Using JAX to develop cycle-skipping resistant full-waveform inversion techniques for ultrasound tomography

Alex Klaben[a], Melanie Singh[a], Nebojsa Duric[a,b], and Rehman Ali[b]

[a]Department of Biomedical Engineering, University of Rochester, Rochester, NY, USA
[b]Department of Imaging Sciences, URMC, Rochester, NY, USA

ABSTRACT

Full-waveform inversion (FWI) is limited by its dependence on accurate starting models and sufficiently low-frequency signals to avoid cycle skipping. This motivates the many cycle-skipping robust alternatives proposed in the literature, including frequency-difference full-waveform inversion (FD-FWI) and adaptive waveform inversion (AWI). In this work, the JAX Python library was used to implement and compare these methods more rapidly than would otherwise be possible. Using JAX, the gradient of an objective function can be found using automatic differentiation, removing the need for analytically derived gradients. FWI, FD-FWI, and AWI were implemented using JAX, and each method was used to reconstruct images of the breast and of the brain through the skull. The data for these images was synthetically created using k-Wave, and a separate time-domain wave-equation solver was used in the reconstruction process. In breast imaging, both FD-FWI and AWI produced more accurate reconstructions than FWI when applied to k-Wave data produced with a 400 kHz center frequency and a 50% bandwidth, with AWI producing the most accurate reconstruction. Relative to FWI, FD-FWI reduced the root-mean-square (RMS) error by 16.3 m/s while AWI reduced the RMS error by 39.4 m/s. In brain imaging, using a 100 kHz center frequency and a 50% fractional bandwidth, FD-FWI and AWI produced more accurate reconstructions than FWI. Relative to FWI, FD-FWI reduced the RMS error by 66.0 m/s while AWI reduced the error by 98.1 m/s. Both accurately recovered the bulk sound speed and two hemorrhages in the brain. Using a 200 kHz center frequency and a 50% fractional bandwidth, FD-FWI increased the RMS error by 20.3 m/s and AWI reduced the RMS error by 85.5 m/s. These results indicate that in many cases FD-FWI is more robust to cycle-skipping than FWI, while AWI is consistently more robust than either FWI or FD-FWI. Additionally, AWI was found to be robust to noise with signal-to-noise ratios of 40 dB or higher.

Keywords: JAX, cycle skipping, full waveform inversion, adaptive waveform inversion, frequency-differencing

1. INTRODUCTION

One of the main limitations of ultrasound tomography (UST) is cycle skipping.[1–3] While UST has been used clinically for breast imaging, applying UST to other contexts has been limited by cycle skipping. For example, UST has the potential to be a radiation-free and cost-effective alternative to CT and MRI scans for brain imaging,[4] but cycle-skipping robust reconstruction methods are needed for this to be possible. Another limit of UST implemented using standard full-waveform inversion (FWI) is the need for low frequency signals that most B-mode ultrasound transducers lack.[1] Many methods for cycle-skipping robust reconstructions have been proposed, but direct comparisons between these methods are lacking in the literature. Using JAX for automatic differentiation allows for quick development of image reconstruction algorithms that utilize different objective functions than FWI. This makes it possible to determine how the proposed methods behave in comparison to each other when tested under the same conditions. Comparing the proposed methods to each other is crucial for determining which methods are suitable for clinical translation, especially in realistic scenarios involving tissue heterogeneity, noise, and modeling errors.

Further author information: (Send correspondence to Rehman Ali)
Rehman Ali: E-mail: Rehman_Ali@URMC.Rochester.edu

Medical Imaging 2026: Ultrasonic Imaging and Tomography, edited by Christian Boehm, Mohammad Mehrmohammadi, Shawn Liangzhong Xiang, Proc. of SPIE Vol. 13931, 139310J · 2026 Published by SPIE
1605-7422 · doi: 10.1117/12.3085595

2. THEORY

2.1 Source Estimation

We assume no a-priori knowledge about the transmit pulses used to produce the data. Instead, during the inversion process we dynamically estimate the shape and amplitude of each transmit pulse from a starting guess (e.g., Gaussian-weighted sinusoid with unit amplitude). Based on this guess, denote $\vec{p}_{sim,i}(t;\vec{s})$ as the simulated time-domain channel data for the ith transmit (vectorization over receivers) and the corresponding measured data $\vec{d}_i(t)$ to the frequency-domain ($\vec{p}_{sim,i}(\omega;\vec{s})$ and $\vec{d}_i(\omega)$, respectively) by taking fast Fourier transforms (FFT) from time t to frequency ω. For each frequency ω and transmit i, we calculate a single scalar complex number $\gamma_i(\omega;\vec{s})$ that, when multiplied to $\vec{p}_{sim,i}(\omega;\vec{s})$, minimizes the error between $\gamma_i(\omega,\vec{s})\vec{p}_{sim,i}(\omega;\vec{s})$ and $\vec{d}_i(\omega)$:

$$\gamma_i(\omega;\vec{s}) = \frac{[\vec{p}_{sim,i}(\omega;\vec{s})]^H \vec{d}_i(\omega)}{[\vec{p}_{sim,i}(\omega;\vec{s})]^H \vec{p}_{sim,i}(\omega;\vec{s})}\,. \tag{1}$$

