My long-term research goal is to develop artificial intelligence that is able to interact with humans through natural language, and that helps humans easily access the information they need. I decided to pursue computer science during my third year of college. Since then, I have completed internships both in engineering and research at several companies including Google and Microsoft Research Lab. I became interested in natural language understanding and the interpretability of machine learning models after working. The former culminated in a project on multilingual question answering (QA) with no QA training data in the language of interest. The latter resulted in a system that accurately estimates uncertainty while achieving state-of-the-art results on depth estimation. The diverse perspectives I encountered during these research projects made me appreciate that encouraging a wider spectrum of people to enter the field is crucial to evolve. As a result, I have led efforts to improve inclusivity and diversity in computer science. For my research, academic, and community achievements, I was awarded the Google's Women Techmakers Scholarship and the Nakajima Foundation Ph.D. fellowship, which guarantees tuition and living expenses funding for the first two years of my Ph.D.

My internship at Google Search in Japan exposed me to the first application of NLP, sparking my interest while developing a transit application for the Google Search iOS App that retrieves transit information for queries such as "How can I get to Narita Airport?". At Google, I saw first hand how these elegant QA interfaces enable millions of users to instantly find information. A key challenge in developing such systems, however, is that resources for many NLP tasks are abundant in English but scarce in languages such as Japanese. When I returned to the University of Tokyo, I sought to address this challenge during a research project advised by professor Tsuruoka, in which I proposed a technique to develop QA systems for languages without QA training data in the language of interest. We did so by incorporating attentive neural machine translation (NMT) models during inference to translate the original language to English with minimal manually annotated data for translation. Given the nature of extractive QA, we also proposed a method to recover the answer in the original language using soft-alignments from the machine translation model instead of back-translating predicted answers. Unlike previous techniques, our method does not require additional annotation for the language of interest and leverages existing resources in English QA. To evaluate our technique, I collected the first extractive QA evaluation sets in Japanese and in French using Amazon Mechanical Turk. On these evaluation sets, our method significantly outperforms baselines that train a QA model on a corpus translated from English to the original language using Google Translate and that back-translate the answer produced by an English model to the original language. I recently submitted this work to NAACL 2019.

During my work on question answering, I became fascinated by the problem of how to interpret machine learning models. Despite the ever-growing success of deep learning models in a variety of domains, including safety-critical ones, they remain largely uninterpretable. How can we accurately quantify the uncertainty in the predictions of these models to understand their behavior? I investigated this question during a research project with Professor Aizawa, in which I formulate uncertainty estimation in prediction as a multi-task learning problem with regression and develop several novel uncertainty loss functions, inspired by variational representations of robust estimation. In addition to quantifying the uncertainty more accurately than previously proposed techniques, our method achieved state-of-the-art performance on several widely used depth estimation benchmarks such as KITTI and NYU Depth Dataset V2. This work ranked 1st among a class of 130 senior-year projects in my department at the University of Tokyo and is in submission to CVPR 2019.

In addition to natural language understanding and machine learning research, I am also passionate about how to use these technologies to create tools that benefit my community. At the University of Tokyo, I developed an iOS application that helps women commute alone safely at

night. The application notifies users of recent crimes nearby, and can automatically send videos and location to family or friends in emergencies. While on exchange at UC Berkeley, I developed a platform that connects local volunteers with organizations that provide emergency response services to increase the efficiency of post-disaster response. In addition to building tools, I also led efforts to improve inclusivity and diversity in computer science. At the University of Tokyo, I organized hackathons for female students to foster interest in computer science. I also served as an invited panelist to the Google Mind the Gap project, which teaches female high school students programming and empowers them to pursue careers in computer science. More recently during an internship at Megaton Labs, I crowd-sourced HappyDB, a collection of 100,000 happy moments expressed in natural language that enables the use of NLP techniques to facilitate research application development in positive psychology. This work was presented at WiNLP 2017 and LREC 2018 and provides the foundation to the AAAI 2019 Affective Content Analysis Shared Task. Because of my projects and services, I have been honored to receive awards at JP Hacks (Japan's largest student hackathon), the Berkeley Social Entrepreneur Contest, as well as the Google Anita Borg Scholarship. I will continue to contribute to my community during my Ph.D.

Currently, I am pursuing my interest in NLP and machine learning during an internship at Microsoft Research Asia, where I am investigating how to incorporate external knowledge into question answering systems while preserving their end-to-end trainability. Recent work on end-to-end neural models for QA has achieved significant performance improvements on a variety of extractive QA tasks. However, analysis of such models shows that they exploit a particular characteristic of these tasks, namely that the questions are largely answerable using surface-level cues. I am interested in developing methods and tasks that incorporate structured knowledge and common sense knowledge to move beyond learning surface-level patterns so that we can build more robust systems for real-world applications. The challenges of incorporating external knowledge into end-to-end QA models are three folds: 1) How do we retrieve the right bits of knowledge from vast amounts of data? 2) Given the relevant and useful knowledge, how do we ground the text to the appropriate knowledge for reasoning? 3) Facts in existing knowledge bases may be precise but are usually sparse, as a result, our methods cannot rely exclusively on such knowledge bases and need to instead reason jointly over knowledge bases and unstructured text. These are some of the research directions I would like to continue to pursue during my Ph.D.

My career goal is to become a professor so that I can continue to pursue my research interests in NLP and machine learning while teaching the next generation of researchers. I am eager to pursue a Ph.D. at the University of Washington for several reasons. First, UW NLP is one of the best NLP groups in the world, with seminal work in a wide breadth of topics such as question answering, common sense knowledge, contextual representations, and semantics. Second, the research produced by UW NLP demonstrates the group's collaborative nature, with partners in other departments as well as in industry labs such as AI2, FAIR, and MSR. Third, I appreciate UW's leadership in diversity. I believe that this combination of collaboration and diversity is the key to bringing new perspectives and generating innovative and important research ideas. At UW NLP, I am especially interested in professor Hajishirzi's research on question answering, professor Yejin Choi's research on common sense knowledge and work on learning and inducing semantics and structure by professors Smith and Zettlemoyer. I was fortunate to visit UW this past September when I had to pleasure of meeting professors Choi and Zettlemoyer as well as Ph.D. students in the group. My conversations further made me realize that I would be love to work with such passionate and brilliant researchers. I am confident that UW is the best place to develop my skills as a researcher, and would love to leverage my experience in research, industry, and community leadership to contribute to UW.