Ketan Suhaas Saichandran

ketansuhaas@gmail.com | Google Scholar | LinkedIn | GitHub | Website

EDUCATION

University of Oxford — Oxford, United Kingdom
DPhil in Statistics (ML for Drug Discovery) | **Expected Graduation: 10/30** — October 2026 – Present
Advisors: ***Prof. Garrett M. Morris, FRSC, Prof. Charlotte Deane, MBE, FRS***

Boston University — Boston, MA, USA
MS in Artificial Intelligence | **GPA: 3.96/4.00** — September 2023 | May 2025
Thesis: Active feature acquisition for efficient & transparent medical diagnoses
Advisor: Dr. Vijaya B. Kolachalama

Indian Institute of Technology Roorkee — Roorkee, India
B.Tech in Electrical Engineering | **CGPA: 8.65/10.00** — July 2019 | July 2023
Undergraduate Thesis: Deep Learning for Segmentation of Cardiac MRI (Dice **0.94**)

RESEARCH INTERESTS

Large language models, including post-training with reinforcement learning (GRPO, RLHF) and reasoning; autonomous AI agents, including computer use agents, multi-agent systems, and AI scientists; generative models, including continuous and discrete diffusion and flow-based models for images and molecules, and language models for chemistry; probabilistic machine learning, including learning under missing data; reinforcement learning; multimodal and vision-language models; and ML for science and health, including drug discovery, computational pathology and active feature acquisition.

RESEARCH APPOINTMENTS

Oxford Protein Informatics Group, University of Oxford — Oxford, United Kingdom
DPhil Student — October 2026 – Present
Advisors: Prof. Garrett M. Morris, FRSC, Prof. Charlotte Deane, MBE, FRS

- Develop Riemannian diffusion and flow-matching models for molecular docking, language models for chemistry, discrete diffusion models, generative models for de novo drug design, and methods for virtual screening of small molecules.

Kolachalama Lab, Boston University — Boston, MA, United States
Staff Scientist, Boston University School of Medicine — June 2025 | January 2026
PI: Dr. Vijaya B. Kolachalama

- Built a multi-agent AI system for automated scientific discovery in Alzheimer's research, testing 100 literature-derived hypotheses on NACC cohort data (*npj Dementia*).
- Built a graph-attention VLM on 40,000+ whole-slide neuropathology images (∼80% caption accuracy).
- Contributed to a multimodal LLM for dementia diagnosis (*Nature Communications*, under review).
- Contributed to contrastive representation learning for instance-adaptive feature acquisition (*TMLR*, under review) and mask-space distributionally robust optimisation for learning with missing features (*ICLR 2027*, under review).

Graduate Researcher — September 2023 | May 2025
PI: Dr. Vijaya B. Kolachalama

- Developed CMIRL, an information-theoretic RL method for active feature acquisition (AFA) in clinical diagnosis (master's thesis; 2–3% accuracy gains; *IEEE TAI*, under review), and contributed to explainability-driven feature ranking for AFA (*ICML 2025*; 1–10% gains across 9 clinical datasets).

Deepti Research Group, Boston University — Boston, MA, United States
Graduate Researcher — September 2024 | May 2025
PI: Dr. Deepti Ghadiyaram

- Developed SCoPE, dynamic coarse-to-fine prompt refinement for text-image alignment in diffusion models, improving on Stable Diffusion for 83% of benchmark samples (CVPR 2025 workshop, oral).

Banaji Implicit Social Cognition Lab, Harvard University — Cambridge, MA, United States
Research Assistant — May 2024 | May 2025
PI: Prof. Mahzarin Banaji

- Showed robust humanlike cognitive dissonance in LLMs across 900+ trials (effect sizes up to $d = 3.775$, $P < 0.001$); contributed to a *PNAS* paper and reply letter.

