The Journal of the Acoustical Society of America

JULY 09 2025

Virtual initialised ray tomography: Towards contact-free realistic ultrasonic bone imaging

Aaron Chung-Jukko ; Peter Huthwaite

Check for updates

J. Acoust. Soc. Am. 158, 276–290 (2025)
https://doi.org/10.1121/10.0036902

View Online
Export Citation

Articles You May Be Interested In

Tunable dual-wavelength erbium ytterbium doped fiber laser assisted by polarization-dependent loss controller

AIP Conf. Proc. (October 2023)

Numerical analysis of pattern shape deformation in UV imprint considering thermal deformation with UV shrinkage and curing

J. Vac. Sci. Technol. B (February 2021)

Linking information theory and thermodynamics to spatial resolution in photothermal and photoacoustic imaging

J. Appl. Phys. (November 2020)

08 September 2025 15:08:20

ASA
LEARN MORE
Advance your science and career as a member of the
Acoustical Society of America

ASA
ACOUSTICAL SOCIETY OF AMERICA

JASA ARTICLE

Virtual initialised ray tomography: Towards contact-free realistic ultrasonic bone imaging

Aaron Chung-Jukko[a] and Peter Huthwaite
Department of Mechanical Engineering, Imperial College London, London, SW7 2AZ, United Kingdom

ABSTRACT:
Ultrasound tomography struggles with high-contrast and/or extended-range problems, such as site-specific *in vivo* bone imaging, which is crucial for assessing fracture risk. Two main obstacles arise. (a) An external region prohibits access to the region of interest (ROI), and (b) high contrast exists between the two regions. These challenges hinder existing algorithms—including bent-ray tomography (BRT). Virtual initialised ray tomography (VIRT) is proposed, which addresses these challenges through (a) placement of virtual transducers on the ROI, facilitating (b) rapid initialisation before BRT inversion. *In silico* validation against BRT with and without prior information shows superior resolution and robustness while improving computational speed. Validation against BRT directly on the ROI shows that VIRT is approaching the resolution limit. Additionally, VIRT performs well under 20 dB signal-to-noise ratio. The ability to solve high-contrast extended-range tomography problems without prior knowledge about the interior of the ROI has many implications. VIRT has the potential to unlock site-specific *in vivo* bone imaging for assessing fracture risk. It may also replace BRT to improve resolution, robustness, and computational speed—especially where the ROI does not cover the entire imaging array.
© *2025 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).* https://doi.org/10.1121/10.0036902
(Received 9 March 2025; revised 19 May 2025; accepted 2 June 2025; published online 9 July 2025)
[Editor: Guillaume Haiat] Pages: 276–290

08 September 2025 15:08:20

I. INTRODUCTION

Ultrasonic tomography is widely applied in diverse fields from medicine to geophysics and engineering. Tomography reconstructs material properties in a region of interest (ROI), which are seldom directly accessible and must not be destroyed in the process. These reconstructions carry critical information and, hence, tomographic algorithms must balance resolution, speed, and robustness. Resolution and speed are straightforward and quantifiable, whereas robustness ensures that an algorithm reliably produces accurate images with limited sensitivity to undesired property variations, minimal manual tuning (e.g., weighting factors), and good convergence properties.

One of the most robust and fastest algorithms is bent-ray tomography (BRT), also known as time-of-flight tomography or travel-time tomography (Huthwaite, 2014b). Current applications of BRT range from corrosion mapping to breast imaging and as a starting background for algorithms with higher resolution but lower robustness (Huthwaite, 2014b; Li *et al.*, 2006; Huthwaite and Simonetti, 2011). However, despite its robust behaviour, BRT struggles with increasing contrast and poor resolution. The latter arises from ignoring diffraction effects and is exacerbated when the ray paths are long. These drawbacks are amplified in high-contrast extended-range tomography problems, where the ROI has a high acoustic contrast with the external region (ER) and is also distant from the array.

An archetype of these problems is *in vivo* ultrasonic imaging of the proximal femur, where bone (ROI) is surrounded by soft tissues (ER), leading to high contrast and a large array diameter. This imaging is needed to determine the risk of hip fractures, the deadliest and costliest type of osteoporotic fracture (Svedbom *et al.*, 2013; Kanis *et al.*, 2003). The current diagnostic gold standard, dual x-ray absorptiometry, measures bone mineral density rather than strength and has high specificity but low sensitivity (World Health Organization, 1994). Ultrasound-based methods are a promising alternative given their ability to uncover bone quality metrics, such as porosity and stiffness (Grimal and Laugier, 2022), but existing tomographic methods "are still too limited today to be used in clinical application" (Lasaygues *et al.*, 2022). These limitations have shifted research towards the suboptimal approach of correlating fracture site properties through ultrasonic characterisation of vastly different peripheral skeletal sites (Bochud and Laugier, 2022; Tran *et al.*, 2022). There are also other examples outside of medicine, such as in nondestructive testing with guided wave tomography, where it is desired to image a local high-contrast inhomogeneity. An imaging algorithm that can successfully perform high-contrast extended-range tomography is desperately needed.

In this paper, we propose a novel algorithm: virtual initialised ray tomography (VIRT). The main goals are to improve the (1) extended-range resolution and (2) high-contrast robustness of BRT, which is validated in the context of bone imaging.

[a]Email: aaron@chung-jukko.com

 ©Author(s) 2025.

https://doi.org/10.1121/10.0036902

Achieving these goals has the potential to significantly advance the accurate diagnosis, monitoring, and treatment of osteoporosis. This advancement can drastically improve the lives of approximately one-third of women and one-fifth of men aged 50 years and older (Svedbom *et al.*, 2013; Kanis *et al.*, 2000) as demographics shift towards older age groups. VIRT will also find applications outside of bone imaging in most scenarios in which BRT is applicable with resolution gains achievable for most problems. The robustness and improved resolution of VIRT offer a solid starting background for high-resolution but computationally intensive and less robust algorithms such as full-waveform inversion.

The theory behind BRT performance in high-contrast extended-range tomography is discussed in Sec. II. Section III explores the key innovations of VIRT: virtualisation of transducers onto the ROI boundary and initialisation step, which improves the high-contrast robustness. In Sec. IV, VIRT is validated with two sample problems, and the results are presented in Sec. V. Further discussion of the implications of VIRT and suggestions for practical implementation are given in Sec. VI.

II. BRT

Before explaining VIRT in detail, we will briefly discuss BRT, which VIRT improves on. The wave equation for a scalar property p (e.g., pressure) in the frequency domain is described by

$$(\nabla^2 + k^2)p = 0, \tag{1}$$

which is the Helmholtz equation. Here, $k = \omega/c_{\text{phase}}$ is the local wavenumber expressed in terms of ω, which is the angular frequency, and c_{phase} is the phase velocity. BRT is based on ray theory, which solves Eq. (1) by taking the infinite frequency assumption ($\omega \to \infty$), valid when perturbations in the field are large compared to the wavelength. This leads to the eikonal equation

$$\frac{1}{c_{\text{phase}}} = |\nabla\tau|, \tag{2}$$

where τ is the time delay for the specific frequency component to travel from a source to any point in the field (Fomel, 1997). In BRT, typically, the wavepacket arrival times τ_{group} are picked, such that an alternative form of $1/c_{\text{group}} = |\nabla\tau_{\text{group}}|$ is used in practice, and we will use c to refer to both velocities, generally, from now on.

The forward problem in BRT is solved by eikonal solvers such as the fast-marching method (FMM; Sethian, 1996). FMM discretises the eikonal equation (infinite frequency version of the wave equation) using the finite-difference approximation, where arrival time is the underlying variable. The resulting equations are then solved by starting from the source (the earliest point in time) and stepping through the grid to update the later times.

The inverse problem in BRT is posed as a minimisation problem in the form of

$$\text{minimise}\, m(\mathbf{c}^{\text{model}}) = \sum_{i,l}\left(\tau_{i,l}^{\text{model}}(\mathbf{c}^{\text{model}}) - \tau_{i,l}^{\text{meas}}\right)^2. \tag{3}$$

Here, m is the misfit, and $\tau_{i,l}^{\text{model}}(\mathbf{c}^{\text{model}})$ and $\tau_{i,l}^{\text{meas}}$ are the modelled and measured arrival times, respectively, between source i and receiver l. This inversion is usually solved iteratively using a conjugate-gradient method that back-traces the misfit onto the field and updates the voxels it passes through (Li *et al.*, 2006; Hormati *et al.*, 2010).

The eikonal equation (2) only uses local velocity c to calculate wavefront propagation, making travel time sensitive only to perturbations on the ray path in ray theory. These perturbations refer to changes in the slowness field, represented as variations in the velocity values c assigned to individual voxels in the model. Diffraction is ignored because the eikonal equation assumes infinite frequency. This also manifests in the fact that the FMM only operates on the first arrival, hence neglecting any later arriving waves appearing in the domain. This means that interference effects are not captured, therefore removing the ability to model diffraction. Considering diffraction effects, the travel time of a finite frequency wave is sensitive to perturbations in the first Fresnel zone (F_1, which is shaped like a banana-doughnut and is widest at the midpoint), which excludes the ray path (Hung *et al.*, 2001). This discrepancy is behind the effect known as wavefront healing and introduces an error whereby ray theory is unable to detect perturbations below size (Hung *et al.*, 2001)

$$\varepsilon_{\text{ray}} \approx F_{1,\max} = \sqrt{\lambda_0 L} = \sqrt{\frac{c_0 L}{f}}. \tag{4}$$

Here, f is the central frequency of the propagated wave, L is the distance between the source and receiver pair, and c_0 and λ_0 are the background sound speed and wavelength, respectively.

The resolution of BRT is limited by this error in ray theory, which carries to the inversion step. The global optimum of Eq. (3) can only achieve a maximum resolution proportional to the intrinsic error in ray theory ε_{ray}. This gives BRT's resolution limit as

$$\varepsilon_{\text{BRT}} \propto \sqrt{\frac{c_0 \max_{i,l}|\mathbf{S}_i - \mathbf{R}_l|}{f}}, \tag{5}$$

assuming that the longest ray path is a straight line (Williamson, 1991; Williamson and Worthington, 1993). Respectively, $\mathbf{S}_i$ and $\mathbf{R}_l$ are the positions of the ith source and lth receiver. For circular arrays, $\max_{i,l}|\mathbf{S}_i - \mathbf{R}_l|$ simplifies to the array diameter D. The challenge in bone imaging and other extended-range tomography problems is that c_0 and D are intrinsic to the problem. Increasing f usually leads to increased attenuation and scattering. This makes it difficult for BRT to achieve high resolution in these contexts.

