Kunjal Panchal

(413) 210 9198 | kpanchal@umass.edu | astuary.github.io/Kunjal/ | github.com/astuary | linkedin.com/in/kunjal-panchal

Skills

Research Areas	Federated Learning, Distributed Training, Large Language Models (LLMs), Multimodal Learning (vision-language, video-language), Agentic Reasoning, Agentic Planning, Personalization, Drift Adaptation, On-device Inference, Edge AI.
Techniques	Model Compression, Quantization, Memory- and Compute-Efficient Training/Inference, Distributed & Parallel ML Algorithms, Test-Time Adaptation, Optimization.
Frameworks	PyTorch, Hugging Face, Flower (Federated Learning Framework), Torch Distributed, PyTorch Mobile, ExecuTorch.

Education

University of Massachusetts Amherst — *Amherst, MA*

Doctor of Philosophy in Computer Science (3.7/4.0 GPA) — *Sep 2021 - May 2026*

- Advisor: Dr. Hui Guan.
- James Kurose Scholar (scholarship given for systems in machine learning project, Spring 2022).
- Jumpstart Fellow (fellowship given to top 5 research proposals by new PhD students, Fall 2021).

University of Massachusetts Amherst — *Amherst, MA*

Master of Science in Computer Science Research Track (3.6/4.0 GPA) — *Sep 2019 - May 2021*

- Advisor: Dr. Adam O'Neill.
- Research Area: Relaxed Cryptography for Digital Signatures and Message Authentication Codes.
- **Courses:** Machine Learning, Computer Vision, Natural Language Processing, Reinforcement Learning, Robotics, Optimization in Computer Science, Advanced Algorithms, Modern Computer Architecture, Advanced Cryptography.

The Maharaja Sayajirao University of Baroda — *Vadodara, India*

Bachelor's in Engineering in Computer Science and Engineering (4.0/4.0 GPA) — *Aug 2015 - Apr 2019*

- Gold medalist in Computer Science and Engineering (2019).
- Student of the Year (class of 2019).

Work Experience

Adobe Research — *San Jose, CA*

Research Scientist/Engineer — *May 2026 - Current (Full-time)*

- Conducting research to design, develop, and optimize machine learning and AI models, for scalable, real-world applications.

Adobe Research — *San Jose, CA*

Research Scientist/Engineer Intern — *May 2025 - Aug 2025 (Full-time); Aug 2025 - Nov 2025 (Part-time)*

- Designed an LLM planning algorithm for design personalization, introducing a novel **belief-shift scoring rule** and preference extractor.
- Designs (ads/posters) generated by the proposed system were chosen by **68% of designers as top-1 (93% top-2)** and rated **4.31/5**, outperforming baselines and aligning with AI-judge rankings.

Adobe Research — *San Jose, CA*

Research Scientist/Engineer Intern — *May 2024 - Aug 2024 (Full-time); Sep 2024 - Nov 2024 (Part-time)*

- Built an **on-device video inference pipeline** (Android Snapdragon) with PyTorch Mobile and quantization, achieving **3× lower memory usage**.
- Decomposed a billion-parameter video-language model into reusable modules, enabling cross-task reuse (captioning, reasoning, retrieval) and **cutting end-to-end latency by 27–33% on smartphones**, with only $\lesssim$2.3% accuracy drop.
- Refactored the visual-language model to support statically-typed forward passes and data-dependent control flows, reducing inference latency by 16.67%. Additionally optimized memory consumption through operator fusion and parameter hoisting techniques.

Adobe Research *San Jose, CA*

Research Scientist/Engineer Intern *May 2023 - Aug 2023 (Full-time); Aug 2023 - Dec 2023 (Part-time)*

- Achieved an increase of 4.74 units for Rouge score and **3.60% for Accuracy@1 improvements** for **few-shot learning in Flan-T5 transformer**, by expanding their capacity to be able to process more in-context example within the same context window length through sub-batching.
- Inched closer to **finetuning-like performance through pure in-context learning** (ICL) by 2.16 units of Rouge score and 3% for Accuracy@1 through mesa-optimization where the transformer acts like an optimizer itself during inference.
- **Improved the cross-domain transfer capabilities** of a transformer (Flan-T5) by 1.68 units for Rouge score and 1.3% for Accuracy@1 through incorporating both cross- and within domain question-answer samples within a limited context window length of 512 tokens.

Adobe Research *Bangalore, India*

Research PhD Intern *May 2022 - Aug 2022*

- Built a federated solution of **personalized recommendation systems** and classifiers for **real-time on-device learning**, by using early-stopping for client-side updates and drift adaptation at server-side, to achieve robustness against concept drift (distribution change with respect to time).
- Presented a **drift-aware adaptive optimization** strategy that can quickly adapt to various concept drift patterns (sudden, incremental, and recurrent), by taking into account historical gradient updates and identifying change in gradient magnitude as drifts, to achieve lowest accuracy drop and fastest recovery from the said drifts.
- Evaluated the proposed algorithm on benchmark computer vision and natural language processing tasks, **achieving the lowest accuracy dip difference** (the lower, the better) of 1.48%-2.99%, while the best performing baselines exhibit 3.15%-9.22%.

SureStart *New York City, NY (Remote)*

AI and Machine Learning Head Mentor of MIT FutureMakers 2022 *Jul 2022 - Aug 2022; Jun 2021 - Aug 2021; Feb 2021 - Mar 2021*

- Led technical sessions in a **6 weeks workshop program on applied deep learning**, as a **head mentor of 50+ U.S.-born students from marginalized communities**, by daily presenting and teaching deep learning concepts.
- Contributed in **curriculum building** to support the **daily discussion sessions** on the nuances of applied deep learning concepts like optimization, generative networks, algorithmic biases, regularization.
- Managed teams of 5 in multiple SureStart programs and guided the teams to build a deep learning based **capstone project addressing real-world challenges** like: (a) finding an optimal trade-off between utility and harmful effects of fracking **(Winner)**, (b) awareness on harmful ingredients in processed food **(Winner)**, (c) marine pollution **(Runner up)**, (d) automotive safety **(Winner)**, and climate change.