PUBLICATIONS & PREPRINTS

1. **Agentic AI for automated hypothesis testing in Alzheimer's disease and related dementias.**
   Saichandran, K. S., Elzokm, K., Guney, O. B., & Kolachalama, V. B. [**npj Dementia '26**]
2. **Active feature acquisition via explainability-driven ranking.**
   Guney, O. B., **Saichandran, K. S.**, Elzokm, K., Zhang, Z., & Kolachalama, V. B. [**ICML '25**]
3. **Progressive prompt detailing for improved alignment in text-to-image generative models.**
   Saichandran, K. S., Thomas, X., Kaushik, P., & Ghadiyaram, D. [**CVPR-W '25 (Oral)**]
4. **Domain-adapted language model using reinforcement learning for various dementias.**
   Kowshik, S. S., et al. [**Nature Communications, under review**]
5. **The Inner Monologue of Language Models: When Reasoning Traces Reveal More Than They Hide.**
   Singla, P., Singh, A., Garg, S., **Saichandran, K. S.**, & Garg, I. [**ACL Findings '26**]
   Mentored collaboration with the IIT Roorkee Vision Language Group
6. **Conditional mutual information-guided reinforcement learning for active feature acquisition.**
   Saichandran, K. S., Guney, O. B., Elzokm, K., & Kolachalama, V. B. [**IEEE TAI, under review**]
7. **Mask-space distributionally robust optimization for learning under stochastic missingness.**
   Xue, C., Guney, O. B., Acar, D. A. E., Elzokm, K., **Saichandran, K. S.**, & Kolachalama, V. B.
   [**ICLR '27, under review**]
8. **Contrastive representation learning for instance-adaptive active feature acquisition.**
   Guney, O. B., Elzokm, K., **Saichandran, K. S.**, & Kolachalama, V. B. [**TMLR, under review**]
9. **Kernels of selfhood: GPT-4o shows humanlike patterns of cognitive dissonance moderated by free choice.**
   Lehr, S. A., **Saichandran, K. S.**, Harmon-Jones, E., Vitali, N., & Banaji, M. R. [**PNAS '25**]
10. **Reply to Cummins et al.: GPT reveals cognitive dissonance that is both irrational and alarmingly humanlike.**
    Lehr, S. A., **Saichandran, K. S.**, Harmon-Jones, E., Vitali, N., & Banaji, M. R. [**PNAS '25**]
11. **A Comparative Analysis of U-Net-based models for Segmentation of Cardiac MRI.**
    Saichandran, K. S. [**arXiv '24**]

POSITIONS OF RESPONSIBILITY

- **Reviewer** at ICML, NeurIPS, and ICLR conferences.

TALKS & PRESENTATIONS

International Conference on Machine Learning (ICML) — Vancouver, Canada
Poster presentation — July 15th, 2025

AI for Content Creation Workshop, Computer Vision and Pattern Recognition (CVPR) — Nashville, TN
Oral & poster presentation — June 12th, 2025

3rd Workshop on Generative Models for Computer Vision, CVPR — Nashville, TN
Poster presentation — June 11th, 2025

Graduate School of Arts & Sciences, Boston University — Boston, MA
Master's Thesis Defense — May 17th, 2025

SELECTED ACHIEVEMENTS

- **Agentic AI Prize for Alzheimer's Disease Research 2025**: Top 10 finalist worldwide (ADDI).
- **YC Startup School 2026**: selected.
- **YC x DeepMind 2026**: selected.
- **ICPC Asia Regionals 2020, 2021**: represented IIT Roorkee in both years.
- **JEE Advanced 2019** (mathematics, physics, chemistry): All India Rank 1640 (FIITJEE Award).
- **JEE Main 2019** (mathematics, physics, chemistry): All India Rank 1390 among 1.2M+ candidates.
- **KVPY Scholar 2018** (mathematics, physics, chemistry): All India Rank 1237.
- **Regional Mathematical Olympiad (RMO) 2018**
- **Indian National Physics Olympiad (INPhO) 2019**: top 0.75% nationally.

