

## ARTICLE OPEN



## Full-waveform inversion imaging of the human brain

Lluís Guasch<sup>1</sup>✉, Oscar Calderón Agudo<sup>1</sup>, Meng-Xing Tang<sup>2</sup>, Parashkev Nachev<sup>3</sup> and Michael Warner<sup>1</sup>

Magnetic resonance imaging and X-ray computed tomography provide the two principal methods available for imaging the brain at high spatial resolution, but these methods are not easily portable and cannot be applied safely to all patients. Ultrasound imaging is portable and universally safe, but existing modalities cannot image usefully inside the adult human skull. We use in silico simulations to demonstrate that full-waveform inversion, a computational technique originally developed in geophysics, is able to generate accurate three-dimensional images of the brain with sub-millimetre resolution. This approach overcomes the familiar problems of conventional ultrasound neuroimaging by using the following: transcranial ultrasound that is not obscured by strong reflections from the skull, low frequencies that are readily transmitted with good signal-to-noise ratio, an accurate wave equation that properly accounts for the physics of wave propagation, and adaptive waveform inversion that is able to create an accurate model of the skull that then compensates properly for wavefront distortion. Laboratory ultrasound data, using ex vivo human skulls and in vivo transcranial signals, demonstrate that our computational experiments mimic the penetration and signal-to-noise ratios expected in clinical applications. This form of non-invasive neuroimaging has the potential for the rapid diagnosis of stroke and head trauma, and for the provision of routine monitoring of a wide range of neurological conditions.

*npj Digital Medicine* (2020)3:28; <https://doi.org/10.1038/s41746-020-0240-8>

## INTRODUCTION

No universally applicable means of imaging the living human brain at high anatomical resolution exists. The modality with the best spatial resolution and tissue contrast, magnetic resonance imaging (MRI), is contraindicated where the presence of magnetic foreign bodies cannot be excluded, and is impractical with claustrophobic, uncooperative, or severely obese patients. Its nearest rival, X-ray computed tomography (CT), involves exposure to harmful ionising radiation. Both require large, expensive, immobile, high-power instruments that are near-impossible to deploy outside specialised environments. The clinical consequences of this are high symptom-to-image times, long inter-scan intervals during serial imaging, and constraints on the range of patients that can be imaged successfully.

Pre-eminent among the many neurological disorders, where patient outcomes are degraded by these restrictions, is stroke: the second most common cause of death worldwide, and the dominant cause of acquired adult neurological disability<sup>1</sup>. Treatment decisions are here critically guided by neuroimaging, ideally performed immediately after symptom onset. Delays of the order of minutes have substantial impact on outcomes, yet the necessity to treat patients only after transport to hospital routinely introduces delays of an hour or more<sup>2</sup>. Accelerating the treatment of stroke by enabling neuroimaging and treatment to be performed at the point of first contact would thus have large population-level impacts on survival and disability. Analogous arguments can be made for improved rapid medical imaging in head trauma, and in routine intraoperative, post-operative, and preventative neurological monitoring, with the potential to impact large numbers of patients worldwide.


We provide in silico proof-of-principle, supported by ex vivo and in vivo laboratory measurements, that the combination of transmitted transcranial ultrasound tomography with a computationally intensive technique originally developed to image the interior of the Earth, can address these clinical needs by providing

portable three-dimensional (3D) quantitative imaging that is less-expensive, faster, and more easily applicable than MRI, and that is safer and has better soft-tissue contrast than CT. This approach results in a three-dimensional, sub-millimetre-resolution, quantitative model of acoustic wave speed within the brain and surrounding tissue, which is capable of distinguishing most of the structures and pathologies to which MRI is sensitive. The combined findings of our in silico and laboratory experiments demonstrate that recording transcranial ultrasound data are feasible, recorded signal-to-noise levels are high, and the data contain information that is sufficient both to build an accurate model of the skull and to reconstruct brain properties. To achieve this, we use two closely related techniques: adaptive waveform inversion (AWI)<sup>3</sup> to build an accurate skull model, and FWI<sup>4</sup> to extract brain properties at high resolution.

Conventional medical ultrasound is fast, safe, portable, and cheap, but is unable to image the adult human brain at high resolution within the skull; the main reasons for this are well understood<sup>5–8</sup>:

1. In both conventional pulse-echo B-mode sonography<sup>9</sup> and time-of-flight ultrasound CT<sup>10</sup>, high frequencies are required in order to obtain high spatial resolution. Scattering and anelastic losses occur within the skull and the brain, and these increase with frequency. At the frequencies used by conventional ultrasound modalities, these signal losses prevent successful imaging of intracranial soft tissue.
2. The contrast in wave speed between the skull and soft tissues, and between the skull and air-and-fluid-filled cavities within it, produces significant refraction, diffraction, and reverberation of ultrasound energy, as it is transmitted through the skull. This significantly distorts and complicates the consequent wavefront, leading to strong aberrations in both phase and amplitude, and to significant spatial and directional variation in the waveform of the transmitted

<sup>1</sup>Department of Earth Science and Engineering, Imperial College London, London SW7 2AZ, UK. <sup>2</sup>Department of Bioengineering, Imperial College London, London SW7 2AZ, UK. <sup>3</sup>Institute of Neurology, University College London, 33 Queen Square, London WC1N 3BG, UK. ✉email: lguasch@ic.ac.uk



**Fig. 1 Experimental geometry.** **a** Three-dimensional array of transducers used for data generation and subsequent inversion. Each transducer acts as both a source and a receiver. The red ellipse shows the location of the two-dimensional array used to generate the data for Fig. 2 and 4. **b** A snapshot in time of the wavefield generated by a source transducer located at the position indicated by the small yellow circle, computed via numerical solution of the 3D acoustic wave equation. The wavefield is dominated by strong reflections from the skull, and by intracranial transmitted energy travelling across the brain; Supplementary Video 1 shows the full wavefield propagating in time. **c** Prototype helmet containing 1024 transducers held rigidly in a 3D-printed framework. In the prototype device, the framework is customised to provide an accurate fit to an individual subject, and filled with water. In portable clinical devices, the sensors move radially, and contact the patient via sonographic gel. **d** Close up of sensor connections in the prototype.

pulse. It is not currently possible to correct these effects with sufficient accuracy using conventional modalities.

3. In pulse-echo sonography, back-scattered reflections are used to generate the image. The bones of the skull differ significantly in wave speed and density from those of surrounding soft tissue. Consequently, the skull generates strong reflections and multiple scattering, and these high-amplitude signals overlie, interfere with and obscure the much-weaker reflections produced by the small impedance contrasts that occur within soft tissue in the brain, leading to low signal and high source-generated noise in intracranial pulse-echo images.
4. Time-of-flight tomography uses a short-wavelength approximation, basing its analysis on the simplified physics of ray theory in which the effects of transmission through a heterogeneous medium are represented by a simple change in travel time. For a finite wavelength, wave transmitted through a medium that is heterogeneous on many scales, such delay times are only sensitive to the properties of the medium averaged over the dimensions of the first Fresnel zone<sup>11</sup>. Consequently, time-of-flight tomography is unable to resolve structure below this scale, and so lacks acceptable resolution at the low frequencies that can be recorded using transcranial ultrasound.


One possible way around these problems is to use natural openings in the skull as acoustic windows, but this approach severely reduces illumination<sup>7</sup>; it is typically limited to neonates through an open fontanelle<sup>12</sup>. In principle, it is also possible to remove, or thin, portions of the skull in order to record conventional high-frequency pulse-echo B-mode reflection images without strong bone reflections, absorption, or distortion. This method has produced promising results in rodents, generating functional ultrasound images that can capture transient changes in blood volume related to brain activity<sup>13,14</sup>, but this invasive approach has obvious limitations in clinical practice.

In this paper, we present the neurological application of full-waveform inversion (FWI)<sup>4</sup>, an imaging method first applied widely in geophysics<sup>15</sup>. FWI is a computationally intensive technique that has been developed to a high level of sophistication by the petroleum industry to image hydrocarbon reservoirs within the Earth<sup>16,17</sup>. The spatial resolution that can be obtained using this technique is much finer than that of time-of-flight tomography. FWI achieves this improved resolution through a

combination of characteristics<sup>18</sup>, of which the most important is that it uses a more-complete description of the physics of wave propagation in heterogeneous media that takes proper account of the finite wavelength of transmitted waves. This description, which involves the full numerical solution of the wave equation, is able to model accurately the effects of sub-Fresnel zone heterogeneity and multiple scattering on the wavefield. FWI combines this more-accurate description of the physics with an appropriate non-linear inversion scheme, and a suitable data acquisition geometry, so that it is able to recover fine-scale heterogeneity throughout the model. Adaptive waveform inversion (AWI)<sup>3,19</sup> is a modification of FWI that is better able to begin from a poor starting model; here we use it as a preconditioner for FWI so that inversion can begin successfully without any a priori information about the skull.

