From Biology to Bytes: Predicting the Path of Ultrasound Waves Through the Human Body

Bradley E. Treeby

Address:
Department of Medical Physics and Biomedical Engineering
University College London
Gower Street
London WC1E 6BT
United Kingdom

Email:
b.treeby@ucl.ac.uk

Jiri Jaros

Address:
Centre of Excellence IT4Innovations
Faculty of Information Technology
Brno University of Technology
Bozetechova 2
612 66 Brno
Czech Republic

Email:
jarosjir@fit.vutbr.cz

Eleanor Martin

Email:
elly.martin@ucl.ac.uk

Ben T. Cox

Email:
b.cox@ucl.ac.uk

Address:
Department of Medical Physics and Biomedical Engineering
University College London
Gower Street
London WC1E 6BT
United Kingdom

Computer simulations are increasingly used to guide ultrasound therapies, but what makes a good model and when can we trust them?

Introduction

The use of ultrasound as a diagnostic imaging tool is well-known, particularly during pregnancy where ultrasound is used to create pictures of developing babies. In recent years, a growing number of *therapeutic* applications of ultrasound have also been demonstrated. The goal of therapeutic ultrasound is to modify the function or structure of biological tissue in some way rather than produce an anatomical image. This is possible because the mechanical vibrations caused by ultrasound waves can affect tissue in different ways, for example, by causing the tissue to heat up or by generating internal forces that can agitate the cells or tissue scaffolding. These ultrasound bioeffects offer enormous potential to develop new ways to treat major diseases. In the last few years, clinical trials of different ultrasound therapies have demonstrated the ability of ultrasound to destroy cells through rapid heating for the treatment of cancer and neurological disorders, target the delivery of anticancer drugs, stimulate or modulate the excitability of neurons, and temporarily open the blood-brain barrier to allow drugs to be delivered more effectively (Konofagou, 2017). These treatments are all completely noninvasive and have the potential to significantly improve patient outcomes.

The fundamental challenge shared by all applications of therapeutic ultrasound is that the ultrasound energy must be delivered accurately, safely, and noninvasively to the target region within the body identified by the doctor. This is difficult because bones and other tissue interfaces can severely distort the shape of the ultrasound beam (see **Figure 1**, ***bottom***, for an example). This distortion can have a significant impact on the safety and effectiveness of therapeutic ultrasound and is one of the major hurdles for the wider clinical acceptance of this exciting technology. In principle, it is possible to predict and correct for these distortions using models of how ultrasound waves travel through the body. However, the underlying physics is complex and typically must consider nonlinear wave propagation through absorbing media with spatially varying material properties. Simple formulas do not exist for this scenario, so models used for studying therapeutic ultrasound are instead based on the numerical solution of the wave equation (or the corresponding constitutive equations). This article is primarily concerned with the development of such models.

Circle of Model Development

One way to consider the development of numerical ultrasound models, and indeed any scientific software that models a physical phenomenon, is described by the circle of *model development* (see **Figure 2**). This has five distinct components:

https://doi.org/10.1121/AT.2019.15.2.36

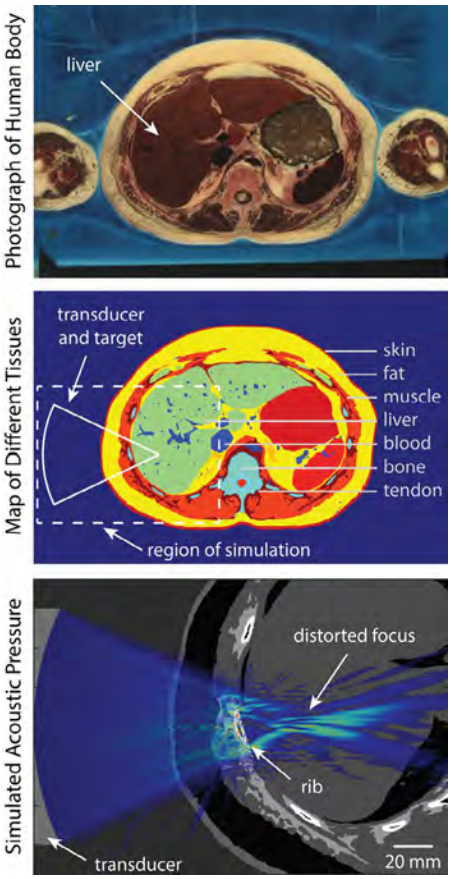

Figure 1. *Predicted distortion of the acoustic field from a therapeutic ultrasound transducer for a liver target when the beam path is occluded by the ribs (**bottom**). The spatially varying acoustic properties for the simulation (**middle**) are derived from cryosection images from the Visible Human Project run by the US National Library of Medicine (**top**). The predicted ultrasound field is calculated using the open-source k-Wave Toolbox (Jaros et al., 2016).*

