This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

UltraBend: Refraction-Aware Propagation Modeling for Speed of Sound Imaging in Ultrasound

Felix Duelmer[1,2][0009−0000−1702−1746], Samuel Beuret[3][0000−0001−6558−0892], Benjamin N. Frey[4][0000−0002−2580−2638], Mohammad Farid Azampour[1,2][0000−0003−4077−1021], Nassir Navab[1,2][0000−0002−6032−5611], and Jeremy J. Dahl[3][0000−0001−9877−452X]

[1] Chair for Computer Aided Medical Procedures (CAMP), Technical University of Munich, Boltzmannstr. 3, 85748 Garching, Germany
[2] Munich Center for Machine Learning (MCML), Munich, Germany
[3] Department of Radiology, School of Medicine, Stanford University, Stanford, CA 94305, USA
[4] Department of Applied Physics, Stanford University, Stanford, CA 94305, USA

Abstract. The local speed of sound (SoS) of human soft tissues is a promising quantitative biomarker for the diagnosis of diseases such as breast cancer or fatty liver disease. One approach to estimate local SoS from pulse-echo ultrasound signals is the autofocusing method, which optimizes for the propagation model that best aligns beamformed signals recorded from different transmit–receive apertures. However, most pulse-echo autofocusing methods still rely on a straight-ray hypothesis, thus neglecting refraction effects occurring in heterogeneous media. To address this shortcoming, we introduce UltraBend, which replaces straight-line travel-time approximations with a refraction-aware bent-ray model for local SoS optimization. We compare our method with a state-of-the-art autofocusing method assuming straight-ray wave propagation and show that UltraBend reduces the mean absolute reconstruction SoS error relative to a straight-ray baseline on all evaluated phantoms, including a reduction from 8.51 to 3.70 m/s on an inclusion phantom. These results suggest that improved modeling of wave propagation can substantially improve pulse-echo SoS imaging.

Keywords: Speed-of-sound estimation · Ultrasound autofocusing · Differentiable beamforming · Refraction-aware ray tracing.

1 Introduction

Relative to other medical imaging modalities, pulse-echo ultrasound is portable, inexpensive, and capable of real-time imaging. In addition to B-mode image formation, pulse-echo ultrasound can be used to estimate the local speed of sound (SoS), a quantitative acoustic property of tissue with promising biomarker value. For example, SoS imaging has been employed for breast cancer detection [11]

and fatty-liver assessment [14]. Beyond diagnostic value, accurate SoS estimates can enhance reconstruction quality in B-mode imaging by reducing aberration effects caused by tissue heterogeneity [4, 9].

Quantitative SoS estimation has been extensively studied in ultrasound computed tomography (UCT), particularly for breast imaging [18, 20]. Pulse-echo SoS imaging is, however, significantly more challenging because it relies solely on backscattered signals, whereas through-transmission signals—which play a pivotal role in UCT—are unavailable. Computed ultrasound tomography in echo mode (CUTE) addresses this challenge by estimating angle-dependent echo phase shifts and using them in a regularized tomographic inversion [2, 8, 15].

A complementary strategy was introduced by Simson et al. [12, 13], who proposed differentiable beamforming for ultrasound autofocusing (DBUA). DBUA estimates local SoS by optimizing the SoS distribution used to compute transmit and receive beamforming delays. The optimization minimizes the common-midpoint phase error (CMPE), which uses residual phase inconsistencies between common-midpoint receive signals as the objective function. However, both CUTE and CMPE-based autofocusing commonly compute propagation times assuming straight-ray wave propagation. This is computationally attractive, but neglects refraction effects, affecting the accuracy of SoS estimation. Indeed, in heterogeneous media, acoustic energy effectively follows curved trajectories governed by the Eikonal equation rather than straight-line segments [6].

Recent refraction-aware methods have shown that more accurate wave propagation models improves pulse-echo SoS imaging and aberration correction [1, 5]. However, these approaches are formulated as CUTE-style local-coherence inversion or as aberration correction from a previously estimated SoS map, rather than as a differentiable CMPE delay model for pulse-echo autofocusing. At the other end of the modeling spectrum, several recent methods have proposed to model wave propagation (including refraction and diffraction effects) using a Fourier split-step angular spectrum method. This added fidelity comes, however, with substantially higher memory and runtime costs during iterative optimization [3, 7, 19].

