Motivation: The effectiveness of neural networks (NNs) has allowed for "deep learning" to have far reaching applications in fields such as drug discovery in pharmacology, protein folding in biology, diagnosis in medicine, computer vision in defense, etc. However, there is a gulf between the incredible practical performance and the fairly weak theoretical guarantees of NNs. It is not well understood how to best analyze NNs, why certain training methods allow NNs to "generalize" outside of their training set well, or how a NN implicitly reasons when placed within a decision-making context like reinforcement learning. When applying AI systems more broadly to critical domains like healthcare or defense, the existence of theoretical guarantees addressing these issues may be paramount for trust to exist between user and system.

Research Career Goal: My research career goal is twofold: (1) I will create novel AI algorithms which **retain theoretical optimality guarantees** while offering performance competitive to similar applications of NNs; (2) I will propose and implement **improvements motivated by recent theoretical advancements** to systems utilizing NNs. Both my research goals directly ameliorate well-described pitfalls of deep learning by introducing a theoretical component to improve and broaden the use of AI.

To achieve this, I will obtain a PhD in Computer Science (CS) then enter into academia or a government research lab to lead a team of researchers describing new theoretical tools for learning guarantees in AI. I will simultaneously teach AI theory at the university level. In addition to receiving mentorship, graduate study will allow me to build knowledge and intuition for the theoretical tools which I intend to create and apply while improving AI systems. Alongside my career as a researcher, I will continue developing programs increasing engagement in STEM with a particular focus on women in CS.

Research Background (Goal 1): In order to prepare myself for creating theoretically grounded alternatives to NNs, I have participated in a variety of research projects which have allowed me to **study useful theoretical techniques** from fields including convex optimization, online learning, and CS theory.

The summer before my senior year, I participated in an NSF REU with Prof. Brendan Juba from Washington University in St. Louis on reinforcement learning (RL), an AI subfield where an agent explores and learns within an environment. A fundamental result in RL theory is an upper bound on the learning speed of an agent, typically in terms of the number of *states* that the environment takes. Yet many practical RL application domains—such as resource allocation tasks or the board game "Go"—have an exponential number of distinct states, which means that *learning could take intractably long*.

Our work presents a practical, theoretically grounded approach for RL in these large and difficult application domains. Previous approaches assume the existence of a *potentially impossible* general polynomial time "oracle planner" for a family of exponential sized environments. In the case the long-term values of states can be described by a basis of linear functions, we *construct* this oracle by showing fast convergence of a cutting plane method with large input. We then prove competitive learning bounds **without reliance on any exponential quantities**. We are efficient in terms of the number of basis functions rather than the number of distinct environment states, where many environments of interest have an extremely large number of states, but may be described or approximated by a sum of local basis functions. Our results have been **submitted to *AISTATS 2021***.

Working remotely due to COVID-19, I started the project with a graduate student and Prof. Juba and met with them twice a week to discuss directions and proofs. In addition to learning effective habits of team-based theoretical collaboration over many hours of discussion, I independently studied, derived, and recorded over 30 pages of proof details. I also wrote a majority of our submitted paper, and gained valuable feedback in framing theory to be accessible and useful to the broader scientific community.

In addition to my work in RL theory to provide alternatives to NNs, I have collaborated with Prof. Nicholas Ruozzi and Prof. Benjamin Raichel at UT Dallas (UTD) to **bridge the gap between traditional and deep RL** with techniques from convex optimization and computational geometry. Further details are given in my research statement. From my second year of university until now, I extended existing convex optimization procedures, wrote and debugged code implementing and testing our algorithm in various RL simulations, and developed a broad knowledge of optimization. I am fortunate to be supported by an NSF REU for Fall 2020, and have presented my work as posters at two undergraduate research competitions.

Although it was frustrating to work for nearly two years with few results, I became driven by the difficulty and eventually learned and developed the grit necessary for the theoretical work that my research career will focus on. I was recently able to obtain exciting preliminary data that further spurred my motivation and passion for the project, and we are now preparing competitive results for publication.

