

ARJUN CHOUDHRY

✉ arjuncho@andrew.cmu.edu ◊ 📍 Pittsburgh, USA

👤 Website ◊ 📄 Google Scholar ◊ 🌐 LinkedIn

RESEARCH INTERESTS

My research focuses on pushing the boundaries of foundation models and their reasoning ability beyond just text and image modalities. I am particularly interested in:

1. Developing pragmatic multimodal foundation models capable of addressing multifaceted challenges across diverse data types including text, images, time series, audio, video, code, and tabular data. [ICML 2024, ICML DMLR 2024, AAAI 2024, NeurIPS DGM4H 2023].
2. Creating efficient inference methods for large-scale foundation models for deployment in edge devices.
3. Designing comprehensive evaluation frameworks and benchmarks for analyzing systemic design flaws and reasoning ability of current state-of-the-art approaches. [NeurIPS 2023, NeurIPS TSALM 2024, ICAIF FM4TS 2024]

EDUCATION

School of Computer Science, Carnegie Mellon University Pittsburgh, USA
Master of Science in Intelligent Information Systems (ML/NLP with Research) Aug 2023 - May 2025 (expected)

Advised by: Prof. Artur Dubrawski

CGPA: 4.22/4.00

Relevant Coursework: Advanced NLP (PhD), Multilingual NLP (PhD), Multimodal ML (PhD), Question Answering (PhD), Advanced Introduction to ML (PhD) (audit)

Delhi Technological University (erstwhile Delhi College of Engineering) New Delhi, India
Bachelor of Technology in Information Technology Aug 2018 - Jun 2022

Advised by: Prof. Seba Susan, Prof. Dinesh Kumar Vishwakarma

CGPA: 9.66/10

Department Rank: 3/138; University Rank: 12/2036

Relevant Coursework: Machine Learning, Natural Language Processing, Computer Vision, Data Warehousing & Data Mining, Artificial Intelligence & Expert Systems, Algorithms Design & Analysis, Database Management Systems, Applied Mathematics

EXPERIENCE

Auton Lab, Carnegie Mellon University Pittsburgh, USA
Graduate Student Researcher - Advised by Dr. Artur Dubrawski Sep 2023 - Present

- Worked on MOMENT, a family of open time series foundation models (TSFMs). Compiled the Time series Pile, a collection of time series datasets for training TSFMs. Introduced an improved evaluation benchmark for TSFMs, focusing on zero-shot and limited supervision settings. MOMENT achieved near state-of-the-art performance on 5 time series downstream tasks. Introduced guidelines for curation of future time series datasets for training TSFMs. Papers accepted at **ICML 2024** and **ICML 2024 DMLR workshop**. [Paper 1, Paper 2]
- Worked on JoLT, a model jointly modeling time series and text for clinical applications using MOMENT and pre-trained text models. Aligned time series and text representations using a Q-former. Performed experiments for medical time series summarization. JoLT outperformed recent image captioning and medical question-answering approaches. Preliminary works accepted at **NeurIPS 2023 DGM4H workshop** and **AAAI 2024 Student Abstract Track (Best Paper Award)**. [Paper 1, Paper 2]
- Currently evaluating the reasoning ability of 7 recent LLMs on time series tasks in a configurable and scalable way. Curated a synthetic dataset of time series QA pairs comprising 5 core time series understanding tasks. Preliminary works accepted at **NeurIPS 2024 TSALM workshop** and **ICAIF 2024 FM4TS workshop**. [Paper 1, Paper 2]
- Currently working on enabling faster inference of foundation models by compressing them to extreme levels by intelligently removing intermediate layers. Achieving exponential decrease in inference time and memory requirements with negligible drop in performance across standard LLM evaluation benchmarks.

Auton Lab, Carnegie Mellon University Pittsburgh, USA
Research Intern - Advised by Dr. Artur Dubrawski Feb 2023 - Aug 2023

- Worked on AQuA, a benchmarking tool for evaluating label error detection approaches and machine learning models trained in the presence of label noise. Surveyed over 60 papers on label error detection, and created a design space to delineate concrete design choices of label error detection models.

