

Divyanshu Aggarwal

☎ +91 837 695 0600 | @ divyanshuggwrw@gmail.com | 🔗 LinkedIn | 📄 Google Scholar | 🐙 GitHub | 🌐 Personal Website | 📍 Delhi, India

EDUCATION

Delhi Technological University (Formerly Delhi College of Engineering)

Delhi, India

Bachelors of Technology

August 2017 - May 2021

Key Courses: Data Structures and Algorithms, Statistics and Probability, Operating Systems, Computer Networking, Database Management Systems, Pattern Recognition, Machine Learning, Natural Language Processing, Computer Vision

PROFESSIONAL EXPERIENCE

Microsoft

India

Applied Scientist 2 | India Applied Sciences Team

Aug 2025 - Present, Full-Time

- **LLM Adaptation and Finetuning:** Finetuning SLMs to enhance the application intelligence of M365 applications by finetuning LLMs and SLMs on synthetically data generated generated using LLMs from private eyes-off datastores. Deploying models and evaluating models and scale and evaluating them on real world customer data.
- **Model Distillation via RL:** Exploring on-policy distillation of blackbox LLMs to build cost efficient purpose and specific SLMs.
- **Research Contribution:** Strengthening innovation ecosystem by curating and disseminating influential papers, leading reading groups, and delivering impactful presentations that drove collaboration, innovation, and continuous learning.

Microsoft Research

India

Research Fellow | Mentored by [Dr. Sunayana Sitaram](#)

Sep 2023 - Jul 2025, Full-Time

- **Improving Cross Lingual Transfer and Language Adaptation in LLMs:** Implemented cross-lingual transfer to adapt pretrained LLMs to newer languages other than english, utilizing *Active Forgetting* and *Modular Deep Learning* methodologies .
- **Multilingual Continual Learning in LLMs:** Addressed the issue of *Catastrophic Forgetting of English Proficiency* in LLMs post *finetuning on multilingual datasets*. Directed cross-team research collaborations, ensuring coordination among various stakeholders.
- **Multilingual PEFT in LLMs:** Analyzed the influence of various *QLoRA* finetuning settings on multilingual downstream tasks.
- **Multilingual Evaluation of LLMs:** Conducted benchmarking of numerous models across diverse languages, tasks, and modalities.
- **Distributed LLM Training:** Managed *100+ GPU cluster* to *optimize LLM training* efforts for pre-training and post-training.

American Express, AI Labs

Gurugram, India

AI Researcher | Advanced NLP Team

Aug 2022 - Aug 2023, Full-Time

- **Finetuning LLMs and SLMs:** Spearheaded the *finetuning of LLaMA models* on large-scale customer datasets, significantly improving response generation quality in customer care workflows. Also designed and deployed a robust *call log classification pipeline* integrating *abstractive summarisation*, and *DeBERTa, RoBERTa* for the classification of texts with high precision and recall. Designed scalable data pipelines on *petabyte-scale Hadoop clusters* to efficiently curate and pre-process training data sets.
- **Research Contribution:** Strengthened Amex AI Labs' research ecosystem by curating and disseminating influential papers, leading reading groups, and delivering impactful presentations that drove collaboration, innovation, and continuous learning.

OPEN SOURCE CONTRIBUTIONS

csinva/imodelsX: Added [batching](#) and [minimum frequency of N-grams](#) to consider in Aug-Gam module.

marin-community/marin: Added continual pretraining script for augmenting the multilingual capabilities of the LLM.

axolotl-ai-cloud/axolotl: Refactored training pipeline to improve integration capabilities making it more flexible and controllable.

adapter-hub/adapters: Fixed modeling mistral file which enabled it to run with flash attention 2.

lm-sys/FastChat: Added Scripts for LoRA Finetuning T5 with DeepSpeed for training on Multiple GPUs

SKILLS SUMMARY

Machine Learning: Python, SQL, Pandas, Numpy, Scikit-learn, Scipy, Pytorch, Pytorch Lightning, deepspeed, AzureML

