

Ayush Mangal

DATA AND APPLIED SCIENTIST · MICROSOFT, HYDERABAD

☎ (+91) 9425981956 | ✉ ayushtues@gmail.com | 🏠 ayushtues.github.io | 📄 ayushtues | 📧 ayushtues | 📧 ayushtues

Education

Indian Institute of Technology Roorkee

Roorkee, India

B.TECH. IN COMPUTER SCIENCE AND ENGINEERING

2018 - 2022

- CGPA : 9.51/10

Publications

Modeling Financial Uncertainty with Multivariate Temporal Entropy-based Curriculums

UAI 2021

RAMIT SAWHNEY, ARNAV WADHWA, **AYUSH MANGAL**, VIVEK MITTAL, SHIVAM AGARWAL, RAJIV RATN SHAH

May 2021

- Developed a novel multi-modal curriculum learning approach: FinCLASS, which evaluates stock affecting signals via entropy-based heuristics and measures their linguistic and price-based complexities in a time-aware, hierarchical fashion to improve training and performance
- Achieved SOTA results on financial tasks like volatility regression, stock price movement prediction on US and Chinese market datasets.
- [Paper] [Presentation] [Video] [Poster] [Github]

Track, Check, Repeat: An EM Approach to Unsupervised Tracking

CVPR 2021

ADAM HARLEY, YIMING ZUO, JING WEN, **AYUSH MANGAL**, SHUBHANKAR POTDAR, RITWICK CHAUDHRY, KATERINA FRAGKIADAKI

Feb 2021

- Developed an unsupervised method for 3D object tracking in RGB-D videos using an expectation maximization framework.
- The basic idea was to use some traditional CV methods to get some pseudo-labels when we were fairly sure about them being right, and using them to train DL based tracking agents, and repeating this process in an EM fashion.
- [Paper] [Video] [Project Page] [Notion]

Multimodal Multi-Task Financial Risk Forecasting

ACM Multimedia 2020 (Oral)

RAMIT SAWHNEY, PUNEET MATHUR, **AYUSH MANGAL**, PIYUSH KHANNA, RAJIV RATN SHAH, ROGER ZIMMERMANN

October 2020

- Developed a multi-task solution using multimodal signals from domain specialized textual features and audio attentive alignment for predicting financial risk and price modelling.
- Tailored a deep multimodal text-audio attention model and optimized volatility, and price movement prediction in a multi-task ensemble formulation.
- [Paper] [Presentation]

DEAP Cache: Deep Eviction Admission and Prefetching for Cache

Submitted to AAAI 2021

AYUSH MANGAL*, JITESH JAIN*, KEERAT KAUR GULLANI*, OMKAR BHALERO*

September 2020

- Led a team consisting of only undergraduate researchers and proposed an end to end deep learning pipeline for automatically learning prefetching, admission and eviction policies for optimizing cache latency.
- Learnt an optimal policy distribution between two orthogonal strategies based on frequency and recency in an online multi-task fashion.
- [Preprint] [Github]

Visual Relationship Detection using Scene Graphs: A Survey

Submitted to ACM CSUR

ANIKET AGARWAL, **AYUSH MANGAL***, VIPUL KUMAR

May 2020

- A student only team research effort in which we presented a detailed survey on the various techniques for scene graph generation, their efficacy to represent visual relationships and how they have been used to solve various downstream tasks.
- Introduced a novel taxonomy to classify existing approaches, with detailed benchmarking and musings about possible future research areas.
- [Preprint]

Experience

Microsoft

Hyderabad, India

DATA AND APPLIED SCIENTIST

June 2022 - Present

- Working in the Bing Travel team, on various tasks related to Computer Vision, creating image ingestion pipelines, deep/shallow models for image quality checks etc.

Rephrase.ai

Bangalore, India

REMOTE RESEARCH INTERN

Dec 2021 - January 2022

- A deep-tech startup that creates personalized videos for marketing using DL [Website]
- Worked on developing a new version of their DL pipeline, conducting various experiments, and the work resulted in new insights about how to potentially enhance their existing pipeline.