- Created an assumption-based taxonomy to streamline decisions for the correct label error detection approach for a given use case. Conducted over 2400 experiments using 4 label error detection approaches, 17 datasets across 4 modalities, 3 noise rates, 7 noise types, and 10 classification models. Paper accepted at **NeurIPS 2023**. [Paper]

Scientific Analysis Group, Defence Research & Development Organization

New Delhi, India

Research Intern - Advised by Mr. Sanjay Kumar [Certificate]

Apr 2022 - Dec 2022

- Developed a Multi-Level Reversible Data Hiding framework for plain and encrypted domains using a gradient and adaptive most significant bits approach. Proposed two modifications: one with overhead information passed separately and the other with overhead information passed within the image. Substantially increased the number of bits that can be embedded in an image compared to prior approaches.

IKB Lab, Université du Québec à Montréal

Montréal, Canada

Research Intern - Advised by Dr. Marie-Jean Meurs [Certificate]

Apr 2022 - Nov 2022

- Proposed an adversarial approach using unlabelled corpora for improved feature extraction in Named Entity Recognition models for limited-resource languages like French. Achieved an improvement of up to 0.05 in F1 score. Student abstract paper accepted at **AAAI 2023**. Full paper accepted at **ECIR 2023**. [Paper 1, Paper 2]
- Worked on a cross-lingual Named Entity Recognition model using multi-lingual language models and domain adaptation for improved generalization between different languages and scripts. Evaluated our approach for extremely low-resource languages like Basque, Czech, and various Indic languages.

University of Technology Sydney | Biometric Research Laboratory, DTU

Sydney, Australia

Research Intern - Advised by Dr. M. Prasad, UTS & Dr. D. Vishwakarma, DTU [Certificate]

Apr 2022 - Oct 2022

- Led three teams of interns from Biometric Research Lab, DTU and UTS on three projects during the internship.
- Proposed a community and K-shell decomposition-based algorithm for ranking seed nodes based on the strength of connections with the adjacent communities for low-budget Influence Maximization. Student abstract paper accepted at **AAAI 2023**. Extension paper **under minor revision in Information Sciences** journal. [Paper 1, Paper 2]
- Proposed a discretized Harris Hawk's Optimization algorithm and a novel Neighbor Scout strategy for Influence Maximization to accelerate the optimization convergence. Paper **published in Applied Soft Computing** journal. [Paper]
- Proved the efficacy of using emotion labels in multi-task settings for improved generalization in cross-domain Fake News Detection. Further proposed an emotion-guided domain adaptive approach for better alignment of domains in cross-domain settings. Student abstract paper published at **AAAI 2023**. Short paper published at **ICON 2022**. [Paper 1, Paper 2]

Biometric Research Laboratory, Delhi Technological University

New Delhi, India

Research Associate & Lead - Advised by Dr. Dinesh Kumar Vishwakarma [Certificate]

Jan 2022 - Jan 2023

- Research Lead for various research projects at the Biometric Research Laboratory, DTU.
- Studied the correlation between fake news and Plutchik's and Ekman's emotion classes in the feature space using PCA. Used these emotion labels in an augmented multi-task setting to improve the detection of fake news and rumors across various domains. Paper **published in IEEE Transactions on Computational Social Systems**. [Paper]
- Worked on temporal analysis of fake news on social networks and its effect on deep learning classifiers' performance. Improving model robustness to minimize performance loss in classifiers over time. Guided one intern on this project.
- Built a Meta-heuristics algorithm and transformer embedding-based framework for real-time object tracking. [GitHub]

Kylo Apps

New Delhi, India

Data Science Intern - Advised by Mr. Anubhav Mittal, CTO [Certificate]

Jan 2022 - Mar 2022

- Built an Optical Character Recognition and medical Named Entity Recognition-based framework for extracting patients' medical results and diagnostic test information from uploaded PDF diagnostic reports in different tabular formats.

