

Armaan A. Abraham

armaan.abraham@hotmail.com
778-240-2138
armaanabraham.com

Langley
British Columbia, Canada

Education

University of California, Los Angeles (UCLA)
B.S., Computer Science, B.S., Biophysics
3.98 GPA, summa cum laude

Jun 2023

AI Safety Fundamentals: Alignment
Online course

Oct 2024 – Present

Professional Experience

Staff Researcher

Jun 2023 – Present

Undergraduate Researcher

Sep 2022 – Jun 2023

Prof. Aaron Meyer's Lab, Bioengineering Department, UCLA

- Independently invented Unaligned Low-rank Tensor Regression with Attention (ULTRA) for applications in single-cell RNA sequencing (scRNA-seq) data.
 - I discovered a novel supervised tensor decomposition approach that somewhat resembles “cell attention” that provides a general solution to answering questions of the form “which gene programs in which subsets of cells explain a given external phenotype?” in scRNA-seq.
 - Currently, I am working on refining (e.g. through regularization) and conducting further tests on this approach.
 - I am supervising an undergraduate student who is assisting in these efforts.
- I am the primary technical consultant for an external drug company, for which I develop and apply a mechanistic binding model to aid in their development of multivalent antibody-cytokine fusions for cancer therapy.
 - I used numerical differencing approaches for gradient approximation to optimize over drug space based on desired functional characteristics inferred by the binding model.
 - I replaced an implementation of the binding model that was the lab standard by my own implementation, which was ~4500x faster (by reimplementing with jax, batching, and switching the root finding to operate over log space). This efficiency was required to use the model within an outer optimization loop as described above.
 - Analytically derived an expression for the number of crosslinking receptors as a function of ligand concentration, relying on combinatorics and equilibrium binding constants.
- (See finished manuscripts below)

Software Engineer Intern

Jun 2022 – Sep 2022

Tesla, Vehicle Software

- Led the design and implementation of the event data recorder (EDR) file parser and report generator for use by vehicle owners, regulators, and safety engineers.
- Worked under pressure and a hard deadline, with the EDR file parser needing to be complete within 8 weeks for Tesla vehicles to continue being sold in China.

- Designed, implemented, and optimized a Redis cache architecture for DBC CAN signals, reducing runtime of vehicle campaign creation from 5+ minutes to 5 seconds.

Lead Software Engineer, Cofounder
Thriftax (thriftax.com)

Mar 2020 – May 2024

- Built a tax filing web application for US nonresidents.
- Assisted in over 5100 tax returns and served over 1100 customers since company inception, primarily composed of international students from UCLA, UC Berkeley, and Stanford.
- Official tax filing partner of AuPairCare, one of the largest Au Pair Agencies in the U.S.
- Designed a PDF generation system that reduces the size of files by 95% and latency of file generation by 90% compared to existing open-source solutions on the web.
- Primarily a solo developer on this project, and thus I needed to develop highly efficient abstractions to be able to implement so much tax logic alone. However, we did have two interns who helped out for a brief period and whom I supervised.

Preprints

Multivalent binding model quantifies antibody species from systems serology (2024) Under review at PLOS Computational Biology.

Armaan A. Abraham, Zhixin Cyrillus Tan, Priyanka Shrestha, Emily R. Bozich, Aaron S. Meyer

Integrative, high-resolution analysis of single cells across experimental conditions with PARAFAC2 (2024) Submitted.

Andrew Ramirez, Brian T. Orcutt-Jahns, Sean Pascoe, Armaan A. Abraham, Breanna Remigio, Nathaniel Thomas, Aaron S. Meyer

Projects

Reinforcement learning of human-interpretable, multimodal tasks from AI feedback ([github](#))

Oct 2024

- I conceptualized and implemented an RL framework whereby a multimodal text+image LLM (e.g. Pixtral) is used to provide preferences for states based on a prompt that is specified by text and/or images. These preferences are then used as the reward signal for the agent, which is trained using standard policy gradient methods (in this case TD3).
- I successfully implemented this approach for the classic pendulum swing-up problem. Pairs of images of states of the pendulum environment were provided to the “judge” LLM which then decided which state was closer to the goal. More complex tasks are a work in progress.
- This may be one approach to addressing some of the scalability issues of RL – namely, how do we robustly design reward signals for objectives that are abstract and/or rely on implicit human concepts?

Sparse autoencoders (SAEs) for analyzing neural network features ([github](#))

Oct 2024

- Implemented and trained an SAE on a simple, 2-layer, 6.3M parameter LLM.
- I was inspired by a tension I noticed in our current understanding of superposition in neural networks between interference minimization and downstream output similarity.
- Examined the effect of feature cooccurrence and downstream output influence on feature alignment and superposition.

RL-based generation of peptide sequences conditioned on fuzzy structural constraints ([github](#))

Apr 2024

- I conceptualized and implemented an RL agent that aims to design proteins satisfying some user-specified structural constraint.
- In this process, the RL agent generates an amino acid sequence incrementally, with its reward signal computed based on how well the structure (from AlphaFold2) of the current prototypical sequence (under construction) meets a user-specified structural design criteria after each incremental generation step.
- The first type of structural constraint I allowed in my implementation was the specification of a desired distance, in Å, between two residues at arbitrary positions in the sequence.
 - This is a work in progress, as although the model training is implemented, the agent still performs poorly.

Sequential Reptile for graph neural network (GNN) transfer learning

Jun 2022

- Adapted neural executor of graph algorithms (from DeepMind) with Sequential Reptile curriculum learning algorithm to improve transfer learning between different graph algorithms (e.g. BFS, Bellman-Ford).
- Completed as part of graduate deep learning class focused on GNNs.
- Outperformed baseline model from original DeepMind paper on key metrics such as MSE, predecessor accuracy, and termination accuracy.

Other

- Implemented GPT-2 from scratch ([github](#)) (Oct 2024).
- Implemented an adaptive computation time (ACT) long short-term memory (LSTM) neural network (see “Adaptive Computation Time for Recurrent Neural Networks” by Graves) to fit fractals (Koch curves) and investigated computation time vs fractal recursion depth (Sep 2024).

Presentations

Microscopists and Modelers (UCLA), Invited oral presentation Oct 2024
“Unaligned Low-rank Tensor Regression with Attention (ULTRA): Generating mechanistic explanations of external phenotypes from scRNA-seq”

Systems Biology for Infectious Diseases Annual Meeting, Selected poster presentation Sep 2024
“Multivalent binding model quantifies antibody species from systems serology”

ImmunologyLA, Selected poster presentation May 2024
“Multivalent binding model quantifies antibody species from systems serology”

Awards / Honors

Highest Honors in Biophysics and Computer Science at UCLA; Upsilon Pi Epsilon International Honor Society Member; Bronze Governor General’s Award (national award for secondary students given on behalf of the Governor General of Canada for highest graduating GPA); University of Toronto national book award (awarded to outstanding high school students across Canada)