Arjun Ramesh Rao

mailarjunrao@gmail.com | arjun.fyi | arjun-rao | arjunra0 | @arjunra0

Education

University of Colorado, Boulder — *Boulder, US*

M.S. in Computer Science — *Aug. 2019 - May 2021*

- Awarded Lloyd Botway Fellowship for Outstanding Master's students

Ramaiah Institute of Technology — *Bangalore, India*

B.E. in Information Science and Engineering — *Aug. 2014 - June 2018*

- **Best Outgoing Student** (Batch of 2014-2018), Dept. of Information Science & Engineering

Industry Experience

Netflix — *Seattle, US*

Software Engineer, Machine Learning — *March. 2024 - Present*

- Working recommendation algorithms and machine learning infrastructure as part of Netflix's Algorithms Engineering organization.
- Technologies Used: Java, Scala, Python, Tensorflow, Pytorch

Microsoft Corp. — *Redmond, US*

Software Engineer II — *June. 2021 - Feb 2024*

- Worked on recommendations and ranking for Microsoft News and Feeds (Microsoft Start).
- **Built the serving stack for a real-time framework that aggregates counting features** for ranking models for various user events like impressions, clicks, likes, etc. The service can aggregate events at the second level and the serving stack has an **end-to-end latency of 300ms (p95)** and currently has on an average **2M+ QPS** from traffic across North America, Europe and Asia. **Was responsible for design & implementation of the feature storage schema, distributed serving components, caching layer, feature parsing layer** for combining document/request/user information, and the design & implementation of the counting feature configuration interface to support rapid prototyping and A/B Experimentation. Served as a **tech lead for the serving path of the service**.
- Served as a **tech lead for the debuggability workstream** for FY2024 Q1-Q2, working with a cross-functional team of engineers distributed across Egypt, China and Redmond. Worked on feature roadmaps, identified gaps and proposed improvements to infrastructure and tooling for reducing the time to debug issues in the recommendation stack.
- Helped architect a new configuration compiler for Experiment Configuration Management that automated several manual steps required to setup ranking A/B experiments and improved time from idea to experiment by several orders of magnitude (**1 hour to 5 minutes**). Played a **principle role in designing the new workflow, SDK interfaces, and technical architecture for the overall project**. Part of it also involved prototyping a UI based solution that won the "Best Technical Effort/Complexity Hack" in an internal Org-level Hackathon from over 100s of projects. The tool is now released and being adopted by various partner teams.
- Improved platform availability by automating configuration management and cleanup to make changes to any set of experiments using code. Built tools to modify configurations of **100s of experiments with few lines of code**. This allowed us to clean up and migrate deprecated config across experiments in the platform in **minutes from weeks/months** it would previously take.
- Served on the Culture Board for the organization for FY23, helping with brainstorming, budget planning and execution of initiatives to improve org-culture, diversity and inclusion. Was directly responsible for the Peer-Recognition program during my term on the culture board
- Lead the Documentation Enhancement Program efforts working with team-level documentation champs to improve documentation across the News & Feeds organization leading to migration of several wikis to a more discoverable internal wiki platform, and improvements to several pieces of documentation.
- Technologies Used: C++, C#, .Net, Python

Microsoft Corp. — *Redmond, US*

Software Engineer Intern — *May 2020 - Aug. 2020*

- Worked on the Core Ranker Team within Microsoft News and Feeds
- Built automated pipelines for continuously evaluating and improving Content Classification models.
- **Built UHRS apps used across 13 markets with 37 tasks** for obtaining crowd sourced training data for document classification. The data collected from these apps helped improve coverage of our document classifier by orders of magnitude.
  We increased coverage from **8% documents with over 0.8% precision**, to **86% documents with over 0.8 precision**
- Built a UHRS app for Side-by-side (SBS) Feed Evaluation and automated judgement collection on a daily basis. Feedback from the app allowed us to increase the SBS **cold start evaluation metric increased from +8.66 to +22.00**
- Technologies Used: C#, Python, .Net Framework, JavaScript, UHRS

Stride.ai Inc.

Bangalore, India

NLP Engineer *July 2018 - May 2019*

- Designed and built re-usable end-to-end components to operationalize solutions for automated document classification and information extraction. Reduced the turn around time to build a new Proof-of-Concept for unseen set of documents from **over 1 month to 1 week**. The components I built included all parts of the stack from UI, Backend, Storage, Model Training, Inference, Feature Parsing and Serving.
- Worked on information extraction use-cases for **identifying key datapoints from documents** by training custom NLP based models and experimenting with custom model architectures. **Trained over 10+ models for various use-cases.**
- **Built a custom PDF Viewer using PDF.js** that allowed user annotation functions and additional capabilities like support for multiple monitors, side-by-side scrolling of multiple documents, cross-browser support, lazy-loading document pages, etc. The viewer is was used in 10+ Projects and reduced latency of loading **100 page pdfs from 5+ seconds to < 1 second** with the new viewer. Additionally it enabled support for ML based data point highlighting, which was previously not possible.

Google Inc.

