Currently, I am editing and adding significant content to various chapters in a book soon to be published from Dr. Song-Chun Zhu, my advisor, and Dr. Ying Nian Wu at UCLA. The book, *Statistical Models for Marr's Paradigm*, is intended for graduate students in artificial intelligence, specifically vision and cognition, and is meant to accompany the *STATS 232A*, *232B*, and *232C* course series at UCLA, which covers statistical modeling, learning, computing, and inference in vision and cognition. I am also completing a Masters thesis in generative modeling in vision under my advisor Dr. Song-Chun Zhu at the Center for Vision, Cognition, Learning, and Autonomy (VCLA) at UCLA. I have contributed to several research projects and a paper *Short-Run MCMC Residual Network toward Energy-Based Model* currently under review by the Conference on Computer Vision and Pattern Recognition (CVPR). My second paper *Generative Recurrent Language Models with MCMC Inference*, which I am co-authoring with a PhD student Bo Pang, will be submitted soon to a top Natural Language Processing (NLP) conference.

I joined the VCLA lab last academic year with the *Vision and Language* group. There I worked closely with the project leader to create models for visual question and answering tasks. Visual question and answering research entails parsing language and video jointly to understand events and answer questions about the events. Our research involved using CNNs and LSTMs with Attention and researching ways to improve hierarchical question-image co-attention models. I joined the *Generative Modeling in Vision and NLP* group this past Fall because I became much more interested in generative modeling throughout my Masters program. A variety of factors contributed to this, including coursework, my interest in missing data problems and vision and language generation, and in general I enjoy learning about the landscape of generative, inference, and energy based models applied to various vision and language tasks. Models like the variational autoencoder (VAE), the generative adversarial network (GAN), and the short-run Markov Chain Monte Carlo (MCMC) model, used in my research at the VCLA lab, each employ unique solutions, quite valuable to learn, for the generative modeling goal of capturing the true distribution of the training data.

I especially enjoyed my opportunity to work on *Short-Run MCMC Residual Network toward Energy-Based Model* with PhD students Erik Nijkamp and Bo Pang. I built and optimized variational autoencoders for different image datasets, implemented dynamic step size approaches for our short-run MCMC model, and made several improvements across various models to optimize synthesized, reconstructed, and interpolated results from VAEs, GANs, and our short-run MCMC models. The short-run MCMC model can be treated as a valid generative model just like a VAE or GAN. For example, unlike a traditional energy based model or MCMC, it is capable of reconstructing observed images and interpolating between images like generator or flow models can. And it produces results that are better than, or rival, VAEs and GANs, while requiring no auxiliary inference network (with its added parameters and additional model specifications) to approximate the posterior distribution of the latent variables, as is necessary with VAEs and GANs.

After submitting this paper, I was subsequently asked by senior PhD Bo Pang to co-author my second paper, *Generative Recurrent Language Models with MCMC Inference*, which is an NLP investigation on generative, recurrent language models that use MCMC inference to generate language. This paper utilizes an image parsing to text description (I2T) framework that generates text descriptions in natural language based on an understanding of image and video content. The l2T framework consists of

four components: an image parsing engine, and and-or-graph for visual knowledge representation, a semantic web, and a text generation engine.

The other key factor that contributed to my interest in generative modeling has been my coursework in *STATS 201C: Advanced Modeling and Inference*, *STATS 202C: Monte Carlo Methods for Optimization*, and *EC ENGR 239AS: Neural Networks and Deep Learning*. The first course I took with Dr. Qing Zhou; for this course, I submitted a final term paper, *Formulation of the Variational Autoencoder and Evidence Lower Bound and an Application to the MNIST Dataset*, to better understand the original variational autoencoder paper. In the second course, instructed by Dr. Mark Handcock, I gained a firmer grasp throughout the quarter of the importance of sampling methods in my research. In the third course, instructed by Dr. Jonathan Kao and based on *CS 231n: Convolutional Neural Networks for Visual Recognition* from Stanford, I also learned a great deal about generative modeling and VAEs and GANs.

After my PhD program, I am very much open to work in academia as a professor. As reflected by my large amount of teaching experience for the SAT, ACT, and GRE exams and various math and science courses, I enjoy teaching a great deal. It is motivating to empathize with students, understand their confusions, and use that to more effectively convey course material. Research as a professor would also provide me the continual satisfaction of carrying out my research intuitions. For now, I am generally open to future opportunities as my primary concern is finding a career that best facilitates my research interests.

I returned to UCLA to study statistics and machine learning after working as a Full Stack Software Engineer in Silicon Valley for almost three years. Since planning to return to school, I have focused on a statistics education for my machine learning research. Before I was even admitted to UCLA for my Masters program, I had completed many courses online using Stanford and MIT online courses and learning platforms such as Coursera and EdX. My success as a Full Stack Software Engineer was due to this self-study and the Hack Reactor program, an advanced software engineering immersive program in San Francisco. I became a paid tutor at Hack Reactor before securing two software engineering positions over the next three years, at Cinemagram in San Francisco and most recently at NatureBox in Redwood City. At Naturebox, I led many full stack projects such as a second version of our API, a customer payment migration from Litle to Stripe, and AmazonPay checkout, amongst other projects.

I have significant experience, both independently through research and through coursework, in generative modeling in computer vision and natural language processing, and my contributions to various research papers at the Center for Vision, Cognition, Learning, and Autonomy has culminated in a research paper *Short-Run MCMC Residual Network toward Energy-Based Model*, under review by CVPR, and a co-authorship for a paper *Generative Recurrent Language Models with MCMC Inference*, soon under review by a top NLP conference. This is in addition to my large contributions to a book *Statistical Models for Marr's Paradigm* by my advisor, Dr. Song-Chun Zhu, and my Masters thesis in generative modeling in vision currently in progress. Additional information about my independent machine learning projects can be found on my website ericmfischer.com and on my Github page github.com/ericmfischer.