(1) deriving equations that describe the underlying physics;
(2) choosing a suitable algorithm (numerical method) to solve these equations;
(3) implementing the numerical scheme as computer code;
(4) defining inputs to the model, for example, the spatial discretization and material properties; and
(5) validating the model (numerically and experimentally).

The process is often iterative (**Figure 2**, ***dashed line***) and is repeated until the model predictions agree with the experimental observations (the level of agreement required depends on the context). For ultrasound therapy, it is critical that each component in the circle of model development is carefully considered before models are used to assist in calculations and predictions because any deficiencies could have serious ramifications for patient safety.

The circle of model development can be used as a general framework to guide developers. Equally, it can be employed by end users to ask, *Is this model appropriate for my application?* For example, considering each of the five components, questions arise.

(1) Do the model equations capture enough of the physics to simulate the phenomenon of interest?
(2) Is the chosen numerical method stable, efficient, and accurate? What are the convergence properties?
(3) Is the computer code fast, fault tolerant, portable, scalable, well-documented, and easy-to-use? Is it regulatory compliant?
(4) What are the model inputs? Are they material properties, physical constants, or tuning parameters, and are they easy to measure or specify? What is the sensitivity of the model to errors in the inputs? What settings do I need to use to ensure accurate results?
(5) Has the model been validated against analytical solutions, experimental measurements, or clinical data? Do I have confidence in using the model to make clinical predictions?

In this article, each component of the circle of model development is discussed in more detail as a general framework for the development of software for modeling therapeutic ultrasound in the human body. Like many endeavors in modern science, adequately addressing all five components requires a team of people with a wide breadth of expertise. This includes physical acoustics, numerical methods, computer science, software engineering, ultrasound metrology, and medical imaging. In many cases, input from regulatory experts, clinicians, and other end users will also be needed.

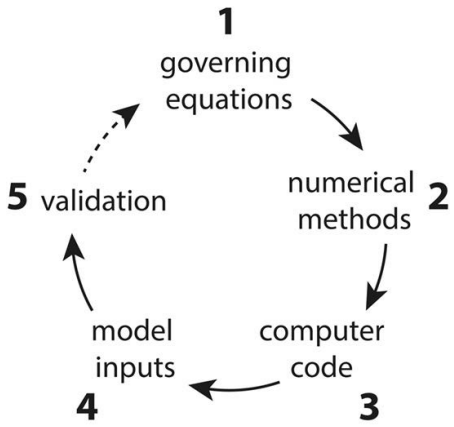

Figure 2. *The circle of model development showing the five key components that should be considered when developing scientific software. See text for explanation.*

Governing Equations

The mathematical expressions that describe the physics of ultrasound wave propagation in tissue are known as *governing equations*. These are typically based on the conservation of momentum (which accounts for the tissue having inertia), the conservation of mass (which accounts for the tissue being compressible), and an equation of state or pressure-density relationship (which encapsulates the thermodynamics of wave propagation). In many branches of acoustics, these equations can be simplified by assuming linear wave propagation in a lossless and homogeneous medium, which leads the widely known *wave equation*. In biomedical ultrasound, this equation is appropriate for studying the spatial distribution of acoustic pressure from therapeutic ultrasound transducers in water at low output levels. However, for modeling therapeutic ultrasound in the human body, the simplifying assumptions mentioned above generally no longer hold.

First, the pressure amplitudes used in biomedical ultrasound are often sufficiently high to give rise to nonlinear effects (this is true for both therapeutic and diagnostic applications of ultrasound). At high acoustic pressures, the stiffness of the tissue depends on how much it is being compressed (material nonlinearity) and the cyclical motion of the medium, due to the acoustic wave, affects the wave speed (convective nonlinearity). Together, these cause the sound speed in the medium to depend dynamically on the local values of the acoustic pressure and particle velocity. For example, during the compressional phase of the wave where the particle velocity is positive (i.e., the medium is being displaced in the same direction as the wave is traveling), the effective sound speed increases and vice versa. This causes a cumulative distortion in the time-domain waveform, which corresponds to the generation of higher frequency harmonics in the frequency domain. From a modeling standpoint, this type of nonlinearity can be captured by retaining second-order terms in the governing equations (Hamilton and Blackstock, 2008).

