Naveen Raman, University of Maryland
Proposed Program of Study

Personal Statement

Improvements in artificial intelligence (AI) and machine learning (ML) algorithms both excite and scare me. On one hand are applications like the kidney exchange program at Facebook, which uses big data to deliver life-changing results, and on the other are racial bias issues surrounding the Correctional Offender Management Profiling for Alternative Sanctions (COMPAS) recidivism algorithm.

My research explored both sides of this dichotomy, and through prior research, I became focused on using intelligent computing to improve the utility and fairness of human systems. My work on ride-pooling fairness at the University of Maryland, online toxicity at Carnegie Mellon, and cancer signatures at Maryland led to a Goldwater Scholarship, and this summer, I worked at MIT Lincoln Labs to improve human-AI collaboration. In graduate school, I intend to study applications of AI and ML for social good in fields such as healthcare and social networks.

I was first introduced to ML through my work with Profs. Max Leiserson and Aravind Srinivasan during my freshman year, where I used non-negative matrix factorization (NMF) algorithms to understand cancer mutations. We used NMF algorithms to discover the mutational signatures present in different cancers. The problem's underspecificty made it difficult to retrieve signatures so we incorporated additional information through graph-based regularization. I used PCA to understand the underlying manifold and tested the effect of patient and cellular data-based graphs, building on manifold information. To assess signature quality, I developed an imputation-based metric that uses cross-validation to determine reconstruction error. The project taught me the fundamentals of ML and excited me about its computational possibilities.

To further explore ML, I participated in the Research Experience for Undergraduates on Software Engineering (REUSE) program at Carnegie-Mellon University (CMU) during the summer after freshman year (2019). I worked at CMU for ten weeks with Professors Bogdan Vasilescu and Christian Kaestner, explored applications of ML to improve software engineering communities, and developed an SVM classifier to detect rudeness and toxicity in software discussions. I first developed a baseline classifier, but it performed poorly due to programming jargon being perceived as toxic. To fix this, I worked with CMU Professor Yulia Tsvetkov to develop a log-odds with Dirichlet prior method that accounted for jargon, improving the F-score of the classifier. The improved classifier helped find trends in community toxicity such as which programming language communities tend to be toxic. We published our findings at *ICSE 2020*, giving me the writing and presentation practice necessary for high-impact research.

I enjoyed the human aspect of software engineering research and desired to explore the full spectrum of human-facing ML research, embarking on a new project with UMD Prof. Jordan Boyd-Graber to improve named entity linking (NEL) algorithms starting fall 2019. To do this, I developed a dataset for the noun phrase linking task, which expands NEL to include anaphoric or indirect references. I developed a baseline noun phrase system that combines a coreference system with entity linking, and using the baseline system, I developed a React-based web app that prompts users to correct mistakes, leading to an improved dataset. To test the impact of noun phrase linking on downstream tasks, I set up a question answering system that relied on entity links. We presented our work at the *MASC-SLL conference 2020*, *HAMLETS workshop at NeurIPS 2021*, and plan to submit our work to *TACL* once data collection is complete.

Simultaneous with my project on entity linking, I discovered my ideal combination of theory and practice through research with Professor John Dickerson on ride-pooling matching algorithms. We developed fairer matching algorithms, which ensure equitable inter-neighborhood service and income inequality minimization. To incorporate fairness, I generalized a state-of-the-art ride-pooling matching algorithm to work for any value function and developed new objective functions. To reduce income inequality, I proposed an income redistribution scheme based on the Shapley values of drivers, which reduced income disparity while incentivizing drivers to maximize effort. We published our work at the *ML For Econ Policy NeurIPS Workshop 2020*, *Undergrad Consortium at AAAI 2021*, and *IJCAI 2021*.

To extend my work on rideshare fairness, I proposed a new project that improved the fairness of rideshare pricing algorithms by incorporating network externalities. Accounting for these externalities, which arise when one agent's utility impacts another, allows for a holistic assessment of prices while discouraging discriminatory pricing. I developed a baseline dynamic programming-based solution and extended this to situations where network externalities are non-linear by combining MDPs with approximate dynamic programming. Doing so gave me practice deploying AI algorithms to find approximate solutions. Our work is under submission at the *AAAI Student Abstract Program 2022.*

After internships in academia (CMU) and industry (software engineering at Facebook), I found a middle ground that combined academic freedom and real-world applicability by working at the MIT Lincoln Labs, a defense lab, this summer. I worked to improve learning-to-defer algorithms, which partition tasks between AI and humans, by incorporating fine-tuning. We used semi-supervised learning to develop fine-tuning algorithms which improve human performance models. I tested them using both a synthetic dataset based on autonomous driving and an image recognition dataset. For the synthetic test, I showed that fine-tuning brings driving time 20% closer to optimal, while for image recognition, I found that fine-tuning helps under in some situations. We published our work at the *Human-Machine Decisions Workshop at NeurIPS 2021*.

