Incorporating Autonomous Vehicles into Rideshare Matching Algorithms

Motivation

Autonomous vehicles use artificial intelligence (AI) to operate with little to no human intervention, potentially saving many lives from car accidents once widely deployed. These vehicles can not only revolutionize personal vehicles, but also revolutionize rideshare fleets, as integrating autonomous vehicles into these fleets improves system performance and increases access to transportation while decreasing the carbon footprint of automobiles.

Integrating autonomous vehicles into rideshare fleets is a technical challenge as it requires modifying current matching algorithms, which determine rider-driver pairs. For example, risk-averse consumers might prefer not to be serviced by autonomous vehicles, and so should not be matched with autonomous vehicles, especially in situations where autonomous vehicles do not perform as well as humans. Matching algorithms need to respect both autonomous vehicle abilities and rider preferences by matching these hesitant riders with human drivers, even if trips take longer. Investigations into improving matching algorithms for rideshare fleets have not been previously explored and can be especially impactful in the future when autonomous vehicles become dominant. Because of this, I propose improving matching algorithms for autonomous rideshare vehicles so they take advantage of human-AI collaboration while respecting rider hesitancy about AI-driven vehicles.

Prior Work

My experience working on rideshare matching algorithms and human-AI collaboration uniquely positions me to work on autonomous vehicle matching algorithms. Matching in rideshare is typically done through a linear program, which finds optimal rider-driver pairs by maximizing a linear objective function subject to constraints. I previously worked to improve the fairness of matching algorithms by developing new fairness-based objective functions for ride-pooling [1], and I can adopt similar techniques to modify matching programs for autonomous vehicles.

Additionally, I can build on my research at Lincoln Labs on human-AI collaboration for autonomous vehicles. My research improved learning-to-defer algorithms, which partition tasks between AI and humans, by using semi-supervised learning to fine-tune for specific people. To test my improved algorithm, I applied it to a synthetic autonomous driving dataset, to see if improved learning-to-defer algorithms lead to quicker rides. I can apply this research to autonomous vehicles to determine when autonomous vehicles outperform humans, which can be incorporated into the matching algorithm.

Research Plan

I plan to incorporate my prior work to design matching algorithms for autonomous vehicles, guiding my research with two questions, both of which modify the linear program for matching.

RQ1: How can we incorporate rider preferences about autonomous vehicles?

Technological hesitancy leads some riders to refuse rides on autonomous vehicles. Serving these riders is important from a fairness perspective and requires adjusting matching policies, which entails changes to the objective function or constraints for linear programs. Constraints on edges can prune matches between hesitant riders and autonomous drivers, but increases the problem's computational complexity and fails to account for non-binary hesitancy values. Non-binary hesitancy values can arise if riders wish to ride in autonomous vehicles only when they are significantly faster or safer than humans. To account for this,, the objective function can be modified to reward matches between hesitant riders and human drivers, weighted by rider reluctance, thereby allowing for non-binary values. In situations where rider preferences are non-explicit, preferences can be learned based on differences in trip ratings between rides driven by autonomous and humans drivers.

To ensure that matches maximize the objective function for the long run, we employ techniques such as Markov decision processes (MDPs), which can estimate future rewards. The value function of an MDP is used to estimate the long-term reward in a state, and values for states from the value function are used with reward values to set weights for the linear program, thereby maximizing long-term reward. Due to the complexity of rider-location combinations, we use approximate dynamic programming techniques and deep learning to approximate the value function [2].

RQ2: How can we effectively deploy AI to situations where humans perform poorly? Assessments of human and AI competency can help develop learning-to-defer algorithms, which determine whether to deploy humans or AI based on conditions like rain and darkness. The output from learning-to-defer algorithms can be combined with hesitancy data in the linear program weights, so more hesitant or risk-averse riders only ride with AI when they significantly outperform humans, even if there's a significant time difference between the two. Developing learning-to-defer algorithms requires labeled data and improved learning algorithms.