Second, biological tissue can strongly attenuate ultrasound waves, particularly waves at megahertz frequencies. The exact mechanisms for the absorption in tissue are complex and occur at both the cellular level (e.g., viscous relative motion and thermal conduction between the cells and their surroundings) and the molecular level (e.g., molecular and chemical relaxation). These processes cause the gradual degradation of acoustic energy into random thermal motion and, consequently, the attenuation of the wave amplitude. In addition to absorption, acoustic energy is also lost due to scattering. This is generally negligible in soft tissue at low megahertz frequencies but can become significant as the wavelength decreases or in highly scattering media such as bone. Overall, the acoustic attenuation in soft biological tissue (which includes both absorption and scattering) has been experimentally observed to follow a frequency power law of the form $\alpha_0 f^y$, where f is the frequency and the power law exponent y is between 1 and 2. This type of behavior can be captured in the governing equations by including a distribution of relaxation processes or by using fractional derivative loss operators (Holm and Nasholm, 2014). A commonly used rule of thumb is that ultrasound in soft biological tissue is attenuated at a rate of 1 dB/MHz/cm.

Third, in biological tissues, medium properties such as the sound speed, mass density, and acoustic absorption coefficient are heterogeneous across multiple scales. At the microscopic level (much smaller than an acoustic wavelength), there are variations in the acoustic properties between individual cells and between cells and other tissue constituents such as blood plasma and the extracellular matrix (the structural scaffolding that holds many cells in place). These differences give rise to diffusive scattering, which is responsible for the speckle pattern characteristic of diagnostic ultrasound images. However, from an ultrasound therapy perspective, this scattering is generally weak and can be accounted for as part of a phenomenological attenuation term in the governing equations.

At a macroscopic level, different structures within an organ, such as blood vessels or regions of fatty and fibrous tissue, can also give rise to scattering. There are also differences at the organ level, again due to variations in the underlying tissue constituents and their structure. For example, tissues with a higher proportion of lipids (e.g., fat) typically have a lower sound speed compared with water at body temperature, whereas tissues with a higher proportion of proteins (e.g., liver) have a higher sound speed. These macroscopic variations in the acoustic properties of tissue can have a significant impact on the propagation of focused ultrasound fields, including changing the shape, position, and amplitude of the focal region (see **Figure 1**, ***bottom***). In some cases, the aberrations are so strong that the focus is completely destroyed. Spatially varying material properties can be included in the governing equations by starting with the conservation equations and retaining the spatial gradients of the material parameters during the linearization process.

The combination of the mass and momentum conservation equations (retaining heterogeneous and nonlinear terms)

and an equation of state (accounting for absorption through a fractional loss operator or sum of relaxation terms) can account for the complex wave behavior seen in biological tissue. This includes scattering, refraction, nonlinear wave steepening, and acoustic absorption following a frequency power law. However, in some cases, additional factors must also be considered, for example, the temperature dependence of the tissue material properties, the motion of organs due to breathing or the cardiac cycle, and/or acoustic cavitation (Maxwell et al., 2012; Gray et al., 2019). Systematically incorporating such extensions into the governing equations in a tissue-realistic manner is not straightforward. For modeling scenarios that involve bones, the generation, propagation, and absorption of shear waves must also be considered. In this case, the mass conservation equation and the equation of state are replaced with a model of viscoelasticity (a generalization of Hooke's law).

Numerical Methods

The techniques used to discretize the governing equations so that they can be solved by a computer are known as *numerical methods*. There are many different types of numerical methods used in acoustics. These include the finite-element method, boundary-element method, finite-difference method, Green's function methods, and spectral methods (Verweij et al., 2014). The most appropriate choice depends on the problem specifics, for example, the distribution of material properties (e.g., homogeneous, piecewise constant, or continuously varying), whether the problem is linear or nonlinear, whether the source is single frequency or broadband, and the scale of the domain of interest. For therapeutic ultrasound (which usually involves nonlinear wave propagation in heterogeneous and absorbing biological tissue), the most common approach is to use computationally efficient collocation methods, such as the finite-difference time domain (FDTD) or pseudospectral time domain (PSTD) methods (Gu and Jing, 2015). These methods can be used to directly solve the governing equations as a system of coupled equations, or the equations can be combined and solved as a generalized wave equation. The former has some advantages for numerically imposing radiation conditions at the edge of the computational domain (such as a perfectly matched layer) and for inputs and outputs that depend on the acoustic particle velocity (including modeling dipole sources and calculating the vector acoustic intensity).