The robustness of BRT is enhanced by its low resolution as the back-traced misfit field is typically smooth and

08 September 2025 15:08:20

monotonic with few local minima (Huthwaite, 2014b). However, high-contrast problems back-trace large misfits, forming many valleys and increasing convergence at local minima during gradient descent. This misconvergence sometimes manifests as ray-like or clustered striated artefacts as rays become unable to escape the valley, such that only voxels directly on the ray path are updated. This effect could be further exacerbated by low sampling. Regularisation techniques, such as total variation and first/second derivative minimisation, aid global convergence but may reduce resolution (e.g., from blurring) and need tweaking. Thus, increased contrast reduces BRT's robustness.

III. VIRT

VIRT addresses BRT's challenges in solving high-contrast extended-range tomography (introduced in Sec. II; see Fig. 1) with two novel steps: virtualisation, which predominantly increases resolution; and initialisation, which increases robustness. Figure 2 shows a flow chart of the VIRT algorithm, where the novel VIRT steps are shown in green and the classical BRT steps are depicted in black. Next, we will discuss virtualisation in Sec. III A and initialisation in Sec. III B.

A. Virtualisation

1. Motivation

The key insight of VIRT is that a flexible, robust, and problem-agnostic algorithm such as BRT is needed only within the unknown (and high-contrast) ROI, as shown in Fig. 3. In contrast, the ER is often quite homogeneous: it consists of various soft tissues in bone imaging and homogeneous waveguide in guided wave tomography. Provided that we have suitable *a priori* information about the low-contrast ER, we can virtualise full-matrix capture (FMC) data from the physical array to the virtual array at the ROI boundary. This boundary corresponds to the minimum convex bounds for concave ROIs or multiple closely clustered ROIs (e.g., corrosion patches).

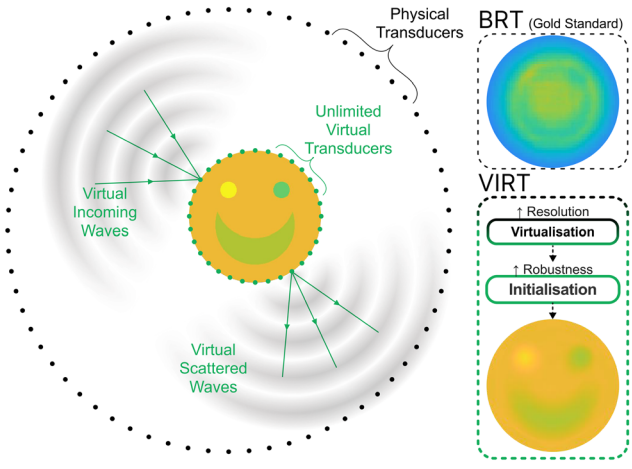


FIG. 1. Graphical representation of VIRT algorithm compared with BRT.

Virtualisation brings three main benefits. First, performing BRT from the virtual transducers gives VIRT's resolution limit as

$$\varepsilon_{\mathrm{VIRT}} \propto \sqrt{\frac{c_{\mathrm{ROI}} \max_{j,k} |\mathbf{S}'_j - \mathbf{R}'_k|}{f}} = \sqrt{\frac{c_{\mathrm{ROI}} D_{\mathrm{ROI}}}{f}}, \tag{6}$$

as calculated from the BRT resolution limit [Eq. (5)]. Here, c_{ROI} is the sound speed within the ROI, $\mathbf{S}'_j$ and $\mathbf{R}'_k$ are the positions of the jth virtual source and kth virtual receiver, respectively, and D_{ROI} is the maximum diameter of the ROI. As D_{ROI} is always smaller than the array diameter D, the resolution is improved in VIRT.

Second, performing BRT from the smaller array of D_{ROI} drastically increases the speed of the eikonal solver when running at the same resolution.

Third, virtualisation increases robustness in the case of high-contrast problems, such as in bone imaging, through a combination of

(1) unlocking the fast and simple initialisation step (discussed in Sec. III B);
(2) separating high (bone) and low (soft tissue) acoustic impedance regions, reducing contrast within the BRT solution area (through strategic ROI boundary selection); and
(3) introducing theoretically unlimited virtual transducers to increase sampling density.

2. Method

Huygens modelled wavefronts as multiple secondary sources, specifically as monopoles, which Miller (1991) corrected to be spatiotemporal dipoles. In a homogeneous background, the spatiotemporal dipole's single-harmonic three-dimensional (3D) field is

$$\phi_{\mathrm{StD,3D}} \propto X_{\mathrm{StD}} \frac{e^{ikr}}{r}\left(-ik(1+\cos\,\theta) + \frac{\cos\,\theta}{r}\right). \tag{7}$$

Here, $X_{\mathrm{StD}} = a_{\mathrm{StD}} e^{-i\omega t}$ is the complex frequency component of the spatiotemporal dipole, k is the wavenumber, r is the radius from the source, and θ is the angle compared to the propagation direction of the spatiotemporal dipole. Thus, the 3D spatiotemporal-dipole wavefield fluctuates with e^{ikr}, decays with $1/r$, and has an angular dependency.

In VIRT, the goal is for the virtual transducers to mimic these realistic secondary sources (spatiotemporal dipoles) and act as conduits between the ROI and ER. This would allow virtual excitation and reception at $\mathbf{S}'_j$ and $\mathbf{R}'_k$, respectively (see Fig. 3).

In the physical FMC dataset, each physical transducer sends the same signal that we want to send from the virtual array. This signal is normally a Hann tone burst, which can be taken as a pulse at $t = 0$. The relevant wavefront for the virtual FMC dataset propagates through the ER to the ROI boundary, interacts with the ROI, and returns through the

08 September 2025 15:08:20

https://doi.org/10.1121/10.0036902

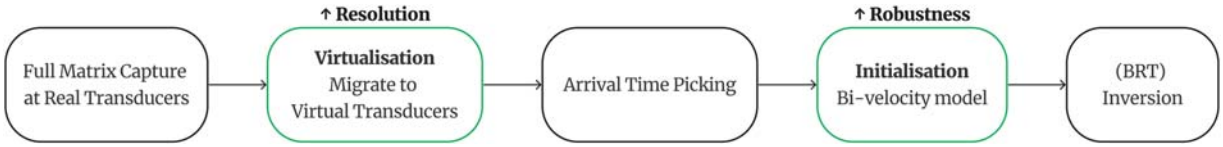


FIG. 2. The steps of VIRT. Novel steps are shown in green, and classical BRT steps are displayed in black.

ER to all physical receivers $\mathbf{R}_l$. In the far field, the incoming wavefront from $\mathbf{S}_i$ arriving at $\mathbf{S}'_j$ results in a complex frequency component at $\mathbf{S}'_j$ in a homogeneous background of

$$X_{\mathbf{S}_i \to \mathbf{S}'_j} = g(k) X_i e^{ikr_{ij}} \frac{1}{A(r_{ij})}, \tag{8}$$

where $X_{a\to b}$ denotes the contribution of signal a on signal b. The complex constant $g(k)$ represents multiple factors within propagation, including the overall behaviour of the data acquisition system (e.g., transducers). $X_i = a_i e^{-i\omega t}$ is the complex frequency component of the wave leaving $\mathbf{S}_i$ and $1/A$ is the decay with distance from geometric spreading ($A_{3D} = r_{ij}$, $A_{2D} = \sqrt{r_{ij}}$), where $r_{ij} = |\mathbf{S}'_j - \mathbf{S}_i|$. The factor $e^{ikr_{ij}}$ accounts for the phase shift from the journey $\mathbf{S}_i \to \mathbf{S}'_j$, which is very similar to e^{ikr} in Eq. (7). This similarity suggests that a wavefield similar to a spatiotemporal dipole can be recreated with sufficient focused wavefronts in the manner of beamforming/delay-and-sum, as depicted in Fig. 4.

We convert the physical FMC data to the virtual FMC data by virtualisation, which defines the virtual signal as

$$X'_j = \sum_i X_{\mathbf{S}_i \to \mathbf{S}'_j} e^{-ikr_{ij}} A(r_{ij}) \frac{\theta_{(i+1)j} - \theta_{(i-1)j}}{2} B(\theta_{ij}), \tag{9}$$

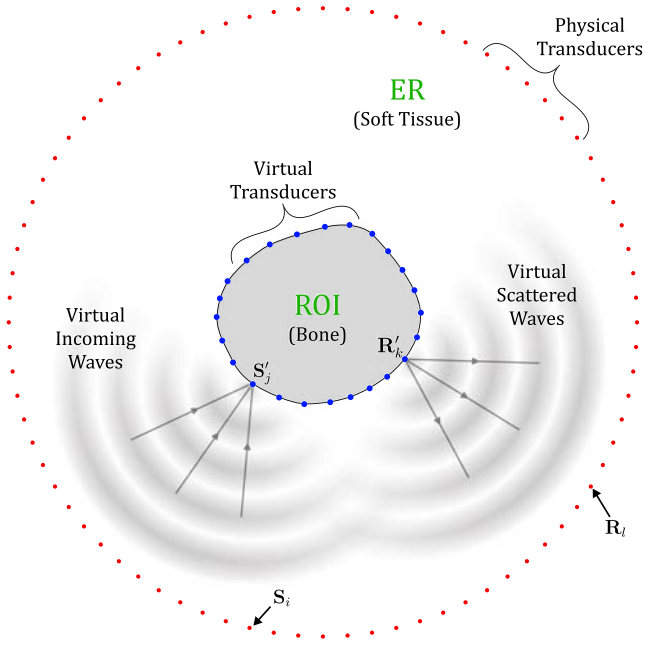


FIG. 3. Visualisation of the VIRT virtualisation step on a bone imaging setup. The ER (soft tissue) prevents placing physical transducers (red) directly on the ROI. VIRT improves the resolution limit by performing BRT from virtual transducers (blue) on the ROI boundary.

which is composed of the incoming waves at $\mathbf{S}'_j$. Here, $e^{-ikr_{ij}}$ and $A(r_{ij})$ are the phase shift and distance weighting factors, respectively, which undo the corresponding effects of ER propagation on the journey $\mathbf{S}_i \to \mathbf{S}'_j$ described in Eq. (8). In the time domain, this phase shift corresponds to a time advance of $\tau_{ij} = r_{ij}/c_{ER} = kr_{ij}/\omega$ in a homogeneous ER with speed c_{ER}. This shift normalises and focuses each incoming wave to give the virtual signal X'_j the same phase as the physical signal X, which is similar to delay-and-sum/beamforming. $(\theta_{(i+1)j} - \theta_{(i-1)j})/2$ accounts for any uneven angular spacing between incoming waves from physical sources $\mathbf{S}_i$, according to the trapezium rule as part of discrete integration.