In graduate school, I plan to study how AI and ML systems can work for social good through applications to healthcare and social networks. ML has the potential to revolutionize healthcare through improvements in clinical prognosis. Predicting future patient outcomes and diseases allows for optimal preventive care, creating the need for algorithms that predict future events from time-series data. This problem is technically difficult due to potentially different train and test distributions and is especially challenging for patients from marginalized communities due to data sparsity and bias. I plan to combat these problems by developing robust learning algorithms that work in the presence of data perturbations, which I can do through modifying objective functions similar to my work on rideshare, allowing for application even when train and test distributions differ. Additionally, I aim to ensure fairness by minimizing error rates across groups and collecting data from marginalized groups, leading to improved prognosis.

Combating healthcare crises involves deciding treatments and interventions in situations where information might be limited. An example of this is spreading HIV prevention information amongst at-risk youth; both problem modeling and information gathering are fraught with difficulty due to uncertainty about network structure. This uncertainty can be tackled through ML, which can learn network structure from partial information, and through game theory, which can decide which nodes or people are important for information diffusion. I plan to improve these algorithms by applying ML and AI techniques, such as transformer-based deep learning models, for other social network-related problems, such as predicting applying interventions for disease spread. To achieve fairness in these situations, I plan on utilizing methods that can ensure interventions are thoroughly tested for all groups before deployment.

After graduate school, I plan to become a professor so I can research applications of ML and also teach. I want my research to have a positive impact, so I plan to focus on social problems such as poverty reduction and healthcare, while working with local organizations.

Proposed Program of Study

User data is critical for machine learning (ML) and artificial intelligence (AI) applications such as targeted advertising and language modeling, leading to the creation of data markets where user data can be bought and sold. These markets largely benefit data trading corporations that own user data, as consumers lack information or control of how their data is used and whether it is sold to third-party companies. This leads to situations where user data is used for causes conflicting with personal preferences, highlighting the lack of control over data usage. While some uses benefit consumers - for example, data sold to medical companies help develop treatments - consumers lack adequate compensation and control over their data.

Consumers need privacy-preserving data markets where individuals are compensated for data sales and privacy preferences are respected. However, developing such a market is difficult as real-life privacy preference data is scarce, and price dependencies between consumers, where the sale of one person's data can lower the value of others' data, makes calculating optimal prices infeasible. While difficult, developing privacy-preserving data markets has enormous social benefits, as it allows users to take control of their data.

Professor Jon Crowcroft's research on privacy preservation and data trading leverages mechanism design and behavioral economics to create privacy-preserving data markets. He analyzed the effects of network externalities, which arise from correlations between users, and concluded that these externalities result in a "race to the bottom" as user data valuations go to zero, hindering the development of privacy-preserving markets (Pal et al.). To counter this, Professor Crowcroft determines prices while maintaining group privacy, finding Nash Equilibria under both a competitive market and an oligopolistic one (Pal et al.).

Two areas for future exploration are realistic privacy preferences and data heterogeneity. In "Preference-Based Privacy Markets," Professor Crowcroft asked whether we can collect realistic human privacy preferences to supplant synthetic data. I propose addressing this question by collecting user preferences for privacy compensation, which identifies user opinions on adequate compensation for privacy loss for different data types. To collect this data, I would create an interface that asks users to quantify their privacy preferences. Using collected data, I would fit a quasilinear utility function for each data type that outputs utility based on privacy and compensation and uses these functions to aggregate user preferences.

Incorporating heterogeneous privacy valuations for different data types requires both a compensation vector of valuations and a matrix representing correlations between data types. These correlations indicate that the sale of one type of data changes the value of another type. One method to determine data prices is to price data types independently. However, this ignores correlations between data types that could lead to market inefficiencies. For example, users might want to increase the value of location data if health data is sold. To counter this, we can iteratively develop approximately optimal prices using reinforcement learning techniques like Markov decision processes. Additionally, heterogeneous user preferences could impact fairness in data market pricing, leading to the need for ensuring fairness through constraints on pricing.

I aim to construct new mechanisms to facilitate privacy markets. I can build on the material I learned from my mechanism design class by using Stackelberg games to construct markets. My experience working with fairness in rideshare gives me background working with markets and pricing, and can also assist with the development of constraints to ensure fairness. I aim to publish two papers for each of the two ideas, the first in a human-computer interaction conference such as CHI, and the second in a journal such as IEEE. By working with Professor Crowcroft, I can learn more about data markets, and combat privacy-related issues.