A significant challenge when modeling biomedical ultrasound is the large distances traveled by the ultrasound waves relative to the acoustic wavelength at the highest frequency of interest. Consider the case of transcranial focused ultrasound surgery, where ultrasound waves are used to destroy a small region of tissue deep in the brain. The domain of interest encompassing the ultrasound transducer and the head is on the order of 30 cm in each direction. For a center frequency of 650 kHz, this distance is on the order of 130 wavelengths at the fundamental frequency and 650 wavelengths at the fifth harmonic. Applying the engineering rule of thumb of 20 points per wavelength (PPW) sometimes used for finite-element and finite-difference methods, this corresponds to a computational grid size of $13{,}000 \times 13{,}000 \times 13{,}000$ grid points (more than 2 trillion degrees of freedom). Simply storing one matrix of this size in single-precision floating-point format would consume 8 terabytes (TB) of memory, and typically several such matrices are needed. To put this into context, the current generation MacBook Air comes equipped with 8 gigabytes (GB) of memory, so 1,000 of them would be required to store one matrix at a cost of more than one million US dollars! Of course, supercomputing is not done using consumer laptops, but the point remains: problems of this nature can become extremely large scale.

So why is the engineering rule of thumb to use 20 PPW when the Nyquist theorem tells us we should only need two? The primary reason is *numerical dispersion*. This is a numerical error in which approximations made in the numerical method cause the modeled ultrasound waves to travel at different speeds depending on their frequency (the dependence of sound speed on frequency is known as dispersion). This dependence means that broadband waves will become increasingly distorted compared with the true solution as they propagate across the computational grid (equivalently, single frequency waves will travel at the wrong speed). This is a particular challenge for the large domain sizes encountered in therapeutic ultrasound (often hundreds of wavelengths), because errors due to numerical dispersion accumulate the further the waves travel.

For finite-difference methods, provided that the numerical scheme mathematically reduces to the governing equations in the limit that the spatial and temporal steps reduce to zero (a condition known as consistency) and that the method is stable (there are standard mathematical and numerical tools for analyzing stability), the *Lax equivalence theorem* tells us that the scheme will be convergent. This means that the numerical solution will approach the exact solution of the governing equations as the size of the spatial and temporal

steps is reduced. In general, this provides a very practical way to test the accuracy of almost any numerical model. Reduce the size of the grid spacing (Δx) and the time step (Δt) or otherwise increase the mesh density, and check to see if the answer remains the same. If not, keep reducing Δx and Δt until the answer no longer changes. This procedure is called a *convergence test* and should always be carried out for every modeling scenario. In general, the output from a numerical model should not be trusted unless convergence has been demonstrated!

Finite-difference methods have been widely used for modeling in acoustics; however, these methods often require very large computational grids to avoid numerical dispersion. To reduce dispersion errors, higher order finite-difference schemes can be implemented that use more neighboring grid points to estimate the spatial and temporal gradients. *Spectral methods* take this idea to the limit and use all of the grid points simultaneously by fitting a finite sum of basis functions to the data. In acoustics, a common choice is to use trigonometric functions, where the fitting is performed by taking a fast Fourier transform (FFT). This is the idea behind the PSTD and *k*-space methods that calculate spatial gradients in the spatial-frequency domain. Although computationally more expensive than the FDTD method for a fixed grid size, these methods can significantly reduce dispersion errors and thus the number of points per wavelength required for accurate solutions (Tabei et al., 2002).

A remaining challenge for collocation methods computed on regular Cartesian grids is the introduction of medium staircasing. This arises because the material properties must be represented at discrete points in the model (think of the intersection of lines on a sheet of graph paper), and in many cases, the material boundaries are not aligned with the grid. This leads to stair-like edges between regions with different material properties that generate spurious acoustic reflections. For the PSTD method, this can be the dominant source of error (Robertson et al., 2017a). Although these errors will reduce with increasing grid density, in some cases, it can be challenging to perform a convergence test because the properties are only known at a fixed resolution (e.g., from a medical image), often of the same order as the acoustic wavelength.

