Kartik Patwari

kpatwari@ucdavis.edu | kartikpatwari | scholar | kartikp7.github.io | kartikp7

Research Interests

Vision Security & Privacy, Multimodal LLMs and Understanding, Synthetic Data Generation, Domain Adaptation

Current Position

- **Research Scientist at MBZUAI Institute of Foundation Models** Mar. 2026 - Present
  Team: World Model, VLM *Sunnyvale, CA*
  - Worked on **multimodal LLM mid-training** and **supervised fine-tuning (SFT)** to improve visual reasoning, instruction following, and multimodal understanding.
  - Built scalable data pipelines for **cleaning, filtering, re-captioning,** and preprocessing of **large-scale SFT data**.
  - Designed **data mixture, difficulty scoring,** and **task balancing** strategies across captioning, VQA, reasoning, math, OCR, and general-knowledge tasks.
  - Performed **benchmark error analysis** on MMStar, MMMU, MathVista, DynaMath, LogicVista, and MMBench to **guide data curation**, re-filtering, knowledge distillation, and mixture reweighting/balancing.

Education

- **Ph.D. Computer Engineering** Mar. 2026
  University of California, Davis
- **M.S. Computer Engineering** Mar. 2024
  University of California, Davis
- **B.S. Computer Engineering (Major), Computer Science (Minor)** Dec. 2020
  University of California, Davis

Select Publications

(*equal contribution) | see Google Scholar for all.

[CVPR'26] **K. Patwari**, N. Vesdapunt, C. Wang, D. Li, C.P. Huynh, N. Zhou, C-N. Chuah, K.K. Fu. Composite-Attribute Person Re-Identification via Pose-Guided Disentanglement. to appear in IEEE/CVF Computer Vision and Pattern Recognition (CVPR), June 2026.

[FG'26] **K. Patwari***, D. Schneider*, X. Sun, C-N. Chuah, L. Lyu, V. Sharma*. Privacy-Complaint Human Data Synthesis in Images. IEEE Conference on Automatic Face and Gesture Recognition (FG), May 2026.

[WACV '26] **K. Patwari***, D. Chen*, Z. Lai, X. Zhu, S. Cheung, C-N. Chuah. Empowering Source-Free Domain Adaptation via MLLM-Guided Reliability-Based Curriculum Learning, to appear in IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), March 2026.

[ICML '24] **K. Patwari***, C-N. Chuah, L. Lyu, V. Sharma*. PerceptAnon: Exploring the Human Perception of Image Anonymization Beyond Pseudonymization for GDPR. International Conference on Machine Learning (ICML), July 2024.

[EuroS&P '22] **K. Patwari**, S. M. Hafiz, H. Wang, H. Homayoun, Z. Shafiq, C-N. Chuah. DNN Model Architecture Fingerprinting Attack on CPU-GPU Edge Devices. IEEE European Symposium on Security and Privacy (EuroS&P), June 2022.

[TMLR '23] A. Chhabra, **K. Patwari**, C. Kuntala, Sristi, D. Sharma, P. Mohapatra (2023). Towards Fair Video Summarization. Transactions on Machine Learning Research, December 2023

Work Experience

- **AI Researcher Intern at Cisco Systems** Sep. 2025 – Dec. 2025
  Team: AI Defense; MLLM Safety and Data *San Jose, CA*
  - Developed **VLM-as-a-judge** pipelines for **image safety relabeling** and **data filtering**.
  - Led **VLM image safety understanding**, my curated and relabeled data improved **F1** by ∼**15%**.
  - Exposed **vision-based prompt injection and jailbreak attacks** in frontier open-source **multimodal LLMs**.
- **Applied Scientist Intern at Amazon** Apr. 2025 – Aug. 2025
  Team: Amazon Ring Devices *Sunnyvale, CA*
  - Designed a new **image+text person re-identification** and retrieval setting by relabeling existing datasets with text supervision from multiple **VLM-as-a-judge** models.
  - Developed a **CLIP-based multimodal retrieval framework** with feature disentanglement and a disentangling loss for conditional image retrieval.
  - Achieved **SOTA** results on person image retrieval benchmarks; paper accepted at **CVPR 2026**.
- **Research Intern at Sony AI** Jun. 2023 – Sep. 2023
  Team: Privacy-Preserving Machine Learning (PPML) *Tokyo, Japan*

- Developed **privacy-preserving anonymization** pipelines for full-body and face images using masking, blurring, inpainting, and diffusion/GAN generation.
- Built **PII detection** and data curation workflows to generate supervision for diffusion-based human data synthesis.
- Research contributed to first-author papers at **ICML 2024** and **FG 2026**.

