My Ph.D. research is in global seismic tomography, the science of making three-dimensional maps of the Earth's deep interior - essentially mapping temperature variations. Thermally driven motions shape the face of our planet and ultimately enable life on Earth. Understanding this "convection machine" is therefore of fundamental interest, and seismic imaging is the most direct approach to the problem. I am developing a method that for the first time allows us to systematically exploit the frequency dependence of seismic body waves.

Our setup is analogous to medical brain imaging. In both cases, numerous wave sources and receivers are placed on the surface of a sphere. If a receiver records a waveform that is distorted compared to what a simple background model predicts, then the wave must have encountered an anomaly somewhere on its way. From a sufficiently dense web of criss-crossing rays, a three-dimensional map of such anomalies can be reconstructed by solving a very large inverse problem. But in contrast to medical imaging, we have no control over our wave sources. Only large earthquakes generate waves powerful enough to travel through the entire body of the Earth; such waves are termed "body waves."

The wavelengths of body waves are of the same order as the dimensions of the objects that we are most interested in. Those include relatively thin former tectonic plates that are subducting back into the mantle and narrow plumes of hot ascending rock that bring up primordial heat. Constrained by computational limitations, seismologists have long been using optical ray theory to model wave propagation. In an obvious modeling mismatch, ray theory assumes that wavelengths are very short compared to the scale of our features of interest; the effect of this mismatch is to blur or render invisible smaller features. New theoretical work by the group at Princeton alleviates this shortcoming by explicitly taking into account the wavelength of seismic waves. A major development effort has yielded computationally efficient tomography software based on this theory.

In order to take full advantage of the more powerful method, data needs to be measured in a much more sophisticated way than the traditional "picking" of a wave's arrival time. The subject of my Ph.D. research is to develop these new measurement techniques and to generate the first data set tailored to the theory. I am also the first to systematically estimate anomalies of wave amplitudes, which is a much more delicate measurement than travel times but delivers complementary information. The central goal is to measure travel times and amplitudes in arbitrary frequency bands instead of just one, thus multiplying the amount of data. Wavelength discriminates between objects of different sizes; short waves are distorted by large and small anomalies, whereas long waves hardly sense small objects.

I spent my first two years developing and making the measurements robust; seismic data comes in very mixed quality, and automating is essential to handle hundreds of thousands of measurements. The great technical challenge is to estimate and remove the strong and complicated individual signature of each earthquake, so that the weak signal of mantle anomalies may emerge. In my third year I did a proof-of-concept study for the Western U.S.; it showed that amplitudes and travel times vary significantly as a function of frequency, and that they follow theoretically predicted patterns. I am currently building a tomography-sized data set for the Western U.S., and I am analyzing its statistics. Soon I will be starting to image the mantle under the Western U.S. while continuing to build a data set with global coverage. The final step of my Ph.D. project will be a global inversion for Earth's mantle structure.