

Navyansh Mahla

📍 Nationality: Indian | 📧 navyanshmahla17@gmail.com | 🌐 navyanshmahla.github.io | 🎂 DOB: 17/05/2004
📍 Address: Hostel 2, Students' Residential Area, IIT-Powai, Powai (400076), Mumbai, Maharashtra

RESEARCH INTERESTS

Generative Models | Privacy-Preserving Machine Learning | Foundation Models | Federated Optimisation | Feature Subset Selection problems in ML

EDUCATION

Indian Institute of Technology, Bombay

Pursuing Minor degree in Machine Learning and Data Science
Nov 2021 – May 2025 (Expected)

Mumbai, Maharashtra

Major in Civil Engineering

Sanskar International School

High School Diploma

Sikar, Rajasthan

July 2021

PUBLICATIONS

- **FESS Loss: Feature-Enhanced Spatial Segmentation Loss for Optimizing Medical Image Analysis**
Charulkumar Chodvadiya*, Navyansh Mahla*, Kinshuk Gaurav Singh, Kshitij Sharad Jadhav
International Symposium on Biomedical Imaging (ISBI 2024)

AWARDS

- **URA01 Undergraduate Research Award:** The Undergraduate Research Award (URA 01) recognizes the exceptional research achievements of undergraduate students at an early stage, highlighting their demonstrable vigor and commitment to scholarly pursuits. [See Certificate](#)

SCHOLASTIC ACHIEVEMENTS

- In the top **1.5%** people among 0.15 million candidates in **JEE Advanced 2021**
- Qualified the Indian Olympiad Qualifier in Physics (**IOQP**) organised by IAPT
- Cleared the Indian Olympiad Qualifier in Astronomy (**IOQA**) organised by IAPT
- Scored 403/450 in the prestigious BITSAT examination organized by Birla Institute of Technology and Science.
- Procured the title of **3rd Runner up** at the International Mental Arithmetic Competition among 5500+ students from more than **30** countries after several rounds testing mental rigour, speed, and memory
- Awarded **MERIT** twice in state-level Mental Arithmetic Competition among **4k+** students

WORK EXPERIENCE

Fractal.ai

AI Research Scientist | Internship

Mumbai, Maharashtra

May 2024 – Present

Working on AI Alignment, security and efficient training for building State-of-the-art Large Language Models (LLMs) for conversational tasks.

Amazon India

AI Researcher | Research Project Collaboration

Mumbai, Maharashtra

Dec 2023 – Present

- Conducting collaborative research with Amazon India under the supervision of Prof. Ganesh Ramakrishnan and Prof. Kshitij Jadhav for data-privacy preservation in medical domain
- Working on problem statements revolving around federated fine-tuning and inference of large language and multi-modal models
- Developing advanced multi-agent client-server based pipelines for decentralised fine-tuning and inference of 100B+ parameter LLMs and LMMs

Koita Centre for Digital Health

Deep Learning Researcher

Mumbai, Maharashtra

July 2023 – Present

- Collaborating research under the supervision of Prof. Ganesh Ramakrishnan and Prof. Kshitij Jadhav in the domain of image processing, language and vision models, PPML, etc. in the problem setting of medical domain
- **Image Processing/Contrastive Learning:** Developed and proposed a novel Feature-Enhanced Spatial Segmentation Loss, integrating contrastive learning with spatial accuracy to enhance precision and feature representation in medical image segmentation, as detailed in a paper submitted to ISBI 2024. Submitted another paper to ICPR 2024 working further on alternative contrastive learning strategies
- **Medical Segment Anything Model:** Implemented a custom version of SAM model for few shot fine-tuning on medical datasets for semantic segmentation and segment classification which is being used for datasets from various hospitals
- **BharatGPT:** Working on foundation multi-modal models for the vision subsystem of BharatGPT. Leading the data team for building high-performance automated data annotation pipelines for labeling Indic-image datasets for the use in the vision subsystem
- **Amazon India:** Collaborating with Amazon India on privacy-preserving strategies for fine-tuning large vision and language models in distributed environments
- **Dataset Unlearning:** Collaborating with various medical institutions for developing novel machine unlearning solutions in medical field
- **Few Shot Memory Efficient Image Segmentation:** Devised a novel strategy for few-shot image segmentation using the Segment Anything Model (SAM). Implemented a novel optimisation algorithm which was able to reduce training memory overhead by **more than 70%** compared to traditional SAM, SAM+LoRA (fine tuning), TransUNet, U-Net, DAE-Former, etc. Submitted the paper to **British Machine Vision Conference (BMVC) 2024**.