Note that $\gamma(\omega;\vec{s})$ is intended to correct for the amplitude and phase at each frequency ω in the ith transmit; however, $\gamma_i(\omega;\vec{s})$ implicitly depends on $\vec{s}$ via the simulation. In practice, we find that this dynamic source adjustment does not introduce any substantial non-uniqueness into the solution. Denote $\vec{p}(\omega;\vec{s})$ as the frequency-domain channel data after applying the frequency-dependent source adjustment:

$$\vec{p}_i(\omega;\vec{s}) = \gamma_i(\omega;\vec{s})\vec{p}_{sim,i}(\omega;\vec{s})\,. \tag{2}$$

We convert $\vec{p}_i(\omega;\vec{s})$ back into the time domain $\vec{p}_i(t;\vec{s})$ by taking an inverse FFT. Both forms of FWI and adaptive waveform inversion (AWI) described in the following sections rely on $\vec{p}_i(t;\vec{s})$ after the source adjustment.

2.2 Time Domain Full-Waveform Inversion

For time domain FWI, the following least-squares objective function is used to reconstruct the image:

$$f_{FWI}(\vec{s}) = \frac{1}{2}\sum_{i=1}^{N}\left(\int_0^{T_{max}} \|\vec{p}_i(t;\vec{s}) - \vec{d}_i(t)\|^2\, dt\right). \tag{3}$$

In this equation, $\vec{p}_i(t;\vec{s})$ is the predicted time-domain channel data from the ith transmit as a function of slowness estimate $\vec{s}$, and $\vec{d}_i(t)$ is the observed time-domain channel data corresponding to the same transmission i. To calculate the value of the objective function f at each iteration of FWI, the contribution associated with each transmit was calculated individually and then summed to find the total value. For each transmit, the channel data associated with that transmit was simulated, source estimation was applied (see the previous section 2.1), and then that contribution to f_{FWI} was computed. Defining this objective function in JAX allowed for its gradient to be calculated via automatic differentiation.

2.3 Frequency-Difference Full-Waveform Inversion

In frequency-difference full-waveform inversion (FD-FWI), low-frequency signals are extrapolated from higher frequency signals. The goal of FD-FWI is to transform high-frequency signals that will produce cycle skipping into low-frequency signals that will not produce cycle skipping. Equation 4 below describes the frequency-differencing transformation, as applied to simulated data

$$p_{ij}^{FD}(\Delta\omega;\vec{s}) = \frac{\int_{\omega_{low}}^{\omega_{high}} p_{ij}(\omega+\frac{1}{2}\Delta\omega;\vec{s})\, p_{ij}^*(\omega-\frac{1}{2}\Delta\omega;\vec{s})\, d\omega}{\epsilon^2 + \int_{\omega_{low}}^{\omega_{high}} p_{ij}(\omega-\frac{1}{2}\Delta\omega;\vec{s})\, p_{ij}^*(\omega-\frac{1}{2}\Delta\omega;\vec{s})\, d\omega}, \tag{4}$$

where we undo the vectorization over the receivers (indexed by j), $\Delta\omega$ is the frequency difference, ϵ^2 is a small regularization value, and $p_{ij}^{FD}(\Delta\omega;\vec{s})$ is a vector of simulated frequency-difference pressure values. Equation 5 describes the frequency-differencing transformation, as applied to observed data

$$d_{ij}^{FD}(\Delta\omega) = \frac{\int_{\omega_{low}}^{\omega_{high}} d_{ij}(\omega+\frac{1}{2}\Delta\omega)\, d_{ij}^*(\omega-\frac{1}{2}\Delta\omega)\, d\omega}{\epsilon^2 + \int_{\omega_{low}}^{\omega_{high}} d_{ij}(\omega-\frac{1}{2}\Delta\omega)\, d_{ij}^*(\omega-\frac{1}{2}\Delta\omega)\, d\omega}. \tag{5}$$

Next we vectorize the frequency-differenced quantities $p_{ij}^{FD}(\Delta\omega;\vec{s})$ and $d_{ij}^{FD}(\Delta\omega)$ in equations 4 and 5 over the receivers (previously indexed by j) as $\vec{p}_i^{FD}(\Delta\omega;\vec{s})$ and $\vec{d}_i^{FD}(\Delta\omega)$, respectively. Similar to FWI, FD-FWI uses the L2 norm of the difference between measured and simulated data as the objective function. However, in FD-FWI, the measured and simulated data have both been transformed through frequency-differencing before the L2 norm is taken:

$$f_{FD}(\vec{s};\Delta\omega) = \frac{1}{2}\sum_{i=1}^{N} \|\vec{p}_i^{FD}(\Delta\omega;\vec{s}) - \vec{d}_i^{FD}(\Delta\omega)\|^2\,, \tag{6}$$

$$f_{FD}^{total}(\vec{s}) = \sum_{\Delta\omega} f_{FD}(\vec{s};\Delta\omega). \tag{7}$$

FD-FWI seeks to update $\vec{s}$ such that the value of f_{FD}^{total} is minimized. For each transmit (indexed by i), the channel data associated with that transmit was simulated, and source estimation was applied (see section 2.1). Frequency differencing was applied to both simulated and measured signals prior to calculating f_{FD} (i.e., the L2 norm of the difference between them). In prior work,[1] frequency differencing was only applied to the measured data, while the forward model was a conventional frequency-domain FWI forward model operating at a single frequency $\Delta\omega$. In this work, FD-FWI applies the same frequency differencing process to both measured and simulate data to eliminate any source of model mismatch. Note that f_{FD} is calculated and summed over multiple frequency differences $\Delta\omega$ to produce the overall objective function f_{FD}^{total}. Defining this objective function f_{FD}^{total} in JAX allowed for its gradient to be calculated via automatic differentiation.