The directivity pattern B varies with the angle θ_{ij}, defined by the normal $\hat{n}_{\mathbf{S}'_j}$ of the ROI boundary at $\mathbf{S}'_j$ as shown in Fig. 4. To mimic a spatiotemporal dipole [Eq. (7)] in the far field ($r \gg 0$),

$$B_{\mathrm{StD,far}} \approx 1 + \cos(\theta_{ij}) \tag{10}$$

gives a cardioid directivity pattern. However, we note that any incoming waves arriving from above the local tangent (i.e., $|\theta| > \pi/2$) could interact with the unknown and potentially high-contrast ROI, breaking Born's approximation.

A well-apodised approximation to a cardioid directivity pattern is

$$B = \max(\cos\,\theta_{ij}, 0)^n, \tag{11}$$

which is the one-sided directivity pattern of a dipole source when $n = 1$. This approximation restricts illumination of the

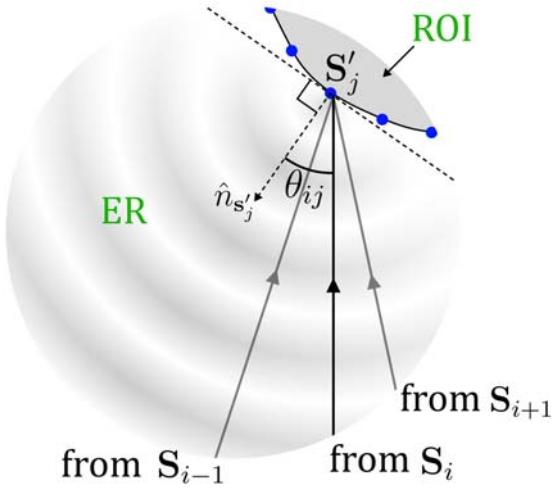


FIG. 4. Visualisation of post-processed virtual excitation. Incoming waves are focused at the jth virtual source.

08 September 2025 15:08:20

virtual source $\mathbf{S}'_j$ to the known and no/low-contrast ER, maintaining Born's approximation. We have, thus, obtained the relationship between all physical sources and all virtual sources.

Although this analysis is for a virtual source, the principle of reciprocity allows for effective signal measurement from a virtual receiver. Instead of an effective radiation pattern into the ROI from the virtual source, the equations now derive the sensitivity pattern of the virtual receiver from the ROI. Therefore, the overall relationship between the physical FMC data set ($\mathbf{S}_i \rightarrow \mathbf{R}_l$) and the virtual FMC data set ($\mathbf{S}'_j \rightarrow \mathbf{R}'_k$) is

$$X'_{\mathbf{S}'_j \rightarrow \mathbf{R}'_k} = \sum_{i,l} X_{\mathbf{S}_i \rightarrow \mathbf{R}_l} \times e^{-ikr_{ij}} A(r_{ij}) \frac{\theta_{(i+1)j} - \theta_{(i-1)j}}{2} B(\theta_{ij}) \times e^{-ikr_{lk}} A(r_{lk}) \frac{\theta_{(l+1)k} - \theta_{(l-1)k}}{2} B(\theta_{lk}), \quad (12)$$

which is easily executable as a frequency-wise matrix multiplication. By taking the inverse Fourier transform of $X'_{\mathbf{S}'_j \rightarrow \mathbf{R}'_k}$ to return to the time domain, we can recover the virtual FMC data from the physical FMC data.

The first point to note is that we have not considered any aberration effects that would occur if the ER were inhomogeneous. As with many focusing approaches, provided that the ER is known or can be approximated in some way, it should be possible to correct for this, yet, such correction will not be evaluated in this paper.

A second point to note is that the migration of FMC data could also be achieved in real time. However, real-time focusing may result in longer data collection times for two reasons. First, any inaccuracies in the *a priori* assumptions (e.g., c_0, ROI boundary location) for the focusing law may require recapturing. Second, separate real-time data capture and processing is required to obtain *a priori* information that is critical to virtualisation. We will not discuss this further in this paper.

A third point to note is that virtual focusing allows one to place unlimited transducers on the ROI boundary. This increases spatial sampling within the ROI boundary, which, in turn, increases the density of ray paths from which inversion occurs. As the distances between rays decrease, rays are less likely to become stuck in local optima, and robustness is increased. This addresses the loss of robustness in BRT resulting from increased contrast (Sec. II) and reduces the reliance on regularisation for convergence.

B. High-contrast initialisation

As mentioned in Sec. II, BRT's robustness drops under high contrast. A large travel-time misfit (m) creates many valleys in the back-traced field. This causes the conjugate-gradient method to converge at local optima. An example is reconstructing a low-speed object, where the fastest ray paths, under the eikonal approximation, bypass the object, causing the reconstructed wave speed within the object to be unrelated to the true value. It is common to initialise the entire domain (ER and ROI) with the sound speed of the ER (c_{ER}), but this only works for low-contrast problems.

Observing that there is usually a clear velocity contrast between the ROI and ER, we can take a bi-velocity model $\mathbf{c}^{\mathrm{bi-velocity}}$ as a first approximation of the true sound speed map such that

$$c(\mathbf{r}) \approx \mathbf{c}^{\mathrm{bi-velocity}}(c_{\mathrm{ROI}}, c_{\mathrm{ER}}) = \begin{cases} c_{\mathrm{ROI}} & \text{if} \quad \mathbf{r} \in \mathrm{ROI}, \\ c_{\mathrm{ER}} & \text{otherwise.} \end{cases} \quad (13)$$

Here, c_{ER} is the speed of the ER, and c_{ROI} is the unknown initial sound speed within the ROI. *A priori* information on c_{ER} allows us to further simplify to $\mathbf{c}^{\mathrm{bi-velocity}}(c_{\mathrm{ROI}}, c_{\mathrm{ER}}) = \mathbf{c}^{\mathrm{bi-velocity}}(c_{\mathrm{ROI}})$.

For BRT, initialisation with the bi-velocity starting model is obtained by solving

$$\text{minimise } m(\mathbf{c}^{\mathrm{bi-velocity}}(c_{\mathrm{ROI}})) = \sum_{i,l} \left(\tau^{\mathrm{model}}_{i,l}(\mathbf{c}^{\mathrm{bi-velocity}}) - \tau^{\mathrm{meas}}_{i,l} \right)^2, \quad (14)$$

which is a more robust version of Eq. (3). The misfit and, hence, $\tau^{\mathrm{model}}_{i,l}(\mathbf{c}^{\mathrm{bi-velocity}})$ (arrival time) must be re-calculated through an eikonal solver (e.g., FMM) for each iteration in any optimisation algorithm. Thus, initialisation iterations are as time-consuming as BRT inversion iterations, causing slow and cumbersome initialisation.

For VIRT, initialisation with the bi-velocity starting model is obtained by solving

$$\text{minimise } m'(\mathbf{c}^{\mathrm{bi-velocity}}(c_{\mathrm{ROI}})) = \sum_{j,k} \left(\tau'^{\mathrm{model}}_{j,k}(\mathbf{c}^{\mathrm{bi-velocity}}) - \tau'^{\mathrm{meas}}_{j,k} \right)^2. \quad (15)$$

Here, m' is the misfit function for virtual FMC data and $\tau'^{\mathrm{model}}_{j,k}$ is the modelled arrival time between $\mathbf{S}'_j$ and $\mathbf{R}'_k$. The measured travel time between $\mathbf{S}'_j$ and $\mathbf{R}'_k$ is $\tau'^{\mathrm{meas}}_{j,k}$. Crucially, this virtual measurement hinges on the virtualisation step (Sec. III A). Given convexity of the ROI boundary, it is trivial to calculate the arrival times $\tau'^{\mathrm{model}}_{j,k}(\mathbf{c}^{\mathrm{bi-velocity}}) = \|\mathbf{R}' - \mathbf{S}'\|/c_{\mathrm{ROI}}$ and, hence, the misfit function. The optimisation in Eq. (15) can, thus, be solved with a golden section search (or other rudimentary algorithms) to rapidly iterate towards c_{ROI}, giving the bi-velocity starting model.

The high-contrast robustness of VIRT is, therefore, increased by this rapid and simple initialisation step. By reducing the initial travel-time misfit, the initialisation step decreases the number and depth of fluctuations in the back-traced misfit field despite high contrast, reducing local minima and increasing robustness. Similar to Sec. III A, this only requires *a priori* information about the exterior but not the unknown high-contrast interior.

08 September 2025 15:08:20

https://doi.org/10.1121/10.0036902

IV. VALIDATION METHODOLOGY

Resulting from the novelty of the proposed algorithm, the problem is distilled into two key challenges for VIRT:

(1) Inaccessibility of the ROI; and
(2) high contrast between the ROI and the surroundings.

If these are unresolved, the first challenge results in a limited resolution and the second challenge results in the BRT step converging to a local rather than global optimum. Our primary goal is to show that VIRT achieves increased resolution and robustness compared to BRT. The secondary goals are to compare VIRT's performance with the theoretical resolution limit [Eq. (6)] and explore its benefits for bone imaging.

In this vein, we will prioritise characterising the algorithm's performance against these metrics rather than the realism of the sample problems posed (Sec. IV B). We have validated VIRT *in silico* and limited ourselves to a two-dimensional (2D) problem, which is discussed in Sec. IV C. Also, we did not fine-tune variables such as frequency and signal choice, which are largely constrained by the problem at hand. Section IV E discusses the practical implementation of VIRT and other compared scenarios. Section V presents the reconstructions and discusses our findings.

A. Compared imaging configurations

We compared five imaging configurations besides the input model (GroundTruth; see Table I): VIRT, BRT, InitBRT, NearBRT and VIRT20dB. VIRT and BRT have been detailed earlier.

