



DWAIT BHATT

 dwaitbhatt.com

 [linkedin.com/in/dwait-bhatt](https://www.linkedin.com/in/dwait-bhatt)

 dhbhatt@ucsd.edu

 (858) 319-6716

SUMMARY

ML PhD student working in robotics at UCSD with 3 years of AI R&D experience at Samsung Research and ISRO. Specifically interested in reinforcement learning and value alignment to develop helpful and collaborative robots.

SKILLS

Relevant Coursework: ML for Robotics, Sensing & Estimation in Robotics, Advanced Computer Vision, Probabilistic Reasoning, Search & Optimization, Information Theory, Deep Generative Models, Linear Systems, Recommender Systems

Programming Languages: Python, Matlab, C++

Libraries: PyTorch, Keras, Numpy, Pandas, Robosuite, OpenCV, Open3D, xArm API

EDUCATION

University of California, San Diego

Ph.D. in Electrical & Computer Engineering - Machine Learning & Data Science

Expected Graduation: June 2029

San Diego, California

University of California, San Diego

M.S. in Electrical & Computer Engineering - Intelligent Systems, Robotics & Control

June 2024

San Diego, California

Birla Institute of Technology and Science, Pilani (BITS Pilani)

B.E. (Hons.) in Electrical & Electronics Engineering

June 2019

Pilani, Rajasthan

EXPERIENCE

Existential Robotics Laboratory

Graduate Researcher

January 2023 - Present

San Diego, California

- Enabled manipulation skill transfer across robot embodiments via latent policies. Co-authored a paper on this work, which is in submission at RA-L.
- Currently working on sim-to-real transfer using progressive neural networks and RL with online and offline data.
- Independently designed and implemented a VR teleoperation system for real-time control of a robotic arm.
- Created a sim-to-real pipeline from scratch for executing robot policies learned in simulation on a physical robot arm. Successfully executed tasks like stacking and pick-and-place on an xArm6 robot with sub-5mm error.

Samsung R&D Institute India

Senior Machine Learning Engineer, On-Device AI

July 2019 - August 2022

Bangalore, Karnataka

- Spearheaded the development of a patented application launch time prediction model, from a python POC to an on-device implementation in C++, which was used to boost CPU for an optimal duration to make application launches 6% faster while extending battery life by 2 hours.
- Automated app launch video analysis using a novel algorithm based on structural similarity between consecutive video frames, saving over 40 man-days annually and significantly enhancing operational efficiency.
- Authored a research paper on this work and presented it at IEEE IJCNN 2022 held in Padua, Italy.
- Guided a team of four undergraduate students through a sequence prediction project as their mentor for Samsung's industry-academia collaboration program, PRISM.

Machine Learning Engineer, On-Device AI

- Designed a reinforcement learning agent that trained on-device to control an Android smartphone's kernel tunables, reducing the decoding time of encrypted models by 10% on Samsung's flagship phones.

Indian Space Research Organisation (ISRO)

Research Trainee

September 2021 - January 2022

Ahmedabad, Gujarat (Remote)

- Implemented a deep reinforcement learning network in PyTorch trained using policy gradients to drop non-essential image patches from CIFAR10 data while maintaining classification accuracy of 90%. This work is being extended to satellite and drone images for efficient poverty estimation from remote sensing imagery.

PROJECTS

Curiosity-driven Robot Reinforcement Learning | *PyTorch*

January 2023 – March 2023

- Introduced curiosity-based intrinsic motivation in a PPO robot policy via Random Network Distillation. Found that naively adding such exploration hurt performance for robotic tasks in contrast to RND's success in playing video games.

Multimodal Material Estimation | *PyTorch*

April 2024 – June 2024

- Predicted object material from images and audio of an object being struck by a drumstick by aligning audio and image embeddings to llama-3 token embeddings using cross-attention. The multimodal approach achieved 20% higher accuracy than using only images, proving that audio cues are crucial for object material estimation.

Musical Accompaniment Generation using Language Modeling | *PyTorch* | UCSD

April 2023 – June 2023

- Designed a novel architecture for accompaniment generation conditioned on a vocal track using language modeling.
- Achieved semantically meaningful generation (FAD 4.07) with a budget of under \$50.

Surface Reconstruction using Locally-Aware NeRFs | *PyTorch*

April 2023 – June 2023

- Tackled the issue of missing local information in the NeRF architecture by introducing features from a Convolutional Occupancy Network. This enabled smoother surface reconstruction of complex geometries and large-scale scenes.