To collect new labeled data, I plan to set up studies that assess human driving under various conditions using speed and safety metrics. These studies observe both humans and AI driving a simulated car, assessing performance through safety metrics, which include braking times, lane departure frequency, and disengagement rate (a measure of how often humans need to intervene for AI). Storing this data is difficult due to its high dimensionality, as capturing environmental conditions requires images and weather information, so I would use compression algorithms.

Previous learning-to-defer algorithms used a ResNet-based deep learning convolutional model to learn when to defer to humans [3]. Strategies to improve these models include the use of downsampling and the incorporation of other deep learning models, such as transformers. I plan to incorporate my work on fine-tuning learning-to-defer algorithms, which account for the particular drivers available using techniques like self-training. These algorithms can be improved through more sophisticated semi-supervised learning techniques, which take advantage of both labeled and unlabeled data, such as incorporating entropy regularization into the loss function. To determine environmental features that make driving difficult, I can employ an item response theory (IRT) model, which explains the relationship between environmental conditions and driving difficulty.

Intellectual Merit

This proposal makes matching algorithms more sensitive to client-side information, allowing them to be better suited for clients (riders). Additionally, improvements to learning-to-defer algorithms improve human-AI collaboration in large-data domains, which has seldom been tackled before and becomes more important due to increases in data availability. This proposal additionally improves algorithms for semi-supervised learning by combining them with fine-tuning, improving system performance and specificity. This is done by developing new objective functions and algorithms, such as combining an entropy regularization term with fine-tuning, which can be applied to other fields.

Broader Impacts

To increase the publicity of research and bring about immediate change, this research will be presented at the Automated Vehicles Symposium run by the National Highway Safety and Traffic Administration (NHSTA). The symposium allows for engagement with traffic experts about designing safer autonomous vehicle-rideshare systems, saving lives by preventing crashes.

Outside of autonomous vehicles, this research can be applied to other matching problems. For example, this research can be extended to assist the National Center for Education Statistics (NCES) by ensuring the proper use of AI in school choice. Research on AI-hesitant individuals can be used to ensure that AI-hesitant parents are not excluded by school choice decisions, creating a more inclusive process. This research can also be used in the medical domain to improve information incorporation for kidney exchange [4] and residency matching [5]. To publicize this research, I could engage in collaborations with the Facebook kidney exchange program, which determines optimal matches between kidney donors and recipients.

References

[1] Raman, Naveen, et al. "Data-Driven Methods for Balancing Fairness and Efficiency in Ride-Pooling." (2020). [2] Shah, Sanket et al. "Neural approximate dynamic programming for on-demand ride-pooling." AAAI. Vol. 34. [3] Mozannar, Hussein et al. "Consistent estimators for learning to defer to an expert."ICML. PMLR, 2020. [4] Dickerson, John P., et al. "Failure-aware kidney exchange." Proceedings of the fourteenth ACM conference on Electronic commerce. 2013. [5] Roth, Alvin E. et al. "The effects of the change in the NRMP matching algorithm." JAMA 278.9 (1997): 729-732.

I have always been interested in stories. Ranging from historical anecdotes to news stories, I found that stories can humanize ideas and elicit empathy, making distant ideas seem close. My interest in computer science came later and was rooted in my fascination with the ability of algorithms to model complex tasks. Although interesting, work in computer science felt removed from reality, as computer programs seemed distant from humans. My perception changed after reading stories about the impacts of Artificial Intelligence (AI) on real-world issues, both positive and negative. 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. These stories convinced me of the momentous impacts AI can have on real people and real social problems, for better or worse, and drove me to develop a focus on researching approaches to AI that account for social good.

This drive has led to projects on a variety of topics throughout my undergraduate career, including work on ride-pooling fairness at the University of Maryland, open-source toxicity at Carnegie Mellon, and human-AI collaboration at MIT Lincoln Labs. By engaging with a variety of Machine Learning (ML)-based projects, I learned that I enjoy public-facing research that combines theory and application. Developing and understanding AI and ML algorithms requires an interdisciplinary approach, and to this end, I've taken classes on education policy, natural resource economics, and ethics. As a computer science PhD student, I plan to continue researching uses for AI and ML, while finding methods to minimize their bias, such as researching AI algorithms and bias mitigations strategies for healthcare.