2.4 Adaptive Waveform Inversion

While FWI uses the L2 norm of the difference between measured and simulated data as the objective function, AWI operates on the filter found by deconvolving the simulated data by the measured data, or vice versa.[3] AWI then aims to push this filter towards a zero time-lag impulse function. We found that working with the filter that deconvolves the simulated data by the measured data is more stable than using the filter that deconvolves the measured by the simulated. The following time-domain convolution relates the filters $v_{ij}(t,\vec{s})$ to the simulated data $p_{ij}(t;\vec{s})$ and the measured data $d_{ij}(t)$:

$$p_{ij}(t;\vec{s}) = d_{ij}(t) * v_{ij}(t;\vec{s})\,. \tag{8}$$

In equation 8, $*$ denotes a time-domain convolution. Thus, each individual receive channel is associated with its own filter that we would like to drive towards a zero-lag unit impulse. In the frequency domain, the convolution in equation 8 becomes multiplication:

$$p_{ij}(\omega;\vec{s}) = d_{ij}(\omega)\, v_{ij}(\omega;\vec{s})\,. \tag{9}$$

The regularized deconvolution in the frequency domain with regularization parameter ϵ is

$$v_{ij}(\omega;\vec{s}) = \frac{p_{ij}(\omega;\vec{s})\, d_{ij}^*(\omega)}{\epsilon^2 + d_{ij}(\omega)\, d_{ij}^*(\omega)}\,. \tag{10}$$

The objective function for AWI is applied to each individual trace. Denote $v_{ij}(t;\vec{s})$ as the filter for transmit i and receive channel j obtained by the regularized deconvolution described in equation 10. We can then introduce a weighting function $T(t)$ that is either minimal or maximal at $t=0$. If we vectorize $\vec{v}_{ij}(t;\vec{s})$ over time t as $\vec{v}_{ij}(\vec{s})$ and diagonalize $T(t)$ into diagonal matrix T, the full AWI objective function $g(\vec{s})$ summed over all the traces is:

$$f_{AWI}(\vec{s}) = \frac{1}{2}\left(\sum_{i=1}^{N}\left(\sum_{j=1}^{N}\frac{\int_{-\infty}^{\infty}|T(t)v_{ij}(t;\vec{s})|^2\,dt}{\int_{-\infty}^{\infty}|v_{ij}(t;\vec{s})|^2\,dt}\right)\right) = \frac{1}{2}\left(\sum_{i=1}^{N}\left(\sum_{j=1}^{N}\frac{\|T\vec{v}_{ij}(\vec{s})\|^2}{\|\vec{v}_{ij}(\vec{s})\|^2}\right)\right)\,. \tag{11}$$

If $T(t)$ is minimal at $t=0$, the objective $f_{AWI}(\vec{s})$ should be minimized, effectively penalizing filter content at large time lag. If $T(t)$ is maximal at $t=0$, the objective $f_{AWI}(\vec{s})$ should be maximized, effectively rewarding filter content at low time lag. We have found that rewarding low time-lag content better captures the bulk sound speed profile than penalizing content at large time lag (see section 3.4 for more details).

3. METHODS

3.1 Synthetic Data

Multistatic data (all receive channels for each single-element transmission) was simulated using the k-Wave toolbox.[5] The channel data was simulated for 100- and 400-kHz-center-frequency transmits in two imaging cases: 1) breast tissue with a high sound speed lesion, 2) a brain with two intraparenchymal hemorrhages through the skull. All channel data was produced using 256 transducer elements in a 22 cm diameter ring, as seen in Figure 1. All transducers had a 50% fractional bandwidth at -6 dB and transmitted a gaussian-weighted sinusoid.

The original channel data was produced by k-Wave. This original data was used for certain image reconstructions as a proof-of-concept, while the same channel data with added white noise was used for other image reconstructions to validate robustness in the presence of noise. To simulate noisy data, an FFT was applied to each channel of the original k-Wave data, a normally distributed random value was added at each frequency for each channel, and then an inverse FFT was applied to return the data to the time-domain. The noise level was determined by modifying the standard deviation of the normal distribution such that the noise-level was below a fixed fraction of the peak signal value in the frequency spectrum as determined by a desired signal-to-noise ratio (SNR). SNR values of 30, 40, and 50 decibels were used to determine the amount of noise added to the k-Wave simulated data.

3.2 General Process for Image Reconstruction

For each image reconstruction method, a JAX-compiled Python function calculates the value of the objective function, given an estimated slowness profile and measured channel data. Inside this function, the predicted channel data is simulated using a JAX-compiled time-domain wave-equation solver (distinct from k-Wave) based on the estimated slowness profile. Then, the objective function is calculated based on this predicted data and the provided measured data. By writing these functions using JAX, the gradient of the objective function with respect to the slowness profile can be found using automatic differentiation. This gradient is then used to make iterative updates to the slowness profile such that the objective function is minimized.

