Short Questions (~300 words)

How did you choose your field and what are the primary expectations for your future career?

The rapid proliferation of artificial intelligence (AI) has led to developments ranging from automated drug development to mental health chatbots, but AI algorithms still struggle to make decisions in multi-agent scenarios, such as automatic matching for rideshare applications. These problems are difficult for AI because agents can act unpredictably or irrationally, causing them to act in globally suboptimal ways.

I was introduced to multi-agent decision-making through two projects at the University of Maryland (UMD) and MIT Lincoln Labs. My first project was with Professor John Dickerson at UMD, where I developed algorithms that match riders and drivers for ride-sharing apps such as Uber. The project taught me about difficulties in multi-agent decision-making, including computational limitations, as the space of rider-driver combinations is exponentially large, and uncertainty, as rider trip requests arrive in real-time making future planning difficult. At the MIT Lincoln Labs, I researched human-AI teaming for learning-to-defer algorithms, which determine whether humans or AI are better for a task. Through this project, I learned about existing issues in multi-agent cooperation, such as conflicting expertise, and about statistical approaches to modeling human-AI interactions.

These projects have led to my current work on matching food rescue volunteers with trips. Food rescue organizations transport excess food from sources such as restaurants and grocery stores to locations such as food banks using volunteers. Determining which volunteers to notify for these trips is challenging, as frequent notifications lead to burnout while scarce notifications lead to food wastage. This project captures the essence of questions I'm interested in tackling including decision-making under imperfect information, and competition between rational agents. Techniques I develop for this problem, such as new variants of bandit algorithms, can be applied to other similar problems, such as organ allocation and residency matching. In addition to performing research on AI decision-making, I'm interested in teaching and mentoring, as outreach is a necessity for any scientist. In the future, I aim to become a professor, so I can combine impactful research with teaching future AI developers the strengths and limitations of AI.

How does your proposed field of study and career constitute an application of the physical sciences or engineering?

Machine learning is a hot dog: just as a hot dog requires the bun, meat, and mustard, machine learning requires statistics, optimization, and engineering. Each piece is necessary for the integrity of the metaphorical hot dog, and only with the combination of the three can algorithms be developed. At its core, machine learning uses statistics to find patterns in data. Statistics captures the relationship between features and labels, and many machine learning algorithms can be viewed from a statistical perspective. Additionally, statistical ideas from learning theory are used to provide guarantees on how many samples are needed so these algorithms converge. Optimization is used within machine learning to find optima for some objective function, which is used as a metric for the performance of a model. Objective functions are maximized through algorithms from optimization, such as gradient descent or simulated annealing. Finally, engineering makes this process usable in the real world and smooths out technical complexities. For example, while theoretical machine learning suggests that non-convex models should not find good minima, engineering efforts develop variants of gradient descent that find these optima despite these limitations.

Multi-agent scenarios excite me because of their interdisciplinary nature, requiring not just machine learning techniques, but also methods from game theory, which is used in combination with optimization and reinforcement learning. For example, one question that arises in multi-agent collaboration scenarios is whether we can design mechanisms so that bad actors don't ruin cooperation, that is, whether various equilibria correspond with social welfare. Finding such equilibria is difficult, as finding Nash Equilibria is a computationally difficult problem, so variants of linear programming are used instead. Finally, game theory is used in combination with reinforcement learning to compute optimal actions for agents under uncertainty.

What are the considerations involved in your choice of graduate school?

Attending Carnegie Mellon University (CMU) allows me to work on a variety of applications and methodologies to better understand how decision-making algorithms perform in multi-agent scenarios. My advisor, Fei Fang, specializes in decision-making and game theory in multi-agent scenarios and has deployed algorithms to domains including food rescue and wildlife conservation. Working with her allows me to both study new methodologies for multi-agent decision-making and work with real-world partners so these algorithms can be deployed.