To address this gap, we introduce UltraBend, a refraction-aware propagation model for CMPE-based SoS optimization. UltraBend replaces straight-line delays with bent-ray travel times estimated from the current SoS map, allowing the optimizer to account for first-order refraction effects. The model remains practicable by approximating dense travel-time fields from traced rays and by computing gradients without storing the full ray-tracing history. We validate our approach on k-Wave simulations with known SoS maps by comparing our approach to a straight-ray DBUA baseline.

2 Method

2.1 Common-Midpoint Phase-Error Autofocusing

We build on common-midpoint phase-error (CMPE) autofocusing, which formulates ultrasound autofocusing as an optimization problem over all possible SoS

distributions [12]. For a candidate SoS map $c(x, z)$, transmit and receive travel times are computed from each array element to a set of image locations. These travel times are used within a differentiable beamformer that synthesizes focused subaperture signals from full synthetic-aperture pulse-echo data. Under diffuse, spatially incoherent scattering, focused transmit–receive subaperture pairs sharing the same common midpoint are expected to have zero relative phase when the assumed propagation model is correct [10,12]. Residual phase differences between such common-midpoint pairs therefore indicate phase aberration caused by an incorrect sound-speed model.

For an image location x_f, let $\mathcal{P}_{\mathrm{CM}}$ denote the set of focused subaperture pairs that share a common midpoint. For $(a, b) \in \mathcal{P}_{\mathrm{CM}}$, we denote by $\hat{\Gamma}^{\mathrm{CM}}_{ab}(x_f; c)$ the estimated complex coherence between the corresponding focused subaperture signals. We form $\mathcal{P}_{\mathrm{CM}}$ using lag-1 subaperture pairs, i.e., adjacent focused subapertures with matched common-midpoint geometry. The CMPE loss is computed as the average absolute phase deviation over all evaluated image locations and common-midpoint pairs:

$$\mathcal{L}_{\mathrm{CMPE}}(c) = \frac{1}{N_x N_{\mathrm{CM}}} \sum_{x_f \in \Omega} \sum_{(a,b) \in \mathcal{P}_{\mathrm{CM}}} \left| \angle \hat{\Gamma}^{\mathrm{CM}}_{ab}(x_f; c) \right|. \tag{1}$$

Here, N_x is the number of evaluated image locations, $N_{\mathrm{CM}} = |\mathcal{P}_{\mathrm{CM}}|$ is the number of common-midpoint pairs, and Ω denotes the region over which the autofocus loss is evaluated.

2.2 Ray-Based Time-of-Flight Model

Let us first describe how UltraBend estimates propagation times from a candidate SoS map. In practice, we consider the slowness map $s(\mathbf{x}) = 1/c(\mathbf{x})$, with $\mathbf{x} = (x, z)$, sampled according to a discrete Cartesian grid. Travel times are evaluated on an autofocus target grid $\Omega_{\mathrm{af}} = \{\mathbf{x}_f\}$ for calculating the CMPE loss. For each transducer element e, we seek the one-way propagation time at each target $\mathbf{x}_f$. In the continuous setting, this corresponds to evaluating the element-wise first-arrival field $T_e(\mathbf{x})$ at $\mathbf{x}_f$: $\tau_e(\mathbf{x}_f; c) = T_e(\mathbf{x}_f)$.

UltraBend samples T_e by marching bent rays through the current slowness field. The ray trajectories are modeled as characteristics of the eikonal equation:

$$\|\nabla T_e(\mathbf{x})\| = s(\mathbf{x}), \tag{2}$$

and integrated using a second-order Runge–Kutta scheme. As the rays propagate, local arrival-time estimates are deposited onto nearby autofocus targets, yielding an approximate travel-time field without solving a separate two-point refracted ray problem for every element–target pair (compare Fig. 1).

