Target-Following Lagrangian Approach to Sound Speed Estimation in Pulse-Echo Ultrasound

Rehman Ali* and Nebojsa Duric*

*Department of Imaging Sciences, URMC, Rochester, NY, 14620 USA

ABSTRACT

Sound speed estimation in pulse-echo ultrasound typically involves minimizing aberrations over a fixed grid of image points. However, the apparent positions of imaging targets depend on the speed of sound. In other words, each image point on a fixed grid may correspond to different imaging targets as the sound speed estimate changes. This source of nonlinearity becomes especially problematic when applying differentiable beamforming over a fixed imaging grid to estimate the speed of sound in a layered medium. Ideally, sound speed reconstruction would be based on a fixed set of imaging targets rather than a fixed set of image points. From a continuum mechanics perspective, working with a fixed image grid takes a strictly Eulerian approach to sound speed estimation. A Lagrangian approach would follow the image targets as the change in sound speed deforms the image from its initial configuration. By using an adaptive grid of image points that follows the image targets as the sound speed changes, we effectively account for both Eulerian and Lagrangian aspects of the problem. Simulations show that this target-following approach reduces the nonlinearity of the inverse problem of sound speed estimation in layered media.

Keywords: Sound speed estimation, differentiable beamforming, continuum mechanics, Lagrangian approach

1. INTRODUCTION

The goal of sound speed estimation in medical ultrasound is to measure the sound speed of human tissue in an accurate and spatially-resolved manner. Most sound speed estimators[1–3] rely on an approximate linear relationship between the spatial distribution of slowness (reciprocal of sound speed) in the tissue and aberration delays (or phases) between partial images at each pixel location. A path-length matrix encodes the line-integrals over the slowness distribution corresponding to each aberration delay measurement and is used to reconstruct the sound speed map using regularized inversion. However, there are underlying nonuniqueness issues caused by the velocity-depth ambiguity.[4] Different initial sound speed guesses can affect the initial placement of imaging targets. Even in the presence of strong regularization, sound speed estimates tend to improve focusing around the initial placement of the imaging targets and remain biased by the initial sound speed guess. In this work, we explicitly demonstrate the underlying nonlinearity of the sound speed estimation process caused by the velocity depth ambiguity over a fixed imaging grid.

Recent work on differentiable beamforming[5,6] makes the nonlinearity of sound speed estimation more explicit. Typically, the estimation of the aberration delay is done separately from the regularized inversion of the path-length matrix. However, in a differentiable beamforming setup, there is no separate measurement of the aberration delays. Instead, an objective function is defined using the differences between the partial images. The gradient of this objective function directly backprojects errors between these partial images along the corresponding ray paths, and a gradient-descent algorithm is used to converge to the final sound speed estimate, while minimizing aberrations in the image. In our previous work,[6] we noted the tendency of the differentiable beamformer to resist changes to the sound speed map that would require imaging targets to move from their initial placements in the image. Bulk and layer-wise sound speed changes require imaging targets to move from their initial locations, while laterally varying sound speed changes tend to improve focusing at current image placements. In order to isolate the nature of this problem, this work applies differentiable beamforming to a layered medium model in which sound speed only varies with depth.

Further author information: (Send correspondence to Rehman Ali)
Rehman Ali: E-mail: Rehman$_{A}li@URMC.Rochester.edu$

Medical Imaging 2026: Ultrasonic Imaging and Tomography, edited by Christian Boehm, Mohammad Mehrmohammadi, Shawn Liangzhong Xiang, Proc. of SPIE Vol. 13931, 1393108 · 2026 Published by SPIE
1605-7422 · doi: 10.1117/12.3085613