Computer Code

Once a numerical method has been developed, this must be turned into *computer code* that can be used to perform simulations. Typically, the high-level goals are to (1) implement the numerical algorithm correctly, (2) maximize performance (e.g., reduce run time), and (3) minimize the computational resources needed (e.g., memory). Unsurprisingly, the development of efficient high-performance computer code is closely connected to a deep understanding of the underlying computer hardware. This is particularly relevant for models of therapeutic ultrasound where the grid sizes are often extremely large and complex calculations such as the FFT are performed (Jaros et al., 2016).

Computational hardware has undergone rapid changes since the first appearance of microprocessors in the late 1960s. Huge increases in performance have been enabled by continual improvements in semiconductor lithography leading to a doubling in the number of transistors on a computer chip approximately every 18 months. During the twentieth century, performance increases were also obtained through increases in transistor switching frequency. These days, however, performance increases are instead driven by increases in parallelization across all levels of processing along with the development of specialized compute units such as graphics processing units (GPUs). This means a modern supercomputing cluster can be highly heterogeneous, consisting of multiple interconnected computers, each potentially containing multiple central processing units (CPUs) and GPUs, where each CPU and GPU has multiple cores, each of which can execute multiple instructions simultaneously on multiple data points! Similarly, there is hierarchy of local and remote memory with different storage capacities and access speeds. Although these details may not be familiar to many acousticians, they are nonetheless important. Effectively programming for such heterogeneous architectures is highly nontrivial and can have a large impact on the performance and tractability of running therapeutic ultrasound simulations (Jaros et al., 2016).

For heterogeneous computer environments, there are two fundamental requirements to consider: data locality and workload balance. Data locality is critical because there is huge difference in the transfer speed (20 times slower) and latency (100 times slower) when accessing data stored on another interconnected computer compared with data stored in local memory (e.g., cache). For the large computational problems encountered in ultrasound, this means the data must be carefully decomposed into different levels of memory so that communication is minimized or overlapped with other useful calculations. Workload balance is critical because different parts of a heterogeneous system can have

significantly different compute power. For example, a GPU can be an order of magnitude faster than a single processor at basic arithmetic but extremely slow at control processes. Naturally, the performance of the code taken as a whole is limited by its slowest component.

There are several other factors that can influence the performance of a computer code. The first is the choice of programming language. Although the general trend in software engineering is towards high-level languages (such as MATLAB, Python, C#, or Java), high-performance computing often still relies on low-level languages (such as C/C++ and Fortran) that allow direct access to the hardware characteristics of the target system. Unfortunately, this also complicates the design, implementation, testing, and portability of the codes. Computing in lower (or mixed) precision can increase the speed of arithmetic operations and reduce memory transfers. However, care must be taken to avoid numerical overflow and round-off errors. For large domain sizes, extracting outputs from the model (for example, saving the time-varying pressure field) can also have a significant impact on performance due to the time taken to write data to disk.

During the design of computer code, many implementation details must also be considered. First, the complexity of scientific software often makes the use of test-driven development crucial. In this paradigm, every component of the software has a corresponding test suite validating its results against a ground truth. Additional tests can also be used to investigate changes in code performance. Second, care must be taken to ensure the code (and the development cycle) is adequately documented. Even smaller software projects can be impossible for other developers to understand without proper documentation. Third, the software development environment must be considered. Generally, modern software projects are managed through software version control systems. These incorporate many useful features, including allowing multiple developers to work simultaneously and providing tools for issue tracking, automated testing, and code release. For software used in a medical context, regulatory standards may also mandate additional documentation, testing, and reporting requirements. Finally, the support of users outside the development team must also be carefully considered, for example, through the publication of a user manual or other articles, the development of worked examples, and a user forum or mailing list.

Model Inputs

Once a computer code has been developed, using it to predict how ultrasound waves travel through the human body requires specification of the *model inputs*. From a physical perspective, this includes the properties of the transducer used to transmit (and in some cases receive) the ultrasound waves and a map of the acoustic material properties and how they vary through the body. From a numerical perspective, this includes the size and distribution of the spatial grid and the number and size of the time steps (or an error threshold if using an adaptive integration scheme). These parameters are normally chosen based on a convergence test as discussed in **Numerical Methods**. Finally, from a computational perspective, this includes which computational hardware is used and how the computational effort is distributed among the available resources. This choice is usually made based on heuristics obtained from previous model runs.