At a ray sample at arc length ℓ along the ray, let $\mathbf{q}(\ell)$ denote the current position, $\mathbf{d}(\ell)$ the current direction, and $\tau^{\mathrm{ray}}(\ell)$ the accumulated time. This time is the path integral of slowness along the traced ray segment $\gamma(\ell)$ from the

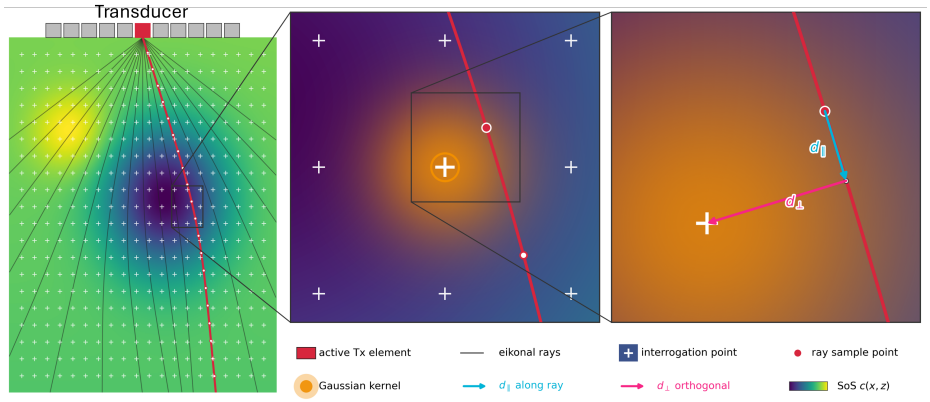

Fig. 1. For each active array element, rays are traced through the current SoS field $c(x, z)$. Nearby autofocus targets accumulate arrival-time estimates via a Gaussian kernel, with along-ray and orthogonal offsets $d_\parallel$ and $d_\perp$ defining the time correction and deposition weight.

element to the current ray position:

$$\tau^{\mathrm{ray}}(\ell) = \int_{\gamma(\ell)} s(\mathbf{x})\, du. \tag{3}$$

u represents the integration variable. For a nearby autofocus target $\mathbf{x}_f$, we define

$$d_\parallel(\ell) = (\mathbf{x}_f - \mathbf{q}(\ell))^{\mathsf{T}}\mathbf{d}(\ell), \qquad d_\perp(\ell) = (\mathbf{x}_f - \mathbf{q}(\ell))^{\mathsf{T}}\mathbf{n}(\ell), \tag{4}$$

where $\mathbf{n}(\ell)$ is the unit normal to the ray. The along-ray offset provides the signed local time shift, while the orthogonal offset contributes the spherical path-length correction used to approximate the arrival time at the target

$$\tau^{\mathrm{loc}}(\mathbf{x}_f, \ell) = \tau^{\mathrm{ray}}(\ell) + \left(d_\parallel(\ell) + \frac{d_\perp(\ell)^2}{2\ell}\right) s(\mathbf{q}(\ell)). \tag{5}$$

The base deposition weight decreases with the distance between the autofocus target and the current ray sample. We use a Gaussian weight based on the offset:

$$w(\mathbf{x}_f, \ell) \propto \exp\left(-\frac{d_\parallel(\ell)^2 + d_\perp(\ell)^2}{2\sigma^2}\right), \tag{6}$$

and set the weight to zero outside a fixed support radius. Here, σ controls how quickly the weight decays with distance from the target. Multiple local arrival estimates are combined with the soft-min weight

$$\alpha(\mathbf{x}_f, \ell) = w(\mathbf{x}_f, \ell) \exp\left[-\beta\left(\tau^{\mathrm{loc}}(\mathbf{x}_f, \ell) - \min_{\ell' \in \mathcal{S}_e} \tau^{\mathrm{loc}}(\mathbf{x}_f, \ell')\right)\right]. \tag{7}$$

Here, $\mathcal{S}_e$ denotes the finite set of sampled arc-length positions along all rays launched from element e. We evaluate α in two tracing passes. The first pass records, for each element–target pair, the earliest supported local arrival time. The second pass traces the same rays again and uses these recorded arrival times to compute the soft-min weights and accumulate the weighted travel-time estimate. In (7), β controls the softness of this replay weighting. Finally, the element-wise travel time is defined as

$$\tau_e(\mathbf{x}_f) = \frac{\sum_{\ell \in \mathcal{S}_e} \alpha(\mathbf{x}_f, \ell)\, \tau^{\mathrm{loc}}(\mathbf{x}_f, \ell)}{\sum_{\ell \in \mathcal{S}_e} \alpha(\mathbf{x}_f, \ell)}. \tag{8}$$

This ray–deposition scheme yields a dense element–target travel-time tensor without explicitly solving a refracted ray path for each element–target pair.