InitBRT represents a realistic best-case scenario for BRT, which is to show that any resolution gains in VIRT are not solely caused by the *a priori* information provided (c_{ER} and ROI boundary). An initialisation step is added before the inversion step in Fig. 2, which solves the more cumbersome Eq. (14) to obtain a starting bi-velocity model. The same FMC data as that for VIRT and BRT are used.

NearBRT represents the inherent resolution limit of BRT at the normally inaccessible boundary by artificially introducing transducers (dipoles) *in silico* at the same locations as virtual transducers in VIRT. It is the only configuration that does not use the FMC data at the physical array. Initialisation with the bi-velocity model is achieved by solving the simple Eq. (15) before inversion.

VIRT20dB represents VIRT's performance under noisy conditions—the addition of noise at 20 dB signal-to-noise ratio (SNR) will be detailed in Sec. IV D.

TABLE I. Abbreviations of compared imaging configurations.

Abbreviation	Description
VIRT	Virtual initialised ray tomography
BRT	Bent-ray tomography
InitBRT	BRT with added initialisation step
NearBRT	BRT from the inaccessible ROI boundary
VIRT20dB	VIRT with added noise at 20 dB SNR

B. Sample problems

The setups of the two sample problems resemble Fig. 3, where the ROI is adjusted for each problem. Although secondary in evaluating VIRT's efficacy, we used realistic values (e.g., speed of sound and density) for biological tissues from IT'IS Foundation (2022), which are highlighted in Table II.

1. Smiley problem (Smiley80)

The first sample problem is that of a 2D bone smiley face embedded within soft tissue (Smiley80), shown in Fig. 5(a). The easily reproducible and analysable geometric shapes aid visualisation of VIRT's resolution gains. Figure 5(c) shows the complete setup, in which the ROI is offset by $x = +5$ mm, $y = +2.5$ mm relative to the centre of the array. The array has a diameter of 80 mm, giving a medium-range but high-contrast problem as a proof of concept. The smaller domain also makes experimentation and iteration much easier.

The detailed geometry and material properties (with reference to Table II) are as follows. The diameter of the circular smiley face is 24 mm. The properties of cortical bone are used on the face, and the properties of muscle are used for the surrounding soft tissue. The eyes are located 4.8 mm above the centre of the smiley and spaced 9.6 mm apart from each other (centre to centre). The mouth is 1.5 mm below the centre of the face and described by a 14.4 mm diameter circle with points over 2.4 mm above its centre reflected. The left eye, right eye, and mouth have a velocity and density of +1 standard deviations (s.d.), −1 s.d., and −0.5 s.d. relative to cortical bone, respectively.

2. Simplified bone model (Bone200)

The second sample problem (Bone200) uses reproducible geometries and simple material properties to aid in characterising VIRT's performance in femoral neck imaging. Figure 5(b) shows a detailed view of the ROI (a simplified ellipsoid bone model) with cancellous bone encased by cortical bone. Figure 5(c) shows the complete setup, in which the ROI is surrounded by soft tissue (the ER), as would be the case *in vivo*.

We based Bone200 on female reference data as women are more likely to suffer from osteoporosis. The asymmetrical periosteum (outer cortical surface) is simplified as an ellipsoid (minor axis, 24.5 mm; major axis, 30.5 mm; minor/major ratio, 0.8), where cortical thickness is averaged from fracture cases and controls of the same age (Zebaze *et al.*, 2005; Bell *et al.*, 1999). The periosteum and endosteum (inner cortical surface) are then constructed with cubic spline interpolation using values from Table III with periodic end conditions.

The material properties include soft tissue for ER, homogeneous cortical bone for the entire cortical region, and a linear gradient of cancellous bone velocity and density of −1 s.d. to +1 s.d. horizontally from left to right (see Table II). This gradient is introduced to differentiate between imaging capabilities for cancellous and cortical bone as both contribute to bone strength. It is set in an

08 September 2025 15:08:20

TABLE II. Properties of biological tissues (IT'IS Foundation, 2022). "±" denotes the standard deviations (s.d.) of the properties.

Tissue	Velocity[a] (ms^{-1})	Density ($kg\ m^{-3}$)	Poisson ratio	Formulation
Cortical bone	3514.9 ± 420.3	1908 ± 133	0.3	Solid
Cancellous bone	2117.5 ± 288.7	1178 ± 149	0.3	Solid
Soft tissue[b]	1588.4 ± 21.6	1090 ± 52	0.5	Fluid[c]

[a]Velocity corresponds to dilational wave velocity. Shear wave velocity was inferred from Poisson ratio for solids.
[b]Muscle is used as the proxy value for soft tissue in the ER for both sample problems.
[c]Shear waves are very slow and highly attenuated in soft tissue.

orthogonal direction to the largely vertical change in cortical thickness to separate the effects. For the cancellous and cortical bones, we decided against a more realistic porous structure to simplify reconstruction comparison. Replacement with empirical values of IT'IS Foundation (2022) was possible as the first arrival wave in cancellous bone, the fast Biot wave, is nondispersive (Hosokawa and Otani, 1997). We have simplified the ER by assuming that all soft tissue (muscle and fat) has the same properties as muscle.

The ER size is of specific concern as this affects the minimum array diameter. *In vivo* femoral neck imaging must be performed at an oblique angle to obtain a cross section along the femoral neck axis and avoid the pelvic bone. This oblique slice starts between the greater trochanter and the iliac crest on one side, ends below the inferior pubic ramus (sit bone), and has a diameter approximately 0.7 times buttocks circumference. The diameter of ER (soft tissue), approximated as a circle, would then be 200–300 mm for women of all ethnicities and age groups from anthropometric data (McDowell *et al.*, 2009). For the current *in silico* study [Fig. 5(c)], we conservatively use a 200 mm diameter for the ER, where the transducers are equally spaced around its boundary. The bone was offset by 20 mm horizontally to account for anatomy.

C. FMC

FMC is executed *in silico* through Pogo (Huthwaite, 2014a), a finite-element modelling software run on graphics processing units (GPUs). We used a displacement-based fluid-solid-coupling model (Simillides *et al.*, 2024). In this model, soft tissue was treated as fluid because of its low shear-wave velocity, whereas bone was treated as solid with a Poisson ratio of 0.3. An unstructured triangular mesh with 30 elements per ER wavelength was used with a 25% denser grid at transducers and boundaries. To largely eliminate boundary reflections such that the simulated domain is a subset of an infinite space, an absorbing boundary using the stiffness-reduction method (SRM; Pettit *et al.*, 2014) was added to the circular propagation space's limits. The propagation space and absorbing boundary were tuned to limit the simulation burden. The absorbing boundary is at least nine wavelengths thick and spaced at least 10 mm from the array.

A 5-cycle, 1 MHz tone burst signal is sent from the transducers. In NearBRT, a single node is inserted at the locations of the VIRT virtual transducers. These transducers transmit the signal as a force in the normal direction of the local ROI boundary, corresponding to a dipole dilational wave. In all other cases, the physical transducers displayed in red in Fig. 5 transmit and receive purely dilational waves.

The spacing between the physical array transducers is equal to half the wavelength of the ER (soft tissue) such that the array is sufficiently sampled (317 for Smiley80 and 792 for Bone200). If fewer transducers were used, it is expected that BRT performance would deteriorate a lot more rapidly than VIRT as the ray-sampling density drops. At minimum, VIRT requires only a pair of physical transducers facing

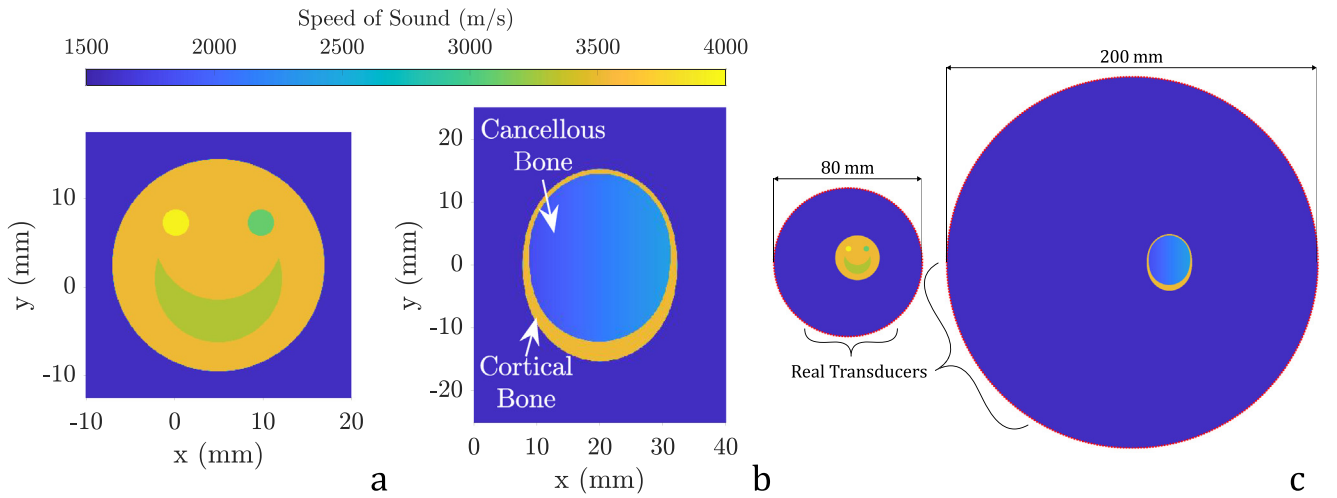


FIG. 5. (a) Smiley80 and (b) Bone200 are displayed in a detailed view. (c) Both problems are shown to scale within the physical array (transducers are in red).

08 September 2025 15:08:20

https://doi.org/10.1121/10.0036902

TABLE III. Values used for construction of the Bone200 example.

Azimuth angle (°)	Diameter (mm)	Thickness (mm)
0	24.5	1.12
45	27.5	0.72
90	30.5	0.75
135	27.5	0.72
180	24.5	1.07
225	27.5	1.97
270	30.5	3.14
315	27.5	2.16

each virtual transducer, although the noise-rejection effect of virtual focusing would drastically decrease given such a low sampling density. This yields a theoretical minimum of 3–4 transducers needed to perform VIRT in the extended range given full-view coverage of a relatively centred ROI from the ER. For NearBRT, the virtual transducers on the ROI boundary are replaced with physical transducers *in silico*.