Considering the physical inputs starting with the ultrasound transducer, the most straightforward way to include this in a numerical model is as an ideal radiator that vibrates uniformly across the transducer surface (Martin et al., 2016). However, in many cases, physical transducers do not act as uniform radiators due to the propagation of surface waves. To improve accuracy, the size and shape of the ideal source can be adjusted until a reasonable match is obtained with an experimental measurement. Alternatively, a full characterization of the spatially (and sometimes temporally) varying pattern of source vibration can be performed using holography or direct measurement with a laser Doppler vibrometer. Holography involves measuring the acoustic pressure over a surface that encloses the source or, in practice, over a surface that captures most of the emitted energy (Sapozhnikov et al., 2015). The measured pressure is then numerically projected back to the source surface (or a nearby plane) to obtain a spatially varying pressure or velocity distribution for input into the model. Holography measurements are normally conducted with the transducer driven at a low level to ensure linear wave propagation, with source conditions at higher drive levels extrapolated based on additional measurements made using a radiation force balance.

Regarding the acoustic material properties, for patient-specific simulations, these are typically obtained from standard medical images such as X-ray computed tomography (CT) or magnetic resonance (MR) imaging. When using CT images, it is possible to convert the grayscale values in the image to the mass density of the tissue based on a stoichiometric calibration of the CT

scanner. The sound speed and absorption coefficient can then be inferred from the mass density, albeit with relatively large uncertainties (Mast, 2000). When using MR images, the different organs within the body must first be segmented, and then values for the tissue properties are assigned (these values are usually taken from previous measurements reported in the literature made using ex vivo tissue samples). Perhaps unsurprisingly, uncertainties in the geometry and properties of the body (particularly the sound speed) form a significant source of error in model-based predictions made for ultrasound therapy (Robertson et al., 2017b).

Model Validation

Checking whether a computer program gives the correct answer under different circumstances is known as *model validation*. Errors can come from any of the four preceding steps in the circle of model development, including invalid assumptions made when developing the governing equations, the wrong choice of grid parameters in the numerical model, mistakes in the computer code or rounding and overflow errors, and inaccuracies in the acoustic material properties or source conditions. The validation of numerical models is an important part of the software design life cycle, particularly in the context of therapeutic ultrasound where software may be used to derive treatment parameters or influence clinical decisions.

Model accuracy is generally tested in several stages, including

(1) performing a convergence test;
(2) comparing with analytical solutions, for example, the scattering of a plane wave by a sphere;
(3) quantitatively comparing the predicted acoustic field against well-controlled laboratory experiments;
(4) quantitatively comparing with experiments conducted using ex vivo tissue or animal models; and
(5) comparing against the outcome of clinical treatments in patients, for example, comparing the volume of ablated tissue after a treatment using high-intensity focused ultrasound.

For therapeutic ultrasound, there are a limited number of relevant analytical solutions, and it is often difficult to quantify model accuracy (or the origins of any discrepancies) using clinical data. Consequently, the bulk of model validation is performed using experimental measurements.

When performing experimental validation, there are two main challenges. The first is precisely replicating the experimental setup in the computer simulation, for example, the characteristics of the ultrasound transducer, the acoustic properties of the medium, the geometry and position of any scattering objects, and the spatial locations at which the pressure is measured. Errors in any of these will lead to discrepancies between the model and measurement. One approach is to use simple geometries and standardized materials with well-known properties. For example, phantoms with precisely known geometries can be created using 3-D printing as seen in **Figure 3** (Robertson et al., 2017b). Unfortunately, performing quantitative validation measurements using more realistic biological specimens such as ex vivo skull samples remains a difficult task.

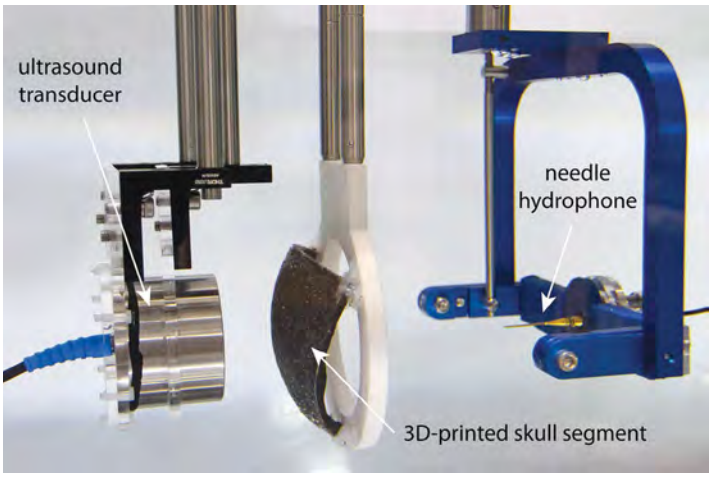

Figure 3. *Measurement of the acoustic field behind a 3-D-printed skull segment* (***center***) *using a focused ultrasound transducer* (***left***) *and a needle hydrophone* (***right***) *in a water tank.*