2.3 Replay Adjoint

UltraBend requires estimating the gradient of the objective function with respect to the local SoS at each point of the medium. A naive differentiation of the ray-tracing process would require storing all intermediate ray states, which quickly becomes infeasible due to memory requirements. We therefore follow the path-replay idea of Vicini et al. [17]: instead of storing the full ray history, we recompute the same traced ray steps during the backward pass and use them to accumulate gradients on the slowness grid. To do so, the CMPE objective is first evaluated from the deposited travel-time tensor, producing adjoints $\bar{\tau}_e(\mathbf{x}_f) = \partial\mathcal{L}/\partial\tau_e(\mathbf{x}_f)$. During replay, the traced ray geometry, deposition weights, first-arrival reference times, and target assignments are kept fixed. The derivative of the soft-min reducer therefore distributes each travel-time adjoint back to the local ray contributions that formed it,

$$\bar{\tau}^{\mathrm{loc}}(\mathbf{x}_f, \ell) = \bar{\tau}_e(\mathbf{x}_f) \frac{\alpha(\mathbf{x}_f, \ell)}{\sum_{\ell' \in \mathcal{S}_e} \alpha(\mathbf{x}_f, \ell')}, \tag{9}$$

where ℓ' indexes ray samples contributing to the same element–target pair. The local adjoint is then propagated through the accumulated path time and the local path-length correction in Eq. (5), giving a total slowness adjoint $\bar{s}(\ell)$ at each replayed ray sample. The final SoS gradient is the bilinear interpolation adjoint followed by the chain rule for $s_i = 1/c_i$,

$$\frac{\partial\mathcal{L}}{\partial c_i} = -\frac{1}{c_i^2} \sum_e \sum_{\ell \in \mathcal{S}_e} B_i(\mathbf{q}(\ell))\, \bar{s}(\ell). \tag{10}$$

Here, B_i denotes the bilinear basis weight for grid node i. Thus the forward pass captures ray bending, while replay supplies practical gradients for SoS updates.

2.4 Optimization and Implementation

We start the optimization process with a homogeneous SoS survey over a predefined range and select the speed with the lowest CMPE value as c_0. The spatial

map is then optimized using the Adam optimizer. The objective combines CMPE with an anisotropic total-variation regularization,

$$\mathcal{L}(c) = \mathcal{L}_{\mathrm{CMPE}}(c) + \lambda_{\mathrm{TV}}\mathrm{TV}(c). \tag{11}$$

In practice we use $\lambda_{\mathrm{TV}} = 0.3$ with a lateral-to-axial weighting of 5:1 on a 1 mm $\times$ 1 mm sound-speed grid. Our implementation combines JAX for the computation of the CMPE loss, and Dr.Jit/Mitsuba for ray tracing and replay differentiation. Ray-time fields were computed with 10,000 ray samples per element, a ray-marching step size of 10 μm, randomized step offsets, a soft-min scale of $\beta^{-1} = 50$ ns, and a Gaussian deposition bandwidth of $\sigma = 0.2\lambda$. CMPE was evaluated with lag-1 common-midpoint pairs, 17-element subapertures, a correlation threshold of 0.9, an F-number aperture mask of 0.75, and 300 Adam iterations. The experiments were run on a workstation with an NVIDIA RTX 6000 GPU. Due to the additional computational complexity of the bent-ray model, UltraBend required 4.53 s per optimization iteration compared with 0.91 s for StraightRay, translating to 22.7 min and 4.5 min runtimes, respectively, for 300 optimization iterations.