An important application of layer-wise sound speed estimation is fat-fraction quantification for fatty liver disease.[7,8] Prior work has used coherence maximization to estimate the optimal beamforming sound speed as a function of depth.[9] In layered media, we can invert the closed-form relationship between the local sound speed and the optimal beamforming sound speed to reconstruct the local sound speed profile as a function of depth. However, we typically maximize coherence over a fixed set of image points, which may correspond to different imaging targets as we vary the beamforming speed of sound. Recent work demonstrates the importance of maximizing the coherence over a fixed set of imaging targets rather than a fixed set of image locations.[10] Over the axial dimension of the image, tracking a fixed set of image targets corresponds to equal spacing in time rather than depth.[11] In general, the image grid should deform both laterally and axially as the sound speed changes. We examine how 1D and 2D deformations of the imaging grid impact the nonlinearity of sound speed estimation. From a continuum mechanics perspective, the adaptive imaging grid induces both Eulerian (i.e., fixed-grid) and Lagrangian (i.e., target-following) terms in a differentiable beamformer. The purpose of this work is to demonstrate a generalized substantive derivative formulation (i.e., with both Eulerian and Lagrangian terms) of sound speed estimation that mitigates the nonlinearity of the inverse problem.

2. BACKGROUND

2.1 Differentiable Delay-and-Sum Beamforming

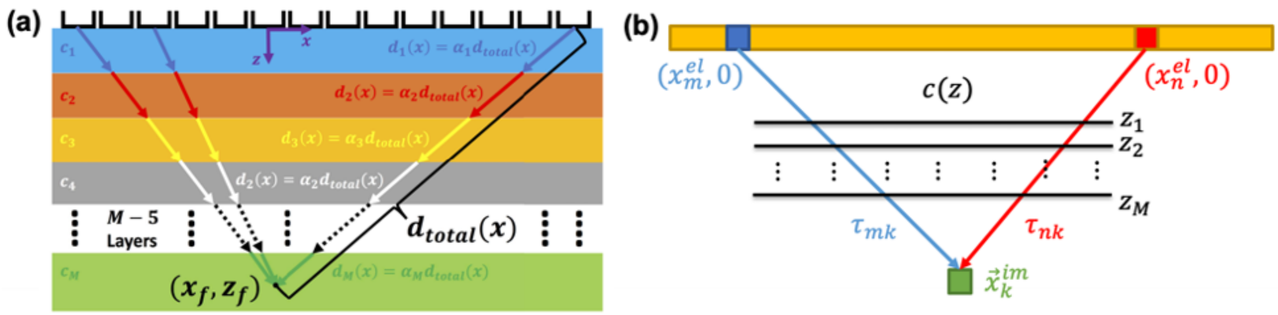

Figure 1. Differentiable Beamforming for Sound Speed Estimation in Layered Media. (a) Layer-wise sound speed model used to calculate beamforming delays. (b) Assumed straight ray paths through layered medium model used to calculate two-way transmit-receive delays.

The differentiable beamforming model[5,6] requires a multistatic synthetic aperture dataset $\tilde{p}_{mn}(t)$ indexed by single-element transmit m and receiver n as a function of time t. For an N-element array, let τ_{lk} refer to the time delay from element l to an image point $\vec{x}_k^{im}$. We assume a layered medium model (Figure 1(a)) in which the sound speed $c(z)$ varies as a function of depth z, which we vectorize as $\vec{c}$. As per Figure 1(b), we can focus the full matrix of transmit and receive signals at each image point $\vec{x}_k^{im}$ by applying focusing delays:

$$p_{mn}(\vec{x}_k^{im}; \vec{c}) = \tilde{p}_{mn}(\tau_{mk}(\vec{c}) + \tau_{nk}(\vec{c})), \tag{1}$$

Summation over receivers yields complex-valued images for each single-element transmit:

$$I_m(\vec{x}_k^{im}; \vec{c}) = \sum_{n=1}^{N} p_{mn}(\vec{x}_k^{im}; \vec{c}). \tag{2}$$

2.2 Coherence Maximization

To estimate the layer-wise sound speed $\vec{c}$, we maximize the coherence factor (CF):[12]

$$CF(\vec{c}) = \frac{\left|\sum_{m=1}^{N} I_m(\vec{x}_k^{im}; \vec{c})\right|^2}{N \sum_{m=1}^{N} \left|I_m(\vec{x}_k^{im}; \vec{c})\right|^2}. \tag{3}$$

The layered medium model, differentiable beamformer, and CF as an objective function were implemented using the automatic differentiation software known as JAX.[13] The conjugate gradient algorithm (50 iterations total) was used to maximize $CF(\vec{c})$ based on gradients calculated by JAX.