Given the varying material properties used and the varying mesh density, a highly conservative value is employed for the time step ($t = 5.76$ ns) to always satisfy the Courant criterion (Courant number ≈ 0.4). Forward simulations of Smiley80 took 25 min and those for Bone200 needed 11 h to run on a computer with an GeForce RTX 2080 Ti (4-core) GPU (Nvidia, Santa Clara, CA). We note that FMC *in vivo* or *ex vivo* does not require FE simulation, simply transmitting and capturing waves from the experimental setup is sufficient.

D. Additive noise (VIRT20dB only)

A control is set up for each sample problem to standardise the signal strength for SNR calculations. The ROI in this control is centred, converted into a perfectly cylindrical shape with the same mean diameter, and filled with cortical bone. The same signal as was used in Sec. IV C is transmitted and subsequently received at the opposite receiver. This received signal is enveloped to give the signal strength and generate bandwidth-matching Gaussian noise. Taking the root mean square of this gives the noise strength. The noise is added to the FMC signal to achieve a SNR of 20 dB.

E. Post-processing

The FMC data from Pogo are transferred into MATLAB [R2023b; (MathWorks, Natick, MA)] for further processing on a dual Intel (Santa Clara, CA) Xeon CPU E5-2667 v4 system (16 cores and 32 parallel threads).

1. Virtualisation (VIRT and VIRT20dB only)

For VIRT, we first migrate the FMC time traces from the physical array to the virtual transducers, which are spaced equally around the ROI boundary. These virtual transducers are oversampled (300 for Smiley80 and 400 for Bone200) as there are no restrictions, and oversampling increases robustness. In preparation for the virtualisation step, a cosine ramp filter is applied to smoothly remove the last 10% of the FMC time traces. The remaining time traces are Fourier transformed into the frequency domain before the lowest 1% of frequencies are filtered out.

Based on the locations of the physical and virtual transducers and the background speed between them, we calculate the corresponding signal correction. The signal corrections for the transmission phase and receiving phase are applied through a page-wise matrix multiplication, where each page corresponds to a single frequency. This is then inverse Fourier transformed back to the time domain to form the virtual FMC time traces. We time window these traces to retain the shorter, relevant portion as the virtual FMC wavefronts traverse only the ROI and not the entire setup. The virtualisation step takes 14 s for Smiley80 and 2 min for Bone200.

The angular weighting B in Eq. (11) was applied with $n = 1$ for Smiley80 and $n = 2$ for Bone200. This was found to make arrival-time picking procedurally easier with no influence on the actual arrival time.

Figure 6(b) shows an intermediate result of the virtual signal received when virtual transducer one is fired. A high-pass filter has already been applied to the virtual signals to suppress low-frequency noise and facilitate arrival picking (discussed in Sec. IV E 2). Comparing Fig. 6(b) (VIRT) with Fig. 6(a) (NearBRT) highlights the strong similarities in the received traces despite the presence of noise before the main arrival (green circle) and the absence of the boundary wave (red oval) in the virtualised signals.

The boundary wave, although visible in NearBRT, is largely irrelevant to subsequent steps and may be exaggerated by the dipole transducers. The primary challenge is the noise preceding the arrival, which is mitigated by filtering, enabling successful arrival-time picking and demonstrating the effectiveness of the virtualisation process.

The excitation patch of the virtual source transducer [Fig. 6(b)] is found to be less focused than the dipole point source [Fig. 6(a)]. This defocusing is exaggerated when n in Eq. (11) is increased (e.g., Bone200) as a smaller portion of transducers dominate in contributions, similar to summing over a more limited opening angle.

2. Arrival-time picking (all cases)

For VIRT, in preparation for arrival-time picking, the virtual time traces are high-pass filtered by means of rolling mean removal to remove low-frequency components. We found that low-frequency noise can arise from the fluid-solid implementation of Pogo: after the arrival of a wave, energy at the nodes may take longer than expected to dissipate. This non-sinusoidal noise follows a geometric decay pattern, hence, could not be cancelled out through interference. The sum of hundreds of these noises, present in the physical FMC data, amplifies the noise in the virtual FMC data to a level that can mask the arrival wavefront.

08 September 2025 15:08:20

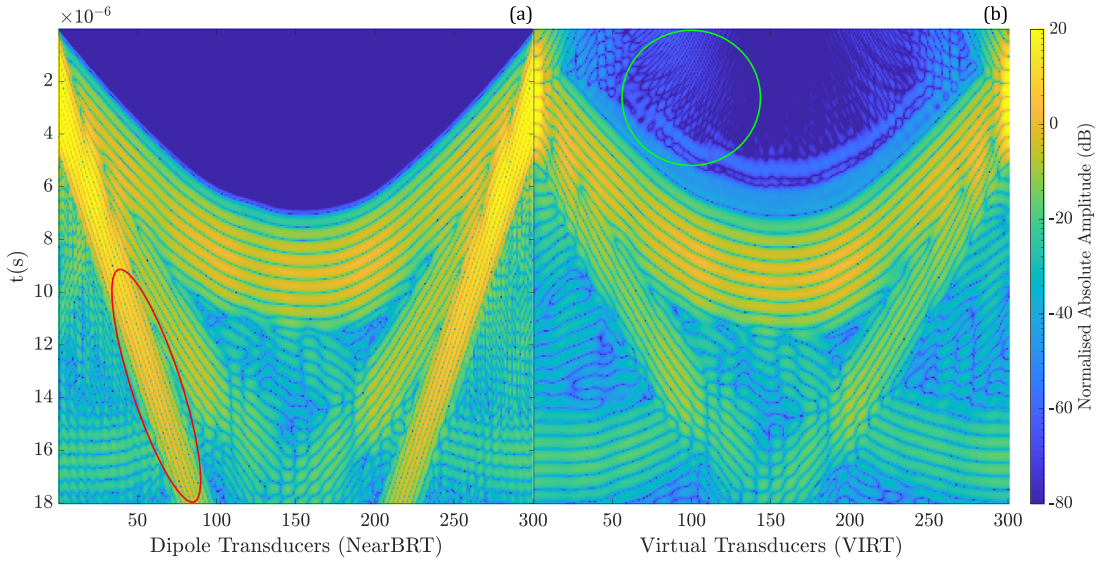


FIG. 6. Intermediate result shows a comparison of received signals at the ROI boundary in (a) NearBRT and (b) VIRT for Smiley80. For VIRT, a high-pass filter is applied by rolling mean removal to suppress low-frequency noise for easier arrival-time picking. The signals are fired from the transducer 1 and normalised around the opposing transducer (150). The two traces show strong similarities overall despite the presence of noise before the main arrival (green circle) and the absence of the boundary wave (red oval) in VIRT.

For all cases, we then pick arrival times through a combination of algorithmic picking (e.g., zero-crossing-based, thresholding-based, and moving-max-based), followed by manual inspection and correction. More sophisticated and fully automated methods exist, but this is not the focus of this paper. We have excluded this element from the processing time calculations.

3. Initialisation (all cases except BRT)

Based on the picked arrival times, the background sound speed is initialised to a bi-velocity starting model (Sec. III B). For VIRT, NearBRT and VIRT20dB, a golden section search between 1500 and 4000 m/s solves the optimisation equation (15) until c_{ROI} converges within 1 m/s, taking a fraction of a second.

For InitBRT, the optimisation equation (14) was used in the golden section search, which required an eikonal solver. We used the same FMM-based eikonal solver as in the forward model of the inversion step (Sec. IV E 4), implemented in C++ (Sethian, 1996). For Smiley80, computation at 0.1 mm resolution took 5.3 min, whereas for Bone200 at 0.5 mm, it took 2.6 min.

4. Inversion (all cases)

Finally, for the BRT inversion step, we used an algorithm based on Li *et al.* (2006) and Hormati *et al.* (2010), which is implemented in C++. Using the bi-velocity starting model, data were reconstructed with minimal, consistent regularisation. Total variation regularisation was applied to both samples with additional first-derivative regularisation in Bone200.

The default resolution is set to 0.1 mm. In some cases, stability issues with the eikonal solver were apparent, and in these cases, or where it is very obvious that the algorithm converged to a local optimum, the inversion step is rerun with a worse resolution until the aforementioned problems are largely resolved. The stopping criterion for the inversion process is when the relative change in the misfit $m(\mathbf{c}^{\mathrm{model}})$ [Eq. (3)] between iterations is less than 0.001%.

V. RESULTS

Smiley80 and Bone200 reconstructions are displayed in Figs. 7 and 8, respectively, where the lines in Figs. 7(a) and 8(a) mark cross section locations, which are shown in

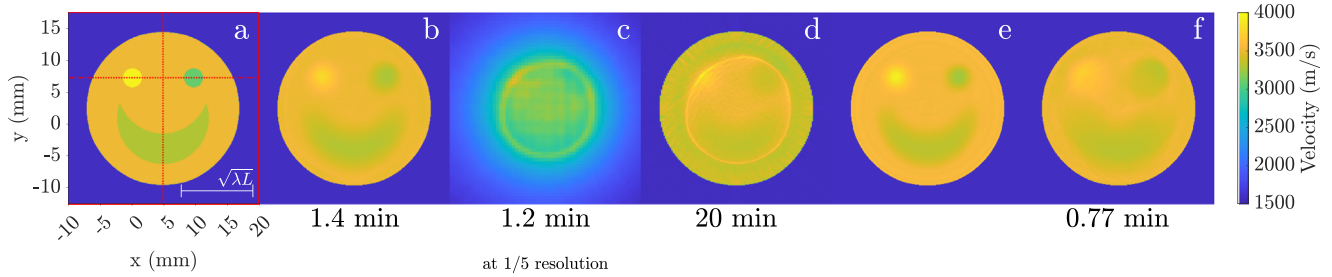


FIG. 7. Smiley80 is depicted by (a) GroundTruth, (b) VIRT, (c) BRT, (d) InitBRT, (e) NearBRT, and (f) VIRT20dB. Processing times are shown below reconstructions of physically implementable algorithms.

