Faster Full-Waveform Inversion

Andreas Fichtner[1], Dirk-Philip van Herwaarden[1], Sölvi Thrastarson[1], Martin van Driel[1], Christian Boehm[1], Michael Afanasiev[1], [1]ETH Zurich, Institute of Geophysics, Zurich, Switzerland

Background, Motivation and Objective

There is an increasing interest in retrieving quantitative information about the mechanical properties of tissues and objects being imaged using ultrasound. To retrieve this quantitative information from the measured data, full-waveform inversion methods are being developed. These methods allow us in principle to exploit the maximum amount of available information, thereby producing images with unprecedented resolution.

While being close to a standard in global seismology and seismic exploration, full-waveform inversion is still seldomly used in medical imaging and non-destructive testing. One of the main reasons is computational cost. Here, we present two methods to significantly reduce the computational cost of full-waveform inversion: (1) wavefield-adapted meshes and (2) adaptive stochastic gradient descent.

Statement of Contribution/Methods

Wavefield-adapted meshes are finite-element meshes that are specifically designed to match the expected geometry of a wavefield. The simplest example are wavefields that are roughly spherical, with shorter-wavelength variations in the propagation direction and longer-wavelength variations perpendicular to the propagation direction. Using this prior expectation, we may design meshes that are approximately spherically symmetric around the source and composed of elements that are stretched in azimuthal direction. In a series of synthetic scenarios we demonstrate that wavefield-adapted meshing may reduce the effective dimensionality of the problem by up to 1, depending on the smoothness of the medium. Thus, an original 2-D problem can almost be solved at the cost of a 1-D problem. Most importantly, a full-waveform inversion workflow based on wavefield-adapted meshes has the potential to reduce the computational requirements by up to an order of magnitude.

Adaptive stochastic gradient descent uses a mini-batch approach to approximate the gradient for all wavefield sources by the gradient for a smaller, quasi-randomly chosen subset of sources. The size of the subset needed to achieve an acceptable gradient approximation is estimated on-the-fly, so that an optimally small number of sources can be used in each iteration. Using the example of a 3-D elastic full-waveform inversion, we demonstrate that the mini-batch approach requires only around 20 % of the computational resources of a traditional full-waveform inversion, where all sources are used in each iteration.

Results/Discussion

In combination, the wavefield-adapted meshes and the adaptive stochastic gradient descent have the potential to accelerate full-waveform inversion substantially, and to allow for the incorporation of higher frequencies and a larger number of sources. Though our synthetic and real-data examples are from seismology where wave propagation is visco-elastic, the translation of our methods to ultrasound seems straightforward.