San Francisco, US

Developer Programs Engineer Intern *June 2017 - Aug. 2017*

- Built tools to help **track code repositories** and generate **consolidated notifications** for events like issues, comments, etc., for faster triage of issues. The tool reduced the number of notifications generated by Github for activity on a popular open source repository from **30+ per day to 1 per day.**
- Part of the project was **released as open source software**, and can be found at github.com/GoogleCloudPlatform/issuetracker
- Technologies Used: Google Cloud Datastore, Google BigQuery, Google App Engine, GoLang, Angular

Skills

Programming Languages **C++**, **Python**, **C# TypeScript**, **Go**, **Java**, **Scala**

Platforms & Frameworks **.NET**, **PyTorch**, **Tensorflow**, **Angular**, **Cloud Computing**, **Android**

Languages **English**, Hindi, Kannada

Academic Experience

University of Colorado, Boulder

Boulder, US

Emotive Computing Lab, Institute of Cognitive Science *Jan. 2020 - May 2021*

- Worked under Prof. Sidney D'Mello at the Emotive Computing Lab.
- Contributed to active research on modeling collaborative problem solving processes and discourse using state of the art natural language processing techniques and multi-modal machine learning.
- Worked on modeling bias in machine learning models for apparent personality prediction in one way behavioral interviews.
- Technologies Used: **Python**, **PyTorch**, **AWS**

Ramaiah Institute of Technology

Bangalore, India

Senior Project - The Milo IDE (miloide.github.io) *Sept. 2017 - April 2018*

- Built a web-based IDE to help students with no prior programming experience learn Machine Learning and Linear Algebra.
- Customized **Google's blockly project**, and designed a visual programming language that supports data science operations.
- Implemented a data explorer with built-in datasets along with support for using custom numeric, image and textual datasets.
- Implemented common ML algorithms using **Tensorflow.js** as blocks and used **D3.js and Plotly.js** for interactive visualizations.
- Presented and published a paper based on a user study with the IDE at IEEE VLHCC 2018 (See Publications).
- Technologies Used: **Node**, **Javascript**

Select Publications

Using artificial intelligence to assess personal qualities in college admissions

Online, USA

Benjamin Lira, Margo Gardner, Abigail Quirk, Cathlyn Stone, **Arjun Rao**, Lyle Ungar, Stephen Hutt, Louis Hickman, Sidney K D'Mello, Angela L Duckworth *Oct. 2023*

- In Science Advances Journal, vol 9, Issue 41 [Link]
- Worked on training and validation of machine learning models to automattically assess college admission essays.

CPSCoach: The Design and Implementation of Intelligent Collaborative Problem Solving Feedback

Tokyo, Japan / Virtual

Angela EB Stewart, **Arjun Rao**, Amanda Michaels, Chen Sun, Nicholas D Duran, Valerie J Shute, Sidney K D'Mello *July. 2023*

- In Artificial Intelligence in Education AIED 2023 Lecture Notes in Computer Science [Link]
- We present the design of a fully-automated system that assesses and provides feedback on collaborative problem solving (CPS) competencies during remote collaborations

Do Speech-Based Collaboration Analytics Generalize Across Task Contexts?

Online, USA

Samuel L Pugh, **Arjun Rao**, Angela EB Stewart, Sidney K D'Mello *March. 2022*

- Proceedings of 12th International Learning Analytics and Knowledge Conference (LAK22), 2021 [PDF]
- We investigated the generalizability of language-based analytics models across two collaborative problem solving tasks, among 95 triads who collaborated over video conferencing.

Say What? Automatic Modeling of Collaborative Problem Solving Skills from Student Speech in the Wild

(Virtual) Paris, France

Samuel L Pugh, Shree Krishna Subburaj, **Arjun Ramesh Rao**, Angela EB Stewart, Jessica Andrews-Todd, Sidney K D'Mello

June. 2021

- Proceedings of the Educational Data Mining Conference 2021 [PDF]
- We investigated the feasibility of using automatic speech recognition (ASR) and natural language processing (NLP) to classify collaborative problem solving (CPS) skills from recorded speech in noisy environments.

Multimodal, Multiparty Modeling of Collaborative Problem Solving Performance

(Virtual) Utrecht, Netherlands

Shree Krishna Subburaj, Angela EB Stewart, **Arjun Ramesh Rao**, Sidney K D'Mello

Oct. 2020

- Proceedings of the 2020 International Conference on Multimodal Interaction, pp. 423-432. [PDF]
- Analyzed data from 101 triads engaged in computer-mediated collaborative problem solving (CPS) in an educational physics game.
- Investigated the accuracy of machine-learned models trained on facial expressions, acoustic-prosodics, eye gaze, and task context information, computed one-minute prior to the end of a game level, at predicting success at solving that level.
- DOI: 10.1145/3382507.3418877

Milo: A visual programming environment for Data Science Education

Lisbon, Portugal

Arjun R Rao, Ayush Bihani, Mydhili K Nair

Oct. 2018

- Proceedings of 2018 IEEE Symposium on Visual Languages and Human-Centric Computing **(VL/HCC'18)**, pp. 211-215. [PDF]
- Designed and implemented a novel **visual programming environment** to help novice students and non-programmers learn **Data Science and ML concepts** using block based programming.
- DOI: 10.1109/VLHCC.2018.8506504

Context Based Approach for Second Language Acquisition

New Orleans, USA

Nihal V Nayak, **Arjun R Rao**

June 2018

- System paper for Duolingo's shared task on Second Language Acquisition Modelling **(SLAM 2018)**. [PDF]
- Published in the Proceedings of the **NAACL-HLT Workshop** on Innovative Use of NLP for Building Educational Applications **(BEA at NAACL 2018)**.
- Trained a logistic regression model to predict the likelihood of a student making a mistake while answering an exercise on Duolingo. Made use of features inspired by research in **code-mixed language learning** where context plays an important role.
- Result: **AUROC scores for English/Spanish = 0.821**, Spanish/English = 0.790 and French/English = 0.812. **2nd best linear model**, finished 9th overall in SLAM 2018