Harsh Khatri

Boston, MA | harshkhatri242@gmail.com | linkedin.com/in/harsh242 | github.com/harsh242bu

Education

Boston University (BU) Sep 2022 – May 2024
Master of Science in Computer Science 4.0/4.0
Masters thesis under Prof. Bryan Plummer

Indian Institute of Technology Bombay (IITB) Jul 2015 – Jun 2019
Bachelor of Technology in Metallurgical Engineering and Materials Science

Technical Skills and Courses

Courses: Robot Learning, Computer Vision, Machine Learning, OOP (Java), Graduate Algorithms.
Languages: Python, Java, C++, JavaScript, ReactJs, SQL.
Tools: OpenCV, PyTorch, PyTorch Lightning, Scikit-Learn, Numpy, Scipy, Vertex AI, OpenAI Gym.
DevOps: AWS EMR, AWS S3, AWS EC2, GCP, Docker, Git, Kubernetes, GKE, CI/CD.
Data Stores: MySQL, MongoDB.

Work Experience

Machine Learning Engineer July 2024 – Present
Logile - The Logic of Retail Boston, MA

- Developed transformer based forecasting algorithm outperforming legacy model, improving forecast performace by 8% while reducing codebase by 83%.
- Leveraged Nvidia Rapids to design and deploy a high-performance scaling solution, resulting in a 20x efficiency boost and a 60% reduction in processing time of legacy forecasting algorithm.
- Enhanced the performance of the legacy algorithm by vectorizing data processing and feature engineering steps.
- Optimized and scaled the legacy forecasting pipeline by architecting and deploying a solution on AWS EMR.

Graduate Research Assistant, IVC Lab Sept 2023 – May 2024
Advised by Prof. Bryan Plummer Boston, MA

- Conducted research for masters thesis on Open-set 3D Object Detection utilizing PointBERT transformer.
- Developed a novel Coarse-to-Fine Grained recognition method discerning confusing categories in the dataset.
- Designed a custom loss function incorporating offline hard negative mining technique for text and image modalities, resulting in 2% improved performance over state-of-the-art method.

Computer Vision Engineer Dec 2023 – Jan 2024
Wicket: Facial Authentication Boston. MA

- Devised a CNN architecture using ResNet to classify facial features achieving 99.5% accuracy on synthetic dataset and 82.3% accuracy on real dataset.
- Developed an architecture employing Stable Diffusion to generate synthetic dataset for facial feature detection.
- Utilized negative prompts to improve the quality of the synthetic data by 33% (Prompt engineering).
- Implemented a filtering mechanism to discern hallucinations and false negatives by extracting features from pretrained CLIP, DINO, and ConvNext models and training a ensemble model to predict human perception decision boundary.

Software Developer Jul 2019 – Aug 2022
Reflexis Systems (Acquired by Zebra Technologies) Pune, India

- Spearheaded a team of 5 developers; identified, prioritized and delegated skill-appropriate tasks using Scrum.
- Developed, maintained, and tested 400,000+ lines of code of Enterprise Performance Product, capable of multidimensional analysis across different verticals using Java, ReactJs and Python.
- Implemented an efficient test automation pipeline utilizing Katalon (Selenium based), Cypress, Jenkins, JUnit and GCP, leading to over 1000 hours of saved developer time previously spent on manual testing.
- Initiated a project to restructure the interactive data visualization dashboard, leading to a performance improvement of 50% using modular code, maximizing API asynchronization, and utilizing hash maps.
- Recognized with an "Exceeds Expectations" rating for exceptional leadership demonstrated in spearheading multiple projects, resulting in a 27% increase in team productivity.
- Initiated development of cross tech knowledge sharing platform and prepared training materials for company.

Projects

Sim-to-Real in 3D Object Detection Sep 2023 – Dec 2023

Advised by Prof. Eshed Ohn-Bar Boston, MA

- Achieved 9% performance improvement over baseline detection models by implementing various methods mentioned below to close the Sim2Real gap in 3D Object Detection.
- Conducted evaluation of cutting-edge 3D object detection methodologies, including PointNet, CurveNet, and GDANet, on real-world OmniObject3D point cloud dataset, effectively demonstrating the Sim2Real gap.
- Proposed a novel random sampling approach for synthetic point cloud data, resulting in a notable 4% enhancement in performance when evaluated on real dataset.
- Enhanced the performance of CurveNet point cloud encoder by integrating language embeddings derived from CLIP, resulting in a significant 3% improvement in accuracy.
- Utilized LLM to generate descriptive captions for classes and augment caption embeddings, resulting in a notable 2% enhancement in model performance.

Autonomous Driving (OpenAI Gym) Sep 2023 – Dec 2023

Advised by Prof. Eshed Ohn-Bar Boston, MA

- Designed an autonomous driving model using: Imitation Learning, Modular Pipeline and Reinforcement Learning.
- Imitation Learning: Trained a custom CNN for the gym car racing achieving 200+ reward, setting the baseline. Optimized CNN architecture by implementing DAGGER algorithm, achieving a 20% increase in performance.
- Modular Pipeline: Engineered a modular self driving pipeline featuring modules—Lane Detection, Waypoint Prediction, Lateral and Longitudinal Control systems (Stanley Controller and PID Controller)—yielding 50% improvement in performance over Imitation Learning.
- Reinforcement Learning: Implemented the PPO (Proximal Policy Optimization) Reinforcement Learning algorithm to achieve a substantial 37% performance enhancement over modular approach.
- Achieved best score in a class of 23 students for autonomous driving with Modular Pipeline.

Phrase Grounding Jan 2023 – May 2023

Advised by Prof. Bryan Plummer Boston, MA

- Compiled and analyzed existing research on phrase grounding, problem of associating natural language phrases or descriptions with corresponding objects in an image.
- Identified and experimented with unique model components such as image and text encoder, multi-modal interaction, loss function, sampling strategies, evaluation metric and regularization method.
- Conducted extensive ablation study by developing and evaluating sampling techniques like Hard Negative, Weighted Negative and Semi-Hard Negative mining.

Hand Gesture Recognition Sep 2022 – Dec 2022

Advised by Prof. Bryan Plummer Boston, MA

- Improved performance by 7% over state-of-the-art gesture recognition by integrating ResNeXt with hand landmark detector using LSTM based integration architecture.
- Conducted an ablation study by training and evaluating vision model architectures like ResNeXt, MobileNet and ShuffleNet for gesture recognition task on Jester dataset (Qualcomm).
- Integrated ResNeXt 101 with Google MediaPipe Hand Landmarker which detects 21 unique landmarks of the hand.

Multi-omics study in COVID-19 patients Dec 2020 – Aug 2021

Advised by Prof. Sanjeeva Srivastava Mumbai, India

- Co-author of original research articles titled *Multi-omics data analysis to assess severity in COVID-19 patients* and *A Multi-omics Longitudinal Study of COVID-19 Patients*.
- Conducted exploratory analysis of mass-spectrometry data and identified 10 significant proteins and 6 metabolites responsible for the severity of COVID-19.
- Generated critical insights by implementing classification models (SVM) and statistical tools (t-test) in Python.

Additional Experiences

- Served as a Vice president of WACM (BU chapter) promoting women in tech, helping people connect, discover opportunities, and build community to advance contributions to computer science.
- Organized computer science department events, promoting community engagement amongst master's students.
- Volunteered to teach a class of 20+ underprivileged students at ASHA NGO, Mumbai, India.
- Completed 108 Surya Namaskar challenge at International Day of Yoga.