I was a writer long before I was a scientist, but it was science that drew me to MIT, tech-heavy startups, and machine learning. Ultimately I want to marry my existence as a scientist to my existence as a writer. This is what will provide me with the skills, insights and motivation to do important and novel research in natural language processing. Although I am currently most interested in affective language, my interest in natural language processing is more broad and includes semantic interpretation, machine learning interpretability and visualization, and generative language systems.

I am fascinated by how machines can understand the affective side of language. In particular I have been thinking a lot about what the literary arts can teach us about natural language processing and emotional language. For instance, poetry highlights technical difficulties in natural language processing as it relies on sophisticated language (e.g. connotation, metaphor, and reference) and less structured syntax. How can we build systems that understand figurative language without relying on manual annotation? In addition, the emotion that some writing expresses cannot be determined by correctly interpreting factual statements. How does that change the way we model language?

I have come to this interest from a winding path. I studied mechanical engineering at MIT as an undergraduate because I was enthralled by how well mathematics could model the real world. I did extensive research on jamming, a technique of applying a vacuum to granular matter such that a structure can be changed from a semi-fluid into a solid. Although the topic of this research is very different from natural language processing, it taught me how to conduct academic research and inspired my interest in graduate studies in general. I demonstrated that the 2D shape metrics circularity and polydispersivity of the granular matter were correlated with the strength of the resulting solid asymptotically, an important discovery for those using this technique in robotics. I presented the results of the paper at the March American Physical Society meeting in 2012 and some of my experimental results were featured in a paper presented at the 2012 IEEE International Conference on Robotics and Automation.

Though I loved this research experience, at the end of my undergraduate degree I was not committed to studying how to build better machines. I went into industry to figure out what interested me more.

After graduation I worked as a product developer on the Mimo Baby Monitor which measures the respiration and movement of an infant. I was quickly wooed by the mass of respiratory data we were collecting. For one project I investigated the sleep cycling of infants, developing an algorithm that could detect infant sleep cycles at an 80% accuracy. This algorithm still today produces a sleep graph shown to parents in the Mimo application. This was my first exposure to the complications of interpreting unstructured data, of building on the research of others, and the power data has to provide useful insights. I began to read about artificial intelligence techniques, learned the programming language R, and took Prof. Ng's online introductory course to machine learning. I looked for a job that would allow me to pursue this more sophisticated analysis.

In my current position as the Research and Development Lead at Soofa, I lead our engineering team in the creation of smart cities products and solutions. We created a pedestrian scanning

sensor for our connected park benches and I am responsible for analyzing the data we collect from cities and parks across the country and communicating this to clients. One area of exploration is finding social linkages between pedestrians by analyzing the name of the Access Points their phones have previously connected to. Based on research the University of Lyon, I have developed a metric that links devices using a statistic similar to term frequency–inverse document frequency. I am currently working on using this to create a location-based community score that measures the strength of connections between people in the community and the potential for unusual collisions.

All the while I have had an ongoing relationship with creative writing. As an undergraduate I won several awards for my writing and a chapbook of poetry I collected entitled 'Poems for MIT Students.' After graduation I continued this pursuit with poetry publications in a literary journal and a newspaper, a chapbook sold at Grolier's Poetry Shop, and the publication of a travel essay in Harvard Bookstore's summer travel anthology. Natural language processing is the intuitive brainchild of a love of words and a love of machine learning.

Although I do not have an undergraduate degree in computer science, my degree in mechanical engineering prepared me with the fundamental mathematical and analytical skills to succeed in a computer science program, as well as the ability to conduct high quality academic research. My time in industry has given me an informal education in computer science and data analysis, working with researchers, stakeholders, and customers. In addition, I've taken several online classes to build on this informal education. I am sure that my previous and perhaps untraditional experiences will only add to my abilities as a researcher.

I am excited to return to an academic environment where intellectual curiosity drives motivation. Over the summer I began reading computer science papers and researching pertinent projects. My first exploration into natural language processing for affect has been the creation of a poetry recommendation engine, an implementation of latent dirichlet allocation which has already produced intriguing results. I have created a web application to interact with the engine that can be found at www.poetry-engine.com.

My unique combination of experience in the arts, engineering research, and data science in tech startups gives me a perspective on computer science research that will be a strength in my pursuit of a Ph.D at Columbia University.