Beyond my advisor, CMU features many professors who work on AI decision-making. I intend to collaborate with Nihar Shah, who works on reviewer dynamics in peer review systems. Peer review provides an opportunity where decision-making algorithms can be used to pair up reviewers with papers and to develop mechanisms so reviewer collusion is disincentivized. Being at CMU also allows me to work with experts in game theory, such as Tuomas Sandholm and Vincent Conitzer, with whom I intend to study cooperation and competition in multi-agent scenarios from a mathematical lens. For example, I could study how various equilibria arise in real-world scenarios, and design algorithms that take advantage of these equilibria.

CMU features experts across other departments who work on multi-agent decision-making using AI. Examples of collaborations I'm interested in pursuing include working with Rema Padman, who works on real-world healthcare decision-making algorithms. I'm interested in joining her ongoing collaboration with the School of Computer Science, which works to identify optimal matches in organ donation, leveraging both healthcare knowledge and AI decision-making. CMU presents the opportunity to work with experts on multi-agent decision-making, both across the theory-practice spectrum and for a variety of applications.

Finally, choose one or two projects that best exemplify your creativity and discuss in more detail what you contributed to them.

Human-AI Collaboration: In the summer of 2021, I worked with Dr. Michael Yee at the MIT Lincoln Labs to improve human-AI collaboration through better learning-to-defer algorithms. Learning-to-defer algorithms determine whether humans or AI should solve a task, but are based on aggregate human performance data rather than fine-tuned for individual humans. Improved learning-to-defer algorithms can lead to better system performance for technologies like autonomous vehicles, saving lives from car accidents. I developed the project from scratch by reading through prior literature and I additionally reached out to other researchers at Lincoln Labs to discuss potential avenues for the project. To start the project, I developed a fine-tuning algorithm using self-training, which leverages unlabeled data for training and is necessary due to labeled data sparsity. For evaluation, I developed a synthetic dataset based on autonomous driving and showed that fine-tuning reduces driving time. I additionally applied my algorithm to deference in the CIFAR10 image classification task by analyzing human performance using dimensionality reduction algorithms such as PCA. I then used deep ResNet models to predict human performance and extended prior work by combining two fine-tuning algorithms with a modified cross-entropy loss function. Our work was published at the *Human-Machine Decisions Workshop at NeurIPS 2021*.

Understanding Concept Bottleneck Models: In 2023, I worked with Professor Mateja Jamnik at the University of Cambridge to understand the properties of interpretable machine learning models. Machine learning models are typically black boxes, making it hard to understand "why" a particular prediction was made. Concept-based interpretability models supposedly fix this issue by making predictions through an intermediate layer that predicts concepts, which extract human-understandable "concepts" from model predictions, however, it was unclear whether these models preserve localities or invariances found in tasks. Models that fail to preserve these localities are unstable, as small perturbations in the input can change predictions drastically. To better understand this phenomenon, I first developed a set of synthetic datasets, which had spatially localized features so only a fraction of the input features were relevant for concept prediction. Using this dataset, I found that models use irrelevant portions of an image to make predictions, so model reasoning fails to match human intuition. I studied this performance across factors like optimizer choice, dataset noise, and model architecture, and found that model size was primarily responsible. To understand why this occurs, I applied a pruning algorithm using layer and weight pruning to investigate if artificially smaller models solve this problem and found that pruned models degrade in accuracy before achieving concept independence. I confirmed these trends apply to real-world datasets by developing an image masking-based experiment, where I showed that masking irrelevant regions still impacted model predictions. Our work demonstrates the instability in current interpretable models and highlights the

need for future work that develops more robust and interpretable models. We submitted our work to the *XAI in Action Workshop at NeurIPs 2023*, and it's currently under review.

Use this page to submit a personal essay that reveals your creativity beyond your research. You may choose to focus on activities since your graduation from high school, but these do not have to be related to the classroom or lab, to STEM, or your future career—they can also be from your more informal pursuits and passions, e.g., your extracurricular and leisure activities. This is an opportunity for you to share your narrative with experiences that demonstrate your creativity, perseverance, resiliency, curiosity, and/or grit. You may use more than the allotted space for your response, and any additional text will be included in your application. (~2 pages?)