2.3 Target-Following Image Grid: 1D Axial Correction

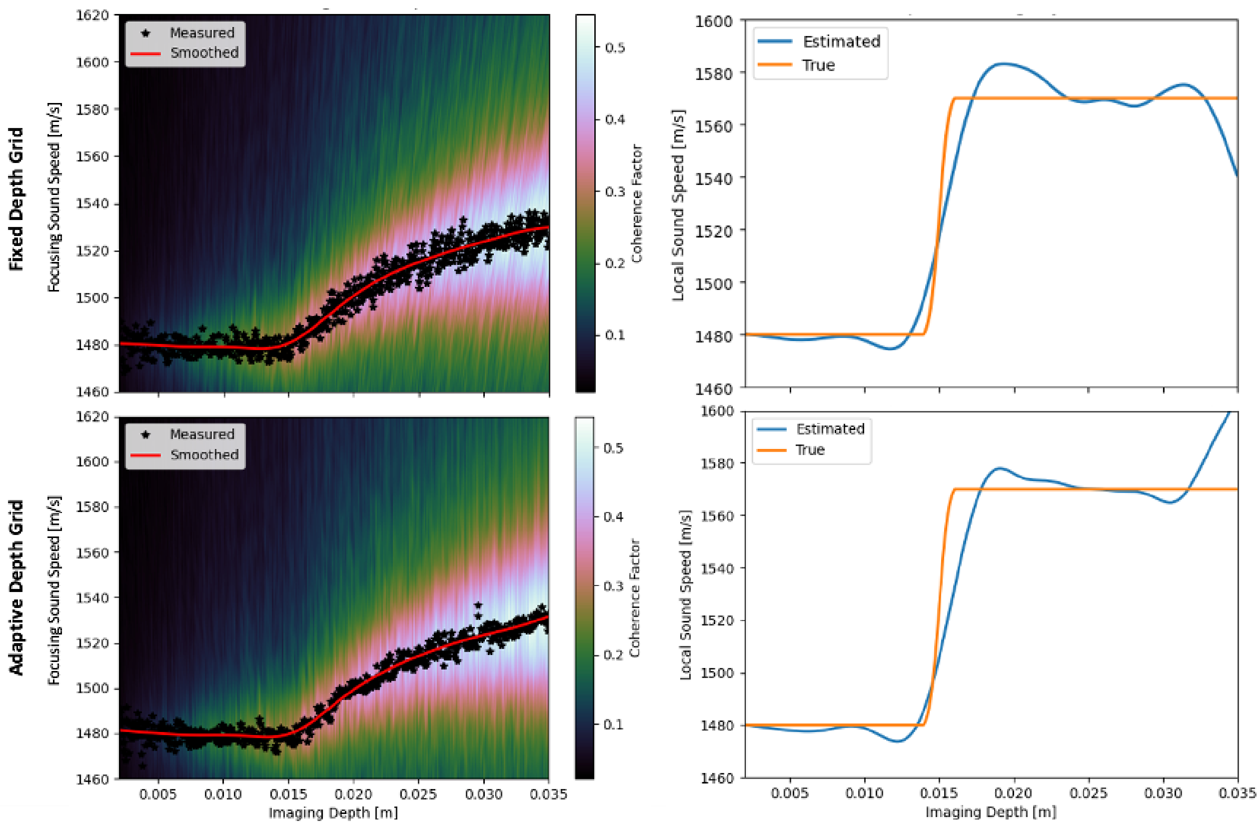

Figure 2. Impact of Fixed vs Adaptive Depth Grid on Sound Speed Estimation in Layered Media. The top row of plots shows slanting contours in the coherence factor as function of fixed depth and focusing sound speed, resulting in high variance in the focusing and local layer-wise estimates of the speed of sound. The bottom row of plots eliminates this slanting effect by using an adaptive depth grid, which reduces the variance in the focusing and local layer-wise estimates of the speed of sound.

On a fixed uniform image grid, image points are placed uniformly in depth according to $z_k = z_0 + k\Delta z$. However, as we change $\vec{c}$, an image target initially placed at a depth z_k may appear deeper or shallower. Ideally, the value of z_k would change with $\vec{c}$ such that it always corresponds to the same image target. This adaptive depth grid would be equally spaced in time rather than depth according to $z_k = z_0 + \sum_{n=1}^{k} c_n \Delta t$. Figure 2 demonstrates the impact of this adaptive grid[10] on layer-wise sound speed estimation.[9] CF as a function of depth and sound speed appears to follow slanted contours. These slanted contours become flat in the adaptive depth grid, reducing the error in the focusing sound speed and the resulting layer-wise reconstruction.[11] The goal of this work is to understand how the same adaptive imaging grid will impact layer-wise sound speed estimation using a differentiable beamformer.