Robot Vision and Learning Lab, University of Toronto

Toronto, Canada

REMOTE RESEARCH INTERN | PROF. FLORIAN SHKURTI

August 2021 - May 22

- Worked on my undergraduate thesis, on developing a new hierarchical reinforcement learning method building on past work using latent skills and world models.

STCI, Microsoft IDC

Hyderabad, India

DATA AND APPLIED SCIENTIST INTERN

June 2021 - July 2021

- Developed a novel end to end ML pipeline for classifying destinations into travel-related categories for Bing Travel.
- Identified what data sources to use from a vast array of in-house data sources, how to leverage each of them, how to combine them to solve the problem best, what model to use and how to evaluate it to iterate towards the final solution.

Machine Learning Department, Carnegie Mellon University

Pittsburgh, USA

REMOTE RESEARCH INTERN | PROF. KATERINA FRAGKIADAKI

June 2020 - Feb 2021

- Helped the researchers in the lab in various ways while developing a method for Unsupervised 3D Multi-Object Tracking.
- Work resulted in a research publication in CVPR 2021, the top conference for Computer Vision.
- Actively participated in various paper reading groups hosted by the lab.

MIDAS Lab, IITD

Delhi, India

REMOTE RESEARCH INTERN | PROF. RAJIV RATN SHAH

March 2020 - May 2020

- Developed a novel multitask multimodal framework for financial risk forecasting from company earning calls.
- Developed a novel attention-based deep learning model to extract multimodal cues from text transcripts, audio features and historical time-series financial data for multitasking stock price prediction and stock volatility regression.
- Developed a novel neural quantitative trading strategy based on the stock price predictions with a volatility-based risk threshold.

Video Analytics Lab, IISc Bangalore

Bangalore, India

WINTER RESEARCH INTERN | PROF. R. VENKATESH BABU

Dec 2019 - Jan 2020

- Worked on developing end to end methods for HDR Deep Deghosting using segmentation models like UNet, ResNet etc.
- Worked on applying Domain Adaptation for processing low-light images based on the ideas of Unsupervised Adversarial Discriminative Domain Adaptation and CycleGAN.

OpenSource

Consistency Models

Diffusers, Huggingface

OPENSOURCE CONTRIBUTION

March 2023 - Present

- Helped in integrating a new family of generative models similar to diffusion models, but with single step inference called Consistency Models into Huggingface Diffusers. [\[Documentation\]](#)[\[Code\]](#)

Kandinsky 2.1

Diffusers, Huggingface

OPENSOURCE CONTRIBUTION

March 2023 - Present

- Helped in integrating a multilingual text-to-image diffusion model, inheriting best practices from DALL-E and Latent Diffusion - Kandinsky 2.1, into Huggingface Diffusers. [\[Documentation\]](#)[\[Code\]](#)

Paper Implementations

Github

OPENSOURCE CONTRIBUTION

Oct 2019 - Present

- Implemented various research papers in Pytorch and JAX
- DDPM [\[Pytorch\]](#) [\[JAX\]](#), ADDA [\[Pytorch\]](#), VAE [\[Pytorch\]](#)

Papers We Read

IIT Roorkee

OPENSOURCE CONTRIBUTION

Oct 2019 - Present

- An open-source repository consisting of insightful summaries of various Tier-1 Deep Learning conferences.
- Crossed 200+ stars on Github. [\[Github Link\]](#)

Extracurricular Activity

Vision And Language Group, ACM IITR

IIT Roorkee

CO-PRESIDENT

Sept. 2019 - May 2022

- Co-president of a student-run group that aims to foster a research-centric Deep Learning Community at IIT Roorkee. [\[Group Website\]](#)

RR with Deku

Youtuber

DL RESEARCH YOUTUBE CHANNEL

Oct. 2020 - Present

- Started a Youtube channel explaining various Deep Learning papers in a detailed manner which even an undergrad could understand [\[Link\]](#)

Blockchain Society IIT Roorkee

IIT Roorkee

FOUNDER

October 2021 - May 2022

- Founded an open group to develop the blockchain, crypto and web 3.0 culture on campus.

ACM IIT Roorkee Student Chapter

IIT Roorkee

CHAIR

March 2021 - May 2022

- Led the overall planning of the group, to foster a generic research culture related to Computer Science at the campus.