Graduate Study Goals: I aspire to study public-facing applications of ML, with an emphasis on social good. Examples include the use of ML for better clinical prognosis across racial groups and investigations into the use of ML in criminal justice applications. I would be excited to pursue these directions at graduate schools such as MIT and Harvard. At MIT, I would love to join the Clinical Machine Learning lab, which researches how ML algorithms can improve clinical prognosis while ensuring the fairness of those algorithms. I would also be excited to join Teamcore at Harvard, which uses AI to design public health interventions. My goal is to produce impactful research during my PhD, and to this end, I plan to work with organizations to assist with their use of ML and ensure its equitable use. Examples include working with the World Wildlife Federation to develop ML algorithms that position security guards to minimize poaching or working with Safe Place for Youth to develop interventions that distribute HIV prevention information. After completing my PhD, I plan on becoming a professor, allowing me the freedom to research, while also being able to teach.

Undergraduate Research Experiences: (1) Cancer Mutations: I was first introduced to ML through my work with Professors Max Leiserson and Aravind Srinivasan, where I used non-negative matrix factorization (NMF) algorithms to understand cancer. We used NMF algorithms to discover the mutational signatures present in cancers, but the problem's underspecificity made it difficult to retrieve high-quality signatures. To combat this, we introduced graph regularization, which enforces similarity constraints between patient signatures. I incorporated patient and cellular graphs to retrieve signatures and I developed an imputation-based metric to assess signature quality. The project allowed me to learn the fundamentals of ML and exciting me about its possibilities. **Broader Impact:** Higher quality signatures through NMF allows for better identification of factors contributing to cancer.

(2) Toxicity in Software Engineering: After my initial research experience, I was curious about uses for ML, and so I participated in the Research Experience for Undergraduates on Software Engineering (REUSE) program at Carnegie-Mellon University (CMU) during the summer after my freshman year (2019). I explored applications of ML to improve software engineering communities, working with Professors Bogdan Vasilescu and Christian Kaestner to develop a Python-based classifier to detect rudeness and toxicity in software discussions. I first analyzed a toxic language dataset and developed a baseline classifier, but I found that 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 classifier's F-score. I used the classifier to find trends in toxicity, such as finding the most toxic programming language communities, and I published and presented a first-author paper at *ICSE 2020*. **Broader Impact:** The toxicity classifier serves as a diagnostic tool, allowing developers to take mitigation steps if their communities are toxic.

(3) Entity Linking: I enjoyed the human aspect of software engineering research and desired to explore applied ML research, embarking on a new project with Prof. Jordan Boyd-Graber to improve named entity linking (NEL) algorithms through data collection. 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 this, I developed a React-based web app that prompts users to correct mistakes, leading to an improved dataset. We presented our work at the *MASC-SLL conference 2020*, *HAMLETS workshop at NeurIPS 2020*, and plan to submit our work to *TACL* once data collection is complete. The project taught me how to develop user interfaces for human-facing studies and improved my web development and data collection skills. **Broader Impact:** We plan to release our dataset after completion so other developers can use noun phrase linking to improve downstream tasks, such as question answering and dialogue.

(4) Fairness in Ride-pooling: I discovered my ideal combination of theory and practice through research with Professor John Dickerson and Sanket Shah, a collaborator at Harvard, on ride-pooling matching algorithms. We aimed to develop fairer matching algorithms, which ensure equitable inter-neighborhood service and income inequality minimization. To incorporate fairness, I first generalized a state-of-the-art ride-pooling matching algorithm for any objective function and developed new objective functions which combined utility and fairness. This increased pickup rates for riders in underserved neighborhoods while minimally impacting profit. To reduce income inequality, I proposed an income redistribution scheme based on the Shapley values of drivers, reducing income disparity while still incentivizing drivers to maximize effort. This led to first-author papers at *ML For Econ Policy NeurIPS Workshop 2020*, *Undergrad Consortium at AAAI 2021*, and *IJCAI 2021*. The project showed me the potential of ML for social good, fueling my desire to work on fairness-related problems in graduate school. **Broader Impact:** Our work incorporates fairness into matching algorithms and applies to other domains, such as food delivery.