2.3.1 Estimating Local Layer-wise Sound Speed from the Depth-wise Focusing Sound Speed

One detail that may not be readily obvious is the inversion of the local layer-wise sound speed from the optimal focusing sound speed as a function of depth. When image points were assumed to be uniformly spaced in depth, the inversion from the optimal focusing sound speed $c_{foc}(z)$ to the local layer-wise sound speed $c(z)$[9] is:

$$\frac{1}{c(z)} = \frac{d}{dz}\left(\frac{z}{c_{foc}(z)}\right) \tag{4}$$

Because the adaptive depth grid is uniformly spaced in time, it may be tempting to apply the following approach[14] as a function of the two-way time τ:

$$c(\tau) = \frac{d}{d\tau}\left(\tau c_{foc}(\tau)\right) \tag{5}$$

Both equations (4) and (5) assume that sound travels on straight ray paths, neglecting refraction. However, Dix inversion[15] explicitly accounts for refraction in a layered medium. In the Dix inversion model, the focusing sound speed $c_{foc}(\tau)$ becomes the root-mean-square (RMS) of the local layer-wise sound speed $c(\tau)$ as a function of τ, resulting in the following inversion formula (see simplified proof based on Snell's Law[11]):

$$c(\tau) = \sqrt{\frac{d}{d\tau}\left(\tau c_{foc}^2(\tau)\right)} \tag{6}$$

In Figure 2, equation (4) was used to reconstruct the local sound speed on the fixed depth grid while equation (6) was used to reconstruct the local sound speed on the adaptive depth grid. The main purpose of Figure 2 is to demonstrate the impact of the 1D adaptive imaging grid on previous methodologies.[9,11] However, this work is primarily concerned with direct layer-wise sound speed estimation in a differentiable beamformer. Differentiable beamforming does not use a depth-wise focusing sound speed to reconstruct a layer-wise sound speed; instead, differentiable beamforming directly reconstructs the layer-wise sound speed by maximizing CF everywhere in the image.

2.4 Full 2D Target-Tracking Image Grid: Eulerian vs Lagrangian Contributions

The gradient calculation in a differentiable beamformer involves taking partial derivatives of the images $I_m(x, z; \vec{c})$ with respect to $\vec{c}$.[6] Over a fixed 2D grid of image points (x, z), the gradient is simply $\frac{\partial I_m(x,z;\vec{c})}{\partial \vec{c}}$. However, if the 2D image grid (x, z) changes with sound speed $\vec{c}$, the gradient calculation will involve the full material derivative $\frac{DI_m(x,z;\vec{c})}{D\vec{c}}$, which includes both Eulerian and Lagrangian terms:

$$\frac{DI_m(x,z;\vec{c})}{D\vec{c}} = \frac{\partial I_m(x,z;\vec{c})}{\partial \vec{c}} + \frac{\partial I_m(x,z;\vec{c})}{\partial x}\frac{\partial x}{\partial \vec{c}} + \frac{\partial I_m(x,z;\vec{c})}{\partial z}\frac{\partial z}{\partial \vec{c}} \tag{7}$$

The Eulerian term $\frac{\partial I_m(x,z;\vec{c})}{\partial \vec{c}}$ only considers the partial derivative with respect to $\vec{c}$, while keeping the image grid (x, z) constant. The remaining Lagrangian terms involve the spatial derivatives of the image as the image grid (x, z) changes with $\vec{c}$. In the axial dimension of the image, the closed-form relationshipe between depth and sound speed ($z_k = z_0 + \sum_{n=1}^{k} c_n \Delta t$) determines $\frac{\partial z}{\partial \vec{c}}$. In general, the 2D shift in the image point $(\frac{\partial x}{\partial \vec{c}}, \frac{\partial z}{\partial \vec{c}})$ for a given sound speed change $\Delta\vec{c}$ is determined by fitting the beamforming delays $\vec{\tau}$ at the new sound speed $\vec{c} + \Delta\vec{c}$ to the original beamforming delays at the initial sound speed $\vec{c}$:[10]

$$(\Delta x, \Delta z) = \underset{(\Delta x, \Delta z)}{\arg\min} \frac{1}{2}\left\|\vec{\tau}(x + \Delta x, z + \Delta z, \vec{c} + \Delta\vec{c}) - \vec{\tau}(x, z, \vec{c})\right\|_2^2. \tag{8}$$

