Aniruddh Sriram

aniruddh.sriram@utexas.edu
aniruddhs24.github.io

Education

University of Texas at Austin, Austin, TX — May 2024
B.S. Computer Science (Turing Scholars), GPA: 3.85 / 4.00
B.S. Mathematics
Minor in Economics

- Honors Thesis: *Contrastive Learning to Improve Retrieval for Real-world Fact Checking*
- Selected Coursework: *Predictive Machine Learning (graduate)*, *Machine Learning Theory*, *Computer Vision*, *Artificial Intelligence*, *Neural Networks*, *Quantum Information Science*, *Linear Algebra*, *Vector Calculus*, *Stochastic Processes*, *Predictive Analytics*, *Mathematical Statistics*, *Real Analysis*

Selected Publications

1. Aniruddh Sriram, Fangyuan Xu, Eunsol Choi, and Greg Durrett. 2024. Contrastive learning to improve retrieval for real-world fact checking. *Proceedings of the Seventh Fact Extraction and VERification Workshop (FEVER) at EMNLP*.
2. Jifan Chen, Grace Kim, Aniruddh Sriram, Greg Durrett, and Eunsol Choi. 2024. Complex claim verification with evidence retrieved in the wild. *Proceedings of NAACL*.
3. Jifan Chen, Aniruddh Sriram, Eunsol Choi, and Greg Durrett. 2022. Generating literal and implied subquestions to fact-check complex claims. *Proceedings of EMNLP*.

Teaching Experience

University of Texas at Austin, Austin, TX

DiRP Mentor — August - December 2023
Led discussion group for undergraduates interested in natural language processing (NLP). Prepared a curriculum and conducted weekly lectures to introduce foundational NLP concepts.

Undergraduate Teaching Assistant — January - May 2022
Led office hours and graded assignments for students enrolled in *DSC 395T: Data Structures & Algorithms*

Work Experience

Voleon Group, Berkeley, CA

Software Engineer — July 2024 - Present
(SECENG) Support quantitative researchers in developing machine learning algorithms for trading

- Develop tools to execute trades based on price forecasts and portfolio goals

Bloomberg LP, New York, NY

Software Engineer Intern — May 2022 - Aug 2022
(PORT DATA Frameworks) Team that manages ETL pipelines for financial data

- Extracted features and designed an algorithm to generate Apache Spark resource configurations using heuristics and KNN with historical runs
- Resulted in a 1.69x speedup for our largest validation job and 27% decrease in unused resources for small jobs

DevRev, Palo Alto, CA (remote)

Machine Learning Intern — June 2021 - Aug 2021

- Developed a word2vec-based model in Python to generate vector representations of developers using GitHub commit data, enabling comparisons via cosine similarity
- Extended to a dual-encoder approach that encodes developers and Github issues onto the same space, enabling intelligent prediction of issue assignees

SELECTED PROJECTS

CUDA Wordle [Code], [PDF]
A CUDA-optimized Wordle solver that uses information-theoretic techniques to update belief distributions about target word

Feedback Analysis [Code]
Tool to analyze textual feedback and produce short, meaningful highlights that characterize sentiment

Binary ISA Classifier [Code]
Automatically predicts the instruction set architecture (ISA) given a binary blob

COV-(AI)D [Code]
Leverage human mobility data to forecast COVID-19 case counts in 50 U.S. states using deep-learning techniques

Machine Learning Articles [Web]
Articles to obtain an intuitive understanding of machine learning topics, emphasizing information I struggled to find when learning these concepts myself

ESLII Lecture Notes [Web]
Collaborated with some friends to read and prepare notes on Elements of Statistical Learning II by Trevor Hastie, Robert Tibshirani, Jerome Friedman