Department of Computer Science & Engineering, Delhi Technological University

New Delhi, India

Research Assistant - Advised by Dr. Sanjay Kumar [Certificate]

Sep 2021 - Jan 2022

- Proposed a Modified Community Diversity algorithm for Influence Maximization incorporating *two-hop* neighborhood information to avoid overlapping of influence (commonly called *Rich-Club effect*) among selected nodes. Paper accepted in **Journal of Intelligent Information Systems**. [Paper]

Department of Information Technology, Delhi Technological University

New Delhi, India

Undergraduate Thesis - Advised by Dr. Seba Susan

Aug 2021 - May 2022

- Proposed TLMOTE, a text-oversampling approach using topic modeling for relevant seed generation and randomized language modeling using three Bi-LSTM models for more meaningful variations between samples generated from similar

seeds. Observed lesser duplicity in samples generated using TLMOTE as compared to LMOTE and DRO. Paper received **Best Student Paper award at FLAIRS 2022**. [Paper, Code, Reproducible Capsule]

- Performed Sentiment Analysis on unlabelled COVID-19 tweets to evaluate the trends in the sentiments associated with Covaxin and Covishield vaccines in India during the second wave. Further linked these variations in the sentiments observed with real-world occurrences using dynamic topic modeling on the tweets. [Paper 1, Paper 2]
- Evaluated the impact of domain-specific transformers models for sentiment analysis on COVID-19 tweets and if introducing synthetic samples using various oversampling approaches can reduce the differences in performance observed. [Paper]

IKB Lab, Université du Québec à Montréal

MITACS Globalink Research Intern - Advised by Dr. Marie-Jean Meurs [Certificate]

Montréal, Canada

May 2021 - Aug 2021

- Worked on Named Entity Recognition (NER) for social media in French and other limited resource languages. Evaluated the impact of language models' pretraining corpora on NER in uni-lingual and multi-lingual setups.
- Trained the XLNet framework on small subsets of the OSCAR French corpus of varying sizes to evaluate the impact of pretraining corpora size on NER.
- Built a Transformer-BiLSTM-CRF framework to evaluate various versions of French and multi-lingual transformers for NER to see how various architectures and pre-training corpora affect downstream performance.

Department of Computer Science & Engineering, Delhi Technological University

Research Intern - Advised by Mrs. Minni Jain [Certificate]

New Delhi, India

Feb 2020 - Sep 2021

- Worked on Word Sense Disambiguation using Fuzzy Graph Centrality Measures in Hindi. Built frameworks for the extraction of fuzzy semantic relation scores like Meronymy, Hyponymy, Hypernymy, and Entailment using the Hindi WordNet, as well as calculation of centrality measures like Betweenness, Closeness, and PageRank.
- Proposed a Multi-Task Learning for Fake News Detection using Emotion Labels as auxiliary information for improved performance. Student Abstract paper published at **AAAI 2022**. [Paper]

Indovision Services Private Limited

Robotic Process Automation Intern - Advised by Mr. Rajiv Mittal, CT&IO [Certificate]

New Delhi, India

Jun 2019 - Jul 2019

- Created an automation workflow using UiPath and C# for extracting employee details from resumes documents and csv files in the company email address and uploading them to the company server. [Report]

SELECTED PUBLICATIONS

For a complete list of my papers, please check [my Google Scholar profile](#).