The second challenge is obtaining accurate, absolute measurements of acoustic pressure. It cannot be assumed that a measurement is the ground truth because there are many factors that give rise to measurement uncertainties. Variations can occur in the transducer output due to fluctuations in the supplied voltage, the electrical impedance, or changes in the water temperature. Errors can also arise from the alignment and positioning of the source and receiver, misalignment of scanning-system axes, and interference from acoustic reflections. Perhaps most importantly, the properties of the hydrophone used can have a significant influence on the measurement. For example, the finite size of the hydrophone detector element can give rise to spatial averaging effects, particularly for the tightly focused fields used in ultrasound therapy. Moreover, for some hydrophones, the frequency-dependent sensitivity is nonuniform in both magnitude and phase, which can result in significant pressure errors if not properly deconvolved

from the measurement data (Wear et al., 2014). Unfortunately, uncertainties in the sensitivity are typically 6-15% depending on frequency, and the angle dependence is not normally known, which ultimately limits the precision of any pressure measurement. These uncertainties should be carefully considered when using experimental data for model validation.

Outlook and Summary

Recent advances in numerical methods and high-performance computing mean that large-scale full-wave simulations of ultrasound propagation in the human body are now within reach. These models have a myriad of uses in ultrasound therapy, including patient selection (determining whether a patient is a good candidate for a particular procedure based on their individual anatomy), treatment verification (determining the cause of adverse events or treatment failures), and model-based treatment planning (determining the best transducer position and sonication parameters to deliver the ultrasound energy). Models are also being increasingly used to characterize clinical equipment and as part of regulatory submissions (the US Food and Drug Administration has recently published guidance on the reporting of computational modeling studies that form part of medical device submissions). One major challenge is obtaining sufficiently accurate maps of material properties. Ultimately, models must balance increasing model complexity (e.g., including shear waves) with the effect of parameter uncertainty on the simulated results. The circle of model development discussed here can be used as a guide for those developing models as well as to aid users in the selection and evaluation of models.

Acknowledgments

This work was supported in part by the Engineering and Physical Sciences Research Council, Swindon, UK.

References

Gray, M. D., Stride, E. P., and Coussios, C. C. (2019). Snap, crackle, and pop: Theracoustic cavitation. *Acoustics Today* 15(1), 19-27. https://doi.org/10.1121/AT.2019.15.1.19

Gu, J., and Jing, Y. (2015). Modeling of wave propagation for medical ultrasound: A review. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 62(11), 1979-1992.

Hamilton, M. F., and Blackstock, D. T. (2008). *Nonlinear Acoustics.* Acoustical Society of America, Melville, NY.

Holm, S., and Näsholm, S. P. (2014). Comparison of fractional wave equations for power law attenuation in ultrasound and elastography. *Ultrasound in Medicine and Biology* 40(4), 695-703.

Jaros, J., Rendell, A. P., and Treeby, B. E. (2016). Full-wave nonlinear ultrasound simulation on distributed clusters with applications in high-intensity focused ultrasound. *The International Journal of High Performance Computing Applications* 30(2), 137-155.

Konofagou, E. E. (2017). Trespassing the barrier of the brain with ultrasound. *Acoustics Today* 13(4), 21-26. https://doi.org/10.1121/AT.2017.13.4.21

Martin, E., Ling, Y. T., and Treeby, B. E. (2016). Simulating focused ultrasound transducers using discrete sources on regular Cartesian grids. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 63(10), 1535-1542.

Mast, T. D. (2000). Empirical relationships between acoustic parameters in human soft tissues. *Acoustics Research Letters Online* 1(2), 37-42.

Maxwell, A., Sapozhnikov, O., Bailey, M., Crum, L., Xu, Z., Fowlkes, B., Cain, C., and Khokhlova, V. (2012). Disintegration of tissue using high intensity focused ultrasound: Two approaches that utilize shock waves. *Acoustics Today* 8(4), 24-37. https://doi.org/10.1121/AT.2012.8.4.24

Robertson, J., Martin, E., Cox, B., and Treeby, B. E. (2017b). Sensitivity of simulated transcranial ultrasound fields to acoustic medium property maps. *Physics in Medicine and Biology* 62(7), 2559-2580.