3 Results

3.1 Experimental Setup

We evaluate UltraBend in controlled simulations with known SoS maps. We consider four synthetic phantoms: a circular inclusion, a layered medium with parallel interfaces, and two checkerboard patterns with different spatial frequencies. These phantoms cover both smooth and piecewise-constant heterogeneity, including cases with oblique interfaces and high-contrast local SoS variations. Pulse-echo full synthetic-aperture (FSA) data are generated using k-Wave [16] using the same parameters as Simson et al. [12]. We compare UltraBend against the straight-ray baseline of Simson et al. [12], which we refer to as StraightRay.

3.2 Arrival-Time Accuracy

Before assessing SoS estimation, we analyze the accuracy of the propagation models by comparing one-way arrival times against k-Wave reference simulations. For each map, k-Wave arrivals are measured for 32 transmit elements

Table 1. One-way time-of-flight residuals relative to k-Wave using a homogeneous mean-offset calibration. Values are mean absolute residual $\pm$ standard deviation [ns] over valid transmit–receiver pairs. Best mean absolute residuals are shown in bold.

Case	StraightRay	UltraBend
Layered	$\mathbf{4.18 \pm 2.38}$	5.34 ± 1.94
Inclusion	2.80 ± 2.77	$\mathbf{2.52 \pm 2.23}$
Checker 2	66.54 ± 135.86	$\mathbf{21.27 \pm 67.35}$
Checker 4	76.14 ± 156.99	$\mathbf{45.66 \pm 141.51}$

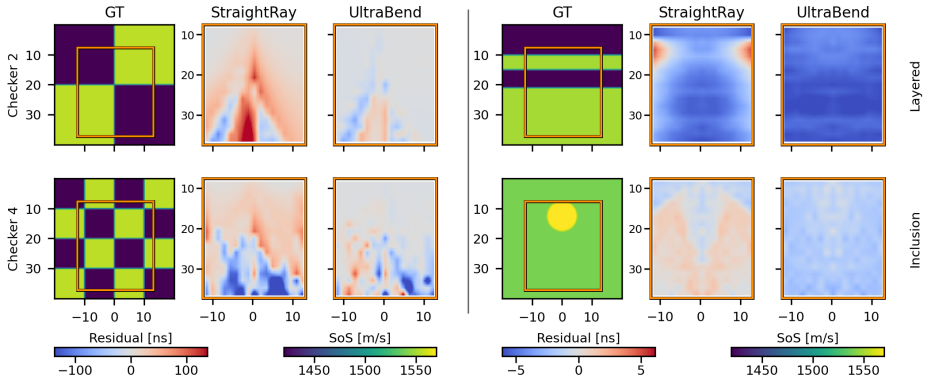

Fig. 2. Arrival-time residual maps after homogeneous timing calibration. Each triplet shows the ground-truth SoS (GT), StraightRay residual, and UltraBend residual. Orange boxes indicate the receiver-grid region. Spatial dimensions are in mm. The checkerboard and layered phantoms use sound speeds of 1420 and 1550 m/s, while the inclusion phantom uses 1540 m/s for the background and 1570 m/s for the inclusion.

and a rectilinear receiver grid. Since absolute k-Wave arrival times include a constant source-wavelet and timing-origin delay, we calculate a single offset estimated from the homogeneous calibration case and subtract it from all maps.

Figure 2 and Table 1 show that UltraBend improves arrival-time accuracy in regions where rays cross oblique interfaces, such as the sides of the layered phantom or the region below the inclusion. In these areas, the StraightRay residual maps show structured artifacts that are substantially reduced by UltraBend. The largest reductions occur in the checkerboard maps, where UltraBend decreases mean absolute residuals from 66.54 to 21.27 ns for Checker 2 and from 76.14 to 45.66 ns for Checker 4. In the smoother heterogeneous cases, UltraBend gives a slightly smaller mean residual for the inclusion phantom, while the layered phantom shows a small remaining bias but lower spread.

3.3 Sound-Speed Reconstruction Accuracy

Table 2. Speed-of-sound reconstruction errors. Values are MAE ± standard deviation of the absolute error in [m/s]. Best MAE values are shown in bold.