TEACHING EXPERIENCE

MPLS Division, University of Oxford — Oxford, United Kingdom
Tutor (upcoming) — Hilary Term 2027
Advanced Topics in Statistical Machine Learning

Faculty of Computing & Data Sciences, Boston University — Boston, MA, United States
Teaching Assistant — January 2024 | April 2024

DS 320: Algorithms for Data Science

INDUSTRY EXPERIENCE

dCortex San Francisco Bay Area, CA, United States
Founding Research Scientist January 2026 | September 2026

- Invented a self-reflecting computer use agent for legacy airline operations software (1990s GUI), designing a code-generating, resolution-agnostic execution pipeline that replaced stochastic LLM decisions with deterministic, verifiable actions — a production feedback loop showed >90% alignment with expert crew controllers.
- Delivered an end-to-end AI automation platform serving multiple airline customers, processing hundreds of airline operations daily through hierarchical multi-agent pipelines with custom constraint solvers, optimizers, and execution schedulers.
- Raised VC funding from Bessemer Venture Partners (backed Anthropic, etc.) as the first founding research member alongside the founding team.

Clairyon CA, United States
AI Engineer May 2025 | December 2025

- Built and deployed predictive diagnostic models integrated with EHR systems via FHIR on AWS with 24/7 availability, working with UCSD Health through Clairyon (founding team from UCSD), scaling to serve 50k+ patients.
- Contributed to a prompt-based AI agents platform with FHIR MCP tools to access EHR data, enabling 20+ specialized automation agents for clinical reports, discharge summaries, quality measures for 10+ diseases/emergencies, and clinical documentation, achieving 3x speedup and 80% accuracy in readmission risk prediction based on EHR data.
- Built an AI-driven pre-visit triage backend that captures and condenses patient history through conversational interaction, reducing intake overhead and increasing patient throughput on a platform already serving 50,000+ patients.

NourishedRx United States
AI Engineer Intern May 2024 | August 2024

- Developed and deployed 3 AI-powered applications and LLM automation workflows, including *AskBetty* chatbot using AWS Bedrock and RAG for personalized health insights, and LLM agents for automated querying of BigQuery and FHIR data sources, serving 1,000+ users.
- Deployed scalable backend infrastructure with AWS Lambda REST APIs and cloud monitoring, ensuring 99.9% uptime and seamless integration across systems.
- Automated Amazon Connect call transcription and summarization via AWS Bedrock LLMs, processing 500+ calls monthly and streamlining clinical documentation and patient note generation, reducing documentation time by 50%.

Slice Bengaluru, Karnataka, India
Software Development Engineer Intern May 2022 | July 2022

- Designed and optimized RESTful APIs in Java Spring Boot for Juspay payment integration, developing 3+ critical endpoints for order creation, status retrieval, and payment authentication, handling 10M+ transactions monthly.
- Ensured high-throughput and low-latency performance, processing millions of transactions with sub-100ms response times and 99.9% uptime.
- Collaborated with front-end, DevOps, and security teams to enhance payment reliability, reducing transaction failures by 15% through improved error handling and retry logic, improving overall system reliability.

OPEN-SOURCE CONTRIBUTIONS

- AWS Strands Agents PR: Contributed to a major PR that introduces new functionalities to fetch and list the prompts from MCP servers.

TECHNICAL SKILLS

- **Programming Languages:** Python, C++, SQL, JavaScript, TypeScript
- **Tools/Software:** Anaconda, VS Code, Git, Docker, Terminal, WandB
- **AI Techniques:** LLM Post-training (GRPO, RLHF), Reinforcement Learning, Reasoning, Agentic AI, Computer Use Agents, Multi-agent Systems, Diffusion & Flow Models (Riemannian, Discrete), Probabilistic ML, VLMs/MLLMs, Active Feature Acquisition, QLoRA, LoRA, CLIP, RAG
- **ML/LLM Stack:** PyTorch, PyTorch Lightning, TensorFlow, Keras, Scikit-learn, OpenCV, Hugging Face, vLLM, LangChain, LlamaIndex, Pinecone, MCP, CUDA, Triton
- **Cloud & Full Stack:** ReactJS, NodeJS, Spring Boot, Gradio, Streamlit, AWS Bedrock, AWS Lambda, GCP Vertex AI, BigQuery, FHIR, Docker, Kubernetes