Research Background (Goal 2): In line with my second goal of improving systems utilizing deep learning, I have pushed myself to participate in a variety of interdisciplinary application domains where NNs are commonplace to gain practical insight and **apply improvements rooted in theory**.

The summer before my junior year, I was a research intern at Johns Hopkins Applied Physics Labs (APL), implementing recent deep learning methods to classified computer vision defense contract work. Collaborating with APL research scientists, I learned firsthand the difficulty in using deep learning for practical problems **where high performance was required**, since NNs infamously do not provide any performance guarantees. This reaffirmed to me the importance of theoretically grounded approaches.

In my second year of university, I also joined the UTD Advanced Networks Research Lab (ANRL) under the supervision of Prof. Jason Jue. With a graduate student, I worked on a project distributing NNs layer-wise over computing nodes in such a way that inference accuracy could still be retained during partial node failure. Motivated by recent theoretical work, I proposed, prototyped, and implemented a system for NN layer distribution using a NN training technique we term *failout*. Failout incorporates simulated inference-time environment failure within NN training. Empirically, this allows NNs to be resistant to *true* inference-time layer failure. Eventually, our work appeared at ***ICML 2020 Workshop on Federated Learning for User Privacy and Data Confidentiality***, and is **currently under submission at *AAAI 2021***.

At ANRL, I also worked on a resource allocation problem called "progressive recovery". After failures such as natural disasters or cyber-attacks, we want to recover communication networks in a prioritized order. With a graduate student and Prof. Jue, we first proved general progressive recovery is NP-hard which **resulted in a conference publication at *IEEE Globecom***. Next, we proved that in many practical cases, optimal orders can be found by searching a smaller solution space. This allowed me to apply deep RL and obtain state-of-the-art results on US network topologies, **leading to a journal publication in *IEEE Journal on Selected Areas in Communication (JSAC), issue on AI/ML for Networking***.

At ANRL I learned how to plan research in the context of our larger lab goal by pushing projects in directions utilizing theoretical guarantees. Furthermore, I practiced communicating theoretical work to a general audience by presenting research on optimization and RL at two UTD graduate seminars, and a poster on progressive recovery at an industry-sponsored, national computer networking conference.

Intellectual Merit: I will start graduate school already focused on two high-impact research goals: **advancing the theoretical knowledge** and **broad proliferation of modern AI systems**. State-of-the-art AI systems with NNs are already being deployed. Yet, NNs are extremely difficult to improve or guarantee performance of, **fundamentally restricting their complete proliferation to sensitive fields** where such guarantees are necessary. The NSF GRFP award will permit me to study both NNs and their alternatives from a theoretical lens. I will continue providing guarantees of successful learning in RL and improving the performance of AI systems using NNs. Using theoretical techniques which I have already gained experience with, I also plan to explore generalization in AI systems more broadly. This will advance both theoretical knowledge and practical insight within already deployed and newly created AI systems.

Personal Background: To prepare for my PhD, I am in the UTD CS honors program which readies a group of 30 students for research through theoretical courses. I am also **pursuing a second degree in math**, and was the only undergraduate in a graduate course on optimization in machine learning (ML). As a result of my research and academic success, I was one of four students **nominated by UTD for the Barry Goldwater scholarship**, one of ten who received a UTD Jonsson Engineering School research award, won a UTD Undergraduate Research Award, and am on a four-year full-ride academic merit scholarship.

In addition to my research interests, I am also passionate about strongly advancing the broader engineering community, partly because I was fortunate to have mentors who shaped my own undergraduate journey. In particular, my goals are to provide resources to prepare students for careers in engineering and research, and to increase the representation of women in CS and STEM more broadly.

