Chinmaya Andukuri

512-919-9174 | andukuri@stanford.edu | LinkedIn | Personal Website | GitHub

Education

Stanford University March 2023 – December 2024
M.S. in Computer Science, concentration in Artificial Intelligence *Stanford, CA*

Stanford University September 2019 – June 2023
B.S. in Mathematical and Computational Science *Stanford, CA*

Technical Skills

Programming Languages: Python, C++, C, SQL, R
Technologies and Frameworks: PyTorch, pandas, HuggingFace transformers, hydra, Weights and Biases, Git/GitHub

Experience

Student Researcher December 2023 – Present
Stanford Artificial Intelligence Laboratory (Computation & Cognition Lab) *Stanford, CA*
- Constructing reusable repositories to study problem-solving and reasoning abilities of language models
- Studying elicitation of preferences by language models through bootstrapping, simulation and self-improvement

Software Engineer Intern June 2023 – August 2023
Capital One *McLean, VA*
- Constructed large language model (LLM) pipeline to provide search capability across company
- Created $6 million in expected savings for HR by embedding >7000 internal documents for semantic search
- Achieved 84% BERTScore F1 similarity between predicted and reference answers on open question-answering tasks

Software Engineer Intern June 2022 – September 2022
Dataherald, YC W21 *Los Angeles, CA*
- Implemented version control system module using Python/Git for MongoDB database with 400+ documents
- Created 20+ self-sufficient data pipelines using Databricks/PostgreSQL to create data visualizations for web app
- Wrote, managed and debugged 50+ MongoDB documents to keep data feeds readily available for clients

Machine Learning Engineer Intern June 2021 – September 2021
AncillaryBox.ai *Arlington, VA*
- Identified lowest-performing points of sale to increase revenue from airline upgrades by tailoring product placement
- Wrote Python scripts to analyze 500,000 rows of customer purchase data for airline products
- Coded k-means clustering, logistic and multivariate regressions to find significant indicators of purchase patterns

Undergraduate Teaching Assistant September 2021 – September 2022
Stanford University (Computer Science Department) *Stanford, CA*
- Communicate complex technical ideas in practice sections and office hours with 15+ students in intro CS courses
- Utilize problem solving skills to evaluate 200+ assignments and exams per quarter in Python and C++

Projects

STaR-GATE: Teaching LMs to Ask Better Clarifying Questions | *VLLM* March 2024 – Present
- Bootstrapped language models to teach themselves to ask better questions
- Achieved 73% win rates as measured by GPT-4 annotator over baseline model
- Developed novel framework for better eliciting human preferences through simulated conversations

manipulativeLMs: Social Reasoning in Language Models | *transformers, LoRA* December 2023 – Present
- Finetuned 7 billion parameter decoder-only language model to improve social reasoning ability
- Constructed 1000-example evaluation benchmark to test manipulative behavior in base- and finetuned- models

Lyric Generation with Transformer-Based RL | *PyTorch, AWS EC2, S3* April 2023 – Present
- Built GPT-2 transformer-based generative deep learning model to produce novel, creative lyrics
- Used reward mechanisms and reinforcement learning with stochastic gradient descent to encourage unique outputs
- Utilized GPUs from remote AWS EC2 instance and S3 bucket to minimize training time and maximize efficiency