In this work, the optimization method used in the image reconstructions was a Limited Memory Broyden–Fletcher–Goldfarb–Shanno algorithm (LBFGS), which was provided by the JAXopt library in Python. JAXopt is an optimization library built on top of JAX. For each image reconstruction, the initial starting model was a constant 1500 m/s sound speed, which corresponds to a constant slowness of $6.667 * 10^{-4}$ s/m. We ran repeated iterations of LBFGS until sufficient convergence was reached. Each sound speed reconstruction was compared to the ground truth using the root-mean-square (RMS) error inside an 18 cm ring centered inside the ring array. The ground truth sound speed profiles can be seen in Figure 1. Sections 3.3 and 3.4 below provide additional details regarding the FD-FWI and AWI reconstruction methods. FWI does not require any further details beyond the general process described here.

3.3 Frequency-Difference Full-Waveform Inversion: Details and Parameter Choices

For FD-FWI, several parameters must be chosen. The range of frequencies to extract from the channel data must be chosen (that is, ω_{high} and ω_{low} in equations 4 and 5). In this work, ω_{high} and ω_{low} were chosen to match the 50% fractional bandwidth at -6 dB of the simulated transducers. For example, with data produced using a 100 kHz center frequency, ω_{high} would be 125 kHz and ω_{low} would be 75 kHz, corresponding to a 50 kHz bandwidth. Note that the frequencies extracted depend on $\Delta\omega$ as well. The highest frequency extracted is $\omega_{high} + \frac{1}{2}\Delta\omega$ and the lowest frequency extracted is $\omega_{low} - \frac{1}{2}\Delta\omega$. It is important that both the measured and predicted channel data have a high SNR for all extracted frequencies. In addition to the range of frequencies to extract from the data, the range of frequency-differences to include in the objective function (equation 7) must also be chosen. Including very low frequency-differences supports the ability of FD-FWI to avoid cycle skipping, while including higher frequency-differences increases the resolution of the image that FD-FWI converges to. An

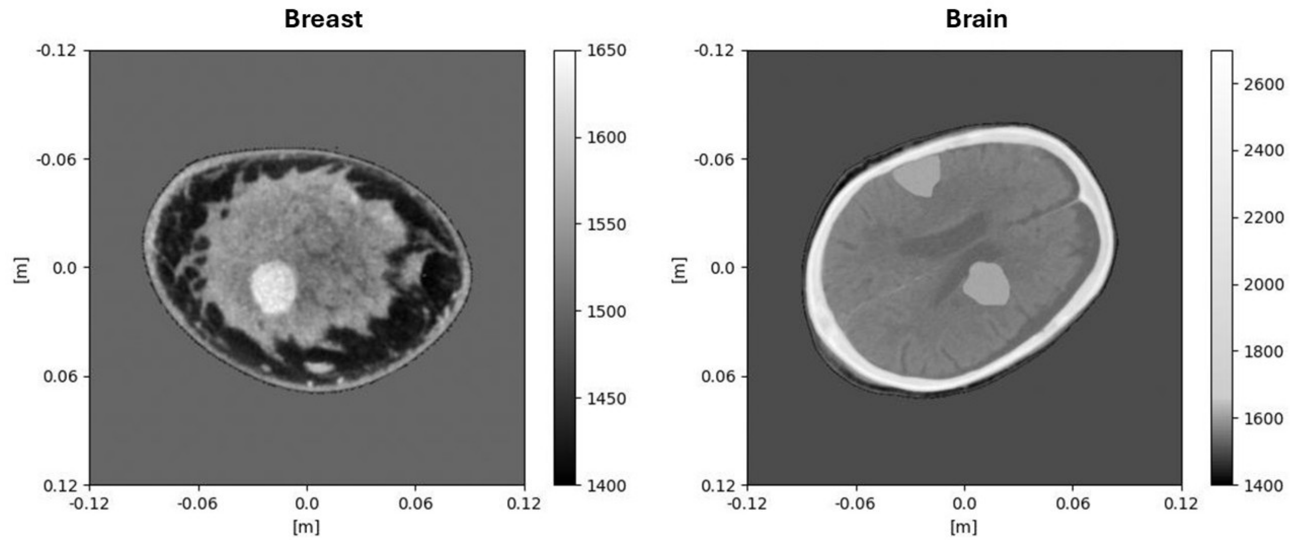

Figure 1. Breast and brain ground truth sound speed profiles, in meters per second. The breast phantom includes a high sound speed lesion. The brain phantom includes the skull, which in this phantom reaches sound speeds as high as 2697 m/s, as well as two intraparenchymal hemorrhages. For both phantoms, the background sound speed was 1500 m/s.

additional consideration is that the greatest frequency-difference $\Delta\omega$ must be small enough such that the highest and lowest frequencies extracted from the measured and predicted data have a high SNR.

For the 400 kHz transmit frequency breast imaging case in this work, 81 evenly-distributed frequency-differences were included in the FD-FWI objective function (equation 7). These frequency-differences ranged from 50 kHz to 150 kHz, while ω_{low} was 300 kHz and ω_{high} was 500 kHz. For the 100 kHz brain imaging case, 81 evenly-distributed frequency-differences were included, ranging from 5 kHz to 35 kHz, with an ω_{low} value of 75 kHz and an ω_{high} value of 125 kHz. For the 200 kHz brain imaging case, 101 evenly-distributed frequency-differences were included, ranging from 5 kHz to 75 kHz, with an ω_{low} value of 150 kHz and an ω_{high} value of 250 kHz. The regularization value ϵ in equations 4 and 5 must be determined as well. In this work, an ϵ value of $1 * 10^{-5}$ was used for all FD-FWI reconstructions.