Marsh McLennan Global Services

RnD Engineer | Internship

Mumbai, Maharashtra

Dec 2022 – Jan 2023

- Spearheaded **research** on smart contract vulnerability tools, deeply analyzing their capabilities and identifying potential weaknesses
- Developed efficient **APIs** that seamlessly integrated with these tools, enabling comprehensive smart contract audits and enhancing team capabilities.
- Led the development of a dedicated smart contract auditing platform with a microservice architecture, ensuring exceptional speed, efficiency, security, and scalability. Enabled evaluation of smart contracts for vulnerabilities, assessing threats, and providing valuable insights for workarounds and mitigations.

Angle.AI

Computer Vision Engineer | Internship

Mumbai, Maharashtra

Sept 2022 – Oct 2022

- Undertook the task of building robust computer vision models to detect defective bearings by identifying missing rivets using **OpenCV** and **CNNs**
- Meticulously developed and fine-tuned the performance code to meet industry standards for efficiency
- Optimised the algorithm using various methods like convex-hull, image-masking and pixel-filtering approach
- The deployed model was successfully integrated into the automated systems of a live manufacturing plant in Mumbai, enhancing their quality control processes.

HACKATHONS & COMPETITIONS

- Placed **180th globally** in the **Amazon ML Challenge 2023**
- **Institute Rank 14** in the data analysis competition organised by **Tower Research Capital**
- **All India Rank 22** and **Institute Rank 1st** for IIT-Bombay in the India Terminal 2023 programming contest organised by **Citadel** and **Citadel Securities**. Also interviewed for **Citadel** based on the performance in the competition
- Awarded the grand prize of **\$500** for building an end-to-end application on **Polygon** blockchain in the **EthIndia** hackathon

8.1 Deep Learning

Contrastive Learning based Medical Image Segmentation | KCDH, IITB (Aug '23 - Nov '23)

Guide: Prof. Kshitij Jadhav

- Engaged in the development of two innovative approaches for medical image segmentation, each featuring novel and groundbreaking loss functions
- Loss functions demonstrate efficiency across a diverse range of medical datasets, encompassing brain tumor data as well as datasets related to various organs
- Submitted two papers for our novel loss functions to **IEEE ISBI** (International Symposium on Biomedical Imaging) 2024 with **first** and **second** co-authorships respectively

Federated Data Augmentation Using Generative Models | KCDH, IITB (Jul '23 - Aug'23)

Guide: Prof. Kshitij Jadhav & Prof. Ganesh Ramakrishnan

- Researched on various different GAN architectures involving CGAN and WGAN to do medical data augmentation under federated settings for multi-client setup
- Utilising federated learning techniques like co-distillation to ensure data privacy among multiple clients

DeepCarlsen | [GitHub](#) | (Feb' 23-Apr' 23)

Independent Research Project

- Developed an end-to-end learning method for chess using **deep belief networks**, achieving a grandmaster-level playing performance.
- Successfully trained the evaluation function of the chess program without any prior knowledge of the game's rules or manual feature extraction.
- Utilized **unsupervised pretraining** and supervised training on large datasets of chess positions, allowing the deep neural network to learn high-level features and make favorable position comparisons.
- Implemented a **deep autoencoder** (Pos2Vec) in **PyTorch** as a nonlinear feature extractor and combined it with fully connected layers to create the DeepChess structure, resulting in a state-of-the-art chess program comparable to manual feature selection methods.
- Optimized model training by implementing wrapper functions using the **Numba** library (Just-In-Time compilation), resulting in significantly improved performance and faster training times for the chess program.

