My collegiate and professional career has been focused around use-inspired artificial intelligence for the fields of humanitarian assistance and disaster response (HADR). For the last three and a half years, I have led a research lab that has focused on computer vision and ground robotics as applied to HADR. We have made a massive impact, changing the way that disaster agencies such as CAL FIRE perform automated building damage assessment after disasters. My work has highlighted the need for the tight integration of efficient machine learning algorithms with HADR domain expertise, along with the proper policies to implement these products for a high degree of success.

I believe that joining the Artificial Intelligence (AI) program at UC Berkeley will enable me to bring a much more rigorous methodology to designing HADR technology of the future. Combined with the expertise of UC Berkeley's Goldman School of Public Policy, we would additionally be able to responsibly deploy this technology.

My research has found tremendous success in solving a central problem every disaster agency has to contend with: understanding the scale of damaged buildings and infrastructure as a disaster occurs. Traditionally, personnel have to be on-site as a disaster is ongoing to begin the assessments. Not only is this extremely dangerous, it is time-consuming and diverts critical personnel from actually combating the disasters. Intuitively, it makes sense that assets such as satellites and unmanned aerial vehicles (UAVs), combined with computer vision, could do this task quicker with a higher degree of accuracy.

In order to solve this problem, I began collaborations with over 40 disaster agencies around the world, such as the UN Disasters Programme, FEMA, India's National Disaster Management Authority, and others. Each agency has a focus on a different type of disaster: wildfires, tsunamis, blizzards, etc. In order to get everyone on the same page, and to create a useful dataset, I worked with each agency to create a comprehensive requirements document. This detailed what agencies looked for when assessing damage, how assessment was currently executed, and what they would want in an automated future capability. After exhaustively combining all feedback, we output two things: 1) a Joint Damage Scale (JDS) that provided an objective scale to assess multiple types of damage from remotely sensed imagery, and 2) an expert-informed data collection and labeling procedure. The JDS was immediately adopted by multiple disaster agencies as their standard for imagery-based damage assessment.

We collected the largest and most-diverse disaster dataset available to date. After building a baseline computer vision model, we organized and launched an international prize challenge called xView 2 (https://xview2.org) in order to spur the research community to exploit the dataset to its maximum potential. To have a successful competition, we had to form partnerships to fund the challenge and transition the outputs, draft policies on how the dataset could be used, and create a system to fairly evaluate competitor submissions. Working in collaboration with the Defense Innovation Unit, we were able to raise $2M to execute the competition across many US Government partners.

Critically, after the competition, I worked with a small team of individuals to productionize the winning models. These were then deployed and used for the 2019 Australian bushfires, the 2020

California wildfire season, and the 2020 FEMA hurricane season, saving countless weeks of time and enabling quicker response and recovery to the disasters.

None of this technology is impactful without strong collaborations and even stronger policies governing the use of the algorithms. Recognizing this issue, I worked to launch bilateral and multilateral domestic and international partnerships focused around AI + HADR. Partnering with the Joint Artificial Intelligence Center, the Department of Energy, and the White House, we created and launched the FirstFive Consortium (https://www.firstfive-ai.org), a public-private partnership aimed at building capability and policies in the world of AI and HADR. Furthermore, I brokered bilateral partnerships between the United Nations Disaster Programme and Carnegie Mellon University, as well as the Defense Innovation Unit and the Australian Geospatial-Intelligence Organisation, to develop more efficient algorithms to be used in isolated parts of the world. Lastly, in order to facilitate the exchange of ideas and discoverability of problems faced by the HADR community, I organized and chaired the AI + HADR series of workshops at NeurIPS 2019 and 2020, bringing together leading AI researchers with the most impactful HADR front-line workers.

It is increasingly clear that efficient ML/AI is going to be a game-changer for the future of the HADR community. UC Berkeley has pioneered methods in both data-efficient and compute-efficient computer vision. Specifically, I believe that my experience and background would pair well with the labs of Drs. Trevor Darrell and Shankar Sastry. Dr. Darrell and I have a good working relationship, collaborating on overarching HADR initiatives and research group meetings. Furthermore, his computer vision research gets to the heart of data efficiency and targets HADR as an obvious transition for the work. Dr. Sastry's lab has a targeted focus on sparse and low-resource AI/ML, and he has been a proponent for increased multi-disciplinary AI work around campus, such as interactions between EECS and Goldman.

With my experience leading international and domestic AI and HADR initiatives and my experience in serving as the principal investigator of a research lab at Carnegie Mellon University, I believe that I bring a mature, targeted approach for research in mixed disciplines. Along with my fundraising acumen (over $5M in sponsored research over three years) and impact in the government and policy sectors, I also bring a burgeoning publication record. Combined with Berkeley's reach and innovative research faculty, having me as a PhD student would expand Berkeley's influence in the world of HADR and public policy.

By focusing on a thesis that combines efficient machine learning, HADR, and policy together, I believe that I can have a lasting impact on the way the growing field of AI + HADR is shaped. Fostering and growing a community of excellence in this area would directly save first responders time, resources, and most importantly, lives. By adapting and expanding my existing work, I can improve my own aptitude as a researcher while simultaneously achieving Berkeley's goals to provide long-term societal benefits through the transition and transmission of advanced knowledge.