08 September 2025 15:08:20

https://doi.org/10.1121/10.0036902

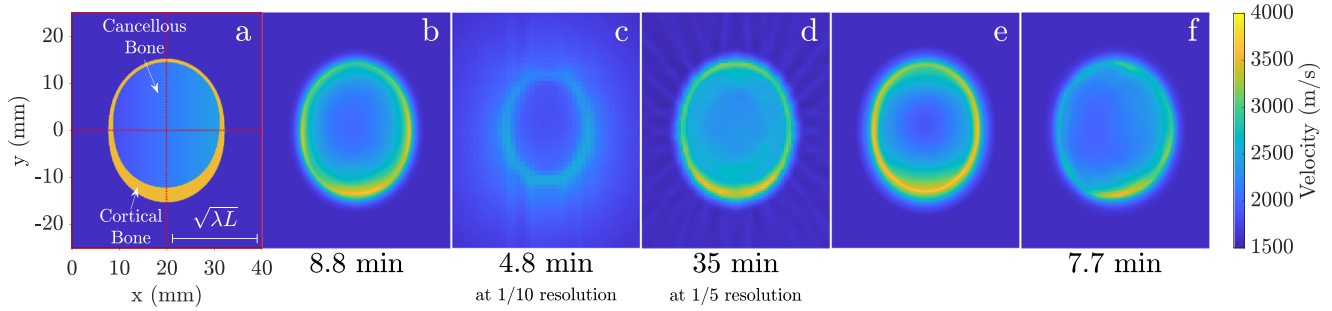


FIG. 8. Bone200 is depicted by (a) GroundTruth, (b) VIRT, (c) BRT, (d) InitBRT, (e) NearBRT, and (f) VIRT20dB. Processing times are shown below reconstructions of physically implementable algorithms.

Figs. 9 and 10. In Smiley80, the horizontal cross section is taken at the height of the eyes, whereas all other cross sections in Smiley80 and Bone200 pass through the centre of the ROI. Figure 11 and Table IV show the mean absolute errors when each reconstruction is compared against the visible regions in Figs. 7(a) and 8(a).

A. BRT

It is clear that VIRT far outperforms BRT from Figs. 7(a)–7(c) and 8(a)–8(c).

The robustness of BRT is far worse than that for VIRT because it was unable to invert at the same resolution. This is exacerbated by the increased array diameter in Bone200. Despite requiring even lower resolution for successful inversion, Fig. 8(c) still shows ray-like artefacts that signify low robustness. Because the first arrival wave consists largely of direct transmissions between physical transducers, there are very few arrival-time pairs that encode information about the ROI. This reduces spatial sampling and encourages convergence to local minima, giving ray-like artefacts.

In terms of resolution, BRT gives a very poor reconstruction of the bulk geometry, which contains significant morphological inaccuracies. The ROIs in Smiley80 and Bone200 shrank by around 40% and 23%, respectively, from the first to the last local maximum in Figs. 9(b) and 10(a). However, we note that this measurement is affected by the poorly defined ROI boundary, where values of the ER and ROI leak into the opposite region. In Fig. 7(b), we can observe that the facial features of Smiley80 are largely fused. Wavefront healing contributes to all of the above problems.

In Fig. 8(b), no significant variation in cortical thickness or cancellous bone is visible. The cancellous bone is observed to be horizontally symmetric in Fig. 10(a). The cross section in Fig. 10(b) shows a *vertical* internal gradient absent in GroundTruth. This suggests that the thicker lower cortical bone with a uniform speed of sound caused BRT to incorrectly predict a higher speed of sound in the cancellous bone. The mean absolute errors of BRT for Smiley80 and Bone200 are much higher than those for all of the other cases (Fig. 11 and Table IV).

B. InitBRT

InitBRT, which is provided with *a priori* information (c_{ER} and ROI boundary) and then undergoes the less efficient initialisation step, performs much better than BRT. Specifically, using InitBRT instead of BRT gives a reduction of 88% and 42% in mean absolute errors for Smiley80 and Bone200, respectively (Fig. 11 and Table IV). However, these reconstructions are still unsatisfactory, arising from the unresolved errors in BRT's ray-based forward model. The discrepancies between BRT's result and GroundTruth show that VIRT is needed to fully exploit the *a priori* information.

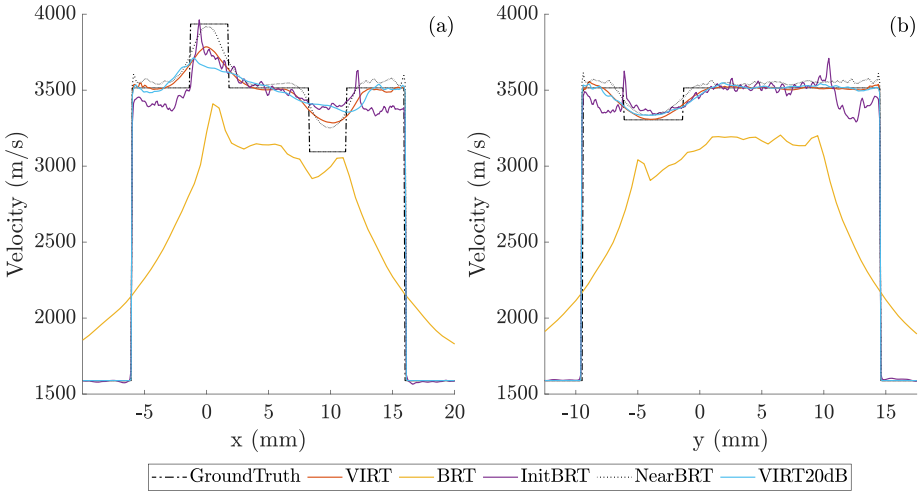


FIG. 9. Cross sections of Smiley80 are shown as (a) horizontal at eye height and (b) vertical at the centre.

08 September 2025 15:08:20

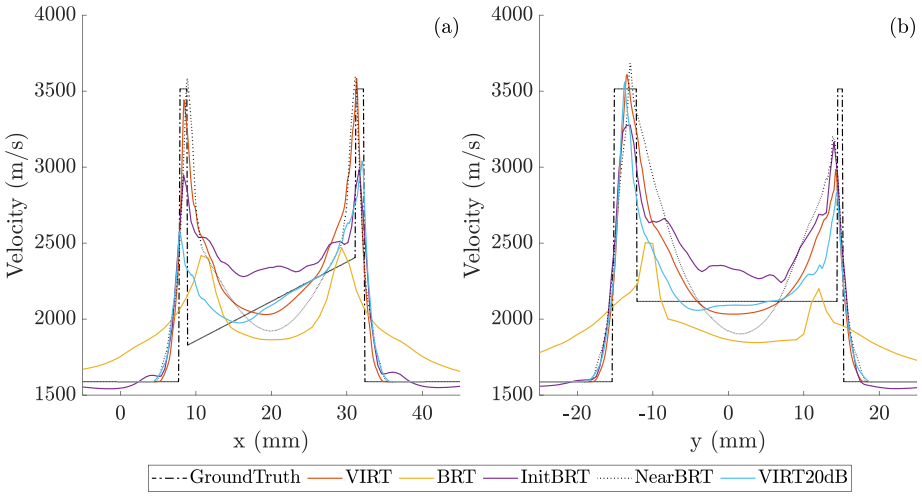


FIG. 10. Cross sections of Bone200 are shown as (a) horizontal and (b) vertical.

InitBRT robustness is still poor. Although Smiley80 could be inverted at the default resolution, a lower resolution was required for initialisation and inversion for the larger Bone200. Smiley80 and Bone200 are better than BRT but worse than VIRT. In addition, ray-like artefacts are visible in Smiley80 and Bone200, which encourage convergence to local minima and signify poor robustness. In Smiley80 [Fig. 8(d)], these ray-like artefacts appear as a ring (diameter ≈ 17 mm) near the ROI boundary, the ROI's centre, and the left eye. In the larger Bone200, the ray-like artefacts are radiating outward into the ER [see Fig. 8(d)]. This is despite the lower resolution and added regularisation in the inversion step (Sec. IV E 4) for Bone200. This suggests that the need for VIRT increases with array diameter.

InitBRT resolution is slightly improved in Smiley80 compared to that in BRT but still poor in Bone200, which has a larger array diameter.

In Smiley80, we note that the aforementioned ring-like artefact in InitBRT [Fig. 7(d)] has a shape and diameter similar to the ROI boundary of BRT's reconstruction [Fig. 7(c)]. This suggests that the error originates from inaccuracies embedded in τ^{model}, calculated from the ray-based forward model, which results in wavefront healing. This artefact masks the poor resolution of Smiley80. In Fig. 9(a), the ring-shaped artefact increases the height of the left eye's reconstruction, whereas the right eye shows only 30% of GroundTruth's velocity change. Both eyes are enlarged with poorly defined boundaries. The ring-shaped artefact also creates a false distinction between the mouth and "chin" by sharply increasing the velocity at the lower edge [Fig. 7(d)]. Instead, Fig. 9(b) shows that the InitBRT mouth is a smudge extending beyond the true mouth.

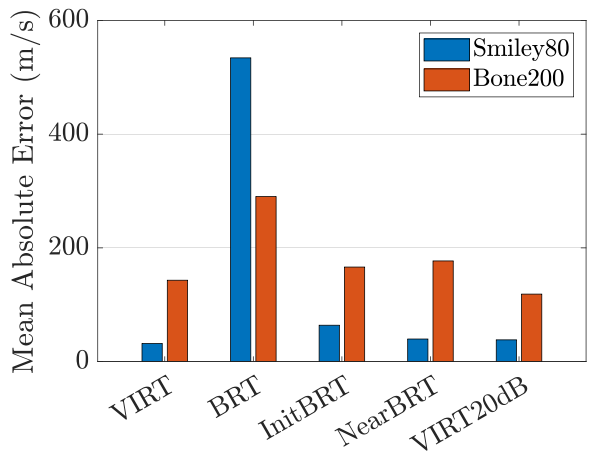


FIG. 11. Error for each sample problem and algorithm.

In Bone200, InitBRT gives a more accurate reconstruction of the cortical than the cancellous bone region. The cortical bone is partially reconstructed, albeit the peaks in Fig. 10 still only reflect between 50% and 80% of true velocity variation, compared to the surroundings. The cancellous bone has values higher than GroundTruth as displayed in Fig. 10, and the horizontal gradient is not reflected by the symmetric cancellous properties in Figs. 8(d) and 10(a). Similar to BRT, InitBRT incorrectly predicts a *vertical* internal gradient in Fig. 10(b), influenced by the change in cortical thickness.

In addition, there appears to be an elliptical ring of increased speed of sound in Fig. 8(d), matching the location of the cortical bone in the BRT result [Fig. 8(c)]. The cross sections (Fig. 10) confirm this with the elliptical ring's clear alignment against the mislocated BRT peaks. This shows that the *a priori* information and initialisation step alone cannot correct for intrinsic errors in the ray-based forward model, especially in extended-range tomography as in bone imaging.