Teaching excites me because of the freedom it gives in how material is presented along with the relationships formed with students over common interests. The freedom present in teaching is similar to the freedom present in story-telling. While story-telling is the art of arranging characters and dialogue in a convincing yet engaging manner, teaching is the act of arranging information in an insightful yet interesting manner. Beyond engagement, teaching needs to reveal the larger trends and ideas that might be hidden from students without guided exploration.

One of my first teaching experiences was serving as a teaching assistant for the programming languages course at the University of Maryland (CMSC330). The class typically features a variety of seemingly unrelated concepts, ranging from web security to lambda calculus to context-free grammar. Typically, once-a-week discussion sections revolve around practicing exercises similar to what was seen in class. While this helps students understand the details of how individual languages or abstractions work, they oftentimes fail to grasp the larger questions connecting different ideas, namely the different factors impacting programming language construction and real-world software engineering projects. To address this issue, I redeveloped the curriculum for the weekly discussion sections so that connections between topics were emphasized. To start, I designed the first week of discussion so it revolved around debugging, which allowed students to gain transferable skills usable across projects and topics. My debugging lesson consisted of a series of exercises where students would fix a faulty program using different techniques, such as code tracing and unit testing. Developing such an exercise is difficult because of the bimodal distribution of problem difficulty; students either tend to find bugs quickly or not at all. To get around this, I created problems of varying difficulty that allowed students to build skills incrementally. In successive weeks, I designed lessons and emphasized connections between the material covered previously and the material covered this week. For example, I connected context-free grammars to lexical analyzers and lambda calculus with the functional programming ideas of map and fold. Developing exercises for challenging topics, such as the programming language Rust, was difficult, requiring me to read through documentation and work through examples so I could better understand the language. Doing so allowed me to better communicate the rationale behind

language construction to the students. Finally, I structured the last discussion section as a big wrap-up story, clearly showing the threads that underlie what was taught throughout the semester.

Teaching my class on "technical interviews for software engineering" at the University of Maryland (CMSC389O) was challenging because it was unclear how to break down technical interviews into smaller, teachable topics. To tackle this, we structured the class so it was clear that technical interview questions only revolved around a few topics, such as dynamic programming and graph algorithms. To maximize student practice, we designed a flipped classroom so students can practice mock interviews with their peers during class while learning the material at home. Additionally, midterms and finals were structured as interviews with course facilitators to maximize interview practice. Finally, we developed a lecture on behavioral interviews, presenting this as the act of story-telling, which linked together the different experiences of a resume.

Through my experience teaching a variety of age groups ranging from elementary schoolers to college students, I've learned that each age group presents a different pedagogical challenge. For example, teaching high school students is challenging because of the need to engage them. Oftentimes, these students are compelled by outside factors, such as parents, to join a club or a class, and so lack enthusiasm. I faced this reluctance when teaching cybersecurity at a local high school. My first task was to gamify learning by creating cybersecurity-related competitions so that students would stay engaged and invested in the material. There were competitions between students over who could fix files the quickest or could identify certain pieces of system information. I additionally worked with my peers to teach programming to these students in a visual manner. For example, we showed them how simple computer programs could draw shapes and patterns, which made them curious about how to extend this to other designs.

I discovered that the most difficult part about teaching reading to elementary school students is the need for resilience. Teaching reading is difficult because of its fundamental and abstract nature; explaining why reading is important to someone who cannot read is difficult. In the beginning, my student would understand how to make each letter sound, yet combining these to form words was a challenge. To get around this, I designed games involving rhyming, so that he would only have to change 1 letter to find new words. Despite the creation of new games, I still had to labor through reading many books with him, leading to sporadic progress. I discovered that there's no silver bullet to teaching reading; trial and error is required.

The skills you learn from teaching are directly connected with research. My experience designing curricula for college courses has sharpened my creativity, and working with younger children

has taught me to be resilient. Additionally, teaching has significantly improved my presentation skills, as different presentation styles are needed for each audience. Beyond this, teaching gave me the skills necessary to distill complex ideas into simple phrases, which is essential for research. By practicing this, I've become better at teaching myself new concepts, which is necessary when performing research.