Research

Mosaic: Runtime-Efficient Multi-Agent Embodied Planning

Kunjal Panchal, Saayan Mitra, Sunav Choudhary, Victor Bursztyn, Somdeb Sarkhel, Hui Guan

In Submission.

The Cost of Avoiding Backpropagation

Kunjal Panchal, Sunav Choudhary, Yuriy Brun, and Hui Guan

In Submission.

ATOM: Efficient On-Device Video-Language Pipelines Through Modular Reuse

Kunjal Panchal, Saayan Mitra, Haoliang Wang, Somdeb Sarkhel, Ishita Dasgupta, Gang Wu, Hui Guan

Published @ ACM MMSys, 2026.

Thinking Forward: Memory-Efficient Federated Finetuning of Language Models

Kunjal Panchal, Nisarg Parikh, Sunav Choudhary, Yuriy Brun, and Hui Guan

Published @ NeurIPS, 2024.

Flash: Concept Drift Adaptation in Federated Learning

Kunjal Panchal, Sunav Choudhary, Koyel Mukharjee, Subrata Mitra, Somdeb Sarkhel, Saayan Mitra, and Hui Guan

Published @ ICML, 2023.

Flow: Per-instance Personalized Federated Learning

Kunjal Panchal, Sunav Choudhary, Nisarg Parikh, Lijun Zhang, and Hui Guan

Published @ NeurIPS, 2023; Preliminary Presentation @ CrossFL, MLSys 2022.

CommunityBots: Creating and Evaluating A Multi-Agent Chatbot Platform for Public Input Elicitation

Zhiqiu Jiang, Mashrur Rashik, **Kunjal Panchal**, Mahmood Jasim, Ali Sarvghad, Pari Riahi, Erica DeWitt, Fey Thurber, and Narges Mahyar

Published @ ACM CSCW 2023.

Leadership / Volunteering

Current	**Peer Reviewer**, NeurIPS '26/'25 (Top Reviewer)/'24, ICML '26/'25, ICLR '26/'25, TNNLS '26, AIStats '26/'25, AAAI '26/'25/'24, ACM MM '25, TSE '25
Oct 2025	**Workshop Organizer**, Short-Form Video Understanding @ ICCV 2025
Oct 2024	**Poster Presenter**, UMass Amherst CS Department Undergrad Research Night
	Research Mentor for Undergraduates, UMass Amherst CS department program to cultivate interest & understanding in research (Dec 2024, Jun 2024, Dec 2023, Jun 2023, Dec 2022)
	Applied Deep Learning Head Mentor, Teaching applied deep learning to undergradutes at SureStart (Volunteer position) (Jun 2024, Jun 2023, Jan 2023, Jun 2022)
Nov 2023	**Panelist and Poster Presenter**, UMass Amherst CS Department Undergrad Research Night
Jan 2022	**Coding Gym Leader**, SureStart winter bootcamp to teach coding interview strategies
Oct 2021	**PhD Applicant Support Program**, Mentoring prospective PhD applicants
Mar 2021	**Machine Learning Mentor**, Virtual AI Learning Program hosted by SureStart
Aug 2020	**Emotion AI Program Mentor**, EMPath Program hosted by Affectiva
Dec 2019	**Campus Leader**, Google Developer Students Club India

Achievements

2022	**James Kurose Scholarship**, Manning College of Info and Comp Sci, UMass Amherst
2021	**CICS Jumpstart Fellowship**, College of Info and Comp Sci, UMass Amherst
2019	**Gold Medalist**, The Maharaja Sayajirao University of Baroda, B.Engg. in Computer Science
2019	**Student of the Year**, The Maharaja Sayajirao University of Baroda, B.Engg. in Computer Science
	National Talent Search Examination, Top 100 in Science and Mathematics in India
	All India Essay Writing Event, Honorable Mention in a state-level essay competition
	Community Science Center, Winner of Conmat Cosmopolitan Tree Garden Award at state-level

Presentations

Career Pathways Seminar Speaker *Spring 2025*

- Delivered an introductory talk on privacy-preserving machine learning and on-device inference to second-year undergraduates, covering industry applications and open research challenges.

Voices of Data Science Poster Presenter *Spring 2023*

- Showcased "Flash: Concept Drift Adaptation in Federated Learning" (ICML '23) in an interdisciplinary poster session hosted across computer science, engineering and social/behavioral science departments. Winner of the poster presentation competition.

Computer Science Department Homecoming Poster Presenter *Fall 2022*

- Presented my research to the department alumni, faculty, dean, and current students, as one of the two presenters.

Computer Science Research Night Poster Presenter *Fall 2022*

- Introduced my lab and research to undergraduate and graduate students looking to understand and participate in the ongoing research works.

Cryptography Honors Seminar Speaker *Fall 2022*

- Discussed federated learning, differential privacy, applications, and why confidentiality of data is important in the world which is shifting towards data-rich artificial intelligence.

AI4ALL Summer Program Speaker *Summer 2021*

- Presented detailed pointers on how to read, understand, write research papers in AI and ML.
- Explained how to figure out unsolved yet solve-able problems, conduct research through creative solutions, evaluate results derived of the proposed approach, and discussed ethics and biases in AI.
- Encouraged 20+ undergraduate students from Boston University, Columbia University, and University of California Berkeley to pursue artificial intelligence research.