I currently serve as the President for the Association for Computing Machinery (ACM) UTD chapter for the 2020 year. Our mission is to prepare undergraduate students for CS careers in industry or academia by providing extracurricular academic resources. Since I joined ACM UTD in Spring 2018, I have helped grow the organization from 30 officers to almost 70, and during COVID-19 created an online support community of more than 1000 engineering students on the platform "Discord".

The semester I joined ACM, I created the ACM Mentor program. The program matched 15 CS underclassmen with older undergraduates to help them derive achievable career goals. I recognized the impact of the program, and recruited additional students to help run it. To date, we have **matched over 120 CS underclassmen with upperclassmen mentors**. I have personally mentored four students over two years, all of whom are interested in pursuing Ph.Ds. I started a reading group covering fundamental ML theory with these students, meeting on the weekends to prepare them for potential research work.

The Fall 2020 semester, ACM started another program called ACM Research. Our goal is to expose early undergraduate students to research in CS. This semester, we have **four research project groups of five students**, one of which is in collaboration with tenured research faculty, on projects related to NNs, networking, and natural language processing. We hope to grow the number of projects and faculty involvement in the coming months. Outside of ACM, I have also **personally recruited and mentored a group of undergraduate students through a research project**. My junior year, with a graduate student and Prof. Jue (UTD), I led a group of five undergraduate students to collaborate on an extension of previous work I had contributed to. Meeting weekly, I taught them deep RL and guided them through the project implementation. We are now collecting competitive results for publication.

ACM UTD also organizes the largest regional hackathon, HackUTD, which in 2019 hosted 750 undergraduate and graduate students from around the US. Attendees were able to interact with industry sponsors, and receive job interviews on site. With over 20 companies in attendance, numerous students received internship or job offers from the event. **We were able to fundraise more than $50k to hold HackUTD 2019**, and hope to grow the yearly event further once it is safe to do so.

One of my proudest accomplishments was **establishing a $30k endowed scholarship** to honor a former department head. Using funds saved from the previous two HackUTDs that I had helped organize, I proposed and signed a $25k donation from ACM UTD to start the "Dr. Gopal Gupta Undergraduate Scholarship for CS established by ACM 2020". *Each year in perpetuity*, one student will receive a $1k scholarship which under Texas law is eligible for an in-state tuition waiver. This could potentially **save each winner more than $10k in yearly tuition fees** if they were an international or out of state student.

I am also passionate about increasing the representation of female-identifying students within engineering. Only 17% of CS undergraduates at UTD identify as female. Even as ACM UTD has more than doubled in size under my leadership, I have proactively maintained that **at least 40% of our almost 200 officers plus program participants (across all our programs) identify as female** through implementing bias reduction in our selection process. Furthermore, I am currently working with pro-bono HR employees to receive more general advice on inclusion within our organization more broadly. Outside of ACM and until COVID-19, I also co-led a program called Empower Through Code to teach STEM concepts to **low-income female students at risk of dropping out of middle school**. We worked with a local non-profit to identify these students, and visited them each week to get them interested in and aware of STEM concepts such as robotics or elementary CS which could directly lead to STEM careers.

Broader Impacts: The NSF GRFP award will allow me to continue to serve communities without early STEM exposure. In graduate school, I will expand Empower Through Code and also develop a curriculum for precollegiate females in STEM based on my own experiences teaching. Additionally, I will leverage my research connections and strong history of undergraduate mentorship in both STEM and research to foster the next generation of women researchers. In particular, I will individually mentor and connect UTD women in STEM with the many research opportunities at UTD and my future graduate institution.

The GRFP will also propel my own research into other scientific domains. My second research goal of applying theoretically grounded improvements to AI systems has applications even outside of CS. As an NSF researcher, I will have access to a rich network of international interdisciplinary collaborations. I will take advantage of these collaborations to continue my work on grounded improvements to AI systems in defense, and also plan to expand into healthcare. My AI research goals and experiences coupled with my leadership demonstrating teamwork and initiative show that I have already developed many skills required to successfully lead a research group while prioritizing the involvement of underrepresented groups.