Provide a concise summary, similar to a resume, in chronological order, with dates, recapitulating significant periods of technical and other creative activity since high school graduation. Omit activities only distantly related to your professional development. Include workshops, summer schools, a general description of all courses of study pursued (e.g. "3 quarters of Differential Equations"), and degrees expected or awarded (dates, institutions, fields). Separate your undergraduate activities from your graduate activities (if/as applicable) with a single dashed line.

2018
----September---
Started school at University of Maryland (UMD), Advanced Cybersecurity Experience for Students (ACES) Honors College for cybersecurity
Worked in Aravind Srinivasan and Max Leiserson's lab to discover mutational processes through non-negative matrix factorization

---October---
Attended StringBio conference in Orlando on string algorithms for computational biology

2019
---January---
Started as a teaching assistant (TA) for a class on web programming, CMSC122

---May to August----
Worked at Carnegie Mellon University for the Research Experience for Undergraduates in Software Engineering (REUSE) program with Bogdan Vasilescu and Christian Kaestner on detecting online toxicity in software engineering communities.

----September---
Started working with John Dickerson at the University of Maryland on rideshare fairness
Started working with Jordan Boyd-Graber at the University of Maryland on improving entity linking through data collection
Started as TA for class on programming languages, CMSC330

2020
---January---
Published research performed at Carnegie Mellon on toxicity at ICSE NIER 2020

Started as TA for class on coding interviews, CMSC389O

----March---
Presented work on entity linking at the MASC-SLL conference on speech and language

----May-August----
Worked at Facebook as a software engineer, developing web UIs for their feed ranking tools

---November----
Published work on entity linking with Jordan Boyd-Graber at HAMLETS workshop at NeurIPS Conference
Published work on rideshare fairness with John Dickerson at Machine Learning for Econ Policy at NeurIPS
Published work on rideshare fairness with John Dickerson at AAAI Undergraduate Consortium

2021
---January---
Became head facilitator for student-taught class on coding interviews (CMSC389O)
Proposed a new project on improving pricing systems for rideshare systems by including network externalities

---January-May---
Took a class on computational linguistics, working on a class project on bias in question answering, published in September at Machine Reading for Question Answering at EMNLP 2021

---April---
Published work on rideshare fairness with John Dickerson at IJCAI 2021

---May-August---
Worked at Lincoln Labs with Michael Yee on improving human-AI collaboration through fine-tuning and semi-supervised learning

---September---
Started working with Partha Lahiri at the University of Maryland on improving edit imputation algorithms

Took a class on deep learning; pursued a project on recurrent networks for denoising, and submitted work at the ECCV conference

--October--
Published work done at Lincoln Labs with Michael Yee at Workshop on Human-Machine Decisions, NeurIPS 2021
Named Phillip Merrill Presidential Scholar (awarded to top seniors at the University of Maryland)

--December--
Named a finalist for the Computing Research Association Undergraduate Researcher of the Year award

2022
--January--
Awarded Churchill Scholarship for a 1-year master's program at the University of Cambridge
Took a class on multilingual Natural Language Processes (NLP), and developed a project that compared the difficulty of training translation algorithms for different languages based on the amount of resources available

--February--
Started working at the World Resource Institute, an environmental non-profit, by collecting and analyzing datasets on electric school bus adoption

--April--
Awarded NSF Fellowship for 3 years of graduate study
Decided to go to Carnegie Mellon for PhD after 1 year at Cambridge

--May--
Graduated from the University of Maryland Magna Cum Laude with High Honors in Computer Science and Math
Named one of the top undergraduate researchers at the University of Maryland (awarded to top 6 undergraduate researchers)

--July--
Helped with reviewing for two conferences: NeurIPS 2022 and CoLING 2022

--Graduate (Cambridge)--
2022
--October--
Started school at the University of Cambridge pursuing an MPhil in Advanced Computer Science
Started research with Professor Mateja Jamnik at the University of Cambridge on Concept Hierarchies