Figure 1 outlines the geometry of the method. Low-frequency ultrasound data are recorded at all available azimuths by surrounding the head with ultrasound transducers in three dimensions. Each transducer is activated separately in turn as a source of ultrasound energy, and the signals that it generates are recorded by every other transducer. Allowing for reverberations to die away after each activation, it takes ~2 seconds in total to acquire a full data set for all 1024 sources each recorded on 1024 receivers. FWI uses predominantly transcranial transmitted energy recorded on the side of the head opposite to the source transducer, but it also extracts information from all other parts of the recorded wavefield including reflections, diffractions, multiple scattering, and guided waves that arrive at any angle at any of the transducers. Unlike conventional ultrasound imaging, FWI does not use focused transducers, focusing arrays or any type of beam forming, either in the experimental configuration or in the computer subsequently.

The paper is organised as follows: we explore our proposed methodology using *in silico* simulations, and present *in vivo* and *ex vivo* laboratory results that support our assumptions. We begin by demonstrating the improvement in resolution provided by FWI in even the simplest case when the model is two-dimensional and the skull has been removed. We follow this by exploring what FWI is able to achieve in ideal circumstances for the intact adult human head in three dimensions; this result demonstrates the resolution and tissue contrast potentially achievable in a clinical setting. We follow this by demonstrating that a combination of AWI and FWI can begin from the simplest of starting models, and we demonstrate the importance of full 3D data acquisition and



**Fig. 2** Inversion of data from a brain outside the skull. **a** A two-dimensional model of acoustic wave speed in the naked brain without the skull. The red ellipse shows the transducer positions; the grey region is masked and held fixed during the inversions. **b** Homogeneous model used to begin inversion. **c** Result of ultrasound computer tomography. The resultant model is accurate but has poor spatial resolution. **d** Result of ultrasound full-waveform inversion. The resultant model is now both accurate and spatially well resolved.

inversion. We present laboratory results, using direct observations of in vivo and ex vivo transcranial ultrasound, to demonstrate that good signal penetration and high signal-to-noise levels are readily achievable. We provide an example of the clinical relevance of our approach by demonstrating the accurate recovery of an intracranial haemorrhage, and discuss clinical applications to stroke and other pathologies. We conclude with an outline of our methodology and algorithms. In Supplementary Information, we discuss how our idealised in silico simulations and conclusions are likely to compare with those obtainable from real-world in vivo observations involving anelastic absorption and noise; we review spatial resolution in pulse-echo, time-of-flight tomography, and FWI methodologies; and we explore some of the practicalities and speed of implementation in likely clinical settings.

## RESULTS

### Resolution with the skull removed

Ray-based time-of-flight tomography and wave equation-based FWI both represent forms of transmission tomography. Figure 2 demonstrates the difference between these two techniques using a simple two-dimensional model of the naked brain without the complicating effects of the skull. Using the model from Fig. 2a, and solving a numerical wave equation, a synthetic data set was generated for transducers located around the brain. Using the homogeneous starting model shown in Fig. 2b, this data set was inverted using both time-of-flight tomography and FWI, to recover the models shown in Fig. 2c, d. Time-of-flight tomography seeks to find the best-fitting model by using geometric ray theory to predict delay times for every source-receiver pair in the data set, whereas FWI seeks to solve the same problem by using the wave equation to predict the detailed variation of acoustic pressure with time recorded at every receiver for every source.

For this numerical experiment, the wavelength in soft tissue of the dominant frequency in the insonifying pulse was  $\sim 3.75$  mm, the shortest wavelength was  $\sim 1.75$  mm, and the minimum diameter of the Fresnel zone for signals that travelled across the model was  $> 20$  mm. For conventional pulse-echo B-mode sonography, spatial resolution depends upon several factors<sup>20</sup>, but when imaging through relatively homogeneous soft tissues with an appropriate transducer configuration and a short pulse length of two cycles, the typical resolution is of the order of the insonifying wavelength at the dominant frequency, or  $\sim 3.75$  mm in the present context. For wave equation-based transmission tomographic methods, such as FWI, the achievable resolution is of the order of half the wavelength of the highest available frequency<sup>16,21</sup>, or  $\sim 0.88$  mm in the present context. For FWI,

neither the pulse duration nor beam width directly affect spatial resolution.

In contrast, well-established theory<sup>11</sup> and numerical experiments<sup>22</sup> show that the maximum spatial resolution that can be achieved, in the far field using ray-based time-of-flight tomography, is of the order of the diameter of the first Fresnel zone. Thus, we would expect that an FWI model would be  $\sim 20$  times better resolved in linear dimensions than the equivalent time-of-flight model. Figure 2 illustrates this behaviour directly. Both methods recover models that are accurate in their locally averaged properties, but the time-of-flight model has only centimetre-scale spatial resolution, whereas the FWI model has millimetre resolution. Note that, in this simple example, the difference in resolution between the two techniques is not related to the presence of the skull, nor to differences in the optimisation scheme—both methods used non-linear least-squares inversion applied to the same input data. A more-detailed discussion of the resolution achievable by pulse-echo sonography, time-of-flight tomography, and FWI is provided within the Supplementary Information.

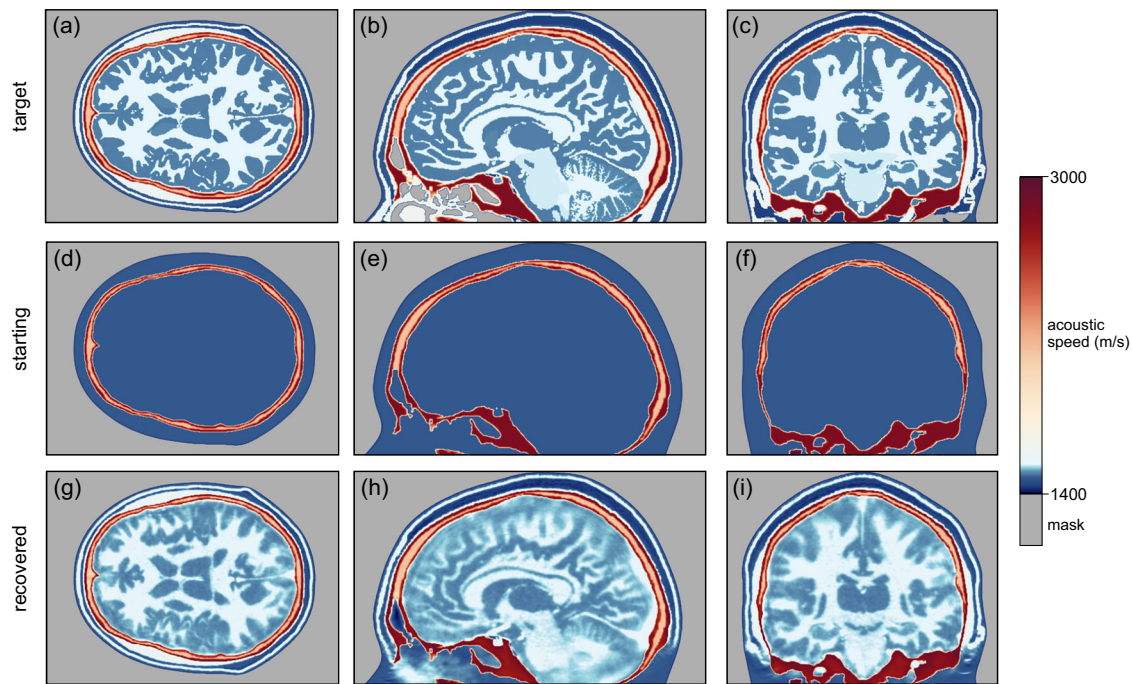
In the absence of the skull, conventional high-frequency pulse-echo sonography would of course be able to recover an accurate and well-resolved image of the naked brain. However, if the skull interposes, then pulse-echo fails entirely to image the brain inside the skull because brain reflections are then significantly distorted, absorbed, and scattered by the skull. Similarly, time-of-flight tomography for the brain inside the intact human skull fails because, at the low frequencies that can be transmitted across the head with acceptable signal-to-noise ratios, spatial resolution is insufficient. FWI does not suffer from either of these problems; it properly accounts for the distorting effects of the skull, and it achieves good spatial resolution even at the low frequencies that can be recorded after transmission through the skull.

### Three-dimensional full-waveform imaging through the skull

FWI has obvious advantages for brain imaging; it does though have two complications of its own: the computational effort required to extract the image from the data in three dimensions is significant, and the method requires a reasonably good starting model in order to proceed to the correct final model.