3. METHODS

3.1 Evaluating the Nonlinearity of the Inverse Problem

A crucial step in the conjugate gradient algorithm is deciding the step size at each iteration (i.e., how far to move along a given direction). The step size that maximizes the objective function along a given direction is determined by a line search, which evaluates the objective function at multiple points along that direction. To avoid these multiple evaluations, we can maximize a quadratic model based on the gradient and Hessian at the point of evaluation. The deviation between the true objective function and its quadratic approximation is a strong indicator of the nonlinearity of the inverse problem. Therefore, we measure the ratio of the step size calculated via line search (in our case, golden-section search) to the approximate step size calculated based on a locally quadratic model. Our prior work[6] notes that when image targets are required to change their depth locations on a fixed grid, the quadratic approximation often results in extremely conservative step sizes (i.e., large step-size ratios). This work shows that an adaptive image grid reduces the discrepancy between step sizes obtained by a true line search and a quadratic approximation (i.e., step size ratio closer to 1).

3.2 Imaging Studies: Simulation and Experiment

The k-Wave[16] simulations used in prior work[9] are reused here to assess sound speed estimation in layered media using a differentiable beamformer with different choices of imaging grid (e.g., fixed, 1D adaptive, 2D adaptive). Figure 2 shows one of three simulations used, where first 15 mm depth into the medium has a sound speed of 1480 m/s, and the remainder of the medium can be 1520, 1540, or 1570 m/s sound speed (1570 m/s shown). Publicly available experimental phantom data,[17] consisting of a slab of chicken was placed on top of an ATS Model 549 phantom (1460 m/s), is shown to verify the findings made from the simulations. In all imaging cases (simulation and experiment), an initial uniform sound speed estimate of 1540 m/s is used.

4. RESULTS AND DISCUSSION

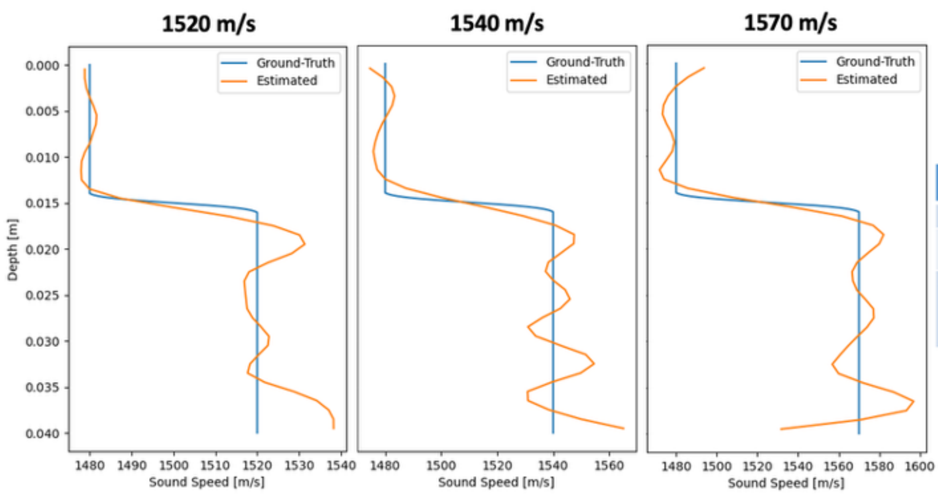

Step Size Ratio at 1st Gradient Descent Step				
		Imaging Grid:		
		Fixed	*1D Adaptive*	*2D Adaptive*
Sound speed in 2nd layer	*1520 m/s*	2558.6	71.0	3.4
	1540 m/s	1868.0	54.7	3.0
	1570 m/s	1656.9	40.0	3.1

Figure 3. Impact of Adaptive Target-Following Grids on Sound Speed Estimation in Simulations of Layered Media. As we move from a fixed to a 1D and a 2D adaptive grid, the step size ratio becomes increasingly closer to 1. Across all three simulations, the layer-wise sound speed estimates are with 10 m/s RMS error from the ground truth. Each sound speed estimate shown is based on the 2D adaptive grid.

