Bhushan Santosh Shah

+1 669 689 9610 | bhsshah@cs.stonybrook.edu | linkedin.com/in/bhushanshah05 | https://github.com/bhushansshah

EDUCATION

State University of New York, Stony Brook **Aug' 2025 - present**
Master of Science (MS), Computer Science

College of Engineering, Pune, Maharashtra, India **Dec' 2021 - May 2025**
Bachelor of Technology (B. Tech), Computer Science and Engineering GPA: (8.66 / 10)

SKILLS

Languages: Python, R, SQL, FastAPI
Libraries: Verl, TRL, LangChain, LangGraph, Lighteval, Liger Kernel, PyTorch, Tensorflow, Keras, Scikit-learn, NLTK, Pandas, NumPy, Seaborn, Matplotlib, OpenCV
Skills: Deep Learning, Machine Learning, Computer Vision, Reinforcement Learning, Data Structures and Algorithms, Artificial Intelligence, Data Science

PATENTS

- **A Generative AI System for Food Safety, Grading, Compliance and Personalized Nutrition Evaluation** **April 2025**
  - IN Patent App. 202521033716 (Patent pending)

PUBLICATIONS

- **Bhushan Shah** et. al., **"Decoding News Bias: Multi Bias Detection in News Articles",** *NLPIR '24: Proceedings of the 2024 8th International Conference on Natural Language Processing and Information Retrieval, 10.1145/3711542.3711601*

EXPERIENCE

Research Assistant, Centre for Visual Computation, SBU **Oct' 2025 - Present**
AR | XR | Computer Vision | Semantic Mapping
- Exploring AR-based terrain and lighting visualization for architectural and landscape design.
- Investigated real-time environmental mapping and semantic understanding for context-aware AR design applications.

AI Intern, IBM **Sep' 2024 – March 2025**
WatsonX | LLM | Zero-Shot Learning | Typesense| AI-based Food Assessment
- Curated **2M+ food products labeled dataset** from OpenFoodFacts with nutrition, ingredient, and health scores.
- Benchmarked GPT-4o, Mistral Large 2, Llama 3.1 405B, and Gemini Flash 2; identified **GPT-4o and Mistral Large 2** as top performers for **scores prediction task**.
- Built a **Generative AI-powered application** allowing users to search the score of food products via product name, URL, or image input.
- **Improved search latency 10× (2s → 0.2s)** by replacing API calls with **Typesense search engine** and improved response quality using **Chain-of-Draft prompting** and **LLM caching**.

SDE Intern, Pattern Technologies **May 2024 - July 2024**
Next.js | React | Tailwindcss | Typescript
- Migrated a product information management app from **React** to **Next.js**, using **SSR to boost scalability and load speed**.
- Integrated backend APIs for real-time updates while updating product information like attributes, images, videos.
- **Built** and **optimized UI components** for product data (attributes, images, videos).

RESEARCH EXPERIENCE

Decoding News Bias: Multi Bias Detection in News Articles **Sep' 2024 - Nov' 2024**
Multi Bias Detection | LLM | NLP
- Built a weakly supervised labeled dataset using GPT-4o-mini, Llama 3.1 405B, Mistral Large 2, and Gemini Flash 2, applying majority voting to minimize model bias.
- Mitigated data imbalance with **StratifiedKFold sampling** and benchmarked multiple models (BERT, DistilBERT, ALBERT, RoBERTa, XLNet), with BERT performing best.
- Demonstrated potential for **news aggregators** to flag biased content and promote fair and responsible reporting.

Optimizing Insertion Operator in Dynamic Ridesharing **April 2023 - June 2023**
Data Structures and Algorithms | Segment Tree | Dynamic Programming
- **Implemented research paper**, "*An Efficient Insertion Operator in Dynamic RideSharing Services*" by Yi Xu et. al., published in IEEE Conference on Data Engineering for dynamic ridesharing applications.
- Leveraged **segment trees** and **dynamic programming** to reduce the time complexity of the Insertion Operator from **$O(n^3)$ to $O(n\log(n))$.**

ACADEMIC PROJECTS

Reasoning Enhancement for Small LLMs **Oct 2025 – Present**
Reasoning | LLMs | LoRA
- Developing **reasoning-optimized** versions of **Qwen (4B/8B)** and **LLaMA-3.1-7B** models using DeepSeek-inspired reasoning process.
- Implementing fine-tuning with **LoRA** and **reasoning adapters**, to improve logical consistency and multi-hop reasoning accuracy.

Video Moment Retrieval **June 2025 – Present**
Transformer | Multimodality | Neural Networks
- **Developing** a system for **retrieving video segments** that semantically match natural language queries.
- Extending the **FlashVTG** architecture by **incorporating audio signals** thus providing additional cues to improve model performance.
- Leveraging the **QVHighlights** dataset containing **10,200** annotated video segments for training and evaluation.