The former requirement has, until recently, limited the applicability of medical FWI to problems that can be usefully solved in two dimensions<sup>23</sup>, and the skull is not even approximately two-dimensional. The advent of large parallel multi-core multi-node compute clusters, of on-demand parallel cloud computing, and of large-memory graphics processing unit (GPU) systems, coupled with improved FWI software and the use of pre-





**Fig. 3 Models of acoustic wave speed.** Transverse (left), sagittal (centre), and coronal (right) sections through the true (top), starting (middle), and recovered (bottom) models. Both the wavefield modelling and waveform inversion are performed in three dimensions. The starting model includes the true model of the skull, but is otherwise homogeneous.

trained supervised deep-learning to accelerate the process, are reducing the computational demands of this method; runtimes and costs continue to reduce year-on-year.

The requirement for a good starting model is straightforward for the soft tissue of the brain where a homogeneous starting model is adequate. For the bones of the skull, however, additional care is required. In this section, we assume that the skull model is known a priori. In the following section, we demonstrate that AWI is capable of building an accurate model of the skull entirely from the ultrasound data without a priori knowledge, and that following this with FWI then recovers a fully resolved model of the brain.

Figure 3 shows transverse, sagittal, and coronal sections through a three-dimensional target model of wave speed, a starting model containing the true skull but otherwise homogeneous, and the model reconstructed using FWI applied to sub-MHz ultrasound data generated by the target model. Supplementary Videos 2, 3, and 4 show the true, starting, and reconstructed models in three dimensions. The colour scale shown in Fig. 3 is designed to highlight heterogeneity within both soft and hard tissues.


The bones of most of the upper cranium are multi-layered, containing the inner and outer tables of denser cortical bone with a high wave speed, surrounding the diploë, which is formed of cancellous bone with lower density and wave speed. This structure, together with the large contrast in properties between the skull and its surrounding soft tissues, provides the principal mechanism for transcranial signal attenuation, with anelastic absorption and elastic mode-conversions playing a less-significant role<sup>24–26</sup>; in the Supplementary Information, we examine explicitly the significance of absorption and demonstrate its limited impact. The model of the skull used in this study included all cavities, foramina, and other structural complications that are present in the adult human head, and that are capable of being captured on the 500  $\mu\text{m}$  grid that we used to represent the model.

The model recovered by FWI, Fig. 3g–i, is in good agreement with the true model, Fig. 3a–c, for both extracranial and intracranial soft tissues. Inside the skull, FWI is able to generate

an accurate and detailed image: grey and white matter match the target tissue properties accurately, both in absolute wave speed and in structure, with sufficient resolution to allow direct identification of cortical folds. Deeper structures such as the corpus callosum, the thalamus, the basal ganglia, and the ventricular system are recovered well. Parts of the venous sinuses have a thickness of 0.8 mm in the true model, as do larger vessels within the brain, and these are recovered in the reconstructed image, demonstrating that we are able to achieve sub-millimetre resolution of the brain and its vascular system using only relatively low frequencies lying below 1 MHz. Parts of the cerebellum and the pons lie inferior to the lowest transducer position in our numerical experiment, but it is still possible to extract sufficient information from the data to image both bodies, although there is a decrease in resolution as illumination is progressively lost in the area close to the base of the skull.

Recovering the brain without an a priori model of the skull

FWI is a local optimisation algorithm that requires an initial model that lies within the basin of attraction of the global solution<sup>15</sup>. The variation in soft-tissue acoustic wave speed is  $\sim \pm 7\%$ , which has values between  $\sim 1400 \text{ ms}^{-1}$  for fat and  $1600 \text{ ms}^{-1}$  for muscle tissue and cartilage<sup>27</sup>. At the frequencies that we use for FWI, such relatively small perturbations are readily retrievable starting from a homogeneous model having a wave speed similar to that of water at  $\sim 1500 \text{ ms}^{-1}$ , as demonstrated in Fig. 2. This is the reason why FWI applied, for example, to breast imaging has been immediately successful<sup>23,28</sup>. In contrast, the variation in wave speed for hard tissue in the cranium is larger at  $\sim \pm 14\%$ , with values between  $\sim 2100 \text{ ms}^{-1}$  for cancellous bone and  $2800 \text{ ms}^{-1}$  for cortical bone<sup>24–26</sup>; the mandible and the vertebrae have even higher wave speeds of  $\sim 3500 \text{ ms}^{-1}$ <sup>27</sup>. These high values are far removed from that of water. Consequently, recovery of the full model of the head, including the bones of the skull, requires a more-sophisticated approach; we show in this section that this can be provided by the AWI algorithm<sup>3</sup>.



**Fig. 4 Inverting from a homogeneous starting model.** **a** Homogeneous starting model with velocity of water. **b** Model recovered using conventional FWI. **c** Model recovered using AWI. **d** Model recovered by FWI following AWI. The colour scale is as shown in Fig. 3. Note that the skull is not present in the starting model, and the final model is well recovered.

Figure 4b shows the failure of an attempt to recover a model of the head using conventional FWI beginning from the purely homogeneous starting model shown in Fig. 4a. This should be compared with Fig. 3g, which shows the analogous result obtained when the starting model contains an accurate model of the skull. The reason for the failure of this attempt is that conventional FWI will not converge to the correct global solution if inversion begins from a model that is too far removed from the true model. Specifically, for successful FWI, the data generated by the starting model must not be shifted in time by more than half a wave cycle at the lowest frequencies that are present in the data. When that condition is not met, the data are “cycle skipped”, and FWI will then typically fail.

AWI is a modification of FWI that is less sensitive to the quality of the starting model, and it is able to move towards the true model even when the data are cycle skipped. It does, however, pay a price for this robustness, and the models that it is able to produce on its own are not normally as well resolved as those that can be produced by FWI. The solution then is to begin with AWI from a simple homogeneous cycle-skipped model, then switch to conventional FWI once AWI has moved the model sufficiently far towards the true model that they are no longer cycle skipped<sup>19</sup>.

Figure 4c shows the result of applying AWI, using the same data and starting from the same homogeneous model as was used to generate Fig. 4b. Now the attempt to recover a model of both the skull and the brain has been reasonably successful, and no a priori model of the skull has been assumed. The AWI model though is not the final result. Figure 4d shows the results subsequently obtained by conventional FWI beginning from a smoothed version of the model previously recovered by AWI. The final model recovered by this combination of methods is now accurate and compares well to the model in Fig. 3g that was recovered using a perfect model of the skull. AWI and FWI together then can fully solve the problem of building a well-resolved accurate model of skull and brain purely from ultrasound data without any a priori knowledge of the skull.

#### Importance of three dimensions

Most three-dimensional medical imaging analyses data initially in two dimensions in order to produce a stack of planes that are combined to form a final 3D image volume. There would be advantages in applying this approach to ultrasound FWI: the computational cost of inverting many 2D slices is lower than that of true 3D inversion, and 2D acquisition systems are simpler to design, build, and operate. However, the structural complexities of the skull, and large contrast with soft tissue, act to distort the wavefronts by refracting and scattering energy out of a 2D plane.

Figure 5a illustrates the detrimental effects of inverting three-dimensional data in only two dimensions. Here, the data being inverted are a dense two-dimensional subset of the full three-

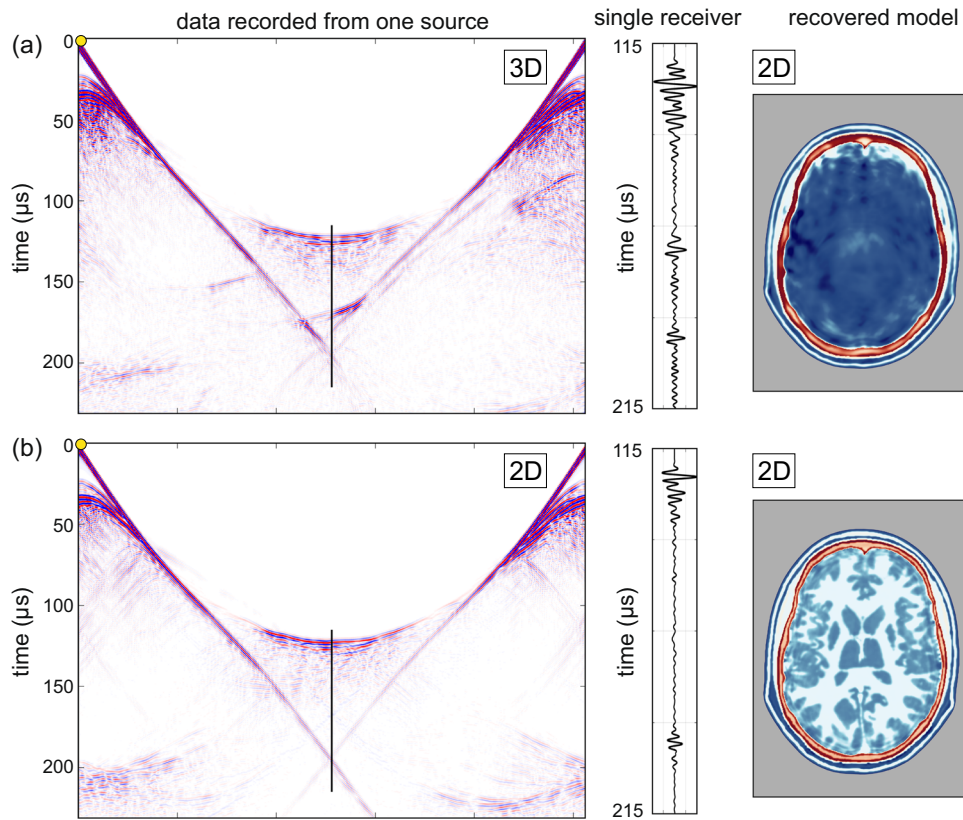
dimensional data used to generate the results shown in Fig. 3. In Fig. 5a, the data to be inverted have been generated by a 3D wave equation applied to a 3D model, but the inversion assumes only a 2D model and uses a 2D wave equation. The inversion is therefore unable to explain energy that has been refracted, reflected, scattered, or guided out of the 2D plane. The model recovered in this case is neither accurate nor useful.