3.4 Adaptive Waveform Inversion: Details and Parameter Choices

In the work we chose $T(t)$ in equation 11 to have a Gaussian shape maximum at $t = 0$ with a full-width at half maximum of 10.2 microseconds. We found that choosing $T(t)$ to have a minimum at $t = 0$ and maximizing the AWI objective function was more effective at capturing the bulk sound speed than choosing $T(t)$ to have a minimum at $t = 0$ and minimizing the AWI objective function. We have found the choice of regularization value ϵ has a large effect on the gradient of the objective function; ϵ must be large enough to minimize the effect of noise on the deconvolution while remaining small enough to take advantage of the deconvolution process. Note that for large ϵ, the deconvolution would simply become the cross-correlation function. In this work, for data that included noise, ϵ was calculated by dividing the peak time-domain signal value by the assumed SNR of the channel data. However, the most accurate way to calculate ϵ would have been to divide the peak signal value in the frequency domain (rather than the time-domain) by the assumed SNR of the channel data. We briefly discuss the implications of this error in the ϵ calculation on the results shown in Figure 4. For data that did not include added noise, ϵ was chosen to be as small as possible without resulting in divide-by-zero errors in equation 10. In this work, for the data without added noise, an ϵ value of $1 * 10^{-8}$ was found to be effectively above the noise floor introduced by floating-point roundoff errors.

4. RESULTS AND DISCUSSION

Figure 2 shows breast image reconstructions produced using FWI, FD-FWI, and AWI, along with the corresponding RMS error values. For these reconstructions, the measured data was produced in k-Wave using a 400

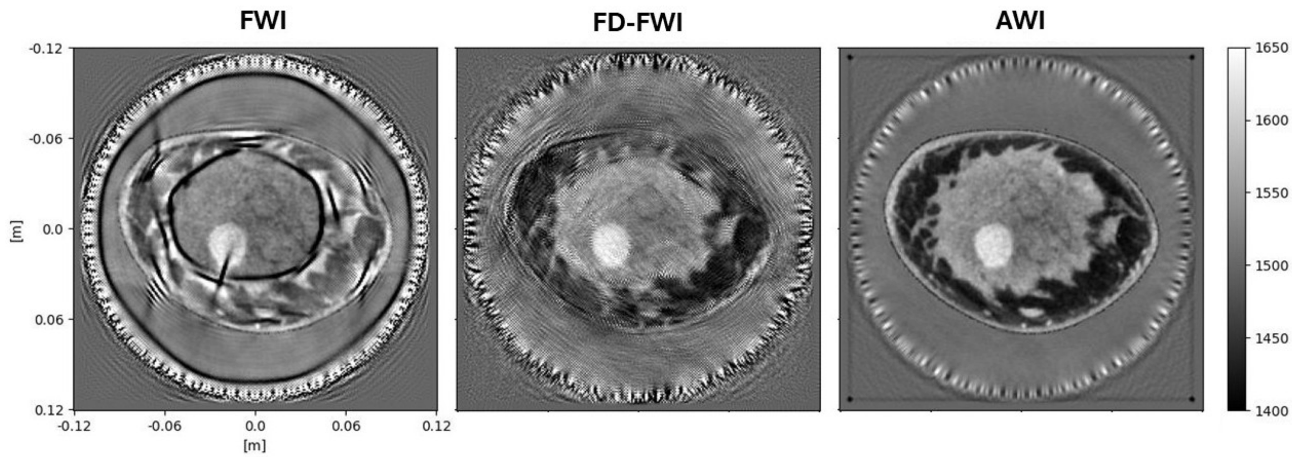

Figure 2. Breast images produced using FWI, FD-FWI, and AWI. The measured k-Wave data for these images was produced using a 400 kHz center frequency and a 50% fractional bandwidth. The RMS error for the FWI image was 44.2 m/s, the RMS error for the FD-FWI image was 27.9 m/s, and the RMS error for the AWI image was 4.8 m/s. Each reconstruction captured the lesion in the breast tissue, but the FWI image includes cycle-skipping artifacts. The AWI reconstruction captured greater detail than the FWI and FD-FWI reconstructions.