Case	StraightRay	UltraBend
Layered	24.73 ± 24.89	$\mathbf{24.10 \pm 24.51}$
Inclusion	8.51 ± 5.59	$\mathbf{3.70 \pm 3.76}$
Checker 2	37.94 ± 28.15	$\mathbf{34.72 \pm 38.33}$
Checker 4	55.83 ± 34.46	$\mathbf{49.25 \pm 44.90}$

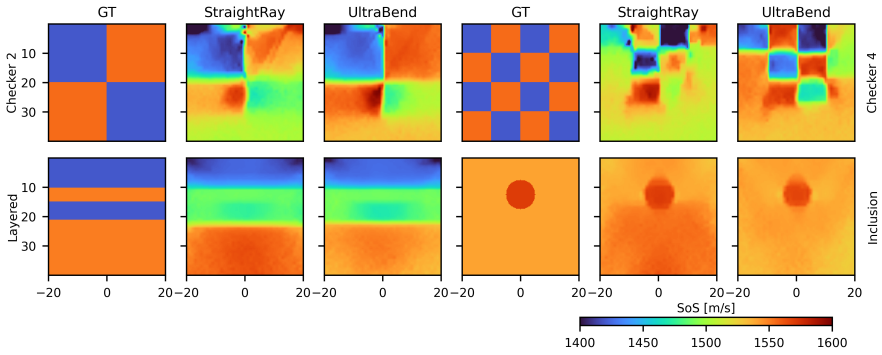

Fig. 3. Speed-of-sound reconstructions from k-Wave simulations with known numerical phantoms. Each triplet shows ground truth (GT), StraightRay, and UltraBend. Spatial dimensions are in mm.

As seen in Figure 3 and Table 2, UltraBend reduces the mean absolute error (MAE) in all four maps. The clearest gain occurs in the inclusion phantom, where MAE decreases from 8.51 to 3.70 m/s, and the absolute-error spread also decreases. The layered phantom shows a smaller but consistent improvement. In the checkerboard phantoms, UltraBend improves the reconstruction of the high-contrast structure and lowers MAE, although the larger standard deviations indicate remaining localized errors near sharp interfaces. Generally, the structure of the SoS maps is preserved more faithfully with UltraBend, and SoS uniformity within homogeneous regions is better enforced.

4 Conclusion

We present UltraBend, a refraction-aware propagation model for CMPE-based ultrasound autofocusing. Our main conclusion is that explicitly modeling acoustic wave propagation using bent rays improves SoS reconstruction in heterogeneous media. In controlled k-Wave simulations, UltraBend reduces one-way arrival-time residuals in regions where propagation was strongly affected by refraction, particularly below inclusions and across high-contrast checkerboard structures. These improvements in travel-time accuracy translate into lower SoS reconstruction error in all four evaluated maps. The strongest improvement is observed in the inclusion phantom, where the MAE decreases from 8.51 to 3.70 m/s. In the checkerboard phantoms, UltraBend also reduces the average error and better preserves high-contrast spatial structure, but at the cost of an increased error spread (28.15 $\rightarrow$ 38.33 m/s for Checker 2 and 34.46 $\rightarrow$ 44.90 m/s for Checker 4), indicating that the remaining error becomes concentrated in fewer, more severe outliers near sharp interfaces and caustics, where the first-arrival ray assumption underlying our soft-min deposition is least valid.

Methodologically, UltraBend provides a practical compromise between straight-ray autofocusing and full-wave inversion. By tracing untargeted bent rays through

the current slowness field and depositing local arrival-time estimates onto nearby autofocus targets, the method incorporates first-order refraction effects without requiring a separate two-point refracted ray solve for every element–target pair. The fixed-replay adjoint further makes this propagation model compatible with iterative SoS optimization without storing the full ray-tracing history.

The present study is limited to controlled simulations and known phantom geometries. Reconstruction errors remain in poorly constrained regions, especially near sharp interfaces and in the image far field, where aperture support is sparse. Future work should therefore focus on experimental validation, more realistic phantoms, improved ray deposition, and regularization strategies that better account for limited angular coverage. Overall, our results suggest that refraction-aware propagation modeling is a promising direction for improving pulse-echo speed-of-sound imaging through differentiable beamforming.