Figure 5b shows the equivalent experiment conducted purely in 2D; in this second case, both the initial data generation and the inversion are two-dimensional. The 2D inversion of 2D data recovers a model that is as accurate as that recovered by 3D inversion of 3D data. Comparing the data and waveforms in Fig. 5a, b demonstrates why 3D FWI of 3D data, and 2D FWI of 2D data both succeed, whereas 2D FWI of 3D data fails entirely. The 2D and 3D data sets show major differences, and it is evident that there must be significant out-of-plane energy present in the 3D data, and this cannot be explained adequately during 2D FWI. As three-dimensional effects will always be present in real data, successful imaging of the brain using transmission FWI will always require 3D data acquisition and 3D inversion in order to correct properly for the three-dimensional distortion of the wavefield produced by the bones of the skull.


#### In vivo laboratory observations

At the sub-MHz frequencies that were used in our numerical experiments, transmission losses in soft tissue are small<sup>27,29–34</sup>, but scattering and anelastic losses in the skull can be important<sup>24–26</sup>. To test the significance of these losses, and to measure signal-to-noise ratios in real transcranial ultrasound, we made in vivo and ex vivo laboratory observations. In both these experiments, the transceiver bandwidth and the source waveform were identical to those used in the in silico simulations. Ultrasound intensities were at all times below the lowest limits recommended by the British Medical Ultrasound Society for continuous adult diagnostic ultrasound<sup>35</sup>. In the in silico and ex vivo experiments, the subject head and the transducers were immersed in water. In contrast, for the in vivo observations, the transducers were held against the scalp, coupled to the subject using sonographic gel. Both in vivo and ex vivo experiments produced similar levels of signal and similar signal-to-noise ratios.

Figure 6 shows transcranial in vivo waveforms acquired in the laboratory, at three orientations, using one of the authors as a subject. The data are displayed un-stacked and un-processed; that is, a single unfocused source was triggered once, and the raw data recorded on a single-channel receiver are displayed without beam forming, dynamic compression or other numerical manipulation. These observations demonstrate unequivocally that transcranial ultrasound can be recorded for the in vivo adult human head, with good signal-to-noise ratios, for the 100–850 kHz bandwidth that we have used through this study using the simplest of acquisition systems. These unfocussed signals have travelled across the head,



**Fig. 5 Model recovery using two-dimensional FWI. a** 3D data inverted using 2D FWI. **b** 2D data inverted using 2D FWI. Left panels show simulated data generated by a single source located at the yellow circle, as recorded on an elliptical array of 512 transducers placed around the head. Centre panels show data recorded by a single receiver located opposite the source; the position of the data shown is indicated by the blue line. Right panels show models recovered using purely two-dimensional FWI. Colour scale as shown in Fig. 3.



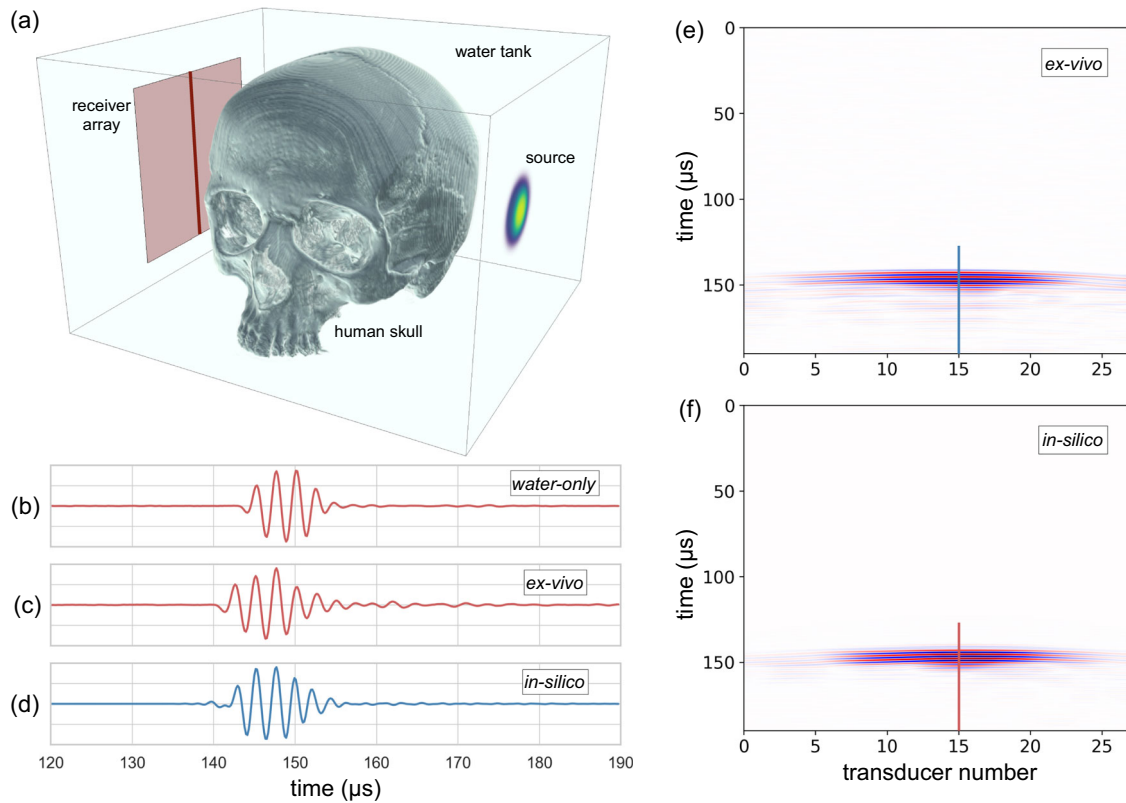
**Fig. 6 In vivo transcranial observations.** Un-processed, un-stacked, transcranial, ultrasound signals, low-pass filtered at 1 MHz, generated using a single source and single receiver, coupled to the subject using sonographic gel, and held firmly against the scalp and hair. The source waveform is identical to that used in both in silico and ex vivo experiments. The sonograms are amplitude normalised; true amplitudes and travel times are approximately proportional to source-receiver separation. **a** Temporal to temporal. **b** Left frontal to right occipital. **c** Frontal to parietal.

successively through gel, hair, skin, external soft tissue, cranium, brain, and other internal soft tissue, cranium, external tissues, and gel, in a variety of orientations and alignments.

We calculated observed signal-to-noise levels by comparing the RMS amplitude recorded before the transmitted signal arrives with that recorded within a 20- $\mu$ s window following the first signal

arrival. Signal-to-noise ratios calculated in this way for the in vivo data, averaged over all orientations, were 39 dB, with a maximum of 10 dB variation with orientation. We demonstrate in the Supplementary Information that FWI remains robust at signal-to-noise ratios below 2 dB. Our direct in vivo observations thus demonstrate observed signal-to-noise levels that are very much





**Fig. 7** Ex vivo and in silico data after transmission across the head. **a** The geometry of the ex vivo laboratory experiment. **b** Data recorded in the laboratory by the central transducer without the skull present. **c** The equivalent ex vivo data recorded with the skull. **d** The equivalent in silico data with the skull. **e** Laboratory ex vivo data recorded on a finely sampled linear array with the skull. **f** The equivalent numerical data simulated in 3D. The physical skull and the numerical model are nominally the same, but differ in detail, and the numerical model does not include the effects of scattering by the physical transducers and their supporting hardware. Figures **b–d** are normalised to their largest amplitude; the un-normalised amplitudes in **b** are about five times larger than those in **c**.

larger than the minimum required for successful FWI, and were obtained using safe levels of incident ultrasound, and employed transducers and electronics that are readily deployable in a clinical setting. It is clear that the signal penetration and signal-to-noise ratios that are likely to be achievable in practical applications at sub-MHz frequencies will be more than sufficient for transcranial neuroimaging using low-frequency ultrasound supported by FWI.