Across all three k-Wave simulations (Figure 3), there is a significant decrease in the step-size ratio as we move from a fixed grid to 1D and 2D adaptive grids. In all three cases, the step-size ratio is greater than 1000 for a fixed grid, less than 100 for a 1D adaptive grid, and between 3 and 3.5 for a 2D adaptive grid. As we increasingly adapt the grid to the expected shifts in the imaging targets, we obtain step size ratios increasingly closer to 1, thereby mitigating the nonlinearity of the inverse problem. In all three cases, the layer-wise sound speed is accurately reconstructed to within 10 m/s root-mean-square error.

In-vitro phantom experiments (Figure 4) show that differentiable beamforming accurately reconstructs the sound speed inside the ATS Model 549 phantom (1460 m/s) and the slab of chicken breast immediately above it (1580-1600 m/s) while correcting the aberrations in the B-mode image. The step size ratio for the first gradient descent step is 9897.2 for a fixed grid, 91.7 for a 1D adaptive grid, and 4.2 for a 2D adaptive grid. The phantom experiment confirms that the step size ratio decreases towards a value of 1 as the choice of imaging grid better follows the imaging targets as the sound speed estimate changes.

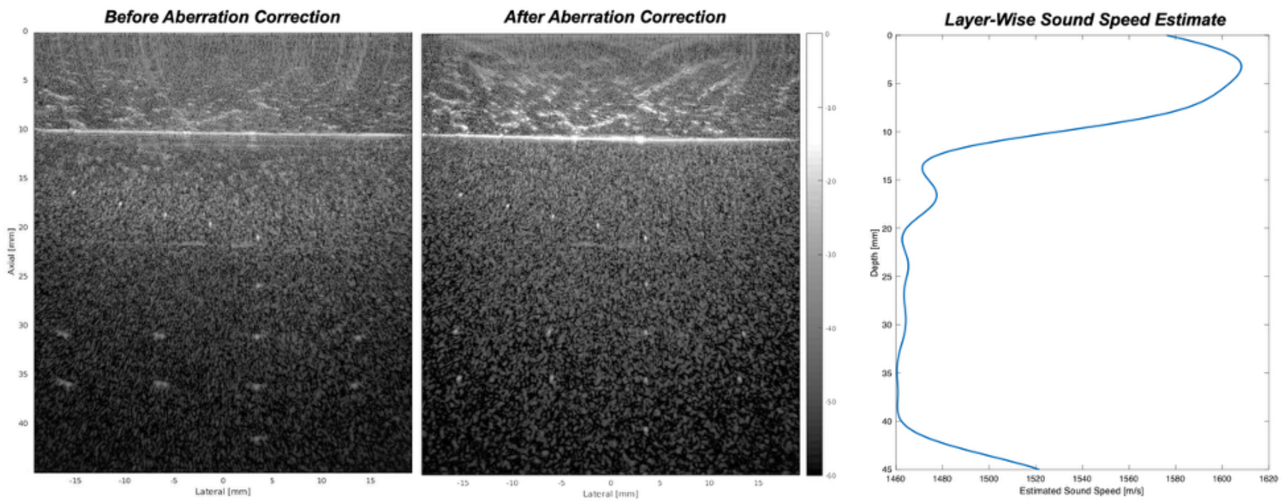

Step Size Ratio at 1st Gradient Descent Step		
Fixed Grid	1D Adaptive Grid	2D Adaptive Grid
9897.2	91.7	4.2

Figure 4. Impact of Adaptive Target-Following Grids on Layer-Wise Sound Speed Estimation in a Two-Layer Sound Speed Phantom Experiment. A slab of chicken breast (1580-1600 m/s) was placed on top of ATS Model 549 phantom (1460 m/s). As we move from a fixed to a 1D and a 2D adaptive grid, the step size ratio becomes increasingly closer to 1. Before aberration correction, the B-mode image (left) is shown based on a beamforming sound speed of 1540 m/s. The sound speed estimate shown (right), based on the 2D adaptive grid, results in the aberration correction shown (middle).