DeepComp | [GitHub](#) | (Mar' 23-May' 23)

Independent Research Project

- Developed a hybrid and robust compression framework called DeepComp that combines an **attention-based autoencoder** with traditional Windows **WinRAR** archiver for efficient compression of numerical and image data formats.
- Proposed an **attention** layer within the autoencoder architecture using **Tensorflow** to minimize reconstruction error during decompression and maintain **spatial consistency** of data points, resulting in improved compression performance.
- Evaluated DeepComp using atmospheric and oceanic data from the National Centers for Environmental Prediction (NCEP), demonstrating its robustness in compressing both numerical and image data types.
- Outperformed traditional Windows archivers by an average of 69% and multilayer autoencoders by 48% in terms of compression ratio, highlighting the superior performance of DeepComp in achieving high compression ratios with acceptable loss.
- Successfully reduced reconstruction loss during data decompression by introducing attention layers, enhancing DeepComp's reconstruction performance compared to multilayer autoencoders.

ImgComp | [GitHub](#) | (Mar' 23-May' 23)

Independent Research Project

- Developed a novel approach for image compression at low bitrates using **PyTorch** by incorporating Convolutional Neural Networks (CNNs) and attention mechanisms into the post-processing stage of traditional image compression decoders.
- Proposed an attention-based CNN architecture consisting of 30 residual blocks with channel attention (CA) and spatial attention (SA) mechanisms to enhance specific regions of the reconstructed image and improve overall compression quality.

- Trained the attention-based post-processing module using Mean Absolute Error (MAE) and Multi-Scale Structural Similarity Index (MS-SSIM) loss functions, aiming to generate visually appealing images while maintaining compression efficiency.
- Achieved remarkable performance with an average Peak Signal-to-Noise Ratio (PSNR) of 32.10 at a bit-rate of 0.15, surpassing other competing methods and demonstrating the effectiveness of the attention-based CNN in improving image compression quality.
- Developed a comprehensive image compression framework based on the Versatile Video Coding (VVC) structure, integrating the proposed post-processing method to enhance coding performance. The framework successfully combines traditional codec techniques with the power of CNNs to achieve impressive compression results.

SRGAN | [GitHub](#) | (Nov' 22-Dec' 22)

Self Project

- Developed and implemented a deep learning model based on the SRGAN architecture to enhance the resolution of low-resolution images, achieving photo-realistic results.
- Utilized advanced techniques such as perceptual loss and adversarial training to improve the quality of the super-resolved images and ensure they closely match the original high-resolution images.
- Conducted rigorous experimentation and optimization to fine-tune the model's hyperparameters and achieve high levels of accuracy and visual fidelity in the super-resolved images.

RAP: Road Accident Prevention | [GitHub](#) | (July' 22-Aug' 22)

Seasons of Code

- Developed a Deep Learning based system to solve the problem of road accidents by identifying the driver drowsiness.
- Used the haar cascade classifier to detect faces and used CNN classifier to predict the status of eyes from live feed
- Implemented the multi-layer sequential model using Keras, Tensorflow to train the data-set of drowsy eyes
- Coded the Convolutional Neural Network in Keras to train the multi-class classification model of drowsy eyes
- Achieved an accuracy level of over 90% for the model's performance over the validation data-set

8.2 Networks, Game Development and Systems

Secure Distributed Cloud Storage | [GitHub](#) | (May' 23-Jul' 23)

Seasons of Code

- Developed a secure and robust distributed storage system for cloud environments, addressing concerns related to data confidentiality, redundancy, and reliability.
- Implemented a decentralized encoding scheme to ensure data redundancy while optimizing storage space, allowing successful data retrieval even in the event of server failures.
- Employed suitable cryptographic mechanisms to maintain data confidentiality, offloading key management responsibilities to cloud key servers to prevent encryption key leakage even in the event of key servers compromise or failure.
- Implemented a proxy re-encryption scheme to enable secure data sharing between users, allowing cloud servers to re-encrypt data on behalf of users without accessing or learning any information about the underlying data.
- Developed the system with a focus on minimizing computational load on users and reducing communication overhead between users and the cloud, ensuring efficient and practical usage of the distributed storage system in real-world scenarios.