#### Ex vivo laboratory observations

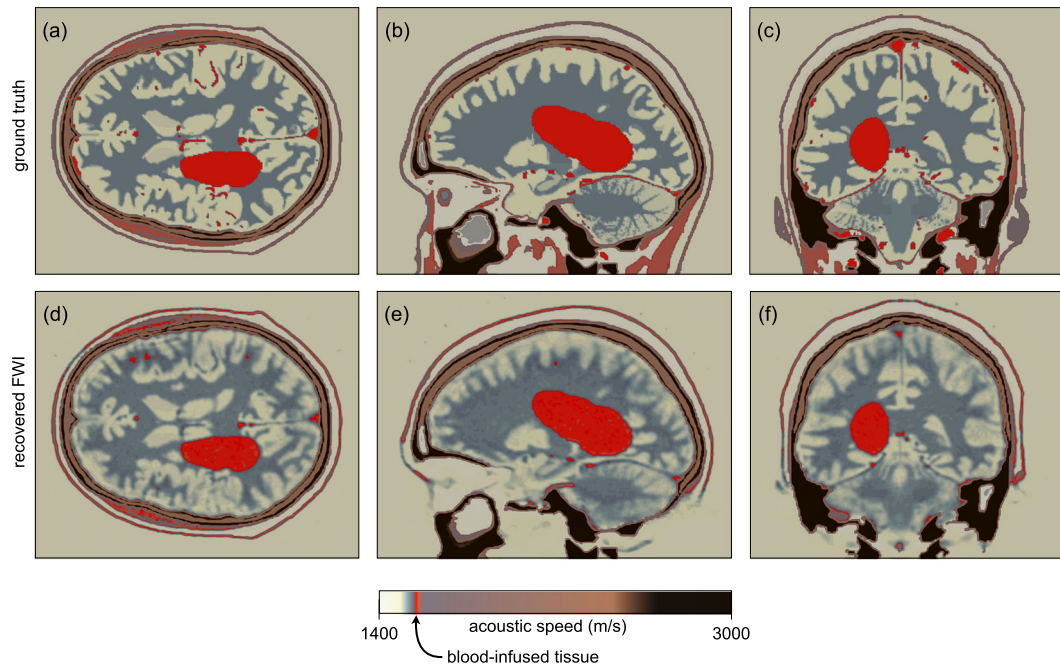
In the ex vivo experiment, we immersed a human skull in water, recording transcranial ultrasound using the same transducers as were used for the in vivo observations. The data recorded, with and without the skull present, are displayed in Fig. 7. The principal signal losses within an in vivo human head occur as the result of anelastic loss within bone, reflection at the inner and outer boundaries of the skull, and scattering within the cranium especially at the boundary between the diploë and cortical bone. The principal wavefront distortions are produced by the large sound speed contrast between the bones of the skull and their surrounding soft tissues, which have a sound speed close to that of water. This ex vivo experiment is designed to capture all those features.

Figure 7b shows the data recorded by a single transducer without the skull present, and Fig. 7c shows data from the same transducer with the skull interposed between the source and receiver. Amplitudes are normalised in these figures so that Fig. 7c is displayed at about five times the gain of Fig. 7b. The noise level in the transcranial ex vivo data set observed in Fig. 7c is low; the signal-to-noise ratio was 34 dB when measured as for the in vivo data. The principal effect of interposing the skull is to advance the

arrival time by  $\sim 2.5 \mu\text{s}$ , which corresponds to one cycle at the 400-kHz dominant frequency, and to reduce the transmitted amplitude by about a factor of five. Assuming a skull thickness of 7 mm, physical properties for wave speed, density, and absorption given in Supplementary Table 1, and a cranium in which lower-velocity lower density diploë is sandwiched between faster, denser cortical bone, the observed delay time is as expected. This delay means that the observed data are cycle skipped with respect to a homogeneous model that contains only water, and it explains why simple FWI fails to converge properly when beginning from such a model, as in Fig. 4a.

Using the same simple model of the cranium to calculate anticipated transmission amplitude losses produced by the skull, the principal effect is that caused by normal-incidence reflections at the water-to-cortical bone interfaces and at the cortical-to-diploë interfaces. Together, these eight reflections reduce transmitted amplitudes to  $\sim 36\%$  of their values in a homogeneous water model. Amplitude loss, produced by anelastic absorption within both the cortical and cancellous layers, reduces amplitudes at 400 kHz to about a further 77%, so that the total transmitted amplitude would be expected to drop to  $\sim 28\%$  as a result of transmission through the skull. In the experimental data, the corresponding amplitude drop is  $\sim 22\%$ , well within the uncertainties of this simple model. That the observed losses are slightly greater than the calculated values is most likely a consequence of small-scale structural complexity within the real skull, which will moderately increase the magnitude of internal scattering.

We repeated the same experiment as shown in Fig. 7c, e in silico, using a model obtained by converting a high-intensity X-ray CT image volume of the ex vivo skull into acoustic sound speed.



**Fig. 8 Recovery of a large haemorrhage by FWI.** **a–c** Slices through a 3D wave-speed model perturbed by a large haemorrhage. **d–f** The same slices through the FWI. The haemorrhage is well recovered by full-waveform inversion at high resolution in all slices. The colour scale has been modified to highlight the haemorrhage.

The conversion from X-ray attenuation to sound speed is not exact, so that the *in silico* data would not be expected to provide an exact match to the laboratory *ex vivo* data. The real experiment will also contain signals scattered by the physical transducers and their supporting infrastructure; we did not attempt to duplicate these additional signals *in silico*. Figure 7 compares the two datasets. It shows that the timing, waveform, absolute amplitude and variation of amplitude with position and time, in the *ex vivo* laboratory data, are well reproduced by the *in silico* simulation, verifying that our skull model and our modelling assumptions are both reasonable.

#### Clinical application to stroke

The application of FWI to neuroimaging has the potential to improve diagnosis in a wide range of neurological pathologies; here, we explore its potential to aid early treatment of stroke, a major cause of death and adult disability worldwide<sup>1</sup>. Stroke has two principal causes: ischaemic stroke is most commonly caused by a blood clot obstructing blood supply to the brain, and haemorrhagic stroke is most commonly caused by bleeding within the brain parenchyma. When the blood supply to the brain is compromised, rapid intervention is required to restore circulatory integrity, halt and reverse tissue damage, and prevent and reduce morbidity, mortality, and disability. Although there are a number of early treatments available, including thrombolysis, mechanical thrombus extraction, and similar interventions<sup>36,37</sup>, their applicability is limited in practice by the requirement for accurate high-resolution brain imaging before these treatments can be deployed<sup>2</sup>.

The indicated treatment for ischaemic stroke is contraindicated for haemorrhagic stroke; brain imaging is therefore required to diagnose and separate these causes. The need for speed is paramount, but MRI is not portable and X-ray CT is barely so. Brain imaging then takes place not when paramedics first reach the patient, not within the ambulance, and not often within accident and emergency units; as a result, relatively few stroke patients receive a brain scan of any kind within the critical first hour, and even fewer receive high-quality MRI<sup>2</sup>. There is then a clear need


for portable, fast, high-resolution, high-fidelity, 3D, brain imaging that can differentiate between ischaemic and haemorrhagic stroke, and differentiate these from other pathologies that can mimic stroke. The development and clinical application of such a method would markedly increase the survival rate and reduce the severity of subsequent disability by enabling much earlier treatment at the point of first patient contact.

To test the viability of this concept, we modified the target model to contain a haemorrhage, Fig. 8a–c. To build this model, we used physical properties for blood-infused soft tissue<sup>38,39</sup>. Using a homogeneous starting model for the brain, and the true model for the skull, we recovered the FWI images in Fig. 8d–f. Figure 9 shows that the haemorrhage can be readily segmented from the three-dimensional model, both in the target and the FWI-recovered models. The target pathology is well recovered in the FWI image, and has good tissue contrast with other features of the brain. The boundaries and exact extent of the pathology are clear; and although it is not shown here, FWI is well able to produce time-lapsed images over a wide variety of time scales from seconds to hours. The spatial resolution of FWI is capable of detecting haemorrhage at all scales down to that of the originating vasculature.