References

1. Ali, R., Brevett, T., Hyun, D., Brickson, L.L., 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)
2. Ali, R., Brevett, T., Zhuang, L., Bendjador, H., Podkowa, A.S., Hsieh, S.S., Simson, W., Sanabria, S.J., Herickhoff, C.D., Dahl, J.J.: Aberration correction in diagnostic ultrasound: A review of the prior field and current directions. Zeitschrift für Medizinische Physik **33**(3), 267–291 (2023)
3. Ali, R., Mitcham, T.M., Doyley, M.M., Duric, N., Dahl, J.J.: Wave-equation migration velocity analysis for multistatic synthetic aperture ultrasound. arXiv preprint arXiv:2604.27428 (2026)
4. Anderson, M., McKeag, M., Trahey, G.: The impact of sound speed errors on medical ultrasound imaging. The Journal of the Acoustical Society of America **107**(6), 3540–3548 (2000)
5. Beuret, S., Besson, A., Hériard-Dubreuil, B., Cohen-Bacrie, C.: A refraction-aware pulse-echo speed-of-sound imaging method for convex transducers. In: 2025 IEEE International Ultrasonics Symposium (IUS) (2025)
6. Cerveny, V.: Seismic ray theory (2001)
7. Hériard-Dubreuil, B., Brenner, E., Rio, B., Lambert, W., Chamming's, F., Fink, M., Aubry, A.: Physics-based learning of the wave speed landscape in complex media. arXiv preprint arXiv:2602.03281 (2026)
8. Jaeger, M., Stähli, P., Martiartu, N.K., Yolgunlu, P.S., Frappart, T., Fraschini, C., Frenz, M.: Pulse-echo speed-of-sound imaging using convex probes. Physics in Medicine & Biology **67**(21), 215016 (2022)
9. Nock, L., Trahey, G.E., Smith, S.W.: Phase aberration correction in medical ultrasound using speckle brightness as a quality factor. The Journal of the Acoustical Society of America **85**(5), 1819–1833 (1989)
10. Rachlin, D.: Direct estimation of aberrating delays in pulse-echo imaging systems. The Journal of the Acoustical Society of America **88**(1), 191–198 (1990)
11. Schweizer, D., Farkas, M., Bezek, C.D., Potempa, A., Leo, C., Kubik-Huch, R.A., Goksel, O.: Pulse-echo imaging of breast speed of sound as a potential biomarker for breast cancer. Ultrasound in Medicine & Biology (2025)

12. Simson, W., Zhuang, L., Frey, B.N., Sanabria, S.J., Dahl, J.J., Hyun, D.: Ultrasound autofocusing: Common midpoint phase error optimization via differentiable beamforming. IEEE Transactions on Medical Imaging (2025)
13. Simson, W., Zhuang, L., Sanabria, S.J., Antil, N., Dahl, J.J., Hyun, D.: Differentiable beamforming for ultrasound autofocusing. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 428–437. Springer (2023)
14. Stähli, P., Becchetti, C., Korta Martiartu, N., Berzigotti, A., Frenz, M., Jaeger, M.: First-in-human diagnostic study of hepatic steatosis with computed ultrasound tomography in echo mode. Communications medicine **3**(1), 176 (2023)
15. Stähli, P., Kuriakose, M., Frenz, M., Jaeger, M.: Improved forward model for quantitative pulse-echo speed-of-sound imaging. Ultrasonics **108**, 106168 (2020)
16. Treeby, B.E., 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. Vicini, D., Speierer, S., Jakob, W.: Path replay backpropagation: differentiating light paths using constant memory and linear time. ACM Transactions on Graphics (TOG) **40**(4), 1–14 (2021)
18. Wiskin, J., Borup, D.T., Johnson, S., Berggren, M.: Non-linear inverse scattering: High resolution quantitative breast tissue tomography. The Journal of the Acoustical Society of America **131**(5), 3802–3813 (2012)
19. Zhuang, L., Beuret, S., Frey, B., Munot, S., Simson, W., Hyun, D., Dahl, J.J.: A wavefield correlation approach to improve sound speed estimation in ultrasound autofocusing. arXiv preprint arXiv:2602.12805 (2026)
20. Zografos, G., Koulocheri, D., Liakou, P., Sofras, M., Hadjiagapis, S., Orme, M., Marmarelis, V.: Novel technology of multimodal ultrasound tomography detects breast lesions. European radiology **23**(3), 673–683 (2013)