J. Lee, K. Sohn, S. Singla, M. Chandraker, J. Xu, J. Liu. *CVPR Workshop, 2024.*

[5] **Open-World Instance Segmentation: Top-Down Learning with Bottom-Up Supervision.** **T. Kalluri**, W. Wang, H. Wang, M. Chandraker, L. Torresani, D. Tran. *CVPR Workshop, 2024.*

[6] **GeoNet: Benchmarking Unsupervised Adaptation across Geographies.** **T. Kalluri**, W. Xu, M. Chandraker. *CVPR, 2023.*

[7] **MemSAC: Memory Augmented Sample Consistency for Large-Scale Domain Adaptation.** **T. Kalluri**, A. Sharma, M. Chandraker. *ECCV, 2022.*

[8] **FLAVR: Flow-Agnostic Video Representations for Fast Frame Interpolation.** **T. Kalluri**, D. Pathak, M. Chandraker, D. Tran. *WACV, 2023.* **[Oral, Best Paper Finalist]**

[9] **Universal Semi-Supervised Semantic Segmentation.** **T. Kalluri**, G. Varma, M. Chandraker, C.V. Jawahar. *ICCV, 2019.*

HONORS & AWARDS

- **Best Paper Finalist**, WACV 2023 for FLAVR (top 12 out of 641 papers, top 2%). *2023*
- **Best Reviewer**, ICLR 2022 and NeurIPS 2022. *2022*
- **IPE Ph.D. Fellowship**, UC San Diego for research in practical ethics in AI. *2020–21*

LEADERSHIP & SERVICE

- **Co-Organizer:** *Workshop on Demographic Diversity in AI*, CVPR 2025 (150+ participants).
- **Main Organizer:** *Workshop & Challenge on Robust CV Across Geographies*, ICCV 2023 (50+ teams).
- **Co-Organizer:** *Multi-Object Tracking & Segmentation in Complex Environments*, ECCV 2022 (100+ teams).
- **Reviewer:** NeurIPS, CVPR, ICCV, ICLR, ECCV, AAAI, TMLR.
- **Open Source:** 550+ GitHub stars across multiple widely adopted codebases and datasets. [Link]

TECHNICAL SKILLS

ML / Training	PyTorch, JAX / FLAX, distributed training (FSDP), SFT / RLHF pipelines, TensorFlow
Data & Infra	Large-scale data curation, evaluation frameworks, Python, C++, SQL, Bash
Tools	Git, OpenCV, Jupyter, NumPy, Pandas, Matplotlib

MENTORSHIP

- **Astuti Sharma**, now Research Engineer at Google (co-authored CVPR 2021 & ECCV 2022).
- **Sreyas Ravichandran**, now Researcher at ServiceNow (co-authored ECCV 2024).
- **Wangdong Xu**, M.S. student, UC San Diego (co-authored CVPR 2023).