#### DISCUSSION

Both X-ray CT and MRI revolutionised medical imaging when they first appeared; three-dimensional transmission and reflection ultrasound tomography using FWI has the same potential for impact across multiple disciplines, and has especial relevance for rapid diagnosis and treatment of stroke. Supplementary Fig. 1 shows that the method is robust against the levels of noise that we observe in realistic *ex vivo* and *in vivo* laboratory experiments, and transcranial ultrasound signals have large amplitudes at the relatively low, sub-MHz, frequencies that are sufficient for successful sub-millimetre resolution. The method overcomes the well-known limitations of conventional pulse-echo ultrasound imaging of the adult human brain, and the related limitations of conventional time-of-flight tomography.





**Fig. 9 Segmented haemorrhage.** **a** Haemorrhage auto-segmented from the true model. **b** Haemorrhage auto-segmented from the model recovered by full-waveform inversion.

It is necessary either to include some prior model of the skull within the starting model before attempting to recover the brain, or to use an advanced form of FWI such as AWI that can converge toward the correct answer from a simple starting model. To account correctly for and remove the distorting effects of the skull, it is essential to acquire and invert the data in three dimensions, and unlike many other imaging techniques, it is not possible to reduce three-dimensional imaging merely to the sum total of a sequence of two-dimensional slices. In Supplementary Fig. 2, we show that, at the low frequencies required by FWI, it is not necessary to include density or anelastic absorption explicitly in the inversion in order to recover a good image, but it may be desirable to do so, both to improve the accuracy of the final image, and to obtain additional independent parameters to aid diagnosis. Other ultrasound methodologies, employing higher frequencies, are more severely impacted by signal attenuation, both in the skull and in soft tissue.

The computational effort required for 3D FWI is considerable. The results shown in Fig. 3 required ~32 hours elapsed time to complete, running on a conventional cluster of 128, CPU-based, 24-core, compute nodes. Our target is to reduce this elapsed time to below 10 minutes; this requires a speed up of ~200 times. The hardware that we used has a peak performance of ~60 tera-flop, so achieving the desired speed up requires hardware capable of operating usefully at a peak of ~12 peta-flop. Individual high-performance GPU-based servers are currently able to achieve speeds in excess of one peta-flop, so that a small array of these would in principle be capable of producing a final model in < 10 minutes. Assuming 2019 prices, amortisation over 3 years, and full utilisation of the hardware, the capital cost of the GPU-server hardware required to do this represents a few tens of dollars to invert a 3D transcranial dataset on a 500- $\mu$ m grid. So, although the computational burden of FWI is high, the cost per patient to achieve a 10-minute turnaround is not high under appropriate circumstances.

The potential value of FWI imaging is threefold. Most importantly, it could improve outcomes in acute neurological disorders such as stroke and head trauma by enabling earlier intervention; the ultimate aim is diagnosis and treatment within minutes of first contact with paramedics. Second, the low cost, high safety, portability, and high resolving power of the technology provides the ability to monitor the brains of patients continuously at the bedside allowing clinicians to intervene, for a range of pathologies, to prevent injury with the speed that the brain demands, acting in rapid response as if the brain image was a simple physiological variable such as blood pressure. And third, the technology can be deployed readily and safely, for prevention and diagnosis, in a wide range of situations where neuroimaging would be desirable but is currently unavailable—for example,

within developing nations with limited health budgets, in remote locations, routinely at contact-sports events, within military deployments, or as part of disaster relief when local infrastructure is compromised.

## METHODS

### In silico model

We used the MIDA 3D numerical model of the human head<sup>40</sup>, at the original sample spacing of 500  $\mu$ m, as the basis to build the sound speed model used in the in silico simulations. Physical properties within the model were derived from the geometry of the segmented model combined with values for acoustic sound speed, density, and absorption for different tissue types from<sup>24–27,29–34</sup>; further details appear in Supplementary Table 1. Most minor tissue types within the model have unmeasured acoustic properties; in these cases, we estimated their values using small perturbations to the properties of other tissue types that appeared analogous in their other physical properties and composition.

The models used to generate the data for all figures except Supplementary Fig. 2 were purely acoustic. The model used for Supplementary Fig. 2 included anelastic absorption, and assumed a linear relationship between attenuation and signal frequency. At the relatively low frequencies used in these simulations, such a model of absorption provides a reasonable approximation to the properties of real tissue<sup>29</sup>.

### In silico modelling

The experimental geometry was restricted to accommodate the application of this technology realistically to human patients in a clinical setting, and therefore no transducers were positioned in front of the face or below the base of the head. Transducers were modelled assuming unfocused single elements. We used 1024 transducers that acted as both sources and receivers, generating just over a million source-receiver records of acoustic pressure, each record lasting 240  $\mu$ s. The source waveform consisted of a three-cycle tone burst having a peak amplitude at a frequency of 400 kHz and a useful bandwidth for FWI extending from ~100–850 kHz. This source waveform is identical to that generated in our laboratory experiments.

The synthetic data were generated by solving the three-dimensional, variable density, isotropic, linear, acoustic wave equation, explicitly in the time domain, using a time-stepping finite-difference algorithm, with an optimised stencil that is nominally tenth-order in space and fourth-order in time. For the two-dimensional data shown in Fig. 4b, the analogous two-dimensional wave equation was solved in which the model, wavefield, sources, and receivers do not vary perpendicular to a two-dimensional plane. For the anelastic data used to generate Supplementary Fig. 2, the visco-acoustic wave equation was solved using the method described in ref. 41.

### Laboratory experiments

The ex vivo laboratory experiment was performed by immersing a formalin-preserved human skull in water, generating an ultrasound pulse on one side of the head, and recording the resultant signals on the other. The skull retained some residual soft tissue, was stored dry, and was

typically immersed for a few tens of minutes during ultrasound measurements. In the ex vivo experiment, sources, and receivers were not in direct contact with the skull; in the in vivo experiment, sources, and receivers were in direct contact with the head, coupled using sonographic gel.

For the in vivo observations, we used single-element flat unfocused Blatek-Industries 400-kHz AT30871 10-mm piezoelectric transducers as source and receiver. For the ex vivo experiments, we used a single-element unfocused Olympus 500 kHz V381 19-mm ultrasound source transducer. In both in vivo and ex vivo experiments, the transducers were driven to generate a three-cycle tone burst centred on 400 kHz at the intensities commonly employed in conventional medical imaging<sup>7</sup>; this waveform matched that used in the in silico experiments. For the ex vivo experiment, we recorded the transmitted signals from the source using an Olympus 500 kHz V301 25-mm transducer located on the opposite side of the skull at a perpendicular distance of 210 mm from the source. A planar receiver array was formed by moving the single receiving transducer successively in 4-mm steps to 729 positions to form a  $27 \times 27$  array measuring  $108 \times 108$  mm. For both the in vivo and ex vivo experiments, the data were generated and recorded using a Verasonics Vantage 256, and the recorded data were low-pass filtered at 1 MHz.

The in vivo experiment employed one of the authors as a participant, for whom informed consent was obtained prior to the experiment. The experiments complied with all relevant ethical regulations; the study was approved by the Joint Research Compliance Office and the Imperial College Research Ethics Committee, both of Imperial College London.

### FWI iteration

During FWI, we inverted the in silico data in the time domain over eight finite-frequency bands, starting at a dominant frequency of  $\sim 100$  kHz, and moving successively to higher frequencies to reach the maximum useable frequency in the data of  $\sim 850$  kHz. This multi-scale approach helps to ensure that the inversion does not become trapped at some local solution produced by inadequacies in the starting model. Within each frequency band, we used ten iterations in total, using a tenth of the data at each iteration, so that data from each source was used just once per frequency band. The highest frequencies present in the data have a half-wavelength in the brain of less than a millimetre, so that we would expect to be able to resolve sub-millimetre structure in the final recovered model.

During AWI, we began the inversion from a homogeneous water model using a single frequency band centred on 100 kHz. Other than the change of misfit function, AWI proceeds in a similar fashion to FWI. Following AWI, we smoothed the resultant model, and used that to begin FWI from 320 kHz, moving successively to the maximum frequency in the data as before.

### FWI algorithm

FWI is an imaging method that seeks to find a model that can numerically reproduce experimental data. It does this by solving a non-linear least-squares local optimisation problem, modifying the model in order to minimise the misfit,  $f$ , defined as (half) the sum of the squares of the differences between the experimental data and an equivalent simulated dataset that is numerically generated using a numerical model of acoustic properties. Thus, we seek to minimise:

$$f = \frac{1}{2}(\mathbf{p} - \mathbf{d})^T(\mathbf{p} - \mathbf{d}) \quad (1)$$

where  $\mathbf{p}$  and  $\mathbf{d}$  represent the predicted (numerically generated) and observed (experimental) data, respectively, organised as vectors that contain concatenated time-series of pressure variations at each recording location for every source location. In this in silico study, the experimental data were themselves generated using a known ground-truth target model. At the initiation of FWI, a starting model  $\mathbf{m}$ , representing an initial estimate of the target of interest and composed of many model parameters  $m_i$ , is used to solve the wave equation and generate the predicted data set  $\mathbf{p}$ . In the account below, the quantities  $f$ ,  $\mathbf{p}$ ,  $\mathbf{A}$ , and  $\mathbf{u}$  each depend upon the assumed model  $\mathbf{m}$ , whereas  $\mathbf{d}$ ,  $\mathbf{s}$ , and  $\mathbf{R}$  do not.