This work ultimately highlights an important consideration in the design of differentiable beamformers for sound speed estimation and aberration correction. A fixed imaging grid causes imaging targets to become cycle-locked to their initial placements in the image. This cycle-locking effect causes locally quadratic models of the objective function to deviate significantly from the true objective function. From the results shown, the step size estimated from a locally quadratic model under a fixed imaging grid can be 1000-10000 times larger than the step size estimated from a numerically exact line search. This cycle locking effect resists changes in the layer-wise sound speed model when relying on locally quadratic models for the step size calculation.

However, an adaptive grid that follows the target reduces the discrepancy between these locally quadratic models and the objective function. A 1D adaptive grid roughly follows the depth placement of the imaging targets while neglecting any lateral displacements, resulting in a step size ratio ranging from 40-100, a significant improvement from 1000-10000 on a fixed grid. A fully 2D adaptive grid that aims to track both lateral and axial displacement in the image reduces the step size ratio to below 5, making efficient use of the Hessian-based step size calculation. Several previous works[4–6,9,17] evaluate aberrations over a fixed set of image points (i.e., an Eulerian perspective), even though the changing sound speed estimate causes targets to shift in the image. These results strongly imply that sound speed estimation and aberration correction should be target-centric rather than image point-centric, highlighting the importance of a Lagrangian target-tracking approach that overcomes the cycle-locking effect that resists sound speed changes in a fixed Eulerian frame of reference.

The Lagrangian terms in this new differentiable beamformer are currently dictated by the mapping between the sound speed change and the expected shift in the imaging target (such as in Equation (8)). Note that there is no unique closed-form mapping between the sound speed change and the expected shifts in the image. For example, an alternative to Equation (8) might be to find the 2D shifts $(\frac{\partial x}{\partial \vec{c}}, \frac{\partial z}{\partial \vec{c}})$ by minimizing the 2-norm length of $\frac{DI_m(x,z;\vec{c})}{D\vec{c}}$; however, this approach is poorly defined for regions in the image with low echogenicity. Another alternative might be to constrain $\frac{\partial z}{\partial \vec{c}}$ (or the axial shifts Δz) based on the closed-form relationship between depth and sound speed ($z_k = z_0 + \sum_{n=1}^{k} c_n \Delta t$) when calculating the lateral shifts Δx in Equation (8). It may also be possible to use deep reinforcement learning with a differentiable beamformer to learn the mapping $(\frac{\partial x}{\partial \vec{c}}, \frac{\partial z}{\partial \vec{c}})$

between the sound speed change and 2D shifts in the image by training on simulated ultrasound channel data. Deep reinforcement learning could aim to enforce the material derivative (7) or directly minimize the step-size ratio as part of its reward function.

5. CONCLUSIONS

We present a target-following approach to layer-wise sound speed estimation in medical pulse-echo ultrasound. This work demonstrates accurate layer-wise sound speed estimation based on an adaptive imaging grid that follows each image target as the sound estimate changes. Adaptive imaging grids that follow the imaging targets are shown to greatly reduce the nonlinearity of the inverse problem of sound speed estimation in layered media. Moving from a fixed to a 1D and a 2D adaptive grid is shown to increasingly diminish the nonlinearity of the inverse problem of sound speed estimation in layered media. In simulations, coherence maximization with a 2D adaptive grid reconstructs the layer-wise sound speed to within 10 m/s root-mean-square error. In a phantom experiment, the same approach leads to sound speed estimates that minimize the aberrations in the image. These results broadly imply that sound speed estimation with a differentiable beamformer is better grounded in a target-centric frame of reference rather than at fixed locations in space.

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. The authors would like to acknowledge Dr. Mohammad Mousavi and Dr. Len Zedel at the Memorial University of Newfoundland for scientific exchanges that led me to Dr. Rifat Ahmed's work on the spatial ambiguity correction[10] that led to this work. Additionally, the authors would like to thank Dr. Rifat Ahmed (formerly at Duke University; now at Vave Health) for his input on the spatial ambiguity correction and future directions in the field.