--December--
Collaborated with Katie Collins and Mateo Espinosa on a project about uncertainty in concept-based models

2023
--May--
Published work on concept uncertainty at the Artificial Intelligence and Ethics Symposium
Presented work on concept hierarchies at the Gordon Research Symposium on Cancer Genetics

--July--
Reviewed for NeurIPS
Started working on distribution shifts project with Bryan Wilder
Started a new project on spurious correlations in concept-based models

--Graduate (Carnegie Mellon University)--
2023
--August--
Started PhD at Carnegie Mellon University (CMU) in Machine Learning
--September--
Submitted work on concept hierarchies to NeurIPs Workshop
Submitted work on interpretable models to NeurIPs Workshop
–October–
Started working with Fei Fang on food rescue matching

---Course Highlights Major----
–Carnegie Mellon–
1 semester in the following computer science graduate classes: Machine Learning
1 semester in the following math classes: Statistics

--Cambridge--
1 semester in the following computer science graduate classes: Machine Learning, Robotics, Machine Learning, Graph Neural Networks
1 semester in the following math classes: Category Theory

--University of Maryland--
1 semester in the following computer science graduate classes: Computational Linguistics, Deep Learning, Mechanism Design, Quantum Computing, Computational Geometry,
1 semester in the following science undergraduate classes: Computer Vision, Cryptography, Databases, Algorithms 2
1 semester in the following math classes: partial differential equations, real analysis, probability theory, statistical theory, numerical methods, advanced linear algebra
1 semester in following lower-level classes: programming languages, algorithms, computer systems, calculus 3, tested out of linear algebra

---Course Highlights, Non-Major---
1 semester in agent-based modeling, natural resource economics, educational policy, physical chemistry, naval ethics, reverse engineering, neuroscience
3 semesters in cybersecurity

Academic Honors

Honors/Fellowships

CS Honors College - Research-based honors college, allowing students to take classes on research, create a senior thesis
High Honors - Completed an outstanding senior thesis in computer science; awarded to a few students every year
Global Fellows Program - Honors program that explores international and policy issues, culminating in a part-time internship
ACES Honors College - Cybersecurity-based honors college that teaches basics of cybersecurity, Linux, bash scripting. Developed cyber security-based research project
CMU REUSE REU Fellowship - Participated in a research program at CMU for 10 weeks, researching toxicity in online communities.
Math Modeling - Received outstanding award for modeling of refugee populations for the SCUDEM math modeling competition
Bloomberg Codecon - Received top 3 awards at the University of Maryland programming competition, and qualified for national competition at Bloomberg headquarters in New York City

List here, in chronological order, any fellowships, scholarships, teaching, or other appointments held since entering college or university. Separate your undergraduate and graduate education intervals (if/as applicable) with a single dashed line. Put an asterisk (*) at the beginning of all lines indicating a national-level award (e.g., award of a National Merit Scholarship; NSF, NASA, or NDSEG Fellowship; election to Phi Beta Kappa, Tau Beta Pi, or Sigma Xi; etc.). Use one line per item whenever feasible (provide basic what, where, and when data). You may use more than the allotted space for your response, and any additional text will be included in your application.

Scholarships

NSF Fellowship - Fellowship that provides 3 years/150K for graduate school study
Churchill Scholarship - Fellowship that provides 1 year fully funded masters at Cambridge; one of the top scholarships for UK study, on par with others like Rhodes, Marshall, and Gates-Cambridge
Goldwater Scholarship - Prestigious scholarship given to top undergraduate researchers nationally
CRA Undergraduate Researcher of the Year Finalist - National award for the best computer science undergraduate researcher; named one of the top 30 undergraduate CS researchers
UMD Undergraduate Researcher of the Year - Given to the top 6 seniors for their research accomplishments
Philip Merill Presidential Scholarship - Given to the top 15 seniors for academic and service record
Iribe Scholarship - Computer Science scholarship worth 11K for academic and research records, given for 2 years (11K each year)
CMNS Summer Scholarship - 1K funding for select students pursuing research internship during the summer
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