The computational cost of numerically solving the governing wave equation in 3D for many sources restricts computationally tractable solutions to iterated local gradient-descent methods. We solve the problem by seeking the direction of steepest descent, on the hyper-surface defined by  $f$ , which has as many dimensions as there are model parameters  $m_i$ . For the 3D model shown in Fig. 3, this hyper-surface had  $\sim 10^9$  dimensions. This gradient-descent algorithm seeks to move from the starting model, by a sequence of small

steps, successively downhill on this hyper-surface, to arrive close to the model that lies at the lowest point on the surface—this is the model that best predicts the observed data in a least-squares sense.

In order to find the direction of steepest descent, the derivative of the misfit with respect to the model parameters is found using the adjoint-state method<sup>15</sup>. The derivative of  $f$  with respect to each of the model parameters  $m_i$  takes the form:

$$\frac{\partial f}{\partial m_i} = \left[ \frac{\partial \mathbf{p}}{\partial m_i} \right]^T (\mathbf{p} - \mathbf{d}) \quad (2)$$

To compute the first derivative of the predicted data with respect to the model parameters  $m_i$ , we start by writing the wave equation as a matrix-vector operation:

$$\mathbf{A}\mathbf{u} = \mathbf{s} \quad (3)$$

where  $\mathbf{A}$  is the wave equation written in a suitable discrete form—here we used high-order finite differences to approximate the 3D anisotropic variable density visco-acoustic wave equation,  $\mathbf{u}$  is the pressure wavefield and  $\mathbf{s}$  is the source. Differentiating this with respect to  $m_i$ , and taking into account that both  $\mathbf{A}$  and  $\mathbf{u}$  depend upon the model parameter  $m_i$ , but that the source  $\mathbf{s}$  does not, gives:

$$\mathbf{A} \frac{\partial \mathbf{u}}{\partial m_i} + \frac{\partial \mathbf{A}}{\partial m_i} \mathbf{u} = \mathbf{0} \quad (4)$$

Assuming that  $\mathbf{A}$  is invertible, which it must be if Eq. 3 has a unique solution, leads to the expression:

$$\frac{\partial \mathbf{u}}{\partial m_i} = -\mathbf{A}^{-1} \frac{\partial \mathbf{A}}{\partial m_i} \mathbf{u} \quad (5)$$

for the variation of the wavefield  $\mathbf{u}$  with the model parameter  $m_i$ . Now, the predicted data  $\mathbf{p}$  are simply a subset of the full wavefield obtained at those locations where we happen to have placed receivers. Thus, we can use a restriction matrix  $\mathbf{R}$  to extract the corresponding data as  $\mathbf{p} = \mathbf{R}\mathbf{u}$ , so that:

$$\frac{\partial \mathbf{p}}{\partial m_i} = -\mathbf{R}\mathbf{A}^{-1} \frac{\partial \mathbf{A}}{\partial m_i} \mathbf{u} \quad (6)$$

Because, again,  $\mathbf{R}$  does not depend on  $m_i$ . The final expression for the gradient is then:

$$\frac{\partial f}{\partial m_i} = -\mathbf{u}^T \frac{\partial \mathbf{A}^T}{\partial m_i} \mathbf{A}^{-T} \mathbf{R}^T (\mathbf{p} - \mathbf{d}) \quad (7)$$

Reading from right to left, this expression implies the following sequence of steps to calculate the gradient:

1. Compute the residual data  $(\mathbf{p} - \mathbf{d})$  by solving the wave equation to generate  $\mathbf{p}$  in the starting model,
2. Inject the residual data into the model at the receiver positions,
3. Solve the wave equation backwards in time using the injected residual data as a virtual source,
4. Scale the resulting wavefield using the differential of the wave equation operator  $\mathbf{A}$  with respect to the model parameters  $m_i$ ,
5. Find the zero lag of the cross-correlation of this wavefield in time with the original forward wavefield  $\mathbf{u}$  at every point in the model.

The result of applying these steps is a gradient vector oriented to point in the direction of maximum increase of  $f$  at the current position in the solution space. The negative of the gradient indicates the direction in which small changes to the model will create the largest decrease in the misfit  $f$ , so the model should be changed in this direction. Typically, the problem will be non-linear and non-convex, and therefore it requires that the model is updated iteratively, and that the starting point is within the basin of attraction of the global minimum. In a practical FWI algorithm, the direction of steepest descent is typically preconditioned in some way to speed convergence; here, we used spatial preconditioning to compensate for illumination variation within the model<sup>16</sup>. A more-complete development, and further details are given in ref. <sup>15,16,18</sup>.

### AWI algorithm

AWI<sup>3</sup> is a form of FWI that has immunity to cycle skipping. In conventional FWI, the algorithm seeks to drive the sample-by-sample difference between the predicted and observed data to zero. In contrast, in AWI the algorithm seeks to drive the ratio between the two data sets to unity. Both approaches aim to drive the predicted data set towards the observed data set, and consequently to drive the recovered model towards the true model. With perfect data, and

perfect algorithms, both approaches will reach the same end point. However, when both methods are implemented using local gradient descent, they follow different paths through the space of possible models in their attempt to reach the true model. In these circumstances, FWI will tend to become trapped in local minima when the predicted data differ from the observed data by more than half a wave cycle, whereas AWI will not.

The ratio used in AWI is not generated sample-by-sample; rather it is a ratio formed frequency-by-frequency after temporal Fourier transform, and the ratio is formed separately for each source-receiver pair. As division in the frequency domain represents deconvolution in the time domain, the AWI algorithm effectively deconvolves one data set by the other and then attempts to drive the result of that deconvolution toward a unit-amplitude delta function at zero temporal lag. The mathematical details are given in ref. <sup>3</sup>, and practical details are given in ref. <sup>19</sup>.

### Algorithm for time-of-flight tomography

The model shown in Fig. 2c was generated using time-of-flight tomography. This method is analogous to FWI, but with two significant differences: the data to be inverted consist of single numbers, one for each source-receiver pair, representing the time taken for energy to travel through the target model from source to receiver; and geometric ray-tracing rather than the full wave equation is used to calculate these travel times. This approach has the advantage that it is computationally more tractable than FWI, and so runs orders of magnitude more quickly; it is also much less likely than FWI to become trapped in local minima. But, it has the disadvantage that it cannot recover structure in the model below the scale of the diameter of the first Fresnel zone.

We solved the tomographic problem by minimising a misfit similar to that shown in Eq. (1), but where **p** and **d** are now vectors containing the travel times rather than the raw observed acoustic pressure data. We address the problem as before by solving a non-linear least-squares local optimisation problem, using gradient descent preconditioned by conjugate gradients. Unlike the more familiar X-ray CT, in ultrasound tomography changes to the model affect the path that energy follows from source to receiver, and it is necessary to include this non-linear effect into the inversion by iterating.

### Reporting summary

Further information on experimental design is available in the Nature Research Reporting Summary linked to this article.

### DATA AVAILABILITY

The true model was generated using the MIDA model described in ref. <sup>40</sup>. The model can be obtained in digital form from the authors of that paper. We populated that model using the physical properties shown in Supplementary Table 1. The synthetic acoustic dataset that we generated from this model is available from the Dryad digital archive (<https://doi.org/10.5061/dryad.nzs7h44n7>).

### CODE AVAILABILITY

The computer code used to generate synthetic data and to perform the inversions is Fullwave3D, described in ref. <sup>16</sup>. This is a large and sophisticated commercial package developed for and licensed to the petroleum industry. The software is available on commercial terms, and can be licensed both for use on private HPC systems and on the commercial public cloud. For those wishing simply to reproduce the results shown here, a more-practical option is likely to be to use one of the several open-source modelling and FWI packages that are now becoming available, or for those with access to commercial geophysical processing software to use one of the growing number of such packages that include facilities for performing FWI. Several of the latter packages are available to university-based researchers, for non-commercial use, on heavily discounted terms.