Robertson, J. L., Cox, B. T., Jaros, J., and Treeby, B. E. (2017a). Accurate simulation of transcranial ultrasound propagation for ultrasonic neuromodulation and stimulation. *The Journal of the Acoustical Society of America* 141(3), 1726-1738.

Sapozhnikov, O. A., Tsysar, S. A., Khokhlova, V. A., and Kreider, W. (2015). Acoustic holography as a metrological tool for characterizing medical ultrasound sources and fields. *The Journal of the Acoustical Society of America* 138(3), 1515-1532.

Tabei, M., Mast, T. D., and Waag, R. C. (2002). A k-space method for coupled first-order acoustic propagation equations. *The Journal of the Acoustical Society of America* 111(1), 53-63.

Verweij, M. D., Treeby, B. E., van Dongen, K. W. A., and Demi, L. (2014). Simulation of ultrasound fields. In D. Panetta and M. Demi (Eds.). *Comprehensive Biomedical Physics*. Elsevier, Amsterdam, pp. 465-500.

Wear, K. A., Gammell, P. M., Maruvada, S., Liu, Y., and Harris, G. R. (2014). Improved measurement of acoustic output using complex deconvolution of hydrophone sensitivity. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 61(1), 62-75.

BioSketches

Bradley E. Treeby is an Engineering and Physical Sciences Research Council (EPSRC) Early Career Fellow, Department of Medical Physics and Biomedical Engineering, University College London, and coauthor of the open-source k-Wave Toolbox. He received the 2017 R. Bruce Lindsay Award from the Acoustical Society of America for "contributions to the modeling of biomedical ultrasound fields."

Jiri Jaros is an associate professor in the Faculty of Information Technology, Brno University of Technology, Czech Republic, leader of the Supercomputing Technologies Research Group, and coauthor of the k-Wave Toolbox. For more than a decade, he has been working in the area of high-performance computing, scientific code development, parallel and distributed algorithms, and numerical simulations.

Eleanor Martin is a research associate in the Department of Medical Physics and Biomedical Engineering, University College London, and an expert in ultrasound metrology. Her current research interests include model validation and the design and characterization of multielement arrays for therapeutic ultrasound.

Ben T. Cox is a reader in the Department of Medical Physics and Biomedical Engineering, University College London, Undergraduate Programme Director of Medical Physics, and coauthor of the k-Wave Toolbox. He teaches ultrasound and biomedical optics, and his research interests are in ultrasound and photoacoustic imaging.

Acoustics Today in the Classroom?

There are now over 250 articles on the *AT* web site (AcousticsToday.org). These articles can serve as supplemental material for readings in a wide range of courses. *AT* invites instructors and others to create reading lists. Selected lists may be published in *AT* and/or placed in a special folder on the *AT* web site to share with others.

If you would like to submit such a list, please include:

- Your name and affiliation (include email)
- The course name for which the list is designed (include university, department, course number)
- A brief description of the course
- A brief description of the purpose of the list
- Your list of *AT* articles (a few from other ASA publications may be included if appropriate for your course). Please embed links to the articles in your list.

Please send your lists to the *AT* editor, **Arthur Popper (apopper@umd.edu)**

Early-Career Acousticians Retreat (EAR)

The Acoustical Society of America (ASA) seeks to engage and foster members by hosting the Early-Career Acousticians Retreat (EAR) 2019! EAR is a two-day workshop for early career professionals in the field of acoustics focused on developing leadership and networking skills for early career professionals in the field of acoustics. The workshop also will allow you to connect and socialize with your fellow early career acousticians as well as more senior members of the Society, learn about mentoring relationships and about the Society, and contribute to the future of ASA.

Registration for EAR 2019 is FREE for up to 30 registrants. FREE Registration includes 3 meals as well as $500 towards lodging and transportation. After 30 participants, registration costs $150.

The workshop will be held at the Hotel del Coronado in San Diego, CA, beginning on Friday, December 6th, 2019 at 3:30 P.M. and ending on Saturday, December 7th, 2019 at 1:00 P.M.

Applicants must be within 10 years of their last degree and not currently a student. Applications are due July 15, 2019 by 5 P.M. EST.

Apply For EAR Here: **forms.gle/Wc4YVKXau28r1xrRA**