My long term goal is to contribute to understanding human intelligence and learning. To this end, I am interested in performing research on applications of graph theory algorithms to computer vision.

I find graph theory lucid and elegant. Graphs can be used to model many phenomena, from social networks, disease contagion, path and roads, and flow of information or commodities through a network. Graph theory algorithms often enforce local consistency while optimizing for global properties, and we can prove theorems about their behavior. While many seemingly simple graph theory problems have NP-hard solutions, there are many questions with elegant and efficient solutions; furthermore, for graph theory problems that are NP-hard, there are often good polynomial time approximation algorithms.

I want to do research in computer vision. Computer vision is particularly interesting because of its potential to benefit many areas of research even outside of computer science. Aside from advancing robotics and computer graphics, computer vision has the potential to benefit fields such as biology, geography, and sociology. These fields frequently collect images for classification and information extraction. In recent years, the amount of data in each of these fields has burgeoned. Satellite data allow geographers to classify and track changes in the types of land on the earth (desert, agriculture, urban, etc). New areas of inquiry in biology were made possible with biomedical imaging that non-invasively produce images of internal aspects of the human body. Along with the proliferation of digital cameras, there is now a wealth of photos and videos online from which social information can be gleaned as well. In all these fields, computer vision is an important information extraction tool. Moreover, understanding vision and building a machine that can "see" the way humans do would indeed improve our understanding of intelligence a great deal.

Surprising as it may be, graph-theoretic algorithms have been successfully applied to solve many computer vision problems. This is because many computer vision problems lend themselves to the same general structure of enforcing local consistency while optimizing for global properties that graph-theoretic algorithms often exhibit. For example, the problems of image smoothing, stereo depth reconstruction, and segmentation can be formulated as solving for the minimum cut of a graph. Each of these computer vision problems can be described as global energy minimization problems.

Currently, I am working with Professor Ramin Zabih to apply graph theory techniques to diffusion imaging tractography in the field of biomedical imaging. The white matter of the brain is a complicated 3D network of composed of short connections–"neural tracts"–among different cortical and subcortical regions. These neural tracts are difficult to image directly. Tractography is a technique to graph the neural tracts of the brain, given diffusion fMRI data. In white matter, water tends to diffuse along the axon, away from the cell body. Diffusion fMRI takes advantage of the anisotropic (directional) nature of water diffusion in white matter to measure the direction and magnitude of water diffusion. Diffusion fMRI data can be gathered in vivo. Given measurements of the direction and magnitude of water diffusion at each voxel (unit of volume in 3D) in the brain, tractography is the technique of laying out neural tracts that reflect the diffusion data. Viewing the diffusion as flowing along paths in a graph with voxels as vertices lends itself to graph theoretical approaches to tractography. Tractography basically decomposes the graph into subgraphs that represent individual neural tract bundles.

Tractography algorithms generally fall into two groups: local and global. Local tractography algorithms usually start at some "seed" voxels in the brain, where the direction of neural tracts is known, and extend tracts from voxel to voxel in the direction of maxi-

mum diffusion. One critical issue of local tractography algorithms is that when the tracts are extended to a voxel where the maximum diffusion direction is not obvious, it's hard to figure out where to extend the tracts next; neural tracts can easily go "off track" from cumulative errors of local tractography algorithms. Fewer global algorithms have been developed for tractography because global algorithms are generally slower than local algorithms. However, global algorithms can work around local ambiguity and allow us to incorporate priors. With global algorithms, we can incorporate prior knowledge of how certain parts of the brain are structured and encourage neural tracts to begin and end at the boundaries of white and grey matter, instead of ending inside white matter as local algorithms are prone to do.

I am looking to apply graph theory approaches to the tractography problem; many graph theory algorithms can be used to quickly solve global optimization problems. Currently, I am exploring the idea of a global tractography algorithm that frames the problem as minimum-cover that will not only work around local ambiguities, but be speedy as well. The minimum-cover problem is an NP-hard global "covering" problem that nevertheless has a fast approximation algorithm; given input size n, an approximation algorithm for minimum-cover can produce results that are within factor $\log(n)$ of optimal. Given the brain diffusion fMRI data, I want to find the simplest set of neural tract bundles that explains ("covers") the diffusion fMRI data by iteratively selecting neural tract bundles minimizing a "cost per voxel" measure. Past work by Pedro Felzenswalb and David McAllester provides an efficient method for implementing the inner-loop of the minimum-cover algorithm in the case where the objects we are looking for are smooth curves; an extension of their work to look for neural tract bundles appears promising. I hope to relate my experience in signal processing, probability, and parallelization to implement a tractography algorithm that runs fast enough to meet clinical constraints (current tractography algorithms can run from a day to a month). By formulating a global tractography algorithm, I hope not only to produce qualitatively correct neural tracts, but also to prove theoretical bounds on the correctness of the neural tracts. Furthermore, the tractography problem occurs outside of analyzing neural tracts in cognitive science. For example, we can analyze bird migration and crowd control with tractography algorithms.

My research on applying graph algorithms to computer vision and biomedical imaging will have real-world consequences and be immensely rewarding intellectually. While I am working on a mathematically interesting problem, I am also able to contribute to understanding human intelligence and learning, one of my long term goals. My research into tractography will improve the reconstruction of images of neural tracts in the brain. My research will have far-reaching consequences enabling cognitive scientists to better study the axon connections in the brain and enabling doctors to better diagnose diseases such as Alzheimer's and stroke.