Received: 16 August 2019; Accepted: 5 February 2020;

Published online: 06 March 2020

### REFERENCES

- Donnan, G. A., Fisher, M., Macleod, M. & Davis, S. M. Stroke. *Lancet* **371**, 1612–1623 (2008).
- Embersson, J. et al. Effect of treatment delay, age, and stroke severity on the effects of intravenous thrombolysis. *Lancet* **385**, 1929–1935 (2014).
- Warner, M. & Guasch, L. Adaptive waveform inversion: theory. *Geophysics* **81**, R42–R45 (2016).
- Tarantola, A. Inversion of seismic reflection data in the acoustic approximation. *Geophysics* **49**, 1259–1266 (1984).
- von Ramm, O. T., Smith, S. W. & Kisslo, J. A. *Ultrasound tomography of the adult brain. Ultrasound in Medicine*. (eds. White, D. & Lyons, E. A.) 261–267 (Springer, Boston, MA, 1978).
- Ylitalo, J., Koivukangas, J. & Oksman, J. Ultrasonic reflection mode computed tomography through a skull bone. *IEEE Trans. Biomed. Eng.* **37**, 1059–1066 (1990).
- Smith, S. W. et al. The ultrasound brain helmet: feasibility study of multiple simultaneous 3D scans of cerebral vasculature. *Ultrasound Med. Biol.* **35**, 329–338 (2009).
- Niesen, W.-D., Rosenkranz, M. & Weiller, C. Bedside transcranial sonographic monitoring for expansion and progression of subdural hematoma compared to computed tomography. *Front. Neurol.* **9**, 374 (2019).
- Hoskins, P. R., Martin, K. & Thrush, A. (eds) in *Diagnostic Ultrasound*, 3rd ed. Physics and Equipment (CRC Press, Boca Raton, Florida, 2019).
- Li, C., Duric, N., Littrup, P. & Huang, L. In-vivo breast sound-speed imaging with ultrasound tomography. *Ultrasound Med. Biol.* **35**, 1615–1628 (2009).
- Williamson, P. R. A guide to the limits of resolution imposed by scattering in ray tomography. *Geophysics* **56**, 202–207 (1991).
- Demene, C. et al. Functional ultrasound imaging of brain activity in human newborns. *Sci. Transl. Med.* **9**, eaah6756 (2017).
- Macé, E. et al. Functional ultrasound imaging of the brain. *Nat. Methods* **8**, 662–664 (2011).
- Yang, G., Pan, F., Parkhurst, C. N., Grutzendler, J. & Gan, W.-B. Thinned-skull cranial window technique for long-term imaging of the cortex in live mice. *Nat. Protoc.* **5**, 201–208 (2010).
- Virieux, J. & Operto, S. An overview of full-waveform inversion in exploration geophysics. *Geophysics* **74**, WCC1–WCC26 (2009).
- Warner, M. et al. Anisotropic 3D full-waveform inversion. *Geophysics* **78**, R59–R80 (2013).
- Routh, P. et al. Impact of high-resolution FWI in the Western Black Sea: revealing overburden and reservoir complexity. *Lead. Edge* **36**, 60–66 (2017).
- Pratt, R. G. Seismic waveform inversion in the frequency domain, Part 1: theory and verification in a physical scale model. *Geophysics* **64**, 888–901 (1999).
- Guasch, L., Warner, M. & Ravaut, C. Adaptive waveform inversion: practice. *Geophysics* **84**, R447–R461 (2019).
- Ng, A. & Swanevelde, J. Resolution in ultrasound imaging. *Contin. Educ. Anaesth. Crit. Care Pain* **11**, 186–192 (2011).
- Dessa, J. X. & Pascal, G. Combined traveltimes and frequency-domain seismic waveform inversion: a case study on multi-offset ultrasonic data. *Geophys. J. Int.* **154**, 117–133 (2003).
- Williamson, P. R. & Worthington, M. H. Resolution limits in ray tomography due to wave behavior – numerical experiments. *Geophysics* **58**, 727–735 (1993).
- Pratt, R. G., Huang, L., Duric, N. & Littrup, P. Sound-speed and attenuation imaging of breast tissue using waveform tomography of transmission ultrasound data. *Proc. SPIE* **6510**, 65104S (2007).
- Fry, F. J. & Barger, J. E. Acoustical properties of the human skull. *J. Acoust. Soc. Am.* **63**, 1576–1590 (1978).
- Pichardo, S., Sin, V. W. & Hynynen, K. Multi-frequency characterization of the speed of sound and attenuation coefficient for longitudinal transmission of freshly excised human skulls. *Phys. Med. Biol.* **56**, 219–250 (2011).
- Pinton, G. et al. Attenuation, scattering, and absorption of ultrasound in the skull bone. *Med. Phys.* **39**, 299–307 (2012).
- Duck, F. A. *Physical Properties of Tissue: A Comprehensive Reference Book* (Academic Press, London, 1990).
- Agudo, O. C., Guasch, L., Huthwaite, P. & Warner, M. 3D imaging of the breast using full-waveform inversion. Proceedings of the International Workshop on Medical Ultrasound Tomography, Speyer, 1–3 Nov. 2017 (KIT Scientific Publishing, Karlsruhe, Germany, 2018).
- Chivers, R. C. & Parry, R. J. Ultrasonic velocity and attenuation in mammalian tissues. *J. Acoust. Soc. Am.* **63**, 940–953 (1978).
- Goss, S. A., Johnston, R. L. & Dunn, F. Comprehensive compilation of empirical ultrasonic properties of mammalian tissues. *J. Acoust. Soc. Am.* **64**, 423–457 (1978).
- Goss, S. A., Frizzell, L. A. & Dunn, F. Ultrasonic absorption and attenuation in mammalian tissues. *Ultrasound Med. Biol.* **5**, 181–186 (1979).
- Goss, S. A., Johnston, R. L. & Dunn, F. Compilation of empirical ultrasonic properties of Mammalian tissues II. *J. Acoust. Soc. Am.* **68**, 93–108 (1980).
- Yongchen, S., Yanwu, D., Jie, T. & Zhensheng, T. Ultrasonic propagation parameters in human tissues. Proceedings of the 1986 IEEE Ultrasonics Symposium, Williamsburg, VA, November 17–19, 1986.
- Culjat, M. O., Goldenberg, D., Tewari, P. & Singh, R. S. A review of tissue substitutes for ultrasound imaging. *Ultrasound Med. Biol.* **36**, 861–873 (2010).



35. The British Medical Ultrasound Society. Guidelines for the safe use of diagnostic ultrasound equipment. (The British Medical Ultrasound Society, London, 2009).
36. Stead, L. G., Gilmore, R. M., Bellolio, M. F., Rabinstein, A. A. & Decker, W. W. Percutaneous clot removal devices in acute ischemic stroke: A systematic review and meta-analysis. *Arch. Neurol.* **65**, 1024–1030 (2008).
37. Joyce, S. et al. A systematic review and meta-analysis of randomized controlled trials of endovascular thrombectomy compared with best medical treatment for acute ischemic stroke. *Int. J. Stroke* **10**, 1168–1178 (2015).
38. Mourad, P. D. & Kargl, S. G. Acoustic Properties of Fluid-Saturated Blood Clots (Technical Report APL-UW TR 2003, Applied Physics Laboratory, University of Washington, 2000).
39. Huang, C.-C., Lin, Y.-H., Liu, T.-Y., Lee, P.-Y. & Wang, S.-H. Review: Study of the blood coagulation by ultrasound. *J. Med. Biol. Eng.* **31**, 79–86 (2011).
40. Iacono, M. I. et al. MIDA: A multimodal imaging-based detailed anatomical model of the human head and neck. *PLOS ONE* **10**, e0124126 (2015).
41. da Silva, N. V., Yao, G. & Warner, M. Semiglobal viscoacoustic full-waveform inversion. *Geophysics* **84**, 11MA–11MA (2019).

## ACKNOWLEDGEMENTS

We gratefully acknowledge assistance in the laboratory from Javier Cudeiro, Thomas Robins, and Carlos Cueto, and from Javier Cudeiro for designing the helmet shown in Fig. 1. Adrian Umpleby was the lead programmer in the project that developed the Fullwave3D package used in this study. This work was funded by the Excellence Fund for Frontier Research, a fund created by Imperial College London to initiate and support research in frontier areas that cross traditional disciplinary boundaries. P.N. is funded by the Wellcome Trust (213038/Z/18/Z) and the UCLH NIHR Biomedical Research Centre.

## AUTHOR CONTRIBUTIONS

All authors designed and performed the research; L.G. and M.W. wrote the main text; all authors agreed the final version of the paper.

## COMPETING INTERESTS

O.C.A. is the author of pending UK patent 1621436, which is relevant to parts of this work. The authors declare no other competing interests.

## ADDITIONAL INFORMATION

**Supplementary information** is available for this paper at <https://doi.org/10.1038/s41746-020-0240-8>.

**Correspondence** and requests for materials should be addressed to L.G.

**Reprints and permission information** is available at <http://www.nature.com/reprints>

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2020