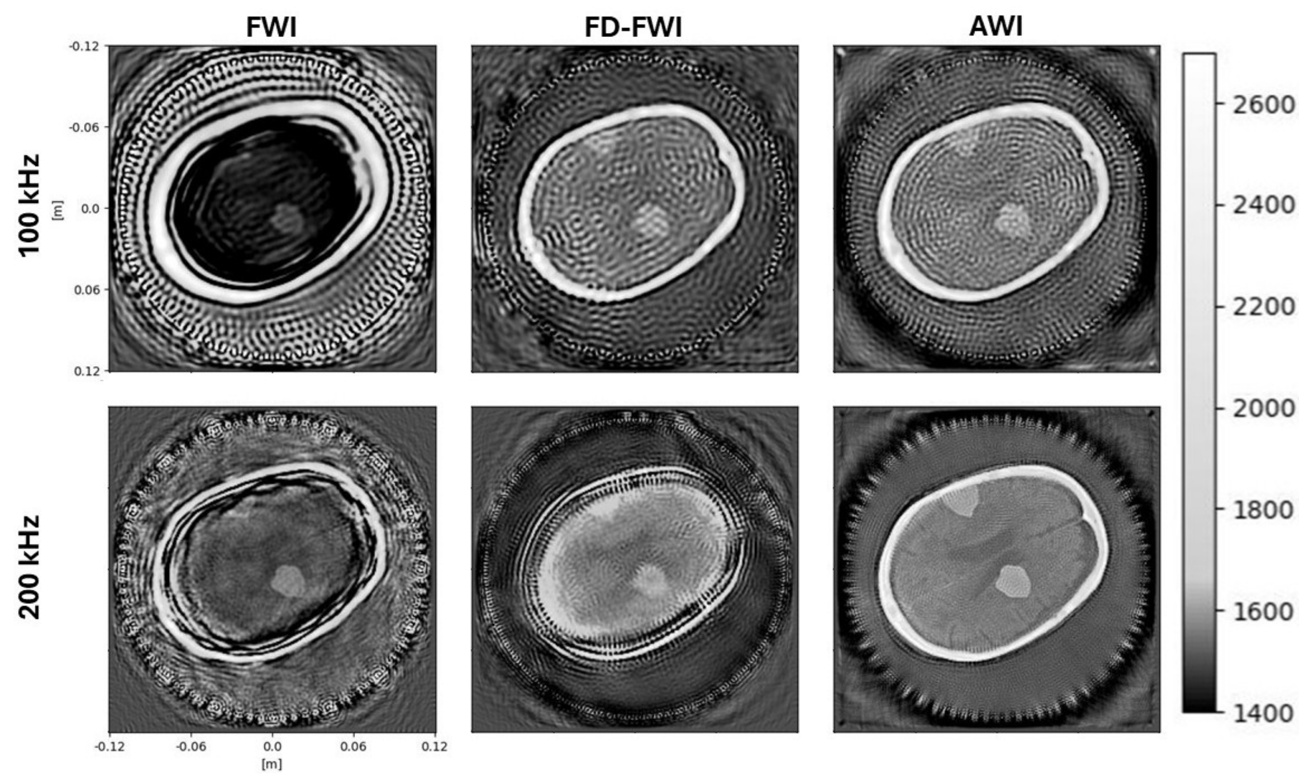

Figure 3. Brain images produced using FWI, FD-FWI, and AWI. The measured k-Wave data was produced using a 100 kHz center frequency (top) or a 200 kHz center frequency (bottom) and a 50% fractional bandwidth. The RMS errors for the 100 kHz center frequency reconstructions are 154.2 m/s for FWI, 88.2 m/s for FD-FWI, and 56.1 m/s for AWI. The RMS errors for the 200 kHz center frequency reconstructions are 120.1 m/s for FWI, 140.4 m/s for FD-FWI, and 34.6 m/s for AWI. In both FWI reconstructions, cycle-skipping artifacts are present and one hemorrhage can be seen while the other is partly or completely obscured by the skull. For FD-FWI, both hemorrhages can be seen for both the 100 kHz and 200 kHz reconstructions, but the skull was only accurately reconstructed while using a 100 kHz center frequency. For AWI, both hemorrhages and the skull were accurately recovered for both center frequencies. The resolution of AWI is much greater for the 200 kHz center frequency data.

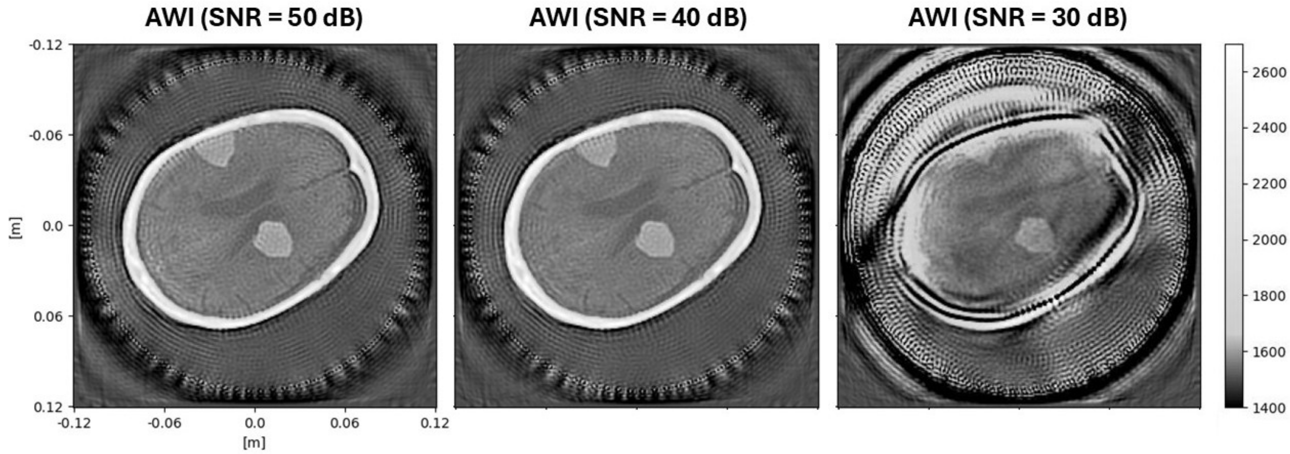

Figure 4. Brain images produced using AWI using k-Wave data that has SNR values of -50, -40, and -30 dB. The measured k-Wave data was produced using a 150 kHz center frequency and a 50% fractional bandwidth. The RMS error for the 50 dB SNR reconstruction is 35.1 m/s, the RMS error for the 40 dB SNR reconstruction is 34.5 m/s, and the RMS error for the 30 dB SNR reconstruction is 148.9 m/s. With an SNR value of 50 dB or 40 dB, AWI accurately produced the hemorrhages and the skull. With an SNR value of 30 dB, AWI accurately reconstructed one hemorrhage, while the other hemorrhage and the skull were inaccurately reconstructed.