InitBRT is much slower than VIRT and BRT, taking 14 times the reconstruction time of VIRT for Smiley80 while requiring the same *a priori* information. VIRT is faster, higher resolution, and more robust. BRT is faster and does not need *a priori* information. Hence, InitBRT lacks

TABLE IV. Mean absolute error for each sample problem and algorithm.

	VIRT	BRT	InitBRT	NearBRT	VIRT20dB
Smiley80	31.8	534	64.0	39.7	38.2
Bone200	143	290	166	177	119

08 September 2025 15:08:20

https://doi.org/10.1121/10.0036902

practical advantages in tomography. However, for high-contrast extended-range tomography, *a priori* knowledge of the ER and ROI boundary is evidently crucial.

C. VIRT and NearBRT

The robustness of VIRT and NearBRT are much better than that for BRT and InitBRT. Ray-like artefacts are absent in Smiley80 and Bone200 [Figs. 7(b) and 7(e) and Figs. 8(b) and 8(e), respectively]. In addition, VIRT and NearBRT can be inverted at the default high resolution without problems.

VIRT and NearBRT are nearly identical and have higher resolutions than BRT and InitBRT (see Figs. 7 and 8). Both appear as slightly blurred versions of GroundTruth, which are noticeable in the eyes and mouth of Smiley80 [Figs. 7(a), 7(b), and 7(e)] and near the cortical bone of Bone200 [Figs. 8(a), 8(b), and 8(e)]. This blurring is expected as the features approach the resolution limit [Eq. (6)], evaluated from the ROI boundary. The similarity between VIRT and NearBRT indicates that VIRT is close to achieving this limit.

For Smiley80, the left eye reconstruction of NearBRT seems much better than that of VIRT [Fig. 9(a)]. However, both NearBRT profiles in Fig. 9 exhibit a systematic velocity increase ($\approx$25 m/s) in the ROI, which is likely caused by the slightly earlier arrival-time picks (e.g., overly sensitive threshold). This is why the mean absolute error for NearBRT appears worse than VIRT in Table IV. Correcting for this systematic error, we can reinterpret that NearBRT performs equally well for each eye and slightly better than VIRT in accuracy of peak height and width. Similar correction of Fig. 9(b) shows that VIRT and NearBRT have nearly identical mouth peak heights, both reaching 100% of GroundTruth depth, where VIRT has a slightly worse peak width. Figures 7(b) and 7(d) show VIRT features are slightly blurrier, especially the smaller features (eyes).

This extra blurring in VIRT is likely the result of two factors. First, the NearBRT transducers were modelled as a single node in Pogo, applying a force normal to the ROI boundary, creating a point-source transducer with an infinite excitation spot. These theoretical transducers cannot be manufactured, representing a theoretical limit. Second, any focusing error during virtualisation would enlarge the already finite excitation spot on the ROI boundary, which is equivalent to the sum of multiple point sources. The combined wavefront from these equivalent sources will arrive at different times than those for the infinite point source modelled by NearBRT, resulting in a blurring effect. This effect is more apparent for smaller features, where feature boundaries are much closer to each other. This also explains the slightly smoother profile in Fig. 9(b) of VIRT compared to that for NearBRT.

In Bone200 (Fig. 8), VIRT and NearBRT reconstructions are much more accurate than those for BRT and InitBRT but slightly blurrier than in Smiley80 (Fig. 7). The first-derivative smoothing regularisation was necessary because of inversion issues (Sec. IV E 4) as ray theory cannot model guided waves travelling circumferentially along the cortical bone.

The cancellous interior is better reproduced by VIRT than NearBRT. The vertical profile in Fig. 10(b) shows NearBRT overshooting the cancellous interior value, whereas the horizontal profile in Fig. 10(a) shows a symmetric cancellous region, lacking the horizontal gradient in GroundTruth. This results in the mean absolute error for VIRT being lower than that for NearBRT (Table IV). This is because NearBRT's first arrivals consist entirely of the guided circumferential lamb-type (CLT) wave, excluding through-transmission waves, which encode cancellous interior information. In VIRT, the central third of the cancellous interior more closely follows the horizontal gradient [Fig. 10(a)] and the non-sloping vertical profile [Fig. 10(b)] of GroundTruth. The enlarged finite excitation spot in VIRT partially suppresses the CLT wave, reducing interference and making through-transmission arrivals easier to identify.

Cortical reconstructions (Fig. 8) of GroundTruth by VIRT and NearBRT are much more accurate than those of BRT and InitBRT. The horizontal profile [Fig. 10(a)] shows that VIRT and NearBRT peak heights match those for GroundTruth almost perfectly for the sides, whereas the vertical profile [Fig. 10(b)] shows a tendency to overshoot for thicker regions and undershoot for thinner regions. A closer inspection of Figs. 8(b) and 8(e) shows that NearBRT is sharper and more consistent, whereas VIRT has variations in thickness and velocity. This may relate to regions where first arrivals transition between through-transmissions and guided waves, indicating a trade-off between information on the cancellous interior and cortical exterior.

The cortical thicknesses in the VIRT and NearBRT reconstructions exceed those for GroundTruth at half-peak height Fig. 10. As mentioned above, the first-derivative regularisation likely contributes to blurring of the boundaries of the periosteum and endosteum.

D. VIRT under noisy conditions

VIRT20dB shows that VIRT's performance remains good under relatively noisy conditions. For Smiley80, Fig. 7(f) depicts a blurrier reconstruction than that for VIRT without noise and NearBRT, but this is still significantly better than BRT and InitBRT reconstructions. This blurriness is also observed in the shallower and more spread out eyes in Fig. 9(a), whereas the mouth is relatively similar to that for VIRT in Fig. 9(b).

For Bone200, the results are better in cancellous than cortical bone. The cortical bone is patchier and less well reconstructed than VIRT without noise. Whereas the peaks of VIRT20dB in the vertical profile [Fig. 10(b)] approximate those of VIRT, in the horizontal profile [Fig. 10(a)], the peaks only reflect 40%–50% of the velocity variations. The poor cortical reconstruction reflects the need to pick stronger and, hence, later through-transmission arrivals as noise overpowers weakened guided wave arrivals.

08 September 2025 15:08:20

However, Fig. 8(f) shows a very clear horizontal gradient, which is also reflected in Fig. 10(a). The vertical profile [Fig. 10(b)] confirms a lack of "hallucinated" gradients as were the cases in BRT and InitBRT, showing true sensitivity to variations in the cancellous bone. The picking of through-transmission waves results in a cancellous reconstruction that is even better than that of VIRT, giving an even lower mean absolute error (Table IV)—a metric biased towards measuring cancellous reconstruction accuracy. Additional VIRT results (not shown) depict that similarly good cancellous reconstructions can be achieved under the non-noisy conditions at the expense of worse cortical reconstructions.

VI. DISCUSSION

The results indicate that VIRT has achieved the goals of higher extended-range resolution and high-contrast robustness. They also reveal that the (1) *a priori* information supplied to VIRT, (2) virtualisation step, and (3) initialisation step are crucial in solving high-contrast extended-range problems.

Guided waves ought to be poorly modelled. The eikonal equation (2) models neither interference nor reflection of partial waves in waveguides, and the assumption $c_{\text{phase}} = c_{\text{group}}$ is inexact for dispersive waves. Despite this, VIRT's Bone200 reconstruction surpassed expectations in terms of precision and similarity. This sets a foundation for further analysis with guided-wave-specific algorithms by approximating the thickness and material properties. Our results suggest that virtual transducers can be employed as a phased array to excite specific guided waves from the ROI boundary given their strength even when not intentionally excited.

Although it is not our focus, arrival-time picking was challenging as a result of artefacts from inexact virtualisation focusing and guided waves in Bone200. The result that VIRT20dB has lower mean absolute error than VIRT in Bone200 suggests that picking through-transmission arrivals results in a more accurate reconstruction of cancellous bone, which occupies a larger region. However, picking the weak, first-arriving guided waves gives a more accurate reconstruction of the cortical bone. As such, cortical or cancellous accuracy could be prioritised according to clinical needs through arrival picking. Alternatively, the two reconstructions may even be fused after inversion.

This work demonstrates that although arrival-time picking is an element of successful BRT imaging (as suggested by Falardeau and Belanger, 2018), reformulating the BRT problem is of vital importance. Placing the array on the ROI boundary and managing the high contrast through initialisation allows BRT to achieve accurate results. VIRT demonstrates that this can be achieved virtually, making it a practical solution when access is limited.

Solutions to similar problems have been developed elsewhere with a notable method from Taskin *et al.* (2020) and Taskin and van Dongen (2020) for redatuming ultrasound data in breast ultrasound tomography. Their approach mathematically decomposes the scattered field into a series of Hankel functions, enabling effective transducers to be placed at arbitrary locations throughout the domain. The method in these papers is more mathematically rigorous, whereas our approach is more physically intuitive. Implementation is simplified by avoiding the need for matrix inversions, regularisation, and the resulting arbitrary tuning parameter. It is suggested that the inversion is effectively bypassed in our virtualisation approach because of the orthogonal nature of the waves propagated and summed, although this has not been mathematically demonstrated. Although our method can establish the initial arrivals, shrinking the domain for full-waveform inversion, as desired in the referenced papers, may require a more rigorous approach to accurately capture the full wavepacket.

In the current work, validation was completed with an *a priori*, homogeneous ER. The resolution and robustness gains of VIRT in this context show a clear demand for VIRT to be extended to an *a priori*, inhomogeneous ER. In addition, we have established that *a priori* information, such as ER velocity and ROI boundary location, is crucial for successful inversion.

Future work includes validation of VIRT in 3D instead of 2D. This is unlikely to be significantly more challenging in a homogeneous ER, provided that A_{3D} is used in focusing [Eq. (12)]. Similar resolution should be achievable.

For *in vivo* femoral neck imaging, high-contrast inhomogeneities in the ER, similar to bone protrusions (e.g., greater trochanter and pelvic bone), may appear near the imaging plane. VIRT is less affected by these high-contrast scatterers as a result of 3D focusing during virtualisation, which limits energy to the protrusions, improving SNR. In any case, we have demonstrated that VIRT is highly tolerant to bandwidth-matched noise as may arise from bone-protrusion scattering. Any affected transducer pairs may also be excluded from the physical FMC data before virtualisation. For assumption-free imaging, the ROI could also be redefined to include scatterers near the femoral neck.

