Vishvesh Trivedi

(551)362-9818 vishvesh106@gmail.com vntrivedi nerdyvisky

Education

New York University, Courant Institute of Mathematical Sciences **Sep 2024 – May 2026**
Master of Science, Computer Science, *GPA: 4.0/4.0*

National Institute of Technology, Surat **Dec 2020 – May 2024**
Bachelor of Technology, Computer Science, *GPA: 9.07/10.0*

Coursework: Data Structures & Algorithms, Programming Languages, Deep Learning, CV, NLP, Operating Systems, Database Management, Computer Architecture, Networks, Probability & Statistics, Mathematics, Data Science

Selected Publications

C=Conference, J=Journal, S=In Submission

[C.1] S. Maniyar*, **V. Trivedi***, A. Mondal, A. Mishra, and C.V. Jawahar (2025). **AI-Generated Lecture Slides for Improving Slide Element Detection and Retrieval**. In *19th International Conference on Document Analysis and Recognition. Cham: Springer Nature* DOI: https://doi.org/10.48550/arXiv.2202.01037

Experience

Bio-DaSH, New York University Langone Health **Nov 2024 – Present**
Data Scientist & ML Researcher *New York, United States*

- Diagnosing performance drift across 5+ AI tools deployed at NYU Langone Health based on temporal metrics monitoring through Kolmogorov-Smirnov, DeLong, and Fermi-Dirac tests. Advised by Prof. Gustavo Stolovitzky.
- Certified EPIC EHR Data Analyst with hands-on experience in using metrics observability tools like Soda and generating population-level insights from a large HIPAA-compliant EPIC EMR dataset of over 300M US patients.
- Contributed to two NSF grant proposals on Health AI policy and won 2nd place, along with a $1300 prize, in the NYU GSAS Threesis Challenge 2025 for a compelling 3-minute talk on the critical need to monitor AI tools in hospitals.

Center for Visual Information Technology, IIIT Hyderabad **Jan 2024 – Aug 2024**
Research Engineer *Hyderabad, India*

- Orchestrated a novel LLM-based pipeline to generate 18,000 high fidelity synthetic slides using university textbooks.
- Constructed a 1050-slide real lecture slides dataset and used LoRA-adapters to train LayoutLMv3, YOLOv8, DETR, LLaVa-1.5, CLIP, models with synthetically generated data to improve SOTA performance across 3 slide based tasks.
- Published findings as an oral presentation at ICDAR 2025 and won the ICDAR 2024 competition in VQA on handwritten documents. Received 200+ downloads for the SynSlides dataset within 2 weeks of release on HuggingFace.

Wells Fargo **May 2023 – Jul 2023**
Summer Analyst *Hyderabad, India*

- Pioneered a web-based AI-powered internal company tool that produces semantic-aware audio-transcriptions of PPT presentations that is 40% faster than screen-readers, and directly impacts 15000 visually impaired WF employees.
- Devised an end-to-end MLOPs pipeline to automate error reporting and reduce redeployment latencies by 25%
- Spearheaded the team showcase event and completed 4 professional certifications on ML-Ops best practices, Data Governance, Agile, and Scrum methodologies. Received full-time return offer but turned down for higher studies.

Projects

Attention-Aware DPO for Reducing Hallucinations in Multi-Image QA [Code] [Website] [Report]
Hugging Face, PyTorch, Python, Bash, HPC, LLM-as-a-judge, Machine Learning, Deep Learning

- Trained LLaVa-1.5 with a novel Attention DPO loss function to increase multi-image VQA accuracy by 8.5%
- Used AdaptVis to optimize model performance at inference and push performance gain to 10% over base model.
- Devised a powerful LLM-as-a-judge using Gemini-2.5-Pro to rate outputs on quantifiable heuristics.

pptGEN: generating synthetic lecture slides on-the-fly as you speak! [Code] [Demo] [Docs]
python-pptx, Langchain, Machine Learning, RAG, Synthetic Data, LoRA, OpenAI, Pydantic, Image Retrieval, Python, Bash

- Built an end-to-end executable app that generates a lecture slide within 15 seconds as the professor speaks.
- Devised a lightweight, modular pipeline that generates coherent multimodal content using LLMs, automatically assigns slide-layout, and preserves presentation style for each slide.

ClinicalML : Traditional Machine Learning vs SOTA LLMs for Clinical Outcome Prediction [Website] [Report]
Python, Pandas/Polars, Numpy, Sci-kit learn, Docker, MySQL, EHR, Machine Learning, Asynchronous Programming

- Engineered a pipeline that extracts low-dimensional drugs and diagnosis from MIMIC-III clinical notes to train classical 4 ML families (Bagging, Boosting, Logistic) for ICU Morality risk and Length of Stay prediction tasks
- Findings indicate only 9% (0.64 vs 0.58 Marco-F1) and 13% (0.38 vs 0.33 Marco-F1) performance gap between SOTA medical LLMs and best ML model for Mortality and Length-of-Stay prediction respectively.

Technical Skills

Languages: Python, C/C++, Java, SQL (Postgres, MySQL), XML, HTML/CSS, JavaScript, TypeScript, Bash/Zsh
Tools/Technologies: AWS, React, REST APIs, Metabase, GCP, Azure, RAG, Docker, GIT, Mongodb, Redis
Frameworks: Sklearn, Pandas, Numpy, Pytorch, TensorFlow, Matplotlib, Tableau, Flask, Django, Streamlit