kHz center frequency, a 50% fractional bandwidth, and the ground truth breast sound speed profile seen in Figure 1. Under these conditions, FWI produced an image corrupted by cycle-skipping artifacts. Both FD-FWI and AWI produced images free of cycle-skipping artifacts. However, while FD-FWI avoided cycle skipping, it produced a much less accurate image than AWI. Relative to FWI, FD-FWI reduced the RMS error by 16.3 m/s while AWI reduced the RMS error by 39.4 m/s.

Figure 3 shows brain image reconstructions produced using FWI, FD-FWI, and AWI, along with the corresponding RMS error values. The measured data for these reconstructions was produced using a 100 kHz or 200 kHz center frequency, a 50% fractional bandwidth, and the ground truth brain sound speed profile seen in Figure 1. Using a 100 kHz center frequency, FWI produced a highly cycle-skipped reconstruction where only one of the two hemorrhages can be seen. FD-FWI and AWI both produced more accurate reconstructions of the skull and hemorrhages, although a ripple artifact can be seen in both images due to the low-frequencies used in the reconstruction. At this center frequency, relative to FWI, FD-FWI reduced the RMS error by 66.0 m/s while AWI reduced the RMS error by 98.1 m/s. Using a 200 kHz center frequency, FWI produced a cycle-skipped reconstruction of the skull, with one hemorrhage easily seen but the other partly obscured by the cycle-skipped skull. At this frequency, FD-FWI failed to converge to an accurate reconstruction of the skull, though both hemorrhages can still be seen. The AWI reconstruction shows both hemorrhages and the shape of the skull clearly, in addition to the sulci of the brain and the cerebral ventricles. Relative to FWI, FD-FWI increased the RMS error by 20.3 m/s and AWI reduced the RMS error by 85.5 m/s. Notably, the ripple artifact seen in the FD-FWI and AWI 100 kHz center frequency reconstructions is greatly reduced for the FD-FWI and AWI 200 kHz center frequency reconstructions due to the higher frequency content present.

A common theme across the brain reconstructions shown in Figure 3 is that the center of the brain converged either more quickly or more accurately than the region of the brain closest to the skull. In both FWI reconstructions and the 200 kHz center frequency FD-FWI reconstruction, the hemorrhage nearest the skull is at least partially obscured, while the more central hemorrhage is easily seen. In both AWI reconstructions and the 100 kHz center frequency FD-FWI reconstruction, the center of the brain converged to its final state in far fewer iterations than the region immediately inside the skull.

Figure 4 shows brain image reconstructions produced using AWI with varying SNR values. For each of these reconstructions, the measured k-Wave data was produced using a 150 kHz center frequency and a 50% fractional bandwidth. For both the 50 dB SNR and 40 dB SNR reconstructions, the resulting reconstruction clearly shows both hemorrhages, the shape of the skull, the sulci, and the cerebral ventricles. With a 30 dB SNR value, the more central hemorrhage can be seen easily, but the other hemorrhage is somewhat obscured. The skull failed to converge to its true shape. The RMS error for the 50 dB SNR reconstruction is 35.1 m/s, the RMS error for the 40 dB SNR reconstruction is 34.5 m/s, and the RMS error for the 30 dB SNR reconstruction is 148.9 m/s. Similarly to the FWI and FD-FWI results in Figure 3, the center of the brain converged more accurately than the region nearest the skull for an SNR of 30 dB. However, these results may have been affected by the methodology used to calculate the ϵ value based on the SNR. The deconvolution step in AWI is performed in the frequency domain rather than the time-domain; therefore, ϵ represents the expected noise floor in the frequency domain. However, the value of ϵ was determined here based on the peak of the time domain signal instead. The immediate degradation of the AWI reconstruction from 40 dB to 30 dB SNR may be the result of an incorrectly calculated ϵ value rather than a true degradation of AWI in response to increased noise.

These results indicate that both FD-FWI and AWI are more robust to cycle-skipping than FWI. However, AWI was consistently able to produce more accurate reconstructions than FD-FWI. In our experience, FWI and FD-FWI reconstructions require much less computation time than AWI reconstructions. Because of this, there may be utility in performing a reconstruction using FD-FWI, then using that reconstruction as a starting model for FWI, in a manner that is similar to the prior work.[1] For example, the FD-FWI breast reconstruction shown in Figure 2 may be close enough to the true sound speed profile where applying an FWI reconstruction would result in a higher resolution image.