* and † indicates equal contribution

Accepted and Published Papers

1. Mononito Goswami, Konrad Szafer*, **Arjun Choudhry***, Yifu Cai, Shuo Li and Artur Dubrawski. *MOMENT: A Family of Open Time Series Foundation Models*. In the 41st International Conference on Machine Learning (ICML) 2024. [PDF]
2. **Arjun Choudhry***, Konrad Szafer*, Mononito Goswami, Yifu Cai, and Artur Dubrawski. *Datasets for Time Series Foundation Models*. In ICML Workshop on Data-centric Machine Learning Research (DMLR) 2024. [Paper]
3. Yifu Cai, **Arjun Choudhry***, Mononito Goswami*, and Artur Dubrawski. *TimeSeriesExam: A Time Series Understanding Exam*. In NeurIPS Workshop on Time Series in the Age of Language Models (TSALM) 2024. **Spotlight Paper. 🏆 Best Paper Award Honorable Mention**. [PDF]
4. Yifu Cai, **Arjun Choudhry***, Mononito Goswami*, and Artur Dubrawski. *Do LLMs Understand Financial Time Series?* In ICAIF Workshop on Foundation Models for Time Series: Exploring New Frontiers (FM4TS) 2024. **Oral Presentation**
5. Yifu Cai, Arvind Srinivasan, Mononito Goswami, **Arjun Choudhry**, Artur Dubrawski. *“JoLT: Jointly Learned Representations of Language and Time-Series for Clinical Time-series Interpretation.”* In the 38th AAAI Conference on Artificial Intelligence (AAAI) 2024 (Student Abstract), **🏆 Best Student Abstract Paper Award** and **3-min presentation contest finalist**. [PDF]
6. Kshitish Ghate, **Arjun Choudhry**, and Vanya Bannihatti Kumar. *“Evaluating Gender Bias in Multilingual Multimodal AI Models: Insights from an Indian Context.”* In ACL Workshop on Gender Bias in Natural Language Processing (GeBNLP) 2024.
7. Mononito Goswami*, Vedant Sanil*, **Arjun Choudhry†**, Arvind Srinivasan†, Chalisa Udompanyawit, Artur Dubrawski. *“AQuA: A Benchmarking Tool for Label Quality Assessment.”* In the 37th Conference on Neural Information Processing Systems (NeurIPS) 2023 Datasets & Benchmarks Track. [PDF]

8. Yifu Cai, Mononito Goswami, **Arjun Choudhry**, Arvind Srinivasan, Artur Dubrawski. “*JoLT: Jointly Learned Representations of Language and Time-Series.*” In the 37th Conference on Neural Information Processing Systems Workshop on Deep Generative Models for Health (DGM4H NeurIPS) 2023. [PDF]
9. **Arjun Choudhry***, Inder Khatri*, Pankaj Gupta, Aaryan Gupta, Maxime Nicol, Marie-Jean Meurs and Dinesh Kumar Vishwakarma. “*Adversarial Adaptation for French Named Entity Recognition.*” In the 45th European Conference on Information Retrieval (ECIR) 2023. [PDF]
10. Inder Khatri*, **Arjun Choudhry***, Aryaman Rao*, Aryan Tyagi, Dinesh Kumar Vishwakarma, and Mukesh Prasad. “*Influence Maximization in Social Networks using Discretized Harris’ Hawks Optimization Algorithm*”. In Applied Soft Computing, Elsevier. 2023. [PDF]
11. Aaryan Gupta*, Inder Khatri*, **Arjun Choudhry** and Sanjay Kumar. “*MCD: A Modified Community Diversity Approach for Detecting Influential Nodes in Social Networks.*” In Journal of Intelligent Information Systems, Springer. 2023. [PDF]
12. **Arjun Choudhry***, Inder Khatri*, Arkajyoti Chakraborty, Dinesh Kumar Vishwakarma and Mukesh Prasad. “*Emotion-guided Cross-domain Fake News Detection using Adversarial Domain Adaptation.*” In the 19th International Conference on Natural Language Processing (ICON) 2022. [PDF]
13. **Arjun Choudhry***, Pankaj Gupta*, Inder Khatri, Aaryan Gupta, Maxime Nicol, Marie-Jean Meurs and Dinesh Kumar Vishwakarma. “*Transformer-based Named Entity Recognition for French Using Adversarial Adaptation to Similar Domain Corpora.*” In the 37th AAAI Conference on Artificial Intelligence (AAAI) 2023 (Student Abstract). [PDF]
14. Arkajyoti Chakraborty*, Inder Khatri*, **Arjun Choudhry***, Pankaj Gupta, Dinesh Kumar Vishwakarma and Mukesh Prasad. “*An Emotion-guided Approach to Domain Adaptive Fake News Detection using Adversarial Learning.*”. In the 37th AAAI Conference on Artificial Intelligence (AAAI) 2023 (Student Abstract). [PDF]
15. Inder Khatri*, Aaryan Gupta*, **Arjun Choudhry***, Aryan Tyagi*, Dinesh Kumar Vishwakarma and Mukesh Prasad. “*CKS: A Community-based K-shell Decomposition Approach using Community Bridge Nodes for Influence Maximization.*” In the 37th AAAI Conference on Artificial Intelligence (AAAI) 2023 (Student Abstract). [PDF]
16. **Arjun Choudhry***, Inder Khatri*, Minni Jain, and Dinesh Kumar Vishwakarma. “*An Emotion-Aware Multi-Task Approach to Fake News and Rumor Detection using Transfer Learning*”. In IEEE Transactions on Computational Social Systems. [PDF]
17. **Arjun Choudhry**, Seba Susan, Anmol Bansal, and Anubhav Sharma. “*TLMOTE: A Topic-Based Language Modelling Approach for Text Oversampling.*” In The International FLAIRS Conference Proceedings, vol. 35, 2022. 🏆 **Best Student Paper Award**. [PDF, Code, Reproducible Capsule]
18. **Arjun Choudhry***, Inder Khatri*, and Minni Jain*. “*An Emotion-Based Multi-Task Approach to Fake News Detection.*” In AAAI Conference on Artificial Intelligence (AAAI) 2022 (Student Abstract). [PDF]

