It is sometimes easy to forget how powerful and complex we are as humans. But our minds are capable of incredible feats, and all within a two-pound biological circuit! It's profound to reflect on, and during my junior year of college I decided to turn my persistent questions into a major and a career so I could think about it full-time.

Of course, understanding the mind will necessarily take an interdisciplinary approach. To this end I have worked in cognitive neuroscience, computational neuroscience, optogenetics, and developmental psychology labs. I've taken classes in artificial intelligence, probability, statistics, cognitive science, and biology. I've read countless papers and books on these topics.

There are many considerations that factored into my decision to transition from a predominantly neuroscientific background to cognitive science. The main motivation, however, is my strong impression that understanding the mind from the top-down, starting at Marr's computational level of analysis, will help propel fields like artificial intelligence and psychology forward. This is an intuition which stems from the fact that algorithms modeled after the physical structure of the brain (e.g., convolutional neural networks) still don't capture features of human thought such as transfer learning, compositionality, or strong inductive biases. Bayesian cognitive science has already shown its strength in these areas and it's exciting to see how closely these models mirror experimentation. Perhaps even more exciting is that addressing these high-level processes has already shed light on deep philosophical questions about the origins of human knowledge, and I believe it will continue to do so. In light of this, computational cognitive science appears uniquely situated to promote progress in A.I., psychology, and philosophy of mind, which is why I'm eager to contribute to the field.

Specifically, I am interested in hypothesis generation. Bayesian cognitive science has been instrumental in modeling human thought based largely on predefined hypothesis spaces. The natural next question is, how do humans form these hypotheses in the first place? Or, framed another way, what types of constraints are there on the algorithm used to search the infinite space of hypotheses? And how is it that humans are able to restructure their ontologies in the absence of additional data, e.g., when performing thought experiments? Bayesian rationality is a powerful filter for selecting hypotheses which fit the data, but we have yet to construct a satisfying account of creative rationality, its generative counterpart.

Addressing this will require computational skill, and as in any scientific endeavor, diligence, independence, and a deep understanding of the field. My experience thus far has prepared me for computational work, including a double major in mathematics and cognitive science. My computer science courses covered artificial intelligence, introducing me to concepts like Monte Carlo methods, MDP's, Causal Bayes Nets, VOI, Markov chains and much more. I also recently completed Andrew Ng's Coursera course on machine learning.

In the Knight Lab at the Helen Wills Neuroscience Institute, I developed my skills in Python, data analysis, and experimental design. I performed literature reviews to help construct the experiments, translated them into code, and collected and analyzed the data. I've continued developing these skills at X where I've designed and evaluated over fifty neural network models for a classification problem. This work constitutes the foundation of a masters-level project which I will publish on next year. I am fluent with Keras, TensorFlow, Scipy, and Numpy and I used my programming skills to contribute to a patent earlier this year which has been filed.

Diligence and independence are harder to quantify, but I believe I exemplify both. For instance, in 2016 I independently researched biosecurity. By the end I had written a 15-page report on the main issues and ways to intervene. I was subsequently offered a grant to pursue this work by a philanthropic organization, which I decided to decline as it conflicted with my new job. This is one example, but most of my work has been largely self-driven, even at places like X. I decide which models to test and which features might be interesting. I'm familiar with working under uncertainty and patiently pushing forward, and I'm excited to continue honing these skills in graduate school.

Since my background has largely been in neuroscience I don't yet have a deep understanding of psychology. However, I'm very excited to develop this! At Berkeley I took Neuropsychology, Social Psychology, Biological Psychology, and Language Acquisition, excelling in all of them. I've read many papers in the field and my experience with literature reviews will give me the opportunity to expand on this.

I would be thrilled to explore my interests at NYU. From what I can tell, the environment in the computational cognitive science cluster is highly interdisciplinary, with many collaborative opportunities, a focus on "cognitive computing," weekly meetings with a range of psychology faculty, and the ability to take courses with top researchers in machine learning. This is, essentially, my dream.

In addition to this, the culture of Professor Gureckis' Computation and Cognition Lab seems ideally suited to help me develop as a graduate student in the ways I'm excited about. For instance, his commitments to being a hands-on advisor and quickly turning ideas into implementations are optimal characteristics for my graduate environment. I am independent and diligent, but I'm also pursuing a Ph.D. to learn. I want to move fast and break things, develop my intuitions as a scientist, and hone my creativity and curiosity. All of which is, I think, best advanced through fast feedback loops supervised by experts in the domain. Also, the lab retreats sound amazing.

If I have the opportunity to pursue my Ph.D. at NYU, I would be particularly excited about working in Professor Gureckis' or Lake's labs. Were I to work with Professor Gureckis, I would be particularly interested in building on the work done in "Grounding Compositional Hypothesis Generation in Specific Instances." One approach I would like to test in this domain is using an explore-exploit tradeoff within a hierarchically structured hypothesis space. This would include a meta-level hypothesis that curtails the space of hypotheses that can be considered and object-level hypotheses which generate the configurations to test. Exploration occurs through the switching of meta-hypotheses and exploit occurs when choosing hypotheses within the current frame. While there are many ways one could instantiate this algorithmically, I'm interested in using bottom-up and analogical constraints to generate meta-hypotheses, stochastic composition for object-level hypotheses, and concepts borrowed from simulated annealing for deciding when to explore. For instance, if one sees that the initial correct koan displays structural similarity, e.g., two green triangles, then the meta-hypothesis could be set to considering only those hypotheses which include two green figures, and the object-level hypotheses could be created using PCFG, relative to this constraint. If there are too many incorrect proposed koans in a row, then the temperature increases, and the meta-hypothesis is switched in correspondence with other salient analogical features.

This proposal is informed by evidence from some of Schulz' work at MIT. Forming meta-hypotheses, irrespective of the algorithmic underpinnings, seems to be a process that people engage in. For instance, children will use information about the structure of the problem to inform their solutions, even when they don't have access to any data. This goal-constrained behavior suggests that people already have an idea about the form their solution will take, i.e., the types of hypotheses they should consider.

If I worked with Professor Lake, I would want to explore topics on structural sparsity, question generation, and Bayesian program learning. I'm especially interested in modeling human creativity in the formation of new concepts, and in line with this I want to develop models which capture distinctly human qualities in areas like art and music. To do so I would use the model previously developed in Lake's paper on program induction, and apply it to music composition, with notes acting as primitives, chords as sub-parts, cells as parts, and so on.

Finally, I would like to touch on my reasons for pursuing graduate school over other options. Working at X has been an amazing experience and I have already learned so much. That said, I notice myself continually drawn to topics outside of work. I read papers on compositional hypothesis generation at lunch and study books like *Surfing Uncertainty* over the weekends. I've been persistently intrigued by philosophical inquiries into cognition and knowledge, questions which are simply not well suited for an industry focused on application.

I know that graduate school will entail challenging, cognitively demanding, and sometimes downright grueling work, and I know that this will sometimes be stressful. However, I'm invigorated by this kind of environment. I love getting absorbed in code for hours on end, reading late into the night, and discussing philosophical and scientific questions at length with my friends. In light of this, the challenges that graduate school offers are more of an opportunity than a hurdle. I feel equipped to take them on and I'm excited to learn and grow my skills as a scientist during my time at NYU and beyond.

Thank you so much for your time.