(5) Network externalities in rideshare pricing: After the conclusion of the fairness project, I worked with Prof. Dickerson to incorporate network externalities into rideshare pricing algorithms. Rideshare firms use dynamic pricing algorithms to vary trip prices in response to demand, but fail to account for network externalities, which can result when the utility for one agent impacts the utility of other agents. I worked with Prof. Dickerson to formalize the problem, then proved the optimality of a dynamic programming algorithm for computing prices under linear utility propagation. I generalized the problem so utility propagation is non-linear and developed an MDP-based solution using approximate dynamic programming and deep learning, which allows for value function estimation and is needed due to the large state space. Our work is under submission at the *AAAI Student Abstract Program 2022.* **Broader Impact:** Our work tackles network externality incorporation in light of fluctuating demand; this feature is not peculiar to rideshare, and can be used for other dynamic pricing algorithms, such as those used by AirBNB for determining room prices.

(6) Human-AI Collaboration: This summer (2021), I worked at the MIT Lincoln Lab with Dr. Michael Yee to improve human-AI collaboration. Learning-to-defer algorithms partition tasks between AI and humans, but are based on aggregate human performance rather than fine-tuned for individuals. To address this, I developed a fine-tuning algorithm using self-training, which improves deep learning models of human performance. I tested these algorithms using both a synthetic dataset based on autonomous driving and the CIFAR10 image recognition dataset. For the synthetic test, I showed that fine-tuning brings driving time 20% closer to optimal, while for the image recognition dataset, I found that fine-tuning improves system performance under certain circumstances. Our work is under submission at the *Human-Machine Decisions Workshop at NeurIPS 2021*. **Broader Impact:** Improved human-AI collaboration improves autonomous vehicles' performance, saving lives through better and safer system performance.

Intellectual Merit: In graduate school, I aim to improve AI algorithms for social good through applications of MDP and multi-armed bandit (MAB)-based simulations to fields such as healthcare and social networks. For these fields, I can improve time-series algorithms, which are used in healthcare to predict patient outcomes, through more robust learning algorithms. Fairness is also important for time-series and other healthcare algorithms, as these algorithms need to give high-quality predictions across demographics groups. In graduate school, I plan to improve the fairness of these algorithms by improving their interpretability, so that the underlying fairness-related reasons can be investigated. I can combine healthcare with social networks to develop graph-based learning algorithms, which can be applied to study disease propagation amongst communities. My research experiences have prepared me to study applications of AI and ML: my work at CMU and Lincoln Labs taught me how to develop classification and regression models, while my experiences working on rideshare have taught me how to develop stochastic simulations using MDPs.

Broader Impacts - Undergraduate - (1) College Park Academy: I enjoy teaching, as it allows me to share my passion for learning with students, while also impacting my community. To pursue this passion, I worked at the College Park Academy charter school and ran a study club where we assisted students with math homework. The program allowed me to introduce potential careers in computer science and STEM to under-represented groups and allowed me to talk with them about why I study computer science.

(2) Maryland Mentor Program: I additionally work with the Maryland Mentor Program to teach reading to English second-language elementary students, setting up phonics activities to help them blend letters. The program allows me to discuss educational issues with other mentors, gaining perspectives on the diversity of educational situations across the country.

(3) Teaching Assistant: At Maryland, I teach two classes: one on programming languages and another on coding interviews. For the latter, as the head teaching assistant, I plan the curriculum and serve as the head teaching assistant, managing class logistics. My favorite part of both classes is holding discussions, as I get to discuss topics I am passionate about and help clarify concepts for students.