In real bone, the endosteum is not a clean smooth boundary as in Bone200 [Fig. 5(b)]. The true endosteum has protruding trabeculae, part of the cancellous bone, which scatter and deflect guided wave energy, effectively suppressing guided waves. Given that the eikonal forward model in the BRT inversion step does not model guided wave effects, the results of VIRT with real cancellous bone may prove to be even better. Moreover, real cancellous bone is not homogeneous, therefore, the fast Biot wave entering radially from the cortical bone may experience lower acoustic contrast, enhancing through-transmission.

In a medical setting, the conservatively chosen 200 mm array for Bone200 serves only the lowest fifth percentile of women of all ethnicities and age groups (McDowell *et al.*, 2009). To cover the 95th percentile, a much larger array diameter of 300 mm is needed. It is expected that the improvements in VIRT over BRT demonstrated here will be even greater *in vivo* because of VIRT's robustness, extended-range resolution, and performance under noisy conditions.

08 September 2025 15:08:20

VIRT could potentially replace classical BRT in numerous biomedical and industrial applications, including breast imaging and guided wave corrosion-patch tomography. In these domains, the acquisition array is customarily oversized to cover the largest ROIs. This leads to a large effective ER:ROI ratio for most ROIs, even without an actual ER obstructing the ROI boundary. This gives VIRT a significant advantage over classical BRT.

VII. CONCLUSIONS

Widely used in tomography, BRT is one of the most robust and fastest tomographic algorithms. However, even BRT struggles with high-contrast extended-range tomography, where transducers have no direct contact with the high-contrast ROI. *In vivo* femoral neck imaging is one such problem and desperately needed by millions of osteoporotic patients to predict hip fracture risk.

We have introduced VIRT, which increases BRT's high-contrast robustness and extended-range resolution. VIRT simplifies the problem by using *a priori* information to segment it into the unknown ROI and the known ER. Two additional steps are introduced: virtualisation and bi-velocity initialisation. The virtualisation step places virtual transducers on the minimum convex boundary around the ROI for BRT-style inversion. We mathematically demonstrated that this unlocks high resolution unattainable with BRT from afar. Virtualisation also increases robustness by (1) boosting ray-sampling density within the ROI, (2) increasing effective SNR, and (3) enabling initialisation, which rapidly estimates the overall properties of the ROI.

VIRT was validated using two 2D *in silico* sample problems: Smiley80 and Bone200. Smiley80 features a smiley-inscribed cylinder with bone-like properties in soft tissue. Bone200 represents an *in vivo* bone imaging setup with a simplified ellipsoid bone model in soft tissue. The ROI has a cortical bone exterior with varying thickness and a cancellous bone interior with varying material properties.

VIRT drastically improves resolution and robustness over BRT and InitBRT (a best case scenario of BRT), while maintaining or improving speed. This improvement becomes pronounced as the ER becomes much larger than the ROI. Furthermore, we show that VIRT approaches resolution achievable from direct access to the ROI boundary (NearBRT) without contact. Minute resolution gains may still be possible, but minor virtualisation errors may actually benefit VIRT in solving the inversion problem. The significant resolution gains of VIRT over BRT are maintained under 20 dB bandwidth-matched noise.

VIRT has demonstrated that high-contrast extended-range tomography can be solved without prior ROI-interior knowledge. Further research may enable VIRT to achieve site-specific, quantitative *in vivo* bone-health imaging, aiding in osteoporosis diagnoses and monitoring.

VIRT also has the potential to replace classical BRT in many other applications as a result of significant resolution gains with large physical array diameters. For further resolution gains, VIRT provides an excellent and trustworthy starting background for algorithms with less robustness and higher computational costs.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

Bell, K. L., Loveridge, N., Power, J., Garrahan, N., Stanton, M., Lunt, M., Meggitt, B. F., and Reeve, J. (**1999**). "Structure of the femoral neck in hip fracture: Cortical bone loss in the inferoanterior to superoposterior axis," J. Bone Miner. Res. **14**(1), 111–119.

Bochud, N., and Laugier, P. (**2022**). *Axial Transmission: Techniques, Devices and Clinical Results* (Springer International, Cham, Switzerland), pp. 55–94.

Falardeau, T., and Belanger, P. (**2018**). "Ultrasound tomography in bone mimicking phantoms: Simulations and experiments," J. Acoust. Soc. Am. **144**(5), 2937–2946.

Fomel, S. (**1997**). "Traveltime computation with the linearized eikonal equation," SEP **94**, 123–131, available at https://sep.sites.stanford.edu/publications/sep-reports/sep-94-1997.

Grimal, Q., and Laugier, P. (**2022**). "Introduction," in *Bone Quantitative Ultrasound. Advances in Experimental Medicine and Biology*, edited by P. Laugier and Q. Grimal (Springer International, Cham, Switzerland), pp. 1–3.

Hormati, A., Jovanović, I., Roy, O., and Vetterli, M. (**2010**). "Robust ultrasound travel-time tomography using the bent ray model," in *Medical Imaging 2010: Ultrasonic Imaging, Tomography, and Therapy*, edited by J. D'hooge and S. A. McAleavey (SPIE, Bellingham, WA), Vol. 7629, p. 76290I.

Hosokawa, A., and Otani, T. (**1997**). "Ultrasonic wave propagation in bovine cancellous bone," J. Acoust. Soc. Am. **101**(1), 558–562.

Hung, S., Dahlen, F. A., and Nolet, G. (**2001**). "Wavefront healing: A banana-doughnut perspective," Geophys. J. Int. **146**, 289–312.

Huthwaite, P. (**2014a**). "Accelerated finite element elastodynamic simulations using the GPU," J. Comput. Phys. **257**, Part A, 687–707.

Huthwaite, P. (**2014b**). "Evaluation of inversion approaches for guided wave thickness mapping," Proc. R. Soc. A **470**, 20140063.

Huthwaite, P., and Simonetti, F. (**2011**). "High-resolution imaging without iteration: A fast and robust method for breast ultrasound tomography," J. Acoust. Soc. Am. **130**(3), 1721–1734.

IT'IS Foundation (**2022**). "Tissue Properties Database V4.1," available at https://itis.swiss/virtual-population/tissue-properties/downloads/database-v4-1/ (Last viewed June 13, 2025).

Kanis, J. A., Johnell, O., Odén, A., Sernbo, I., Redlund-Johnell, I., Dawson, A., Laet, C. D., and Jönsson, B. (**2000**). "Long-term risk of osteoporotic fracture in Malmö," Osteoporos. Int. **11**(8), 669–674.

Kanis, J. A., Odén, A., Johnell, O., Laet, C. D., Jönsson, B., and Oglesby, A. (**2003**). "The components of excess mortality after hip fracture," Bone **32**(5), 468–473.

Lasaygues, P., Espinosa, L., Bernard, S., Petit, P., and Guillermin, R. (**2022**). *Ultrasound Computed Tomography* (Springer International, Cham, Switzerland), pp. 227–250.

Li, S., Mueller, K., Jackowski, M. P., Dione, D. P., and Staib, L. H. (**2006**). "Fast marching method to correct for refraction in ultrasound computed tomography," in *3rd IEEE International Symposium on Biomedical Imaging: Nano to Macro*, April 6–9, Arlington, VA (IEEE, New York), pp. 896–899.

McDowell, M. A., Fryar, C. D., and Ogden, C. L. (**2009**). "Anthropometric reference data for children and adults: United States, 1988–1994," Technical Report (National Center for Health Statistics, Hyattsville, MD), available at https://permanent.fdlp.gov/gpo57963/sr11_249.pdf (Last viewed June 13, 2025).

Miller, D. A. B. (**1991**). "Huygens's wave propagation principle corrected," Opt. Lett. **16**(18), 1370–1372.

08 September 2025 15:08:20

Pettit, J., Walker, A., Cawley, P., and Lowe, M. (**2014**). "A stiffness reduction method for efficient absorption of waves at boundaries for use in commercial finite element codes," Ultrasonics **54**(7), 1868–1879.
Sethian, J. A. (**1996**). "A fast marching level set method for monotonically advancing fronts," Proc. Natl. Acad. Sci. U.S.A. **93**(4), 1591–1595.
Simillides, Y., Huthwaite, P., Kalkowski, M. K., and Lowe, M. J. (**2024**). "A displacement-based finite element formulation for solving elastic wave problems in coupled fluid-solid media on a GPU," Comput. Struct. **299**, 107369.
Svedbom, A., Hernlund, E., Ivergård, M., Compston, J., Cooper, C., Stenmark, J., McCloskey, E., Jönsson, B., Kanis, J. A., and the EU review panel of the IOF (**2013**). "Osteoporosis in the European Union: A compendium of country-specific reports," Arch. Osteoporos. **8**(1–2), 137.
Taskin, U., van der Neut, J., Gemmeke, H., and van Dongen, K. W. A. (**2020**). "Redatuming of 2-D wave fields measured on an arbitrary-shaped closed aperture," IEEE Trans. Ultrason. Ferroelectr. Freq. Control **67**(1), 173–179.
Taskin, U., and van Dongen, K. W. A. (**2020**). "3D redatuming for breast ultrasound," in *Medical Imaging 2020: Physics of Medical Imaging*, edited by H. Bosmans and G.-H. Chen (SPIE, Bellingham, WA), p. 113125H.
Tran, T. N. H. T., Xu, K., Le, L. H., and Ta, D. (**2022**). *Signal Processing Techniques Applied to Axial Transmission Ultrasound* (Springer International, Cham, Switzerland), pp. 95–117.
Williamson, P. R. (**1991**). "A guide to the limits of resolution imposed by scattering in ray tomography," Geophysics **56**(2), 202–207.
Williamson, P. R., and Worthington, M. H. (**1993**). "Resolution limits in ray tomography due to wave behavior; numerical experiments," Geophysics **58**(5), 727–735.
World Health Organization (**1994**). "Assessment of fracture risk and its application to screening for postmenopausal osteoporosis: Report of a who study group [meeting held in Rome from 22 to 25 June 1992]," Technical Report, WHO Technical Report Series 843, available at https://iris.who.int/handle/10665/39142 (Last viewed June 13, 2025).
Zebaze, R. M., Jones, A., Welsh, F., Knackstedt, M., and Seeman, E. (**2005**). "Femoral neck shape and the spatial distribution of its mineral mass varies with its size: Clinical and biomechanical implications," Bone **37**(2), 243–252.

08 September 2025 15:08:20