Prabhav Singh

+1-(669)-255-1483 | psingh54@jhu.edu | linkedin.com/in/prabhavsingh | Github | Google Scholar
3900 N Charles Street, Baltimore MD, 21218

Education

Johns Hopkins University — Baltimore, MD
Masters of Science (Thesis) - Computer Science — *Expected May 2026*

- GPA: **3.9/4**, Specialization: Human Language Technologies
- Advisors: Prof. Jason Eisner, Prof. Jesus Villalba
- Selected Coursework: *Natural Language Processing, Advanced Statistical Self Supervised NLP, Statistical ML, Information Extraction from Speech, Replicable Theory of ML*

Delhi University — Delhi, India
Bachelor of Engineering (Honors) - Electrical Engineering — *August 2022*

- GPA: **3.83/4**
- Advisor: Prof. KPS Rana

Industry Experience

Machine Learning Engineering Intern — June 2025 – Present
Insors — *San Jose, CA*

- Working on building multi-agent pipelines to increase zero-touch rate for production APaaS pipelines for 2 clients.
- Involved in finetuning a two-stage hierarchical pipeline for tax bracket and chapter identification for products for an E-Commerce giant to improve efficiency by 40%.

Founding Machine Learning Engineer — January 2023 – August 2024
Sapper.ai — *Bangalore, India*

- Responsible for developing **glazeDONUT** - An OCRLess document extraction transformer. The development was an **Industry first**. Reduced extraction cost by **65% per page and improved accuracy by 30%**.
- Drove the ML efforts for 4 Clients to Production. The projects today contribute to **60% of the firm's revenue**.

Data Scientist - I — July 2022 – January 2023
Ola — *Bangalore, India*

- Worked on building and engineering data pipelines as part of the EDP team (Data Platform) for **customer facing real-time analytics** for multiple verticals at Ola.
- Migrated Ola's nationwide ride allocation ML Models to a ***KubeFlow Platform*** to enable MLOps. and Managed over **2PB Data migration** from Azure to AWS to **reduce costs by 30%**.

Research Experience

Graduate Research Assistant — August 2024 – Present
CLSP, JHU — *Baltimore, MD*

- Currently working with **Prof. Jason Eisner** at CLSP, JHU. Working on developing AnnotationArena, a novel approach to make it easier, faster, and cheaper to spin up LLM-based annotation systems, drive their performance upward, and evaluate them.
- Also worked with **Prof. Jesus Villalba** on a novel method for speaker diarization by **replacing heuristic clustering** in inference stage with learned speaker counts via a multitask multimodal approach.
- Involved in the team that was selected for a submission to **NIST SRE 2024**. Implemented a novel quality filtering multimodal approach to low-resource speaker recognition.

Undergraduate Research Assistant — October 2018 – April 2022
APC Lab, Delhi University — *Delhi, India*

- Worked under Prof. KPS Rana and Prof. Vineet Kumar at the APC Lab (NSIT) during my undergraduate degree, where I mostly worked on multimodal methods in NLP.
- Published 3 peer-reviewed journal papers with the lab working on fields like multimodal emotion recognition, benchmarking ensemble techniques for fake news detection and faster approaches for extractive summarization.

Publications

- **P. Singh**, J. Villalba (2025). EmoJudge: LLM Based Post-Hoc Refinement for Multimodal Speech Emotion Recognition. In Proc. of INTERSPEECH 2025 *[Accepted]*
- **P. Singh**, J. Villalba, N. Dehak (2025). Count Your Speakers! Multitask Learning for Multimodal Speaker Diarization. In Proc. of INTERSPEECH 2025 *[Accepted]*
- T. Thebaud, A. Favaro, **P. Singh** et al. (2024). Multimodal Emotion Recognition Harnessing the Complementarity of Speech, Language, and Vision. *In Proc. of the 26th International Conference on Multimodal Interaction (ICMI '24)*, ACM, New York, NY, USA, 684–689 *First Position, EVAC, ICMI'24* ***Link***
- **P. Singh**, R. Srivastava, K. Rana, & V. Kumar (2023). SEMI-FND: Stacked ensemble based multimodal inferencing framework for faster fake news detection. *Expert Systems With Applications, 215, 119302.* ***Link***
- R. Srivastava, **P. Singh**, K. Rana, & V. Kumar (2022). A topic modeled unsupervised approach to single document extractive text summarization. *Knowledge Based Systems, 246, 108636.* ***Link***
- **P. Singh**, R. Srivastava, K. Rana, & V. Kumar. (2021). A multimodal hierarchical approach to speech emotion recognition from audio and text. *Knowledge-Based Systems, 229, 107316.* ***Link***

Preprints, Workshops & Posters

- J. Villalba, **P. Singh**, J. Borgstrom et al. (2024). The JHU-MIT Submission to NIST 2024 Speaker Recognition Evaluation (SRE24), *Accepted & Presented at NIST SRE24*
- **P. Singh**, H. Shi, J. Eisner (2025). Active Learning and Feature-Acquisition with LLMs and Humans *Best Poster Award, MASC-SLL 2025* ***Link***

Skills

Languages: Python, C++, SQL (Postgres), NoSQL (MongoDB)
ML Areas: Statistical, Language Processing, Speech Processing, Transformers, Deep Learning, LLM-Evaluation
Infrastructure, Cloud & Frameworks: Flink, Kubernetes, Docker, PyTorch, Tensorflow, MXNet, MLFlow, KubeFlow, AWS, Microsoft Azure

Teaching and Services

Graduate Course Assistant Jan 2025 – May 2025
NLP for Self-Supervised Learning *JHU, Baltimore MD*

- In addition to the general duties as a TA, I also contributed in writing questions for the 3 quizzes conducted in the class.

Activities and Awards

- ***Premier Research Excellence Award*** (2021 & 2022): Awarded by Delhi University for publishing 3 papers in a journal with IF ≥ 8. Included Semester Scholarship for Fall 2022 and Fall 2021.
- ***3rd Position (Gridlock 2020)***: Awarded by Flipkart National Hackathon for developing a fashion trend predictor for the E-Commerce Giant in India.
- ***Member and Instructor - IEEE Delhi University***: Conducted multiple classes for students interested in technology, specifically in topics like Python Programming and AI.