Graduate Broader Impact Goals: I intend to continue my passion for teaching by working with organizations such as StreetCode at Stanford and Middle School Bridge at CMU. StreetCode works with communities of color to teach basic coding skills and Middle School Bridge assists high school freshmen with math. These programs allow me to practice teaching and also give back to my local community. In addition to teaching, I want my research to improve societal issues, such as healthcare and criminal justice, and to improve ML fairness.

Future Scientific and Broader Impacts Goals: My research on AI for social good relates to other fields within computer science. My work on fairness in AI most directly relates to work in natural language processing (NLP) on the bias of language models. As a result, I could engage in collaborations with NLP researchers to address concerns over fairness in AI and NLP. This could additionally relate to my interest in using AI for healthcare, as I could work on biased language in healthcare.

AI is additionally connected with ethics, particularly data privacy, and so I am interested in engaging in collaborations with computer security researchers to develop systems that allow users to maintain privacy. One example of this is data markets, which combine game theory, economics, and learning algorithms so agents can buy and sell data. My previous experience working at the intersection of economics and AI has given me the skills to improve these markets by developing mechanisms so privacy preservation is incentivized. In general, my previous research experiences with AI and ML, and my interdisciplinary background, give me the tools to perform high-impact research in graduate school.

List any significant academic honors, fellowships, scholarships, publications and presentations. (less than 16000 characters total).

Scholarships
Goldwater Scholarship - Prestigious scholarship given to top undergraduate researchers nationally
Philip Merill Presidential Scholarship - Given to top 15 seniors for academic and service record
Iribe Scholarship - Computer Science scholarship worth 11K for academic and research record, given for 2 years (11K each year)
Presidential Scholarship - 4 year scholarship worth 20K given to rising freshmen
Capital One Scholarship - Computer Science scholarship worth 1K for academic and research record
Corporate Partners Scholarship - Computer Science scholarship worth 2K for academic and research record

Honors/Fellowships
CMU REUSE REU Fellowship - One of twenty students accepted into ten week research program at CMU, where I researched toxicity in online communities
Global Fellows Program - Accepted into an honors program that explores international and policy issues. Currently taking a class on science diplomacy, and will work a part-time internship in the government next semester.
ACES Honors College - Part of the cyber security-based honors college, where I took three courses on cybersecurity, Linux, and bash-scripting, and two courses on reverse engineering and cybersecurity accounting. For one class, developed cyber security-based research project which explored the impact of file obfuscation on attackers' ability to retrieve information. s
CS Honors College - Accepted into research-based honors college, where I took classes that taught me how to conduct research and got the opportunity to create a senior thesis
Math Modelling - Competed in SCUDEM math modeling competition in a group of three, where I was tasked to develop a model of interactions between refugee populations in migrant camps. Received the outstanding award for the paper.
Bloomberg Codecon - Competed in the local Bloomberg Codecon challenge, a programming competition at the University of Maryland, and placed in the top-3. As a result, qualified to compete in national codecon at Bloomberg headquarters in New York City.

Publications

1. Data-Driven Methods for Balancing Fairness and Efficiency in Ride-Pooling at IJCAI 2021
2. Investigating methods of balancing inequality and efficiency in Ride Pooling at AAAI Undergraduate Consortium 2021
3. What more can entity linking do for Question Answering at HAMLETS NeurIPS 2020 Workshop
4. Eliciting Bias in Question Answering Models through Ambiguity at MRQA EMNLP 2021 Workshop
5. Stress and Burnout in Open Source: Toward Finding, Understanding, and Mitigating Unhealthy Interactions at ICSE 2020
6. A Muffin-Theorem Generator at Fun with Algorithms (FUN) 2018

Presentations
Improving Entity Linking through Quizbowl - MASC-SLL 2020

Under Submission
Paper at Workshop on Human-Machine Decisions, NeurIPS 2021
Paper at AAAI Student Abstract, 2022