NLP: LM-Evaluation-Harness, Transformers, Datasets, Huggingface PEFT, HF Accelerate, Adapters, Axolotl, LLM Training

OS and Platforms: Linux, macOS, Unix, Bash, Git, GitHub, Azure, AWS, Docker, CI/CD, ML-OPS, Distributed Systems

Software Eng: Data Structures, Algorithms, Low-Level Design, System Design, FastAPI, OOP, Kubernetes

PUBLICATIONS

Exploring Continual Fine-Tuning for Enhancing Language Ability in Large Language Model

*Divyanshu Aggarwal**, Sankarshan Damle*, Navin Goyal, Satya Lokam and Sunayana Sitaram

 [paper](#)

ACL 2026

Improving Cross Lingual Transfer by Pretraining with Active Forgetting

Divyanshu Aggarwal, Ashutosh Sathe and Sunayana Sitaram

 [paper](#)

EMNLP 2025

MAPLE: Multilingual Evaluation of Parameter Efficient Finetuning of Large Language Models

Divyanshu Aggarwal, Ashutosh Sathe, Ishaan Watts and Sunayana Sitaram

 [paper](#)

ACL 2024

Improving Consistency in LLM Inference using Probabilistic Tokenization

*Ashutosh Sathe**, *Divyanshu Aggarwal** and Sunayana Sitaram

 [paper](#)

NAACL 2025

MEGAVERSE: Benchmarking LLMs Across Languages, Modalities, Models and Tasks

Sanchit Ahuja, *Divyanshu Aggarwal*, Varun Gumma Kalika Bali and Sunayana Sitaram

 [paper](#)

NAACL 2024

IndicXNLI: Evaluating Multilingual Inference for Indian Languages

Divyanshu Aggarwal, Vivek Gupta and Anoop Kunchukuttan

 [paper](#)

EMNLP 2022

Evaluating Inter-Bilingual Semantic Parsing for Indian Languages

Divyanshu Aggarwal, Vivek Gupta and Anoop Kunchukuttan

 [paper](#)

NLP4ConvAI at ACL 2023

XInfoTabS: Evaluating Multilingual Tabular Natural Language Inference <i>Bhavnick Minhas, Anant Shankhdhar, Vivek Gupta, Divyanshu Aggarwal and Shuo Zhang</i>		 paper FEVER at ACL 2022
Fine-tuning Distributional Semantic Models for Closely-Related Languages <i>Kushagra Bhatia, Divyanshu Aggarwal and Ashwini Vaidya</i>		 paper VarDial at EACL 2021
A Review of Deep Learning Techniques for Protein Function Prediction <i>Divyanshu Aggarwal and Yasha Hasija</i>		 paper IEEE INCET 2022
ACADEMIC RESEARCH EXPERIENCE		
AI4bharat, IIT Madras <i>Research Collaborator Advisor: Dr. Anoop Kunchukuttan</i>		Remote 2022 - 2023
• Indic LLMs: Conducted comprehensive evaluations and fine-tuning of large language models for Indic languages.		
University of Utah <i>Research Collaborator Advisor: Dr. Vivek Gupta</i>		Remote 2021 - 2022
• Multi-Lingual NLI and Tabular NLI: Research project on multi lingual NLI, Tabular NLI and Semantic Parsing. Collaborated with researchers from University of Utah, Bloomberg London and Microsoft Research. The Research works are published and under review at several prestigious NLP conferences including ACL, NAACL, EMNLP		
Indian Institute of Technology, Delhi <i>Summer Fellowship Research Intern Advisor: Dr. Ashwini Vaidya</i>		Delhi, India 2020 - 2021
• Multi Lingual Semantic Similarity Dataset, Baselines and Analysis: Multi lingual NLP for Indian Languages .Work published at VarDail Workshop at EACL 2021 Conference.		
ACADEMIC SERVICE		
Conferences: EMNLP 2024, NAACL 2024		
Workshops: NLP4ConvAI@ACL2024, ACM AIMLSys Workshop 2023, Continual FOMO Neurips 2024, FITML Neurips 2024		