- **Research Engineer Intern at Sony** Jul. 2022 – Sep. 2022
  Team: Sony Semiconductor Solutions (SSS) – Imaging & Sensing *Tokyo, Japan*
  - Evaluated **3D reconstruction** pipelines from images, including SfM, MVS, depth estimation, and mesh generation.
  - Benchmarked classical and learning-based reconstruction methods on custom imaging datasets.
  - Adapted **SOTA deep reconstruction** methods for integration into an internal imaging and sensing pipeline.

Ongoing Research

- **Video Anonymization for Privacy-Preserving Video Pretraining** *Sep. 2025 – Present*
  SonyAI, UC Davis | Role: Project Lead
  - Proposed a **diffusion-refined video anonymization** pipeline for **privacy-preserving large-scale video pretraining.**
  - Built **video preprocessing** and **evaluation** workflows to obfuscate identity while preserving **motion, scene, and activity cues.**
  - Benchmarked **utility and privacy** using activity recognition, temporal consistency, person re-ID, and DP training
  - Under submission at **ECCV 2026.**
- **LLM Interpretability for Dementia Analysis using Visual Context** *Nov. 2025 – Present*
  UC Davis | Role: Project Lead
  - Studied **LLM SFT** for transcript-based dementia detection using linguistic feature guidance and visual context.
  - Analyzed **token-level attributions** to evaluate alignment with clinically meaningful linguistic cues.
  - Under submission to **ARR March Cycle.**
- **Reliable Medical Vision-Language Model Alignment** *Oct. 2025 – Present*
  UC Davis | Role: Collaborator
  - Collaborating on **preference optimization** methods to **improve factuality and visual grounding** in medical VLMs.
  - Evaluating medical VQA and radiology report generation under robustness- and reliability-focused settings
  - Under submission at **ECCV 2026**.
- **From Queries to Clones: Vision Encoder Stealing Attacks** *Sep. 2025 – Feb. 2026*
  UC Davis | Role: Project Lead
  - Benchmarked **black-box encoder stealing** attacks against CLIP, DINO, and task-specific vision encoders.
  - Studied distribution shift, query budgets, and cross-modal text guidance for foundation encoder theft
  - Under review at **TMLR**.

Technical Skills

- **Relevant Courses:** Machine Learning, Vision and Language Research, ML Hardware, Image Processing
- **Programming & Tools:** Python, C/C++, CUDA, Docker, Git, Jupyter, Conda, Latex
- **Programming/Frameworks:** PyTorch, PyTorch3D, HuggingFace, OpenCilk, OpenCV, OpenMP, Scikit-Learn

Teaching / Mentoring

- **Lead Teaching Assistant** *Fall 2022 - 2024; Winter 2023 - 2025*
  EEC 193/174AY: Applied ML Senior Design University of California, Davis
  - Developed assignments for image classification, object detection & tracking, segmentation & inpainting.
  - Gave lectures on ViT, security & privacy in ML, model compression & optimization.
  - Mentoring & leading teams in projects related to computer vision, scene understanding, autonomous driving.
  - Won ECE Best TA and Outstanding Graduate Student Teaching Awards.

Professional Service

- **Reviewer** | NeurIPS 2026 | CVPR 2026 | AAAI 2026 | AISTATS 2026,2025 | ICIP 2026 | VISION Workshop 2025,2024 | DataCV Workshop 2025 | ACM Computing Surveys 2024 | IEEE IoT Journal 2024

Certifications

- **NVIDIA** Fundamentals of Accelerated Data Science *March 2022*

AWARDS

- **WACV 2026 Doctoral Consortium** *Jan 2026*
  IEEE/CVF Winter Conference on Applications of Computer Vision
- **Outstanding Graduate Student Teaching Award** *June 2025*
  Graduate Studies, UC Davis
- **ECE Best Teaching Assistant Award** *May 2024*
  Electrical and Computer Engineering (ECE), UC Davis
- **Smita Bakshi Digital Learning and Teaching Award** *May 2024*
  Electrical and Computer Engineering (ECE), UC Davis
- **Advanced to Candidacy (AC) Fellowship** *April 2024*
  Electrical and Computer Engineering (ECE), UC Davis
- **EuroS&P Conference Student Grant** *May 2022*
  IEEE EuroS&P 2022, Genoa
- **ECE Outstanding Senior Design Project Award** *June 2020*
  Electrical and Computer Engineering (ECE), UC Davis