CodeWars V3 | [GitHub](#) | (Nov' 22-Dec' 22)

Web and Coding Club (WnCC)

- Developed the backend of a multiplayer strategy game in C++ that runs on automated scripts
- Utilised the C++ Boost libraries for socket communication and SFML libraries for graphic integration
- Developed optimized multithreaded asynchronous servers to reduce latency

Ray Tracing Engine | [GitHub](#) | (Jul' 23-Aug' 23)

Self Project

- Developed a simple ray tracing engine in C++
- Implemented fundamental ray tracing techniques, including ray-object intersection, shading models, reflection, refraction, and shadows, resulting in realistic and visually compelling rendering of 3D scenes.
- Implemented shaders for different types of material surfaces, including diffuse, specular, transparent, and reflective materials, adding realism and visual variety to the rendered scenes produced by the ray tracing engine.
- Incorporated anti-aliasing techniques into the ray tracing engine, reducing visual artifacts such as jagged edges and pixelation, resulting in smoother and more visually appealing images.

8.3 Algorithmic Design

VCSync | [GitHub](#) | (May' 23-Jun' 23)

Self Project

- Designed and developed a custom version control system from scratch using Python
- Implemented core functionalities such as commit tree, allowing for efficient tracking and management of project history, enabling users to visualize the commit history and easily navigate through different versions of the codebase.
- Incorporated advanced features of version control systems, including git stash, enabling users to temporarily save changes without committing them, and pull/push operations for seamless collaboration and synchronization of code across multiple repositories.
- Leveraged graph theory algorithms to design and implement the commit tree functionality in the version control system, enabling efficient and intuitive visualization of the project's history and branching structure.
- Utilized cryptographic functions to hash the source code files, reducing their size while ensuring data integrity and enabling efficient storage and retrieval of code snapshots in the version control system.

SKILLS

Programming: C, C++, Python, PyTorch, Tensorflow, Keras, NodeJS, React, MERN, Cython, CUDA

Technologies: Git, Docker, Linux, Bash, VSCode, Kubernetes, AWS

POSITION OF RESPONSIBILITY

Department Academic Mentor

- Selected as one of the DAMP mentors for the 30 membered department's academic team after rigorous screening interview rounds among 100+ students.
- Responsible for guiding junior students regarding anything related to academics, career and personal life.

Convener: Web and Coding Club

- Selected out of 200+ applicants to be part of a team representing the interests of 10k+ students
- Organized Internship Preparation Program and sessions on DBMS/OOPs/OS helping a total of 400+ students.
- Conducted a Git GitHub workshop and avidly catered to the doubts of participants.
- Conducted an Open Source session for GSoC enthusiasts to make them aware regarding open-source development with some tips in general.

System Administrator: Hostel 2

- Part of the 20 membered **Network Committee** of IIT-Bombay responsible for the smooth functioning and maintenance of the network infrastructure of 18 hostels, 50+ labs and department buildings
- Identify, **debug**, and effectively fix issues faced by the residents using practical **networking knowledge** of the core infrastructure of IITB
- Efficiently **coordinate** and **liaise** between the residents and the Computer Centre to facilitate network **repairs**

Founding Manager, AI Community IIT-Bombay

- Founder of the official AI Community of IIT-Bombay catering to more than 3000 undergraduates and 1000 post graduates
- Building the Institute's first ChatGPT-like chatbot called InstiGPT leading a team of 10 AI and software engineers
- Collaborating with HDFC bank on 3 machine learning projects aiming to revolutionise the insurance industry by automating manual processes
- Managing 4 projects and a team of more than 40 members to build cutting-edge AI projects which are scalable to a large audience

TEACHING EXPERIENCE

Mentor: Seasons of Code 2023

- Mentored 20 mentees for the DeepCarlsen project of SoC 2023.
- Guided the students to create a deep learning-based chess engine using various different approaches as studied in several research papers.
- Helped mentees clear their doubts and guided them in reading several research papers
- Mentored the students in debugging their code and optimizing their model for best performance.

RELEVANT COURSEWORK

Relevant Major coursework: Calculus I-II, Linear Algebra, Differential Equations I-II, Introduction to Electrical and Electronics Circuits

Minor coursework & MOOC: Discrete Computational Structures, Introduction to Object-Oriented Programming, Data Structures and Algorithms, Natural Language Processing, Deep Learning, Mathematical models for systems and control, Introduction to Machine Learning, Optimisation