Under Submission

1. Aaryan Gupta*, Inder Khatri*, **Arjun Choudhry***, Pranav Chandok, Dinesh Kumar Vishwakarma and Mukesh Prasad. “*A Spreader Ranking Algorithm for Extremely Low-budget Influence Maximization in Social Networks using Community Bridge Nodes.*” Under minor revision at Information Sciences, Elsevier. [PDF]

ACHIEVEMENTS

- NeurIPS TSALM Workshop 2024 Best Paper Award Honorable Mention.
- AAAI 2024 **Best Student Abstract Paper Award 2024**
- Delhi Technological University **Research Excellence Award 2024**, recognizing my research contributions in my undergraduate degree.
- Florida Artificial Intelligence Research Society **Best Student Paper Award 2022**
- Awarded the **Vector Scholarship for Artificial Intelligence 2023** upon acceptance into the Masters program at University of Toronto.
- **MITACS Globalink Research Award 2021**.
- Got a **top 0.58 percentile** rank in the **Joint Entrance Examination Mains 2018 (JEE Mains 2018)** out of over 1,100,000 students across India.
- Stood **1st** among students from all over India in the **Next Genius Critical Thinking Olympiad 2015** Grade 10 category. Offered **full scholarship** worth **Rs 7,000,000 (€65,000)** to attend the IB Diploma Program at the prestigious **Carlsbad International School, Czech Republic**.

INVITED TALKS

- Guest speaker at NeuralAI, DTU. Gave a talk on “*Getting Started with Artificial Intelligence Research as an Undergraduate*” to over 200 freshman and sophomore students - March 2022.
- Invited as a speaker at NeuralAI, DTU. Gave a talk on “*Undergraduate Research Internship Opportunities and How to Crack Them*” to over 250 sophomore and junior students - June 2022.

TEACHING AND MENTORSHIP EXPERIENCE

- Research mentor to a senior undergraduate student at DTU for his B.Tech thesis on Brain-Computer Interface.
- **Teaching Assistant** for Fall 2024 for the course **11631 Data Science Seminar** at LTI, CMU, under **Prof. Marteen Sap**.

ACADEMIC SERVICE

- Served as a reviewer for **NeurIPS TSALM Workshop 2024, ICLR DPFM Workshop 2024, ICML DMLR Workshop 2024, AAAI Spring Symposium on Clinical Foundation Models 2024, IEEE Transaction on Computational Social Systems 2022, and Information Sciences journal 2022.**

SKILLS

Languages	Python, C, C++, Java, SQL
Frameworks	PyTorch, TensorFlow, Pandas, Scikit-Learn, NLTK, SpaCy, Matplotlib, NumPy Flask, HuggingFace
Tools	Git, MySQL, MongoDB, Unreal Engine, UiPath, CI/CD, Linux, GCP, AWS, CUDA