

Amit Rand

(857)-998-9070 | amit.rand@ucla.edu | github.com/randamit123 | amitrاند.com | Los Angeles, CA

EDUCATION

University of California, Los Angeles June 2026
B.S. in Computer Science and Applied Mathematics *Los Angeles, CA*
Awards: Samuelli Research Scholar, Research Scholar Spotlight, Research Travel Sponsorship, Dean's Honors
Coursework: Data Structures & Algorithms, Linear Algebra, Statistics & Probability, Real & Numerical Analysis, Computer Vision, Scientific Machine Learning*, Reinforcement Learning*, Deep Learning*, Generative Models* (Graduate Coursework*)

PUBLICATIONS

Physics-Informed K-Space Diffusion for Accelerated MRI Reconstruction 2026
• **Amit Rand**, et al. *IEEE Transactions on Medical Imaging* [In preperation.]

Diffusion-based k-space inpainting for improved 5D free-running CMR reconstruction 2026
• Thomas Coudert, **Amit Rand**, et al. *ISMRM-ISMRT Annual Meeting and Exhibition*, Cape Town, South Africa.

Beyond Conventional Transformers: A Medical X-ray Attention Block for Improved Diagnosis 2025
• **Amit Rand***, Hadi Ibrahim*. *NeurIPS 2025 Imageomics Workshop*, San Diego, CA.

TECHNICAL SKILLS

Deep Learning: PyTorch, Generative Models, Complex-Valued Networks, Deep Networks
Scientific Computing: NumPy, Pandas, Matplotlib, Seaborn, Pydantic, VTK/ITK/VMTK (Medical/Geometric Data)
Robotics & RL: SB3, OpenAI Gym, OpenCV, Reinforcement Learning, Imitation Learning, Motion Planning, Manipulation
Systems & Languages: Python, C++, Java, Bash, SQL, AWS (S3, EC2, Lambda, etc.), Docker, Git, Weights & Biases

RESEARCH EXPERIENCE

Undergraduate Researcher July 2024 – Present
Cardiovascular Imaging Research Lab, UCLA *Los Angeles, CA*
• Developing measurement-conditioned diffusion and flow-based generative models for k-space inpainting on multi-coil cardiovascular MRI, utilizing 3D spatiotemporal blocks and specialized CUDA projection layers.
• Optimized complex-valued diffusion architectures with U-Net/Transformer backbones to ensure hard data consistency in 2D MRI reconstruction, achieving state-of-the-art interpolation fidelity and accuracy for globally predictable k-space data.
• Designed physics-informed auxiliary losses (structured low-rank, inter-coil consistency) and symmetry-based input augmentations in the Fourier domain, improving reconstruction fidelity (+3 dB PSNR, +0.10 SSIM, 50% NMSE reduction) and demonstrating the impact of inductive biases in generative models.
• First author on physics-informed deep learning research in preperation for IEEE Transactions on Medical Imaging and accepted to ISMRM, supported by the Samuelli Research Scholarship and recognized Research Scholar Spotlight.

Research Scientist Intern April 2025 – June 2025
AI/ML Research Accelerator (Graph AI Group), Leidos *Santa Clara, CA*
• Investigated GNN embedding techniques to enhance feature clustering and decision-making in Reinforcement Learning environments, emulating dynamic graphs to optimize agent performance in non-stationary spaces.
• Integrated RAG with supervised fine-tuning (SFT) to build enterprise-grade generative reasoning solutions for internal corporate upskilling, bridging the gap between structured graph data and stochastic RL policies.

PROFESSIONAL EXPERIENCE

Software-ML Development Engineer Intern June 2025 – Sep 2025
Amazon, Project Kuiper *Redmond, WA*
• Delivered PyTorch models projected to reduce antenna calibration downtime by 20–30% across thousands of satellites.
• Developed physics-aware transformers, reducing calibration time from 3 hours to mere inference via AWS SageMaker.
• Architected automated AWS retraining pipelines to replace expensive thermal testing with high-fidelity surrogate simulation.

GenAI Technical Advisor Intern Feb 2025 – May 2025
Scale AI *Los Angeles, CA*
• Refined Chain-of-Thought (CoT) and fine-tuning methodologies for LLMs to optimize long-horizon reasoning performance.
• Engineered terabyte-scale preprocessing pipelines, lifting model throughput by 70% via optimized feature engineering.

Software Development Engineer Intern Aug 2024 – Oct 2024
Q.ai (Acquired by Apple for 2B, Backed by Google) *Cupertino, CA*
• Deployed deep learning models and statistical EDA in Python to drive a 20% accuracy boost for proprietary systems.
• Built real-time Dash/Plotly dashboards with SQLAlchemy ETL to monitor 10M+ files across 5 sites and 50 operators.
• Optimized backend data ingestion and UX for core collection teams, ensuring high-quality integrity for model training.