REFERENCES

[1] Podkowa, A. S. and Oelze, M. L., "The convolutional interpretation of registration-based plane wave steered pulse-echo local sound speed estimators," *Physics in Medicine & Biology* **65**(2), 025003 (2020).

[2] Stähli, P., Kuriakose, M., Frenz, M., and Jaeger, M., "Improved forward model for quantitative pulse-echo speed-of-sound imaging," *Ultrasonics* **108**, 106168 (2020).

[3] Schweizer, D., Rau, R., Bezek, C. D., Kubik-Huch, R. A., and Goksel, O., "Robust imaging of speed of sound using virtual source transmission," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **70**(10), 1308–1318 (2023).

[4] Ali, R., Mitcham, T. M., Singh, M., Doyley, M. M., Bouchard, R. R., Dahl, J. J., and Duric, N., "Sound speed estimation for distributed aberration correction in laterally varying media," *IEEE transactions on computational imaging* **9**, 367–382 (2023).

[5] Simson, W., Zhuang, L., Frey, B. N., Sanabria, S. J., Dahl, J. J., and Hyun, D., "Ultrasound autofocusing: Common midpoint phase error optimization via differentiable beamforming," *IEEE Transactions on Medical Imaging* (2025).

[6] Ali, R., Mitcham, T., Singh, M., Bouchard, R., Doyley, M., Dahl, J., and Duric, N., "Iterative sound speed tomography for distributed aberration correction," in [*2023 IEEE International Ultrasonics Symposium (IUS)*], 1–4, IEEE (2023).

[7] Imbault, M., Faccinetto, A., Osmanski, B.-F., Tissier, A., Deffieux, T., Gennisson, J.-L., Vilgrain, V., and Tanter, M., "Robust sound speed estimation for ultrasound-based hepatic steatosis assessment," *Physics in Medicine & Biology* **62**(9), 3582 (2017).

[8] Telichko, A. V., Ali, R., Brevett, T., Wang, H., Vilches-Moure, J. G., Kumar, S. U., Paulmurugan, R., and Dahl, J. J., "Noninvasive estimation of local speed of sound by pulse-echo ultrasound in a rat model of nonalcoholic fatty liver," *Physics in Medicine & Biology* **67**(1), 015007 (2022).

[9] Ali, R., Telichko, A. V., Wang, H., Sukumar, U. K., Vilches-Moure, J. G., Paulmurugan, R., and Dahl, J. J., "Local sound speed estimation for pulse-echo ultrasound in layered media," *IEEE transactions on ultrasonics, ferroelectrics, and frequency control* **69**(2), 500–511 (2021).

[10] Ahmed, R. and Trahey, G. E., "Spatial ambiguity correction in coherence-based average sound speed estimation," *IEEE transactions on ultrasonics, ferroelectrics, and frequency control* **71**(10), 1244–1254 (2024).

[11] Ali, R., Hyun, D., and Dahl, J. J., "Application of common midpoint gathers to medical pulse-echo ultrasound for optimal coherence and improved sound speed estimation in layered media," in [*2020 IEEE International Ultrasonics Symposium (IUS)*], 1–4, IEEE (2020).

[12] Mallart, R. and Fink, M., "Adaptive focusing in scattering media through sound-speed inhomogeneities: The van cittert zernike approach and focusing criterion," *The Journal of the Acoustical Society of America* **96**(6), 3721–3732 (1994).

[13] Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., et al., "Jax: Autograd and xla," *Astrophysics Source Code Library* , ascl–2111 (2021).

[14] Jakovljevic, M., Hsieh, S., Ali, R., Chau Loo Kung, G., Hyun, D., and Dahl, J. J., "Local speed of sound estimation in tissue using pulse-echo ultrasound: Model-based approach," *The Journal of the Acoustical Society of America* **144**(1), 254–266 (2018).

[15] Dix, C. H., "Seismic velocities from surface measurements," *Geophysics* **20**(1), 68–86 (1955).

[16] 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), 021314–021314 (2010).

[17] Ali, R., Brevett, T., Hyun, D., Brickson, L. L., and Dahl, J. J., "Distributed aberration correction techniques based on tomographic sound speed estimates," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **69**(5), 1714–1726 (2022).