Teaching/Volunteering
Technights (Teach local middle school students about computer science) - Fall 2023 - Present
Graduate Academic Support Program (Help applicants with their PhD application materials) - Fall 2023 - Present
AI Mentorship (Working with Undergraduates to learn about research) - Fall 2023 - Present
Student Community Action (Working with local children as a mentor) - Fall 2022 - Spring 2023
CMSC330 TA (Programming Languages) - Fall 2019-Spring 2022
CMSC389O TA (Coding Interviews) - Spring 2020 - Present, head TA Spring 2021-Spring 2022
CMSC122 TA - Spring 2019
Maryland Mentors Program, volunteer at local elementary schools - Fall 2020 - Spring 2022
College Park Academy Volunteer (volunteering at local charter school) - Fall 2018 - Spring 2021, assisted with cybersecurity club, math tutoring
College Mentors (helping middle schoolers with career pathways) - Spring 2021

First, please list the most significant research projects that you have pursued, in chronological order.

1) 1) A Muffin-Theorem Generator at Fun with Algorithms (FUN) 2018, PI: Bill Gasarch, May 2016-August 2016
2) 2) Discovering mutational signatures through non-negative matrix factorization, PI: Aravind Srinivasan and Max Leiserson, May 2018-August 2019
3) 3) "Stress and Burnout in Open Source: Toward Finding, Understanding, and Mitigating Unhealthy Interactions" at ICSE 2020, PI: Bogdan Vasilescu and Christian Kaestner, May 2019-August 2019
4) 4) "Investigating methods of balancing inequality and efficiency in Ride Pooling" at AAAI Undergraduate Consortium 2021, PI: John Dickerson, September 2019 - January 2021
5) 5) "Data-Driven Methods for Balancing Fairness and Efficiency in Ride-Pooling" at ML For Economics Policy NeurIPS 2020 Workshop, PI: John Dickerson, September 2019 - January 2021
6) 6) "What more can entity linking do for Question Answering at HAMLETS" NeurIPS 2020 Workshop, "Entity Linking Beyond the Easy Stuff: Annotating Oblique Entity Mentions" submitted but not accepted at EMNLP 2022, PI: Jordan-Boyd Graber, September 2019-Current
7) 7) "Data-Driven Methods for Balancing Fairness and Efficiency in Ride-Pooling", presented/published at IJCAI 2021, PI: John Dickerson, September 2019 - January 2021
8) 8) "Eliciting Bias in Question Answering Models through Ambiguity" at MRQA EMNLP 2021 PI: Jordan-Boyd Graber, January 2021-August 2021
9) 9) "Improving fairness in rideshare dynamic pricing algorithms", PI: John Dickerson, January 2021-May 2022,
10) 10) "Improving human-AI deference algorithms through fine-tuning", PI: Michael Yee, done while at Lincoln Labs, May 2021-August 2021, Workshop on Human-Machine Decisions at NeurIPS
11) 11) "Generalizing from Easy to Hard Examples in Image Denoising using Recursive Networks", PI: Soheil Feizi, September 2021-March 2022, submitted but not accepted to ECCV 2022
12) 12) "Human Uncertainty in Concept-Based AI Systems", PI: Adrian Weller, Accepted at Artificial Intelligence, Ethics, and Society Conference
13) 13) "Concept Hierarchies for Concept Learning" - MPhil Thesis Project at Cambridge, PI: Mateja Jamnik, Presented at Gordon Research Seminar on Cancer Genetics and Epigenetics; Under Review at NeurIPS XAI in Action Workshop 2023; October 2022-September 2023

14) 14) "Spurious Correlations in Concept-Based Models", PI: Mateja Jamnik, Under review at NeurIPs XAI in Action Workshop 2023; July 2023 - Present
15) 15) "Restless Multi-Armed Bandits in Online Food Rescue Notification Matching", PI: Fei Fang, October 2023 - Present