Although not shown in this work, AWI does fail to converge to an accurate sound speed profile given a high enough center frequency. For example, AWI would fail to converge to an accurate brain image given measured k-Wave data produced with a 300 or 400 kHz center frequency. Despite AWI being designed to overcome cycle skipping in the presence of high-frequency content, there are practical limitations to how high the center frequency can be before AWI breaks down. The maximum frequency where AWI will converge to an accurate brain image depends on the number of transmit and receive locations. For example, the 200 kHz center frequency brain image produced using AWI in Figure 3 would fail to converge to the true sound speed profile if only 64 transmit locations were used, instead of the 256 used in this work. However, if only 64 transmit locations were used for the 100 kHz center frequency AWI brain reconstruction, the quality of the resulting image would be very similar. This relationship requires further investigation, along with the role that other parameters play in the ability of AWI to reconstruct images given high frequency data. For example, it is possible that the shape of $T(t)$ in equation 11 affects the behavior of AWI when applied to high frequency data.

In this work, we demonstrate that automatic differentiation software can be used to execute FWI and its cycle-skipping robust alternatives without requiring analytically derived gradients. This allows for the development of image reconstruction algorithms to be much more efficient. Analytically deriving gradients can be difficult and error-prone for the complicated objective functions seen in cycle-skipping robust reconstruction methods. Furthermore, the use of automatic differentiation makes modifying the objective function extremely quick; any modifications to the objective function are reflected in the entire optimization process automatically. The efficiency gained from automatic differentiation is particularly useful for comparing multiple reconstruction methods, where many variations of the objective function may need to be implemented.

In addition to the FD-FWI and AWI methods discussed in this work, there are many other cycle-skipping resistant full-waveform inversion techniques that have been proposed in the literature. There is a group of techniques related to AWI, which use a weighting function $T(t)$ to drive the optimization of the slowness profile. The objective function in AWI relies on deconvolutions of the predicted data by the measured data, or vice versa, then normalizing the resulting filters.[3] Related methods are based on deconvolutions without normalization,[6] cross-correlations with normalization,[7] and cross-correlations without normalization.[8] Another group of techniques is related to optimal transport, where each trace of the pressure data is converted to a cumulative distribution function, and then the misfit between measured and predicted cumulative distributions is calculated.[9] These methods seek to minimize the misfit between cumulative distribution functions. Sources of variation within this group of techniques include the method used to convert a pressure trace to a cumulative distribution function and the method of calculating the misfit (e.g., the quadratic Wasserstein metric, the 1-Wasserstein metric, and the normalized integration method).[9]

5. CONCLUSIONS

This work demonstrates that automatic differentiation software like JAX can simplify the implementation of proposed UST reconstruction methods. Automatic differentiation allows for much quicker prototyping of image reconstruction methods. The AWI and FD-FWI implementations created with JAX were shown to produce cycle-skipping robust image reconstructions even in situations where standard FWI is corrupted by cycle skipping; however, AWI consistently yielded a more accurate and more robust sound speed reconstruction than FD-FWI in every scenario. In the future, JAX could be used to develop prototypes of additional methods that have been proposed for cycle-skipping robust UST.

ACKNOWLEDGMENTS

This work was supported by the National Institute of Biomedical Imaging and Bioengineering (NIBIB) from the National Institute of Health (NIH) under grants F32-EB034589 and K99-EB037080 (Contact PI: Rehman Ali). Alex Klaben was supported by the Schwartz Discover Grant and the Researh Presentation Award from the Office of Undergraduate Research at the University of Rochester. Furthermore, Alex Klaben's attendance was made possible via the SPIE Medical Imaging 2026 Student Conference Support.

REFERENCES

[1] Ali, R., Mitcham, T., Owolabi, I., McConnell, S., and Duric, N., "Frequency-differencing strategy to kickstart full-waveform inversion without cycle skipping," *JASA Express Letters* **5**, 012001 (01 2025).

[2] Boehm, C., Krischer, L., Ulrich, I., Marty, P., Afanasiev, M., and Fichtner, A., "Using optimal transport to mitigate cycle-skipping in ultrasound computed tomography," in [*Medical Imaging 2022: Ultrasonic Imaging and Tomography*], Bottenus, N. and Ruiter, N. V., eds., **12038**, 1203809, International Society for Optics and Photonics, SPIE (2022).

[3] Warner, M. and Guasch, L., "Adaptive waveform inversion: Theory," *Geophysics* **81**, R429–R445 (09 2016).

[4] Lluís, G., Calderón, A. O., Meng-Xing, T., Parashkev, N., and Warner, M., "Full-waveform inversion imaging of the human brain," *NPJ Digital Medicine* **3** (12 2020).

[5] Treeby, B. E. and Cox, B. T., "k-Wave: MATLAB toolbox for the simulation and reconstruction of photoacoustic wave fields," *Journal of Biomedical Optics* **15**(2), 021314 (2010).

[6] [*A Deconvolution-based Objective Function For Wave-equation Inversion*], *SEG International Exposition and Annual Meeting* **2011 SEG Annual Meeting** (09 2011).

[7] Van Leeuwen, T. and Mulder, W., "Velocity analysis based on data correlation," *Geophysical Prospecting* **56**(6), 791–803 (2008).

[8] Van Leeuwen, T. and Mulder, W. A., "A correlation-based misfit criterion for wave-equation traveltime tomography," *Geophysical Journal International* **182**, 1383–1394 (09 2010).

[9] Yang, Y. and Engquist, B., "Analysis of optimal transport and related misfit functions in full-waveform inversion," *Geophysics* **83**, A7–A12 (11 2017).